

Formal Analysis of a Proof-of-Stake Blockchain

Wai Yan Maung Maung Thin
National University
of Singapore,
Singapore

Naipeng Dong
National University
of Singapore,
Singapore

Guangdong Bai
Griffith University,
Australia

Jin Song Dong
National University
of Singapore, Singapore
Griffith University, Australia

Abstract—Blockchain technology relies on consensus algorithms to resolve conflicts in Byzantine environments. New blockchain algorithms are rapidly designed and implemented without a properly conducted formal analysis and verification. In this paper, we conducted a study on Tendermint which is a proof-of-stake consensus algorithm. We verified that the consensus protocol is deadlock-free and is able to reach consensus when at least 2/3 of the network is in agreement. We also proved that a minority set of nodes that compose more than 1/3 of the network is enough to censor the majority of the network and prevent the network from reaching consensus and conclude that the algorithm has some shortcomings on availability.

I. INTRODUCTION

Introduced by Bitcoin in 2008 as a means to provide a trustworthy service among non-trusting peers without a central governing body, the blockchain technology has long evolved from its traditional use as a decentralized ledger for digital currencies. Many platforms such as Ethereum [1] and Ripple [2] are already using blockchains to support smart contracts which allow the execution of arbitrary codes in a decentralized manner [3]. The technology can also be easily applied to several fields such as Internet of Things (IoT) [4] or applications such as data privacy tools [5] and reputation systems [6].

The growth of the ecosystem and the adoption of the blockchain technology in various domains and applications have created a competitive environment where newer protocols and algorithms are being rapidly introduced and implemented without a formal analysis and verification process. The lack of related published research work implies that many of the widely used protocols and algorithms have yet to undergo a formal analysis that verifies the security properties and the correctness of the system. This process is extremely important, especially for the consensus algorithms that mandate how to resolve conflicting claims in the absence of a central authority and define the consistency, performance, scalability and liveness of the system.

Depending on the nature of the blockchain, the focus of the consensus algorithm is also different. For instance, Tendermint [7] favours consistency over liveness (availability) whereas Casper [8], the consensus algorithm of Ethereum sacrifices consistency to achieve plausible liveness. Previous research works on Tendermint blockchain [9] did not focus on the formal analysis of the consensus protocol. This paper aims to fill the gap by providing a formal analysis of the Tendermint consensus protocol in CSP# [10] using the PAT [11] model

checker. We prove that the consensus protocol is deadlock-free and is able to reach consensus when at least 2/3 of the network is in agreement.

II. BACKGROUND

A. Blockchain Architecture

A blockchain is a sequence of blocks, each of which maintains the hash value of itself and the link to its previous block (with the exception of the *genesis* block which is the first block in the chain). This design architecture ensures that the contents for an existing block (for example, transactions in a block) in the chain cannot be updated without modifying the hash values of the following blocks. Contents of a block widely vary based on the application. A typical block consists of a block header and a block body. The header is used to store the metadata of the previous and the current block, such as the signatures and the hash values. The body is used as a payload and can be used to store anything, ranging from messages and transactions to smart contracts.

A blockchain network is a set of nodes responsible for keeping track of the blockchain and validating and appending new blocks to the blockchain. Based on the difference in functionality, nodes can be categorized into two - miners and validators. Miners store only the last few blocks of the blockchain and are responsible for creating new blocks. Validators store a full copy of the blockchain and ensure that the new block submitted by the miners is valid and conforms to the rules of the blockchain. Based on the architecture of the blockchain, it is possible for a node to be both a miner and a validator at the same time.

B. Consensus in Blockchains

The nodes in the blockchain network operate under the fact that there is no central authority governing over them. In a perfect scenario, all the members of the network always agree on the same new block to be appended to the blockchain and there exists only a single blockchain in the whole network. In reality, however, nodes might get disconnected from the network or even act maliciously in Byzantine environments on purpose. Thus, a fault-tolerant consensus protocol is required which is agreed by all the nodes in the blockchain to resolve any potential conflict.

Consensus protocols differ based on the aspects or the attributes of the node that are being assessed. Proof-of-work (PoW) requires the node to provide the computing power




Fig. 1. State machine of the consensus protocol, adapted from *Tendermint: Consensus without mining* [7]

to solve a mathematical problem in order to append a new block to the blockchain. The most well-known user of PoW consensus is Bitcoin where miners have to solve cryptographic hash puzzles as proof of work. In Proof-of-stake (PoS), a node has to stake something it owns, usually in the form of a cryptocurrency. If a malicious node tries to manipulate the blockchain and the other nodes detect it, the locked-up stakes get slashed or rewards are withheld. Delegated proof-of-stake (DPoS) is where the nodes of the network vote for a set of nodes to be the delegators.

In a PoW system, the nodes are rewarded for performing an operation that is agreed by a majority of the nodes in the system. The caveat here is that participants are not punished for performing a malicious operation. As a result, PoW systems cannot deter the participants from performing a selfish mining [12] or participating in a 51% attack. In order to solve this problem, newer generations of blockchains (Ethereum, Tendermint, etc.) have started to use proof-of-stake as the consensus algorithm. In a PoS system, the participants are rewarded for performing a non-malicious activity just like a PoW system but they are also held accountable and are punished for any malicious operation.

III. TENDERMINT

Tendermint is a proof-of-stake consensus protocol that is Byzantine fault tolerant. Participants in the protocol are called validators. There is no concept of miners in Tendermint and thus, validators are also responsible for the creation of new blocks. The height of the chain increases every time a block is added to the chain. Validators are chosen in a round-robin manner to become the proposer who is in charge of creating and proposing a block for the current round. Validators are required to post a *bond transaction* that will lock a set amount of his coins (stakes) for a set duration. If the validator is found to be involved in any malicious activity within this duration, it can be punished by slashing away its deposited stake. After this duration, stakes is unlocked and returned to the validator.

A. Consensus Algorithm

The consensus algorithm consists of five steps - *Propose*, *Prevote*, *Precommit*, *Commit* and *NewHeight*.

In the *Propose* step, the proposer broadcasts a proposal to its peers. A proposal includes the block, the signatures of the validators who have validated the block, the signature of the proposer as well as the round and the height information. If the proposer has already locked on a block during the precommit of the previous round, that block will be used for proposal.

Otherwise, a new block will be created. All nodes will gossip the *proposal* to their neighbouring peers during this period.

In the *Prevote* step, each validator will vote for a block and gossip it to the neighbours. A vote consists of the hash of the voted block, the signature of the voter, type of vote - whether it is a *prevote* or a *precommit* plus the round and the height information. The block to be included is chosen in the following order - (1) a locked proposed block from prior rounds and (2) a valid acceptable block from the current proposal. If neither is available, a special NIL prevote is broadcast to the neighbours. All nodes will gossip all *prevotes* for the round to their neighbouring peers.

In the *Precommit* step, the validator checks if it has received more than $2/3$ of prevotes for an acceptable block. If there is one, the validator releases the existing lock, instead locks onto this block and signs and broadcasts a precommit vote for this block. The validator also packages the prevotes for the locked block into a proof-of-lock which will be used to create the block in the next *Proposal*. In the case where there are fewer than $2/3$ prevotes, the validator will neither sign nor lock on any block. During this period, all nodes will gossip all *precommits* for the round to all neighbouring peers. At the end of *Precommit*, if the node has received more than $2/3$ of precommits for a particular block, it will proceed to *Commit* step. Otherwise, it transits to the *Propose* step of the next round.

In the *Commit* step, two parallel conditions must be fulfilled before the consensus algorithm can cycle back to *Propose* step. First and foremost, the node must have received the block from one of its peers so that it can sign and broadcast a commit to other peers. Second, the node must wait until at least $2/3$ commits of the block are received by the network. Once these are satisfied, the node will set the *CommitTime* to current time and move on to *NewHeight* step where the nodes will stay for a fixed duration. The purpose is to allow the nodes to wait for additional commits of the committed block which were not received in *Precommit* due to network latency issues. After the set duration is up, the algorithm starts again from *Propose*. At anytime during the consensus process, if a node receives more than $2/3$ commits for a particular block, it will immediately enter the *Commit* step.

IV. FORMAL VERIFICATION

Our Tendermint consensus model is built using CSP# [10] and is verified using PAT [11] model checker. CSP# integrates the high-level modeling operators of CSP with low-level procedural codes in C# language and supports custom data structures which is extensively used in our model to represent the properties and the behaviors of the blocks and the blockchain. We define a set of assertion rules to capture the security properties that we want to verify against the model. We run the verification through PAT which either gives us verification of satisfied assertions or the traces of the unsatisfied assertions. The complete model and the runtime for each verification are available at the GitHub repository[13].

In this paper, we focused on the consensus protocol of the blockchain and as such we have abstracted the structure and the content of the block. We modelled the channel communication among nodes only for consensus. Data exchange among the nodes for the internal contents of the blocks such as transactions and contracts is considered to be outside of the scope of this work. We ran two sets of verification - one with 3 validators in the network and the other with 4 validators (where 1 node represents $< 1/3$ of the network; 2 nodes represent $\geq 1/3$ and $\leq 2/3$ of the network and; 3 nodes represent $> 2/3$ of the network). The number of nodes in the network is limited to 4 to reduce the number of states in the system. This is the minimum amount of nodes required to prove the properties of the model while keeping the states small enough for PAT to verify. The following assumptions have been made in the model.

- 1) All the nodes in the network are connected to each other.
- 2) Existing nodes will not leave the network and no new nodes will join the network.
- 3) All nodes have the same weight or voting power and thus have the same stake.
- 4) There is no latency in the network.

A. Modelling

We define a blockchain as a set of nodes N where each node $x \in N$ executes a sequence of processes in parallel. Since there is no network latency in our model, *NewHeight* is not required and *Commit* does not need to wait until the node receives the block. We also split the *PreparePOL* from *Precommit* since preparing proof-of-lock happens at the end of *Precommit* and not during. At the end of each round, the *NextRound()* process is executed to select the next proposer.

$BlockChain() = (\parallel x : \{0..N - 1\} \bullet (Propose(x); Prevote(x); Precommit(x); PreparePOL(x); Commit(x))); NextRound());$

where $P; Q$ represents process P followed by process Q and $P \parallel Q$ represents synchronous processes P and Q .

The *Propose(x)* process is defined as follows. If a node is a *proposer* for this round, it will broadcast the proposal to other nodes using the communication channels. Otherwise, it will wait for a proposal from the proposer.

The *Prevote(x)* process is a sequence of two processes, prefixed by event *propose_end* which is used to represent the end of the *Propose* step. The first process *Prevote_x()* validates the proposal received from *Propose(x)* and votes accordingly. The second process *BroadcastPrevotes()* represents the broadcasting and the receiving of votes among the nodes.

The malicious node behavior is simulated in process *Prevote_x()*. An honest node validates the block as is whereas a malicious node with the intent to overwrite the proposed block broadcasts a different block from the one it received. A malicious node with the intent to vote an invalid block is simulated as voting for a duplicate block that already exists in the chain.

The *Precommit(x)* process is similar to *Prevote(x)*, defined by another sequence of two processes. The first process *Precommit_x()* retrieves the first block with at least $2/3$ confidence from the votes received from *Prevote(x)* step. The second process *BroadcastPrecommits()* is similar to *BroadcastPrevotes()*.

The *PreparePOL(x)* process prepares the proof of lock, gathers the signatures of the validators who voted for this block and stores them inside the block.

The *Commit(x)* process adds the precommitted block having at least $2/3$ consensus to the chain. The node adds the block to the chain only if it is not *NIL*. This marks the end of one round of consensus.

B. Properties

We aim to verify the following set of properties that are claimed.

- **Deadlockfree-ness.** A deadlock-free model means that at any point in time, no node should be waiting for another node to broadcast or receive a proposal or a vote, continue with validating of a block, or committing a block to its chain, even in Byzantine environments. We can use PAT to query for the deadlocks in the system.
- **Ability to reach consensus.** All the nodes in the network must be able to agree on one block and commit the same block to their own copy of the chain at the end of a round.
- **Overwriting proposed block.** This is an attack where a set of nodes collude to reject the proposed block and propose a new block of their own. For a blockchain to be resistant to this attack, all nodes must agree on the same block to be added to the chain after reaching consensus and the block must be identical to the originally proposed block.
- **Invalid blocks.** The network must be able to reject invalid blocks submitted by malicious nodes that do not conform to the rules of the blockchain. This is simulated by malicious nodes trying to add a duplicate (existing) block to the chain.
- **Censorship attacks.** Censorship attack is an attack where malicious nodes in the network refuse to broadcast or vote a valid block in order to censor a particular content of the block or censor the node itself. The blockchain is immune to this attack if other nodes can reach consensus even with the absence of malicious nodes in the voting process. The assertions used to check for the reachability of consensus can be reuse for this verification.

V. RESULTS AND ANALYSIS

We will be verifying the properties of the model under normal conditions as well as against specific attacks in Byzantine environments. We define nine more variations of *Blockchain()* with different percentage of malicious nodes in the system with different threat factors.

- P1 $BlockChainWithMinorityOverwrite() = SimulateMalicious (MINORITY, OVERWRITE_VOTING); BlockChain();$
P2 $BlockChainWithHalfOverwrite() = SimulateMalicious (HALF, OVERWRITE_VOTING); BlockChain();$
P3 $BlockChainWithMajorityOverwrite() = SimulateMalicious (MAJORITY, OVERWRITE_VOTING); BlockChain();$

P4	<i>BlockChainWithMinorityInvalid()</i> = <i>SimulateMalicious</i> (<i>MINORITY, INVALID_BLOCK_VOTING</i>); <i>BlockChain()</i> ;
P5	<i>BlockChainWithHalfInvalid()</i> = <i>SimulateMalicious</i> (<i>HALF, INVALID_BLOCK_VOTING</i>); <i>BlockChain()</i> ;
P6	<i>BlockChainWithMajorityInvalid()</i> = <i>SimulateMalicious</i> (<i>MAJORITY, INVALID_BLOCK_VOTING</i>); <i>BlockChain()</i> ;
P7	<i>BlockChainWithMinorityCensor()</i> = <i>SimulateMalicious</i> (<i>MINORITY, NO_VOTING</i>); <i>BlockChain()</i> ;
P8	<i>BlockChainWithHalfCensor()</i> = <i>SimulateMalicious</i> (<i>HALF, NO_VOTING</i>); <i>BlockChain()</i> ;
P9	<i>BlockChainWithMajorityCensor()</i> = <i>SimulateMalicious</i> (<i>MAJORITY, NO_VOTING</i>); <i>BlockChain()</i> ;

Table I presents the verification results of properties of the model in normal and Byzantine environments. The model is **deadlock-free (T1)** at all times but does not always reach **consensus (T2)**. It is important to note that not being able to reach consensus does not always mean the system is vulnerable. Based on the threat model, it might be better for the network to agree on no block than to agree on the manipulated block. As such, consensus should be analyzed with respect to each individual threat model A1-A3.

TABLE I
RESULTS FOR DEADLOCK-FREENESS AND CONSENSUS REACHABILITY

	T1	T2	A1	A2	A3
P0 <i>BlockChain</i>	✓	✓	✓	✓	✓
P1 (<i>overwrite</i> ≤ 1/3)	✓	✓	✓		
P2 (1/3 < <i>overwrite</i> < 2/3)	✓	✗	✓		
P3 (<i>overwrite</i> ≥ 2/3)	✓	✓	✗		
P4 (<i>invalid</i> ≤ 1/3)	✓	✓		✓	
P5 (1/3 < <i>invalid</i> < 2/3)	✓	✗		✓	
P6 (<i>invalid</i> ≥ 2/3)	✓	✗		✓	
P7 (<i>no_vote</i> ≤ 1/3)	✓	✓			✓
P8 (1/3 < <i>no_vote</i> < 2/3)	✓	✗			✗
P9 (<i>no_vote</i> ≥ 2/3)	✓	✗			✗

A1: Proposed block overwrites We can see that process *BlockChain* (P0), *BlockChainWithMinorityOverwrite* (P1) and *BlockChainWithHalfOverwrite* (P2) are immune to overwrites. We have also learnt from Table I that *BlockChainWithHalfOverwrite* (P2) failed to reach consensus. By failing to reach consensus, the network did not commit a block during this round. This is in contrast to *BlockChainWithMajorityOverwrite* (P3) where the manipulated block voted by the majority (≥ 2/3) is added to the network instead of the proposed block.

A2: Invalid block insertions All four variations of the *BlockChain* (P0, P4, P5, P6) can withstand the invalid block insertions. The only difference is that *BlockChainWithMinorityInvalid* (P4) and *BlockChain* are able to reach consensus whereas the other two are not. *BlockChainWithMinorityInvalid* (P4) is able to insert a valid block to the chain but no block is inserted to the chain in the case of the other two.

A3: Censorship attacks. As for the censorship attacks, we find that the blockchain becomes vulnerable once the number of malicious nodes increases past 1/3. We have stated before that it takes ≥ 2/3 of the network to collude to vote on a same non-proposed block to add it to the blockchain. However, if the intent is not for overwriting but just for censorship or disruption of the service, this can be achieved with merely > 1/3 of

the network just by having them not participating in the voting process. It is also possible that nodes get disconnected from the network without any malicious intent. In both cases, the chain will be prevented from adding new blocks until ≥ 2/3 of the network participates in the voting. Since the proposer is chosen by round robin, it is possible for the malicious minority nodes to stay offline until it is their turn to vote or until the next *Propose* step begins, effectively censoring the majority of the network from publishing anything.

VI. CONCLUSIONS

In this paper, we discussed the need for the formal analysis and verification of blockchain protocols. We formally modelled a simple consensus algorithm for a proof-of-stake blockchain based on Tendermint and verified that the algorithm is deadlock free. We also proved that it takes ≥ 2/3 of the network to reach consensus and it takes > 1/3 of the network to censor the majority of the nodes from publishing a new block. It is possible to update the algorithm that chooses the proposer to ignore disconnected nodes in the previous round but this will only mitigate the problem. Alternative solutions to this are to implement a time lock consensus protocol [14] or to punish the nodes that are offline regardless of their intent, similar to how a node is punished by participating in malicious voting in proof-of-stake.

Acknowledgement: This work is partially supported by NRF2015NCR-NCR003-003.

REFERENCES

- [1] V. Buterin, "Ethereum white paper," *GitHub repository*, 2013.
- [2] D. Schwartz, N. Youngs, and A. Britto, "The ripple protocol consensus algorithm," *Ripple Labs Inc White Paper*, vol. 5, 2014.
- [3] M. Swan, *Blockchain: Blueprint for a new economy*. O'Reilly Media, Inc., 2015.
- [4] Y. Zhang and J. Wen, "An IoT electric business model based on the protocol of bitcoin," in *18th International Conference on Intelligence in Next Generation Networks (ICIN)*. IEEE, 2015, pp. 184–191.
- [5] G. Zyskind, O. Nathan, and A. S. Pentland, "Decentralizing privacy: Using blockchain to protect personal data," in *Security and Privacy Workshops (SPW)*. IEEE, 2015, pp. 180–184.
- [6] M. Sharples and J. Domingue, "The blockchain and kudos: A distributed system for educational record, reputation and reward," in *European Conference on Technology Enhanced Learning*. Springer, 2016, pp. 490–496.
- [7] J. Kwon, "Tendermint: Consensus without mining," *Retrieved May*, vol. 18, p. 2017, 2014.
- [8] V. Buterin and V. Griffith, "Casper the friendly finality gadget," *arXiv preprint arXiv:1710.09437*, 2017.
- [9] E. Buchman, "Tendermint: Byzantine fault tolerance in the age of blockchains," Ph.D. dissertation, The University of Guelph, 2016.
- [10] J. Sun, Y. Liu, J. S. Dong, and C. Chen, "Integrating specification and programs for system modeling and verification," in *Third IEEE International Symposium on Theoretical Aspects of Software Engineering (TASE)*. IEEE, 2009, pp. 127–135.
- [11] J. Sun, Y. Liu, J. S. Dong, and J. Pang, "Pat: Towards flexible verification under fairness," in *International Conference on Computer Aided Verification*. Springer, 2009, pp. 709–714.
- [12] I. Eyal and E. G. Sirer, "Majority is not enough: Bitcoin mining is vulnerable," in *International conference on financial cryptography and data security*. Springer, 2014, pp. 436–454.
- [13] W. Y. M. M. Thin, N. Dong, G. Bai, and J. S. Dong, "Formal analysis of a proof-of-stake blockchain," <https://github.com/waiyan1612/blockchain-analysis>, 2018.
- [14] V. Buterin, "The problem of censorship," 2015. [Online]. Available: <https://blog.ethereum.org/2015/06/06/the-problem-of-censorship/>