



HAL
open science

Ethereum Proof-of-Stake under Scrutiny

Ulysse Pavloff, Yackolley Amoussou-Guenou, Sara Tucci-Piergiovanni

► **To cite this version:**

Ulysse Pavloff, Yackolley Amoussou-Guenou, Sara Tucci-Piergiovanni. Ethereum Proof-of-Stake under Scrutiny. [Technical Report] CEA DILS. 2022. hal-03821290v1

HAL Id: hal-03821290

<https://hal.science/hal-03821290v1>

Submitted on 28 Oct 2022 (v1), last revised 11 Sep 2023 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Ethereum Proof-of-Stake under Scrutiny

Ulysse Pavloff  

Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France

Yackolley Amoussou-Guenou

Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France

Sara Tucci-Piergiovanni 

Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France

Abstract

Ethereum has undergone a recent change called *the Merge*, which made Ethereum a Proof-of-Stake blockchain shifting closer to BFT consensus. Ethereum, which wished to keep the best of the two protocols designs (BFT and Nakamoto-style), now has an involved consensus protocol as its core. The result is a blockchain being possibly produced in a tree-like form while participants try to finalize blocks. Several attacks jeopardizing liveness have been found in this new setting. The Ethereum community has responded by creating a patch. We discovered a new attack on the patched protocol. To support our analysis, we propose a new formalization of the properties of liveness and availability of the Ethereum blockchain, and we provide a pseudo-code. We believe this formalization to be helpful for other analyses as well. Our results yield that the Ethereum Proof-of-Stake has probabilistic liveness, influenced by the parameter describing the time frame allowed for validators to change their mind about the current main chain.

2012 ACM Subject Classification Theory of computation → Distributed algorithms; Computer systems organization → Dependable and fault-tolerant systems and networks

Keywords and phrases Ethereum Proof-of-Stake, Liveness, Availability, Bouncing attack

Digital Object Identifier 10.4230/OASICS...

1 Introduction

Ethereum has recently undergone a major change in its protocol, successfully passing from proof-of-work to proof-of-stake. The change underpins an entirely new consensus protocol, which brings Byzantine fault-tolerance to Ethereum. The main design goal is to keep using a Nakamoto-style consensus, i.e., a protocol that constantly creates blocks in a tree-like form and selects a branch as the current chain using a fork-chain rule. However, a new mechanism (called finality gadget) incrementally finalizes blocks in the chain as opposed to pure Nakamoto-style consensus. A *finalized block* is a block that is voted by at least two-thirds of validators¹. In a system with less than one-third of Byzantine validators, a finalized block is never revoked.

Interestingly, this design aims at guaranteeing, at the same time, the availability of the chain and consistency (uniqueness) of a finalized prefix. Note that classical BFT consensus protocols re-adapted to blockchains such as Tendermint [6] for the Cosmos blockchain [15] or Tenderbake [5] for Tezos blockchains [14], finalize one block at the time: for each height of the blockchain only one block is ever added. On the other hand, Ethereum PoS builds a common prefix, but the suffix can change: for a given height of the blockchain, different blocks can be seen at that height over time. The advantage of this approach is to always make progress, regardless of Byzantine behavior and network partitions, while classical

¹ To become a validator, one needs to 'stake' an amount of 32 Eth (the native cryptocurrency of the blockchain).



BFT consensus protocols stop producing blocks during asynchronous periods and attacks. Ethereum Proof-of-Stake tries then to provide consistency without renouncing availability. This seems to be in striking contrast with the CAP theorem [13], which states that it is impossible to guarantee progress and consistency in the case of network partitions. The caveat here is that Ethereum maintains a data structure where the prefix is consistent and finalized only when possible, while the suffix can grow without being consistent. The resulting protocol, however, is quite involved. For this reason, we aim to provide a formal ground for analysis of the Ethereum Proof-of-Stake (PoS) protocol.

We first provide a novel formalization of the Ethereum Proof-of-Stake blockchain properties, which are a combination of properties of Nakamoto-style blockchains and BFT-style ones. We then provide a formalization of the protocol itself through pseudo-code. Formalisation allows to analyse the protocol in terms of its properties.

In this paper, we focus on the liveness of the finalization process, where liveness is the capacity to always finalize new blocks. Our analysis reveals a new possible attack on the protocol's liveness. Indeed previous work has already pointed out the risks of a so-called *bouncing attack* on liveness [17]. A bouncing attack is an attack that prevents the chain from being finalized because the main chain selected through the fork chain rule continually bounces between two alternative branches. After the attack identification, the Ethereum community responded by implementing a patch to the protocol to prevent this attack. The patch aims at mitigating bouncing by forcing validators to stick to a chain after a while. Interestingly, we managed to find a bouncing attack on the patched version of the protocol. We found that the bouncing can be repeated over time but with decreasing probability of success. This shows that the liveness of the patched protocol is probabilistic. Note that the attack is plausible in a Byzantine environment since it only relies on the Byzantine validator's capacity to withhold votes and to release them at the right time to make honest validators change their mind on the chosen chain.

The paper is organised as follows: Section 2 elaborates on the system model while Section 3 defines the properties essential to our formalization. In Section 4, we explain and formalize the Ethereum PoS protocol providing all the materials necessary for its understanding, including pseudo-code. Section 5 presents a new liveness attack still possible in the current version of the Ethereum PoS protocol. We present the related works in Section 6, and we conclude in Section 7.

2 System model

We consider a system composed of a finite set Π of processes called *validators*². There are a total of n validators. Each validator has an associated public/private key pair for signing and can be identified by its public key. We assume that digital signatures cannot be forged. Validators have synchronized clocks³. Time is measured by periods of 12 seconds called *slots*, a period of 32 slots is called an *epoch*.

Network

Processes communicate by message passing. We assume the existence of an underlying broadcast primitive, which is a best effort broadcast. This means that when a correct process

² At the implementation level, validators are the processes with ETH staked that allow them to vote as part of the consensus protocol.

³ Clocks can be offset by at most τ , this way, the offset can be captured as part of the network delay.

broadcasts a value, all the correct processes eventually deliver it. Messages are created with a digital signature.

We assume a *partially synchrony model* [9], where after some unknown Global Stabilization Time (GST), the system becomes synchronous, and there is a finite known bound Δ on the message transfer delay. Note that even if we have synchronized clocks, having an asynchronous network before GST still makes the system partially synchronous.

Fault Model

Validators can be *correct* or *Byzantine*. Correct validators follow the protocol, while Byzantine ones may arbitrarily deviate from the protocol⁴. We denote by f the number of Byzantine validators, with $f < n/3$.

3 Blockchain properties

In our analysis, we will continuously use the term Ethereum Proof-of-Stake as [24] to name the new protocol of Ethereum⁵. To begin our analysis of Ethereum Proof-of-Stake, we start by defining the terms and properties that we will investigate.

Similarly to [3] we formalize the data structure implemented by blockchain as a *BlockTree*. Indeed the blockchain takes the form of a tree in which every node is a block pointing to its unique parent, and the tree's root is the *genesis block*. Among the different branches of the BlockTree, the protocol indicates a unique branch, or chain, to build upon with a so-called fork choice rule (e.g., the longest chain rule in Bitcoin). The selected chain is called the *candidate chain*.

► **Definition 1 (Candidate chain).** We call *candidate chain* the chain designated as the one to build upon by the fork choice rule. Considering the view of the chain of an honest validator i , i 's associated candidate chain is noted C_i .

The blocks in the candidate chain can be finalized or not.

► **Definition 2 (Finalized block).** A block is finalized for a validator i if and only if the block cannot be revoked, i.e., it permanently belongs to the candidate chain C_i .

Note: It stems from the definition that all the predecessors of a finalized block are finalized.

► **Definition 3 (Finalized chain).** The finalized chain is the chain constituted of all the finalized blocks.

Note: The finalized chain C_{fi} is always a prefix of any candidate chain C_i .

To analyze the protocol, one needs to analyze the capability of the Ethereum Proof-of-Stake protocol to construct a consistent blockchain (safety), to allow validators to add blocks despite network partitions and failures (availability), and to make constant progress on the finalization of new blocks (liveness). Safety, availability, and liveness are expressed as follows:

⁴ Since in this paper we are only interested in the Consensus protocol, we only characterize validator's behavior. For clients submitting transactions, as in any blockchain, we assume they can be Byzantine.

⁵ Other appellation such as Ethereum 2.0 or Consensus Layer can be found, we have chosen to stick with Ethereum Proof-of-Stake.

XX:4 Ethereum Proof-of-Stake under Scrutiny

► **Definition 4 (Safety).** A blockchain is consistent or *safe* if for any two correct validators i and j , having a finalized chain C_{f_i} and C_{f_j} , respectively, then C_{f_i} 's is the prefix of C_{f_j} or viceversa.

► **Definition 5 (Availability).** A blockchain is *available* if the following two conditions hold: (1) any correct validator is able to append a block to the candidate chain in bounded time, regardless of the failures of other validators and the network partitions; (2) the candidate chain is eventually growing, i.e., given a block b_k added to its candidate chain at a distance d from the genesis block b_0 , where the distance is the number of blocks separating b_k from b_0 , then eventually a block b_l will be added to the candidate chain at a distance $d' > d$.

► **Definition 6 (Liveness).** A blockchain is *live* if the finalized chain is ever growing.

The fundamental difference between the finalized and the candidate chain lies in the fact that blocks of the finalized chain can never be revoked, while the candidate chain can change from one branch to another in the tree so that a suffix of blocks of the previously selected branch might be revoked. Availability, on the other hand, guarantees that adding blocks to the candidate chain is a wait-free operation whose time to complete does not depend on network failures or Byzantine behaviors. Availability also implies that blocks are constantly added in such a way that the height of the candidate chain eventually grows. This property avoids the pathological scenario in which all the blocks are added to the genesis block to form a star.

As in any distributed system, blockchains are faced with the dilemma brought by the CAP-Theorem [13]. This theorem states that distributed systems cannot satisfy these three properties at the same time: *consistency*, *availability*, and *partition tolerance*. Indeed, if network partitions occur, either the system remains available at the expense of consistency, or it stops making progress until the network partition is resolved to guarantee consistency. This means that no blockchain can simultaneously be available and safe. However, by maintaining the candidate and the finalized chain at the same time, Ethereum Proof-of-stake aims to offer both safety and availability. The candidate chain aims to be available but without guaranteeing consistency all the time, while the finalized chain falls on the other side of the spectrum, guaranteeing consistency without availability. Therefore, the finalized chain will finalize blocks only when it is safe to do so where the candidate chain will still be available during network partitions (caused by network failures or attacks). The only caveat here is that finalized chain grows by finalizing blocks of the candidate chain, which means that the properties of the two chains are interdependent. In particular to assure liveness it is necessary that the candidate chain steadily grows. This interdependence is a source of vulnerability that must be thoroughly analyzed.

4 Ethereum Proof-of-Stake Protocol

4.1 Overview

The Ethereum Proof-of-Stake (PoS) protocol design is quite involved. We identify, similarly to [19], the objectives underlying its design as follows: (i) finalizing blocks and (ii) having an available candidate chain that does not rely on block finality to grow. To this end, the Ethereum Proof-of-Stake protocol combines two blockchain designs: a Nakamoto-style protocol to build the tree of blocks containing the transactions and a BFT finalization protocol to progressively finalize blocks in the tree. The objective is to keep the blockchain creation

process always available while guaranteeing the finalization of blocks through Byzantine-tolerant voting mechanisms. The finalization mechanism is a *Finality Gadget* called *Casper FFG*, and the fork choice rule to select candidate chains is *LMD GHOST*.

Before introducing how the fork choice rule and the finality gadget work together, we will introduce the following basic concepts: (i) slots, epochs, and checkpoints, which set the pace of the protocol allowing validators to synchronize together on the different steps, (ii) committees formation and assignment of roles to validators as proposers and voters for each slot, and (iii) the different types of votes the validators must send in order to grow and maintain the candidate chain as well as the finalized chain.

In this section, we focus on providing a formal version of the protocol following the specification given by the Ethereum Foundation [11]. Every implementation of the protocol must be compliant with the specification. Note that a description of an initial plan of the protocol is proposed by Buterin and al. in [7]. In this paper, we describe and formalize, through pseudo-code, the current implementation of the protocol.

4.1.1 Slots, Epochs & Checkpoints

In proof of work protocols, such as originally in [16], the average frequency of the block creation is predetermined in the protocol, and the mining difficulty changes to follow that pace. On the contrary, in Ethereum Proof-of-Stake, each validator uses the hypothesis of synchronized clocks to propose blocks at regular intervals. More specifically, in the protocol, time is measured in *slots* and *epochs*. A slot lasts 12 seconds. Slots are assigned with consecutive numbers; the first slot is slot 0. Slots are encapsulated in *epochs*. An epoch is composed of 32 slots, thus lasting 6 minutes and 24 seconds. The first epoch (epoch 0) contains from slot 0 to slot 31; then epoch 1 contains slot 32 to 63, and so on. These slots and epochs allow associating the validators' roles to the corresponding time frame. An essential feature of epochs is the *checkpoint*. A checkpoint is a pair block-epoch (b, e) where b is the block of the first slot⁶ of epoch e .

4.1.2 Validators & Committees

Validators have two main roles: *proposer* and *attester*. The proposer's role consists in proposing a block during a specific slot⁷. This role is pseudo-randomly given to 32 validators by epoch (one for each slot). The attester's role consists in producing an attestation sharing the validator's view of the chain. This role is given once by epoch to each validator.

During every epoch, each validator is assigned to one committee. A committee C_j is a set of validators. One validator belongs to exactly one committee, i.e., $\forall j \neq k, C_j \cap C_k = \emptyset$. Each committee is associated with a slot. During this slot, each member of the committee will have to perform an *attestation* to indicate its view of the chain.

In short, during an epoch, validators are all attesters once and have a small probability of being proposers ($32/n$).

⁶ In the event of an epoch without a block for the first slot, the block used for the checkpoint is the last block in the candidate chain, belonging then to a previous epoch. On the contrary, if the proposer of the first slot proposes multiple blocks, this will make multiple checkpoints for the other validators to choose from using the fork choice rule.

⁷ The current protocol specifications [11] indicate that correct validators should send their block proposition during the first third of their designated slot.

4.1.3 Vote & Attestation

There are two types of votes in Ethereum Proof-of-Stake, the *block vote*⁸ and the *checkpoint vote*⁹. The message containing these two votes is called *attestation*. During an epoch, each validator must make one attestation. The attestation ought to be sent during a specific slot. This slot depends on the committee of which the validator is a member. The two types of votes, checkpoint vote and block vote, have very distinct purposes. The checkpoint vote is used to finalize blocks to grow the finalized chain. The block vote is used to determine the candidate chain. Although validators cast their two types of votes in one attestation, an important distinction must be made between the two. Indeed, the two types of votes do not require the same condition to be taken into account. The checkpoint vote of an attestation is only considered when the attestation is included in a block. In contrast, the block vote is considered one slot after its emission, whether it is included in a block or not. The code associated with the production of attestations is described in Algorithm 3 at Subsection 4.2. We then describe in Algorithm 6 how the reception of attestations is handled.

4.1.4 Finality Gadget

The finality gadget is the mechanism that aims at finalizing blocks. The finality gadget grows the finalized chain disregarding the block production. This decoupling of the finality mechanism and the block production permits block availability even when the finalizing process is slowed down. This differs from protocols like Tendermint [6], where a new block can be added to the chain only after being finalized.

The finality gadget works at the level of epochs. Instead of finalizing blocks one by one, the protocol uses checkpoint votes to finalize entire epochs. We now present in more detail how the finality gadget of Ethereum PoS grows the finalized chain.

Recall that to be taken into account, a checkpoint vote needs to be included in a block. The vote will then influence the behavior of validators regarding this particular branch. Thus, in Algorithm 9 of Subsection 4.2 the function `countMatchingCheckpointVote` only count the matching checkpoint vote of attestation included in a block. More specifically, this function takes a source and a target as an argument and only count checkpoint vote with matching source and target.

Justification

The justification process is a step to achieve finalization¹⁰. It operates on checkpoints at the level of epochs. Justification occurs thanks to checkpoint votes. The checkpoint vote contains a pair of checkpoints: the checkpoint *source* and the checkpoint *target*. We can count with `countMatchingCheckpointVote` the sum of balances of the validator's checkpoint votes with the same source and target. If validators controlling more than two-thirds of the stake make the same checkpoint vote, then we say there is a *supermajority link* from the checkpoint source to the checkpoint target. The checkpoint target of a supermajority link is said to be *justified*.

More formally, a checkpoint vote is in the form of a pair of checkpoints: $((a, e_a), (b, e_b))$, also noted $(a, e_a) \rightarrow (b, e_b)$. For the checkpoint vote $(a, e_a) \rightarrow (b, e_b)$ we call (a, e_a) the

⁸ Also called GHOST vote.

⁹ Also called FFG vote.

¹⁰ The genesis checkpoint (i.e., the checkpoint of the first epoch) is the exception to this rule: it is justified and finalized by definition.


checkpoint source and (b, e_b) the checkpoint target. The checkpoint source is necessarily of an earlier epoch than the checkpoint target, i.e., $e_a < e_b$. In line with [7] we say there is a *supermajority link* from checkpoint (a, e_a) to checkpoint (b, e_b) if validators controlling more than two-thirds of the stake cast an attestation with checkpoint vote $(a, e_a) \rightarrow (b, e_b)$. In this case, we write $(a, e_a) \xrightarrow{J} (b, e_b)$ and the checkpoint (b, e_b) is justified.

Finalization

The finalization process aims at finalizing checkpoints, thus growing the finalized chain. Checkpoints need to be justified before being finalized. Let us illustrate the finalization process with the two scenarios that can lead to finalization. The first case presents the main scenario in the synchronous setting. It shows how a checkpoint can be finalized in two epochs, the least amount of epochs needed for finalization.

Case 1: The scenario is depicted in Figure 1.

1. Let $A = (a, e)$ and $B = (b, e + 1)$ be checkpoints of two consecutive epochs such that $A = (a, e)$ is justified.
2. A supermajority link occurs between checkpoints A and B where A is the source and B the target. This justifies checkpoint B . Hence, we can write: $(a, e) \xrightarrow{J} (b, e + 1)$ or equivalently $A \xrightarrow{J} B$.
3. This leads to A being finalized.



■ **Figure 1** The figure depicts the finalization scenario of **case 1** with the 3 steps from top to bottom. We represent a checkpoint with a hexagon, a justified checkpoint with a double hexagon, and a finalized checkpoint with double hexagon coloured. The arrow between two checkpoints indicates a supermajority link.


The second case illustrates the scenario in which two consecutive checkpoints are justified but not finalized. This means that the current highest justified checkpoint (e.g., B in Figure 2) was not justified with a supermajority link having the previous checkpoint A as its source. Then occurs a new justification with the source and target being at the maximum distance (2 epochs) for the source to become finalized. An important note is that there is no limit on the distance between two checkpoints for justification to be possible. This limit only exists for finalization.

Case 2: The scenario is depicted in Figure 2.

1. Let $A = (a, e)$, $B = (b, e + 1)$ and $C = (c, e + 2)$ be checkpoints of consecutive epochs such that A and B are justified. There is no supermajority link between A and B , A cannot be finalized as in Case 1 above.

XX:8 Ethereum Proof-of-Stake under Scrutiny


2. Now, a supermajority link occurs between checkpoints A and C where A is the source and C the target. This justifies checkpoint C , i.e., $A \xrightarrow{J} C$.
3. This leads to A being finalized.



■ **Figure 2** The figure depicts the finalization scenario of **case 2** with the 3 steps from top to bottom. We represent a checkpoint with a hexagon, a justified checkpoint with double hexagon, and a finalized checkpoint with a coloured double hexagon. The arrow between two checkpoints indicates a supermajority link.

These two cases illustrate the fact that for a checkpoint to become finalized, it needs to be the source of a supermajority link between justified checkpoints. Once a checkpoint is finalized, all the blocks leading to it (including the block in the pair constituting the checkpoint) become finalized. We now describe the condition for a checkpoint to be finalized more formally. Let (a, e_A) and (b, e_B) be two checkpoints such that $e_a < e_b$ and $e_b - e_a \leq 2$.¹¹ The checkpoint (a, e_A) is finalized if the following conditions are respected:

- Checkpoint (a, e_A) is justified.
 - There exists a supermajority link $(a, e_a) \xrightarrow{J} (b, e_b)$.
 - If $e_b - e_a = 2$ then the checkpoint in between at epoch $e_a + 1 (= e_b - 1)$ must be justified.
- The importance of the last condition is illustrated by the Figure 3. In practice, as explained in [7], for the checkpoint (A, e_A) to become finalized, it must not be more than 4 epochs old. At the implementation level, checkpoints more than 4 epochs old are not considered for finalization. This is illustrated by the last 4 conditions of Algorithm 9 in Subsection 4.2.



■ **Figure 3** This illustrate the case of two checkpoints respecting all the conditions but the one that stipulates that . We represent a checkpoint with a hexagon and a justified checkpoint with double hexagon. The arrow between two checkpoints indicates a supermajority link.


¹¹ This last condition necessitating the two checkpoints to be at most 2 epochs away from each other is also called *2-finality* [7].

4.1.5 Fork choice rule & Block proposition

The fork choice rule is the mechanism that allows each validator to determine the candidate chain depending on their view of the BlockTree and the state of checkpoints. Ethereum PoS fork choice rule is LMD GHOST. The LMD GHOST fork choice rule stems from the Greedy Heaviest-Observed Sub-Tree (GHOST) rule [25] which considers only each participant’s most recent vote (Latest Message Driven). During an epoch, each validator must make one *block vote* on the block considered as the head of the candidate chain according to its view. To determine the head of the candidate chain, the fork choice rule does the following:

1. Go through the list of validators and check the last block vote of each.
2. For each block vote, add a weight to each block of the chain that has the block voted as a descendent. The weight added is proportional to the stake of the corresponding validator.
3. Start from the block of the justified checkpoint with the highest epoch and continue the chain by following the block with the highest weight at each connection. Return the block without any child block. This block is the head of the candidate chain.

The actual implementation is written in Algorithm 7 in Subsection 4.2. This algorithm is similar to the one already presented in [7]. Albeit each *block vote* being for a specific block, the fork choice rule considers all the chains leading to that block. This reflects the fact that a vote for a block is a vote for the chain leading to that block. The Figure 4 offers an explanation with an example of how attestations influence the fork choice rule. At each chain intersection, the fork choice rule favors the chain with the most attestations.



■ **Figure 4** Fork choice rule example observed from a validator i 's point of view. We represent block votes with blue circles. Block votes point to specific blocks indicating the block considered as the *headblock* of the candidate chain at the moment of the vote. Each block has a number representing the value attributed by the fork choice rule algorithm (cf. Algorithm 7) to determine the candidate chain - we assume for this example that each validator has the same stake of 1. At the end of slot 4, i 's fork choice rule gives the block of slot 4 as C_i 's head. At the end of slot 5, the fork choice rule designates the block of slot 5 as the head of the candidate chain.

4.2 Pseudo Code

In this section, we dive into a practical understanding of the mechanism behind the Ethereum PoS protocol. According to the specifications [11] and various implementations (such as Prysm [22] and Teku [8]) we formalize the main functions of the protocol through pseudo-codes for a better understanding and for analysis purposes.

XX:10 Ethereum Proof-of-Stake under Scrutiny

Each validator p runs an instance of this particular pseudo code. For instance, when a validator p proposes a block, he broadcasts the following message: $\langle PROPOSE, (slot, \text{hash}(\text{headBlock}_p), \text{content}) \rangle$, where $slot$ is the slot at which the proposer proposes the block, the hash of the headBlock_p is the hash of the block considered to be the head of the candidate chain according to the fork choice rule (see Algorithm 7), and content contains data used for pseudo randomness, among other things that we will not detail here. We instead focus on the consensus protocol.

We describe in the following paragraphs the variables and functions used in the pseudo-code and the goal of these functions.

Variables.

During the computation, each variable takes a value that is subjective and may depend on the validator. We indicate with p the fact that the value of variables on each process. The variable tree_p is considered to be a graph of blocks with each block linked to its predecessor and thus represents the view of the blockchain (more precisely, tree_p represents the view of all blocks received by the validator since the genesis of the system). Each tree_p starts with the genesis block. role_p corresponds to the different roles a validator can have, which is possibly none (i.e., for each slot, the validator can be proposer, attester, or have no role). role_p is a list containing the role(s) of the validator for the current slot. The slot_p is a measure of time. In particular, a slot corresponds to 12 seconds. $\text{slot}_p \in \mathbb{N}$. Slot 0 begins at the time of the genesis block and is incremented every 12 seconds. headblock_p is the head of the candidate chain according to p 's local view and the fork choice rule. A checkpoint C is a pair block-epoch that is used for finalization. C has two attributes which are *justified* and *finalized* which can be true or false (e.g., if C is only justified $C.\text{justified} = \text{true}$ and $C.\text{finalized} = \text{false}$). $\text{lastJustifiedCheckpoint}_p$ is the *justified* checkpoint with the highest epoch. $\text{currentCheckpoint}_p$ is the checkpoint of the current epoch. The list attestation_p is a list of size n (i.e., the total number of validators). This list is updated only to contain the latest messages of validators (of at most one epoch old). CheckpointVote_p is a pair of checkpoints, so a pair of pairs, used to make a checkpoint vote. Let us stress out the fact that all those variables are local, and at any time, two different validators may have different valuations of those variables.

Functions.

We describe the main functions of the protocol succinctly before giving the pseudo code associated and a more detailed explanation:

- `validatorMain` is the primary function of the validator, which launches the execution of all subsidiary functions.
- `sync` is a function that runs in parallel of the `validatorMain` function and ensures the synchronization of the validator. It updates the slot, the role(s) and processes justification and finalization at the end of the epoch and when a new validator joins the system.
- `getHeadBlock` applies the fork choice rule. This is the function that indicates the head block of the candidate chain.
- `justificationFinalization` is the function that handles the justification and finalization of checkpoints.

We depict in Algorithm 1 the main of the validator. This function initializes all the values necessary to run a validator. In this paper, we consider the selection of validators already made to focus on the description of the consensus algorithm itself. The main starts

a routine called `sync` to run in parallel. Then there is an infinite loop that handles the call to an appropriate function when a validator needs to take action for its role(s).

■ **Algorithm 1** Main code for a validator p

```

1: procedure VALIDATORMAIN( )
2:    $tree_p := nil$  ▷ The tree represents the linked received blocks
3:    $role_p := []$  ▷  $role_p$  can be ROLE_PROPOSER and/or ROLE_ATTESTER when it
   is not empty
4:    $slot_p := 0$  ▷  $slot_p \in \mathbb{N}$ 
5:    $lastJustifiedCheckpoint_p := (0, genesisBlock)$  ▷ A checkpoint is a tuple (epoch,
   block)
6:    $attestation_p := []$  ▷ List of latest attestations received for each validator
7:    $validatorIndex_p := \text{index of the validator}$  ▷ Each validator has a unique index
8:    $listValidator := [p_0, p_1, \dots, p_{N-1}]$  ▷ A list of the validators index
9:    $balances := []$  ▷ A list of the balances of the validators, their stake
10:
11:  start sync( $tree_p, slot_p, attestation_p, lastJustifiedCheckpoint_p, role_p, balances$ )
12:
13:  while true do
14:    if  $role_p \neq \emptyset$  then
15:      if  $ROLE\_PROPOSER \in role_p$  then
16:        prepareBlock()
17:      if  $ROLE\_ATTESTER \in role_p$  then
18:        prepareAttestation()
19:       $role_p \leftarrow []$ 
20:    else
21:      no role assigned ▷ No action required

```

The roles performed by the validator when proposer or attester are defined in Algorithm 2 and Algorithm 3, respectively. The proposer of a block does the three following tasks:

1. Get the head of its candidate chain to have a block to build upon;
2. Sign a predefined pair to participate in the process of pseudo-randomness;
3. Broadcast a new block built on top of the head of the candidate chain.

The attestation is composed of three parts: the slot, the block vote, and the checkpoint vote. The validator uses the fork choice rule presented in Algorithm 7 to obtain the block chosen for the block vote. Algorithm 7 and the one stemming from it, Algorithm 8, have already been defined in [7]. We restated them here for the sake of completeness. For the checkpoint vote, an honest validator should always vote for the current epoch as the target and take the justified checkpoint with the highest epoch (i.e., $lastJustifiedCheckpoint$) as the source. In order to broadcast this attestation, the attester must wait for one of two things, either a block has been proposed for this slot or 1/3 of the slot (i.e., 2 seconds) has elapsed.

The synchronization of the validator p is handled by the function `sync` described in Algorithm 4. This algorithm allows the validator to update its view of the blockchain, in particular, the current slot, the list of attestations, the last justified checkpoint, the validator's role, and the balances of all validators. This function also starts two other routines which are `sync Block` and `sync Attestation` corresponding to Algorithm 5 and Algorithm 6, respectively. These routines are used to handle the broadcast of proposers and attesters. In

XX:12 Ethereum Proof-of-Stake under Scrutiny

■ Algorithm 2 broadcast block

```
1: procedure PREPAREBLOCK( )
2:    $headBlock_p := \text{getHeadBlock}( )$ 
3:    $randaoReveal := \text{sign}(slot, \text{epochOf}(slot))$ 
4:   broadcast  $\langle PROPOSE, (slot, \text{hash}(headBlock_p), randaoReveal, content) \rangle$ 
```

■ Algorithm 3 Broadcast Attestation

```
1: procedure PREPAREATTESTATION( )
2:   waitForBlockOrOneThird( )  $\triangleright$  wait for a new block in this slot or  $\frac{1}{3}$  of the slot
3:    $headBlock_p := \text{getHeadBlock}( )$ 
4:    $currentCheckpoint_p := (\text{first block of the epoch}, \text{epochOf}(slot))$ 
5:    $CheckpointVote_p := (\text{lastJustifiedCheckpoint}_p, currentCheckpoint_p)$ 
6:   broadcast  $\langle ATTEST, (slot_p, \underbrace{\text{hash}(headBlock_p)}_{\text{block vote}}, \underbrace{CheckpointVote_p}_{\text{checkpoint vote}}) \rangle$ 
```

both functions, upon receiving a block or an attestation, the validator p verifies that it is valid thanks to the `isValid` function. It is important to note that upon receiving a block, a validator can update the last justified checkpoint only if the current epoch has not started more than 8 slots ago. This particular condition is what the patch has brought to prevent a liveness attack (see Section 5).

■ Algorithm 4 Sync

```
1: procedure SYNC( $tree, slot, attestation, role, lastJustifiedCheckpoint,$ )
2:   start syncBlock( $slot, tree$ )
3:   start syncAttestation( $attestation$ )
4:   repeat
5:      $previousSlot \leftarrow slot$ 
6:      $slot \leftarrow \lfloor \text{time in seconds since genesis block} / 12 \rfloor$ 
7:     if  $previousSlot \neq slot$  then  $\triangleright$  If we start a new slot
8:        $roleSlotDone \leftarrow \text{false}$ 
9:       if  $\text{computeProposerIndex}(\text{getSeed}(\text{current epoch})) = validatorIndex$  then
10:         $role_p \leftarrow \text{ROLE\_PROPOSER}$ 
11:        if  $slot \pmod{32} = 0$  then  $\triangleright$  First slot of an epoch
12:          justificationFinalization( $tree, lastJustifiedCheckpoint$ )
13:          getCommitteeAssignment( $\text{next epoch}, validatorIndex, listValidator$ )
14:   until validator exit
```

Algorithm 9 can be considered the most intricate one. This algorithm is responsible for the justification or finalization of the checkpoints at the end of an epoch. To do so, it counts the number of checkpoint votes with the same source and target. If this number corresponds to more than $2/3$ of the stake of all validators then the target is considered justified for the validator running this algorithm. The last four conditions concern finalization. They verify among the last four checkpoints which one fulfills the conditions to become finalized. The conditions to become finalized are formally described in Subsection 4.1 and can be summarized by: the checkpoint must be the source of a supermajority link, and all the checkpoints between the source and target included must be justified.

Algorithm 5 Sync Block

```

1: procedure SYNCBLOCK( $slot, tree$ )
2:   upon  $\langle PROPOSE, (slot_i, \text{hash}(\text{headBlock}_i), \text{randaoReveal}_i, \text{content}_i) \rangle$  from
   validator  $i$  do
3:      $block := \langle PROPOSE, (slot_i, \text{hash}(\text{headBlock}_i), \text{randaoReveal}_i, \text{content}_i) \rangle$ 
4:     if isValid( $block$ ) then
5:       add  $block$  to  $tree$ 
6:       if  $slot \pmod{32} \leq 8$  then
7:         update justified checkpoint if necessary

```

Algorithm 6 Sync Attestation

```

1: procedure SYNCATTESTATION( $attestation$ )
2:   upon  $\langle ATTEST, (slot_i, \text{headBlock}_i, \text{checkpointEdge}_i) \rangle$  from validator  $i$  do
3:      $attestation_i := \langle ATTEST, (slot_i, \text{headBlock}_i, \text{checkpointEdge}_i) \rangle$ 
4:     if isValid( $attestation$ ) then
5:        $attestation[i] \leftarrow attestation_i$ 

```

5 Liveness attack

In this section, we describe a liveness attack called *bouncing attack* that delays finality in a partially synchronous network after GST. Previous works also exhibit liveness attacks against the protocol using the intertwining of the fork choice rule and the finality gadget [17, 19]. To prevent this type of attack, the protocol now contains a “patch” [23] that was suggested on the Ethereum research forum [18]. We show that the implemented patch is insufficient and this type of attack is still possible if certain conditions are verified. This is a probabilistic liveness attack against the protocol of Ethereum Proof-of-Stake. Our attack can happen with less than 1/3 of Byzantine validators, as discussed in Subsection 5.3.

5.1 Bouncing Attack

The *Bouncing Attack* [17] describes a liveness attack where the suffix of the chain changes repetitively between two candidate chains, thus preventing the chain from finalizing any checkpoint. The Bouncing Attack exploits the fact that the candidate chains should start from the justified checkpoint with the highest epoch. It is possible for Byzantine validators to divide honest validator’s opinion by justifying a new checkpoint once some honest validators have already cast their vote (made an attestation) during the asynchronous period before GST.

The bouncing attack becomes possible once there is a *justifiable* checkpoint in a different branch from the one designated by the fork choice rule with a higher epoch than the current highest justified checkpoint. A *justifiable checkpoint* is a checkpoint that can become justified only by adding the checkpoint votes of Byzantine validators. If this setup occurs, the malicious validators could make honest validators start voting for a different checkpoint on a different chain, leaving a justifiable checkpoint again for them to repeat their attack and thus making validators *bounce* between two different chains and not finalizing any checkpoint. Hence the name Bouncing attack.

Let us illustrate the attack with a concrete case. We show in Figure 5 an oversimplified case with only 10 validators, among which 3 are Byzantines. To occur, the attack needs to

XX:14 Ethereum Proof-of-Stake under Scrutiny

■ Algorithm 7 Get Head Block

```
1: procedure GETHEADBLOCK( )
2:    $block \leftarrow$  block of the justified checkpoint with the highest epoch
3:   while  $block$  has at least one child do
4:      $block \leftarrow \underset{b' \text{ child of } block}{\text{arg max}} \text{ weight}(tree, Attestation, b')$ 
5:     (ties are broken by hash of the block header)
6:   return  $block$ 
```

■ Algorithm 8 Weight

```
1: procedure WEIGHT( $tree, Attestation, block$ )
2:    $w \leftarrow 0$ 
3:   for every validator  $v_i$  do
4:     if  $\exists a \in Attestation$  an attestation of  $v_i$  for  $block$  or a descendant of  $block$  then
5:        $w \leftarrow w + \text{stake of } v_i$ 
6:   return  $w$ 
```

have a justifiable checkpoint with a higher epoch than the last justified checkpoint. We reach this situation before GST, which is presented in the left part of the figure. After reaching GST, Byzantine validators wait for honest validators to make a new checkpoint justifiable. When a new checkpoint is justifiable, the Byzantine validators cast their vote to justify another checkpoint, as shown by the right part of the figure. This will lead honest validators to start voting for the left branch, thus reaching a situation similar to the first step allowing the bouncing attack to continue. The repetition of this behavior is the bouncing attack. We emphasize this example in more detail in Figure 6 by detailing the sequence of votes allowing a “bounce” to occur and leaving a justifiable checkpoint on the other branch.

5.2 Implemented patch

The explication of the patch is described on the Ethereum research forum [18]. The solution that was found to mitigate the bouncing attack is to engrave in the protocol the fact that validators cannot change their minds regarding justified checkpoints after a part of the epoch has passed.

The goal of the solution proposed is to prevent the possibility of justifiable checkpoints being left out by honest validators. To prevent honest validators from leaving a justifiable checkpoint, the patch must stop validators from changing their view of checkpoints before more than $1/3$ of validators have cast their checkpoint vote. This condition stems from the fact that we reckon the proportion of Byzantine validators to be at most $1/3 - \epsilon$. To apply this condition, the patch designates a number of slots after which honest validators cannot change their view of checkpoints. Indeed since validators are scattered equally among the different slots to cast their vote (in attestations) in a specific time frame, stopping validators from changing their view after a number of slots is equivalent to stopping changing their view after a certain proportion of validators have voted. This does look like a solution to prevent Byzantine validators from influencing honest validators into forsaking a checkpoint now *justifiable* for them.

To enforce this behavior, called the “fixation of view”, the protocol has a constant j called `SAFE_SLOTS_TO_UPDATE_JUSTIFIED` in the code (cf. Algorithm 5 in Subsection 4.2). This

Algorithm 9 Justification and Finalization

```

1: procedure JUSTIFICATIONFINALIZATION(tree, lastJustifiedCheckpoint)
2:   source  $\leftarrow$  lastJustifiedCheckpoint
3:   target  $\leftarrow$  the current checkpoint
4:   nbCheckpointVote  $\leftarrow$  countMatchingCheckpointVote(source, target)
5:    $\triangleright$  justification process:
6:   if nbCheckpointVote  $\geq \frac{2}{3}$  * total balance of validators then
7:     target.justified = true
8:     lastJustifiedCheckpoint  $\leftarrow$  target
9:    $\triangleright$  finalization process:
10:  A, B, C, D  $\leftarrow$  the last 4 checkpoints  $\triangleright$  With D being the current checkpoint.
11:  if A.justified  $\wedge$  B.justified  $\wedge$  (A  $\xrightarrow{J}$  C) then
12:    D.finalized = true  $\triangleright$  Finalization of A
13:  if B.justified  $\wedge$  (B  $\xrightarrow{J}$  C) then
14:    B.finalized = true  $\triangleright$  Finalization of B
15:  if B.justified  $\wedge$  C.justified  $\wedge$  (B  $\xrightarrow{J}$  D) then
16:    B.finalized = true  $\triangleright$  Finalization of B
17:  if C.justified  $\wedge$  (C  $\xrightarrow{J}$  D) then
18:    C.finalized = true  $\triangleright$  Finalization of C

```

constant is the number of slots¹² until when validators can change their view of the justified checkpoints. In the patch introducing this constant j , they mention a possible attack called *splitting attack*. As they point out, the splitting attack relies on a “last minute delivery” family of strategies whereby releasing a message late enough, some validators will consider it too late while others will not. This could thus split the validators into two different chains, not being able to conciliate their view before the end of the epoch. They consider the assumption of attackers being able to send a message at the right time to split honest validators too strong, we discuss in Section 6 whether it is the case or not. In 5.3, we present a new attack inspired by the splitting attack with more realistic assumptions.


5.3 Probabilistic Bouncing attack - why the patch is not enough

In this part, we present our novel attack against the protocol of Ethereum Proof-of-Stake.

5.3.1 Attack Conditions

Our attack takes place during the synchronous period and uses the power of *equivocation* of Byzantine processes. Equivocation is caused by a Byzantine process that sends a message only to a subset of validators at a given point in time and potentially another message or none to another subset of validators. The effect is that only a part of the validators will receive the message on time. More in detail, the bounded network delay is used by a Byzantine validator to convey a message to be read on a specific slot by some validators and read on the next slot by the other validators. Note that if a protocol is not tolerant to equivocation, then it is not BFT (Byzantine-Fault Tolerant), since equivocation is the typical action possible for Byzantine validators.

¹²At the time of writing this paper, $j = 8$ [11].



■ **Figure 5** A bouncing attack presented in 3 steps. We have 10 validators, of which 3 are Byzantines. Checkpoints are represented by hexagonal shapes. A double hexagon signifies that the checkpoint is justified. The number inside each hexagon corresponds to the number of validators who made a checkpoint vote with this checkpoint as target. **1st step:** We start in a situation where two checkpoints are justified ($7 > \frac{2n}{3}$) and there are two concurrent chain. We are at the end of the third epoch in which honest validators have divided their vote on each side. **2nd step:** We have reached GST at the beginning of the fourth epoch and 4 honest validators have already voted (rightfully so). **3rd step:** Here is the moment Byzantine validators take action and release their checkpoint vote for the concurrent chain, thus justifying the previously forsaken checkpoint and thereby changing the highest justifying checkpoint. By repeating this process, the bouncing attack can continue indefinitely.


5.3.2 Attack Description and Analysis

Let $\beta \leq f/n$ be the fraction of Byzantine validators in the system. We start by explaining the attack on the Ethereum PoS protocol before generalizing our analysis. The attack setup is the following. First, as in the traditional bouncing attack, we start in a situation where the network is still partially synchronous. A fork occurs and results in the highest justified checkpoint being on chain A at epoch e , and a justifiable checkpoint at epoch $e + 1$ on chain B . Assume now that GST is reached, the attack can proceed¹³ as follows:

1. Since GST is reached, the network is fully synchronous. Chain A is the candidate chain for all validators.
2. Just before validators must stop updating their view concerning justified checkpoint (i.e., before reaching the limit of 8 slots in the epoch corresponding to the condition line 6 in Algorithm 5), a Byzantine proposer proposes a block (cf. Algorithm 2) on chain B . This block contains attestation with enough checkpoint votes to justify the justifiable checkpoint that was left by honest validators. The attestations included in the block are attestations of Byzantine validators that were not issued in the previous epoch when they were supposed to. The block must be released just in time, that is, right before the end of slot j , so that $(2/3 - \beta)$ of the validators change their view of the candidate chain to be active on chain B while the rest of honest validator continue on chain B . This is possible due to the patch preventing validators from changing their mind after j slots.
3. Repeat the process.

An important aspect to consider in the attack is the probability for Byzantine validators to become proposers. This is an important part since without the role of proposer, validators are

¹³Note that before GST, no algorithm can ensure liveness since communication delays may not be bounded.



■ **Figure 6** This figure presents a detailed version of the bouncing attack. In this example, we have a total of 100 validators, of which 23 are Byzantines. Checkpoints are represented by hexagonal shapes. A double hexagon signifies that the checkpoint is justified. A block in a checkpoint corresponds to the block associated with that checkpoint. The number inside each hexagon (hovering a block) corresponds to the number of validators who made a checkpoint vote with this checkpoint as target. We distinguish between two sorts of checkpoint votes, the Byzantine ones, which are bi-color rectangles, and the honest ones, which are uni-color rectangles. We compile the 3 steps of Figure 5 in 2 with more information on how justification's turning point is accomplished because of the Byzantine agents. **First step:** We begin from a situation where epoch $e - 1$ just ended and we now reach GST. Notice that the candidate chain is chain A because the checkpoint with the highest epoch is on chain A but not chain B. **Second step:** In this step, the checkpoint vote released during epoch e can change the last justified checkpoint to change the candidate chain for chain A to chain B. Byzantine validators released their checkpoint vote from the previous epoch during epoch e . They send their last checkpoint vote at slot 23 once the checkpoint of epoch e on chain A has reached 44, thus becoming justifiable (i.e., not yet justified but with enough votes so that Byzantine validators can justify it). This trigger the candidate chain to change from chain A to chain B starting the *bounce*.

not legitimate to propose blocks and thus cannot add new attestations containing checkpoint votes on top of the concurrent chain.¹⁴ The probability of being selected to be a proposer directly impacts how long the bouncing attack can continue. In the following theorem, we establish the probability of a bouncing attack lasting for a specific number of epochs.

► **Theorem 7.** *The probability of a bouncing attack occurring during k epochs after GST and a favorable setting is:*

$$P(\text{bouncing } k \text{ times}) = (1 - \alpha^j)^k, \quad (1)$$

with $\alpha \in [0, 1]$ the proportion of honest validators and j the number of slots before locking a choice for the justification.

Proof. We denote by α the proportion of honest validators and j the number of slots before locking the choice for justification. We want to know the probability of delaying the finality for k epochs. Once we assume a setup condition sufficient to start a bouncing attack, the attack continues until it becomes impossible for Byzantine validators to cast a vote to justify the justifiable checkpoint. To cast their vote, Byzantine validators need one of the j first slots of the concurrent chain to have a Byzantine validator as proposer. Considering the probability of choosing between each validator, the chance for a Byzantine validator to be a proposer for one of the first j slots is $(1 - \alpha^j)$, with α being the proportion of honest validators. For k epochs, we take this result to the power of k . ◀

As we can see, the probability of the bouncing attack to continue for k epochs depends on two factors: α the proportions of honest validators which cannot be controlled and j the number of slots before which validators are allowed to switch branches. Reducing j to 0 would prevent the bouncing attack from happening (probability falls to 0), but it would mean that validators are never allowed to change their view of the candidate chain. This naive solution would allow irreconcilable choices between the set of validators and prevent any new checkpoint from being justified, which is a more severe threat to the liveness of Ethereum PoS.

Reducing the number of slots where validators can change their view of the blockchain implies that different views cannot reconcile quickly. Or at least, the window of opportunity for doing so gets smaller. In theory, the proportion of Byzantine validators necessary to perform this attack is $1/n$. This is because we assume a favorable setup and that Byzantine validators can send messages so that only a wished portion of honest validators receives it on time. Our analysis focuses on the course of action of the attackers during the attack rather than the course of action necessary for it to appear.

6 Related Works

Blockchain protocol analyses can be divided into two main categories: those that specify or formalize protocols and those that identify vulnerabilities of the protocols. Our work is at the junction of the two categories because we both formalize the protocol to permit its analysis, and we present a novel attack.

¹⁴Note that Byzantine validators can not use their role of proposer during the previous epoch to release a block with the right attestations because it might not be the last block of the epoch. Indeed because a part of honest validators is on the concurrent chain, they also add blocks to it. For the checkpoint votes contained in the Byzantine attestations to justify the justifiable checkpoint, it must be on the same chain as the attestation making the checkpoint justifiable in the first place.

The category of specification and formalization includes the “white papers” (e.g., Bitcoin [16], Ethereum [26]). It also contains academic papers providing formal specifications and demonstrating the properties guaranteed by the protocols. For example, [12] formally describes, analyzes the Bitcoin protocol, and proves its security guarantees, [1] does the same for the Tendermint’s protocol (the consensus protocol of the Cosmos blockchain[6]). Our work lies in this category of formalization by formalizing the Ethereum Proof-of-Stake protocol. For what concerns protocol formalization of Ethereum Proof-of-Stake, [7] is the first to propose a draft specification of the Ethereum PoS protocol and related properties. However, that specification is outdated and not complete. Our paper provides a complete formalization of the consensus mechanisms with respect to the current implemented code, along with a novel specification of its properties.

A famous example of papers identifying vulnerabilities is [10], which presents the selfish mining attack on Bitcoin. [10] shows that in Bitcoin (and proof-of-work in general), miners can benefit from deviating from the prescribed protocol by withholding blocks for a while at the expense of honest miners.

[2] points out a liveness vulnerability of the Tendermint protocol. [21] presents an attack where nodes can reorganize the Tezos’ Emmy+ chain and then do a double-spend attack. Our work follows the line of research focusing on flaws of Ethereum Proof-of-Stake [19, 20, 24]. Neu et al. [19] exhibit a balancing attack, highlighting the shortcomings of a consensus mechanism being separated into two layers (finality gadget, fork choice rule). Mitigation against this attack was proposed, but [20] overcame this mitigation with a new balancing attack. [24] presented reorg attacks revealing that validators with the role of proposer could gain from disturbing the protocol by releasing their block late. Our paper presents another flaw regarding the liveness of the current Ethereum Proof-of-stake protocol, on which some attacks do not seem feasible anymore, and thus insists on the importance of finding new ways to conciliate availability and finality.

Nakamura [18] presents an attack called *splitting attack* in which the adversary sends messages to split the set of validators. However, to achieve this attack, Nakamura assumes that the adversary needs to control and play with network delays. This is a strong assumption and can be considered unrealistic. More recently, [24] showed through experiments that attackers can predict the proportion of validators receiving a given message within a specific time frame with sufficient accuracy. This contradicts Nakamura’s claim that the attack necessitates the adversary to control the network delay. In this work, we present a form of the splitting attack based on the weaker assumption that the adversary knows the network delay (in line with [24]) but does not control it. Moreover, our attack is repeated, hence the name bouncing, being a threat to the liveness.

Outside of these two categories lie works that provide formal ground for blockchains. [4] describes the types of finality a blockchain can achieve, [3] proposes a formalization of blockchains and their evolutions as BlockTrees. We rely on the definition of BlockTree and finality to express the Ethereum protocol properties.

7 Conclusion

We described a framework for formal analysis of the Ethereum Proof-of-Stake (PoS) protocol. Our contribution provides a formal description of Ethereum PoS, which can be used for future analysis and comparison with future versions of the protocol. We proposed a novel distinction between the definition of liveness and availability. This distinction is crucial to pinpoint the difference between Nakamoto-style and BFT consensus. It makes possible a

comparison between the two. We outlined an attack against the liveness of the protocol, showing probabilistic liveness to the protocol under this attack. This work explores whether it is possible to combine Nakamoto-style and BFT to ensure safety and liveness. Another aspect is that our analysis did not consider the protocol's rewards and incentives. We leave this part, as well as analyzing rational behavior in the protocol, as future work.

References

- 1 Yackolley Amoussou-Guenou, Antonella Del Pozzo, Maria Potop-Butucaru, and Sara Tucci Piergiovanni. Dissecting tendermint. *Networked Systems*, pages 166–182, 2019.
- 2 Yackolley Amoussou-Guenou, Antonella Del Pozzo, Maria Potop-Butucaru, and Sara Tucci Piergiovanni. Correctness of tendermint-core blockchains. In *22nd International Conference on Principles of Distributed Systems, OPODIS 2018, December 17-19, 2018, Hong Kong, China*, pages 16:1–16:16, 2018.
- 3 Emmanuelle Anceaume, Antonella Del Pozzo, Romaric Ludinard, Maria Potop-Butucaru, and Sara Tucci-Piergiovanni. Blockchain abstract data type. *Cryptology ePrint Archive*, Paper 2018/561, 2018.
- 4 Emmanuelle Anceaume, Antonella Del Pozzo, Thibault Rieutord, and Sara Tucci Piergiovanni. On finality in blockchains. In *25th International Conference on Principles of Distributed Systems, OPODIS 2021, December 13-15, 2021, Strasbourg, France*, pages 6:1–6:19, 2021.
- 5 Lacramioara Astefanoaei, Pierre Chambart, Antonella Del Pozzo, Thibault Rieutord, Sara Tucci Piergiovanni, and Eugen Zalinescu. Tenderbake - A solution to dynamic repeated consensus for blockchains. In *4th International Symposium on Foundations and Applications of Blockchain 2021, FAB 2021, May 7, 2021, University of California, Davis, California, USA (Virtual Conference)*, pages 1:1–1:23, 2021.
- 6 Ethan Buchman, Jae Kwon, and Zarko Milosevic. The latest gossip on bft consensus, 2018.
- 7 Vitalik Buterin, Diego Hernandez, Thor Kamphofner, Khiem Pham, Zhi Qiao, Danny Ryan, Juhyeok Sin, Ying Wang, and Yan X Zhang. Combining ghost and casper, 2020.
- 8 Consensys. Teku consensus client, 2022. URL: <https://github.com/ConsenSys/teku/tree/bada6df06edd1f8b5d727a011f440d9825e17d99>.
- 9 Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, pages 288–323, 1988.
- 10 Ittay Eyal and Emin Sirer. Majority is not enough: Bitcoin mining is vulnerable. *Communications of the ACM*, pages 95–102, 2018.
- 11 Ethereum Foundation. Consensus specifications github, 2022. URL: <https://github.com/ethereum/consensus-specs/tree/ae89e4e6158344a5ab736d834e2239efd431ef5f/specs>.
- 12 Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In *Advances in Cryptology - EUROCRYPT 2015*, pages 281–310, 2015.
- 13 Seth Gilbert and Nancy Lynch. Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, page 51–59, 2002.
- 14 L.M. Goodman. Tezos — a self-amending crypto-ledger, 2014.
- 15 Jae Kwon and Ethan Buchman. Cosmos, 2016.
- 16 Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. 2008.
- 17 Ryuya Nakamura. Analysis of bouncing attack on ffg. URL: <https://ethresear.ch/t/analysis-of-bouncing-attack-on-ffg/6113>.
- 18 Ryuya Nakamura. Prevention of bouncing attack on ffg. URL: <https://ethresear.ch/t/prevention-of-bouncing-attack-on-ffg/6114>.
- 19 Joachim Neu, Ertem Nusret Tas, and David Tse. Ebb-and-flow protocols: A resolution of the availability-finality dilemma. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 446–465, 2021.
- 20 Joachim Neu, Ertem Nusret Tas, and David Tse. Two attacks on proof-of-stake ghost/ethereum. 2022.

- 21 Michael Neuder, Daniel J. Moroz, Rithvik Rao, and David C. Parkes. Defending against malicious reorgs in tezos proof-of-stake. In *AFT '20: 2nd ACM Conference on Advances in Financial Technologies, New York, NY, USA, October 21-23, 2020*, pages 46–58, 2020.
- 22 Prysm. Code consensus client, 2022. URL: <https://github.com/prysmaticlabs/prysm/tree/cafe0bd1f81298ee34b450ab8d7e74b0036e9803>.
- 23 Specification Pull Request. Bouncing attack patch, 2019. URL: <https://github.com/ethereum/consensus-specs/pull/1465>.
- 24 Caspar Schwarz-Schilling, Joachim Neu, Barnabé Monnot, Aditya Asgaonkar, Ertem Nusret Tas, and David Tse. Three attacks on proof-of-stake ethereum. 2021.
- 25 Yonatan Sompolinsky and Aviv Zohar. Secure high-rate transaction processing in bitcoin. In *Financial Cryptography and Data Security - 19th International Conference, FC 2015, San Juan, Puerto Rico, January 26-30, 2015, Revised Selected Papers*, 2015.
- 26 Gavin Wood. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, pages 1–32, 2014.