



1

NIST Internal Report NIST IR 8460 ipd

2

3

State Machine Replication and Consensus with Byzantine Adversaries

4

5

6

Michael Davidson

7

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.IR.8460.ipd>

8



9

10

NIST Internal Report NIST IR 8460 ipd

11

12

State Machine Replication and Consensus with Byzantine Adversaries

13

14

15

Michael Davidson
*Computer Security Division
Information Technology Laboratory*

16

17

18


This publication is available free of charge from:
<https://doi.org/10.6028/NIST.IR.8460.ipd>

19

20

April 2023

21



22

U.S. Department of Commerce
Gina M. Raimondo, Secretary

23

24

National Institute of Standards and Technology

25

Laurie E. Locascio, NIST Director and Under Secretary of Commerce for Standards and Technology

26 Certain commercial equipment, instruments, software, or materials, commercial or non-commercial, are
27 identified in this paper in order to specify the experimental procedure adequately. Such identification does
28 not imply recommendation or endorsement of any product or service by NIST, nor does it imply that the
29 materials or equipment identified are necessarily the best available for the purpose.

30 There may be references in this publication to other publications currently under development by NIST in
31 accordance with its assigned statutory responsibilities. The information in this publication, including
32 concepts and methodologies, may be used by federal agencies even before the completion of such
33 companion publications. Thus, until each publication is completed, current requirements, guidelines, and
34 procedures, where they exist, remain operative. For planning and transition purposes, federal agencies may
35 wish to closely follow the development of these new publications by NIST.

36 Organizations are encouraged to review all draft publications during public comment periods and provide
37 feedback to NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at
38 <https://csrc.nist.gov/publications>.

39 **NIST Technical Series Policies**
40 [Copyright, Use, and Licensing Statements](#)
41 [NIST Technical Series Publication Identifier Syntax](#)

42 **Publication History**
43 Approved by the NIST Editorial Review Board on YYYY-MM-DD [Will be added upon final publication]

44 **How to cite this NIST Technical Series Publication:**
45 Davidson M (2023) State Machine Replication and Consensus with Byzantine Adversaries. (National
46 Institute of Standards and Technology, Gaithersburg, MD), NIST IR 8460 ipd.
47 <https://doi.org/10.6028/NIST.IR.8460.ipd>

48 **NIST Author ORCID iDs**
49 Michael Davidson: 0000-0002-4862-5697

50 **Public Comment Period**
51 April 26, 2023 - September 1, 2023

52 **Submit Comments**
53 ir8460-comments@nist.gov

54 **All comments are subject to release under the Freedom of Information Act (FOIA).**

55 **Abstract**

56 The objective of state machine replication (SMR) is to emulate a centralized service in a
57 distributed, fault-tolerant fashion. To this end, a set of mutually distrusting processes must
58 agree on the execution of client-submitted commands. Since the advent of Bitcoin, the
59 idea of SMR has received significant attention. This document surveys both classical and
60 more modern research on SMR and details many of the most significant permissioned and
61 permissionless algorithms, their performance, and security considerations.

62 **Keywords**

63 atomic broadcast; Bitcoin; blockchain; Byzantine Fault Tolerance; consensus; cryptocur-
64 rency; distributed ledger technology; Ethereum; state machine replication.

65 **Note to Reviewers**

66 Figure 22a © IFCA 2014.

67 Figure 28 © IFCA 2015.

68 Figure 37 © IFCA 2020.

69 Figure 42 © IFCA 2018.

70 Figure 51 © IFCA 2021.

71 Figure 56 © IFCA 2020.

74 **Table of Contents**

75	Executive Summary	1
76	0.1. Purpose and Scope	1
77	0.2. Notes on Terminology	2
78	0.3. Document Structure	3
79	1. Introducing the Problems	4
80	1.1. The Byzantine Generals Problem	4
81	1.2. Broadcast Problems and Byzantine Agreement	5
82	1.3. State Machine Replication (SMR)	7
83	1.4. The Adversary	9
84	1.5. Timing Assumptions	10
85	1.6. Permissioned vs. Permissionless	13
86	2. System Components	14
87	2.1. Data Structures for Distributed Ledgers	15
88	2.2. Sybil-Resistance Mechanism	15
89	2.3. Leader Election and Committee Selection	16
90	2.4. Fork-Choice or Chain Selection Rules	17
91	2.5. Networking	18
92	2.6. Incentive Mechanism	19
93	2.7. Cryptographic Primitives	19
94	2.8. State Machine	21
95	2.8.1. UTXO vs. Account Model	22
96	2.8.2. Changing the Rules	23
97	3. Scaling and "Decentralization"	28
98	3.1. A Note on Decentralization	28
99	3.2. Full Nodes and Light Clients	30
100	3.3. Scalability Challenges and Block Sizes	33
101	4. Practical Byzantine Fault Tolerance (PBFT)	39
102	4.1. PBFT View Change	41
103	4.2. PBFT Security	42
104	4.3. Zyzzyva and Speculative Execution	43

105	4.4. A Permissioned DAG: Blockmania	45
106	5. Modern High-Performance Blockchains	48
107	5.1. Streamlined Blockchains	48
108	5.2. PiLi and PaLa	49
109	5.3. HotStuff	51
110	5.3.1. Sync HotStuff	54
111	5.4. Further Optimizing Latency	55
112	6. Asynchronous BFT	55
113	6.1. HoneyBadgerBFT	56
114	6.1.1. Mostéfaoui et al.'s Asynchronous Binary Agreement Protocol . . .	58
115	6.1.2. Reducing HoneyBadgerBFT's latency with BEAT	60
116	6.1.3. Improving ACS Performance with Dumbo	61
117	6.2. An Asynchronous Permissioned DAG: Hashgraph	64
118	7. Miscellaneous Permissioned BFT	66
119	7.1. Fairly Ordering Transactions	66
120	7.2. Accountability Against Malicious Replicas	68
121	7.3. Specially Designated Roles for Replicas	69
122	7.4. Deterministic Longest Chain Protocols	69
123	7.5. Flexible BFT	71
124	7.6. View Change Algorithms	73
125	8. Localizing Trust Over Incomplete Networks With Open Membership	76
126	8.1. Stellar	77
127	8.1.1. FBAS Background	77
128	8.1.2. Stellar Consensus Protocol (SCP)	78
129	8.1.3. SCP Security	80
130	8.2. Ripple	82
131	8.3. Cobalt	84
132	8.3.1. Background	84
133	8.3.2. Broadcast in Incomplete Networks	86
134	9. Proof of Work: The Basics	88
135	9.1. Proof of Work and Sybil Resistance	88

136	9.1.1. Mining Pools	91
137	9.1.2. Hardware: ASICs and ASIC Resistance	93
138	9.1.3. Mining Centralization in Practice	97
139	9.2. Difficulty Adjustment Algorithms	99
140	9.3. Attacks Against Mining Pools: Pool-Hopping and Block Withholding . . .	100
141	9.4. Selfish Mining	103
142	10. Nakamoto Consensus	106
143	10.1. Theory of Nakamoto Consensus	108
144	10.1.1. Nakamoto Consensus With Chains of Variable Difficulty	111
145	10.1.2. Additional Analyses of Nakamoto Consensus	112
146	10.2. Violating the Nakamoto Consensus Security Assumptions	115
147	10.2.1. Network Delay and Block Propagation	115
148	10.2.2. Majority Hash Rate Attacks (51% Attacks)	118
149	10.2.3. Hash Function Collisions	120
150	10.3. (More) Attacks Against Nakamoto Consensus	120
151	11. More Proof-of-Work Protocols	123
152	11.1. Nakamoto Consensus Protocol Adjustments	124
153	11.1.1. Weak Blocks and Pre-Consensus	124
154	11.1.2. Bitcoin-NG	125
155	11.1.3. Tie-Breaking Schemes	126
156	11.1.4. DECOR+	127
157	11.1.5. Publish or Perish	128
158	11.1.6. NC-Max	129
159	11.2. Greedy Heaviest-Observed Sub-Tree (GHOST)	130
160	11.3. FruitChains	131
161	11.4. Parallel Chain Approaches	133
162	11.4.1. Prism	134
163	11.5. Proof-of-Work DAGs	135
164	11.5.1. Inclusive Blockchains and Conflux	136
165	11.5.2. SPECTRE and Phantom	138
166	11.5.3. Tangle	144

167	11.5.4. Meshcash	147
168	11.6. Proof of Work for Committee Selection	150
169	11.6.1. Hybrid Consensus	150
170	11.6.2. Solida	151
171	12. Proof of Stake: The Basics	154
172	12.1. Early Attempts at Proof of Stake	155
173	12.1.1. Nothing-at-Stake and Costless Simulation	159
174	12.1.2. Long-Range Attacks, Posterior Corruption, and Weak Subjectivity	161
175	12.1.3. Leader Election, Anonymity, and Security Against Adaptive Adver-	
176	saries	164
177	12.2. Leader Predictability and Security	168
178	12.3. Wealth Concentration, Block Rewards, and Centralization	173
179	13. Proof-of-Stake Protocols	178
180	13.1. Chain-Based Proof of Stake	178
181	13.1.1. Chains of Activity	178
182	13.1.2. Snow White	180
183	13.1.3. Ouroboros Family: Praos and Genesis	182
184	13.1.4. DFINITY	185
185	13.2. Ethereum 2.0	187
186	13.3. DAG-based Proof of Stake	189
187	13.3.1. Fantômette	189
188	13.3.2. Avalanche	194
189	13.3.3. Parallel Chains	197
190	13.4. BFT-Based Proof of Stake	198
191	13.4.1. Tendermint	198
192	13.4.2. Algorand	200
193	14. Hybrid and Alternative Sybil-Resistance Mechanisms	204
194	14.1. Proof of Space	204
195	14.1.1. Spacemint	205
196	14.1.2. Chia	206
197	14.2. Proof of Activity	209
198	14.3. Checkpoints and Finality Gadgets	210

199	14.3.1. Ad Hoc Finality Layers and Reorg Protection	210
200	14.3.2. Casper the Friendly Finality Gadget (FFG)	212
201	14.3.3. More Finality Gadgets and Checkpointing Protocols	216
202	15. Sharding	220
203	15.1. Intra-Shard Consensus	221
204	15.2. Identity Registration, Committee (Re)configuration, and Epoch Randomness	222
205	15.3. Cross-Shard Transaction Processing	226
206	15.4. A Different Approach: Monoxide	229
207	15.5. Fraud Proofs and Data Availability	230
208	16. Interoperability	235
209	16.1. Cross-Chain Communication, Fair Exchange, and Atomic Swaps	235
210	16.2. Bootstrapping Methods: Merged Mining and Proof of Burn	239
211	16.3. Sidechains, Relays, and Asset Transfer	241
212	16.3.1. Permissionless Sidechains	243
213	17. Networking	246
214	17.1. Networking for Permissionless Systems	246
215	17.1.1. Peer Discovery	247
216	17.1.2. Neighbor Selection	248
217	17.1.3. Communication Strategy	249
218	18. State Machines	250
219	18.1. Virtual Machine Design	250
220	18.1.1. Concurrency in Smart Contracts	253
221	18.1.2. Zero-Knowledge Proofs and Verifiable Computation	255
222	18.1.3. Delegating Execution	256
223	18.2. Layer 2 Protocols	259
224	18.2.1. Payment and State Channels	259
225	18.2.2. Plasma and Rollups	263
226	19. Incentives	265
227	19.1. Block Rewards: Subsidies and Transaction Fees	265
228	19.1.1. The Mining Gap and (the Absence of a) Block Subsidy	266
229	19.2. State Machines, Incentives, and Security	269

230	19.3. Alternative Transaction Fee Protocols	273
231	References	277

232 **List of Tables**

233	Table 1. Percentages of eligible tokens actively staked	176
-----	---	-----

234 **List of Figures**

235	Fig. 1. Blockchain vs. DAG	16
236	Fig. 2. Forking and the Longest Chain Rule	18
237	Fig. 3. UTXO transaction graph	23
238	Fig. 4. Hard forks and soft forks	24
239	Fig. 5. Simplified Payment Verification (SPV)	31
240	Fig. 6. PBFT normal case operation	40
241	Fig. 7. Zyzzyva's speculative execution	44
242	Fig. 8. Blockmania state machine examples	47
243	Fig. 9. Streamlet finalization rule	49
244	Fig. 10. Pipelining in Chained HotStuff	53
245	Fig. 11. Chained HotStuff justification	54
246	Fig. 12. HoneyBadgerBFT's ACS structure	57
247	Fig. 13. ACS structure of Dumbo protocols	62
248	Fig. 14. Hashgraph strongly seeing example	65
249	Fig. 15. Committee reconfiguration attack	76
250	Fig. 16. Federated voting stages	78
251	Fig. 17. Cascade effect in federated voting	79
252	Fig. 18. FBAS Quorum intersection insufficient for safety	80
253	Fig. 19. Ripple "support" example	84
254	Fig. 20. Bitcoin mining and AsicBoost	90
255	Fig. 21. Hardware hashrate asymmetry	96
256	Fig. 22. Selfish mining strategy	103
257	Fig. 23. Compact Blocks	117
258	Fig. 24. Double-spend probabilities	121
259	Fig. 25. Finney attack example	122
260	Fig. 26. Vector76 attack example	123
261	Fig. 27. NC-Max block propagation mechanism	129
262	Fig. 28. GHOST fork choice rule	130
263	Fig. 29. FruitChains architecture	132
264	Fig. 30. Prism structure	135
265	Fig. 31. Conflux example	138
266	Fig. 32. SPECTRE voting example.	140
267	Fig. 33. Phantom example	143
268	Fig. 34. Parasite chain attack against the Tangle	146
269	Fig. 35. Large weight attack against the Tangle	146

270	Fig. 36. Meshcash example	149
271	Fig. 37. Stake shift for select cryptocurrencies	162
272	Fig. 38. Predictable bribe attacks against proof of stake	171
273	Fig. 39. Undetectable Nothing-at-Stake Attack	171
274	Fig. 40. Latest Message Driven (LMD) GHOST	188
275	Fig. 41. Ethereum 2.0 Randoa architecture	190
276	Fig. 42. Spacemint grinding defense	206
277	Fig. 43. Chia design	207
278	Fig. 44. Chia grinding attack	208
279	Fig. 45. Casper FFG attacks	215
280	Fig. 46. Sharding architecture	220
281	Fig. 47. Wormhole shard allocation	226
282	Fig. 48. Cross-shard transactions	227
283	Fig. 49. Invalid shard state transition	231
284	Fig. 50. Data availability attack	232
285	Fig. 51. Cross-chain communication	236
286	Fig. 52. Atomic swap	238
287	Fig. 53. Merged mining	240
288	Fig. 54. Sidechain pegging methods	244
289	Fig. 55. Execute-Order-Validate	258
290	Fig. 56. Lightning Network channel closing	261
291	Fig. 57. Lightning Network payment	262
292	Fig. 58. Mining gap	267
293	Fig. 59. Undercutting attack	268
294	Fig. 60. Front-running and Miner Extractable Value	270

295 **Acknowledgments**

296 The author thanks Tyler Diamond, Frederic de Vault, and Andrew Regenscheid for their helpful
297 comments.

298

299 **Executive Summary**

300 Since the deployment of Bitcoin on January 3rd, 2009, and its description by the pseudony-
301 mous Satoshi Nakamoto in 2008 [1], research and development of new, practical state ma-
302 chine replication (SMR) systems have surged. It has been stated that Bitcoin provided a
303 solution to the "Byzantine Generals Problem." While not strictly true, it is a useful starting
304 point for this document's analysis of consensus algorithms, state machine replication, and
305 distributed ledger technology (DLT).

306 More generally, the goal of these types of problems is to allow a set of mutually distrusting
307 processes (e.g., computer processes) to agree on the outcome of some deliberation despite
308 the possibility that some of them are faulty or even malicious. In essence, the goal is to
309 provide some service to clients that emulates a centralized service while operating as a dis-
310 tributed server. There are a variety of ways to formulate this problem (several are described
311 in Section 1), but they all require some notion of agreement between the distributed pro-
312 cesses. Research in this area began in the early 1980s when the Byzantine Agreement [2]
313 and Byzantine Generals [3] problems were formulated. The first algorithms to solve these
314 problems were extremely inefficient, and real-world usage did not become plausible until
315 the celebrated Practical Byzantine Fault Tolerance (PBFT) algorithm was invented in 1999
316 [4].

317 Until the advent of Bitcoin, it was believed that all distributed agreement algorithms re-
318 quired a fixed set of identifiable participants known in advance. Bitcoin was the first pro-
319 tocol to demonstrate that consensus can be maintained across distributed processes in an
320 open network with free entry and no fixed identifiers. Today, the terms *permissioned* and
321 *permissionless* are used to describe the difference between the two models.

322 Bitcoin also popularized the idea of a *blockchain* as the data structure over which the dis-
323 tributed processes maintain agreement. Since then, the concept has been generalized to
324 *distributed ledgers* more broadly. A blockchain is an ordered list of client-submitted com-
325 mands, or transactions, that modify the system state. Because all participants execute the
326 same agreed-upon commands in the same order, participants are able to maintain a common
327 view of the execution of a protocol-defined state machine. In addition to cryptocurrencies,
328 state machine replication and the *smart contracts* they enable have been suggested for use
329 in trade settlement, finance, identity management, supply chains, healthcare, Internet of
330 Things (IoT), and other industries.

331 **0.1. Purpose and Scope**

332 This document is intended to serve as an advanced treatment of consensus algorithms, state
333 machine replication, and distributed ledger technology. It may also function as a reference
334 for consensus algorithms as it contains fairly detailed descriptions of a variety of algorithms
335 that may be useful in different scenarios. The reader is expected to already have a high-
336 level understanding of distributed ledger technology, such as that provided by NIST IR

337 8202, *Blockchain Technology Overview* [5].

338 This document first discusses the properties required of distributed consensus systems and
339 the many kinds of subprotocols used to implement them in a variety of system models.
340 Many algorithms, both permissioned and permissionless, are then described in detail. The
341 discussion on permissioned consensus starts with the classic PBFT algorithm but focuses
342 heavily on techniques that have been developed more recently and improve performance or
343 enable security in more challenging environments. Permissionless algorithms are divided
344 into categories based on the Sybil-resistance mechanism employed – that is, proof of work
345 (PoW), proof of stake (PoS), and alternative mechanisms. There is extensive discussion on
346 the unique security issues that arise in each case, the architectural reasons they exist, and
347 techniques that can be used to mitigate them. Finally, a variety of more advanced topics are
348 discussed, including scalability methods such as sharding and "layer 2" technologies, inter-
349 operability, state machine design, networking, and how incentives impact system security.

350 0.2. Notes on Terminology

351 The distributed systems literature is rife with synonyms and inconsistent or imprecise use
352 of terms. When the word "consensus" appears in this document, it is meant as a general-
353 ization that captures all of the agreement problems described, including various broadcast
354 problems, Byzantine Agreement, and state machine replication. The term "broadcast" –
355 when not being used to describe broadcast problems specifically – is meant to convey the
356 idea of simultaneously transmitting a message to multiple peers. In addition, the ill-defined
357 term "decentralized," which is used frequently in the literature, is discussed in more detail
358 in section 3.1.

359 Several groups of synonymous words appear in this document. Most of the time, termi-
360 nology from the original source paper was used. For example, the terms "node," "replica,"
361 "process," "validator," and "miner" are used as synonyms but often in slightly different
362 contexts, such as an entity that participates in consensus. Many protocols have leaders,
363 which may be called the "primary" or "block producer." When a node is not the leader, it
364 may be a "secondary" or "backup." A "malicious" node may be considered "Byzantine,"
365 "faulty," "corrupt," or "dishonest," whereas honest nodes are sometimes called "correct."
366 When nodes eventually agree on a value, it is sometimes said that they "decide," "output,"
367 or "accept" the value.

368 The term "blockchain" is defined in NIST IR 8202, *Blockchain Technology Overview*:

369 Blockchains are distributed digital ledgers of cryptographically signed trans-
370 actions that are grouped into blocks. Each block is cryptographically linked to
371 the previous one (making it tamper evident) after validation and undergoing a
372 consensus decision. As new blocks are added, older blocks become more dif-
373 ficult to modify (creating tamper resistance). New blocks are replicated across
374 copies of the ledger within the network, and any conflicts are resolved auto-

375 matically using established rules. [5]

376 While this is a good description of some specific blockchains, the term is used here to
377 only mean "a chain of blocks," capturing the fact that blocks are cryptographically linked
378 together in a list. This makes blockchains a particular data structure within a broader set of
379 distributed ledgers.

380 0.3. Document Structure

381 The rest of this document is organized as follows:

- 382 • **Sections 1-3** are introductory material. **Section 1** introduces the formal problems
383 that are solved by the protocols described throughout the document and some of the
384 model assumptions under which the problems can be solved. **Section 2** describes the
385 various sub-protocols and components that are often used in designing distributed
386 ledger systems for SMR. **Section 3** discusses the trade-offs between maintaining
387 "decentralization" and the scalability of DLT systems.
- 388 • **Sections 4-8** describe protocols for permissioned consensus. **Section 4** describes
389 Practical Byzantine Fault Tolerance (PBFT), the first system design scalable enough
390 to be used in practice. **Section 5** describes more modern, high-performance consen-
391 sus algorithms. **Section 6** discusses algorithms that are designed for asynchronous
392 networks where messages may be arbitrarily delayed. **Section 7** surveys a variety
393 of extra properties that one might desire from a permissioned consensus algorithm.
394 **Section 8** describes protocols where participants may select their own quorums of
395 trusted replicas and need not be aware of the existence of every replica in the net-
396 work.
- 397 • **Sections 9-11** discuss protocols that use proof of work (PoW) as a Sybil-resistance
398 mechanism. **Section 9** discusses aspects common to most PoW protocols. **Section**
399 **10** describes Nakamoto Consensus – the protocol used by Bitcoin – in detail. **Section**
400 **11** describes a wide variety of PoW consensus designs.
- 401 • **Sections 12-13** discuss protocols that use proof of stake (PoS) as a Sybil-resistance
402 mechanism. **Section 12** provides a historical overview of PoS and the security issues
403 that need to be considered as part of PoS algorithm design. **Section 13** describes a
404 variety of specific PoS protocols.
- 405 • **Section 14** discusses protocols that use alternative Sybil-resistance algorithms, such
406 as proof of space, as well as hybrid mechanisms.
- 407 • **Sections 15-19** cover a variety of more advanced topics related to the design of DLT
408 for SMR. **Section 15** discusses the technical details of one of the more promising
409 scalability methods: sharding. **Section 16** discusses interoperability between sys-
410 tems. **Section 17** covers topics related to the network layer of these protocols, such

411 as how a node discovers new peers and communication strategies. **Section 18** dis-
412 cusses state machine design considerations and some "layer 2" scaling protocols that
413 can be built on an underlying replicated state machine. **Section 19** considers a variety
414 of incentive-related security issues that can arise in these systems.

415 1. Introducing the Problems

416 1.1. The Byzantine Generals Problem

417 The *Byzantine Generals Problem* (BGP) is equivalent to the *Byzantine Broadcast* (BB)
418 problem, which is the more commonly used term in the distributed systems literature. This
419 problem was introduced in [3] as *interactive consistency*.

420 Consider a city besieged by several divisions of the Byzantine army. Each division is led
421 by its own general, and the generals communicate with each other via messenger. The
422 generals must agree upon a common strategy: either attack or retreat. Unfortunately for the
423 Byzantine army, some of the generals may be disloyal and actively working to sabotage the
424 agreement. A solution to the BGP, then, is an algorithm that ensures that:

- 425 • Every loyal general decides upon the same strategy.
- 426 • If the number of traitors is small, the traitors cannot cause the loyal generals to decide
427 on a "bad" plan – that is, a plan they would not have otherwise agreed to in the first
428 place.

429 Stated differently, let there be n total generals (or computer processes), up to f of which
430 may be disloyal (or malicious/faulty). One general in particular, the commanding general,
431 sends an order to $n - 1$ lieutenant generals, such that the following interactive consistency
432 conditions are maintained:

- 433 • **IC1:** All loyal lieutenants obey the same order.
- 434 • **IC2:** If the commanding general is loyal, then all loyal lieutenants obey the com-
435 mander's order.

436 The same paper proved that the BGP is solvable only if more than $\frac{2}{3}$ of the generals are
437 loyal when using only "oral" messages (unsigned messages) and solvable for any number
438 of generals/traitors with "unforgeable written messages" (signed messages) when there is
439 a known fixed upper bound on how long it takes to send a message from one general to
440 another (synchrony). That is, in a synchronous network without using digital signatures,
441 if f generals are traitors, then no algorithm will work without $n > 3f$ total generals. In a
442 synchronous network with signed messages, the problem is solvable with $n > f$ generals.

443 1.2. Broadcast Problems and Byzantine Agreement

444 A variety of broadcast problems exist, including the BGP. In each case, there is a des-
445 ignated sender with an input value that they would like to distribute to the remainder of
446 the processes. An algorithm that solves a broadcast problem will have certain properties,
447 such as the "interactive consistency" requirements described in Section 1.1. In particular,
448 there are requirements around *consistency* (also known as *agreement*), *validity*, *integrity*,
449 and *termination*. The consistency, validity, and integrity properties are the same for most
450 broadcast problems, but the termination property differs. The following are the properties
451 of a Byzantine Broadcast (BB) algorithm, which is equivalent to BGP and has the strictest
452 termination requirement:

- 453 • **Consistency/Agreement:** If two honest replicas decide v and v' , then $v = v'$.
- 454 • **Validity:** If the sender is honest and begins with input value v , then all honest nodes
455 decide v .
- 456 • **Integrity:** Every honest process delivers at most one value, and if it delivers v , then
457 some process must have broadcast v .
- 458 • **Termination:** Every honest process eventually decides some value.

459 Notice that property IC1 from the previous section is equivalent to consistency, while IC2
460 is equivalent to the validity property. By relaxing the termination requirement, different
461 broadcast problems can be described as follows:

- 462 • *Reliable broadcast* (RB or RBC) requires that either all honest processes eventually
463 decide, or no honest process decides. That is, RBC may not terminate if the sender is
464 faulty, but if any honest party obtains an output, all other honest parties must as well.
- 465 • *Byzantine consistent broadcast* (BCB) allows some honest parties to decide without
466 requiring all honest parties to do so.
- 467 • *Terminating reliable broadcast* (TRB) requires correct processes to agree on the re-
468 ceipt of a message or agree that the sender is faulty (so termination must occur, but
469 correct processes need not get a value out of it).

470 Reliable broadcast is commonly used as an underlying communication primitive for more
471 complex distributed protocols, such as multi-party computation (MPC) or consensus. Mes-
472 sages "delivered" by the broadcast algorithm are then used as input messages in an MPC or
473 consensus protocol execution. Note that reliable broadcast does not guarantee agreement
474 over the order of messages, only that the messages are delivered at all.

475 While broadcast algorithms are only mentioned a few times in this document, they have
476 many similarities to the algorithms discussed throughout while being fairly easy to under-
477 stand. Bracha's broadcast is among the most celebrated of reliable broadcast algorithms

478 and was originally proposed in the late 1980s as a technique to convert crash fault toler-
479 ant consensus protocols into ones capable of resisting malicious adversaries, often called
480 *Byzantine* adversaries [6].

- 481 1. The designated sender with starting value v sends a message, $(INIT, v)$, to all pro-
482 cesses (here, $n = 3f + 1$). Then each process executes the following three steps.
- 483 2. Wait until the receipt of either one $(INIT, v)$ message, $2f + 1$ $(ECHO, v)$ messages,
484 or $f + 1$ $(READY, v)$ messages for some v . Then send $(ECHO, v)$ to all processes.
- 485 3. Wait until the receipt of $2f + 1$ $(ECHO, v)$ messages or $f + 1$ $(READY, v)$ messages
486 for some v , including messages received in the previous step. Then send $(READY, v)$
487 to all processes.
- 488 4. Wait until the receipt of $2f + 1$ $(READY, v)$ messages for some v , including ones
489 received in the prior steps. Then deliver v .

490 Bracha's broadcast algorithm is a solution to the reliable broadcast problem so long as $n \geq$
491 $3f + 1$ and protocol messages eventually reach their destination. Consider the following
492 arguments:

- 493 1. Let p and q be two correct processes, and assume that process p is the first to send a
494 $(READY, v)$ message, while process q is the first to send a $(READY, v')$ message.
495 Assume (by contradiction) that $v \neq v'$. Since process p is the first to have sent
496 $(READY, v)$, it must have seen at least $2f + 1$ $(ECHO, v)$ messages. Similarly, pro-
497 cess q must have seen at least $2f + 1$ $(ECHO, v')$ messages. In total, $4f + 2$ $ECHO$
498 messages were sent. Because $n \geq 3f + 1$, at least $f + 1$ replicas must have sent both
499 $(ECHO, v)$ and $(ECHO, v')$ messages. However, this implies that an honest replica
500 sent both an $(ECHO, v)$ and an $(ECHO, v')$ message. Because honest replicas do
501 not equivocate (they only send one message of each type), this is a contradiction.
502 Therefore, $v = v'$.
- 503 2. Now assume that process p delivers a value v , and process q delivers a value v' . This
504 means that process p saw at least $2f + 1$ $(READY, v)$ messages, of which at least
505 $f + 1$ are from honest replicas. Analogously, process q must have seen at least $f + 1$
506 $(READY, v')$ messages from honest replicas. By argument 1, this implies that $v = v'$.
- 507 3. If an honest process p delivers the value v , then every other honest process will
508 eventually deliver v . Because p accepts v , it must be the case that p has seen $2f + 1$
509 $(READY, v)$ messages, and at least $f + 1$ of those were from honest processes. This
510 implies that every other process will see at least $f + 1$ $(READY, v)$ messages and then
511 send their own $(READY, v)$ message (because honest processes cannot equivocate).
512 Therefore, at least $n - f \geq 2f + 1$ processes send a $(READY, v)$ message, so all honest
513 processes will eventually see $2f + 1$ $(READY, v)$ messages and deliver v .
- 514 4. If the designated sender p is honest and broadcasts v , all honest processes eventually

515 decide v . Every honest process will receive the sender's $(INIT, v)$ message and then
516 send an $(ECHO, v)$ message. Then every honest process q will receive $n - f \geq 2f + 1$
517 $(ECHO, v)$ messages from other honest processes (and at most f other messages
518 from faulty processes). Process q will then send a $(READY, v)$ message. In the final
519 step of the protocol, all honest processes will receive $n - f \geq 2f + 1$ $(READY, v)$
520 messages (and at most f ready messages from faulty processes). Therefore, honest
521 process q will deliver v .

522 Putting this all together, argument 4 shows that if the designated sender p is honest and
523 broadcasts v , all honest replicas will deliver v . On the other hand, if the designated sender
524 p is faulty and some honest process q delivers v , then all honest processes deliver v by
525 argument 3. Otherwise (p is faulty and q does not accept v), no honest process will accept
526 any value. This proves that Bracha's algorithm achieves reliable broadcast.

527 *Byzantine agreement* (BA), introduced in [2], is a closely related problem to broadcast
528 but with one crucial difference: instead of having a dedicated sender who disseminates a
529 value to the rest of the network, every process starts with an initial value. This redefined
530 problem has many synonyms in the literature, including *consensus*, *total order broadcast*,
531 and *atomic broadcast*. Often, the term "consensus" is used specifically for one instance
532 of BA (*single-shot BA*), while "atomic broadcast" is used in repeated or sequential BA,
533 where the agreed upon values are also in an agreed upon order. This document uses the
534 terms interchangeably but primarily focuses on repeated BA, which is needed for state
535 machine replication. When BA is used to agree on a single bit, it is called *binary Byzantine*
536 *agreement*, whereas it is called *multi-value Byzantine agreement* (MVBA) when there are
537 more than two possible choices.

538 The consistency, integrity, and termination properties are the same in BA and BGP, but the
539 validity property necessarily changes to reflect the lack of a designated sender. Specifically,
540 the validity property of BA is that if all correct processes propose the same value v , then
541 any correct process must decide v . There can also be *weak validity* (in contrast to the
542 "strong" validity just mentioned): for each correct process, its output must be the input of
543 some correct process. In either case, validity may also require the decided value to satisfy
544 an external predicate, like a valid digital signature or other (state machine) rules. In this
545 document, most protocols aim for strong validity.

546 1.3. State Machine Replication (SMR)

547 State machine replication is built off of atomic broadcast, which guarantees that each
548 agreed-upon value is also in an agreed-upon order. These values are typically called *trans-*
549 *actions* or commands. Processes send these transactions to each other, which are then
550 placed in a total ordering via atomic broadcast and then executed by the processes in that
551 order. Transactions operate on some global state and transform the state via a determinis-
552 tic program. SMR guarantees lock-step execution of identical commands and agreement
553 over state by all honest processes. The SMR approach was first described in [7] but was

554 popularized in [8].

555 Distributed ledger technology (DLT) and blockchains are specific examples of state ma-
556 chine replication. Given a blockchain protocol, one can derive the SMR by having replicas
557 execute the blockchain protocol, having honest nodes broadcast all transactions they see to
558 each other, and – for Bitcoin and many related protocols – removing some number of trail-
559 ing *blocks* (sets of transactions). Such a distributed ledger provides the following properties
560 [9]:

- 561 • **Persistence:** Once a transaction is at least k blocks deep into the ledger of an honest
562 node (where k is a security parameter), it will be included in the same permanent
563 position in the ledger of every honest node with overwhelming probability.
- 564 • **Liveness:** All transactions from honest clients will eventually be at a depth of more
565 than k blocks in an honest node’s blockchain.

566 A distributed ledger can be shown to satisfy persistence and liveness if it satisfies the *com-*
567 *mon prefix*, *chain quality*, and *chain growth* properties, which are discussed in more detail
568 in Section 10.1. Informally, the common prefix property says that the blockchains of two
569 honest nodes differ only in their last k blocks from the chain tip. Chain quality is the prop-
570 erty that "enough" of the blocks that wind up in the blockchain were proposed by honest
571 nodes. Finally, chain growth means that the blockchains accepted by honest nodes contin-
572 uously grow at a certain pace.

573 It may be tempting to think that state machine replication can be realized by simply exe-
574 cuting a BA protocol repeatedly in serial, but this is false due to an issue with how BA’s
575 validity property is defined [10, 11]. In BA, validity ensures that if all honest replicas had
576 a particular input value, then that is the value of the output as well. This does not suffice
577 for SMR, where each replica maintains a local buffer for transactions they have received
578 but have not agreed upon yet (often called the *mempool*), and it cannot be guaranteed that
579 honest parties will initiate BA with the same transaction set. As a result, achieving validity
580 would likely come at the expense of liveness, as it would be challenging for honest replicas
581 to agree on anything other than empty blocks. This is not a problem for distributed ledgers,
582 which can base their validity on enforcing the predicate that whatever state machine rules
583 exist must be followed for a block to be considered valid.

584 One of the primary use cases for SMR is payments, or secure asset transfer. For this
585 specific use case, enforcing a total order on transactions is not strictly necessary; some
586 form of reliable broadcast is sufficient [12]. As long as user accounts are associated with
587 distinct owners, then the owner of an account determines the order of transfers out of their
588 account with no need to agree on the ordering with other users. Others only need to verify
589 that the owner’s decisions maintain any needed causal relations among accounts (e.g., that
590 the order does not create a scenario where an account transfers out more units of an asset
591 than it possesses at the time). These causal relations establish a *partial ordering* rather
592 than a total ordering. Several payment systems based on this idea have been proposed, and

593 they are usually highly performant by eliminating the requirement of agreeing on a fully
594 linearized ordering of transactions [13–19].

595 Relying on a partial ordering has trade-offs. Most notably, without enforcing a total order-
596 ing on transactions, the majority of complex smart contracts (e.g., on-chain cryptocurrency
597 exchanges) are no longer possible. If transactions are partially ordered, they may attempt
598 to access and modify the same system state concurrently, which would lead to inconsisten-
599 cies. Another consequence of replacing BA with reliable broadcast is that the termination
600 property of reliable broadcast is not guaranteed if the designated sender is faulty. In this
601 case, it means that if the spender in a transaction submits two conflicting payments to the
602 network, it is possible that neither transaction is ever executed. This is not a problem if
603 one assumes that a client who attempts to double-spend in this way is malicious. However,
604 in practice, this behavior may be non-malicious, in which case an innocent user may have
605 their funds frozen permanently as honest nodes fail to agree on a transaction to include in
606 the ledger. For example, when the Bitcoin network is congested, some clients will use a
607 technique called *replace-by-fee* (RBF) to reissue a transaction with a higher fee attached
608 with the intention of "jumping the line" and having miners include the transaction in the
609 blockchain more quickly. A payment system with partial ordering is unable to handle this
610 scenario.

611 1.4. The Adversary

612 When discussing the security of various consensus protocols, it is important to consider the
613 powers that an adversary has available to disrupt the protocol. It is generally assumed that
614 the adversary in a distributed system controls f parties and has the ability to coordinate
615 them. For instance, in a protocol that has $n = 10$ parties and can tolerate up to $f = 3$ faulty
616 ones, a single adversary controls all faulty parties (or faulty parties are in collusion). This
617 assumption is favorable to the attacker and conservative for the protocol because – in the
618 real world – the faulty parties could be controlled by multiple adversaries. That would
619 correspond to a situation where, say, each of the three faulty parties were hacked by a
620 different adversary. If the protocol is secure against the combined adversary, it would also
621 be secure against a less coordinated set of adversaries.

622 The adversary can be either *static* or *adaptive*. In the static model, the adversary corrupts its
623 f parties at the beginning of the protocol execution, and those specific parties remain cor-
624 rupt for the duration of the execution. Alternatively, an adaptive adversary can observe the
625 protocol execution and corrupt parties "on the fly." That is, one can think of the adversary
626 as having a "corruption budget" of f , which can be used throughout the protocol execution
627 and chosen to the adversary's advantage based on the messages seen. Attaining provable
628 security against an adaptive adversary is significantly more challenging than against a static
629 adversary.

630 Various models exist regarding what specific powers the adversary has over the corrupted
631 replicas. The two most common types of faults that a protocol may be designed to tolerate

632 are *crash faults* and *Byzantine faults*. In a crash fault, the faulty party simply stops par-
633 ticipating in the protocol. This corresponds to victims of denial-of-service attacks as well
634 as more benign crashes. It is assumed that these parties never recover. In this document,
635 the focus is on Byzantine faults and Byzantine fault-tolerant (BFT) protocols as opposed to
636 crash fault-tolerant (CFT) ones. A Byzantine fault is any arbitrary fault, including deliber-
637 ately malicious ones, such as sending equivocating messages to try to disrupt the protocol.
638 Other possibilities exist that are outside of the scope of this document, such as *omission*
639 *faults* (not sending messages when supposed to), and crash failures *with* recovery. Faults of
640 each type can also be caused by software bugs rather than any particular malicious action.

641 Protocol designers often focus on two types of participants: 1) the honest ones, who faith-
642 fully execute the protocol as specified no matter what, and 2) the Byzantine ones, who
643 are controlled by the adversary and deliberately attempt to arbitrarily subvert the protocol
644 (whether this is beneficial to the adversary or not). However, there may be cases in which
645 honest behavior is irrational and not in the best interest of the participants. The assumption
646 that a particular fraction of the participants behave honestly may be unlikely to hold in
647 practice. Traditional security models typically ignore this.

648 In contrast, the BAR model [20] considers three types of players: 1) Byzantine players; 2)
649 altruistic players, who always behave honestly even if it is irrational for them to do so; and
650 3) rational players, who will act selfishly and deviate from the protocol in order to increase
651 their utility but will not arbitrarily deviate. This is a more challenging model for a protocol
652 to be secure against because it must also be in the best interest of players to behave honestly
653 based on some utility function. This involves the use of game-theoretic techniques that are
654 beyond the scope of this document. It is also possible to design protocols without any
655 altruistic players [21].

656 1.5. Timing Assumptions

657 Among the most important factors in evaluating a consensus protocol are the assumptions it
658 makes on the timing of message delivery. Naturally, more conservative timing assumptions
659 make the protocol more robust but impose stricter requirements on its design and may
660 negatively impact performance. In each case, it is assumed that if an honest party sends a
661 message to another honest party, it is eventually received, even though it may take a while or
662 arrive out of order (perhaps an adversarially chosen order). Here, time is usually measured
663 in terms of rounds of protocol execution. The three most common network timing models
664 are *synchronous*, *asynchronous*, and *partially synchronous*.

- 665 • **Synchronous:** There is a fixed known upper bound, Δ , on the time it takes for mes-
666 sages to be sent from one processor to another. If an honest replica sends a message
667 to another honest replica in round r , then the recipient will have seen the message by,
668 at latest, the beginning of round $r + \Delta$. In synchronous networks, there can be a *rush-*
669 *ing adversary* who acts last in each round of the protocol and can see all messages
670 sent by honest parties before deciding what to do with their control over corrupted

671 participants. Under synchrony, agreement is possible when the majority of partici-
672 pants are honest ($n \geq 2f + 1$), while broadcast simply requires that $f < n$, assuming
673 a PKI and digital signatures.

674 • **Asynchronous:** There is no fixed bound on the time it takes for messages to be
675 sent from one honest processor to another. There may or may not be guaranteed
676 delivery eventually, but if not, consensus is impossible due to the important "FLP
677 impossibility" result [22]. If messages are guaranteed to be delivered eventually, then
678 secure protocols can be designed. The optimal resilience for asynchronous consensus
679 protocols is $n \geq 3f + 1$, and safety and liveness are both maintained so long as this
680 holds. Progress occurs as messages arrive rather than based on fixed rounds.

681 The celebrated FLP impossibility result shows that no asynchronous and determin-
682 istic protocol can achieve consensus with even one faulty processor, and this fault
683 need not be Byzantine. A simple crash can prevent consensus due to the inability to
684 ensure that the protocol will terminate. This result can be circumvented in two ways:
685 by using randomization as part of the algorithm [23, 24] or by providing probabilistic
686 termination (with probability 1) [25]. Probability 1 does not imply that all executions
687 terminate, but that the set of non-terminating executions is extremely improbable in
688 the limit. The difference between probabilistic termination and using randomiza-
689 tion is that the protocol itself does not incorporate randomness in probabilistically
690 terminating protocols.

691 Asynchronous protocols may be even easier to implement than synchronous or par-
692 tially synchronous protocols because they do not require dealing with explicit time-
693 outs. The implementation can be entirely message-driven.

694 • **Partially synchronous:** This model was introduced in [26], which presented two
695 conceptions of the idea:

696 1. There is a fixed but unknown upper bound, Δ , on the time it takes for mes-
697 sages to be sent from one process to another. This is sometimes called *semi-*
698 *synchronous* or the *bounded delay* model.

699 2. There is a fixed, known upper bound, Δ , for message transmission, but this
700 bound is only guaranteed to hold after some unknown time called the *Global*
701 *Stabilization Time* (GST). That is, there is a period of asynchrony (where mes-
702 sages may be lost) followed by a period of synchrony.

703 Since messages can be lost prior to GST, they must be re-sent every round to ensure
704 that they are eventually received. This implies that there may be unbounded commu-
705 nication costs prior to GST [27]. In the bounded delay variant, messages cannot be
706 lost, so only a single message transmission is needed.

707 This model decouples the safety/consistency and liveness properties of the system. If
708 the underlying network is in fact asynchronous, then a partially synchronous protocol

709 will continue to maintain safety but lose liveness and stall. The optimal resilience for
710 partially synchronous protocols is $n \geq 3f + 1$, and safety is maintained so long as
711 this holds.

712 Synchronous consensus algorithms are desirable for their superior security bounds com-
713 pared to asynchronous and partially synchronous protocols ($n \geq 2f + 1$ for synchronous
714 but $n \geq 3f + 1$ otherwise). However, synchronous consensus protocols are inconvenient
715 for at least two other reasons. First, latency is dependent on the estimated network delay
716 bound Δ , which creates a trade-off between security and latency (*responsive* protocols can
717 solve this; see Section 5.2). Second, synchronous protocols are unable to tolerate network
718 partitions, which is problematic for long-running protocols over the internet.

719 The *sleepy model* is meant to capture the beneficial security bounds of synchronous proto-
720 cols while making the system more partition-tolerant [28]. In other models, honest nodes
721 are assumed to be online throughout the entirety of the execution. Once a node goes offline,
722 it is considered faulty forever, even after coming back online. In contrast, the sleepy model
723 allows nodes to be "alert" (online) or "asleep" (offline), where asleep nodes can wake up
724 and become honest again. Synchronous protocols have difficulties because if an honest
725 player is offline for a sufficiently long time ($> \Delta$), they will reject honest messages when
726 they come back online. Asynchronous protocols require a threshold of honest validators to
727 respond to a proposal, but there may not be that many honest parties online at the time. The
728 sleepy model is similar to synchrony in that alert nodes have a network with a known delay
729 parameter and similar to asynchrony in that nodes are allowed to go offline indefinitely and
730 receive all pending messages upon coming back online. The sleepy model allows consen-
731 sus as long as a majority of the alert nodes are honest. That is, it has the same security
732 bound as synchronous systems when everyone is online and scales down as replicas go of-
733 fline. Even if only a tiny fraction of the system's nodes are alert, the protocol will continue
734 to make progress; that is, the sleepy model aims to maintain liveness despite an arbitrary
735 number of nodes going offline.

736 A related model, *weak synchrony*, has also been proposed as a way of dealing with nodes
737 temporarily going offline [29]. In contrast to the sleepy model, weakly synchronous proto-
738 cols favor consistency rather than liveness. When the network is partitioned, the minority
739 partition stops making progress but avoids the risk of deciding on values or blocks that are
740 inconsistent with the majority partition. The weak synchrony assumption is that the ma-
741 jority of nodes are both honest and online in each round, but the set of honest and online
742 nodes need not be the same in each round. This is a generalization of the synchronous
743 model, where the honest and online set must contain every honest node in every round.
744 Most synchronous consensus protocols can adopt slight adjustments in order to be secure
745 under weak synchrony [30].

746 Not every problem is solvable in any given model. For example, Byzantine broadcast only
747 works under synchrony, not partial synchrony or asynchrony, even if at most one replica is
748 faulty. The termination property cannot be satisfied because replicas do not know whether

749 the sender sent them a message at all, and thus do not know whether to keep waiting for
750 it or use a default value. The sender can be faulty and simply not send a message. To get
751 around this, one must weaken the termination property, as is done for reliable broadcast.

752 1.6. Permissioned vs. Permissionless

753 After the advent of Bitcoin, it became important to draw a distinction between the classical
754 consensus protocols designed since the 1980s (permissioned) and the new style of protocol
755 (permissionless).

756 In the *permissioned* model of consensus, there are n replicas, up to f of which may be un-
757 der the control of the adversary. Both n and the identity of the replicas are known to every
758 participant. Communication between participants typically takes place over authenticated
759 channels, in which case the existence of a public key infrastructure (PKI) is generally as-
760 sumed. At a minimum, every replica needs to agree with every other replica about the
761 set of public keys used in the system. The permissioned model corresponds to classical
762 consensus algorithms, which are discussed in Sections 4 through 7.

763 *Permissionless* systems differ from classical ones in four key ways, according to Pass and
764 Shi [31]:

- 765 1. There is no access control mechanism that determines which nodes can join the sys-
766 tem, and nodes can freely join or leave the system at any time.
- 767 2. Nodes are not aware of the other protocol participants a priori. In particular, commu-
768 nication is not over authenticated channels, so message senders are not authenticated.
- 769 3. The protocol itself may be unaware of how many nodes are participating in its exe-
770 cution.
- 771 4. The number of nodes involved in the system can grow or shrink over time.

772 Pass and Shi go on to prove that several of the limitations of permissionless consensus
773 are required in order to work in such a challenging environment [31]. First, a Sybil-
774 resistance mechanism is needed in order to maintain consensus when communication is
775 not authenticated (Sybil-resistance mechanisms are introduced in Section 2.2; the proof-
776 of-work mechanism is assumed in [31]). This is because an adversary needs to have their
777 messages rate-limited, regardless of whether nodes can join freely, everyone knows how
778 many nodes there are, and without message delays. Second, in order to allow nodes to
779 freely join the system after it has been set up, proofs of work must be performed continu-
780 ously throughout the lifetime of the system. If proofs of work ever cease, then new nodes
781 can be tricked into preferring a simulated execution. Not all Sybil-resistance mechanisms
782 can fully accomplish this free entry property (in particular, proof-of-stake systems require
783 workarounds for this, as described in Section 12.1.2). To maintain this free entry condi-
784 tion, there must be an honest majority in control of the Sybil-resistance resource. This is

785 a standard consistency argument, which is discussed more in Section 10.1. Finally, to tol-
786 erate uncertainty in the number of participants, there must be a known upper bound in the
787 network delay. That is, a synchrony assumption is needed. This is true because the network
788 can be partitioned in half, and if the adversary can cause an arbitrarily large message delay,
789 there is always some delay that will cause a consistency violation. Honest nodes are unable
790 to tell whether the other side of the partition exists at all or if there is just a message delay.
791 This is discussed more thoroughly in Sections 10.1 and 10.2.1.

792 The permissionless model also has setup assumptions. In particular, permissionless net-
793 works require a trusted setup to create the *genesis block* – a data structure that encodes the
794 initial state of the system. This is to prevent precomputation attacks, where the entity that
795 creates the system creates a hidden blockchain in advance which can be used to gain an
796 advantage in consensus. Famously, Satoshi Nakamoto included a newspaper headline in
797 the Bitcoin genesis block in order to prove that this attack had not taken place.

798 2. System Components

799 Distributed ledger systems are usually composed of a variety of different subprotocols or
800 components. Not every system will use each of the components described here. Further, a
801 single subprotocol may be responsible for multiple aspects of the system simultaneously.
802 Some of these subprotocols include:

- 803 • Data structures over which consensus is maintained
- 804 • Sybil-resistance mechanisms
- 805 • Leader election and/or committee selection
- 806 • Fork-choice or chain-selection rules
- 807 • Networking components
- 808 • Incentive mechanisms
- 809 • Cryptographic primitives
- 810 • The state machine itself

811 This section will introduce these ideas, but most will be explained in much more detail later
812 in the document. Note that many systems in the literature describe only one or two of the
813 components listed above but can often be adapted to alternative situations. For example,
814 some fork-choice rules will be analyzed for a single Sybil-resistance mechanism but could
815 be paired with a different mechanism in practice.

816 2.1. Data Structures for Distributed Ledgers

817 There are a variety of ways to organize the linearized/ordered transaction log in SMR. The
818 simplest possibility is an ever-expanding list of transactions. It is typically more efficient
819 to batch transactions into "blocks" (a group of transactions) rather than to handle them
820 individually, which increases throughput at the expense of latency. These blocks can be
821 "chained" together via collision-resistant cryptographic hash functions in order to form a
822 blockchain. In other words, a blockchain is a chain of blocks that each reference a hash
823 of the earlier blocks. In this way, each block is like a vote or commitment to the entire
824 chain before it. In computer science terms, this is a singly linked list. Generally, there is
825 a *block header* that includes metadata like the hash reference to the previous block and a
826 commitment to the content of the block itself.

827 Some systems may incorporate multiple blockchains that operate in parallel and may pro-
828 vide different functionalities. Some systems employ a sharded architecture, where multiple
829 blockchains exist in parallel but are coordinated via another blockchain. See Section 15
830 for more information on sharding. Other systems use directed acyclic graphs, or DAGs
831 [32]. Technically, a blockchain is a simple DAG. However, the term is typically used in
832 this space to describe systems where each block may point to multiple earlier blocks in-
833 stead of just one. In a number of systems, there is a primary pointer to a previous block
834 plus several so-called *uncle blocks*, which are produced around the same time but not in
835 the "main" blockchain. Sometimes, the DAG is formed over individual transactions rather
836 than blocks, which may result in only a partial ordering of transactions instead of a total
837 ordering. There may also be a subprotocol for extracting a total order of transactions from
838 the DAG.

839 The motivation for using a DAG or parallel blockchains is generally to improve latency and
840 throughput (under optimistic assumptions) compared to the original blockchain structure.
841 This is because proof-of-work blockchains require block distribution times to be a small
842 fraction of the maximum network delay (see Section 10.2.1). Proof-of-work protocols that
843 use DAGs are discussed in Section 11.5, and proof-of-stake DAGs are discussed in Section
844 13.3.

845 Note that throughout this document, the terms "blockchain" and "chain" are sometimes
846 used generically to refer to any ledger. An example of a blockchain and a DAG are shown
847 in Figure 1.

848 2.2. Sybil-Resistance Mechanism

849 A Sybil attack is a scenario in which a single real-world entity controls multiple in-protocol
850 participants while making it look as though they are controlled by several different entities.
851 Designers of permissioned networks need not be concerned with handling Sybil attacks
852 because all replicas are already aware of each others' identities. However, in permissionless
853 networks, a single malicious entity can represent itself as multiple distinct identities within

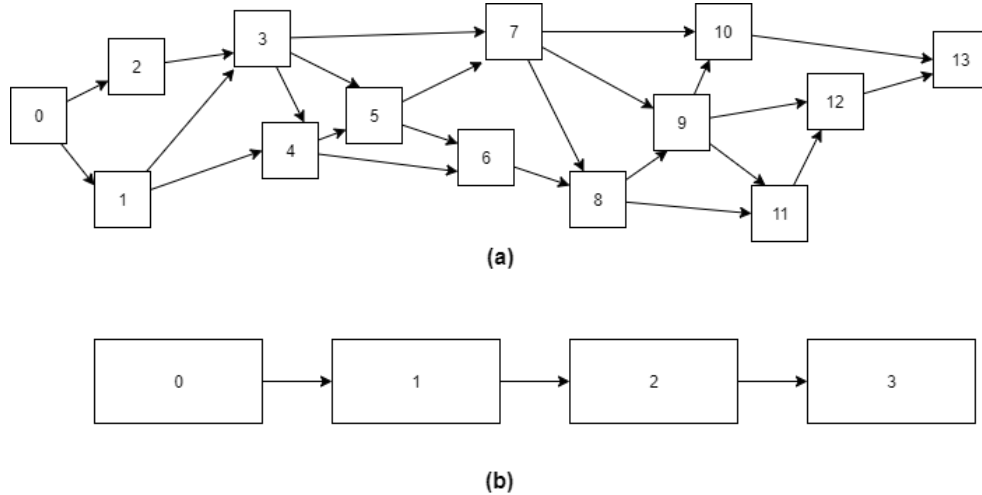


Fig. 1. Blockchain vs. DAG. Time moves from the left to the right. (a) A DAG where each block references two previous blocks except for the genesis block and the ones immediately after. In some systems, one of the pointers is the primary reference, while the other pointer is to an uncle block. (b) A simple blockchain.

854 a system. Therefore, permissionless systems require some method of addressing Sybil
855 attacks.

856 Sybil attacks were introduced in [33]. Without a trusted central identification authority,
857 Sybil attacks are always possible absent unrealistic or challenging assumptions, like re-
858 source parity among entities. However, permissionless distributed ledger systems can use
859 some type of scarce resource to mitigate Sybil attacks. The two most commonly proposed
860 Sybil-resistance mechanisms are proof of work and proof of stake. The scarce resource
861 used under proof of work is computational effort, while the resource in proof of stake is
862 virtual currency units that are native to the system. Proof of work and proof of stake have
863 very different properties, which are discussed in detail in Sections 9 and 12. These differ-
864 ences can have material impacts on the security of a system, so designers cannot simply
865 treat them as substitutes. There are other less common mechanisms as well, including proof
866 of space (see Section 14.1) and those that utilize trusted execution environments (TEEs)
867 [34, 35].

868 2.3. Leader Election and Committee Selection

869 Most consensus protocols require periodically electing a leader or a committee to perform
870 particular tasks, such as proposing blocks of transactions to the rest of the network or voting
871 on whether to accept a block that has been proposed. Other algorithms are leaderless,
872 though this is less common.

873 In permissioned networks, the leader election subprotocol can be as simple as rotating
874 through each validator in round-robin fashion. A typical method would be to label each of

875 the n validators with a number in $\{0, \dots, n - 1\}$, and for each round r , the leader is validator
876 $r \bmod n$. Alternatively, the leader could be chosen pseudorandomly based on a keyed hash
877 function that hashes to the domain $\{0, \dots, n - 1\}$. In some cases, the elected leader may
878 remain the leader for many rounds, which can improve the network's throughput but at the
879 cost of poor load balancing and fairness.

880 Leader election in permissionless networks is typically more complicated due to the lack
881 of stable, known identities. In proof-of-work systems like Bitcoin (and its consensus algo-
882 rithm, dubbed Nakamoto Consensus), the elected leader in a given round is the first node
883 to solve a moderately hard puzzle. Proof of stake often uses advanced cryptographic prim-
884 itives like verifiable random functions to aid in leader election, as described in Section
885 2.7. The process relies on uniformly random sampling from the participants in the network
886 based on their share of the resource used for Sybil resistance. That is, a proof-of-work
887 miner with 20% of the total computational power deployed on the network should have a
888 20% chance of being elected for each block. Leader election is typically tied to the mint-
889 ing mechanism for cryptocurrencies: the elected leader collects the newly created rewards
890 (however, they can be decoupled [36]). As a result, leader election tends to have important
891 implications for the incentive compatibility of a system.

892 If the leader is Byzantine, the outcome depends heavily on the specific protocol in use. In
893 permissioned systems, as well as a few permissionless ones, a *view change* subprotocol
894 is needed to securely move to the next leader (a *view* is a phase of the protocol where
895 a particular replica acts as the leader; see Section 7.6). This is critical for maintaining
896 liveness because a faulty leader can stall progress by not proposing a block when it is their
897 turn. Progress in SMR is guaranteed once all correct processes synchronize to the same
898 round and the leader of that round is correct. In most permissionless systems, a malicious
899 leader can also weaken liveness by not including transactions in blocks and can sometimes
900 use their status as leader to try to launch other attacks.

901 2.4. Fork-Choice or Chain Selection Rules

902 If a node is confronted with multiple valid ledgers, it has several methods for deciding
903 which one to adopt, such as following the blockchain that has the most total proof of work
904 (Nakamoto Consensus, see Section 10), selecting the heaviest subtree (GHOST, see Section
905 11.2), or accepting the result of a Byzantine agreement algorithm run by a duly elected
906 committee of validators.

907 Nakamoto Consensus uses a simple fork-choice rule called the *longest chain rule*. When
908 presented with two valid blockchains, a node in a system that uses Nakamoto Consensus
909 will prefer the chain that has the largest cumulative amount of proof of work supporting
910 it. For instance, if presented with the two chains from the bottom panel of Figure 2, a
911 node will adopt the chain that ends with Block 5, assuming that each block has the same
912 amount of work. Any alternative blocks that are discarded are called *stale blocks* or *orphan*
913 *blocks*. If the node had adopted the other chain beforehand, the switch is called a *blockchain*

914 *reorganization*, or *reorg* for short. In the top panel of Figure 2, the disconnected Block 2
915 and Block 4 are stale, and the shorter chain in the bottom panel that ends with Block 4
916 is stale. Note that transaction recipients can choose to wait for as many blocks as they
917 want (called *confirmations*) before accepting a payment and providing goods or services,
918 and waiting longer can offer more confidence that a block containing the payment will not
919 become stale.




Fig. 2. Forking and the Longest Chain Rule.

920 2.5. Networking

921 In permissioned networks, validators will typically have pairwise point-to-point channels
922 with all other validators, though not all systems are fully connected in this way. For ex-
923 ample, in some systems (e.g., PBFT), every honest replica sends messages to every other
924 replica for each round of the protocol. In other systems, replicas send messages to a leader,
925 who then aggregates them and disseminates them to the rest of the replicas. Some systems
926 even use a more complicated tree-like network topology for communication [37].

927 In permissionless settings, parties usually "diffuse" messages via gossip. Messages re-
928 ceived by a party are then forwarded on to their peers with no specific source or destination
929 for the message and without using authenticated channels. There is, thus, an implicit echo-
930 ing assumption that messages are forwarded when heard. For permissionless networks, a
931 more thorough treatment of possible networking components is provided in Section 17.1.

932 2.6. Incentive Mechanism

933 Historically, incentives were rarely considered in deployments of SMR. The replicas were
934 typically run by a single organization that merely wanted additional fault tolerance. More
935 recently, however, these systems have been deployed by a consortium of organizations or
936 over open networks. In these cases, it is important that validators are incentivized to behave
937 honestly rather than to try to disrupt the network.

938 In permissioned networks, honest behavior may be enforceable by appealing to legal sys-
939 tems. Permissionless systems do not have this luxury, so incentives must be considered an
940 explicit part of the design. These incentives frequently involve some *token* native to the
941 network (e.g., bitcoin in the Bitcoin network or ether in the Ethereum network). Typically,
942 appending a new block to the blockchain gives the leader a reward that includes newly
943 created cryptocurrency (the *block subsidy*) and transaction fees. In Bitcoin and many other
944 systems, this block reward comes in a special transaction called the *coinbase* transaction,
945 which is located first in a block and is the only transaction that can create new tokens rather
946 than just transfer existing tokens. Due to the inconsistencies that can arise due to forking,
947 the coinbase transaction has a *maturity window* in which funds cannot be spent from it (in
948 Bitcoin, the maturity window is 100 blocks, or slightly under 17 hours).

949 Incentives are discussed in more detail in Section 19.

950 2.7. Cryptographic Primitives

951 This document assumes that the reader is familiar with basic cryptographic primitives like
952 digital signatures and collision-resistant hash functions, which are used in all of the pro-
953 tocols described here. Beyond this, more advanced cryptography is applied judiciously
954 in systems that implement state machine replication. This can be in either the atomic
955 broadcast/consensus protocol or the state machine itself. Cryptographic techniques used in
956 blockchain systems are surveyed in [38].

957 Hash functions are used in blockchain systems in a number of areas, such as for proof
958 of work, generating *addresses* from public keys (concise representations of a public key
959 or payment destination), and as building blocks for more advanced primitives, such as
960 Merkle trees. Digital signatures are used to authorize transactions, implement authenticated
961 channels, and vote in a variety of permissioned and proof-of-stake systems.

962 *Threshold cryptography* is frequently employed to improve fault tolerance or performance.
963 Threshold signatures are the most common and are used both as a form of multi-factor au-
964 thentication for signing transactions and to reduce the communication complexity of con-
965 sensus by aggregating the individual signatures of several validators. Some protocols use
966 threshold decryption to hide the contents of transactions from validators and prevent cen-
967 sorship or favoritism before transactions are ordered, only allowing them to be decrypted
968 after being committed to the transaction log.

969 Cryptographic *accumulators* appear in nearly every blockchain system. Accumulators are
970 compact ways of representing sets, which – at a minimum – allow concise proofs of mem-
971 bership or inclusion in the set. Block headers typically include an accumulator that commits
972 to the set of transactions included in the block itself. The most common example of this is
973 a *Merkle tree*, where each leaf of the tree is a transaction included in the block, and the root
974 of the tree is included in the block header. Among other things, this allows *light clients*
975 (see Section 3.2), which do not fully verify the ledger, to verify that transactions they care
976 about have been added to the blockchain. In some systems, accumulators are also used
977 to represent the global state of the system’s state machine. That is, the block header may
978 include an accumulator that includes all accounts and their balances.

979 Cryptographic *commitment* protocols are another common occurrence, though they are
980 more commonly used in the state machine rather than in consensus. Commitments are
981 a digital version of a sealed envelope – the party that commits to a value cannot alter the
982 contents after the fact (it is already in the envelope), and no other party can discern the
983 value until the committing party reveals it (no one can read the contents of the envelope
984 until it has been opened). Cryptographic commitments are sometimes used to provide con-
985 fidentiality for the amount sent in a transaction. This is often done by using commitments
986 as a component of a *range proof*, which shows that a value is in a particular range and
987 which can prove that a transaction did not create assets "out of thin air" to spend.

988 *Zero-knowledge proofs* – and particularly, *zkSNARKs* – are used in multiple ways in blockchain
989 systems [39]. A zero-knowledge proof is a protocol that enables a prover to convince a ver-
990 ifier that a particular statement is true without the verifier learning anything other than the
991 truth of the statement. *zkSNARKs* allow a prover to convince a verifier that the results
992 of a computation on secret inputs are correct without the verifier needing to execute the
993 computation or learn anything about it. This can be used to add privacy for payments but is
994 also used in several advanced state machine designs, as discussed in Section 18.1.2. While
995 *zkSNARKs* are powerful, they tend to require a trusted setup to generate parameters for
996 the system, and an entity that knows the random values used for the setup has the power to
997 create fraudulent proofs.

998 For the purposes of this document, the most important cryptographic primitives are those
999 that aid in distributed randomness generation, which is most importantly used for leader
1000 election where an agreed-upon source of randomness must be shared among replicas. These
1001 schemes include common coin protocols, verifiable random functions (VRFs), and verifi-
1002 able delay functions (VDFs).

1003 A *common coin* is a randomness source that is observable by all participating processes
1004 but unpredictable for an adversary. The common coin abstraction can be realized using
1005 threshold signature schemes or *verifiable secret sharing*. Secret sharing allows a party to
1006 distribute a secret value to other parties such that a threshold of them are required in order to
1007 reconstruct the secret, while verifiable secret sharing uses commitments to ensure that the
1008 party who distributed the values has done so correctly. Sometimes, a *weak common coin* is

1009 sufficient, where weak means that there is a constant probability that the functionality will
1010 return different values to different processes.

1011 *Verifiable random functions* (VRF), introduced in [40], are pseudorandom functions that
1012 provide publicly verifiable proofs of correctness. They can be thought of as a public key
1013 version of a keyed hash function. For a fixed key pair and input value, a VRF will produce
1014 a unique pseudorandom and verifiable output, even if the key pair was chosen adversarially.
1015 The VRF proof is usually the signature over some input data, and the pseudorandom output
1016 is a hash of the signature. A verifier will ensure that the signature is valid and a preimage
1017 of the output. For this to work, the signature must be *unique*, which means that only a
1018 single valid signature exists for a given message and key. Due to this requirement and the
1019 ease of constructing threshold implementations, BLS signatures are most frequently used
1020 for this purpose [41]. VRFs are used most frequently for leader election in proof-of-stake
1021 systems, such as Ouroboros Praos, Algorand, and Fantôme (see Sections 13.1.3, 13.4.2,
1022 and 13.3.1, respectively).

1023 Finally, *verifiable delay functions* (VDFs) are relatively new primitives, which are functions
1024 that require a significant amount of sequential computation but where the correctness of the
1025 result is easy to verify. VDFs can be thought of as a time delay that is imposed upon the
1026 generation of output for a pseudorandom number generator. The delay prevents malicious
1027 actors from influencing the output of the generator because all of the inputs to the generator
1028 are finalized before the delay ends. VDFs are used for randomness in leader election, may
1029 be helpful against certain attacks on proof-of-stake and proof-of-space systems, and can be
1030 used to limit the frequency with which an adversary can send messages or vote in consensus
1031 [42]. VDFs improve upon VRFs in that a single honest participant is required instead of a
1032 non-colluding honest majority (requiring a single honest participant is sometimes called the
1033 *anytrust* model). The idea of VDFs was first formalized in [43], but improved versions were
1034 found soon after [44, 45]. Due to being relatively new, there are still some uncertainties
1035 regarding their security, including the possibility of specialized hardware and tuning the
1036 time parameter properly.

1037 The VDF from [44] works by choosing a time parameter T , a finite abelian group G of
1038 unknown order, and a hash function H with a domain that consists of the elements of G .
1039 For an input x , let $g = H(x)$. The VDF is evaluated by computing $y = g^{2^T}$. Repeated
1040 squaring does not reveal any information about the output until the final squaring, and the
1041 computation must be done in serial (because the order of G is unknown). Unfortunately,
1042 generating a group of unknown order requires a trusted setup.

1043 2.8. State Machine

1044 The state machine is the set of rules that replicas enforce while transitioning the state of
1045 the system via client-submitted transactions. Though introduced here, Section 18 provides
1046 a more detailed discussion of the design space for state machines.

1047 A wide variety of rules are possible, but they typically include ensuring that transactions
1048 have valid signatures of the clients who are authorized to act on the portion of the state
1049 involved, as well as preventing "double-spending." In the context of a cryptocurrency, this
1050 means that the owner of the coins signed the transaction and that there is not a conflicting
1051 transaction that spends the same coins included in the ledger. Other rules may include a
1052 maximum allowed block size, restrictions on the timestamps in block headers, that blocks
1053 and transactions are syntactically well-formed, and that all transactions in a block are com-
1054 mitted to in the block header.

1055 In addition to these basic rules, the state machine may include a means of executing various
1056 types of computations, which are encoded as *smart contracts*. The state machine often
1057 provides a programming language and execution environment where developers can write
1058 programs (smart contracts) that are executed by all of the network's replicas, who then
1059 agree on the result and modify the system's state accordingly. In some systems, such as
1060 Ethereum, the execution environment is essentially *Turing-complete*, so it can execute any
1061 arbitrary deterministic program subject to the system's *block gas limit*. Each operation
1062 in the state machine has an associated *gas cost*, where *gas* is a unit of measurement that
1063 corresponds to the amount of effort or resources consumed to execute the operation. The
1064 gas limit is analogous to a maximum block size, but instead of limiting the total size of
1065 transactions, it bounds the total computational effort.

1066 2.8.1. UTXO vs. Account Model

1067 There are two common models of how the state machine's state is represented: the *UTXO*
1068 *model* and the *account model*. A UTXO, short for *unspent transaction output*, is an im-
1069 mutable object associated with a particular spending condition. By fulfilling this spending
1070 condition, a client becomes authorized to destroy the UTXO and create new ones from it
1071 (so long as other state machine rules are followed, like not creating more new coins than
1072 were spent). In the UTXO model, each transaction takes some UTXOs as inputs, destroys
1073 them completely, and creates new UTXOs as output with different spending conditions that
1074 reflect the new owner. This often includes "change" outputs that return funds to the sender.

1075 In contrast, accounts are mutable objects associated with a balance that can increase or
1076 decrease as funds are moved in and out of the account. The UTXO model creates a directed
1077 graph of transaction outputs that move between owners (shown in Figure 3), whereas the
1078 account model is like a database of the current system state. One can think of UTXOs
1079 as individual dollar bills with arbitrary denominations, and accounts are more like regular
1080 bank accounts with an identifier and a balance. In either case, the UTXO or account may
1081 have programmable spending conditions.

1082 Each model has particular advantages and disadvantages. In the UTXO model, transactions
1083 may be more private due to the ease of creating new addresses for each output, whereas all
1084 of the transactions to and from a particular account in the account model are automatically
1085 linked. Due to the immutability of UTXOs, disk access to check system state is more par-

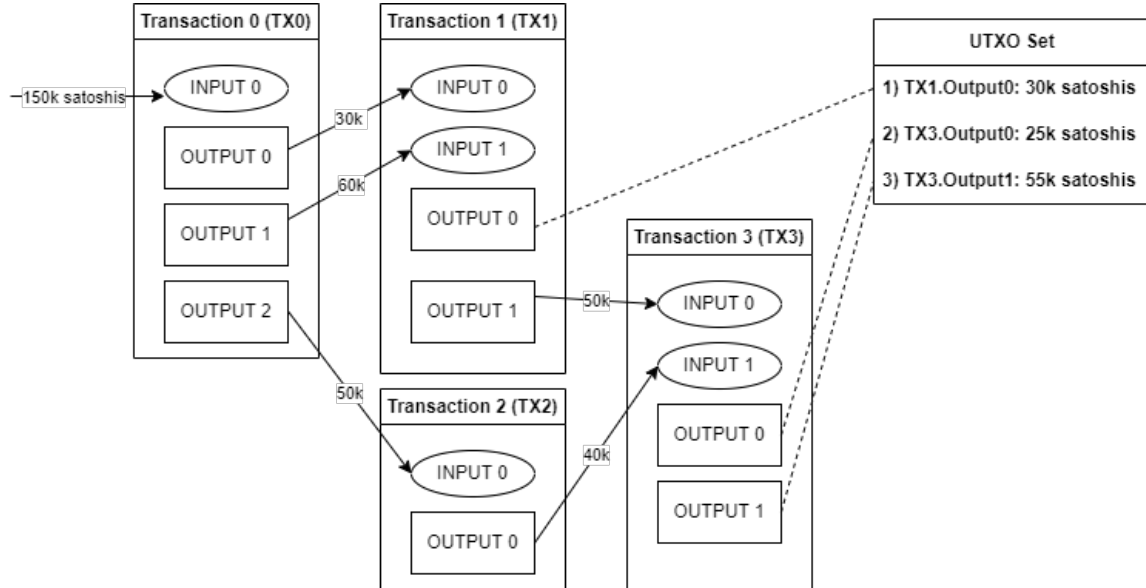


Fig. 3. UTXO transaction graph. Outputs of some transactions become inputs to later ones. Each transaction depicted includes a 10k satoshi fee, where satoshis are the smallest atomic unit of currency in the Bitcoin network.

1086 allelizable than in the account model, where the accounts can be modified. On the other
1087 hand, the account model has the performance advantage of typically smaller transactions.
1088 In the account model, every transaction has a single "input" and "output." With UTXOs,
1089 some transactions may become extremely large if there are many recipients or if they re-
1090 quire simultaneously spending many UTXOs that have a small value. The account model
1091 is also simpler for light clients; in the UTXO model, a light client must keep track of each
1092 UTXO owned by the client and update this for every transaction. One of the biggest disad-
1093 vantages faced by the account model is that transactions require nonces in order to prevent
1094 replay attacks. The nonces must be sequential, which can impact transaction processing
1095 if an account issues many transactions in short periods of time (they must be processed
1096 sequentially, regardless of network lag and the attached fees).

1097 Perhaps the most straightforward advantage of the account model is that accounts are able
1098 to maintain a persistent state, which can make complex smart contract programming much
1099 simpler than in the UTXO model, where UTXOs do not maintain state. The statelessness
1100 of UTXOs makes it challenging to program smart contracts with multiple phases. In the
1101 account model, a smart contract exists at a static address, which can store state and be
1102 referenced without updating.

1103 2.8.2. Changing the Rules

1104 All validators must agree on the state machine rules. However, the rules themselves may
1105 change over time. Broadly speaking, there are two types of rule changes, or *forks*: *hard*

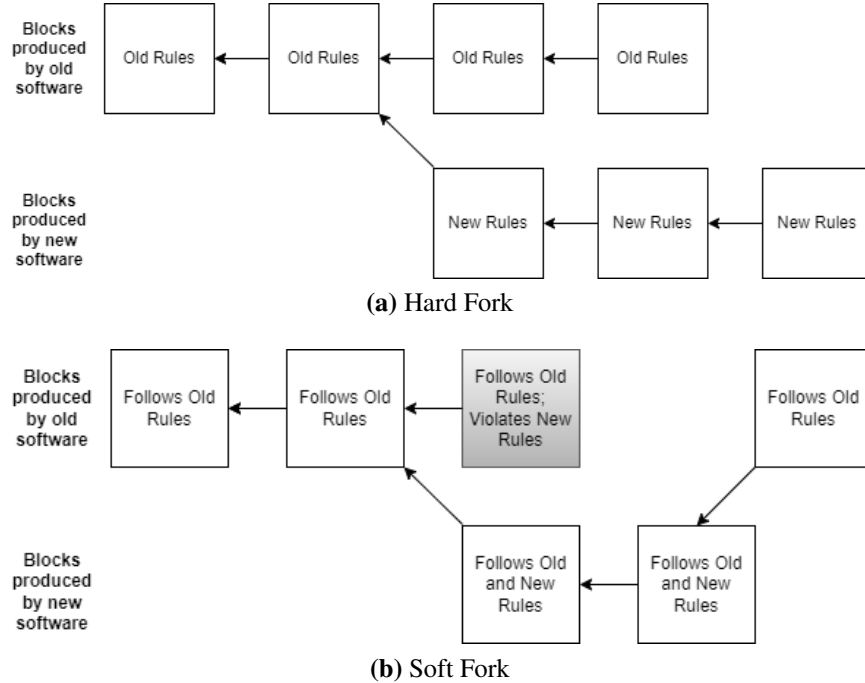


Fig. 4. Hard forks and soft forks.

1106 *forks*, which make blocks or transactions valid when they were previously considered in-
1107 valid, and *soft forks*, which make blocks or transactions invalid when they were previously
1108 considered valid. Hard forks are not backward compatible, whereas soft forks are. The
1109 different fork types also have different security ramifications, different practical consid-
1110 erations regarding the mechanics of implementing the fork, and – according to some –
1111 different ethical considerations. Hard forks are depicted in Figure 4a, and soft forks are
1112 shown in Figure 4b.

1113 Hard forks can be thought of as the creation of an entirely new system, where a successful
1114 hard fork results in a smooth transition from the old system to the new system. Technically,
1115 a hard fork always results in two separate systems, but a smooth hard fork with (near)
1116 universal consent from the community may result in the original chain being abandoned.
1117 Nodes that fail to update their software may be more easily attacked around the time of the
1118 hard fork, so hard forks require near-unanimous consent and for users to all update their
1119 software within a specific period of time. For example, a node that has not upgraded can be
1120 presented with a chain of blocks under the old rules by an attacker. Since those blocks will
1121 not be reorged, the attacker can trick the recipient into accepting coins that the remainder
1122 of the network will not accept. In Figure 4a, a node that follows the old rules and accepts
1123 transactions with only two confirmations can be easily defrauded.

1124 In contrast, soft forks tend to be more secure than hard forks because they do not require
1125 every user to update their software near-simultaneously, although those actively participat-

1126 ing in consensus, such as miners and staking validators, should update prior to the soft fork
1127 activating. For a soft fork to happen smoothly, the majority of the mining power (or stake)
1128 should recognize and enforce the soft fork; a supermajority would be ideal. If some miners
1129 are running software that enforces the old rules, users could receive "fake" confirmations
1130 in blocks that are valid under the old rules but invalid under the new rules and, thus, soon
1131 to be reorged. This is the scenario created by the shaded block in Figure 4b; if a user's
1132 node follows the old rules and accepts transactions with a single confirmation, they can be
1133 defrauded. As a result, non-upgraded nodes lose some security during a soft fork.

1134 Both hard forks and soft forks can potentially result in permanent chain splits, but it is
1135 far less likely to occur with soft forks and can be avoided completely through the choice
1136 of activation mechanism. Permanent chain splits create considerable complications for
1137 users. First, the values that accrue from network effects are reduced by splitting the relevant
1138 communities in two. Second, unless the fork changes the rules for how transactions are
1139 constructed, some transactions that are valid on one chain will also be valid on the other
1140 in case of a split, which creates the opportunity for *replay attacks*. An attacker can take a
1141 transaction from chain A, broadcast it again to the network maintaining chain B, and have
1142 it confirmed. If a chain split occurs and a cryptocurrency exchange only supports chain
1143 A, then an attacker who initiates a withdrawal from the exchange on chain A can copy
1144 the transaction, broadcast it over the network managing chain B, and receive an equivalent
1145 amount of currency B that the exchange never intended to give up.

1146 Chain splits cause problems for blockchain service providers, such as wallets and ex-
1147 changes. Exchanges must decide whether to support one or both sides of the split, and
1148 users of the exchange may not have the latitude to choose which chain to follow. As a
1149 result, a user who wanted to stay on the old chain but used an exchange that only follows
1150 the new chain could potentially lose their funds. Chain splits also pose a problem for light
1151 clients. Because light clients do not validate the state transitions of the network, they are
1152 unaware of what rules are being followed. That means that they will simply follow the ma-
1153 jority, which can be dangerous. For example, if a majority of the Bitcoin hash rate decided
1154 to increase inflation, full nodes would recognize the block as invalid, but light clients would
1155 simply follow the majority of the hash rate. In addition to the explicit danger of attacks dur-
1156 ing this period, if a significant portion of the community were to engage in commerce on
1157 the attacking/inflation chain, then a social consensus may form around the chain despite its
1158 invalidity. Contentious hard forks that are likely to result in a split are especially dangerous
1159 to light clients, which make up the bulk of typical users. More information on light client
1160 security compared to fully verifying nodes can be found in Section 3.2.

1161 There are several methods for performing forks that primarily depend on whether they are
1162 user-activated or miner-activated. In a *user-activated soft fork* (UASF), the software im-
1163 plementation begins enforcing the new, more restrictive rules at an agreed-upon time, or
1164 "flag day." In a *miner-activated soft fork* (MASF), the new rules are enforced only after an
1165 extended period of miner signaling. Miners can set a particular bit in their block headers
1166 if they support the rule change, and after a threshold percentage of blocks (usually 75%

1167 to 95%) within a given time period signal support, the new rules can be enforced (usually
1168 after a short delay to give time for more nodes to update their software once the change is
1169 "locked in"). Only a UASF can cause a chain split, but MASFs can be chaotic and result
1170 in many short forks if the threshold is set too low. One issue with miner activation is that
1171 miner signaling is in no way a guarantee that those miners will actually enforce the new
1172 rules. For example, the Bitcoin community debated on whether to increase the maximum
1173 block size, and supporters of an increase released a (hard fork) client called Bitcoin XT
1174 that enforced a block size increase, while the dominant client – Bitcoin Core – did not. An
1175 anonymous developer then released the NoXT client, which mimicked Bitcoin XT's be-
1176 havior and signaling but actually followed rules compatible with the Bitcoin Core software
1177 [46]. If some miners were running this client, the fork would be activated prematurely to
1178 the detriment of those who wanted to enforce the new rules.

1179 User activation has its own problems. The idea of a UASF is that if the majority of eco-
1180 nomically significant users of the system adopt it, then the miners will be forced by their
1181 own self-interest to adopt it as well. As a result, the UASF may be able to get around the
1182 resistance of miners if they are acting in opposition to the broader community of users.
1183 Again, the primary issue is that it is likely to result in a chain split without the support of a
1184 majority of the miners. In fact, without majority mining support, a UASF looks a lot like a
1185 hard fork with all of the risks that hard forks entail, all of the limitations of soft forks, and
1186 the possible need for a hard fork to protect against the miners that do not accept the fork
1187 rules [47].

1188 For any fork to be successful, nodes need to coordinate with each other regarding both the
1189 time to lock in as well as – after a delay – the time to activate the rule change. In Bitcoin and
1190 similar systems, there are several notions of time that can be used for this coordination: the
1191 block timestamp, the *median time past* (MTP), and the block height. The block timestamp
1192 is unsuitable for use because it does not increase monotonically, which is a requirement
1193 for coordination. The MTP is defined as the median timestamp of the block and its 10
1194 predecessors. Bitcoin has rules about the timestamps that force MTP to monotonically
1195 increase though it is not monotonic in the face of reorgs. In all but the most exceptional of
1196 circumstances, block height will be monotonic even across reorgs. The advantage of using
1197 MTP to coordinate is that the timing can be selected in order to minimize the chance that
1198 the fork occurs at particular times. This may be useful in order to time the fork to happen
1199 while, say, engineers at important companies may be at work instead of asleep. On the
1200 other hand, MTP can be manipulated by miners in a number of ways through their control
1201 of the block timestamps, and a majority of miners can use this mechanism to completely
1202 skip over relevant activation periods [48]. Block heights are safe from this risk but are less
1203 likely to match a desired time for the humans running the nodes or the developers writing
1204 the code that schedules activation.

1205 Many in the blockchain community believe that there are ethical considerations with re-
1206 gard to the style of fork employed. This goes beyond the nature of the specific rule changes
1207 themselves (e.g., a rule that allows a privileged party to steal from others) but rather to the

1208 meta-problem of the methodology of forking. Intuitively, in a distributed network com-
1209 posed of many heterogeneous individuals with potentially competing interests, any poten-
1210 tial change in the rules (other than, say, a fix for a catastrophic software bug) is likely to
1211 be opposed by at least one person. As a result, if the change occurs, there may be a per-
1212 ception that this person or people were coerced into accepting the change. A representative
1213 of this view is Ethereum creator Vitalik Buterin, who believes that both types of forks are
1214 coercive, but that soft forks are more coercive than hard forks:

1215 There is an essential difference between hard forks and soft forks: hard forks
1216 are opt-in, whereas soft forks allow users no "opting" at all. In order for a user
1217 to join a hard forked chain, they must personally install the software package
1218 that implements the fork rules, and the set of users that disagrees with a rule
1219 change even more strongly than they value network effects can theoretically
1220 simply stay on the old chain...In the case of soft forks, however, if the fork
1221 succeeds the unforked chain does not exist. Hence, soft forks clearly institu-
1222 tionally favor coercion over secession, whereas hard forks have the opposite
1223 bias...If I had to guess why, despite these arguments, soft forks are often billed
1224 as "less coercive" than hard forks, I would say that it is because it feels like
1225 a hard fork "forces" the user into installing a software update, whereas with
1226 a soft fork users do not "have" to do anything at all. However, this intuition
1227 is misguided: what matters is not whether or not individual users have to per-
1228 form the simple bureaucratic step of clicking a "download" button, but rather
1229 whether or not the user is coerced into accepting a change in protocol rules
1230 that they would rather not accept. And by this metric, as mentioned above,
1231 both kinds of forks are ultimately coercive, and it is hard forks that come out
1232 as being somewhat better at preserving user freedom. [49]

1233 As alluded to by Vitalik, others take the position that hard forks are more coercive or
1234 at least more problematic from the standpoint of the developers who must write code that
1235 implements the hard fork. While most users agree that hard forks may be necessary in some
1236 cases, some believe that *controversial* hard forks should be avoided. For example, Bitcoin
1237 Core developer Pieter Wuille, in the context of a debate about changing the maximum block
1238 size in Bitcoin, expressed a distaste for hard forks that lack near-universal agreement in the
1239 Bitcoin community:

1240 [The responsibilities of the Bitcoin Core developers] include[] participating in
1241 discussions about consensus changes, but not the responsibility to decide on
1242 them – only to implement them when agreed upon. It would be irresponsible
1243 and dangerous to the network and thus the users of the software to risk forks,
1244 or to take a leading role in pushing dramatic changes. Bitcoin Core developers
1245 obviously have the ability to make any changes to the codebase or its releases,
1246 but it is still up to the community to choose to run that code...Bitcoin Core is
1247 not running the Bitcoin economy, and its developers have no authority to set
1248 its rules...Worse, intervening in consensus changes would make the ecosystem

1249 more dependent on the group taking that decision, not less. So to point out
1250 what I consider obvious: if Bitcoin requires central control over its rules by a
1251 group of developers, it is completely uninteresting to me. Consensus changes
1252 should be done using consensus, and the default in case of controversy is no
1253 change. [50]

1254 Essentially, this view suggests that it is not the role of developers to decide what Bitcoin is
1255 or should be and that developers have a responsibility to not "force" controversial changes
1256 to the system. To do so would elevate the developers to a more powerful position than they
1257 ought to have and would go against the decentralized ethos of the community.

1258 Over the lifetime of a long-lived system, it is likely that a hard fork will be necessary, and
1259 soft forks are simply not preventable when supported by a majority of miners so perhaps
1260 the ethics are less relevant than practical concerns. In either case, no users are forced into
1261 running particular code on their machine. Ultimately, users decide the rules because users
1262 decide which software to run. The "coercion" in either case is in the eyes of the beholder –
1263 no one is forced to use any particular rules, but they may be "forced" to use particular rules
1264 if they want to remain compatible with their counterparties.

1265 3. Scaling and "Decentralization"

1266 3.1. A Note on Decentralization

1267 The term *decentralization* is used frequently when discussing modern replicated state ma-
1268 chines. However, there is no single, accepted definition of what decentralization means. As
1269 a result, its use often confounds more than clarifies. It is worth investigating a few of the
1270 many proposed definitions of decentralization in order to grasp the diversity of views.

1271 Balaji Srinivasan and Leland Lee propose a metric they call the (*minimum*) *Nakamoto Co-*
1272 *efficient* to measure the decentralization of a system [51]. The idea behind this metric is to
1273 first create a list of the essential subsystems of the decentralized system under analysis, de-
1274 termine the number of entities that an adversary would need to compromise in order to take
1275 effective control over each subsystem, and take the minimum as a measure of the system's
1276 decentralization. The system is more decentralized when more entities must be corrupted
1277 to control an essential subsystem.

1278 There are many plausible subsystems, though it is unclear which ought to be considered
1279 "essential." Subsystems may include the distribution of mining rewards, the number of ex-
1280 changes, the volume traded on exchanges, the number of software clients, the number of
1281 developers per client, the number of full nodes being run (and their distribution by legal
1282 jurisdiction), the distribution of asset ownership, the fraction of users that hold their own
1283 private keys instead of delegating to a custodian, the fraction of users who validate the
1284 ledger with their own node instead of trusting another entity, the number of businesses run-
1285 ning economically significant nodes (i.e., who validate many incoming transactions), the

1286 number and distribution of hardware manufacturers in proof-of-work systems, the number
1287 of mining pools, and so on.

1288 In addition to the challenge of determining which of these subsystems are essential, it can
1289 be difficult to determine what the proper threshold of corrupted entities would need to be for
1290 some subsystems. For instance, it is unclear what percentage of exchange volume would
1291 suffice to consider that subsystem captured or centralized. In addition, reasonable people
1292 may disagree on the benefits of having multiple clients instead of a single client with more
1293 developer attention. In other cases, the subsystems may take on a different significance
1294 depending on the network in question. For example, the concentration of coin ownership is
1295 more important in a proof-of-stake network than in a proof-of-work network. If the wealth
1296 within a system is highly concentrated, then a government or other powerful adversary may
1297 only need to target a few large holders in order to acquire a large enough fraction of the
1298 asset and cause a market crash.

1299 Paul Sztork proposed an alternative measure of decentralization: the *cost of node-option*
1300 (CONOP) [52]. He points out that "[t]he process of 'money' is 'knowing you've been
1301 paid'" and that a process is more decentralized when it happens more locally. As a result,
1302 he says that "'decentralized money' is the local cost of knowing you've been paid: the
1303 cost of running a full node." This can be generalized to arbitrary computation possible in
1304 state machine replication: decentralized computation can be measured by the cost required
1305 to validate the correctness of computations. Actually running a fully validating node is
1306 not strictly necessary for everyone. Instead, the relevant consideration is the cost required
1307 to start and operate a node should they desire to do so. Ultimately, CONOP is intuitive
1308 because if it is extremely easy to set up and run a full node, the network remains difficult to
1309 shut down or prevent access to, even if there are not many nodes running at any given time.
1310 On the other hand, if only a handful of entities have the resources to run a fully validating
1311 node, the system can easily be shut down or succumb to undue influence. In this case, most
1312 individuals would be required to trust a third party and would not be able to locally verify
1313 for themselves whether the payment or computation happened correctly.

1314 A final decentralization metric proposed in [53] is (m, ϵ, δ) -decentralization. This is de-
1315 fined as a state of the system such that at least m participants run nodes that participate in
1316 consensus, and "the ratio between the total resource power of nodes run by the richest and
1317 δ -th percentile participants is less than or equal to $1 + \epsilon$ " [53]. Since this metric is only
1318 concerned with the centralization of the Sybil-resistance aspect of the protocol, it is more
1319 constrained in scope than the Nakamoto Coefficient or CONOP. The intuition for this met-
1320 ric is that it is preferable to have a large number of participants involved in consensus, and
1321 the participants ought to have a roughly even distribution of power. The paper derives four
1322 necessary conditions for a system to have "full decentralization" according to the authors'
1323 conception of the term:

- 1324 1. Any nodes with resource power must earn rewards (and there must be at least m of
1325 them).

- 1326 2. Participants should find it more profitable to run their own nodes than to delegate
1327 their resources to another participant.
- 1328 3. It should not be more profitable for an entity to run several nodes than to run a single
1329 one.
- 1330 4. "[T]he ratio between the resource power of the richest and δ -th percentile nodes
1331 converges in probability to a value of less than $1 + \epsilon$ " [53]. That is, the advantage of
1332 the most well-resourced node is bounded.

1333 Crucially, the third condition is not possible without a trusted third party, leading to the pa-
1334 per's conclusion that "full decentralization" is impossible in permissionless systems. With-
1335 out some form of trusted identification authority, it is unclear how to design a system where
1336 it costs more for one entity to run multiple nodes than it costs for multiple individuals to
1337 each run their own node.

1338 Each of these conceptions of decentralization highlights a different but important aspect of
1339 system security and architecture. The Nakamoto Coefficient highlights the fault tolerance
1340 of the system in an intuitive way: a system is more decentralized when it can tolerate a
1341 greater number of faulty or malicious entities. Practical challenges to using this as a metric
1342 include a lack of clarity regarding which subsystems should be included and the resultant
1343 difficulty of comparing different systems when their component subsystems differ in im-
1344 portance. In contrast, CONOP has nothing to do with fault tolerance. Instead, CONOP
1345 highlights the barriers to entry for users to be able to take advantage of the security proper-
1346 ties of the distributed system. In this conception, a more decentralized system is one where
1347 a greater number of users have the ability to fully validate the replicated state machine –
1348 that is, to be an *equal* peer in a peer-to-peer network rather than dependent on an external
1349 trusted party. Finally, (m, ϵ, δ) -decentralization highlights the impact that wealth or re-
1350 source inequalities outside of the system have on the relative power of participants within
1351 the system. In an open, permissionless system, these inequalities are likely to be reflected
1352 in the distribution of the Sybil-resistance resource. In this case, a system is more decentral-
1353 ized when the resource disparities between more powerful nodes and less powerful ones
1354 are smaller.

1355 In this document (NIST IR 8460), the term "decentralization" is used sometimes but spar-
1356 ingly. It does not refer to any specific meaning above but should be clear based on the
1357 context in which it is written.

1358 3.2. Full Nodes and Light Clients

1359 There are multiple distinct operating modes that a user might choose when interacting with
1360 a replicated state machine. These modes can be broken into two categories: *full nodes*,
1361 which process and validate the complete historical ledger of transactions, and *light clients*,
1362 which may validate a portion of the ledger or a subset of the rules but require placing some
1363 amount of trust in other entities.




Fig. 5. Simplified Payment Verification (SPV). A light client with a chain of block headers can verify the inclusion of Transaction 3 in Block 1 when provided with *Hash2*, *Hash01*, and *Hash4567*.

1364 There are different types of light clients, but the original one proposed in the Bitcoin white
1365 paper uses *Simplified Payment Verification* (SPV) [1]. SPV clients connect to full nodes
1366 but do not download the entire blockchain. Instead, they only maintain the chain of block
1367 headers for the best chain they are aware of (in Bitcoin and similar systems, this is the chain
1368 with the most cumulative proof of work). Because the block header contains a Merkle root
1369 that commits to every transaction contained in the corresponding block, it is possible to
1370 create proofs of transaction inclusion that scale logarithmically with the number of trans-
1371 actions in the block. This is shown in Figure 5. While this process does not demonstrate
1372 that the transaction is valid, it proves that the transaction is included in the canonical chain
1373 if connected to at least one honest full node. This implies that the transaction can be trusted
1374 to be executed and not reverted if the adversary is beneath the security threshold of the
1375 system and the block has enough confirmations. Alternatively, it shows how much work
1376 would be required to revert the transaction and perform a double-spend.

1377 The full nodes that serve the SPV client need to know which transactions the client cares
1378 about in order to provide a Merkle inclusion proof for it. To this end, the SPV client
1379 constructs a *Bloom filter* – a compact, probabilistic data structure that allows one to test
1380 whether an item is a member of a set – that includes the receiving addresses that the user
1381 wants monitored. The Bloom filter includes false positives, such that some of the addresses
1382 are not ones that the client cares about. In theory, this should substantially improve pri-
1383 vacy for SPV users compared to simply providing a list of addresses. In practice, however,
1384 Bloom filters provide almost no additional privacy for light clients, and users should un-
1385 derstand that they are – in essence – telling a service provider about all of their transactions
1386 (and possibly associating them with an IP address) [54]. Perhaps for this reason, many
1387 light clients actually do simply provide a list of addresses to a wallet service provider and
1388 completely sacrifice transaction privacy.

1389 An alternative model of light client improves privacy at the cost of significantly increased
1390 bandwidth (though still substantially less than a full node). *Compact block filters* are com-
1391 pressed representations of the contents of blocks using Golomb-Rice coding and allow a
1392 light client to check whether the block contains transactions that are relevant to the client
1393 [55, 56]. If so, the light client downloads the full block in order to extract the relevant trans-
1394 action. In addition to downloading a subset of complete blocks, the filters impose greater
1395 resource requirements than SPV but scale better for the full nodes that serve data. For each
1396 block, a full node only needs to create the filter once and can send the same filter to any
1397 light client who asks. In contrast, each SPV user provides its own Bloom filter, and the full
1398 node needs to perform considerable disk I/O to check for transactions for each of them.

1399 Regardless of the method used, light clients have degraded security properties compared to
1400 full nodes, which is the price paid for convenience and low resource usage. First, as men-
1401 tioned previously, light clients reduce user privacy compared to full nodes. Even compact
1402 block filters reveal the subset of blocks in which the user is the recipient of a transaction.
1403 Second, nothing prevents the full node server from lying by omission and failing to tell the
1404 light client that they have received a transaction. While this can be mitigated by connecting
1405 to multiple full nodes, many wallets connect to only a single server, and there is a risk that
1406 the full nodes are coordinating a Sybil attack and are not independent. Light clients need
1407 to trust that there are available sockets to connect to these honest full nodes, which full
1408 nodes lack an incentive to provide and which has not always been the case in practice [57].
1409 Compact block filters also mitigate the problem of lying by omission by allowing the light
1410 client to check the validity of the filter itself to see if the full node is lying. Third, light
1411 clients are more susceptible to certain attacks, like the vector76 attack described in Section
1412 10.3.

1413 More importantly, light clients have no way of knowing whether the chain that they are
1414 following is valid in the first place. Light clients do not maintain the full system state that
1415 would be required to validate transactions. As a result, light clients can be deceived by
1416 adversaries in ways that full nodes cannot. The only thing that light clients have assurance
1417 on – assuming that they are connected to at least one honest full node – is that they are aware
1418 of the chain that satisfies the network’s chain selection rule, ignoring the requirement that
1419 the chain be valid. As a result, light clients can be deceived into trusting an invalid state
1420 if the security threshold of the system is violated temporarily, if the adversary does not
1421 exceed the threshold but gets lucky, or if the adversary can partition the victim from the
1422 rest of the network. As a result, a miner can double spend a limitless amount of funds
1423 against a light client, whereas they can only double spend what they already own against
1424 a full node. Similarly, light clients do not recognize hard forks if the fork does not change
1425 the semantics of the block header and may, therefore, be "forced" to accept a rule change
1426 that the user of the light client would not otherwise approve of.

1427 A possible solution to this issue is *fraud proofs*, which are "alerts" that a full node can
1428 provide to a light client to convince the light client that a block is invalid. There are many
1429 subtle challenges to implementing fraud proofs, which are elaborated on in Section 15.5.

1430 However, even with a working fraud proof implementation, an adversary may still be able
1431 to partition the light client, filter out fraud proof messages, and launch the same attacks,
1432 which would not work against a full node.

1433 In short, while running a full node may be dramatically more resource-intensive than a light
1434 client, there are clear benefits. This includes substantially better privacy and requiring far
1435 less trust in third parties to be honest about the state of the ledger.

1436 3.3. Scalability Challenges and Block Sizes

1437 One of the most frequently discussed issues in the blockchain community is the scalability
1438 of the underlying technology. In particular, debates about the appropriate means to scale
1439 the Bitcoin network led to the high-profile August 2017 hard fork that created a new asset
1440 – Bitcoin Cash – with a larger maximum block size than Bitcoin proper. New replicated
1441 state machine projects routinely advertise themselves as being highly scalable. Usually, the
1442 claim is that a project can handle more transactions per second (TPS) than its competitors.

1443 Scalability in the context of replicated state machines can mean several different things.
1444 As mentioned above, it is often used in reference to increasing the system’s maximum
1445 transaction throughput, which is usually measured in TPS. Another scalability metric is
1446 the latency of transaction commitment, where a more scalable system has lower latency.
1447 Finally, it can also be used in reference to the total number of validators able to participate
1448 in consensus, such that a more scalable system is one that can function with a greater
1449 number of participating validators.

1450 Some of the confusion regarding scalability is a result of these different meanings of the
1451 term. This problem is greatest when comparing a permissioned system to a permission-
1452 less one based on TPS but can confuse a comparison between two permissionless systems
1453 as well. Despite being common, these comparisons are unscientific. In a permissioned
1454 network, one can fix the number of validators – n – of which f are byzantine, define a stan-
1455 dard transaction, measure TPS in a controlled environment, and make comparisons *to other*
1456 *permissioned systems* while holding the number of validators constant (as well as other en-
1457 vironmental factors). In a permissionless network, n naturally fluctuates. Furthermore, it
1458 is often the case that systems that scale to a larger number of validators are not capable of
1459 tolerating as high a level of throughput. That is, there is a trade-off between maximizing n
1460 and TPS, and increasing n is just as legitimate of a scaling goal as increasing TPS. There-
1461 fore, TPS does not provide a scientific way of comparing permissionless networks where n
1462 is not constant.

1463 There is no single, correct answer to how many consensus-participating nodes a system
1464 should have, and the trade-offs between types of scalability imply that this is highly sub-
1465 jective. The answer depends on the goals of the system designer, the system’s architecture,
1466 and the preferences of the system’s users. In a permissioned system, the existence of more
1467 replicas implies that more organizations can participate in the consortium, so the relevant

1468 number likely relates to the structure of the industry in question. Because the replicas are
1469 run by organizations that may be able to afford expensive hardware and bandwidth, permis-
1470 sioned systems can focus more on maintaining high TPS while assuming that substantial
1471 resources are provisioned at each node.

1472 Permissionless systems are more complicated, in part because it is common to have fully
1473 validating nodes that do not directly participate in consensus (i.e., by mining or staking). In
1474 a certain sense, the only node that one needs to care about is their own as long as the security
1475 threshold of the system is not violated. However, the network itself likely requires hundreds
1476 of nodes to provide sufficient diversity in terms of legal jurisdiction and geographic region
1477 to make it challenging for a moderately powerful adversary to disrupt the network. At the
1478 same time, many thousands of nodes are likely needed in order to serve the population of
1479 light clients and make the network useful to a large number of people. In permissionless
1480 networks, the intention is to allow anyone to participate anonymously. As such, it must
1481 target users with more average-to-worst-case hardware and bandwidth.

1482 Running a full node requires a number of computing resources, some of which present con-
1483 siderable scaling challenges. Consider the job of a full node in a typical blockchain system,
1484 such as Bitcoin. To join the network, a new node must download the entire blockchain his-
1485 tory all the way back to the genesis block. Then the entire chain must be fully validated
1486 (although there are some possible shortcuts that can reduce security in order to improve
1487 initial synchronization time) to build and evolve the system state as it goes. This usually
1488 involves checking that at least one digital signature per transaction, as well as checking
1489 that it is consistent with the system state, which is held either in memory or on disk. Upon
1490 synchronizing with the network, the node must constantly verify transactions and blocks as
1491 they are gossiped over the network while relaying this data over several network connec-
1492 tions as quickly as possible.

1493 The historical blockchain can be stored on a hard disk drive (HDD) because what occurred
1494 in the past does not need to be accessed frequently. However, enough nodes need to main-
1495 tain this data in order to serve it to new nodes who join the network and want to perform
1496 full validation. The history does not require quick access, but the state of the system does
1497 (in Bitcoin, this is the UTXO set). In order to validate transactions and blocks as quickly as
1498 possible, the state should be stored entirely in system memory. When this is not possible,
1499 some of the state must reside in more persistent storage. However, frequently accessing
1500 this state can result in a significant disk I/O burden. For example, Ethereum's system state
1501 is tens of gigabytes, making it all but impossible to run a node without a solid state drive
1502 (SSD), which provides quick random access to the system state. A node using an HDD
1503 is unlikely to ever remain in sync with the rest of the network. Unfortunately, RAM is
1504 expensive, and the state size can – in the worst case – expand as quickly as the maximum
1505 block size. To verify transactions, a node will heavily utilize a CPU. Even the simplest
1506 transactions almost always require verifying a digital signature, which is a fairly expensive
1507 operation. Worse, specially constructed transactions can be extremely time-consuming to
1508 validate (see Section 19.2 for a more detailed description of this issue), and in permission-

1509 less systems, it is necessary to consider worst-case algorithmic complexity to avoid attacks.
1510 Finally, this all takes place over a network, and sufficient networking resources must be de-
1511 ployed in order to remain in consensus. Bandwidth is important for all nodes (upstream
1512 bandwidth in particular), and latency is critically important for miners. The node must be
1513 connected to several other full nodes to exchange data and may also need connection slots
1514 and other resources to serve light clients.

1515 Perhaps the greatest scaling challenge is with the initial blockchain download and syn-
1516 chronization (IBD), which inherently becomes a bigger and bigger problem over time. In
1517 Bitcoin, IBD could well be impractical today were it not for a custom cryptographic library
1518 – *libsecp256k1* – that verifies signatures dramatically faster than OpenSSL, which was used
1519 before [58]. After exhausting the low-hanging fruits of performance improvements, IBD
1520 performance will likely take longer and longer. This can create a potential DevOps prob-
1521 lem, as it is typical to resynchronize a node any time validation code changes. If it takes
1522 a month to synchronize, development will be significantly hampered. Far worse, if IBD
1523 becomes impossible or sufficiently challenging, the free entry condition of permissionless
1524 networks may be violated.

1525 Performance is especially important for miners and, more broadly, those who participate in
1526 any synchronous consensus protocol. The latency of block propagation is a major contrib-
1527 utor to the centralization of power among consensus-participating nodes and is discussed
1528 in detail in Section 10.2.1. In short, larger miners are more likely to win in block "races"
1529 when two miners mine blocks at around the same time. With higher latency – such as
1530 that caused by larger blocks – more forks will occur, providing the larger miner with dis-
1531 proportionate rewards over time. Some argue that this is acceptable or even desirable and
1532 represents a healthy competition that results in miners upgrading their hardware and pro-
1533 visioning greater bandwidth access, which can lead to an increase in TPS. Bitcoin Core
1534 contributor Peter Todd responded to this by pointing out that requiring miners to dedicate
1535 more resources to survive actually reduces decentralization and makes those TPS worth
1536 less:

1537 What's tricky is designing a Bitcoin protocol that creates the appropriate in-
1538 centives for mining to remain decentralized, so we get good value for the large
1539 amount of money being sent to miners. I've often likened this task to building
1540 a robot to go to the grocery store to buy milk for you. If that robot doesn't have
1541 a nose, before long store owners are going to realise it can't tell the difference
1542 between unspoilt and spoilt milk, and you're going to get ripped off paying
1543 for a bunch of spoiled milk. Designing a Bitcoin protocol where we expect
1544 "competition" to result in smaller miners in more geographically decentralized
1545 places to get outcompeted by larger miners who are more geographically cen-
1546 tralized gets us bad value for our money. Sure it's "self-correcting", but not in
1547 a way that we want. [59]

1548 Block propagation latency is not the only incentives-related factor that can lead to cen-

1549 tralization in the absence of a sufficiently small maximum block size. For instance, in a
1550 system with no single hard limit on the maximum block size (where miners communicate
1551 their own desired limits instead), consensus becomes unstable and can lead to lengthy forks
1552 that benefit miners of larger blocks at the expense of those who can only process smaller
1553 blocks [60]. More to the point, miners who are able to handle larger blocks can form a "car-
1554 tel" of sorts that increases the block size and makes smaller miners unprofitable [60]. A
1555 similar cartel-forming result is likely to hold if the maximum block size is sufficiently high,
1556 even if there is agreement on it. The break-even cost for a miner to include a transaction
1557 in a block decreases as the mining pool size increases because orphan risk decreases. As a
1558 result, larger pools will be able to process more transactions and collect more transaction
1559 fees from users, forcing smaller miners out of business as the mining difficulty adjusts.

1560 Another major problem with increasing the on-chain throughput of a permissionless system
1561 is that miners become more and more incentivized to completely skip validation as the
1562 burden of validation increases. Miners would prefer to just assume that a block they receive
1563 is valid and start mining on top of it, which will provide a profitability advantage by starting
1564 work on the next block more quickly. This is especially significant if transaction fees are not
1565 a major portion of the miner reward, which is more likely to be the case with larger block
1566 sizes (an increased supply of block space leads to a decrease in the price of transaction
1567 inclusion). Unfortunately, this can result in (portions of) the network accepting invalid
1568 blocks, which happened on the Bitcoin network on July 4th, 2015 [61]. Six blocks were
1569 built on top of an invalid block because enough mining pools failed to validate blocks
1570 they received at the time. While full nodes merely experienced a slowdown in the growth
1571 of the honest chain, light clients could have been subject to attacks during that period,
1572 depending on which nodes they were connected to. If mining nodes had not been the only
1573 ones misconfigured or if non-mining full nodes were not the dominant type of node on the
1574 network, the situation could have been catastrophic, as light clients would have followed
1575 the invalid chain by default.

1576 A potential solution to this problem is the idea of fraud proofs, which was introduced in
1577 the previous section (and detailed in Section 15.5). If there were a way for full nodes to
1578 send an alert to light clients that proved that a block was invalid, light clients could be
1579 prevented from following an invalid chain. Unfortunately, even if fraud proofs could be
1580 easily deployed, they would not be a panacea. In order to construct a fraud proof, a full
1581 node needs access to the data that proves fraud. However, a malicious miner can construct
1582 an invalid block with an otherwise valid block header and simply refuse to publish the
1583 fraudulent part of the block. Until the fraudulent data is made available, light clients will
1584 follow the invalid chain if they are aware of it. In addition, light clients need to actually
1585 receive and validate the fraud proof, but this may be prevented by a Sybil attack. If the
1586 network is sufficiently centralized or light clients do not connect to many different full
1587 nodes, an explicit Sybil attack may not even be necessary.

1588 While fraud proofs would no doubt be beneficial, it is concerning to consider what might
1589 happen socially if fraud is not immediately detected, especially in a high-TPS system. If

1590 there are not "enough" full nodes (where "enough" is ill-defined), it becomes far more
1591 challenging to coordinate a response to fraud. Important ecosystem participants may be
1592 able to force undesirable rule changes onto the network and suppress fraud proofs for long
1593 enough that substantial commerce may occur on the fraudulent chain. If, say, a day passed
1594 before the fraud was noticed, the community would need to roll back a day's worth of
1595 honest commerce to correct the fraud, which might be sufficiently damaging on its own
1596 that users simply decide to follow the invalid chain anyway. To roll back the chain and
1597 continue using the old rules, the community would need to quickly bootstrap a new set of
1598 nodes capable of handling the high throughput of the chain. If this task is too challenging,
1599 then the fraud is likely to persist.

1600 A possible conclusion that one can draw from this is that the benefits of permissionless
1601 systems may fail to hold at sufficiently high TPS. When resource requirements become
1602 prohibitive for typical end users, the population of full nodes will decrease and may come
1603 to be dominated by a handful of large businesses. These few nodes may end up running
1604 on centralized cloud services and create a strong risk of correlated failures that hamper the
1605 availability of the system. Market concentration also makes the system more susceptible to
1606 censorship or coercive rule changes imposed by external actors, such as governments. As
1607 a result, a high-throughput permissionless chain may end up losing the free entry condition
1608 described in Section 1.6. With few synchronized nodes, there would be little incentive for
1609 them to share the blockchain with new nodes hoping to join the network. In fact, there
1610 would be strong reasons for them to avoid doing so: not only would this impose a signifi-
1611 cant bandwidth cost, but preventing new nodes from synchronizing is an effective way to
1612 prevent business competitors from arising. Further, new nodes themselves would have a
1613 very challenging time getting synchronized even if current nodes did serve the blockchain.
1614 They would need to pay for their own data center provisioning for heavy bandwidth and
1615 computation. It is no stretch of the imagination to think that, at this point, identification
1616 requirements could be imposed on validators, and the system could devolve into a de facto
1617 permissioned network.

1618 Despite all of these risks and problems, the scalability potential for permissionless ledgers
1619 is not all bad. Even if block sizes must be constrained, the maximum block size need not
1620 remain fixed forever (though a hard fork is required in order to increase it). As the under-
1621 lying computing resources needed to run a node improve and become cheaper over time,
1622 greater throughput becomes possible at the same cost for full validation. In addition, the
1623 more (publicly reachable) full nodes there are, the cheaper it becomes to run one. This
1624 is because much of the work that a full node performs involves serving others (including
1625 IBD), relaying blocks and transactions, and serving light clients. More nodes help spread
1626 the burden of these resource-hogging functions. Better networking subprotocols can be
1627 deployed in order to reduce the latency of block propagation and its centralizing effects.
1628 Some of these protocols are already in use and described in Section 10.2.1. For some ap-
1629 plications, scalability can be improved by eschewing a total ordering for transactions, such
1630 as with some DAG-based protocols or the reliable broadcast payment schemes mentioned

1631 near the end of Section 1.3. Payment channels and state channels can be used to reduce the
1632 amount of data stored on-chain and the bandwidth used to distribute such data, as described
1633 in Section 18.2.1. These technologies allow transactions and smart contracts to be executed
1634 by only the participants involved instead of the entire network. There are a number of ways
1635 to further scale computation, such as adding concurrency to the state machine (see Section
1636 18.1.1). Techniques such as optimistic rollups allow most nodes to skip the execution of
1637 transactions and rely on rationality assumptions, while zk-rollups allow nodes to skip exe-
1638 cution and instead check an easy-to-verify proof of correct execution (see Section 18.2.2).
1639 Sharding may be used to distribute the load of computation, bandwidth, and storage among
1640 smaller sets of validating nodes, as described in Section 15. Clever use of cryptographic
1641 accumulators can reduce the system state to a few kilobytes or less at the cost of increased
1642 bandwidth consumption, which may be able to resolve the problem of unbounded state
1643 growth. An example of this is the Utreexo proposal, which relies on dynamic hash-based
1644 accumulators [62].

1645 The most significant problem is bootstrapping a new node with the IBD process. The sys-
1646 tem can support a greater amount of activity without imposing a further burden on the
1647 initial synchronization process to the extent that many of the above scaling technologies
1648 can reduce the demand for block space by keeping transactions off-chain, as is done with
1649 payment channels and rollups. Technologies like Utreexo that keep the state size small can
1650 improve the synchronization process by allowing it to happen entirely in RAM, obviating
1651 the need for slower disk queries. The Mimblewimble protocol, described in Section 18.1,
1652 can make the burden of IBD scale with the state size rather than the complete transaction
1653 history, which can be a significant difference in practice. More exotic cryptographic con-
1654 structions like recursive SNARKs can be used to make IBD near instant, as is done in the
1655 Mina protocol, which maintains a constant-sized blockchain of less than 22 KB [63].

1656 There are other ways to mitigate the burden of IBD, such as by having either the node
1657 software or blockchain include commitments to the system state. While this does not fun-
1658 damentally solve the IBD scaling problem, it allows new nodes to become useful more
1659 quickly, though at a lower security level that is roughly on par with light clients. For exam-
1660 ple, some Ethereum clients have a "Fast Sync" mode that takes advantage of state commit-
1661 ments contained in block headers [64]. Fast Sync downloads the full blockchain but skips
1662 the execution of transactions prior to a specified *launch block*, assuming that transaction
1663 execution has been performed correctly up to that point. The node then contacts its peers
1664 to request a snapshot of the system state immediately prior to the launch block and verifies
1665 that the hash of the state matches the state commitment in the block header at that point.
1666 Afterward, the node performs standard IBD from the launch block toward the chain tip and
1667 executes transactions normally.

1668 The "assumeutxo" proposal for Bitcoin is similar but with two major differences [65]. First,
1669 Bitcoin does not commit to the system state anywhere in the ledger, so a hash of the state is
1670 hard-coded into the client software for a block height that is sufficiently far in the past that a
1671 significant amount of work has been proven since that block height. The node must acquire

1672 the state snapshot itself out of band. Second, while completing the synchronization from
1673 the snapshot's block height to the chain tip, a background process starts from the genesis
1674 block and executes the complete blockchain up to the assumed valid point to ensure that the
1675 state is correct, at which point the security model becomes identical to that of a full node.

1676 4. Practical Byzantine Fault Tolerance (PBFT)

1677 The celebrated Practical Byzantine Fault Tolerance (PBFT) algorithm by Castro and Liskov
1678 [4] was the first state machine replication algorithm to possess good enough performance
1679 to be used in the real world. Because many other permissioned consensus algorithms have
1680 a similar structure to PBFT, this section will describe the system in its entirety. The algo-
1681 rithm is secure under the partially synchronous network model and is optimally resilient,
1682 remaining secure against f faulty processes so long as $n \geq 3f + 1$. Roughly, the system
1683 works as follows:

- 1684 1. A client makes a request to the leader/primary replica.
- 1685 2. The primary broadcasts the request to the secondary replicas.
- 1686 3. Replicas execute the request and reply to the client with the result.
- 1687 4. The client accepts the result after receiving $f + 1$ replies from different replicas with
1688 the same result.

1689 In more detail, a client begins with a REQUEST message to the primary with the client
1690 ID, the command issued, and a timestamp. When the primary receives the request, they
1691 broadcast it and initiate a 3-phase commit (3-PC) process with the rest of the replicas:
1692 pre-prepare, prepare, and commit. The first two phases – pre-prepare and prepare – totally
1693 order the requests that clients sent in the same view even if the primary replica is faulty. The
1694 second two phases – prepare and commit – guarantee that committed requests are totally
1695 ordered *across* different views. The normal-case operation of the algorithm can be seen in
1696 Figure 6 and works as follows:

- 1697 1. **Pre-prepare:** The primary broadcasts a signed PRE-PREPARE message with a se-
1698 quence number, a view number, and a cryptographic digest of the request message.
1699 Replicas accept the message if the signature is valid, the digest matches the message,
1700 the view is correct, and the sequence number does not match an already accepted
1701 sequence number in that view. Because this phase relies solely on the primary and
1702 lacks redundancy, packet loss during this phase has the greatest impact on transaction
1703 confirmation latency [66].
- 1704 2. **Prepare:** If a replica accepts a PRE-PREPARE message, it broadcasts a signed PRE-
1705 PARE message to the rest of the replicas, which includes the view number, the se-
1706 quence number, the request digest, and the replica's ID. Other replicas accept this
1707 message if the signature is valid, the sequence number is correct, and the view num-




Fig. 6. PBFT normal case operation. In this case, C is the client, replica 0 is the primary, and replica 3 is faulty. [4]

1708 ber matches the replica’s current view. A given replica is said to be *prepared* if they
 1709 have seen a REQUEST, a corresponding PRE-PREPARE, and $2f + 1$ PREPARE
 1710 messages from different secondary replicas that match the PRE-PREPARE (includ-
 1711 ing its own). A replica that is prepared on a given request is also sometimes said to
 1712 be *locked* on that request. Once locked on a request in a given view, a replica will
 1713 only vote for that request in later views unless it "unlocks" from the request, which
 1714 would occur if it finds out that $2f + 1$ replicas are not locked on that request in that
 1715 view or higher. This unlocking occurs when entering the view change subprotocol,
 1716 as described in Section 4.1.

1717 3. **Commit:** When a replica is prepared, it broadcasts a signed COMMIT message
 1718 that includes the view and sequence numbers, the replica ID, and the request dig-
 1719 est. Replicas accept COMMIT messages if the signature is valid and the additional
 1720 data matches. A replica is said to be *committed-local* if it is prepared and has ac-
 1721 cepted $2f + 1$ COMMITs from different replicas (including itself) that match the
 1722 PRE-PREPARE for a given request. A request is said to be *committed* if $f + 1$ non-
 1723 faulty replicas are prepared on that message. The commit phase enforces that if a
 1724 non-faulty replica is committed-local on a given request, then the request is com-
 1725 mitted globally. A set of $2f + 1$ COMMIT messages from different replicas is sometimes
 1726 called a *commit certificate* or *quorum certificate*.

1727 Replicas execute the request when they are committed-local on it and when their state
 1728 reflects the sequential execution of all requests with lower sequence numbers. Replicas
 1729 respond to the client with a REPLY message with the current view number, the timest-
 1730 amp of the request, the replica ID, and the result of executing the operation. The client accepts
 1731 after seeing $f + 1$ matching replies from distinct replicas. This typical execution of the
 1732 protocol has communication complexity of $O(n^2)$ when the proposer is honest and the
 1733 network is synchronous because every replica must communicate with every other replica.

1734 The protocol described above creates an append-only totally ordered log of client-issued

1735 transactions that grows without bound. For performance reasons, it would be beneficial
1736 if replicas could discard old transactions instead of storing them permanently. Further, if
1737 a problem arises that causes a replica to fall out of sync with other replicas, it would be
1738 desirable to have a recovery procedure to acquire the missing state or lost messages. A
1739 checkpointing subprotocol is used to safely delete old transactions while creating a "proof
1740 of correctness" that allows a replica to trust that the state provided during recovery is the
1741 agreed-upon state of the remainder of the honest replicas. For safety, replicas cannot delete
1742 messages until they know that the associated transactions have been executed by at least
1743 $f + 1$ honest replicas and that it can prove this during the view-change subprotocol.

1744 A *checkpoint* is generated periodically, such as after a constant number of requests have
1745 been executed. When a replica generates a checkpoint, it broadcasts a signed CHECK-
1746 POINT message that includes the most recent sequence number and a digest of the state.
1747 Replicas store these messages until they receive $2f + 1$ CHECKPOINT messages for the
1748 same sequence number and digest, at which point the checkpoint is considered *stable*, and
1749 these $2f + 1$ messages constitute a proof of correctness for the checkpoint. When a replica
1750 has a proof of correctness for a checkpoint, the message log for requests up to that sequence
1751 number can be discarded. This method involves taking a snapshot of the system state. The
1752 system may halt for a few seconds while replicas save their state and stop processing re-
1753 quests. This can be mitigated by having replicas stagger when they take state snapshots
1754 [67].

1755 4.1. PBFT View Change

1756 The view change subprotocol provides liveness by allowing the state machine to make
1757 progress even when the leader is faulty. At most, liveness can be impeded by f faulty
1758 primaries in a row. The subprotocol ensures that replicas agree on the sequence number of
1759 requests that commit locally in different views at different replicas.

1760 Secondary replicas start a timer whenever they receive a request (and double the timer
1761 length if the view change fails for view $v + 1$ before attempting another view change to
1762 $v + 2$). If the timer expires, the replica initiates a view change and stops accepting messages
1763 within the old view. It sends a signed VIEW-CHANGE message with the new view number,
1764 the sequence number of the last stable checkpoint and its correctness proof, and the set of
1765 valid PRE-PREPARE/PREPARE messages for requests that have not been committed yet
1766 in the old view. When the new presumptive primary receives this message from $2f + 1$
1767 replicas, it broadcasts a signed NEW-VIEW message with the new view number, the set
1768 of valid VIEW-CHANGE messages received, and a set of PRE-PREPARE messages with
1769 the new view number. At this point, the primary moves to the new view, and replicas
1770 accept the new view if the signature is valid, the view number is correct, and the set of
1771 PRE-PREPAREs is valid. The secondary replicas broadcast PREPARE messages for each
1772 of these and move into the next view.

1773 If a replica receives $f + 1$ valid VIEW-CHANGE messages for views that are not the

1774 replica’s current view, it will broadcast a VIEW-CHANGE for the lowest view in the set
1775 (whether its timer has expired or not) to prevent it from starting a new view change too late.

1776 The communication complexity of the view change subprotocol is $O(n^3)$. The cubic mes-
1777 sage complexity comes from requiring the new primary to broadcast a NEW-VIEW mes-
1778 sage with quadratic size that contains $2f + 1$ commit certificates, where each commit cer-
1779 tificate contains $2f + 1$ messages. Because there can be up to f leader failures, even under
1780 synchrony, PBFT has worst-case complexity of $O(fn^3)$ or $O(n^4)$. More information on
1781 view change protocols can be found in Section 7.6.

1782 4.2. PBFT Security

1783 While rigorous security proofs are out of scope for this document, it is important to un-
1784 derstand why algorithms like PBFT are secure, at least informally. PBFT has optimal
1785 resilience for a partially synchronous (or asynchronous) protocol, such that the number of
1786 replicas $n \geq 3f + 1$. If up to f replicas are faulty or have their messages delayed, an honest
1787 replica must be able to proceed to the next step of the protocol after having communicated
1788 with only $n - f$ replicas. It is possible that those f missing replicas were actually honest,
1789 but network asynchrony has delayed their messages. This implies that f out of the $n - f$
1790 communicating replicas may in fact be faulty. Safety requires that a replica must hear from
1791 more honest replicas than faulty ones, so that $n - 2f > f$, which implies that $n > 3f$.

1792 Many BFT protocols make their security arguments based on *Byzantine quorums* [68] (a
1793 generalization of which is discussed in Section 7.5). The idea is that a set of replicas can
1794 be divided into a collection of subsets of replicas, called quorums, such that each pair of
1795 quorums intersects at a minimum of one honest replica. In theory, each quorum can act
1796 independently on behalf of the system, and *quorum intersection* guarantees that operations
1797 performed by distinct quorums maintain consistency. To see this, suppose two different
1798 transactions at the same position gain $\frac{2n}{3}$ votes (tx_1 and tx_2). Then a set of $\frac{2n}{3}$ distinct
1799 replicas (Q_1) voted for tx_1 , and another set of distinct replicas (Q_2) voted for tx_2 . Then
1800 $|Q_1 \cap Q_2| \geq \frac{2n}{3} + \frac{2n}{3} - n = \frac{n}{3}$. By assumption, fewer than $\frac{n}{3}$ replicas are corrupt, so an
1801 honest replica is in the set $\{Q_1 \cap Q_2\}$ and voted for both transactions at the same position.
1802 However, this is ruled out by the invariant that an honest replica will only vote for one
1803 transaction at any given position.

1804 Across multiple views, a locking mechanism provides safety at the end of the prepare
1805 phase. If tx_1 is committed in a view, then a quorum of replicas must have locked on tx_1
1806 in that view. If that quorum contains an honest replica that unlocks from tx_1 , then another
1807 quorum must be claiming to not be locked on tx_1 . The intersection of these two quorums
1808 contains at least one honest replica, but this honest replica would need to falsely claim that
1809 it is not locked on tx_1 , which is a contradiction. This demonstrates why 3-PC is necessary
1810 instead of 2-PC. A 2-phase commit would fail to achieve safety because a replica cannot
1811 guarantee that it will be prepared or locked by a sufficient number of honest replicas. That
1812 is, the replica would not know that $f + 1$ honest replicas are prepared until receiving $2f + 1$

1813 votes in the commit step. Without this assurance, two different requests could be committed
1814 at the same sequence number, violating safety. In fact, the "delegated BFT" algorithm used
1815 in the NEO blockchain system was essentially a two-phase version of PBFT with safety
1816 violations across view changes, and fixing the system involved adding the commit step
1817 back in [69, 70].

1818 As discussed earlier, PBFT's liveness is ensured by the view change subprotocol. Liveness
1819 can be framed in terms of *quorum availability*, which requires a full quorum of $2f + 1$ hon-
1820 est replicas available to respond to an honest leader within a view. Assuming that message
1821 delays do not grow exponentially, quorum availability is achieved via three mechanisms:

- 1822 1. To avoid initiating a view change too quickly, replicas use a timer that grows ex-
1823 ponentially with failed attempts at view changes. Exponential growth creates longer
1824 and longer periods in which replicas can synchronize their views and achieve quorum
1825 availability during a period of network synchrony.
- 1826 2. To avoid initiating a view change too late, after receiving $f + 1$ VIEW-CHANGE
1827 messages from other replicas, a replica will broadcast their own VIEW-CHANGE
1828 message even if the timer has not run out. This helps bring a lagging replica up to
1829 speed more quickly once they know at least one other honest replica wants to change
1830 views, and it prevents faulty replicas from forcing view changes too frequently.
- 1831 3. Faulty leaders cannot indefinitely hinder progress because there can only be faulty
1832 leaders for f views in a row at most.

1833 In PBFT, view changes only occur when leaders appear unresponsive or malicious to at
1834 least one honest replica. This has an interesting consequence for liveness in that transaction
1835 censorship cannot reliably be proven by a client. If a primary refused to order transactions
1836 from a particular client (or based on any other criteria), the client would be unable to prove
1837 this and force a view change. Therefore, despite formal liveness guarantees, it is possible
1838 for a primary to censor requests.

1839 A well-known, open-source PBFT library is BFT-SMaRt, which includes a slight change to
1840 the view change algorithm to make the system more modular [71]. Another implementation
1841 designed to be compatible with permissioned versions of Ethereum is called Istanbul BFT.
1842 However, the original version had both safety and liveness issues [72, 73], which led to an
1843 improved "IBFT 2.0" [74] upon being fixed. An alternative fixed IBFT is the one deployed
1844 in the Quorum system [75].

1845 4.3. Zyzzyva and Speculative Execution

1846 When none of the replicas are faulty, it is possible to significantly improve the performance
1847 of BFT protocols. Replicas can execute commands speculatively, assuming that none of
1848 the other replicas will experience a fault.

1849 Zyzzyva is an algorithm that pioneered this approach [76]. Replicas (speculatively) execute

1850 requests just after receiving the sequence number of the request from the primary and then
 1851 reply to the client. The client verifies consistency by checking that it received $3f + 1$ replies
 1852 with the same result. At this point, the replicas do not have a guarantee of consistency,
 1853 but the client does. In Zyzzyva, the client drives the consensus process by informing the
 1854 replicas if they detect an issue. When a client receives conflicting responses from replicas,
 1855 the client sends proof (two signed but conflicting messages) to the replicas in order to
 1856 initiate a view change. The replicas then "roll back" to a safe state prior to the inconsistent
 1857 execution and change the primary. This approach improves best-case latency and message
 1858 complexity but risks substantially worse performance when the primary is faulty.

1859 In more detail, the client sets a timer when it makes its request, and if there is a time-out
 1860 where the client only receives between $2f + 1$ and $3f$ consistent replies, it broadcasts a
 1861 commit certificate to the replicas and starts another timer. Upon receiving $2f + 1$ acknowl-
 1862 edgements of the commit certificate, the client considers the request complete. If there is
 1863 a time-out before receiving $2f + 1$ local commit acknowledgements, the client broadcasts
 1864 the request to every replica. The replicas then start a timer, ask the primary to order the
 1865 request, and initiate a view change if they do not receive word from the primary that the
 1866 request was ordered. A diagram of the message flow for Zyzzyva can be seen in Figure 7.

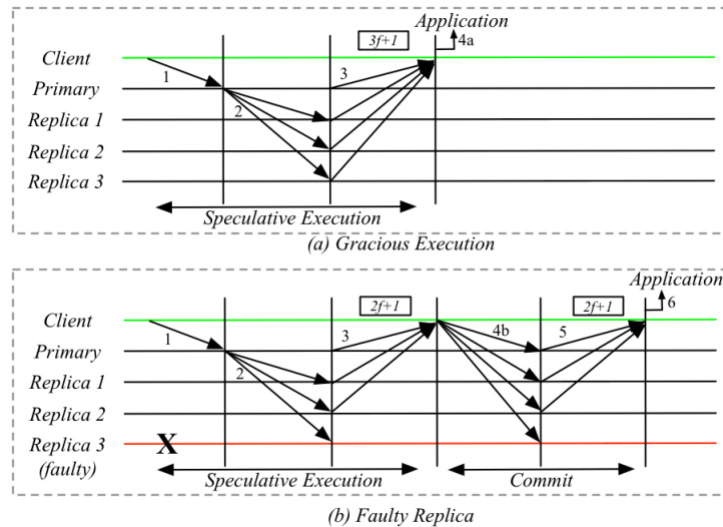


Fig. 7. Zyzzyva's speculative execution. In panel (a), all replicas reply to the client with consistent results in a timely manner. In panel (b), replica 3 is faulty, so an extra phase of commit acknowledgement is required. [77]

1867 Because Zyzzyva removes one of PBFT's phases (there is a pre-prepare phase and a second
 1868 phase), the view change requires adjusting. To safely change views without three phases,
 1869 an honest replica must not abandon a view until it knows every other correct replica will. To
 1870 this end, another phase is added to the view change: a correct replica broadcasts a complaint
 1871 when it suspects that the primary is Byzantine. Any replica that sees $f + 1$ complaints
 1872 knows to commit to a view change. This shifts costs from the agreement subprotocol to

1873 the view change subprotocol (so frequent view changes are performance-intensive). An
1874 additional adjustment is necessary to ensure that the replicas have a commit certificate
1875 when they are unanimous: replicas include all order requests (that is, the PRE-PREPAREs)
1876 since the last stable checkpoint in their view change messages, and a new honest primary
1877 extends the log with all requests that occurred in at least $f + 1$ of the $2f + 1$ view change
1878 messages they received. Note that the original protocol has a safety violation in the view
1879 change subprotocol that can be triggered by a faulty primary, which was only found 10
1880 years later [78].

1881 A related protocol, AZyzyva, does not suffer from the safety violation in Zyzyva [79].
1882 It is essentially the same as Zyzyva in the optimistic "common case," but it falls back on
1883 PBFT when detecting asynchrony or failures. That is, if a client does not receive $3f + 1$
1884 matching replies, it alerts the replicas, and the replicas then send signed message histories
1885 to the client. Upon receiving $2f + 1$ of these, clients switch to the backup mechanism:
1886 send these unforgeable histories to the replicas, and the replicas use PBFT to order a pre-
1887 specified number of requests, including the requests from the signed histories. This requires
1888 more steps than Zyzyva on the slow/recovery path and takes longer to switch back to the
1889 fast path, but it dramatically simplifies the codebase without a safety issue.

1890 SACZyzyva uses hardware-assisted trusted monotonic counters to improve upon Zyzyva
1891 [80]. SACZyzyva inherits the optimal resilience of Zyzyva and eliminates the need for
1892 non-speculative fallback while only requiring a single replica – the primary – to use a
1893 trusted monotonic counter at any given time. There must be $f + 1$ replicas with a trusted
1894 component to ensure that there is at least one correct replica that can be primary. The
1895 main idea is that the trusted monotonic counter value is attached to a message so that it
1896 is detectable if the sender equivocates due to the existence of a hole in the set of counter
1897 values. The primary uses the counter to bind a sequence of consecutive counter values to
1898 incoming requests to order them without communication between replicas.

1899 4.4. A Permissioned DAG: Blockmania

1900 Blockmania is a partially synchronous BFT algorithm that effectively embeds PBFT-like
1901 messages inside of a DAG [81]. Essentially, there is a leaderless version of a PBFT state
1902 machine embedded inside the block content in the DAG, and interpreting the DAG allows
1903 for recreating a PBFT execution transcript without requiring the transmission of additional
1904 messages. This is possible because replicas "gossip about gossip," or tell each other every-
1905 thing they learn from every other replica (a similar approach is used in Hashgraph, which
1906 is discussed in Section 6.2). Every honest replica produces a (single) block in each round
1907 (hence, being leaderless). To form the DAG, honest replicas include references to all valid
1908 blocks they have seen, including contradictory ones, when they create their own blocks.
1909 Compared to PBFT, this approach reduces the worst-case communication complexity from
1910 $O(n^4)$ to $O(n^2)$. This performance improvement is because COMMIT, VIEW-CHANGE,
1911 and NEW-VIEW messages normally need evidence sent with them, but correctly interpret-

1912 ing the DAG recreates this evidence implicitly.

1913 In Blockmania, blocks are referenced as (p, k) , where p is the creator of the block, and k is
1914 its sequence number. The block's contents can include both a list of transactions and a list of
1915 references to all valid blocks received from other parties. Each replica also stores a current
1916 view number, $view_p$, and a list of input and output messages, in_p and out_p . Assuming no
1917 timeouts, the normal case operation of the protocol decides on position (p, k) by using a
1918 protocol similar to PBFT:

- 1919 1. To propose a block B as a value for position (p, k) , replica p broadcasts a PRE-
1920 PREPARE (p, k, B, v) message, where v is the view number, which begins at 0.
- 1921 2. When a replica receives the first PRE-PREPARE message for view v , if the recipient's
1922 $view_p = v$, it broadcasts a PREPARE (p, k, B, v) message and adds the PREPARE and
1923 PRE-PREPARE messages to in_p . This is the only block that replica p will prepare in
1924 this view.
- 1925 3. Replicas listen for PREPARE (p, k, B, v) messages and add them to in_p when $view_p \geq$
1926 v . Once the replica has $2f + 1$ PREPARE (p, k, B, v) messages and the associated
1927 PRE-PREPARE in in_p , it broadcasts COMMIT (p, k, B, v) .
- 1928 4. Replicas listen for COMMIT (p, k, B, v) messages, adding them to in_p when $view_p \geq$
1929 v . Once the replica has $2f + 1$ COMMIT (p, k, B, v) messages in in_p , the replica
1930 considers B to be decided at position (p, k) . Note that B may be empty, or \perp .

1931 However, timeouts can occur. This triggers a view change, which happens as follows:

- 1932 1. The replica increases $view_p$ by one and broadcasts a VIEW-CHANGE $(p, k, view_p, S)$
1933 message, where S is the set of all PREPARE and PRE-PREPARE messages support-
1934 ing a block B that replica p is prepared on for position (p, k) . If the replica has not
1935 locked on a value at that position, $S = \emptyset$. At this point, replica p stops participating
1936 in prior views except for potentially receiving more COMMIT messages.
- 1937 2. Replicas wait for VIEW-CHANGE $(p, k, view_p, S)$ messages and add them to in_p if
1938 $v > view_p$. When $2f + 1$ of these messages have been seen, the replica updates $view_p$
1939 to v and broadcasts NEW-VIEW $(p, k, view_p, V)$, where V is the set of $2f + 1$ VIEW-
1940 CHANGE messages.
- 1941 3. When a replica sees the first NEW-VIEW (p, k, v, V) message where $v \geq view_p$ and
1942 $v > 0$, the replica sets $view_p = v$. The replica then checks if any VIEW-CHANGE
1943 messages included in V commit to a block B . If this is the case, the message is
1944 interpreted as a PRE-PREPARE (p, k, B, v) message; if not, the replica interprets it as
1945 a PRE-PREPARE (p, k, \perp, v) message.
- 1946 4. Replicas respond to the implied PRE-PREPARE accordingly and continue on with
1947 the protocol as normal.

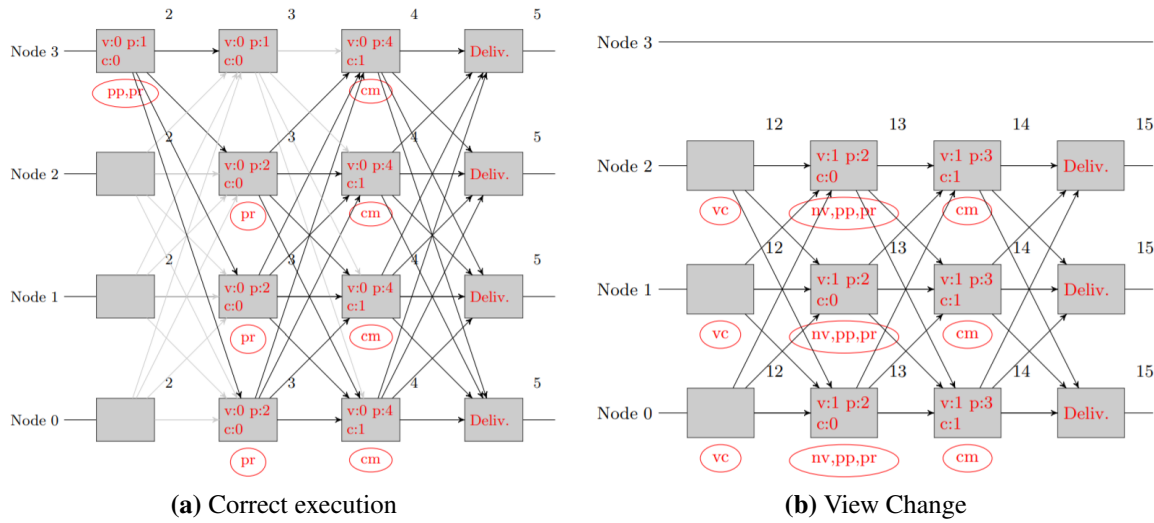


Fig. 8. Blockmania state machine interpretation example per block for position (3,2). Each state machine includes a view number (v) and a count for prepare (p) and commit (c) messages received by the block in red. The *out* buffer is in the red circle below the blocks (' pp ' for PRE-PREPARE, ' pr ' for PREPARE, ' cm ' for COMMIT, ' vc ' for VIEW-CHANGE, and ' nv ' for NEW-VIEW). [81]

1948 Throughout this process, all sent messages are implicitly included in out_p . The protocol
 1949 does not actually get executed via sending those messages directly. Instead, the protocol is
 1950 inferred from the block graph, where each block is interpreted as including a set of PBFT
 1951 state machines for positions that have not yet been decided. Denote a block at position
 1952 (p, k) as $B_{(p,k)}$. Blocks are associated with the union of messages in the *out* sets of all state
 1953 machines contained in the block. When a block $B_{(p,k)}$ includes a reference to another block
 1954 $B_{(p',k')}$, a replica interprets this as replica p' sending replica p the messages that are included
 1955 in the the *out* buffer of $B_{(p',k')}$. Those messages are then used to make progress in the PBFT
 1956 state machines of block $B_{(p,k)}$ based on the ordered sequence of block references. New
 1957 messages are added to the block's *out* buffer as the various state machines are interpreted
 1958 while validating the block. When first attempting to decide on position (p', k') , the replica
 1959 inserts a PRE-PREPARE($p', k', B_{(p',k')}, v = 0$) message, which is then included in its *out*
 1960 buffer. Eventually, when a PBFT state machine embedded in the DAG decides, then the
 1961 replica interpreting the state machine considers it decided as well. Once decisions are made
 1962 for all n blocks for round k , a total ordering of transactions can be derived in some agreed
 1963 upon way, such as the included fee.

1964 See Figure 8 for an example of interpreting the DAG. Note that only the blocks are broad-
 1965 cast; the material in red is only interpreted from the block but never sent as a separate net-
 1966 work communication. Figure 8a is a good execution, and Figure 8b shows a view change
 1967 when replica 3 is faulty.

1968 An advantage of interpreting a PBFT state machine rather than fully executing PBFT is that

1969 replicas can simply propagate blocks via gossip instead of having a complicated network-
1970 ing stack that handles the various message types. This can continue seamlessly even when
1971 view changes are occurring. In addition to the quadratic worst-case performance improve-
1972 ment over PBFT, Blockmania has low overhead due to interpreting blocks as messages
1973 themselves. Note that any deterministic BFT algorithm, not just PBFT, can be embedded
1974 in a communication DAG in this way [82].

1975 5. Modern High-Performance Blockchains

1976 Recently, a line of work has dramatically improved the simplicity and performance of BFT
1977 algorithms via pipelining and eliminating the need for a separate view change algorithm,
1978 among other innovations.

1979 5.1. Streamlined Blockchains

1980 The so-called "streamlined" blockchain is a generalization of a class of newer SMR con-
1981 sensus algorithms that are simple and highly performant [83, 84]. The simplicity of this
1982 approach makes it ideal for introducing more specific algorithms that have similar benefits,
1983 such as HotStuff, Casper-FFG, Tendermint, PiLi, and PaLa.

1984 Classical BFT protocols had what might be called a "normal mode" and a "recovery mode."
1985 The normal mode is potentially simple: a designated proposer proposes a block by signing
1986 it, other replicas vote on it by signing it, and it becomes finalized upon receiving $\frac{2n}{3}$ dis-
1987 tinct votes. This assures consistency when fewer than $\frac{n}{3}$ nodes are Byzantine, even if the
1988 proposer is Byzantine, regardless of network assumptions. However, a Byzantine proposer
1989 can stall liveness by not proposing anything or sending conflicting proposals, which neces-
1990 sitates a "recovery mode" (view change) to address liveness. Unfortunately, changes to the
1991 normal mode must be made to make the recovery mode work.

1992 This motivates modern protocols that dispense with the view change entirely. Streamlined
1993 blockchains, such as the Streamlet protocol described below, provide an opportunity to
1994 elect a new leader in every epoch, so view changes are baked into the normal mode [84].
1995 They support multiple proposer-election policies, including the more equitable policy of
1996 rotating every block and the more stable and performant policy in which rotation only
1997 occurs upon suspicion of misbehavior. Optimizations can be applied to the streamlined
1998 approach to improve upon the constant performance parameters presented here, which are
1999 not tight. Rather, the presentation is meant to be simple and easy to understand.

2000 A partially synchronous version of the Streamlet protocol is presented here. Note that there
2001 is an implicit assumption that whenever a replica observes a new message, it echoes it
2002 to everyone else. In the following, a block is considered *notarized* upon receiving more
2003 than $\frac{2n}{3}$ votes from distinct replicas, and a chain is considered notarized when the blocks it
2004 includes are all notarized. The protocol follows a *Propose-Vote-Finalize* paradigm:




Fig. 9. Streamlet Finalization Rule. All blocks are notarized. The chain prefix up to the epoch-6 block is finalized (“ $\perp - 2 - 5 - 6$ ”) once a replica has seen blocks 5, 6, and 7 built atop one another. There cannot be another block notarized at the same height as epoch-6, so chain “ $\perp - 1 - 3$ ” cannot grow further. [84]

- 2005 1. **Propose:** The primary replica for the current epoch proposes a new block that ex-
2006 tends from the longest notarized chain it has seen and breaks ties arbitrarily.
- 2007 2. **Vote:** Replicas vote for the first proposal they see from the epoch’s leader if the
2008 proposed block is appended to one of the longest notarized chains that the replica has
2009 seen. They do so by signing the proposed block.
- 2010 3. **Finalize:** When there are three adjacent blocks with consecutive epoch numbers in
2011 any notarized chain, replicas finalize the prefix of the chain up to the second of the
2012 three blocks (see Figure 9).

2013 To make the Streamlet protocol synchronous (and thus honest majority), this paradigm is
2014 tweaked by 1) decreasing the notarization requirement from $> \frac{2n}{3}$ to $> \frac{n}{2}$, 2) adjusting the
2015 finalization rule to require six consecutive epochs in a notarized chain in order to finalize
2016 the prefix of the chain with the last five blocks removed, and 3) requiring that the six blocks
2017 in consecutive epochs do not have conflicting notarizations at the same lengths at the time
2018 of finalization.

2019 5.2. PiLi and PaLa

2020 The protocols in this section follow the same *Propose-Vote-Finalize* paradigm with some
2021 adjustments in order to attain *optimistic responsiveness*. *Responsiveness* is the property
2022 that transactions are processed without reliance on synchrony. That is, processing depends
2023 on the actual network delay δ , which may be significantly less than worst-case delay:
2024 $\delta \ll \Delta$. Synchronous protocols usually have slow confirmation latency because Δ must
2025 be set conservatively for safety reasons. The idea of optimistic responsiveness – in which
2026 a protocol is responsive under certain good conditions – was introduced in [85], which
2027 showed that a $\frac{3n}{4}$ supermajority of honest and online replicas are necessary to attain this
2028 property for any honest majority protocol.

2029 The PiLi protocol is secure against rushing adversaries in a weakly synchronous network
2030 model, which provides resilience to node churn [86]. It attains optimistic responsiveness
2031 when $\frac{3n}{4}$ honest nodes remain online and the proposer is among them. If $\frac{n}{2}$ honest nodes are
2032 online for the length of the confirmation delay, transactions are confirmed in an expected
2033 constant number of synchronous rounds. Consistency is guaranteed for all honest nodes,
2034 even those that drop offline sometimes, as long as the total number of offline and corrupt
2035 nodes are less than $\frac{n}{2}$. This makes the protocol more robust than classical synchronous
2036 consensus. Compared to the synchronous Streamlet protocol, PiLi gains these advantages
2037 at the expense of worse finalization latency.

2038 In the PiLi protocol, a block is considered notarized upon receiving votes from the majority
2039 of voters. The finalization rule is that if a notarized chain ends with 13 consecutive epochs,
2040 the trailing eight blocks and the chain prefix are considered final. Define a *normal block*
2041 to be one where its epoch number is one greater than its parent's epoch number and a *skip*
2042 *block* to be one with an epoch number that is a multiple of 16 and at least 16 epochs apart
2043 from its parent.

2044 In each epoch, the leader proposes a block that extends the *freshes*t notarized chain it has
2045 seen (the freshest chain is the one whose chain tip has the higher epoch number). Replicas
2046 then vote on the first valid proposal they see, given that 1) the block extends from a parent
2047 block of an epoch no older than the freshest notarized chain that the replica had seen by
2048 the beginning of the previous epoch, and 2) if the block is not a skip block itself, then
2049 the replica has not observed any notarizations for conflicting blocks with the same epoch
2050 number for every block in the chain since the last skip block in the chain that the proposed
2051 block extends from.

2052 Further, at the beginning of each epoch, replicas set a timer. If the timer goes off, replicas
2053 send each other timeout messages and advance to the next epoch upon receiving timeout
2054 messages from the majority of replicas. Alternatively, when an $epoch_e$ block receives votes
2055 from a $\frac{3n}{4}$ supermajority of voters, an $epoch_{e+1}$ block can be proposed and immediately
2056 voted on by the leader of epoch $e + 1$, thus achieving optimistic responsiveness.

2057 For weak synchrony, replicas must only believe in "no-conflict" (that is, condition 2 above)
2058 if the majority of nodes confirm that belief rather than just using their own view. Voting on
2059 a block acts as an attestation that the replica has not seen recent conflicts, so a notarization
2060 on a block means that many replicas have not seen conflicts. Skip blocks prevent a denial-
2061 of-service attack where corrupt replicas double-vote, preventing honest nodes from voting
2062 due to conflict. Refusing to vote acts as a "complaint" that temporarily halts chain growth,
2063 but progress continues when an honest node produces a skip block during a period of
2064 synchrony.

2065 A related protocol is the partially synchronous PaLa protocol, which tolerates $\frac{n}{3}$ corruptions
2066 [87]. For PaLa, notarization requires $\frac{2n}{3}$ votes. If two blocks are built on top of each other
2067 in back-to-back epochs, the first of them is finalized. The protocol description defines the

2068 terms "second" and "minute" using the constant $c = 6$, where a second is $c\Delta$ and a minute
2069 is $c^2\Delta$.

2070 In each epoch e , a node is elected to be the proposer. If the proposer's chain ends with a
2071 block from epoch $e - 1$, they immediately propose a new block. If their chain ends with a
2072 block from an epoch earlier than $e - 1$ (that is, the prior proposer's block was not included),
2073 they wait for one second to potentially receive a fresher chain before proposing a new block.
2074 Whenever a replica during epoch e receives an $epoch_e$ block, it votes for the block if 1) it is
2075 consistent with the existing chain, 2) at least as fresh as the chain they saw at the beginning
2076 of epoch e , and 3) it has not previously signed a block for this epoch.

2077 A replica advances to epoch e when it is currently in an epoch earlier than e and either of the
2078 following occur: it sees a notarized chain for epoch $e - 1$, or it receives signed $timeout(e)$
2079 messages from more than $\frac{2n}{3}$ replicas. Replicas set a timer for one minute upon entering
2080 epoch $e - 1$ before broadcasting a $timeout(e)$ message.

2081 Because the replicas must wait to collect notarizations for each block of the entire chain
2082 before voting on the next block, performance can be hindered significantly in high latency
2083 environments, even when there is considerable bandwidth available. This motivates the
2084 "doubly pipelined" variant of PaLa, which rotates leaders less frequently and has better
2085 performance. Here, replicas vote even if not all notarizations have been collected for the
2086 ancestor chain so long as the number of blocks at the tip that are not notarized is bounded
2087 by a security parameter k (with the above scheme having $k = 1$), and finalization requires
2088 k consecutively notarized blocks. For the current epoch, a proposer can repeatedly propose
2089 blocks so long as there are fewer than k proposed but not notarized blocks at the end of the
2090 chain.

2091 PaLa allows committee reconfiguration to be done for half of the committee at a time,
2092 splitting the reconfiguration among two consecutive sets of k blocks to maintain safety
2093 when half of the committee has switched, assuming that each committee-half has fewer
2094 than $\frac{n}{3}$ corruptions.

2095 5.3. HotStuff

2096 HotStuff is a state-of-the-art high-performance BFT algorithm that utilizes pipelining while
2097 remaining secure in partially synchronous networks [88]. A synchronous version will be
2098 introduced later in this section [89], and an asynchronous version called VABA has also
2099 been proposed [90]. A variant of HotStuff is slated to be used in the Libra cryptocurrency
2100 (later renamed as Diem) [91].

2101 The protocol has optimistic responsiveness, $O(n)$ communication complexity when the pro-
2102 poser is honest, $O(n)$ communication complexity for view changes, and $O(fn)$ worst-case
2103 complexity with f leader failures. That is, compared to PBFT, the view change procedure
2104 has a quadratic reduction in communication complexity. The cost of having a new leader
2105 drive consensus is no greater than that for the current leader, which supports frequent leader

2106 changes. These improvements come at the cost of adding an extra phase to each view. If
2107 the new pre-commit phase were removed, then liveness could be permanently stalled by an
2108 attacker. Note that there are similar protocols, such as Tendermint and Casper FFG, that
2109 make a different design choice and sacrifice optimistic responsiveness in order to avoid
2110 adding the extra phase.

2111 The protocol defines a "quorum certificate" (QC) – sometimes also called a "commit cer-
2112 tificate" – as $n - f$ votes on a proposal. These quorum certificates serve as proof that a
2113 sufficient number of replicas agree on something and may have different names depending
2114 on what they are used for.

2115 The PBFT view change involves a proposer broadcasting proof from $2f + 1$ replicas about
2116 their highest commit certificates, each of which includes $2f + 1$ signatures. HotStuff im-
2117 proves upon this by a linear factor by including only one commit certificate and adding
2118 a rule that replicas only unlock a commit vote if they receive a commit certificate from a
2119 higher phase. An additional linear improvement to communication complexity comes from
2120 using aggregated threshold signatures instead of individual signatures and sending them to
2121 the leader instead of broadcasting them to everyone.

2122 When entering a new view or after a timeout occurs, replicas transmit NEW-VIEW mes-
2123 sages that include *prepareQC*, which is the highest QC for which a replica has voted to
2124 pre-commit. Replicas also store *lockedQC*, which is the highest QC for which the replica
2125 has voted to commit, and use the *SafeNode()* predicate to decide whether a proposal is safe
2126 to accept. The predicate is true if either of the following holds: the proposal extends from
2127 the currently locked node (to maintain safety), or the proposal has a higher view number
2128 than the current *lockedQC* (to maintain liveness). The basic partially synchronous HotStuff
2129 algorithm is described here before pipelining is added:

- 2130 1. **Prepare:** As leader, wait for $n - f$ NEW-VIEW messages, select the highest *prepareQC*
2131 included, and denote it *highQC*. The leader creates a proposal that extends from this
2132 block and broadcasts a PREPARE message with the proposal and *highQC* justifying
2133 its safety. Upon receiving a PREPARE message, other replicas decide to accept it or
2134 not using *SafeNode()* and send a PREPARE message with a (partial) signature on
2135 the proposal to the leader.
- 2136 2. **Pre-commit:** When the leader receives $n - f$ PREPARE votes on their proposal,
2137 they aggregate them into a *prepareQC* that is then broadcast in a PRE-COMMIT
2138 message. Replicas respond to the leader with a signed PRE-COMMIT message on
2139 the proposal.
- 2140 3. **Commit:** When the leader receives $n - f$ PRE-COMMIT votes on their proposal,
2141 they aggregate them into a *precommitQC* that is then broadcast in a COMMIT mes-
2142 sages, and replicas respond with their COMMIT vote. At this point, replicas set their
2143 *lockedQC* to the *precommitQC* they received.




Fig. 10. Pipelining in Chained HotStuff. A quorum certificate is simultaneously used in different phases of the protocol to improve efficiency. [88]

2144 4. **Decide:** When the leader receives $n - f$ COMMIT votes, they combine them into
 2145 a *commitQC* and broadcast it in a DECIDE message. Upon receiving a DECIDE
 2146 message, a replica considers the proposal linked to the *commitQC* to be committed,
 2147 executes the commands, and moves to the next view.

2148 The pipelined version of HotStuff is dubbed "Chained HotStuff" and works by having one
 2149 epoch's commit step be executed during the next epoch's prepare phase. The idea is de-
 2150 picted in Figure 10.

2151 Chained HotStuff improves upon the basic algorithm above by initiating a new view on
 2152 every prepare phase. This not only allows pipelining but also reduces the scheme to two
 2153 message types: NEW-VIEW and GENERIC. During the prepare phase of view v , the leader
 2154 collects other replicas' votes into a *genericQC*, which is then sent to the next views' leader
 2155 as part of delegating the next phase of the protocol. Instead of following through with a
 2156 pre-commit phase, the new leader initiates a new prepare phase with its own proposal for
 2157 view $v + 1$. That is, this second prepare phase is the prepare phase for view $v + 1$ while
 2158 simultaneously satisfying the pre-commit phase for view v . This continues such that the
 2159 prepare phase in view $v + 2$ acts as the pre-commit phase for view $v + 1$ as well as the
 2160 commit phase for view v .

2161 Each block contains a quorum certificate which may or may not point to the direct par-
 2162 ent block of the chain. When a block b^* has a QC that refers to (and thus *justifies*)
 2163 a direct parent, it is called a "One-Chain." That is, $b^*.QC.block = b^*.parent$. Define
 2164 $b'' = b^*.QC.block$. Then block b^* is a Two-Chain if it is a One-Chain and $b''.QC.block =$
 2165 $b''.parent$. It forms a Three-Chain if b'' forms a Two-Chain. There may also be gaps, which
 2166 represent leader failures. Assume the chain of justifications is $b \leftarrow b' \leftarrow b'' \leftarrow b^*$. When
 2167 b^* is a One-Chain, the prepare phase of b'' has succeeded. When b^* is a Two-Chain, then
 2168 the pre-commit of block b' has succeeded. When b^* is a Three-Chain, the commit phase of
 2169 block b has succeeded, and b is committed and final. This is illustrated in Figure 11.

2170 In more detail, Chained HotStuff works as follows. For each view v , the leader of the view
 2171 will update their *highQC* based on NEW-VIEW messages received and update *genericQC*
 2172 if *highQC* is higher. They then create a proposal b^* and begin the prepare phase by sending
 2173 a GENERIC message.




Fig. 11. Chained HotStuff justification. The blocks for v_4 , v_5 , and v_6 form a Three-Chain, whereas v_8 fails to form a One-Chain. v_8 also shows a quorum certificate that justifies a block that is not the direct parent of the block proposed in that view. Instead, the parent is a dummy block that represents the leader failure in v_7 . [88]

2174 Each replica then checks the predicate $SafeNode(b^*)$, and if it evaluates to true, they send
 2175 a **GENERIC** vote on b^* to the leader of view $v + 1$. If $b^*.parent = b''$, the replica updates
 2176 *genericQC* to $b^*.QC$. Here, b^* is a One-Chain, so the prepare phase for b'' has succeeded,
 2177 and the replica begins to pre-commit on b'' . If $b^*.parent = b''$ and $b''.parent = b'$, the
 2178 replica updates *lockedQC* to $b''.QC$. This means that b^* is a Two-Chain, so the replica
 2179 enters the commit phase of b^* 's grandparent, b' . If $b^*.parent = b''$, $b''.parent = b'$, and
 2180 $b'.parent = b$, then the replica executes commands through b and responds to clients. That
 2181 is, when b^* is a Three-Chain, the replica enters the decide phase on b^* 's great-grandparent,
 2182 b . The leader of the next view waits for all messages until they see $n - f$ votes on a proposal
 2183 and then updates their own *genericQC*.

2184 5.3.1. Sync HotStuff

2185 There is also a HotStuff variant intended for synchronous and weakly synchronous net-
 2186 works [89]. This variant has nearly optimal latency of $2\Delta (+O(\delta))$, optimistic responsive-
 2187 ness with $> \frac{3n}{4}$ honest replicas (which requires larger QCs than those used above), and
 2188 tolerates up to $\frac{n}{2}$ corruptions. Sync HotStuff removes the need for lock-step execution
 2189 that PiLi has and thus dramatically reduces finalization latency from what would otherwise
 2190 be between 40Δ and 65Δ . The key trick is to move the synchronous waiting periods off
 2191 the critical path, and the only step that requires waiting $O(\Delta)$ time is to check for leader
 2192 equivocation before committing.

2193 Leaders propose blocks, and then replicas vote on the proposal as soon as they have seen
 2194 it so long as they have not seen an equivocating proposal. Replicas then start a timer for
 2195 2Δ . The proposal is committed if the timer goes off, the view has not changed, and there
 2196 has not been equivocation. Replicas begin timers for subsequent heights without waiting
 2197 for the previous height to commit, so they may have numerous commit timers running
 2198 simultaneously at different heights. View changes are initiated upon detecting equivocation
 2199 or the leader failing to propose a block within 2Δ .

2200 For weakly synchronous networks, the above is changed so that the timer to commit block k
 2201 begins after receiving $f + 1$ proposals for block $k + 1$ (which has the certificate for k). When
 2202 timeouts occur, replicas broadcast commit messages for block k , and replicas commit after
 2203 seeing $f + 1$ commit messages.

2204 In Sync HotStuff, if a replica votes at time t , the commit is considered safe at time $t + 2\Delta$.
2205 Waiting for 2Δ provides safety because the replica's vote reaches all honest replicas by
2206 $t + \Delta$, at which point they will not vote for an equivocating block. If another replica voted
2207 for an equivocating block before $t + \Delta$, the first replica would have seen it by $t + 2\Delta$. This
2208 also implies that all honest replicas will have voted by $t + \Delta$, so a QC will form by $t + 2\Delta$,
2209 ensuring safety.

2210 An earlier version of the protocol suffered from a "force-locking" attack that removed
2211 safety from the synchronous version and liveness from all versions but has since been fixed
2212 [92].

2213 5.4. Further Optimizing Latency

2214 It is possible to obtain optimistic responsiveness with lower latency than the protocols
2215 described above using protocols similar to PiLi and Sync HotStuff [93]. The challenge of
2216 optimizing for latency is in handling the situation where it is unclear whether the fast-path
2217 optimistic conditions are met. If it is clear that the optimistic conditions hold, there are
2218 latency-optimized protocols to use for that setting.

2219 When an optimistically responsive commit rule and a slower synchronous commit rule ex-
2220 ist in synchronous protocols, the sum of the latencies of the two rules must be at least 2Δ ,
2221 such that when a faster optimistic commit rule is used, the synchronous commit rule must
2222 be correspondingly slower to make up for it [93]. For example, if the optimistic commit
2223 path requires 0.5Δ to commit, then the synchronous path requires at least 1.5Δ to commit.
2224 In Sync HotStuff, the synchronous slow-path latency is 2Δ , while the fast path commit la-
2225 tency is 2δ , so it is not fully optimal. Further, switching between the optimistic and the
2226 slow paths imposes its own delays of at least Δ . The protocols described in [93] achieve op-
2227 timistic responsiveness and optimal latency simultaneously by removing the explicit switch
2228 between fast and slow paths and making it such that replicas need not know which path they
2229 are on; both progress simultaneously. Alternatively, the Apollo SMR protocol has an opti-
2230 mistic finalization latency of $(f + 1)\delta$ and commits a block every δ time units [94] while
2231 maintaining linear communication complexity in the best case. It achieves this by having a
2232 block committed whenever $f + 1$ blocks are built on top of it, which obviates the need to
2233 detect equivocation.

2234 6. Asynchronous BFT

2235 The protocols described up to this point in the document all fail to maintain liveness when
2236 messages may be arbitrarily delayed and synchronous ones further lose their safety. Even
2237 when partially synchronous timing assumptions hold in practice, performance degrades
2238 rapidly in an unpredictable network, and throughput fails to quickly recover after a tempo-
2239 rary network partition. In contrast, asynchronous protocols are able to continue confirming
2240 transactions at network speed as protocol messages arrive, regardless of the message delay

2241 or ordering. These protocols do not rely on timeout mechanisms, which may make them
2242 easier to implement.

2243 To see how liveness fails under partial synchrony, one may consider the following adver-
2244 sarial network scheduler attacking an execution of PBFT, which results in zero throughput.
2245 Assume that a single replica has suffered a crash fault. The adversarial scheduler induces
2246 message delays whenever the leader is honest, causing an eventual view change on time-
2247 outs and thus moving to the next leader. At some point, the crashed replica becomes the
2248 next leader, and the scheduler quickly delivers all messages between honest replicas. Un-
2249 fortunately, because the leader is crashed, there is still no progress made. The existence
2250 of synchronous periods where messages are delivered is insufficient to maintain liveness
2251 because they occur at times when the algorithm is unable to make use of them. An asyn-
2252 chronous protocol, on the other hand, would make progress and confirm transactions during
2253 these synchronous periods.

2254 Because the timeout interval (usually) increases exponentially in partially synchronous pro-
2255 tocols, even a limited version of this network attack would dramatically slow the network's
2256 recovery from a partition. If a replica is crashed, the network is partitioned for a length of
2257 $2^D\Delta$, and the scheduler heals the network when the crashed node is supposed to become
2258 leader, then there is a $2^{D+1}\Delta$ delay before starting a new view change despite the network
2259 being synchronous. Again, an asynchronous protocol would begin confirming transactions
2260 immediately during this period.

2261 6.1. HoneyBadgerBFT

2262 The first practical asynchronous BFT algorithm designed was the leaderless HoneyBad-
2263 gerBFT, which utilizes randomization to overcome FLP impossibility [95]. HoneyBad-
2264 gerBFT provides security against static adversaries but can be secure against adaptive ad-
2265 versaries by substituting in different cryptographic primitives at the cost of worse perfor-
2266 mance [96]. The primary novelty is a reduction from atomic broadcast to the asynchronous
2267 common subset (ACS) primitive, which was improved from $O(n^2)$ to $O(1)$ (the prior state-
2268 of-the-art was from [97]) but also uses threshold decryption for censorship resistance and
2269 improving liveness.

2270 An ACS protocol allows each replica to propose a value and then ensures that every replica
2271 outputs a common vector that includes the input values from at least $n - 2f$ correct replicas.
2272 One can create a state machine replication protocol by simply having each replica propose
2273 a set of transactions as input into the ACS primitive with the output being the union of the
2274 transactions. ACS has the following requirements:

- 2275 • Agreement – If an honest replica outputs v , then every replica outputs v .
- 2276 • Validity – If an honest replica outputs a set v , then $|v| \geq n - f$ and v contains the
2277 inputs of at least $n - 2f$ honest replicas.




Fig. 12. HoneyBadgerBFT’s ACS structure. There are n instances of reliable broadcast followed by n instances of asynchronous binary agreement. [98]

- **Totality** – If $n - f$ honest replicas receive an input, then all honest replicas produce an output.

ACS protocols typically involve two phases: 1) reliable broadcast to distribute each replica’s proposed values and 2) asynchronous binary Byzantine agreement (sometimes abbreviated as ABA or ABBA) to decide on a bit vector that indicates which reliable broadcasts have completed. The structure of HoneyBadgerBFT’s ACS protocol can be seen in Figure 12.

The HoneyBadgerBFT algorithm is composed of a variety of subprotocols, including an expected constant-time ABA protocol, threshold signatures to produce the common coins required for the ABA protocol, a bandwidth-efficient reliable broadcast protocol, and threshold decryption. The protocol idea is simple: if the maximum block size is B and there are n replicas, each replica will randomly choose $\frac{B}{n}$ transactions from their mempools and encrypt them. This encrypted set is then passed into the ACS protocol, which outputs the common subset of transactions. Each replica then performs its share of decryption on the common subset and broadcasts its decryption share. Upon receiving $f + 1$ decryption shares, replicas fully decrypt the transaction set. The block that results is the canonically sorted set of transactions.

Note how efficiency is improved by having honest replicas try to propose disjoint subsets. By having replicas pick transactions randomly, it is expected that each transaction is only proposed by a single replica. The addition of threshold decryption prevents some types of transaction censorship because replicas will not know which transactions were proposed by whom. A trade-off to this is that a faulty replica may propose invalid transactions in the ACS, which would not be detected until after they are included in a block. This means that after transactions are ordered, the invalid transactions must be removed from the block or at least remain unexecuted.

For the reliable broadcast subprotocol, HoneyBadgerBFT actually employs a more sophis-

2303 ticated primitive called asynchronous verifiable information dispersal (AVID). This primi-
2304 tive was first presented in [99] and combines asynchronous reliable broadcast and erasure
2305 coding. The use of erasure coding significantly reduces communication complexity for
2306 large messages compared to Bracha's broadcast protocol described in Section 1.2.

2307 6.1.1. Mostéfaoui et al.'s Asynchronous Binary Agreement Protocol

2308 Mostéfaoui et al. introduced a core component of HoneyBadgerBFT and some related
2309 protocols: asynchronous binary agreement (ABA), where agreement occurs over a single
2310 bit [100]. Note that the "broadcast" protocols described here do not meet the standard
2311 definition used in Section 1.2 because each party transmits a value.

2312 The protocol is built in layers with the most basic component being a binary value (BV)
2313 broadcast protocol, which eliminates from consideration any value that was broadcast only
2314 by malicious parties. Each correct process p_i BV-broadcasts a binary value and eventu-
2315 ally obtains a set of binary values stored in bin_values_i . Let $witness(v)$ be the number
2316 of different processes from which $B_VAL(v)$ was received, and let $bin_values_i = \emptyset$. The
2317 operation $BV_BroadcastMSG(v_i)$ consists of multicasting $B_VAL(v_i)$ and then returning.
2318 Then, when p_i receives $B_VAL(v)$, it does the following:

- 2319 1. If p_i has not yet multicast $B_VAL(v)$ to other replicas, and if $witness(v) \geq f + 1$, then
2320 p_i multicasts $B_VAL(v)$. Note that a process echoes a value like this only once.
- 2321 2. If $witness(v) \geq 2f + 1$ and $v \notin bin_values_i$, then p_i locally delivers the value by
2322 setting $bin_values_i \leftarrow bin_values_i \cup \{v\}$.

2323 The next layer is strong binary value (SBV) broadcast, which synchronizes processes such
2324 that if a single value v is delivered to an honest process, then v is delivered to all honest pro-
2325 cesses. The operation $SBV_BroadcastMSG(v_i)$ begins by invoking $BV_BroadcastMSG(v_i)$
2326 and waiting until bin_values_i is non-empty.

2327 When the waiting stops, bin_values_i may not be at its final value. Nevertheless, p_i then
2328 multicasts a message $AUX(w)$, where $w \in bin_values_i$. If there is more than one value in
2329 bin_values_i , any value w will do. Then p_i waits until there exists a set $view_i$ such that its
2330 values belong to bin_values_i and come from $AUX()$ messages received from $n - f$ distinct
2331 processes. Finally, the algorithm returns $view_i$.

2332 The third layer is a double-synchronized binary value (DSBV) broadcast algorithm. The
2333 goal of this algorithm is to replace the potentially distinct values v and w that are broadcast
2334 by honest replicas by at most one of v or w , plus potentially the default value denoted \perp . As
2335 with SBV, this returns a $view_i$, but it can now contain a default value \perp and – at most – only
2336 one of the two possible binary values. DSBV broadcast reduces the power of Byzantine
2337 replicas such that no view delivered by an honest replica includes values broadcast only by
2338 Byzantine replicas (from BV broadcast) and that if a view delivered to an honest replica
2339 contains only the single value v , then v will be included in the views of all honest replicas

2340 (from SBV broadcast). The operation $DSBV_BroadcastMSG(v_i)$ works as follows:

- 2341 1. $view_i[0] \leftarrow SBV_BroadcastSTAGE[0](v_i)$.
- 2342 2. If $view_i[0] = \{v\}$, then $aux_i \leftarrow v$. Else, $aux_i \leftarrow \perp$.
- 2343 3. $view_i[1] \leftarrow SBV_BroadcastSTAGE[1](aux_i)$.
- 2344 4. Return $view_i[1]$.

2345 That is, DSBV broadcast consists of two stages of SBV broadcast. When the first SBV
2346 broadcast returns, it is possible that each of the two binary values are represented in $view_i$.
2347 This occurs when honest replicas input different values to their SBV broadcast operation. In
2348 line (2) of the algorithm, replicas filter the output from the first SBV broadcast to determine
2349 their input into the second instance. While two possible values – say, a and b – were
2350 possible inputs into the DSBV broadcast (and therefore the first SBV broadcast), the second
2351 SBV broadcast will only include one of those inputs (which must have been proposed by
2352 an honest replica) and \perp .

2353 The final ABA algorithm utilizes this DSBV broadcast operation and a *weak common coin*.
2354 Common coin protocols result in some random output that can be observed by all partic-
2355 ipants and is unpredictable to an adversary. For a weak common coin protocol, there is a
2356 constant probability that the replicas involved will see different random values returned by
2357 the functionality. In the description of the ABA protocol that follows, the weak common
2358 coin is invoked by calling $CCRandom()$.

2359 The ABA algorithm begins when a process p_i invokes $propose(v_i)$, where $v_i \in \{0, 1\}$. The
2360 operation begins by setting est_i (p_i 's current estimate of the decision) to v_i and setting
2361 the round number, r_i , to zero. Processes proceed in asynchronous rounds, where each
2362 round is two phases, and each phase is made up of a single DSBV broadcast instance. Let
2363 $view_i[r_i, 1..2]$ be the local results of the DSBV broadcast instance at round r_i . The ABA
2364 algorithm repeats the following sequence indefinitely:

- 2365 • $r_i \leftarrow r_i + 1$.
- 2366 • Phase 1: Help replicas agree on a single value.
 - 2367 1. $view_i[r_i, 1] \leftarrow DSBV_BroadcastPHASE[r_i, 1](est_i)$.
 - 2368 2. $b_i[r_i] \leftarrow CCRandom()$.
 - 2369 3. If $view_i[r_i, 1] = \{v\}$ and $v \neq \perp$, then set $est_i \leftarrow v$. Otherwise, $est_i \leftarrow b_i[r_i]$.
- 2370 • Phase 2: Help replicas recognize when they have reached a round where all honest
2371 replicas will forever maintain the same estimate.
 - 2372 1. $view_i[r_i, 2] \leftarrow DSBV_BroadcastPHASE[r_i, 2](est_i)$.
 - 2373 2. If $view_i[r_i, 2] = \{v\}$, then decide v if not yet done.

2374 3. Else if $view_i[r_i, 2] = \{v, \perp\}$, then $est_i \leftarrow v$.

2375 4. Else if $view_i[r_i, 2] = \{\perp\}$, then skip to the next round.

2376 As presented, the ABA algorithm is non-terminating. To make it guarantee termination,
2377 whenever a process decides, it also terminates and multicasts a message $TERM[r](v)$ to
2378 tell other processes it is done. That is, replace "decide v if not yet done" with "multicast
2379 $TERM[r_i](v)$ and return v ." It also requires mild tweaks to the SBV and BV broadcast
2380 algorithm to accommodate waiting for these messages.

2381 6.1.2. Reducing HoneyBadgerBFT's latency with BEAT

2382 Duan et al. present BEAT – a family of five asynchronous BFT protocols with different
2383 trade-offs, all of which outperform HoneyBadgerBFT in latency and some of which im-
2384 prove throughput [101]. These protocols follow a similar structure as HoneyBadgerBFT
2385 but use more efficient labeled threshold cryptography and erasure coding schemes to im-
2386 prove performance.

2387 Of the five algorithms, BEAT0, BEAT1, and BEAT2 are most relevant to this document
2388 because they are for general state machine replication. In contrast, BEAT3 and BEAT4
2389 are "BFT storage" algorithms that do not require replicas to maintain the full system state.
2390 BFT storage only allows read and write operations to a key-value store, and the state is
2391 erasure-coded, so it cannot support on-chain smart contracts.

2392 The base algorithm, BEAT0, uses labeled threshold encryption to make transactions uniquely
2393 identifiable. This makes it easier for replicas to ignore duplicate transactions that had been
2394 input to the ACS subprotocol. Further, instead of HoneyBadgerBFT's use of threshold sig-
2395 natures to derive the common coin, BEAT0 uses a more efficient threshold pseudorandom
2396 function. Finally, it uses more efficient erasure-coding. The other BEAT algorithms use
2397 these improvements as well.

2398 BEAT1 replaces the AVID reliable broadcast primitive used in HoneyBadgerBFT with
2399 Bracha's broadcast. This induces a trade-off because it increases asymptotic communi-
2400 cation complexity. However, in situations where there is little contention and small batches
2401 (i.e., small blocks), Bracha's broadcast has significantly lower latency than AVID.

2402 BEAT2 uses Bracha's broadcast as well. Additionally, it optimistically moves the threshold
2403 encryption to the client instead of the replicas. This substantially reduces latency because
2404 the threshold encryption is one of the biggest drivers of latency in HoneyBadgerBFT. The
2405 trade-off is that there is a weaker version of liveness in which servers may be able to censor
2406 specific clients. However, full liveness is restored by using an anonymous communication
2407 network like Tor.

2408 The asymptotic communication complexity with n replicas ordering a set of transactions
2409 of size B is $O(nB)$ in HoneyBadgerBFT and BEAT0. BEAT1 and BEAT2 worsen this to

2410 $O(n^2B)$. BEAT3 and BEAT4, which are not covered in this document, reduce this complex-
2411 ity significantly to $O(B)$ by having agreement occur on a constant-sized checksum instead
2412 of all data in the block as well as some other techniques.

2413 6.1.3. Improving ACS Performance with Dumbo

2414 HoneyBadgerBFT's asynchronous common subset protocol, as shown in Figure 12, re-
2415 quires n instances of reliable broadcast (RBC) to be run so that each replica can propose
2416 their own set of transactions, followed by n instances of asynchronous binary agreement
2417 to decide on each of those inputs. In more detail, whenever a replica delivers a value that
2418 was broadcast by its peer P_i , the replica starts the i -th ABA instance with an input of 1.
2419 When any honest replica receives 1 from $n - f$ ABA instances, it inputs 0 to the remaining
2420 instances and moves on.

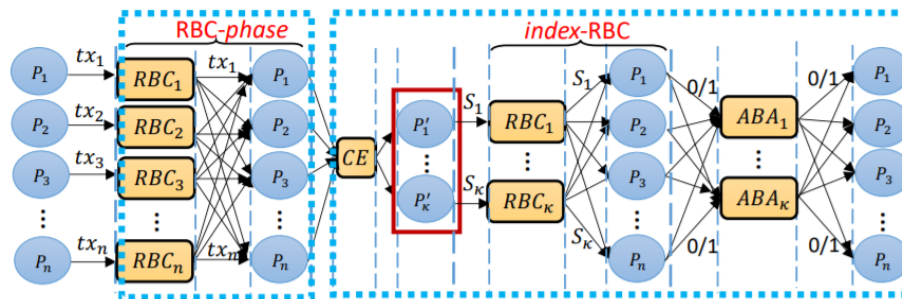
2421 Because the ABA protocol is randomized, the expected number of rounds of each ABA in-
2422 stance is constant. However, running n instances at the same time can blow up the expected
2423 number of rounds considerably. This requirement of executing n ABA instances is a major
2424 performance bottleneck for the protocol, particularly because the slowest ABA instance
2425 determines the running time to finish the ACS execution. This problem is exacerbated by
2426 having different ABA instances start at different times, depending on when each instance
2427 of reliable broadcast delivers its value.

2428 The Dumbo family of protocols improves upon HoneyBadgerBFT's performance by mod-
2429 ifying the ACS protocol to require substantially fewer instances of the ABA subprotocol to
2430 be run [98]. Two particular improvements are suggested:

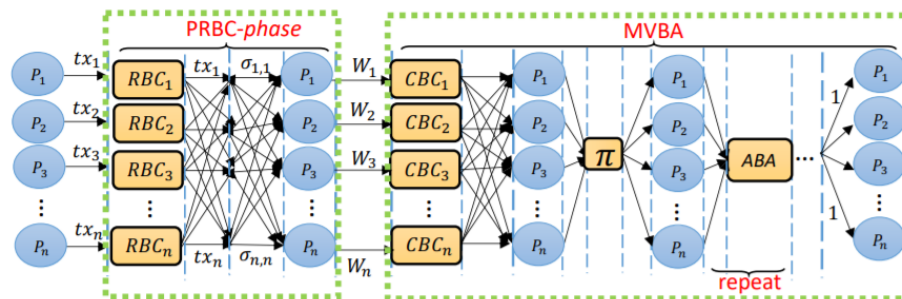
- 2431 1. Rather than having an ABA instance per replica, a small set of κ "aggregators" can
2432 be chosen to lead the reliable broadcast instances, such that only κ instances of ABA
2433 must be run in the Dumbo1 protocol.
- 2434 2. Careful use of multi-valued validated Byzantine agreement (MVBA) can further re-
2435 duce the number of ABA instances required to a constant (an expected three consec-
2436 utive rounds of ABA).

2437 MVBA was rejected in HoneyBadgerBFT because it has high communication complexity
2438 when run over values that may be of large size, such as a block of transactions. However, a
2439 proper strategy would allow MVBA to be run on smaller values, as is done in Dumbo2. Part
2440 of the trick involves a new primitive – provable reliable broadcast (PRBC) – that outputs
2441 a short proof that at least one honest peer has received the input to the reliable broadcast.
2442 The first improvement can be seen in Figure 13a, and the second is displayed in Figure 13b.
2443 Note that the insights from BEAT0-BEAT2 can be applied to improve Dumbo as well.

2444 Dumbo1 requires two phases of reliable broadcast. The first phase is the same as Honey-
2445 BadgerBFT, where replicas broadcast their input sets of transactions. At this point, if there
2446 were an honest leader, they could wait until $n - f$ RBC instances have completed and then



(a) ACS of Dumbo1



(b) ACS of Dumbo2

Fig. 13. ACS structure of Dumbo protocols. RBC is reliable broadcast, ABA is asynchronous binary agreement, CE is committee election, PRBC is provable reliable broadcast, MVBA is multi-valued validated Byzantine agreement, CBC is consistent broadcast (a relaxation of reliable broadcast), and π is a permutation. [98]

2447 simply tell the other replicas what the results were. By selecting κ leaders, the protocol can
2448 ensure that there is at least one honest leader with high probability. These κ leaders then
2449 begin a second round of RBC called index-RBC, where they broadcast the set of indices of
2450 the $n - f$ values that the leader has already received. That is, the inputs into the index-RBC
2451 instances are just small subsets of indices, denoted S_i , rather than the actual transactions.
2452 The final phase is to run κ ABA instances, where an honest replica will input 1 to the i -th
2453 instance if it has already seen S_i and all of the corresponding values from the first round of
2454 reliable broadcasts. At least one honest replica will receive all of the input values that cor-
2455 respond to S_i , so reliable broadcast ensures that all honest replicas will eventually deliver
2456 those values as well. When an honest replica outputs 1 for the i -th ABA instance, all honest
2457 replicas are eventually guaranteed to output 1 as well. The overhead for this procedure is κ
2458 additional RBC instances and a single coin tossing for committee election (CE in 13a), but
2459 the ABA savings dominate this overhead.

2460 While Dumbo1 reduces the number of ABA instances needed, it is still necessary for κ in-
2461 stances to be executed. In theory, this is still $\kappa - 1$ "too many" instances. If it were possible
2462 to identify a single, correct input vector, this could be further reduced. This motivates the
2463 use of MVBA in Dumbo2. In MVBA, each replica submits an input value, and the output
2464 is a single value that satisfies a predefined predicate. Since MVBA is inefficient when there
2465 are large inputs, and the inputs to the ACS protocol may indeed be large, it is necessary to
2466 figure out how to safely reduce the size of the MVBA input. This is done through the clever
2467 use of a short "indicator" value, denoted W_i in Figure 13b, which is used as the MVBA input
2468 instead of the much larger original data. The MVBA protocol will output a single indicator
2469 that honest replicas use to work backward to select the corresponding reliable broadcast
2470 instance. Unfortunately, honest replicas may end up outputting the indicator from a Byzantine
2471 replica. This problem is addressed through the use of a new primitive, *provable* reliable
2472 broadcast (PRBC), which provides a succinct proof that at least one honest replica received
2473 the input. PRBC can be realized through the use of threshold signatures on the RBC index:
2474 when honest replicas receive a value from a sender, they multicast a signature share on the
2475 index and can construct the full signature as the proof upon seeing $f + 1$ signature shares
2476 for the same index.

2477 The indicator value will include a set of these proofs and their corresponding RBC in-
2478 dices. The predicate used for MVBA requires that $n - f$ of these indices exist and that their
2479 corresponding proofs are valid. As a result, an honest replica knows that the transactions
2480 corresponding to the included RBC indices were received by enough replicas to know that
2481 at least one (other) honest replica received them and that all honest replicas will eventually
2482 receive them. Therefore, if an honest replica initiated MVBA after seeing $n - f$ PRBC
2483 instances complete, all honest replicas simply need to wait to receive the values from the
2484 PRBC phase.

2485 The Dumbo protocols can also be combined with an optimistic fast path to substantially
2486 improve latency when network conditions are good [102]. In this case, a much faster
2487 consensus protocol is executed, and Dumbo is used as a fallback mechanism in a way that

2488 is similar to a view change.

2489 6.2. An Asynchronous Permissioned DAG: Hashgraph

2490 The protocols above allow a set of replicas to maintain a sequentially ordered log of trans-
2491 actions or batches of transactions in the form of a blockchain. This section explores Hash-
2492 graph – an asynchronous BFT protocol that operates over a DAG. Like HoneyBadgerBFT,
2493 Hashgraph is a leaderless protocol that uses randomization to achieve consensus with prob-
2494 ability 1 [103]. The protocol is similar to Blockmania (Section 4.4) in that it involves
2495 "gossip about gossip" and utilizes "virtual voting" to interpret the resulting communication
2496 graph (called a *hashgraph* throughout this section). That is, replicas tell each other about
2497 the communications that they have received from other replicas, and the communications
2498 themselves are used to infer how the replicas would vote.

2499 In the Hashgraph algorithm, an *event* is a vertex in the hashgraph structure that is signed
2500 by its creator. When Bob gossips to Alice and sends her his current view of the hashgraph,
2501 an event is created and signed by Alice. This event contains the hashes of Alice's most
2502 recent event (called its *self-parent*), Bob's most recent event prior to the gossiping, and
2503 a timestamp. It can also include transactions that Alice proposes to the network. The
2504 algorithm occurs in rounds, and the first event that any given replica creates in a round is
2505 called a *witness*. Only witness events are considered in virtual voting. A *famous witness* is
2506 a witness that – according to the hashgraph – was received by a supermajority of replicas
2507 by the start of the next round. Byzantine agreement is sought over the witnesses rather than
2508 all events, and transactions are ordered based on whether witnesses are famous or not.

2509 A few more terms must be defined before presenting the algorithm. An event x is an
2510 *ancestor* of event y when x is y or when x can be reached from y through any sequence of
2511 parent relationships. Further, x is a *self-ancestor* of y if x is y or when x can be reached
2512 from y by a sequence of self-parent relations. A pair of events (x, y) – is a *fork* if x and
2513 y were created by the same replica, but neither is a self-ancestor of the other. Forks reflect
2514 inconsistencies in the hashgraph. An event x can *see* event y when y is an ancestor of x ,
2515 and the ancestors of x do not include a fork by the creator of y . Finally, an event x *strongly*
2516 *sees* event y when x can see y , there exists a set of events by more than $\frac{2n}{3}$ replicas such that
2517 x can see every event in the set, and every event in the set can see y . Figure 14 shows an
2518 example of strongly seeing.

2519 The security of Hashgraph can be argued based on quorum intersection. In particular, if
2520 a pair of events (x, y) is a fork, and x is strongly seen by event z in hashgraph A , then y
2521 will not be strongly seen by any event in any hashgraph B that is consistent with A . The
2522 Hashgraph protocol itself involves replicas running two loops in parallel:

- 2523 1. Pick a random replica, gossip to them all known events, and create another event to
2524 record this gossip.
- 2525 2. Receive gossip from another replica, create a new event, and then call three subpro-




Fig. 14. Hashgraph strongly seeing. Time advances from the bottom of the hashgraphs to the top. The yellow event at the top of each hashgraph can strongly see an orange event on the bottom row. In each hashgraph, the orange event is an ancestor of four (or five) red events by different block creators, each of which is an ancestor of the yellow event. If the four orange events and both parents of the yellow event were created in round r , then the yellow event is created in round $r + 1$ because it strongly sees more than $\frac{2n}{3}$ witnesses created by different replicas in round r . [103]

2526 cedures: *divideRounds()*, *decideFame()*, and *findOrder()*.

- 2527 • *divideRounds()*: For every event x , the replica sets r as the maximum round
 2528 of the parents of x . Then, if x can strongly see more than $\frac{2n}{3}$ of the witnesses
 2529 from round r , the replica sets x 's round to be $r + 1$. Otherwise, it is round r .
 2530 Regardless, the replica then determines whether x is a witness by checking if x
 2531 has no self-parent or if x 's round is greater than the round of x 's self-parent.

- 2532 • *decideFame()*: This is the step where Byzantine agreement occurs. For each
 2533 witness that has been identified, the replica checks whether it is a famous wit-
 2534 ness. Given a witness x in round r , each witness from round $r + 1$ votes (im-
 2535 plicitly) that x is famous if it can see it. If more than $\frac{2n}{3}$ agree that x is a famous
 2536 witness, agreement has been found and the election is over. This continues for
 2537 as many rounds as necessary to reach agreement. In normal rounds, the witness
 2538 votes in line with the majority of the witnesses it strongly sees from the prior
 2539 round. There are also periodic rounds where the vote is determined via a coin
 2540 flip in order to get around FLP impossibility when networks are asynchronous.

- 2541 • *findOrder()*: At this point, all witnesses from round r have had their fame de-
 2542 cided. The replica removes any famous witness with the same creator as any
 2543 other in that set. The unique famous witnesses that remain are used to totally
 2544 order events. An event has a "received round" of r if that is the first round in
 2545 which every unique famous witness is a descendent of the event, and the fame
 2546 of each witness was decided by round r . The timestamp of the event is taken
 2547 as the median of the timestamps of the events where each replica first received
 2548 it. Events are then sorted by received round with ties broken by the median
 2549 timestamp and further ties broken based on the signature on the event XORed

2550 with the signatures of the unique famous witnesses in the same received round.

2551 One of the alleged advantages of the Hashgraph protocol is "fairness" in transaction order-
2552 ing (see Section 7.1). In particular, it is claimed that it would be challenging for an attacker
2553 to "manipulate which of two transactions will be chosen to be first in the consensus order"
2554 [103]. Unfortunately, it fails to achieve this, as there is a method that allows an adversary to
2555 determine the order of transactions by exploiting the use of the median of the timestamps
2556 [104].

2557 Other asynchronous DAG-based protocols that interpret a communication graph in this
2558 way exist, including Aleph [105] and DAG-Rider [106]. Aleph improves upon Hashgraph
2559 by using a more efficient binary agreement subprotocol than the *decideFame()* procedure
2560 above [105]. DAG-rider improves upon both Hashgraph and Aleph by removing the need
2561 for signatures to maintain safety and by providing *eventual fairness*, such that all block
2562 proposals by honest replicas are guaranteed to eventually be included in the ledger [106].

2563 7. Miscellaneous Permitted BFT

2564 This section demonstrates some additional properties attainable by or techniques useful
2565 for BFT algorithms. These properties go beyond agreement and liveness to provide extra
2566 guarantees or functionality. For example, protocols can prevent replicas from manipulating
2567 the order of transactions, detect and expel replicas that behave maliciously, or dynamically
2568 reconfigure the set of replicas participating in consensus. Other protocols give certain repli-
2569 cas special roles to enhance performance, allow clients to have flexibility with respect to
2570 beliefs about network assumptions or the number of faults, or favor availability over con-
2571 sistency during network partitions (while maintaining deterministic finality). The details of
2572 the protocols here are omitted.

2573 7.1. Fairly Ordering Transactions

2574 There are situations in which the ability to choose the order of transactions included in the
2575 ledger provides an unwarranted or undesirable advantage to a replica. A major problem in
2576 traditional finance is *front-running*, where a participant uses the knowledge of somebody
2577 else's transactions for selfish advantage in their own dealings. For example, Bob may
2578 submit an order to his broker Alice to purchase 1000 shares of stock in company A. Before
2579 Alice fills the order, she can submit her own order to purchase 1000 shares and immediately
2580 turn around and sell the 1000 shares to Bob at a higher price for a guaranteed profit. For
2581 many distributed applications, the privileged position of block producers that choose the
2582 order of transactions makes them a potential adversary. Some consensus protocols attempt
2583 to reduce the ability of replicas to perform this and other kinds of manipulation by reducing
2584 the power of replicas to order transactions [104, 107–109].

2585 An early attempt to achieve fairness in transaction ordering was the Helix protocol [108],
2586 which used threshold cryptography to limit the ability of replicas to censor transactions

2587 proposed by each other. Threshold decryption is used to create a randomness beacon that
2588 is then used to elect the next PBFT-esque committee in an unpredictable way and to force
2589 replicas to randomly sample the choice of transactions they include. Because clients en-
2590 crypt transactions before sending them to replicas, the replicas lack information that would
2591 be useful to exploit ordering. However, this allows clients to arbitrarily spam the network
2592 with invalid transactions, so it may require identity verification for clients.

2593 More recently, the Aequitas family of protocols [104] attempted to provide a property
2594 dubbed *order fairness*, such that if sufficiently many replicas receive transaction tx_1 be-
2595 fore a different transaction tx_2 , then all honest replicas will output tx_1 before tx_2 . While
2596 it was proven that this property is impossible *within* blocks, the Aequitas protocols can
2597 provide this property *between* blocks. That is, if enough replicas receive tx_1 before tx_2 ,
2598 then tx_2 will appear in either the same block as tx_1 or a later block. The protocols rely
2599 on a primitive – FIFO-broadcast – that guarantees that broadcasts are delivered by honest
2600 replicas in the same order in which they were originally broadcast. Replicas use FIFO-
2601 broadcast to gossip their local transaction ordering and then later come to agreement on the
2602 set of replicas whose local orderings should be considered for a given transaction before
2603 ultimately finalizing the global ordering based on these local orderings.

2604 A variety of similar fairness-related definitions can be found in [107], which presents a
2605 set of protocol extensions or "widgets" that provide fairness. They can be added as a
2606 preprocessing step to blockchain protocols at the expense of increased latency. As with
2607 Aequitas, true fairness is impossible because it would require blocks to be of potentially
2608 unlimited size (this is a trade-off that the Aequitas protocols accept). A variety of more
2609 relaxed notions are possible, including that fairness be provided with a fixed probability or
2610 that if all honest parties saw a transaction tx_1 by time T and another transaction tx_2 after
2611 time T , then tx_1 will be committed before tx_2 .

2612 Another proposal, Pomp \bar{e} , decouples the transaction ordering process from the agreement
2613 process in order to obtain an alternative version of fairness [109]. The replicas involved in
2614 consensus express their preferences regarding the ordering of transactions, and given these
2615 sets of preferences, some possible total orderings are not considered valid. Specifically,
2616 if replicas base their ordering preferences on the time they first see a transaction, then the
2617 protocol guarantees the following: if the lowest timestamp that any honest replica assigns
2618 to tx_2 is higher than the highest timestamp that any honest replica assigns to tx_1 , then tx_1
2619 will precede tx_2 in the ledger. Unlike the order fairness property from [104], this ordering
2620 property is expressed based on the preferences of honest replicas rather than some fraction
2621 of all of the replicas, which may include Byzantine ones. This restriction results in pro-
2622 tocols with improved fault tolerance compared to the Aequitas protocols, which require at
2623 least $4f + 1$ replicas in order to tolerate f failures.

2624 In addition to providing some notion of fairness in transaction ordering, other possible no-
2625 tions of fairness can exist in permissioned systems. For instance, protocols with leaders
2626 tend to have an uneven distribution of effort for replicas since the leader does a dispropor-

2627 tionate amount of work. Rotating leaders or using leaderless protocols can help. Along
2628 these lines, [110] proposes a protocol designed to fairly balance the processing load based
2629 on past performance. It is similar to PBFT but with multiple simultaneous leaders. This
2630 system partitions client transactions among replicas according to each replica's process-
2631 ing ability, which prevents replicas from facing "unfairly" large resource burdens. Another
2632 possible notion of fairness arises if the protocol provides rewards to replicas, such as trans-
2633 action fees. For example, FairLedger is a protocol that – assuming that otherwise-honest
2634 replicas are rational – ensures that each replica receives fair shares of the fees [111]. Along
2635 similar lines, [112] proposes the notion of *strongly fair validity*: if n replicas are involved in
2636 an instance of Byzantine agreement, then the probability that a particular replica's proposal
2637 is accepted by the honest replicas is lower-bounded by $\frac{1}{n} - \epsilon$, where ϵ is negligible. *Weakly*
2638 *fair validity* captures the same idea but over repeated blocks, whereas strongly fair validity
2639 is with respect to a single-shot BA instance.

2640 7.2. Accountability Against Malicious Replicas

2641 Another useful property – and one that is occasionally used to make proof-of-stake proto-
2642 cols more robust – is accountability. In an *accountable* BA protocol, honest nodes that are
2643 not in agreement can exchange sufficient information to provably identify at least $\frac{n}{3}$ mali-
2644 cious nodes if $f \geq \frac{n}{3}$. This property relies on the idea that malicious equivocation can be
2645 detected due to the existence of two signatures from the same key on conflicting messages.
2646 When misbehavior can be detected by honest parties and individual responsibility can be
2647 assigned, it may be possible to relax some security assumptions while maintaining system
2648 security [113]. Accountability can also be used to help make BFT protocols more secure
2649 when some participants are rational rather than altruistic [114].

2650 A protocol that provides this is Polygraph, which increases the communication complexity
2651 of the BA protocol it builds off of by a linear factor [115]. Polygraph relies on the leaderless
2652 Democratic BFT protocol proposed in [116] and used in the Red Belly Blockchain project
2653 [117]. It is also used in the Long-Lasting Blockchain (LLB) protocol in order to tolerate
2654 considerably more than the typical f Byzantine failures [118]. When there are more than
2655 f Byzantine replicas, a fork can be created and detected via conflicting signed messages.
2656 When LLB detects a fork, there is a recovery procedure that merges the conflict instead of
2657 discarding one of the conflicting blocks.

2658 Reference [119] shows how to provide an alternative notion of accountability in a variety
2659 of BFT algorithms, including PBFT and HotStuff. Specifically, if $f < \frac{n}{3}$ is the maximum
2660 number of Byzantine faults that the protocol tolerates but $t > f$ Byzantine faults occur, then
2661 this notion of accountability guarantees that the protocol can detect at least $f + 1$ Byzantine
2662 faults, where $2f + 1 - t$ honest replicas can testify to that effect, and proof of malfeasance
2663 only requires communicating with one of those honest replicas.

2664 7.3. Specially Designated Roles for Replicas

2665 Instead of requiring every replica to perform the same actions throughout a protocol ex-
2666 ecution, some protocols may assign special roles to some replicas in order to enhance
2667 performance. A simple example of this was seen in the HotStuff protocol (see Section 5.3),
2668 where replicas sent their signature shares directly to the leader for aggregation instead of
2669 multicasting them to every replica. The use of threshold signatures in this way is fairly
2670 common because it can reduce the communication complexity from quadratic to linear.

2671 This technique is also employed in SBFT, where the aggregating party is called a "collector"
2672 [120]. In addition, SBFT uses the optimistic fast path from Zyzzyva, where the client
2673 can be viewed as the collector (see Section 4.3), but makes it more resilient by adding
2674 redundancy such that more than c faulty replicas are required to leave the fast path. A
2675 trade-off is that including this redundancy requires more replicas to maintain the same
2676 degree of fault tolerance ($n \geq 3f + 2c + 1$). SBFT recommends $c \leq \frac{f}{8}$ and uses $c + 1$
2677 collectors who rotate in round-robin fashion. SBFT further reduces client communication
2678 via threshold signatures. Instead of waiting for $f + 1$ replies from replicas, an "execution
2679 collector" gathers the replies into a single signature over the result to send to the client.

2680 Another technique, employed in the Proteus protocol, is to elect a subset of c replicas with
2681 $c \ll n$ as a "root committee," which is responsible for executing a BFT algorithm among
2682 themselves [121]. The block proposed by the committee is then validated by the remainder
2683 of the replicas, and if valid, the replicas sign it and return it to the root committee. When
2684 the root committee sees $2f + 1$ signatures, they commit the block and send the signatures
2685 to the remainder of the replicas who will then commit as well. Proteus ensures stable per-
2686 formance regardless of the number of failures. View changes do not just change the leader
2687 but rather the full committee of c replicas. Compared to the typical $O(n^2)$ communication
2688 complexity of PBFT (and $O(n^4)$ for view changes), Proteus has a complexity of $O(c^2 + cn)$
2689 for normal and view change modes. The root committee tolerates up to $\frac{2c}{3}$ failures because
2690 the remainder of the replicas initiate a view change of the committee if the committee fails
2691 to propose a block.

2692 7.4. Deterministic Longest Chain Protocols

2693 Another style of BFT algorithm is the deterministic, longest-chain protocol, such as the
2694 "proof of authority" (PoA) algorithms Aura and Clique, which have been used for permis-
2695 sioned deployments of Ethereum. This kind of protocol favors availability over consistency
2696 during network partitions, not unlike Bitcoin and other "longest chain" protocols (see Sec-
2697 tion 10 on Nakamoto Consensus) [122]. This means that the chain can fork temporarily.
2698 Unlike with Bitcoin, the protocols mentioned here are deterministic, so the schedule of
2699 block proposers is known in advance. In other BFT protocols, leader election and voting
2700 are separate processes, but deterministic longest chain protocols combine them. When a
2701 replica is elected leader, it has the opportunity to vote, and honest replicas vote on the most
2702 "popular" ledger seen so far.

2703 Aura assumes that replicas have synchronized clocks, and the next leader is elected at reg-
2704 ular intervals. In Aura, the leader broadcasts a signed block to the other replicas, the leader
2705 is rotated in round-robin style, and future leaders build on the longest chain that they have
2706 seen. A block is considered final after a sufficient number of the replicas sign blocks on
2707 the same chain extending it. This uses only two rounds of communication as opposed to
2708 PBFT's three. Clique allows multiple simultaneous block proposers, resolves forks using
2709 the GHOST protocol (see Section 11.2), and only requires a single round of communication
2710 to commit. Early variants of both of these algorithms were subject to various attacks, in-
2711 cluding the "cloning" attack that allowed double-spending during network partitions [123].
2712 Note that using GHOST as the fork-choice rule in a permissioned system may be dan-
2713 gerous due to the *balance attack*, where an adversary keeps the network partitioned and
2714 prevents transactions from being finalized by keeping the two partitioned blockchains at
2715 similar lengths for a period of time [124, 125].

2716 In [126], the security of a variant of Aura is proven when $n \geq 3f + 1$ in synchronous
2717 networks while simultaneously pointing out a vulnerability in an existing implementation.
2718 The attack prevents consistency by having the network constantly switch back and forth
2719 between equal-length chains. Prior to this analysis, it had been wrongly believed that the
2720 protocol was secure as long as the majority was honest. Randomized variants of longest
2721 chain protocols like Nakamoto Consensus have a better security margin (honest majority)
2722 and better confirmation times, but deterministic variants are secure against a more powerful
2723 adversary capable of "after-the-fact message removal," where "the adversary can observe
2724 what an honest node i wants to send in some round r , adaptively corrupt node i , erase the
2725 message it originally wanted to send, and then insert arbitrary corrupt messages on behalf
2726 of node i in round r " [126].

2727 Ouroboros-BFT is a deterministic longest chain protocol similar to Aura and secure when
2728 $n \geq 3f + 1$ [127]. It is synchronous and, thus, not optimally resilient for synchronous
2729 protocols (though it is optimally resilient for deterministic longest chain protocols). In the
2730 *covert* setting, where an adversary does not want to create evidence of their misbehavior
2731 (the accountability property discussed in Section 7.2 is relevant here), Ouroboros-BFT can
2732 tolerate $n \geq 2f + 1$. The clock advances in slots every few seconds. Each time the clock
2733 advances to a new slot, replicas receive some transactions and blockchain candidates from
2734 the network and 1) update their mempool with the new transactions, 2) update their local
2735 preferred blockchain via the longest chain rule, and 3) check whether they are the current
2736 round-robin slot leader. If so, they extend their longest chain with a new block and diffuse
2737 it to the other replicas. A block is finalized when it has a slot timestamp more than $3f + 1$
2738 slots in the past.

2739 Ouroboros-BFT has a few advantages and trade-offs compared to algorithms like PBFT and
2740 Zyzzyva. Ouroboros-BFT is simpler in that replicas perform the same steps in each slot
2741 independently of their view, whereas replicas in PBFT execute different steps depending
2742 on their current state. Synchrony assumptions differ, and PBFT maintains consistency with
2743 unbounded delays. When the network delay bound Δ is violated in Ouroboros-BFT, the

2744 chain forks, but replicas will regain a consistent view when the delay bound is respected
2745 again. Ouroboros-BFT provides speculative execution immediately, while the longest chain
2746 mechanism orders transactions over time, and communication complexity is optimistically
2747 $\Theta(n)$ per round and $\Theta(nf)$ in the worst case. Transactions have guaranteed liveness after
2748 $5f + 2$ rounds, but a speculative outcome is produced in only two rounds. Zyzzyva has the
2749 same optimistic performance but requires the client to be actively involved in consensus,
2750 unlike PBFT and Ouroboros-BFT.

2751 7.5. Flexible BFT

2752 Most BFT algorithms provide safety and liveness for all users under certain conditions,
2753 such as a synchronous network with the majority of replicas being honest. Flexible BFT
2754 separates the fault model from the protocol design itself, allowing flexible beliefs about
2755 the network and number of faults. A client may specify a fault threshold and a maximum
2756 message delay, and safety and liveness will be maintained for any client with correct be-
2757 liefs (and any two clients with correct beliefs will be in agreement with each other) [128].
2758 Replicas commit transactions the way other BFT algorithms do: only clients commit, and
2759 they may do so at different times. The flexibility provided here is akin to the recipient de-
2760 ciding how many confirmations to wait for in Bitcoin. A \$10 million transaction deserves
2761 a lengthier confirmation period than a \$10 one. Further, clients may update their beliefs
2762 based on observation, such that if they notice more votes on conflicting values or lengthier
2763 message delays, they may become more conservative.

2764 Flexible BFT also introduces a new *alive-but-corrupt* (a-b-c) fault model, where an ad-
2765 versary will try to violate the safety of the protocol but, failing that, will not try to disrupt
2766 liveness. This model may be justified when violating safety could reward an adversary with
2767 more money (e.g., via double-spending) but violating liveness may not (because keeping
2768 the service running may allow the adversary to collect transaction fees). This relaxation of
2769 the fault model allows for protocols that remain secure despite a combination of a-b-c and
2770 Byzantine faults in excess of $\frac{n}{3}$ under partial synchrony and $\frac{n}{2}$ under synchrony. To achieve
2771 these properties, Flexible BFT uses two key techniques:

- 2772 1. Replicas run a partially synchronous protocol, but clients assume synchrony bounds
2773 for committing, which allows flexible assumptions regarding Δ . That is, the timing
2774 assumptions for replicas and for clients are distinct.
- 2775 2. Replicas will use a particular quorum size while executing the protocol, but clients
2776 choose their own quorum sizes before committing. This allows clients to have diver-
2777 gent beliefs about the security threshold.

2778 The use of quorums in proving the security of BFT protocols is discussed in Section 4.2.
2779 *Quorum intersection* both within and across views can prove safety by demonstrating that
2780 at least one honest replica would have needed to misbehave to violate safety. *Quorum avail-*
2781 *ability* can be used to prove liveness by showing that a sufficiently large quorum contains

2782 no Byzantine replicas, and thus there are enough honest replicas to respond to an honest
2783 leader.

2784 Flexible BFT separates out the quorums needed for different parts of the algorithm: locking
2785 on a value or forming a certificate requires q_{lck} , while unlocking requires q_{ulck} . The quorum
2786 that clients need for learning certificate uniqueness is q_{unq} , and the quorum required for
2787 safely committing is q_{cmt} . Clients require a q_{unq} fraction of votes in the first round and a
2788 q_{cmt} fraction of votes in the second round to commit.

2789 Flexible BFT can ensure quorum intersection within a view by ensuring that every q_{lck}
2790 quorum intersects with every q_{unq} quorum at a minimum of one honest replica, which
2791 requires that the fraction of faulty replicas is less than $q_{lck} + q_{unq} - 1$. This guarantees
2792 that if a client commits a value, it is the only value with a certificate in this view. To
2793 ensure quorum intersection across views, every q_{ulck} quorum must intersect with every
2794 q_{cmt} quorum at a minimum of one honest replica, which requires that the fraction of faulty
2795 replicas to be less than $q_{ulck} + q_{cmt} - 1$. This guarantees that when a client commits a
2796 value, replicas that have locked on that value will not later unlock it. To ensure quorum
2797 availability within a view and thus liveness, the fraction of Byzantine replicas must be less
2798 than $1 - \max(q_{unq}, q_{cmt}, q_{lck}, q_{ulck})$.

2799 Balanced quorum sizes are optimal, such that $q_{lck} = q_{ulck} = q_r$ and $q_{unq} = q_{cmt} = q_c$. Ad-
2800 ditionally, $q_c \geq q_r$, because the q_r votes that replicas use to lock can also be used in the
2801 q_{cmt} quorums by clients. Here, q_r is the quorum that replicas need to lock a value (q_r is
2802 chosen by the system designer), and q_c is the quorum that clients need for safety (chosen by
2803 clients, along with Δ). Taken together, *flexible quorum intersection* holds when the fraction
2804 of faulty replicas is $< q_c + q_r - 1$, and *flexible quorum availability* holds when the fraction
2805 of Byzantine replicas is $\leq 1 - q_c$. For example, if $q_r = 0.7$ and $q_c = 0.75$, a client can
2806 maintain security when the fraction of Byzantine replicas is 0.25 and the fraction of a-b-c
2807 replicas is up to 0.2.

2808 Other protocols have similar goals to Flexible BFT in that they attempt to accommodate a
2809 wider variety of assumptions simultaneously. For example, the protocols in Section 8 can
2810 allow replicas with different trust assumptions to maintain agreement. The Heterogeneous
2811 Paxos protocol allows different clients to set their own mixed failure tolerances for crash
2812 and Byzantine faults using differently trusted sets of replicas, essentially combining some
2813 of the benefits of Flexible BFT and the protocols from Section 8 [129]. The Highway
2814 protocol allows flexible assumptions regarding the number and types of faulty replicas and
2815 then uses these assumptions to establish different confidence thresholds for the finality of
2816 blocks [130]. Other protocols provide similar flexibility to Flexible BFT, but rather than
2817 using PBFT as the basis for the algorithm, ones with better communication complexity like
2818 HotStuff and Streamlet are used instead [131].

2819 7.6. View Change Algorithms

2820 In partially synchronous permissioned consensus protocols with leaders, *view changes* are
2821 the mechanism used to maintain liveness despite a malicious leader. A *view* can be consid-
2822 ered a phase of the protocol where a particular replica acts as the leader and is generally
2823 represented as a monotonically increasing integer. A view change algorithm, sometimes
2824 called a *pacemaker* or *synchronizer*, has every replica start at view zero. The algorithm al-
2825 lows replicas to signal that they wish to advance to the next view but only actually advances
2826 the replica to a new view when the synchronizer emits a notification to the higher-level con-
2827 sensus protocol.

2828 The challenge of view synchronization is that during good periods, progress can be main-
2829 tained by the f Byzantine replicas combined with the lowest latency group of $f + 1$ honest
2830 replicas. In this case, it is possible that only the quickest $f + 1$ honest replicas recognize
2831 that a block has been agreed upon and advance to the next round of the consensus proto-
2832 col, while the f slowest-but-honest replicas fall behind. Later, if the f Byzantine replicas
2833 stop behaving as though they were honest, the view change mechanism is required in order
2834 to get the f slow replicas to "catch up" to the same view as the $f + 1$ quicker ones. The
2835 *Byzantine view synchronization problem* is solved by algorithms with the following two
2836 properties [132]:

- 2837 1. **View Synchronization:** All honest nodes must execute the same view with an honest
2838 leader for a sufficiently long time, and there must be an infinite number of views with
2839 an honest leader.
- 2840 2. **Synchronization Validity:** A synchronizer will only signal a new view if an honest
2841 node wished to advance to it.

2842 A naive approach to view synchronization is to simply double the duration of each view,
2843 which ensures that there will be a sufficiently long period where all honest nodes share
2844 the same view. This has the advantage of requiring no communication between parties but
2845 is clearly impractical due to the unbounded latency that it introduces into the system. A
2846 more practical approach is the broadcast-based view change mechanism used in PBFT (see
2847 Section 4.1), which has high communication complexity but $O(\Delta)$ expected latency and
2848 $O(f\Delta)$ worst-case latency. Using this approach, when a replica sees messages suggesting
2849 that $f + 1$ replicas wish to enter the same view, it relays its own wish to advance to that
2850 view using reliable broadcast.

2851 These algorithms leave room for improvement, and several works have proposed new view
2852 synchronizers with better performance [132–134]. For example, the Cogsworth synchron-
2853 izer matches the latency and communication complexity of the broadcast variant but, in
2854 some scenarios, can improve the communication complexity by a linear factor [132]. In-
2855 stead of having every replica multicast messages to every other replica, they send a message
2856 directly to the leader of the view to which the replica wishes to advance. An honest leader
2857 then broadcasts a single message with an aggregated threshold signature with only linear

2858 communication complexity. If the leader of a view is Byzantine instead, then replicas may
2859 need to time out and attempt each of the $f + 1$ subsequent leaders until they find an honest
2860 leader who will help relay the aggregated signature.

2861 The Cogsworth protocol involves four types of messages, including two sent from replicas
2862 to leaders and another two sent from leaders to the rest of the replicas. Replicas send to
2863 prospective leaders ($WISH, v$) and ($VOTE, v$) messages when they want to move to the
2864 next step of the protocol, where v is the view number in question. The two that leaders
2865 send are "time certificate" and "quorum certificate" messages, each with one aggregated
2866 threshold signature: the leader sends (TC, v) upon receiving $f + 1$ ($WISH, v$) messages
2867 or sends (QC, v) upon receiving $2f + 1$ ($VOTE, v$) messages. When a replica wishes to
2868 advance to another view v , it sends ($WISH, v$) to the leader of view v . If the leader of
2869 view v receives $f + 1$ of these messages, they broadcast (TC, v). However, if 2Δ time
2870 passes after the replica sends its first ($WISH, v$) without receiving a corresponding (TC, v)
2871 message, it sends ($WISH, v$) to the leader of view $v + 1$. The replica continues to wait in
2872 2Δ increments, sending ($WISH, v$) to the leader of view $v + 2$ until it eventually receives a
2873 (TC, v) message. Upon receiving (TC, v), replicas send ($VOTE, v$) to the leader of view v ,
2874 even if they had not already sent a ($WISH, v$) message. The replica waits 2Δ time to receive
2875 a (QC, v) message and contacts subsequent leaders as above if time runs out. In this case,
2876 the replica sends leaders both ($VOTE, v$) and (TC, v). Replicas enter view v immediately
2877 upon receiving (QC, v) from a leader.

2878 A follow-up work proposed an alternative view change algorithm built off of Cogsworth
2879 [133]. One issue with Cogsworth is that its expected linear communication complexity
2880 with benign failures becomes quadratic with Byzantine failures in the average case. The
2881 algorithm in [133] maintains an expected linear message complexity, even with Byzantine
2882 failures. To achieve this, the algorithm modifies Cogsworth by adding another phase to the
2883 algorithm and having replicas' signed messages to leaders include the identity of the leader
2884 it is intended for.

2885 One problem with both of the above algorithms is that they rely on a modified variant of par-
2886 tial synchrony where there is no clock drift, and messages sent before GST cannot be lost
2887 and will necessarily arrive by time $GST + \Delta$. This is potentially problematic because prior
2888 to GST, if clocks diverge, then the replicas' notion of the duration of a view can become
2889 desynchronized. Further, messages that may be needed in order to get replicas to be in the
2890 same view can be lost. Worse, the above algorithms require potentially unbounded memory
2891 space in order to store each ($WISH, v$) message that falls below the required threshold for
2892 relaying, as well as maintaining a copy of every message it sends in order to enable it to be
2893 sent again. The FastSync synchronizer was designed to work under true partial synchrony
2894 with unbounded clock drift and the possibility of message loss before GST while running in
2895 bounded space [134]. FastSync works almost exactly like Bracha's broadcast (see Section
2896 1.2), but replicas are not constrained to only acting on identical messages. Instead, replicas
2897 maintain only the highest ($WISH, v$) message received from each process and can act on
2898 sets of $WISH$ messages for non-identical views.

2899 One of the risks of the leader-based view-by-view progression of partially synchronous
2900 protocols is that leaders can be targeted in adaptive denial-of-service attacks in order to in-
2901 duce asynchrony and eliminate liveness. A leader-based view abstraction was designed in
2902 [135], which takes any view-by-view consensus protocol, wraps it in an API, and builds an
2903 asynchronous SMR system from it. This reduces performance in optimistic cases while dra-
2904 matically improving liveness while under attack or while poor networking conditions hold.
2905 Separately, an asynchronous view change algorithm has been proposed as a replacement
2906 for the synchronizers discussed above. When combined with HotStuff, the resulting SMR
2907 system has linear communication complexity while the network is synchronous, quadratic
2908 complexity when the network is asynchronous, and always maintains liveness [136].

2909 Strongly related to the idea of view changes is the ability to dynamically reconfigure con-
2910 sensus committees, such that replicas may join or leave the system over time. A variety of
2911 proposals exist that can provide such a functionality [137–142]. Reconfiguration can even
2912 be bundled with accountability (described in Section 7.2) in order to immediately change
2913 the committee upon detecting Byzantine behavior [142]. A system that is run by a sin-
2914 gular organization may not need to dynamically reconfigure the consensus committee, but
2915 when the replicated state machine is operated by a consortium of different organizations or
2916 individuals, the ability to add new members and remove others becomes more important.

2917 To enable reconfiguration, the initial consortium members and their public keys are stored
2918 in the system’s genesis block. In the scheme presented in [137], transactions that recon-
2919 figure the committee are included in separate *reconfiguration blocks*. Each normal block
2920 (the ones that execute client-issued transactions) includes a pointer to the previous recon-
2921 figuration block in the chain, so verifiers have access to the public keys needed to validate
2922 the signatures over proposed blocks. The primary security challenge while dealing with re-
2923 configuration is preventing faulty replicas active in previous views from causing problems
2924 in later ones, as shown in Figure 15. The attack is analogous to long-range attacks against
2925 proof of stake, which is discussed in Section 12.1.2.

2926 One way to resolve this problem is to separate a replica’s “permanent” key pair from their
2927 “consensus” key pair, which is what is used for block signing. For each new configura-
2928 tion that a replica is involved in, it generates a new consensus key pair and discards the
2929 prior one (discarding the old key pair can be done using the “forgetting” scheme described
2930 in [143]). Securely discarding the old consensus key pair ensures that the faulty replica
2931 cannot recover the key and sign blocks corresponding to an old configuration should they
2932 become corrupted after a reconfiguration (as is done by replicas *B* and *C* in Figure 15). Re-
2933 configuration blocks store the identities of the replicas and each of their consensus public
2934 keys in the new configuration.

2935 For a new replica to join the system, it first asks the current replicas for permission. The
2936 current replicas choose to accept or reject the new replica, and if they accept, they send
2937 a signed reply message that includes the replica’s intended consensus public key for the
2938 next configuration. Upon receiving $n - f$ signed acceptance messages, the prospective new

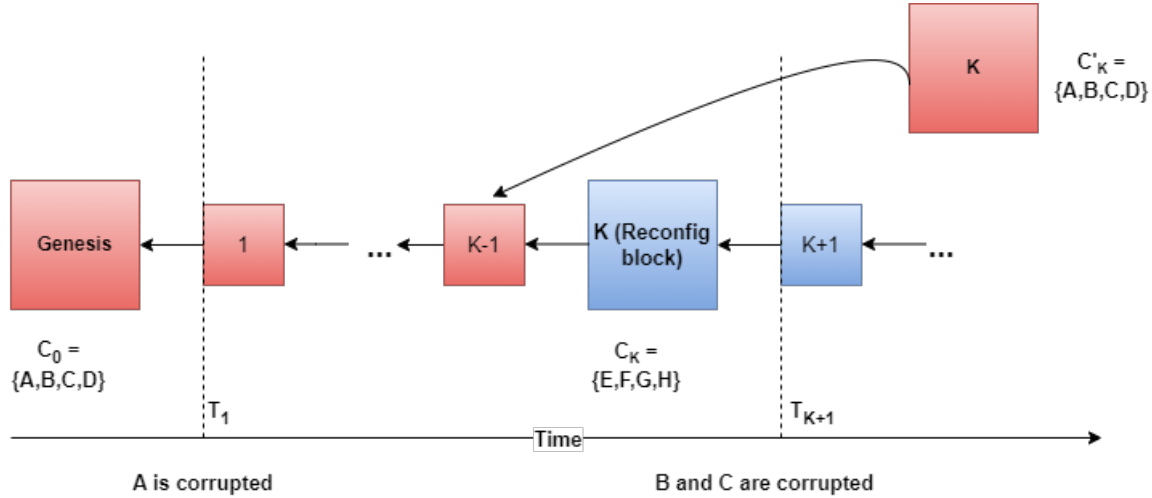


Fig. 15. Committee reconfiguration attack. While the threshold of faulty replicas is respected in each view, replicas A , B , and C work together to create a fork by extending the blockchain without the reconfiguration block at height K . This is possible even though none of the original replicas remain in the committee. If a client, aware of the initial configuration C_0 , is offline during reconfiguration and then reconnects, they will contact the committee members from C_0 . If this happens at any time after T_{K+1} , the corrupted quorum will convince the client to follow the wrong chain starting from height K .

2939 replica creates a reconfiguration transaction that bundles these signatures together. When
 2940 this transaction is included in a reconfiguration block and executed, the new replica can
 2941 begin participating. Replicas can leave the system by themselves or by being kicked out by
 2942 the rest of the replicas. If leaving by choice, the exiting replica gathers consensus public
 2943 keys from existing replicas for a new configuration, bundles them together, and issues a
 2944 reconfiguration transaction. If the replica is being kicked out, then existing replicas issue
 2945 special reconfiguration transactions that ask the network to remove the target replica and
 2946 provide a new consensus public key for the following configuration. Upon seeing $n - f$
 2947 of these transactions that remove the same replica, a new reconfiguration is generated without
 2948 the replica.

2949 8. Localizing Trust Over Incomplete Networks With Open Membership

2950 The consensus protocols described in Sections 4 through 7 require that all n replicas in
 2951 the system be fully aware of each others' identities. Further, they must all trust each other
 2952 equally despite real-world relationships that may have varying levels of trust. Prior works
 2953 have explored how to loosen restrictions on the knowledge of other replicas when net-
 2954 works are not fully connected [144, 145]. However, these BFT-CUP (Consensus with Un-
 2955 known Participants) protocols still require that the unknown replicas be equally trusted as
 2956 as the known ones; that a fixed set of consensus replicas exist, each with a unique and

2957 Sybil-resistant ID; and that all replicas are aware of the maximum failure threshold f . It is
2958 also possible to design permissioned agreement algorithms that are unaware of the precise
2959 values of n and f but maintain optimal resiliency [146]. Another line of work introduces
2960 asymmetric trust assumptions, where the replicas in the system do not adhere to a sin-
2961 gle global trust assumption [147, 148]. In these systems, each replica can choose which
2962 combinations of other processes it trusts and which may be considered faulty.

2963 The protocols described later in this section combine some of the benefits of these sys-
2964 tems. The key feature of these Federated Byzantine Agreement Systems (FBAS), some-
2965 times called Federated Byzantine Quorum Systems (FBQS), is that each replica chooses
2966 its own group or quorum of trusted replicas to believe without needing to be aware of the
2967 existence of all other replicas or trusting them to the same degree. These systems represent
2968 a middle ground between permissioned and permissionless, though in practice, they are far
2969 closer to permissioned networks in both operation and trust levels. Section 4.2 contains
2970 some background information on Byzantine quorums.

2971 8.1. Stellar

2972 The Stellar Consensus Protocol (SCP) was introduced by Mazieres in [149] and further
2973 described by the Stellar Development Foundation in [150]. A generalization of Stellar’s
2974 quorum system with security proofs is provided in [151].

2975 8.1.1. FBAS Background

2976 Stellar and similar protocols allow each user of the system to unilaterally define their own
2977 sets of trusted replicas, called *quorum slices*. A single organization may have multiple
2978 slices, any one of which suffices to convince them of a statement in case the others fail. A
2979 replica should choose their quorum slices such that they believe that

- 2980 • If every member of a slice agrees about the state of the system, then they are correct
2981 about the state, and
- 2982 • At least one of its slices will be available to provide information about the state of
2983 the system in a timely manner.

2984 Let S be the set of replicas from which a set of messages originated. Every replica specifies
2985 its quorum slices in every message it sends, and it considers the set of messages to have
2986 reached the quorum threshold when every member of S has a slice included in S . A quorum
2987 is a non-empty set S of replicas that includes at least one quorum slice of each non-faulty
2988 member. If S is unanimous, the agreement requirements for all of its members are satisfied.

2989 Two honest replicas v_1 and v_2 are considered *intertwined* when every quorum of v_1 inter-
2990 sects every quorum of v_2 in at least one honest replica. In an FBAS, agreement is only
2991 ensured between intertwined replicas. A set of replicas I is *intact* if I is a quorum whose
2992 replicas are uniformly honest, such that every pair of members of I is intertwined even

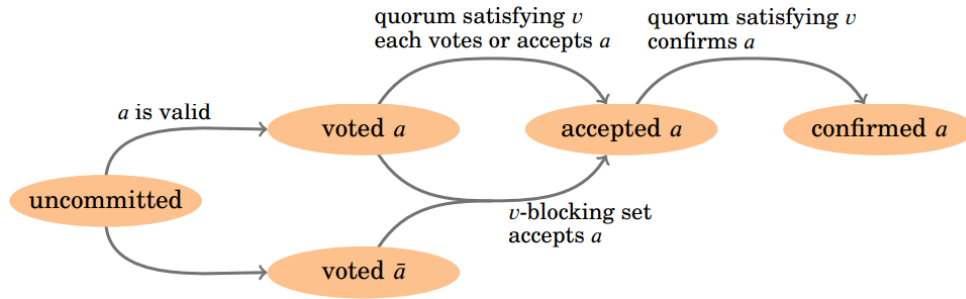


Fig. 16. Federated voting stages. [149]

2993 if every replica outside of I is faulty. An intact set I cannot be harmed by the actions of
 2994 non-intact replicas, and the union of two intact sets that intersect is an intact set. Stated
 2995 differently, intact sets are partitions of the honest replicas, where each partition maintains
 2996 safety and liveness under certain conditions but where different partitions may have diver-
 2997 gent outputs.

2998 A v -blocking set is a set of replicas that intersects every quorum slice of v . A v -blocking
 2999 set B of faulty replicas can block progress by replica v because liveness requires that v has
 3000 at least one quorum slice comprised solely of honest replicas. If B unanimously votes for
 3001 a value x , then v knows that either x is true or v is not intact. A full quorum is required in
 3002 order for v to know that x will not be contradicted by intertwined replicas. This requirement
 3003 adds an additional round of communication to Federated Byzantine Agreement protocols
 3004 compared to standard BFT protocols. Ultimately, there are three levels of confidence that
 3005 a replica can have regarding consensus on a particular value: uncommitted, *accepted* (i.e.,
 3006 safe to assume among intact replicas), and *confirmed* (i.e., safe to assume among inter-
 3007 twined replicas). If replicas broadcast the values that they accept and a full quorum accepts
 3008 a value, these values will propagate through intact sets due to the *cascade theorem*:

3009 If I is an intact set, Q is a quorum of any member of I , and S is any superset
 3010 of Q , then either $S \supseteq I$ or there is a member $v \in I$ such that $v \notin S$ and $I \cap S$ is
 3011 v -blocking. [150]

3012 8.1.2. Stellar Consensus Protocol (SCP)

3013 The Stellar Consensus Protocol is a partially synchronous protocol, where each single-shot
 3014 attempt at achieving consensus on a value is called a *ballot*, and ballots employ increasing
 3015 timeouts in order to synchronize replicas on the same ballot. Ballots are typically called
 3016 views in other protocols. For each ballot n , a process called *federated voting* occurs on
 3017 both PREPARE and COMMIT statements for a value x . Federated voting is the main
 3018 subprotocol that provides agreement, and its stages can be seen in Figure 16.

3019 A replica v will vote for any valid statement x that does not contradict its other outstanding
 3020 votes and accepted statements by broadcasting a signed vote message. It then accepts x

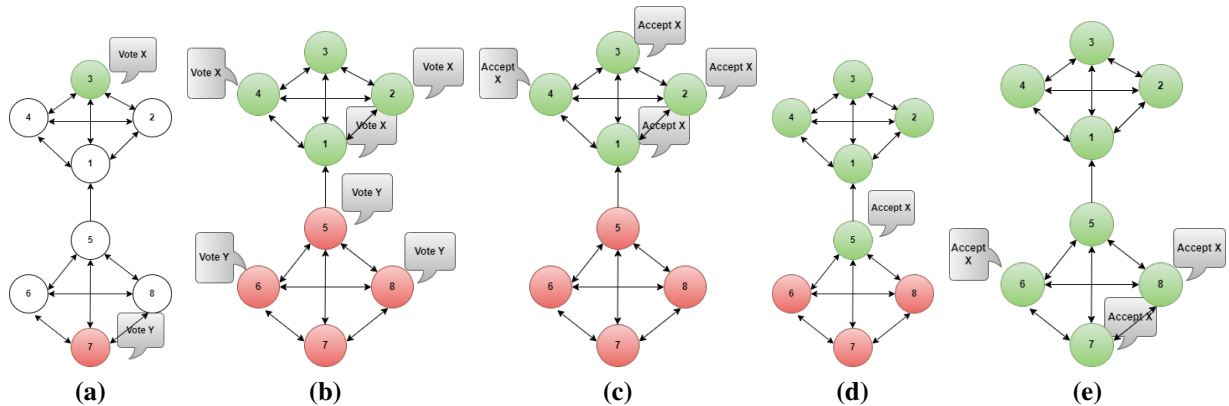


Fig. 17. Cascade effect in federated voting. Each replica has a single quorum slice indicated by arrows to members of the slice. All trust relationships are bidirectional with the exception of replica 5 including replica 1 in their quorum slice, but replica 1 does not reciprocate. (a) Contradictory statements X and Y are introduced. (b) Replicas vote for valid statements. (c) Replicas 1, 2, 3, and 4 accept X after their quorum $\{1, 2, 3, 4\}$ unanimously votes for X . (d) Set $\{1\}$ is 5-blocking, so replica 5 accepts X and overrules its previous vote for Y . (e) Set $\{5\}$ is 6-, 7-, and 8-blocking, so replicas 6, 7, and 8 each accept X and overrule their previous votes for Y .

3021 if x is consistent with other accepted statements and either 1) v is a member of a quorum
 3022 where each replica either votes for x or accepts x , or 2) a v -blocking set accepts x even
 3023 if v did not vote for x . In this case, if v had previously cast votes that contradicted x ,
 3024 those votes are overruled and forgotten. Replica v then broadcasts an accept message for
 3025 x and confirms x when it is in a quorum that unanimously accepts x . Figure 17 shows an
 3026 example execution of federated voting that demonstrates the cascading effect described by
 3027 the cascade theorem.

3028 Replicas will begin ballot n by initiating federated voting on a $\text{PREPARE}(n, x)$ message. If
 3029 a previous PREPARE message was successfully confirmed via federated voting, the replica
 3030 will choose x as the confirmed PREPARE of the highest ballot. Otherwise, it sets x as the
 3031 output of a *nomination* subprotocol. A replica attempts federated voting on $\text{COMMIT}(n, x)$
 3032 if and only if the replica successfully confirms a $\text{PREPARE}(n, x)$ in ballot n . If this suc-
 3033 ceeds, then the message is committed. However, if the replicas are unable to confirm a
 3034 COMMIT message in the current ballot, they will employ increasing timeouts in order to
 3035 synchronize on a particular ballot.

3036 The nomination subprotocol involves federated voting over messages of the form
 3037 $\text{NOMINATE}(x)$. Once a replica confirms a single NOMINATE message, it stops voting
 3038 to nominate new values, so the set of nominated values is finite. This process creates an
 3039 evolving, deterministic combination of all values in confirmed NOMINATE messages. In-
 3040 tact replicas will eventually converge on the same set of values at some point arbitrarily late
 3041 in the protocol execution. To reduce the number of values nominated, only a leader who




Fig. 18. FBAS Quorum intersection insufficient for safety. Quorums $\{v_1, v_2, v_3, v_7\}$, $\{v_4, v_5, v_6, v_7\}$, and $\{v_7\}$ intersect at v_7 , but because v_7 is Byzantine, agreement is not guaranteed. [149]

3042 has not already voted for a NOMINATE message may introduce a new x value. To tolerate
 3043 failures, the set of leaders grows as timeouts occur.

3044 The leader election mechanism employed in Stellar uses two hash functions, H_0 and H_1 ,
 3045 where $H_i(m) = SHA256(i || b || r || m)$, b is the block number, and r is the leader election round
 3046 number. Define $priority(v) = H_1(v)$. Define $weight(u, v) \in [0, 1]$ as the fraction of replica
 3047 u 's quorum slices containing replica v . When u is selecting a new leader, it only consider
 3048 its neighbors: $neighbors(u) = \{v \mid H_0(v) < 2^{256} * weight(u, v)\}$. Replica u starts with an
 3049 empty set of leaders, and it adds the replica v in $neighbors(u)$ with the highest $priority(v)$
 3050 at each round. When the set is empty for a round, replica u instead adds the replica v with
 3051 the lowest value of $\frac{H_0(v)}{weight(u, v)}$.

3052 The normal-case operation of Stellar involves six messages: two to vote and accept a NOM-
 3053 INATE, two to accept and confirm a PREPARE, and two to accept and confirm a COMMIT.

3054 8.1.3. SCP Security

3055 Quorum intersection is a necessary but insufficient condition for the safety of the SCP pro-
 3056 tocol. This can be seen in Figure 18, where the replica shared across quorums is malicious.
 3057 Instead, SCP requires that quorum intersection continues to hold, even after deleting all of
 3058 the faulty replicas from the trust graph.

3059 The federated voting procedure of SCP guarantees that no two members of an intertwined
 3060 set confirm contradictory messages because the two quorums would share an honest replica
 3061 that would not accept contradictory messages. However, a split vote can result in a state-
 3062 ment becoming permanently stuck waiting for a quorum and not being confirmed. An
 3063 intact set would not become stuck if a replica in it confirms a statement because replicas
 3064 will vote on a value when a v -blocking set accepts it, and the cascade theorem ensures that
 3065 this eventually spreads to the remainder of the intact set. Once the members of the intact
 3066 set vote unanimously, confirmation is guaranteed.

3067 To see why SCP itself maintains safety, consider an intertwined set S . The safety of fed-
 3068 erated voting ensures that, for a given ballot, a maximum of one value can be confirmed

3069 prepared by members of S . This implies that at most one value can be confirmed com-
3070 mitted in a given ballot as well. Now assume that $\text{COMMIT}(n, x)$ is confirmed. This
3071 implies that $\text{PREPARE}(n, x)$ was also confirmed. $\text{PREPARE}(n, x)$ would contradict any
3072 $\text{COMMIT}(n', x')$ of different values from earlier ballots ($n' < n$), so federated voting guar-
3073 antees that no other value was decided by members of S in any earlier ballot number.
3074 Therefore, SCP provides safety when quorum intersection holds in the quorums where
3075 misbehaving replicas are removed from the trust graph.

3076 In addition to a sufficiently long period of network synchrony, liveness requires that a quo-
3077 rum exists and remains available even after deleting the Byzantine replicas. Specifically,
3078 a replica remains live only if it has at least one quorum slice comprised solely of honest
3079 replicas. The set of faulty replicas must not be v -blocking for any honest v in the system.
3080 As mentioned in Section 8.1.1, liveness can be guaranteed for intact sets.

3081 If replicas do not start at the same time or have become desynchronized for any reason,
3082 timeouts alone will not suffice to achieve quorum availability within a ballot. Replicas
3083 begin their timers for ballot n only once they are part of a quorum where each replica is at
3084 ballot n or later, which prevents members of intact sets from staying too far ahead of the
3085 rest of the set. Additionally, if a replica v ever notices a v -blocking set at a later ballot, v
3086 will immediately skip to the lowest ballot such that this is no longer the case, regardless of
3087 timers. These mechanisms combined with the cascade theorem will help replicas that fall
3088 behind to catch up to the same ballot once the network experiences synchrony.

3089 Assume that a ballot n is synchronous for a "long enough" time and that I is an intact set. By
3090 ballot $n + 1$, all replicas in I will have confirmed the same (possibly empty) set of PREPARE
3091 messages, P . If $P = \emptyset$, the nomination subprotocol will converge on a value x . If P is not
3092 empty, let x be the value from the PREPARE message with the highest ballot number in P .
3093 In either case, replicas in I will perform federated voting on $\text{PREPARE}(n + 1, x)$ in ballot
3094 $n + 1$, so there will be a decision on x if ballot $n + 1$ is also synchronous. This demonstrates
3095 that termination is guaranteed if the following conditions hold: the network is synchronous
3096 for two consecutive ballots, and the faulty members of honest replicas' quorum slices fail
3097 to interfere during those ballots.

3098 Liveness, therefore, assumes that every replica in a quorum of a member of I must become
3099 synchronized or not send any messages at all for a sufficiently long period, which may
3100 require members of I to adjust their quorum slices. In an FBAS, replicas can unilaterally
3101 adjust their quorum slices at any time, making recovery from liveness failures substantially
3102 easier in an FBAS than in a typical closed BFT system. In a closed BFT system, consensus
3103 must be achieved on the replica reconfiguration events themselves, as described in Section
3104 7.6, and this is especially challenging when the system has lost liveness.

3105 Because quorum slices are user-configured, there is no guarantee that the emergent struc-
3106 ture actually satisfies the security requirements regarding quorum intersection and avail-
3107 ability. In fact, the actual Stellar network in January 2019 was sufficiently centralized that
3108 if a mere two replicas – both owned by the Stellar Foundation – were hacked or knocked

3109 offline, a liveness failure would have ensued, though improvements have been made since
3110 then [152]. Another line of work has shown that determining whether all quorums of a
3111 given FBAS intersect is actually an NP-complete problem, but there are some heuristic
3112 algorithms that can increase confidence that a particular configuration is secure [153–156].

3113 Despite these heuristics, [155] points out that many of the likely ways that an FBAS would
3114 end up being configured revolve around a "top tier," or small group of replicas that are
3115 essential to providing safety and liveness to the whole system. The membership in this
3116 top tier cannot change without either the active involvement of existing top tier replicas or
3117 a loss of safety guarantees, which calls into question how "open" membership truly is in
3118 practice.

3119 8.2. Ripple

3120 The Ripple protocol was originally introduced in [157], but more up-to-date analyses are
3121 provided in [158, 159], which demonstrate that the original Ripple protocol has problems
3122 with both liveness and safety. Ripple-affiliated researchers have suggested a different algo-
3123 rithm – Cobalt (Section 8.3) – that could be used as an alternative for the network [160].

3124 Not unlike Stellar, the Ripple protocol is designed to guarantee consistency even with only
3125 partial agreement on who participates in consensus. Each user who wants to participate
3126 defines a *unique node list* (UNL), which is the set of replicas that the user will trust for
3127 making decisions about the state of the network. Safety is determined by the intersection
3128 of UNLs between pairs of correct replicas. The original protocol authors believed that
3129 the minimum overlap necessary for security was 20%, but a later analysis suggested it
3130 was 40%, and finally [158] showed that it is actually more than 90%, which is a difficult
3131 condition to satisfy. This is the primary motivation for Cobalt, which only requires a 60%
3132 overlap. Further, if there is not universal agreement on the participants, the network can
3133 get stuck and lose liveness even with 99% UNL agreement and no faulty nodes, requiring
3134 manual intervention to continue making progress. In contrast, Cobalt can make progress
3135 during asynchrony.

3136 In the following description of Ripple’s consensus protocol, let the size of a UNL for replica
3137 P_i be n_i , and define a quorum, q_i , to be 80% of the UNL size. That is, replica P_i can tolerate
3138 no more than 20% of its UNL being faulty. Additionally, the term "ledger" is often used
3139 synonymously with "block" or "state" in the description. The consensus algorithm involves
3140 three steps: deliberation (proposing), validation (voting), and preferred branch (fork-choice
3141 rule). The algorithm is similar to GHOST (Section 11.2) or leaderless versions of longest-
3142 chain permissioned protocols (Section 7.4).

3143 In the deliberation phase, replicas iteratively propose sets of transactions to the other repli-
3144 cas in their UNL over the course of several rounds. In round r of deliberation, a replica
3145 will only include transactions present in at least $threshold(r)$ of the recently received pro-
3146 posals from their UNL, where the thresholds increase from $0.5 \rightarrow 0.65 \rightarrow 0.70 \rightarrow 0.95$ as r

3147 increases. The increasing thresholds are intended to prevent slow replicas from preventing
3148 consensus. When a replica sees a quorum q_i of its trusted nodes agree on the transaction
3149 set, it applies the transactions to its ledger L , broadcasts a *validation* (vote) $V_{L,i}$ on the set,
3150 and begins a new round of deliberation.

3151 Replicas will only issue a validation on one block with a given sequence number. More
3152 formally, replica P_i only issues a validation $V_{L,i}$ for a block L if its height or sequence
3153 number $seq(L)$ is greater than that of any previous block validated by P_i . If the replica
3154 determines that it is not on the preferred branch, it will switch to the preferred one but will
3155 not issue validations until it catches up to the sequence number it was on prior to switching.
3156 In the validation phase, replicas listen for validations from other replicas in their UNL. If
3157 replica P_i sees a quorum q_i of validations for block L , then P_i sets the new fully validated
3158 tip to L .

3159 The preferred branch phase is the chain-selection or fork-choice rule, and it is used to
3160 determine how to make progress when the network is not synchronous. The preferred
3161 branch selection algorithm only switches to a different branch when it knows that enough
3162 replicas have committed to that chain of blocks such that an alternative chain cannot have
3163 more support. Let *lastVals* be the set of most recently validated ledgers (blocks). The three
3164 types of support used in the algorithm are:

3165 1. **Tip support:** The number of trusted replicas whose most recently validated block is
3166 L

$$3167 \quad \text{supp}_{tip}(L) = |\{V_{L',i} \in \text{lastVals} : L = L'\}|$$

3168 2. **Branch support:** The number of trusted replicas whose most recently validated
3169 block is either L or descended from L

$$3170 \quad \text{supp}_{branch}(L) = \text{supp}_{tip}(L) + |\{V_{L',i} \in \text{lastVals} : L \in \text{ancestors}(L')\}|,$$

3171 where $\text{ancestors}(L')$ traces back to the genesis block

3172 3. **Uncommitted support** on sequence number s : The number of trusted replicas whose
3173 most recent validated block has either a sequence number lower than s or lower than
3174 that of the highest block L validation that the replica has personally broadcast

$$3175 \quad \text{uncommitted}(s) = |\{V_{L',i} \in \text{lastVals} : seq(L') < \max(s, seq(L))\}|$$

3176 To find the preferred ledger, a replica walks down the blockchain starting from the block L
3177 that is the common ancestor of the most recently validated blocks. The replica then selects
3178 the child block $L' \in \text{children}(L)$ with the highest $\text{supp}_{branch}(L')$ that would still have the
3179 most support even if all $\text{uncommitted}(seq(L'))$ picked a conflicting fork. If no child of L
3180 satisfies this requirement, then L is the preferred ledger. However, if L' does exist, then
3181 the process is repeated on the children of L' . To break ties on forked blocks, the one with
3182 the higher hash value is selected. If this process leads to a block that is an ancestor of
3183 the replica's current working ledger, they keep that ledger as preferred, since they do not

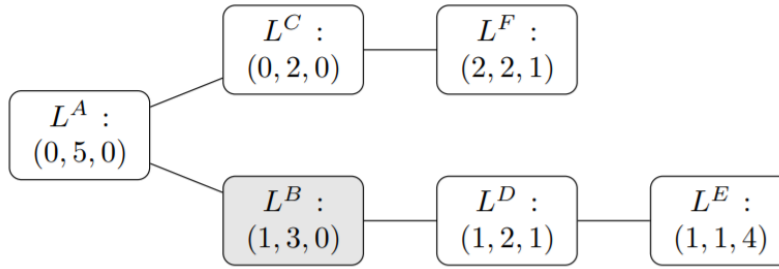


Fig. 19. Ripple "support" example. Each ledger is labeled with the tuple $(supp_{tip}, supp_{branch}, uncommitted)$ from the view of a replica that last validated L^F and has five UNL members. Two trusted replicas had last validated L^F , and one each had last validated L^B , L^D , and L^E . The preferred branch is L^D . [158]

3184 yet know that they are on the wrong branch. See Figure 19 for an example of the support
 3185 definitions being used to choose a preferred branch.

3186 8.3. Cobalt

3187 Cobalt is another algorithm intended for use in open networks with individualized trust as-
 3188 sumptions, and it is slated to be implemented in the Ripple network [160]. As mentioned in
 3189 Section 8.2, Cobalt lowers the UNL overlap requirement from 90% to 60% while provid-
 3190 ing liveness in any network that satisfies this overlap bound between every pair of honest
 3191 replicas. Further, the overlap condition for safety is *local*, so two replicas with sufficient
 3192 overlap with each other cannot arrive at inconsistent ledger states, regardless of the over-
 3193 laps between other pairs of nodes. That is, poorly configured replicas are unable to harm
 3194 properly configured replicas. This makes evaluating whether the network is in a safe con-
 3195 dition easier in comparison to Stellar and the original Ripple protocol, where evaluating
 3196 whether a configuration is safe is an NP-complete problem. Unlike the earlier Ripple pro-
 3197 tocol and SCP, Cobalt maintains liveness in asynchronous networks. The Cobalt algorithm
 3198 is similar to Bracha's broadcast (see Section 1.2) but generalized to incomplete networks
 3199 and combined with a few other techniques.

3200 8.3.1. Background

3201 To generalize classic BFT algorithms and techniques to operate in a setting with incomplete
 3202 networks, some additional concepts must be introduced. Define the *extended UNL* for
 3203 replica P_i , denoted UNL_i^∞ , to recursively include the UNLs from all honest replicas in
 3204 UNL_i^∞ . The extended UNL is P_i 's view of the whole network and includes any replica that
 3205 may have some impact on P_i . Similar to the idea of quorums in Ripple, replicas have a set of
 3206 *essential subsets*, ES_i , such that $UNL_i = \bigcup_{E \in ES_i} E$. If $S \in ES_i$ for P_i , define $n_S = |S|$, and two
 3207 parameters t_S (the maximum tolerable number of Byzantine replicas in S for safety), and
 3208 q_S (the number of correct replicas needed to guarantee liveness) that satisfy the following:

- 3209 1. $0 \leq t_S, q_S \leq n_S$.
- 3210 2. $t_S < 2q_S - n_S$. This means that any two subsets of q_S replicas must intersect at an
3211 honest replica unless the number of Byzantine replicas exceeds the t_S threshold. This
3212 guarantees consistency.
- 3213 3. $2t_S < q_S$. This ensures that liveness holds for replicas with $S \in ES_i$.

3214 Replicas can configure q_S and t_S for each $S \in ES_i$ on an individual basis, and the above
3215 conditions hold when $n_S > 3t_S + 1$ and $q_S = n_S - t_S$.

3216 Replicas P_i and P_j are considered *linked* if there is an essential subset $S \in ES_i \cap ES_j$ such
3217 that fewer than t_S replicas in S are Byzantine. P_i and P_j are *fully linked* if there is some
3218 essential subset $S \in ES_i \cap ES_j$ such that at least q_S replicas in S are honest, at most t_S
3219 replicas in S are Byzantine, and $t_S \leq n_S - q_S$. Linkage is important for consistency while
3220 full linkage helps provide liveness. A replica is said to be *healthy* if it is honest and no
3221 more than $\min\{t_S, n_S - q_S\}$ replicas in each of its essential subsets are not healthy. If a
3222 replica is healthy, then even the influence of Byzantine replicas cannot cause it to accept or
3223 broadcast arbitrary messages. Similarly, a replica is *unblocked* if it is healthy and at most
3224 $\min\{t_S, n_S - q_S\}$ replicas in each of its essential subsets are not unblocked. Blocked nodes
3225 can be arbitrarily prevented from terminating. Replica P_i is *strongly connected* if every pair
3226 of healthy replicas in UNL_i^∞ are fully linked with each other and *weakly connected* if it is
3227 fully linked with every healthy replica in UNL_i^∞ .

3228 A replica P_i sees *strong support* for a message M if P_i receives M from q_S replicas in every
3229 essential subset $S \in ES_i$. Similarly, P_i sees *weak support* for a message M if P_i receives M
3230 from $t_S + 1$ replicas in some essential subset $S \in ES_i$. The security of Cobalt can be derived
3231 based on the following two principles, which suffice to translate normal BFT techniques
3232 that work in complete networks to ones that can operate securely in incomplete networks:

- 3233 1. If two replicas are fully linked, then their essential subsets will overlap to the point
3234 where if one of the replicas sees strong support for a message, the other replica will
3235 eventually see weak support.
- 3236 2. If two replicas are linked, then their essential subsets will overlap to the point where
3237 the replicas cannot simultaneously see strong support for messages that contradict
3238 each other.

3239 The Cobalt protocol itself combines a number of subprotocols that are each adapted to
3240 work in incomplete networks. In particular, it uses an adapted form of reliable broadcast
3241 (RBC) and a multi-value Byzantine agreement (MVBA) that itself uses asynchronous bi-
3242 nary Byzantine agreement (ABA). In essence, proposers distribute their proposals using a
3243 variant of reliable broadcast called democratic RBC (DRBC), which is described in detail
3244 in the next section, and then use MVBA to agree on the assignment of a single proposal for
3245 each slot. The ABA algorithm is adapted from Mostéfaoui's protocol described in Section
3246 6.1.1. An extra messaging round is added to Cobalt's version in order to guarantee that

3247 the ABA’s consistency holds as a local property, which is a significant advantage in open
3248 networks. The MVBA algorithm is out of scope for this document.

3249 8.3.2. Broadcast in Incomplete Networks

3250 The broadcast problem and Bracha’s solution to it are described in 1.2. Cobalt generalizes
3251 the problem and solution to incomplete networks. A solution to the broadcast problem in
3252 this environment with designated sender B_i has the following properties:

- 3253 1. **Consistency:** If an honest replica accepts a message M , no honest replica linked to
3254 it ever accepts an alternative message $M' \neq M$.
- 3255 2. **Reliability:** If a replica is strongly connected and any healthy replica in its extended
3256 UNL accepts a message, then all unblocked replicas in the extended UNL eventually
3257 accept the message.
- 3258 3. **Validity:** If B_i is honest and broadcasts a message M , then any healthy replica that
3259 accepts a message must accept M .
- 3260 4. **Non-triviality:** If B_i is honest and can broadcast to every honest replica in the net-
3261 work, then every unblocked replica will eventually accept the message.

3262 A solution to the democratic reliable broadcast problem has two additional properties that
3263 replace non-triviality, and allows replicas to elect to support or oppose a message.

- 3264 1. **Democracy:** If a healthy and weakly connected replica accepts a message M , then
3265 there is an essential subset $S \in ES_i$ such that the majority of all honest replicas in S
3266 supported M .
- 3267 2. **Censorship-resilience:** If B_i can broadcast to every honest replica in the network and
3268 all honest replicas support M , then every unblocked replica will eventually accept M .

3269 The reliable broadcast protocol below begins by having the designated sender B_i multicast
3270 $INIT(M)$ to everyone who will listen. Then all replicas P_j perform the following steps,
3271 where an empty parenthesis implies an arbitrary message.

- 3272 1. If a replica receives an $INIT(M)$ message directly from B_i , it multicasts $ECHO(M)$
3273 if it has not already sent $ECHO()$.
- 3274 2. If a replica sees weak support for $ECHO(M)$, it multicasts $ECHO(M)$ if it has not
3275 already sent $ECHO()$.
- 3276 3. If a replica sees strong support for $ECHO(M)$, it multicasts $READY(M)$ if it has not
3277 already sent $READY()$.
- 3278 4. If a replica sees weak support for $READY(M)$, it multicasts $READY(M)$ if it has not
3279 already sent $READY()$.

3280 5. If a replica sees strong support for $\text{READY}(M)$, it accepts M .

3281 This RBC protocol can be modified to convert it into a DRBC protocol by having replicas
3282 only send $\text{ECHO}(M)$ messages if they actually support the proposal M . Even if a replica
3283 opposes the proposal, it must participate in the READY phase as normal.

3284 The ECHO phase of the protocol is used to guarantee consistency, while the READY phase
3285 is used to provide reliability. To see why the protocol provides consistency, note that a
3286 replica only accepts a message M if it sees strong support for $\text{READY}(M)$. If two honest
3287 replicas are linked and both accept a message, then by principle (2) above, the messages
3288 must be the same.

3289 Proving reliability is trickier due to having much stronger network assumptions than con-
3290 sistency. If a replica P_k is strongly connected, and two healthy replicas P_i and P_j in its
3291 extended UNL send $\text{READY}(M)$ and $\text{READY}(M')$ messages, then $M = M'$. Steps 3 and 4
3292 of the protocol have an honest P_i send $\text{READY}(M)$ if it saw strong support for $\text{ECHO}(M)$ or
3293 weak support for $\text{READY}(M)$. For P_i to see weak support for $\text{READY}(M)$, then a healthy
3294 replica in $\text{UNL}_i \subseteq \text{UNL}_k^\infty$ must have already sent $\text{READY}(M)$ before P_i . This means that
3295 there was some healthy replica – say, $P_{i'}$ – that was first to send $\text{READY}(M)$, and they did
3296 so because they saw strong support for $\text{ECHO}(M)$. Therefore, one can assume that two
3297 healthy replicas exist, $P_{i'}, P_{j'} \in \text{UNL}_k^\infty$ such that $P_{i'}$ saw strong support for $\text{ECHO}(M)$ and
3298 $P_{j'}$ saw strong support for $\text{ECHO}(M')$. By the definition of strongly connected, $P_{i'}$ and $P_{j'}$
3299 are fully linked, so $M = M'$.

3300 The insight from the previous paragraph can be used to prove reliability. Note that every
3301 pair of healthy replicas in UNL_k^∞ are fully linked. If $P_i \in \text{UNL}_k^\infty$ accepts M , then by principle
3302 (1), all unblocked replicas in UNL_k^∞ will see weak support for $\text{READY}(M)$. The previous
3303 paragraph shows that a healthy replica in P_k 's extended UNL cannot have already sent a
3304 $\text{READY}(M')$ for $M' \neq M$. Step 4 of the protocol ensures that a healthy replica in this
3305 extended UNL will eventually send $\text{READY}(M)$. If $P_j \in \text{UNL}_k^\infty$, then all healthy replicas
3306 in $\text{UNL}_j \subseteq \text{UNL}_k^\infty$ eventually send $\text{READY}(M)$. If P_j is unblocked, it will eventually see
3307 strong support for $\text{READY}(M)$ and thus accept M .

3308 It is clear that the protocol provides validity: healthy replicas do not send $\text{ECHO}(M)$ with-
3309 out having received either $\text{INIT}(M)$ from the designated sender or $\text{ECHO}(M)$ from another
3310 healthy replica. Since the sender is honest, it sends $\text{INIT}(M)$, and no healthy replica will
3311 send an $\text{ECHO}(M')$ with $M' \neq M$. Similar logic holds for the READY messages, implying
3312 that no healthy replica will see enough $\text{READY}(M')$ messages to accept M' . Non-triviality
3313 is even simpler: every replica can receive $\text{INIT}(M)$ from the designated sender, so all
3314 healthy replicas will send $\text{ECHO}(M)$ and eventually $\text{READY}(M)$, such that all unblocked
3315 replicas eventually accept M .

3316 The proof of censorship-resilience for the DRBC variant is identical to the proof of
3317 non-triviality for the standard incomplete networks variant. To see why the democracy
3318 property holds, let P_k be a healthy and weakly connected replica that accepts M . Let

3319 $P_i \in UNL_k^\infty$ be the first healthy replica to have sent $READY(M)$. Then P_i saw strong
3320 support for $ECHO(M)$. Weak connectivity implies that P_i and P_k are (fully) linked, so
3321 there exists an essential subset $S \in ES_k$ where at least $q_S - t_S$ honest replicas in S send
3322 $ECHO(M)$ and no more than $n_S - q_S$ honest replicas in S did not send an $ECHO(M)$.
3323 Because $t_S < 2q_S - n_S$ (the second assumption on the parameters, from Section 8.3.1),
3324 $q_S - t_S > q_S - (2q_S - n_S) = n_S - q_S$, so a majority of honest replicas in S supported M .

3325 9. Proof of Work: The Basics

3326 The concept of proof of work was first suggested by Dwork and Naor as a way to counter
3327 email spam by adding a small cost to sending emails [161]. The term "proof of work" was
3328 later coined by Jakobsson and Juels [162]. The proof-of-work scheme used in Bitcoin was
3329 inspired by Adam Back's "Hashcash" [163]. More recently, "resource burning" has been
3330 studied in general, which includes proof of work but also proof of space, which is discussed
3331 in Section 14.1 [164].

3332 In the context of a replicated state machine, a proof-of-work consensus algorithm will con-
3333 sist of a specific proof-of-work puzzle, a difficulty adjustment algorithm, a fork-choice rule
3334 (and implicitly, a data structure to work on, like a blockchain or DAG), and an incentiviza-
3335 tion scheme. This section marks a shift in this document's discussion from permissioned
3336 to permissionless consensus algorithms and will discuss the basic functionality of proof of
3337 work and how it operates as a Sybil-resistance mechanism for distributed ledger systems.

3338 9.1. Proof of Work and Sybil Resistance

3339 The actual mechanics behind constructing a proof of work are quite simple. Some proof-
3340 of-work function must be defined, and a cryptographic hash function is typically used. A
3341 puzzle difficulty is chosen (in the consensus context, this is done with a difficulty adjust-
3342 ment algorithm, discussed in Section 9.2), that then determines the range of hash function
3343 outputs that would constitute a successful proof.

3344 To find a proof of work, entities try to find partial hash collisions using the given func-
3345 tion, difficulty target, and specified message format. The message needs to include any-
3346 thing that should be "tied" to the proof. For example, in the context of email spam, the
3347 message should include the contents of the email to be sent. For a proof-of-work cryp-
3348 tocurrency like Bitcoin, the proof must cover the entire contents of the block to be mined
3349 and the previous block hash (if a DAG is used instead of a blockchain, then multiple pre-
3350 vious block hashes may be included). Given a proof-of-work function (H), a difficulty
3351 target to satisfy ($Target$), a previous block hash to mine on top of ($prevHash$), and a set
3352 of transactions to mine into a block (txs), a miner will repeatedly try different nonces until
3353 $H(prevHash||txs||nonce) < Target$.

3354 A good proof-of-work function has several important properties. First, the puzzle must be
3355 moderately hard, have this hardness be tunable, and be very fast to verify. Partial hash

3356 inversions are a good example of this, where the difficulty of the puzzle can be tuned by
3357 changing the *Target*, and verification is as fast as a single hash query. The function must
3358 also be *memoryless*, such that the time it takes to solve the puzzle does not depend on how
3359 much time has already elapsed or the history of attempts to solve the puzzle. That is, if a
3360 miner spends five minutes trying to mine a block, they are no closer to having found it than
3361 when they began. This is why verifiable delay functions by themselves are inadequate for
3362 proof of work. An adversary with a single processor with sequential processing speeds just
3363 slightly faster than honest parties will almost inevitably solve every puzzle first. Contrast
3364 this with a cryptographic hash function, where guessing nonces can be performed in parallel
3365 and where previous guesses do not bring the miner closer to finding a solution to the puzzle.

3366 Finally, the most fundamentally important property of a good proof-of-work function is that
3367 there should be no strategies or shortcuts that aid in solving the puzzle. A cryptographic
3368 hash function (or at least one modeled as a random oracle) is a good example because there
3369 is no strategy better than arbitrarily guessing nonces. It is this property that ensures that
3370 any verified proof-of-work solution implies that the prover invested a sufficient amount of
3371 computational effort. This is what makes proofs of work useful tools in consensus: by
3372 imposing a cost on participants who want to send messages, parties running the consensus
3373 protocol have opportunities to synchronize their local views regarding the state of the sys-
3374 tem [165]. An adversary with a shortcut would be able to out-compete other miners and
3375 ultimately centralize the system under their control.

3376 An example that vividly demonstrates the significance of potential shortcuts is AsicBoost,
3377 a technique that improves the efficiency of Bitcoin mining by about 20% [166]. Bitcoin's
3378 mining algorithm involves performing a double SHA-256 hashing of the 80-byte block
3379 header. SHA-256 operates on 64-byte chunks of the message at a time, so two chunks must
3380 be hashed. The Bitcoin block header – including how it is broken down into chunks for
3381 mining – is displayed in Figure 20a. The two chunks are processed in the manner shown in
3382 Figure 20b.

3383 The outer loop of the mining process – shown in green in Figure 20b – preprocesses the
3384 first chunk using the expander and compressor functions of SHA-256 and results in an
3385 output called the *midstate*. The inner loop – shown in red in Figure 20b – preprocesses
3386 the second chunk analogously but then takes the midstate from the outer loop as input to
3387 a second round of SHA-256 expansion and compression. A *work item* in Bitcoin consists
3388 of the midstate generated from preprocessing the first chunk, as well as the tail, timestamp,
3389 and bits fields from the second chunk. The nonce is incremented in each run through the
3390 inner loop, requiring another two expansions and two compressions. The performance
3391 gains from AsicBoost come from reusing the first expansion of the second chunk across
3392 multiple work items. To do so, the miner needs multiple block header candidates that share
3393 the same tail, timestamp, and bits fields.

3394 Notice how the Merkle root, which commits to every transaction included in the block,
3395 spans both chunks of the block header. The AsicBoost technique exploits collisions in

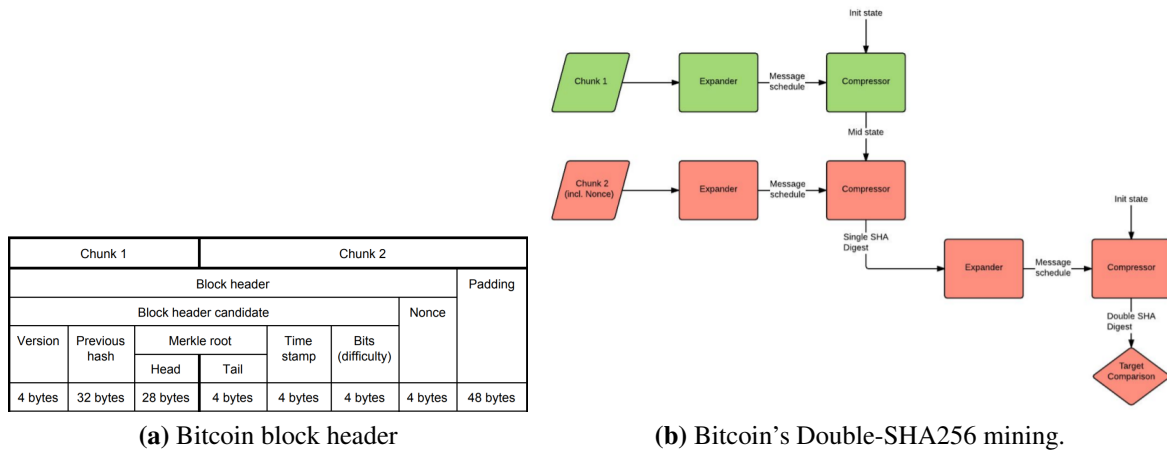


Fig. 20. Bitcoin mining and AsicBoost. (a) Bitcoin block header, as used for mining. (b) Bitcoin's double SHA-256 mining. The outer loop is in green, while the inner loop is in red. [166]

3396 the last four byte "tail" of the Merkle root. A miner employing the technique will find a
 3397 set of Merkle roots that collide in these bytes by varying and/or rearranging the order of
 3398 transactions included in the block (finding four byte collisions is not too challenging due
 3399 to the birthday paradox). That is, the miner will have a large set of differing first chunks
 3400 (and their corresponding midstates) that remain valid when paired with a particular second
 3401 chunk. By reusing the current value of the second chunk with Merkle roots colliding in the
 3402 tail, a miner can simply swap out midstates to their precomputed values while skipping the
 3403 first expansion of the inner loop. As a result, the inner loop only has three large operations
 3404 instead of four.

3405 Suppose that a miner finds three Merkle roots where the final four bytes collide, and call the
 3406 resulting midstates that arise from the first chunk *A*, *B*, and *C*. The resulting mining loop
 3407 will begin with the miner setting the nonce to zero and using midstate *A*, requiring four large
 3408 operations. The miner then keeps the same nonce but swaps in midstate *B*, which allows
 3409 skipping the first expansion and only performing three operations. If this fails, computing
 3410 using midstate *C* also only requires three operations. Then the miner sets the nonce to one
 3411 and resets the midstate to *A*, computes the inner loop using four operations, and continues
 3412 in this way, swapping midstates before incrementing the nonce.

3413 One of the major reasons why the discovery of AsicBoost was a significant event in the Bit-
 3414 coin community was due to patents. The technique itself was patented, but if the technique
 3415 could not be used universally, it would provide a likely insurmountable mining advantage
 3416 to the patent-holder. The AsicBoost patent is now held under the Blockchain Defensive
 3417 Patent License, which obligates any participating entity to share their own mining-related
 3418 patents. As a result, this is less likely to be an issue in the future. Further, the version of
 3419 AsicBoost described above is no longer possible on Bitcoin due to the Segregated Witness

3420 soft fork performed in 2017. An overt version of AsicBoost that simply adjusts the version
3421 number in the first chunk instead of finding Merkle root tail collisions is still possible.

3422 9.1.1. Mining Pools

3423 Rewards for mining are provided to miners when new blocks are found. In Bitcoin, new
3424 blocks are found approximately every 10 minutes or 144 per day. Other networks produce
3425 blocks more frequently, though network latency places fundamental bounds on how quickly
3426 blocks can be produced (this latency issue is discussed in detail in Section 10.2.1). Because
3427 there are relatively few opportunities to mine blocks and thus collect rewards, it can take an
3428 exceedingly long time for small miners to find a block once there are a significant number
3429 of miners on the network. As a result, small miners could easily go out of business before
3430 mining a single block, and the variance in rewards for these miners is substantial.

3431 This high variance in mining rewards is the primary motivation for *mining pools*, where a
3432 variety of hardware operators cooperate in mining and share the rewards among themselves.
3433 This allows smaller, more frequent mining payouts, which makes mining viable to a much
3434 wider variety of entities, including those with fewer resources. Pools operate by setting a
3435 difficulty level that is a small fraction of the difficulty of mining an actual block, and proofs
3436 of work at these lower difficulty levels are called *shares*. Participating miners submit shares
3437 to the pool operator in exchange for a portion of the pool's reward income (various reward
3438 schemes and attacks against them are discussed in Section 9.3). Because there are far more
3439 shares than blocks, there are many more opportunities for small miners to collect payouts.

3440 Mining pool operators run full nodes and assemble block templates to send to the miners,
3441 providing each miner with a particular nonce range to search in. Miners themselves simply
3442 operate the hardware. They do not necessarily validate blocks or otherwise participate
3443 in the network. The protocol that pool servers use to communicate with miners is called
3444 Stratum, which lacks authentication and has resulted in some miners losing their payouts
3445 due to network-layer attacks [167].

3446 It is a common misunderstanding that mining pools and miners are effectively the same
3447 thing, even though individual miners can easily switch pools at will. If a small set of min-
3448 ing pools are responsible for mining the vast majority of blocks, many people conclude that
3449 only a handful of entities control the network, oversimplifying the nature of "control." Un-
3450 der the current Stratum protocol, there is *some* truth to this. By assembling block templates
3451 and thus choosing which transactions to include in a block, mining pools can launch certain
3452 attacks or otherwise engage in nefarious behavior. This is one of the primary motivations
3453 for the ongoing development of Stratum v2, which eliminates these risks by letting miners
3454 select transactions while still pooling rewards [168]. Currently, due to the privileged po-
3455 sition of mining pools as the entities that select transactions, the following behaviors are
3456 indeed possible:

- 3457 • A mining pool can censor the inclusion of transactions that it does not like or that

3458 regulators tell them to. If a majority of the computational power of the network is not
3459 directed at censorship, however, it will only cause increased latency for transaction
3460 confirmation. Further, pools lose out on transaction fees by censoring, so even if the
3461 majority of the hash rate is malicious, there is an incentive for miners to deviate from
3462 the censoring cartel in order to capture more fee revenue.

3463 • Pools can attempt chain reorganizations for double-spending. This is unlikely to suc-
3464 ceed without a majority of the computational power of the network involved, and
3465 the miners themselves may have strong incentives not to go along with it. Addi-
3466 tionally, as miners detect an attack in progress, they can migrate to another pool.
3467 If a successful double-spend would decrease the value of a miner's virtual assets or
3468 the hardware used to mine them, there is a significant incentive to deviate from the
3469 double-spending cartel.

3470 • Pools can use their position to influence the rules of the network, using the miners
3471 as leverage. For example, if the pool operator believes that the rules of the net-
3472 work should change to allow higher throughput, they can mine empty blocks and
3473 allow a backlog of transactions to accumulate in order to increase fees and frustrate
3474 users. Pools can also direct hashpower to networks with alternative rules (but the
3475 same proof-of-work algorithm) without the consent of miners. For instance, miners
3476 who thought they were mining Bitcoin can be made to mine Bitcoin Cash instead.
3477 Similarly, one of the more common mechanisms employed in changing the rules of
3478 the network is miner signaling: miners set a bit in their block headers to signify ap-
3479 proval. Pools can thus signal on behalf of miners who may or may not agree. As
3480 with the other possible attacks, miners can switch pools to those more aligned with
3481 their values (including, perhaps, profitability) once they detect that their current pool
3482 is behaving in ways that they do not approve of. To be clear, neither mining pools nor
3483 miners can arbitrarily change the rules of the network, even with a majority of the
3484 hash rate. Blocks that are invalid under the "old" rules will be rejected and ignored
3485 by the rest of the network.

3486 An interesting consequence of having only a handful of major mining pools operating si-
3487 multaneously is that many pool operators are known, identified entities. The implications
3488 of this are mixed: it should be easier for pool operators to coordinate nefarious activity or
3489 be coerced into performing attacks, but they can also be held more accountable for mis-
3490 behavior due to a desire to maintain their reputation. Again, miners themselves can easily
3491 switch pools, so a pool that acts against the best interests of its miners may not survive for
3492 long.

3493 Even if a (small) majority mining cartel is formed, deviating from the cartel can be highly
3494 profitable both for pool operators and the miners themselves. If the deviation pushes the
3495 cartel's influence back under 50%, then participants in the cartel gain no rewards, and the
3496 deviating party (as well as those who never participated in the first place) temporarily face
3497 less competition and thus collect greater rewards. Any pool operator representing enough

3498 hash rate to reduce the cartel back to a minority has a strong incentive to be the one who
3499 deviates. If the cartel agreement is only between pools, miners within those pools are
3500 likely to switch to other pools once the attack is detected. The miners themselves face
3501 electricity and other operational costs to participate in the attack but get no benefit from it
3502 (while risking the value of their assets, as described above). As a result of these factors,
3503 an attack by one or more mining pools is likely to only succeed in the short term rather
3504 than establishing long-term control over the contents of the blockchain. There are many
3505 more miners than pools, and miners may operate anonymously, so organizing a majority
3506 cartel of miners is far more challenging in practice. For these reasons, it is a significant
3507 oversimplification to suggest that the network is controlled by very few entities.

3508 Due to the perceived centralization of mining into pools, there have been efforts to remove
3509 or reduce their power, including the aforementioned Stratum v2 protocol. Another effort is
3510 P2Pool, a more distributed mining pool. P2Pool miners create their own blockchain that
3511 is composed of their lower difficulty shares, and the difficulty is set such that shares are
3512 expected to be found every 30 seconds (or 20 times more frequently than blocks). The
3513 P2Pool sharechain holds 8,640 shares at a time, and payouts are performed proportionally
3514 to a miner's fraction of those shares. To enforce this, miners verify that any shares they
3515 build off of include the appropriate payouts in the coinbase transaction. Unfortunately,
3516 P2Pool has not been able to gain a significant amount of hash rate, in large part because the
3517 latency problems that arise in proof of work are magnified substantially, resulting in a very
3518 high rate of stale shares.

3519 A more extreme way to neutralize the power of mining pools is to design a proof-of-work
3520 algorithm that entirely precludes pools from existing profitably in the first place. Specif-
3521 ically, *non-outsourcable puzzles* have the property that if a pool operator is able to out-
3522 source work to a miner, the miner is capable of stealing the reward in a way that does not
3523 implicate themselves [169]. There are two major problems with this approach:

- 3524 1. In many cases, it is possible to devise smart contracts where miners must submit col-
3525 lateral that can be seized by the pool operator should a miner steal the block reward
3526 [170]. This essentially recreates pools but with higher overhead and complexity.
- 3527 2. By reducing the variance of payouts, pools provide a valuable service. Many small
3528 mining operations simply would not be viable in the first place if it were not for
3529 pools. If non-outsourcable puzzles successfully eliminated mining pools, it is quite
3530 likely that there will be far fewer miners overall, resulting in even more centralization
3531 of mining.

3532 9.1.2. Hardware: ASICs and ASIC Resistance

3533 One of the more significant debates among cryptocurrency enthusiasts is whether it is de-
3534 sirable for the proof-of-work algorithm to be mined with specialized hardware, such as
3535 custom-built application-specific integrated circuits (ASICs) or general-purpose hardware

3536 like CPUs and GPUs.

3537 *ASIC resistance* is a property of a proof-of-work algorithm that stipulates that it would be
3538 challenging to build specialized hardware that is capable of solving proofs of work much
3539 more efficiently than general-purpose hardware. Proponents of ASIC resistance argue that
3540 the wider availability of CPUs and GPUs results in a far lower barrier to entry for small min-
3541 ers while providing a mechanism for hobbyists to acquire cryptocurrency without needing
3542 to identify themselves to a centralized exchange. Large segments of the population already
3543 possess general-purpose hardware or can acquire it without revealing their specific interest
3544 in cryptocurrency. Contrast this with ASICs, which tend to cost hundreds or thousands of
3545 dollars per machine and which leave no plausible deniability for their purpose. ASICs can
3546 be intercepted at national borders or otherwise stopped by hostile governments or shipping
3547 companies, and their manufacturers can be identified and targeted. Additionally, ASIC
3548 manufacturers themselves can become important players in the "politics" of the networks
3549 they produce hardware for, which can include advocating for controversial rule changes
3550 and selectively selling machines to preferred partners.

3551 Many approaches have been used in attempts to gain the ASIC resistance property, but the
3552 overwhelming majority of them have failed within a few years. Once there is sufficient
3553 money at stake – that is, when a proof-of-work cryptocurrency becomes valuable enough
3554 – it seems that hardware manufacturers find a way to build an efficient ASIC for the al-
3555 gorithm. It is unclear to what extent long-term ASIC resistance is even possible in the
3556 first place. One of the more promising attempts at ASIC resistance is called RandomX,
3557 which has most prominently been used for the Monero cryptocurrency since November
3558 2019 [171]. RandomX is designed to be mined on CPUs. The algorithm relies on a pseu-
3559 dorandom key that is periodically extracted from the blockchain. This key is used to help
3560 generate random programs in a very general, low-level instruction set where any random
3561 string is a valid program. This is translated into machine code and executed in a way that
3562 uses as many components of the CPU as possible before being hashed into a final result.
3563 This use of many CPU components should complicate the design of an ASIC, which is pro-
3564 grammed to only execute a very specific task. Only time will tell if ASICs are eventually
3565 built for RandomX as well.

3566 On the other side of the debate are those who feel that proof of work is more secure when
3567 specialized hardware is widely used for mining. In this way of thinking, it is good for the
3568 security of a given network if its proof-of-work algorithm is the dominant application of a
3569 particular piece of hardware. An ASIC-resistant algorithm is likely to have a huge quantity
3570 of (unused) *potential* hash rate and may be more likely to have the security assumptions
3571 of the network (e.g., honest majority of hash rate) be violated. Every CPU or GPU in the
3572 world could potentially be turned on to attack the network. This is not solely a concern
3573 for ASIC-resistant algorithms. If two cryptocurrencies are mined using the same proof-
3574 of-work algorithm (e.g., Bitcoin and Bitcoin Cash), the less valuable one will suffer a
3575 security deficit for the same reason. In addition, ASICs – by only being useful for a limited
3576 purpose – require miners to have a vested interest in the network. If the network is attacked,

3577 ASIC owners will lose out on their initial hardware investment, which cannot be repurposed
3578 except to mine an alternative cryptocurrency using the same algorithm that may also be
3579 attacked by the same hardware.

3580 Because many believe that long-term ASIC resistance is futile, attempting to provide ASIC
3581 resistance merely imposes barriers to entry for hardware manufacturers. This makes it even
3582 easier for the market to become concentrated, perhaps with a single major manufacturer that
3583 can build ASICs and mine with them in secret. A related risk is that frequently hard-forking
3584 to eliminate ASICs from the network creates strong incentives for collusion between de-
3585 velopers and hardware manufacturers and, thus, introduces a strong point of centralization
3586 around the development of the mining algorithm. A related concern is that the designer of
3587 an ASIC-resistant algorithm that will initially be executed mostly on CPUs can develop an
3588 optimized GPU implementation that they can run in secret, providing a significant mining
3589 advantage.

3590 Some of these security claims have been investigated in the literature. For instance, when
3591 multiple cryptocurrencies use the same proof-of-work algorithm but one has much higher
3592 difficulty than the other, rational miners from the dominant chain may be incentivized to
3593 migrate to the minority chain to perform a 51% attack [172]. Several works have contem-
3594 plated what happens when multiple cryptocurrencies share the same proof-of-work algo-
3595 rithm and some portion of miners automatically switch between networks in order to mine
3596 the more profitable coin [173, 174]. The primary finding of these works is that there exists
3597 a hash rate equilibrium between the competing networks based solely on the market price
3598 of their respective coins and that miners allocating their hash rate off-equilibrium merely
3599 creates a profitable arbitrage opportunity for other miners. As a result, miners "loyal" to the
3600 minority chain – that is, miners who will continue mining on the lower difficulty network
3601 – will mine alone if they exceed the equilibrium allocation. The security of the lower value
3602 network cannot be improved by allocating more hash rate below the equilibrium value, and
3603 loyal miners would centralize the network above the equilibrium value by causing other
3604 miners to migrate back to the stronger network. The only way to alter the equilibrium is to
3605 change the relative prices of the respective assets, which is non-trivial.

3606 The relationship between a ledger's security and the price of its native asset appears to have
3607 better dynamics under certain conditions that are far more likely to hold when ASICs are
3608 involved in mining rather than just GPUs and CPUs. A combination of high fixed costs to
3609 buy hardware and set up a mining operation plus a low salvage value for the hardware itself
3610 leads to asymmetric hash rate changes in response to price changes [175, 176]. Specifically,
3611 when a network relies on ASICs, miners will deploy more hardware as prices increase but
3612 do not decrease to the same degree when prices decline. By responding less to adverse
3613 exchange rate shocks, ledgers secured by ASICs are less likely to suffer double-spend
3614 attacks after a price decline. Empirically speaking, Ethereum's hash rate responded to price
3615 symmetrically at first, but after ASICs were developed for Ethash (Ethereum's formerly
3616 ASIC-resistant proof-of-work algorithm), hash rate changes in response to price became
3617 more asymmetric [176]. This beneficial asymmetry is reduced if the ASIC is transferable

3618 to another cryptocurrency, though it is likely to still exist if the secondary network is small
3619 relative to the network they are actually mining on. A paper by Garratt and van Oordt
3620 explains the economic reasoning behind these insights [175]:

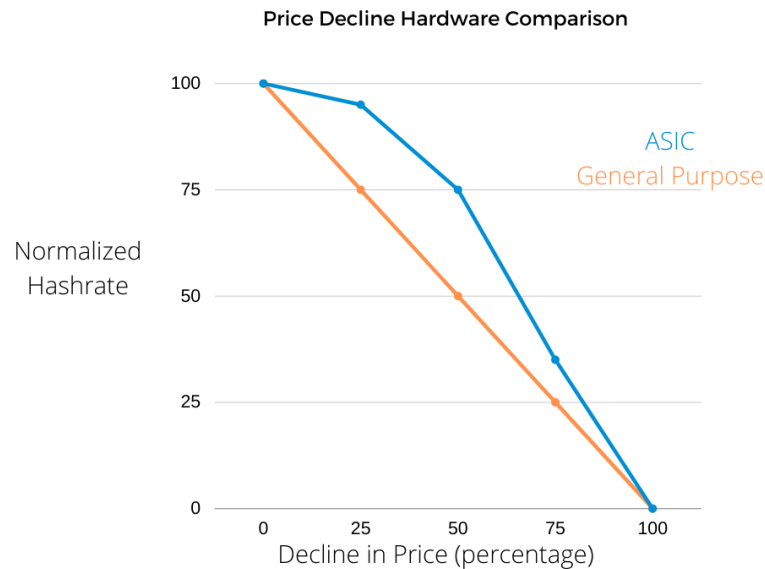


Fig. 21. In the face of an unanticipated price decline, the behavior of miners depends on the type of hardware they use. Because miners using general-purpose hardware can leave freely, they leave in proportion to the decline in the block reward. In contrast, miners who have already invested in ASICs, which cannot be used for other purposes, may continue to mine.

- 3621 • Without fixed costs, the decline in the hash rate is proportional to the decline in the
3622 exchange rate due to free entry and exit into mining. With fixed costs and small
3623 exchange rate declines, miners continue mining to recover some of these fixed costs
3624 and will do so until it becomes more profitable to sell the hardware for scrap. ASICs,
3625 which are useless for anything other than mining on a particular ledger, have very
3626 low scrap value. This difference is shown in Figure 21.
- 3627 • With fixed costs and low scrap value, launching a double-spending attack results in a
3628 loss in the present value of future mining rewards due to the attack's negative impact
3629 on the asset's price. Without fixed costs, the only costs of attempting a double-
3630 spending attack are the potential loss of mining rewards during the attack if it fails
3631 and the variable costs of mining for the duration of the attack.
- 3632 • With fixed costs and low scrap value, the determination of whether a double-spending
3633 attack is profitable or not is path-dependent. At any exchange rate, the hash rate will
3634 be higher if the current exchange rate is less than a previous peak because mining op-
3635 erations will expand as the price approaches its peak. The price decline implies a de-

3636 creased present value of future mining rewards, such that profitable double-spending
3637 attacks are more likely if the current price is the result of a decline from a previous
3638 peak. The principle here holds regardless of the hardware used but presents a more
3639 substantial security risk for general-purpose hardware due to the first point above.

3640 • When multiple cryptocurrencies are mined with the same hardware and an attack on
3641 one of the networks is not expected to impact the price of other cryptocurrencies that
3642 can be mined with the same equipment, double-spending attacks are more likely to
3643 be profitable than if only a single cryptocurrency were available. Whether an attack
3644 can be expected to reduce the price of alternative cryptocurrencies using the same
3645 hardware likely depends on their relative sizes. It is hard to imagine a successful
3646 attack on the larger cryptocurrency not causing a decline in value for the smaller
3647 one because it demonstrates the insecurity of the smaller network. An attack on the
3648 smaller one is less likely to have the same impact and is unlikely to significantly
3649 affect the expected return from a given piece of hardware.

3650 These arguments strongly suggest that there is a security advantage to a network when it
3651 is the dominant application of a given piece of hardware. However, the wider availability
3652 of general-purpose hardware offers other advantages that are difficult to measure but worth
3653 considering.

3654 9.1.3. Mining Centralization in Practice

3655 Some believe that – due to economies of scale – proof of work in the long term is likely
3656 to result in a concentration of power among an oligopoly consisting of a handful of large
3657 miners. For instance, in the model presented in [177], even relatively small economies
3658 of scale resulted in few miners operating simultaneously, and a miner with an $X\%$ cost
3659 advantage over any other miner will have at least $X\%$ of the total hash rate. As a result,
3660 a world in which a miner is able to gain a significant cost advantage over the competition
3661 is likely to result in a centralized mining network. That said, the model assumes that
3662 all potential miners decide whether to make investments in equipment simultaneously at
3663 particular factor prices. In the real world, these decisions are not made simultaneously, and
3664 prices change, likely disturbing the equilibrium analysis.

3665 It is difficult to obtain data on the distribution of hash rate among miners because block
3666 reward payouts go to the pools rather than to the miners operating the hardware. That said,
3667 Romiti et al. performed an empirical analysis of Bitcoin mining shares and inferred that
3668 the addresses paid out from coinbase transactions likely belong to miners of the pool [178].
3669 Their data came from three large mining pools and found that fewer than 20 distinct entities
3670 collect more than half of the rewards within each pool and that rewards from multiple
3671 pools often go to the same entity. This may sound like a substantial degree of mining
3672 concentration, but many of these payout addresses belong to exchanges, which dramatically
3673 complicates the analysis. Every distinct miner who directly sends their payouts to the same
3674 exchange will be considered a single entity under this methodology. As such, it is difficult

3675 to draw conclusions about how many actual miners operate on the network and their relative
3676 computational power.

3677 Regardless of the concentration among miners, it is clear that the majority of miners will
3678 tend to form only a relatively small number of pools. As described above in Section 9.1.1,
3679 the centralization into relatively few pools does indeed present risks. For example, majority
3680 attacks have been conducted against many small and medium-sized networks where that
3681 ledger was not the dominant application for the relevant mining hardware. Further, even
3682 Bitcoin had a brief period in 2014 where a mining pool – GHash.io – had more than half the
3683 total hash rate of the network, though no attacks occurred at that time. On a more optimistic
3684 note, the ability of miners to allocate their hashes between multiple pools tends to lead to
3685 the decentralization of pools as larger pools increase their fees [179]. Large pools raise
3686 fees in order to decrease their hash rate because their positive impact on the difficulty level
3687 harms them more. In practice, distributing hash power this way is easy for miners. Tools
3688 exist to help and have been experimentally shown to increase a miner’s Sharpe ratio (i.e.,
3689 average excess reward over its standard deviation/volatility) by 260% compared to more
3690 passive miners [180]. Additionally, the adoption of Stratum v2 is likely to go a long way
3691 toward mitigating these pool concentration risks in the future.

3692 In addition to miners and pools, hardware manufacturing and design is another potential
3693 avenue of centralization for proof-of-work cryptocurrencies. Designing ASICs requires
3694 specialized knowledge and capital, and there have historically been very few SHA-256
3695 ASIC manufacturers (and even fewer for other proof-of-work algorithms). This market
3696 concentration extends to the foundries responsible for manufacturing the integrated cir-
3697 cuits. That said, in the years since the first Bitcoin ASIC was produced in 2013, ASIC
3698 replacement cycles have grown substantially longer [181]. Early designs were obsolete
3699 within a few months, but designs have caught up with Moore’s Law, and ASICs now often
3700 last years before more efficient designs make them unprofitable. This has opened up new
3701 competition in ASIC design, resulting in far more competition in the Bitcoin ASIC indus-
3702 try [181]. As it becomes harder and harder to squeeze out further optimizations for these
3703 chips, ASIC production is expected to become further commoditized, reducing the power
3704 of manufacturers.

3705 Besides general concentrations of power, there may be other concerns regarding geographic
3706 centralization in certain regions or political jurisdictions. If miners or mining pools are
3707 concentrated in relatively few geographical areas, it becomes substantially easier for the
3708 governments of those jurisdictions to disrupt the network. It also creates certain risks that
3709 might not otherwise exist, such as flooding in a region popular for mining, which can take
3710 a large fraction of the network hash rate offline simultaneously.

3711 There are a number of reasons why miners may cluster into relatively few geographical
3712 regions. First, miners are inclined to set up their facilities wherever electricity is cheapest,
3713 since it is the greatest variable cost for mining and thus one of the simplest ways to gain
3714 an advantage over the competition. In practice, a number of governments offer generous

3715 electricity subsidies, which tends to attract miners. Second, the incentives of mining pools
3716 are such that they benefit most by propagating their blocks to more than 50% of the hash
3717 rate as quickly as possible but not necessarily 100% for network latency-related reasons, as
3718 discussed in Section 10.2.1. This latter incentive is particularly relevant when governments
3719 or ISPs engage in network censorship that may slow down the distribution of blocks. If
3720 the majority of the hash rate is located in such a jurisdiction, then miners outside of that
3721 jurisdiction face a distinct mining disadvantage.

3722 However, there are several factors that can mitigate the risks of geographic concentrations
3723 of miners. For one thing, there are significant operational challenges with seizing control
3724 over all of the deployed physical machines within a country. It is relatively easy to relocate
3725 machines, which is often done to take advantage of different electricity prices at different
3726 times of year. In addition, the hardware is already spread out across a variety of regions
3727 throughout the country. Seizing the vast majority of this equipment would be very chal-
3728 lenging to do quickly and discreetly, and once word has gotten out that such seizures are
3729 happening, the community can plan for emergency actions, such as a hard fork to change
3730 the proof-of-work algorithms and render those machines useless. It would be far easier for
3731 the government to target the pools instead of the miners, but again, it would be challenging
3732 to do so covertly. Mining pools can quickly be set up in other locations around the world,
3733 and it is easy for miners to switch pools. Ultimately, when Stratum v2 is more widely used
3734 by these pools, the attack vector will be closed.

3735 9.2. Difficulty Adjustment Algorithms

3736 Since proof-of-work cryptocurrencies lack a central authority responsible for determin-
3737 ing who can mine and with what machines, there will inevitably be a varying amount of
3738 hash power deployed to the network as the exchange rate and mining profitability changes.
3739 Nevertheless, in order to maintain a specific monetary policy and provide a better user ex-
3740 perience, new blocks need to be found at a relatively consistent pace (e.g., roughly every
3741 10 minutes in Bitcoin), regardless of the amount of hash power deployed. In these systems,
3742 the *difficulty adjustment algorithm* (DAA) is responsible for adjusting the difficulty of the
3743 mining puzzle in response to changes in the network's total hash rate in order to keep block
3744 production consistent. Without a DAA, an increasing hash rate would result in blocks being
3745 found more and more frequently, inflating the currency more quickly and making payments
3746 less predictable and secure for users.

3747 Although the primary purpose of a DAA is to maintain consistent inter-block arrival times
3748 over the long term despite fluctuations in hash rate, there are a variety of possible consid-
3749 erations that may go into its design. The algorithm should ideally avoid sudden difficulty
3750 changes when the hash rate remains constant, discourage wild oscillations from the feed-
3751 back between hash rate and difficulty, and avoid exceptionally long intervals between new
3752 blocks. Generally speaking, DAAs face a trade-off between being more stable or being
3753 more quickly responsive to changes in hash rate. More stable algorithms, such as Bitcoin's

3754 DAA, result in more blocks that are "inappropriately" cheap or expensive because they
3755 measure hash rate less accurately and are more disruptive to users if the hash rate drops
3756 suddenly, but they are more secure and easier to reason about [182–184]. In Bitcoin, the
3757 difficulty is adjusted every 2,016 blocks (about two weeks) proportional to the degree to
3758 which blocks were found too quickly or slowly. For instance, if it takes four weeks to
3759 mine 2,016 blocks, the difficulty for the following 2,016 block epoch will be cut in half.
3760 Algorithms that adjust the difficulty more frequently will maintain more consistent block
3761 times but are easier for miners to game for unfair advantage [185]. Common DAAs include
3762 variants of simple moving averages and exponential moving averages that adjust after every
3763 block, though there is a wide design space.

3764 In a world where more than one proof-of-work ledger exists, there is a risk that difficulty
3765 changes in one network cause a behavioral shift among miners that affects another network.
3766 For example, an interaction between the DAAs of Bitcoin and Bitcoin Cash (which share
3767 the same hardware) caused wild fluctuations and instability in the Bitcoin Cash hash rate
3768 until a hard fork changed the Bitcoin Cash DAA [186]. More generally, *coin-hopping*
3769 *attacks* are possible between networks that share the same hardware [187]. At the beginning
3770 of a difficulty epoch, an adversarial or strategic miner can switch to mining on an alternative
3771 network and then switch back to the original network at the beginning of the following
3772 difficulty epoch when the difficulty is lower due to the adversary's withdrawn hash rate.
3773 This allows the miner to mine at a lower average difficulty and cost than competing miners
3774 that do not hop between coins. In some circumstances, it may be profitable to perform
3775 similar strategic deviations for a single network, mining during one difficulty epoch and
3776 turning machines off during the next, or using only a portion of the miner's available hash
3777 rate during some epochs in order to lower the difficulty for future epochs [188, 189].

3778 **9.3. Attacks Against Mining Pools: Pool-Hopping and Block Withholding**

3779 Mining pools allow miners to lower the variance in their revenue by allowing them to
3780 split the block reward based on how many shares each miner submits to the pool operator.
3781 An intuitive scheme for doling out a pool's rewards to its constituent miners would be to
3782 define a round as the period between two blocks being found by a given pool and then
3783 have individual miner payouts be proportional to the number of shares that each miner
3784 submitted during the round. Unfortunately, while this proportional reward allocation may
3785 be simple, it is also insecure due to the *pool-hopping attack* [190]. In fact, any pool reward-
3786 sharing scheme where the profitability of mining depends on the current state of the pool is
3787 potentially subject to pool-hopping.

3788 In a pool-hopping attack, a miner will join a pool, mine within it for some period of time,
3789 and then leave for another pool in such a way that they receive a disproportionate amount of
3790 revenue at the expense of other miners who remained in the pool. Under the proportional
3791 reward scheme, the payout for every share is equal to the block reward divided by the
3792 number of shares submitted in the round. As more time passes within a round, more shares

3793 are submitted, and thus the reward per share decreases. Therefore, an individual miner
3794 benefits if they submit shares to the pool during its shorter rounds but mine in other pools
3795 during longer rounds. This can be approximated simply by mining within a given pool for
3796 a short period and then switching to another pool if the round has not completed yet. As
3797 a result of this attack, honest miners who continuously participate in the pool can, in the
3798 worst case, receive 43% less than their "fair share." Unsurprisingly, mining pools no longer
3799 use the proportional payout scheme.

3800 Another attack that can be launched against mining pools is the *block withholding attack*,
3801 where a malicious miner submits valid shares to the pool operator but does not submit
3802 blocks when a full proof of work is found [190]. Block withholding can be especially nasty
3803 when mining pool operators are also miners. With multiple competing mining pools, the
3804 Nash equilibrium is a mixed strategy where miners use some of their hash rate to infiltrate
3805 competing pools to withhold blocks [191, 192]. This is analogous to the famous prisoner's
3806 dilemma in game theory: while everyone would be best off if all parties refrained from
3807 block withholding, each party benefits from their own malicious deviation. The result is
3808 that all parties will deviate and withhold blocks, and all parties will be worse off. If mining
3809 pools are attacking each other in this way, the total potential network hash rate will not be
3810 fully utilized, and the network will be less secure than it otherwise would be. If the network
3811 hash rate is static during a difficulty epoch, the attack only becomes profitable after the next
3812 difficulty adjustment. Until then, block withholding lowers the revenue for both the attacker
3813 and the victims, though it harms the victims to a greater degree. This has occurred in the
3814 real world, such as when the Eligius mining pool suffered a 300 bitcoin loss in June 2014
3815 due to block withholding [193]. However, the need to wait for a difficulty epoch before
3816 it becomes profitable and the low returns from the attack combine to make it relatively
3817 uncommon. A related but much stronger attack – *fork after withholding* – operates the
3818 same as block withholding, except that when the attacker detects that a competing block
3819 from a different pool is found, they release the withheld block in order to create a deliberate
3820 fork [194].

3821 It is possible to mitigate the threat of block withholding through an adjustment to the pool
3822 payout scheme. For example, the "incentive compatible" scheme described in [195] makes
3823 block withholding unprofitable. Let D be the expected number of shares found per full
3824 proof-of-work solution (i.e., D is the ratio between the full difficulty and the share diffi-
3825 culty). When a miner within a pool finds a block and submits it to the pool operator, D is
3826 compared with the number of shares submitted in that round, S . If $S \geq D$, then rewards are
3827 proportional. On the other hand, if $S < D$, each share receives a fixed reward of $\frac{1}{D}$, while
3828 the remainder of the available block reward goes to the miner who submitted the full solu-
3829 tion. By providing a larger reward for the miner who found the block, they are incentivized
3830 to report it. Unfortunately, this reward scheme is also susceptible to pool-hopping, as with
3831 the proportional scheme, and is therefore not appropriate for real-world use. Most existing
3832 mining pools follow one of three styles of alternative reward-sharing mechanisms:

3833 1. *Pay-per-share (PPS)*: Each submitted share immediately credits the miner with a
3834 fixed payout, and the pool operator keeps the full block reward when blocks are
3835 found. This minimizes variance-related risks for the miner [196] but turns the pool
3836 operator into a financier who takes on the full mining variance risk. During short
3837 rounds, the pool operator wins, whereas they lose money during long rounds in which
3838 many shares are submitted. Due to this risk, PPS pools tend to charge miners higher
3839 fees. PPS is fully immune to pool-hopping attacks.

3840 2. *Pay-per-last-N-shares (PPLNS)*: This method eschews the idea of rounds, and in-
3841 stead, the pool operator maintains a queue of the most recent N shares submitted to
3842 it. Whenever a miner within the pool finds a block, payouts are proportional to the
3843 number of shares each miner found that remains on the queue. A randomized version
3844 of PPLNS is also possible, where instead of a queue in which the oldest share is re-
3845 moved when each new share is found, a random one is removed [197]. While mining
3846 reward variance is higher for PPLNS than PPS, it has the lowest variance for schemes
3847 where the mining pool does not risk running out of funds and having a deficit with
3848 respect to the miners [196].

3849 This payout scheme is less susceptible to pool-hopping than the proportional method
3850 but is not hopping-proof. Miners can benefit from joining a pool immediately before
3851 the difficulty is about to decrease and leave a given pool when the difficulty is about
3852 to increase [190]. PPLNS is fairly robust, but there are situations where miners can
3853 benefit from strategic deviations, such as hoarding a certain number of shares and
3854 only submitting them to the pool upon finding a block. This can prevent old shares
3855 from exiting the queue [198].

3856 3. *Share-scoring methods*: This includes the original Slush method and the superior
3857 geometric method that it inspired, which is fully resistant to pool-hopping. These
3858 schemes give each newly submitted share a score that depends on how much time
3859 has elapsed since the beginning of the current round: as more time passes, the score
3860 increases. At the end of the round, payouts are proportional to the total score rather
3861 than the total number of shares. Early shares are worth more than late shares when
3862 using the proportional method, but the increasing scores over time help to counter
3863 this effect by reducing the value of early shares. The geometric method includes a
3864 variable fee that is higher earlier in the round but decays throughout and is parame-
3865 terized such that the score given to any new share is always the same relative to both
3866 existing and future scores. This removes any advantage that a miner could gain by
3867 timing when to enter or leave the pool [190].

3868 Note that the victim of a block withholding attack depends on the payout scheme used. If
3869 PPS is used, then the pool operator is the victim because they do not get to collect the block
3870 reward yet must still pay a constant amount to each miner for their shares. In methods that
3871 do not have operator risk, such as proportional, PPLNS, and geometric, the other miners
3872 within the same pool suffer when block withholding attacks are performed.

<pre> 1 on Init 2 public chain ← publicly known blocks 3 private chain ← publicly known blocks 4 privateBranchLen ← 0 5 Mine at the head of the private chain. 6 on My pool found a block 7 Δ_{prev} ← length(private chain) – length(public chain) 8 append new block to private chain 9 privateBranchLen ← privateBranchLen + 1 10 if Δ_{prev} = 0 and privateBranchLen = 2 then 11 publish all of the private chain 12 privateBranchLen ← 0 13 Mine at the new head of the private chain. 14 on Others found a block 15 Δ_{prev} ← length(private chain) – length(public chain) 16 append new block to public chain 17 if Δ_{prev} = 0 then 18 private chain ← public chain 19 privateBranchLen ← 0 20 else if Δ_{prev} = 1 then 21 publish last block of the private chain 22 else if Δ_{prev} = 2 then 23 publish all of the private chain 24 privateBranchLen ← 0 25 else 26 publish first unpublished block in private block. 27 Mine at the head of the private chain. </pre>	<pre> DEFINITIONS: effectiveState := sumWork(privateChain) – sumWork(mainChain) ifLose := effectiveState – nextMainChainBlockDifficulty INIT: effectiveState ← 0 ifLose ← –(nextMainChainBlockDifficulty) ON SELFISH MINER FINDS BLOCK: append new block to private chain and continue mining on private chain effectiveState ← effectiveState + newPrivateBlockDifficulty ifLose ← ifLose + newPrivateBlockDifficulty ON OTHER MINERS FIND BLOCK: append new block to public, main chain effectiveState ← effectiveState – newPublicBlockDifficulty ifLose ← ifLose – newPublicBlockDifficulty if ifLose ≤ 0 and len(privatechain) > 0 and effectiveState > 0 then publish private chain, overtake main chain, and mine on top of new public chain tip else if effectiveState = 0 and len(privatechain) > 0 then publish private chain and enter race else if ifLose > 0 then continue mining on private chain else if len(privatechain) = 0 then mine on top of new public chain tip end if </pre>
(a) Bitcoin DAA	(b) Variable DAA

Fig. 22. Selfish mining strategy. With a DAA that adjusts the difficulty every block, the *effectiveState* variable captures the difference in total work between the selfish miner’s private chain and the main chain. For Bitcoin’s DAA, it is sufficient to keep track of the difference in blocks instead of work. The *ifLose* variable represents what the difference in work would be if the honest miners won the next block. It is used to address situations in which the selfish miner currently has the most work chain but would give up if honest miners win the next block. [185, 199]

3873 9.4. Selfish Mining

3874 Satoshi Nakamoto believed that in the Bitcoin system, it is in miners’ rational interest
 3875 to behave honestly so long as the adversary does not control more than half of the net-
 3876 work’s computational power [1]. This was later shown to be false due to the *selfish mining*
 3877 phenomenon, where strategically withholding blocks from the network can be more prof-
 3878 itable than honest mining, even with only a minority of the network’s hash rate [199–201].
 3879 While the original attack was described assuming constant difficulty blocks, the strategy
 3880 was expanded to account for variable difficulty blocks in [185]. The attack is relevant to
 3881 Nakamoto Consensus (see Section 10) but also applies to longest-chain proof of stake and
 3882 several other proof of work protocols.

3883 Selfish mining works by forcing honest miners to waste their energy mining blocks on a
 3884 chain that is "destined" to be reorganized and become stale. When the attacker success-
 3885 fully mines blocks, they keep them private and only reveal them to the rest of the network
 3886 when the honest chain is closer to being caught up with the adversarial chain. The attack
 3887 allows the selfish miner to gain a disproportionate share of the mining rewards (i.e., it re-
 3888 duces chain quality and improves the adversary’s *relative revenue*) and becomes profitable
 3889 sometime after the difficulty level drops to reflect the reduced block production rate. The
 3890 detailed strategy is listed in Figure 22a for the case of Bitcoin’s DAA (or any DAA where
 3891 the difficulty remains constant for an epoch before changing) and in Figure 22b when the
 3892 difficulty changes with every new block.

3893 The attack is parameterized by two variables, α and γ , with $0 \leq \alpha, \gamma \leq 1$. α is the fraction of
3894 the total hash rate that is engaged in selfish mining, and γ represents the fraction of honest
3895 miners that will add blocks to the chain on top of the selfish miner's block during a block
3896 race (and thus relates to the adversary's control over the network). γ can be considered a
3897 measure of the rushing capabilities of the adversary in that it captures the advantage that
3898 the selfish miner has when reacting to blocks found by honest peers.

3899 Whether the selfish mining strategy is rational or not depends on the values of α and γ . The
3900 hash rate where a selfish miner can improve their relative revenue above and beyond what
3901 they would earn from mining honestly is given in the following inequalities:

$$\frac{1 - \gamma}{3 - 2\gamma} < \alpha < \frac{1}{2} \quad (1)$$

3902 If $\gamma = \frac{1}{2}$, selfish mining is profitable when $\alpha \geq \frac{1}{4}$, and if $\gamma = 0$, selfish mining is profitable
3903 when $\alpha \geq \frac{1}{3}$. The $\gamma = \frac{1}{2}$ scenario corresponds to the situation in which miners – upon
3904 seeing two equal-work chain tips – choose randomly between them. Further, if $\gamma = 1$,
3905 selfish mining is beneficial to the attacker at any hash rate. In the optimized attack from
3906 [202], the required adversarial hash rate is decreased from 25% to 23.21% when $\gamma = \frac{1}{2}$.

3907 Nayak et al. proposed three additional "stubborn" variants of the strategy, which may
3908 perform better than selfish mining under some conditions [203]:

- 3909 1. *Lead stubborn*: When a selfish miner has a lead of two blocks, but the honest miners
3910 find a block and cut the lead to one, the selfish miner publicizes their longer chain to
3911 win the race. In contrast, a lead-stubborn miner will only publicize the first block to
3912 "match" the honest chain and try to cause a race.
- 3913 2. *Equal-fork stubborn*: In normal selfish mining, when there is a block race between
3914 equal length forks and the attacker mines a block, they reveal it immediately to win
3915 the fork. In equal-fork stubborn mining, they would withhold the new block and just
3916 have a one block lead again.
- 3917 3. *Trail stubborn (T_j -stubborn)*: When a selfish miner falls behind the honest chain,
3918 they simply adopt the honest chain. A T_j -stubborn miner will continue mining on
3919 their private chain until they are $j + 1$ blocks behind the honest chain.

3920 Selfish mining is one of the most well-studied attacks in the permissionless ledger litera-
3921 ture and has been investigated under a wide variety of scenarios. For example, heteroge-
3922 neous network connectivity among honest miners helps a selfish mining attacker because γ
3923 rapidly increases as the variance in block propagation delay increases [204]. In [205], self-
3924 ish mining was found to be far more profitable with larger network delays, but technologies
3925 that speed up block propagation – such as compact blocks and relay networks – effectively
3926 eliminate this difference for blocks of moderate size (compact blocks and relay networks
3927 are discussed in Section 10.2.1).

3928 Most selfish mining research assumes the existence of a single adversary, but a line of
3929 work investigates what happens with multiple selfish miners simultaneously operating on
3930 the same network [206–209]. Unsurprisingly, the security of the network is further re-
3931 duced with multiple selfish miners. For example, with two independent selfish miners, the
3932 threshold for profitably selfish mining drops from 25% to 21.48% [206]. As the number of
3933 simultaneous selfish miners increases, the profitability threshold for each miner decreases
3934 [207, 209]. A necessary condition for n selfish miners to simultaneously benefit from the
3935 attack is that each of their individual hash rates α_i is in the range $\frac{1}{n+1} < \alpha_i < \frac{1}{n-2}$, such that
3936 up to seven adversaries can operate with 12% of the hash rate each [209].

3937 The profitability of selfish mining depends on the network against which it is applied. The
3938 DAA is a significant factor. For example, in systems where the difficulty is able to rapidly
3939 adjust downward when blocks are slow, selfish mining tends to be more profitable. When
3940 the DAA only looks at a relatively short period of time as the basis for the adjustment,
3941 selfish miners can dramatically increase their profits by manipulating the timestamps they
3942 include in block headers [185].

3943 Ethereum, in particular, has received the most attention outside of Bitcoin [210–212].
3944 Ethereum’s consensus algorithm, GHOST (discussed in Section 11.2), differs from Bit-
3945 coin’s in that it includes the existence of stale blocks ("uncles") as part of the DAA as
3946 well as the reward scheme. Most significantly, Ethereum miners are rewarded even if
3947 their blocks do not make it into the main chain so long as they are referenced as uncle
3948 blocks shortly after being mined. As a result, the threshold for selfish mining profitabil-
3949 ity is reduced because even when the attacker loses a block race, they are likely to still
3950 get some reward for taking the risk. Using the observed uncle block ratio from Ethereum
3951 in December 2017, [211] showed that the profitability threshold for selfish mining was
3952 $\alpha = 0.185 \pm 0.012$. Another model found that selfish mining was profitable with $\alpha > 0.163$,
3953 and that beneath this value, the selfish miner loses far less than they would by attacking Bit-
3954 coin [212].

3955 Despite being perhaps the most significant theoretical attack on many permissionless sys-
3956 tems (aside from the majority attack), significant selfish mining attacks have not been ob-
3957 served in the real world. There are a number of reasons why this may be the case:

- 3958 • Executing the selfish mining strategy requires specific software for implementation.
3959 The expertise to develop this is not widespread and the software itself is likely to be
3960 bug-prone and challenging to test. Due to the costs imposed by proof of work, a bug
3961 could cause tens or hundreds of thousands of dollars in losses for the attacker.
- 3962 • Most models assume that the exchange rate for the asset being mined selfishly re-
3963 mains constant throughout the duration of the attack and beyond. However, if selfish
3964 mining was discovered on a network, one might reasonably expect it to decrease the
3965 exchange rate of the asset in question, which can dramatically reduce the profitability
3966 of the attack (especially if the relevant hardware cannot be repurposed).

- 3967 • For larger networks, like Bitcoin, it would be prohibitively expensive for individ-
3968 ual miners to acquire enough computational power to profitably perform the attack.
3969 While mining pools can perform selfish mining, it is likely to be discovered quickly,
3970 at which point the individual miners can defect to honest pools. Unfortunately, these
3971 miners are incentivized to remain in the pool and take advantage of the profits that
3972 come from selfish mining so long as the exchange rate remains constant (but this
3973 assumption is questionable). Selfish mining detection is facilitated by the common
3974 behaviors of individual miners belonging to multiple pools, as well as pools moni-
3975 toring each other.
 - 3976 • If enough computational power can be acquired to perform selfish mining, it is often
3977 easier and perhaps more profitable to double-spend with a majority attack. This is
3978 likely the case for small networks where markets exist to rent significant portions of
3979 the hash rate.
 - 3980 • The parameter γ is important to the profitability of the strategy, but estimating it is
3981 challenging, and it is likely to change throughout the duration of the attack.
 - 3982 • Selfish mining is not profitable until sometime after the difficulty adjusts downward
3983 to reflect the adversary "dropping off" the network. Most research assumes that no
3984 new miners are enticed to (honestly) mine on the network based on this lower diffi-
3985 culty, which may not be realistic. If this assumption turns out to be false, the selfish
3986 miner may never recoup their losses from prior to the difficulty adjustment.
- 3987 On the other hand, selfish mining immediately lowers the profitability of honest miners. If
3988 this induces some of them to quit mining, selfish mining becomes even more profitable for
3989 the attacker [213].

3990 10. Nakamoto Consensus

3991 Satoshi Nakamoto first proposed what has been dubbed "Nakamoto Consensus" in 2008,
3992 and it ushered in a new era of consensus algorithm research and innovation [1]. The term
3993 "Nakamoto Consensus" is used frequently and inconsistently in the consensus literature.
3994 Sometimes, it is used merely to describe the longest chain rule without specifying what
3995 mechanism is used to prevent Sybil attacks. However, that is not the definition used here.
3996 In this document, Nakamoto Consensus is defined by two things:

- 3997 1. Proof of work is used as the Sybil-resistance mechanism and for randomized leader
3998 election.
- 3999 2. The longest chain rule (LCR) is used as the fork-choice rule to form a blockchain.
4000 Technically, the LCR is a misnomer because the "longest" chain is defined as the
4001 heaviest one or the one that proves the most work rather than the chain with the
4002 greatest number of blocks.

4003 Nakamoto Consensus is most prominently used in Bitcoin but is also one of the more
4004 commonly deployed consensus models overall. In the permissionless design space, it is by
4005 far the most well-understood protocol with formal security proofs under a wide variety of
4006 models. Having been deployed in Bitcoin since January 2009, Nakamoto consensus has
4007 the most real-world battle testing of any permissionless protocol. The assumptions upon
4008 which its security is derived are simpler than most:

- 4009 1. **Honest majority.** The majority of the work that can be applied to the network is
4010 being used by honest participants rather than malicious ones. See Section 10.2.2 for
4011 more details.
- 4012 2. **Bounded network delay.** Messages that are sent by participants must propagate to
4013 the rest of the network within some bounded period of time. See Section 10.2.1 for
4014 more details.
- 4015 3. **Collision-resistant hash function.** A collision-resistant cryptographic hashing func-
4016 tion is required in order to maintain the integrity of the blockchain itself and ensure
4017 that each block acts as a commitment or vote on all preceding blocks in the chain.

4018 Nakamoto Consensus is also very efficient in terms of minimizing communication over-
4019 head. Other protocols often require the exchange of additional information regarding which
4020 validators have already seen which blocks or allow duplicate or invalid transactions inside
4021 of blocks that are only later filtered out. In addition, the simple blockchain structure – as
4022 opposed to multiple parallel chains or DAGs – ensures that transactions are globally or-
4023 dered as soon as blocks are generated. This reduces transaction confirmation latency and
4024 is compatible with all smart contract programming models. In contrast, using a different
4025 structure requires some care to ensure that the transactions are totally ordered so as to be
4026 suitable for smart contracts. As such, it would not make sense to describe other, more
4027 recent consensus algorithms as strict improvements over Nakamoto Consensus but rather
4028 as having trade-offs. Proof-of-work variants that stay close to Nakamoto Consensus are
4029 described in Section 11.1, while a wider variety are discussed throughout Section 11.

4030 The mechanics of Nakamoto Consensus are simple. All nodes begin with the *genesis block*,
4031 a common reference string that initializes the state of the system. Miners then search for
4032 proof-of-work solutions to a random puzzle that includes the genesis block as an input (as
4033 well as a list of transactions and some other metadata). A miner who finds a proof-of-work
4034 solution is elected leader and produces a block, which is then published and gossiped over
4035 a peer-to-peer network. When other nodes see this block, they validate the proof of work
4036 and the transactions, forward the block to their peers, and begin trying to find solutions to
4037 the next puzzle using the new block as input. Over time, this builds a chain of blocks that
4038 extend from the genesis block, adding transactions to the ledger and updating the system
4039 state.

4040 Due to random chance, network delays, or malice, it is possible for more than one com-
4041 peting proof of work to be found at the same height in the blockchain. For instance, if

4042 Alice and Bob both find blocks at height one immediately after the genesis block, then it
4043 is unclear whether miners should consider legitimate and thus build new blocks on top of
4044 Alice's chain or Bob's chain. In this case, there is a tie, and miners can choose arbitrarily
4045 (though they typically prefer the one they saw first). At some point, a miner will extend
4046 the chain from either Alice's block or Bob's block, and then that chain will have the most
4047 work backing it. This breaks the tie, and all honest nodes will switch to the chain with the
4048 new block because honest nodes follow the chain with the most work supporting it. This
4049 process continues indefinitely with honest miners building blocks that extend the longest
4050 (most work) chain they are aware of and results in a growing ledger of transactions.

4051 **10.1. Theory of Nakamoto Consensus**

4052 Nakamoto Consensus (and sometimes just Bitcoin) is frequently brought up as a solution to
4053 the Byzantine Generals Problem (see Section 1.1), thus making it a broadcast algorithm. At
4054 other times, it is presented as an agreement algorithm. Despite these frequent contentions, it
4055 is not strictly true that Bitcoin and Nakamoto Consensus perfectly fit into these paradigms.

4056 In particular, the "validity" property of BGP and BA does not hold for Nakamoto Consen-
4057 sus. Recall that BGP validity guarantees that if the leader is honest and begins with input
4058 value v , then all honest nodes decide v . BA validity guarantees that if all honest nodes
4059 propose the same value, then any correct node must decide v . There is also weak validity,
4060 which requires that each honest node's output must be the input of some honest node. The
4061 randomized proof-of-work process interferes with achieving these notions of validity. Un-
4062 less the adversary has a negligible fraction of the total computing power of the network,
4063 they will eventually be first to produce a proof of work, at which point all honest nodes will
4064 immediately switch to and extend the chain from this adversarial block and abandon the
4065 original "honest" input.

4066 Similarly, the probabilistic guarantees provided by Nakamoto Consensus as a result of
4067 its proof-of-work process can be framed as providing agreement but not termination or,
4068 alternatively, termination but not agreement [214]. Ultimately, Nakamoto Consensus sat-
4069 isfies "eventual consistency" or "probabilistic consistency," where the probability of agree-
4070 ment on a particular block at a particular position increases exponentially as additional
4071 blocks extend the chain's tip from the block in question. One can, therefore, consider
4072 Nakamoto Consensus to satisfy agreement but not termination: if the blockchain protocol
4073 executes forever, then agreement will hold probabilistically with probability 1. Alterna-
4074 tively, Nakamoto Consensus may satisfy termination but not agreement: if x blocks exist
4075 and the system runs for a finite duration of k blocks, then consensus is achieved for block
4076 x , but the probability that agreement is satisfied is less than 1 and grows exponentially with
4077 k .

4078 By now, there is a considerable amount of literature proving that Nakamoto Consensus and
4079 Bitcoin provide particular security properties under a variety of assumptions [9, 165, 182–
4080 184, 215–230]. Ultimately, Nakamoto Consensus creates a state machine that operates

4081 over a distributed ledger that satisfies *persistence* and *liveness* with security parameter k ,
4082 sometimes called the "Bitcoin Backbone."

- 4083 • **Persistence.** Once a transaction is at least k blocks "deep" in the blockchain of
4084 an honest node, then with overwhelming probability, the same transaction will be
4085 included at the same position of every other honest node's blockchain and remain
4086 there permanently.
- 4087 • **Liveness.** Transactions sent by honest clients will eventually be at a depth of more
4088 than k blocks in honest nodes' blockchains.

4089 The first security proof for Nakamoto Consensus in synchronous networks is from [9],
4090 where persistence and liveness follow from the *common prefix* and *chain quality* properties.
4091 It was then shown in [215] that an additional *chain growth* property was, in fact, required
4092 for liveness.

- 4093 • **Common prefix.** For a security parameter k , the blockchains of two honest nodes
4094 differ only in the last k blocks from the tip. That is, if one were to "chop off" the
4095 final k blocks of two honest nodes' chains, one of the resulting subchains would be
4096 a prefix of the other. A related requirement is *future self-consistency*, such that at
4097 any two points in time, with high probability in the security parameter T , the chain
4098 of an honest node differs only within the last T blocks (i.e., alternating between two
4099 completely different chains is not allowed).
- 4100 • **Chain quality.** For parameters μ and T , it holds that for any T consecutive blocks in
4101 the blockchain of an honest party, the ratio of honest blocks is at least μ . Stated differ-
4102 ently, it should be the case that "enough" of the blocks that end up in the blockchain
4103 were proposed by honest nodes.
- 4104 • **Chain growth.** For parameters τ and s and any honest party P with chain C in a
4105 given view, it holds that for any s rounds, there are at least $\tau * s$ blocks added to the
4106 chain of P . In other words, the blockchains of honest nodes must continuously grow
4107 at a certain pace over time.

4108 The model used in [9] assumed a fixed set of n equally powerful miners (but with n un-
4109 known), where each party can make the same number of queries, q , to a hash function
4110 modeled as a random oracle. This is considered a "flat" model because each miner is equal,
4111 but one can imagine differences as being represented by real-world entities controlling vari-
4112 able numbers of players. The adaptive, rushing adversary controls $f < \frac{n}{2}$ of those parties
4113 and can thus make $f * q$ proof-of-work queries per round. In this environment, if $\frac{f}{n-f} < 1$,
4114 honest parties will have blockchains with large common prefixes when pruning the most
4115 recent k blocks from the tip of the chain, except for some small probability that drops ex-
4116 ponentially with k . Further, the blockchains of honest parties will include blocks from both
4117 honest and malicious parties, but so long as the majority of the hash power is honest, the
4118 ratio of blocks produced by honest parties compared to malicious parties is bounded by

4119 $\frac{f}{n-f}$. This chain quality is fairly low unless there is a large majority of honest miners (see
4120 Section 9.4 on selfish mining for one adversarial strategy that reduces chain quality). For
4121 example, if an adversary controls $\frac{1}{3}$ of the hash power, they may end up contributing nearly
4122 half the blocks. Finally, if the network is unable to remain sufficiently synchronized (i.e.,
4123 if the message latency between honest miners approaches the expected time that it takes
4124 to find proof-of-work solutions), then maintaining the common prefix property requires a
4125 larger and larger hash rate (super)majority.

4126 The analysis in [216] provides security proofs in the " Δ -bounded delay" model. This model
4127 comports with the standard definition of partial synchrony presented in [26], where there
4128 exists a fixed, unknown upper bound Δ on message delay. However, it does not fit the game-
4129 based description that immediately followed in [26], where the protocol designer designs
4130 a protocol and the adversary then chooses Δ . Interestingly, this makes systems that depend
4131 on Nakamoto Consensus, like Bitcoin, behave like asynchronous or partially synchronous
4132 protocols while still technically being synchronous.

4133 It is instructive to consider why it is more challenging to address network delays in a proof-
4134 of-work model than in most permissioned systems. To transform a synchronous permis-
4135 sioned protocol into a Δ -bounded delay one, honest replicas can simply wait Δ time steps
4136 before responding to any messages. In the permissionless setting using proofs of work, the
4137 adversary can increase its use of computational resources by a factor of Δ to try to solve
4138 puzzles while the honest miners "wait."

4139 In the context of Nakamoto Consensus, the delay must not only be bounded but sufficiently
4140 bounded. An arbitrary but finite delay will not suffice. If the delay were truly arbitrary (as
4141 it would be in an asynchronous network), an adversary with a small fraction of the hash rate
4142 could simply delay the arrival of honest parties' messages long enough to guarantee that
4143 the adversary creates an even longer chain than the ones acknowledged by honest nodes,
4144 forcing honest nodes to adopt the adversarial chain. Without any adversarial hash power,
4145 it is still possible for this network adversary to cause common prefix violations. However,
4146 in a model where the adversary can delay messages with only some probability less than 1,
4147 Nakamoto Consensus can remain secure even when delays may be quite a bit larger [217],
4148 which may motivate the use of satellite or radio technologies to broadcast blocks in ways
4149 that are more difficult for adversaries to disrupt.

4150 Even under partial synchrony, if ρ is the fraction of the hash power controlled by the ad-
4151 versary, n is the number of miners, and p is a mining hardness parameter that captures
4152 the probability that a single random oracle hash query generates a valid proof of work, the
4153 common prefix property (and thus persistence) will not hold when $p > \frac{1}{n\rho\Delta}$.

4154 As a result, the mining puzzle's difficulty must be appropriately set as a function of the
4155 maximum network delay in order to maintain security. The primary result from [216] is
4156 that so long as $\rho < \frac{1}{2}$, then for any delay bound Δ , there exists a sufficiently small p such
4157 that Nakamoto Consensus maintains consistency. The consistency proof relies on a con-
4158 cept dubbed *convergence opportunities*, which provide honest nodes the opportunity to

4159 synchronize themselves and converge upon the same blockchain. A convergence opportu-
4160 nity occurs when:

- 4161 1. There is a period of Δ rounds where no honest miners mine a block.
- 4162 2. A round occurs where a single honest miner mines a block.
- 4163 3. Another period of Δ rounds occur without any honest miners mining a block.

4164 Note that in step 2, only a single honest miner can mine a block during that time. If mul-
4165 tiple honest miners mine competing blocks, there is no convergence opportunity because
4166 honest miners may be split over which block they ought to mine on top of (later works
4167 tightened the security proof by relaxing this requirement [225, 226]). To prove the security
4168 of Nakamoto Consensus, [216] shows that convergence opportunities occur with sufficient
4169 frequency and that attackers can only prevent them by mining a block that is accepted by
4170 the honest nodes during this time. An adversarial block that prevents a convergence op-
4171 portunity cannot have been mined and withheld particularly long before the time period in
4172 question because the honest mining majority will quickly produce and distribute a longer
4173 chain, making the adversarial block irrelevant. When the adversarial hash rate is suffi-
4174 ciently bounded, the number of adversarial blocks will be outnumbered by the frequency
4175 of convergence opportunities, guaranteeing that honest nodes will eventually converge on
4176 a consistent chain.

4177 10.1.1. Nakamoto Consensus With Chains of Variable Difficulty

4178 The above analyses assumed that the mining difficulty remained constant throughout the
4179 execution of the protocol. In actuality, the difficulty adjusts with the addition or removal
4180 of computing power so as to maintain an average target block interval (for instance, in
4181 Bitcoin, the target block interval is 10 minutes or 600 seconds). Several works extend the
4182 prior analyses to distributed ledgers where the difficulty changes over time [182–184].

4183 In [182], the authors extend the synchronous model from [9] to account for Bitcoin’s dif-
4184 ficulty adjustment algorithm, which recalibrates the difficulty parameter after every 2,016
4185 blocks (approximately two weeks) in order to account for changes in hash rate during that
4186 time. The new model does away with the fixed number of miners n and instead allows
4187 the number of miners to vary. There are bounds on how quickly the number of players
4188 can change over time while maintaining consistency, and if these bounds are respected and
4189 miners do not deviate too far from expectation, the block production rate should remain
4190 sufficiently stable to maintain the common prefix and chain quality properties. The most
4191 crucial result from the analysis is that the length of a difficulty epoch needs to be suffi-
4192 ciently large in order to bound the probability of attacks, which helps to justify the lengthy
4193 epochs used in Bitcoin (while raising questions as to the security of more rapidly adjusting
4194 DAAs). Further, Bitcoin’s DAA utilizes "dampening filters" such that the difficulty adjusts
4195 by no more than a factor of four in either direction, and this dampening is required in order
4196 to prevent attacks against the common prefix property [200].

4197 Some limitations from [182] are addressed in other papers that account for variable dif-
4198 ficulty [183, 184]. In particular, [182] relies on two key assumptions: 1) a synchronous
4199 network prevents the adversary from delaying network messages between honest parties,
4200 and 2) the adversary cannot adaptively choose how many players – and thus how much
4201 hash rate – is deployed in any given round but must instead schedule this in advance.
4202 Clearly, these assumptions do not hold in the real world, but [183, 184] prove the se-
4203 curity of Nakamoto Consensus with an adaptive work schedule and Δ -bounded message
4204 delays. Note that changes in computing power must not occur too quickly or else the min-
4205 ing difficulty will not be appropriately set relative to the network delay. The only other
4206 condition required in the proof from [183] is that the initial difficulty is set appropriately to
4207 the network delay and available computational power. That is, if the difficulty is properly
4208 calibrated from the start and does not change too quickly, Nakamoto Consensus remains
4209 secure.

4210 The difficulty adjustment algorithm employed in [183] differs slightly from that of Bitcoin.
4211 In particular, for an additional security parameter κ_0 , the difficulty calculation "chops off"
4212 the first and last κ_0 blocks of the epoch, which helps maintain consistency and avoids the
4213 need to deal with epoch boundaries in the analysis. In addition, the timestamp validity
4214 rules for a given block differ slightly from that of Bitcoin: all timestamps must strictly
4215 increase, and honest nodes reject blocks where the timestamp is in the future. Assuming
4216 that the chain quality property holds, these rules ensure that an honestly produced block
4217 occurs periodically, forcing adversarial blocks to be located between two nearby honest
4218 blocks and thus preventing adversarial blocks' timestamps from deviating too far from the
4219 actual time. The analysis in [184] uses the real Bitcoin DAA but with similarly modified
4220 timestamp rules.

4221 10.1.2. Additional Analyses of Nakamoto Consensus

4222 There are quite a few additional security analyses of Nakamoto Consensus and Bitcoin that
4223 often use different models. For example,

- 4224 • Bitcoin was proven secure with *universal composability* (UC) in [218] in networks
4225 with bounded delays. A UC-model proof is important because it means that a Bitcoin-
4226 like ledger can be arbitrarily composed with other protocols that may rely on it, such
4227 as running multi-party computations on top of a blockchain.
- 4228 • The Bitcoin Backbone was proven secure in the quantum random oracle model in
4229 [220, 231], assuming a sufficiently bounded quantum adversary.
- 4230 • In [221], the Bitcoin Backbone was proven secure when different miners have differ-
4231 ent hash rates (i.e., it avoids the "flat" model where each miner is considered to have
4232 an equal amount of hash power).
- 4233 • The Bitcoin Backbone remains secure even when some of the nodes "prune" their
4234 blockchains by removing old blocks from their local ledger [224].

- 4235 • A simple proof of the security of Nakamoto Consensus in the continuous time model
4236 (instead of using discrete rounds) was provided in [229].
- 4237 • The trade-off between confirmation latency and security was formalized in [232],
4238 which provides guidance for optimizing throughput by choosing appropriate param-
4239 eters for block size and expected block times.
- 4240 • Permissionless longest-chain protocols, including Nakamoto Consensus, were proven
4241 secure in a generalization of the synchronous communication model, where message
4242 delays are independent, identically distributed, and thus potentially unbounded [233].
- 4243 • Some works have attempted to remove the random oracle assumption from their se-
4244 curity proofs. For example, [223] formalizes a primitive called "signatures of work,"
4245 which enables honest majority consensus in permissionless settings analogously to
4246 how digital signatures allow honest majority consensus in permissioned settings. The
4247 "key" for these "signatures" is the most recent block hash seen by a miner, so the key
4248 ensures timeliness of the work. Unfortunately, the only known instantiations of the
4249 primitive rely on random oracles themselves. A standard model proof of security
4250 without the use of a random oracle for the Bitcoin Backbone protocol was provided
4251 in [165].
- 4252 • Several works have tightened the consistency bounds for longest-chain protocols like
4253 Nakamoto Consensus [222, 225, 226, 230]. One reason why earlier bounds were not
4254 tight was due to the negative influence on security that occurs when multiple honest
4255 miners mine blocks close to each other (see the discussion on convergence opportu-
4256 nities in Section 10.1), and [226] uses proof techniques that mitigate the impact of
4257 this. The resulting consistency bound – where r_a is the expected number of adver-
4258 sarial proofs of work generated per unit time, r_h is the expected number of honest
4259 proofs of work, and the maximum message delay is Δ_0 – is

$$r_a < \frac{1}{\Delta_0 + \frac{1}{r_h}} \quad (2)$$

4260 The analysis in [222] produced security bounds in terms of the expected number of
4261 network delays until a block is mined, c (e.g., if it takes 10 seconds before all miners
4262 receive and validate a block and the expected block interval is 600 seconds, $c = 60$).
4263 In this case, if h is the honest fraction of the hash rate and f is the adversarial fraction,
4264 consistency is maintained so long as

$$c > \frac{2h}{\ln(\frac{h}{f})} \quad (3)$$

- 4265 • A particularly interesting result is that the Bitcoin Backbone can maintain its primary

4266 security properties (i.e., chain quality, chain growth, and common prefix) even in
4267 situations where a dishonest majority temporarily exists so long as there is an honest
4268 majority in expectation over time [227, 228]. This would capture situations where,
4269 for example, a mining pool has suffered a denial of service that creates a temporary
4270 dishonest majority of miners that will be rectified when the pool comes back online.

4271 One more analysis that deserves special attention is [219], which uses the game-theoretic
4272 rational protocol design framework to show that – assuming a particular class of incentives
4273 within the Bitcoin protocol – the honest majority assumption can be replaced by the weaker
4274 assumption that miners seek to maximize their profits. The utility function that accounts
4275 for these incentives incorporates the block subsidy, transaction fees, and mining costs but
4276 does not include any incentives external to the protocol, such as exploiting the transaction
4277 ordering for front-running or shorting the bitcoin asset in order to profit from a decrease in
4278 its price. A number of important conclusions can be drawn from this:

- 4279 • In the less realistic model where transaction fees are excluded from the analysis,
4280 Bitcoin is incentive-compatible, and even large coalitions of miners lack an incentive
4281 to deviate from the protocol so long as the price of Bitcoin is high enough for mining
4282 to be profitable.
- 4283 • When transaction fees are included, the system remains incentive-compatible so long
4284 as fees are sufficiently large. However, the same is not true for arbitrary transaction
4285 fee distributions. If some fees are much higher than others, an attacker may try to
4286 build two chains in parallel with the high fee transactions and maintain the fork long
4287 enough to spend those fees on both chains. On the other hand, if an honest majority of
4288 computing power exists, the only profitable deviation by the adversary is to withhold
4289 high-fee transactions from the rest of the miners in order to claim them without the
4290 risk of a competitor collecting those fees. As a result, the security assumption for
4291 Bitcoin can be relaxed from needing an honest majority to needing many high fee
4292 transactions and only requiring an honest majority when fees are low.
- 4293 • As the block subsidy in Bitcoin is reduced over time and miners increasingly rely
4294 on transaction fees for revenue, there must be a significant transaction fee backlog
4295 in order to maintain incentive compatibility. Specifically, the total reward for mining
4296 blocks must be non-decreasing. That is, after a block is mined, other miners must
4297 have enough transaction fees available in their mempools to be incentivized to move
4298 the chain forward instead of creating a fork that includes more fees. Ultimately, this
4299 motivates having a maximum block size limit that is small enough that miners cannot
4300 build blocks that claim "too much" of the outstanding transaction fees. This issue is
4301 discussed in more detail in Section 19.1.1.

4302 10.2. Violating the Nakamoto Consensus Security Assumptions

4303 Roughly speaking, there are three major ways that Nakamoto Consensus can fail. The
4304 following sections look at each of these concerns in turn:

- 4305 1. The network delay is not sufficiently bounded such that blocks take "too long" to
4306 propagate.
- 4307 2. The majority of the computing power deployed to the network can be malicious.
- 4308 3. An adversary is able to create hash function collisions.

4309 10.2.1. Network Delay and Block Propagation

4310 Section 10.1 described how an adversary who can arbitrarily delay the transmission of
4311 blocks to honest miners would be capable of causing consensus failure. Whether adversarially
4312 induced or simply a result of typical network latency, any delay in block propagation is
4313 detrimental to Nakamoto Consensus. Specifically, latency in block propagation can cause
4314 the blockchain to fork. Two different miners may produce blocks at the same height of the
4315 chain at approximately the same time because they were both unaware of the competing
4316 block. Whenever forks occur, honest miners temporarily work on different problems, so
4317 the actual work that secures the chain is reduced.

4318 The creation of new blocks is a Poisson process, which means that the times between
4319 two blocks being mined follow an exponential distribution. Thus, the probability of an
4320 unintentional fork in the blockchain is $1 - e^{-\frac{x}{T}}$, where x is the propagation delay and T is
4321 the expected block interval. For example, with Bitcoin's 600 second block interval and
4322 a 10 second propagation delay, there is approximately a 1.7% chance of conflicting forks
4323 appearing on the network when a block is found, so forks would be expected every 60
4324 blocks on average. If the block interval were one minute instead, there would be more
4325 than a 15% chance of forks or forks every 6 or 7 blocks. According to an early study
4326 on Bitcoin's block propagation, it took 12.6 seconds for a block to propagate to 90% of
4327 the network [234]. However, that was prior to the deployment of techniques that reduced
4328 latency, such as compact blocks and the Fast Internet Block Relay Engine (FIBRE). Since
4329 then, accidental forks have become extremely rare.

4330 The concerns regarding frequent forks relate to more than just maintaining consistency
4331 among nodes. Due to the permissionless nature of Nakamoto Consensus, adequate incen-
4332 tives are required for the system to work properly. When forks occur, only one of the
4333 competing blocks will actually remain in the chain for the long term, while the others are
4334 discarded and become stale, leaving their creators with no reward. Unfortunately, this im-
4335 plies that frequent forks can lead to mining power becoming concentrated among a smaller
4336 and smaller set of larger and larger entities. For example, consider a fork that occurs be-
4337 tween a miner with 35% of the network's hash rate and a miner with only 5%. Both send
4338 their block out to the rest of the network and immediately begin mining on top of their own

4339 block. In this case, the larger miner instantly has an additional 30% of the network hash
4340 rate supporting their block compared to the small miner and thus has a 30% lower stale
4341 block rate. If it takes 30 seconds for a block to propagate to the rest of the network, there
4342 would be an approximately 4.88% chance of a conflict, so the larger miner gains an overall
4343 1.46% revenue advantage. Difficulty adjustments make mining a low-margin business, so
4344 this revenue advantage can be substantial. In other words, big miners will create larger
4345 and larger blocks in order to squeeze out the competition over time (in the absence of a
4346 sufficiently small maximum block size limit).

4347 Similar centralization issues arise based not just on the relative size of the miners but also on
4348 their relative network connectivity [235–237]. Miners who have low latency connections
4349 to other miners have the same types of advantages during forks that large miners do – the
4350 remaining miners are more likely to build off of the well-connected miner’s block because
4351 they will see it first. In addition, better latency improves a miner’s ability to include high fee
4352 transactions in their blocks before the transaction has fully propagated through the network.

4353 A related concern is that miners are not strictly incentivized to get their blocks propagated
4354 to 100% of the network hash rate as quickly as possible but rather somewhere above 50%
4355 and less than 100% (see Section 9.4 on selfish mining, which is a related problem). This
4356 implies that crafting big blocks that reach the majority quickly but slowly propagate to the
4357 rest can harm small miners. It also implies that networking impediments like the Great
4358 Firewall of China are problematic. If the majority of the hash rate is on one side of the
4359 firewall, that side has a significant advantage over the miners on the other side. Miners can
4360 intentionally slow down block propagation in a few ways, such as stuffing blocks full of
4361 transactions that have not been forwarded to the rest of the network or including transac-
4362 tions in blocks that are deliberately slow to validate (see Section 19.2 for more information
4363 on slow to validate transactions).

4364 Improving block propagation significantly reduces security issues with Nakamoto Consen-
4365 sus. One way of doing this is to increase the number of network connections that each node
4366 has, though that comes with the cost of increased bandwidth consumption [238]. There are
4367 also a variety of techniques that can be used to reduce block propagation time. Only com-
4368 pact blocks and the FIBRE relay network are described here because they are used on the
4369 Bitcoin network, though a variety of other block transmission mechanisms have been pro-
4370 posed (e.g., Xthin, Xthinner, Graphene, BlockTorrent, Falcon relay network, etc.). While
4371 important, these techniques have some limitations:

- 4372 • They rely on the assumption that nodes have relatively synchronized mempools and
4373 perform poorly when this is not the case. That is, if there are many transactions
4374 that are included in a block but have not propagated to all of the miners, the extra
4375 network round trips and bandwidth of transmitting those transactions will reduce the
4376 effectiveness of compact blocks and increase latency and orphan rates. As block sizes
4377 increase and expected block intervals decrease, more transactions will exist that have
4378 not propagated fully.

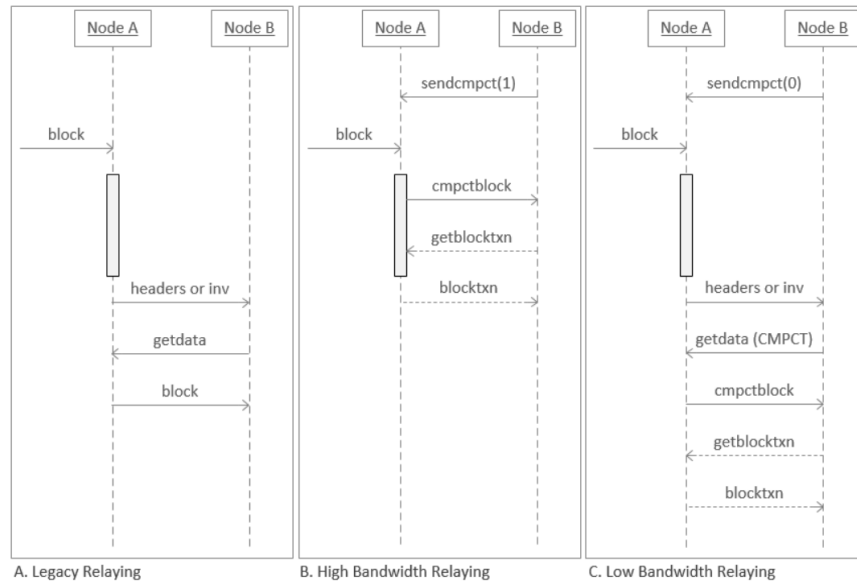


Fig. 23. Compact block message flow in which node A receives a block and sends it to node B. [239]

- 4379 • While they may reduce latency and bandwidth for nodes that are currently online,
4380 nodes that were offline for a period of time or are being booted up for the first time
4381 must still download the complete block. Therefore, these techniques do not improve
4382 the performance of an initial synchronization and thus do not represent a complete
4383 scaling "solution."
- 4384 • An adversary who does not want their block quickly propagated to all of the miners
4385 may use these techniques solely for personal advantage by, for example, using com-
4386 pact blocks to send blocks to a favored portion of the network and sending full blocks
4387 or nothing at all to a disfavored portion.

4388 The compact blocks technique involves transmitting only the 80-byte block header, short-
4389 ened transaction identifiers (txid) that are hardened to prevent denial-of-service attacks, the
4390 coinbase transaction, and a small selection of full transactions that the sending peer predicts
4391 the receiving peer may not have seen yet [239]. The receiving peer then tries to reconstruct
4392 the block themselves using the information provided and their own mempool and then re-
4393 quests any missing transactions. There's also a high-bandwidth mode where the receiving
4394 peer has a few of their peers send new blocks without asking first, which increases band-
4395 width (because the node may receive the block multiple times simultaneously) but reduces
4396 latency. The message flow can be seen in Figure 23. Typically, these can reduce what
4397 would have been a 1 MB block transmission to approximately 20 KB of data over the wire.

4398 The short txids are created by taking a SHA-256 hash of the block header and nonce, using
4399 the fast hashing algorithm SipHash on the txid and some of the output of the SHA-256 hash,

4400 and removing the two most significant bytes of the result. The short txid is the remaining six
4401 bytes. Without this hashing procedure, it would be possible to create short txid collisions
4402 too easily, allowing attackers to prevent the scheme from working. While a 48-bit hash
4403 is not sufficient to prevent intentional collisions, the use of the block hash as a key to the
4404 hashing algorithm prevents an attacker from predicting the actual keys that will be used
4405 when the adversarially created transactions are included in a block. Additionally, peers
4406 have a 64-bit nonce that they share with each other and are unique per connection. These
4407 are mixed in to prevent even the block creator from controlling where collisions occur. The
4408 use of SHA-256 for this mixing is much slower than SipHash but is only performed once
4409 per block, so it does not add much overhead.

4410 Another way that block propagation has been sped up is with the use of FIBRE,¹ which uses
4411 UDP instead of TCP and encodes (compact) blocks using forward error correction (FEC).
4412 TCP requires receiving the network packets that include the block in order, whereas UDP
4413 allows data to be consumed as fast as the network allows. FEC helps by better handling
4414 dropped packets. A FIBRE network is essentially an allowlist of miners who prioritize ex-
4415 tremely low latency in block transmission. Because it consists of a curated or permissioned
4416 set of nodes, the network operator can theoretically perform censorship. However, anyone
4417 can set up a FIBRE network, so it is not conceptually dependent on a single centralized
4418 entity. Miners who utilize a relay network are substantially more likely to win in a block
4419 race than those who do not participate [240].

4420 10.2.2. Majority Hash Rate Attacks (51% Attacks)

4421 For Nakamoto Consensus (and other pure proof-of-work consensus mechanisms), a secu-
4422 rity assumption has been violated if the majority of the hash rate is malicious. An entity
4423 that controls the majority of the computational power deployed to the network has com-
4424 plete control over the content of blocks and can thus double-spend or censor transactions
4425 with certainty given enough time. This does not give the adversary the ability to change
4426 the rules arbitrarily, but it does allow them to decide which subset of valid transactions are
4427 added to the ledger and in which order.

4428 One of the more unfortunate aspects of majority attacks is that they are largely self-funding.
4429 The majority attacker can replace any blocks created by honest miners with their own
4430 and thus claim 100% of the block rewards. Once an attacker acquires the majority of the
4431 computational power deployed on the network, incentive compatibility requires that the
4432 block rewards must be large relative to the benefits of attacking the chain. This condition
4433 is difficult to maintain unless "(i) the mining technology used to run the blockchain is both
4434 scarce and non-repurposable, and (ii) any majority attack is a 'sabotage' in that it causes a
4435 collapse in the economic value of the blockchain" [241]. Stated differently, to dissuade a
4436 majority attacker, the proofs of work and the hardware used to create them must be useless
4437 outside of the network in question, and an attack must cause a decrease in the purchasing

¹<https://bitcoinfibre.org/>

4438 power of the underlying asset.

4439 Part of the security argument against majority attacks is that a rational attacker is not going
4440 to harm the system that pays them. Of course, this does not hold in what is known as a
4441 *Goldfinger attack* – when an attacker expects to gain utility by causing the asset price to
4442 crash, perhaps because they have shorted the asset or are a central bank that fears currency
4443 competition.

4444 This rationality argument is also inapplicable if a temporary majority can be created by
4445 renting hash rate. This is easier when the difficulty is low, perhaps because the network
4446 is still young or when the proof-of-work algorithm is ASIC-resistant and does not require
4447 specialized hardware. One strategy for renting computational power would be to form a
4448 negative fee mining pool that would have higher payouts than honest pools, which entices
4449 otherwise honest miners to the attacker's pool [242]. This is also possible using bribery
4450 via out-of-band payments, which require more trust, or *whale transactions* [243], which
4451 are high-value transactions intended to bribe miners onto a particular side of a fork. To
4452 create whale transactions, the attacker needs to first move some money to an address in
4453 the first block of their preferred fork so that spending from that address cannot happen on
4454 the other side of the fork. The attacker then creates a transaction with j outputs that are
4455 spendable by anyone but "timelocked" so that they can only be claimed in a series of j
4456 escalating blocks from the fork point. Miners who mine on this fork can then claim these
4457 bribes in the blocks they mine, but the bribes are never paid if the attacker's fork does not
4458 take over. As an alternative to this anyone-can-spend output method, a stream of high-fee
4459 transactions could work but would be less effective than the timelocks. Larger miners are
4460 more likely to switch to the whale transaction branch than smaller miners because the fork
4461 is then more likely to succeed, and larger miners may collect a higher proportion of blocks
4462 on the forked chain. It is also possible to use smart contracts deployed on other systems in
4463 order to facilitate bribery attacks [244–247].

4464 In an extension to the model in [241], the ability for 51% attack victims to retaliate with
4465 their own rented hash rate or whale transaction bribes – which appears to have occurred on
4466 the Bitcoin Gold blockchain in February 2020 – disincentivizes majority attacks in the first
4467 place [248]. This result makes the relatively weak assumptions that the cost of the attacks
4468 and counterattacks increase over time and that the victim would suffer a cost in reputation
4469 for being the victim of a double-spending attack. This assumption is likely true for entities
4470 like cryptocurrency exchanges, which are also most likely to be the victims of majority
4471 attacks. However, this argument does not hold if the attacker is merely trying to censor
4472 transactions instead of double-spend. The simplest defense against censorship attacks is
4473 to have high transaction fees, which increases the opportunity cost of censorship to the
4474 attacker.

4475 10.2.3. Hash Function Collisions

4476 If there were efficient ways to produce hash collisions, then proof of work would not be
4477 effective as a Sybil-resistance mechanism, and consensus would trivially fail. In addition,
4478 the hash references to previous blocks in the blockchain would no longer commit to the
4479 entirety of the chain up to that point and would fail to properly commit to the data in
4480 the blockchain itself. For example, an adversary capable of finding hash collisions would
4481 be able to arbitrarily change the content of blocks at any point in the blockchain without
4482 actually breaking the chain itself and could thus arbitrarily manipulate the state of the
4483 system to their advantage. As a consequence, none of the data in the blockchain would be
4484 tamper-resistant or have any integrity, and it would be impossible for nodes to agree on the
4485 state of the system.

4486 10.3. (More) Attacks Against Nakamoto Consensus

4487 Even when the typical security assumptions for Nakamoto Consensus hold, there are still
4488 possible attacks against users of the system. In particular, attackers with less than half
4489 of the network's hash rate can still attempt to double-spend or censor transactions. The
4490 Bitcoin white paper [1] considered double-spending with less than half the computational
4491 power assuming a fixed hash rate and constant difficulty, although the calculations were
4492 flawed and later corrected by Rosenfeld in [249]. The updated probabilities can be found
4493 in Figure 24.

4494 The first step of the attack is to mine a block that includes a transaction in which the at-
4495 tacker sends some of their funds to themselves. They withhold this block from the network
4496 and then broadcast a transaction that sends those same funds to a merchant. The merchant
4497 waits for k confirmations before giving the attacker whatever they purchased, and the at-
4498 tacker continues mining on their secret branch of the chain during this time. If the attacker
4499 manages to build a longer chain, they broadcast it to the network, which then adopts that
4500 chain as the best one. This overwrites the transaction sending funds to the merchant, and
4501 the attacker can walk away with both their original funds and the goods they did not pay
4502 for.

4503 A piece of Bitcoin folk wisdom is that merchants should wait for six confirmations before
4504 accepting a transaction. This number was selected by bounding the risk of a double-spend
4505 to 0.1%, assuming that the attacker has amassed no more than 10% of the network's hash
4506 power. However, these numbers are arbitrary, and as the attacker's computational power
4507 gets closer to 50%, the number of required confirmations blows up toward infinity. On the
4508 other hand, small valued transactions are likely safe after just one or two confirmations. To
4509 be profitable (rather than just probable), the attacker should incorporate a stopping thresh-
4510 old for when their private chain falls too far behind and success is unlikely [250].

4511 The attack can also be run multiple times concurrently to steal from multiple merchants at
4512 once. This makes it harder for a merchant to calculate how long to wait based on the value

q	1	2	3	4	5	6	7	8	9	10
2%	4%	0.237%	0.016%	0.001%	≈ 0	≈ 0	≈ 0	≈ 0	≈ 0	≈ 0
4%	8%	0.934%	0.120%	0.016%	0.002%	≈ 0	≈ 0	≈ 0	≈ 0	≈ 0
6%	12%	2.074%	0.394%	0.078%	0.016%	0.003%	0.001%	≈ 0	≈ 0	≈ 0
8%	16%	3.635%	0.905%	0.235%	0.063%	0.017%	0.005%	0.001%	≈ 0	≈ 0
10%	20%	5.600%	1.712%	0.546%	0.178%	0.059%	0.020%	0.007%	0.002%	0.001%
12%	24%	7.949%	2.864%	1.074%	0.412%	0.161%	0.063%	0.025%	0.010%	0.004%
14%	28%	10.662%	4.400%	1.887%	0.828%	0.369%	0.166%	0.075%	0.034%	0.016%
16%	32%	13.722%	6.352%	3.050%	1.497%	0.745%	0.375%	0.190%	0.097%	0.050%
18%	36%	17.107%	8.741%	4.626%	2.499%	1.369%	0.758%	0.423%	0.237%	0.134%
20%	40%	20.800%	11.584%	6.669%	3.916%	2.331%	1.401%	0.848%	0.516%	0.316%
22%	44%	24.781%	14.887%	9.227%	5.828%	3.729%	2.407%	1.565%	1.023%	0.672%
24%	48%	29.030%	18.650%	12.339%	8.310%	5.664%	3.895%	2.696%	1.876%	1.311%
26%	52%	33.530%	22.868%	16.031%	11.427%	8.238%	5.988%	4.380%	3.220%	2.377%
28%	56%	38.259%	27.530%	20.319%	15.232%	11.539%	8.810%	6.766%	5.221%	4.044%
30%	60%	43.200%	32.616%	25.207%	19.762%	15.645%	12.475%	10.003%	8.055%	6.511%
32%	64%	48.333%	38.105%	30.687%	25.037%	20.611%	17.080%	14.226%	11.897%	9.983%
34%	68%	53.638%	43.970%	36.738%	31.058%	26.470%	22.695%	19.548%	16.900%	14.655%
36%	72%	59.098%	50.179%	43.330%	37.807%	33.226%	29.356%	26.044%	23.182%	20.692%
38%	76%	64.691%	56.698%	50.421%	45.245%	40.854%	37.062%	33.743%	30.811%	28.201%
40%	80%	70.400%	63.488%	57.958%	53.314%	49.300%	45.769%	42.621%	39.787%	37.218%
42%	84%	76.205%	70.508%	65.882%	61.938%	58.480%	55.390%	52.595%	50.042%	47.692%
44%	88%	82.086%	77.715%	74.125%	71.028%	68.282%	65.801%	63.530%	61.431%	59.478%
46%	92%	88.026%	85.064%	82.612%	80.480%	78.573%	76.836%	75.234%	73.742%	72.342%
48%	96%	94.003%	92.508%	91.264%	90.177%	89.201%	88.307%	87.478%	86.703%	85.972%
50%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%

Fig. 24. Probability of successful double spend as a function of the attacker's hash rate q and number of confirmations. These probabilities assume that the attacker has already mined and withheld a single block that includes the self-paying double-spend transaction. [249]

4513 of a transaction, and they may conservatively have to consider the value of an entire block.
 4514 Further, an attacker with arbitrarily low hash rate can profitably implement a double-spend
 4515 by combining it with selfish mining. In this case, the attacker mines in secret for a double-
 4516 spend, but when it is unlikely that they will be able to reorg a sufficient number of blocks,
 4517 they switch to selfish mining and publish their secret blocks to get the block rewards.

4518 Sometimes, for small transactions, a merchant may accept a transaction without any con-
 4519 firmations (colloquially called "0-conf"), even though the protocol provides no security
 4520 guarantees for these transactions even against adversaries that do not control any computa-
 4521 tional power. Assuming that most spenders are honest, this might be acceptable from a risk
 4522 management perspective. However, when the attacker does control computational power,
 4523 additional attacks are possible, including the *Finney attack* and *vector76 attack*. Both of
 4524 these attacks can be generalized in order to attack merchants who require confirmations as
 4525 well [251].

4526 The Finney attack was proposed by cryptographer Hal Finney and starts with the attacker
 4527 mining a block that contains a transaction sending their funds to another address that the
 4528 attacker controls. Instead of broadcasting the block to the network, the attacker broadcasts a
 4529 transaction from the original address to the merchant, collects the goods from the merchant,
 4530 and finally broadcasts the pre-mined block to double-spend [252]. The Finney attack can
 4531 also be used when the merchant requires k confirmations before accepting a transaction. In

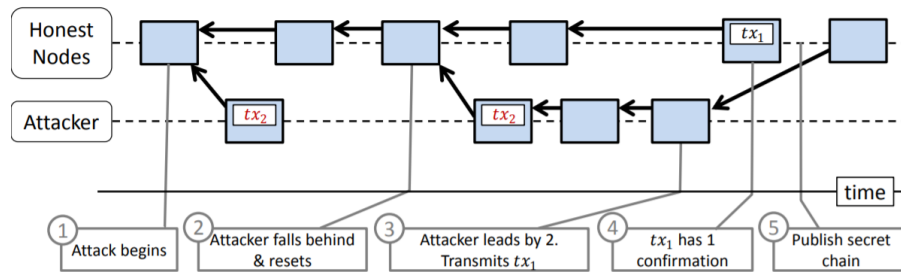


Fig. 25. Finney attack example. [251]

4532 this case, the adversary works on a secret chain that embeds the double-spend transaction,
4533 tx_2 , which conflicts with the merchant transaction, tx_1 . The attacker attempts to create $k + 1$
4534 blocks more than the honest miners but aborts if the honest miners ever get ahead of the
4535 attacker, potentially restarting at a later block in the chain. Once the attacker has a $k + 1$
4536 block lead, they broadcast tx_1 , wait for the honest miners to produce enough blocks for
4537 tx_1 to have k confirmations, and then broadcast their private chain to overtake the network
4538 and reverse tx_1 . An example of a 1-conf generalized Finney attack appears in Figure 25.
4539 Importantly, if the attacker can choose when they want to submit a transaction, the pre-
4540 mining stage can be attempted repeatedly until the attack can be executed with a near-
4541 guarantee of success.

4542 The vector76 attack requires the adversary to have a direct connection to the victim [253].
4543 The attacker mines a block that includes a transaction sending funds to the merchant but
4544 does not broadcast the block or the transaction to the rest of the network. The attacker then
4545 waits until another block is mined at the same height by the honest miners before sending
4546 the pre-mined block directly to the victim (hopefully before the victim sees the other block)
4547 and collecting the goods. Finally, the attacker broadcasts a conflicting transaction that
4548 refunds themselves. The network adds this double-spend transaction to a block, and the
4549 original pre-mined attack block becomes stale.

4550 There is a generalized version of this attack that works against victims that do not relay
4551 blocks they have received, like light clients [251]. This is one reason why light clients
4552 are less secure than full nodes and why merchants using light clients should wait for more
4553 confirmations than they otherwise would. For example, say that a merchant only uses a light
4554 client for validation and requires k confirmations to accept a transaction. The attacker starts
4555 pre-mining a secret branch that includes a transaction to the merchant, tx_1 , in the first block
4556 of their private chain. They continue mining on this branch until tx_1 has k confirmations.
4557 If the attacker's branch is longer than the honest branch, they show the merchant their
4558 pre-mined chain, collect the goods, and stop mining on the private branch. Finally, the
4559 attacker broadcasts a conflicting transaction, tx_2 , to the network, which is unaware of the
4560 attack chain that is known only to the attacker and the merchant. Eventually, the honest
4561 network will build a longer branch than the private one that includes tx_2 , thus reversing
4562 tx_1 . Figure 26 shows an example of this attack executed against a merchant requiring two

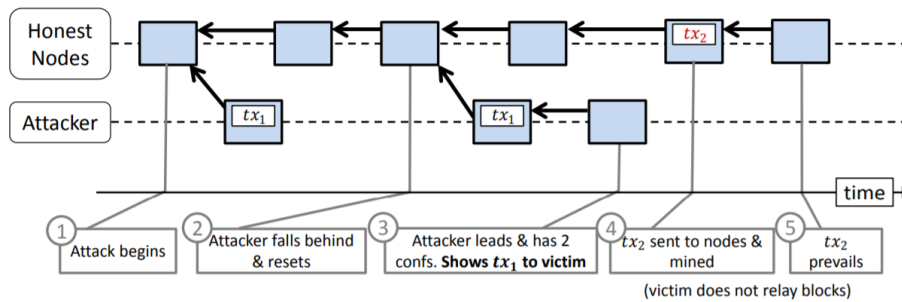


Fig. 26. Vector76 attack example. [251]

4563 confirmations.

4564 So far, this section has only considered ways in which a minority hash rate attacker can
4565 double-spend against adversaries who require only a small number of confirmations. An
4566 adversary could also attempt to censor transactions via *feather-forking* [254]. In a feather-
4567 forking attack, the miner with a minority of the hash rate commits to not mining on top
4568 of any chain that contains an address or transaction that they would like to censor, at least
4569 for several blocks. If the attacker makes it credibly known to other miners that they are
4570 following this policy, then other miners are less likely to include the address or transaction
4571 in question for fear that their block might be reorged. If the attacker is able to coax more
4572 than half of the total hash rate to censor, they have succeeded.

4573 11. More Proof-of-Work Protocols

4574 Nakamoto Consensus is the paradigmatic proof-of-work consensus algorithm, but many
4575 others have been proposed or implemented. This section begins by considering some more
4576 modest adjustments to Nakamoto Consensus. Most of these protocols attempt to address
4577 one of its primary flaws: the degradation of security that occurs when block transmission
4578 latency is too high. Several also attempt to improve upon its low chain quality in order
4579 to resist selfish mining attacks. The GHOST fork-choice rule is used to allow blocks to
4580 be produced more quickly and reduce the time until transaction settlement. FruitChains
4581 modify the blockchain structure in order to achieve better chain quality guarantees.

4582 Other protocols run Nakamoto Consensus over multiple separate blockchains in paral-
4583 lel, which can substantially reduce the latency of transaction settlement. By giving these
4584 blockchains different roles, the Prism construction is able to increase throughput and im-
4585 prove latency. Several protocols utilize DAGs instead of linear blockchains. These ap-
4586 proaches address Nakamoto Consensus's block transmission issue by largely doing away
4587 with the concept of stale blocks in the first place and allowing all competing blocks to ul-
4588 timately be incorporated into the ledger. This comes with increased complexity in protocol
4589 design and analysis and, in some cases, only results in a partial rather than total ordering of
4590 transactions, making them unsuitable for generic smart contracts.

4591 Finally, there are protocols that use proof of work as a Sybil-resistance mechanism to con-
4592 trol who may participate in classical permissioned BFT consensus. These protocols allow
4593 responsive transaction commitment, where settlement occurs at the actual speed of network
4594 communication rather than based on the maximum network delay. This advantage comes
4595 at the cost of increased trust assumptions that are more in line with proof-of-stake protocols
4596 than the other proof-of-work ones discussed here.

4597 Note that many of the protocols presented in this section have a proof-of-stake "equivalent"
4598 or can be used as subprotocols in proof-of-stake systems. The security ramifications would
4599 need to be considered in their respective environments, but the protocols are included here
4600 because they were originally proposed and analyzed in the proof-of-work context. This is
4601 the case for GHOST, FruitChains, Parallel Chains, and several of the DAG protocols.

4602 **11.1. Nakamoto Consensus Protocol Adjustments**

4603 Recall that the primary security challenge with Nakamoto Consensus relates to the network
4604 latency of block transmission (see Section 10.2.1). A variety of algorithms attempt to work
4605 around this limitation.

4606 **11.1.1. Weak Blocks and Pre-Consensus**

4607 Weak blocks are a way to slightly modify Nakamoto Consensus in order to help synchro-
4608 nize mempools between miners. The idea is that there are two separate proof-of-work
4609 difficulty targets: the typical one for mining blocks and a substantially easier one for min-
4610 ing weak blocks (analogous to mining pool shares). For example, the weak block difficulty
4611 may be set to $\frac{1}{10}$ of the normal block difficulty such that, in expectation, there are 10 weak
4612 blocks found in between normal blocks.

4613 Like normal blocks, weak blocks carry transactions. This helps miners know which trans-
4614 actions other miners have seen in the meantime. Essentially, blocks are transmitted piece
4615 by piece during the standard block interval rather than all at once in a burst when blocks
4616 are found. This helps minimize block propagation delays and allows for higher through-
4617 put without increasing the risk of stale blocks. Weak blocks also allow "pre-consensus"
4618 on transactions that have yet to be confirmed, slightly increasing the security of accepting
4619 transactions with no confirmations. Weak blocks require additional overhead and are only
4620 effective if the majority of miners participate. Further, non-mining full nodes may not be
4621 incentivized to broadcast weak blocks because they use bandwidth but only benefit miners
4622 directly.

4623 Specific weak block proposals include Subchains [255] and Flux [256]. The Subchains
4624 idea is merely an addition to Nakamoto Consensus, whereas Flux actually modifies the
4625 consensus algorithm. The Subchains protocol works as follows, starting from the point
4626 where a miner sees and accepts a regular block:

- 4627 1. The miner begins working on the next block, including a hash pointer to the block that

4628 was just accepted, and finds a proof-of-work that satisfies the weak block difficulty.
4629 The miner broadcasts this weak block to the network.

4630 2. Miners verify the weak block. If valid, miners adjust the coinbase transaction's re-
4631 ward to include the new transaction fees, add any new transactions to their block
4632 candidate that they desire, and compute a new Merkle root for the block header.
4633 They then continue looking for a valid nonce to find another proof of work. Denote
4634 the new coinbase transaction, the newly included transactions, the Merkle root, and
4635 the previous block hash as a Δ -block.

4636 3. When miners compute new proofs of work that create weak blocks, they broadcast
4637 only the Δ -block and the hash of the previous weak block. That is, miners essentially
4638 cooperate to build the next block incrementally by building a subchain.

4639 4. Eventually, a miner will find a regular block that satisfies the more difficult target.
4640 When they do so, they broadcast only their Δ -block and the hash of the previous weak
4641 block. Note that this should require significantly less bandwidth than broadcasting a
4642 full block, thus improving latency.

4643 If there are multiple competing subchains, honest miners build on top of the longest sub-
4644 chain that they are aware of. If conflicting transactions exist, any transaction included in a
4645 subchain is prioritized over competing transactions that only exist in their mempool.

4646 Flux modifies the fork-choice rule to include the work from weak blocks in the total work
4647 when deciding which chain is "longest" so that a lengthy chain of weak blocks can reorg a
4648 normal block. Flux also changes the incentives to include revenue sharing among those who
4649 mine weak blocks. When building a subchain of weak blocks, miners must include the prior
4650 weak block miner addresses in their coinbase transaction in order to be considered valid,
4651 proportional to the number of weak blocks found. This revenue-sharing scheme reduces
4652 the variance of the mining reward and may reduce the profitability of selfish mining. On
4653 the other hand, miners who have already created weak blocks in the subchain lose rewards
4654 when new weak blocks are added, so they may try to create forks in the subchain or not
4655 propagate other miners' weak blocks. Additionally, once a long subchain exists, miners
4656 may no longer find it profitable to mine until the next regular block is found.

4657 11.1.2. Bitcoin-NG

4658 Bitcoin-NG ("next generation") was one of the earliest proposals for modifying Nakamoto
4659 Consensus to reduce block transmission latency [257]. The primary idea is to separate
4660 the leader election process from transaction serialization, which is achieved by having *key*
4661 *blocks* and *microblocks*. Key blocks are like normal Bitcoin blocks from a consensus per-
4662 spective but do not include transactions and that specify a public key for the leader. Mi-
4663 croblocks are produced between key blocks and contain transactions. Mining a key block
4664 makes the public key holder the leader of the next epoch, and they are entitled to mine
4665 microblocks at a fixed rate (microblocks do not require proof of work but rather are signed

4666 by the specified public key). Bitcoin-NG uses the longest chain rule for the chain of key
4667 blocks but changes the incentives: the block subsidy goes to whomever finds the key block,
4668 but transaction fees are split so that 40% goes to the leader (and thus microblock producer)
4669 and 60% goes to the miner who finds the next key block.

4670 Microblocks do not include a proof of work, so a malicious leader can split the chain (of
4671 microblocks) for free. To prevent this, Bitcoin-NG uses *poison transactions* to invalidate
4672 the rewards of malicious leaders. Poison transactions include the header of the first mi-
4673 croblock of a conflicting fork to prove that fraud has occurred. This special transaction can
4674 be included in any key block during the coinbase maturity window of the malicious leader's
4675 key block and rewards the leader who includes it with a fraction of the forfeited funds.

4676 Because key blocks do not include transactions, they are small and should propagate through
4677 the network quickly, so increasing transaction throughput does not increase the risk of a key
4678 block becoming stale. However, there are a number of trade-offs, including making SPV
4679 light clients impractical, eliminating the ability of individual miners to choose which trans-
4680 actions are in a block (only mining pool leaders can do so), and potentially encouraging
4681 denial-of-service attacks against leaders by revealing their identity before they have pro-
4682 duced microblocks.

4683 Several works have revisited Bitcoin-NG's incentives [258, 259]. In [258], it is shown that
4684 the reward split should provide only $\frac{3}{11}$ of the reward to the leader instead of 40% because
4685 the original analysis ignored the possibility that a miner in one epoch also gets the key
4686 block for the next epoch. Several additional selfish mining-related attacks are shown in
4687 [259], and the microblock architecture increases the profitability of selfish mining when
4688 the attacker has more than 35% of the network's hash rate.

4689 11.1.3. Tie-Breaking Schemes

4690 When multiple blocks are mined simultaneously, miners must decide which block to mine
4691 on top of. In Bitcoin, nodes accept the first block they see until one side of the fork has
4692 more work than the other. In Ethereum and Bitcoin-NG, nodes choose uniformly at random
4693 which side of the fork to prefer (UTB). In the following sections, additional schemes called
4694 DECOR+ and Publish or Perish (PoP) suggest other tie-breaking methods (DECOR+ is
4695 discussed in more detail in Section 11.1.4, and PoP is detailed in Section 11.1.5). An early
4696 DECOR+ proposal broke ties using the chain tip with the smallest hash (SHTB) [260]. The
4697 updated DECOR+ uses an unpredictable deterministic tie-breaking scheme (UDTB) where
4698 the winner is chosen based on a pseudorandom function taking all competing blocks as
4699 inputs [261]. PoP compares chain "weights," and blocks published after their competitors
4700 do not contribute weight, while blocks that incorporate links to their parents' competitors
4701 have higher weights. Thus, blocks that are kept secret until competing blocks are published
4702 will contribute to neither or both branches and thus confer no advantage in winning the
4703 block race [262].

4704 Each of these choices has its own security ramifications, particularly with respect to self-
4705 ish mining resistance (see Section 9.4 for background). By accepting the first seen block,
4706 Bitcoin allows the honest computational power to become split such that sufficiently well-
4707 connected adversaries can selfish mine profitably at arbitrarily low hash rates. However, if
4708 the attacker is not well-placed on the network, it is the tie-breaking method most resistant
4709 to selfish mining. UTB and UDTB have nearly identical resistance to selfish mining, and
4710 neither outperforms Nakamoto Consensus when $\gamma \leq 0.5$. By not taking into account the
4711 time that a block was received, both UTB and UDTB allow an attacker who mines "from
4712 behind" to still win the block race with a tie [263]. SHTB has the lowest chain quality
4713 of all because a miner who finds a block with a particularly low hash can continue self-
4714 ish mining privately and feel confident that they will win block races when they occur, so
4715 they can strategically deviate when the odds are in the attacker's favor. SHTB also suffers
4716 from the same mining from behind issue as UTB and UDTB. When $\gamma = 0$, none of these
4717 protocols have better chain quality than Nakamoto Consensus for $\alpha < 0.39$. However, PoP
4718 begins to outperform Nakamoto Consensus when $\alpha \geq 0.4$ [263]. The poor chain quality
4719 of Nakamoto Consensus and these other tie-breaking variants is caused by an informa-
4720 tion asymmetry, where the attacker has more information than the honest miners regarding
4721 network connectivity [263].

4722 11.1.4. DECOR+

4723 DECOR+ (deterministic conflict resolution) is a revenue-sharing mechanism intended to
4724 incentivize miners to converge on the same block during block races [260, 261]. It tries
4725 to share the reward among all miners who mined a block at the same height (uncles) in
4726 order to allow faster block production. DECOR+ was designed such that if all nodes have
4727 access to the same blockchain state, any conflicts will be quickly resolved in a deterministic
4728 fashion that is incentive-compatible for miners [260]. This prevents honest miners from
4729 splitting their hash rate between two forks for long. Originally, it was supposed to break
4730 ties using the block with the highest fees and then the smallest hash if fees were tied, but
4731 fees are gameable by miners (e.g., they can include their own transactions with arbitrary
4732 fees paid to themselves, and users can also pay fees to miners out-of-band). In [261], a
4733 different selection function was proposed: hash all block headers and then take the XOR-
4734 sum, modulo the number of competing blocks. This way, the miner cannot compute their
4735 block in a way that gives them a higher chance of winning.

4736 A number of possible reward functions were described in [261], each of which accounts
4737 for uncle blocks in order to incentivize the inclusion of uncle references in blocks. To
4738 bound the maximum money supply, the reward function can bound the number of uncles
4739 for which a reward is provided. For instance, the reward function may put the L competing
4740 block headers in order by block hash and reward the N lowest ones. There can also be a
4741 punishment fee if a miner does not follow the deterministic selection function. For example,
4742 let X be the deterministically selected block, Z the block selected by the miner, Y the mined
4743 block, r the total block reward, i the inclusion reward per uncle, and p the punishment fee.

4744 If $X \neq Z$, let $r_y = \frac{r(1-p)}{N+1}$. Otherwise, let $r_y = \frac{r}{N+1}$. Then the reward array is:

$$[r_{i_1}, \dots, r_{i_N}, r_{i_{N+1}}, \dots, r_{i_{N+L}}, R_Y] := \left[\frac{r}{N+1}, \dots, \frac{r}{N+1}, 0, \dots, 0, r_Y + iN \right] \quad (4)$$

4745 By punishing miners who deviate from the deterministic block of choice, DECOR+ is more
4746 resistant to selfish mining and double-spending but at the cost of being more susceptible
4747 to feather-forking censorship [263]. This is because a malicious miner has an easier time
4748 decreasing the income of honest miners.

4749 11.1.5. Publish or Perish

4750 Publish or Perish (PoP) was proposed as a defense against selfish mining and makes it so
4751 that withholding blocks does not provide a miner with an advantage in winning block races
4752 [262]. In this scheme, a node will consider a block *in time* if the height of the block is larger
4753 than the node's local best chain height or if its height matches the local best chain height
4754 and is received within a bounded length of time (corresponding to a typical network delay)
4755 from seeing the first block of that height. PoP uses a nonstandard definition of "uncle"
4756 blocks, where an "uncle" must be in time, and the height of a block's "uncle" must be one
4757 less than the height of the block. Miners are encouraged to include references to these
4758 "uncles" in the blocks that they mine.

4759 The fork-choice rule nodes follow in PoP uses the *weight* of a chain, which is the number
4760 of in-time blocks plus the number of in-time "uncles" referenced in the in-time blocks. It
4761 also uses a security parameter k (where $k = 3$ is suggested by the authors), which manages
4762 a trade-off between selfish mining resistance and the ability of nodes to quickly recover
4763 from network partitions. Specifically, when a block race occurs, the fork-choice rule is as
4764 follows:

- 4765 1. If one chain is longer than another by at least k blocks, then follow the longest chain.
- 4766 2. If the difference is less than k blocks, then follow the chain with the most weight.
- 4767 3. If chains are tied for highest weight, then follow a random one.

4768 If a selfish miner withholds a block and keeps it secret until after a competing block has
4769 already been published, then the selfish miner's block will not contribute to the weight of
4770 the attacker's chain. Alternatively, if the attacker's block is published around the same time
4771 as the competing honest block, then the next honest block that is produced can include an
4772 "uncle" reference to the previously secret block, which increases the weight of the honest
4773 chain. In either case, the selfish miner does not gain an advantage in the block race by
4774 withholding. Because of the trade-off between selfish mining resistance and tolerating
4775 network partitions, merchants in a system that uses the PoP fork-choice rule may want
4776 to wait at least k blocks before considering a transaction of non-trivial value sufficiently
4777 confirmed.

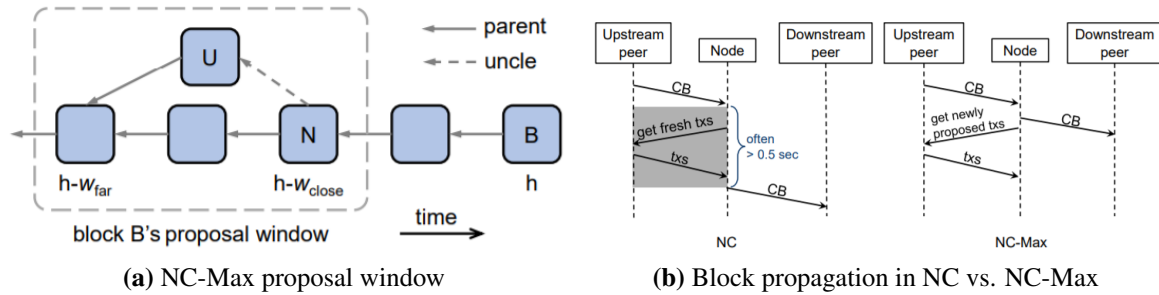


Fig. 27. NC-Max block propagation mechanism. Panel (a) shows the NC-Max proposal window with $w_{far} = 4$ and $w_{close} = 2$. If a transaction is proposed in block U and committed in B , the transaction fee goes to the miner of block N . Panel (b) shows how compact blocks can propagate more quickly using NC-Max compared to Nakamoto Consensus. [264, 265]

4778 **11.1.6. NC-Max**

4779 NC-Max attempts to remove the latency-related scalability bottleneck of Nakamoto Con-
 4780 sensus by separating transaction synchronization from transaction confirmation [264, 265].
 4781 It builds off of the compact blocks idea (see Section 10.2.1) and attempts to work around
 4782 the issue of requiring more round-trip communication to handle transactions that have not
 4783 propagated through the network when blocks are found. It does so via a two-step process
 4784 of transaction proposal and transaction commitment.

4785 In NC-Max, miners include references to uncle blocks (stale blocks in the same difficulty
 4786 epoch) within the blocks that they find. The part of the block that contains transactions
 4787 is called the transaction commitment zone. Each block also has a *transaction proposal*
 4788 *zone* that includes txpids, or the first few bytes of the transaction ID, and a transaction is
 4789 considered proposed if it is included here (or if it was proposed in an uncle block that was
 4790 referenced). The transactions referenced in the transaction proposal zone need not be valid
 4791 transactions for the block itself to be valid. A new validity rule is added for transactions
 4792 in the commitment zone at height h : it must have been proposed during the *proposal window*
 4793 in a block of height $h - w_{far}$ to $h - w_{close}$. An example is shown in Figure 27a. Note that
 4794 the coinbase transaction is excluded from this mechanism.

4795 Nodes forward compact blocks that include the transactions in the proposal zone to their
 4796 peers once they themselves have reconstructed the commitment zone. They should have all
 4797 of these transactions after receiving the blocks in the proposal window. In the meantime,
 4798 the node requests any proposed transactions they have not yet seen from their peers. As
 4799 a result, the extra round-trips required to receive these transactions do not impact block
 4800 propagation time. A comparison of compact block propagation in Nakamoto Consensus
 4801 and NC-Max is in Figure 27b.

4802 NC-Max also adjusts the reward distribution and the difficulty adjustment algorithm. When
 4803 a transaction is committed, the fee is split 70%-30% between the miner who commits it and

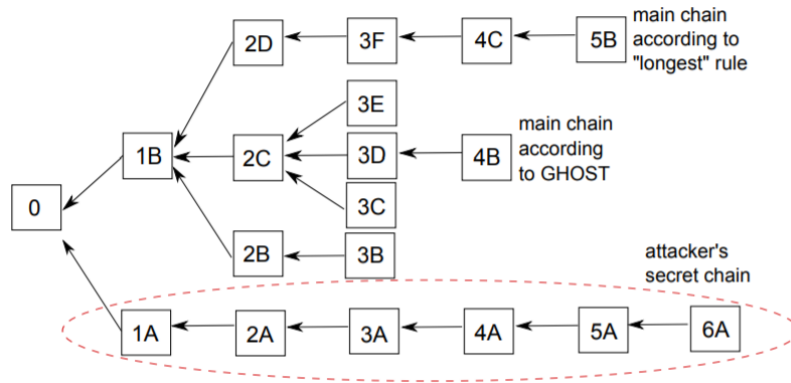


Fig. 28. GHOST fork choice rule, assuming all blocks have the same difficulty. An attacker chain would overtake the "longest" chain but not the GHOST chain, which consists of blocks 0, 1B, 2C, 3D, and 4B. In this example, the subtree that begins at block 1B has 12 blocks, whereas its competitor at 1A only has six. The subtree of block 2D has four blocks, for 2C has five blocks, and for 2B has only two blocks. The subtree beginning at block 3D has two blocks, as opposed to only one block for 3E and 3C. Every block added to the subtree of block 1B, regardless of block height, contributes to the security of block 1B. [266]

4804 the miner who proposes it [265]. Uncle miners do not collect any reward. The details of
4805 the difficulty adjustment are out of scope for this document but take into account the uncle
4806 blocks from the prior epoch in order to target a pre-specified stale block rate.

4807 11.2. Greedy Heaviest-Observed Sub-Tree (GHOST)

4808 One of the earliest proposed fork-choice rules aside from the longest chain rule was GHOST,
4809 a variant of which was later used in the Ethereum network [266]. It allows the expected
4810 block interval to be much shorter than Nakamoto Consensus because the proofs of work
4811 used for stale blocks are counted when deciding which chain is canonical. In other words,
4812 blocks that do not make it into the main blockchain still contribute to the total work for the
4813 chain. Contrast this with Nakamoto Consensus, where stale blocks contribute nothing to
4814 the security of the chain.

4815 In a GHOST blockchain, blocks contain an extra field that is used to reference uncle blocks.
4816 This creates a DAG of blocks and block references but still only chooses a linear chain of
4817 blocks by using the information contained in the DAG. When a fork exists, the GHOST
4818 algorithm greedily selects the heaviest subtree of blocks that begin from any fork point,
4819 starting from the genesis block. See Figure 28 for an example of the GHOST fork-choice
4820 rule. As one may infer from the example, GHOST leads to a small weakening in the chain
4821 growth property compared to Nakamoto Consensus, but this does not adversely impact the
4822 security of the chain.

4823 GHOST has been proven secure in a synchronous network with constant difficulty and
4824 has superior liveness (when not under attack) though worse chain quality compared to

4825 Nakamoto Consensus [267]. The consistency property of GHOST was also proven in
4826 [230], but contrary to many peoples' expectations, it has the same consistency bounds
4827 as Nakamoto Consensus. Therefore, GHOST is unable to handle a significantly higher
4828 throughput with the same security from a consistency perspective. It will degrade in secu-
4829 rity just as Nakamoto Consensus does (though it can still tolerate shorter block intervals).

4830 GHOST's liveness can suffer due to an attack that requires less than half of the network's
4831 computational power [230]. The adversary attempts to maintain forks for as long as possi-
4832 ble and thus delay transaction confirmation. When a fork occurs, the adversary attempts to
4833 keep both subtrees balanced by using their computational power on whichever side of the
4834 fork requires it. When the block interval is low compared to block propagation time, [230]
4835 showed that forks can be maintained for 10 or more blocks with non-negligible probabili-
4836 ty. GHOST is more susceptible to this attack because blocks mined by the adversary can
4837 be withheld from the network for longer than they can in Nakamoto Consensus and still
4838 contribute to the subtree as uncle blocks.

4839 Due to its use in Ethereum, GHOST has been subject to more real-world testing than any
4840 fork-choice rule other than the longest chain rule of Nakamoto Consensus. However, be-
4841 cause it is not as simple of a rule, it makes other aspects of the system – such as incen-
4842 tivization and the choice of difficulty adjustment algorithm – more complicated and less
4843 well-studied. This can lead to security issues. For instance, Ethereum's variant of GHOST
4844 provides incentives for miners of uncle blocks, and an old version of the protocol allowed
4845 a strategy that would create undesirable inflation and allowed greedy miners to gain at the
4846 expense of coin-holders [268]. This issue resulted in a change to Ethereum's difficulty ad-
4847 justment algorithm. Additionally, GHOST with uncle rewards is more susceptible to selfish
4848 mining than Nakamoto Consensus [210–212]. As α increases, the revenue of both selfish
4849 and honest miners increases due to uncle rewards, which may lead to greater inflation of
4850 the supply of ether. Rewarding uncle blocks may also encourage block withholding and
4851 fork-after-withholding attacks [269, 270].

4852 11.3. FruitChains

4853 Unlike GHOST, which reduces chain quality, the FruitChains protocol was designed in or-
4854 der to improve chain quality and reduce the efficacy of selfish mining [271]. Specifically,
4855 FruitChains are δ -approximately fair: any honest miners that control a ϕ fraction of com-
4856 putational power are guaranteed (with high probability) to get at least a $(1 - \delta)\phi$ -fraction
4857 of the rewards in any $\Omega(\frac{\kappa}{\delta})$ length portion of the chain, with κ as a security parameter. This
4858 prevents any adversarial minority of the hash rate from improving their revenue more than
4859 a factor of $(1 + 3\delta)$ compared to honest mining.

4860 The FruitChains protocol has miners simultaneously mine normal blocks as well as *fruits*,
4861 which have a lower difficulty. As with Nakamoto Consensus, miners follow the chain of
4862 blocks with the most work. However, unlike Nakamoto Consensus, blocks contain fruits
4863 instead of transactions, and the fruits are where transactions are recorded. A fruit must be




Fig. 29. FruitChains architecture. Assume the recency parameter $T_0 = 3$. In this case, fruits F_3 and F_4 are valid fruits included in block B_4 , because they point to only the two most recent blocks. Fruit F_1 could not be included in block B_4 , because it isn't sufficiently recent. Fruit F_2 also cannot be included, because it points to a stale block.

4864 sufficiently "recent" to count; that is, it must include a reference to a block not too far in
4865 the past from the block that contains the fruit (say, within T_0 blocks). This prevents fruits
4866 mined long in the past from being included. It also prevents fruit-withholding attacks,
4867 where the attacker builds up a bank of fruits and then mines them all into the chain in
4868 a short period of time, causing a large fraction of fruits to be adversarial. To construct
4869 a linearized log of transactions, an ordered sequence of distinct fruits must be extracted
4870 from the chain, including only the first fruit if duplicates exist. Then the transactions are
4871 extracted in order from this chain of fruits, also removing duplicates as needed. The block
4872 reward and transaction fees for a block are evenly distributed to the miners of a constant-
4873 length number of blocks preceding the block in question. The FruitChains architecture is
4874 shown in Figure 29.

4875 The simultaneous mining of fruits and blocks utilizes a "2-for-1" proof-of-work technique
4876 described in [9], which allows a single random oracle $H()$ to operate as two independent
4877 oracles, $H_0()$ and $H_1()$. In other words, miners can attempt to compute proofs of work
4878 for two different schemes for the cost of one oracle query without the ability to use their
4879 compute power to favor one process over the other. An example would be to use the m
4880 most significant bits of the hash function's output for one scheme (e.g., mining blocks) and
4881 the n least significant bits for another (mining fruits).

4882 A detailed security analysis of FruitChains was provided in [263]. In FruitChains, mining
4883 a block provides no direct reward (except a share of the rewards from future blocks), which
4884 has significant security implications. For example, a selfish miner has no incentive to pub-
4885 lish blocks when neither the secret chain nor the honest chain have reached T_0 blocks since
4886 the fork point – the attacker's fruits that were mined before the T_0 -th block – will also ap-
4887 pear in the honest chain. If the attacker is able to win a block race of at least T_0 blocks, they
4888 can invalidate all honest fruits, so double-spending is more profitable against FruitChains

4889 than against Nakamoto Consensus. Increasing T_0 makes this more challenging and reduces
4890 the incentive to selfish mine, but the safe transaction confirmation delay increases linearly
4891 with T_0 . By making fruit mining increasingly less difficult compared to block mining (a
4892 larger fruit-to-block ratio), selfish mining becomes less profitable but at the expense of
4893 having more repeated transactions and fruits that need to be removed from the ledger, thus
4894 wasting bandwidth. On the positive side, FruitChains are more censorship-resistant than
4895 Nakamoto Consensus because the attacker would need to overwrite T_0 blocks to invalidate
4896 honest fruit.

4897 11.4. Parallel Chain Approaches

4898 One way to reduce the latency of transaction settlement is to run multiple blockchains
4899 in parallel and then use some procedure to combine the contents of the separate chains
4900 [272, 273]. To that end, [272] generalizes the "2-for-1" mining technique of [9] (and
4901 described in Section 11.3) into an m -for-1 mining scheme. Then, given an underlying
4902 blockchain protocol, the technique can be used to execute m instances of it in parallel while
4903 the same mining operation is used across all m instances without allowing a deviant miner
4904 to focus their computational power onto a particular chain. A generic technique for then de-
4905 terministically combining these m (nearly) independent ledgers into a single *virtual ledger*
4906 was described in [273]. Care must be taken when handling the design of the difficulty ad-
4907 justment algorithm for schemes that employ parallel chains because naive adaptations of
4908 common DAAs are insecure in this context [274].

4909 Nakamoto Consensus has relatively high transaction settlement latency due to the need to
4910 wait for enough blocks to confirm a transaction and stabilize an agreed-upon chain prefix,
4911 where the possibility of a common prefix violation decreases exponentially in the num-
4912 ber of blocks. This settlement latency is directly related to and limited by the latency of
4913 block propagation on the underlying network. By combining m ledgers like this in paral-
4914 lel, settlement time can be reduced by up to a $\Theta(m)$ multiplicative factor when including a
4915 transaction in each ledger. This is done by using a ranking algorithm for each underlying
4916 ledger and then combining these ranks by an exponential sum of the ranks of the transaction
4917 from each individual chain. This allows for a trade-off between transaction fees and set-
4918 tlement time: clients can issue a transaction on a single ledger and pay one transaction fee
4919 with settlement times roughly on par with the settlement time of the underlying blockchain
4920 or alternatively pay up to m transaction fees to achieve the multiplicative $\Theta(m)$ reduction
4921 in settlement time.

4922 The details of how this ranking and combining are performed are beyond the scope of
4923 this document, but roughly, the ranking for a blockchain using Nakamoto Consensus could
4924 be the timestamp embedded in the block that contains the transaction in question. The
4925 combined rank is then essentially an average of the ranks in the individual blockchains,
4926 which amplifies the exponential rate at which transactions are settled:

$$e^{-\frac{\text{combinedrank}(tx)}{L}} = \frac{1}{m} * \sum_{i=1}^m e^{-\frac{\text{rank}_i(tx)}{L}}, \quad (5)$$

4927 where L is a parameter that is proportional to the security parameter of the system.

4928 The following subsection presents the Prism algorithm, which is a concrete instantiation of
4929 a parallel chains approach rather than the abstract and generic one presented here.

4930 11.4.1. Prism

4931 Unlike the generic parallel chains construction discussed above, Prism optimizes through-
4932 put in addition to latency [275, 276]. It is capable of achieving optimal throughput by
4933 taking advantage of the network’s full communication bandwidth and near-optimal trans-
4934 action settlement latency of approximately the network propagation delay.

4935 In a longest chain protocol, blocks perform several functions: they elect leaders, add trans-
4936 actions into the ledger, and vote for their ancestor blocks via parent link relationships.
4937 Prism separates these roles by using three separate types of blocks: proposer blocks for
4938 leader election, transaction blocks for transaction inclusion, and voter blocks to confirm
4939 transactions. Whenever a block is mined, it is randomly sortitioned into one of the three
4940 types of blocks and, if it is a voter block, further sortitioned into one of m voter chains
4941 ($m = 1000$ is suggested). This sortition process ensures that miners are unable to choose
4942 which type of block they mine. Rather, they simultaneously mine for each chain and only
4943 learn the type of the mined block after a valid proof of work is found. Proposer blocks
4944 contain a list of references to transaction blocks, as well as a single reference to a par-
4945 ent proposer block. As with Nakamoto Consensus, honest miners will mine on top of the
4946 longest proposer chain they are aware of. However, it is the voter chains that determine the
4947 final sequencing of proposer blocks (and thus elected leaders).

4948 Define the *level* of a proposer block as its distance from the genesis block of the proposer
4949 chain and the *height* of the proposer chain as the maximum level containing any proposer
4950 blocks. Voter blocks include a reference to a proposer block in order to cast a vote on it that
4951 is subject to two requirements: 1) the voter block is in the longest chain of its respective
4952 voter tree, and 2) each voter chain votes for exactly one proposer block at each level. The
4953 leader sequence, then, is the proposer block at each level with the highest number of votes
4954 among all proposer blocks at that level with ties broken by the smallest hash of the proposer
4955 blocks. With the proposer blocks ordered, their references to transaction blocks create
4956 an agreed-upon ordering of transaction blocks and thus transactions. The ordered list of
4957 transactions must then be sanitized in order to remove invalid or duplicate transactions.
4958 See Figure 30 for Prism’s structure.

4959 Ultimately, the security of the system is provided by the voter trees that give confirma-
4960 tions/votes to the proposer blocks. Changing an elected leader requires reversing a suffi-
4961 cient number of voter blocks, and each vote is secured by following the longest chain rule

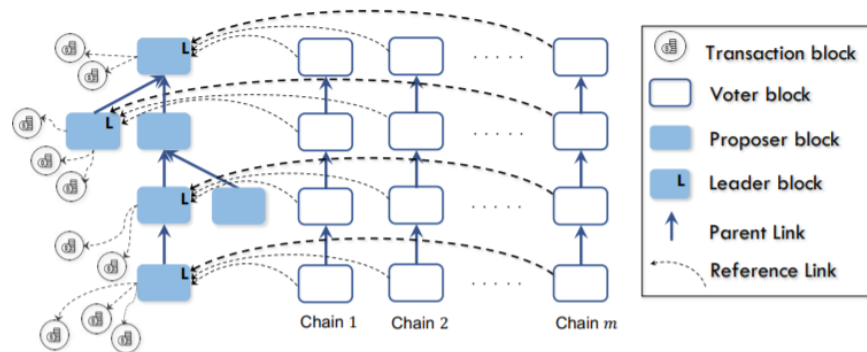


Fig. 30. Prism structure. [276]

4962 in its voter tree. Due to the large number of voter chains and the global hash rate being split
4963 among them, it is possible to minimize the amount of forking in each voter chain. So long as
4964 as the voter chains are secure, the proposer chain will be secure as well. As with the more
4965 abstract parallel chain protocol above, the existence of many voter chains substantially im-
4966 proves the latency of the transaction settlement.

4967 Prism has security proofs in synchronous networks in a couple of different models [277,
4968 278], and it can be made capable of handling smart contracts [279]. On the other hand, the
4969 protocol as described does not have a clear way of supporting light clients, which would be
4970 especially important for a protocol designed to maximize throughput. More research must
4971 also be performed to develop a suitable incentive system for Prism.

4972 11.5. Proof-of-Work DAGs

4973 One of the more common design decisions in the distributed ledger space is to utilize DAGs
4974 rather than singularly linked blockchains. The primary motivation – as with most of the pro-
4975 tocols described above – is to increase the transaction throughput of the system by reducing
4976 the negative security ramifications of block propagation latency. DAG-based ledgers are
4977 similar to GHOST (Section 11.2) in that blocks incorporate references to more than one
4978 previous block but go further by including those blocks directly into the ledger rather than
4979 merely using them as input to determine a particular chain to follow. While not a perfect
4980 analogy, each block produced acts as a confirmation for all prior blocks it references rather
4981 than just one. As a result of this architectural change, blocks can be produced dramatically
4982 more frequently, which allows miners to maintain eventual consistency with incomplete
4983 information about the state of the DAG.

4984 On the other hand, it also complicates the design and analysis of these systems. For ex-
4985 ample, [280] investigates the fairness (i.e., whether a miner’s rewards are proportional to
4986 their hash power) and efficiency (i.e., the fraction of transactions broadcast to the network
4987 that are included in the ledger after a certain period of time) of some DAG-based protocols.
4988 They found that – unlike Nakamoto Consensus – several DAG-based protocols lack fair-

4989 ness even when all miners are honest because fairness is inherently limited by the number of
4990 block pointers that can be included in a given block. Both fairness and efficiency may break
4991 down when miners have varying connectivity to the network, and perversely, large miners
4992 may benefit from having reduced connectivity. While DAG-based schemes are promising,
4993 extreme care must be taken to ensure that they can remain secure and incentive-compatible
4994 in the real world.

4995 11.5.1. Inclusive Blockchains and Conflux

4996 One of the early proposals for a DAG-based ledger was the Inclusive protocol [281], a vari-
4997 ant of which was later employed in the Conflux system [282, 283]. Inclusive was designed
4998 to tolerate larger and more frequent blocks to enable higher throughput without penalizing
4999 miners who are poorly connected to the network. The protocol is called Inclusive because
5000 it includes transactions from all blocks in the DAG into the final ledger. Unlike some other
5001 DAG protocols, Inclusive creates only a single main chain, but that chain incorporates
5002 transactions from blocks that are not ultimately accepted into the main chain. In an Inclu-
5003 sive protocol, miners include a reference to every chain tip that they are aware of in their
5004 own blocks rather than just one. The DAG formed by these blocks and references allows
5005 the "simulation" of any chain selection rule, including the longest chain rule and GHOST.
5006 The first listed reference should be for the block that would be the preferred chain tip based
5007 on the system's chain selection rule.

5008 The algorithm performs a postorder traversal of the block DAG while incorporating valid
5009 transactions from off the main chain into the linearized ledger of accepted transactions. For
5010 a canonical chain $C = B_1, B_2, \dots, B_L$ of blocks B_i , the Inclusive rule will place all the blocks
5011 of the system in a particular order, including those outside of C . Let $past(B)$ be the set of
5012 blocks reachable from B in the DAG. The ordering operates as follows: for $B_i \in C$, insert
5013 before B_i all blocks in the set $past(B_i) \setminus past(B_{i-1})$. The included set of blocks are sorted
5014 topologically with ties broken in favor of the lowest block hash. When blocks are ordered
5015 this way, some invalid or duplicate transactions will exist, which are then removed from
5016 the ledger of transactions.

5017 Fees are given to the miner of the block that a transaction is included in, even if the block
5018 is not in the canonical chain. Say a miner mines a block B , and let $T(B)$ be the set of
5019 transactions included in the ledger from block B . Depending on the block's location in the
5020 DAG (specifically, how quickly the block is referenced by a canonical chain block), the
5021 block producer may only get a fraction of the total fees included in $T(B)$.

5022 Let $before(B)$ be the latest block from the main chain that is reachable from B and $after(B)$
5023 the earliest block from the main chain from which B can be reached. If $after(B)$ does not
5024 exist, it is considered a "virtual block" with height infinity, representing the location of
5025 the next block that a miner with the same view of the ledger would produce. When B
5026 is in the canonical chain, $before(B) = after(B) = B$. Let $gap(B) := after(B).height -$
5027 $before(B).height$ describe the delay in a block's publication with respect to the canonical

5028 chain. The fraction of the fees collected by the miner of B will be a weakly decreasing
5029 function of $gap(B)$. Main chain blocks, as well as blocks that are relatively synchronized
5030 with the main chain, receive the full reward. As $gap(B)$ increases, the fraction of the fees
5031 that the miner collects decreases until a certain cutoff point where they no longer receive
5032 a reward. For example, the reward function may work as follows, where $fee(tx)$ is the
5033 transaction fee for transaction tx :

- 5034 • If $0 \leq gap(B) \leq 3$, the miner collects the full reward, $\sum_{tx \in T(B)} fee(tx)$.
- 5035 • If $3 < gap(B) < 10$, the miner collects $\frac{10-gap(B)}{7} * \sum_{tx \in T(B)} fee(tx)$.
- 5036 • If $gap(B) \geq 10$, the miner gets nothing.

5037 This type of fractional fee scheme has security trade-offs. Providing fees for off-chain
5038 blocks is what allows poorly connected and smaller miners to continue receiving signifi-
5039 cant enough rewards instead of being competed out of existence by larger miners. However,
5040 to the extent that miners of off-chain blocks receive rewards, malicious behaviors are en-
5041 couraged. An attacker who tries and fails to double-spend still receives some revenue from
5042 their off-chain blocks, and thus the attack is subsidized. As a result, double-spending is eas-
5043 ier in Inclusive than it is in, say, Nakamoto Consensus and thus requires waiting for more
5044 confirmations for the same level of security. In addition, Inclusive provides no defense
5045 against selfish mining.

5046 The authors of [281] argue that under a number of game theoretic models, miners will at-
5047 tempt to include transactions that minimize collisions with other blocks rather than merely
5048 include those with the highest fees. The higher performance of Inclusive stems from
5049 this collision-avoiding transaction selection policy because collisions waste bandwidth and
5050 other resources. Unfortunately, this may make miners less likely to broadcast transactions
5051 with high fees to the network because they would prefer to keep those fees to themselves
5052 without risking collisions.

5053 The Conflux protocol is distinct from Inclusive but highly related [282, 283]. The primary
5054 difference is that Conflux blocks include two distinct types of references to other chain
5055 tips instead of treating all references equally, as Inclusive does. By distinguishing between
5056 *parent edges* and *reference edges*, Conflux’s safety follows more directly from the safety
5057 of the GHOST fork-choice rule that it employs (see Section 11.2 to review this rule). In
5058 Conflux, the parent edges act as votes on the proper chain history, whereas reference edges
5059 demonstrate only that the referenced blocks were produced before the block that referenced
5060 them. When a miner mines a block, it sets the parent edge to be the chain tip that follows
5061 GHOST, and the sequence of parent edges create a canonical chain called the *pivot chain*.
5062 In other words, miners compute the pivot chain based on GHOST, then set the tip of the
5063 pivot chain as their parent edge, and any other chain tip off of the pivot chain is set as a
5064 reference edge. Each pivot chain block establishes a new *epoch*, where an epoch contains
5065 every block in the DAG that is reachable from the pivot chain block and is not in a prior




Fig. 31. Example of local Conflux DAG state. The pivot chain is composed of the yellow blocks, which are used to partition the DAG into epochs. The block total order is: Genesis, A, B, C, D, F, E, G, J, I, H, and K. The transaction total order is: T_{x_0} , T_{x_1} , T_{x_2} , T_{x_4} , with T_{x_3} and a duplicate T_{x_4} excluded. [282]

5066 epoch.

5067 Conflux then totally orders the blocks and uses the block ordering to totally order the trans-
 5068 actions. To totally order the blocks, the blocks are first grouped together by epoch – all
 5069 blocks in an epoch come before any block in the following epoch. Within each epoch,
 5070 blocks are ordered based on a topological ordering that follows the reference edges, break-
 5071 ing ties in some deterministic way (e.g., the smallest hash). The total ordering of trans-
 5072 actions follows logically from the ordering of the blocks, with invalid and duplicate trans-
 5073 actions removed. Figure 31 shows an example of Conflux’s procedure for creating a total
 5074 ordering of transactions from the block DAG.

5075 Conflux includes an additional rule that is intended to mitigate the threat of certain liveness
 5076 attacks against GHOST [283, 284]. The structure of the past-subgraph of the DAG is used
 5077 to detect when an attack is underway, at which point the fork-choice rule is adjusted to
 5078 assign blocks an adaptive weight: the weight is h with probability $\frac{1}{h}$, or zero otherwise.
 5079 When no attack is detected, the weight is one. The adaptive weight should help miners
 5080 converge on a single chain by disrupting the ability of an adversary to balance the weights
 5081 of each chain against each other.

5082 11.5.2. SPECTRE and Phantom

5083 SPECTRE (or “Serialization of Proof-of-work Events: Confirming Transactions via Recur-
 5084 sive Elections”) is a DAG-based algorithm that was designed to allow for very high block
 5085 creation rates without reducing the honest majority consensus bound (i.e., regardless of
 5086 network conditions, an attacker requires a majority of the computational power to prevent
 5087 agreement) [285]. Unlike the proof-of-work algorithms considered so far, SPECTRE does
 5088 not lead to a total ordering of transactions but rather a partial order, so it cannot be used for
 5089 arbitrary smart contracts. The algorithm outputs a pairwise ordering between blocks, so it
 5090 is possible that cycles may exist in the ordering. If a total ordering is required, the Phantom
 5091 protocol, which was inspired by SPECTRE, should be used instead [286, 287]. However,
 5092 SPECTRE’s performance stems largely from not running a complete consensus algorithm.

5093 As a result, miners need not be concerned with how well-synchronized other miners are.
5094 This allows very low latencies for transaction acceptance. As block creation rates increase,
5095 this latency reduces approximately to the propagation delay for reaching a large amount of
5096 honest nodes. To understand how these algorithms operate, a few definitions are needed:

- 5097 • $past(B, G)$ represents blocks that were provably created before block B in the DAG
5098 G . Once block B is mined, the set $past(B, G)$ does not change.
- 5099 • $future(B, G)$ represents blocks that were provably created after block B in the DAG
5100 G .
- 5101 • $cone(B, G)$ is the set of blocks in the DAG G that have been ordered with respect to
5102 B . That is, $cone(B, G) := past(B, G) \cup future(B, G) \cup \{B\}$.
- 5103 • $anticone(B, G)$ is the set of blocks in the DAG G that are not directly ordered compared to
5104 B . That is, $anticone(B, G) := G \setminus past(B, G) \cup future(B, G) \cup \{B\}$.
- 5105 • $tips(G)$ is the set of chain tips, or the blocks without any incoming edges in G .
- 5106 • $virtual(G)$ is a non-existent, hypothetical block that satisfies $past(virtual(G)) = G$.
5107 It represents the next block that a miner would create if their view of the DAG was
5108 G .

5109 As with other DAG protocols, miners include references to all known chain tips inside
5110 their blocks. The block DAG is then interpreted in order to extract a partial ordering of
5111 transactions that everyone can agree on. Naturally, if a block $X \in past(Y)$, then X precedes
5112 Y , or $X \prec Y$. Similarly, if $tx_1 \in X$ and $tx_2 \in Y$, $tx_1 \prec tx_2$. For a pair of blocks $(X, Y) \in G$,
5113 all other blocks $B \in G$ are interpreted as votes on the pairwise ordering of X and Y . The
5114 voting rules of SPECTRE correspond to a generalization of Nakamoto Consensus's longest
5115 chain rule applied to DAGs. Specifically, for a block $B \in G \cup virtual(G)$, voting uses the
5116 following rules (an example of these voting rules being applied to a DAG can be seen in
5117 Figure 32):

- 5118 1. If $B \in future(X)$ but $B \notin future(Y)$, then B 's vote is that $X \prec Y$. This rule gives
5119 votes to blocks that were published quickly instead of being withheld.
- 5120 2. If $B \in future(X) \cap future(Y)$, then B 's vote is the same vote as $virtual(past(B))$.
5121 Ties are broken arbitrarily in some agreed upon way. This rule, as well as rule 4,
5122 gives more votes to blocks that are already supported by the majority in order to help
5123 nodes quickly converge on the same precedence relations.
- 5124 3. If $B \notin future(X) \cup future(Y)$, then B 's vote will match that of the majority of blocks
5125 in $future(B)$. This rule counters pre-mining attacks where a block is withheld for a
5126 long time.
- 5127 4. If $B = virtual(G)$, then B 's vote will match that of the majority of blocks in G .

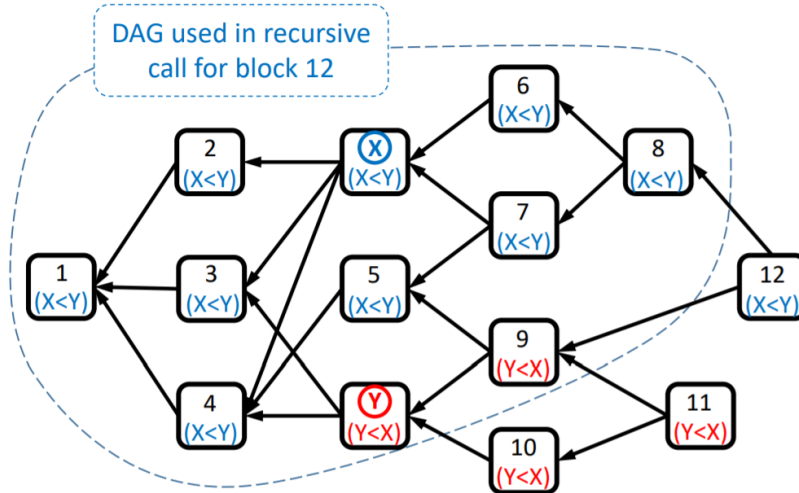


Fig. 32. SPECTRE voting example. Blocks X and 6-8 vote $X \prec Y$ because they see X but not Y in their $past()$. Blocks Y and 9-11 vote $Y \prec X$ for the same reason. Blocks 6-11 vote based on rule 1 in the text, and blocks X and Y vote using rule 5. Blocks 1-5 vote $X \prec Y$ because they see more $X \prec Y$ than $Y \prec X$ votes in their $future()$. Blocks 1-5 use rule 3 to determine their votes. Because block 12 is in $future(X) \cap future(Y)$, it votes according to a recursive call on the DAG that does not include blocks 10, 11, or 12, which are not in its $past()$. Block 12 votes using rule 2. Finally, a miner with this view of the DAG would create a block that references blocks 11 and 12, and its vote would be that $X \prec Y$ based on rule 4. [285]

5128 5. If $B \in \{X, Y\}$, then B 's vote is such that for any block $Z \in past(B)$, $Z \prec B$, and for
 5129 any block $Z' \notin past(B)$, $B \prec Z'$.

5130 For a merchant to consider a transaction tx as confirmed, all of the transaction's inputs must
 5131 be confirmed. Two additional conditions must hold, assuming tx is contained in a block B .
 5132 If there are conflicting transactions in $anticone(B)$, then the blocks where those conflicting
 5133 transactions reside must be preceded by B . Finally, any conflicting transactions in $past(B)$
 5134 must have been rejected.

5135 The voting procedure described above is fairly unintuitive, but a few key ideas can help
 5136 show why SPECTRE is secure. First, if a block is seen by honest miners, then those
 5137 miners will (directly or indirectly) reference it, such that it ends up in the past sets of newly
 5138 created honest blocks. By rule 5, blocks support other blocks in their past. This implies
 5139 that an attacker who withholds their blocks will lose votes. Second, when a block X has
 5140 the majority of votes compared to a potentially conflicting block Y , this majority quickly
 5141 becomes amplified, allowing miners to converge and making it challenging for an attacker
 5142 to reverse the precedence relation. This is because, by rules 2 and 4, new blocks will vote
 5143 the same way as the majority of blocks in their past. Third, by rule 1, blocks created in the
 5144 past will vote based on which competing block is in its future. This incentivizes miners to
 5145 reference recently created blocks in order to solicit those blocks' votes. A block created
 5146 by an attacker that does not reference recent blocks (perhaps because it was withheld in

5147 order to try to double-spend) will lose votes compared to the honest miners. Finally, blocks
5148 from the past will counter-balance pre-mining attacks where blocks are withheld. By rule
5149 3, blocks produced by honest miners who are unaware of the pre-mined block will vote
5150 in line with the majority of blocks in their future set. To see why this matters, consider
5151 an attacker who attempts to double-spend by pre-mining a long secret chain and sending
5152 a conflicting transaction to a merchant, which ends up in a publicly known block. Honest
5153 blocks produced while the double-spending block was withheld will vote with the majority
5154 of blocks in their future set. As long as the attacker does not have the majority of the
5155 computational power, it becomes exponentially more likely as time passes that past blocks
5156 will support the block paying a merchant.

5157 Unfortunately, SPECTRE has a number of implementation complexities. The handling of
5158 difficulty adjustments and transaction fees are out of scope for this document. Further-
5159 more, there are practical limits to how many blocks can be referenced by any given block
5160 and, thus, the extent to which block creation rates can be increased before the overhead
5161 of the references themselves becomes prohibitive. That said, the throughput and latency
5162 improvements are still substantial.

5163 The Phantom protocol is based off of the same ideas as SPECTRE but uses some additional
5164 techniques to enforce a total ordering of transactions. Thus, it is suitable for a broader
5165 variety of settings at the cost of higher transaction settlement latency [286, 287]. Note
5166 that the original Phantom protocol from [286] suffered from a liveness issue where a low-
5167 hash-rate attacker could delay transaction confirmation indefinitely, as pointed out in [282].
5168 This is fixed in [287]. Technically, Phantom requires solving an NP-hard problem, and
5169 GHOSTDAG is the algorithm that approximates an "ideal" Phantom, but this section will
5170 use the more well-known name Phantom for both.

5171 Honest miners in Phantom (as well as other systems) are expected to broadcast their blocks
5172 as quickly as they are found. They are also expected to reference as many chain tips as
5173 they are aware of and ideally be well-connected to each other. If the majority of miners
5174 are honest, this should result in a cluster of blocks that are well-connected to each other.
5175 Adversarial miners may, in contrast, withhold blocks temporarily or create new blocks
5176 that do not reference publicly visible chain tips in order to overwrite other blocks. These
5177 malicious behaviors are detectable based on the structure of the DAG because they will not
5178 be as well connected to the majority honest cluster (although, as with Nakamoto Consensus,
5179 this malicious behavior cannot be distinguished from a network partition).

5180 In more detail, let Δ be the upper bound on the network's propagation delay, and assume
5181 that an honest miner found a block B at time t . Then all blocks that were broadcast by
5182 time $t - \Delta$ will be in $past(B)$, and B will be in the past set of any honestly produced blocks
5183 after $t + \Delta$. Recall that $anticone(B)$ is the set of blocks that are not referenced by B and
5184 that do not reference B . This implies that $anticone(B)$ will contain only malicious blocks
5185 (ignoring network partitions) and a small number of honest blocks produced in the 2Δ -
5186 sized interval $[t - \Delta, t + \Delta]$. Due to the nature of proof of work, this number of blocks will

5187 be bounded by some k with high probability. That is, the parameter k is a function of the
 5188 network delay. The idea behind Phantom is to recognize and select the largest cluster of
 5189 blocks in the DAG by observing these anticones. Specifically, Phantom attempts to solve
 5190 the *maximum k -cluster subDAG problem*: given a DAG $G = (V, E)$, Phantom attempts to
 5191 output a maximally sized subset $S^* \subset V$ such that $|anticone(B) \cap S^*| \leq k$ for all $B \in S^*$.

5192 Phantom generalizes Nakamoto Consensus to a DAG, where every block produced ulti-
 5193 mately ends up in the ledger, but adversarially produced blocks should appear later in the
 5194 total ordering. This is done by extracting a well-connected cluster of blocks that are pre-
 5195 sumed to be honest by the honest majority assumption and then totally ordering blocks in
 5196 a way that prioritizes blocks inside of the cluster. Phantom uses a greedy algorithm that
 5197 approximates a solution to the maximum k -cluster subDAG problem, where blocks within
 5198 the cluster are called Blue blocks, and ones outside of the cluster are Red. The algorithm,
 5199 $OrderDAG(G, k)$, outputs a set of Blue blocks and an ordered list of all blocks in G :

- 5200 1. If $G = \{genesis\}$, then return $[\{genesis\}, \{genesis\}]$.
- 5201 2. For $B \in tips(G)$ do: $[BlueBlocks_B, OrderedBlocks_B] \leftarrow OrderDAG(past(B), k)$. That
 5202 is, use recursion on the chain tips in order to find the best tip.
- 5203 3. Inherit the Blue set of the best tip, and add this tip to the Blue set and the end of the
 5204 ordered list.
 - 5205 • $B_{max} \leftarrow argmax\{|BlueBlocks_B| : B \in tips(G)\}$ (break ties according to lowest
 5206 hash)
 - 5207 • $BlueBlocks_G \leftarrow BlueBlocks_{B_{max}}$
 - 5208 • $OrderedBlocks_G \leftarrow OrderedBlocks_{B_{max}}$
 - 5209 • add B_{max} to $BlueBlocks_G$
 - 5210 • add B_{max} to the end of $OrderedBlocks_G$
- 5211 4. For $B \in anticone(B_{max}, G)$ do (in some topological ordering): If $BlueBlocks_G \cup$
 5212 $\{B\}$ is a k -cluster, then add B to $BlueBlocks_G$. Either way, add B to the end of
 5213 $OrderedBlocks_G$.
- 5214 5. Return $[BlueBlocks_G, OrderedBlocks_G]$.

5215 To summarize, the DAG is colored recursively, so the Blue set will include all of the Blue
 5216 blocks from the chain tip with the largest Blue set in its past (denoted B_{max}). New Blue
 5217 blocks are then added from blocks that lie outside of $past(B_{max})$ but not before checking
 5218 whether the k -clustering would be violated by their inclusion (step 4). The ordering works
 5219 similarly and begins by inheriting the ordering from B_{max} . Then B_{max} is next, and then
 5220 blocks that lie outside of $past(B_{max})$ are topologically ordered in some agreed-upon way.
 5221 An example of this algorithm is in Figure 33. Once blocks are totally ordered, transac-




Fig. 33. Phantom example, constructing $BlueBlocks_G$ with parameter $k = 3$. For each block, the circle near it represents its "score," or the number of Blue blocks in its past set. The algorithm begins at the highest scoring tip, $B_{max} = M$, and greedily selects its predecessors: K , H , D (arbitrarily breaking the tie between C , D , and E), and the genesis block. This creates a chain from B_{max} back to genesis. The block V is a hypothetical "virtual" block that references all tips of the DAG. The set $BlueBlocks_G$ begins empty and is constructed recursively as follows. (1) Visit D , and add the genesis block to $BlueBlocks_G$ since it is the only block in $past(D)$. (2) Visit H , and add blocks C , D , and E to $BlueBlocks_G$ because they are in $past(H)$. (3) Visit K , and add blocks H and I to $BlueBlocks_G$. Block B is in $past(K)$, but $anticone(B)$ has 4 Blue blocks and thus is not included. (4) Visit M , and add K to $BlueBlocks_G$. Block F is not added because, like block B , $anticone(F)$ has more than k Blue blocks. (5) Visit the virtual block, V , and add M to $BlueBlocks_G$. Block L is not added because $anticone(L)$ includes C , H , K , and M , which exceeds k (there is an error in the image, and L should not be colored Blue). Block J is not added because the inclusion of J would add another Blue block to to $anticone(I)$, which already contains blocks C , D , and H . [287]

5222 tions are totally ordered in the natural way, removing duplicates and invalid transactions.
 5223 Note that, similar to Inclusive (Section 11.5.1), miners are incentivized to randomize their
 5224 transaction selection in order to maximize fees, though this may also encourage them to
 5225 withhold high-fee transactions.

5226 It is instructive to compare Phantom to other protocols. Recall that Δ is the upper bound
 5227 on network propagation delay, but as with Nakamoto Consensus, its value is unknown.
 5228 That said, it is assumed to be smaller than a constant Δ_{max} , which is used to derive the
 5229 hard-coded parameter k . This parameter is the maximum number of blocks that may not be
 5230 referenced by each other that can be created by the full mining network over the course of a
 5231 single delay (and a larger k requires increasing the waiting time for transaction settlement).
 5232 Phantom's security model differs from that of SPECTRE primarily due to the reliance on
 5233 Δ_{max} , which must be used explicitly in Phantom but is what allows a total order to be
 5234 established.

5235 Let the block creation rate of a system be λ ($\lambda = \frac{1}{600}$ blocks per second in Bitcoin, for
5236 example). In Nakamoto Consensus, the security threshold of the system decreases toward
5237 zero as $\Delta\lambda$ increases. That is, the less time there is to synchronize between blocks being
5238 found, the lower the security threshold. In contrast, as long as $\Delta \leq \Delta_{max}$, Phantom's se-
5239 curity threshold is at least $\frac{1}{2} * (1 - \epsilon)$ for some small ϵ . As a result, Phantom can tolerate
5240 much higher block creation rates and throughput while maintaining security under honest
5241 majorities. Put differently,

- 5242 • Nakamoto Consensus assumes that $\Delta\lambda \ll 1$.
- 5243 • The parallel chains approach with m chains, described in Section 11.4, assumes that
5244 $\frac{\Delta\lambda}{m} \ll 1$.
- 5245 • Phantom assumes $\Delta\lambda \ll k$.

5246 11.5.3. Tangle

5247 The other DAG-based protocols discussed in this section construct DAGs where the vertices
5248 are blocks, but some designs – including the Tangle – use transactions as vertices instead
5249 [288]. A primary motivation for the Tangle structure is to enable fee-less transactions by
5250 allowing typical clients to operate as miners. Instead of submitting transaction fees, each
5251 transaction is accompanied by a small proof of work that includes validating and approving
5252 two other transactions within the Tangle.

5253 The most important concept in a Tangle-based system is the *tip selection strategy*, or
5254 how a client chooses which two transactions at the "tip" of the DAG to reference and ap-
5255 prove. This policy cannot be imposed by the network, so it is imperative that an incentive-
5256 compatible and secure default exists. The security assumption behind the Tangle is that
5257 there must be a large enough inflow of transactions posted to the network that are gener-
5258 ated by honest clients to outweigh the computational ability of an adversary. Honest clients
5259 must frequently issue transactions for this to happen and for the Tangle to work in a permis-
5260 sionless environment without centralized coordination [289]. If this assumption holds, the
5261 Tangle can maintain a partial order (not a total order) over transactions. As with SPECTRE
5262 (Section 11.5.2), this makes it unsuitable for generic smart contracts.

5263 Transactions in the Tangle have a *weight* associated with them, which corresponds to the
5264 amount of work performed to issue it. The *own weight* of a transaction specifically refers
5265 to the work performed to issue a transaction, which is normalized to one in this document.
5266 Transactions also have a *cumulative weight*, which is the transaction's own weight plus the
5267 sum of the own weights of every transaction that approves of it (directly or indirectly). A
5268 transaction's *score* is its own weight plus the sum of the own weights of all transactions
5269 directly or indirectly approved by it. Stated differently, a transaction's cumulative weight
5270 is the sum of the own weights of the transaction's future set and itself, while the transac-
5271 tion's score corresponds to the own weights of all transactions in its past set and itself. A

5272 transaction's *height* is the length of the longest path from it to the genesis transaction, and
5273 its *depth* is the length of the longest (reverse) path from it to some chain tip.

5274 The originally proposed tip-selection algorithm is dubbed the *Markov chain Monte Carlo*
5275 algorithm (MCMC). It starts by choosing some locations in the Tangle and then performing
5276 random walks toward chain tips. Specifically, where H_{tx} is the cumulative weight of a
5277 transaction tx in the Tangle, and using W and α as parameters, a client selects chain tips to
5278 approve as follows:

- 5279 1. Consider all transactions in a particular range of the Tangle, $[W, 2W]$, as possible
5280 starting locations. Randomly choose N of these transactions.
- 5281 2. Perform independent, biased random walks from these N locations based on prior
5282 approvals, such that the walk moves from tx_1 to tx_2 only if tx_2 directly referenced tx_1 .
5283 Let the set of transactions that directly reference tx_1 be denoted TX . The probability
5284 of transitioning from tx_1 to tx_2 is $\frac{e^{-\alpha(H_{tx_1} - H_{tx_2})}}{\sum_{tx \in TX} e^{-\alpha(H_{tx_1} - H_{tx_2})}}$.
- 5285 3. Once two of these random walks end at chain tips, select those tips. However, if a
5286 random walk ends up at a chain tip "too quickly," it may be considered a *lazy tip* and
5287 discarded. A lazy tip is one that approved of an old transaction in order to avoid the
5288 effort of verifying transactions. Lazy tips have a low probability of being selected
5289 because the cumulative weight of the lazy tip and any others would be substantial.

5290 The parameter α used in MCMC is important. Higher values of α (closer to one) are more
5291 "deterministic" and provide a better defense against various adversarial strategies because
5292 they increase the chance of the random walk moving to high-scoring tips. A lower value
5293 of α (closer to zero) makes the system more stable with respect to transaction confirmation
5294 times. A high α will result in many stale tips that need to be reattached to the Tangle in
5295 order to achieve confirmation.

5296 Some potential attacks on the Tangle include *parasite chain* attacks, *large weight* attacks,
5297 and *splitting* attacks. The parasite chain attack, displayed in Figure 34, is a classic double-
5298 spend attack. The attacker builds a subtangle in secret while occasionally referencing the
5299 main Tangle in order to inflate the secret subtangle's score. When the attacker has less
5300 computational power than the honest portion of the network, the parasite chain attack is
5301 challenging for the same reason that selecting lazy tips is unlikely: the attacker's subtangle
5302 is likely to have lower cumulative weight, so the random walks will likely remain on the
5303 main Tangle. Other techniques for detecting and mitigating parasite chain attacks have
5304 been proposed as well [290].

5305 In a large weight attack, as shown in Figure 35, an adversary attempts to double-spend by
5306 putting an especially large amount of work into the double-spending transaction in order to
5307 outweigh the honest portion of the Tangle. To protect against this attack, there should be
5308 an upper bound on the own weight of any given transaction.




Fig. 34. Parasite chain attack against the Tangle. The two red circles indicate a double-spend attempt by the attacker. [288]

5309 Finally, a splitting attack is one in which the adversary attempts to divide the Tangle into
5310 two incompatible branches and then keep them balanced over time until they can spend
5311 the same funds on both sides of the split. The attacker issues conflicting transactions near
5312 the beginning of the split so that the two sides cannot be reconnected. If honest miners
5313 are divided between the two subtangles, a low-hash-rate attacker can attempt to mine on
5314 whichever side of the split is required to maintain balance. Starting the random walks at
5315 transactions with greater depths in the Tangle makes this attack more challenging to pull
5316 off. In addition, the splitting attack is much easier with low α . When higher, random walks
5317 quickly converge toward the subtangle with higher cumulative weight.

5318 In addition to the explicit attacks on the Tangle, there are potential game-theoretic issues
5319 relating to the incentives of the participants. Intuitively, in order to encourage more clients
5320 to approve of one's transactions, one would want to place their transactions in the heavier
5321 subtangles because that is where the random walks are more likely to end up. This sug-
5322 gests that incentives are aligned properly, but an alternative strategy would be to simply
5323 remember the last 10 or so transactions that were gossiped and approve two of those when
5324 issuing a transaction instead of storing the full Tangle.² Assuming that most other clients
5325 are honest, then the most recently seen transactions are likely to be where random walks
5326 end up. If the majority of clients were to use this strategy, then transactions would con-
5327 tinue to be approved and everything would appear normal, but the actual security bound




Fig. 35. Large weight attack against the Tangle. [288]

²Suggested in <https://twitter.com/AlexSkidanov/status/1130505820930695169>.

5328 decreases toward zero. An attacker can cause a deep fork by quickly posting multiple (in
5329 this case, more than 10) transactions in a row on a different subtangle, at which point all of
5330 the lazy clients will immediately switch to it and begin approving them.

5331 A number of works have explored the ramifications of alternative tip-selection algorithms
5332 [291–294]. Simple algorithms – such as uniformly random tip selection and the use of
5333 unbiased random walks – are highly susceptible to parasite chain attacks, while using bi-
5334 ased random walks (MCMC) with large α protects against attacks [291]. A modification to
5335 the MCMC algorithm that makes parasite chain attacks more difficult while maintaining a
5336 lower α is to incorporate the derivative of the cumulative weight with respect to time in the
5337 random walk probabilities [292]. Another proposal, dubbed G-Iota, attempts to maintain a
5338 higher α while mitigating the issue that this leads to more honest chain tips becoming stale
5339 [293]. Finally, the E-Iota proposal attempts to reduce the number of random walks that
5340 need to be executed while maintaining the best security guarantees of MCMC and G-Iota
5341 [294]. The algorithm is parameterized by $p_1 < p_2 < 1$. Any time the client needs to select
5342 new tips, it generates a random number r . If $r < p_1$, the client uses N uniform random
5343 walks for tip selection. If $p_1 \leq r < p_2$, it uses N biased random walks with a low α value.
5344 If $r \geq p_2$, it uses N biased random walks with high α .

5345 11.5.4. Meshcash

5346 Meshcash is a modular consensus algorithm that combines aspects of proof of work with
5347 asynchronous binary Byzantine agreement (ABA) protocols (see Section 6.1.1) [295]. It
5348 involves a slower proof-of-work protocol, dubbed the "tortoise," that provides eventual
5349 consensus, as well as a quicker "hare" protocol. If the hare protocol succeeds, agreement
5350 occurs quickly, but if it fails, the tortoise protocol will ensure eventual consistency. The
5351 Meshcash protocol provides the following minimal security guarantees, which are not tight:

- 5352 • If an adversary controlling less than $\frac{1}{3}$ of the system's computational power is unable
5353 to disrupt the hare protocol, then the Meshcash protocol achieves consensus.
- 5354 • The tortoise protocol achieves consensus against adversaries who control less than
5355 $\frac{1}{15}$ of the computational power, regardless of the outcome of the hare protocol.

5356 Both the ABA protocol and the tortoise protocol rely on a *weak common coin* (as defined in
5357 Section 2.7). This is implemented via the proof-of-work algorithm with an expected block
5358 time of T (e.g., $T = 600$ seconds in Bitcoin). For a player P beginning at time t , the weak
5359 common coin protocol works as follows:

- 5360 1. P waits until time $t + T$ and keeps track of the set of valid blocks received between
5361 time t and $t + T$, denoted S_P .
- 5362 2. P sorts the blocks in S_P by their hashes. The weak common coin output is the least
5363 significant bit of the smallest hash block in S_P .

5364 The Meshcash ledger is constructed as a layered DAG, where each block belongs to a
5365 particular layer and references blocks from earlier layers (the number of layers that can
5366 be referenced by miners depends on the hare protocol). The hare protocol is used for
5367 consensus on blocks from more recent layers, while the tortoise protocol orders blocks in
5368 the more distant past. Honest miners must remain relatively synchronized with respect to
5369 the layers they are participating in, so the layer counter is incremented whenever a miner
5370 sees a threshold of valid blocks in a given layer. The hare protocol has two additional
5371 requirements:

- 5372 1. Blocks on layer i are valid only if there are at least T_{min} valid blocks on layer $i - 1$ in
5373 their past sets. This prevents miners from pre-mining blocks in future layers.
- 5374 2. For a block B in layer i , honest miners with layer counters in the range $[i + t, i +$
5375 $s]$ agree on the validity of B and will continue to do so throughout the interval
5376 $[start_{i+t}, start_{i+s+1})$, where $start_i$ is the time that the first honest miner entered layer
5377 i . This property is called *limited $[t, s]$ -consistency*.

5378 A variety of suitable hare protocols are possible. In the hare protocol, blocks on a given
5379 layer are used to elect a committee that runs a traditional ABA protocol (e.g., the one
5380 described in Section 6.1.1) off-chain in order to get agreement on the blocks within the
5381 layer and then to append signatures to these blocks to attest to their validity. An ABA
5382 instance is run for each block in the layer. A block claiming to be in layer i is valid if and
5383 only if it has valid signatures from the majority of layer- i committee members.

5384 The tortoise protocol is derived from a hare protocol Π with output interval $[t, s]$. In addition
5385 to any requirements from Π , the tortoise protocol requires that honestly produced blocks
5386 include references to every chain tip (*view references*) as well as references to all valid
5387 blocks in layers $[i - s, i - t]$ (*voting references*). Additionally, given a network propagation
5388 delay upper bound of Δ , miners include a *coin bit*, a *before coin bit*, and an *early block bit*.
5389 The coin bit is the result of the weak common coin protocol that begins at time $start_i + \Delta$.
5390 The before coin bit is set to indicate whether the block was produced before the common
5391 coin protocol concluded in order to abstain from voting when the coin bit matters. Finally,
5392 the early block bit is set if the block was produced less than Δ time after the layer began
5393 so that these blocks abstain from voting if late blocks from the prior layer would make a
5394 difference.

5395 The tortoise protocol's block voting procedure, an example of which is in Figure 36, is
5396 performed on any block B in layer $i' < i - s$. All valid blocks in the range $[i' + 1, i - 1]$ vote
5397 on B if they were received by $start_i + \Delta$. Votes are weighted by the proof-of-work difficulty
5398 of the blocks, and B is valid if this weighted sum is positive. A block Q in layer $j > i'$ votes
5399 on block B using the following rules and a protocol-defined threshold θ :

- 5400 1. When $j < i' + t$, Q votes zero because Q was generated before the hare protocol had
5401 a chance to achieve consensus on B .




Fig. 36. Example of Meshcash block voting. Here, $s = 4$ and $t = 3$. The blocks in layers $[i - s, i - t]$ vote 0 by rule 1. The blocks in layers $[i - 2, i - 1]$ vote according to rule 2. Three of these blocks have voting edges to B , and two do not. Block A considers block B valid by rule 3 because the sum of block votes in its past in favor of B is positive. This assumes that, weighted by the proof-of-work difficulty, this sum is greater than θ . [295]

- 5402 2. When j is in the range $[i' + t, i' + s]$, Q votes 1 if it has a voting reference to B and -1
 5403 otherwise.
- 5404 3. When $j > i' + s$, Q will vote in agreement with the weighted sum of blocks in its past
 5405 set so long as this sum is outside of the range $[-\theta, \theta]$. However, if this weighted sum
 5406 is in the range $[-\theta, \theta]$, then Q 's vote will be determined by the coin bit and before
 5407 coin bit:
- 5408 • If the before coin bit is set to 1, then Q votes 0.
 - 5409 • If the before coin bit is set to 0 and the coin bit is set to 1, then Q 's vote is 1.
 - 5410 • If the before coin bit is set to 0 and the coin bit is set to 0, then Q 's vote is -1.

5411 This consensus protocol is compatible with any way of allocating rewards, but the authors
 5412 suggest that transaction fees be distributed to all miners who created blocks in recent layers
 5413 in order to decrease the incentive for miners to keep high-fee transactions secret. This
 5414 would also make Meshcash more resistant to selfish mining, not unlike FruitChains (see
 5415 discussion in Section 11.3). Specifically, the total fees from layer i would be split among
 5416 the miners of the prior k blocks proportionally to the number of blocks in each layer. By
 5417 splitting the reward in this way, double-spending attacks become subsidized, but censorship
 5418 becomes more challenging.

5419 11.6. Proof of Work for Committee Selection

5420 An alternative approach to permissionless consensus is to use proof of work as a Sybil-
5421 resistance measure for electing a committee of miners to act as replicas in a standard per-
5422 mitted BFT protocol. Some of the best advantages of proof of work are removed when
5423 using schemes that elect committees in this way. In particular, this introduces the *posterior*
5424 *corruption* issue that separates proof-of-stake protocols from many proof-of-work ones,
5425 making the protocols *weakly subjective* (see Section 12.1.2). Essentially, this means that a
5426 new node joining the network (or returning online after an extended absence) must acquire
5427 a copy of the chain from a trusted party. If past committees are compromised, they can
5428 create alternative histories that would be believed by these newly online nodes. In addition,
5429 these protocols are less safe against bribery attacks because once a node is in the commit-
5430 tee, no additional proof-of-work resource expenditure is needed for it to issue malicious
5431 statements.

5432 On the other hand, these protocols are able to provide *responsiveness*, or the ability to
5433 confirm transactions at the speed of the actual network delay, rather than the worst-case
5434 delay (see Section 5.2). For a permissionless proof-of-work system to be responsive, [296]
5435 proved that four conditions must hold:

- 5436 1. The protocol must know an upper bound on the network delay, Δ .
- 5437 2. There must be a non-responsive "warmup" period, after which transaction confirma-
5438 tion can become responsive.
- 5439 3. There must be some "stickiness" to honest nodes. That is, it must take some time for
5440 an adversary to corrupt a node or knock it offline. These protocols can only be secure
5441 against mildly adaptive adversaries, not fully adaptive ones.
- 5442 4. Fewer than $\frac{1}{3}$ of the nodes can be corrupt.

5443 11.6.1. Hybrid Consensus

5444 The Hybrid Consensus algorithm uses proof of work to agree on a rotating committee, and
5445 the committee uses a traditional BFT algorithm to select transactions [296]. The protocol
5446 utilizes an underlying blockchain protocol, such as Nakamoto Consensus or FruitChains,
5447 though FruitChains is preferable (see Section 11.3 for discussion on FruitChains). Miners
5448 who have successfully mined blocks in the recent past are elected to a committee that then
5449 runs the classic PBFT algorithm among themselves (PBFT is described in Section 4).

5450 In more detail, Hybrid Consensus works as follows, where λ is the common prefix con-
5451 sistency parameter (all nodes agree on the contents of the blockchain, except for possibly
5452 the trailing λ blocks). A new committee is elected whenever the underlying blockchain
5453 adds an additional λ blocks. To elect the committee safely, the trailing $\Theta(\lambda)$ blocks must
5454 not be considered because they are potentially unstable. The most recent λ blocks' min-
5455 ers within the common prefix are the new committee members. This means that the same

5456 miner may occupy multiple spots in the committee proportional to the number of blocks
5457 they have mined in the stable portion of the chain. Prior works that attempted to elect BFT
5458 committees using proof of work failed to remove the trailing $\Theta(\lambda)$ blocks, which made
5459 them insecure because there was no agreement on the members of the committee itself. By
5460 the chain growth property, new committees are elected at regular, somewhat predictable
5461 intervals.

5462 Finally, consider chain quality. If the underlying blockchain has at least $\frac{2}{3}$ -chain quality,
5463 then it can guarantee that for each period of λ consecutive blocks, at least $\frac{2}{3}$ of them will be
5464 mined by honest replicas. This is necessary to ensure that at least $\frac{2\lambda}{3}$ of the BFT committee
5465 members are honest. If Nakamoto Consensus were used for the underlying blockchain
5466 protocol, then $\frac{3}{4}$ of the hash rate would need to be honest in order to achieve $\frac{2}{3}$ -chain quality
5467 (see discussion on selfish mining in Section 9.4). This motivates the use of FruitChains as
5468 the underlying ledger because it has a higher chain quality than Nakamoto Consensus and,
5469 thus, allows near optimal resilience.

5470 As described, if an old committee ever surpasses the adversarial corruption threshold, Hy-
5471 brid Consensus is subject to the issue of posterior corruptions, where new nodes will be
5472 unable to tell which of two conflicting chain forks is the correct one. To protect against
5473 this, whenever a committee is switched out and a new committee is elected, at least $\frac{1}{3}$ of
5474 the old committee will sign a hash of the ledger (or rather, the portion of the ledger that the
5475 committee worked on) and include those signatures in the blockchain. This prevents old
5476 committees from equivocating.

5477 Transitioning smoothly between multiple consensus committees is also non-trivial. The old
5478 committee must undergo some stopping procedure that will overlap with the new commit-
5479 tee's reign. Because of this overlap, the new committee's output should be deferred until
5480 the old committee has fully completed its term. The stopping procedure involves sending
5481 special signed stop instructions until more than $\frac{1}{3}$ of the committee have done so, at which
5482 point transactions are ignored.

5483 11.6.2. Solida

5484 Solida is another proposal that combines aspects of proof of work with classical BFT con-
5485 sensus [297]. Unlike with Hybrid Consensus, Solida does not use proof of work to es-
5486 tablish an underlying blockchain but rather as a more generic Sybil-resistance mechanism
5487 for leader election. A modified version of PBFT (Section 4) is used to commit blocks of
5488 transactions or *reconfiguration events* for the consensus committee into the ledger. In order
5489 to join the committee, miners must find a proof of work for a computational puzzle, and the
5490 existing committee tries to commit the miner's public key, the proof of work, and the sys-
5491 tem state into the ledger. The miner then joins the committee, pushing the oldest member
5492 of the committee out of it.

5493 The normal PBFT leader cannot be in charge of these reconfiguration events because a

5494 Byzantine leader can stall the network until the adversary can generate more proofs of
5495 work and become over-represented on the committee. Instead, it is the successful miner
5496 who leads the attempt to elect themselves onto the committee. The new member immedi-
5497 ately becomes the new PBFT leader for committing transactions. The challenge with this
5498 approach is how to address contention issues – that is, when more than one miner finds a
5499 proof of work at around the same time. Solida addresses this with a ranking scheme, such
5500 that only higher ranked miners can interrupt the reconfiguration of lower ranked miners.

5501 Solida is essentially composed of three subprotocols: the steady state protocol to commit
5502 transactions into *slots*, denoted s ; the view change protocol to replace faulty leaders; and the
5503 reconfiguration protocol. Any given configuration c can have multiple *lifespans*, denoted
5504 e , and each lifespan can have multiple views v . Each Solida leader $L(c, e, v)$ is ranked in c ,
5505 e , v order and uses the following rules:

- 5506 1. When reconfiguration events are committed, each committee member switches to a
5507 new configuration by incrementing c and setting $e = v = 0$.
- 5508 2. When a new proof of work is found under the current configuration, each committee
5509 member increments e and sets $v = 0$.
- 5510 3. If a timeout occurs and a faulty leader is detected, each committee member incre-
5511 ments v . For some hash function H in a system with n committee members, let
5512 $l = H(c, e) + v \bmod n$. Then $L(c, e, v)$ is the l -th member of the existing committee.

5513 The steady state protocol for committing blocks of transactions is nearly identical to PBFT
5514 with two changes: 1) the original pre-prepare, prepare, and commit phase messages include
5515 extra contextual information, particularly the (c, e, v) tuple; and 2) an additional "notify"
5516 step occurs at the end. After committing a block, replicas send a NOTIFY message that
5517 includes a commit certificate and move to slot $s + 1$. Upon receiving a NOTIFY message,
5518 replicas commit, broadcast their own NOTIFY message, and move to slot $s + 1$. The view
5519 change protocol requires more substantial adjustments and works as follows when the net-
5520 work delay upper bound is Δ :

- 5521 1. Upon advancing to the next slot s , replicas set a timer and initiate a view change
5522 if it reaches 4Δ before the replica has committed anything in slot s . It does so by
5523 broadcasting a VIEW-CHANGE(c, e, v) message to the rest of the committee. If a
5524 replica sees $2f + 1$ matching VIEW-CHANGE(c, e, v) messages and has not already
5525 advanced to a higher view, it forwards the set of VIEW-CHANGE(c, e, v) messages
5526 to the new leader $L(c, e, v + 1)$. The replica then listens for a NEW-VIEW message
5527 from the new leader, and if they do not receive it within 2Δ time, they broadcast a
5528 VIEW-CHANGE($c, e, v + 1$) message.
- 5529 2. When $L(c, e, v + 1)$ receives $2f + 1$ matching VIEW-CHANGE messages, they broad-
5530 cast a NEW-VIEW($c, e, v + 1$) message (which includes the set of VIEW-CHANGE
5531 messages) and enter $(c, e, v + 1)$. When replicas receive this message and are not

5532 already in a view higher than $(c, e, v + 1)$, they begin the new view and start another
5533 timer. If the timer hits 8Δ and no slot has been committed, replicas give up on the
5534 current leader and broadcast another VIEW-CHANGE message.

5535 3. When entering view (c, e, v) , a replica sends a STATUS message to $L(c, e, v)$ that
5536 contains the slot number $s - 1$, the value committed in that slot (denoted h), the
5537 corresponding commit certificate C , the value accepted in slot s (denoted h'), and
5538 the corresponding prepare certificate P . When the leader receives $2f + 1$ STATUS
5539 messages, they concatenate them into a status certificate, S . The leader then finds the
5540 STATUS message that corresponds to the highest last-committed slot s^* , breaking
5541 ties using the message that contains the highest ranked prepared value in slot $s^* + 1$.
5542 Denote the commit certificate and prepare certificate from this message as C^* and
5543 P^* , respectively.

5544 4. The new leader broadcasts a REPROPOSE message that includes $s^* + 1$, h' , S , C^* ,
5545 and P^* . Receiving C^* allows replicas to commit slot s^* . In addition, the inclusion
5546 of S and P^* prove that h' is a safe value for slot $s^* + 1$, so h' is repropose for that
5547 slot. If the message is valid, replicas commit slot s^* if they have not already and then
5548 move into the prepare phase of the steady state protocol for slot $s^* + 1$.

5549 The reconfiguration protocol requires miners to submit a proof-of-work solution to a configuration-
5550 specific puzzle, $puzzle(c)$. The puzzle difficulty is periodically updated in order to maintain
5551 an expected average reconfiguration interval that is analogous to the block interval in other
5552 systems. Naively, this would allow an adversary to gain an advantage over the honest por-
5553 tion of the network by withholding their puzzle solutions analogously to selfish mining. To
5554 bound the advantage that withholding provides an attacker, $puzzle(c + 1)$ includes any set
5555 of $f + 1$ commit certificates from the NOTIFY messages in the most recent reconfigura-
5556 tion. This ensures that adversaries learn a puzzle at most 2Δ earlier than honest miners.
5557 The reconfiguration protocol operates as follows:

5558 1. A successful miner broadcasts their puzzle solution to the committee members. Com-
5559 mittee members then enter view $(c, e + 1, 0)$ with the miner as the new leader and start
5560 a timer. If no slot is committed by the time the timer hits 8Δ , replicas initiate a new
5561 view change.

5562 2. In the new lifespan, replicas perform the same actions that they would in the STATUS
5563 step of the view change protocol.

5564 3. Let h^* be the committed value in the highest committed slot s^* in the status certifi-
5565 cate and h' be the highest ranked prepared value in slot $s^* + 1$. The new external
5566 leader will then take action based on whether h^* and h' are blocks of transactions or
5567 reconfiguration events:

5568 (a) If h^* is a reconfiguration event that advances to configuration $c + 1$, the leader
5569 broadcasts the commit certificate C and gives up joining the committee. An

5570 alternative leader has already completed reconfiguration, so the miner begins
5571 working on $puzzle(c + 1)$.

5572 (b) If h^* is a block of transactions and h' is a reconfiguration event to $c + 1$, then
5573 the leader broadcasts a REPROPOSE message that includes $(c, e, v = 0, s^* + 1, h', S, C^*, P^*)$ and terminates the reconfiguration protocol.
5574

5575 (c) If h^* is a block of transactions and h' is empty, then the leader attempts to move
5576 to slot $s^* + 1$ by broadcasting a REPROPOSE message that includes $(c, e, v =$
5577 $0, s^* + 1, h, S, C^*, P^*)$, where h is a reconfiguration event that would allow the
5578 leader to join the committee. Should h be committed, then in slot $s^* + 2$, replicas
5579 advance to the next configuration that includes the miner as a replacement for
5580 the oldest committee member.

5581 (d) If h^* and h' are both blocks of transactions, then the leader broadcasts a RE-
5582 PROPOSE message that includes $(c, e, v = 0, s^* + 1, h', S, C^*, P^*)$ in order to
5583 repropose h' for slot $s^* + 1$. Next, the leader tries to join the committee by get-
5584 ting a reconfiguration event, h , into slot $s^* + 2$. They do this with a PROPOSE
5585 message that includes $(c, e, v = 0, s^* + 2, h)$. Should h be committed, then in
5586 slot $s^* + 3$, replicas advance to the next configuration that includes the miner as
5587 a replacement for the oldest committee member.

5588 Except for case 3a, committee members handle receipt of the REPROPOSE message iden-
5589 tically to how they would during the view change protocol.

5590 This approach differs markedly from the Hybrid Consensus protocol described earlier. One
5591 benefit of this approach is that the identities of committee members need not be buried
5592 under several blocks' worth of work and, thus, are publicly exposed for a shorter period
5593 of time (though they are still exposed, so Solida cannot be secure against fully adaptive
5594 adversaries either). Note that as the adversarial hash rate approaches 33%, the number of
5595 committee members required to maintain safety blows up toward infinity, and the system
5596 is unable to recover if a committee ever has $f + 1$ Byzantine replicas in it.

5597 12. Proof of Stake: The Basics

5598 Proof of stake is the first major permissionless Sybil-resistance mechanism introduced af-
5599 ter proof of work and is, in part, motivated by the desire to reduce the high electricity
5600 consumption of proof of work while achieving similar goals. This section introduces the
5601 idea and provides the context necessary for understanding the more advanced protocols de-
5602 scribed in Section 13. First, some of the earliest proof-of-stake systems are discussed and
5603 then used to demonstrate security issues that are unique to proof of stake and how they are
5604 addressed in more mature protocols. Next, leader election mechanisms for provably secure
5605 instantiations of proof of stake are introduced. The leader election process in proof of stake
5606 is fundamentally different from that of proof of work, leading to security issues that result

5607 from being able to know in advance when a block producer will be elected. Finally, the
5608 block reward mechanism is investigated with a focus on how different reward schemes can
5609 lead to or avoid wealth concentration and, thus, centralization.

5610 There are two broad styles of proof-of-stake consensus algorithm. Chain-based schemes
5611 are modeled after Nakamoto Consensus with a major difference: using proof of work, a
5612 block is constructed before it is determined whether the miner has permission to broadcast
5613 it, but in proof of stake, permission to broadcast a block is provided prior to the block
5614 actually being constructed [298]. BFT-based schemes effectively operate as permissionless
5615 generalizations of traditional, permissioned BFT consensus. With some minor adjustments,
5616 many of the permissioned BFT algorithms covered earlier in this document can be made to
5617 work in a proof-of-stake context. Chain-based schemes have greater fault tolerance since
5618 they are secure so long as the majority of the stake is honest. In contrast, BFT-based proof
5619 of stake inherits the security threshold of the underlying BFT scheme, which is typically
5620 $\frac{1}{3}$. Chain-based schemes maintain availability during network partitions, while BFT-based
5621 ones sacrifice availability for consistency. Unlike chain-based schemes, BFT-based proof
5622 of stake can result in low latency while finalizing transactions. The specific properties, of
5623 course, depend on the design of the proof-of-stake algorithm.

5624 12.1. Early Attempts at Proof of Stake

5625 The first network to use proof of stake was Peercoin, proposed in 2012 [299]. The network
5626 retained an element of proof of work in order to distribute the coins fairly, but the fork-
5627 choice rule used for consensus replaces total work accrued with total *coin age destroyed*.
5628 Coin age is defined as the number of units of currency times the length of time that the units
5629 have been held without use. Peercoin blocks include a special transaction where the block
5630 producer consumes their coin age by sending funds to themselves. The first input from this
5631 transaction is called the *kernel*, and the kernel is hashed and checked against a target value
5632 similarly to how this is performed in proof of work. The target is calculated per unit of coin
5633 age spent in the kernel such that consuming more coin age proportionally increases the
5634 chance of being able to produce a block. As with Nakamoto Consensus, this process can
5635 result in multiple blocks being produced at the same time, but instead of choosing which
5636 fork to stay on based on the greatest total work, Peercoin validators prefer the fork with the
5637 highest coin age destroyed, where every transaction in a block contributes to the consumed
5638 coin age.

5639 This system includes protection against some very basic attacks. For instance, to prevent
5640 users from moving their stake from one output to another in order to increase their chance
5641 of producing a block, the coin age computation requires a minimum age of one month and
5642 is considered zero below this. Centralized, developer-signed checkpoints are then periodi-
5643 cally issued in order to ensure that every node agrees on transactions older than one month,
5644 which is necessary for verifying the kernel. A winning block producer could potentially
5645 use their single proof of stake to create many valid blocks for use in a denial-of-service

5646 attack. To prevent this, nodes collect all kernels and associated timestamps they have seen
5647 and ignore any blocks with the same (kernel, timestamp) tuple as a previously received
5648 block.

5649 After Peercoin, a second generation of proof-of-stake networks launched that entirely es-
5650 chew proof of work, including BlackCoin and NXT. BlackCoin’s consensus, PoS v2, is
5651 similar to Peercoin but removes coin age and makes several other small changes [300]. The
5652 use of coin age has the unfortunate side-effect of encouraging users to stay offline most of
5653 the time and only open their wallets to stake once every month or so. Further, when a user
5654 has a large amount of coin age built up, they can produce new blocks almost immediately
5655 and more easily execute double-spend attacks. Another observation was that of predictabil-
5656 ity – the input into the hashing algorithm for leader election does not prevent an adversary
5657 from precomputing future proofs of stake, which allows them to know in advance when
5658 they might produce several blocks in a row which facilitates double-spend attacks (this is
5659 discussed in detail in Section 12.2). PoS v2 mitigates this by including a *stake modifier*
5660 in the hash calculation that changes relatively frequently and allows for shorter precompu-
5661 tation windows. Finally, BlackCoin considerably restricts the timestamp rules for blocks,
5662 such that a node only accepts blocks with timestamps less than 15 seconds ahead of local
5663 time, while the timestamp granularity is changed from one second to 16 second intervals.
5664 This gives potential block producers less freedom to produce more blocks by trying to hash
5665 kernels with different timestamps.

5666 The previous point is important. If a staker has a large degree of freedom in selecting the
5667 input to the hash function used in leader election, they can perform a *stake grinding* attack,
5668 where they iterate through these potential inputs until they find one that is especially advan-
5669 tageous to them. An attacker can go through the history of the blockchain and – wherever
5670 their stake was selected to be the next block producer – modify the next block header or
5671 kernel over and over until it finds one that helps elect them as the next block’s producer as
5672 well. By changing the granularity of timestamps from one second to 16 seconds, Black-
5673 Coin reduced the number of grinding attempts by a factor of 16. Later, a set of incremental
5674 improvements was made to PoS v2, leading to the creation of PoS v3 [301]. The staking
5675 process for PoS v3 works as follows and loops forever:

- 5676 1. Nodes check their system clocks and set the (potential) block timestamp to the system
5677 time modulo 16.
- 5678 2. The network difficulty is computed, and the local difficulty target is computed by
5679 multiplying network difficulty by the number of coins held in a particular UTXO.
- 5680 3. The node iterates through each UTXO in its wallet. For each UTXO, the node com-
5681 putes a SHA-256 hash of several pieces of data, including the previous block’s stake
5682 modifier, some data from the UTXO, and the timestamp.
- 5683 4. These hashes are compared to the local difficulty targets for each UTXO, and if the
5684 value of a hash is less than the target, the node can create a new block. The block

5685 hash is then signed by the public key specified in the special staking transaction.

5686 5. If no hash satisfies the target, then the node waits 16 seconds and tries again with the
5687 new timestamp.

5688 Unlike BlackCoin and Peercoin, NXT uses an account model instead of UTXOs [302]. The
5689 entire NXT supply was distributed up front in a presale and then encoded into the genesis
5690 block. That is, there is no inflationary block subsidy, and block producer rewards consist
5691 solely of transaction fees. The NXT leader election protocol is called "forging" (instead of
5692 mining) and is described below.

5693 As with Peercoin and BlackCoin, potential block producers hash some input in the hopes
5694 that the hash is less than a particular target value, which changes from block to block in a
5695 manner analogous to proof-of-work difficulty adjustment algorithms (Section 9.2) in order
5696 to maintain 60-second block intervals on average. This target is computed individually for
5697 each account and determined in part by a *base target* value common to all users. Let S be
5698 the average block interval for the last three blocks, T_p the base target for the previous block,
5699 and T_b the base target being calculated for the current block. Then,

$$T_b = \begin{cases} \frac{T_p * \min(S, 67)}{60} & S > 60 \\ T_p - \frac{T_p * 0.64 * (60 - \max(S, 53))}{60} & S \leq 60 \end{cases}$$

5700 Here, the constants 67 and 53 bound the size of target adjustments, while 0.64 is included
5701 to allow block intervals to be shortened more rapidly than they can increase, which helps
5702 reduce the incidence of extremely lengthy block intervals. Each account can then compute
5703 its individual target value, T , as $T = T_b * S_p * B_e$, where S_p is the number of seconds that
5704 have passed since the previous block and B_e is the effective balance of the account. Only
5705 coins that have been in the account for at least 1440 blocks (one day) count toward the
5706 effective balance for staking.

5707 Each block includes a *generation signature*, which is a 32-byte hash output. Stakers will
5708 concatenate the previous block's generation signature with their public key and hash it
5709 to form the next block's potential generation signature. The first eight bytes of this hash
5710 are interpreted as an integer called the *hit*. As with proof of work, if the hit is less than
5711 the target, the account holder is authorized to produce a block. Unlike with proof of work,
5712 generation signatures make it possible to predict which account will produce the next block
5713 with fairly high accuracy. The target value increases with every second passed since the
5714 previous block, so even if most accounts are offline, a block will eventually be produced. If
5715 multiple blocks are produced around the same time, nodes prefer the fork with the highest
5716 cumulative difficulty (difficulty being the inverse of the target value). The current block
5717 difficulty, D_{cb} , is calculated as $D_{cb} = D_{pb} + \frac{2^{64}}{T_b}$, where D_{pb} is the previous block's difficulty.
5718 In addition, NXT eschews centralized checkpoints and instead just maintains a local rule
5719 that nodes will never reorg more than 720 blocks.

5720 Another early proof-of-stake variant, *delegated proof of stake* (DPoS), continues to be
5721 widely used despite not being formally studied to the degree that many other protocols
5722 have been. Technically, a variety of proof-of-stake protocols, including some chain-based
5723 ones, allow nodes to delegate their stake to another party, which allows the delegator to
5724 contribute to the network's security even while they are offline. Here, DPoS is used to refer
5725 to a specific type of protocol where only elected delegates participate directly in consensus
5726 as block producers, but all users have the ability to delegate their stake in order to elect a
5727 committee of delegates. The first DPoS system, BitShares, went live in July 2014 [303].
5728 Since then, it has been used in a variety of projects, including EOS, Lisk, Steemit, and
5729 Tron, as well as some sidechains (sidechains are discussed in Section 16.3).

5730 The mechanics of DPoS are relatively simple: users with stake in the system vote for a
5731 committee of n delegates who execute a permissioned BFT algorithm among themselves.
5732 Stake delegation usually involves signing a message from the public key that holds the
5733 stake, which assigns that user's stake to another user. In many systems, stakeholders can
5734 delegate their stake to several potential delegates at a time. Users are ordered by the amount
5735 of stake assigned to their keys. The top n become the delegates that make up the committee,
5736 and more act as backups who are ready to take over should votes change. This delegation
5737 occurs on a continuous basis, so the delegates themselves can be swapped in and out at any
5738 time based on changes to the stake delegation distribution. The actual consensus algorithm
5739 executed by the committee may differ, but one possibility is that leaders rotate in round-
5740 robin fashion, where a block is finalized after $\frac{2n}{3}$ delegates have voted for it by building
5741 blocks on top of it. Validators follow the longest chain but will not reorg past a finalized
5742 block.

5743 By having only a small number of delegates participate directly in consensus, DPoS can
5744 maintain high transaction throughput because the few delegates can run expensive, high-
5745 performance servers. In addition, it has relatively low latency due to finalization, unlike
5746 most proof of work or chain-based proof-of-stake systems. These performance benefits
5747 come at the cost of the increased security risks related to centralization. Delegates are
5748 likely to be easier to locate and conduct denial-of-service attacks against because they are
5749 likely to be recognized, fairly static entities who operate nodes in data centers, possibly
5750 with public IP addresses, and the time slots in which they are supposed to produce blocks
5751 are publicly known in advance. This also makes it easier for block producers to collude and
5752 perform censorship, and if throughput is high and validation is expensive, it can be hard to
5753 vote them out and find replacements.

5754 In DPoS systems, there is a distinct risk of having factions or cartels form, resulting in a
5755 form of oligopoly among delegates. In most DPoS models, delegates share block rewards
5756 with stakeholders who delegate to them. One implication of this is that stakeholders are
5757 likely to delegate only to those potential delegates who are highly likely to be elected, which
5758 implies that delegates who have already been elected to the committee are more likely to
5759 be voted for and remain in place. To maintain their position in power, these delegates may
5760 organize into factions that vote for each other and then impose a requirement that stake-

5761 holders must delegate to every member of the cartel in order to receive their "kickback"
5762 payment. This is exacerbated by low voter turnout combined with highly unequal distri-
5763 butions of stake, resulting in a handful of big players dominating. This occurred in Lisk,
5764 where two large factions controlled a combined 85% of the 101 delegates, with one alone
5765 controlling more than half of the total delegates [304]. Similarly, Steemit's top delegates
5766 are rarely replaced, and the support of only a few large stakeholders is sufficient to keep
5767 them in place [305].

5768 Social choice theory suggests that there may be some risks to the use of delegation in proof-
5769 of-stake systems. If some potential delegates are malicious but honest stakeholders are un-
5770 able to tell whether a potential delegate is malicious or not, it is possible for malicious ones
5771 to gain disproportionate power in some circumstances [306]. In addition, delegation favors
5772 a minority view, which could amplify the power of a malicious minority of stake [307].
5773 On the other hand, there may exist delegation schemes that avoid the over-representation
5774 of minority views on a committee and make it as expensive as possible to elect a certain
5775 threshold of delegates (say, $\frac{1}{3}$ of the committee) [308]. More research is needed to gain a
5776 better understanding of how delegation impacts the security of proof of stake.

5777 12.1.1. Nothing-at-Stake and Costless Simulation

5778 The most obvious and fundamental difference between proof of work and proof of stake
5779 with respect to their ability to act as Sybil-resistance mechanisms is the *nothing-at-stake*
5780 problem, sometimes called *costless simulation*. In proof of work, a significant amount of
5781 computation is required in order to determine who has the right to produce a block, whereas
5782 it takes only a handful of hashing operations to elect a leader under proof of stake. As a
5783 result, it costs (practically) nothing for a block producer to append a new block simulta-
5784 neously to as many chains as possible. This costlessness, combined with the reward for
5785 creating blocks, implies that block producers may update the ledger at any opportunity,
5786 even if the update would perpetuate disagreement. That is, if a proof-of-stake block pro-
5787 ducer is presented with multiple competing chains, they do not need to commit to one of
5788 them as miners must when using proof of work but rather can produce blocks on both
5789 chains for free, giving them the best chance of having a block they produced end up in
5790 the canonical chain. As a result, validators may struggle to converge on a single chain as
5791 everyone builds on every chain they see.

5792 Costless simulation does not present an insurmountable problem for a proof-of-stake sys-
5793 tem, and there are some who argue that it hardly presents a real problem for consensus at
5794 all [309]. In short, the argument states that equivocating block producers undermine the
5795 utility of the system, which should lead to a loss in the exchange value of their stake, and
5796 thus there are indeed real costs to producing blocks on multiple chains simultaneously. In
5797 order to make sure this condition is sufficient to induce an incentive to maintain consensus,
5798 [309] recommends that:

- 5799 1. There should be a minimum stake requirement in order to participate as a block

5800 producer. A small stakeholder only undermines their own wealth to a very small
5801 degree by trying to delay consensus through taking advantage of costless simulation,
5802 whereas larger stakeholders undermine their own wealth to a larger degree.

5803 2. Block rewards should be sufficiently low. Under proof of stake, stakeholder incen-
5804 tives come from both the value of the initial coin holdings as well as the block reward.
5805 By keeping the block reward component small relative to the initial stake, maintain-
5806 ing coin value by honestly participating in consensus becomes more important than
5807 short-term rewards from block inclusion. As the block reward decreases, the mini-
5808 mum stake requirements can become looser.

5809 Consider a scenario in which block producers receive no explicit reward for publishing
5810 blocks. If the value of the underlying stake is lowered by delaying consensus, then taking
5811 advantage of costless simulation indeed imposes a cost with no offsetting reward. If other
5812 stakeholders are honestly following the longest chain rule, then a malicious costless simu-
5813 lator undermines their own wealth both by refusing to add a block to the longest chain and
5814 by adding blocks to a shorter chain (and not being compensated for doing so). As a result,
5815 following the longest chain rule is an equilibrium for stakeholders when there is no block
5816 reward. With a block reward, this equilibrium continues to hold so long as the minimum
5817 stake is high enough [309].

5818 In practice, the most common solution to nothing-at-stake instability is to require potential
5819 stakers to submit a bond that may be *slashed* (i.e., seized and destroyed by the protocol) if
5820 malicious behavior is detected. For example, in Ethereum 2.0 (Section 13.2), participating
5821 in consensus requires depositing 32 ETH into a smart contract that is capable of verifying
5822 that a particular stakeholder signed equivocating blocks and then seizing the stake from
5823 them. Another possible solution is the idea of "virtual ASICs," which mimics proof of work
5824 in a proof-of-stake context – stakeholders purchase virtual mining machines and power
5825 them with virtual electricity, inducing an operational cost for block production [310].

5826 The early schemes described in the previous section did not take steps to address the
5827 nothing-at-stake problem, with the possible exception of NXT. NXT has no block sub-
5828 sidy, so as long as transaction fees are sufficiently small, consensus should be stable based
5829 on the argument presented in [309]. None of the other protocols imposed a minimum stake
5830 requirement or eliminated block rewards, so they are likely susceptible to costless simula-
5831 tion attacks. Delegated proof of stake is especially susceptible to costless simulation since
5832 the value of the stake used to elect delegates is not fully owned by the delegates themselves.

5833 For example, these protocols are more likely to fall victim to bribery attacks than their
5834 proof-of-work equivalents. Consider a merchant who waits for six confirmations before
5835 providing a good to a customer. After six confirmations, a malicious customer could pub-
5836 licly announce the intention to create a fork that would revert those six blocks and offer
5837 a bribe to any stakeholders who would sign blocks on the attacker's branch [311]. Unlike
5838 with proof of work, stakeholders who collude with the attacker in this way risk nothing if

5839 the attack fails. Proof-of-work miners have real-world costs imposed on them while mining
5840 on the attacker's fork, whereas there is nothing at stake for stakeholders (except for causing
5841 consensus delay, as described above).

5842 Proof of work also acts as a rate limiter of how quickly consensus-related messages can be
5843 generated in a way that proof of stake cannot. A bug shared by at least five proof-of-stake
5844 cryptocurrencies using PoS v3 resulted in improper validation of stake and allowed nodes
5845 without any stake at all to perform resource exhaustion attacks against any victim node,
5846 completely filling RAM or disk space [312].

5847 A concept strongly related to costless simulation is that of *stake shift*. Ideally, leader elec-
5848 tion would be based on the most up-to-date stake distribution possible in order to provide
5849 the maximum amount of Sybil resistance. Unfortunately, this distribution cannot be as
5850 recent as desired for two main reasons. First, in chain-based systems such as the ones de-
5851 scribed above, there is no agreement on the most recent few blocks, so they cannot be used
5852 as part of the stake distribution. More generally, even in BFT-based systems where con-
5853 sensus is achieved on a single block before proceeding to the next, the most recent stake
5854 distribution cannot be used because the distribution must be fully determined before an
5855 adversary can acquire any information on the randomness used to sample from the stake
5856 distribution. As a result, for all known proof-of-stake systems, there is a gap between the
5857 current stake distribution that establishes the actual incentives for participants and the stake
5858 distribution that is in fact used for leader election. This is problematic because a stakeholder
5859 can divest themselves of their stake between the time that their stake is used in leader elec-
5860 tion and the time they are entitled to produce a block. When it comes time to produce the
5861 block, they no longer have anything at stake to prevent bad behavior. This suggests that an
5862 additional security requirement for proof-of-stake ledgers is that stake cannot change hands
5863 "too quickly" [313, 314].

5864 The stake distribution in deployed proof-of-stake protocols typically has a lag of one to
5865 10 days, with chain-based protocols like Snow White and the Ouroboros family being
5866 on the longer side, while BFT-based protocols like Algorand have lags on the lower end
5867 of that range [315]. The degree of stake shift that occurs during this lag should count
5868 directly as part of the adversarial stake when evaluating the security margin of proof-of-
5869 stake consensus. More formally, if σ is the stake shift, α is the fraction of adversarial
5870 stake, and T is the normal security threshold of the system (i.e., $\frac{1}{2}$ for most chain-based
5871 systems and $\frac{1}{3}$ for most BFT-based ones), a proof-of-stake protocol maintains safety if
5872 $\alpha < (1 - \epsilon) * T - \sigma$ for some $\epsilon > 0$. Stake shifts for some more established cryptocurrencies
5873 can be seen in Figure 37 and appear to be a few percentage points for typical lags.

5874 12.1.2. Long-Range Attacks, Posterior Corruption, and Weak Subjectivity

5875 Most of the early proof-of-stake protocols described in Section 12.1 included a mechanism
5876 for establishing *checkpoints* in the ledger, such that a node will never reorganize the chain
5877 to overwrite the checkpointed block. This is due to the possibility of *long-range attacks*,

Lag (in days)	BTC			BCH			LTC			ZEC		
	Mean	Median	Std Dev	Mean	Median	Std Dev	Mean	Median	Std Dev	Mean	Median	Std Dev
1	0.013	0.010	0.0098	0.013	0.011	0.0102	0.014	0.011	0.0123	0.014	0.012	0.0102
2	0.020	0.017	0.0129	0.020	0.017	0.0134	0.022	0.017	0.0177	0.023	0.020	0.0146
3	0.026	0.022	0.0155	0.026	0.023	0.0161	0.030	0.023	0.0219	0.031	0.027	0.0181
4	0.031	0.027	0.0177	0.032	0.027	0.0183	0.036	0.029	0.0255	0.038	0.034	0.0211
5	0.036	0.031	0.0196	0.037	0.032	0.0203	0.042	0.034	0.0289	0.045	0.040	0.0238
6	0.040	0.035	0.0213	0.041	0.036	0.0221	0.048	0.039	0.0319	0.051	0.047	0.0262
7	0.045	0.039	0.0229	0.045	0.039	0.0238	0.053	0.044	0.0347	0.058	0.053	0.0286
8	0.049	0.042	0.0244	0.050	0.043	0.0253	0.058	0.048	0.0374	0.063	0.059	0.0308
9	0.053	0.045	0.0257	0.053	0.046	0.0267	0.063	0.052	0.0399	0.069	0.065	0.0328
10	0.056	0.049	0.0270	0.057	0.050	0.0281	0.068	0.057	0.0423	0.074	0.070	0.0346
11	0.060	0.052	0.0282	0.061	0.053	0.0293	0.073	0.060	0.0446	0.079	0.075	0.0364
12	0.063	0.055	0.0294	0.064	0.056	0.0305	0.077	0.064	0.0469	0.084	0.081	0.0380
13	0.067	0.058	0.0305	0.068	0.059	0.0317	0.082	0.068	0.0490	0.089	0.085	0.0395
14	0.070	0.061	0.0316	0.071	0.062	0.0329	0.086	0.072	0.0510	0.094	0.090	0.0410

Fig. 37. Empirical measurements of stake shift for some high market capitalization cryptocurrencies. While these assets do not use proof of stake, there is little reason to believe that fund movements would differ based on the Sybil-resistance mechanism. Stake shift appears to be smaller on longer-running, more established networks. [315]

5878 where an adversary creates a fork very deep in the chain in an attempt to overwrite the
 5879 canonical chain. Because producing blocks in a proof-of-stake system is free, a malicious
 5880 stakeholder suffers no additional costs for attempting to fork the chain from further back in
 5881 time (even all the way back to the genesis block), as opposed to forking near the chain tip.
 5882 Checkpoints limit how far back such a reorg can go.

5883 Consider a proof-of-stake system without checkpoints where participating nodes do not
 5884 have synchronized system clocks. A malicious stakeholder could create a competing blockchain
 5885 that forks the chain far into the past. If other stakeholders are honest, however, they will
 5886 not build on the malicious chain unless it is also the longest one, which is unlikely if the
 5887 attacker has a minority of the stake. However, without synchronized clocks, the adversary
 5888 can continue to construct blocks that appear to be from the future, ultimately creating the
 5889 longest chain and succeeding in the attack regardless of the adversarial fraction of stake.
 5890 As a result, proof-of-stake protocols require synchronized clocks, so most schemes have
 5891 an external dependency on a clock synchronization protocol, such as the Network Time
 5892 Protocol (NTP) [316].

5893 Assuming that clocks are synchronized but checkpoints are not in use, such a long-range
 5894 attack is extremely unlikely to succeed without a majority of the stake backing the attack
 5895 chain. It is generally assumed that acquiring the majority of the network's stake at any point
 5896 in time would be prohibitively expensive. However, due to the possibility of *posterior cor-*
 5897 *ruption*, this may not be the case. Former stakeholders who have since traded away their
 5898 stake can collude to extend the ledger from any point at which they did have the majority
 5899 of stake in the past. This can be rational because it is costless to build an alternate chain
 5900 and is not detrimental to them because they do not have current stake in the success of the

5901 network. Thus, old private keys have some positive value and may be kept after spend-
5902 ing the associated stake and then sold to malicious parties. In case of majority posterior
5903 corruption, the honest chain can be overwritten from a point deep in the blockchain.

5904 Another variant of long-range attack is *stake bleeding* [317], which exploits the ability
5905 of transactions that are valid on the canonical chain to be equally valid on a competing
5906 chain. This allows a long-range attacker to use the history of transactions from the canon-
5907 ical ledger on their own attack chain, collect their transaction fees and block rewards, and
5908 increase their amount of adversarial stake on the attack chain. As a result, a minority stake
5909 attacker can become a majority attacker, particularly if the network has been running for a
5910 long time. A fix that prevents stake bleeding – though one with high overhead – is to make
5911 transactions context-aware by requiring that they include a recent block hash and are only
5912 valid in chains where that block hash exists before the transaction is included.

5913 The most common and simplest fix to implement for each of these long-range attacks is
5914 the aforementioned checkpointing. Unfortunately, checkpointing fundamentally changes
5915 the system model to one that is *weakly subjective* rather than objective. When a system
5916 is objective, a new node (or one that has been offline for an extended period of time)
5917 that properly implements the protocol and receives the full set of blocks or other relevant
5918 consensus messages will fully agree with the rest of the network regarding the current state.
5919 However, a subjective system can have stable states in which nodes do not agree, thus
5920 necessitating a social context (e.g., reputation) in order to work. In a weakly subjective
5921 system, the new node can come to agree with the rest of the network so long as it has a
5922 properly implemented protocol, the complete set of consensus-related messages and blocks,
5923 and – crucially – a sufficiently recent (say, N blocks) state that is known to be valid; that is,
5924 a checkpoint.

5925 There are two ways in which a node can acquire a sufficiently recent known-valid state:
5926 1) being online frequently (at a minimum, at less than N block intervals) and witnessing
5927 the checkpoint block as an active participant of the network or 2) by getting the checkpoint
5928 from a trusted party if the node has been offline for more than N blocks (or is first joining
5929 the network). As a result, the security model for new nodes and those that have been offline
5930 for a while is fundamentally different because they are required to get a checkpoint from a
5931 trusted party. Without a trusted checkpoint, these nodes would be unable to differentiate the
5932 canonical chain from alternative valid chains. This must be contrasted with (most) proof-
5933 of-work schemes that can be objectively validated without introducing a trusted component
5934 because total work can be evaluated without having been online.

5935 When a new node attempts to connect to a proof-of-stake blockchain for the first time (or
5936 the first time in more than N blocks), it is that user's responsibility to verify a recent state
5937 out-of-band by checking a block explorer website, asking businesses they would like to
5938 interact with, or asking a friend for a recent block hash. In fairness to proof of stake, there
5939 is always a certain degree of subjectivity in terms of downloading a software client from a
5940 trusted source because the software must properly implement the protocol for the network

5941 that the user wants to connect to. That is, a user who downloads software that falsely
5942 claims to implement the Bitcoin protocol can accept a non-canonical, attacker-controlled
5943 chain. That said, code signing can help by allowing a user to verify that they are running
5944 the desired software for a given network.

5945 In addition to checkpoints, a number of long-range attack mitigations are possible [318].
5946 As mentioned earlier, context-aware transactions can prevent stake bleeding but do not fix
5947 posterior corruption. Another mitigation is to use *key-evolving signatures*. Key-evolving
5948 signatures allow signatures to be valid only for short time periods while also allowing the
5949 private key to evolve as these periods advance despite maintaining the same public key.
5950 In the erasure model, where honest nodes are assumed to securely delete their old private
5951 keys from local systems, key-evolving signatures can address posterior corruption but not
5952 stake bleeding. These kinds of signatures were suggested for use in Ouroboros Praos [319].
5953 Long-range attacks can also be addressed by a change in the fork-choice rule, as is done
5954 in Ouroboros Genesis. Roughly speaking, stakeholders using the new rule will prefer the
5955 chain that has a greater density of blocks for some period of time after the block where the
5956 chains diverge. As discussed in Section 13.1.3, this is not entirely without trade-offs. It
5957 may also be possible to protect against long-range attacks using verifiable delay functions,
5958 which can be used to prevent an adversary from producing blocks from a time when they
5959 were not online but with a significantly reduced security margin [320].

5960 Finally, note that security deposits are not a solution to long-range attacks. While they can
5961 address short-term nothing-at-stake issues, users need to be able to eventually withdraw
5962 their stake for other uses. Once the security deposit is withdrawn, the ability to hold the
5963 stakeholder accountable is eliminated, and they can freely engage in long-range attacks
5964 without penalty.

5965 12.1.3. Leader Election, Anonymity, and Security Against Adaptive Adver- 5966 saries

5967 In the early proof-of-stake protocols described above, two different types of leader election
5968 are performed. Most of them use a Bitcoin-esque process of hashing some data, called
5969 the kernel, and checking to see if the hash is beneath an agreed-upon target. The kernel
5970 structure itself should be designed to reduce or eliminate opportunities for stake grinding.
5971 With this style of leader election, offline users do not contribute stake to the security of the
5972 system, so the majority of *online* stake must be honest for security to hold. DPoS leader
5973 election simply establishes a list of delegates based on stake delegated to them and assigns
5974 leaders in round-robin fashion. This is deterministic and results in leaders being known in
5975 advance with high probability.

5976 These leader election processes leave much to be desired. Luckily, as proof-of-stake pro-
5977 tocols are studied further, the security properties of the leader election subprotocol have
5978 become better understood. This section provides a brief overview of some possible im-
5979 provements in proof-of-stake leader election but is not intended to be comprehensive. More

5980 details on a variety of these and other schemes are provided in Section 13.

5981 The first major alternative is a method called *follow the satoshi*, which is used in the Chains
5982 of Activity protocol described in Section 13.1.1. In Bitcoin, the smallest possible atomic
5983 unit of cryptocurrency is called a *satoshi*. Given a random index into the set of all exist-
5984 ing satoasis, one can find the block in which that particular satoshi was minted and then
5985 trace the transaction flow from that block to the present to find the public key of whoever
5986 currently owns that satoshi. To generate the random seed, each block producer includes
5987 a uniformly random bit in their blocks, and at the end of an epoch, the seed is a concate-
5988 nation of these bits. The seed is used for the follow-the-satoshi elections two epochs later
5989 (that is, after skipping an epoch). This procedure results in leaders who are publicly known
5990 in advance and, thus, is insecure against adaptive adversaries. To address the problem of
5991 offline stake, when an output is selected by follow the satoshi three times without a block
5992 being produced, that output can no longer produce blocks until they have spent their coins.
5993 This helps maintain consistent block production even when a significant amount of stake
5994 may be offline.

5995 Snow White (Section 13.1.2) uses two separate phases to determine each epoch's leaders
5996 and the random seed. Each phase lasts for roughly κ blocks, with κ a security parameter.
5997 As with Chains of Activity, Snow White block producers include random data in their
5998 blocks, which are used in the second phase. In the first phase, the stake distribution used
5999 for leader election is determined as the distribution that existed 2κ blocks in the past. The
6000 random seed is extracted from the blocks in $chain[-2\kappa : -\kappa]$. As long as a single block is
6001 produced by an honest user during this κ -block epoch, the seed can be used securely (this
6002 assumption only holds if the majority of the stake is honest). Because the stake distribution
6003 is set in stone κ blocks prior to the seed being known, an adversary is unable to adaptively
6004 select their public key to perform stake grinding by sending funds to themselves at public
6005 keys that would give them an advantage. However, once the adversary knows the next seed
6006 used for leader election, they have a one epoch delay in which to try to adaptively corrupt
6007 the leaders for the epoch in which the seed will be used. The actual mechanics of leader
6008 election are similar to the kernel hash ones above, where a hash of the seed, public key, and
6009 timestamp are checked against a difficulty target.

6010 Ouroboros Praos was the first proof-of-stake consensus algorithm to achieve security against
6011 fully adaptive adversaries (Section 13.1.3), which it accomplished through the use of a ver-
6012 ifiable random function (VRF). The VRF provides the property that only the block produc-
6013 ers themselves are aware that they have been elected to produce a block in a given time
6014 slot, so other validators – and the adversary – do not know who the leader will be until they
6015 have received a signed block with a valid VRF proof from the leader. At that point, it is
6016 already too late to perform an adaptive corruption or denial-of-service attack against the
6017 leader. Although secure against adaptive adversaries, this method still results in stakehold-
6018 ers knowing locally and in advance when they will have the right to produce a block. This
6019 predictability has security ramifications that are discussed in Section 12.2.

6020 In addition to the leader election mechanisms mentioned above and in Section 13, a variety
6021 of schemes have been proposed in the literature. Of note is the idea of *single secret leader*
6022 *election* (SSLE), which has some advantages over VRFs [321, 322]. In particular, while
6023 VRFs have the benefit of hiding the identity of potential block producers until the blocks
6024 are announced, they can result in there being no leaders elected to a given time slot or
6025 multiple leaders being elected within the same time slot. In the latter case, multiple leaders
6026 can cause undesirable forks in the chain. An SSLE scheme has the following properties:

- 6027 • *Uniqueness*: Each election results in exactly one leader chosen.
- 6028 • *Fairness*: If there are N registered users in the system, then each user has a $\frac{1}{N}$ chance
6029 of being elected in a given time slot. Further, a single honest user participating in
6030 the protocol should suffice to prevent a set of malicious participants from biasing the
6031 results of the election.
- 6032 • *Unpredictability*: If the adversary does not control the elected leader, the adversary
6033 cannot learn which user was elected.

6034 Unfortunately, concrete protocols for SSLE tend to rely on very advanced cryptographic
6035 primitives, such as indistinguishability obfuscation, threshold fully homomorphic encryp-
6036 tion, and public key encryption with keyword search. Another, more practical one is secure
6037 under the decisional Diffie-Hellman assumption and uses random shuffles but requires re-
6038 laxing a security property. SSLE protocols are promising, but more work should be done
6039 to improve their performance and make them secure using less exotic cryptography.

6040 The leader election schemes discussed thus far provide neither privacy nor anonymity to
6041 the leader. That is, in each case, the leader must reveal their public key to the network in
6042 order to create a block. Similarly, the above schemes also provide some information on
6043 the quantity of stake owned by the block producer. It is desirable to have leader election
6044 protocols that can hide such information, and some work has been done toward this end.
6045 For example, in the Ouroboros family of protocols, the Ouroboros Cryptsinous proposal
6046 is designed so that leader election can occur even in a system where privacy-preserving
6047 techniques are used to hide the number of coins associated with each public key, although
6048 it does not provide anonymity for the chosen leader and their public key [323].

6049 Other proposals focus on providing anonymity for proof-of-stake leader election [324,
6050 325]. Note how in proof of work, the leader election process does not depend on the
6051 identity of the miner because they can include a fresh public key in every block they mine,
6052 and the proof is just a valid puzzle solution. Proof of stake, on the other hand, cannot sep-
6053 arate the identity of the selected leader from their proof of eligibility. That said, advanced
6054 techniques can be used to hide this identity, even if it is used as part of the proof.

6055 The approach in [324] has the ledger store commitments to the total number of coins as-
6056 sociated with a public key rather than the amount itself, and stakeholders use NIZK proofs
6057 to show that they were elected properly. By itself, this would still leak information based

6058 on the frequency with which a given public key is elected and thus provide clues as to
6059 the balance of an account. A new *anonymous VRF* (AVRF) primitive addresses this. In
6060 an AVRF, multiple verification keys exist for the same private key, and given two proofs
6061 for different inputs under different verification keys, an eavesdropper cannot tell whether
6062 the proofs were generated using the same private key. The proof statement for the private
6063 leader election takes the list of all accounts as an input, proves knowledge of a private key
6064 corresponding to a public key in the list, and proves that the stake in that particular ac-
6065 count won the election. Unfortunately, the approaches used in both [323] and [324] fail to
6066 account for information leaks at the network level and can have their anonymity attacked
6067 by an adversary who can control the network delay facing targeted parties. By influencing
6068 stakeholders' local views of the network, the adversary can infer stake quantities based on
6069 how their targets publish blocks and which views those blocks are consistent with [326].

6070 An alternative approach is described in [325], which proposes an anonymous variant of
6071 Algorand's leader election scheme (further detailed in Section 13.4.2). Each party P_i is
6072 identified by a public key pk_i . All roles that one could be a leader for are given a tag, tag ,
6073 that may include things like the protocol round, the specific step or role in the round, and the
6074 random seed. P_i checks whether they are eligible for tag by signing tag with their private
6075 key sk_i , creating signature σ_i . This signature is input into a hash function H , generating
6076 $y = H(\sigma_i)$. If this output is below a threshold T , P_i has been selected, and proves this with
6077 (y, σ_i) . Verifying this requires pk_i and, thus, is not anonymous.

6078 To anonymize this, [325] utilizes a trapdoor permutation, TRP . Each P_i has a public value
6079 associated with it for each possible tag, $V_i = H(i|tag)$, as well as a public key $TRP.pk_i$ for
6080 a trapdoor permutation f . Party P_i checks whether they are a leader for tag by using their
6081 private trapdoor key $TRP.sk_i$ to calculate $v_i = f_{TRP.sk_i}^{-1}(V_i)$ and then checks if v_i is less than
6082 the relevant threshold. If so, they compute a zk-SNARK that proves they know a preimage
6083 of one of the V_i that makes them eligible.

6084 To summarize this section, there are certain desirable properties that proof-of-stake leader
6085 election should have:

- 6086 • The selection function should resist stake grinding. That is, stakeholders should have
6087 minimal flexibility when it comes to computing their eligibility query. This often
6088 necessitates having restrictive rules on timestamps.
- 6089 • The stake distribution used for leader election should be fixed prior to generating
6090 a random seed to prevent the specific form of stake grinding that comes from an
6091 adversary adaptively choosing their public keys.
- 6092 • Ideally, the selection function should be privately evaluated, such that only the stake-
6093 holder knows that they are elected before they announce this to the network. This
6094 can be done by making the selection function depend on the stakeholder's secret key,
6095 most commonly using a VRF. This is necessary for preventing adaptive corruptions.

- 6096 • The selection function should be fair in that a stakeholder’s chance of being elected
6097 leader is proportional to their share of the total system stake. Part of this involves
6098 minimizing the bias that an adversary can force into the random seed. This bias can
6099 be sufficiently bounded by allowing block producers to include random data in their
6100 blocks and having enough blocks in an epoch such that – with high probability – at
6101 least one block producer is honest.
- 6102 • Selecting multiple leaders in a time slot can lead to forks, so it is better to have only
6103 a single leader per slot. VRFs can lead to multiple leaders, but SSLE protocols can
6104 solve this.
- 6105 • Protecting the anonymity of the block producer would be ideal, but most existing
6106 leader election procedures do not provide this property.

6107 12.2. Leader Predictability and Security

6108 Every proof-of-stake scheme discussed thus far suffers, to some extent, from the problem
6109 of predictability. When proof of work is used, a miner does not know that they are a block
6110 producer until the moment they have succeeded in generating a proof. In contrast, a proof-
6111 of-stake system necessarily results in stakeholders being aware of their right to produce a
6112 block a minimum of one block in advance. This advance knowledge of leadership slots
6113 provides important information to an adversary, which can help them more confidently
6114 engage in selfish mining/staking, double-spending, and bribery attacks [327, 328]. Before
6115 explaining these attacks, some additional definitions are useful. In the following, a block
6116 denoted $Pred(B)$ is the predecessor of block B , and $Pred^D(B)$ is the D -th predecessor block
6117 of B .

- 6118 • **D -locally predictable:** A unit of stake s has this property if, at block A , the owner of s
6119 knows that they will be able to produce a block B after D blocks. That is, after seeing
6120 A , the owner of s knows that they can produce a block B such that $Pred^D(B) = A$.
- 6121 • **D -globally predictable:** A unit of stake s has this property if, at block A , every
6122 stakeholder knows that the owner of s will be able to produce a block B after D
6123 blocks. That is, after seeing A , all participants know that the owner of s can produce
6124 a block B such that $Pred^D(B) = A$.
- 6125 • **D -recent:** This is a negation of local predictability. A unit of stake s has this property
6126 if, at block A , the owner of s does not know whether or not they will have the right
6127 to produce a block B such that $Pred^D(B) = A$. The validity of block B depends on
6128 some information contained in the most recent D predecessors of B .

6129 All proof-of-stake protocols, including BFT-based ones, are (at least) 1-locally predictable.
6130 This is necessarily the case because all information used as input into the leader election
6131 procedure must be agreed upon prior to the leader election. Additionally, for any D and
6132 any block A , a given unit of stake is either D -locally predictable or D -recent. Protocols

6133 with recency have security risks in that competing chains may use different random seeds
6134 for leader election, which can make it more challenging to detect certain malicious pro-
6135 tocol deviations. As a result, there are security trade-offs between protocols with recency
6136 compared to those with local predictability.

6137 As mentioned above, predictability facilitates more advanced versions of selfish mining,
6138 double-spending, and bribery. In protocols with global predictability, selfish mining and
6139 double-spending attacks can be conducted in a way that is guaranteed to succeed. With
6140 local predictability, these attacks are not entirely risk-free, but the adversary is still provided
6141 with a significant statistical edge due to knowledge of the attacker's future blocks. Luckily,
6142 global predictability is not necessary. The use of VRFs allow leader election to be privately
6143 evaluated and publicly verified. That said, globally predictable protocols do exist, so the
6144 security concerns relating to them cannot simply be dismissed.

6145 Predictable selfish mining and double-spending operate almost identically, so the following
6146 description of selfish staking applies to both (selfish mining was introduced in Section 9.4).
6147 Let block A be the current best chain tip on a protocol with D -global predictability and t be
6148 the current timestamp, and let $S()$ be a score function for a chain (e.g., $S(B)$ may be equal to
6149 the number of predecessor blocks in a longest chain protocol). Globally predictable selfish
6150 mining works as follows [327]:

- 6151 1. For all $k \in \{1, \dots, D\}$, let t'_k be the earliest time that the adversary is entitled to produce
6152 a block B such that $S(B) > S(A) + k$ and where the adversary is the block producer
6153 of all blocks between A and B .
- 6154 2. For all $k \in \{1, \dots, D\}$, let t_k^* be the earliest time that the rest of the network is able to
6155 produce a block B such that $S(B) > S(A) + k$ and where the adversary is the block
6156 producer for none of the blocks between A and B .
- 6157 3. At time t , if a k exists such that $t'_k < t_k^*$, the adversary should immediately stop pub-
6158 lishing blocks until t'_k , at which point they publish B and every block along the path
6159 from A to B from the first step. If multiple k exist where this holds, the strongest
6160 attack is to use the largest one.

6161 This works exactly like selfish mining in the proof-of-work case, except that all of the
6162 risks are eliminated for the attacker. An adversary only withholds blocks from the net-
6163 work when they know for sure that the attack will succeed. Predictable double-spending
6164 works the same way, except for the details of including the specific transactions on each
6165 side of the fork. With local predictability, the attacker cannot perform the second step but
6166 can choose to perform the attack when they know they will produce a disproportionately
6167 large number of blocks in a certain window to maximize their chance of success. These
6168 prediction attacks were studied as applied to the Tezos blockchain, where it was found that
6169 it is often profitable to create two-block reorgs for selfish mining and that opportunities for
6170 predictable double-spending could be frequent. An adversary with 36% of the stake could
6171 expect to perform a predictable 20-block reorg once per year, and a 40% stakeholder could

6172 expect an opportunity for a 20-block predictable double-spend on a daily basis [329, 330].

6173 Using a VRF for leader election, as Ouroboros Praos does, provides security against adap-
6174 tive adversaries. However, the adaptive adversary model used in this and other protocols
6175 fails to account for bribery attacks, such as the one depicted in Figure 38a [328]. Consider
6176 a proof-of-stake protocol with D -local predictability, where a merchant – and potential vic-
6177 tim – considers a transaction final when the block it is included in has κ confirmations.
6178 First, consider the typical case where $D > \kappa$. An adversary can announce to the world that
6179 they would like to bribe block producers to contribute to an adversarial chain, and block
6180 producers can use their VRF proofs to demonstrate their eligibility. If the adversary can get
6181 $\kappa + 1$ eligibility proofs, they send a transaction to the victim merchant that is included in
6182 the next block. The attacker allows the honest chain to grow by κ blocks until the merchant
6183 provides whatever the adversary purchased. With goods in hand, the adversary creates a
6184 double-spend transaction and uses the $\kappa + 1$ eligibility proofs to publish a longer chain.
6185 For this attack to work, it is only necessary that $\kappa + 1$ out of the next 2κ block producers
6186 participate, but each of those block producers can have an infinitesimally small amount of
6187 stake. In addition, the bribed block producers have plausible deniability, as they need not
6188 sign conflicting blocks.

6189 A naive solution to this problem would be for merchants to increase their required number
6190 of confirmations to something beyond the prediction window ($\kappa > D$). Note, however, that
6191 the security of these proof-of-stake protocols depends on generating random seeds with
6192 minimal bias. If the majority of block proposers in an epoch are bribed, the adversary can
6193 control the random seed to perform stake grinding and enlarge the prediction window into
6194 future epochs.

6195 While BFT-based proof of stake may appear to be largely immune to this kind of attack,
6196 this is not necessarily true. Consider Algorand, where Byzantine agreement must occur
6197 over every block, each step of the agreement protocol has a separate committee, and the
6198 committees are 1-locally predictable. An adversary who can bribe $\frac{2}{3}$ of the committee at
6199 any given step can sign a block that conflicts with the one that honest stakeholders agreed
6200 on in an earlier step, as shown in Figure 38b. Because Algorand has relatively small,
6201 constant-sized committees, this $\frac{2}{3}$ supermajority may represent only a small fraction of the
6202 total stake.

6203 Revisiting the leader election schemes discussed in Section 12.1.3, one can see that both
6204 Chains of Activity and Snow White have global predictability, while Ouroboros Praos and
6205 Algorand are locally predictable, though with very different predictability windows. Pro-
6206 tocols with (short) recency face a different security obstacle: the undetectable nothing-at-
6207 stake attack, shown in Figure 39.

6208 Let A be the highest scoring chain tip that an adversary is aware of in some D -recent
6209 protocol, and assume that the adversary has spread their coins across many public keys.
6210 To perform an undetectable nothing-at-stake attack, the adversary finds a block A' , where
6211 $S(A')$ is maximal over all blocks that are not descendants of $Pred^D(A)$ (i.e., A' is the "next

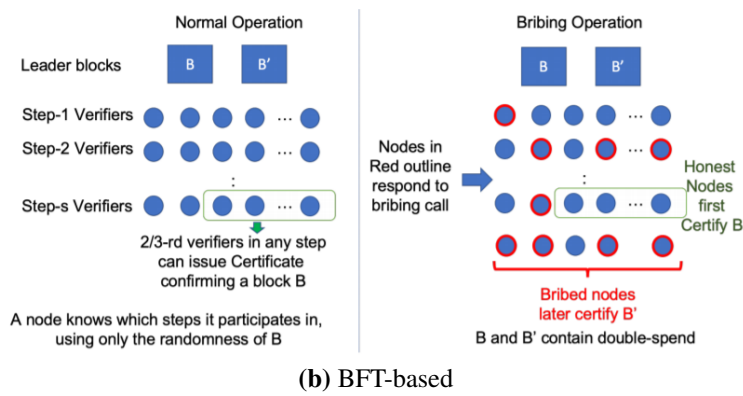
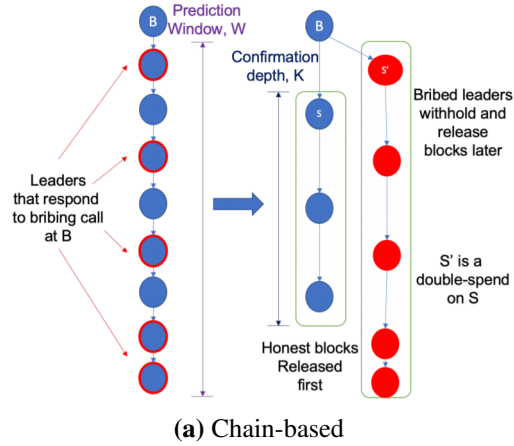


Fig. 38. Predictable bribe attacks against proof of stake. [328]




Fig. 39. Undetectable nothing-at-stake attack. The attacker creates both B and B' . [327]

6212 best" block after A). Next, the attacker simultaneously tries to produce blocks B and B' on
6213 top of A and A' , respectively. Should they successfully produce block B , they broadcast it.
6214 If they successfully produce B' , they broadcast B' as well, so long as doing so would not
6215 cause a provable deviation. The two types of provable deviation occur when the same unit
6216 of stake is used to:

- 6217 1. Sign two blocks, B and B' , that have the same timestamp or
- 6218 2. Sign two blocks, B and B' , where the timestamp of B' is greater than that of B
6219 ($t_{B'} > t_B$), but $S(B) > S(\text{Pred}(B'))$. This is a provable deviation because if $S(B) >$
6220 $S(\text{Pred}(B'))$, an honest stakeholder would not produce B' at time t_B .

6221 Because the protocol is D -recent, whether a valid B' can be produced on top of A' using a
6222 given unit of stake depends on blocks between A' and $\text{Pred}^D(A')$. As a result, each unit of
6223 stake provides another opportunity to produce a block on top of A' . Checkpointing protects
6224 against this.

6225 The existence of this nothing at stake attack concretely lowers the security bounds of proof-
6226 of-stake protocols with recency. For example, a protocol called Nakamoto-PoS has recency
6227 and is secure against adversaries with less than $\frac{1}{1+e}$ ($\approx 26.9\%$) of the total stake [328]. In
6228 Nakamoto-PoS, the random seed is updated with every block and consists of the hash of the
6229 predecessor block with leader election using a VRF that takes the timestamp and previous
6230 block hash as input. Given that the randomness depends on the block and lacks consensus,
6231 the adversary can take advantage of the nothing-at-stake phenomenon to lower the security
6232 threshold. At the same time, this minimizes predictability, such that stakeholders only
6233 know if they can produce a block upon seeing the previous one. Note that, for predictable
6234 proof-of-stake systems like Snow White and Ouroboros Praos where the same random seed
6235 is reused for an extended period, asymptotic consistency guarantees match that of Bitcoin
6236 and Nakamoto Consensus, despite the nothing-at-stake problem [331, 332].

6237 Interestingly, [328] showed how to adjust the Nakamoto-PoS protocol to create a family
6238 of chain-based proof-of-stake protocols that decouple the prediction window D from the
6239 security parameter, κ . This allows for dramatic reductions in the predictability window
6240 compared to algorithms like Ouroboros Praos and Snow White, where the epoch length
6241 (and thus prediction window) is proportional to κ . This family of protocols are defined by
6242 two parameters: c and s . The block hash used for randomness is updated every c blocks
6243 (so standard Nakamoto-PoS has $c = 1$), and only coins that are s blocks deep count toward
6244 the stake distribution used for leader election. Instead of following the longest chain rule,
6245 stakeholders use a fork-choice rule inspired by Ouroboros Genesis (Section 13.1.3). When
6246 comparing two chains, the stakeholder finds the block where the two chains diverge and
6247 prefers the chain where s blocks were produced after the fork most quickly based on their
6248 reported timestamps. If fewer than s blocks have been produced after the fork, then the
6249 longest chain rule is used. As c increases, the fraction of adversarial stake that the proto-
6250 col can tolerate increases toward 50%, though predictability increases as well. This family

6251 of protocols can achieve any level of predictability by setting c to the desired predictabil-
6252 ity window while simultaneously being secure for any confirmation depth κ by using an
6253 appropriate value for s .

6254 Recency is also problematic for proof-of-stake protocols that use the GHOST fork-choice
6255 rule (Section 11.2) because an adversary with little stake can dramatically increase the size
6256 of a particular subtree by trying to produce new blocks on top of every existing block in
6257 the subtree, magnifying the nothing-at-stake attack [327]. This subtree will grow expo-
6258 nentially faster than the number of blocks produced by honest stakeholders, forcing honest
6259 participants to accept the attacker's subtree.

6260 12.3. Wealth Concentration, Block Rewards, and Centralization

6261 In a proof-of-work cryptocurrency, miners face electricity costs and other operational ex-
6262 penditures, which often requires them to sell the cryptocurrency on the open market to fund
6263 their operations. This has the beneficial side-effect of distributing the coins to the public in
6264 a relatively "fair" way. However, when proof of stake is used, these costs no longer exist.
6265 As a result, there is a perception – real or imagined – that proof of stake faces the problem
6266 of wealth concentration, or "the rich getting richer." This section explores the consequences
6267 and likelihood of such a wealth concentration issue.

6268 In a proof-of-stake network, the resource used to provide Sybil resistance is the native as-
6269 set of the system. As a result, the block reward schedule (i.e., the monetary policy of the
6270 network) has potential implications for the security of the network beyond that of similar
6271 proof-of-work schemes. It is conceivable that under some reward schemes, large stakehold-
6272 ers are more likely to capture greater and greater shares of the total stake, whereas others
6273 do not encourage centralization in this way. Whether a particular reward scheme induces
6274 wealth concentration depends on the specific consensus algorithm. There may be differ-
6275 ences when delegation to stake pools is allowed, if offline stakeholders exist, if a minimum
6276 bond must be posted in order to stake, and so on. Different studies investigate different
6277 models, and this remains a relatively understudied area, so care must be taken when trying
6278 to draw conclusions regarding the problem of wealth concentration.

6279 In one study on this issue, Fanti et al. define the notion of *equitability*, which quantifies
6280 the degree to which a stakeholder's share of the total stake at some point in the future can
6281 change under a given reward function compared to their initial investment [333]. When
6282 a reward function is more equitable, the variability in this stake fraction should be small.
6283 The model used in [333] assumes that all stake is always online and available to produce a
6284 block when elected and that all rewards are reinvested into staking. When this is the case,
6285 they found that:

- 6286 • The rich do get richer because larger stakeholders have lower variance in their invest-
6287 ment returns than small ones, which allows their stake to compound. Allowing stake
6288 pools can lower this variance and, thus, the impact of compounding wealth.

- 6289 • A large initial stake should be allocated in order for rewards to be more equitable.
6290 That is, the block reward should be small relative to the initial stake.
- 6291 • The most equitable reward function is a geometric reward function where a constant
6292 fraction of the total stake is emitted at each block because rewards are smallest when
6293 the total system stake is smallest. Rewards then grow proportionally to the total
6294 system stake, thus bounding the wealth-compounding impact of each block. This
6295 implies an ever-increasing inflation rate. A constant reward function resulted in low
6296 equitability, and a decreasing reward function was even worse.
- 6297 • Strategic behavior, such as selfish staking, exacerbates equitability issues more under
6298 proof of stake than under proof of work due to this compounding effect (even before
6299 considering predictability issues, as discussed in the previous section). Even the geo-
6300 metric reward function failed to mitigate the impact of compounding in the presence
6301 of selfish stakers.

6302 Rosu and Saleh arrive at very different conclusions under a similar model but where the
6303 time horizon is infinite rather than finite [334]. They argue that while it is true that larger
6304 stakeholders are more likely to be elected and thus increase their share of the total stake,
6305 their share is decreased by an even larger amount when they are not elected. The result
6306 of these two competing forces is that a stakeholder's share is not expected to change over
6307 time from its initial value. While a geometric reward may minimize the variance of a
6308 stakeholder's share of the total stake over a finite period, it has high variance over large time
6309 periods and leads to wealth concentration in the limit. Contrary to [333], Rosu and Saleh
6310 predict that constant and decreasing reward functions induce a stable wealth concentration.
6311 Both analyses agree that a large initial stake total results in more stability for stakeholder
6312 shares.

6313 Irresberger also investigated the evolution of share distributions resulting from different
6314 reward functions [335]. In this model, not all coins are actively staked, and participants
6315 must choose to either stake their coins or have them available for spending. Under these
6316 circumstances, the geometric reward schedule (called "constant" in [335]) results in greater
6317 centralization of stake than a dynamic reward schedule that targets a specific participation
6318 rate in consensus. That is, the least centralizing reward scheme is one that picks a certain
6319 number of stakers and increases the rewards when there are fewer or decreases rewards
6320 when there are surplus stakers. However, the difference between the geometric rewards and
6321 dynamic rewards was small as long as the geometric inflation rate remained low (1-5%).
6322 When block rewards decrease over time, the number of staking nodes tends to decrease as
6323 well and lead to significant wealth centralization.

6324 The idea of having a dynamic reward schedule that adjusts in order to target a specific
6325 number of block producers has other advantages. In particular, because the capital used
6326 for staking is fairly liquid, the return on investment to staking must be competitive with
6327 alternative on-chain uses for that capital, such as collateralized lending protocols (i.e., de-

6328 centralized finance [DeFi]) [336, 337]. If yields from these alternatives become higher than
6329 the yield from staking, "bank runs" can occur, where a significant number of stakers "un-
6330 stake" their coins at around the same time, leading to a collapse in the security level of the
6331 system [336]. This is an especially significant concern in systems with decreasing block
6332 rewards or low inflation levels. A dynamically adjusting staking yield or a withdrawal de-
6333 lay for "unstaking" a participant's coins (as is done in Ethereum 2.0, discussed in Section
6334 13.2) can help prevent this collapse. Interestingly, [337] found that both wealth concentra-
6335 tion effects and the security risk from higher yield lending can be mitigated using *staking*
6336 *derivatives*. Consider a (fictional) asset, STK, that is used in a proof-of-stake system. A
6337 staking derivative is a synthetic asset, say, sSTK, that a staker can acquire while borrowing
6338 against their staked STK tokens. For the stakeholder to recover their STK and any block
6339 rewards they have earned from staking, they must buy back the STK with sSTK through an
6340 on-chain smart contract. Unfortunately, the positive results in terms of wealth distribution
6341 and security only hold when there is a non-negligible risk of having one's stake slashed,
6342 which is also undesirable.

6343 Thus far, reward schemes have only been discussed for chain-based proof-of-stake schemes,
6344 where only a single stakeholder ultimately produces a block in a given time slot. For BFT-
6345 based proof-of-stake schemes, where an entire committee of stakeholders are involved in
6346 producing each block, incentive design becomes even trickier, as some of the committee
6347 members may be faulty and undeserving of reward. It is possible to design a fair reward
6348 scheme for this kind of system, where fairness means that any honest stakeholder should
6349 receive a fraction of the total reward equal to their fraction of the total stake, but only if the
6350 network is synchronous [338]. Without synchrony, it is not possible to determine whether a
6351 committee member is Byzantine or experiencing network delays, so rewarding stakehold-
6352 ers based on merit is not possible. Weighted voting – where the past behavior of faulty
6353 processes that do not send messages is used to assign them less influence in the consensus
6354 process – can further improve fairness but also induces centralization [339].

6355 Another complication with some of this research on wealth concentration is the assumption
6356 that every stakeholder is actively staking and online at all times. There are several reasons
6357 why this is unlikely to be the case. First, without some form of stake delegation or stake
6358 pools, most individuals are not likely to keep an application up and running at all times due
6359 to sheer inconvenience. Second, being a block producer will usually require running a fully
6360 validating node, which may be performance intensive and contribute to the inconvenience
6361 of participation. Third, even with delegation or stake pools, it is likely that some (perhaps
6362 small) portion of users will fail to delegate or join a pool. In particular, it is likely to
6363 be small stakeholders who fail to overcome the friction involved because their expected
6364 absolute returns to staking are small. Finally, individuals who are less financially well off
6365 are likely to need or want greater liquidity from their assets, so these individuals may not
6366 be able to stake in systems where they must post a bond to become a block producer. Of
6367 course, only those who actually stake will get block rewards, so to the extent that some are
6368 excluded from staking, it is likely the case that wealth will tend to grow more concentrated

Network	Reward	Adj. Reward	Staked Value	Stake Ratio
Ethereum	4.79%	5.01%	\$35.8 B	15.62%
Cardano	3.23%	0.16%	\$9.5 B	68.04%
Solana	6.47%	-1.05%	\$8.3 B	72.83%
BNB Chain	2.66%	8.3%	\$7.5 B	15.35%
Avalanche	7.96%	2.32%	\$4.5 B	61.27%
Polygon	7.34%	3.89%	\$4.1 B	39.93%
Polkadot	14.34%	6.66%	\$4.0 B	47.03%
Tron	3.77%	1.69%	\$2.6 B	42.97%
Cosmos	23.1%	4.53%	\$2.6 B	66.97%
Internet Computer (Dfinity)	7.41%	-2.4%	\$1.8 B	73.2%
NEAR	9.24%	4.19%	\$1.0 B	45.34%
Tezos	5.32%	0.92%	\$774 M	72.41%

Table 1. Percentages of eligible tokens actively staked as of April 4, 2023. Adjusted reward is the annualized reward rate adjusted by the expected inflation of the network supply. [340]

6369 within these systems. Table 1 shows the percentage of eligible tokens that are actively
6370 staked in some of the proof-of-stake networks with the greatest total value staked as of
6371 April 4, 2023.

6372 This lack of universal staking suggests that stake pools may help reduce wealth concen-
6373 tration effects that could otherwise result from proof of stake (though care must be taken
6374 to avoid "address malleability" attacks if doing so [341]). Of course, this creates the need
6375 for reward schemes that properly incentivize the creation of a decentralized network of
6376 stake pools, which is a non-trivial problem. In fact, a "fair" scheme that rewards pools pro-
6377 portionally to their size/stake leads to an equilibrium with a single dominating stake pool
6378 [342]. Instead, [342] proposes a "cap-and-margin" reward scheme that has a Nash equilib-
6379 rium resulting in the desired number of pools forming if participants are rational. If k stake
6380 pools are desired, the reward function caps the rewards for a pool once it reaches a size
6381 greater than $\frac{1}{k}$ -th of the total stake in the system so that there is no longer an incentive for
6382 users to stake with a large pool. The reward function is also influenced by the amount of
6383 stake contributed by the leader of the pool, such that the pool leader will get higher returns
6384 by contributing more of their own stake to the pool. This is intended to discourage Sybil
6385 attacks, where a single entity with low stake is a leader for many different pools. There is a
6386 parameter that controls a trade-off between this Sybil-resistance property and the *egalitari-*
6387 *anism* of the system, such that more Sybil-resistant rewards imply a less egalitarian reward
6388 scheme.

6389 The egalitarianism of a system is a metric proposed in [343] based on the idea that a certain
6390 investment in capital to become a block producer should generate returns proportional to
6391 the capital invested. That is, in expectation, wealthier investors should not be able to gain
6392 disproportionate rewards compared to poorer investors. Egalitarianism is distinct from the

6393 equitability metric mentioned earlier in this section. Equitability captures the idea that over
6394 a series of blocks, a stakeholder's share of the total stake should not vary significantly from
6395 where it began. Egalitarianism, on the other hand, is about minimizing the variation in
6396 expected returns given an initial capital distribution. Stated differently, the randomness
6397 involved in equitability is over an execution of the protocol and how stake shares evolve,
6398 whereas the randomness involved in egalitarianism is over the distribution of wealth at the
6399 beginning of the protocol execution. Stake pools likely improve equitability by eliminating
6400 the penalty that comes from being offline but reduce the level of egalitarianism.

6401 Proof of stake is more egalitarian than ASIC-resistant proof of work, which is more egalitarian
6402 than proof of work using ASICs [343]. This reflects the greater economies of scale
6403 that are possible under proof of work. In this respect, proof of stake has an advantage over
6404 proof of work in the likelihood of centralization, though egalitarianism does not quite capture
6405 the "rich getting richer" phenomenon the way equitability does. One reason for this is
6406 that the wealthy can afford to stake a much greater percentage of their assets than individuals
6407 with fewer resources (who need more liquidity). Since the egalitarianism metric is the
6408 return on capital invested and not the ability to invest capital in the first place, it does not
6409 account for this. Ultimately, the two Sybil-resistance mechanisms are simply different, and
6410 there is not a long enough history of deployment of either mechanism to determine whether
6411 one or the other is "better" at avoiding centralization. A few points of comparison are in
6412 order:

- 6413 • In both proof of stake and proof of work, incentives matter for security and de-
6414 centralization. However, the specific details of the incentives under proof of stake
6415 have greater significance and are less well understood. For instance, preventing the
6416 nothing-at-stake issue is easier when block rewards are small, but the research pre-
6417 sented in this section is inconclusive regarding the inflation schedule.
- 6418 • Proof of stake is far more dependent on the initial stake distribution for security than
6419 proof of work is. Achieving a good distribution is challenging both for technical
6420 reasons as well as a lack of understanding of what a "good" distribution would be
6421 in the first place. One way to overcome this is to have a proof-of-work phase where
6422 the cryptocurrency is first distributed "fairly" before switching to proof of stake. In
6423 practice, it is common for the cryptocurrency to be auctioned off when the system is
6424 initiated, but this is an inherently centralized process.
- 6425 • It is usually assumed that the Sybil-resistance resource (work or stake) is acquired
6426 by honestly paying for it. In proof of stake, this should increase the exchange rate
6427 of the asset and the cost of the attack (this price increase may raise the cost of the
6428 attack by 20-40%, depending on several factors [344]). However, there are ways
6429 for well-positioned malicious parties to acquire the resource for free or cheap. In
6430 proof of work, the physical hardware used for mining can be stolen (or the owners
6431 of the hardware coerced). Depending on the situation, the attacker may still need
6432 to pay for operational costs, such as electricity. In proof of stake, the private keys

6433 used for staking can be hacked or stolen (including from vulnerable smart contracts),
6434 at which point there are no ongoing costs to use the stake. Identifying targets for
6435 such attacks in proof of work can involve monitoring electricity usage to detect large
6436 mining farms. In proof of stake, some large stakeholders can be easily identifiable:
6437 exchanges are likely to be large stakeholders and have the personal information of
6438 other large stakeholders. That said, it is likely easier to stake at significant scale
6439 without being detected than it is to mine undetected. On the other hand, the need for
6440 hardware in proof-of-work systems may add considerable friction for attackers that
6441 may result in the attacker being detected.

6442 Apart from these differences, the consequences of a concentration of wealth or consensus
6443 power in these networks can differ. For example, both Nakamoto Consensus and chain-
6444 based proof of stake can recover from adversaries that gain a temporary majority of the
6445 Sybil-resistance resource, whereas BFT-based proof of stake is unable to recover [228]. In
6446 chain-based proof of stake, the spike in adversarial stake must be shorter than an epoch
6447 length in order to recover. Here, "recover" means that the ledger properties of persistence
6448 and liveness return after some period of insecurity. Of course, with proof of stake, any
6449 temporary majority can become a permanent majority if the majority behaves in any way
6450 that does not result in their stake being slashed.

6451 Slashing is one of the biggest advantages of proof of stake compared to proof of work.
6452 Consider a situation where a powerful adversary manages to acquire a majority of the
6453 Sybil-resistance resource and has the capital to maintain this advantage for a while. For a
6454 proof-of-work system to survive, the proof-of-work algorithm must be changed in a hard
6455 fork. If the network used an ASIC-resistant algorithm when it forked, the attacker can use
6456 the same equipment to attack the hard-forked chain. If the algorithm was mined by ASICs,
6457 then a hard fork would resolve the attack by making all of the attacker's hardware useless,
6458 but this nuclear option destroys all of the honest miners' hardware as well. Proof of stake
6459 with slashing is better here because for the most serious kinds of misbehavior – such as
6460 a block producer equivocating and producing conflicting blocks at the same height – the
6461 protocol can specifically target the adversary's stake for slashing. This is like burning down
6462 an entire mining farm in proof of work but avoids collateral damage. Slashing requires extra
6463 care in situations where the network has a hard fork that causes a chain split because two
6464 blocks may then be signed at the same height on both sides of the fork and slashed.

6465 13. Proof-of-Stake Protocols

6466 13.1. Chain-Based Proof of Stake

6467 13.1.1. Chains of Activity

6468 An early chain-based proof-of-stake proposal is the Chains of Activity (CoA) algorithm
6469 that was designed for the UTXO model [311]. Unlike many of the proof-of-stake protocols
6470 described later in this document, the security of CoA is heuristic rather than proven. It

6471 borrows the *follow-the-satoshi* idea from the proof-of-activity hybrid algorithm described
6472 in Section 14.2 and introduced in Section 12.1.3. To "follow" a satoshi, one takes an index
6473 of the atomic unit of cryptocurrency, finds the block in which that unit was minted, and then
6474 traces the transaction graph from its minter to the public key of whoever currently owns that
6475 unit. Given some suitable procedure for securely generating that initial index, validators
6476 are able to agree on a leader. The CoA protocol utilizes the following parameters:

- 6477 • There are 2^κ satoxis minted in the system.
- 6478 • $w \geq 1$ is a subgroup length.
- 6479 • $L = \kappa * w$ is a group length .
- 6480 • $comb()$ is a function with L bit domain and κ bit range. It can be as simple as
6481 concatenating the inputs, but a variety of functions are possible.
- 6482 • G_0 is the minimum block interval time.
- 6483 • T_0 is the double-spending safety bound.
- 6484 • C_0 is the minimal stake amount.
- 6485 • C_1 is an award amount such that $0 \leq C_1 < C_0$.

6486 The protocol itself can be described by assuming that some leader with a publicly known
6487 identity has been elected via some mechanism, and they construct and sign block B_i . Each
6488 block B_i is associated with a deterministically generated, uniformly distributed bit b_i (e.g.,
6489 the least significant bit of a hash of B_i). For two blocks B_i and B_j , the time gap between
6490 them must be at least $|j - i - 1| * G_0$, and a newly created block is considered invalid if the
6491 timestamp is too far in the future compared to the local clock's time. It is possible that a
6492 leader elected for a given slot is inactive, in which case their slot would be empty and have
6493 no block associated with it.

6494 An epoch in CoA consists of L valid blocks B_{i_1}, \dots, B_{i_L} being created. After an epoch is
6495 completed, the network forms a κ -bit seed $S^{B_{i_L}} = comb(b_{i_1}, \dots, b_{i_L})$. The seed is then used
6496 to derive the identities of the block producers two epochs later (i.e., skipping L blocks)
6497 using follow-the-satoshi. Specifically, if the next L blocks are $B_{i_L+1}, B_{i_L+2}, \dots, B_{i_L+L}$, then
6498 the stakeholder elected to create block B_{i_L+L+z} is the one found by following the satoshi
6499 with index $H(i_L, z, S^{B_{i_L}})$ for $z \in \{1, \dots, L\}$. If a node is ever presented with multiple chains,
6500 the chain with the most blocks is canonical.

6501 CoA uses *slashing* to penalize bad behavior by forfeiting the stake of a malicious validator
6502 when misbehavior is detected. It may be the case that the satoshi used in leader election
6503 belongs to an unspent output with $c < C_0$ coins, or below the minimum. In this case, the
6504 leader must submit an additional signature proving ownership of the difference, $C_0 - c$
6505 coins, before creating their block, and at least C_0 coins must be available to be slashed. The
6506 relevant outputs may not be spent for another T_0 blocks after the newly created one. The

6507 stakeholder who creates a block B_i may equivocate by creating another block at the same
6508 height, B'_i , but any other leader elected in the next T_0 blocks can submit the conflicting
6509 signatures as a special transaction in their block to forfeit C_0 of the equivocator's coins,
6510 awarding the leader who caught them with C_1 of the forfeited coins.

6511 There is also a "three strikes" rule to exclude inactive participants. If an output $txout_0$
6512 was selected by follow-the-satoshi and has not created a block three times in a row, then
6513 $txout_0$ is no longer allowed to participate in the creation of new blocks (once the coins in
6514 $txout_0$ are spent, the new UTXO can be used for staking). Should follow-the-satoshi select
6515 $txout_0$, the slot is skipped and no block is produced. This helps CoA maintain liveness and
6516 performance even when users lose their coins or are offline for long periods.

6517 The leader election process in CoA leaves a few things to be desired from a security per-
6518 spective. First, as discussed in Section 12.2, having leader assignments known in advance
6519 is a security risk. An adversary who knows that they are the leader for several consecu-
6520 tive blocks may be able to perform predictable double-spends that succeed with certainty.
6521 Further, as discussed in Section 12.1.3, CoA results in leaders being *publicly* known in
6522 advance, so it cannot be secure against adaptive adversaries. This combination – but es-
6523 pecially the predictability – can facilitate bribery attacks, though the slashing mechanism
6524 partially mitigates this risk. For a large epoch length L , this predictability problem becomes
6525 worse. On the other hand, a small L would make it easier for a malicious coalition to in-
6526 fluence future leader elections. Finally, CoA requires a form of checkpointing to prevent
6527 long-range attacks.

6528 13.1.2. Snow White

6529 The first provably secure proof-of-stake system designs were Snow White [313, 314] and
6530 Ouroboros Praos (Section 13.1.3). Snow White is secure against *weakly adaptive* adver-
6531 saries, where corruption takes place with some delay, whereas Praos is secure against fully
6532 adaptive adversaries. Snow White follows the longest chain rule and is secure if there is
6533 an honest majority of *online* stake but requires longer confirmation times than Nakamoto
6534 Consensus (simulations show 34% to 43% more blocks are required). From a networking
6535 perspective, Snow White follows the "sleepy consensus" model, which is similar to syn-
6536 chrony but allows nodes to go offline and start back up without losing security guarantees
6537 [28].

6538 To address the issues of stake grinding and adaptive public key selection (to adversarially
6539 bias the random seed), Snow White uses a "two-lookback" mechanism, where each epoch's
6540 committee and random seed are generated in two separate phases with each phase lasting
6541 for roughly κ blocks, where κ is a security parameter. In the first phase, the prefix of
6542 the chain with the trailing 2κ blocks removed is used to determine the next consensus
6543 committee. In the second phase, randomness is extracted from the blocks in $chain[-2\kappa :$
6544 $-\kappa]$ to generate a random seed used for leader election in the current epoch. By forcing the
6545 committee to be selected κ blocks before the seed, an adversary cannot adaptively choose

6546 their public key based on the random seed to perform stake grinding.

6547 When an honest validator creates a block, they include some random data in it. As long as a
6548 single honest block is produced in $chain[-2\kappa : -\kappa]$, randomness can be extracted from the
6549 chain in order to create the agreed upon random seed. This means that κ should be set such
6550 that the chain quality property can ensure at least one honest block over the κ block period.
6551 Snow White uses the FruitChains idea (Section 11.3) to improve chain quality. This setup
6552 still allows an adversary to bias the seed slightly by using strategically chosen data in its
6553 own blocks at the end of an epoch, but they cannot gain a significant long-term advantage
6554 as long as the same seed is reused for leader election a sufficient number of times in a row.

6555 The Snow White protocol frames its lookback parameters by clock time rather than blocks.
6556 Here, 2ω is a lookback parameter for determining the next committee, and it must be
6557 sufficiently large that alert nodes will have a common prefix in their local chains by the
6558 start of an epoch. Similarly, ω is the lookback parameter for determining the next epoch's
6559 random seed. Let $extractpks()$ be a function that takes a chain and determines a committee
6560 from it as a set of public keys. Similarly, $extractseed()$ takes a chain, outputs the random
6561 seed used for leader election, and may be as simple as concatenating the random data
6562 included in each block of the chain prefix.

6563 The e -th epoch takes place during the time interval $[start(e), end(e))$, and lasts for a dura-
6564 tion of T_{epoch} time steps. An alert validator determines the committee for epoch e by finding
6565 the last block of its local chain with a timestamp no later than $start(e) - 2\omega$. Denote the in-
6566 dex of this block in the chain as L_0 . The committee is the output of $extractpks(chain[:L_0])$.
6567 Similarly, at any time $t \in [start(e), end(e))$, alert validators determine the random seed by
6568 finding the last block in its local chain whose timestamp is no greater than $start(e) - \omega$
6569 and denote its index as L_1 . The random seed is the output of $extractseed(chain[:L_1])$. At
6570 any time step t , a validator determines if it has been elected the block proposer by checking
6571 whether $H^{seed_e}(pk, timestamp) < D_p$, where D_p is a difficulty parameter, H is a random
6572 oracle, and pk is the public key of an elected committee member. In proof of stake, if a
6573 public key has m units of cryptocurrency associated with it, then the block proposer check
6574 can be $H^{seed_e}(pk||i, timestamp) < D_p$ for $i \in \{1, \dots, m\}$. The difficulty parameter is set such
6575 that a committee member is elected in a given time step with probability p . If elected,
6576 the validator extends its longest chain by signing a new block. Block timestamps must be
6577 strictly increasing, and validators reject blocks with timestamps in the future. The security
6578 of this scheme depends on the lookback parameters being sufficiently far in the past that
6579 alert nodes will share common prefixes up to at least block L_1 and that the parameters are
6580 sufficiently spaced apart such that $seed_e$ cannot be predicted until sufficiently long after the
6581 committee is elected.

6582 Snow White uses checkpoints to handle posterior corruption attacks. New nodes or those
6583 offline for a long time need to be able to determine the correct history to believe in. They
6584 can do this by contacting a list of nodes (perhaps provided to them upon startup) – the
6585 majority of which are *alert* (i.e., honest and online) – and trusting the majority view in

6586 terms of stake. Further, alert validators will always reject any chain sent to them that would
6587 reorg "too many" blocks. Specifically, any chain that would revert $\kappa_0 = \frac{\kappa}{2}$ blocks is rejected
6588 by alert validators.

6589 Given a sufficient posterior corruption window, even a majority of corrupt stake *before*
6590 that window should not be able to harm consensus, which implies that security requires
6591 that cryptocurrency not change ownership "too quickly" (the stake shift issue described in
6592 Section 12.1.1). The posterior corruption window, $W > \omega$, also relates to security against
6593 weakly adaptive adversaries. When an adversary knows the next seed used for leader elec-
6594 tion, they must not be able to adaptively corrupt leaders for at least one epoch length. Of
6595 course, such adaptive corruptions are possible in the real world. Security requires that for
6596 any committee, the alert stake must outnumber the adversarial stake for a time window of
6597 W . The protocol ensures that if a validator is corrupted after time $t + W$, they can no longer
6598 influence any of the chain's history before time t .

6599 13.1.3. Ouroboros Family: Praos and Genesis

6600 Ouroboros Praos was the first proof-of-stake protocol to be proven secure against fully
6601 adaptive adversaries [319]. Praos uses a VRF for adaptively secure leader election, which
6602 maintains unpredictability even against adversarially chosen keys. To address posterior
6603 corruption issues, Praos uses forward secure key-evolving signatures, where old keys must
6604 be securely erased after a short period (validators are trusted to perform this erasure, even
6605 though it may not be rational). This prevents an adversary from generating signatures for
6606 messages that were issued in the past. As with Nakamoto Consensus, Praos follows the
6607 longest chain rule and was proven secure in the bounded delay model assuming an honest
6608 majority of stake. Unlike with Nakamoto Consensus, Praos validators will not reorg more
6609 than k blocks of the chain, where k is a security parameter.

6610 Praos divides time into discrete *slots*, only some of which have blocks created in them.
6611 Empty slots allow honest validators to remain synchronized. This is analogous to how
6612 Nakamoto Consensus requires the expected block time to be significantly longer than the
6613 time it takes to propagate the block across the network. The canonical longest chain will
6614 have at most one block per slot, though there may be multiple leaders elected during a
6615 particular slot. The use of a VRF ensures that only slot leaders themselves are aware that
6616 they are a leader in a given slot, and other validators are not until they receive signed blocks
6617 that include a valid VRF proof from the legitimate leader. The leader election process in
6618 Ouroboros uses the following parameters:

- 6619 • R is the number of slots in each epoch.
- 6620 • f is the *active slot coefficient*, or the chance that anyone is elected leader of a partic-
6621 ular slot.
- 6622 • ϕ_f is the function that – given a stake fraction α_i for user i – outputs the probability
6623 of being elected leader: $\phi_f(\alpha_i) = 1 - (1 - f)^{\alpha_i}$.

- 6624 • k is the maximum reorg depth.
- 6625 • η_j is the random seed for epoch j .
- 6626 • s is the number of slots to check for chain density in Ouroboros Genesis.

6627 The protocol is defined over a sequence of $L = ER$ slots, $S = \{sl_1, \dots, sl_L\}$, consisting of
6628 E epochs with R slots in each epoch. For each new epoch e_j , the stakeholder distribution
6629 used to select leaders is drawn from the most recent block with timestamp up to $(j - 2)R$
6630 in the local canonical chain. Using the stake distribution as of two epochs in the past helps
6631 ensure agreement and, as with Snow White, prevents an adversary from adaptively choos-
6632 ing their public keys to aid in stake grinding. To generate the random seed for epoch e_j , a
6633 validator takes every block from the chain belonging to the first two-thirds of epoch e_{j-1} ,
6634 concatenates the block randomness VRF outputs and proofs into a value v , and computes
6635 the new epoch randomness η_j as $H(\eta_{j-1} || j || v)$. Stated differently, validators use the first
6636 $\frac{2R}{3}$ blocks of the prior epoch to determine the current epoch's randomness using the entropy
6637 contained in each block from the leader election process itself.

6638 With the stake distribution and epoch randomness determined, slot leader assignment works
6639 as follows. A stakeholder U_i is independently selected as leader for slot sl_j with probability
6640 p_i depending only on their relative stake, α_i . The active slot coefficient, f , determines
6641 the relationship between the probability and the relative stake. Specifically, $p_i = \phi_f(\alpha_i) =$
6642 $1 - (1 - f)^{\alpha_i}$. Since $\phi_f(1) = f$, f is the probability that someone with all of the stake is the
6643 slot leader. This function is unaffected by whether a party splits up their stake into multiple
6644 virtual identities or not. The threshold for a stakeholder U_i for epoch e_j with a stake of α_i^j is
6645 $T_i^j = 2^{L_{VRF}} \phi_f(\alpha_i^j)$, where L_{VRF} is the length of the VRF output. Validators check whether
6646 they are slot leader by evaluating the VRF with the secret key associated with their stake,
6647 taking the slot number and epoch randomness η as input. They are elected to be a leader
6648 if the VRF output is beneath the threshold. The procedure described here is very similar to
6649 Snow White but protects against fully adaptive corruptions. It still allows some mild stake
6650 grinding via the hashing of VRF outputs, which an adversary may influence, but the bias
6651 that this allows is bounded with suitable parameter choices.

6652 With leader election out of the way, the rest of consensus is simple: validators follow and
6653 build off of the longest valid chain they are aware of but refuse to overwrite more than
6654 k blocks of the chain. When a node is offline for a long time or is bootstrapped for the
6655 first time, it must acquire the chain or a checkpoint from a trusted party. The Ouroboros
6656 Genesis protocol was proposed in an effort to remove this trusted component and allow
6657 a new node to bootstrap themselves securely from the genesis block without needing to
6658 acquire a trusted checkpoint [345].

6659 Ouroboros Genesis is nearly identical to Praos but modifies the chain selection rule. Let
6660 the local best chain be C_{max} . For each potential new chain C_i , if C_i forks from C_{max} at most
6661 k blocks, and $len(C_i) > len(C_{max})$, then set $C_{max} = C_i$. If, however, C_i forks from C_{max}
6662 by more than k blocks, validators apply a different rule. Let j be the highest slot index

6663 where C_i and C_{max} have a block in common and s be a parameter for the rule. Then, if
6664 $len(C_i[0 : j + s]) > len(C_{max}[0 : j + s])$, set $C_{max} = C_i$. That is, for some s -slot period after
6665 the point where the two chains diverge, validators compare the density of the two chains
6666 and prefer the one with more blocks shortly after the fork. During this short period after
6667 the fork, the two chains still share a view of how leader election should occur in the post-
6668 fork slots. Intuitively, an honest majority at that time will work on the chain that ends up
6669 becoming canonical, so the canonical chain should be the one that contains more blocks
6670 during that period.

6671 This new chain selection rule is not without risk and may require re-parameterizing the
6672 system in ways that may be detrimental to security. This is because this new chain selection
6673 rule allows the following attack:

- 6674 1. An adversary acquires a substantial amount of stake but still less than half. This
6675 adversary runs a modified version of the software that does not erase old private keys
6676 after evolving them.
- 6677 2. At the start of each epoch, the adversary checks which slots they will be eligible for.
6678 If the number of slots they are selected for suggests that there is a significant chance
6679 of producing more blocks than everyone else combined during some s -slot period of
6680 the epoch, the adversary will observe the production of blocks but avoid creating any.
6681 When the s -slot period is over, the adversary will have near certainty that they could
6682 have overwritten the chain during that period. Let the first slot of this s -slot period
6683 be denoted sl_j .
- 6684 3. The adversary may then wait any arbitrary period of time after the requirements
6685 from the prior step have been met. At any time, they may create an alternative chain
6686 forking from sl_j . This alternative chain will include every block that the adversary
6687 is entitled to (they can do this because they did not delete their keys). Because the
6688 chain density in the s slots after sl_j is higher in this attack chain, the network will
6689 reorg around the attack chain (even, say, 10 years after the fork occurred).

6690 This attack will be possible if there is any adversary who can acquire more blocks than
6691 everyone else for any s -slot period in the lifetime of the system. It would permanently
6692 eliminate the security of the system in a way that may be undetectable to honest parties
6693 until after the fact. Therefore, the security of the system relies on having a parameterization
6694 that would make it vanishingly unlikely for a party possessing a “realistic” amount of stake
6695 to ever be in a position where they can reorg the rest of the network during any s -slot
6696 period. In particular, this requires the product $s * f$ – that is, the number of blocks that
6697 the system expects to be produced over s slots – to be large. Increasing f is equivalent to
6698 lowering the expected block time, which reduces the security margin due to latency issues
6699 (see Section 10.2.1). Alternatively, increasing s necessitates increasing the epoch length
6700 R . Longer epochs mean that chain-based predictability attacks become more severe (see
6701 Section 12.2), and there is a longer period for stake shift to reduce the security margin.

6702 Ignoring this extra attack, Ouroboros Genesis comes extremely close to providing the same
6703 security guarantees as Bitcoin. However, it is still subject to posterior corruption if the
6704 erasure model does not hold, requires a secure initial distribution of stake, and relies on
6705 more advanced cryptographic primitives.

6706 There are several other protocols in the Ouroboros family, but the details of their oper-
6707 ation are out of scope for this document. Ouroboros Chronos extends Genesis in order
6708 to remove the need for globally synchronized clocks and, thus, remove dependence on
6709 the Network Time Protocol (NTP) that proof-of-stake protocols typically require [346].
6710 Ouroboros Cryptsinous is a design that ensures that provably secure proof-of-stake leader
6711 election can occur in a system that includes privacy-preserving transactions such that it is
6712 not clear from the ledger how many coins are associated with each public key [323].

6713 13.1.4. DFINITY

6714 The DFINITY consensus protocol is composed of four layers: identity registration, a ran-
6715 domness beacon, the blockchain, and the notarization layer [347]. By combining these
6716 layers, DFINITY is able to offer some of the benefits of both chain-based and BFT-based
6717 proof-of-stake systems.

6718 Identity registration requires potential validators to make a deposit of stake associated with
6719 their public key so that a known and agreed upon number of parties exists in order to per-
6720 form the necessary threshold cryptography utilized by the protocol and so that misbehavior
6721 can be punished by slashing these deposits. These registered keys are then used to set up the
6722 randomness beacon by creating a VRF out of the BLS threshold signature scheme. A group
6723 of validators must run a distributed key generation (DKG) protocol to create a group pub-
6724 lic key for verifying the random outputs. To compute the random seed for round r , $seed_r$,
6725 members of the group provide partial signatures over the round number and $seed_{r-1}$. The
6726 partial signatures are aggregated into the final signature, which is then hashed to produce
6727 $seed_r$.

6728 Since many users are expected to participate, the scalability of the system requires that
6729 only subsets of the registered users are involved in generating random beacon values (as
6730 well as participating in notarization, which is described later). A random seed is used as
6731 part of a cryptographic shuffle to select a list of m groups or committees. The committees
6732 run the DKG protocol, and the resulting public keys can be put in the genesis block. The
6733 round r committee used for the beacon and notarization, G^j , is the j -th listed group, where
6734 $j = seed_r \bmod m$. The members of group G^j create the random seed used to select group
6735 G^{r+1} .

6736 In practice, it is unlikely that all identities will be registered before the system starts up, and
6737 allowing dynamic participation is desirable. To do so, special registration and deregistration
6738 transactions can be used to join or exit the system. Time is divided into epochs, where the
6739 first block of the epoch contains a summary of all the registrations or deregistrations from

6740 the prior epoch. In this new block, a random seed is used for a cryptographic shuffle as with
6741 the static case, and new groups can perform DKG and issue another special transaction to
6742 register their aggregated group public key. These newly registered groups can participate
6743 and be selected two epochs after they have been registered.

6744 The blockchain itself is constructed using the *probabilistic slot protocol* and *notarization*.
6745 These layers assume that a group G^r and random beacon value $seed_r$ has been selected
6746 for round r . Any user $U \in G^r$ may produce a block in round r , but the probabilistic slot
6747 protocol establishes a *priority* among group members and favors high priority blocks. A
6748 cryptographic shuffle is performed using $seed_r$ to establish a ranked order among group
6749 members, where a lower rank means a higher priority. There is also a *weight* function
6750 $w(x) = 2^{-x}$, and a block B 's weight is $w(B's\ rank)$. The weight of a chain is the sum of the
6751 weight of its blocks. Similar to other longest chain protocols, the heaviest chain wins in
6752 DFINITY.

6753 As validators produce blocks in a round, they are sent to the group in charge of notarization
6754 for that round, which may be the same group. The notary waits for a protocol-specified
6755 *BlockTime* to receive proposed blocks. After *BlockTime*, notary group members will sign
6756 all of the highest priority valid blocks that they saw in that time and continue to sign the
6757 highest priority blocks until the next round begins. The next round begins after observing a
6758 notarization for the current round, where a notarization is an aggregated threshold signature
6759 on the block. The notarized block is broadcast to the network, and validators will update
6760 their local chains, end the round, and have the random beacon output a seed for the next
6761 round. Because the notary is optimistic instead of a full consensus algorithm, it can notarize
6762 two conflicting blocks at the same height. This is resolved through the chain selection rule's
6763 weight function. The role of notarization is to make it impossible for an adversary to build a
6764 secret chain of notarized blocks, thus addressing nothing-at-stake issues and selfish mining.
6765 It also establishes finality for blocks, such that a finalized block can never be undone.

6766 The finalization procedure involves gathering notarized block proposals and placing them
6767 into buckets based on their round. Once a notarized block is found in round r , finalization is
6768 scheduled for the blocks of round $r - 1$ in T time units. At that time, the heaviest common
6769 prefix ending at round $r - 1$ is finalized. If at round r , all valid, notarized blocks B_{r-1} point
6770 to the same predecessor, B_{r-2} , validators finalize B_{r-2} and its predecessors. Each observer
6771 can specify their own T ; it need not be common.

6772 The security proof for DFINITY assumes a synchronous network and is secure for an hon-
6773 est majority of stake against mildly adaptive adversaries. If Δ is the maximum network de-
6774 lay, then secure parameterization of DFINITY requires that $BlockTime \geq 3\Delta$ and $T \geq 2\Delta$.
6775 The DFINITY protocol was further analyzed in [348], which pointed out that if the highest
6776 priority member of a group is Byzantine, the communication complexity of the protocol
6777 is unbounded because they can send an unlimited number of highest-priority blocks, and
6778 honest validators are required to vote on all of them. To fix this, whenever an honest val-
6779 idator sees two blocks from the same block proposer, the conflicting signatures are used as

6780 proof of misbehavior and lower the effective priority of the block. This fix increases DFIN-
6781 ITY's worst-case latency to 17Δ . In the optimistic case that the actual network delay is
6782 small compared to Δ , the expected latency is 8Δ . When the actual delay is Δ , the expected
6783 latency is 14Δ without equivocation.

6784 13.2. Ethereum 2.0

6785 The Ethereum network has recently undergone a transition from a pure proof-of-work sys-
6786 tem using GHOST as the fork-choice rule (Section 11.2) to a hybrid system that combines
6787 proof-of-work and a proof-of-stake finality layer (Casper FFG, described in Section 14.3.2)
6788 and, finally, to a pure proof-of-stake system dubbed Ethereum 2.0 [349]. This section may
6789 not completely match the "final" design for Ethereum's proof-of-stake system since it is
6790 under continual development and subject to change.

6791 The architecture of Ethereum 2.0 includes a *beacon chain* that handles random number gen-
6792 eration and coordinates separate *shard* chains. Sharding is a scalability technique in which
6793 each shard is its own separate blockchain with its own state (described in further detail in
6794 Section 15). It allows nodes to only validate a fraction of the system's total transaction
6795 throughput. The beacon chain stores validator information and establishes consensus over
6796 data related to the state of each shard while mediating cross-shard communications. To
6797 become a beacon chain validator, users must deposit 32 ETH into a staking contract so that
6798 their stake may be slashed in case of misbehavior. This amount was chosen to balance com-
6799 peting factors: larger deposits prevent people from participating, whereas smaller deposits
6800 increase the overhead of verifying the chain. The beacon chain divides time into 12 second
6801 slots, and each epoch has 32 slots (6.4 minutes). An "eek" is 2,048 epochs (≈ 9 days). In
6802 each epoch, validators make *attestations*, which consist of:

- 6803 • A hash of what the validator considers the chain tip
- 6804 • A hash of what the validator believes is the correct shard block to include
- 6805 • The "source" and "target" hashes from Casper FFG (see Section 14.3.2 for details)
- 6806 • A signature from the validator over the above data

6807 These attestations are used to come to consensus using a combination of a variant of the
6808 GHOST fork-choice rule called Latest Message Driven (LMD) GHOST and Casper FFG
6809 [350]. LMD GHOST is a proof-of-stake variant of GHOST where, at each fork, nodes pre-
6810 fer the fork that contains more total support from validators (based on the stake-weighted
6811 sum of attestations) while counting only the most recent message from each validator. The
6812 combined protocol (dubbed "Gasper") finalizes blocks, keeps track of the latest justified
6813 checkpoint using Casper FFG, and uses the LMD GHOST rule to determine the chain tip
6814 by treating the latest justified checkpoint as the root of the chain. The LMD GHOST rule
6815 is shown in Figure 40.

6816 Validators are rewarded or punished based on their attestations. Correct attestations are




Fig. 40. LMD GHOST. The number in each block is the weight of a block when each vote has weight one (circles represent votes). The blue chain, ending with the block of weight three, is the canonical chain. [350]

6817 rewarded, whereas missed slots result in penalties. The interest rate for correct voting de-
6818 pends on how many validators are participating. Fewer staking validators will raise the
6819 interest rate to entice more validators to secure the chain. If penalties accrue to the point
6820 where a validator has less than 16 ETH at stake, they are ejected from the validator set. Val-
6821 idators who equivocate by voting for conflicting blocks are slashed, lose some fraction of
6822 their deposit (substantially more than the penalty for missing a slot), are removed from the
6823 set of validators, and have the remainder of their deposit frozen for an extra four eeks be-
6824 fore being allowed to withdraw their funds. During this waiting period after being slashed,
6825 the validator is penalized proportionally to how many other validators are slashed during
6826 that same period. This is to discourage correlated failures. Isolated, honest mistakes will
6827 not be penalized as heavily as active attacks using a large portion of the total stake. It also
6828 means that smaller validators take on less risk than larger ones and discourages validators
6829 from joining large stake pools.

6830 Validators may also voluntarily leave the system and are allowed to withdraw their funds
6831 after 256 epochs, or about one day. If a large number of validators try to exit at once, a
6832 queue will form and they will exit over time. This queuing prevents an adversary from
6833 creating many validators, performing some malicious action, and then exiting before they
6834 can be slashed. Validator registration is similarly rate-limited.

6835 Leader election in Ethereum 2.0 requires generating unpredictable random seeds using a
6836 commit-reveal mechanism inspired by a process called RANDAO. Ultimately, the random
6837 seed is used both to determine block proposers and to assign validators to committees.
6838 There is a beacon chain committee per slot, where one validator is the block proposer and
6839 the rest of the committee members issue attestations for the slot. There are also committees

6840 of validators responsible for handling each particular shard chain. A cryptographic shuffle
6841 is performed on the validator list using the seed as input, and committees are consecutive
6842 slices of the resulting list.

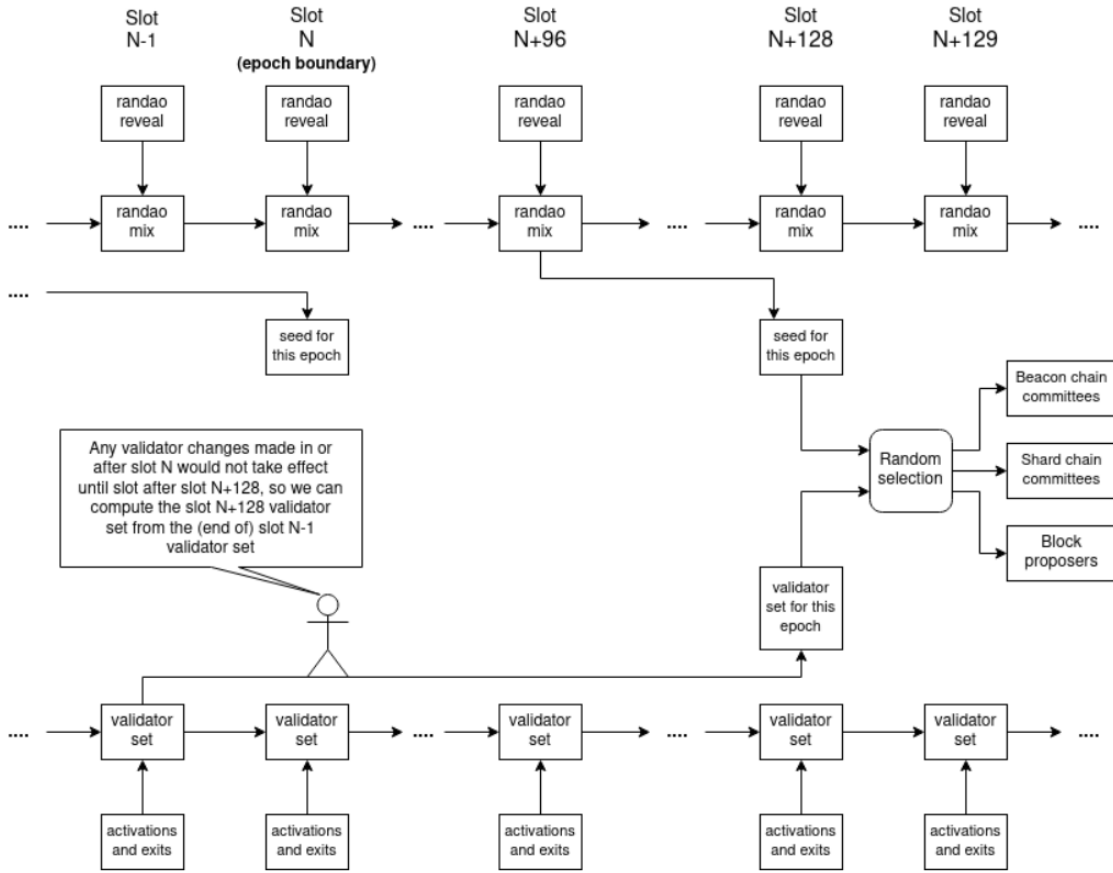
6843 To prevent stake grinding via adversarial key selection, the validator set is fixed four epochs
6844 in advance. There is only a single block proposer per slot, and this proposer includes a
6845 *randao mix* into their block. The randao mix is a hash that is XORed into the seed to update
6846 it every slot. Specifically, it is the hash of a BLS signature over the current epoch. This
6847 hash is unknown to other validators ahead of time, but there is only one valid submission
6848 due to the uniqueness property of BLS signatures. The randao mix from the beginning of
6849 epoch e is used to calculate the seed in epoch $e + 1$. Note that only the last proposer in an
6850 epoch has the ability to bias the seed by one bit, which they can do based on their decision
6851 to publish a block or not, potentially sacrificing their reward by withholding a block. This
6852 process is shown in Figures 41a and 41b.

6853 To prevent even this single bit of bias that an adversary can induce, the output of the randao
6854 mix can be input into a verifiable delay function (VDF). The VDF's delay is parameterized
6855 to be longer than the time window where a validator could benefit from influencing the
6856 random seed, or at least one epoch. This prevents the final block proposer from being able
6857 to know the eventual seed quickly enough to decide whether to withhold their own randao
6858 mix or not. The Ethereum 2.0 randomness generation process favors liveness over being
6859 unbiased because of the possibility of using a VDF as well as the economic penalty that
6860 comes from sacrificing the block it would take to manipulate the seed. As a result, the
6861 beacon chain can continue producing pseudorandom numbers even when many validators
6862 are partitioned from the network or offline. Contrast this with, say, DFINITY, where a
6863 group may fail to reach the threshold needed to produce the aggregate BLS signature, thus
6864 stalling the chain until the network is healed.

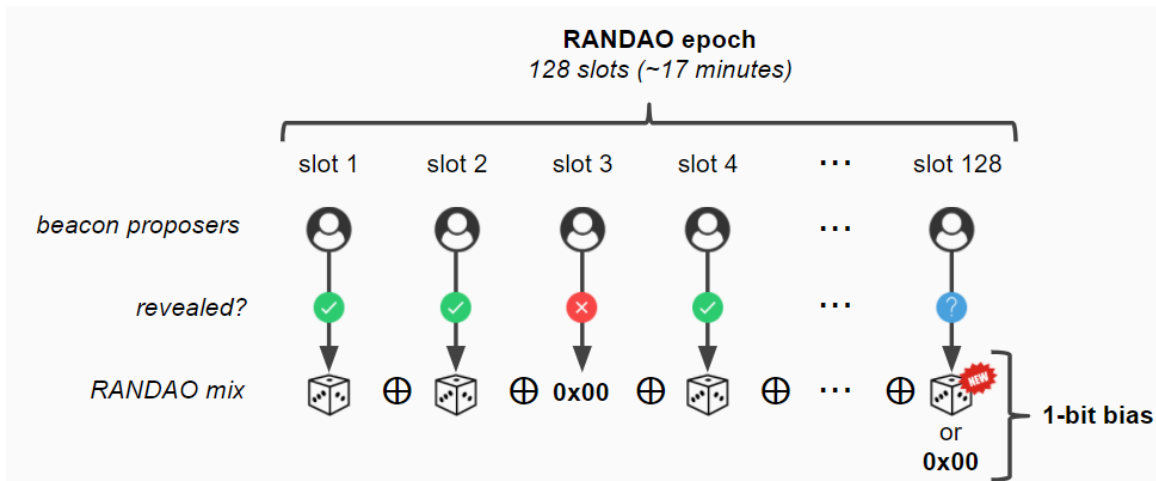
6865 13.3. DAG-based Proof of Stake

6866 13.3.1. Fantôme

6867 Fantôme is a DAG-based protocol that was designed to be game-theoretically secure un-
6868 der the BAR model (see Section 1.4) so long as the non-altruistic stake is less than $\frac{1}{3}$ of the
6869 total stake [353]. Specifically, a coalition of less than $\frac{1}{3}$ of participants are unable to sig-
6870 nificantly increase their rewards if the rest of the participants are altruistic, and a coalition
6871 of up to one-fourth of the participants being Byzantine are unable to lower the payoff for
6872 altruistic participants. It employs several cryptographic tools, such as publicly verifiable
6873 secret sharing (PVSS), a verifiable random function (VRF), and a verifiable delay func-
6874 tion (VDF). In Fantôme, the DAG is formed using two kinds of references: each block
6875 has a single parent block, denoted B_{prev} , but can also reference a set of other chain tips,
6876 which they call *leaves* and denote B_{leaf} . Fantôme is inspired by the Phantom protocol,
6877 and its description utilizes much of the same language employed in Section 11.5.2. Some
6878 additional definitions are used as well:



(a) Ethereum 2.0 randomness



(b) RANDAO in Ethereum 2.0

Fig. 41. Ethereum 2.0 Randao architecture. Note that 128 slots would take 25.6 minutes rather than 17, based on the proposed slot duration of 12 seconds. [351, 352]

- 6879 • $Ancestors(B)$ is the set of all blocks that are direct or indirect parents of B .
- 6880 • $DirectFuture(B)$ is the set of blocks that directly reference B (that is, blocks where
6881 $B \in B'_{leaf}$ for any B'). Contrast this with $future(B)$, which includes blocks that indi-
6882 rectly reference B as well.
- 6883 • The *biggest common prefix DAG* (BCPD) is the largest subDAG that more than half
6884 of the players agree on.
- 6885 • *Double* is the set of all blocks that use the same proof-of-stake leader election eligi-
6886 bility proof but different contents. That is, these are equivocating blocks.

6887 Fantôme uses a leader election protocol called Caucus that is designed to be secure
6888 against fully adaptive adversaries and is similar to Algorand’s cryptographic sortition (Sec-
6889 tion 13.4.2). To be considered during the leader election process, participants are required
6890 to post a security deposit that may be slashed if misbehavior is detected. Caucus has four
6891 steps:

- 6892 1. **Commit.** Participants must commit to their VRF secret key, sk , by issuing a special
6893 *commit transaction*, tx_{com} , that specifies their corresponding public key pk . These
6894 commitments are appended to a list of commitments, c , that is part of the system’s
6895 public state, $state_{pub}$. After being added to c , a participant must wait for some fixed
6896 number of protocol rounds, x_{wait} , before they may be elected as leader. The added
6897 waiting time maintains unpredictability and prevents the adversary from grinding
6898 through adversarially chosen keys. Denote the round that a participant issued tx_{com}
6899 as rnd_{joined} .
- 6900 2. **Update.** This step is run only on the first round, where $rnd = 1$. A certain threshold
6901 of participants must have successfully committed to the leader election. At this point,
6902 the participants run a coin-tossing protocol (using a PVSS) to generate a random
6903 value, R_1 , and then update the system state to $state_{pub} = (c, R_1)$.
- 6904 3. **Reveal.** Participants check their own eligibility in each round where $rnd > 1$. Let
6905 $y_{rnd} = VRF_{sk}(R_{rnd})$, n_{rnd} be the number of eligible participants that have waited
6906 enough rounds since their commitment round, and $target = \frac{H_{max}}{n_{rnd}}$. A participant is
6907 eligible if $H(y_{rnd}) < target$. If a participant discovers that they are eligible, they is-
6908 sue a special *reveal transaction*, tx_{rev} , that includes y_{rnd} and $p_{rnd} = p_{sk}(R_{rnd})$, where
6909 p_{rnd} is the VRF proof.
- 6910 4. **Verify.** When a participant sees a transaction tx_{rev} from participant i , they check
6911 whether $VerifyVRF(R_{rnd}, y_{rnd}, p_{rnd}) = 1$ and that $rnd_{joined} > rnd - x_{wait}$. If so, then
6912 the public randomness is updated to $R_{rnd+1} \leftarrow R_{rnd} \oplus y_{rnd}$, and the verification is
6913 considered a success. Otherwise, the public state remains unchanged.

6914 In some rounds, there will not be any elected leader, or an adversary is elected but
6915 does not issue a tx_{rev} . If no participant reveals tx_{rev} in a given round, the public

6916 randomness is updated as $R_{rnd} \leftarrow VDF(R_{rnd})$. If an honest participant finds that they
6917 are not eligible in the reveal phase, they immediately begin computing the VDF. If
6918 they have not seen a valid tx_{rev} by the time the VDF is computed, they re-check their
6919 eligibility using the new VDF output as R_{rnd} . This maintains liveness.

6920 There may also be more than one valid tx_{rev} . In this case, the winning leader can be
6921 the participant whose y_{rnd} is the lowest.

6922 The Caucus protocol may be modified in order to improve fairness by ensuring that the
6923 same participant is not elected leader a disproportionate number of times. When computing
6924 the *target*, the system may use $\frac{n_{rnd}+1}{2}$ instead of n_{rnd} . In this case, the verify step would
6925 add a rule that the participant was not leader in the prior $\frac{n_{rnd}-1}{2}$ rounds.

6926 The basic idea behind Fantôme is that participants use their stake to bet on the block that
6927 they believe has the highest *score*, where a block is considered a bet on its ancestor set,
6928 while simultaneously demonstrating that they are well-connected and honest by including
6929 other leaf block references.

6930 Fantôme employs a distributed checkpointing scheme in order to finalize blocks. Begin-
6931 ning from the genesis block, a candidate for finalization is a block that bets on the genesis
6932 block or specifies the genesis block as its parent. Candidate blocks have a *rank*, where this
6933 first set of blocks (with the genesis block as parent) belongs to rank $rk = 1$. In a system
6934 with n participants, a block B is a *witness* for candidate block C if $\frac{2n}{3}$ participants have bet
6935 on C when considering the set $past(B) \cup \{B\}$. In this case, C is said to be *justified*. If, in
6936 turn, a block B_2 is a witness of B , it is said to be a *second witness* of C . When a block B
6937 bets on a second witness with rank rk , then B is considered a candidate for rank $rk + 1$. As
6938 described, this procedure requires a fixed set of participants and would require adjustment
6939 to allow participants to join or leave during the checkpointing process.

6940 Fantôme's fork-choice rule calculates the score of a leaf block B as the sum of the number
6941 of references made by each block in $past(B)$, except for equivocating blocks (those in the
6942 set *Double*). To decide which block to bet on, a validator computes the score of each block
6943 and sets the parent reference as the highest scoring leaf block. More formally, given a DAG
6944 G :

- 6945 1. $w \leftarrow \emptyset$
- 6946 2. For $B \in Leaves(G)$ do: for $B' \in past(B) \setminus Double$ do: $w[B'] = |B'[B_{leaf}]|$
- 6947 3. $CW \leftarrow argmax_{B \in leaf(G)} w(B)$
- 6948 4. In case of a tie, choose the block with the smallest hash: $B \leftarrow argmin_{B \in CW} H(B)$, and
6949 return B .

6950 After computing the winning block B based on the fork-choice rule above, the validator
6951 checks their most recent second witness block and, to avoid long-range attacks, ensures

6952 that there is a candidate block associated with it in $Ancestors(B)$. They then check whether
6953 they are eligible using the VRF and, if so, produce a block with B as the parent and every
6954 other leaf they see as a reference. The block B is considered valid if:

- 6955 1. $B_{prev} = ForkChoiceRule(past(B))$. That is, B bets on the block that the fork choice
6956 rule requires based on B 's view of the DAG.
- 6957 2. The creator of block B successfully proves their eligibility.
- 6958 3. If a witness block exists in $past(B)$, then there is also a witness block in $Ancestors(B)$.
6959 In other words, a participant may not bet on a non-justified block if they are aware of
6960 a justified one.
- 6961 4. If B bets on a second witness block B_s , then B also bets on a block in $past(B_s)$. That
6962 is, $Ancestors(B) \cap past(B_s) \neq \emptyset$.

6963 Incentivization in Fantôme uses a function inspired by Phantom that labels each block as
6964 a winner, a loser, or neutral. If the fork-choice rule selects a block B , then B and blocks in
6965 $Ancestors(B)$ are labeled as winners. A block that bets on a winner is neutral, and winning
6966 and neutral blocks in DAG G form $BlueSet_G$. Similar to Phantom, for a parameter k , a block
6967 B where $|anticone(B) \cap BlueSet_G| \leq k$ is labeled neutral. If $|anticone(B) \cap BlueSet_G| > k$,
6968 then the block is a loser. Specifically, the labeling function with input a DAG G and output
6969 a labeling of the DAG M is:

6970 1. Set:

- 6971 • $B \leftarrow ForkChoiceRule(G)$
- 6972 • $BlueSet_G \leftarrow BlueSet_G \cup \{B\}$
- 6973 • $M(B) = winner$

6974 2. For $B_i \in Ancestors(B)$ do:

- 6975 • $BlueSet_G \leftarrow BlueSet_G \cup \{B_i\}$
- 6976 • $M(B_i) = winner$
- 6977 • For $B_j \in DirectFuture(B_i) \setminus Ancestors(B)$ do:
 - 6978 – $BlueSet_G \leftarrow BlueSet_G \cup \{B_j\}$
 - 6979 – $M(B_j) = neutral$

6980 3. For $B_i \in G \setminus BlueSet_G$ do:

- 6981 • If $|anticone(B_i) \cap BlueSet_G| \leq k$, then:
 - 6982 – $BlueSet_G \leftarrow BlueSet_G \cup \{B_i\}$

6983 – $M(B_j) = \textit{neutral}$

6984 • else: $M(B_i) = \textit{loser}$

6985 4. Return M

6986 Rewards are calculated over a period of some number of blocks, at which point the above la-
6987 beling function is applied to the BCPD. To encourage participants to reference each others'
6988 blocks and have the DAG be well-connected, the reward that a block provides is propor-
6989 tional to the number of leaves it references: $\textit{rwd}(B) = |B_{\textit{leaf}}| * c$, for some protocol-defined
6990 constant c . There are also two constants for punishing misbehavior or poor connectivity:
6991 blocks labeled as losers will penalize their creator by \textit{pun} , and malicious behaviors (like
6992 equivocation or not referencing one's own blocks) are punished by \textit{bigpun} .

6993 Fantôme's reward scheme is incentive-compatible, so participants have the highest utility
6994 by behaving honestly. Recall that the validity of a block depends on it being the one chosen
6995 via the fork-choice rule based on its past set. In addition, rewards are proportional to the
6996 number of leaves referenced in the block, and the score of a block also corresponds to the
6997 number of leaves it references. Because the participant must follow the fork-choice rule,
6998 the only way they can pick another block is by reducing the number of leaves it references.
6999 However, this makes the block more likely to be labeled as a loser (as it will have a larger
7000 anticone) and decreases its reward if it is not. Participants can also withhold their blocks,
7001 but doing so means they will not be referenced as frequently and are less likely to be labeled
7002 as winners.

7003 13.3.2. Avalanche

7004 Avalanche is a leaderless protocol that provides a partial ordering of transactions via a
7005 transaction DAG [354]. It is typically associated with proof of stake but is agnostic to the
7006 Sybil-resistance mechanism (there is even a proposal for using IP prefixes for this purpose
7007 [355]). Avalanche – or rather, the Snowball protocol that it is based off of – operates rather
7008 differently from other algorithms described in this document. Nodes arrive at decisions
7009 based on repeatedly sampling other nodes for their opinion and then being steered toward
7010 a common view as rounds progress. This structure provides considerable efficiency advan-
7011 tages: there is only $O(1)$ communication cost per round over an expected $O(\log n)$ rounds
7012 for a given node.

7013 In fact, [354] presents a family of "Snow" protocols that follow a similar structure. Of
7014 particular importance is the Snowball protocol, because Avalanche is actually composed of
7015 multiple Snowball instances organized over a transaction DAG, which reduces costs from
7016 logarithmic in the number of nodes to constant and results in very low latencies for trans-
7017 action settlement. As with several other DAG-based protocols, the algorithm is framed as
7018 deciding between two different colors: red and blue. Essentially, Snowball can be consid-
7019 ered a synchronous binary agreement protocol where the agreement property is relaxed to
7020 be probabilistic.

7021 In Snowball, a node will repeatedly sample a constant number k of other nodes on the
7022 network until it has sufficient confidence that the network is consistent. There are two
7023 relevant security parameters – α and β – that are decision-making thresholds employed in
7024 the algorithm. Nodes also maintain some internal state for keeping track of their confidence
7025 in deciding on a color. Specifically, $d[R]$ and $d[B]$ are counters that are incremented each
7026 time a round of sampling results in their respective color receiving a threshold α of votes
7027 (where $\alpha > \lfloor \frac{k}{2} \rfloor$). An additional counter, cnt , is incremented each time a round of sampling
7028 fails to change the node’s preference to the other color and is reset to one whenever the
7029 color preference switches. For a node to change its color preference, its confidence count
7030 in the alternate color must exceed the confidence count in its currently preferred color. That
7031 is, if the node currently prefers red, then it will switch to blue only when $d[B] > d[R]$, and
7032 vice versa. A node will only consider a decision final on their current color when it has at
7033 least β consecutive rounds where the color preference was unchanged (i.e., when $cnt \geq \beta$).
7034 More formally, Snowball works as follows in order to decide between red and blue, where
7035 a node takes $col_0 \in \{R, B, \perp\}$ as input:

- 7036 1. Initialize state.
 - 7037 • $col := col_0$
 - 7038 • $lastcol := col_0$
 - 7039 • $cnt := 0$
 - 7040 • $d[R] := 0$
 - 7041 • $d[B] := 0$
- 7042 2. While undecided, do:
 - 7043 (a) If $col = \perp$, continue
 - 7044 (b) $K := Sample(k)$
 - 7045 (c) $P := [Query(v, col)$ for $v \in K]$
 - 7046 (d) $maj := false$
 - 7047 (e) For $col' \in \{R, B\}$, do:
 - 7048 i. If $P.COUNT(col') \geq \alpha$, then:
 - 7049 • $maj := true$
 - 7050 • $d[col'] ++$
 - 7051 ii. If $d[col'] > d[col]$, then: $col := col'$
 - 7052 iii. If $col' \neq lastcol$, then:

- 7053 • $lastcol := col'$
- 7054 • $cnt := 1$
- 7055 iv. Else if $col' = lastcol$, then: $cnt ++$
- 7056 v. If $cnt \geq \beta$ then ACCEPT(col').
- 7057 (f) If $maj = false$, then: $cnt := 0$

7058 Snowball is able to provide probabilistic consistency guarantees with failure probability
7059 ϵ (not unlike Nakamoto Consensus). Even if the network begins in a fully bivalent state
7060 (equal preference for each color), the random perturbations that occur due to sampling
7061 will cause the preferences to drift toward one color or the other by some amount δ . As
7062 δ increases, it becomes exponentially less likely that the minority value will overtake the
7063 majority. Consistency requires properly setting the parameters k , α , and β . The liveness
7064 guarantee provided by Snowball is such that if the number of adversarial nodes $f \leq O(\sqrt{n})$,
7065 then the protocol terminates with probability $\geq 1 - \epsilon$ in $O(\log n)$ rounds. The number of
7066 rounds increases as the adversary controls more nodes, requiring an exponentially increas-
7067 ing number as f approaches $\frac{n}{2}$ (as with Nakamoto Consensus).

7068 Avalanche uses Snowball internally to decide between *conflict sets*, or transactions that
7069 spend the same funds. When there are no double-spends on the network, a conflict set
7070 contains just a single transaction. When clients issue transactions, they specify one or more
7071 parent transactions which form the edges of a DAG. The DAG improves efficiency because
7072 voting for a particular transaction also counts as a vote on all of its ancestor transactions
7073 back to genesis.

7074 Let tx be a transaction that an Avalanche node is trying to decide on. The node will query
7075 a sample of its peers to ask about tx . Other nodes will only support tx if every trans-
7076 action reachable from tx in the DAG is the preferred transaction in their respective con-
7077 flict sets. Stated differently, a node will only vote yes on tx if $\forall tx_j \in ConflictSet(tx_i)$,
7078 $\forall tx_i \in past(tx) \cup \{tx\}$, tx_i is preferred over tx_j . Should a threshold α of yes votes be re-
7079 ceived, the transaction collects a *chit*. A node's confidence in a transaction is the number
7080 of chits in the future set of that transaction, or transactions from which it is reachable in
7081 the DAG. As such, confidence increases as the DAG expands. Ties are broken in favor of
7082 the transaction that a node saw first. The transaction recipient decides when to accept the
7083 transaction based on a threshold of β consecutive chits in its favor.

7084 This procedure is similar to that employed in the Tangle algorithm (Section 11.5.3). A
7085 crucial difference is that in Avalanche, confidence in a transaction is not based solely on the
7086 structure of the DAG but rather the accumulation of chits. This makes Avalanche resistant
7087 to certain attacks that Tangle-based systems are subject to, like parasite chain attacks, where
7088 the adversary creates large subDAGs to overpower the honest portion of the graph. Another
7089 protocol that is very similar to Avalanche is Fast Probabilistic Consensus (FPC) [356].
7090 FPC assumes that nodes have access to some form of trusted randomness beacon, perhaps

7091 through a distributed random number generation process. A round of the FPC protocol
7092 works in three steps with a different k queries to different nodes in each round. Replicas
7093 maintain a counter that is incremented in each round where the consensus value remains
7094 the same, and the protocol terminates for the node when this exceeds another threshold
7095 value. The three steps in a round are:

- 7096 1. Each node samples the beliefs of k other nodes.
- 7097 2. After sampling, the nodes run some kind of distributed randomness generation pro-
7098 tocol to determine a (moving) threshold X_r where $0.5 < X_r < 1$.
- 7099 3. Each node then chooses 1 if the number of nodes in their sample was greater than
7100 $k * X_r$ and 0 otherwise.

7101 The first round has a separate initial threshold value for step two, but what is described
7102 above is for all other rounds. Because the random value is chosen *after* the sampling is per-
7103 formed, even a nearly omniscient adversary (knowing everything except the future random
7104 value) will not know which threshold to attempt to sway the sample toward. In particular,
7105 this helps with speedy termination relative to Snowball when an adversary attempts to in-
7106 terfere. The security threshold for FPC is atypical and approximately 38%, but the reasons
7107 are beyond the scope of this document.

7108 13.3.3. Parallel Chains

7109 The idea of composing a ledger out of m separate ledgers was discussed in the proof-of-
7110 work context in Section 11.4. In that case, it was used to reduce the settlement latency
7111 of transactions. A similar approach for proof of stake was proposed in [357], where the
7112 composition of m chains is used to achieve nearly optimal transaction throughput. Unlike
7113 the proof-of-work case, this scales throughput rather than latency, and transactions can only
7114 appear on a single chain instead of multiple chains. Each chain shares a common genesis
7115 block, and the leader election procedure is executed independently for each chain but using
7116 the shared stake distribution across the entire system.

7117 The idea from [357] can be used with other chain-based proof-of-stake protocols but is
7118 described as the composition of m separate chains that operate using a modified version of
7119 Ouroboros Praos (Section 13.1.3) that also adopts the Inclusive rule described in Section
7120 11.5.1. The use of the Inclusive rule means that this scheme runs m parallel DAGs, though
7121 a single linear chain is formed for each of the m DAGs by including blocks that would
7122 have otherwise been rejected by the longest chain rule. A combining procedure is used to
7123 take these m separate DAGs and form a single linearized ledger of transactions. Unlike
7124 many DAG protocols, this approach keeps separate graphs that grow more slowly and fork
7125 infrequently. Typical DAGs increase throughput by producing blocks rapidly and accepting
7126 frequent forks.

7127 Validators include references in their blocks to other blocks off of the main chain other than

7128 the specific parent reference, as is typical in DAGs. For a main chain $C = B_1, B_2, \dots, B_L$ of
7129 blocks B_i , the Inclusive rule includes blocks outside of C into a single chain as follows: for
7130 $B_i \in C$, insert before B_i all blocks in $past(B_i) \setminus past(B_{i-1})$, where those blocks are sorted
7131 topologically and ties are broken via block hash.

7132 The parallel chains approach is to then perform a comparable procedure across all m chains,
7133 C_1, \dots, C_m . All blocks in all chains are first put in order based on their respective slot index
7134 with ties broken based on their chain number $c \in \{1, \dots, m\}$. The result is a sequence of
7135 blocks B_1, B_2, \dots, B_L . Then, $\forall i \in \{1, \dots, L\}$, insert prior to B_i all blocks from $\bigcup_{j \in \{1, \dots, m\}} S_j \cap$
7136 $\{past(B_i) \setminus past(B_{i-1})\}$ sorted topologically (and breaking ties by block hash). Here, S_i is
7137 the portion of chain C_i that is stable (i.e., with blocks at the end chopped off to maintain
7138 a common prefix). Finally, based on this block ordering, transactions must be sanitized in
7139 order to remove those that are invalid.

7140 13.4. BFT-Based Proof of Stake

7141 13.4.1. Tendermint

7142 The Tendermint algorithm, which is an adaptation of the PBFT algorithm to proof of stake,
7143 is described in [358]. The original version had flaws that were found and fixed in [359],
7144 and the latest version was analyzed and proven secure in [360]. The novelties of Tender-
7145 mint as compared to PBFT are the use of gossip for networking and the elimination of
7146 the separate view change algorithm (similar to some of the algorithms from Section 5). A
7147 new leader is elected in each round as part of the normal processing rather than using the
7148 additional subprotocol described in Section 4.1. This improves Tendermint's worst-case
7149 communication complexity from PBFT's $O(n^4)$ to $O(n^3)$ because processes locally keep
7150 track of potentially decided values rather than exchanging all of the messages they have
7151 already delivered. Instead of requiring $\frac{2n}{3}$ validator signatures, as in PBFT, Tendermint re-
7152 quires $\frac{2}{3}$ of the total stake to sign off on blocks. For ease of exposition, this will sometimes
7153 be described as seeing $2f + 1$ signatures or messages. Unlike most other proof-of-stake
7154 algorithms, Tendermint's reliance on PBFT makes it secure under partial synchrony rather
7155 than only in synchronous networks. Unlike PBFT, validators are not necessarily connected
7156 over a complete network but rather communicate via gossip over a peer-to-peer overlay
7157 network.

7158 Leader election follows a stake-weighted round-robin process, such that block proposers
7159 are selected proportionally to their stake. In Tendermint, a single block is created at a given
7160 chain height, but multiple rounds may be required to generate this block if replicas cannot
7161 agree on a value under a given leader. Each consensus instance is described by a chain
7162 height and round number, and all messages exchanged will include those values. Ten-
7163 dermint uses PROPOSAL, PREVOTE, and PRECOMMIT messages (which correspond to
7164 their PBFT equivalents of PRE-PROPOSE, PROPOSE, and COMMIT, respectively), and
7165 each of these message phases has a corresponding timeout.

7166 The PROPOSAL message includes the full block proposed, but the PREVOTE and PRE-
7167 COMMITs just carry a block hash in order to save bandwidth. After receiving a PRO-
7168 POSAL for value v , validators send a PREVOTE for it if valid (or a PREVOTE with a
7169 special *nil* value if invalid or timed out). Upon receiving PREVOTE messages for v signed
7170 by keys associated with more than $\frac{2}{3}$ of the system's stake, validators send a PRECOMMIT
7171 message for v (or PRECOMMIT with a *nil* value on timeout or if the replica has not seen a
7172 valid block with the given hash). Similarly, after receiving $2f + 1$ matching PRECOMMITs
7173 in a given round, a correct process decides on that value. If multiple rounds are necessary at
7174 a given chain height, a correct round leader will propose the same value as the prior round
7175 leader if that value is valid. So far, this is just standard PBFT, but Tendermint also includes
7176 several locking variables in order to maintain safety across different rounds without a view
7177 change mechanism:

- 7178 • *lockedValue* is the most recent value with respect to a round number for which a
7179 PRECOMMIT has been sent.
- 7180 • *lockedRound* is the last round in which the validator sent a PRECOMMIT that is not
7181 *nil*.
- 7182 • *validValue* is the most recent possible decision value. *validValue* is the last value
7183 that a validator delivered $2f + 1$ times and can differ from *lockedValue*.
- 7184 • *validRound* is the last round in which *validValue* was updated.

7185 Leaders send PROPOSAL messages that include their local *validValue* if their *validValue* \neq
7186 *nil*. These messages include *validRound*, so other processes are informed about the last
7187 round in which the proposer observed *validValue* as a possible decision value. Other val-
7188 idators then check the PROPOSAL against their own local *lockedValue*. A PROPOSAL
7189 is accepted if it is valid and either the round number in the PROPOSAL \geq *lockedRound*
7190 or the PROPOSAL matches their *lockedValue*. Otherwise, they send a PREVOTE with *nil*
7191 value, which they also do if there is a timeout and they have yet to send a PREVOTE for
7192 the current round. This is *timeoutPropose*, which is triggered when a new round begins.

7193 After receiving the $2f + 1$ PREVOTES on a valid PROPOSAL, a validator "locks" on that
7194 value (they set *lockedValue* and *lockedRound* before sending a PRECOMMIT if still in
7195 the prevote step and set *validValue* and *validRound* if in the prevote or precommit steps).
7196 Otherwise they send a PRECOMMIT to *nil*, which is also sent if a timer expires without
7197 having sent a PRECOMMIT in their current round. This is *timeoutPrevote*, and its timer
7198 begins when an honest validator sends a PREVOTE or receives $2f + 1$ PREVOTES in a
7199 round. The block is committed if the replica sees a valid PROPOSAL and $2f + 1$ matching
7200 PRECOMMITs, and validators move on to the next height of the chain. At any point, if a
7201 validator sees $f + 1$ messages of the same kind for a round greater than the local round, the
7202 validator advances to that round's proposal step.

7203 Processes "unlock" a block only when the block is committed or when seeing $2f + 1$ PRE-

7204 VOTES on a conflicting block (note that a vulnerability, discussed below, existed at this
7205 step until corrected in [359]). When locked on a value, an honest validator only sends
7206 messages for that value. There is also a timer to prevent blocking on the final precommit
7207 step, *timeoutPrecommit*, where the timer begins upon receiving any $2f + 1$ PRECOMMIT
7208 messages in the current round, regardless of whether the values conflict. If it expires, they
7209 move to the next round.

7210 The original Tendermint protocol had a potential safety violation during the unlocking
7211 process. Say that a validator is locked on a block B . If that validator sees $2f + 1$ support for
7212 block B' , they set *lockedValue* = *nil*. However, the validator must ensure that $B \neq B'$ at this
7213 step so that the following scenario does not occur: the validator locks on B in round r and
7214 then sees $2f + 1$ support for B again in round $r' > r$, which then unlocks B but fails to lock
7215 it again. In this case, some validators may commit to B , while others commit to $B' \neq B$.

7216 Tendermint is unable to have light clients work the "standard" way that clients would work
7217 for PBFT. Clients do not query validators directly and, thus, cannot check for $f + 1$ match-
7218 ing answers. Furthermore, the proof-of-stake validator set is constantly changing, and a
7219 light client needs an accurate view of this set in order to check the validity of signatures on
7220 blocks. An approach to creating Tendermint light clients was described in [361].

7221 A variant of the Tendermint algorithm, Tenderbake, was designed to be secure even with
7222 bounded message buffers [138]. As it stands, validators need to store messages for a po-
7223 tentially unbounded number of rounds. Tenderbake has the additional advantage of termi-
7224 nating more quickly than Tendermint, but it loses the optimistic responsiveness that comes
7225 with partial synchrony.

7226 13.4.2. Algorand

7227 Algorand uses proof of stake and *cryptographic sortition* (using a VRF) to elect a com-
7228 mittee of validators who then execute a novel single-shot Byzantine agreement algorithm
7229 called BA^* [362]. A primary benefit of the algorithm is that forks are (effectively) impossi-
7230 ble and, thus, transactions are finalized (almost) immediately upon inclusion in a block. Al-
7231 gorand's liveness relies on a "strong synchrony" assumption, where most honest users will
7232 see the messages of most other honest users within a known time bound. Safety requires
7233 "weak synchrony," where the network can be asynchronous for a lengthy but bounded in-
7234 terval, but must then be followed by a sufficiently long synchronous period (confusingly,
7235 this is distinct from the weakly synchronous model introduced in Section 1.5). A variant of
7236 the Algorand protocol can also be modified to be secure in the sleepy model, though with
7237 a performance penalty [363].

7238 Algorand can scale to a large number of validators by running each step of BA^* with a
7239 committee of only a subset of validators large enough to ensure that no selected committee
7240 would exceed the security threshold over the lifetime of the system (with high probabili-
7241 ty). Messages from the BA^* execution are gossiped across the network instead of sent

7242 point-to-point between elected committee members. The BA \star algorithm does not require
7243 participants to maintain any private state other than their private keys, and participants are
7244 expected to only send a single message, which allows them to be immediately replaced
7245 after sending their message. Each step of BA \star has a new set of committee members. Com-
7246 bined with the VRF for sortition, this makes the system adaptively secure and resistant to
7247 targeted denial of service on committee members.

7248 The sortition process provides each selected member with a *priority*, which can be com-
7249 pared between participants. Since multiple validators may be selected to propose blocks,
7250 priority determines which one should be adopted. Validators initialize BA \star with the highest
7251 priority block they have seen, and BA \star then executes in repeated steps. Each step begins
7252 with sortition to determine the committee members, and elected members broadcast a proof
7253 of their selection over the network. These steps are then repeated until there are enough
7254 participants in the committee to reach consensus. Each validator checks whether they were
7255 elected for the next step as soon as the previous step ends.

7256 The sortition process uses a *role* parameter that specifies the particular roles that a partic-
7257 ipant may play in the consensus process (e.g., block producer, committee member in step
7258 two, etc.). For any given role, an expected number τ of participants are selected. A single
7259 user may be elected as multiple sub-users for a given role based on their stake (i.e., an en-
7260 tity with a large stake may count as multiple members of a single committee). Let *seed* be
7261 a random seed to be discussed shortly, w be the number of atomic units of cryptocurrency
7262 owned by a particular user, W be the total stake in the system, $p = \frac{\tau}{W}$ be the probability of
7263 a particular unit of stake being selected, and the probability that exactly k of the user's w
7264 stake units are selected is binomially distributed and denoted as $B(k; w, p)$. Users compute
7265 $(y, \pi) = \text{VRF}_{sk}(seed || role)$ and then use the pseudorandom value y to determine how many
7266 sub-users they control that were elected. To do so, the interval $[0, 1)$ is divided into subin-
7267 tervals $I^j = [\sum_{k=0}^j B(k; w, p), \sum_{k=0}^{j+1} B(k; w, p))$ for $j \in \{0, 1, \dots, w\}$. If y_{len} is the length of the
7268 output of the VRF, then a user has j sub-users selected if the value $\frac{y}{2^{y_{len}}} \in I^j$. The priority
7269 of a block produced by sub-user i for $i \in \{1, \dots, j\}$ is equal to $H(y || i)$.

7270 In Algorand, a new block is appended to the blockchain in each round, and each round r has
7271 a corresponding random seed, $seed_r$. When selected to be a block proposer in round $r - 1$,
7272 validator i also proposes a seed to be used for round r as $(seed_r, \pi) = \text{VRF}_{sk_i}(seed_{r-1} || r)$
7273 and includes the seed in their proposed block. Once there is agreement on a block for round
7274 $r - 1$, all validators will agree on the seed to be used for round r . If a proposed block does
7275 not contain a valid seed (or the block is otherwise invalid), participants treat the block as
7276 though it were empty and compute $seed_r = H(seed_{r-1} || r)$ for some hash function H .

7277 To be secure, sk_i must be chosen sufficiently far in advance of its use in determining a seed,
7278 such that the seed will remain pseudorandom even if i is malicious. To prevent the adversary
7279 from manipulating the sortition process, the seed used for sortition itself is changed only
7280 every R rounds, such that the round r seed used in sortition is actually $seed_{r-1-(r \bmod R)}$.
7281 For some timing parameter b , round r 's sortition uses the sk_i 's and their associated stake

7282 from the last block created b -time before block $r - 1 - (r \bmod R)$. Security requires that
7283 there be at least one honestly produced block in this b -timed interval. By keeping b large,
7284 network adversaries are less likely to be able to keep the network partitioned long enough
7285 to produce all empty blocks and control the seed. On the other hand, a longer b leads to
7286 higher stake shift and correspondingly lowers the security margin.

7287 Algorand selects up to $\tau_{proposer}$ proposers via sortition, so multiple blocks may be produced
7288 in a round. Validators wait for some short period of time (on the order of 10 seconds) to
7289 see blocks in a round before beginning BA^* – keeping the highest priority block they have
7290 seen in a round – and ignore blocks with a lower priority. When the time is up, they start
7291 BA^* with the highest priority block as input or an empty block if none were seen.

7292 BA^* has two phases, the first of which has two steps, and the second of which has at least
7293 two steps. In the first phase, one of the proposed blocks is chosen to compete against an
7294 empty block in the second phase, which runs binary Byzantine agreement. Each step has
7295 an expected number of participants τ_{step} , except for the final step which has an expected
7296 τ_{final} . The voting threshold at each step is then either $T * \tau_{step}$ or $T * \tau_{final}$, where $T > \frac{2}{3}$ is
7297 the same at all steps except potentially the final step.

7298 The *CommitteeVote()* procedure runs sortition for the "committee" role of a given round
7299 and step and then votes on a block hash taken as input. The *CountVotes()* procedure
7300 takes the votes for the current round and step and outputs the block hash of the first
7301 proposed block to gather the appropriate threshold of votes. At each step, validators run
7302 *CommitteeVote()* on some block hash, and everyone waits for a specified amount of time
7303 to count votes before timing out if enough votes are not accumulated.

7304 In the first step of the first phase, committee members vote for the block hash that BA^* was
7305 initially passed. In the second step, members vote for either a block hash that received at
7306 least $T * \tau_{step}$ or an empty block if the voting threshold is not reached. At the end of this
7307 phase, honest participants will have no more than a single non-empty block to consider in
7308 the next phase. An empty block is more likely to be chosen if the highest-priority block
7309 proposer was malicious and equivocated or if the network is not synchronous.

7310 In the second phase, a binary BA algorithm is executed in order to have validators agree
7311 on a choice between an empty block and the hopefully non-empty block output from the
7312 first phase. Should a validator receive the threshold of votes needed for a block hash in
7313 any step of this phase, they will vote for that block hash if they are elected as a committee
7314 member for the next step. On the other hand, the network may not be synchronous, or
7315 an adversary may allow a particular validator to see enough votes to decide on a block
7316 while preventing other validators from seeing enough votes before they time out. In this
7317 case, the algorithm must ensure that votes in the next step do not result in a different block
7318 being decided by different validators, so members' next step votes are of a specific value
7319 that could have been returned in a given step, and each step can only return one particular
7320 value. In addition, whenever a participant returns a block from the protocol, they continue
7321 to vote for it for the following three steps in order to help that value gain enough votes in

7322 future steps. Specifically, given a parameter $MAXSTEPS$, the binary Byzantine agreement
7323 protocol works as follows while $step < MAXSTEPS$:

7324 1. Check for agreement on $block_hash$:

7325 (a) $CommitteeVote()$

7326 (b) $v \leftarrow CountVotes()$

7327 (c) If $v = TIMEOUT$, then $v \leftarrow block_hash$

7328 (d) Else if $v \neq empty_hash$:

7329 i. For $step < s' \leq step + 3$, call $CommitteeVote()$

7330 ii. If $step = 1$, call $CommitteeVote(FINAL)$

7331 iii. Return v

7332 (e) $step++$

7333 2. Check for agreement on $empty_hash$:

7334 (a) $CommitteeVote()$

7335 (b) $v \leftarrow CountVotes()$

7336 (c) If $v = TIMEOUT$, then $v \leftarrow empty_hash$

7337 (d) Else if $v = empty_hash$:

7338 i. For $step < s' \leq step + 3$, call $CommitteeVote()$

7339 ii. Return v

7340 (e) $step++$

7341 3. Use a common coin:

7342 (a) $CommitteeVote()$

7343 (b) $v \leftarrow CountVotes()$

7344 (c) If $v = TIMEOUT$ and $CommonCoin() = 0$, then $v = block_hash$

7345 (d) If $v = TIMEOUT$ and $CommonCoin() = 1$, then $v = empty_hash$

7346 (e) $step++$

7347 With an honest block proposer and a synchronous network, most committee members will
7348 begin with the same block and return on the first step. However, during a network partition,
7349 it is possible for honest users to return different blocks in the protocol as described so far.

7350 Algorand solves this through the notions of *tentative consensus* and *final consensus*. After
7351 the binary BA algorithm returns, there is one last voting step in BA \star in which votes are
7352 counted and a block is decided if validators see $T * \tau_{final}$ votes for it. Final consensus
7353 exists if enough validators return at the first step and then see enough votes from other
7354 committee members demonstrating that they did too. If the highest priority block proposer
7355 was honest and the network is synchronous, then final consensus is achieved after four
7356 interactive steps. A particularly lucky adversary may cause this to take an expected 13
7357 steps, which requires the adversary to be the highest priority proposer for the round and to
7358 control a large fraction of committee members at each step.

7359 An adversary who controls the network can prevent a subset of validators from reaching
7360 either final or tentative consensus for an arbitrary number of steps, and each step increases
7361 the adversary's chance of getting consensus on an empty block. To prevent this, BA \star has a
7362 bounded number of steps, *MAXSTEPS*, before halting and requiring a recovery procedure.
7363 For an attacker to keep the validators from agreeing, they need to know how a participant
7364 will vote after they have a timeout during vote counting. The *CommonCoin()* procedure
7365 makes this more difficult. Each user sets their common coin to be the least significant bit
7366 of the lowest committee member hash (produced from the VRF) seen during the step.

7367 When consensus on a block is only tentative, it is possible for multiple forks to exist, which
7368 can inhibit liveness. To remedy this, users passively monitor all BA \star votes to keep track
7369 of all forks and periodically use sortition to propose one of the forks for finalization. If a
7370 validator is chosen via sortition, they propose an empty block that extends their preferred
7371 fork of the chain. Other validators then wait for the highest priority proposal, check that it
7372 extends their longest local chain, and invoke BA \star on that proposed block.

7373 Finally, Algorand's design does not require checkpointing or having parties regularly be
7374 online. The BA \star process results in a *certificate* that can be used to prove consensus on a
7375 block. Specifically, a certificate is composed of the votes from the last step of the binary
7376 BA algorithm (but not the "final" step that comes after).

7377 14. Hybrid and Alternative Sybil-Resistance Mechanisms

7378 14.1. Proof of Space

7379 While proof of work and proof of stake get most of the attention when it comes to Sybil-
7380 resistance mechanisms in permissionless environments, another promising avenue is *proof*
7381 *of space* (PoSpace) [364]. In a proof-of-space system, users must prove that they are ded-
7382 icating disk space rather than computational resources through a challenge that requires a
7383 significant amount of memory or disk space to solve.

7384 The proof-of-space protocol introduced in [364] is an interactive protocol between a prover
7385 and a verifier that has an initialization phase and an execution phase. The initialization
7386 phase results in the prover storing some data and the verifier storing a short commitment to
7387 the data. In the execution phase, the verifier sends a challenge to the prover, who responds

7388 with an answer that requires reading some portion of the original data stored on disk. In a
7389 consensus setting, there is no designated verifier, so anyone with a view of the public ledger
7390 must be able to verify the proof.

7391 The security ramifications of proof of space may be similar to that of ASIC-resistant proof
7392 of work in some ways (see Section 9.1.2). Because hard drives are widely available and
7393 distributed, proof of space allows people with typical consumer-grade hardware to partic-
7394 ipate in consensus using their idle resources. Further, storage is never "used up," so hard
7395 drives used for proving storage can be used for useful storage later. Disk space can be
7396 repurposed and thus has high salvage value, like a CPU or GPU that can be used for other
7397 purposes when not participating in consensus. On the other hand, proof of space requires
7398 few ongoing costs, unlike proof of work. The cost of participation is the upfront cost of
7399 acquiring storage plus the opportunity costs of using that storage to participate in consensus
7400 instead of doing something else with it. Stated differently, the vast majority of the expenses
7401 incurred from proof of space are capital expenditures rather than operational expenditures.
7402 In this respect, proof of space is more similar to proof of stake.

7403 14.1.1. Spacemint

7404 Spacemint is an academic proof of concept of a cryptocurrency that uses proof of space
7405 instead of proof of work [365]. Adapting proof of space to the cryptocurrency setting re-
7406 quires addressing several challenges, including two that are familiar from proof of stake:
7407 grinding attacks and costless simulation to mine multiple chains simultaneously (see Sec-
7408 tion 12.1.1).

7409 In a grinding attack, an adversary who mines block B_i can (cheaply) try out multiple block
7410 hashes and use one that gives them an advantage in constructing block B_{i+1} . Spacemint
7411 decouples the hash chain that includes the proofs of space from a separate signature chain
7412 for transactions, which prevents adversaries from manipulating the block hash based on the
7413 included transactions. The transactions must be bound to the proof chain, so the purpose of
7414 the signature chain is to prevent past transactions from being altered when new blocks are
7415 added. This architecture is displayed in Figure 42.

7416 Spacemint addresses the costless simulation issue by defining a proof's *quality* and making
7417 it fixed for a given time step in order to eliminate the benefit of trying many chains. The
7418 probability that a proof has the best quality should be equal to that miner's fraction of the
7419 total space being used to mine the chain. Miners who publish multiple conflicting blocks
7420 can then be punished by having their reward forfeit and giving half of it to the miner who
7421 includes proof of equivocation. Because the proof-of-space miner does not have stake
7422 that can be taken, this punishment is limited and less effective than it is in proof-of-stake
7423 systems. The penalty only deters attacks like selfish mining, where the potential gain may
7424 be less than the potential penalty, but is unlikely to discourage double-spending attacks.

7425 The proof-of-space challenge for block B_i is the hash of block $B_{i-\delta}$, where δ is parameter-




Fig. 42. Spacemint grinding defense. The proof chain is denoted PC, and the signature chain is denoted SC. If an honest miner mines the i -th block and does not equivocate, then past transactions cannot be changed. If an adversary wanted to change transactions in the j -th block, $j < i$, while maintaining the same proof chain, they would need to compute new signatures for all of the blocks in between, which requires the corresponding secret keys. [365]

7426 ized such that $B_{i-\delta}$ is a block that was mined a short while in the past. However, δ cannot
7427 be set arbitrarily high because miners can precompute their answers for δ time steps and
7428 then not need to access their storage for some time. This would potentially allow a miner to
7429 reuse the same storage space for multiple space commitments because they could perform
7430 the initialization procedure multiple times. Therefore, a related security requirement is that
7431 initialization must be time consuming (at least δ blocks' worth of time).

7432 A final consideration in Spacemint is that miners must commit to the storage space they
7433 intend to mine with in advance. This requires a special space commitment transaction
7434 that specifies the miner's public key and their space commitment used for verification. If
7435 miners were not required to commit this in advance on the blockchain, then they would be
7436 able to reuse the same space for multiple commitments due to properties of the proof-of-
7437 space scheme from [364]. An unfortunate side-effect of this is that if an adversary acquired
7438 the majority of the storage space used to secure the chain, they could censor new space
7439 commitment transactions and maintain their majority indefinitely.

7440 14.1.2. Chia

7441 The Chia network consensus algorithm is inspired by Spacemint, but adds a verifiable
7442 delay function (VDF) and fixes some of its weaknesses. The description in this section
7443 corresponds to the original Chia "Green Paper" [366], but the Chia network as deployed
7444 may differ slightly from the presentation here.

7445 Chia addresses grinding attacks by combining two interlinked chains in a manner similar
7446 to Spacemint; an ungrindable *trunk* contains the proof-of-space and VDF output, while the
7447 *foliage* chain contains transactions, timestamps, and signatures that bind the foliage to the
7448 trunk. The proof-of-space challenge is derived from the last VDF output. Specifically, the
7449 challenge is the hash of a BLS signature over the VDF output, which is unpredictable to an
7450 adversary. The uniqueness of BLS signatures, where there is only one valid signature for a
7451 given public key and message pair, prevents grinding over different challenges in the trunk.
7452 This architecture is shown in Figure 43.




Fig. 43. Chia design. A block $\beta = (\sigma, \tau)$ – where σ is a proof of space and τ is a VDF output – is in the ungrindable trunk chain. The foliage blocks α_j contain transactions, a signature on the previous foliage block (to create a blockchain linked by signatures), and a signature on the proof of space (to bind the foliage to the trunk). Because the challenges for the proof of space and VDF come from previous trunk values, grinding on various values of data in the foliage does not give the attacker an advantage. [366]

7453 Chia uses a proof-of-space algorithm where the initialization procedure is non-interactive,
7454 unlike the one from [364] used in Spacemint. This removes the need to pre-register the
7455 storage commitment in a special transaction, which Spacemint uses to prevent grinding
7456 attacks.

7457 Chia alternates between proofs of space and VDF executions, where the VDF is used to
7458 protect against long-range attacks and remove the need for synchronized clocks. In Chia
7459 parlance, the entities that compute the proofs of space are called *farmers*, and the ones
7460 that compute the VDF are called *time lords*. In a proof-of-space system, unlike proof of
7461 work or proof of stake, any farmer can generate a valid proof during any time slot, so
7462 a mechanism is required to determine which block would "win." To prevent unbounded
7463 bandwidth consumption, proofs that are unlikely to win are disallowed in the first place.
7464 Chia adopts the idea of assigning a quality to each proof of space (proportional to the
7465 amount of space used to construct it) and agreeing on the best quality one as the legitimate
7466 chain extension, as was done in Spacemint. Here, quality is the hash of the proof, and the
7467 best quality proof is the one with the lowest valued hash. This mechanism is augmented
7468 by requiring a VDF output in order to consider the block valid. The VDF is parameterized
7469 such that the time it takes the time lords to compute it is linear in the quality of the proof of
7470 space (specifically, the VDF time parameter is the quality times the current difficulty level),
7471 which leads to having the best quality proofs completed first, as desired.

7472 The use of the VDF opens up the possibility for another subtle grinding vector that could
7473 be used to launch double-spending attacks, which is addressed by slightly modifying the
7474 difficulty adjustment from the one used by Bitcoin. The time parameter for the VDF de-
7475 pends on the difficulty level, which is computed from the timestamps located in the foliage.
7476 Grinding on the timestamp in the foliage can lead to grinding against the VDF output in the
7477 trunk if the new difficulty kicks in immediately following the block whose timestamp was
7478 used to recompute the difficulty, as is done in Bitcoin. This threat is prevented by having
7479 the difficulty adjust only after several additional blocks have been constructed. An example
7480 of this grinding attack is shown in Figure 44.

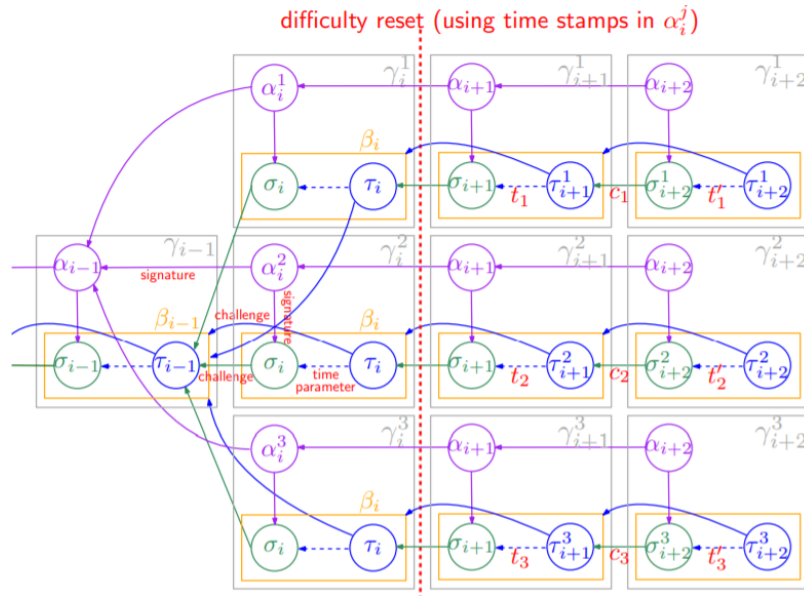


Fig. 44. Chia difficulty grinding. The foliage blocks α_i^1 , α_i^2 , and α_i^3 each have different timestamps. Because they occur right before the difficulty reset at block i , each of the three chains has a slightly different difficulty parameter to be used for blocks at height $i+1$ or greater. The adversary can run three VDF executions in parallel, and there will be three different outputs at height $i+1$, causing each chain to have a different proof-of-space challenge, c_1 , c_2 , or c_3 for blocks at height $i+2$. [366]

7481 Chia does not follow Spacemint in using punishment to discourage costless simulation
 7482 attacks where farmers extend multiple chains at once. Instead, it embraces this "double-
 7483 dipping" and allows honest farmers to engage in the practice as well. If honest farmers
 7484 do not double-dip, then at least $\approx 73.1\%$ of the total allocated space must be supplied by
 7485 honest farmers. Chia improves this bound by adding a local security convention where
 7486 honest farmers extend the first $\kappa > 1$ chains at every chain depth. It is recommended that
 7487 $\kappa = 3$, which reduces the security margin to requiring $\approx 61.5\%$ of space to be supplied by
 7488 honest farmers. A higher value of κ improves the security margin (Chia would be secure
 7489 with an honest majority if farmers were required to extend every block) but requires that
 7490 honest farmers compute a factor of κ more proofs and makes agreement take longer.

7491 While the chain quality and chain growth properties were demonstrated in the original Chia
 7492 paper, the crucial common prefix property was not proven until [225]. Chia's security anal-
 7493 ysis differs from that of chain-based proof of stake, despite both suffering from a variant
 7494 of the nothing-at-stake problem. This difference arises because in, say, Ouroboros Praos
 7495 and Snow White, the same randomness is used across multiple blocks, whereas Chia's
 7496 randomness is independent at each block. This gives an adversary in Chia an independ-
 7497 ent opportunity for winning leader election at every block, which magnifies the attacker's
 7498 power by a factor of e .

7499 If one were to naively adopt an equivalent of a longest chain rule and apply it to a proof-
7500 of-space system, one might have honest farmers agree on the chain where the sum of the
7501 qualities of the included proofs is highest. This opens up the possibility of long-range
7502 attacks because it is possible to generate a heavier chain as long as the adversary has enough
7503 space to beat the *average* of the space used across the entire honest chain. If the value of
7504 farming Chia blocks increased significantly or the cost of storage decreased significantly
7505 after the chain began, beating the average might be much easier than overpowering the
7506 amount of space currently allocated to Chia farming. Spacemint mitigates this by only
7507 considering the quality of more recent blocks when comparing the quality of chains. This
7508 is not a problem for Chia because the VDF prevents an adversary who forks the chain far
7509 in the past from catching up to the honest chain.

7510 14.2. Proof of Activity

7511 Proof of activity (PoA) is a hybrid proof-of-work and proof-of-stake system intended to
7512 diversify the types of entities securing the blockchain, particularly as transaction fees be-
7513 come a more important component of the block reward [367]. Stakers pick transactions to
7514 include in blocks, which makes it harder for miners to censor transactions using feather-
7515 forking or majority attacks. PoA requires that a majority of the online stake is honest in
7516 order to maintain security.

7517 Leader election is performed using the "follow-the-satoshi" idea, similar to how it is done
7518 in the Chains of Activity protocol (Section 13.1.1). A pseudorandom seed is used to select
7519 uniformly from the total set of *satoshis* (atomic cryptocurrency units) based on when it
7520 was minted. That satoshi is then traced through the transaction graph to the public key
7521 that currently controls it, which determines the leader. Note that this procedure can cause
7522 some challenges when using more advanced scripting capabilities than simple public key
7523 to public key financial transfers.

7524 The PoA protocol begins by having miners mine a block header that does not include any
7525 transactions and broadcasting the empty block to the network upon success. The hash
7526 of this block header is then used as the seed to deterministically derive N pseudorandom
7527 stakeholders (with $N = 3$ suggested). The follow-the-satoshi procedure is invoked N times
7528 using $H(\text{header_hash}||\text{previous_block_hash}||i)$ for $i \in \{1, \dots, N\}$. All online stakeholders
7529 check whether they are one of the N selected stakeholders for that block. The first $N - 1$
7530 stakeholders sign the hash of the empty header and broadcast it to the network. The N -
7531 th stakeholder creates and broadcasts a *wrapped block* that extends the empty header by
7532 including transactions, the $N - 1$ other signatures, and their own signature over all of this
7533 information. Nodes continue to follow the longest chain rule, as in Nakamoto Consensus.
7534 The transaction fees are shared between the block's miner and the N stakeholders.

7535 If not all of the N stakeholders are online, another miner will mine another header at the
7536 same height of the chain, which will derive a different N stakeholders. The use of proof of
7537 work limits the ability of an attacker to bias the randomness used for the stake election. To

7538 incentivize a certain level of online stake, the N -th stakeholder can include in their block
7539 all of the other empty blocks mined at the same height but without the required signatures.
7540 Nodes can count how many empty blocks were mined during a difficulty retargeting win-
7541 dow. If too many were mined, then stakeholders can get a higher fraction of the block
7542 reward during the next window, and vice versa.

7543 By combining PoW and PoS, an adversary may need to acquire both a significant amount
7544 of computational power *and* stake to pull off attacks. Specifically, if the follow-the-satoshi
7545 procedure is a random oracle, then an attacker with α fraction of the online stake needs
7546 more than $(\frac{1}{\alpha} - 1)^N$ times the hash power of honest miners in order to perform typical
7547 majority attacks, like censorship and double-spending. An implication of this is that an
7548 adversary with the majority of the online stake will actually have their computational power
7549 magnified substantially. On the other hand, with $N = 3$ and only one-third of the online
7550 stake, an adversary requires eight times more computational power than the honest nodes.
7551 The inclusion of proof of stake increases the susceptibility of PoA to bribery attacks, which
7552 become more severe as N increases. An attacker may be able to bribe stakeholders to
7553 withhold their signatures and only needs to bribe a single stakeholder per block.

7554 14.3. Checkpoints and Finality Gadgets

7555 The *CAP theorem* states that a distributed system can only maintain consistency or avail-
7556 ability in the face of a network partition but not both. Many of the protocols discussed
7557 in this document, most notably Nakamoto Consensus, favor availability over consistency.
7558 When the network of miners is partitioned, transactions can still be put into blocks on both
7559 sides of the partition (availability), but these transactions may conflict and need to be recon-
7560 ciled when the partition ends. Eventually, the conflict is resolved, so Nakamoto Consensus
7561 has *eventual consistency*. However, even this notion of eventual consistency leaves a small
7562 possibility that a block may be reverted in the future.

7563 Contrast this with the idea of *finality*. Permissioned systems, some committee-based proof-
7564 of-work systems, and BFT-based proof of stake tend to have finality, where a finalized
7565 block can never be reverted under any circumstances. As a trade-off, this means that the
7566 system halts during a network partition. Finality is a desirable property because any ap-
7567 plication that relies on the underlying state machine can truly treat a finalized transaction
7568 as final rather than needing to be prepared to revert its execution. Furthermore, it is not
7569 too difficult to notice when the system halts for an extended period, so manual intervention
7570 may be able to fix the availability issues fairly quickly. This section discusses protocols for
7571 adding finality to permissionless blockchain systems, including those that otherwise favor
7572 availability.

7573 14.3.1. Ad Hoc Finality Layers and Reorg Protection

7574 The earliest form of ad hoc checkpointing were the hard-coded checkpoints that Satoshi
7575 included in early versions of the Bitcoin software. Typically, during a software release,

7576 Satoshi would pick a block hash from a few months earlier, embed that hash into the node
7577 software, and enforce a rule that the node will never accept a chain that does not include the
7578 hard-coded checkpoints. While this technically provided finality for transactions through
7579 the latest checkpoint block, it was generally too far in the past to be of practical use. The
7580 real benefit of hard-coded checkpoints in early Bitcoin was to prevent denial-of-service
7581 attacks that stemmed from the very low difficulty of mining blocks at that time. A modestly
7582 resourced attacker could generate long chains of low difficulty blocks to fill up a node's
7583 disk or prevent them from being able to synchronize the chain. This threat was later solved
7584 when the Bitcoin software was updated to perform *headers-first synchronization*, where the
7585 node verifies a complete chain of headers and their proofs of work before downloading the
7586 contents of the blocks themselves.

7587 Hard-coded checkpoints like this provide finality in a way that is easy to reason about. It
7588 also grants software developers considerable centralized power. While this is not much
7589 of an issue if the checkpoints are from the distant past, as in early Bitcoin, hard-coded
7590 checkpoints are only practically useful for finality if issued for the recent past. In this
7591 case, it may allow developers to unfairly determine which chain might be preferred as the
7592 canonical one.

7593 Other proof-of-work cryptocurrencies have added more sophisticated forms of finality, or
7594 *reorg protection*, primarily to mitigate the threat of majority hash rate attacks. For example,
7595 several of the more popular Bitcoin Cash clients include a default rule that they will not
7596 perform a reorg of more than 10 blocks at any time, even if the alternative chain has greater
7597 proof of work. The justification for this extra rule was that, because only a small minority
7598 of the available double-SHA-256 computational power was mining on Bitcoin Cash, the
7599 network was highly susceptible to a 51% attack, and cryptocurrency exchanges wanted to
7600 be protected from this risk. If the exchange ran a client enforcing this rule, then they could
7601 accept deposits after 10 confirmations knowing that even a well-resourced attacker could
7602 not revert. As a side benefit, checkpoints also reduce the profitability of selfish mining.

7603 That said, these "moving checkpoints" are not without risk. Most obviously, this rule cre-
7604 ates a race condition that can cause accidental chain splits. If the network is partitioned
7605 for a few hours, or if an adversary mines a 10-block private chain and publishes it at the
7606 right time, there will be a permanent chain split that requires manual intervention to fix. In
7607 addition, the new rule removes one of the biggest advantages of Nakamoto Consensus by
7608 making the system weakly subjective (see Section 12.1.2), such that a bootstrapping node
7609 (or one that has been offline for a few hours) can no longer be assured that the most-work
7610 chain is the correct one and must acquire it from a trusted source.

7611 In addition, the finality benefits may be illusory. For instance, an exchange that is temporar-
7612 ily partitioned from the network may have finality for a customer deposit on the particular
7613 chain that the exchange is following, but if the network splits, *the finalized chain might*
7614 *not be the one that the exchange wants to follow, particularly if most of the network is*
7615 *following a different chain.* In this case, an exchange that would have quickly rejoined the

7616 rest of the network at the conclusion of the attack instead has a false sense of security and
7617 must manually intervene in order to join the rest of the network while still suffering from a
7618 double-spend attack.

7619 Reorg protection protocols of this nature were formalized in [368], where an additional
7620 "front-running" attack was proposed that prevents liveness when there is a rushing ad-
7621 versary with a hash rate majority. The adversary mines their own private chain without
7622 adopting any honest blocks. Whenever a new block is published that extends the longest
7623 honest chain, the adversary publishes a block while keeping the remainder of their private
7624 chain hidden. Because the adversary has the majority of the computational power, they
7625 can counter every block, and because they are rushing, adversarial blocks will always win.
7626 In this case, every block on the canonical chain will be adversarial. While the intentions
7627 of reorg protection assume that there is an adversarial hash rate majority, it is not trivial
7628 for an adversary to be rushing on a permissionless network, so this attack is not easy to
7629 pull off. The checkpointing mechanisms presented in [368] resist this attack by operating
7630 as an unpredictable randomness beacon that "refreshes" the execution at each checkpoint,
7631 preventing the adversary from using blocks mined before the checkpoint was issued so that
7632 they do not maintain a persistent advantage.

7633 Another interesting implementation of reorg protection is the *ChainLocks* system deployed
7634 on the Dash cryptocurrency network [369]. Dash was a pure proof-of-work Nakamoto
7635 Consensus cryptocurrency until ChainLocks was added, making it a hybrid with proof
7636 of stake. Any Dash user with 1000 Dash can become a *masternode* and provide extra
7637 services to the network in exchange for some of the block reward. For each block that is
7638 published, a set of a few hundred masternodes is selected, which attempt to finalize the
7639 block by issuing a ChainLock on it, where a ChainLock is a BLS threshold signature. Each
7640 selected masternode submits a signature share for the first block that they see at each chain
7641 height, and if enough of these masternodes sign the same block, the threshold signature
7642 is created and broadcast to the rest of the network as a ChainLock. Any node that has
7643 seen a ChainLock for a given block will not reorg past that block. Under normal network
7644 conditions, this will likely result in blocks being finalized almost immediately and make
7645 it more difficult for an adversary to cause a chain split. This also gives the masternodes
7646 a significant amount of power that can be wielded to an adversary's advantage, but the
7647 adversary would need to own a very large amount of Dash (well over half) or be able to
7648 bribe other masternode owners in order to take advantage of it.

7649 14.3.2. Casper the Friendly Finality Gadget (FFG)

7650 The Ethereum network began as a pure proof-of-work cryptocurrency that followed the
7651 GHOST fork-choice rule but is currently undergoing a multi-phase process to become a
7652 pure proof-of-stake network colloquially called Ethereum 2.0, as described in Section 13.2.
7653 One step of the process is to adopt a hybrid approach of maintaining the underlying proof-
7654 of-work chain and augmenting it with a proof-of-stake finality gadget called Casper FFG

7655 [370, 371]. The intention is to combine Casper FFG and the GHOST fork-choice rule
7656 ("Gaspar"), as described in [350], which can work in a pure proof-of-stake system as well.

7657 Casper FFG uses any block proposal mechanism (e.g., leader election via proof of work)
7658 and uses proof of stake to finalize blocks, called checkpoints, such that a checkpoint can
7659 never be reverted. Checkpoints are created 100 blocks apart and form a checkpoint tree
7660 or chain. Casper provides *accountable safety* in that it is impossible for two conflicting
7661 checkpoints to be finalized unless more than $\frac{1}{3}$ of the deposited stake violates a *slashing*
7662 *condition* (accountability in BFT protocols is discussed in Section 7.2) and can thus be
7663 punished. Casper also provides *plausible liveness*, such that if $\geq \frac{2}{3}$ of the stake is honest, it
7664 is possible to continue finalizing new checkpoints without having any validators violate a
7665 slashing condition.

7666 To be entitled to vote in Casper FFG, a potential validator must put up a stake deposit
7667 as collateral to be slashed in case of misbehavior. A vote consists of a signed message
7668 containing hashes of *source* and *target* checkpoints, denoted s and t , as well as the heights
7669 of s and t in the checkpoint tree, $h(s)$ and $h(t)$. The target checkpoint must be a descendant
7670 of the source in the checkpoint tree. Casper uses the following definitions:

- 7671 • A *supermajority link* is an ordered pair of checkpoints $a \rightarrow b$ where at least $\frac{2}{3}$ of
7672 validators have cast votes with source a and target b . Supermajority links may skip
7673 over checkpoint blocks.
- 7674 • A checkpoint c is considered *justified* if it is the genesis block or if there exists a
7675 supermajority link $c' \rightarrow c$ where c' is justified.
- 7676 • A checkpoint c is considered *finalized* if it is the genesis block or if the following
7677 conditions hold: c is justified, there exists a supermajority link $c \rightarrow c'$, there is no
7678 conflict between c and c' , and $h(c') = h(c) + 1$.
- 7679 • The *dynasty* of a block b is the number of finalized checkpoints in the chain from the
7680 genesis block to the parent of block b .

7681 Justification and finalization are analogous to the first and second voting rounds in a typical
7682 BFT protocol (i.e., prepare and commit, respectively, in PBFT). Recall that Casper provides
7683 accountable safety, meaning that the only way for two conflicting checkpoints to become
7684 finalized is for at least $\frac{1}{3}$ of the deposited stake to violate a slashing condition. If a validator
7685 violates one of these conditions, an honest validator can prove it and submit a special
7686 transaction to the blockchain that forfeits the malicious validator's deposit and rewards the
7687 honest validator with some portion of it. A validator may not publish two distinct votes,
7688 $(s_1, t_1, h(s_1), h(t_1))$, and $(s_2, t_2, h(s_2), h(t_2))$ where either of the following is true:

- 7689 • $h(t_1) = h(t_2)$. That is, a validator may not vote for the same target height twice.
- 7690 • $h(s_1) < h(s_2) < h(t_2) < h(t_1)$. That is, a validator may not vote within the span of its
7691 other votes.

7692 Honest validators in Casper FFG will follow the chain that contains the justified checkpoint
7693 of the greatest height and never revert past a known finalized checkpoint. Where ties exist,
7694 chains are prioritized based on the underlying proof-of-work scheme, be it Nakamoto
7695 Consensus or GHOST.

7696 Slight adjustments must be made in order for Casper FFG to safely handle dynamic val-
7697 idator sets. To join the set of validators, a user sends a deposit message. If it is included
7698 in a block with dynasty d , then this validator joins the validator set at the first block with
7699 dynasty $d + 2$, and this starting dynasty is denoted as $DS(v)$. To exit, withdrawal mes-
7700 sages are handled similarly, with the ending dynasty denoted $DE(v)$. When $DE(v)$ begins,
7701 the deposit is locked for some period of time so that the exiting validator can be slashed
7702 for misbehavior. Define the *forward validator set* and *rear validator set* for dynasty d as
7703 $V_f(d) \equiv \{v : DS(v) \leq d < DE(v)\}$ and $V_r(d) \equiv \{v : DS(v) < d \leq DE(v)\}$, respectively.
7704 To handle these validator sets, Casper redefines a supermajority link (s, t) for a target in
7705 dynasty d such that at least $\frac{2}{3}$ of both $V_f(d)$ and $V_r(d)$ have published votes $s \rightarrow t$. The
7706 finalization of checkpoint c then has an additional requirement: the votes for the two super-
7707 majority links $c \rightarrow c'$ and the one that justifies c must be included in c' 's blockchain prior
7708 to the child of c' or before block $h(c') * 100 + 1$. This change prevents safety violations in
7709 the pathological case where two grandchild blocks of a finalized checkpoint have different
7710 dynasties and evidence of slashing condition violations are included in one chain but not
7711 the other. This situation is displayed in Figure 45a.

7712 A final component of Casper FFG is the *inactivity leak*, which is intended to maintain
7713 liveness even if more than a third of validators are partitioned from the network or crash
7714 at the same time. In this case, supermajority links can no longer be created, preventing
7715 the finalization of additional checkpoints. For this reason, when validators fail to vote for
7716 checkpoints, they are slowly penalized by having their deposit reduced until connected
7717 validators become a supermajority again. This mechanism, shown in Figure 45b, can result
7718 in conflicting finalized blocks and cause chain splits. In practice, this would require the
7719 network partition to last for about three weeks [371].

7720 There is a liveness attack against Casper FFG called the *bouncing attack*, which allows an
7721 attacker with less than $\frac{1}{3}$ of the total stake to prevent finalization by "bouncing" between
7722 justifying one side of a fork or another [372]. The attack is possible when there is a latest
7723 justified checkpoint C and *justifiable* checkpoint C' such that C' is from a later epoch than C
7724 and conflicts with C . A justifiable checkpoint is one in which the attacker has enough votes
7725 to justify it but has not published them. A potential fix is to only allow the latest justified
7726 checkpoint to change during the first third of an epoch, otherwise marking it as pending
7727 and reevaluating when the epoch ends. A comparable attack is also possible against the
7728 Gasper protocol that combines Casper FFG with GHOST [373, 374]. The bouncing attack
7729 is shown in Figure 45c.

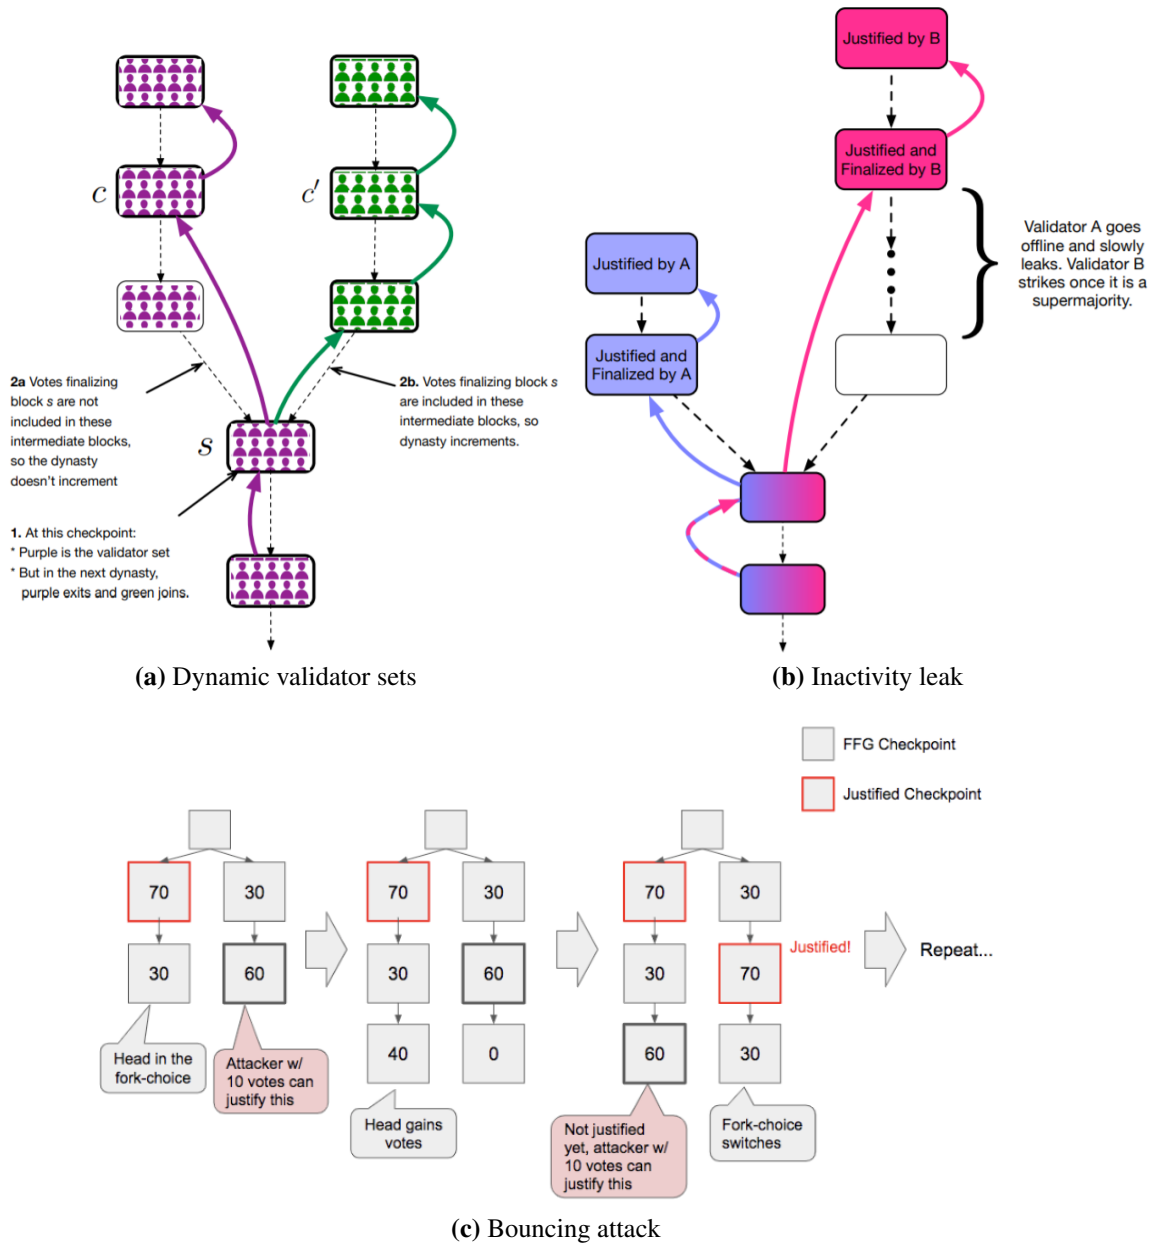


Fig. 45. Casper FFG attacks. (a) The validator sets finalizing checkpoints c and c' are disjoint, so no one gets slashed despite violating the slashing condition that c and c' are at the same height. (b) The checkpoint on the left can be finalized immediately, but a network partition prevents some validators from seeing the relevant votes. If the partition lasts for enough time, the stake deposits of those voters who support the left checkpoint will continuously deplete until a supermajority link can be formed on the right side. (c) Bouncing attack on Casper FFG. [370, 372]

7730 14.3.3. More Finality Gadgets and Checkpointing Protocols

7731 Several other checkpointing protocols have been proposed [373, 375–379], and this section
7732 compares a few at a high level.

7733 Winkle is a very different checkpointing mechanism that was designed to thwart long-
7734 range attacks in account-based proof-of-stake systems but can be used as a more general
7735 checkpointing mechanism as well [377]. In Winkle, clients issue stake-weighted votes for
7736 checkpoints with their transactions by including a hash of the most recent block they are
7737 aware of. A checkpoint is issued whenever at least $\frac{2}{3}$ of the total stake has voted in favor of
7738 a block. Based on actual transaction data, a checkpoint on Ethereum would take between
7739 50 days and a year to finalize; on Bitcoin, it would take between four months and three
7740 years. To speed up the checkpointing process, Winkle also allows stake to be delegated, in
7741 which case a new checkpoint could be issued every few hours or days.

7742 Most similar to Casper FFG is a finality gadget called GRANDPA, which is deployed on
7743 the Polkadot network [376]. The biggest difference is that participants in GRANDPA vote
7744 on the block of the greatest height they are aware of rather than a block at a predetermined
7745 chain height, and the vote transitively applies to all blocks preceding the one voted on.
7746 Roughly, the highest block with a supermajority of votes becomes finalized. It operates
7747 very similarly to PBFT, where there are two rounds of voting to finalize the chain and
7748 explicit timeouts to begin a new round. Because it is partially synchronous, finalization
7749 stops if there is a network partition, though blocks continue to be produced. Unlike with
7750 Casper FFG, there is no inactivity leak to continue finalization during a partition. After
7751 the partition ends, finalization only needs to happen once near the chain tip rather than for
7752 everything that came before. Like Casper FFG, GRANDPA provides accountable safety
7753 and allows for the ability to slash participants' stake for misbehaving.

7754 Finality layers for eventually consistent blockchains were formalized in [375], where the
7755 Afgjort finality layer was proposed. In the model provided, a finality layer must:

- 7756 • Form a chain of finalized blocks;
- 7757 • Have all parties agree on the finalized blocks;
- 7758 • Ensure that the last finalized block does not fall too far behind the last block of the
7759 underlying block proposal mechanism (i.e., the finalized chain should grow about as
7760 fast as the underlying chain); and
- 7761 • Require that all finalized blocks have at some point been on the chain accepted by at
7762 least k honest parties measured as a fraction of stake or computation, depending on
7763 the Sybil-resistance mechanism used for the underlying blockchain (k -support).

7764 The Afgjort protocol can speed up finality by an order of magnitude or so compared to
7765 eventually consistent algorithms like Nakamoto Consensus under good conditions and can
7766 "turn off" to avoid safety violations under bad conditions, such as an adversary trying to

7767 disrupt the protocol. In particular, blocks are declared final once they are in the common
7768 prefix of the underlying blockchain, but blocks are no longer finalized when there are be-
7769 tween $\frac{n}{3}$ and $\frac{n}{2}$ adversarial nodes, though the underlying blockchain remains live.

7770 Naively, a simple way to finalize a block at height d would be to run a typical Byzantine
7771 agreement algorithm and consider the agreed upon block as final. Unfortunately, the valid-
7772 ity property of most BA algorithms only guarantees that if every honest party starts with
7773 the same input, that is the agreed-upon value. However, if honest parties start with different
7774 inputs, then the agreed-upon value may be arbitrary. This implies that if the BA algorithm
7775 is executed before the common prefix of all honest parties includes height d , then an arbi-
7776 trary block absent from the chains of all honest parties may become finalized and violate
7777 the k -support property.

7778 Afgjort addresses this issue in a rather intuitive way: wait until the same block at height
7779 d is in the common prefix of all honest parties. When a member of the finalization com-
7780 mittee has a valid chain up to height $d + 1$, they vote on the block at height d in some
7781 BFT algorithm that requires unanimity to succeed, and the block is finalized if successful.
7782 Otherwise, committee members will continue attempting to finalize the block until they
7783 do succeed, where the i -th attempt occurs when they have a chain of height $d + 2^i$. This
7784 exponential backoff guarantees that if the underlying blockchain is secure, d is eventually
7785 in the common prefix.

7786 Of course, while all honest parties may have d in their common prefix, adversarial nodes
7787 can always vote to prevent unanimity. There are two Afgjort variants that solve this is-
7788 sue. The first is more efficient and secure but requires the additional security assumption
7789 of *bounded dishonest chain growth*, which states that a chain only adopted by dishonest
7790 parties grows more slowly than the chains of honest parties. This assumption will not hold
7791 if the underlying blockchain uses proof of work but has a difficulty adjustment algorithm
7792 that adjusts rapidly. If it does hold, however, Afgjort simply requires that participants vot-
7793 ing for the block at depth d also send the 2^i following block hashes in order to justify
7794 their vote. Participants then run a subprotocol, Freeze, that quickly settles on a uniquely
7795 justified block or \perp , followed by binary Byzantine agreement to agree on that block. The
7796 randomness in the binary agreement algorithm comes from a VRF. The protocol satisfies
7797 all desired properties above with $\frac{n}{3}$ -support as long as fewer than $\frac{n}{3}$ corrupted parties are in
7798 the finalization committee. The alternative protocol that does not rely on the bounded dis-
7799 honest chain growth assumption adds an extra filtering step to the beginning of the Freeze
7800 subprotocol that attempts to remove votes that came from dishonest parties. It does so by
7801 ignoring any votes that were supported by fewer than $f + 1$ committee members. As a
7802 consequence, this version only has 1-support, a much weaker property.

7803 Afgjort must also ensure that finalization does not fall too far behind the underlying chain's
7804 growth. To this end, block producers include a pointer to the most recently finalized block
7805 and committee member signatures in their block to attest to this finalization. When the
7806 chain grows too quickly, there will be blocks produced where the pointer differs from the

7807 actual most recently finalized block, and the next block to finalize can be adjusted ahead to
7808 account for this.

7809 Afgjort improves upon GRANDPA in a number of ways. GRANDPA only achieves 1-
7810 support as opposed to Afgjort's $\frac{n}{3}$ -support. GRANDPA also relies on a leader, which
7811 may impact liveness if the leader is corrupted or suffers from a denial of service. Fi-
7812 nally, GRANDPA includes explicit fixed timeouts, which prevent it from being responsive.
7813 Afgjort can finalize blocks based on the actual network delay rather than needing to wait
7814 for timeouts to occur.

7815 More recently, a family of so-called *snap-and-chat* protocols, inspired by Flexible BFT
7816 (Section 7.5), has been proposed [373, 378]. Snap-and-chat protocols output two ledgers
7817 instead of one: a dynamically available ledger that can remain secure with an unknown
7818 number of nodes that may go offline temporarily (such as a blockchain using Nakamoto
7819 Consensus) and a finalized ledger. When network conditions are poor, the dynamically
7820 available ledger is live but potentially unsafe, while the finalized ledger is safe but may not
7821 be live. When the network becomes synchronous again, the dynamically available ledger
7822 reconciles any inconsistencies while the finalized ledger catches up. Similar to Flexible
7823 BFT, snap-and-chat protocols support clients with differing beliefs: clients who priori-
7824 tize safety under network partitions can follow the finalized chain, while those who prefer
7825 availability can follow the other. When the network heals, all clients will agree on a sin-
7826 gle history. A major difference between Flexible BFT and snap-and-chat protocols is that
7827 Flexible BFT only guarantees consistency when clients have the same beliefs, whereas
7828 snap-and-chat maintains a common prefix regardless of client beliefs.

7829 Snap-and-chat protocols use an off-the-shelf dynamically available protocol, Π_{LC} (LC for
7830 longest chain), and an off-the-shelf partially synchronous BFT protocol, Π_{BFT} , and run
7831 the two subprotocols in parallel. If the underlying BFT protocol provides accountability,
7832 then snap-and-chat can provide accountable safety as well [374]. Transactions are input
7833 to Π_{LC} , which outputs a growing ledger LOG_{LC} . Nodes periodically take snapshots of
7834 LOG_{LC} , which are then input to Π_{BFT} , which spits out a growing ledger of these snap-
7835 shots, LOG_{BFT} , in an attempt to finalize some of the transactions. The finalized ledger,
7836 LOG_{fin} , is formed by concatenating the snapshots in LOG_{BFT} and then removing duplicate
7837 transactions. To create the dynamically available ledger, LOG_{da} , the finalized LOG_{fin} is
7838 prepended to LOG_{LC} and sanitized again. This sanitization breaks typical light clients, but
7839 a more complicated light client design is presented in [378].

7840 Clients who want availability follow the LOG_{da} ledger, not LOG_{LC} . These two ledgers are
7841 equivalent under favorable network conditions, but under less favorable conditions, LOG_{fin}
7842 and LOG_{LC} may diverge. In this case, by prepending LOG_{fin} to LOG_{LC} , the finalized
7843 ledger is guaranteed to be a prefix of LOG_{da} and remain safe so long as Π_{BFT} remains
7844 safe. This ensures that all clients eventually agree on a single history regardless of the
7845 chain they follow. A malicious party could attempt to break safety by using a fake snapshot
7846 with unconfirmed transactions as input to Π_{BFT} . To prevent this, Π_{BFT} requires a slight

7847 adjustment, where honest nodes refuse to accept the finalization of snapshots that the node
7848 does not believe are confirmed in their view of Π_{LC} .

7849 Around the same time that the snap-and-chat construction was proposed, [379] described
7850 a protocol with similar goals called the checkpointed longest chain. Like snap-and-chat,
7851 the checkpointed longest chain allows clients the flexibility to choose whether they prefer
7852 guaranteed finality or availability. Unlike snap-and-chat, however, the checkpointed longest
7853 chain construction provides *coupled validity*, wherein finalized blocks exist on a single lin-
7854 ear chain, so one can verify a transaction based on the information from the blockchain
7855 leading up to the block that the transaction is included in. Because of this, block proposers
7856 know which transactions are valid when they propose blocks, which assures them that the
7857 transaction fees included in the block will be received. It also enables simple light clients
7858 who merely verify that a transaction is included in a block while making the network more
7859 efficient by not processing invalid or duplicate transactions. In [379], four desirable prop-
7860 erties of a checkpointing protocol are described, where the properties hold under network
7861 synchrony:

- 7862 1. **Safety:** Even during periods of asynchrony, honest participants checkpoint the same
7863 block in an iteration of the protocol, and this checkpoint lies on the same chain as all
7864 prior checkpoints.
- 7865 2. **Recency condition:** A newly checkpointed block must have been at an exact depth
7866 k in the chain of an honest user within the recent past. If arbitrarily old checkpoints
7867 could be issued, many honestly mined blocks could be overwritten, so this property
7868 bounds the number of honest blocks that may be reorged.
- 7869 3. **Gap in checkpoints:** There must be a large enough interval between checkpoints
7870 that the recency condition holds. This limits the frequency with which honest blocks
7871 may be overwritten by checkpoints.
- 7872 4. **Conditional liveness:** If all honest nodes share a common prefix in all but a few
7873 blocks, then a new checkpoint will be issued within a bounded time.

7874 The checkpointing subprotocol used in the checkpointed longest chain protocol is just a
7875 slight modification of Algorand’s BA^* protocol (described in Section 13.4.2) extended
7876 from single-shot BA to a repeated version. BA^* satisfies all four properties, but many
7877 alternative BFT algorithms do not. For example, PBFT and HotStuff fail the recency con-
7878 dition, which provides an adversary with a powerful attack: lock on a block privately while
7879 the leader of a view, but wait to finalize it until it is not in the longest chain anymore (de-
7880 spite being a descendent of the most recent checkpoint). This attack would cause many
7881 honest blocks to be invalidated.

7882 Of the above listed properties, Afgjort only requires safety, though it adds two other desired
7883 properties. Afgjort requires that the finalized chains held by honest nodes keep up with
7884 the underlying blockchain’s chain growth, which is the opposite of having a sufficiently




Fig. 46. Sharding architecture. The system is partitioned into groups that each maintain separate shards with their own distinct ledger and transactions. A cross-shard transaction protocol is needed in order for shards to interact safely in parallel. In some sharded systems, a single node may participate in multiple shards, displayed here as multicolored circles. [396]

7885 large gap in checkpoints, as required by the checkpointed longest chain. Both protocols
7886 provide $\frac{n}{3}$ -support. The properties of GRANDPA include safety, the recency condition, and
7887 conditional liveness, but it does not satisfy the gap in checkpoints property. GRANDPA
7888 tries to finalize blocks close to the tip of the chain, where blocks may not yet be in a
7889 common prefix of honest nodes. Because honest nodes might have different views of blocks
7890 near the tip, it is likely that more honest blocks must be discarded when they are found to
7891 conflict with blocks that are ultimately finalized.

7892 15. Sharding

7893 Sharding is a scaling methodology from traditional database systems. In the context of
7894 state machine replication, the goal of sharding is to divide the required network communi-
7895 cation, state data storage, and transaction processing computations across different subsets
7896 of validating nodes. In this way, a given node need not maintain or process the entire log
7897 of the state machine. There are many academic and commercial sharding proposals, such
7898 as Elastico [380], OmniLedger [381], RapidChain [382], Monoxide [383], Zilliqa [384],
7899 NEAR Protocol [385], Ethereum 2.0, and more [386–393]. The basic idea of sharding is
7900 shown in Figure 46. A sharding protocol that implements state machine replication typi-
7901 cally contains the following subcomponents [394, 395]:

- 7902 • Identity establishment and committee selection, which may include registering public
7903 keys or IP addresses and submitting proof-of-work solutions or other Sybil-resistance
7904 measures
- 7905 • An overlay setup for committees so that committee members may communicate with
7906 one another
- 7907 • Intra-committee consensus to agree on a set of transactions within a shard
- 7908 • Cross-shard transaction processing to maintain atomicity when a transaction requires
7909 reading or writing the state of multiple shards

- 7910 • Epoch reconfiguration, which generally requires unpredictable and unbiased ran-
7911 domness

7912 Essentially, a sharding protocol will take a registered set of validators and securely sort
7913 them into committees responsible for validating particular shards. Each committee executes
7914 a consensus algorithm to agree on a transaction log within a particular shard. Transactions
7915 that take place across shards must be handled with extra care. Periodically, the committees
7916 need to be re-shuffled in order to prevent an adversary from taking over individual shards.
7917 Depending on the scheme, there may be some fixed number of shards, or more shards may
7918 be added as new validators join the network. Designing a scalable and secure sharding
7919 protocol is non-trivial, and each one of the above components has its own challenges and a
7920 wide design space.

7921 There is a limit to how much sharding can improve the scalability of state machine repli-
7922 cation. In particular, sharding cannot improve scalability in the face of a fully adaptive
7923 adversary but can scale from $O(n)$ to $O(\frac{n}{\log n})$ against mildly adaptive adversaries (i.e., cor-
7924 ruptions are determined at the beginning of each epoch and static throughout), where n is
7925 the number of validators in all shards [397]. This improvement requires the system to pro-
7926 vide a mechanism to create succinct proofs that the epoch's state updates were valid (e.g.,
7927 issuing checkpoints or some form of verifiable computation) at the end of every epoch.
7928 Scalability is greatly impacted by how many transactions take place across shards and how
7929 many shards they involve. If validators are required to keep track of the information in the
7930 corresponding shards of all cross-shard transactions, then the sharding protocol does not
7931 (asymptotically) improve scalability [397].

7932 15.1. Intra-Shard Consensus

7933 The participants are split into shards that each independently run their own consensus al-
7934 gorithm. One of the primary challenges in sharding design is to ensure that validators are
7935 partitioned into committees such that all committees satisfy the security threshold of the
7936 underlying consensus mechanism with overwhelming probability. Access to unpredictable
7937 randomness is critical to this process and is discussed in more detail in Section 15.2. Be-
7938 cause there are fewer validators per shard, it is more challenging to prevent individual
7939 shards from being taken over by an adversary than it is to secure an unsharded network.

7940 Most sharding proposals use a permissioned BFT algorithm to maintain consensus within
7941 an individual shard. One of the first sharding systems proposed, Elastico, simply uses
7942 PBFT. Unfortunately, because PBFT's communication complexity is $O(n^4)$ in the worst
7943 case, shards are unable to support many validators. This leads to an unacceptably high
7944 failure probability for individual shards at the claimed security level, which tolerates up to
7945 $\frac{1}{4}$ of the computational power of the network being Byzantine (where proof of work is used
7946 to register identities). Elastico was tested with shards of only 100 validators. Based on the
7947 cumulative binomial distribution, the probability of having at least 34 Byzantine validators
7948 assigned to a shard is 2.76% per shard.

7949 OmniLedger improves upon this by suggesting a more scalable BFT protocol called Byz-
7950 CoinX. In ByzCoinX, the communication pattern is changed so that instead of every valida-
7951 tor being pairwise connected, validators are assigned evenly to groups within each shard.
7952 The protocol leader assigns one validator in each group to be a group leader in charge of
7953 communicating with the protocol leader on behalf of the group members, and new group
7954 leaders are chosen if they do not respond quickly enough. Within each group, signatures
7955 are aggregated to reduce communication complexity, and once the protocol leader gets re-
7956 sponses from $\frac{2}{3}$ of group leaders, they can move on to the next phase of the consensus
7957 protocol. This pattern allows far more validators to participate in a shard and is, thus, able
7958 to attain security when up to $\frac{1}{4}$ of the computational power is Byzantine.

7959 RapidChain is able to improve this threshold to tolerate up to $\frac{1}{3}$ of the computational power
7960 being adversarial by using a synchronous BFT protocol (secure under an honest majority
7961 of participants) for intra-shard consensus. While this does allow a shard to remain secure
7962 with fewer validators, cross-shard transactions can cause consensus failure if the synchrony
7963 assumption is violated.

7964 The division into small committees presents other security issues. The incentives for val-
7965 idators to behave properly within a shard are understudied, and for a simple scheme where
7966 rewards are equally split among participating validators, the Nash equilibrium includes not
7967 participating in committee tasks like validation and message passing [398]. Worse, if the
7968 adversary is capable of bribing committee members, it can be substantially cheaper to bribe
7969 and corrupt enough validators on a particular shard than to take over the entire system. If
7970 an invalid state transition occurs in one shard, invalid state may propagate to other shards
7971 via cross-shard transactions, which can allow actions like spending funds that do not exist.
7972 Fraud proofs, discussed in Section 15.5, are an attempt to remedy this.

7973 15.2. Identity Registration, Committee (Re)configuration, and Epoch Ran- 7974 domness

7975 Like the other systems described in this document, some kind of Sybil-resistance mecha-
7976 nism is required to participate in a sharded ledger system. Many systems employ a global
7977 ledger (sometimes called the beacon chain, identity chain, reference chain, etc.) to keep
7978 track of these identities and other shard-related metadata.

7979 In Elastico, the number of committees grows linearly with the total computational power
7980 deployed on the network. In the first step of every epoch, validators generate an identity by
7981 completing a proof of work that covers a public key, IP address, and the epoch's random
7982 seed. An identity is assigned to a committee of c validators based on the least significant
7983 bits of the proof-of-work solution. Upon creating an identity, a validator must figure out
7984 which other validators are assigned to the same shard in order to establish point-to-point
7985 connections with them. To reduce the bandwidth of communicating these identities, par-
7986 ties contact a *directory committee* composed of the first c identities created within the same
7987 epoch according to that party's view and receive at least $\frac{2c}{3}$ lists of c identities. The valida-

7988 tor then considers the union of these lists to be their committee. This leads to discrepancies
7989 in validators' knowledge of who else is in the same committee, but these discrepancies
7990 are bounded. Each shard runs its intra-committee consensus algorithm on disjoint transac-
7991 tion sets and then sends the valid transactions to the *final committee*. The final committee
7992 merges the transactions from each shard, which is then broadcast to the network.

7993 The last step of the epoch requires the final committee to generate and broadcast the next
7994 epoch's random seed using a commit-and-XOR protocol. Each member of the final com-
7995 mittee generates an r -bit random string, R_i , and sends the hash $H(R_i)$ to the other members.
7996 The committee executes a consensus instance to agree on a set S of at least $\frac{2c}{3}$ hashes, and
7997 then S is broadcast to the network. Next, each final committee member reveals their R_i
7998 preimage to the complete network. At this point, every node will have received between
7999 $\frac{2c}{3}$ and $\frac{3c}{2}$ random strings, discarding ones that do not match their hash. In the next epoch,
8000 users determine their random seed by XORing any $\frac{c}{2} + 1$ of the R_i . Since nodes may choose
8001 different R_i , nodes must include the set of random strings in their identity so that others may
8002 verify that they match the commitments in S .

8003 There are a number of weaknesses to this scheme, which are improved upon in later de-
8004 signs. Firstly, the commit-and-XOR randomness generation process can be biased by at-
8005 tackers. In addition, validators will frequently switch shards, which reduces scalability;
8006 switching only some nodes per epoch would help. This is compounded by the need to store
8007 and propagate every transaction in the system from every shard. Only the transaction vali-
8008 dation and execution process is improved, but neither storage nor bandwidth improve from
8009 sharding this way. As a result, there are no cross-shard transactions in Elastico. Finally,
8010 Elastico requires a trusted setup to generate the initial epoch randomness, which must be
8011 revealed to all parties at the same time.

8012 OmniLedger improves upon Elastico in a number of ways. Identities are generated sim-
8013 ilarly and then committed to a global *identity blockchain*, where the commitment must
8014 occur the epoch prior to a validator participating in the system. In each epoch, a new,
8015 bias-resistant random seed is generated using a VRF in a manner similar to that used in
8016 Algorand (Section 13.4.2). The VRF output is used to elect a leader to run a distributed
8017 random beacon protocol called RandHound, which is then used to create the epoch's ran-
8018 dom seed.

8019 More specifically, let e be the epoch, $config_e$ be the list of registered validators from
8020 the identity blockchain, and v be the view. Then, when epoch e begins, each validator i
8021 computes $ticket_{i,e,v} = VRF_{sk_i}("leader" || config_e || v)$ and broadcasts it to the network. After
8022 waiting for Δ time, validators accept the lowest valued ticket as the leader for RandHound.
8023 Validators wait another Δ time for the leader to initiate RandHound, incrementing the view
8024 number if the leader fails. Eventually, RandHound will output a random seed and correct-
8025 ness proof, which is broadcast to all registered validators. The random value is used to
8026 permute the registered validator list, which is then divided into roughly equal-sized com-
8027 mittees. To maintain liveness during epoch changes, shards are reconfigured only in small

8028 batches of at most $\frac{1}{3}$ of the shard size at a time.

8029 Crucial to OmniLedger’s performance is the idea of *state blocks*, which commit to the
8030 complete state of the shard and are analogous to checkpoints in classic BFT systems like
8031 PBFT. At the end of each epoch, the state block is appended to the shard’s chain and points
8032 to the previous epoch’s state block. When a validator switches shards, they do not execute
8033 every transaction included but rather bootstrap their state based on the state block. This
8034 enables transactions from older epochs to be pruned from the shard chain.

8035 OmniLedger still requires a trusted setup to generate the initial random seed for the VRF.
8036 RapidChain solves this problem with a secure bootstrapping subprotocol that uses verifiable
8037 secret sharing (VSS). During bootstrapping, the initial set of RapidChain participants agree
8038 on a set of $O(\sqrt{n})$ nodes to be the *root group*. The root group generates a random seed
8039 using VSS to be used to elect a *reference committee*, C_R , of size $O(\log n)$. The reference
8040 committee then creates k more shards of similar size by hashing participants’ identities to
8041 the range $[0,1)$ and partitioning this space into k regions. To create the initial root group,
8042 RapidChain participants are divided uniformly at random into groups of size $O(\sqrt{n})$ using
8043 a deterministic process. Each group creates its own random seed using VSS. Within each
8044 group, all participants hash the seed and their public key. A small constant number of
8045 the lowest hashes from each group are elected and make up the root group. To inform
8046 everyone else, these hashes are gossiped to other groups with at least half of the signatures
8047 of the group.

8048 Identity registration in RapidChain is performed using proof of work that incorporates the
8049 prior epoch’s randomness. Upon solving a proof-of-work puzzle, a node submits it to the
8050 reference committee to be included in a *reference block*, which contains the active identities
8051 for the next epoch, the shards they are mapped to, and the next epoch’s randomness. The
8052 reference block is then sent to each committee.

8053 The epoch reconfiguration protocol in RapidChain is based on the *Cuckoo rule*, which
8054 is safe because it only allows a constant number of validators to switch their committee
8055 per epoch. This makes RapidChain resistant to attacks where the adversary strategically
8056 attempts to join or leave the network in order to concentrate their power within a target
8057 shard. As a result, committees can be reconfigured more frequently than with OmniLedger.
8058 The Cuckoo rule works as follows: new identities that join the network are mapped to a
8059 random position $r \in [0,1)$. For some constant x , identities within the interval $(r-x, r+x)$
8060 are moved to new random positions. Partitioning nodes into k groups of $O(\log n)$ has been
8061 proven secure against an adversary with $\leq \frac{1}{2} - \frac{1}{k}$ of the computational power. That is, each
8062 committee will have a bounded number of Byzantine nodes.

8063 The issue of shard (re)allocation was studied more formally in [399], where the Worm-
8064 hole protocol was proposed. This work describes two performance metrics for sharding
8065 allocation:

8066 1. *Self-balance*: Ideally, nodes would be uniformly distributed among shards. Other-

8067 wise, shards with fewer nodes will have weaker fault tolerance, while shards with
8068 more nodes will have worse performance. While it may not be possible to attain
8069 perfect load-balancing in a permissionless network due to nodes freely joining and
8070 leaving, the optimal load balance can be achieved by having a random subset of
8071 nodes move to other shards. Self-balance, then, is the ability of the sharding protocol
8072 to recover when the load becomes imbalanced.

8073 2. *Operability*: When nodes move from one shard to another, they must synchronize
8074 their state to match that of the blockchain in the new shard. This process can have
8075 high overhead and reduce these nodes' availability while synchronizing. Operability
8076 measures the cost of performing this relocation to another shard.

8077 Unfortunately, it is impossible for a shard allocation protocol to simultaneously satisfy both
8078 optimal self-balance and optimal operability. Existing protocols tend to fall into extremes
8079 on either self-balance or operability. For example, Elastico and OmniLedger optimize self-
8080 balance at the expense of operability, while Monoxide (Section 15.4) has optimal operabil-
8081 ity but lacks self-balance. RapidChain does not fall into either extreme but has inferior per-
8082 formance compared to Wormhole. Luckily, there is another property, *non-memorylessness*,
8083 which allows a shard allocation protocol to parameterize between the two metrics. A proto-
8084 col is non-memoryless if each shard allocation does not rely solely on current and incoming
8085 system states but also takes into account prior system states. Wormhole has this property
8086 and, thus, allows a system designer to parameterize this trade-off instead of picking one to
8087 optimize.

8088 Wormhole (Figure 47) assumes the existence of an external randomness beacon and uses
8089 the beacon output as input to a VRF. The VRF output determines shard allocation. Each
8090 randomness beacon output updates the system state used for non-memorylessness. How-
8091 ever, the more prior system states that are used, the more information it takes to prove
8092 membership in a shard. Furthermore, the shard membership proof in epoch r depends on
8093 the proof in epoch $r - 1$ in a recursive manner. Wormhole prevents this proof from be-
8094 coming unbounded in size by using a parameter, w , that controls how often old proofs are
8095 discarded. An *era* consists of w epochs, where each epoch corresponds to a new beacon
8096 output. Each node must have one non-memory-dependent allocation epoch per era in or-
8097 der to discard old proofs. If all nodes were to discard their historical proofs and have this
8098 non-memory-dependent allocation in the same epoch, operability would be significantly
8099 impacted as many nodes would change their shards simultaneously. Instead, Wormhole
8100 randomly assigns nodes to non-memory-dependent allocation epochs. When w is larger,
8101 nodes joining the system will need to execute the VRF more times, making it more expen-
8102 sive for an adversary to join a target shard by trying to join the system repeatedly.

8103 Wormhole uses an operability parameter, op , to manage the trade-off between operability
8104 and self-balance. When op is large, there is only a low probability that a node will be
8105 moved to another shard. To determine which shard a node is allocated to in memory-
8106 dependent epoch r , nodes consider the most recent VRF outputs since their latest non-

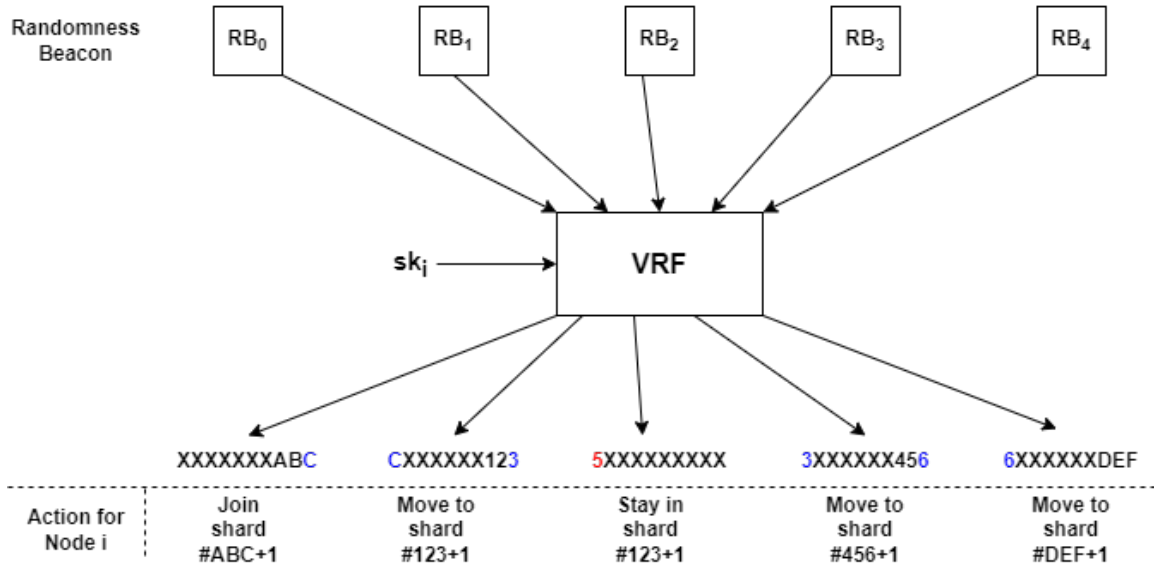


Fig. 47. Wormhole shard allocation. Simplified example in hexadecimal with parameters $op = 4$, $w = 5$, and $m = 16^3$ shards. Epoch 0 is the most recent non-memory-dependent epoch for node i , so they join a shard based solely on the VRF output. In the other $w - 1$ epochs, node i compares the op most significant bits of the new VRF output with the op least significant bits of the output that most recently assigned them to a new shard. Matches are shown in blue, and misses are in red.

8107 memory-dependent epoch, x . Let m be the number of shards and $\{h_x, \dots, h_r\}$ be the set
 8108 of VRF outputs in memory. Then, for $j \in [x + 1, r]$, nodes check whether the op most
 8109 significant bits of h_j match the op least significant bits in the VRF output from the most
 8110 recent epoch where a node changed shards. If so, the node switches to shard $(h_r \bmod m) + 1$,
 8111 and if not, the node remains on the same shard.

8112 15.3. Cross-Shard Transaction Processing

8113 In order for sharding to improve a system's throughput, it is necessary for transactions to
 8114 be partitioned across shards such that not all nodes need to process every transaction or
 8115 maintain the full system state. As a result, many transactions are likely to take place across
 8116 multiple shards: if Alice's account is on shard A and Bob's account is on shard B, then
 8117 Alice and Bob will not be able to transact without impacting the state of both shards. For
 8118 this to be safe, sharding systems require an *atomic commit* for the state changes that occur
 8119 across shards; either both accounts are updated, or neither is. The most common algorithm
 8120 for solving this problem is the two-phase commit (2PC) algorithm, and the protocols in this
 8121 section follow that paradigm.

8122 OmniLedger presents a UTXO-based cross-shard transaction subprotocol called Atomix.
 8123 In order to remove the need for direct shard-to-shard communication, Atomix tasks the

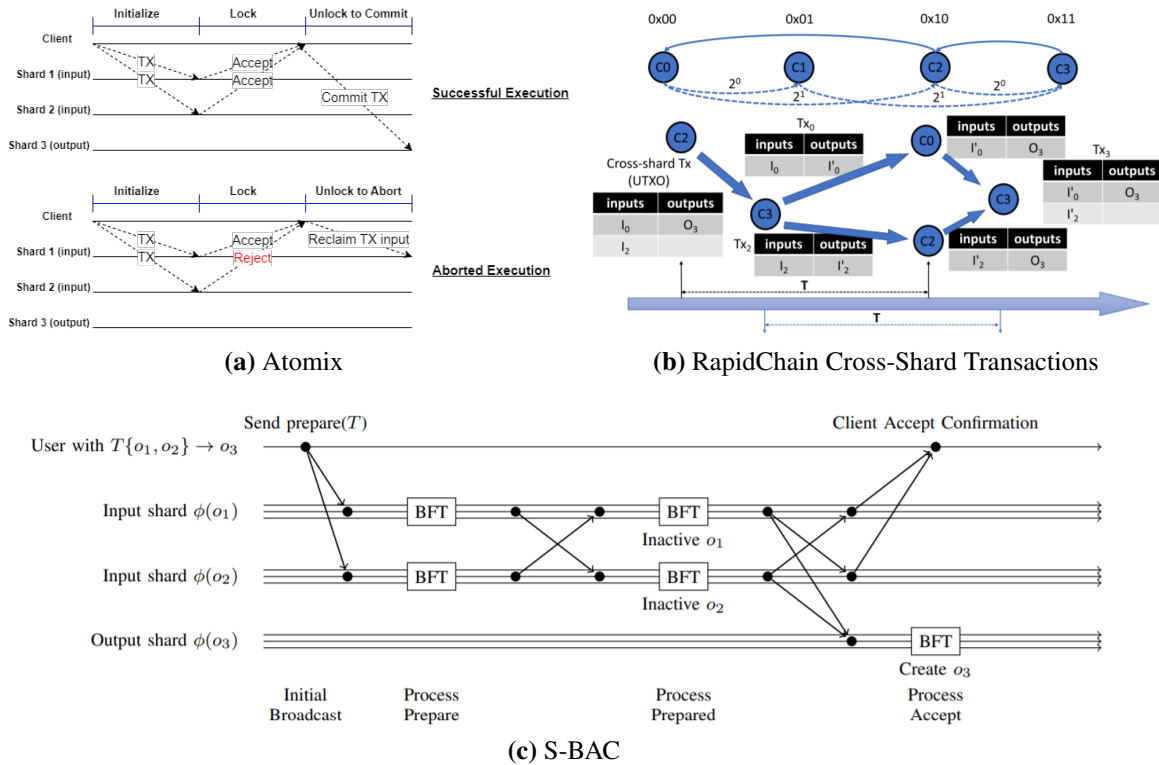


Fig. 48. Cross-shard transactions. (a) Atomix. (b) RapidChain. The top panel shows each shard maintaining a routing table with $O(\log n)$ other shards to improve communication efficiency. Committee C_0 can reach C_3 via C_2 for transactions beginning with 0x11. The bottom shows a cross-shard transaction being split into three pieces. (c) S-BAC. Transaction T has two inputs, o_1 and o_2 , and one output o_3 . The client sends T to all nodes in the input shards. A designated leader in each shard sends either a $prepared(accept, T)$ or $prepared(abort, T)$ message to the nodes within their shard. The leader of each shard determines whether all shards are in a state of $proposed(accept, T)$ or if there are any in $proposed(abort, T)$, handles the $accept(T, *)$ messages, and sends the decision to the client. [386, 396]

8124 transacting client with the responsibility of driving the process forward but allows other
 8125 parties to help if a transaction is stalled. Atomix is a three-step process of initialization,
 8126 locking, and unlocking, shown in Figure 48a. Atomix is initialized by having the client
 8127 create a cross-shard transaction and gossiping it to each shard responsible for the transac-
 8128 tion inputs. If the transaction is valid, it is included in a shard block, and a Merkle proof
 8129 of inclusion is the transaction’s proof of acceptance. If the transaction is invalid, a proof
 8130 of rejection is created instead by setting a bit in the block. The client gathers proofs of
 8131 acceptance from each input shard and can communicate them to the output shards who can
 8132 then generate the needed UTXOs. Alternatively, with a proof of rejection, the client can
 8133 abort the transaction to unlock their funds on the input shards.

8134 RapidChain adopts a different approach to cross-shard transactions in the UTXO model.

8135 To reduce the amount of communication used between shards, a Kademlia-like overlay
8136 network is used to route cross-shard transactions to the appropriate shards, as shown in
8137 Figure 48b. Cross-shard transactions in RapidChain are broken into multiple transactions.
8138 For instance, for a two-input one-output transaction with inputs I_1 and I_2 from shards one
8139 and two and output O in shard three, the transaction in Figure 48b will be executed as three
8140 different sub-transactions:

- 8141 • tx_1 consumes input I_1 and creates output I'_1 belonging to shard three.
- 8142 • tx_2 consumes input I_2 and creates output I'_2 belonging to shard three.
- 8143 • tx_3 consumes inputs I'_1 and I'_2 and creates output O .

8144 In effect, tx_1 and tx_2 transfer I_1 and I_2 to the output shard, which are then spent in tx_3
8145 to create the intended output O . Each sub-transaction occurs on a single shard, and the
8146 committee for the output shard will route the input transactions to their relevant shard
8147 committees. If, say, tx_2 were to fail while tx_1 was successfully committed, the owner of
8148 UTXO I_1 instead uses I'_1 in a future transaction.

8149 Another approach, S-BAC (Sharded Byzantine Atomic Commit), was proposed for the
8150 sharded smart contract platform Chainspace [386]. S-BAC is similar to Atomix but eschews
8151 the client-driven model in favor of one that explicitly runs a Byzantine agreement algorithm
8152 combined with atomic commit, as shown in Figure 48c. The client sends their transaction,
8153 T , to the input shards, which internally execute PBFT in order to agree on a tentative de-
8154 cision to accept or abort the transaction. Those tentative decisions are then broadcast to
8155 the other shards involved in the transaction as $prepared(abort, T)$ or $prepared(accept, T)$
8156 messages that include signatures from shard members proving they decided a particular
8157 way. If the transaction was locally accepted, the input is considered locked. Shard mem-
8158 bers listen for responses from the other shards involved. If all responses support accept-
8159 ing transaction T , then it is committed, but if any shards want to abort, the transaction is
8160 aborted. The shards then exchange a round of $accept(commit, T)$ or $accept(abort, T)$ mes-
8161 sages based on their decision and send them to the client as well. Once the transaction is
8162 committed, the output shards generate new outputs, and the inputs are consumed. If the
8163 transaction is aborted, input shards unlock the inputs.

8164 Both S-BAC and Atomix were susceptible to replay attacks that do not require violating
8165 Byzantine thresholds [400]. It is crucial for a system design for cross-shard transactions to
8166 defend itself against replay attacks in order to prevent invalid state from being incorporated
8167 into the ledger. The attacks had two causes, which informed their fixes:

- 8168 1. The input shards do not have a way of knowing that particular protocol messages
8169 received correspond to a specific instance of a transaction, so old messages can be
8170 replayed. To fix this, sequence numbers are added to transactions.
- 8171 2. In some cases, the output shards are only involved in the later unlocking phase of the
8172 protocol and therefore have no knowledge of the transaction context that is available

8173 to the input shards. To fix this, output shards create dummy objects in the earlier
8174 locking phase, which makes them input shards as well.

8175 A concern related to cross-shard transactions is how to partition the system state across
8176 shards. Because cross-shard transactions have higher latency and require more effort from
8177 the network, dividing the state in such a way that accounts that frequently interact with
8178 each other are located on the same shard would improve the system's performance. On the
8179 other hand, partitioning the state this way can lead to very memory-inefficient mappings
8180 between shards and the state they are responsible for. Most schemes partition the state
8181 based on a simple mapping of account prefixes to shards, which is essentially random.
8182 Assuming that there are 100 shards, one would expect 99% of transactions to be cross-
8183 shard. With only 10 shards, this still results in 90% of transactions being cross-shard. A
8184 few alternative mapping strategies that can reduce the fraction of cross-shard transactions
8185 while being relatively memory-efficient are explored in [401]. This includes ideas such as
8186 using graph clustering on accounts that historically have interacted with each other (though
8187 this is an NP-hard optimization problem), clustering the most heavily used accounts and
8188 assigning the rest randomly, and clustering based on the most frequently used accounts
8189 during a recent time period.

8190 **15.4. A Different Approach: Monoxide**

8191 Monoxide proposes a very different approach to sharding from what has been discussed so
8192 far [383]. Unlike the above examples, identities in Monoxide are established only once,
8193 and there is no committee reconfiguration. Monoxide uses accounts rather than UTXOs,
8194 and accounts are partitioned based on their most significant bits. Each individual shard
8195 uses proof of work and the GHOST fork-choice rule (Section 11.2) for consensus. Every
8196 node in Monoxide must be a light client of every shard and maintain a chain of block
8197 headers for each. Nodes also maintain a distributed hash table (DHT) for routing cross-
8198 chain transactions and for peer discovery.

8199 The most noteworthy aspect of Monoxide is its use of *chu-ko-nu mining*, which is similar to
8200 the concept of merged mining (see Section 16.2). Chu-ko-nu mining allows miners to create
8201 blocks in multiple shards simultaneously with a single proof of work, allowing honest
8202 miners to amplify their hash rate and prevent malicious miners from concentrating their
8203 computational power into taking over single shards. If every honest miner takes advantage
8204 of this and mines on every shard simultaneously, then Monoxide attains security assuming
8205 an honest majority of the hash rate. Miners will gather valid transactions from all shards
8206 and include a Merkle root of the block headers for each shard in their proof of work.

8207 While this mechanism increases the adversarial threshold required to attack Monoxide
8208 compared to other sharding protocols, it is imperfect. A miner needs to verify all trans-
8209 actions in order to participate in all of the shards and achieve the honest majority security
8210 bound, but this eliminates the scalability benefits for those miners. The result is likely to
8211 be severe centralization pressure among miners, since well-resourced miners can mine on

8212 more shards and collect more transaction fees.

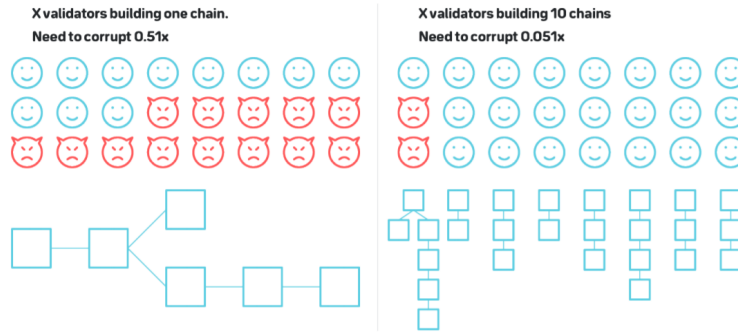
8213 Monoxide also takes a different approach to cross-shard transactions, bypassing the over-
8214 head of the locking and unlocking operations used in other protocols. Instead, Monoxide
8215 accepts that transaction atomicity only holds eventually with high probability. As a result, a
8216 credit can happen on one shard before the corresponding debit on a different shard has been
8217 fully settled. Transactions are validated in the shard of the payer and then verified in the
8218 shard of the payee using the header chain of the payer's shard and a *relay transaction*.
8219 The relay transaction exists on the payer's shard and provides any metadata required to
8220 validate the original transaction using only the shard's header chain. A miner in the payee
8221 shard then verifies that the relay transaction is stable and includes it in a block on that shard
8222 as well.

8223 While interesting, Monoxide has some challenges. In addition to the mining centralization
8224 pressure, it may be challenging to handle transaction fees, particularly for cross-shard trans-
8225 actions. Similarly, keeping track of inflation of the money supply is difficult or impossible
8226 without completely validating all transactions on all shards. Finally, the requirement to act
8227 as a light client for all shards makes it so that Monoxide does not scale asymptotically,
8228 though it may offer a significant constant factor improvement.

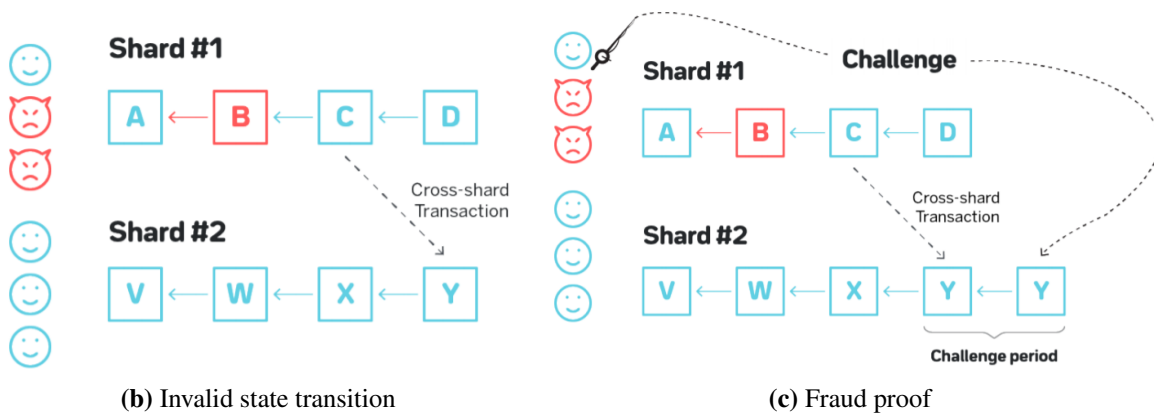
8229 **15.5. Fraud Proofs and Data Availability**

8230 In a sharded blockchain system, the first challenge to resolve is the so-called *1% attack*,
8231 where an adversary with a small fraction of the total network resources (e.g., work, stake,
8232 validated identities, etc.) can concentrate them in a single shard and exceed its Byzantine
8233 threshold, as seen in Figure 49a. These types of attacks are typically prevented using a
8234 form of random sampling to assign validators to shards in an unpredictable way, such as
8235 the mechanisms described in Section 15.2. Unfortunately, reshuffling has high overhead.
8236 Each time a validator is assigned to a new shard, it must download the state of that shard
8237 in order to validate a new block, which can take an extended period of time. This prevents
8238 reshuffling from happening too frequently. Naively, this would require downloading and
8239 executing all shard blocks that have been produced since the last time the node was assigned
8240 to that shard, in which case sharding has provided no scalability benefit. One step toward
8241 mitigating this is to have every shard block's header commit to the state of the shard, so
8242 nodes need not store the complete system state at all times. However, it does not ensure
8243 that the state at the tip of the chain is correct, which would require processing all of the
8244 shard blocks anyway.

8245 On the other hand, the frequency with which shard reallocation is performed relates to how
8246 adaptive of an adversary the system can tolerate. When reallocations are infrequent, the
8247 adversary is given a lot of time to (adaptively) target and corrupt validators on individual
8248 shards, reintroducing the risk of single-shard takeover attacks. In most of the systems
8249 described above, the shards implicitly trust each other. That is, nodes in shard A simply
8250 assume that no invalid state is ever committed to shard B. This may not be realistic in the



(a) Single shard takeover attack



(b) Invalid state transition

(c) Fraud proof

Fig. 49. Invalid shard state transition after a single shard takeover attack. An invalid state transition occurred in block B, perhaps crediting an account with undeserved tokens. A cross-shard transaction moves these counterfeit coins to another shard. Any honest node on the corrupted shard can generate a fraud proof and submit it to the victim shard to demonstrate that the state transition was invalid and should be rolled back. [385]

8251 real world, where adversaries may indeed be adaptive and can even bribe validators. In this
 8252 case, the security of the protocol decreases linearly in the number of shards. When invalid
 8253 state is committed to a shard, it can then influence the rest of the system via cross shard
 8254 transactions, as shown in Figure 49b.

8255 To provide security against fully adaptive or bribing attackers, a sharding system should
 8256 provide validators with the ability to quickly see if a proposed shard block is valid or
 8257 not. Security in this model is possible using *fraud proofs*, where a node supplies proof
 8258 of an invalid state change, and adding an *any-trust assumption* on the validators that were
 8259 assigned to the shard in the last round. That is, as long as there is a single honest and
 8260 available validator assigned to the shard, it can construct a fraud proof in response to an
 8261 invalid block being created on their shard, as shown in Figure 49c. A fraud proof may
 8262 contain the invalid transaction and some Merkle tree data verifiable from the shard block
 8263 header’s state commitment. In other words, a fraud proof will contain the relevant portions

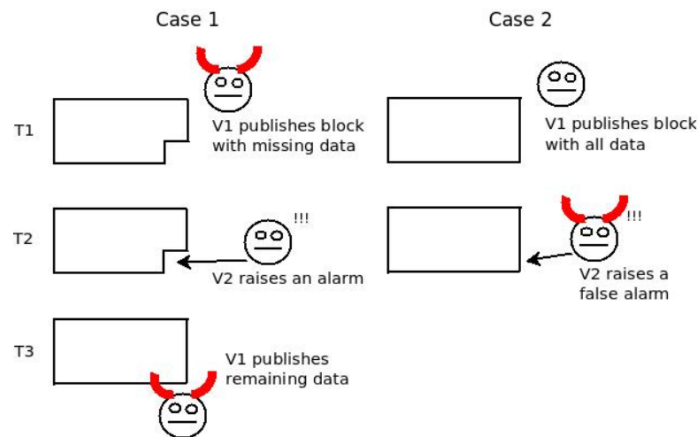


Fig. 50. Data availability attack. The left and right sides are indistinguishable after T_3 . [402]

8264 of the state needed to process the block, as well as intermediate hashes of the Merkle tree
8265 that prove that the provided state is in fact the state that the block claims to be using [402].
8266 For example, for an invalid transaction that debits 100 units of currency from an account
8267 possessing only 50 units, the fraud proof would include the transaction itself, an assertion
8268 that the account's balance before the transaction was only 50 units, and a Merkle proof
8269 showing that this account balance was committed in the state root of the block header.

8270 For this to work, there must be some bounded challenge period where an honest validator
8271 has sufficient time to download and process the block, prepare a fraud proof if it is invalid,
8272 and send it across the network. While the challenge period must be long enough to allow
8273 invalid shard blocks to be caught, longer periods delay the settlement of cross-shard trans-
8274 actions. Furthermore, this creates a new attack vector where malicious nodes send invalid
8275 fraud proofs. This can be mitigated by requiring stake deposits from validators; if a node
8276 sends an invalid fraud proof, they lose their deposit. Similarly, if a node submits a valid
8277 fraud proof, the creator of the invalid block loses their deposit, and a portion of it can go to
8278 the node who created the proof as a finder's fee.

8279 One more major problem remains. Consider an adversary who controls a supermajority
8280 of the validators on a shard. This adversary creates an invalid block and signs off on its
8281 validity through the validator keys they control. Now, the adversary announces the block
8282 by broadcasting its header but withholds some or all of the block data itself. In this case,
8283 the honest validators on the shard are unable to produce a fraud proof and can only claim
8284 that block data is being withheld. Unfortunately, after an honest node makes this claim, the
8285 adversary can immediately publish the data, and other nodes will be unable to distinguish
8286 between a situation where the block producer maliciously withheld the block or a malicious
8287 validator raised a false alarm, as shown in Figure 50.

8288 This *speaker-listener fault equivalence* prevents the system from punishing false data avail-
8289 ability alarms or rewarding valid ones. As a result, under adversarial conditions, all valida-

8290 tors would have to download all shard blocks, eliminating any scalability gains. To see this,
8291 suppose an attacker executes this attack (Case 1 in Figure 50). If the expected return from
8292 raising an alarm were positive, this would encourage malicious validators to submit false
8293 alarms frequently and pocket the proceeds. If the expected reward were neutral, then this
8294 false alarm method provides a free denial of service to force everyone to download all shard
8295 blocks. If the expected reward were negative, then raising the alarm would be altruistic and
8296 irrational. An adversary could simply outlast the altruists and then launch data availability
8297 attacks while honest nodes have no recourse [402].

8298 This attack is why *data availability proofs* are required for effective sharding protocols
8299 in models that assume fully adaptive or bribing attackers. These proofs can be generated
8300 using erasure codes, where a block M chunks in size is expanded into N chunks, $N > M$,
8301 such that any M chunks are sufficient to recover the original data. Block headers include a
8302 Merkle root committing to these N chunks. Light clients can then contact fully validating
8303 nodes to check whether the majority of the N chunks are available. If the majority is indeed
8304 available, this implies one of three possibilities:

- 8305 1. The full block is available and valid with a correctly constructed erasure code. In this
8306 case, the light client should accept the block.
- 8307 2. The full block is available, and the erasure code is correctly constructed, but the block
8308 is invalid. In this case, the light client would expect an honest full node to construct
8309 a fraud proof and broadcast it to the network shortly.
- 8310 3. The full block is available, but the erasure code is not properly constructed. In this
8311 case, the light client would expect an honest full node to construct and broadcast a
8312 special fraud proof that demonstrates that the erasure code is incorrect.

8313 Essentially, after conducting the data availability check and having it pass, a light client can
8314 wait for the duration of some challenge period to see a fraud proof and treat the block as
8315 valid if none are forthcoming. A clever attacker may try to beat this by releasing individual
8316 chunks of data as clients ask for them. Since light clients will only sample chunks proba-
8317 bilistically, if there are not enough light clients performing this sampling, an attacker can
8318 indeed pull this off and trick them into thinking the block is fully available when it is not.
8319 This necessitates an additional security assumption: there are enough light clients making
8320 data availability queries that it is overwhelmingly likely that the intersection of their re-
8321 quests covers enough erasure-coded data to recover the full block. Erasure coding forces
8322 the adversary to withhold a much greater fraction of the block than they otherwise could
8323 have to perform malice, while the data availability queries make it harder for the adversary
8324 to get away with it.

8325 The adversary still has one more trick up their sleeve: they can honestly answer requests
8326 from many light clients but stop responding before the clients can verify enough of the
8327 block's availability. In this case, the adversary can still trick some of the light clients into
8328 thinking the block is available when it is not. To address this, clients should send their data

8329 availability queries through an anonymous network like Tor, so the adversary cannot tell
8330 which queries came from the same client.

8331 A relatively efficient proof system using sparse Merkle trees to represent the system state
8332 and 2-dimensional Reed-Solomon erasure codes for data availability is presented in [403].
8333 This system encodes k chunks of data as $2k$ erasure-coded chunks and then provides the
8334 following guarantees under synchrony with maximum network delay Δ :

- 8335 • *Soundness*: If an honest light client is convinced that a block is available, then there
8336 exists at least one honest full node with the complete block within some delay $k * \Delta$.
- 8337 • *Agreement*: If an honest light client is convinced that a block is available, then every
8338 light client will consider it available within some delay $k * \Delta$.

8339 The fraud proofs from [403] grow logarithmically in size with the size of the block and
8340 state, while the availability proofs grow proportional to the square root of the block size.
8341 Similar schemes using a new primitive called Coded Merkle Trees have been proposed
8342 [404–406].

8343 By combining fraud and data availability proofs, the security of the system can be assured
8344 by assuming a sufficient number of light clients querying for data availability, as well as
8345 a single honest full node per shard who is capable of producing and broadcasting a fraud
8346 proof within the challenge period. If the adversary is capable of linking validator IDs to
8347 IP addresses (perhaps using Sybil nodes on the network and performing traffic analysis),
8348 they can violate the any-trust assumption by launching denial-of-service attacks against
8349 validators who refuse to collude.

8350 Recall that when a validator changes shards, they need to synchronize their state in order to
8351 process and verify new shard blocks. If the validator needed to fully execute all of the shard
8352 blocks since they were last validating the shard, this could take hours or days. However,
8353 given a system of fraud and data availability proofs, and assuming that there are sufficient
8354 light clients to verify data availability and the any-trust assumption within the shard holds,
8355 this synchronization process can be sped up dramatically: download the new state, verify
8356 that it matches the state commitment in the block header, and wait for a single challenge
8357 period to see if anyone submits a fraud proof.

8358 Despite the speedup, this still leaves new nodes unavailable for some period of time after
8359 being reassigned to a new shard. One mechanism that can mitigate the consequences of
8360 this is to abstract different protocol roles to different parties. For instance, the management
8361 of data and execution may be split between nodes with different roles: block proposers,
8362 data availability notaries, and transaction executors [407]. Block proposers build a chain of
8363 shard blocks, and notaries attest to the availability of data from the shard blocks. The block
8364 proposers do not perform any state-dependent validation. Executors take the shard chain
8365 agreed upon by the block proposers and execute the transactions, skipping over any that are
8366 invalid and signing the state. This would allow executors to stay on one shard while block

8367 proposers and notaries can be shuffled and reassigned more frequently. Unfortunately,
8368 this has several important trade-offs, such as complicating the design of light clients and
8369 making block proposers ignorant of the validity of the transactions they include in their
8370 blocks, which makes it harder to keep the system design incentive-compatible.

8371 This delay issue can be fully resolved with the idea of *stateless clients* [408]. Roughly,
8372 these would be clients that only maintain the state root, not the state itself. Instead of
8373 downloading and maintaining the full system state, stateless clients download a block wit-
8374 ness composed of Merkle proofs for all of the pieces of state accessed in the block. This
8375 allows the node to compute new state roots without needing a full copy of the state. Block
8376 witnesses computed this way are quite large themselves, making this scheme impractical.
8377 Recent cryptographic advancements, such as Verkle trees, may be able to reduce witness
8378 sizes sufficiently to enable stateless clients [409].

8379 16. Interoperability

8380 Due to the proliferation of distributed ledgers that have been and will be deployed, some de-
8381 gree of interoperability is desirable. This can allow for atomically exchanging assets across
8382 multiple ledgers, creating "wrapped" representations of assets that can exist on alternative
8383 ledgers, and other cross-chain smart contracts [410, 411].

8384 16.1. Cross-Chain Communication, Fair Exchange, and Atomic Swaps

8385 A basic requirement to have interoperability among distributed ledgers is for the relevant
8386 ledgers be able to "communicate" with each other. In particular, a ledger should be able
8387 to validate that certain events have occurred in another ledger in order to interact with it.
8388 More formally, an entity P monitoring ledger L_X and entity Q monitoring ledger L_Y should
8389 be able to synchronize with each other such that Q writes TX_Q to L_Y if and only if P has
8390 written TX_P to L_X [412]. In other words, the goal of cross-chain communication is to
8391 ensure that either both of TX_Q and TX_P are written to their respective ledgers or neither is.
8392 An example is shown in Figure 51.

8393 Cross-chain communication is very similar to the atomic commit problem mentioned in the
8394 context of cross-shard transactions in Section 15.3, where relevant processes must agree on
8395 whether a transaction was committed or aborted. More specifically, *non-blocking atomic*
8396 *commit* (NB-AC) requires that all correct processes decide on a common outcome even
8397 when other processes fail. Cross-chain communication is essentially the NB-AC problem,
8398 but participants can have Byzantine failures instead of just crash failures.

8399 In fact, it has been proven that cross-chain communication reduces to the *fair exchange*
8400 *problem* [412]. In a fair exchange, two parties want to exchange goods, and either both
8401 items are transferred successfully or neither are. In the analog world, this can be thought of
8402 as the problem of a customer simultaneously handing cash to the cashier while the cashier
8403 hands the customer a purchased product. In the digital realm, this is often associated with

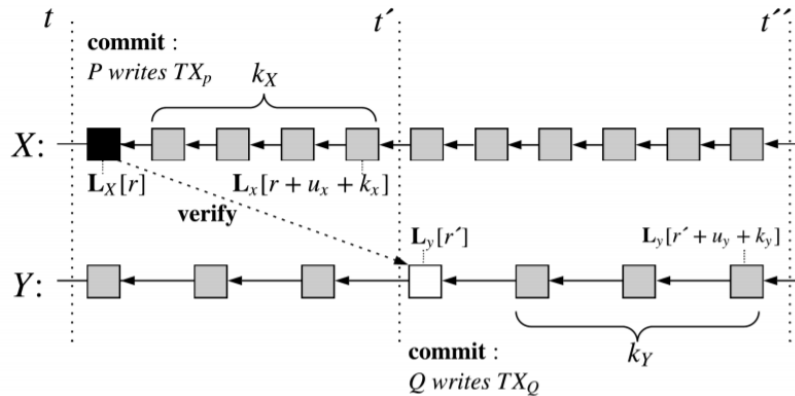


Fig. 51. Cross-chain communication example. Systems X and Y maintain ledgers L_X and L_Y , respectively. A node Q writes TX_Q if and only if P has written TX_P . In this example, transactions are committed in the next block of the ledger once broadcast (liveness delays $u_X = u_Y = 0$). For safety, the persistence delays for systems X and Y are $k_X = 4$ and $k_Y = 3$. [412]

8404 either selling access to files or trading digital signatures. Unfortunately, it has long been
 8405 known that fair exchange is impossible without a trusted third party in asynchronous net-
 8406 works [413]. Cross-chain communication, therefore, requires either:

- 8407 • Synchrony between participants such that one ledger can verify the existence of a
 8408 transaction on another ledger within a known, bounded amount of time or
- 8409 • A trusted third party, which can be abstracted as a separate replicated state machine.

8410 A line of research has produced protocols for fair exchange that use smart contracts as the
 8411 trusted third party and often rely on incentives and the rationality of the participants for
 8412 security [414–418]. These protocols punish malicious users who try to violate the fairness
 8413 of the exchange such that a rational participant will not do so. However, this does not
 8414 ensure correct execution of the exchange, even if it prevents honest parties from being
 8415 harmed economically. While these protocols tend to operate on a single ledger, similar
 8416 concepts are frequently used in interoperability protocols.

8417 One of the earliest interoperability proposals in the distributed ledger space is the *atomic*
 8418 *swap*, where an asset from one blockchain system is traded for an asset on a separate
 8419 blockchain system without using a trusted intermediary. A specific mechanism for doing so
 8420 using *hashed timelocked contracts* (HTLCs) was proposed by TierNolan [419], generalized
 8421 and formalized in [420], and machine verified in [421]. The atomic swap protocol provides
 8422 the following guarantees:

- 8423 • When all parties execute the protocol honestly, all swaps will occur.
- 8424 • If any parties deviate from the protocol, honest parties will not end up worse off.

- 8425 • No party or coalition of parties is incentivized to deviate from the protocol.

8426 An example of a cross-chain atomic swap is shown in Figure 52. Assume that Alice pos-
8427 sesses some of an asset, ACoin, that resides on blockchain *A*, while Bob possesses BCoin
8428 on blockchain *B*. Alice and Bob can perform an atomic swap as follows:

- 8429 1. Alice generates a random number s and keeps it secret until later.
- 8430 2. Alice publishes a smart contract with a hash lock $H(s)$ and time lock t_A to blockchain
8431 *A*. If Bob can show s before t_A , the ACoin is transferred to Bob; otherwise, Alice can
8432 reclaim it after t_A .
- 8433 3. Seeing the published value $H(s)$, Bob publishes a smart contract on blockchain *B*
8434 using hash lock $H(s)$ and time lock t_B with $t_B < t_A$. If Alice provides s before t_B , then
8435 BCoin is transferred to Alice; otherwise, Bob can reclaim it after t_B .
- 8436 4. Prior to t_B , Alice will publish s and acquire BCoin from Bob's smart contract. At this
8437 point, s is revealed on blockchain *B*.
- 8438 5. Prior to t_A , Bob publishes s to acquire ACoin from Alice's smart contract.

8439 Many other atomic swap protocols have been proposed since. For example, JugglingSwap
8440 works even when the elliptic curves used for transaction signing differ across blockchains
8441 [424]. Another alternative, AC³WN, allows participants to publish their smart contracts
8442 simultaneously and thus substantially reduce latency for the swap [425]. Other proposals
8443 focus on specific cryptocurrency pairs with unique challenges, such as swapping Bitcoin
8444 and Monero [426, 427]. Because Monero does not include a scripting language, different
8445 techniques are required. There are also protocols for more general atomic commerce, such
8446 as auctions, arbitrage, brokering deals, and escrow-based payments [428, 429]. Atomic
8447 debt instruments [430] and options contracts [422] have also been proposed. Finally, there
8448 are a variety of cross-chain communication protocols that allow arbitrary smart contracts
8449 to function across multiple blockchains [431–436].

8450 A known issue with most atomic swap protocols is that they behave like a free American
8451 call option for one of the parties to the swap [437–440]. An American call option is a
8452 contract that allows a party to purchase some quantity of an asset at an agreed-upon price
8453 at any point prior to the agreed-upon expiry time. The party that initiates the atomic swap is
8454 the "buyer" of the American call option but does not need to pay a premium for this option.
8455 This can cause problems if the exchange rate between the swapped assets changes during
8456 the period that the swap is executed. The option buyer can abort the swap if the change
8457 in exchange rate is unfavorable to them. Adding an extra phase where a premium is paid
8458 helps mitigate this issue [440].

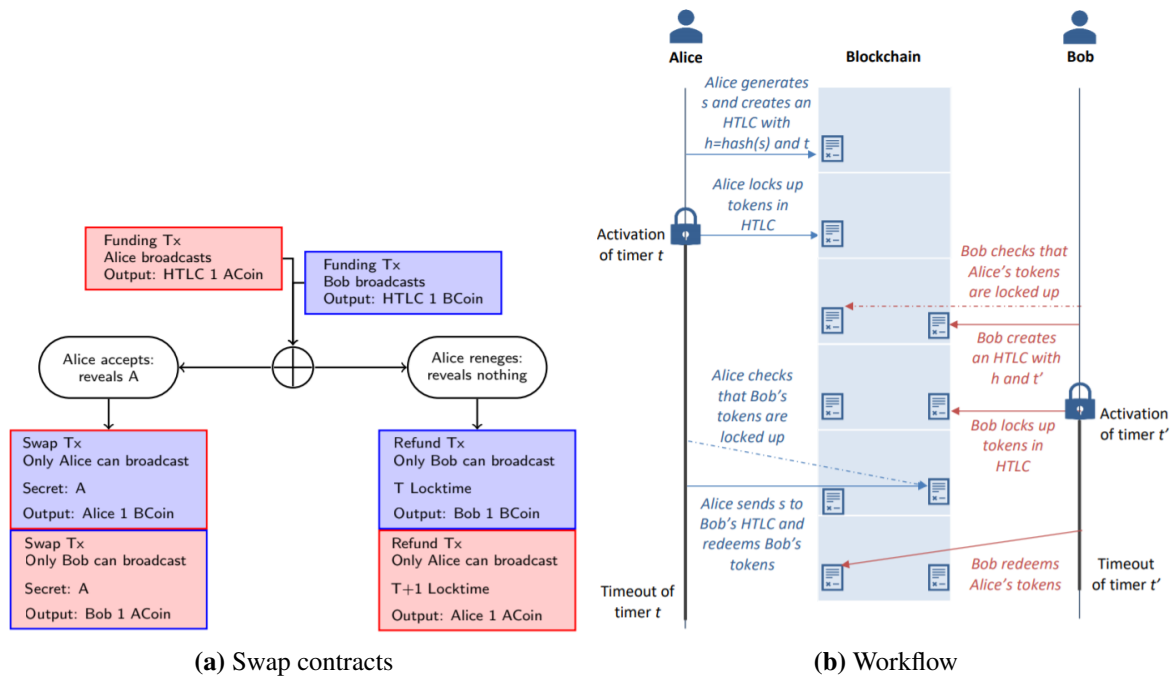


Fig. 52. Atomic swap. (a). Each box is a transaction. Boxes with red fill color take place on the ACoin blockchain, whereas those with blue fill color are for the BCoin blockchain. Alice can publish transactions with red borders, while Bob can publish transactions with blue borders. A box directly beneath another box represents a transaction that should be published in response to the one above it. (b). Another representation of the atomic swap workflow. [422, 423]

8459 16.2. Bootstrapping Methods: Merged Mining and Proof of Burn

8460 Merged mining is a specialized form of interoperability where a miner mines on multi-
8461 ple blockchain networks simultaneously. It is primarily a mechanism for bootstrapping a
8462 new "child" proof-of-work blockchain by piggybacking on the security of a more mature
8463 "parent" chain.

8464 When using merged mining, a valid proof-of-work solution for the parent chain is consid-
8465 ered valid for the child, which allows the child chain to inherit the security of the parent
8466 chain. Miners construct a block template for the child chain and include a hash of this
8467 block in their template for the parent chain. Miners work off of this parent chain template
8468 and search for proofs of work at either the (lower) difficulty level of the child chain or the
8469 (higher) difficulty of the parent chain. When a valid proof of work at the parent difficulty
8470 level is found, it is a normal valid block in that network, and the child chain block hash
8471 is ignored by participants in the parent network. If a valid proof of work is found at the
8472 lower child difficulty level, a block on the child chain is created and will include the parent
8473 chain block header, the coinbase transaction where the child hash is included, and a proof
8474 that the coinbase transaction is included in the parent block. This block is submitted to the
8475 child chain's network, which accepts the proof of work. Because the parent chain header
8476 cannot be created without including the child chain's hash, this still proves that work was
8477 performed while constructing the child block. Of course, a block solved at the parent dif-
8478 ficulty level will also satisfy the child difficulty level. An example of the merged mining
8479 procedure is shown in Figure 53.

8480 There are a number of prominent blockchain networks that employ merged mining, includ-
8481 ing Namecoin and Rootstock on Bitcoin and Dogecoin on Litecoin [442, 443]. According
8482 to [443], some 90% of Bitcoin miners are merged mining as of October 2020, which is
8483 a dramatic increase compared to a few years prior. At that time, Rootstock was the most
8484 popular chain for merged mining, with 40 to 50% of Bitcoin miners participating.

8485 Merged mining has security implications for both the parent and child chains. Merged
8486 mining includes a certain amount of overhead, primarily to perform validation of the child
8487 chain. Larger miners on the parent chain are more likely to be able to afford this overhead
8488 and thus claim the extra rewards from mining the child chain, which may lead to centraliza-
8489 tion pressure for the parent chain [444, 445]. An alternative merged mining proposal called
8490 *blind merged mining* may be able to address the centralization concerns for the parent chain
8491 [446]. This will be discussed more in the next section in the context of the Drivechains pro-
8492 tocol [447, 448]. Because validation costs are the primary cost of participating in merged
8493 mining, the miners are also less likely to validate the child chain. As a result, the security
8494 of the child chain may not be as significant as the hash rate would suggest. In practice,
8495 it has frequently been the case that the hash rate on merged mined child chains fluctuates
8496 significantly and is highly concentrated [444].

8497 There are additional risks to parent chains, which can present interesting problems given
8498 that a parent chain cannot prevent being merged-mined by child chains. This issue is com-

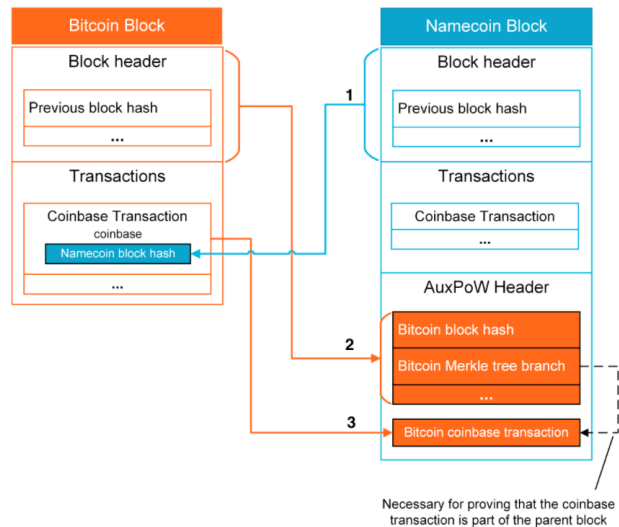


Fig. 53. Merged mining example, where Namecoin is the child chain to Bitcoin. In the first step, a Namecoin block is constructed, and its hash is placed in the coinbase transaction of a Bitcoin block. If a proof of work is found at Bitcoin's difficulty level, the Bitcoin block header, the coinbase transaction, and a Merkle proof of the coinbase transaction are included in a special portion of the Namecoin block. [441]

8499 pounded by the advent of multi-merged-mined cryptocurrencies that have multiple parent
8500 chains. A miner for this cryptocurrency could attack a subset of the parent chains at no
8501 additional cost [444]. Parent chains can also be attacked when there are (social) disputes
8502 over the rules that a ledger should follow, though it would be challenging to pull off such
8503 an attack [449].

8504 Another technique that is frequently used for interoperability and to bootstrap new ledgers
8505 is *proof of burn*, or the verifiable destruction of a cryptocurrency [450]. To "burn" an asset
8506 involves sending it from a source ledger to a provably unusable address. The proof of
8507 burn itself consists of the transaction that burned the asset, the block header for the block
8508 containing the burn transaction, a (Merkle) proof demonstrating that the burn transaction
8509 was included in that block, and a proof that the submitted block is in the "best" source
8510 blockchain (e.g., the chain with the most proof of work) and is stable. This technique is
8511 used in asset transfer using relay smart contracts, as described in the next section, but also
8512 sometimes as a Sybil-resistance mechanism for a separate consensus algorithm.

8513 Related to proof of burn is another bootstrapping mechanism employed in the Stacks
8514 blockchain called proof of transfer (PoX), where the security of the Stacks ledger is an-
8515 chored to Bitcoin [451]. In PoX, instead of burning units of the base cryptocurrency (Bit-
8516 coin), validators transfer those coins to owners of the new cryptocurrency (e.g., Stacks) via
8517 a predetermined set of Bitcoin addresses as a reward for producing blocks on the new chain.
8518 A Stacks block producer commits to their proposed block in a Bitcoin transaction, and one
8519 of the various block commit transactions within a single Bitcoin block is chosen via a VRF

8520 in proportion to the amount of Bitcoin transferred. Stacks block producers then follow the
8521 longest chain rule. This mechanism ensures that the entire history of Stacks forks is public
8522 unless the adversary launches a majority attack against the underlying Bitcoin chain. As a
8523 result, network participants can react to reorg attempts (e.g., by increasing the number of
8524 required confirmations) in order to make it far more difficult to launch profitable attacks
8525 [452].

8526 16.3. Sidechains, Relays, and Asset Transfer

8527 A *sidechain*, first proposed in [445], is a blockchain whose currency is "pegged" to a
8528 "main chain" but is responsible for its own consensus security, independent of the main
8529 chain. A sidechain protocol, therefore, includes a main-chain consensus mechanism, a
8530 sidechain consensus mechanism, and cross-chain communication. A sidechain may be
8531 centrally administered, federated (i.e., permissioned), or permissionless, though permis-
8532 sionless sidechains have known issues and are relatively untested in the real world. A
8533 sidechain may have a two-way peg, in which case sidechain assets can be transferred back
8534 to the main chain or a one-way peg, where sidechain assets permanently remain on the
8535 sidechain. Several desirable properties of pegged sidechains were described in [445]:

- 8536 • Assets should be able to be transferred with minimal counterparty risk.
- 8537 • An asset on a sidechain should be able to be transferred back to the main chain by
8538 the current holder of the asset and no one else, including previous holders of the
8539 sidechain asset (if there is a two-way peg).
- 8540 • Asset transfers should not have failure modes that result in the loss of funds or the
8541 fraudulent creation of assets.
- 8542 • Participants in the system should not need to monitor sidechains that they are not
8543 actively involved with.
- 8544 • Sidechains ought to be firewalled from each other and the main chain. That is, if
8545 there is a bug or blockchain reorganization in a sidechain, asset inflation or the theft
8546 of assets should be localized to the sidechain in question.

8547 Examples of federated sidechains include private Ethereum sidechains [453] and the Liquid
8548 sidechain on Bitcoin [454, 455]. A federated sidechain may employ various security fea-
8549 tures, such as having separate federations for managing sidechain consensus and moving
8550 assets from the sidechain to the main chain, as well as using hardware security modules for
8551 transaction and consensus signing.

8552 A related interoperability concept is that of relays, which operate as smart contracts on
8553 one ledger that act as light clients for another ledger. The relay smart contract is sup-
8554 plied with block headers from the alternative chain that it needs to validate. Examples
8555 include BTC Relay, allowing Bitcoin verification on Ethereum [456]; PeaceRelay, which

8556 creates a two-way bridge between Ethereum and Ethereum Classic [457]; Dogetherium,
8557 allowing Dogecoin on Ethereum [458]; and XCLAIM, which can work with a variety of
8558 ledgers [459]. A scalability issue with this simple type of relay is that there is an increasing
8559 overhead for submitting an up-to-date block header as the length of time since the prior
8560 submission increases. All intermediate headers must be submitted and validated as well.

8561 While most relays accept the full chain of block headers from the alternative chain, some
8562 designs are more advanced. For example, FlyClient is a more efficient way of performing
8563 light validation using only a logarithmic number of block headers [460]. Other relays use
8564 advanced cryptographic primitives, like SNARKs, to produce concise proofs of valid chain
8565 progression [461, 462]. Another proposal uses cryptoeconomic incentives to optimistically
8566 accept block headers instead of validating them and allows submitted headers to be chal-
8567 lenged. Submitters post a bond to the relay smart contract that is forfeit if a successful
8568 challenge proves the header invalid [463].

8569 In addition to relay smart contracts, there are relay chains that are purpose-built for in-
8570 teroperability, such as Polkadot [464] and Cosmos [465]. These schemes have a central
8571 blockchain that acts as a communication hub between the various interoperable blockchains
8572 within their respective networks, not unlike how sharding works (Section 15). The two
8573 projects differ in several important ways [466]:

- 8574 • The security models of each project differ. In Polkadot, the interoperating blockchains
8575 inherit the security of the global relay chain. As a result, relay chain validators have
8576 the final say over state changes in the individual "shards" (called "parachains" in
8577 Polkadot parlance). In contrast, individual chains in the Cosmos network are in-
8578 dependent and responsible for securing themselves. This makes the chains more
8579 "sovereign" by allowing them to be run independently but also requires them to boot-
8580 strap their own security.
- 8581 • The governance of the respective projects differ. To become a parachain in Polkadot,
8582 one must acquire a significant quantity of Polkadot's native DOT cryptocurrency and
8583 stake them. Polkadot can only support a particular number of parachains. Cosmos
8584 lacks fixed membership rules, allowing anyone to establish a new blockchain within
8585 the network with its own custom governance process.
- 8586 • Polkadot allows arbitrary message passing between participating parachains, includ-
8587 ing interoperable smart contract calls. Cosmos is particularly focused on asset trans-
8588 fer between chains. Architecturally, this manifests itself in the "tight coupling" of
8589 Polkadot's parachains, achieved by using fraud and data availability proofs to handle
8590 invalid blocks on parachains (see Section 15.5). Cosmos lacks this tight coupling, so
8591 users of participating blockchain *A* need to "trust" the relay chain used to communi-
8592 cate with blockchain *B*. In Cosmos, a malicious validator cannot corrupt a blockchain
8593 that they do not belong to. In Polkadot, on the other hand, the entire system is im-
8594 pacted if invalid parachain blocks are not challenged.

- 8595 • The two projects use different consensus algorithms. Polkadot uses the GRANDPA
8596 algorithm, which scales to a potentially large number of validators [376]. The Cos-
8597 mos relay chain, or hub, uses Tendermint consensus, which handles fewer validators
8598 but provides quick finality [358, 360].

8599 16.3.1. Permissionless Sidechains

8600 Sidechains that utilize a permissionless consensus algorithm are more challenging to de-
8601 sign than permissioned sidechains, despite being the original motivation for the sidechain
8602 concept in the first place. As originally conceived, the state machines of sidechain and
8603 main chain would need to be capable of acting as light clients for other blockchains. In
8604 particular, they would need to be able to validate SPV proofs in order to verify that trans-
8605 fers from the main chain to the sidechain (and vice versa) were included and stable in the
8606 canonical chain [445].

8607 For this type of sidechain, there would be special outputs on the main chain that are desig-
8608 nated as belonging to the sidechain. To transfer coins from the main chain to the sidechain,
8609 a user will lock their coins in one of these special outputs for a certain *confirmation period*.
8610 The confirmation period exists to allow enough proof of work to be generated to prevent
8611 denial-of-service attacks against another waiting period that follows. Once the confirma-
8612 tion period is over, the user can create a transaction on the sidechain that references the
8613 special main-chain output and uses an SPV proof to demonstrate that the locking was per-
8614 formed correctly and generate corresponding sidechain coins. Once this is done, the user
8615 must wait for a *contest period* during which the newly converted assets are unspendable on
8616 the sidechain. This contest period is used to prevent a double-spending attack vector where
8617 previously locked coins on the main chain are transferred to a non-sidechain address due
8618 to a majority hash rate attack. During the contest period, any sidechain user may publish a
8619 reorganisation or reorg proof showing that there exists a main chain with greater aggregate
8620 proof of work and where the block containing the locked output no longer exists. This in-
8621 validates the main-chain-to-sidechain transfer. This two-way pegging mechanism is shown
8622 in Figure 54, where it is contrasted with federated and centralized pegs.

8623 There are several challenges with this sidechain mechanism. First, it leads to the risk of
8624 centralization of mining in the same way that merged mining does. Larger miners are more
8625 able to cope with the demands of validating the sidechain and are thus more likely to collect
8626 transaction fees on the sidechain. Over time, it is conceivable that the extra revenue source
8627 can drive smaller miners out of business. Second, the security model seems to rely on the
8628 ability to publish reorg proofs on the sidechain. However, an adversary capable of perform-
8629 ing a lengthy reorg on the main chain is likely capable of doing the same on the sidechain,
8630 in which case they can censor the inclusion of a reorg proof. This kind of sidechain, there-
8631 fore, may allow a majority hash rate attacker to steal all of a sidechain's coins for free.
8632 Third, SPV proofs grow linearly with the size of the chain, so main-to-sidechain and side-
8633 to-main-chain transfers will involve increasingly large transactions. Luckily, these proofs

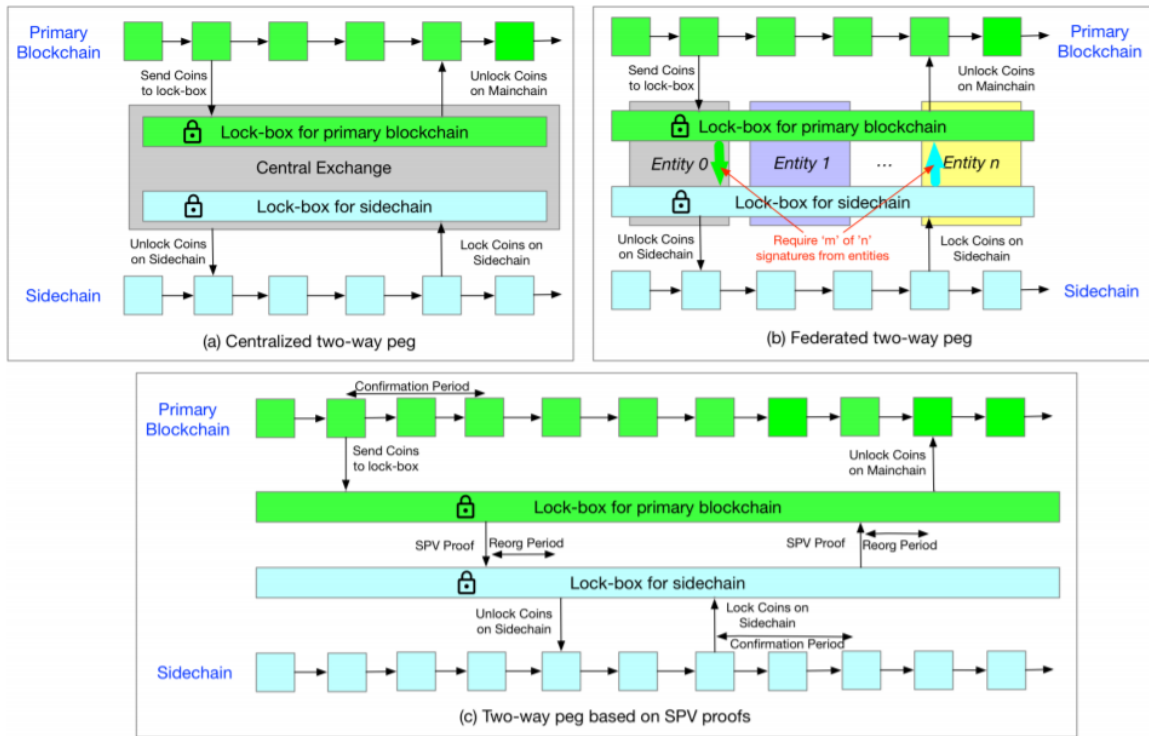


Fig. 54. Sidechain pegging methods. [467]

8634 can be made logarithmic using modified light client protocols, such as FlyClient [460].
8635 An alternative technology that can make proofs logarithmic is the idea of *non-interactive*
8636 *proofs of proofs of work* (NIPoPoWs), which use skip lists to take advantage of blocks that
8637 prove more work than necessary for the difficulty level, which occur with relatively pre-
8638 dictable frequency [468, 469]. Unfortunately, NIPoPoWs are currently unable to handle
8639 chains where the difficulty level is variable, making them less practical than FlyClient.

8640 Proof of stake sidechains have also been proposed [470]. In this case, instead of proving
8641 that a sufficient amount of work occurred during the confirmation period, a proof includes
8642 the set of signatures from the elected slot leaders whose length is proportional to the com-
8643 mon prefix confirmation security parameter. One of the challenges related to bootstrapping
8644 a proof-of-stake sidechain is that it would be highly vulnerable to attack before a significant
8645 amount of stake is dedicated to it, but validators may not want to commit their stake to the
8646 sidechain while it is vulnerable. This is addressed by the idea of merged-staking, where
8647 stakeholders from the main chain can signal their awareness of a sidechain in order to be
8648 eligible for block creation on that sidechain but without needing to transfer the stake itself.

8649 Another proposal for proof-of-work sidechains is called Drivechains [447, 448]. A Drivechain
8650 is like a "plugin" for a main chain, and a main-chain full node is a required component of
8651 running a sidechain full node. In the context of Bitcoin's state machine rules, all coins in
8652 a particular Drivechain are locked in an "anyone can spend" address on the Bitcoin chain,

8653 and a new soft-forked rule is added to prevent these coins from being spent on the Bitcoin
8654 chain unless additional conditions are met.

8655 Drivechains replace the requirement of submitting chains of block headers for each trans-
8656 action with a more explicit form of miner voting. Transfers from the main chain to the
8657 sidechain no longer require a lengthy confirmation period and can be done nearly instantly
8658 upon deposit in the "anyone can spend" address. On the other hand, sidechain-to-main-
8659 chain transfers use a much slower mechanism. When sidechain users want to withdraw
8660 their coins back to the main chain, they create a sidechain withdrawal transaction that burns
8661 the sidechain coins. Over an extended period of time, these withdrawal transactions accu-
8662 mulate on the sidechain and are eventually aggregated into a large main-chain transaction
8663 from the locked "anyone can spend" address to the withdrawal addresses specified in the
8664 sidechain withdrawal transactions. The transaction ID of this aggregated withdrawal trans-
8665 action can then be embedded in the coinbase transaction on the main chain and represents
8666 a withdrawal attempt. For some pre-specified number of blocks following the inclusion of
8667 this withdrawal attempt, miners vote on whether the withdrawal is valid. If the miner vote
8668 passes a threshold in favor of the withdrawal, the actual withdrawal transaction is included
8669 in a block. Due to the lengthy withdrawal period, users may tend to perform atomic swaps
8670 with main-chain coins rather than withdraw directly.

8671 This security model allows a majority coalition of main-chain miners to steal the full con-
8672 tents of a sidechain by creating and voting to approve any withdrawal transaction of their
8673 choosing. Because the main chain is unaware of the sidechain state, users of the sidechain
8674 may be required to organize themselves into a user-activated soft fork that prevents this
8675 theft. Realistically, this would only be plausible with relatively popular sidechains.

8676 A final piece of the Drivechain puzzle is the use of *blind merged mining* to reduce the
8677 centralizing impact that other permissionless sidechains may have [446]. The underlying
8678 centralization concern is that the sidechain fees are only available to those entities with the
8679 resources to run a fully validating sidechain node. Blind merged mining allows main-chain
8680 miners to opt out of running a sidechain node but without the opportunity cost of losing
8681 sidechain fees. It does this by providing a mechanism for main-chain miners to "sell" the
8682 service of finding a sidechain block in a trustless way. Miners on the sidechain know that
8683 a valid sidechain block with hash H is worth, say, X sidechain coins in transaction fees
8684 and are thus willing to pay up to (approximately) X main-chain coins to have H included
8685 in the main chain's coinbase transaction (and thus acquire X sidechain coins themselves).
8686 A main-chain miner can be unaware of the sidechain and still accept payment of X main-
8687 chain coins to include H in a block they are constructing at no cost to themselves. Should
8688 H be invalid according to the sidechain rules, the sidechain can simply ignore it, and the
8689 main-chain miner keeps the main-chain payment anyway. The sidechain miner can keep a
8690 small amount for themselves, but this should create a competitive market with main-chain
8691 bribes close to X .

8692 17. Networking

8693 In any distributed system, processes are required to communicate over a network. In the
8694 majority of permissioned systems, every node will have direct, point-to-point communi-
8695 cation channels with every other node (see Section 8 for some examples of exceptions).
8696 Permissionless networks, on the other hand, have a more complex design space.

8697 17.1. Networking for Permissionless Systems

8698 Neudecker and Hartenstein identify five goals for the networking layer of a permissionless
8699 system: 1) performance, 2) low cost of participation, 3) anonymity, 4) resistance to denial-
8700 of-service attacks, and 5) topology hiding [471]. These goals sometimes conflict with each
8701 other. For example, high performance and resistance to denial-of-service attacks tend to
8702 require more data to be transmitted, while anonymity, low cost, and topology hiding tend to
8703 require transmitting less data over the network. All five of these goals have implications for
8704 system security. For example, if the performance of block dissemination is not extremely
8705 fast in Nakamoto Consensus, the security bound decreases, as discussed in Section 10.2.1.
8706 Low cost of participation is critical to permissionless systems in general lest full node
8707 ownership not be well distributed (see Section 3.2).

8708 Network-layer adversaries may engage in passive or active attacks to achieve certain goals.
8709 An adversary may want to deanonymize users of the system by linking network informa-
8710 tion to application-layer information. For example, the adversary may attempt to discern
8711 which IP address a particular transaction was first broadcast from. Alternatively, the ad-
8712 versary may try to exploit the network to deny service to certain nodes in order to provide
8713 the attacker with a monetary reward. For example, an attacker could induce a network
8714 partition to increase γ for selfish mining purposes (see Section 9.4), and fully isolating a
8715 node through an eclipse attack can increase the chance of a successful double-spend. These
8716 attacks all benefit from knowledge of the network topology, so an adversary may try to map
8717 the topology as an initial phase of their attack.

8718 Methods for topology inference have been studied for Bitcoin [472], Ethereum [473], Zcash
8719 [474], Monero [475], and probably other systems. Various methods of connecting IP ad-
8720 dresses to transactions have been suggested for the Bitcoin network (similar methods would
8721 work for comparable systems), some of which find or exploit knowledge of the topology
8722 [476–478]. Networking side-channels have also been found to deanonymize users of pri-
8723 vate cryptocurrency systems Zcash and Monero [479].

8724 Network-layer attackers may take advantage of nuances in the networking protocols of
8725 permissionless systems in order to force a victim node to connect to the network solely
8726 through attacker-controlled nodes, thus allowing an attacker to control the victim’s view
8727 of the blockchain. Specific eclipse attacks have been found against both the Bitcoin [480]
8728 and Ethereum networks [481, 482]. A security requirement for a node on a permissionless
8729 network is to have at least one honest peer connection, but eclipse attacks can violate this

8730 assumption and make targeted double-spending attacks much easier.

8731 An attacker may also attempt to cause more significant network partitions. In practice,
8732 Bitcoin nodes are concentrated among relatively few autonomous systems, so an adversary
8733 could cause network partitions by performing BGP hijacking of relatively few IP prefixes
8734 [483, 484]. On the other hand, if a significant number of mining pool nodes are hidden
8735 behind Tor, BGP hijacking can be made significantly more difficult [485]. BGP hijacking
8736 is easily detected, but far stealthier network partitioning attacks have also been proposed
8737 [486]. Due to there being IP addresses that host nodes from multiple cryptocurrency net-
8738 works, it may be possible to attack several networks at once by attacking a small subset
8739 of nodes [487]. There are designs for relay networks that are resistant to BGP hijacking
8740 attacks, such as SABRE [488]. Additionally, blocks can be transmitted over satellite or
8741 radio networks in case of network partitions.

8742 More esoteric attacks are also possible, such as *timejacking* [489]. Some systems, such as
8743 Bitcoin, have nodes adjust their clocks based on the median timestamp reported to them
8744 by their peers. This allows a Sybil attack variant where malicious nodes connect, report
8745 incorrect times, and move the clock of a victim node. This can get them to do things
8746 like temporarily believe that a block is invalid and cause temporary network partitions. In
8747 tandem with a software bug, this can also cause permanent chain splits that require manual
8748 intervention to fix [490].

8749 17.1.1. Peer Discovery

8750 The first thing that a new node needs to do upon startup is find a set of reachable IP ad-
8751 dresses in order to create outgoing connections to peers. Usually, the software for run-
8752 ning a node includes some hard-coded "seed" nodes that are DNS servers run by respected
8753 community members. This is inevitably a more centralized process than is desirable, as
8754 malicious seed nodes can eclipse a victim node by providing only IP addresses under their
8755 control (if the user contacts only a single seed node). A malicious seed node can also use
8756 their position to strategically impact the network topology. Finally, if the seed nodes are
8757 unavailable, new nodes are unable to start, which may make them a target for denial-of-
8758 service attacks.

8759 To mitigate these risks, there should ideally be a relatively large number of seed nodes be-
8760 ing operated and available. When a new node starts up for the first time, it should contact
8761 multiple different seed nodes and establish outgoing connections using IP addresses sup-
8762 plied by each of them. Topology hiding can be improved at the expense of increased cost
8763 of participation if clients only connect to a small portion of the IP addresses supplied by
8764 the seeds.

8765 Once a node has an established connection, it can also discover new peers from the existing
8766 ones, either by asking for more IP addresses or receiving them unsolicited. Here, there is
8767 a tension between the quality of the connection and topology hiding. When IP addresses

8768 are announced, it provides information that can be used as part of an eclipse attack or to
8769 infer the network topology, especially if nodes announced their neighbors. On the other
8770 hand, announced IP addresses should be reachable nodes, which suggest strategies such as
8771 announcing the IP addresses of recent or current connections. Because nodes frequently
8772 leave and rejoin the network, there is a risk of announcing IP addresses that are old and
8773 unreachable.

8774 Accounting for this node churn is important. The overwhelming majority of nodes will
8775 leave and rejoin the network multiple times over the course of a couple of months, the av-
8776 erage exceeds four churns per node per day, and churn increases block propagation time by
8777 weakening the performance of compact blocks (discussed in Section 10.2.1) [491]. Con-
8778 siderable churn has been observed on the Ethereum network as well [492]. Luckily, large
8779 networks with more than 4,000 nodes are fairly resistant to the negative effects of churn
8780 [493].

8781 **17.1.2. Neighbor Selection**

8782 Peer discovery provides a node with a set of IP addresses to establish outgoing connections
8783 to. With this set, the node must decide which specific peers to connect to, as well as how
8784 to handle incoming connection requests. Given a list of IP addresses, a node can either try
8785 connecting to random ones or use their knowledge about the peers to inform the decision.
8786 This knowledge can include information based on the IP address itself, such as the AS or
8787 region where the IP address is located. Alternatively, it can use information based on past
8788 experience with the peer or information from trusted external sources. A Bitcoin node, for
8789 example, limits the number of outgoing connections it establishes within a given IP address
8790 range. This forces adversaries to control a larger and more diverse set of IP addresses in
8791 order to perform an eclipse attack. Bitcoin nodes also utilize past observations from their
8792 own connections to block peers that misbehave.

8793 Another component of the neighbor selection policy is deciding how many outgoing and in-
8794 coming connections to establish. A larger number of connections can improve block prop-
8795 agation speed and make the network more robust to eclipse and partitioning attacks. On the
8796 other hand, more connections increase the cost of participation. Outgoing connections are
8797 more trustworthy than incoming ones because even weak adversaries can establish a sig-
8798 nificant number of incoming connections to an honest peer. As a result, denial-of-service
8799 resistance increases to a greater degree from more outgoing connections than additional
8800 incoming ones. That said, a large number of incoming connection slots is desirable so that
8801 other peers can establish outgoing connections and to prevent an adversary from using up
8802 the available incoming connection slots.

8803 The node software must also decide whether to maintain established connections for as
8804 long as possible or to strategically disconnect from some neighbors and replace them with
8805 others. Maintaining connections for longer makes it easier for an adversary to infer the
8806 network topology but more challenging to exploit it using their own Sybil nodes. It is

8807 also possible for the neighbor selection protocol to strategically generate certain network
8808 topologies. For instance, creating clusters of nodes based on geographic proximity can
8809 speed up information propagation but makes it easier to infer the network topology and
8810 launch various attacks.

8811 Finally, nodes can monitor their own connections for performance and security and dis-
8812 connect from peers that send too much or too little information, are slow to respond, or
8813 otherwise behave unexpectedly. While this can improve resistance to attacks, adversaries
8814 can occasionally take advantage of these mechanisms to aid in other attacks. For example,
8815 an attack on the Bitcoin network was proposed where nodes connecting over Tor can be
8816 eclipsed by taking advantage of peer monitoring [494].

8817 **17.1.3. Communication Strategy**

8818 While peer discovery and neighbor selection determine the network structure, nodes must
8819 also have a communication strategy to determine how best to disseminate information.
8820 Messages may be treated differently based on their content or the context they occur in. For
8821 example, in Bitcoin, a received block is immediately relayed (after being verified), whereas
8822 transactions are relayed with a random delay. Similarly, a node may relay transactions
8823 differently when created by the user of the node itself as opposed to transactions received
8824 from the network. As a final example, nodes may not participate in transaction relay at all
8825 when they are still performing their initial block download to synchronize with the network.
8826 Using message content and context in this way can help balance competing priorities.

8827 One dimension of the communication strategy is whether messages should be pushed in an
8828 unsolicited fashion or whether messages should be announced first and only sent directly
8829 if requested by the peer. Pushing messages results in them being disseminated across the
8830 network more quickly than the announce-and-request method because information is sent
8831 in half a round trip instead of 1.5 round trips. On the other hand, the announce-and-request
8832 method uses much less bandwidth than the unsolicited push by avoiding duplicate mes-
8833 sages. Nodes must also decide whether to flood a message to all of its neighbors or gossip
8834 the message to a (possibly random) subset of connections. Flooding uses more bandwidth
8835 than gossip but is also more robust against attacks. A node might, for instance, pick a few
8836 peers to push unsolicited blocks to in order to speed up block propagation while following
8837 announce-and-request for blocks sent to other peers (and transactions sent to all peers) in
8838 order to keep bandwidth consumption modest.

8839 Nodes must also determine when to send messages to their peers. They can, for instance,
8840 add a random delay before forwarding messages, which can improve topology hiding and
8841 anonymity while slowing down performance. They can also choose to aggregate several
8842 messages together to send simultaneously, which reduces the bandwidth overhead. For
8843 instance, a node might forward five transactions at a time. This has other potential benefits,
8844 such as making the code cleaner by maintaining only a single outgoing message queue
8845 per connection. Allowing messages to accumulate has an unclear impact on propagation

8846 speed but can improve it on average if new outgoing messages are added to a queue that
8847 was already scheduled to be sent (assuming they would otherwise be sent with a random
8848 delay). Message accumulation also has an unclear impact on anonymity and topology
8849 hiding. Message transmission times depend on information about received messages that
8850 an attacker may not be privy to. On the other hand, it can open up new attack vectors, such
8851 as inferring when a peer received a certain message from another peer using information
8852 from the aggregated set of messages.

8853 18. State Machines

8854 Distributed ledger systems utilize consensus in order to execute a state machine, where a
8855 set of servers execute commands on behalf of clients in an agreed-upon order to modify
8856 the state of the system. The state machine is the set of rules that validating servers enforce
8857 while transforming the state based on these client-submitted commands, called transac-
8858 tions. While a wide variety of rules are possible, some types of rules appear frequently. For
8859 instance, a transaction that modifies a piece of the state of the system should include a valid
8860 signature from a client that is authorized to act on that portion of the state. Other common
8861 rules include establishing a maximum block size that block proposers can create, requiring
8862 a particular syntax for blocks and transactions, and ensuring that all transactions included
8863 in a block are committed to in the block header.

8864 The idea of state machines was introduced in Section 2.8. There are two primary models
8865 that distributed ledgers use to organize and manage their state: the UTXO model and the
8866 account model, discussed in detail in Section 2.8.1. Some unique challenges with respect to
8867 changing the rules of distributed ledgers were discussed in Section 2.8.2. This section will
8868 provide some additional details regarding state machine design. A variety of alternative
8869 state machines with different properties are possible. In addition, "off-chain" or "second
8870 layer" systems that take advantage of the properties of the underlying state machine to
8871 improve scalability are introduced.

8872 18.1. Virtual Machine Design

8873 The majority of distributed ledger systems include a runtime environment or virtual ma-
8874 chine that executes platform-independent code in a low-level programming language (ex-
8875 ceptions exist, such as Monero). Transactions act as miniature programs or function calls
8876 that modify a portion of the state of the ledger by executing a series of *opcodes*, or in-
8877 structions. Some systems – like Bitcoin – use a set of opcodes with limited functionality,
8878 while others – like Ethereum – use a Turing-complete instruction set capable of executing
8879 any general program. In many cases, higher-level languages exist that are friendlier to the
8880 programmer and compile down to bytecode made up of virtual machine-specific opcodes.

8881 Bitcoin uses the Bitcoin Script language, a Forth-like, stack-based scripting system with no
8882 loops. The original implementation had many more opcodes, but a large number of them

8883 were deprecated early on in Bitcoin's history due to fears of denial of service and other
8884 attacks. The instruction set includes a number of "no operation" opcodes, OP_NOP, which
8885 allow new opcodes to be added via soft fork. In Bitcoin's UTXO model, outputs contain a
8886 *scriptPubKey* field that includes an encumbrance, or a condition to satisfy in order to spend
8887 the output. When an input spends an output, it uses the input's *scriptSig* field to include
8888 data and opcodes that satisfy the spending condition in the *scriptPubKey*. To validate an
8889 input being spent, an empty stack is created, and the program to execute consists of the
8890 *scriptSig* prepended to the *scriptPubKey* of the output being spent. This combined script
8891 is executed from left to right and is considered valid if nothing triggers failure during the
8892 script execution and the top stack item when the script completes is True (or non-zero).

8893 One of the most common, standard transaction types in Bitcoin is the pay-to-
8894 public-key-hash (P2PKH) transaction, which transfers some bitcoin from one ad-
8895 dress to another (where an address is a hash of a public key). In this case,
8896 the *scriptPubKey* will look like the following, where bracketed items are data to
8897 be pushed to the stack: OP_DUP OP_HASH160 <Owner's address> OP_EQUALVERIFY
8898 OP_CHECKSIG. The corresponding *scriptSig* consists of two push operations: <Owner's
8899 signature> <Owner's public key>. As a result, the combined script to be executed is
8900 the following: <Owner's signature> <Owner's public key> OP_DUP OP_HASH160
8901 <Owner's address> OP_EQUALVERIFY OP_CHECKSIG.

8902 Execution proceeds from left to right, so the first thing that happens is that the signature is
8903 pushed to the stack, and then the owner's public key is pushed to the stack above it. The
8904 OP_DUP operation duplicates the top item on the stack (in this case, the public key). Next,
8905 the OP_HASH160 opcode takes the public key at the top of the stack and replaces it with the
8906 RIPEMD-160 hash of the SHA-256 hash of the public key. The owner's address is then
8907 pushed to the top of the stack. From top to bottom, the stack now consists of the address
8908 supplied in the output, the hash of the public key given in the input, the public key itself,
8909 and the signature from the input. The OP_EQUALVERIFY opcode checks that the top two
8910 items on the stack are equal and fails if they are not, in which case it ensures that the public
8911 key supplied in the *scriptSig* is a preimage for the address given in the *scriptPubKey*.
8912 Assuming that they match, the top two stack items are removed, leaving just the public key
8913 on top of the signature. Finally, OP_CHECKSIG takes the public key and signature, pops
8914 them off the stack, and pushes True if the signature is valid.

8915 To check whether the signature is valid, the execution environment needs to know what
8916 message is being signed. To this end, the signature includes a one byte *SIGHASH flag* that
8917 determines which portions of the transaction are used to form the message. This provides
8918 some flexibility for users, so that only some parts of the transaction are signed, allowing
8919 others to modify the transaction if desired. For instance, it may be desirable to have the
8920 message only include a single output, allowing other users to contribute additional inputs
8921 to the transaction so long as the specified output is included too.

8922 In contrast to Bitcoin, Ethereum uses the Turing-complete Ethereum Virtual Machine

8923 (EVM) as its execution environment and uses the account model rather than the UTXO
8924 model. The EVM is capable of executing any deterministic program but subject to the
8925 EVM's *gas limit*. Gas is a unit of measurement that is intended to correspond to the com-
8926 putational effort required to execute an opcode, so each opcode has a particular gas cost.
8927 The gas limit is analogous to a maximum block size but for computational effort instead
8928 of space. There are two types of accounts in Ethereum: contract accounts and externally
8929 owned accounts (EOAs). EOAs are typical user accounts controlled by a private key that
8930 allows the user to transfer the ether currency or interact with contract accounts. Special
8931 transactions are used to deploy smart contracts by creating contract accounts that include
8932 the smart contract's bytecode. These contract accounts are then controlled by the EVM
8933 bytecode when an EOA triggers a function call on the smart contract.

8934 The EVM's global state is a mapping between addresses and the associated account's state.
8935 Each account has an ether balance, a nonce, a hash of the code deployed to that address,
8936 and a hash of the account's storage. There is also a machine state that includes a program
8937 counter pointing to the next instruction to execute, the remaining gas available, a stack,
8938 and memory. To invoke a smart contract, a user signs a transaction where the *recipient*
8939 field contains the smart contract address, and the *data* field contains information on the
8940 function to be called and arguments to pass to it. A transaction's execution may fail and
8941 throw an exception, in which case gas fees are still charged to the account and given to the
8942 transaction's miner, but any state transitions the transaction would have caused are reverted.

8943 It is extremely challenging for developers to write smart contracts directly as low-level
8944 bytecode. Instead, developers will typically use a high-level language, such as Solidity
8945 or Vyper, and compile it down to EVM bytecode before deploying. Other high-level lan-
8946 guages have been developed for other distributed ledger platforms, such as the Move lan-
8947 guage created for the Diem blockchain. Move was designed to handle the problem of con-
8948 servation, or ensuring that fund transfers preserve the total monetary supply in the system
8949 [495]. Solidity handles this naturally by associating each account with a balance that can
8950 only be modified by special instructions. For tokens built on Ethereum, the EVM cannot
8951 maintain this supply conservation property without encoding it in the token smart contract
8952 itself, whereas the Move language maintains this property even for custom tokens.

8953 A number of projects use WebAssembly (WASM) as their runtime environment for smart
8954 contracts. Due to its support on most modern web browsers, distributed ledger clients will
8955 be able to run in the browser more easily. WASM instructions can be directly mapped
8956 to machine instructions, so it should be highly performant. In the future, Ethereum will
8957 likely migrate from the EVM to eWASM, or a restricted subset of WebAssembly de-
8958 signed for Ethereum. Specifically, eWASM is the same as WASM but with floating point
8959 non-determinism removed, gas metering added, and an Ethereum-specific interface [496].
8960 While eWASM is expected to ultimately improve performance, current implementations
8961 have variable but relatively poor execution speeds [497].

8962 A final point on Turing-completeness is in order. An interesting side-effect of having a

8963 Turing-complete instruction set with gas metering is that it makes it unsafe to conduct
8964 some kinds of soft forks due to opening up a denial-of-service vector [498]. On the bright
8965 side, this also prevents some malicious soft forks, such as smart contract censorship. Con-
8966 sider a soft fork that made any transaction that operates on a particular contract invalid.
8967 Once the soft fork is deployed, an attacker can broadcast many transactions to the network
8968 that perform a variety of challenging computations before invoking the censored smart con-
8969 tract. Miners that run the soft fork code would have to execute these transactions before
8970 finding them invalid, but due to the soft fork, they would not be able to claim gas fees as
8971 compensation, nor would the attacker pay fees. Since this costs the attacker nothing, they
8972 can amplify the attack by setting high gas prices, encouraging miners to waste more com-
8973 putational resources. Static analysis could be used to see whether the transaction interacts
8974 with the censored address, but this address could be obfuscated, and static analysis is itself
8975 computationally intensive. Turing-completeness implies that, due to the halting problem,
8976 transactions must be executed in order to determine how the computation will unfold.

8977 A different approach is taken in the Mimblewimble protocol, where the scripting language
8978 is removed entirely [499–501]. Mimblewimble’s design prioritizes improved scalability for
8979 newly joining nodes while also providing a privacy boost for users. When a node joins a
8980 Mimblewimble network, the bandwidth and computational effort required to synchronize
8981 with the network is roughly proportional to the size of the UTXO set, whereas the effort
8982 is proportional to the entire transaction history for Bitcoin and most other ledgers. This is
8983 accomplished by allowing *transaction cut-through*. Consider a ledger following the UTXO
8984 model. If a transaction spends output TXO_1 and creates TXO_2 , and a second transaction
8985 spends TXO_2 and creates TXO_3 , this is equivalent to a single alternative cut-through trans-
8986 action that spends TXO_1 and creates TXO_3 . In Bitcoin, cut-through is impossible once
8987 transactions are included in the ledger, but the Mimblewimble ledger is ultimately a sin-
8988 gular aggregated transaction with a set of outputs equivalent to the full UTXO set. With
8989 no scripting language, Mimblewimble lacks a virtual machine, but it still performs state
8990 machine replication when combined with a consensus algorithm.

8991 18.1.1. Concurrency in Smart Contracts

8992 In state machine replication, where all transactions are totally ordered, it is natural to think
8993 in terms of sequential execution of transactions. However, one of the best ways to improve
8994 performance in computing is to exploit parallelism to make better use of available CPU
8995 cores. The primary challenge in adding concurrency to state machine execution is that the
8996 runtime environment requires determinism to ensure that validators remain in agreement
8997 with one another. It may nevertheless be worthwhile to attempt to solve this problem
8998 since some estimates suggest that if all available concurrency opportunities in Ethereum
8999 were exploited, there would be a factor of six improvement in execution speed [502]. A
9000 variety of approaches to adding concurrency to smart contracts have been explored [503–
9001 509]. Solana, for example, is a prominent smart contract platform where transactions must
9002 specify in advance the portions of system state that are accessed, which allows validators

9003 to execute non-conflicting transactions in parallel [510].

9004 An early approach used a locking mechanism to divide transactions across several threads
9005 of execution [503]. Whenever a transaction attempts to access a portion of the system's
9006 state, the thread attempts to acquire a lock for it and will not execute the transaction without
9007 having a lock on the relevant state. Because this is not a deterministic process, a miner
9008 needs to keep track of the history of which threads acquired which locks. This history
9009 creates a happens-before graph of the executed transactions, which must be transmitted to
9010 validators as a part of the block. Validators can use this graph to spawn the required threads
9011 and assign transactions to them.

9012 Another approach utilizes speculative execution in order to achieve concurrency in a lock-
9013 free manner [504]. First, all transactions are executed concurrently, assuming that there
9014 are no conflicts. When conflicts do occur, the transaction is aborted and put into a sep-
9015 arate "bin" for sequential transactions. Once the speculative concurrent execution of all
9016 transactions has completed, the sequential transactions are executed. As a result, when a
9017 transaction causes a conflict, it must be executed twice, introducing some overhead.

9018 A third approach takes advantage of software transactional memory systems (STMs)
9019 [506, 507]. In contrast with the use of locks, transactional memory assumes that simul-
9020 taneous access to state will not cause conflict; that is, it is optimistic. Therefore, threads
9021 do not need to wait for each other in order to access the same portion of the state. Dur-
9022 ing transaction execution, miners attempt to update the state but do not fully commit these
9023 changes until later. When a state change is about to be committed, the runtime environ-
9024 ment checks for conflicts based on prior commits, and the changes are reverted in case of
9025 conflict. This process is non-deterministic in the same way as using locks, so the miner
9026 generates a happens-before graph to transmit to validators in a similar way. This method
9027 results in a lot less waiting time when multiple transactions access the same piece of state
9028 but in ways that do not cause conflicts.

9029 One can also parallelize existing ledger systems by using some basic static analysis. For
9030 example, [505] defines the notion of *strongly swappable* transactions: transactions TX_1 and
9031 TX_2 are strongly swappable if the state variables that each transaction accesses are disjoint.
9032 To exploit this, any method of static analysis can be used to over-approximate the portions
9033 of state accessed in a transaction. For UTXO-based ledgers, such as Bitcoin, a simple
9034 check on the inputs and outputs is sufficient. Strong swappability then implies a partial
9035 ordering of transactions that validators can deduce themselves and that is guaranteed to be
9036 equivalent to a serial execution of transactions.

9037 Finally, [508] describes a clever solution that utilizes a variant of locking, concurrency
9038 delegation, and static analysis. Each state variable is augmented with a taint value. A state
9039 variable that has not been accessed during the current execution is considered untainted,
9040 but once a thread attempts to read or write to it, it is considered tainted with a value that
9041 identifies the thread. This is a relaxation of the idea of locking such that long wait times are
9042 avoided. Instead, if a thread attempts to access a tainted variable, the execution immediately

9043 stops. The thread then forwards the conflicting transaction to the thread that tainted the
9044 variable, which is more likely to be able to execute it. After the transaction is delegated
9045 in this way, it may still not be able to be executed, in which case it is added to a queue of
9046 sequential transactions to be executed later.

9047 This delegation process requires that transactions are initially distributed to threads in some
9048 way, ideally while minimizing conflicts. To this end, transactions are analyzed based on
9049 static information to provide "hints" without needing to execute the transactions. For ex-
9050 ample, the sending address in a transaction constitutes a good hint. A programmer can
9051 then annotate transactions with these hints, which are beneficial despite being imperfect. A
9052 primary thread uses this information to distribute transactions across worker threads, who
9053 use tainting and delegation as described above. If a tainted transaction cannot be executed
9054 by the thread it was delegated to, the worker thread sends it back to the primary. Workers
9055 let the primary thread know that they have completed execution, at which point the work-
9056 ers are shut down, and the sequential queue is executed. Finally, the primary thread labels
9057 transactions based on which worker thread they were executed by or whether they were
9058 sequential, providing that information to validators and allowing deterministic execution.

9059 **18.1.2. Zero-Knowledge Proofs and Verifiable Computation**

9060 A common technique in distributed ledger protocols is to employ zero-knowledge proofs,
9061 frequently zk-SNARKS, in order to prove that some computation was performed correctly
9062 while revealing no additional information. Depending on the exact techniques used, this
9063 can enhance privacy as well as improve the scalability of distributed ledgers by having
9064 validators verify short proofs of computation instead of performing all computations them-
9065 selves.

9066 An early application of this is the Zerocash protocol, which has been deployed in the Zcash
9067 network [511]. In Zcash, zk-SNARKS are used to create "shielded" transactions that hide
9068 the transaction's inputs, outputs, and amount. The zk-SNARK proves to the rest of the
9069 network that the transaction adhered to the state machine's rules. The proof demonstrates
9070 that the amount of Zcash transmitted in the inputs is equal to the amount of Zcash in the
9071 outputs less a fee and that the spender possesses the requisite private key to spend funds.
9072 Similar techniques can be used to add privacy to smart contracts on Ethereum or other
9073 Turing-complete platforms. Zether, for example, uses a different form of zero-knowledge
9074 proof (called Bulletproofs) to hide transaction amounts while performing fund transfers to
9075 and from smart contracts [512]. This enables things like sealed-bid auction smart contracts.

9076 Other schemes make the smart contract computations themselves more private rather than
9077 just payments [513–515]. For example, the Hawk system can be used to hide the inputs
9078 to smart contracts from everyone except a special party called the manager, which may be
9079 implemented via multi-party computation [513]. Zexe hides not just the inputs and outputs
9080 of a function to be executed but also the function itself as well as the internal state of the
9081 smart contract that executes it [514]. To see why this is valuable, one can imagine that

9082 every token on Ethereum is built with a Zerocash-like scheme that hides the transaction
9083 details of token transfers. In this case, transactions would still leak which particular tokens
9084 were involved in a transaction. An eavesdropper may not know how much of a token was
9085 transferred, but they could still deduce the token smart contract that was used. In Zexe, a
9086 shared execution environment is created in which transactions reveal no information about
9087 the offline computations performed for a smart contract, are unlinkable to other transactions
9088 by the same user and/or the same type of computation, and can be verified in constant
9089 time regardless of the complexity of the computation. Users perform the computations
9090 themselves (or, in an extension to the protocol, delegate the computations to someone else)
9091 on plaintext inputs, encrypt the inputs and outputs to the computation, and combine it with a
9092 zk-SNARK before submitting the transaction to the network. Validators only need to verify
9093 that the zk-SNARKs included in the transactions are valid but do not need to re-execute the
9094 computations.

9095 Another approach, dubbed smartFHE, uses fully homomorphic encryption to allow smart
9096 contracts where users maintain privacy over their inputs and outputs to the contract function
9097 [515]. That is, it allows computation to be performed on encrypted data supplied by one
9098 or more users, where the results of the computation remain private to those users as well.
9099 Fully homomorphic encryption allows users to submit to a smart contract a list of encrypted
9100 inputs to a function to execute as well as a zero-knowledge proof that the submitted cipher-
9101 texts are well-formed. After verifying the proof, miners will execute the requested function
9102 calls on the ciphertexts. In contrast to Zexe, the private computations in smartFHE are
9103 performed on-chain, and the smart contract's code is public. smartFHE does not provide
9104 anonymity to users interacting with the smart contracts.

9105 18.1.3. Delegating Execution

9106 Most distributed ledgers require all participants to execute every transaction. Combined
9107 with a consensus algorithm that guarantees agreement over a total ordering of these trans-
9108 actions, this is a natural way to achieve state machine replication. However, requiring that
9109 every node executes every transaction hinders the scalability of the network and reduces the
9110 privacy of the computation. It would be better if only a subset of participants were required
9111 to execute transactions rather than every node on the network.

9112 TrueBit is a scheme that uses smart contracts to outsource computation and provides eco-
9113 nomic incentives such that rational participants will execute the computation correctly
9114 [516]. As a result, network validators need not execute these computations, while the com-
9115 plexity of the computations can be significantly greater than the base layer's gas limit would
9116 typically allow. A user who would like a computation executed deposits some funds to pay
9117 for the result. Anyone can act as a solver or verifier in the system by submitting a deposit
9118 to the contract, though solvers are matched to computational tasks at random. To provide
9119 proper incentives for verifiers in TrueBit, potential verifiers need to believe that there is a
9120 real chance of finding a flawed computation from a solver. As a result, TrueBit periodically

9121 requires incorrect solutions to be submitted. Once a solver is randomly assigned to a task,
9122 they calculate both a correct and incorrect result to the computation and send commitments
9123 to both to the TrueBit contract. The contract determines whether a "forced error" should
9124 occur for this computation, in which case the solver opens the commitment to the incorrect
9125 result (otherwise, they open the correct result). Verifiers then check the solution and issue
9126 a challenge if they disagree with the solver. The verifier receives a large payment if the
9127 solution is erroneous and the forced error regime is in effect. The system will accept the
9128 result as accurate if no verifier challenges it, but if a challenge exists, a *verification game* is
9129 played.

9130 Assume that the requested computation runs in t time steps and requires at most s bits of
9131 state at any given point in the computation. Both the solver and the challenger create a
9132 mapping of each time step of the computation to its internal state at that time. Let $c > 1$
9133 be a parameter of the verification game that determines a trade-off between the number of
9134 game rounds and the amount of communication required with the underlying ledger.

9135 The verification game runs a loop that attempts to determine where in the computation
9136 the discrepancy arises. At the start of the loop, the solver picks c equally spaced state
9137 configurations based on the current range of the disputed computation. For each of these
9138 c configurations, the solver creates a Merkle tree with s leaves that correspond to each
9139 bit of the internal state and publishes the roots to the TrueBit contract. The challenger
9140 submits $i \leq c$ to the contract, where i is the first time step in the list where they disagree
9141 with the state. The TrueBit contract verifies that c Merkle roots have been submitted and
9142 that $1 \leq i \leq c$ and has the relevant party lose if a check fails. The loop begins again but
9143 only using the state configurations between the $(i - 1)$ -th and i -th configurations from the
9144 prior round of the loop. Eventually, this loop converges to the first disputed step over the
9145 whole computation, and this step is then verified by the smart contract. If this is the e -th
9146 time step, the solver will submit to the contract paths from time $e - 1$ and e 's Merkle roots
9147 to the leaves that contain the relevant Turing-machine variables needed to check the one
9148 computational step. The winner of the verification game is determined by whether these
9149 paths are valid and the computation step was executed properly.

9150 Another technique, employed in Hyperledger Fabric, is to decouple the ordering of transac-
9151 tions and their execution [517]. For most ledgers, a consensus algorithm first orders batches
9152 of transactions that are then executed serially, updating the state when execution completes.
9153 In contrast, Fabric uses the *execute-order-validate* (EOV) paradigm, where transactions are
9154 first executed, then ordered by the consensus algorithm, and finally validated for consis-
9155 tency to remove conflicts before updating the state. These approaches are compared in Fig-
9156 ure 55. Because potentially invalid transactions will be ordered, throughput can be harmed
9157 by the inclusion of duplicate or conflicting transactions [518]. To mitigate the performance
9158 impact of this, client authentication and other techniques may be used.

9159 Because not all parties need to care about every smart contract on the system, each contract
9160 in Fabric specifies a set of *endorsers* that execute the transactions and an endorsement

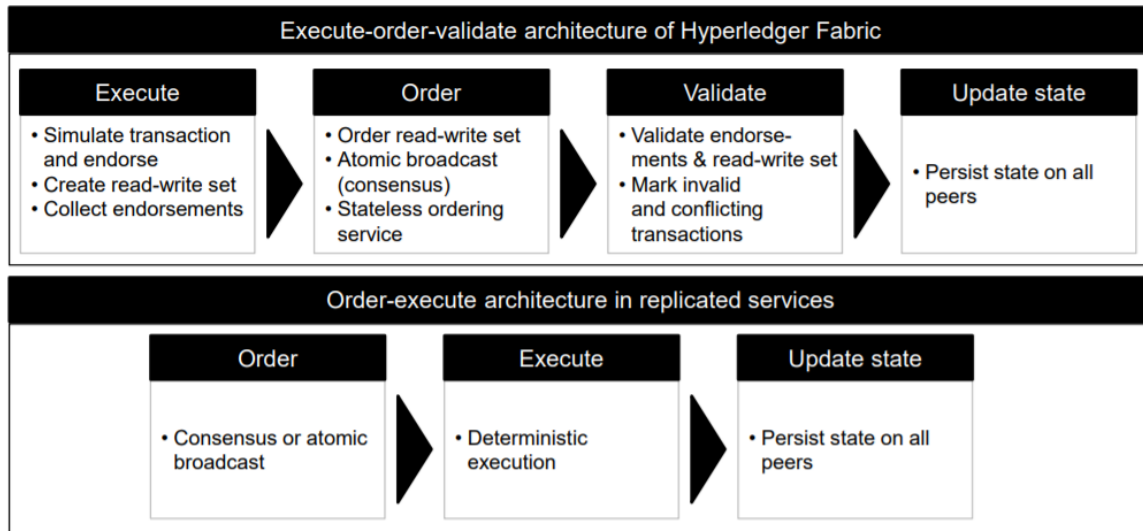


Fig. 55. Comparison between execute-order-validate and order-execute. [519]

9161 policy that provides the threshold that must agree or any specific peers to trust. In the
9162 execution phase, clients sign transactions and send them to the relevant endorsers. The
9163 endorsers create a *read set* and *write set* of the pieces of state that were read from or
9164 written to by the transaction and reply to the client with a signed endorsement of these sets
9165 and the result. Clients collect endorsements until the endorsement policy is met. The client
9166 bundles up these endorsements and creates a signed transaction to send to the ordering
9167 service, which uses a consensus algorithm to totally order transactions. In the validation
9168 phase, replicas verify that transactions meet their endorsement policies and are marked
9169 invalid if they fail. Replicas then sequentially check transactions for conflicts in their read
9170 and write sets and are marked as invalid if conflicts arise. Finally, the state is updated by
9171 the valid transactions, the block is committed, and clients are notified.

9172 Arbitrum is a mechanism that allows a smart contract to be defined as a virtual machine
9173 where computation is performed off-chain, and both the code and the data for the smart
9174 contract remain off-chain for improved privacy [520]. When a contract is deployed, the
9175 creator selects a set of VM managers who are in charge of executing the code. As long as
9176 as one VM manager is honest, security is maintained. Transactions are sent to the VM
9177 managers, who execute them and sign the result. If there is unanimous agreement among
9178 VM managers, their signatures are published to the underlying ledger, and replicas are
9179 only required to verify the signatures instead of performing the requested computation. If
9180 parties behave irrationally, a TrueBit-like verification game is played that resolves disputes
9181 by examining the execution of a single instruction.

9182 ACE is a scheme that is intended to allow highly complex contracts that may take min-
9183 utes to execute on top of an underlying ledger, such that computations themselves occur
9184 off-chain and asynchronously [521]. As with Arbitrum, contract issuers appoint a set of

9185 service providers to execute the code. However, ACE allows flexible thresholds instead of
9186 requiring unanimous agreement and allows different smart contracts to safely interact with
9187 each other in a composable way. In ACE, blocks include separate ordering and result sec-
9188 tions. Miners will pre-order transactions, and service providers execute transactions from
9189 this section off-chain when their contracts are called. Service providers sign and broadcast
9190 the results of the execution over the network and create a special state change transaction.
9191 If a sufficient threshold of service providers signs off on the state change, miners include it
9192 in the results section of a block, thereby committing the transaction. By having each execu-
9193 tion result reference the preceding contract state, consistency can be maintained even when
9194 the threshold is less than half of the designated service providers. A similar scheme can
9195 also be implemented on Bitcoin and other networks with less expressive scripting support
9196 [522].

9197 A different approach employed in LazyLedger is to have the blockchain order transactions
9198 and make them available but shift the burden of executing and validating transactions to the
9199 clients who care about them [523]. For consensus participants, verifying a block consists
9200 solely of verifying data availability, either by downloading the whole block or by using
9201 probabilistic techniques like those discussed in Section 15.5. The LazyLedger blockchain
9202 just stores messages for smart contracts, but the contract logic itself is off-chain, can be
9203 written in any language or environment, and can be changed without a hard fork. This ar-
9204 chitecture allows the users of a smart contract to ignore messages related to other contracts
9205 and only execute or validate state transitions that they care about. While this can dramati-
9206 cally improve the scalability of computation, there is a risk of denial-of-service attacks
9207 on smart contracts if malicious clients create many invalid transactions, so LazyLedger
9208 may work better in a permissioned system. Another limitation is that it is challenging to
9209 construct light clients for smart contracts using this architecture.

9210 **18.2. Layer 2 Protocols**

9211 Generally speaking, distributed ledgers scale somewhat poorly because every single node
9212 on the network is required to process every transaction. This consumes large amounts of
9213 storage and bandwidth, and verifying a large number of signatures can be taxing on a CPU.
9214 It is preferable for transactions to be processed locally by the parties involved rather than
9215 by every single node. This is where "layer 2" techniques can be useful. Smart contracts can
9216 rely on a small amount of global state in order to execute many transactions "off-chain,"
9217 using the ledger itself only for settlement and dispute resolution.

9218 **18.2.1. Payment and State Channels**

9219 The most basic second layer scaling technology is the *payment channel*. Once set up, the
9220 participants can send a potentially limitless number of transactions back and forth between
9221 each other without needing to touch the underlying ledger except for final settlement. In-
9222 stead of issuing a separate transaction for each morning coffee for a month, a customer

9223 could open a payment channel with the coffee shop on-chain at the beginning of the month,
9224 purchase each coffee off-chain, and close the channel at the end of the month. In this case,
9225 instead of having 30 different transactions processed globally, only two are required. The
9226 same idea can be applied to smart contracts in general rather than just payments, in which
9227 case it is called a *state channel*.

9228 A variety of payment channel constructions exist [524, 525]. This section describes the
9229 basics of the Lightning Network (LN) payment channel construction [526, 527]. An LN
9230 channel consists of the following transactions (where the notation is from [524]):

9231 • **Funding:** This transaction opens the channel and deposits funds into it, similar to a
9232 prepaid debit card. It is signed by both parties and creates a 2-of-2 multisig UTXO.

9233 • **Commitment:** These transactions spend from the multisig output and thus require
9234 signatures from both channel participants, who receive their counterparty's signature
9235 in advance of usage. These transactions have two outputs: the first sends coins to the
9236 broadcaster, and the second sends coins to the counterparty. The counterparty can
9237 spend the second output by signing it. The first output has two potential spending
9238 conditions:

9239 1. A signature from the broadcaster plus a *relative lock time* such that the commit-
9240 ment transaction has depth λ in the ledger before being spendable

9241 2. A signature from the counterparty plus a preimage $S_{X,j}$ of revocation hash
9242 $h_{X,j} = H(S_{X,j})$

9243 • **Revocable Delivery:** These transactions send coins to their broadcaster and require
9244 the signature of the broadcaster and for the commitment transaction they spend from
9245 to have a depth of λ blocks in the ledger.

9246 • **Delivery:** These transactions immediately send the counterparty their share of coins
9247 when signed by the counterparty that did not broadcast their commitment transaction.

9248 • **Breach Remedy:** These transactions can only be broadcast after a revoked commit-
9249 ment transaction is included on-chain. They require a signature from the counterparty
9250 that did not broadcast the revoked commitment transaction, as well as the preimage
9251 $S_{X,j}$ of the revocation hash from the commitment transaction. The breach remedy
9252 transaction has no relative lock time, so the counterparty can immediately take all of
9253 the channel's coins.

9254 A channel can be closed either cooperatively or unilaterally. If both parties are online and
9255 cooperative, they can simply agree upon a transaction that sends each party their correct
9256 share. There is also a dispute process for when one of the parties is not online: either party
9257 can broadcast their most recent commitment transaction and revocable delivery transaction
9258 to recover their funds, while the counterparty then broadcasts their delivery transaction to
9259 receive their own share. If either party publishes an outdated commitment transaction, the

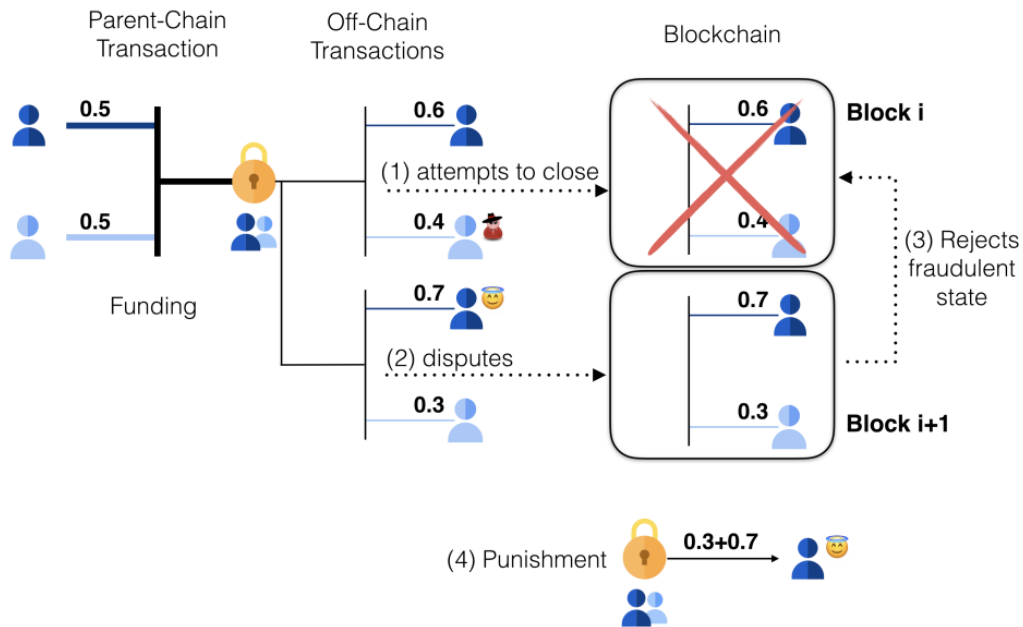


Fig. 56. Lightning network channel closing. In this example, the channel is updated twice off-chain. A malicious channel party tries to close the channel with an outdated state favorable to them by broadcasting the corresponding commitment transaction. The malicious party must wait λ blocks before their portion is spendable. The honest counterparty sees the malicious transaction confirmed on the blockchain before the timelock expires, and submits a breach remedy transaction to claim all the funds in the channel - including the coins of the malicious party. [528]

9260 other party can punish them by taking all of the coins in the channel with their breach
9261 remedy transaction. An example of this is shown in Figure 56.

9262 Payment channels would have limited utility if they only allowed transactions between di-
9263 rect channel partners. This would require consumers to open channels with every business
9264 they frequent. Luckily, the Lightning Network allows payments to flow across a network
9265 of open payment channels: if Alice has an open channel with Bob, and Bob has an open
9266 channel with Carol, then Alice can pay Carol using Bob as an untrusted intermediary. This
9267 process can be seen in Figure 57 and utilizes HTLCs, as described in Section 16.1. The
9268 final recipient chooses a preimage, and sends its hash to the original sender. The parties
9269 along the payment path set up HTLCs with each other using this as the hashlock, with
9270 timelocks getting shorter with each payment channel the path crosses. The preimage is
9271 then passed backwards along the path in order to allow each payment to complete.

9272 This payment method utilizing HTLCs is susceptible to the *wormhole attack*, which allows
9273 colluding users to steal transaction fees from intermediaries [530]. If Alice and Carol
9274 colluded (or were, in fact, a single entity), steps 6 and 7 in Figure 57 could be skipped, and




Fig. 57. Lightning Network payment. The notation $HTLC(Alice, Bob, y, \alpha, t)$ means that: (1) If timeout t expires, Alice can get back the α coins she locked, and (2) If Bob reveals a value x such that $H(x) = y$ before timeout t , then Alice will pay α coins to Bob. [529]

9275 Carol could directly share the secret preimage with Alice instead of Bob. In this case, Bob's
9276 timelock will be left to expire, but Alice can complete the payment afterward because her
9277 timelock is longer. In the end, Bob will not collect his fee, leaving Alice and Carol with
9278 extra funds. This vulnerability stems from the use of the same preimage across the entire
9279 payment but can be solved using a cryptographic primitive called adaptor signatures. An
9280 adaptor signature is essentially a promise that the publication of an agreed-upon signature
9281 will reveal a secret value.

9282 While payment channels are useful for payments, more general use cases may require the
9283 use of state channels, which are like payment channels but for arbitrary smart contracts.
9284 In many cases, these take advantage of Turing-complete state machines, such as the EVM.
9285 However, Bitcoin-compatible state channels have also been proposed [531]. These use
9286 adaptor signatures to create generalized state channels capable of executing any operation
9287 that the underlying network (e.g., Bitcoin) supports but off-chain. Turing-completeness is
9288 also sometimes required for certain types of payment channels, such as Perun virtual chan-
9289 nels, where once set up, intermediaries are no longer required to be online for payments
9290 between endpoints [532]. Virtual channels were later extended to the UTXO model using
9291 only signatures and timelocks in order to be Bitcoin-compatible [533, 534]. Perun was
9292 extended to a multi-hop environment and for generalized state channels in [535] (which in-
9293 troduced state channels contemporaneously with [536]). These virtual state channels were
9294 further extended to allow contracts with more than two parties in [537].

9295 General state channels operate very similarly to payment channels. Alice and Bob can
9296 execute smart contract code off-chain by deploying a state channel, G . First, both parties
9297 agree to an initial state G_0 of G and then exchange signatures on the tuple $(G_0, 0)$, where
9298 the second item is a sequence number. The smart contract is then executed over time by
9299 exchanging further tuples of an agreed-upon state and sequence number. Let (G_s, s) be the
9300 most recent mutually signed state and sequence number. If Alice wants to call a function
9301 f of the smart contract using input x , she will execute the function call locally using the
9302 current state G_s . Next, she sends Bob the signed tuple $(G_{s+1}, s+1)$ along with f and x . If
9303 Bob believes the computation was accurate, he replies with his signature over $(G_{s+1}, s+1)$.
9304 The state channel can be closed cooperatively using both parties' signatures over the most

9305 recent state. If one party is uncooperative, the contract will accept a new mutually signed
9306 state with a higher sequence number during a dispute period, allowing an honest user to
9307 ensure that the most recent state is committed on-chain.

9308 An early state channel proposal is Sprites [538], which requires a form of globally shared
9309 state that Bitcoin does not provide. Sprites, when used for payments, implements HTLCs in
9310 a globally shared smart contract such that the expiration time and hash-lock can be enforced
9311 in the single smart contract instead of along each channel individually, so collateral need
9312 not be locked for as long along a payment path. If an LN payment is transferred across a
9313 path of L channels, and Δ is the amount of time before a transaction can be committed on
9314 chain, then collateral must be locked for $\Theta(L\Delta)$ time. Sprites reduces the collateral cost
9315 to $\Theta(L + \Delta)$ by using a smart contract (dubbed the PreimageManager) that logs statements
9316 along the lines of "the preimage R of hash $H = H(R)$ was recorded on the ledger before time
9317 $T_{expiration}$." Sprites contracts have a dispute handler that will query the PreimageManager
9318 contract to determine whether R was revealed on time, ensuring that disputed channels
9319 close in a consistent way using the single time $T_{expiration}$.

9320 One of the more significant shortcomings of payment and state channels is that participants
9321 must periodically go online and check the underlying ledger to see if their channel coun-
9322 terparty has published an old state. If so, they need to initiate a dispute in order to prevent
9323 being defrauded. To reduce the impact of this requirement, the idea of *watchtowers* has
9324 been proposed [539–543]. Watchtowers act as an online monitor that can handle disputes
9325 on behalf of a channel party while the party is offline. Watchtower constructions have a vari-
9326 ety of trade-offs: some scale poorly, may not be incentive-compatible, may harm privacy,
9327 or may require more complicated channel constructions to work properly. Another alterna-
9328 tive is to design state channels that remain secure under asynchrony, such as Brick, which
9329 only allows unilateral channel closure after getting approval from an external committee
9330 [544]. Note that this section only scratched the surface of the issues involved in payment
9331 and state channels. Most known attacks and privacy issues have not been discussed here
9332 nor were the complexities of network topology and payment routing explored.

9333 18.2.2. Plasma and Rollups

9334 Other second layer technologies include Plasma and an idea called "rollups." Like
9335 sidechains (Section 16.3), these schemes involve running a separate blockchain system
9336 in parallel to a primary, parent ledger. Unlike sidechains, however, the security of transac-
9337 tions on these chains is derived from the underlying parent chain, and they are not required
9338 to run and secure their own consensus algorithm.

9339 A Plasma chain is a blockchain whose security is anchored to its parent chain but where
9340 security is maintained by having users submit fraud proofs in order to resolve disputes
9341 (fraud proofs are discussed in Section 15.5). The entity or entities running the Plasma
9342 chain submit block headers – including a commitment to the state of the Plasma chain – to
9343 a parent chain smart contract (e.g., on Ethereum). If anyone detects invalid state transitions

9344 on the Plasma chain, they can dispute the block by submitting a fraud proof along with a
9345 bond to prevent misbehavior. If the fraud proof is valid, the Plasma chain is rolled back
9346 past the offending block, and the block creator is penalized.

9347 As with payment and state channels, withdrawing funds from a Plasma chain requires a
9348 delay. In this case, the delay is to allow other parties to submit fraud proofs that challenge
9349 the blocks in which those withdrawals were processed. Plasma chains operate in the UTXO
9350 model, and an exiting party submits a bitmap of UTXOs to withdraw funds. Anyone else
9351 may submit a bonded fraud proof that challenges the original bitmap to prove that some of
9352 those funds have already been spent. Several Plasma variants have been proposed, often
9353 differentiating themselves based on their withdrawal mechanism [545–548]. For example,
9354 the original Plasma [545] and the simplified "Minimum Viable Plasma" proposal [546]
9355 require each party to individually submit an exit transaction, and they are prioritized based
9356 on how old the relevant UTXOs are. This prevents invalid withdrawal requests from being
9357 processed before valid ones that are less likely to come "out of nowhere." Alternatively, the
9358 "More Viable Plasma" proposal [547] prioritizes withdrawals based on the youngest input
9359 referenced in the exit transaction.

9360 The constructions discussed so far have a variety of problems. Besides the user experi-
9361 ence issue of needing to wait for a full challenge period before withdrawals are processed,
9362 Plasma users need to keep track of and verify the entire Plasma chain in order to detect ma-
9363 licious behavior to initiate an exit in the first place. Further, it may be the case that all users
9364 of a Plasma chain need to exit at once if the Plasma chain operator becomes unavailable
9365 and does not serve relevant data to users. This "mass exit" scenario requires the entire state
9366 of the Plasma chain to be dumped onto the parent ledger, which could cause significant
9367 congestion and prevent fraud proofs from being processed on time.

9368 An alternative construction, Plasma Cash, resolves some of these issues [548]. Plasma Cash
9369 reduces users' data-checking requirements by using non-fungible tokens (NFTs) and sparse
9370 Merkle Trees. With this, users need only keep track of their own coins rather than the whole
9371 chain. The recipient of transactions on a Plasma Cash chain is responsible for checking
9372 that the coins being spent have a valid history on the chain based on proofs supplied by
9373 the spender. A limitation of Plasma Cash is that it only works with fixed denominations,
9374 requiring each NFT to be spent in full in any given transaction. Unfortunately, both Plasma
9375 Cash and the original (fungible) Plasma are unable to prevent an adversary from forcing
9376 honest users to take some involuntary, potentially expensive (in terms of transaction fees)
9377 on-chain action, be it a mass exit for fungible Plasma or non-constant sized exits for honest
9378 users on Plasma Cash [549].

9379 Rollups are an alternative to Plasma that address the data availability problem (and, thus,
9380 mass exits) by taking all of the transactions that occur on the rollup chain and committing
9381 some of the transaction metadata to the parent chain. In the Ethereum context, this is done
9382 by taking this "rolled up" data and posting it in a transaction's *calldata* field, which is a
9383 read-only portion of Ethereum transactions that supply function call arguments. Calldata is

9384 vastly cheaper than typical blockchain storage, making this approach efficient. By ensuring
9385 that this transaction data is always available, user liveness requirements and data availabil-
9386 ity assumptions can be dispensed with, thus improving substantially over Plasma (and state
9387 channels). Because transaction data is available on-chain and consensus is maintained over
9388 that rollup data, any user can process the full rollup when desired and thus detect fraud or
9389 initiate withdrawals. Further, by resolving the data availability problem, assets need not be
9390 mapped to owners, which allows rollups to be general-purpose. Instead of just handling
9391 payments, the full EVM can be run inside of a rollup [550].

9392 Two types of rollups have been proposed: optimistic rollups and zk-rollups. In an opti-
9393 mistic rollup, all blocks are simply assumed to be valid unless proven otherwise, and fraud
9394 proofs can be submitted to challenge an invalid block. This means that, as with Plasma,
9395 withdrawals require a delay such that challenges can be submitted and the rollup chain can
9396 be rolled back to deal with fraud. For zk-rollups, each rollup block comes with a SNARK
9397 that proves that the included transactions were properly executed. As such, zk-rollups al-
9398 low withdrawals nearly instantly, but they create a burden on a chain operator, who must
9399 create an expensive proof. Rollups are likely to figure prominently in Ethereum’s attempts
9400 to scale.

9401 **19. Incentives**

9402 In distributed ledger systems – especially permissionless environments – an important chal-
9403 lenge is to have the incentives of the various participants aligned such that rational partic-
9404 ipants behave honestly. Incentives have been discussed throughout this document already,
9405 but this section will introduce some of the more nuanced aspects of how incentives interface
9406 with the security of the network.

9407 **19.1. Block Rewards: Subsidies and Transaction Fees**

9408 The most important and frequently discussed source of incentives in distributed ledger sys-
9409 tems is the block reward that is earned when a block is produced. The block reward consists
9410 of one or both of a block subsidy (newly minted cryptocurrency units) and transaction fees,
9411 which clients pay in order to prioritize their transaction for inclusion in a block. Generally
9412 speaking, block rewards should be allocated proportionally to hash rate or stake in order to
9413 maintain desirable properties, such as being resistant to Sybil attacks and miner collusion
9414 [551, 552].

9415 The block subsidy is a result of the monetary policy of the network and provides an intuitive
9416 method of distributing newly minted coins. In the Bitcoin network, for example, the block
9417 subsidy began as 50 bitcoin per block but is halved every four years until a total of 21 mil-
9418 lion bitcoin exist, which should occur by approximately year 2140. In contrast, Ethereum’s
9419 initial supply of 72 million ether was distributed in a sale and placed in the genesis block,
9420 while the initial block subsidy was 5 ether per block (ignoring uncle blocks). This has been

9421 reduced several times via hard fork. Some proof of stake coins will distribute the entire
9422 supply in the genesis block and have no subsidy. Others, such as Monero, reduce the is-
9423 suance up to a point and then have a *tail emission*, where a constant quantity of monero is
9424 minted per block. This is sometimes called a *disinflationary* monetary policy.

9425 As of 2023, block subsidies dominate the block reward for most cryptocurrencies, and
9426 transaction fees make up a trivially small portion of the total rewards in networks other
9427 than Bitcoin and Ethereum. However, since block subsidies in most networks decrease
9428 over time, transaction fees are expected to become a more and more important compo-
9429 nent of the block reward. The implications of having transaction fees become the primary
9430 means of funding network security are worth exploring. If transactions are being broadcast
9431 at approximately the same pace that they are put into blocks, miners who order transac-
9432 tions based on the order they are seen will generate revenue substantially less than those
9433 who process transactions greedily based on fees [553]. As a result, if a miner is deciding
9434 between several conflicting transactions to include in their block, they are strongly incen-
9435 tivated to mine the transactions with the highest fee rate. This justifies certain policies,
9436 such as *replace-by-fee*, where clients can replace transactions in other nodes' mempools by
9437 bidding up their transaction fee enough to cover the additional bandwidth that the trans-
9438 action will consume. This also suggests that the censorship resistance for transactions in
9439 most cryptocurrencies is primarily a result of the transaction fees provided because the fee
9440 is the opportunity cost to miners for censoring the transaction. A related issue is that miners
9441 may have little incentive to propagate high-fee transactions because they would prefer to
9442 claim the fees for themselves [554]. However, doing so increases block propagation time
9443 and thus increases the risk that the miner's block becomes stale.

9444 One of the earliest insights into Bitcoin's fee economics was that if there is no maximum
9445 block size limit, transaction fees will fall toward zero, resulting in negligible security for
9446 the system [555]. This is generally true if there is an abundance of block space. If there
9447 is not a competitive market for this space, then fees will tend toward zero. Unless there is
9448 a persistent block subsidy (or sufficiently limited block size), the system may experience
9449 instability and low security. This is discussed in more detail in the next section.

9450 **19.1.1. The Mining Gap and (the Absence of a) Block Subsidy**

9451 The block subsidy provides a fairly consistent reward for miners, whereas transaction fees
9452 will vary from block to block based on the supply of and demand for block space. When
9453 the block subsidy dominates transaction fee income, the reward for any given block is
9454 nearly constant. On the other hand, as the subsidy becomes dominated by transaction fees,
9455 the variance in block rewards becomes significant. This variance may cause instability in
9456 consensus by encouraging miners to abandon the longest chain rule and perform adversarial
9457 mining strategies instead.

9458 Carlsten et al. investigated the rational behavior of miners in a regime with no block sub-
9459 sidy, transaction fees accruing in the mempool at a constant rate, no latency in transaction




Fig. 58. Mining gap. The potential block reward is equal to the average transaction fee (a) multiplied by elapsed time (t) since the prior block was mined. [556]

9460 or block propagation, and – crucially – miners having the space to include all available
9461 transactions in the next block if desired [556]. That is, the maximum allowed block size
9462 is significantly larger than what is in the mempool at any given time. In this scenario, it
9463 is possible for a *mining gap* to form: immediately after a block is mined, the mempool is
9464 empty and there are no transaction fees to put into the next block, so it makes sense for
9465 miners to shut down their machines until there is enough fee income to be worth the elec-
9466 tricity and other operating costs. The idea behind the mining gap is illustrated in Figure
9467 58.

9468 This mining gap also incentivizes miners to *undercut* each other instead of moving the
9469 chain forward. If a block mined at height X claimed all available transaction fees in the
9470 mempool, then it makes little sense for competing miners to build off of that block at height
9471 $X + 1$. Instead, they can gain an advantage by mining a different block at height X that only
9472 claims a portion of the total transaction fees. This leaves revenue available for other miners
9473 to be incentivized to mine at height $X + 1$ on top of the adversary's block. Undercutting
9474 becomes an even more significant problem when nodes have heterogeneous internet con-
9475 nection speeds [237]. Miners may undercut each other more and more aggressively, forking
9476 with smaller and smaller fees and leaving more and more remaining available for others to
9477 mine on top of. This aggressive undercutting can lead to transactions failing to become
9478 confirmed as the chain fails to advance forward. An example of undercutting is shown in
9479 Figure 59.

9480 Further, selfish mining is more profitable in this regime. A selfish miner's block on average
9481 will tend to be larger and include more transactions and, thus, rewards than when block
9482 subsidies exist and most block rewards are similar. This is because when the selfish miner
9483 has a long lead, the blocks that they mine are disproportionately large since it tends to take
9484 longer for the adversary to mine a block than the rest of the network. During this time,
9485 more transaction fees accumulate, and the selfish miner can still include transactions that
9486 were included in competing blocks on the public chain.

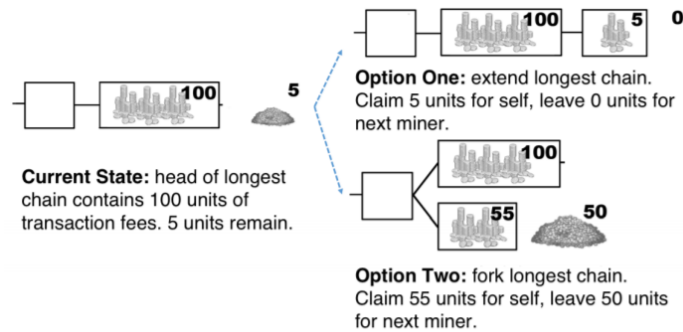


Fig. 59. Undercutting attack. [556]

9487 The mining gap was further studied in [557], where the block subsidy was not reduced
9488 to zero but where transaction fees were nevertheless a significant component of the block
9489 reward. To this end, they define a metric – EBRR – that is the ratio between the expected
9490 base reward and the expected accumulated fees. Here, the base reward includes whatever
9491 transaction fees remain in the mempool after the most recent block has been mined. The
9492 expected accumulated fees is the expected block interval multiplied by the rate at which
9493 new transaction fees are added to the mempool. A mining gap is expected to form when
9494 the EBRR is approximately six but can be avoided if sufficiently high. An interesting result
9495 of that paper is that the consequences of the mining gap differ depending on the size of the
9496 miner. Because large miners are more likely to mine the next block and get a reward, they
9497 are willing to wait longer before turning their machines back on and cut electricity expenses
9498 while still having a good chance of winning. On the other hand, very small miners with
9499 low hash rate have almost no chance of winning unless they mine continuously and must
9500 therefore pay the electricity costs in order to maintain what little chance they have. In
9501 extreme scenarios, this could reduce mining utilization by up to 90%. There are several
9502 potential solutions to the issues related to the mining gap and undercutting attacks, each of
9503 which is intended to reduce the variance in block rewards:

- 9504 1. The monetary policy in the network can include a block subsidy that is large enough
9505 to dominate transaction fees in perpetuity. This solution is simple but results in per-
9506 manent inflation of the underlying asset.
- 9507 2. The maximum block size can be sufficiently constrained such that, given the demand
9508 for block space, there is a nearly perpetual transaction backlog in the mempool. This
9509 option may increase transaction commit latency and fees for users. This was empir-
9510 ically shown to be effective in [558], where miners earn close to their fair share of
9511 rewards and avoid undercutting so long as they leave enough fees available in the
9512 mempool for the next miner.
- 9513 3. The reward scheme can be designed such that the transaction fees for transactions
9514 in a given block are shared among miners of future blocks. Unfortunately, this has

9515 several drawbacks that will be discussed in Section 19.3. Most significantly, it fails
9516 to solve the problem in the common case where out-of-band payments can be made
9517 to miners for transaction inclusion.

9518 4. In some cases, the state machine may provide mechanisms that ensure that a transac-
9519 tion is only valid if included after particular blocks, such as the `lock_time` field in
9520 Bitcoin transactions. By default, Bitcoin wallets set the transaction `lock_time` to the
9521 height of the next block that they expect to be created. That is, if the wallet is aware
9522 of the chain tip being at block height X , the new transaction will only be considered
9523 valid in blocks of height $X + 1$ or higher.

9524 Finally, while the mining gap issue does not exist for proof-of-stake ledgers due to the lack
9525 of operational expenses like electricity, undercutting attacks are still a concern.

9526 19.2. State Machines, Incentives, and Security

9527 There are security ramifications to the interactions between the consensus algorithm and
9528 the state machine that is being replicated through it. The particulars of the state machine
9529 and the types of transactions that clients want executed can cause consensus instability
9530 through a variety of mechanisms. In this section, three issues will be discussed:

- 9531 1. Miner extractable value
- 9532 2. Mispriced computations
- 9533 3. The Verifier's Dilemma

9534 *Miner extractable value* (MEV) is the total reward that a block producer can gain by ma-
9535 nipulating the order of transactions in a particular time frame [559, 560]. Examples include
9536 front-running attacks (introduced in Section 7.1) and the undercutting attacks mentioned in
9537 the previous section. In many cases, MEV is captured by parties that are not block pro-
9538 ducers themselves, but miners are best positioned to capitalize on the existence of MEV
9539 given their privileged position in transaction ordering (they are essentially a "rushing ad-
9540 versary"). The ability to choose the order of transactions in the ledger is valuable in itself
9541 and can provide an additional source of revenue that exists exogenously to the consensus
9542 protocol. If MEV is significant enough, it may be used as a way of subsidizing chain reorg
9543 attacks. Empirical evidence suggests that MEV is a significant and growing phenomenon
9544 [559–563]. It has appeared in decentralized exchanges (DEXes), collectible games, gam-
9545 bling, name services, and initial coin offerings (ICOs) and was used by a prominent mining
9546 pool, F2Pool, to front-run the Status ICO [560]. While still a small minority, thousands
9547 of Ethereum blocks contain more MEV than honest block rewards [559, 563]. There are a
9548 variety of ways of exploiting MEV, and the most well-known are shown in Figure 60.

9549 Front-running comes in two varieties: destructive and cooperative. When destructive front-
9550 running occurs, an attacker's transaction is placed in front of a victim transaction in a



Fig. 60. Front-running and Miner Extractable Value. Attacker transactions are denoted T_A and marked in red, while the victim transaction is denoted T_V and is marked in blue. When destructive front-running occurs, the victim's transaction becomes invalid and fails to execute.

9551 block, causing the victim transaction to become invalid. In cooperative front-running, the
 9552 displaced victim transaction remains valid and executed but only after the attacker's trans-
 9553 action executes. Back-running is the opposite: the attacker wishes to have their transaction
 9554 placed immediately after a particular target transaction. Finally, an attacker may try to clog
 9555 the network with transactions in order to push a victim transaction into a future block. Con-
 9556 tracts that have deadlines may create an incentive to clog the blockchain with transactions
 9557 in order to suppress another user from having their transaction confirmed on time. For ex-
 9558 ample, many of the layer 2 schemes discussed in Section 18.2 require parties to challenge
 9559 fraudulent transactions within a given time limit.

9560 These MEV building blocks may be profitable to exploit on their own, or they can be com-
 9561 bined into more advanced attacks. For example, in a *sandwich attack*, a trader on a DEX
 9562 will both front and back-run a victim transaction, T_V , simultaneously. The adversary listens
 9563 on the network until a transaction arrives in their mempool that – after being executed –
 9564 is expected to cause a change in price of an asset on the exchange. Say T_V is expected to
 9565 increase the price of an asset. The attacker will first broadcast transaction TA_1 , which buys
 9566 the asset and uses it to cooperatively front-run T_V . This way, the asset is purchased before
 9567 T_V raises the price. Next, the adversary broadcasts TA_2 , which sells the asset and back-runs
 9568 it immediately behind T_V .

9569 In addition to DEXes, common financial applications built into smart contracts are lending

9570 and debt protocols. Typically, borrowers are required to over-collateralize their debt by
9571 locking, say, 150% of the value that the borrower wants to be loaned. For example, a user
9572 may lock \$1,500 of ether in a smart contract that then provides them \$1,000 worth of a
9573 stablecoin. If the collateral value decreases below \$1,500 and the borrower fails to deposit
9574 more collateral, the original collateral is made available to liquidators to purchase at a
9575 discount in order to repay the debt. This can be exploited in multiple ways by liquidators.
9576 For instance, if the execution of a block creates a liquidation opportunity, the adversary
9577 can create transaction T_A that liquidates the collateral and keep bidding up its fee in order
9578 to destructively front-run competing liquidators. Alternatively, the adversary may spot
9579 a transaction T_V by an off-chain pricing oracle that creates a liquidation opportunity by
9580 adjusting the on-chain market price between the collateral asset and the borrowed asset. In
9581 this case, they can back-run T_V with their liquidation transaction. This has the advantage
9582 of avoiding a transaction fee bidding war.

9583 Addressing MEV is challenging, especially in open blockchains where the ledger is public.
9584 Developers should be aware of MEV when designing applications and ideally design the
9585 application in such a way that transaction ordering is unimportant. Another possibility is to
9586 adjust the consensus layer so as to remove the ability of miners to arbitrarily order transac-
9587 tions. This is the approach taken in [564], which extends the Aequitas protocol mentioned
9588 in Section 7.1 to the permissionless setting. Finally, cryptography may be used to reduce
9589 an adversary's visibility of transactions in order to reduce the information available for
9590 the adversary to exploit. For example, HoneyBadgerBFT (Section 6.1) orders encrypted
9591 transactions, so an adversary would need to take advantage of meta-information in order
9592 to profitably front-run. Recently, a new cryptographic primitive called "multi-party timed
9593 commitments" has been proposed as a way for application developers to address front-
9594 running [565].

9595 Another consensus security issue that arises due to the state machine is that the cost of
9596 executing certain operations may not accurately reflect the amount of computational effort
9597 or other resources required for the execution. Not every operation performed in the virtual
9598 machine has identical costs. Verifying a signature, for instance, is more computationally
9599 intensive than performing a simple addition, and the (gas) price of opcodes should reflect
9600 the costs of running them as accurately as possible. Finding accurate costs is a challenging
9601 task, particularly because these costs may not remain static over time or consistent across
9602 machines. The primary risk of mispriced state machine instructions is that they can en-
9603 able network-wide denial-of-service attacks by creating transactions that are exceedingly
9604 difficult to verify.

9605 The pricing of EVM opcodes has been studied empirically, and problems have been iden-
9606 tified [566–568]. The main factor that seems to make gas prices improperly tuned is the
9607 role of state storage for smart contracts. The system state is too large to store in memory
9608 on most systems, so disk access is often required when a transaction must read or write to
9609 smart contract state. Not only is this a slow operation, but there can be high variability in
9610 its execution time depending on whether the data is cached already or if using an HDD or

9611 SSD. This implies that high-end miners are advantaged relative to smaller hobbyist miners
9612 [566]. The mispricing opens up the possibility of malicious contracts that are extremely
9613 slow to verify. For example, [567] showed how to construct an Ethereum block that would
9614 take 93 seconds to verify – multiples longer than the expected block interval.

9615 Opcode pricing was actually exploited on the live Ethereum network in September 2016,
9616 when an attacker constructed a malicious smart contract that required reading large amounts
9617 of state but where those operations were severely underpriced. This attack vector was fixed
9618 by increasing the gas cost of a variety of storage-accessing opcodes in EIP 150 [569]. A
9619 similar vulnerability existed on the live network but was fixed in April 2021 [570]. Bitcoin’s
9620 more limited scripting support provides only marginal protection from this. It is possible
9621 to construct Bitcoin transactions that take an exceedingly long time to verify as well [571,
9622 572]. One mitigation employed in Bitcoin is that developers have created standardized
9623 transaction templates for the most common types of transactions, and nodes have a policy
9624 where they refuse to propagate non-standard transactions. This means that in practice, the
9625 attacker would need to be a miner or collude with a miner.

9626 A third security issue due to incentive alignment problems from state machines is called the
9627 *Verifier’s Dilemma* [573]. When an honest node receives a newly mined block, it will verify
9628 that the block is valid before forwarding it to its peers and potentially mining on top of it
9629 himself. Unfortunately, this makes nodes susceptible to resource exhaustion attacks. For
9630 example, it is possible to design smart contracts that are slow to execute but that result in
9631 predictable state changes without the designer needing to execute it [574]. Fully verifying
9632 the block also imposes a delay for a miner, who would prefer to assume that the block is
9633 valid and immediately start working on the next block. That is, skipping verification can
9634 provide a revenue advantage for a miner by allowing them to find the next block more
9635 quickly. As a result, some miners may not validate blocks, which opens up the possibility
9636 of invalid blocks being accepted as part of the canonical consensus chain.

9637 More broadly, the incentive to verify the correctness of transactions may not be sufficiently
9638 strong. In particular, as blocks grow larger and computational complexity increases, val-
9639 idators are more and more likely to skip verification to gain some advantage. If Ethereum
9640 were to raise the block gas limit significantly, non-verifying miners could gain substantially
9641 at the expense of those who verify [575]. Other research has shown that if validation takes
9642 20% of the expected block time, then a non-validating miner with 33% of the hash rate
9643 can mine between 53% and 68% of main chain blocks, depending on network connectiv-
9644 ity/delay [576]. Bitcoin’s more limited scripting and the use of standardized transactions
9645 help mitigate this issue, but these are not complete solutions. In fact, on July 4, 2015, the
9646 Bitcoin network had a temporary six-block fork built on top of an invalid block, which
9647 suggests that a large segment of mining pools were not validating blocks that they received
9648 at the time [61].

9649 In permissioned systems, the incentives will typically be exogenous to the system instead
9650 of denominated in a native token. The validator benefits because they can provide a better

9651 service and/or reduce the cost of providing the service. The costs for validators can be
9652 imposed legally instead of being in the form of work to prove or a capital investment in
9653 stake that can be slashed. However, this has not been tested legally in the real world, and
9654 complications such as jurisdictional questions and the possibility of accidentally produc-
9655 ing invalid blocks (e.g., via a software bug) may make it non-trivial to impose those legal
9656 costs. If legal penalties cannot be relied upon, it can be rational for otherwise "honest" per-
9657 mitted validators to accept invalid blocks under some circumstances, especially when
9658 blocks are very large or computationally intensive [577, 578].

9659 Addressing the Verifier's Dilemma and the potential disincentive to validate blocks in gen-
9660 eral is challenging. Fundamentally, this requires that the cost of verifying blocks and trans-
9661 actions be substantially limited through having modest maximum block sizes or gas limits.
9662 Additionally, some of the alternative state machine techniques discussed in Sections 18.1.2
9663 and 18.1.3 can limit the problem.

9664 **19.3. Alternative Transaction Fee Protocols**

9665 The current transaction fee mechanism used in Bitcoin and most existing cryptocurrency
9666 networks is a *multi-unit first-price auction* for block space. Clients attach a fee to their
9667 transactions, which is paid to the miner of the block that includes the transaction (in Bitcoin,
9668 the fee is determined implicitly as the difference between the input amounts and output
9669 amounts). Miners maximize their revenue by choosing the highest fee transactions (or
9670 more accurately, the highest fee *rate*) from the mempool up to the block size limit.

9671 Ideally, clients would place bids that honestly represent the value that the user would get
9672 from transaction inclusion. Unfortunately, first-price auctions lead to strategic fee selection
9673 when bidding for block space. Instead of bidding the true value, the appropriate strategy
9674 for users is to deliberately underbid and then use techniques like replace-by-fee to increase
9675 their bid as needed. In practice, this can lead to rapid increases in fees as congestion
9676 increases and having those fees be "sticky" even when congestion goes back down. This
9677 makes both fees and the latency of transactions unpredictable to clients, which creates
9678 a negative user experience while consuming more bandwidth due to re-broadcasting the
9679 transactions. Another issue with the current scheme is that if there is insufficient demand
9680 for block space, clients can offer arbitrarily small fees above the marginal risk of the block
9681 becoming stale due to increased propagation delay, which creates a race to the bottom
9682 where miners may not be paid sufficiently to secure the network.

9683 These issues led to several proposals for alternative fee mechanisms inspired by generalized
9684 *multi-unit second-price auctions* [579, 580]. In a second-price auction, the winning bidder
9685 only pays the bid of the second-highest bidder. In a generalized second-price auction (or
9686 K -th priced auction), K items are sold to the K -highest bidders, who each pay the $K + 1$ -th
9687 highest bid. This type of auction is known to lead to the dominant strategy where users'
9688 bids reflect their true value for the items. In the context of distributed ledgers, however,
9689 this type of auction is impossible to enforce because there is no consensus on the $K + 1$ -

9690 th highest bid, which will only exist in miners' mempools. These two new fee protocols
9691 charge users the K -th bid, which is actually observable. In both cases, the advantage to be
9692 gained from strategic bidding goes toward zero as the number of clients issuing transactions
9693 increases.

9694 While both new protocols operate similarly, they are optimized for different goals. The
9695 LSZ protocol from [579] is designed to decouple the fee market from the block size, such
9696 that an increase in the block size limit (or a decrease in demand for block space) will not
9697 lead to plummeting fees. That is, LSZ is intended to maximize fee revenue to improve
9698 security but comes at the expense of the social welfare of users. On the other hand, the
9699 BEOS protocol from [580] attempts to maximize social welfare and thus results in lower
9700 fees for users, as well as better throughput and latency for transactions than [579].

9701 The LSZ protocol has clients specify in each transaction the maximal fee that user is willing
9702 to pay for transaction inclusion. Miners have full latitude to decide which transactions to
9703 put in blocks, but all transactions will end up paying the the same fee, which is equivalent
9704 to the lowest fee included in that block. For example, if the set of fees available in a
9705 miner's mempool is $\{5, 2, 1, 1\}$, the miner will maximize their revenue by including only
9706 the transaction that pays five units of fees. Had the miner included all of those transactions
9707 in their block or just the first two, they would only end up claiming four units of fees.

9708 The BEOS protocol differs from LSZ in two primary ways:

- 9709 1. Under BEOS, miners are required to completely fill their blocks with transactions. If
9710 demand is insufficient, then the miner must stuff their block with "artificial" transac-
9711 tions, where the miner sends funds between addresses under their control.
- 9712 2. The transaction fees are shared between miners over a period of B blocks. When a
9713 miner successfully mines a block, they are paid the average fee collected over the
9714 most recent B blocks, including their own.

9715 Using the example above, the miner would be required to include all of the transactions
9716 available up to the block size limit and would thus generate four units of fees, which are
9717 then shared with the next B winning miners. This results in higher throughput and more
9718 stable fees for miners but lower overall mining rewards than either LSZ or the existing
9719 first-price auction mechanism. This may result in a lower amount of computational effort
9720 or stake used to secure the network. Reward sharing reduces the risk for miners to fork
9721 the ledger in order to try to double-spend but can improve censorship resistance by making
9722 it harder for miners to invalidate other miners' rewards. That said, it also weakens the
9723 censorship resistance benefit to users that comes with paying higher fees. Note that reward
9724 sharing is important in this scheme to reduce profits from miner manipulation of the fee
9725 mechanism. In the running example, assuming no reward sharing and a requirement of
9726 four transactions included in a block, a miner could gain by including three self-paying
9727 transactions with a fee of five units each. This would result in five units of fee revenue
9728 instead of four.

9729 While these mechanisms both have desirable features, they have several problems. First,
9730 these schemes fail to account for the benefits miners can gain through MEV, as discussed
9731 in the previous section. Second, these mechanisms may be challenging to implement, par-
9732 ticularly for UTXO ledgers. Because fees are typically implied by the difference between
9733 the inputs and the outputs, some mechanism is required in order to provide a refund for the
9734 surplus transaction fees included in their bid. This would require a substantial change to
9735 the system architecture and the way transactions are constructed. Third, and most impor-
9736 tantly, these generalized second-price auction schemes do not have a mechanism to address
9737 out-of-band payments to miners or miner collusion with transaction senders. The BEOS
9738 scheme, in particular, creates an explicit incentive for side-dealing by sharing the reward.
9739 Miners are inclined to accept side payments from clients who want to transact in exchange
9740 for accepting transactions with low fees because miners cannot be forced to share the out-
9741 of-band payment. These bribes can be implemented easily in Bitcoin by including an extra
9742 output in a transaction that has a *scriptPubKey* set to OP_TRUE and is, thus, immediately
9743 spendable by anyone. Whoever mines the transaction that includes this output can then
9744 spend the output to themselves within the same block, accepting the bribe.

9745 Not only can transactors bribe miners out-of-band, but miners can also collude with trans-
9746 actors to raise fees. If a transaction sender would have normally submitted a low fee of
9747 f_{low} , a miner can ask them to include a higher transaction fee of f_{high} instead and then pay
9748 the transactor $f_{high} - \frac{f_{low}}{2}$ [581]. This allows the miner to capture some of the revenue that
9749 would otherwise be lost by excluding transactions (in LSZ) or filling the block with dummy
9750 transactions (in BEOS, if rewards were not shared).

9751 Both first- and second-price auctions leave much to be desired in the context of fee pro-
9752 tocols. In addition to the issues already discussed above, an ideal fee mechanism would
9753 take into account the marginal social costs, or externalities, imposed on nodes for process-
9754 ing transactions. Transaction inclusion provides a private benefit to the sender but imposes
9755 computation, bandwidth, and storage costs on every node. To internalize these externalities,
9756 the fee mechanism ought to result in fees commensurate to the resources used. This can be
9757 done by either setting a quantity limit, like a maximum block size, or by fixing a minimum
9758 price. In traditional economic theory, if the marginal social cost of resource consumption
9759 is fixed but the marginal private benefit is decreasing, it is more efficient to set a price for
9760 the resource instead of setting quantity limits. However, when marginal social costs are
9761 increasing, quantity limits are superior to price setting [581]. An informal analysis of the
9762 marginal social costs of transactions suggests a decreasing marginal cost at low transaction
9763 throughput and rapidly increasing marginal costs with higher throughput (due to decreased
9764 security), which "suggests that a flat per-weight-unit in-protocol transaction fee, coupled
9765 with a hard limit at the point where the marginal social cost starts rapidly increasing, is
9766 superior to a pure weight limit-based regime" [581]. This reasoning ultimately led to the
9767 fee mechanism described in EIP-1559, which is currently deployed in several networks,
9768 including Ethereum and Filecoin [582].

9769 EIP-1559 breaks the total transaction fee down into two components: an algorithmically
9770 computed *base fee* that is burned and a user-supplied "tip" that is collected by the miner
9771 of the transaction. The base fee does not depend on the transactions included in a block
9772 but rather is determined by the preceding blocks. A target block size, s_{target} , is selected
9773 at the start, and the maximum block size is double this amount. Whenever the size of the
9774 most recent block is greater than s_{target} , the base fee is adjusted upward; if it is smaller
9775 than the target, the base fee decreases. Specifically, if r_{pred} is the base fee of the preceding
9776 block and s_{pred} is the size of the preceding block, then the current block's base fee $r_{cur} :=$
9777 $r_{pred} * (1 + \frac{1}{8} * \frac{s_{pred} - s_{target}}{s_{target}})$. Transactions specify a tip and a fee cap, and transactions will
9778 not be included in the chain unless the fee cap exceeds r_{cur} .

9779 A game-theoretic analysis of the EIP-1559 mechanism was performed in [583]. Rational
9780 miners are disincentivized from including fake transactions in their blocks, and the mecha-
9781 nism provides no avenues for miners and users to profitably collude using side payments.
9782 Furthermore, the mechanism itself does not lead to a decrease in security against selfish
9783 mining or double-spending, except insofar as the fee burn reduces the total hash rate. As
9784 a result, the inflationary block subsidy takes on increased importance. In addition, the
9785 optimal bidding strategy for users – except during sudden demand spikes – is to set their
9786 transaction's fee cap to their true value of having the transaction included. A separate anal-
9787 ysis of the dynamics of EIP-1559 found that with steady transaction demand, there can be
9788 chaotic periods with alternating full and empty blocks. This is far more likely to occur if
9789 user valuations for what to tip are similar to each other and, thus, have low variance [584].
9790 Another variant of EIP-1559 that eliminates the first-price auction optimal bidding issue
9791 during demand spikes has been proposed as well [585].

9792 To review, first-price auctions, as currently used in most permissionless ledger systems,
9793 make reasoning about fees challenging for users. The optimal bid for the user depends on
9794 bids offered by other users at the same time. Second-price auctions can be manipulated by
9795 miners by stuffing blocks with their own transactions instead of real user ones. However,
9796 these strategic issues do not arise in fixed-price sales, so establishing a base fee indepen-
9797 dent of transactions in the current block is a reasonable way of easing the burden of fee
9798 estimation on users. If the user values transaction inclusion more than the base fee, they
9799 bid the base fee plus enough to compensate the miner for the marginal costs of transaction
9800 inclusion. The base fee must be burned or otherwise withheld from the block's miner to
9801 prevent side dealing, and it must adjust dynamically as demand for block space changes.
9802 There must be agreement over the base fee, so a proxy for demand is used (e.g., in the
9803 case of EIP-1559, a variable block size). If the entirety of the fee was burned, then miners
9804 would have no incentive to include transactions in their blocks at all. The extra tip resolves
9805 this issue while also providing a way for users to signal how much they value transaction
9806 inclusion.

9807 **References**

- 9808 [1] Nakamoto S (2008) Bitcoin: A peer-to-peer electronic cash system.
9809 [2] Pease M, Shostak R, Lamport L (1980) Reaching agreement in the presence of faults.
9810 *Journal of the ACM (JACM)* 27(2):228–234.
9811 [3] LAMPORT L, SHOSTAK R, PEASE M (1982) The byzantine generals problem.
9812 *ACM Transactions on Programming Languages and Systems* 4(3):382–401.
9813 [4] Castro M, Liskov B, et al. (1999) Practical byzantine fault tolerance. *OSDI*, Vol. 99,
9814 pp 173–186.
9815 [5] Yaga D, Mell P, Roby N, Scarfone K (2018) Blockchain technology overview (Na-
9816 tional Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency
9817 or Internal Report (IR) 8202. <https://doi.org/10.6028/NIST.IR.8202>.
9818 [6] Bracha G (1987) Asynchronous byzantine agreement protocols. *Information and*
9819 *Computation* 75(2):130–143.
9820 [7] Lamport L (1984) Using time instead of timeout for fault-tolerant distributed
9821 systems. *ACM Transactions on Programming Languages and Systems (TOPLAS)*
9822 6(2):254–280.
9823 [8] Schneider FB (1990) Implementing fault-tolerant services using the state machine
9824 approach: A tutorial. *ACM Computing Surveys (CSUR)* 22(4):299–319.
9825 [9] Garay J, Kiayias A, Leonardos N (2015) The bitcoin backbone protocol: Analysis
9826 and applications. *Annual International Conference on the Theory and Applications*
9827 *of Cryptographic Techniques* (Springer), pp 281–310.
9828 [10] Blum E, Katz J, Loss J (2020) Network-agnostic state machine replication. *arXiv*
9829 *preprint arXiv:200203437* .
9830 [11] Garay J, Kiayias A (2020) Sok: A consensus taxonomy in the blockchain era. *Cryp-*
9831 *tographers' Track at the RSA Conference* (Springer), pp 284–318.
9832 [12] Guerraoui R, Kuznetsov P, Monti M, Pavlovič M, Seredinschi DA (2019) The con-
9833 sensus number of a cryptocurrency. *Proceedings of the 2019 ACM Symposium on*
9834 *Principles of Distributed Computing*, pp 307–316.
9835 [13] Sliwinski J, Wattenhofer R (2019) Abc: Asynchronous blockchain without consen-
9836 sus. *arXiv preprint arXiv:190910926* .
9837 [14] Baudet M, Danezis G, Sonnino A (2020) Fastpay: High-performance byzantine fault
9838 tolerant settlement. *arXiv preprint arXiv:200311506* .
9839 [15] Collins D, Guerraoui R, Komatovic J, Monti M, Xygkis A, Pavlovic M, Kuznetsov P,
9840 Pignolet YA, Seredinschi DA, Tonkikh A (2020) Online payments by merely broad-
9841 casting messages (extended version). *arXiv preprint arXiv:200413184* .
9842 [16] Auvolat A, Frey D, Raynal M, Taïani F (2020) Money transfer made simple. *hal-*
9843 *02861511v2f* .
9844 [17] Cohen S, Keidar I (2021) Tame the wild with byzantine linearizability: Reliable
9845 broadcast, snapshots, and asset transfer. *arXiv preprint arXiv:210210597* .
9846 [18] Georgiades Y, Streit R, Garg V (2021) Who needs consensus? a distributed mon-
9847 etary system between rational agents via hearsay. *arXiv preprint arXiv:210407574*

- 9848 .
- 9849 [19] Kuznetsov P, Pignolet YA, Ponomarev P, Tonkikh A (2021) Permissionless and asyn-
9850 chronous asset transfer [technical report]. *arXiv preprint arXiv:210504966* .
- 9851 [20] Aiyer AS, Alvisi L, Clement A, Dahlin M, Martin JP, Porth C (2005) Bar fault
9852 tolerance for cooperative services. *Proceedings of the twentieth ACM symposium on*
9853 *Operating systems principles*, pp 45–58.
- 9854 [21] McMEnamin C, Daza V, Pontecorvi M (2020) Achieving state machine replication
9855 without honesty assumptions. *arXiv preprint arXiv:201210146* .
- 9856 [22] Fischer MJ, Lynch NA, Paterson MS (1985) Impossibility of distributed consensus
9857 with one faulty process. *Journal of the ACM (JACM)* 32(2):374–382.
- 9858 [23] Ben-Or M (1983) Another advantage of free choice: Completely asynchronous
9859 agreement protocols (extended abstract),". *Proceedings of the 2nd ACM Annual*
9860 *Symposium on Principles of Distributed Computing, Montreal, Quebec*, pp 27–30.
- 9861 [24] Rabin MO (1983) Randomized byzantine generals. *24th Annual Symposium on*
9862 *Foundations of Computer Science (sfcs 1983)* (IEEE), pp 403–409.
- 9863 [25] Bracha G, Toueg S (1985) Asynchronous consensus and broadcast protocols. *Jour-*
9864 *nal of the ACM (JACM)* 32(4):824–840.
- 9865 [26] Dwork C, Lynch N, Stockmeyer L (1988) Consensus in the presence of partial syn-
9866 chrony. *Journal of the ACM (JACM)* 35(2):288–323.
- 9867 [27] Spiegelman A (2020) In search for a linear byzantine agreement. *arXiv preprint*
9868 *arXiv:200206993* .
- 9869 [28] Pass R, Shi E (2017) The sleepy model of consensus. *International Conference on*
9870 *the Theory and Application of Cryptology and Information Security* (Springer), pp
9871 380–409.
- 9872 [29] Guo Y, Pass R, Shi E (2019) Synchronous, with a chance of partition tolerance.
9873 *Annual International Cryptology Conference* (Springer), pp 499–529.
- 9874 [30] Kim J, Mehta V, Nayak K, Shrestha N (2021) Making synchronous bft protocols
9875 secure in the presence of mobile sluggish faults. *IACR Cryptology ePrint Archive*
9876 :603.
- 9877 [31] Pass R, Shi E (2017) Rethinking large-scale consensus. *2017 IEEE 30th Computer*
9878 *Security Foundations Symposium (CSF)* (IEEE), pp 115–129.
- 9879 [32] Wang Q, Yu J, Chen S, Xiang Y (2020) Sok: Diving into dag-based blockchain
9880 systems. *arXiv preprint arXiv:201206128* .
- 9881 [33] Douceur JR (2002) The sybil attack. *International workshop on peer-to-peer systems*
9882 (Springer), pp 251–260.
- 9883 [34] Bowman M, Das D, Mandal A, Montgomery H (2021) On elapsed time consensus
9884 protocols. *IACR Cryptol ePrint Arch* :086.
- 9885 [35] Yandamuri S, Abraham I, Nayak K, Reiter MK (2021) Communication-efficient bft
9886 protocols using small trusted hardware to tolerate minority corruption. *IACR Cryptol*
9887 *ePrint Arch* :184.
- 9888 [36] Deuber D, Döttling N, Magri B, Malavolta G, Thyagarajan SAK (2018) Minting
9889 mechanisms for blockchain-or-moving from cryptoassets to cryptocurrencies. *IACR*

- 9890 *Cryptol ePrint Arch* .
- 9891 [37] Neiheiser R, Matos M, Rodrigues L (2021) The quest for scaling bft consensus
9892 through tree-based vote aggregation. *arXiv preprint arXiv:210312112* .
- 9893 [38] Wang L, Shen X, Li J, Shao J, Yang Y (2019) Cryptographic primitives in
9894 blockchains. *Journal of Network and Computer Applications* 127:43–58.
- 9895 [39] Partala J, Nguyen TH, Pirttikangas S (2020) Non-interactive zero-knowledge for
9896 blockchain: A survey. *IEEE Access* 8:227945–227961.
- 9897 [40] Micali S, Rabin M, Vadhan S (1999) Verifiable random functions. *40th annual sym-*
9898 *posium on foundations of computer science (cat. No. 99CB37039)* (IEEE), pp 120–
9899 130.
- 9900 [41] Boneh D, Lynn B, Shacham H (2001) Short signatures from the weil pairing. *In-*
9901 *ternational conference on the theory and application of cryptography and information*
9902 *security* (Springer), pp 514–532.
- 9903 [42] Dryja T, Liu QC, Narula N (2020) A lower bound for byzantine agreement and
9904 consensus for adaptive adversaries using vdfs. *arXiv preprint arXiv:200401939* .
- 9905 [43] Boneh D, Bonneau J, Bünz B, Fisch B (2018) Verifiable delay functions. *Annual*
9906 *international cryptography conference* (Springer), pp 757–788.
- 9907 [44] Pietrzak K (2018) Simple verifiable delay functions. *10th innovations in theoretical*
9908 *computer science conference (itcs 2019)* (Schloss Dagstuhl-Leibniz-Zentrum fuer
9909 Informatik).
- 9910 [45] Wesolowski B (2019) Efficient verifiable delay functions. *Annual International Con-*
9911 *ference on the Theory and Applications of Cryptographic Techniques* (Springer), pp
9912 379–407.
- 9913 [46] (2015) Not bitcoin xt. Accessed on April 13, 2021. Available at [https://github.com](https://github.com/xtbit/notbitcoinxt)
9914 [/xtbit/notbitcoinxt](https://github.com/xtbit/notbitcoinxt).
- 9915 [47] Daian P (2017) On soft fork security. Accessed on April 4, 2021. Available at [https:](https://pdaian.com/blog/on-soft-fork-security/)
9916 [//pdaian.com/blog/on-soft-fork-security/](https://pdaian.com/blog/on-soft-fork-security/).
- 9917 [48] Folkson M (2021) Should block height or mtp or a mixture of both be used in a
9918 soft fork activation mechanism?. Accessed on April 13, 2021. Available at [https:](https://bitcoin.stackexchange.com/questions/103854/should-block-height-or-mtp-or-a-mixture-of-both-be-used-in-a-soft-fork-activatio)
9919 [//bitcoin.stackexchange.com/questions/103854/should-block-height-or-mtp-or-a-m](https://bitcoin.stackexchange.com/questions/103854/should-block-height-or-mtp-or-a-mixture-of-both-be-used-in-a-soft-fork-activatio)
9920 [ixture-of-both-be-used-in-a-soft-fork-activatio](https://bitcoin.stackexchange.com/questions/103854/should-block-height-or-mtp-or-a-mixture-of-both-be-used-in-a-soft-fork-activatio).
- 9921 [49] Buterin V (2017) Hard forks, soft forks, defaults and coercion. Accessed on April 13,
9922 2021. Available at https://vitalik.ca/general/2017/03/14/forks_and_markets.html.
- 9923 [50] Wuille P (2015) Bitcoin core and hard forks. Accessed on April 13, 2021. Available
9924 at <https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2015-July/009515.html>.
- 9925 [51] Srinivasan BS, Lee L (2017) Quantifying decentralization. Accessed on April 20,
9926 2021. Available at [https://news.earn.com/quantifying-decentralization-e39db233c](https://news.earn.com/quantifying-decentralization-e39db233c28e)
9927 [28e](https://news.earn.com/quantifying-decentralization-e39db233c28e).
- 9928 [52] Sztork P (2015) Measuring decentralization. Accessed on April 20, 2021. Available
9929 at <https://www.truthcoin.info/blog/measuring-decentralization/>.
- 9930 [53] Kwon Y, Liu J, Kim M, Song D, Kim Y (2019) Impossibility of full decentralization
9931 in permissionless blockchains. *Proceedings of the 1st ACM Conference on Advances*

- 9932 *in Financial Technologies*, pp 110–123.
- 9933 [54] Gervais A, Capkun S, Karame GO, Gruber D (2014) On the privacy provisions of
9934 bloom filters in lightweight bitcoin clients. *Proceedings of the 30th Annual Com-*
9935 *puter Security Applications Conference*, pp 326–335.
- 9936 [55] Osuntokun O, Akselrod A, Posen J (2017) Client side block filtering. Accessed on
9937 April 22, 2021. Available at [https://github.com/bitcoin/bips/blob/master/bip-0157.](https://github.com/bitcoin/bips/blob/master/bip-0157.mediawiki)
9938 [mediawiki](https://github.com/bitcoin/bips/blob/master/bip-0157.mediawiki).
- 9939 [56] Osuntokun O, Akselrod A (2017) Compact block filters for light clients. Accessed
9940 on April 22, 2021. Available at [https://github.com/bitcoin/bips/blob/master/bip-015](https://github.com/bitcoin/bips/blob/master/bip-0158.mediawiki)
9941 [8.mediawiki](https://github.com/bitcoin/bips/blob/master/bip-0158.mediawiki).
- 9942 [57] Hearn M (2012) [bitcoin-development] roadmap to getting users onto spv clients.
9943 Accessed on April 27, 2021. Available at [https://lists.linuxfoundation.org/pipermai](https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2012-December/002083.html)
9944 [l/bitcoin-dev/2012-December/002083.html](https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2012-December/002083.html).
- 9945 [58] Research B (2019) Bitcoin’s initial block download. Accessed on April 22, 2021.
9946 Available at <https://blog.bitmex.com/bitcoins-initial-block-download/>.
- 9947 [59] Todd P (2015) [bitcoin-development] fwd: Block size increase requirements. Ac-
9948 cessed on April 22, 2021. Available at [https://lists.linuxfoundation.org/pipermail/bi](https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2015-June/008407.html)
9949 [tcoin-dev/2015-June/008407.html](https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2015-June/008407.html).
- 9950 [60] Zhang R, Preneel B (2017) On the necessity of a prescribed block validity consensus:
9951 Analyzing bitcoin unlimited mining protocol. *Proceedings of the 13th International*
9952 *Conference on emerging Networking EXperiments and Technologies*, pp 108–119.
- 9953 [61] Bitcoinorg (2015) Some miners generating invalid blocks. Available at [https://bitcoi](https://bitcoin.org/en/alert/2015-07-04-spv-mining)
9954 [n.org/en/alert/2015-07-04-spv-mining](https://bitcoin.org/en/alert/2015-07-04-spv-mining).
- 9955 [62] Dryja T (2019) Utreexo: A dynamic hash-based accumulator optimized for the bit-
9956 coin utxo set. *IACR Cryptol ePrint Arch* :611.
- 9957 [63] Bonneau J, Meckler I, Rao V, Shapiro E (2020) Mina: Decentralized cryptocurrency
9958 at scale. Available at [https://minaprotocol.com/wp-content/uploads/technicalWhit](https://minaprotocol.com/wp-content/uploads/technicalWhitepaper.pdf)
9959 [epaper.pdf](https://minaprotocol.com/wp-content/uploads/technicalWhitepaper.pdf).
- 9960 [64] karalabe (2015) eth/63 fast synchronization algorithm. Accessed on April 28, 2021.
9961 Available at <https://github.com/ethereum/go-ethereum/pull/188>.
- 9962 [65] jamesob (2019) assumeutxo. Accessed on April 28, 2021. Available at [https://gith](https://github.com/bitcoin/bitcoin/issues/15605)
9963 [ub.com/bitcoin/bitcoin/issues/15605](https://github.com/bitcoin/bitcoin/issues/15605).
- 9964 [66] Loruenser T, Rainer B, Wohner F (2021) Towards a performance model for byzan-
9965 tine fault tolerant (storage) services. *arXiv preprint arXiv:210104489* .
- 9966 [67] Bessani A, Santos M, Felix J, Neves N, Correia M (2013) On the efficiency of
9967 durable state machine replication. *2013 {USENIX} Annual Technical Conference*
9968 *({USENIX}{ATC} 13)*, pp 169–180.
- 9969 [68] Malkhi D, Reiter M (1998) Byzantine quorum systems. *Distributed computing*
9970 *11(4):203–213*.
- 9971 [69] Wang Q, Yu J, Peng Z, Bui VC, Chen S, Ding Y, Xiang Y (2020) Security analysis
9972 on dbft protocol of neo. *International Conference on Financial Cryptography and*
9973 *Data Security* :20–31.

- 9974 [70] Wang Q, Li R, Chen S, Xiang Y (2021) Formal security analysis on dbft protocol of
9975 neo. *arXiv preprint arXiv:210507459* .
- 9976 [71] Bessani A, Sousa J, Alchieri EE (2014) State machine replication for the masses with
9977 bft-smart. *2014 44th Annual IEEE/IFIP International Conference on Dependable
9978 Systems and Networks (IEEE)*, pp 355–362.
- 9979 [72] Saltini R, Hyland-Wood D (2019) Correctness analysis of ibft. *arXiv preprint
9980 arXiv:190107160* .
- 9981 [73] Saltini R (2019) Ibft liveness analysis. *2019 IEEE International Conference on
9982 Blockchain (Blockchain) (IEEE)*, pp 245–252.
- 9983 [74] Saltini R, Hyland-Wood D (2019) Ibft 2.0: A safe and live variation of the ibft
9984 blockchain consensus protocol for eventually synchronous networks. *arXiv preprint
9985 arXiv:190910194* .
- 9986 [75] Moniz H (2020) The istanbul bft consensus algorithm. *arXiv preprint
9987 arXiv:200203613* .
- 9988 [76] Kotla R, Alvisi L, Dahlin M, Clement A, Wong E (2007) Zyzzyva: speculative
9989 byzantine fault tolerance. *Proceedings of twenty-first ACM SIGOPS symposium on
9990 Operating systems principles*, pp 45–58.
- 9991 [77] Wang G (2021) Sok: Understanding bft consensus in the age of blockchains. *Cryp-
9992 tology ePrint Archive* :911.
- 9993 [78] Abraham I, Gueta G, Malkhi D, Alvisi L, Kotla R, Martin JP (2017) Revisiting fast
9994 practical byzantine fault tolerance. *arXiv preprint arXiv:171201367* .
- 9995 [79] Guerraoui R, Knežević N, Quéma V, Vukolić M (2010) The next 700 bft protocols.
9996 *Proceedings of the 5th European conference on Computer systems*, pp 363–376.
- 9997 [80] Gunn LJ, Liu J, Vavala B, Asokan N (2019) Making speculative bft resilient with
9998 trusted monotonic counters. *2019 38th Symposium on Reliable Distributed Systems
9999 (SRDS) (IEEE)*, pp 133–13309.
- 10000 [81] Danezis G, Hrycyszyn D (2018) Blockmania: from block dags to consensus. *arXiv
10001 preprint arXiv:180901620* .
- 10002 [82] Schett MA, Danezis G (2021) Embedding a deterministic bft protocol in a block
10003 dag. *arXiv preprint arXiv:210209594* .
- 10004 [83] Shi E (2019) Streamlined blockchains: A simple and elegant approach (a tutorial
10005 and survey). *International Conference on the Theory and Application of Cryptology
10006 and Information Security (Springer)*, pp 3–17.
- 10007 [84] Chan BY, Shi E (2020) Streamlet: Textbook streamlined blockchains. *IACR Cryptol
10008 ePrint Arch* 2020:88.
- 10009 [85] Pass R, Shi E (2018) Thunderella: Blockchains with optimistic instant confirmation.
10010 *Annual International Conference on the Theory and Applications of Cryptographic
10011 Techniques (Springer)*, pp 3–33.
- 10012 [86] Chan TH, Pass R, Shi E (2018) Pili: A simple, fast, and robust family of blockchain
10013 protocols (Cryptology ePrint Archive, Report 2018/980, 2018. <https://eprint.iacr.org> . . .),
- 10014
10015 [87] Chan THH, Pass R, Shi E (2018) Pala: A simple partially synchronous blockchain.

- 10016 *IACR Cryptol ePrint Arch* 2018:981.
- 10017 [88] Yin M, Malkhi D, Reiter MK, Gueta GG, Abraham I (2018) Hotstuff: Bft consensus
10018 in the lens of blockchain. *arXiv preprint arXiv:180305069* .
- 10019 [89] Abraham I, Malkhi D, Nayak K, Ren L, Yin M (2019) Sync hotstuff: Simple and
10020 practical synchronous state machine replication. *IACR Cryptology ePrint Archive*
10021 :270.
- 10022 [90] Abraham I, Malkhi D, Spiegelman A (2019) Asymptotically optimal validated asyn-
10023 chronous byzantine agreement. *Proceedings of the 2019 ACM Symposium on Prin-*
10024 *ciples of Distributed Computing*, pp 337–346.
- 10025 [91] Baudet M, Ching A, Chursin A, Danezis G, Garillot F, Li Z, Malkhi D, Naor O,
10026 Perelman D, Sonnino A (2019) State machine replication in the libra blockchain.
10027 *The Libra Assn, Tech Rep* .
- 10028 [92] Momose A, Cruz JP (2019) Force-locking attack on sync hotstuff. *IACR Cryptol*
10029 *ePrint Arch* :1484.
- 10030 [93] Abraham I, Nayak K, Ren L, Shrestha N (2020) On the optimality of optimistic
10031 responsiveness. *IACR Cryptol ePrint Arch* :458.
- 10032 [94] Bhat A, Bandarupalli A, Bagchi S, Kate A, Reiter MK (2021) Apollo—optimistically
10033 linear and responsive smr. *IACR Cryptol ePrint Arch* :180.
- 10034 [95] Miller A, Xia Y, Croman K, Shi E, Song D (2016) The honey badger of bft protocols.
10035 *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communica-*
10036 *tions Security*, pp 31–42.
- 10037 [96] Liu C, Duan S, Zhang H (2020) Epic: Efficient asynchronous bft with adaptive secu-
10038 rity. *2020 50th Annual IEEE/IFIP International Conference on Dependable Systems*
10039 *and Networks (DSN) (IEEE)*, pp 437–451.
- 10040 [97] Cachin C, Kursawe K, Petzold F, Shoup V (2001) Secure and efficient asynchronous
10041 broadcast protocols. *Annual International Cryptology Conference (Springer)*, pp
10042 524–541.
- 10043 [98] Guo B, Lu Z, Tang Q, Xu J, Zhang Z (2020) Dumbo: Faster asynchronous bft pro-
10044 tocols. *IACR Cryptology ePrint Archive* :841.
- 10045 [99] Cachin C, Tessaro S (2005) Asynchronous verifiable information dispersal. *24th*
10046 *IEEE Symposium on Reliable Distributed Systems (SRDS'05) (IEEE)*, pp 191–201.
- 10047 [100] Mostéfaoui A, Moumen H, Raynal M (2015) Signature-free asynchronous binary
10048 byzantine consensus with $t < n/3$, $o(n^2)$ messages, and $o(1)$ expected time. *Journal*
10049 *of the ACM (JACM)* 62(4):1–21.
- 10050 [101] Duan S, Reiter MK, Zhang H (2018) Beat: Asynchronous bft made practical. *Pro-*
10051 *ceedings of the 2018 ACM SIGSAC Conference on Computer and Communications*
10052 *Security*, pp 2028–2041.
- 10053 [102] Lu Y, Lu Z, Tang Q (2021) Bolt-dumbo transformer: Asynchronous consensus as
10054 fast as pipelined bft. *arXiv preprint arXiv:210309425* .
- 10055 [103] Baird L (2016) The swirls hashgraph consensus algorithm: Fair, fast, byzantine
10056 fault tolerance. *Swirls, Inc Technical Report SWIRLDS-TR-2016 1*.
- 10057 [104] Kelkar M, Zhang F, Goldfeder S, Juels A (2020) Order-fairness for byzantine con-

- 10058 sensus (IACR Cryptology ePrint Archive, 2020: 269),
10059 [105] Gagol A, Leśniak D, Straszak D, Świetek M (2019) Aleph: Efficient atomic broad-
10060 cast in asynchronous networks with byzantine nodes. *Proceedings of the 1st ACM*
10061 *Conference on Advances in Financial Technologies*, pp 214–228.
10062 [106] Keidar I, Kokoris-Kogias E, Naor O, Spiegelman A (2021) All you need is dag.
10063 *arXiv preprint arXiv:210208325* .
10064 [107] Kursawe K (2020) Wendy, the good little fairness widget. *arXiv preprint*
10065 *arXiv:200708303* .
10066 [108] Asayag A, Cohen G, Grayevsky I, Leshkowitz M, Rottenstreich O, Tamari R, Yakira
10067 D (2018) A fair consensus protocol for transaction ordering. *2018 IEEE 26th Inter-*
10068 *national Conference on Network Protocols (ICNP)* (IEEE), pp 55–65.
10069 [109] Zhang Y, Setty S, Chen Q, Zhou L, Alvisi L (2020) Byzantine ordered consensus
10070 without byzantine oligarchy. *14th {USENIX} Symposium on Operating Systems De-*
10071 *sign and Implementation ({OSDI} 20)*, pp 633–649.
10072 [110] Avarikioti Z, Heimbach L, Schmid R, Wattenhofer R (2020) Fnf-bft: Exploring per-
10073 formance limits of bft protocols. *arXiv preprint arXiv:200902235* .
10074 [111] Lev-Ari K, Spiegelman A, Keidar I, Malkhi D (2019) Fairledger: A fair blockchain
10075 protocol for financial institutions. *arXiv preprint arXiv:190603819* .
10076 [112] Kuo PC, Chung H, Chao TW, Cheng CM (2020) Fair byzantine agreements for
10077 blockchains. *IEEE Access* 8:70746–70761.
10078 [113] Küsters R, Rausch D, Simon M (2020) Accountability in a permissioned blockchain:
10079 Formal analysis of hyperledger fabric (full version). *IACR Cryptology ePrint Archive*
10080 .
10081 [114] Ranchal-Pedrosa A, Gramoli V (2021) Agreement in the presence of disagreeing
10082 rational players: The huntsman protocol. *arXiv preprint arXiv:210504357* .
10083 [115] Civit P, Gilbert S, Gramoli V (2019) Polygraph: Accountable byzantine agreement.
10084 *IACR Cryptol ePrint Arch* :587.
10085 [116] Crain T, Gramoli V, Larrea M, Raynal M (2018) Dbft: Efficient leaderless byzantine
10086 consensus and its application to blockchains. *2018 IEEE 17th International Sympo-*
10087 *sium on Network Computing and Applications (NCA)* (IEEE), pp 1–8.
10088 [117] Crain T, Natoli C, Gramoli V (2018) Evaluating the red belly blockchain. *arXiv*
10089 *preprint arXiv:181211747* .
10090 [118] Ranchal-Pedrosa A, Gramoli V (2020) Blockchain is dead, long live blockchain!
10091 accountable state machine replication for longlasting blockchain. *arXiv preprint*
10092 *arXiv:200710541* .
10093 [119] Sheng P, Wang G, Nayak K, Kannan S, Viswanath P (2020) Bft protocol forensics.
10094 *arXiv preprint arXiv:201006785* .
10095 [120] Golan-Gueta G, Abraham I, Grossman S, Malkhi D, Pinkas B, Reiter MK, Seredin-
10096 schi DA, Tamir O, Tomescu A (2019) Sbft: A scalable and decentralized trust in-
10097 frastructure. *2019 49th Annual IEEE/IFIP International Conference on Dependable*
10098 *Systems and Networks (DSN)* :568–580.
10099 [121] Jalalzai MM, Busch C, Richard GG (2019) Proteus: A scalable bft consen-

- 10100 sus protocol for blockchains. *2019 IEEE International Conference on Blockchain*
10101 *(Blockchain)* (IEEE), pp 308–313.
- 10102 [122] De Angelis S, Aniello L, Baldoni R, Lombardi F, Margheri A, Sassone V (2018) Pbft
10103 vs proof-of-authority: applying the cap theorem to permissioned blockchain. *Pro-*
10104 *ceedings of the Second Italian Conference on Cyber Security, Milan, Italy, February*
10105 *6th - to - 9th, 2018* 2058.
- 10106 [123] Ekparinya P, Gramoli V, Jourjon G (2019) The attack of the clones against proof-of-
10107 authority. *arXiv preprint arXiv:190210244* .
- 10108 [124] Natoli C, Gramoli V (2017) The balance attack or why forkable blockchains are
10109 ill-suited for consortium. *2017 47th Annual IEEE/IFIP International Conference on*
10110 *Dependable Systems and Networks (DSN)* (IEEE), pp 579–590.
- 10111 [125] Ekparinya P, Gramoli V, Jourjon G (2018) Double-spending risk quantifica-
10112 tion in private, consortium and public ethereum blockchains. *arXiv preprint*
10113 *arXiv:180505004* .
- 10114 [126] Shi E (2019) Analysis of deterministic longest-chain protocols. *2019 IEEE 32nd*
10115 *Computer Security Foundations Symposium (CSF)* (IEEE), pp 122–12213.
- 10116 [127] Kiayias A, Russell A (2018) Ouroboros-bft: A simple byzantine fault tolerant con-
10117 sensus protocol. *IACR Cryptol ePrint Arch* :1049.
- 10118 [128] Malkhi D, Nayak K, Ren L (2019) Flexible byzantine fault tolerance. *Proceedings*
10119 *of the 2019 ACM SIGSAC Conference on Computer and Communications Security*,
10120 pp 1041–1053.
- 10121 [129] Sheff I, Wang X, van Renesse R, Myers AC (2021) Heterogeneous paxos. *24th Inter-*
10122 *national Conference on Principles of Distributed Systems (OPODIS 2020)* (Schloss
10123 Dagstuhl-Leibniz-Zentrum für Informatik).
- 10124 [130] Kane D, Fackler A, Gagol A, Straszak D, Zamfir V (2021) Highway: Efficient con-
10125 sensus with flexible finality. *arXiv preprint arXiv:210102159* .
- 10126 [131] Xiang Z, Malkhi D, Nayak K, Ren L (2021) Strengthened fault tolerance in byzan-
10127 tine fault tolerant replication. *arXiv preprint arXiv:210103715* .
- 10128 [132] Naor O, Baudet M, Malkhi D, Spiegelman A (2019) Cogsworth: Byzantine view
10129 synchronization. *arXiv preprint arXiv:190905204* .
- 10130 [133] Naor O, Keidar I (2020) Expected linear round synchronization: The missing link
10131 for linear byzantine smr. *arXiv preprint arXiv:200207539* .
- 10132 [134] Bravo M, Chockler G, Gotsman A (2020) Making byzantine consensus live (ex-
10133 tended version). *arXiv preprint arXiv:200804167* .
- 10134 [135] Spiegelman A, Rinberg A (2019) Ace: Abstract consensus encapsulation for liveness
10135 boosting of state machine replication. *arXiv preprint arXiv:191110486* .
- 10136 [136] Gelashvili R, Kokoris-Kogias L, Spiegelman A, Xiang Z (2021) Be prepared
10137 when network goes bad: An asynchronous view-change protocol. *arXiv preprint*
10138 *arXiv:210303181* .
- 10139 [137] Bessani A, Alchieri E, Sousa J, Oliveira A, Pedone F (2020) From byzan-
10140 tine replication to blockchain: Consensus is only the beginning. *arXiv preprint*
10141 *arXiv:200414527* .

- 10142 [138] Aștefanoaei L, Chambart P, Del Pozzo A, Tate E, Tucci S, Zălinescu E (2020)
10143 Tenderbake—classical bft style consensus for public blockchains. *arXiv preprint*
10144 *arXiv:200111965* .
- 10145 [139] Hao X, Yu L, Zhiqiang L, Zhen L, Dawu G (2018) Dynamic practical byzantine fault
10146 tolerance. *2018 IEEE Conference on Communications and Network Security (CNS)*
10147 (IEEE), pp 1–8.
- 10148 [140] Kuznetsov P, Tonkikh A (2020) Asynchronous reconfiguration with byzantine fail-
10149 ures. *arXiv preprint arXiv:200513499* .
- 10150 [141] Vizier G, Gramoli V (2018) Comchain: Bridging the gap between public and con-
10151 sortium blockchains. *2018 IEEE International Conference on Internet of Things*
10152 *(iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE*
10153 *Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (Smart-*
10154 *Data)* (IEEE), pp 1469–1474.
- 10155 [142] Freitas de Souza L, Kuznetsov P, Rieutord T, Tucci-Piergiorganni S (2021) Ac-
10156 countability and reconfiguration: Self-healing lattice agreement. *arXiv preprint*
10157 *arXiv:210504909* .
- 10158 [143] Martin JP, Alvisi L (2004) A framework for dynamic byzantine storage. *Internat-*
10159 *ional Conference on Dependable Systems and Networks, 2004* (IEEE), pp 325–334.
- 10160 [144] Alchieri EA, Bessani AN, da Silva Fraga J, Greve F (2008) Byzantine consensus
10161 with unknown participants. *International Conference On Principles Of Distributed*
10162 *Systems* (Springer), pp 22–40.
- 10163 [145] Alchieri EAP, Bessani A, Greve F, da Silva Fraga J (2016) Knowledge connectiv-
10164 ity requirements for solving byzantine consensus with unknown participants. *IEEE*
10165 *Transactions on Dependable and Secure Computing* 15(2):246–259.
- 10166 [146] Khanchandani P, Wattenhofer R (2021) Byzantine agreement with unknown partici-
10167 pants and failures. *arXiv preprint arXiv:210210442* .
- 10168 [147] Cachin C, Tackmann B (2019) Asymmetric distributed trust. *arXiv preprint*
10169 *arXiv:190609314* .
- 10170 [148] Cachin C, Zanolini L (2020) Asymmetric byzantine consensus. *arXiv preprint*
10171 *arXiv:200508795* .
- 10172 [149] Mazieres D (2015) The stellar consensus protocol: A federated model for internet-
10173 level consensus. *Stellar Development Foundation* .
- 10174 [150] Lokhava M, Losa G, Mazières D, Hoare G, Barry N, Gafni E, Jove J, Malinowsky
10175 R, McCaleb J (2019) Fast and secure global payments with stellar. *Proceedings of*
10176 *the 27th ACM Symposium on Operating Systems Principles*, pp 80–96.
- 10177 [151] García-Pérez Á, Schett MA (2019) Deconstructing stellar consensus (extended ver-
10178 sion). *arXiv preprint arXiv:191105145* .
- 10179 [152] Kim M, Kwon Y, Kim Y (2019) Is stellar as secure as you think? *2019 IEEE Eu-*
10180 *ropean Symposium on Security and Privacy Workshops (EuroS&PW)* (IEEE), pp
10181 377–385.
- 10182 [153] Lachowski Ł (2019) Complexity of the quorum intersection property of the federated
10183 byzantine agreement system. *arXiv preprint arXiv:190206493* .

- 10184 [154] Gaul A, Khoffi I, Liesen J, Stüber T (2019) Mathematical analysis and algorithms
10185 for federated byzantine agreement systems. *arXiv preprint arXiv:191201365* .
- 10186 [155] Florian M, Henningsen S, Scheuermann B (2020) The sum of its parts: Analysis of
10187 federated byzantine agreement systems. *arXiv preprint arXiv:200208101* .
- 10188 [156] Bracciali A, Grossi D, de Haan R (2019) Decentralization in open quorum systems.
10189 *arXiv preprint arXiv:191108182* .
- 10190 [157] Schwartz D, Youngs N, Britto A, et al. (2014) The ripple protocol consensus algo-
10191 rithm. *Ripple Labs Inc White Paper 58* .
- 10192 [158] Chase B, MacBrough E (2018) Analysis of the xrp ledger consensus protocol. *arXiv*
10193 *preprint arXiv:180207242* .
- 10194 [159] Amores-Sesar I, Cachin C, Mičić J (2020) Security analysis of ripple consensus.
10195 *arXiv preprint arXiv:201114816* .
- 10196 [160] MacBrough E (2018) Cobalt: Bft governance in open networks. *arXiv preprint*
10197 *arXiv:180207240* .
- 10198 [161] Dwork C, Naor M (1992) Pricing via processing or combatting junk mail. *Annual*
10199 *International Cryptology Conference* (Springer), pp 139–147.
- 10200 [162] Jakobsson M, Juels A (1999) Proofs of work and bread pudding protocols. *Secure*
10201 *information networks* (Springer), pp 258–272.
- 10202 [163] Back A, et al. (2002) Hashcash-a denial of service counter-measure .
- 10203 [164] Gupta D, Saia J, Young M (2020) Resource burning for permissionless systems.
10204 *arXiv preprint arXiv:200604865* .
- 10205 [165] Garay J, Kiayias A, Panagiotakos G (2020) Blockchains from non-idealized hash
10206 functions (Cryptology ePrint Archive 2019/315, <https://eprint.iacr.org/2019/315>),
- 10207 [166] Hanke T (2016) Asicboost-a speedup for bitcoin mining. *arXiv preprint*
10208 *arXiv:160400575* .
- 10209 [167] Recabarren R, Carbutar B (2017) Hardening stratum, the bitcoin pool mining pro-
10210 tocol. *Proceedings on Privacy Enhancing Technologies 2017(3):57–74*.
- 10211 [168] StopAndDecrypt (2019) Betterhash: Decentralizing bitcoin mining with new hash-
10212 ing protocols. Available at [https://medium.com/hackernoon/betterhash-decentraliz-
10213 ing-bitcoin-mining-with-new-hashing-protocols-291de178e3e0](https://medium.com/hackernoon/betterhash-decentralizing-bitcoin-mining-with-new-hashing-protocols-291de178e3e0).
- 10214 [169] Miller A, Kosba A, Katz J, Shi E (2015) Nonoutsourcable scratch-off puzzles to
10215 discourage bitcoin mining coalitions. *Proceedings of the 22nd ACM SIGSAC Con-
10216 ference on Computer and Communications Security*, pp 680–691.
- 10217 [170] Chepurnoy A, Saxena A (2020) Bypassing non-outsourcable proof-of-work
10218 schemes using collateralized smart contracts. *IACR Cryptol ePrint Arch 2020:44*.
- 10219 [171] tevador (2019) Randomx. Available at <https://github.com/tevador/RandomX>.
- 10220 [172] Han R, Sui Z, Yu J, Liu J, Chen S (2019) Fact and fiction: Challenging the honest
10221 majority assumption of permissionless blockchains. *IACR Cryptol ePrint Arch* .
- 10222 [173] Kwon Y, Kim H, Shin J, Kim Y (2019) Bitcoin vs. bitcoin cash: Coexistence or
10223 downfall of bitcoin cash? *2019 IEEE Symposium on Security and Privacy (SP)*
10224 (IEEE), pp 935–951.
- 10225 [174] Bissias G, Böhme R, Thibodeau D, Levine BN (2020) Pricing security in proof-of-

- 10226 work systems. *arXiv preprint arXiv:201203706* .
- 10227 [175] Garratt R, van Oordt MR (2019) Why fixed costs matter for proof-of-work based
10228 cryptocurrencies. *Available at SSRN* .
- 10229 [176] Mueller P (2020) Cryptocurrency mining: Asymmetric response to price movement.
10230 *Available at SSRN 3733026* .
- 10231 [177] Arnosti N, Weinberg SM (2018) Bitcoin: A natural oligopoly. *arXiv preprint*
10232 *arXiv:181108572* .
- 10233 [178] Romiti M, Judmayer A, Zamyatin A, Haslhofer B (2019) A deep dive into
10234 bitcoin mining pools: An empirical analysis of mining shares. *arXiv preprint*
10235 *arXiv:190505999* .
- 10236 [179] Cong LW, He Z, Li J (2019) Decentralized mining in centralized pools. *The Review*
10237 *of Financial Studies* .
- 10238 [180] Chatzigiannis P, Baldimtsi F, Griva I, Li J (2019) Diversification across mining
10239 pools: Optimal mining strategies under pow. *arXiv preprint arXiv:190504624* .
- 10240 [181] Elmandjra Y, Hsue D (2020) Bitcoin mining the evolution of a multibillion dollar
10241 industry. Available at [https://research.ark-invest.com/hubfs/1_Download_Files_A](https://research.ark-invest.com/hubfs/1_Download_Files_A_RK-Invest/White_Papers/ARKInvest_031220_Whitepaper_BitcoinMining.pdf)
10242 [RK-Invest/White_Papers/ARKInvest_031220_Whitepaper_BitcoinMining.pdf](https://research.ark-invest.com/hubfs/1_Download_Files_A_RK-Invest/White_Papers/ARKInvest_031220_Whitepaper_BitcoinMining.pdf).
- 10243 [182] Garay J, Kiayias A, Leonardos N (2017) The bitcoin backbone protocol with chains
10244 of variable difficulty. *Annual International Cryptology Conference* (Springer), pp
10245 291–323.
- 10246 [183] Chan TH, Ephraim N, Marcedone A, Morgan A, Pass R, Shi E (2017) Blockchain
10247 with varying number of players.
- 10248 [184] Garay JA, Kiayias A, Leonardos N (2020) Full analysis of nakamoto consensus in
10249 bounded-delay networks. *IACR Cryptol ePrint Arch* :277.
- 10250 [185] Davidson M, Diamond T (2020) On the profitability of selfish mining against multi-
10251 ple difficulty adjustment algorithms. *IACR Cryptol ePrint Arch* :94.
- 10252 [186] Werner SM, Ilie DI, Stewart I, Knottenbelt WJ (2020) Unstable throughput: When
10253 the difficulty algorithm breaks. *arXiv preprint arXiv:200603044* .
- 10254 [187] Meshkov D, Chepurnoy A, Jansen M (2017) Short paper: Revisiting difficulty
10255 control for blockchain systems. *Data Privacy Management, Cryptocurrencies and*
10256 *Blockchain Technology* (Springer), pp 429–436.
- 10257 [188] Goren G, Spiegelman A (2019) Mind the mining. *Proceedings of the 2019 ACM*
10258 *Conference on Economics and Computation*, pp 475–487.
- 10259 [189] Fiat A, Karlin A, Koutsoupias E, Papadimitriou C (2019) Energy equilibria in proof-
10260 of-work mining. *Proceedings of the 2019 ACM Conference on Economics and Com-*
10261 *putation*, pp 489–502.
- 10262 [190] Rosenfeld M (2011) Analysis of bitcoin pooled mining reward systems. *arXiv*
10263 *preprint arXiv:11124980* .
- 10264 [191] Luu L, Saha R, Parameshwaran I, Saxena P, Hobor A (2015) On power splitting
10265 games in distributed computation: The case of bitcoin pooled mining. *2015 IEEE*
10266 *28th Computer Security Foundations Symposium* (IEEE), pp 397–411.
- 10267 [192] Eyal I (2015) The miner’s dilemma. *2015 IEEE Symposium on Security and Privacy*

- 10268 (IEEE), pp 89–103.
- 10269 [193] wizkid057 (2014) Eligius: 0% fee btc, 105% pps nmc, no registration, cppsrb. Avail-
10270 able at <https://bitcointalk.org/?topic=441465.msg7282674>.
- 10271 [194] Kwon Y, Kim D, Son Y, Vasserman E, Kim Y (2017) Be selfish and avoid dilem-
10272 mas: Fork after withholding (faw) attacks on bitcoin. *Proceedings of the 2017 ACM*
10273 *SIGSAC Conference on Computer and Communications Security*, pp 195–209.
- 10274 [195] Schrijvers O, Bonneau J, Boneh D, Roughgarden T (2016) Incentive compatibil-
10275 ity of bitcoin mining pool reward functions. *International Conference on Financial*
10276 *Cryptography and Data Security* (Springer), pp 477–498.
- 10277 [196] Roughgarden T, Shikhelman C (2021) Ignore the extra zeroes: Variance-optimal
10278 mining pools. *International Conference on Financial Cryptography and Data Secu-*
10279 *rity* (Springer), pp 233–249.
- 10280 [197] Katz J, Lazos P, Marmolejo-Cossío FJ, Zhou X (2021) Rpplms: Pay-per-last-n-shares
10281 with a randomised twist. *arXiv preprint arXiv:210207681* .
- 10282 [198] Zolotavkin Y, García J, Rudolph C (2017) Incentive compatibility of pay per last
10283 n shares in bitcoin mining pools. *International Conference on Decision and Game*
10284 *Theory for Security* (Springer), pp 21–39.
- 10285 [199] Eyal I, Sirer EG (2014) Majority is not enough: Bitcoin mining is vulnerable. *In-*
10286 *ternational conference on financial cryptography and data security* (Springer), pp
10287 436–454.
- 10288 [200] Bahack L (2013) Theoretical bitcoin attacks with less than half of the computational
10289 power (draft). *arXiv preprint arXiv:13127013* .
- 10290 [201] Kiayias A, Koutsoupias E, Kyropoulou M, Tselekounis Y (2016) Blockchain mining
10291 games. *Proceedings of the 2016 ACM Conference on Economics and Computation*,
10292 pp 365–382.
- 10293 [202] Sapirshstein A, Sompolinsky Y, Zohar A (2016) Optimal selfish mining strategies
10294 in bitcoin. *International Conference on Financial Cryptography and Data Security*
10295 (Springer), pp 515–532.
- 10296 [203] Nayak K, Kumar S, Miller A, Shi E (2016) Stubborn mining: Generalizing selfish
10297 mining and combining with an eclipse attack. *2016 IEEE European Symposium on*
10298 *Security and Privacy (EuroS&P)* (IEEE), pp 305–320.
- 10299 [204] Göbel J, Keeler HP, Krzesinski AE, Taylor PG (2016) Bitcoin blockchain dynam-
10300 ics: The selfish-mine strategy in the presence of propagation delay. *Performance*
10301 *Evaluation* 104:23–41.
- 10302 [205] Gervais A, Karame GO, Wüst K, Glykantzis V, Ritzdorf H, Capkun S (2016) On
10303 the security and performance of proof of work blockchains. *Proceedings of the 2016*
10304 *ACM SIGSAC conference on computer and communications security* (ACM), pp 3–
10305 16.
- 10306 [206] Bai Q, Zhou X, Wang X, Xu Y, Wang X, Kong Q (2019) A deep dive into blockchain
10307 selfish mining. *ICC 2019-2019 IEEE International Conference on Communications*
10308 *(ICC)* (IEEE), pp 1–6.
- 10309 [207] Leelavimolsilp T, Tran-Thanh L, Stein S (2018) On the preliminary investigation of

- 10310 selfish mining strategy with multiple selfish miners. *arXiv preprint arXiv:180202218*
10311 .
- 10312 [208] Marmolejo-Cossío FJ, Brigham E, Sela B, Katz J (2019) Competing (semi-) selfish
10313 miners in bitcoin. *Proceedings of the 1st ACM Conference on Advances in Financial*
10314 *Technologies*, pp 89–109.
- 10315 [209] Zhang S, Zhang K, Kemme B (2020) Analysing the benefit of selfish mining with
10316 multiple players. *2020 IEEE International Conference on Blockchain (Blockchain)*
10317 (IEEE), pp 36–44.
- 10318 [210] Grunspan C, Pérez-Marco R (2019) Selfish mining in ethereum. *arXiv preprint*
10319 *arXiv:190413330* .
- 10320 [211] Ritz F, Zugenmaier A (2018) The impact of uncle rewards on selfish mining in
10321 ethereum. *2018 IEEE European Symposium on Security and Privacy Workshops*
10322 *(EuroS&PW)* (IEEE), pp 50–57.
- 10323 [212] Feng C, Niu J (2019) Selfish mining in ethereum. *2019 IEEE 39th International*
10324 *Conference on Distributed Computing Systems (ICDCS)* (IEEE), pp 1306–1316.
- 10325 [213] Shibuya Y, Yamamoto G, Kojima F, Shi E, Matsuo S, Laszka A (2021) Selfish min-
10326 ing attacks exacerbated by elastic hash supply. *arXiv preprint arXiv:210308007* .
- 10327 [214] Gramoli V (2020) From blockchain consensus back to byzantine consensus. *Future*
10328 *Generation Computer Systems* 107:760–769.
- 10329 [215] Kiayias A, Panagiotakos G (2015) Speed-security tradeoffs in blockchain protocols.
10330 *IACR Cryptol ePrint Arch* :1019.
- 10331 [216] Pass R, Seeman L, Shelat A (2017) Analysis of the blockchain protocol in asyn-
10332 chronous networks. *Annual International Conference on the Theory and Applica-*
10333 *tions of Cryptographic Techniques* (Springer), pp 643–673.
- 10334 [217] Wei P, Yuan Q, Zheng Y (2018) Security of the blockchain against long delay at-
10335 tack. *International Conference on the Theory and Application of Cryptology and*
10336 *Information Security* (Springer), pp 250–275.
- 10337 [218] Badertscher C, Maurer U, Tschudi D, Zikas V (2017) Bitcoin as a transaction ledger:
10338 A composable treatment. *Annual International Cryptology Conference* (Springer),
10339 pp 324–356.
- 10340 [219] Badertscher C, Garay J, Maurer U, Tschudi D, Zikas V (2018) But why does it work?
10341 a rational protocol design treatment of bitcoin. *Annual international conference on*
10342 *the theory and applications of cryptographic techniques* (Springer), pp 34–65.
- 10343 [220] Cojocararu A, Garay JA, Kiayias A, Song F, Wallden P (2019) The bitcoin backbone
10344 protocol against quantum adversaries. *IACR Cryptol ePrint Arch* :1150.
- 10345 [221] Ni P, Li H, Pan D (2020) Analysis of bitcoin backbone protocol in the non-flat model.
10346 *Science China Information Sciences* 63(3):1–14.
- 10347 [222] Zhao J (2019) An analysis of blockchain consistency in asynchronous networks:
10348 Deriving a neat bound. *arXiv preprint arXiv:190906587* .
- 10349 [223] Garay JA, Kiayias A, Panagiotakos G (2020) Consensus from signatures of work.
10350 *Cryptographers’ Track at the RSA Conference* (Springer), pp 319–344.
- 10351 [224] Duong T, Chepurnoy A, Zhou HS (2018) Multi-mode cryptocurrency systems. *Pro-*

- 10352 *ceedings of the 2nd ACM Workshop on Blockchains, Cryptocurrencies, and Con-*
10353 *tracts*, pp 35–46.
- 10354 [225] Dembo A, Kannan S, Tas EN, Tse D, Viswanath P, Wang X, Zeitouni O (2020)
10355 Everything is a race and nakamoto always wins. *arXiv preprint arXiv:200510484* .
- 10356 [226] Gazi P, Kiayias A, Russell A (2020) Tight consistency bounds for bitcoin. *IACR*
10357 *Cryptol ePrint Arch* .
- 10358 [227] Avarikioti G, Käppeli L, Wang Y, Wattenhofer R (2019) Bitcoin security under tem-
10359 porary dishonest majority. *International Conference on Financial Cryptography and*
10360 *Data Security* (Springer), pp 466–483.
- 10361 [228] Badertscher C, Gazi P, Kiayias A, Russell A, Zikas V (2020) Consensus redux:
10362 Distributed ledgers in the face of adversarial supremacy. *IACR Cryptol ePrint Arch*
10363 :1021.
- 10364 [229] Ren L (2019) Analysis of nakamoto consensus. *IACR Cryptol ePrint Arch* :943.
- 10365 [230] Kiffer L, Rajaraman R, Shelat A (2018) A better method to analyze blockchain con-
10366 sistency. *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Com-*
10367 *munications Security*, pp 729–744.
- 10368 [231] Cojocar A, Garay J, Kiayias A, Song F, Wallden P (2020) Post-quantum security
10369 of the bitcoin backbone and quantum multi-solution bernoulli search. *arXiv preprint*
10370 *arXiv:201215254* .
- 10371 [232] Li J, Guo D, Ren L (2020) Close latency–security trade-off for the nakamoto con-
10372 sensus. *arXiv preprint arXiv:201114051* .
- 10373 [233] Sankagiri S, Gandlur S, Hajek B (2021) The longest-chain protocol under random
10374 delays. *arXiv preprint arXiv:210200973* .
- 10375 [234] Decker C, Wattenhofer R (2013) Information propagation in the bitcoin network.
10376 *IEEE P2P 2013 Proceedings* (IEEE), pp 1–10.
- 10377 [235] Silva P, Vavříčka D, Barreto J, Matos M (2020) Impact of geo-distribution and min-
10378 ing pools on blockchains: A study of ethereum. *50th IEEE/IFIP International Con-*
10379 *ference on Dependable Systems and Networks (DSN)* .
- 10380 [236] Xiao Y, Zhang N, Lou W, Hou YT (2020) Modeling the impact of network
10381 connectivity on consensus security of proof-of-work blockchain. *arXiv preprint*
10382 *arXiv:200208912* .
- 10383 [237] Jain A, Siddiqui S, Gujar S (2021) We might walk together, but i run faster: Network
10384 fairness and scalability in blockchains. *arXiv preprint arXiv:210204326* .
- 10385 [238] Shahsavari Y, Zhang K, Talhi C (2020) A theoretical model for block propagation
10386 analysis in bitcoin network. *IEEE Transactions on Engineering Management* .
- 10387 [239] Corallo M (2016) Compact block relay. Available at [https://github.com/bitcoin/bip](https://github.com/bitcoin/bips/blob/master/bip-0152.mediawiki)
10388 [s/blob/master/bip-0152.mediawiki](https://github.com/bitcoin/bips/blob/master/bip-0152.mediawiki).
- 10389 [240] Otsuki K, Banno R, Shudo K (2020) Quantitatively analyzing relay networks in
10390 bitcoin. *2020 IEEE International Conference on Blockchain (Blockchain)* (IEEE),
10391 pp 214–220.
- 10392 [241] Budish E (2018) The economic limits of bitcoin and the blockchain (National Bureau
10393 of Economic Research),

- 10394 [242] Bonneau J, Felten EW, Goldfeder S, Kroll JA, Narayanan A (2016) Why buy when
10395 you can rent? bribery attacks on bitcoin consensus. *Citeseer* .
- 10396 [243] Liao K, Katz J (2017) Incentivizing blockchain forks via whale transactions. *Inter-*
10397 *national Conference on Financial Cryptography and Data Security* (Springer), pp
10398 264–279.
- 10399 [244] Teutsch J, Jain S, Saxena P (2016) When cryptocurrencies mine their own business.
10400 *International Conference on Financial Cryptography and Data Security* (Springer),
10401 pp 499–514.
- 10402 [245] McCorry P, Hicks A, Meiklejohn S (2018) Smart contracts for bribing miners. *Inter-*
10403 *national Conference on Financial Cryptography and Data Security* (Springer), pp
10404 3–18.
- 10405 [246] Winzer F, Herd B, Faust S (2019) Temporary censorship attacks in the presence of
10406 rational miners. *2019 IEEE European Symposium on Security and Privacy Work-*
10407 *shops (EuroS&PW)* (IEEE), pp 357–366.
- 10408 [247] Judmayer A, Stifter N, Zamyatin A, Tsabary I, Eyal I, Gazi P, Meiklejohn S, Weippl
10409 ER (2019) Pay-to-win: Cheap, crowdfundable, cross-chain incentive manipulation
10410 attacks on cryptocurrencies. *IACR Cryptol ePrint Arch* :775.
- 10411 [248] Moroz DJ, Aronoff DJ, Narula N, Parkes DC (2020) Double-spend counterattacks:
10412 Threat of retaliation in proof-of-work systems. *arXiv preprint arXiv:200210736* .
- 10413 [249] Rosenfeld M (2014) Analysis of hashrate-based double spending. *arXiv preprint*
10414 *arXiv:14022009* .
- 10415 [250] Grunspan C, Pérez-Marco R (2019) On profitability of nakamoto double spend.
10416 *arXiv preprint arXiv:191206412* .
- 10417 [251] Sompolinsky Y, Zohar A (2016) Bitcoin’s security model revisited. *arXiv preprint*
10418 *arXiv:160509193* .
- 10419 [252] Finney H (2011) Best practice for fast transaction acceptance - how high is the risk?.
10420 Available at <https://bitcointalk.org/index.php?topic=3441.msg48384#msg48384>.
- 10421 [253] vector76 (2011) Fake bitcoins?. Available at <https://bitcointalk.org/index.php?topic=36788.msg463391#msg463391>.
- 10422
- 10423 [254] Miller A (2013) Feather-forks: enforcing a blacklist with sub-50% hash power.
10424 Available at <https://bitcointalk.org/index.php?topic=312668.0>.
- 10425 [255] Rizun PR (2016) Subchains: A technique to scale bitcoin and improve the user
10426 experience. *Ledger* 1:38–52.
- 10427 [256] Zamyatin A, Stifter N, Schindler P, Weippl ER, Knottenbelt WJ (2018) Flux: Revis-
10428 iting near blocks for proof-of-work blockchains. *IACR Cryptol ePrint Arch* :415.
- 10429 [257] Eyal I, Gencer AE, Siler EG, Van Renesse R (2016) Bitcoin-ng: A scalable
10430 blockchain protocol. *13th {USENIX} symposium on networked systems design and*
10431 *implementation ({NSDI} 16)*, pp 45–59.
- 10432 [258] Yin J, Wang C, Zhang Z, Liu J (2018) Revisiting the incentive mechanism of bitcoin-
10433 ng. *Australasian Conference on Information Security and Privacy* (Springer), pp
10434 706–719.
- 10435 [259] Niu J, Wang Z, Gai F, Feng C (2020) Incentive analysis of bitcoin-ng, revisited.

- 10436 *arXiv preprint arXiv:200105082* .
- 10437 [260] Lerner SD (2015) Decor+ hop: A scalable blockchain protocol.
- 10438 [261] Camacho P, Lerner SD (2016) Decor+ lami: A scalable blockchain protocol.
- 10439 [262] Zhang R, Preneel B (2017) Publish or perish: A backward-compatible defense
10440 against selfish mining in bitcoin. *Cryptographers' Track at the RSA Conference*
10441 (Springer), pp 277–292.
- 10442 [263] Zhang R, Preneel B (2019) Lay down the common metrics: Evaluating proof-of-
10443 work consensus protocols' security. *2019 IEEE Symposium on Security and Privacy*
10444 (*SP*) (IEEE), pp 175–192.
- 10445 [264] Zhang R, Zhang D, Wang Q, Xie J, Preneel B (2020) Nc-max: Breaking the through-
10446 put limit of nakamoto consensus. *IACR Cryptol ePrint Arch* :1101.
- 10447 [265] Zhang R, Zhang D, Wang Q, Wu S, Xie J, Preneel B (2022) Nc-max: Breaking
10448 the security-performance tradeoff in nakamoto consensus. *29th Annual Network and*
10449 *Distributed System Security Symposium, NDSS 2022, San Diego, California, USA,*
10450 *April 24-28, 2022* (The Internet Society). Available at [https://www.ndss-symposium](https://www.ndss-symposium.org/ndss-paper/auto-draft-255/)
10451 [.org/ndss-paper/auto-draft-255/](https://www.ndss-symposium.org/ndss-paper/auto-draft-255/).
- 10452 [266] Sompolinsky Y, Zohar A (2015) Secure high-rate transaction processing in bitcoin.
10453 *International Conference on Financial Cryptography and Data Security* (Springer),
10454 pp 507–527.
- 10455 [267] Kiayias A, Panagiotakos G (2017) On trees, chains and fast transactions in the
10456 blockchain. *International Conference on Cryptology and Information Security in*
10457 *Latin America* (Springer), pp 327–351.
- 10458 [268] Lerner SD (2016) Uncle mining, an ethereum consensus protocol flaw. Available at
10459 [https://bitslog.com/2016/04/28/uncle-mining-an-ethereum-consensus-protocol-fla](https://bitslog.com/2016/04/28/uncle-mining-an-ethereum-consensus-protocol-flaw/)
10460 [w/](https://bitslog.com/2016/04/28/uncle-mining-an-ethereum-consensus-protocol-flaw/).
- 10461 [269] Chang SY, Park Y, Wuthier S, Chen CW (2019) Uncle-block attack: Blockchain
10462 mining threat beyond block withholding for rational and uncooperative miners. *In-*
10463 *ternational Conference on Applied Cryptography and Network Security* (Springer),
10464 pp 241–258.
- 10465 [270] Werner SM, Pritz PJ, Zamyatin A, Knottenbelt WJ (2019) Uncle traps: harvest-
10466 ing rewards in a queue-based ethereum mining pool. *Proceedings of the 12th EAI*
10467 *International Conference on Performance Evaluation Methodologies and Tools*, pp
10468 127–134.
- 10469 [271] Pass R, Shi E (2017) Fruitchains: A fair blockchain. *Proceedings of the ACM Sym-*
10470 *posium on Principles of Distributed Computing*, pp 315–324.
- 10471 [272] Fitzi M, Gazi P, Kiayias A, Russell A (2018) Parallel chains: Improving throughput
10472 and latency of blockchain protocols via parallel composition. *IACR Cryptol ePrint*
10473 *Arch* :1119.
- 10474 [273] Fitzi M, Gazi P, Kiayias A, Russell A (2020) Ledger combiners for fast settlement.
10475 *IACR Cryptol ePrint Arch* :675.
- 10476 [274] Wang X, Muppurala VV, Yang L, Kannan S, Viswanath P (2021) Securing parallel-
10477 chain protocols under variable mining power. *arXiv preprint arXiv:210502927* .

- 10478 [275] Bagaria V, Kannan S, Tse D, Fanti G, Viswanath P (2019) Prism: Deconstructing
10479 the blockchain to approach physical limits. *Proceedings of the 2019 ACM SIGSAC*
10480 *Conference on Computer and Communications Security*, pp 585–602.
- 10481 [276] Yang L, Bagaria V, Wang G, Alizadeh M, Tse D, Fanti G, Viswanath P (2019) Prism:
10482 Scaling bitcoin by 10,000 x. *arXiv preprint arXiv:190911261* .
- 10483 [277] Li J, Guo D (2020) Continuous-time analysis of the bitcoin and prism backbone
10484 protocols. *arXiv preprint arXiv:200105644* .
- 10485 [278] Li J, Guo D (2020) Liveness and consistency of bitcoin and prism blockchains: The
10486 non-lockstep synchronous case. *2020 IEEE International Conference on Blockchain*
10487 *and Cryptocurrency (ICBC) (IEEE)*, pp 1–9.
- 10488 [279] Wang G, Wang S, Bagaria V, Tse D, Viswanath P (2020) Prism removes consensus
10489 bottleneck for smart contracts. *arXiv preprint arXiv:200408776* .
- 10490 [280] Birmpas G, Koutsoupias E, Lazos P, Marmolejo-Cossío FJ (2020) Fairness and effi-
10491 ciency in dag-based cryptocurrencies. *International Conference on Financial Cryp-*
10492 *tography and Data Security (Springer)*, pp 79–96.
- 10493 [281] Lewenberg Y, Sompolinsky Y, Zohar A (2015) Inclusive block chain protocols. *In-*
10494 *ternational Conference on Financial Cryptography and Data Security (Springer)*, pp
10495 528–547.
- 10496 [282] Li C, Li P, Zhou D, Xu W, Long F, Yao A (2018) Scaling nakamoto consensus to
10497 thousands of transactions per second. *arXiv preprint arXiv:180503870* .
- 10498 [283] Li C, Li P, Zhou D, Yang Z, Wu M, Yang G, Xu W, Long F, Yao ACC (2020) A de-
10499 centralized blockchain with high throughput and fast confirmation. *2020 {USENIX}*
10500 *Annual Technical Conference ({USENIX}{ATC} 20)*, pp 515–528.
- 10501 [284] Li C, Long F, Yang G (2020) Ghastr: Breaking confirmation delay barrier in
10502 nakamoto consensus via adaptive weighted blocks. *arXiv preprint arXiv:200601072*
10503 .
- 10504 [285] Sompolinsky Y, Lewenberg Y, Zohar A (2016) Spectre: A fast and scalable cryp-
10505 tocurrency protocol. *IACR Cryptol ePrint Arch* :1159.
- 10506 [286] Sompolinsky Y, Zohar A (2018) Phantom: A scalable blockdag protocol. *IACR*
10507 *Cryptol ePrint Arch* :104.
- 10508 [287] Sompolinsky Y, Zohar A (2020) Phantom and ghostdag: A scalable generalization
10509 of nakamoto consensus, .
- 10510 [288] Popov S (2016) The tangle.
- 10511 [289] Bramas Q (2018) The stability and the security of the tangle. *HAL archives-ouvertes*
10512 *hal-01716111* .
- 10513 [290] Penzkofer A, Kusmierz B, Capossele A, Sanders W, Saa O (2020) Parasite chain
10514 detection in the iota protocol. *arXiv preprint arXiv:200413409* .
- 10515 [291] Kusmierz B, Sanders W, Penzkofer A, Capossele A, Gal A (2019) Properties of the
10516 tangle for uniform random and random walk tip selection. *2019 IEEE International*
10517 *Conference on Blockchain (Blockchain) (IEEE)*, pp 228–236.
- 10518 [292] Cullen A, Ferraro P, King C, Shorten R (2019) Distributed ledger technology for iot:
10519 Parasite chain attacks. *arXiv preprint arXiv:190400996* .

- 10520 [293] Bu G, Gürçan Ö, Potop-Butucaru M (2019) G-iota: Fair and confidence aware tan-
10521 gles. *IEEE INFOCOM 2019-IEEE Conference on Computer Communications Work-*
10522 *shops (INFOCOM WKSHPs)* (IEEE), pp 644–649.
- 10523 [294] Bu G, Hana W, Potop-Butucaru M (2019) Metamorphic iota. *arXiv preprint*
10524 *arXiv:190703628* .
- 10525 [295] Bentov I, Hubáček P, Moran T, Nadler A (2017) Tortoise and hares consensus:
10526 the meshcash framework for incentive-compatible, scalable cryptocurrencies. *IACR*
10527 *Cryptol ePrint Arch* :300.
- 10528 [296] Pass R, Shi E (2017) Hybrid consensus: Efficient consensus in the permission-
10529 less model. *31st International Symposium on Distributed Computing (DISC 2017)*
10530 (Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik).
- 10531 [297] Abraham I, Malkhi D, Nayak K, Ren L, Spiegelman A (2016) Solida: A
10532 blockchain protocol based on reconfigurable byzantine consensus. *arXiv preprint*
10533 *arXiv:161202916* .
- 10534 [298] Lewis-Pye A, Roughgarden T (2020) Resource pools and the cap theorem. *arXiv*
10535 *preprint arXiv:200610698* .
- 10536 [299] King S, Nadal S (2012) Ppcoin: Peer-to-peer crypto-currency with proof-of-stake.
10537 *self-published paper* .
- 10538 [300] Vasin P (2014) Blackcoin’s proof-of-stake protocol v2. Accessed on March 17, 2021.
10539 Available at <https://blackcoin.co/blackcoin-pos-protocol-v2-whitepaper.pdf>.
- 10540 [301] Earls J (2017) The missing explanation of proof of stake version 3. Accessed on
10541 March 17, 2021. Available at [http://earlz.net/view/2017/07/27/1904/the-missing-e-](http://earlz.net/view/2017/07/27/1904/the-missing-explanation-of-proof-of-stake-version)
10542 [xplanation-of-proof-of-stake-version](http://earlz.net/view/2017/07/27/1904/the-missing-explanation-of-proof-of-stake-version).
- 10543 [302] Nxt whitepaper. Accessed on March 17, 2021. Available at [https://nxtdocs.jelurida](https://nxtdocs.jelurida.com/Nxt_Whitepaper)
10544 [.com/Nxt_Whitepaper](https://nxtdocs.jelurida.com/Nxt_Whitepaper).
- 10545 [303] Delegated proof of stake (dpos). Accessed on March 18, 2021. Available at [https:](https://how.bitshares.works/en/master/technology/dpos.html)
10546 [//how.bitshares.works/en/master/technology/dpos.html](https://how.bitshares.works/en/master/technology/dpos.html).
- 10547 [304] Günther S (2018) Lisk - the mafia blockchain. Accessed on March 18, 2021. Avail-
10548 able at <https://medium.com/coinmonks/lisk-the-mafia-blockchain-47248915ae2f>.
- 10549 [305] Li C, Palanisamy B (2019) Incentivized blockchain-based social media platforms:
10550 A case study of steemit. *Proceedings of the 10th ACM Conference on Web Science*,
10551 pp 145–154.
- 10552 [306] Gersbach H, Mamageishvili A, Schneider M (2021) Vote delegation and misbehav-
10553 ior. *arXiv preprint arXiv:210208823* .
- 10554 [307] Gersbach H, Mamageishvili A, Schneider M (2021) Vote delegation favors minority.
10555 *arXiv preprint arXiv:210208835* .
- 10556 [308] Cevallos A, Stewart A (2020) A verifiably secure and proportional committee elec-
10557 tion rule. *arXiv preprint arXiv:200412990* .
- 10558 [309] Saleh F (2020) Blockchain without waste: Proof-of-stake. Available at SSRN
10559 [3183935](https://ssrn.com/abstract=3183935) .
- 10560 [310] Ganesh C, Orlandi C, Tschudi D, Zohar A (2020) Virtual asics: Generalized proof-
10561 of-stake mining in cryptocurrencies. *IACR Cryptol ePrint Arch* :791.

- 10562 [311] Bentov I, Gabizon A, Mizrahi A (2016) Cryptocurrencies without proof of work.
10563 *International conference on financial cryptography and data security* (Springer), pp
10564 142–157.
- 10565 [312] Kanjalkar S, Kuo J, Li Y, Miller A (2019) Short paper: I can’t believe it’s not stake!
10566 resource exhaustion attacks on pos. *International Conference on Financial Cryptog-*
10567 *raphy and Data Security* (Springer), pp 62–69.
- 10568 [313] Bentov I, Pass R, Shi E (2016) Snow white: Provably secure proofs of stake. *IACR*
10569 *Cryptol ePrint Arch* :919.
- 10570 [314] Daian P, Pass R, Shi E (2019) Snow white: Robustly reconfigurable consensus and
10571 applications to provably secure proof of stake. *International Conference on Finan-*
10572 *cial Cryptography and Data Security* (Springer), pp 23–41.
- 10573 [315] Stütz R, Gaži P, Haslhofer B, Illum J (2020) Stake shift in major cryptocurrencies:
10574 An empirical study. *International Conference on Financial Cryptography and Data*
10575 *Security* (Springer), pp 97–113.
- 10576 [316] Mills DL (2017) *Computer network time synchronization: the network time protocol*
10577 *on earth and in space* (CRC press).
- 10578 [317] Gaži P, Kiayias A, Russell A (2018) Stake-bleeding attacks on proof-of-stake
10579 blockchains. *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)*
10580 (IEEE), pp 85–92.
- 10581 [318] Deirmentzoglou E, Papakyriakopoulos G, Patsakis C (2019) A survey on long-range
10582 attacks for proof of stake protocols. *IEEE Access* 7:28712–28725.
- 10583 [319] David B, Gaži P, Kiayias A, Russell A (2018) Ouroboros praos: An adaptively-
10584 secure, semi-synchronous proof-of-stake blockchain. *Annual International Confer-*
10585 *ence on the Theory and Applications of Cryptographic Techniques* (Springer), pp
10586 66–98.
- 10587 [320] Deb S, Kannan S, Tse D (2020) Posat: Proof-of-work availability and unpredictabil-
10588 ity, without the work. *arXiv preprint arXiv:201008154* .
- 10589 [321] Boneh D, Eskandarian S, Hanzlik L, Greco N (2020) Single secret leader election.
10590 *IACR Cryptol ePrint Arch* :25.
- 10591 [322] Catalano D, Fiore D, Giunta E (2021) Efficient and universally composable single
10592 secret leader election from pairings. *IACR Cryptol ePrint Arch* :344.
- 10593 [323] Kerber T, Kiayias A, Kohlweiss M, Zikas V (2019) Ouroboros cryptsinous: Privacy-
10594 preserving proof-of-stake. *2019 IEEE Symposium on Security and Privacy (SP)*
10595 (IEEE), pp 157–174.
- 10596 [324] Ganesh C, Orlandi C, Tschudi D (2019) Proof-of-stake protocols for privacy-aware
10597 blockchains. *Annual International Conference on the Theory and Applications of*
10598 *Cryptographic Techniques* (Springer), pp 690–719.
- 10599 [325] Baldimtsi F, Madathil V, Scafuro A, Zhou L (2020) Anonymous lottery in the proof-
10600 of-stake setting. *IACR Cryptol ePrint Arch* :533.
- 10601 [326] Kohlweiss M, Madathil V, Nayak K, Scafuro A (2021) On the anonymity guarantees
10602 of anonymous proof-of-stake protocols. *IACR Cryptol ePrint Arch* :409.
- 10603 [327] Brown-Cohen J, Narayanan A, Psomas A, Weinberg SM (2019) Formal barriers to

- 10604 longest-chain proof-of-stake protocols. *Proceedings of the 2019 ACM Conference*
10605 *on Economics and Computation*, pp 459–473.
- 10606 [328] Bagaria V, Dembo A, Kannan S, Oh S, Tse D, Viswanath P, Wang X, Zeitouni
10607 O (2019) Proof-of-stake longest chain protocols: Security vs predictability. *arXiv*
10608 *preprint arXiv:191002218* .
- 10609 [329] Neuder M, Moroz DJ, Rao R, Parkes DC (2019) Selfish behavior in the tezos proof-
10610 of-stake protocol. *arXiv preprint arXiv:191202954* .
- 10611 [330] Neuder M, Moroz DJ, Rao R, Parkes DC (2020) Defending against malicious reorgs
10612 in tezos proof-of-stake. *Proceedings of the 2nd ACM Conference on Advances in*
10613 *Financial Technologies*, pp 46–58.
- 10614 [331] Blum E, Kiayias A, Moore C, Quader S, Russell A (2020) The combinatorics of the
10615 longest-chain rule: Linear consistency for proof-of-stake blockchains. *Proceedings*
10616 *of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (SIAM),
10617 pp 1135–1154.
- 10618 [332] Kiayias A, Quader S, Russell A (2020) Consistency of proof-of-stake blockchains
10619 with concurrent honest slot leaders. *arXiv preprint arXiv:200106403* .
- 10620 [333] Fanti G, Kogan L, Oh S, Ruan K, Viswanath P, Wang G (2019) Compounding of
10621 wealth in proof-of-stake cryptocurrencies. *International Conference on Financial*
10622 *Cryptography and Data Security* (Springer), pp 42–61.
- 10623 [334] Rosu I, Saleh F (2020) Evolution of shares in a proof-of-stake cryptocurrency. *HEC*
10624 *Paris Research Paper No FIN-2019-1339* .
- 10625 [335] Irresberger F (2018) Coin concentration of proof-of-stake blockchains. *Leeds Uni-*
10626 *versity Business School Working Paper* .
- 10627 [336] Chitra T (2019) Competitive equilibria between staking and on-chain lending. *arXiv*
10628 *preprint arXiv:200100919* .
- 10629 [337] Chitra T, Evans A (2020) Why stake when you can borrow? *Available at SSRN*
10630 *3629988* .
- 10631 [338] Amoussou-Guenou Y, Del Pozzo A, Potop-Butucaru M, Tucci-Piergiovanni
10632 S (2019) On fairness in committee-based blockchains. *arXiv preprint*
10633 *arXiv:191009786* .
- 10634 [339] Leonardos S, Reijsbergen D, Piliouras G (2019) Weighted voting on the blockchain:
10635 Improving consensus in proof of stake protocols. *2019 IEEE International Confer-*
10636 *ence on Blockchain and Cryptocurrency (ICBC)* (IEEE), pp 376–384.
- 10637 [340] (2021) Proof of stake crypto assets ranking and yields. Accessed on April 4, 2023.
10638 Available at <https://www.stakingrewards.com/proof-of-stake>.
- 10639 [341] Karakostas D, Kiayias A, Larangeira M (2020) Account management in proof of
10640 stake ledgers. *IACR Cryptol ePrint Arch* :525.
- 10641 [342] Brünjes L, Kiayias A, Koutsoupas E, Stouka AP (2018) Reward sharing schemes
10642 for stake pools. *arXiv preprint arXiv:180711218* .
- 10643 [343] Karakostas D, Kiayias A, Nasikas C, Zindros D (2019) Cryptocurrency egalitarian-
10644 ism: A quantitative approach. *International Conference on Blockchain Economics,*
10645 *Security and Protocols (Tokenomics 2019)* (Schloss Dagstuhl-Leibniz-Zentrum fuer

- 10646 Informatik).
- 10647 [344] Catalini C, Jagadeesan R, Kominers SD (2020) Markets for crypto tokens, and secu-
10648 rity under proof of stake. *Available at SSRN* .
- 10649 [345] Badertscher C, Gaži P, Kiayias A, Russell A, Zikas V (2018) Ouroboros genesis:
10650 Composable proof-of-stake blockchains with dynamic availability. *Proceedings of*
10651 *the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pp
10652 913–930.
- 10653 [346] Badertscher C, Gazi P, Kiayias A, Russell A, Zikas V (2019) Ouroboros chronos:
10654 Permissionless clock synchronization via proof-of-stake. *IACR Cryptol ePrint Arch*
10655 :838.
- 10656 [347] Hanke T, Movahedi M, Williams D (2018) Dfinity technology overview series, con-
10657 sensus system. *arXiv preprint arXiv:180504548* .
- 10658 [348] Abraham I, Malkhi D, Nayak K, Ren L (2018) Dfinity consensus, explored. *IACR*
10659 *Cryptol ePrint Arch* :1153.
- 10660 [349] Ethereum 2.0 specifications. Available at <https://github.com/ethereum/eth2.0-specs>.
- 10661 [350] Buterin V, Hernandez D, Kamphofner T, Pham K, Qiao Z, Ryan D, Sin J, Wang Y,
10662 Zhang YX (2020) Combining ghost and casper. *arXiv preprint arXiv:200303052* .
- 10663 [351] Buterin V Vitalik’s annotated ethereum 2.0 spec. Available at [https://github.com/eth](https://github.com/ethereum/annotated-spec/blob/master/phase0/beacon-chain.md)
10664 [ereum/annotated-spec/blob/master/phase0/beacon-chain.md](https://github.com/ethereum/annotated-spec/blob/master/phase0/beacon-chain.md).
- 10665 [352] Drake J (2018) Ethereum 2.0 randomness using a verifiable delay function (vdf).
10666 Available at [https://docs.google.com/presentation/d/1VKBMcEMaY6PQjDJK4y](https://docs.google.com/presentation/d/1VKBMcEMaY6PQjDJK4yAFtn0vD41eYXW05YA3vdTBsI/edit#slide=id.g4677b9389c_0_732)
10667 [AFtn0vD41eYXW05YA3vdTBsI/edit#slide=id.g4677b9389c_0_732](https://docs.google.com/presentation/d/1VKBMcEMaY6PQjDJK4yAFtn0vD41eYXW05YA3vdTBsI/edit#slide=id.g4677b9389c_0_732).
- 10668 [353] Azouvi S, McCorry P, Meiklejohn S (2018) Betting on blockchain consensus with
10669 fantomette. *arXiv preprint arXiv:180506786* .
- 10670 [354] Rocket T, Yin M, Sekniqi K, van Renesse R, Siler EG (2019) Scalable
10671 and probabilistic leaderless bft consensus through metastability. *arXiv preprint*
10672 *arXiv:190608936* .
- 10673 [355] Auvolet A, Bromberg YD, Frey D, Tañani F (2021) Basalt: A rock-solid foundation
10674 for epidemic consensus algorithms in very large, very open networks. *arXiv preprint*
10675 *arXiv:210204063* .
- 10676 [356] Popov S, Buchanan WJ (2019) Fpc-bi: Fast probabilistic consensus within byzantine
10677 infrastructures. *arXiv preprint arXiv:190510895* .
- 10678 [357] Fitzi M, Gazi P, Kiayias A, Russell A (2020) Proof-of-stake blockchain protocols
10679 with near-optimal throughput. *IACR Cryptol ePrint Arch* :37.
- 10680 [358] Buchman E, Kwon J, Milosevic Z (2018) The latest gossip on bft consensus. *arXiv*
10681 *preprint arXiv:180704938* .
- 10682 [359] Amoussou-Guenou Y, Del Pozzo A, Potop-Butucaru M, Tucci-Piergiovanni S
10683 (2018) Correctness and fairness of tendermint-core blockchains. *arXiv preprint*
10684 *arXiv:180508429* .
- 10685 [360] Amoussou-Guenou Y, Del Pozzo A, Potop-Butucaru M, Tucci-Piergiovanni S
10686 (2019) Dissecting tendermint. *International Conference on Networked Systems*
10687 (Springer), pp 166–182.

- 10688 [361] Braithwaite S, Buchman E, Khoffi I, Konnov I, Milosevic Z, Ruetschi R, Widder J
10689 (2020) A tendermint light client. *arXiv preprint arXiv:201007031* .
- 10690 [362] Gilad Y, Hemo R, Micali S, Vlachos G, Zeldovich N (2017) Algorand: Scaling
10691 byzantine agreements for cryptocurrencies. *Proceedings of the 26th Symposium on*
10692 *Operating Systems Principles*, pp 51–68.
- 10693 [363] Goyal V, Li H, Raizes J (2021) Instant block confirmation in the sleepy model. *In-*
10694 *ternational Conference on Financial Cryptography and Data Security* .
- 10695 [364] Dziembowski S, Faust S, Kolmogorov V, Pietrzak K (2015) Proofs of space. *Annual*
10696 *Cryptology Conference* (Springer), pp 585–605.
- 10697 [365] Park S, Kwon A, Fuchsbauer G, Gaži P, Alwen J, Pietrzak K (2018) Spacemint:
10698 A cryptocurrency based on proofs of space. *International Conference on Financial*
10699 *Cryptography and Data Security* (Springer), pp 480–499.
- 10700 [366] Cohen B, Pietrzak K (2019) The chia network blockchain.
- 10701 [367] Bentov I, Lee C, Mizrahi A, Rosenfeld M (2014) Proof of activity: Extending bit-
10702 coin’s proof of work via proof of stake. *ACM SIGMETRICS Performance Evaluation*
10703 *Review* 42(3):34–37.
- 10704 [368] Karakostas D, Kiayias A (2020) Securing proof-of-work ledgers via checkpointing.
10705 *IACR Cryptol ePrint Arch* :173.
- 10706 [369] Block A (2018) Chainlocks. Available at [https://github.com/dashpay/dips/blob/mas-](https://github.com/dashpay/dips/blob/master/dip-0008.md)
10707 [ter/dip-0008.md](https://github.com/dashpay/dips/blob/master/dip-0008.md).
- 10708 [370] Buterin V, Griffith V (2017) Casper the friendly finality gadget. *arXiv preprint*
10709 *arXiv:171009437* .
- 10710 [371] Buterin V, Reijersbergen D, Leonardos S, Piliouras G (2019) Incentives in ethereum’s
10711 hybrid casper protocol. *2019 IEEE international conference on blockchain and cryp-*
10712 *tocurrency (ICBC)* (IEEE), pp 236–244.
- 10713 [372] Nrryuya (2019) Analysis of bouncing attack on ffg. Available at [https://ethresear.ch](https://ethresear.ch/t/analysis-of-bouncing-attack-on-ffg/6113)
10714 [/t/analysis-of-bouncing-attack-on-ffg/6113](https://ethresear.ch/t/analysis-of-bouncing-attack-on-ffg/6113).
- 10715 [373] Neu J, Tas EN, Tse D (2020) Ebb-and-flow protocols: A resolution of the
10716 availability-finality dilemma. *arXiv preprint arXiv:200904987* .
- 10717 [374] Neu J, Tas EN, Tse D (2021) The availability-accountability dilemma and its reso-
10718 lution via accountability gadgets. *arXiv preprint arXiv:210506075* .
- 10719 [375] Dinsdale-Young T, Magri B, Matt C, Nielsen JB, Tschudi D (2019) Afgjort: A par-
10720 tially synchronous finality layer for blockchains, .
- 10721 [376] Stewart A, Kokoris-Kogia E (2020) Grandpa: a byzantine finality gadget. *arXiv*
10722 *preprint arXiv:200701560* .
- 10723 [377] Azouvi S, Danezis G, Nikolaenko V (2019) Winkle: Foiling long-range attacks in
10724 proof-of-stake systems. *IACR Cryptol ePrint Arch* :1440.
- 10725 [378] Neu J, Tas EN, Tse D (2020) Snap-and-chat protocols: System aspects. *arXiv*
10726 *preprint arXiv:201010447* .
- 10727 [379] Sankagiri S, Wang X, Kannan S, Viswanath P (2020) The checkpointed longest
10728 chain: User-dependent adaptivity and finality. *arXiv preprint arXiv:201013711* .
- 10729 [380] Luu L, Narayanan V, Zheng C, Baweja K, Gilbert S, Saxena P (2016) A secure

- 10730 sharding protocol for open blockchains. *Proceedings of the 2016 ACM SIGSAC Con-*
10731 *ference on Computer and Communications Security*, pp 17–30.
- 10732 [381] Kokoris-Kogias E, Jovanovic P, Gasser L, Gailly N, Syta E, Ford B (2018) Om-
10733 niledger: A secure, scale-out, decentralized ledger via sharding. *2018 IEEE Sympo-*
10734 *sium on Security and Privacy (SP)* (IEEE), pp 583–598.
- 10735 [382] Zamani M, Movahedi M, Raykova M (2018) Rapidchain: Scaling blockchain via
10736 full sharding. *Proceedings of the 2018 ACM SIGSAC Conference on Computer and*
10737 *Communications Security*, pp 931–948.
- 10738 [383] Wang J, Wang H (2019) Monoxide: Scale out blockchains with asynchronous con-
10739 sensus zones. *16th {USENIX} Symposium on Networked Systems Design and Im-*
10740 *plementation ({NSDI} 19)*, pp 95–112.
- 10741 [384] Team Z, et al. (2017) The zilliqa technical whitepaper. Retrieved September 16:2019.
- 10742 [385] Skidanov A, Polosukhin I (2019) Nightshade: Near protocol sharding design. URL:
10743 <https://nearprotocol.com/downloads/Nightshade.pdf> :39.
- 10744 [386] Al-Bassam M, Sonnino A, Bano S, Hrycyszyn D, Danezis G (2017) Chainspace: A
10745 sharded smart contracts platform. *arXiv preprint arXiv:170803778* .
- 10746 [387] Li S, Yu M, Avestimehr S, Kannan S, Viswanath P (2018) Polyshard: Coded shard-
10747 ing achieves linearly scaling efficiency and security simultaneously. *arXiv preprint*
10748 *arXiv:180910361* .
- 10749 [388] Das S, Krishnan V, Ren L (2020) Efficient cross-shard transaction execution in
10750 sharded blockchains. *arXiv preprint arXiv:200714521* .
- 10751 [389] Rana R, Kannan S, Tse D, Viswanath P (2020) Free2shard: Adaptive-adversary-
10752 resistant sharding via dynamic self allocation. [2005.09610](https://arxiv.org/abs/2005.09610).
- 10753 [390] Wang C, Raviv N (2020) Low latency cross-shard transactions in coded blockchain.
10754 *arXiv preprint arXiv:201100087* .
- 10755 [391] Hellings J, Hughes DP, Primero J, Sadoghi M (2020) Cerberus: Minimalistic multi-
10756 shard byzantine-resilient transaction processing. *arXiv preprint arXiv:200804450* .
- 10757 [392] Androulaki E, De Caro A, Elkhiyaoui K, Gorenflo C, Sorniotti A, Vukolic M
10758 (2020) Multi-shard private transactions for permissioned blockchains. *arXiv preprint*
10759 *arXiv:201008274* .
- 10760 [393] David B, Magri B, Matt C, Nielsen JB, Tschudi D (2021) Gearbox: An efficient uc
10761 sharded ledger leveraging the safety-liveness dichotomy. *IACR Cryptol ePrint Arch*
10762 :211.
- 10763 [394] Wang G, Shi ZJ, Nixon M, Han S (2019) Sok: Sharding on blockchain. *Proceedings*
10764 *of the 1st ACM Conference on Advances in Financial Technologies*, pp 41–61.
- 10765 [395] Liu Y, Liu J, Salles MAV, Zhang Z, Li T, Hu B, Henglein F, Lu R (2021) Building
10766 blocks of sharding blockchain systems: Concepts, approaches, and open problems.
10767 *arXiv preprint arXiv:210213364* .
- 10768 [396] Yu G, Wang X, Yu K, Ni W, Zhang JA, Liu RP (2020) Survey: Sharding in
10769 blockchains. *IEEE Access* 8:14155–14181.
- 10770 [397] Avarikioti G, Kokoris-Kogias E, Wattenhofer R (2019) Divide and scale: Formaliza-
10771 tion of distributed ledger sharding protocols. *arXiv preprint arXiv:191010434* .

- 10772 [398] Manshaei MH, Jadliwala M, Maiti A, Fooladgar M (2018) A game-theoretic analysis
10773 of shard-based permissionless blockchains. *IEEE Access* 6:78100–78112.
- 10774 [399] Han R, Yu J, Zhang R (2020) Analysing and improving shard allocation protocols
10775 for sharded blockchains. *IACR Cryptol ePrint Arch* :943.
- 10776 [400] Sonnino A, Bano S, Al-Bassam M, Danezis G (2019) Replay attacks and de-
10777 fenses against cross-shard consensus in sharded distributed ledgers. *arXiv preprint*
10778 *arXiv:190111218* .
- 10779 [401] Mizrahi A, Rottenstreich O (2020) Blockchain state sharding with space-aware rep-
10780 resentations. *IEEE Transactions on Network and Service Management* .
- 10781 [402] Buterin V (2018) A note on data availability and erasure coding. Available at <https://github.com/ethereum/research/wiki/A-note-on-data-availability-and-erasure-coding>.
10782
10783
- 10784 [403] Al-Bassam M, Sonnino A, Buterin V (2018) Fraud and data availability proofs:
10785 Maximising light client security and scaling blockchains with dishonest majorities.
10786 *arXiv preprint arXiv:180909044* .
- 10787 [404] Yu M, Sahraei S, Li S, Avestimehr S, Kannan S, Viswanath P (2020) Coded merkle
10788 tree: Solving data availability attacks in blockchains. *International Conference on*
10789 *Financial Cryptography and Data Security* (Springer), pp 114–134.
- 10790 [405] Mitra D, Taus L, Dolecek L (2020) Concentrated stopping set design for coded
10791 merkle tree: Improving security against data availability attacks in blockchain sys-
10792 tems. *arXiv preprint arXiv:201007363* .
- 10793 [406] Sheng P, Xue B, Kannan S, Viswanath P (2020) Aced: Scalable data availability
10794 oracle. *arXiv preprint arXiv:201100102* .
- 10795 [407] Buterin V (2017) On sharding blockchains. *Sharding FAQ* .
- 10796 [408] Buterin V (2017) The stateless client concept. Available at <https://ethresear.ch/t/the-stateless-client-concept/172>.
10797
- 10798 [409] Kuszmaul, J.: Verkle trees. In: Verkle Trees, pp. 1–12 (2019).
10799 <https://math.mit.edu/research/highschool/primes/materials/2018/Kuszmaul.pdf>.
- 10800 [410] Belchior R, Vasconcelos A, Guerreiro S, Correia M (2020) A survey on blockchain
10801 interoperability: Past, present, and future trends. *arXiv preprint arXiv:200514282* .
- 10802 [411] Robinson P (2020) Consensus for crosschain communications. *arXiv preprint*
10803 *arXiv:200409494* .
- 10804 [412] Zamyatin A, Al-Bassam M, Zindros D, Kokoris-Kogias E, Moreno-Sanchez P, Ki-
10805 ayias A, Knottenbelt WJ (2019) Sok: communication across distributed ledgers.
10806 *Cryptology ePrint Archive* :1128.
- 10807 [413] Pagnia H, Gärtner FC (1999) On the impossibility of fair exchange without a trusted
10808 third party (Citeseer),
- 10809 [414] Dziembowski S, Ekeley L, Faust S (2018) Fairswap: How to fairly exchange digital
10810 goods. *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Com-*
10811 *munications Security*, pp 967–984.
- 10812 [415] Ekeley L, Faust S, Schlosser B (2020) Optiswap: Fast optimistic fair exchange. *Pro-*
10813 *ceedings of the 15th ACM Asia Conference on Computer and Communications Se-*

- 10814 *curity*, pp 543–557.
- 10815 [416] Shook JM, Simon S, Mell P (2019) A smart contract refereed data retrieval protocol
10816 with a provably low collateral requirement. *IACR Cryptol ePrint Arch* :541.
- 10817 [417] Hall-Andersen M (2019) Fastswap: Concretely efficient contingent payments for
10818 complex predicates. *IACR Cryptol ePrint Arch* :1296.
- 10819 [418] Janin S, Qin K, Mamageishvili A, Gervais A (2020) Filebounty: Fair data exchange.
10820 *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*
10821 (IEEE), pp 357–366.
- 10822 [419] TierNolan (2013) Alt chains and atomic transfers. Available at <https://bitcointalk.org/index.php?topic=193281.msg2224949#msg2224949>.
- 10823 [420] Herlihy M (2018) Atomic cross-chain swaps. *Proceedings of the 2018 ACM symposium on principles of distributed computing*, pp 245–254.
- 10824 [421] van der Meyden R (2019) On the specification and verification of atomic swap smart
10825 contracts. *2019 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)* (IEEE), pp 176–179.
- 10826 [422] Liu JA (2018) Atomic swaptions: cryptocurrency derivatives. *arXiv preprint*
10827 *arXiv:180708644* .
- 10828 [423] Lesavre L, Varin P, Yaga D (2021) Blockchain networks: Token design and management overview (National Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency or Internal Report (IR) 8301. <https://doi.org/10.6028/NIST.IR.8301>.
- 10829 [424] Shlomovits O, Leiba O (2020) Jugglingswap: Scriptless atomic cross-chain swaps. *arXiv preprint* *arXiv:200714423* .
- 10830 [425] Zakhary V, Agrawal D, Abbadi AE (2019) Atomic commitment across blockchains. *arXiv preprint* *arXiv:190502847* .
- 10831 [426] Gugger J (2020) Bitcoin–monero cross-chain atomic swap. *IACR Cryptol ePrint Arch* :1126.
- 10832 [427] Hoenisch P, Soriano del Pino L (2021) Atomic swaps between bitcoin and monero. *arXiv preprint* *arXiv:210112332* .
- 10833 [428] Herlihy M, Liskov B, Shrira L (2019) Cross-chain deals and adversarial commerce. *arXiv preprint* *arXiv:190509743* .
- 10834 [429] van Glabbeek R, Gramoli V, Tholoniati P (2019) Cross-chain payment protocols with success guarantees. *arXiv preprint* *arXiv:191204513* .
- 10835 [430] Black M, Liu T, Cai T (2019) Atomic loans: Cryptocurrency debt instruments. *arXiv preprint* *arXiv:190105117* .
- 10836 [431] Liu Z, Xiang Y, Shi J, Gao P, Wang H, Xiao X, Wen B, Hu YC (2019) Hyperservice: Interoperability and programmability across heterogeneous blockchains. *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pp 549–566.
- 10837 [432] Fynn E, Bessani A, Pedone F (2020) Smart contracts on the move. *arXiv preprint* *arXiv:200405933* .
- 10838 [433] Robinson P, Hyland-Wood D, Saltini R, Johnson S, Brainard J (2019) Atomic cross-

- 10856 chain transactions for ethereum private sidechains. *arXiv preprint arXiv:190412079*
10857 .
- 10858 [434] Robinson P, Ramesh R (2020) General purpose atomic crosschain transactions.
10859 *arXiv preprint arXiv:201112783* .
- 10860 [435] Nissl M, Sallinger E, Schulte S, Borkowski M (2020) Towards cross-blockchain
10861 smart contracts. *arXiv preprint arXiv:201007352* .
- 10862 [436] Ghaemi S, Rouhani S, Belchior R, Cruz RS, Khazaei H, Musilek P (2021)
10863 A pub-sub architecture to promote blockchain interoperability. *arXiv preprint*
10864 *arXiv:2101:12331* .
- 10865 [437] Xu J, Ackerer D, Dubovitskaya A (2020) A game-theoretic analysis of cross-chain
10866 atomic swaps with htles. *arXiv preprint arXiv:201111325* .
- 10867 [438] Han R, Lin H, Yu J (2019) On the optionality and fairness of atomic swaps. *Proceed-*
10868 *ings of the 1st ACM Conference on Advances in Financial Technologies*, pp 62–75.
- 10869 [439] Rueegger J, Machado GS (2020) Rational exchange: Incentives in atomic cross
10870 chain swaps. *2020 IEEE International Conference on Blockchain and Cryptocur-*
10871 *rency (ICBC)* (IEEE), pp 1–3.
- 10872 [440] Xue Y, Herlihy M (2021) Hedging against sore loser attacks in cross-chain transac-
10873 tions. *arXiv preprint arXiv:210506322* .
- 10874 [441] Zamyatin A (2016) *Merged Mining: Analysis of Effects and Implications*. Ph.D. the-
10875 sis. PhD thesis. Master’s thesis, Vienna University of Technology, 2016 (cit. on . . . ,
10876 .
- 10877 [442] Merged mining introduction. Available at [https://tlu.tarilabs.com/merged-mining/
10878 merged-mining-scene/MergedMiningIntroduction.html](https://tlu.tarilabs.com/merged-mining/merged-mining-scene/MergedMiningIntroduction.html).
- 10879 [443] Research B (2020) The growth of bitcoin merge mining. Available at [https://blog.b
10880 itmex.com/the-growth-of-bitcoin-merge-mining/](https://blog.bitmex.com/the-growth-of-bitcoin-merge-mining/).
- 10881 [444] Judmayer A, Zamyatin A, Stifter N, Voyiatzis AG, Weippl E (2017) Merged min-
10882 ing: Curse or cure? *Data Privacy Management, Cryptocurrencies and Blockchain*
10883 *Technology* (Springer), pp 316–333.
- 10884 [445] Back A, Corallo M, Dashjr L, Friedenbach M, Maxwell G, Miller A, Poelstra A,
10885 Timón J, Wuille P (2014) Enabling blockchain innovations with pegged sidechains .
- 10886 [446] Sztork P, CryptAxe (2019) Bip 301. Available at [https://github.com/bitcoin/bips/bl
10887 ob/master/bip-0301.mediawiki](https://github.com/bitcoin/bips/blob/master/bip-0301.mediawiki).
- 10888 [447] Sztork P (2015) Drivechain - the simple two way peg. Available at [http://www.tru
10889 tcoin.info/blog/drivechain/](http://www.truthcoin.info/blog/drivechain/).
- 10890 [448] Sztork P, CryptAxe (2017) Drivechain documentation – hashrate escrows. Available
10891 at [https://github.com/drivechain-project/docs/blob/master/bip1-hashrate-escrow.
10892 md](https://github.com/drivechain-project/docs/blob/master/bip1-hashrate-escrow.md).
- 10893 [449] Weippl E (2018) Pitchforks in cryptocurrencies: Enforcing rule changes through of-
10894 fensive. *Data Privacy Management, Cryptocurrencies and Blockchain Technology:
10895 ESORICS 2018 International Workshops, DPM 2018 and CBT 2018, Barcelona,
10896 Spain, September 6-7, 2018, Proceedings* (Springer), Vol. 11025, p 197.
- 10897 [450] Karantias K, Kiayias A, Zindros D (2019) Proof-of-burn. *IACR Cryptol ePrint Arch*

- 10898 2019:1096.
- 10899 [451] Ali M, Blankstein A, Freedman MJ, Galabru L, Gupta D, Nelson J, Soslow J, Stanley
10900 P (2020) Pox: Proof of transfer mining with bitcoin. *Blockstack PBC* .
- 10901 [452] Nelson J (2021) What kind of blockchain is stacks?. Available at <https://stacks.org/s>
10902 [tacks-blockchain](https://stacks.org/s/tacks-blockchain).
- 10903 [453] Robinson P (2018) Requirements for ethereum private sidechains. *arXiv preprint*
10904 *arXiv:180609834* .
- 10905 [454] Dilley J, Poelstra A, Wilkins J, Piekarska M, Gorlick B, Friedenbach M (2016)
10906 Strong federations: An interoperable blockchain solution to centralized third-party
10907 risks. *arXiv preprint arXiv:161205491* .
- 10908 [455] Nick J, Poelstra A, Sanders G (2020) Liquid: A bitcoin sidechain .
- 10909 [456] Btc relay. Available at <https://github.com/ethereum/btcrelay>.
- 10910 [457] Peace relay. Available at <https://github.com/loiluu/peacerelay>.
- 10911 [458] Teutsch J, Straka M, Boneh D (2019) Retrofitting a two-way peg between
10912 blockchains. *arXiv preprint arXiv:190803999* .
- 10913 [459] Zamyatin A, Harz D, Lind J, Panayiotou P, Gervais A, Knottenbelt W (2019)
10914 Xclaim: Trustless, interoperable, cryptocurrency-backed assets. *2019 IEEE Sym-*
10915 *posium on Security and Privacy (SP)* (IEEE), pp 193–210.
- 10916 [460] Bünz B, Kiffer L, Luu L, Zamani M (2020) Flyclient: Super-light clients for cryp-
10917 tocurrencies. *2020 IEEE Symposium on Security and Privacy (SP)* (IEEE), pp 928–
10918 946.
- 10919 [461] Garoffolo A, Kaidalov D, Oliynykov R (2020) Zedoo: a zk-snark verifiable cross-
10920 chain transfer protocol enabling decoupled and decentralized sidechains. *arXiv*
10921 *preprint arXiv:200201847* .
- 10922 [462] Westerkamp M, Eberhardt J (2020) zkrelay: Facilitating sidechains using zksnark-
10923 based chain-relays. *Contract* 1(2):3.
- 10924 [463] Frauenthaler P, Sigwart M, Spanring C, Schulte S (2020) Testimonium: A cost-
10925 efficient blockchain relay. *arXiv preprint arXiv:200212837* .
- 10926 [464] Burdges J, Cevallos A, Czaban P, Habermeier R, Hosseini S, Lama F, Alper HK,
10927 Luo X, Shirazi F, Stewart A, et al. (2020) Overview of polkadot and its design con-
10928 siderations. *arXiv preprint arXiv:200513456* .
- 10929 [465] Goes C (2020) The interblockchain communication protocol: An overview. *arXiv*
10930 *preprint arXiv:200615918* .
- 10931 [466] Koh J (2019) 5 differences between cosmos & polkadot. Available at [https://medi-
10932 um.com/@juliankoh/5-differences-between-cosmos-polkadot-67f09535594b](https://medium.com/@juliankoh/5-differences-between-cosmos-polkadot-67f09535594b).
- 10933 [467] Wang G (2021) Sok: Exploring blockchains interoperability. *IACR Cryptol ePrint*
10934 *Arch* :537.
- 10935 [468] Kiayias A, Miller A, Zindros D (2017) Non-interactive proofs of proof-of-work.
10936 *IACR Cryptol ePrint Arch* 2017(963):1–42.
- 10937 [469] Kiayias A, Zindros D (2019) Proof-of-work sidechains. *International Conference on*
10938 *Financial Cryptography and Data Security* (Springer), pp 21–34.
- 10939 [470] Gaži P, Kiayias A, Zindros D (2019) Proof-of-stake sidechains. *2019 IEEE Sympo-*

- 10940 *sium on Security and Privacy (SP)* (IEEE), pp 139–156.
- 10941 [471] Neudecker T, Hartenstein H (2018) Network layer aspects of permissionless
10942 blockchains. *IEEE Communications Surveys & Tutorials* 21(1):838–857.
- 10943 [472] Deshpande V, Badis H, George L (2018) Btmap: Mapping bitcoin peer-to-peer
10944 network topology. *2018 IFIP/IEEE International Conference on Performance Eval-*
10945 *uation and Modeling in Wired and Wireless Networks (PEMWN)* (IEEE), pp 1–6.
- 10946 [473] Wang T, Zhao C, Yang Q, Zhang S (2020) Ethna: Analyzing the underlying peer-to-
10947 peer network of the ethereum blockchain. *arXiv preprint arXiv:201001373* .
- 10948 [474] Daniel E, Rohrer E, Tschorsch F (2019) Map-z: Exposing the zcash network in
10949 times of transition. *2019 IEEE 44th Conference on Local Computer Networks (LCN)*
10950 (IEEE), pp 84–92.
- 10951 [475] Cao T, Yu J, Decouchant J, Luo X, Verissimo P (2020) Exploring the monero peer-
10952 to-peer network. *Financial Cryptography and Data Security 2020, Sabah, 10-14*
10953 *February 2020* .
- 10954 [476] Delgado-Segura S, Bakshi S, Pérez-Solà C, Litton J, Pachulski A, Miller A, Bhat-
10955 tacharjee B (2019) Txprobe: Discovering bitcoin’s network topology using orphan
10956 transactions. *International Conference on Financial Cryptography and Data Secu-*
10957 *rity* (Springer), pp 550–566.
- 10958 [477] Miller A, Litton J, Pachulski A, Gupta N, Levin D, Spring N, Bhattacharjee B (2015)
10959 Discovering bitcoin’s public topology and influential nodes. Available at [http://ww](http://www.cs.umd.edu/projects/coinscope/coinscope.pdf)
10960 [w.cs.umd.edu/projects/coinscope/coinscope.pdf](http://www.cs.umd.edu/projects/coinscope/coinscope.pdf).
- 10961 [478] Koshy P (2013) Coinseer: A telescope into bitcoin .
- 10962 [479] Tramèr F, Boneh D, Paterson K (2020) Remote side-channel attacks on anony-
10963 mous transactions. *29th {USENIX} Security Symposium ({USENIX} Security 20)*,
10964 pp 2739–2756.
- 10965 [480] Heilman E, Kendler A, Zohar A, Goldberg S (2015) Eclipse attacks on bitcoin’s
10966 peer-to-peer network. *24th {USENIX} Security Symposium ({USENIX} Security*
10967 *15)*, pp 129–144.
- 10968 [481] Marcus Y, Heilman E, Goldberg S (2018) Low-resource eclipse attacks on
10969 ethereum’s peer-to-peer network. *IACR Cryptol ePrint Arch* :236.
- 10970 [482] Henningsen S, Teunis D, Florian M, Scheuermann B (2019) Eclipsing ethereum
10971 peers with false friends. *2019 IEEE European Symposium on Security and Privacy*
10972 *Workshops (EuroS&PW)* (IEEE), pp 300–309.
- 10973 [483] Apostolaki M, Zohar A, Vanbever L (2017) Hijacking bitcoin: Routing attacks on
10974 cryptocurrencies. *2017 IEEE Symposium on Security and Privacy (SP)* (IEEE), pp
10975 375–392.
- 10976 [484] Saad M, Cook V, Nguyen L, Thai MT, Mohaisen A (2019) Partitioning attacks on
10977 bitcoin: colliding space, time, and logic. *2019 IEEE 39th International Conference*
10978 *on Distributed Computing Systems (ICDCS)* (IEEE), pp 1175–1187.
- 10979 [485] Saad M, Anwar A, Ravi S, Mohaisen D (2021) Hashsplit: Exploiting bitcoin asyn-
10980 chrony to violate common prefix and chain quality. *IACR Cryptol ePrint Arch* :299.
- 10981 [486] Tran M, Choi I, Moon GJ, Vu AV, Kang MS (2020) A stealthier partitioning at-

- 10982 tack against bitcoin peer-to-peer network. *IEEE Symposium on Security and Privacy*
10983 (*S&P*).
- 10984 [487] Paphitis A, Kourtellis N, Sirivianos M (2021) A first look into the structural proper-
10985 ties and resilience of blockchain overlays. *arXiv preprint arXiv:210403044* .
- 10986 [488] Apostolaki M, Marti G, Müller J, Vanbever L (2018) Sabre: Protecting bitcoin
10987 against routing attacks. *arXiv preprint arXiv:180806254* .
- 10988 [489] Boverman A (2011) Timejacking & bitcoin. Available at [http://culubas.blogspot.c](http://culubas.blogspot.com/2011/05/timejacking-bitcoin_802.html)
10989 [om/2011/05/timejacking-bitcoin_802.html](http://culubas.blogspot.com/2011/05/timejacking-bitcoin_802.html).
- 10990 [490] Company EC (2020) New releases: 2.1.1 and hotfix 2.1.1-1. Available at [https:](https://electriccoin.co/blog/new-releases-2-1-1-and-hotfix-2-1-1-1/)
10991 [//electriccoin.co/blog/new-releases-2-1-1-and-hotfix-2-1-1-1/](https://electriccoin.co/blog/new-releases-2-1-1-and-hotfix-2-1-1-1/).
- 10992 [491] Imtiaz MA, Starobinski D, Trachtenberg A, Younis N (2019) Churn in the bit-
10993 coin network: Characterization and impact. *2019 IEEE International Conference*
10994 *on Blockchain and Cryptocurrency (ICBC)* (IEEE), pp 431–439.
- 10995 [492] Kiffer L, Salman A, Levin D, Mislove A, Nita-Rotaru C (2021) Under the hood of
10996 the ethereum gossip protocol. *Proceedings of the Financial Cryptography and Data*
10997 *Security (FC'21) St George's, Grenada* .
- 10998 [493] Motlagh SG, Mišić J, Mišić VB (2020) Impact of node churn in the bitcoin network.
10999 *IEEE Transactions on Network Science and Engineering* .
- 11000 [494] Biryukov A, Pustogarov I (2015) Bitcoin over tor isn't a good idea. *2015 IEEE*
11001 *Symposium on Security and Privacy* (IEEE), pp 122–134.
- 11002 [495] Blackshear S, Dill DL, Qadeer S, Barrett CW, Mitchell JC, Padon O, Zohar Y (2020)
11003 Resources: A safe language abstraction for money. *arXiv preprint arXiv:200405106*
11004 .
- 11005 [496] Schwartz C (2019) Ethereum 2.0: A complete guide. ewasm.. Available at [https:](https://medium.com/chainsafe-systems/ethereum-2-0-a-complete-guide-ewasm-394cac756baf)
11006 [//medium.com/chainsafe-systems/ethereum-2-0-a-complete-guide-ewasm-394cac7](https://medium.com/chainsafe-systems/ethereum-2-0-a-complete-guide-ewasm-394cac756baf)
11007 [56baf](https://medium.com/chainsafe-systems/ethereum-2-0-a-complete-guide-ewasm-394cac756baf).
- 11008 [497] Zheng S, Wang H, Wu L, Huang G, Liu X (2020) Vm matters: A comparison
11009 of wasm vms and evms in the performance of blockchain smart contracts. *arXiv*
11010 *preprint arXiv:201201032* .
- 11011 [498] Hess T, Keefer R, Siner EG Emin Siner (2016) Ethereum's dao wars soft fork is a
11012 potential dos vector. Available at [https://hackingdistributed.com/2016/06/28/ethere](https://hackingdistributed.com/2016/06/28/ethereum-soft-fork-dos-vector/)
11013 [um-soft-fork-dos-vector/](https://hackingdistributed.com/2016/06/28/ethereum-soft-fork-dos-vector/).
- 11014 [499] Jedusor TE (2016) Mumblewimble. Available at [https://scalingbitcoin.org/papers/mi](https://scalingbitcoin.org/papers/mimblewimble.txt)
11015 [mblewimble.txt](https://scalingbitcoin.org/papers/mimblewimble.txt).
- 11016 [500] Poelstra A (2016) Mumblewimble. Available at [https://download.wpsoftware.net/bi](https://download.wpsoftware.net/bitcoin/wizardry/mimblewimble.pdf)
11017 [tcoin/wizardry/mimblewimble.pdf](https://download.wpsoftware.net/bitcoin/wizardry/mimblewimble.pdf).
- 11018 [501] Fuchsbaauer G, Orrù M, Seurin Y (2019) Aggregate cash systems: A cryptographic
11019 investigation of mumblewimble. *Annual International Conference on the Theory and*
11020 *Applications of Cryptographic Techniques* (Springer), pp 657–689.
- 11021 [502] Reijnsbergen D, Dinh TTA (2020) On exploiting transaction concurrency to speed up
11022 blockchains. *arXiv preprint arXiv:200306128* .
- 11023 [503] Dickerson T, Gazzillo P, Herlihy M, Koskinen E (2019) Adding concurrency to smart

- 11024 contracts. *Distributed Computing* :1–17.
- 11025 [504] Saraph V, Herlihy M (2019) An empirical study of speculative concurrency in
11026 ethereum smart contracts. *arXiv preprint arXiv:190101376* .
- 11027 [505] Bartoletti M, Galletta L, Murgia M (2020) A theory of transaction parallelism in
11028 blockchains. *arXiv preprint arXiv:201113837* .
- 11029 [506] Anjana PS, Kumari S, Peri S, Rathor S, Somani A (2019) An efficient framework
11030 for optimistic concurrent execution of smart contracts. *2019 27th Euromicro Inter-*
11031 *national Conference on Parallel, Distributed and Network-Based Processing (PDP)*
11032 (IEEE), pp 83–92.
- 11033 [507] Anjana PS, Kumari S, Peri S, Rathor S, Somani A (2021) Optsmart: A space efficient
11034 optimistic concurrent execution of smart contracts. *arXiv preprint arXiv:210204875*
11035 .
- 11036 [508] Paimani K (2021) Sonicchain: A wait-free, pseudo-static approach toward concur-
11037 rency in blockchains. *arXiv preprint arXiv:210209073* .
- 11038 [509] Sheff I, Wang X, Myers AC, van Renesse R (2018) A web of blocks. *arXiv preprint*
11039 *arXiv:180606978* .
- 11040 [510] Yakovenko A (2020) Sealevel - parallel processing thousands of smart contracts.
11041 Available at [https://medium.com/solana-labs/sealevel-parallel-processing-thousan-](https://medium.com/solana-labs/sealevel-parallel-processing-thousands-of-smart-contracts-d814b378192)
11042 [ds-of-smart-contracts-d814b378192](https://medium.com/solana-labs/sealevel-parallel-processing-thousands-of-smart-contracts-d814b378192).
- 11043 [511] Sasson EB, Chiesa A, Garman C, Green M, Miers I, Tromer E, Virza M (2014)
11044 Zerocash: Decentralized anonymous payments from bitcoin. *2014 IEEE Symposium*
11045 *on Security and Privacy* (IEEE), pp 459–474.
- 11046 [512] Bünz B, Agrawal S, Zamani M, Boneh D (2020) Zether: Towards privacy in a smart
11047 contract world. *International Conference on Financial Cryptography and Data Se-*
11048 *curity* (Springer), pp 423–443.
- 11049 [513] Kosba A, Miller A, Shi E, Wen Z, Papamanthou C (2016) Hawk: The blockchain
11050 model of cryptography and privacy-preserving smart contracts. *2016 IEEE sympo-*
11051 *sium on security and privacy (SP)* (IEEE), pp 839–858.
- 11052 [514] Bowe S, Chiesa A, Green M, Miers I, Mishra P, Wu H (2020) Zexe: Enabling decen-
11053 tralized private computation. *2020 IEEE Symposium on Security and Privacy (SP)*
11054 (IEEE), pp 947–964.
- 11055 [515] Solomon R, Almashaqbeh G (2021) smartfhe: Privacy-preserving smart contracts
11056 from fully homomorphic encryption. *IACR Cryptol ePrint Arch* :133.
- 11057 [516] Teutsch J, Reitwießner C (2019) A scalable verification solution for blockchains.
11058 *arXiv preprint arXiv:190804756* .
- 11059 [517] Androutaki E, Barger A, Bortnikov V, Cachin C, Christidis K, De Caro A, Enyeart D,
11060 Ferris C, Laventman G, Manevich Y, et al. (2018) Hyperledger fabric: a distributed
11061 operating system for permissioned blockchains. *Proceedings of the thirteenth Eu-*
11062 *roSys conference*, pp 1–15.
- 11063 [518] Chacko JA, Mayer R, Jacobsen HA (2021) Why do my blockchain transactions fail?
11064 a study of hyperledger fabric (extended version). *arXiv preprint arXiv:210304681* .
- 11065 [519] Guggenberger T, Sedlmeir J, Fridgen G, Luckow A (2021) An in-depth in-

- 11066 investigation of performance characteristics of hyperledger fabric. *arXiv preprint*
11067 *arXiv:210207731* .
- 11068 [520] Kalodner H, Goldfeder S, Chen X, Weinberg SM, Felten EW (2018) Arbitrum: Scal-
11069 able, private smart contracts. *27th {USENIX} Security Symposium ({USENIX} Se-*
11070 *curity 18)*, pp 1353–1370.
- 11071 [521] Wüst K, Matetic S, Egli S, Kostianen K, Capkun S (2020) Ace: Asynchronous
11072 and concurrent execution of complex smart contracts. *Proceedings of the 2020 ACM*
11073 *SIGSAC Conference on Computer and Communications Security*, pp 587–600.
- 11074 [522] Wüst K, Diana L, Kostianen K, Karame G, Matetic S, Capkun S (2019) Bitcon-
11075 tracts: Adding expressive smart contracts to legacy cryptocurrencies. *IACR Cryptol*
11076 *ePrint Arch :857*.
- 11077 [523] Al-Bassam M (2019) Lazyledger: A distributed data availability ledger with client-
11078 side smart contracts. *arXiv preprint arXiv:190509274* .
- 11079 [524] McCorry P, Möser M, Shahandasti SF, Hao F (2016) Towards bitcoin payment net-
11080 works. *Australasian Conference on Information Security and Privacy* (Springer), pp
11081 57–76.
- 11082 [525] Decker C, Russell R, Osuntokun O (2018) eltoo: A simple layer2 protocol for bit-
11083 coin. *White paper: <https://blockstream.com/eltoo.pdf>* .
- 11084 [526] Poon J, Dryja T (2016) The bitcoin lightning network: Scalable off-chain instant
11085 payments. Available at <http://lightning.network/lightning-network-paper.pdf>.
- 11086 [527] Kiayias A, Litos OST (2020) A composable security treatment of the lightning net-
11087 work. *2020 IEEE 33rd Computer Security Foundations Symposium (CSF)* (IEEE),
11088 pp 334–349.
- 11089 [528] Gudgeon L, Moreno-Sanchez P, Roos S, McCorry P, Gervais A (2020) Sok: Layer-
11090 two blockchain protocols. *Financial Cryptography and Data Security: 24th Inter-*
11091 *national Conference, FC 2020, Kota Kinabalu, Malaysia, February 10–14, 2020*
11092 *Revised Selected Papers 24* (Springer), pp 201–226.
- 11093 [529] Aumayr L, Abbaszadeh K, Maffei M (2022) Thora: Atomic and privacy-preserving
11094 multi-channel updates. *IACR Cryptol ePrint Arch 2022:317*.
- 11095 [530] Malavolta G, Moreno-Sanchez P, Schneidewind C, Kate A, Maffei M (2019) Anony-
11096 mous multi-hop locks for blockchain scalability and interoperability. *26th Annual*
11097 *Network and Distributed System Security Symposium, NDSS 2019*.
- 11098 [531] Aumayr L, Ersoy O, Erwig A, Faust S, Hostakova K, Maffei M, Moreno-Sanchez P,
11099 Riahi S (2020) Generalized bitcoin-compatible channels. *IACR Cryptol ePrint Arch*
11100 *:476*.
- 11101 [532] Dziembowski S, Eckey L, Faust S, Malinowski D (2019) Perun: Virtual payment
11102 hubs over cryptocurrencies. *2019 IEEE Symposium on Security and Privacy (SP)*
11103 (IEEE), pp 106–123.
- 11104 [533] Aumayr L, Ersoy O, Erwig A, Faust S, Hostáková K, Maffei M, Moreno-Sanchez P,
11105 Riahi S (2020) Bitcoin-compatible virtual channels. *IACR Cryptol ePrint Arch :554*.
- 11106 [534] Jourenko M, Larangeira M, Tanaka K (2020) Lightweight virtual payment channels.
11107 *International Conference on Cryptology and Network Security* (Springer), pp 365–

- 11108 384.
- 11109 [535] Dziembowski S, Faust S, Hostáková K (2018) General state channel networks. *Pro-*
11110 *ceedings of the 2018 ACM SIGSAC Conference on Computer and Communications*
11111 *Security*, pp 949–966.
- 11112 [536] Coleman J, Horne L, Xuanji L (2018) Counterfactual: Generalized state channels.
11113 Available at <https://14.ventures/papers/statechannels.pdf>.
- 11114 [537] Dziembowski S, Eckey L, Faust S, Hesse J, Hostáková K (2019) Multi-party virtual
11115 state channels. *Annual International Conference on the Theory and Applications of*
11116 *Cryptographic Techniques* (Springer), pp 625–656.
- 11117 [538] Miller A, Bentov I, Bakshi S, Kumaresan R, McCorry P (2019) Sprites and state
11118 channels: Payment networks that go faster than lightning. *International Conference*
11119 *on Financial Cryptography and Data Security* (Springer), pp 508–526.
- 11120 [539] Avarikioti G, Laufenberg F, Sliwinski J, Wang Y, Wattenhofer R (2018) Towards
11121 secure and efficient payment channels. *arXiv preprint arXiv:181112740* .
- 11122 [540] Khabbazian M, Nadahalli T, Wattenhofer R (2019) Outpost: A responsive
11123 lightweight watchtower. *Proceedings of the 1st ACM Conference on Advances in*
11124 *Financial Technologies*, pp 31–40.
- 11125 [541] McCorry P, Bakshi S, Bentov I, Meiklejohn S, Miller A (2019) Pisa: Arbitration
11126 outsourcing for state channels. *Proceedings of the 1st ACM Conference on Advances*
11127 *in Financial Technologies*, pp 16–30.
- 11128 [542] Avarikioti Z, Litos OST, Wattenhofer R (2020) Cerberus channels: Incentivizing
11129 watchtowers for bitcoin. *International Conference on Financial Cryptography and*
11130 *Data Security* (Springer), pp 346–366.
- 11131 [543] Mirzaei A, Sakzad A, Yu J, Steinfeld R (2021) Fppw: A fair and privacy preserving
11132 watchtower for bitcoin. *IACR Cryptol ePrint Arch* :117.
- 11133 [544] Avarikioti G, Kogias EK, Wattenhofer R (2019) Brick: Asynchronous state channels.
11134 *arXiv preprint arXiv:190511360* .
- 11135 [545] Poon J, Buterin V (2017) Plasma: Scalable autonomous smart contracts. *White paper*
11136 :1–47.
- 11137 [546] Buterin V (2018) Minimum viable plasma. Available at <https://ethresear.ch/t/minimum-viable-plasma/426>.
- 11138 [547] Jones B, Fichter K (2018) More viable plasma. Available at <https://ethresear.ch/t/more-viable-plasma/2160>.
- 11139 [548] Konstantopoulos G (2019) Plasma cash: Towards more efficient plasma construc-
11140 tions. *arXiv preprint arXiv:191112095* .
- 11141 [549] Dziembowski S, Fabianski G, Faust S, Riahi S (2020) Lower bounds for off-chain
11142 protocols: Exploring the limits of plasma. *IACR Cryptol ePrint Arch* 2020:175.
- 11143 [550] Buterin V (2021) An incomplete guide to rollups. Available at [https://vitalik.ca/gen-](https://vitalik.ca/general/2021/01/05/rollup.html)
11144 [eral/2021/01/05/rollup.html](https://vitalik.ca/general/2021/01/05/rollup.html).
- 11145 [551] Leshno J, Strack P (2019) Bitcoin: An impossibility theorem for proof-of-work
11146 based protocols. Available at *SSRN: https://ssrn.com/abstract=3487355* .
- 11147 [552] Chen X, Papadimitriou C, Roughgarden T (2019) An axiomatic approach to block
11148
11149

- 11150 rewards. *Proceedings of the 1st ACM Conference on Advances in Financial Tech-*
11151 *nologies*, pp 124–131.
- 11152 [553] Siddiqui S, Vanahalli G, Gujar S (2020) Bitcoin: Achieving fairness for bitcoin
11153 in transaction fee only model. *Proceedings of the 19th International Conference on*
11154 *Autonomous Agents and MultiAgent Systems*, pp 2008–2010.
- 11155 [554] Babaioff M, Dobzinski S, Oren S, Zohar A (2012) On bitcoin and red balloons.
11156 *Proceedings of the 13th ACM conference on electronic commerce*, pp 56–73.
- 11157 [555] Houy N (2014) The economics of bitcoin transaction fees. *GATE WP 1407*.
- 11158 [556] Carlsten M, Kalodner H, Weinberg SM, Narayanan A (2016) On the instability of
11159 bitcoin without the block reward. *Proceedings of the 2016 ACM SIGSAC Conference*
11160 *on Computer and Communications Security*, pp 154–167.
- 11161 [557] Tsabary I, Eyal I (2018) The gap game. *Proceedings of the 2018 ACM SIGSAC*
11162 *conference on Computer and Communications Security*, pp 713–728.
- 11163 [558] Gong T, Minaei M, Sun W, Kate A (2020) Undercutting bitcoin is not profitable.
11164 *arXiv preprint arXiv:200711480* .
- 11165 [559] Daian P, Goldfeder S, Kell T, Li Y, Zhao X, Bentov I, Breidenbach L, Juels A (2019)
11166 Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in
11167 decentralized exchanges. *arXiv preprint arXiv:190405234* .
- 11168 [560] Eskandari S, Moosavi S, Clark J (2019) Sok: Transparent dishonesty: front-running
11169 attacks on blockchain. *International Conference on Financial Cryptography and*
11170 *Data Security* (Springer), pp 170–189.
- 11171 [561] Torres CF, Camino R, State R (2021) Frontrunner jones and the raiders of the
11172 dark forest: An empirical study of frontrunning on the ethereum blockchain. *arXiv*
11173 *preprint arXiv:210203347* .
- 11174 [562] Qin K, Zhou L, Gervais A (2021) Quantifying blockchain extractable value: How
11175 dark is the forest? *arXiv preprint arXiv:210105511* .
- 11176 [563] Zhou L, Qin K, Cully A, Livshits B, Gervais A (2021) On the just-in-time discovery
11177 of profit-generating transactions in defi protocols. *arXiv preprint arXiv:210302228* .
- 11178 [564] Kelkar M, Deb S, Kannan S (2021) Order-fair consensus in the permissionless set-
11179 ting. *IACR Cryptol ePrint Arch* :139.
- 11180 [565] Doweck Y, Eyal I (2020) Multi-party timed commitments. *arXiv preprint*
11181 *arXiv:200504883* .
- 11182 [566] Yang R, Murray T, Rimba P, Parampalli U (2019) Empirically analyzing ethereum’s
11183 gas mechanism. *2019 IEEE European Symposium on Security and Privacy Work-*
11184 *shops (EuroS&PW)* (IEEE), pp 310–319.
- 11185 [567] Perez D, Livshits B (2019) Broken metre: Attacking resource metering in evm. *arXiv*
11186 *preprint arXiv:190907220* .
- 11187 [568] Baird K, Jeong S, Kim Y, Burgstaller B, Scholz B (2019) The economics of smart
11188 contracts. *arXiv preprint arXiv:191011143* .
- 11189 [569] Buterin V (2016) Eip 150. Available at <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-150.md>.
- 11190
- 11191 [570] Holst Swende M, Szilagyi P (2021) Dodging a bullet: Ethereum state problems.

- 11192 Accessed on May 24, 2021. Available at [https://blog.ethereum.org/2021/05/18/eth_11194 \[571\] Lerner SD \(2013\) New bitcoin vulnerability: A transaction that takes at least 3 min- 11195 utes to be verified by a peer. Available at <https://bitcointalk.org/?topic=140078>. 11196 \[572\] maaku \(2015\) alternatives-to-block-size-as-aggregate-resource-limits. Available at 11197 \[https://diyhpl.us/wiki/transcripts/scalingbitcoin/alternatives-to-block-size-as-aggre- 11199 \\[573\\] Luu L, Teutsch J, Kulkarni R, Saxena P \\(2015\\) Demystifying incentives in the con- 11200 sensus computer. *Proceedings of the 22nd ACM SIGSAC Conference on Computer 11201 and Communications Security*, pp 706–719. 11202 \\[574\\] Pontiveros BBF, Torres CF, et al. \\(2019\\) Sluggish mining: Profiting from the veri- 11203 fier’s dilemma. *International Conference on Financial Cryptography and Data Se- 11204 curity* \\(Springer\\), pp 67–81. 11205 \\[575\\] Alharby M, Lunardi RC, Aldweesh A, van Moorsel A \\(2020\\) Data-driven model- 11206 based analysis of the ethereum verifier’s dilemma. *arXiv preprint arXiv:200412768* 11207 . 11208 \\[576\\] Das S, Awathare N, Ren L, Ribeiro VJ, Bellur U \\(2020\\) Better late than 11209 never; scaling computations in blockchain by delaying execution. *arXiv preprint* 11210 *arXiv:200511791* . 11211 \\[577\\] Amoussou-Guenou Y, Biais B, Potop-Butucaru M, Tucci-Piergiorganni S 11212 \\(2019\\) Rationals vs byzantines in consensus-based blockchains. *arXiv preprint* 11213 *arXiv:190207895* . 11214 \\[578\\] Amoussou-Guenou Y, Biais B, Potop-Butucaru M, Tucci-Piergiorganni S \\(2020\\) *Ra- 11215 tional Behavior in Committee-Based Blockchains*. Ph.D. thesis. CEA List; LIP6, 11216 Sorbonne Université, CNRS, UMR 7606; HEC Paris, . 11217 \\[579\\] Lavi R, Sattath O, Zohar A \\(2019\\) Redesigning bitcoin’s fee market. *The World Wide 11218 Web Conference*, pp 2950–2956. 11219 \\[580\\] Basu S, Easley D, O’Hara M, Siler E \\(2019\\) Towards a functional fee market for 11220 cryptocurrencies. Available at SSRN 3318327 . 11221 \\[581\\] Buterin V \\(2018\\) Blockchain resource pricing. Available at \\[https://ethresear.ch/t/dra- 11222 ft-position-paper-on-resource-pricing/2838\\]\\(https://ethresear.ch/t/draft-position-paper-on-resource-pricing/2838\\). 11223 \\[582\\] Buterin V, Connor E, Dudley R, Slipper M, Norden I \\(2019\\) Fee market change for 11224 eth 1.0 chain. Available at \\[https://github.com/ethereum/EIPs/blob/master/EIPS/ei- 11225 p-1559.md\\]\\(https://github.com/ethereum/EIPs/blob/master/EIPS/eip-1559.md\\). 11226 \\[583\\] Roughgarden T \\(2020\\) Transaction fee mechanism design for the ethereum 11227 blockchain: An economic analysis of eip-1559. *arXiv preprint arXiv:201200854* 11228 . 11229 \\[584\\] Leonardos S, Monnot B, Reijsbergen D, Skoulakis S, Piliouras G \\(2021\\) Dynamical 11230 analysis of the eip-1559 ethereum fee market. *arXiv preprint arXiv:210210567* . 11231 \\[585\\] Ferreira MV, Moroz DJ, Parkes DC, Stern M \\(2021\\) Dynamic posted-price mech- 11232 anisms for the blockchain transaction-fee market. *arXiv preprint arXiv:210314144* 11233 .\]\(https://diyhpl.us/wiki/transcripts/scalingbitcoin/alternatives-to-block-size-as-aggre- 11198 gate-resource-limits/\)](https://blog.ethereum.org/2021/05/18/eth_11193 _state_problems/)