Consensus Protocols III

Murat Osmanoglu

Voting-Based Protocols - Quorum

- Quorum, adaptation of Ethereum to permissioned blockchain, developed by JP Morgan as an enterprise platform
- two consensus protocols commonly used in Quorum: a variant of Raft and IBFT

Voting-Based Protocols - Quorum

- Quorum, adaptation of Ethereum to permissioned blockchain, developed by JP Morgan as an enterprise platform
- two consensus protocols commonly used in Quorum: a variant of Raft and IBFT

Raft for Blockchain

- leader combines the transactions it receives into a new block and sends it to other nodes
- others check the block, and if it is valid, send a message indicating they agree on the block
- if majority of the nodes send such message, leader considers that block to be committed and sends a message to others to inform them
- leader and other nodes append the new block to their chain
- different than the original Raft, leaders can remove offline followers or candidates from the committee

Voting-Based Protocols - Quorum

Istanbul BFT

- introduced by Moniz [16] in 2020

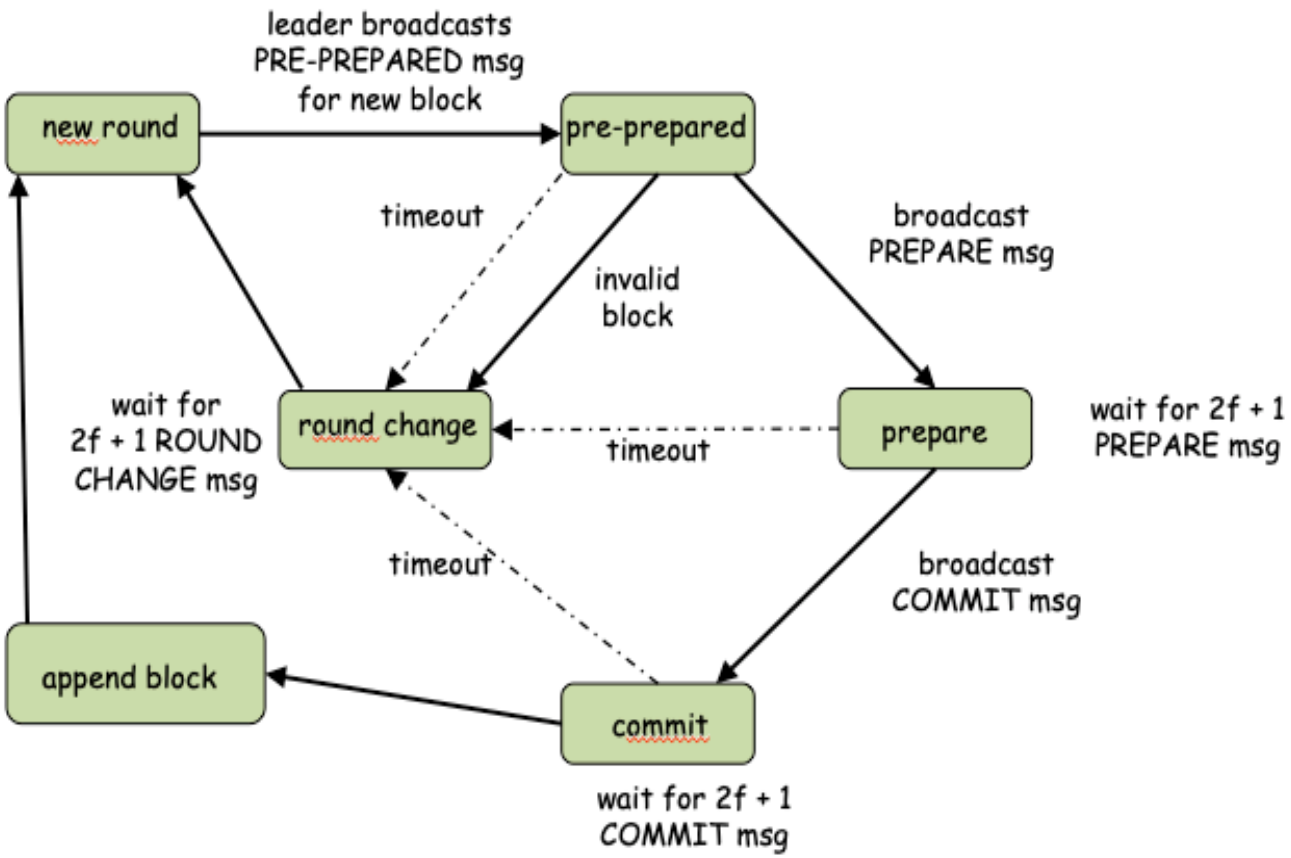

Voting-Based Protocols - Quorum

Istanbul BFT

- introduced by Moniz [16] in 2020

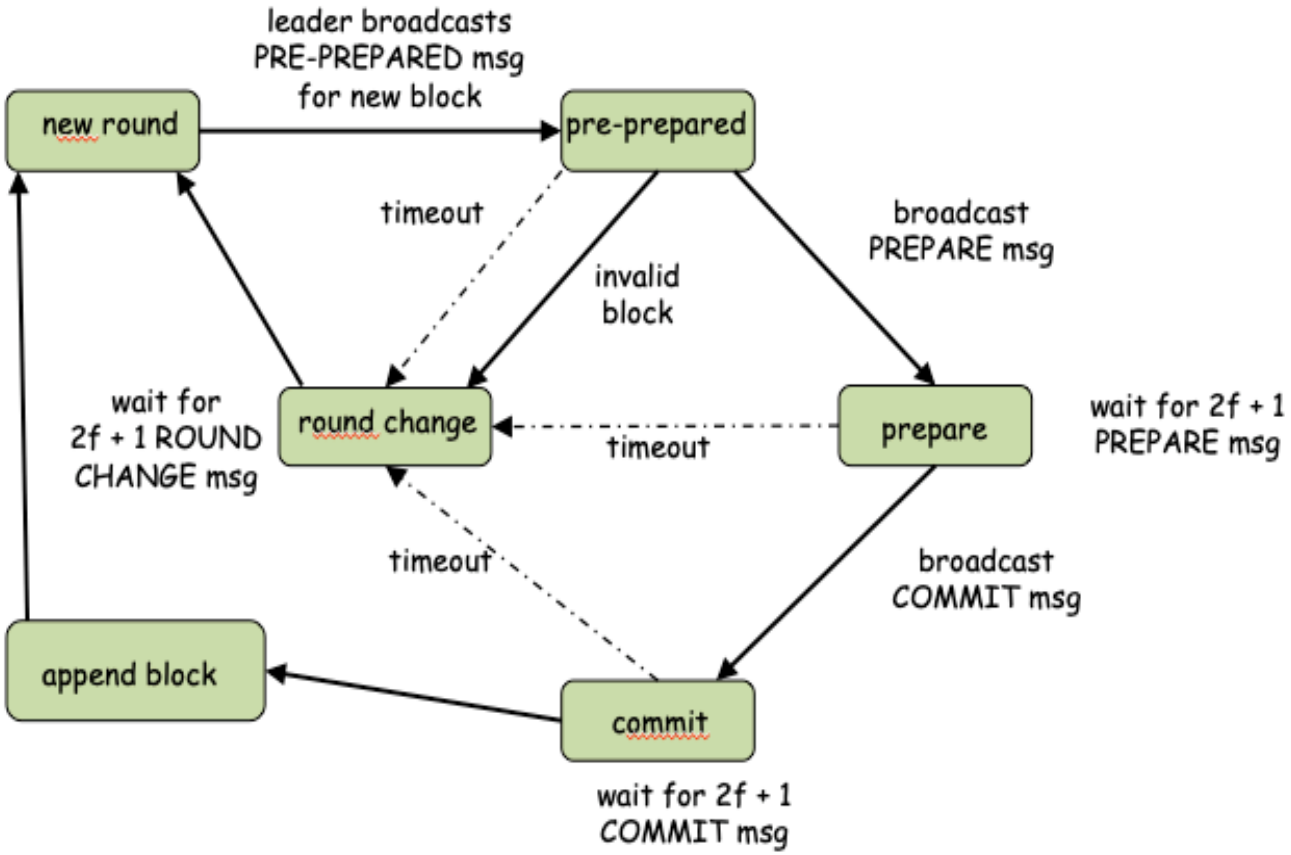

- validators initiate round change when round change time expires, or invalid prepare message received

Voting-Based Protocols - Hyperledger

- similar to Quorum, Hyperledger is an open source enterprise blockchain platform initiated by Linux Foundation and supported by IBM, Intel vs.
- Hyperledger includes different frameworks employing different consensus protocols (Raft for Fabric, Tendermint for Burrow, RBFT for Indy, vs.)

RBFT

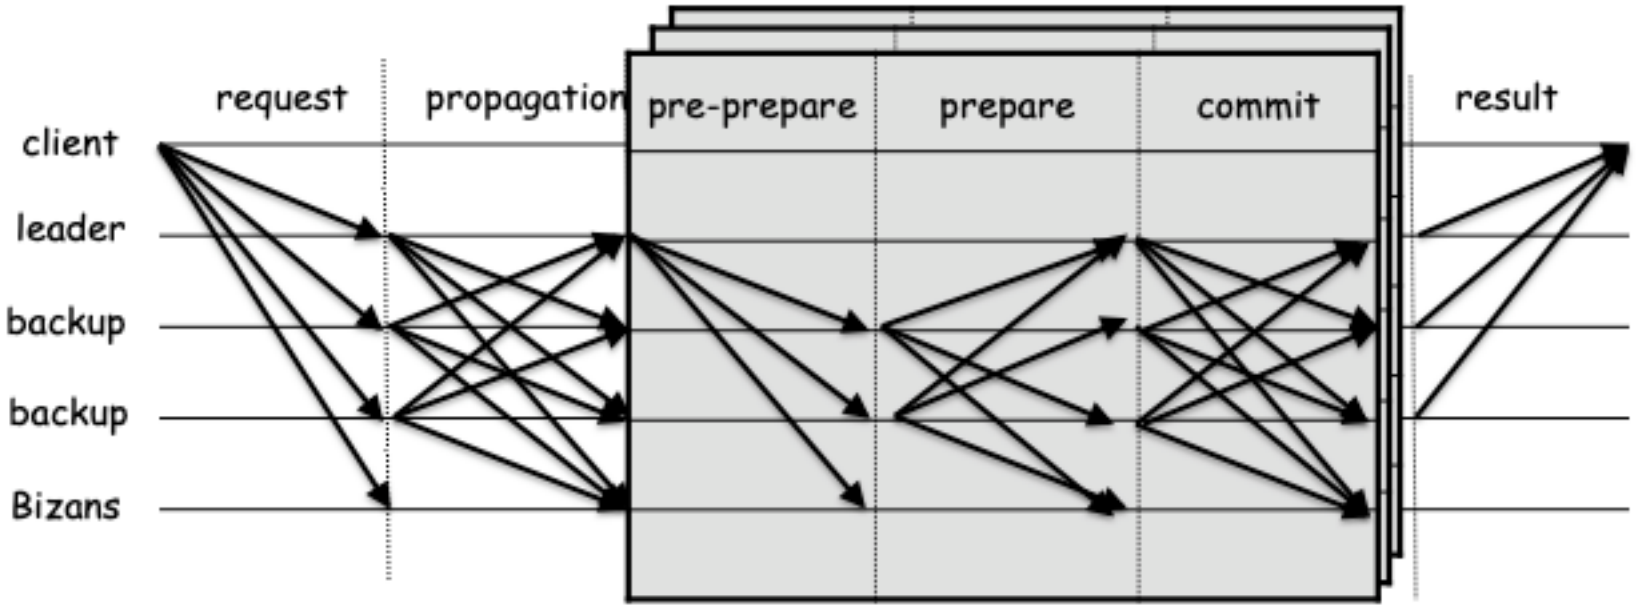

Voting-Based Protocols

protocol	leader selection	fault tolerance	fault type	delay	throughput
Raft	PoW	minority of nodes	crash	1.5 sn	750 tps with 3 nodes
IBFT	PoW	less than 1/3 of the nodes	Byzantine	5 sn	600 tps with 20 tps
RBFT	PoS	less than 1/3 of the nodes	Byzantine	?	10 tps with 50 nodes

Hybrid Consensus Protocols

- hybrid consensus protocols take the best of both worlds :
  - leaders chosen through a lottery-based election (establishing trust in the wild)
  - blocks approved by a committee of nodes before being appended to the chain (providing deterministic block finalization)

Hybrid Consensus Protocols - PoW + BFT

Peercensus

- introduced by Decker et al. [17] in 2016
- blocks creators chosen through Proof-of-Work algorithm
- blocks approved by a committe of nodes before being appended to the chain (a variant of PBFT – chain of agreement)

Hybrid Consensus Protocols - PoW + BFT

Peercensus

- introduced by Decker et al. [17] in 2016
- blocks creators chosen through Proof-of-Work algorithm
- blocks approved by a committe of nodes before being appended to the chain (a variant of PBFT – chain of agreement)
- how the members of the committee chosen?

Hybrid Consensus Protocols - PoW + BFT

Peercensus

- introduced by Decker et al. [17] in 2016
- blocks creators chosen through Proof-of-Work algorithm
- blocks approved by a committe of nodes before being appended to the chain (a variant of PBFT – chain of agreement)
- how the members of the committee chosen?
  - block creators join to the committee
  - the one created the last block will be the leader in the next view

Hybrid Consensus Protocols - PoW + BFT

Peercensus

- block creator sends the new block to primary
- primary validates the block, assigns it the current timestamp, and initiates 3-phase PBFT (pre-prepare, prepare, commit) by sending the new block to the members of the committee
- at the end of commit phase, each member appends the block to its chain

Hybrid Consensus Protocols - PoW + BFT

Peercensus

- block creator sends the new block to primary
- primary validates the block, assigns it the current timestamp, and initiates 3-phase PBFT (pre-prepare, prepare, commit) by sending the new block to the members of the committee
- at the end of commit phase, each member appends the block to its chain
- all members send a ping message to each other to check whether they are online
  - if some member offline, they initiate a leave operation to get this member out of the committee (liveness)

Hybrid Consensus Protocols - PoW + BFT

Peercensus

- if more than 2/3 of the committee honest, peercensus provides liveness and safety
- message complexity O(K) where K is the size of the committee (large K causing scalability problem)
- rewards distributed to the committee instead of just block creators

Hybrid Consensus Protocols - PoS + BFT

Tendermint

- introduced by Kwon [18] in 2014
- leaders who create the blocks determined in a round-rabin fashion from the committee with the frequency in proportion to their deposit
- Tendermint protocol used to finalize blocks by executing a variant of PBFT among the members of the committee

Hybrid Consensus Protocols - PoS + BFT

Tendermint

- introduced by Kwon [18] in 2014
- leaders who create the blocks determined in a round-rabin fashion from the committee with the frequency in proportion to their deposit
- Tendermint protocol used to finalize blocks by executing a variant of PBFT among the members of the committee
- how the members of the committee chosen?
  - nodes deposit some money to join the committee
  - they can leave the committee, but they need to wait for some blocks after their withdrawal message included in the chain

Hybrid Consensus Protocols - PoS + BFT

Tendemint

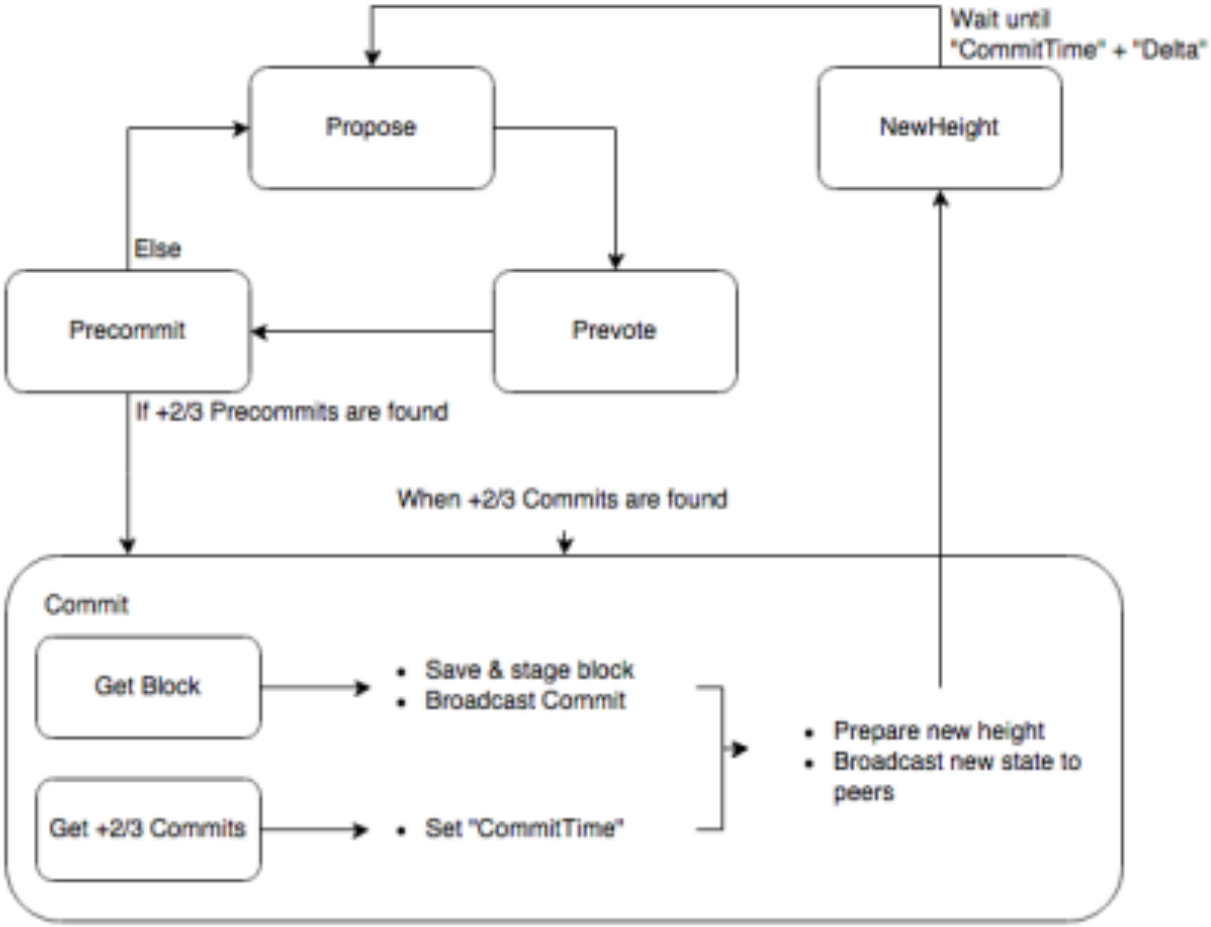

Hybrid Consensus Protocols - PoS + BFT

Tendermint

- when a validator signs two different blocks with same height, its deposit will be destroyed
- fees distributed among validators
- Tendermint has NewHeight before the next round that enables the commits of slower validators to be included in the blockchain
- if there are less than 1/3 Byzantine voting power, Tendemint provides safety and liveness

Hybrid Consensus Protocols - PoS + BFT

Casper

- introduced by Buterin and Griffith [19] in 2019
- leaders who create the blocks elected similar to traditional BFT
- Casper protocol used to finalize blocks by selecting a unique chain in every k blocks

Hybrid Consensus Protocols - PoS + BFT

Casper

- introduced by Buterin and Griffith [19] in 2019
- leaders who create the blocks elected similar to traditional BFT
- Casper protocol used to finalize blocks by selecting a unique chain in every k blocks
- how the members of the committee chosen?
  - nodes deposit some money to join the committee
  - they can leave the committee, but they need to wait for some blocks after their withdrawal message included in the chain
  - if they leave, they cannot rejoin to the committee

Hybrid Consensus Protocols - PoS + BFT

Casper

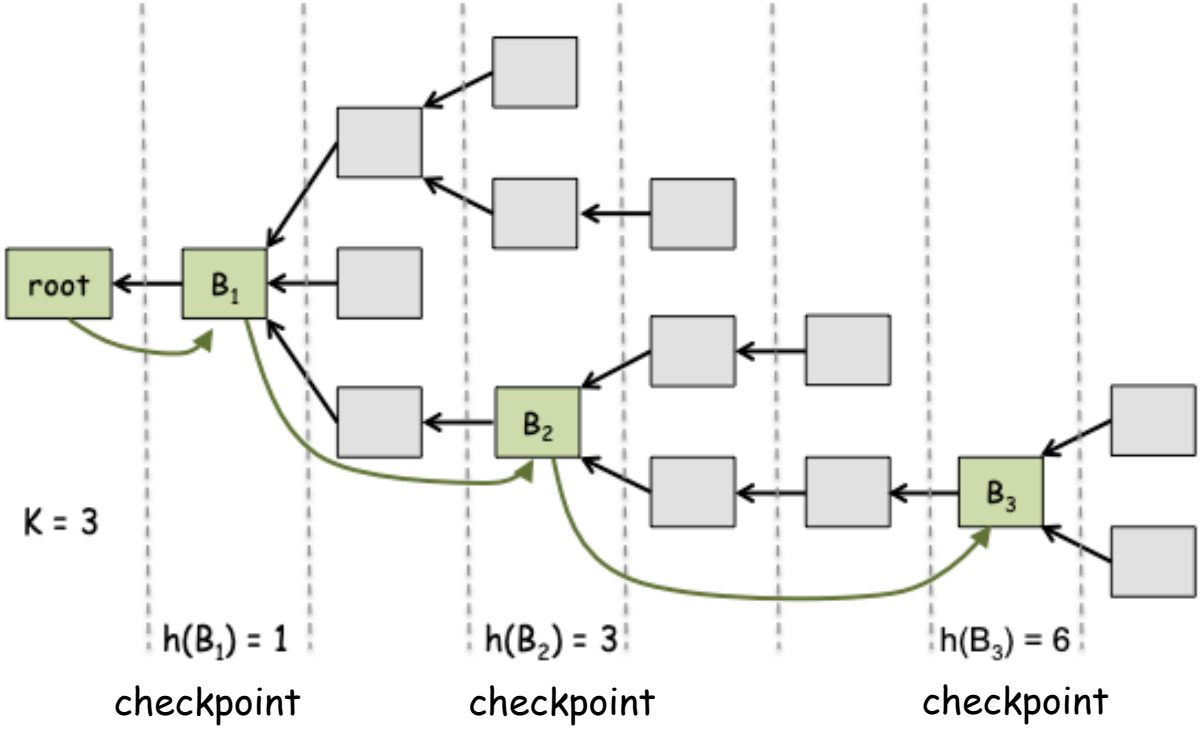

Hybrid Consensus Protocols - PoS + BFT

Casper

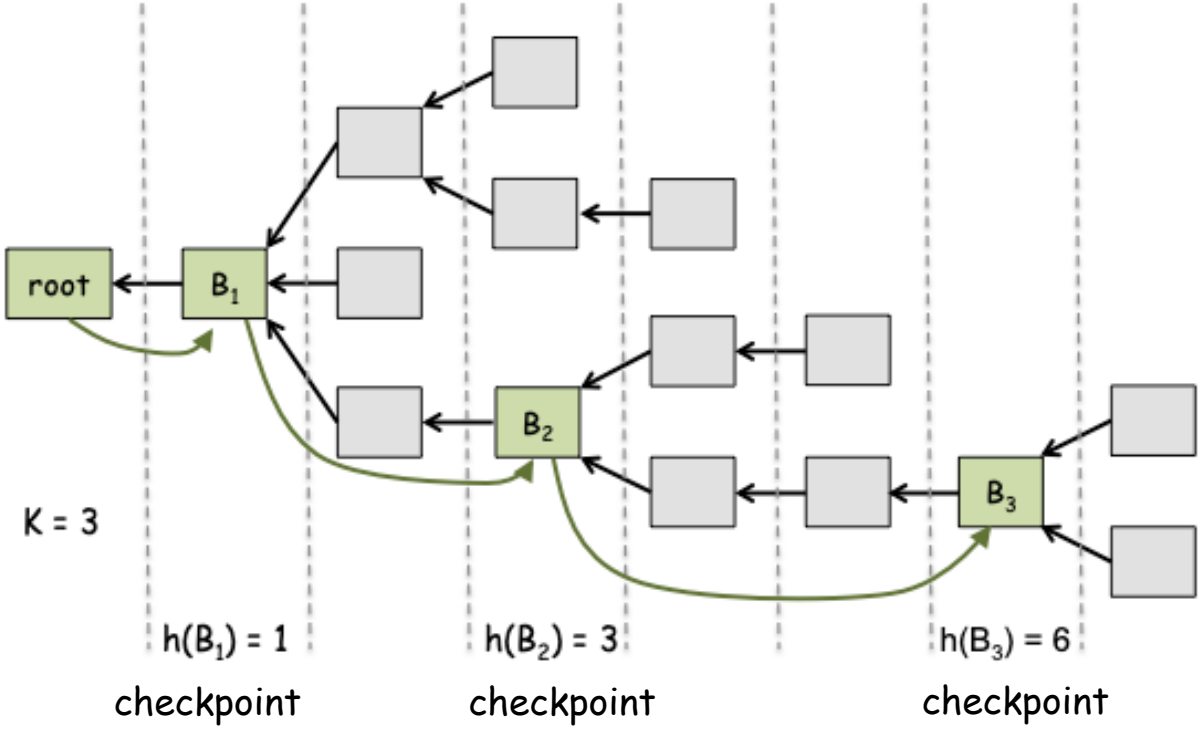

- at every checkpoint, validator v shares his vote as [v, s, t, h(s), h(t)] where s is the hash of the approved checkpoint, t is the hash of the checkpoint that is a descendant of s, h(.) is the height of a checkpoint

Hybrid Consensus Protocols - PoS + BFT

Casper

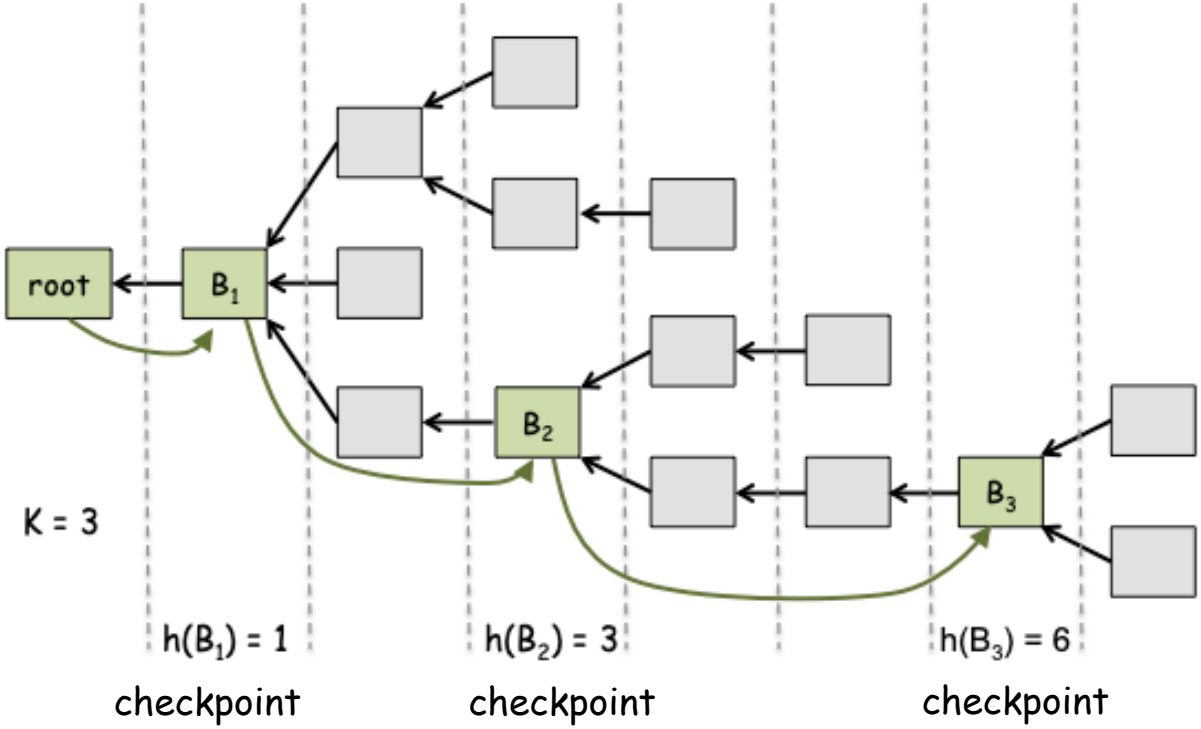

- at every checkpoint, validator v shares his vote as [v, s, t, h(s), h(t)] where s is the hash of the approved checkpoint, t is the hash of the checkpoint that is a descendant of s, h(.) is the height of a checkpoint
- a checkpoint s is called approved if there is a supermajority link s' ➔ s where s' is approved
- a supermajority link is a pair of checkpoints (s, t) such that at least 2/3 of validators (by deposit) have shared votes

Hybrid Consensus Protocols - PoS + BFT

Casper

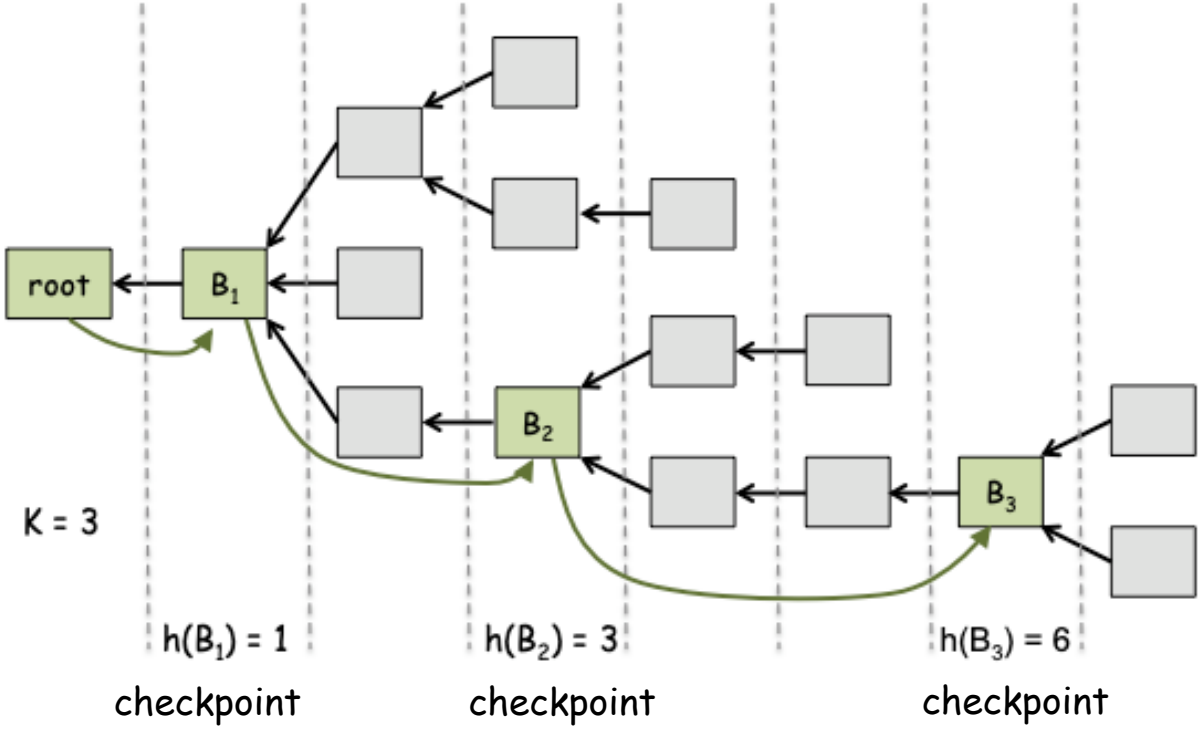

- at every checkpoint, validator v shares his vote as [v, s, t, h(s), h(t)]
- if s is not ancestor of t in the tree, the vote is not valid
- if public key of validator v not in validator set, the vote is not valid

Hybrid Consensus Protocols - PoS + BFT

Casper

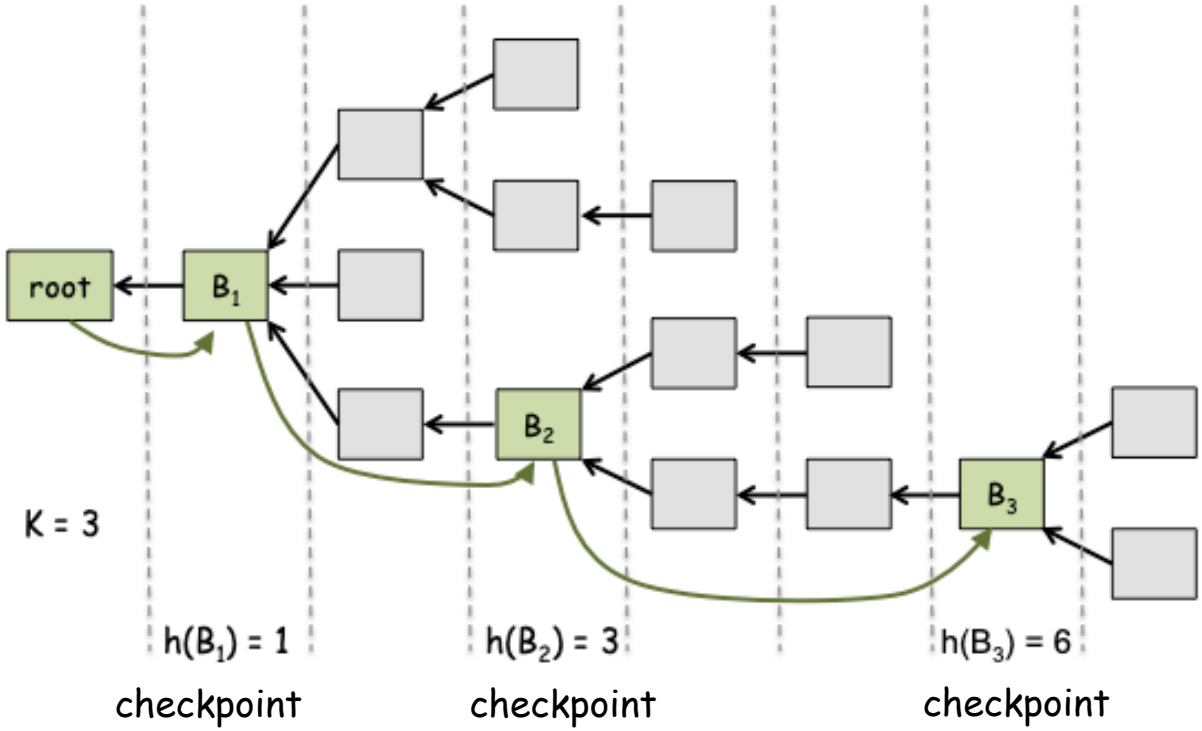

- at every checkpoint, validator v shares his vote as [v, s, t, h(s), h(t)]
- if v shares two different votes [v, s_1, t_1, $h(s_1)$, $h(t_1)$] and [v, s_2, t_2, $h(s_2)$, $h(t_2)$] s.t.
  - $h(t_1) = h(t_2)$
  - $h(s_1) < h(s_2) < h(t_2) < h(t_1)$

  its deposit slashed

Hybrid Consensus Protocols - PoS + BFT

Casper

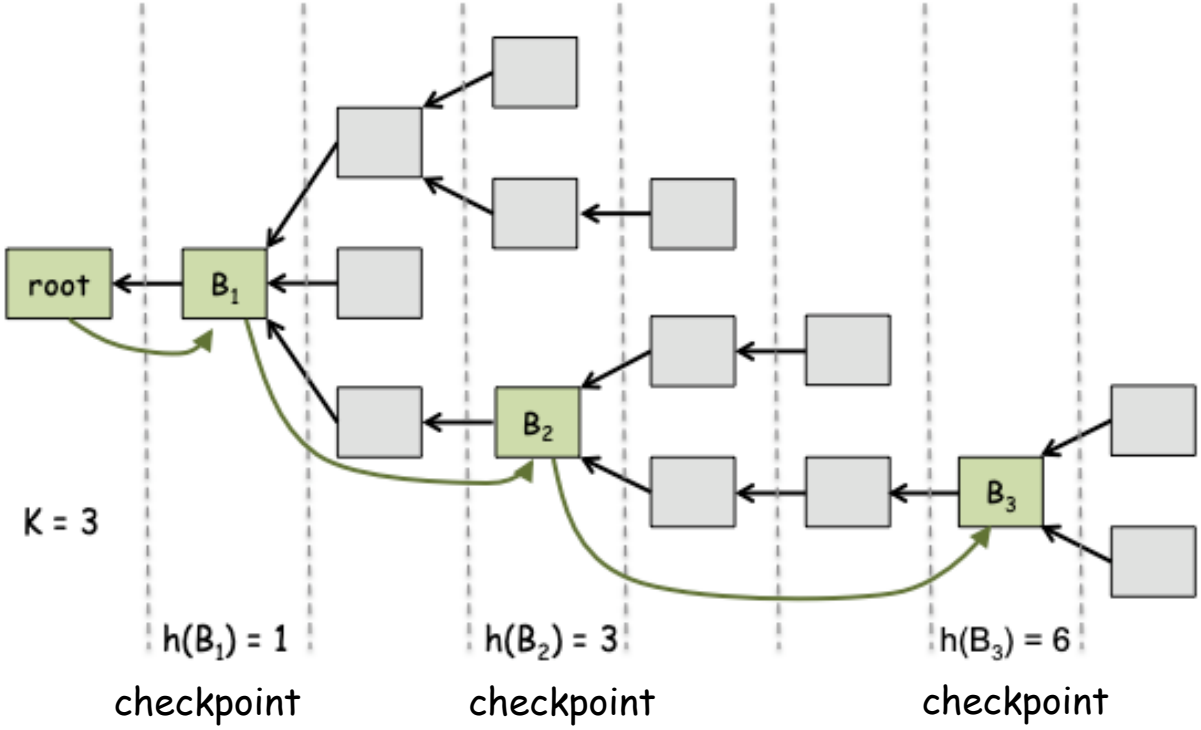

- the protocol provides safety (two conflicting checkpoints not finalized) and liveness (supermajority links always added to get new finalized checkpoints) if the validators holding more than 2/3 of voting power follow the protocol

Hybrid Consensus Protocols - PoS + BFT

EOSIO

- protocol enables players to delegate their stake to others
- at the beginning of each round, 21 nodes chosen depending on the stake delegated to them in order to form the committee of that round
- they assigned to time slots of 6 sec, and produce blocks from the transactions shared in that particular time slot
  (each round takes 126 sec)
- blocks produced at every 0.5 second
  (a member of the committee can produce at most 12 blocks)

Hybrid Consensus Protocols - PoS + BFT

EOSIO

- protocol enables players to delegate their stake to others
- at the beginning of each round, 21 nodes chosen depending on the stake delegated to them in order to form the committee of that round
- they assigned to time slots of 6 sec, and produce blocks from the transactions shared in that particular time slot
  (each round takes 126 sec)
- blocks produced at every 0.5 second
  (a member of the committee can produce at most 12 blocks)
- after producing blocks, block producers execute a BFT-type protocol to validate and append blocks to the chain
  (if 15 members sign a block, it can be considered as valid)

Hybrid Consensus Protocols - PoS + BFT

EOSIO

- each tx includes the hash of the last block added to the chain
  - to prevent tx to be added to alternative chain
  - to inform network about which chain holding the stake of a particular player

Hybrid Consensus Protocols - PoS + BFT

EOSIO

- each tx includes the hash of the last block added to the chain
  - to prevent tx to be added to alternative chain
  - to inform network about which chain holding the stake of a particular player
- there is no fee in EOSIO

Hybrid Consensus Protocols - PoS + BFT

EOSIO

- each tx includes the hash of the last block added to the chain
  - to prevent tx to be added to alternative chain
  - to inform network about which chain holding the stake of a particular player
- there is no fee in EOSIO
- block producers rewarded with newly minted tokens (total annual increase in token supply not exceeding 5%)

Hybrid Consensus Protocols - PoS + BFT

- for Tendermint and Casper, if the size of the committee is too big, it will create scalability problem

  (in that case, BFT protocols will generate too many messages which will be more than the network handles)

- EOSIO having only 21 nodes to execute BFT protocol (lacks of security analysis)

Hybrid Consensus Protocols

protocol	leader selection	committee formation	message complexity	fault tolerance	throughput	reward
Peercensus	PoW	PoW	large	less than 1/3 of the committee	?	fresh coin + fee
Tendermint	round-rabin*	PoS	large	less than 1/3 of total deposit	350 tps for 16 nodes	fee
Casper	PoS	PoS	large	less than 1/3 of total deposit	?	fee
EOSIO	round-rabin	DPoS	small	less than 1/3 of the committee	9656 tps	fresh coin

Consensus Protocols for Blockchain

lottery-based protocol	voting-based protocols
block finalization is probabilistic	block finalization is deterministic (all the nodes contribute the block validation process)
less messages (scales well)	message load (scales poorly)
mostly preferred in permissionless setting	mostly preferred in permissioned setting
focusing on leader election to establish trust in the wild	focusing on block voting
can tolerate Byzantine faults controlling minority of total hashing power or coin etc	can tolerate Byzantine faults less than n/3

Consensus Protocols for Blockchain

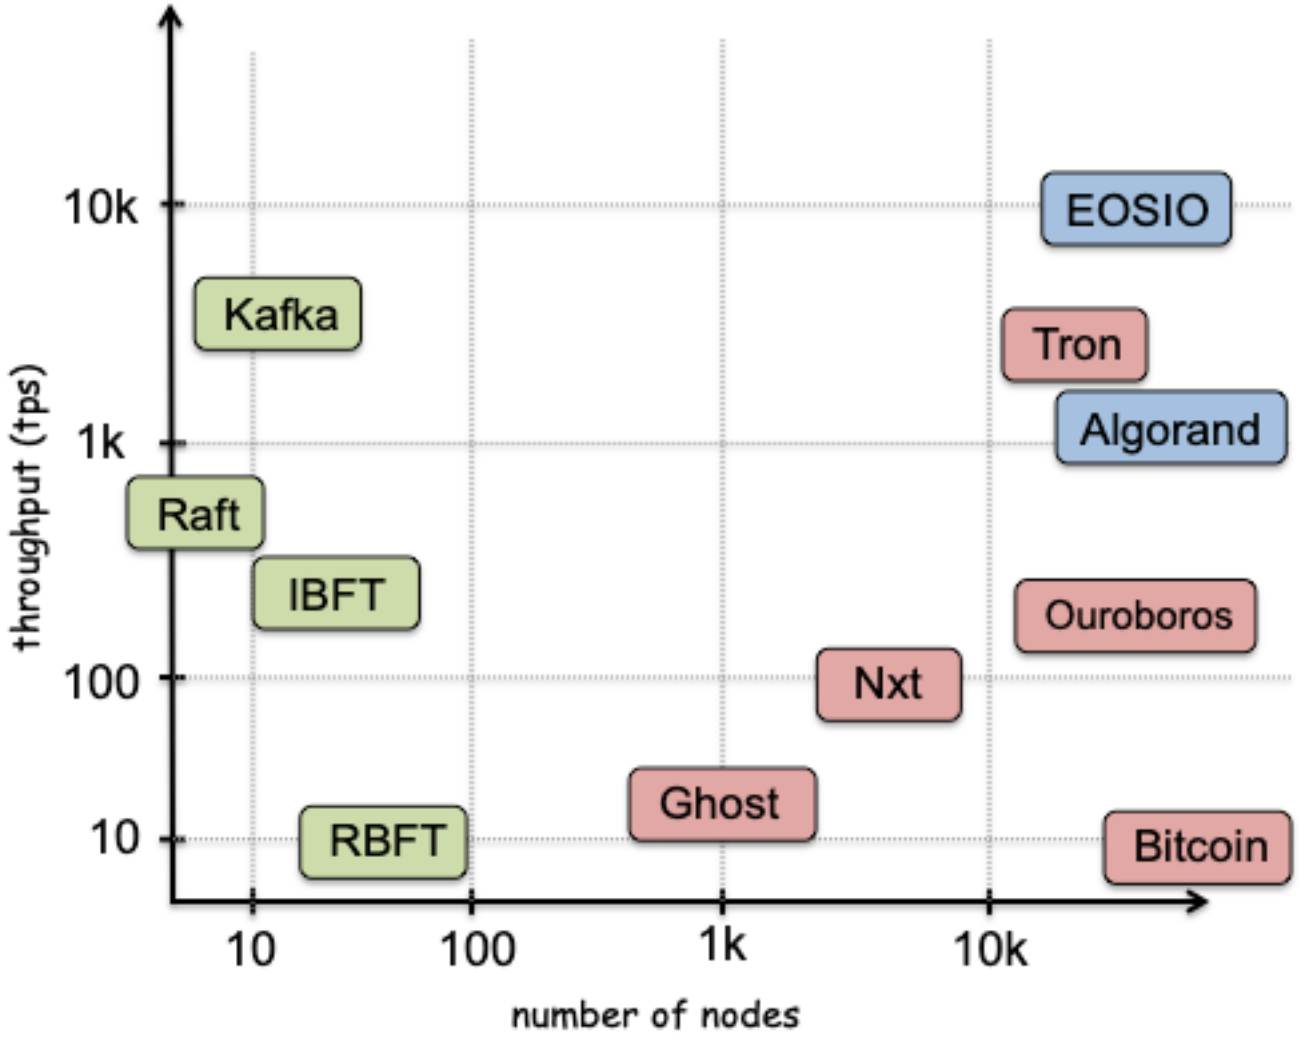

References

1. S. Bano, A. Sonnino, M. Al-Bassam, S. Azouvi, P. McCorry, S. Meiklejohn and G. Danezis, "SoK: Consensus in the age of blockchains," in Proc. ACM Conf. Adv. Financial Technol. (AFT), New York, NY, USA, 2019, pp. 183–198.

2. Oki, B., and Liskov, B. "Viewstamped Replication: A New Primary Copy Method to Support Highly-Available Distributed Systems", in Proc. of ACM Symposium on Principles of Distributed Computing (1988), pp. 8–17.

3. B. Liskov and J. Cowling, "Viewstamped replication revisited," MITCSAIL: Computer Science and Artificial Intelligence Laboratory, Boston, MA, USA, Tech. Rep. TR2012-021, 2012.

4. D. Ongaro and J. K. Ousterhout, "In Search of an Understandable Consensus Algorithm", in Proc. USENIX Annual Technical Conference, Philadelphia, PA, USA, 2014, pp. 305–320.

5. M. Castro and B. Liskov, "Practical Byzantine fault tolerance," in Proc. USENIX Symp. Oper. Syst. Design Implement. (OSDI), Feb. 1999, pp. 173–186.

References

6. P.L. Aublin, S. B. Mokhtar, and V. Quema, "Redundant Byzantine Fault Tolerance", in Proc. 33rd International Conference on Distributed Computing Systems (ICDCS), Philadelphia, PA, USA, 2013, pp. 297–306.

7. J. A. Garay, A. Kiayias, and N. Leonardos. "The bitcoin backbone protocol: Analysis and applications", in Proc. 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, 2015, pp. 281–330.

8. S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System", 2008. Accessed on: June 10, 2021. [Online]. Available: http://bitcoin.org/bitcoin.pdf

9. J. Bonneau, "Why buy when you can rent? bribery attacks on bitcoin consensus", in Proc. 20th Financial Cryptography and Data Security, Barbados, 2016, pp. 19–26

10. "Visanet: The technology behind visa". Accessed on: June 13, 2021. [Online]. Available: https://usa.visa.com/dam/VCOM/download/corporate/media/visanet-technology/visa-netbooklet.pdf

References

11. K. Croman, C. Decker, I. Eyal, A. E. Gencer, A. Juels, A. Kosba, A. Miller, P. Saxena, E. Shi, and E. G. Sirer, "On scaling decentralized blockchains," in Proc. 20th Financial Cryptography and Data Security, Barbados, 2016, pp. 106–125.

12. Y. Sompolinsky and A. Zohar, "Secure high-rate transaction processing in Bitcoin", in Proc. 19th Financial Cryptography and Data Security, San Juan, Puerto Rico, 2015, pp. 507–527.

13. I. Bentov, A. Gabizon, and A. Mizrahi, "Cryptocurrencies without proof of work", in Proc. 20th Financial Cryptography and Data Security, Barbados, 2016, pp. 142–157.

14. A. Kiayias, A. Russell, B. David, and R. Oliynykov, "Ouroboros: A provably secure proof-ofstake blockchain protocol", in Proc. 37th Annual International Cryptology Conference (CRYPTO 2017), Santa Barbara, CA, USA, 2017, pp. 357–388.

15. T. Kerber, A. Kiayias, M. Kohlweiss, and V. Zikas, "Ouroboros crypsinous: Privacy-preserving proof-of-stake", in Proc. IEEE Symposium on Security and Privacy, San Francisco, CA, USA, 2019, pp. 157–174.

References

16. H. Moniz, "The Istanbul BFT Consensus Algorithm", 2020. Accessed on: June 15, 2021. [Online]. Available: https://arxiv.org/pdf/2002.03613.pdf

17. C. Decker, J. Seidel, and R. Wattenhofer, "Bitcoin Meets Strong Consistency", *in Proc. the 17th International Conference on Distributing Computing and Networking*, Singapore, 2016, pp. 1–10.

18. E. Buchman, "Tendermint: Byzantine fault tolerance in the age of blockchains", M.S. thesis, University of Guelph, Canada, 2016.

19. V. Buterin and V. Griffith, "Casper the Friendly Finality Gadget", 2017. Accessed on: June 18, 2021. [Online]. Available: https://arxiv.org/pdf/1710.09437.pdf