

A Simple Single Slot Finality Protocol For Ethereum

Francesco D’Amato
Ethereum Foundation
francesco.damato@ethereum.org

Luca Zanolini
Ethereum Foundation
luca.zanolini@ethereum.org

Abstract

Currently, Gasper, the implemented consensus protocol of Ethereum, takes between 64 and 95 slots to finalize blocks. Because of that, a significant portion of the chain is susceptible to reorgs. The possibility to capture MEV (Maximum Extractable Value) through such reorgs can then disincentivize honestly following the protocol, breaking the desired correspondence of honest and rational behavior. Moreover, the relatively long time to finality forces users to choose between economic security and faster transaction confirmation. This motivates the study of the so-called single slot finality protocols: consensus protocols that finalize a block in each slot and, more importantly, that finalize the block proposed at a given slot within such slot.

In this work we propose a *simple, non-blackbox* protocol that combines a synchronous dynamically available protocol with a partially synchronous finality gadget, resulting in a consensus protocol that can finalize one block per slot, paving the way to *single slot finality* within Ethereum. Importantly, the protocol we present can finalize the block proposed in a slot, within such slot.

1 Introduction

Traditional Byzantine consensus protocols, such as PBFT [5], are devised in a partial synchronous network model [8], in the sense that they always guarantee safety, but they guarantee liveness only after GST. In this setting, however, participants in the protocol are fixed, known in advance, and without possibility to go *offline*.

Dynamic participation (among systems’ participants) has lately become an essential prerequisite for developing permissionless consensus protocols. This concept, initially formalized by Pass and Shi via their *sleepy model*, [16] encapsulates the ability of a system to handle participants joining or leaving during a protocol execution. In particular, a consensus protocol that preserves safety and liveness while allowing dynamic participation is called *dynamically available*.

One problem of such protocols, as a result of the CAP theorem [9][11], is that they do not tolerate network partitions; no consensus protocols can both satisfy liveness (under dynamic participation) and safety (under temporary network partitions). Simply put, a consensus protocol (for state-machine replication) cannot produce a single chain that concurrently offers dynamic availability and guarantees transaction finality in case of asynchronous periods or network partitions. Because of that, dynamically available protocols studied so far are focused on a synchronous model [6][12][13].

To overcome this impossibility result, Neu *et al.* [15] introduce a family of protocols, referred to as *ebb-and-flow* protocols, which operate under two confirmation rules, and outputting two chains, one a prefix of the other. The first confirmation rule defines what is known as the *available chain*, which provides liveness under dynamic participation (and synchrony). The second confirmation rule defines the *finalized chain*, and provides safety even under network partitions. Interestingly, such family of protocols also captures the nature of the Ethereum consensus protocol, Gasper [4], in which the available chain is output by (the confirmation rule of) LMD-GHOST [18], and the finalized chain by the (confirmation rule of the) *finality gadget* Casper FFG [3]. However, the (original version of) LMD-GHOST is actually not secure [15] even in a context of full-participation.

Motivated by finding a (more secure) alternative to LMD-GHOST, and following the ebb-and-flow approach, D’Amato *et al.* [6] devise a synchronous dynamically available consensus protocol, *Goldfish*, that,

combined with a generic (partially synchronous) finality gadget, implements an ebb-and-flow protocol. Moreover, *Goldfish* is reorg resilient: blocks proposed by honest validators are guaranteed inclusion in the chain. However, *Goldfish* is brittle to temporary asynchrony [7], in the sense that even a single violation of the bound of network delay can lead to a catastrophic failure, jeopardizing the safety of *any* previously confirmed block, resulting in a protocol that is not practically viable to replace LMD-GHOST in Ethereum. In other words, *Goldfish* is not *asynchrony resilient*.

To cope with the problem of *Goldfish*, D’Amato and Zanolini [7] propose RLMD-GHOST, a provably secure synchronous protocol that does not lose safety during *bounded* periods of asynchrony and which tolerates a weaker form of dynamic participation, offering a trade-off between dynamic availability and asynchrony resilience. Their protocol results appealing for practical systems, where strict synchrony assumptions might not always hold, contrary to what is generally assumed with standard synchronous protocols.

In this work we build upon the work of D’Amato and Zanolini [7], and we devise a protocol that combines RLMD-GHOST with a partially synchronous finality gadget. In particular, we give the following contributions. We devise a secure and reorg-resilient ebb-and-flow protocol [15] as a potential substitute for the current Ethereum consensus protocol, Gasper [4], which can finalize (at most) one block per slot. In particular, our protocol can finalize the block proposed in the current slot, within such slot, paving the way to *single slot finality* [2] protocols for practical use within Ethereum. Finally, we expand upon the *generalized sleepy model* [7] introduced by D’Amato and Zanolini[7], adjusting it to accommodate a partially synchronous setting. We refer to the resulting model as the *generalized partially synchronous sleepy model*. This enhanced model not only extends the original sleepy model, first presented by Pass and Shi [16], but it also introduces stronger and more generalized constraints related to the corruption and sleepiness power of the adversary. Furthermore, our model integrates the concept of partial synchrony, setting it apart from the model proposed by D’Amato and Zanolini [7]. Our security results will be proven within this extended model. The remainder of this work is structured as it follows. In Section 3 we present our system model. Prerequisites for this work are presented in Section 4; we recall RLMD-GHOST as originally presented by D’Amato and Zanolini [7], state its properties, and show a class of protocols, called *propose-vote-merge* protocols, that groups together (a variant of) LMD-GHOST, (a variant of) *Goldfish*, and RLMD-GHOST under an unique framework. Protocol specifications are described in Section 5. In particular, we show how to slightly modify RLMD-GHOST to interact with a finality gadget, and then present the full protocol. In Section 6 we formally prove the properties that our protocol satisfy. Finally, in Section 7 we enable our protocol to finalize the block proposed in the current slot through *acknowledgments*, messages sent by participants in the consensus protocol, but only relevant to external observers. Conclusions are drawn in Section 8.

2 Related works

Pass and Shi [16] introduced the *sleepy model of consensus*, which models a distributed system where the participants can be either online or offline, meaning their participation is dynamic. This differs from the standard models in the literature that assume honest participants are always online and execute the assigned protocol. Dynamic participation became a key requirement to devise consensus protocols, as it adds a more robustness to systems that allow participants to go offline, while preserving safety and liveness of such *dynamically available* protocols.

Neu et al [15] introduce the *partially synchronous sleepy model* and define the objectives of the Ethereum consensus protocol, Gasper [4], through the concept of an *ebb-and-flow protocol*. A secure ebb-and-flow protocol produces both a dynamically available ledger and a finalized ledger, that is always safe and live after $\max\{\text{GST}, \text{GAT}\}$. In the context of Gasper, the dynamically available ledger is defined by LMD-GHOST [18] and the finalized ledger by Casper [3].

However, under a deeper analysis, Neu *et al* [15] show that LMD-GHOST is not dynamically available, by presenting an attack to its liveness. D’Amato *et al.* [6] introduce *Goldfish*, a simplified variant of LMD-GHOST, aiming at solving some problems related to LMD-GHOST [15, 14], that results in a synchronous dynamically available protocol in the partially synchronous sleepy model that, composed with a generic finality gadget, implements an ebb-and-flow protocol. *Goldfish* however is brittle to temporary asynchrony, in the sense that even a single violation of the bound of network delay can lead to a catastrophic failure, jeopardizing the safety of *any* previously confirmed block.

D’Amato and Zanolini [7] introduce the *generalized sleepy model*. This model takes up from the original

sleepy model presented by Pass and Shi [16] and extends it with more generalized and stronger constraints in the corruption and sleepiness power of the adversary. This allow to explore a broad space of dynamic participation regimes which fall between complete dynamic participation and no dynamic participation. Moreover, they introduce RLMD-GHOST, a generalization of (variants of) Goldfish and LMD-GHOST, that offers a trade-off between resilience to temporary asynchrony and dynamic availability. RLMD-GHOST represents a middle ground between LMD-GHOST, an asynchrony resilient but not dynamically available protocol, and Goldfish, a dynamically available but not asynchrony resilient protocol. RLMD-GHOST is resilient to bounded asynchrony *up to a vote expiry period*, and satisfies an appropriate notion of dynamic availability in the generalized sleepy model.

3 Model and Preliminary Notions

3.1 System model

We consider a set of n *validators* v_1, \dots, v_n that communicate with each other through exchanging messages. Every validator is identified by a unique cryptographic identity and the public keys are common knowledge. Validators are assigned a protocol to follow, consisting of a collection of programs with instructions for all validators. A validator that follows its protocol during an execution is called *honest*. Each validator has a *stake*, which we assume to be the same for every validator. If a validator v_i fails to serve the role assigned to it or tries to deliberately deviate from the protocol, i.e., v_i is *Byzantine*, and a proof of this misbehavior is given, it loses a part of its stake proportional to the severity of the fault (v_i gets *slashed*). We assume the existence of a probabilistic poly-time adversary \mathcal{A} that can choose up to f validators to corrupt over an entire protocol execution. Corrupted validators stay corrupted for the remaining duration of the protocol execution, and are thereafter called *adversarial*. The adversary \mathcal{A} knows the the internal state of adversarial validators. The adversary is *adaptive*: it chooses the corruption schedule dynamically, during the protocol execution.

We assume that a best-effort gossip primitive that will reach all validators is available. In a protocol, this primitive is accessed through the events “sending a message through gossip” and “receiving a gossiped message.” Moreover, we assume that messages from honest validator to honest validator are eventually received and cannot be forged. This includes messages sent by Byzantine validators, once they have been received by some honest validator v_i and gossiped around by v_i .

Time is divided into discrete *rounds*. We consider a partially synchronous model in which validators have synchronized clocks but there is no a priori bound on message delays. However, there is a time (not known by the validators), called *global stabilization time* (GST), after which message delays are bounded by Δ rounds. Moreover, we define the notion of *slot* as a collection of 4Δ rounds. The adversary \mathcal{A} can decide for each round which honest validator is *awake* or *asleep* at that round [16]. Asleep validators do not execute the protocol and messages for that round are queued and delivered in the first round in which the validator is awake again. Honest validators that become awake at round r , before starting to participate in the protocol, must first execute (and terminate) a *joining protocol* (Section 4), after which they become *active*. All adversarial validators are always awake, and are not prescribed to follow any protocol. Therefore, we always use active, awake, and asleep to refer to honest validators. As for corruptions, the adversary is adaptive also for sleepiness, i.e., the sleepiness schedule is also chosen dynamically by the adversary. Moreover, there is a time (not known by the validators), called *global awake time* (GAT), after which all validators are always awake.

We assume that every message has an *expiration period* η [6][7]. More specifically, for a given slot t and a constant $\eta \in \mathbb{N}$ greater than or equal to 0, the *expiration period* for slot t is the interval $[t - \eta, t - 1]$. Only messages sent within this time frame influence the behavior of the protocol at slot t . Furthermore, during each protocol execution slot, only the most recent messages sent by validators are considered.

We require that, for some fixed parameter $1 \leq \tau \leq \infty$, the following condition, referred by D’Amato and Zanolini [7] as τ -*sleepiness at slot t* , holds for any slot t after GST:

$$|H_{t-1}| > |A_t \cup (H_{t-\tau, t-2} \setminus H_{t-1})| \tag{1}$$

with H_t , A_t , and $H_{s,t}$ are the set of active validators at round $4\Delta t + \Delta$, the set of adversarial validators at round $4\Delta t + \Delta$, and the set of validators that are active *at some point* in slots $[s, t]$, i.e., $H_{s,t} = \bigcup_{i=s}^t H_i$

(if $i < 0$ then $H_i := \emptyset$), respectively. Note that $f = \lim_{t \rightarrow \infty} |A_t|$. In other words, we require the number of active validators at round $4\Delta(t-1) + \Delta$ to be greater than the number of adversarial validators at round $4\Delta t + \Delta$, together with the number of validators that were active at some point between rounds $4\Delta(t-\tau) + \Delta$ and $4\Delta(t-2) + \Delta$, but not at round $4\Delta(t-1) + \Delta$.

Intuitively, this condition is designed to work with a protocol that applies expiration to its messages, with the period set as $\eta = \tau$. The messages taken into consideration at slot t originate from slots $[t - \tau, t - 1]$. Among these, the only messages sent by honest validators that can be relied upon come from H_{t-1} . However, unexpired messages from honest validators, who were inactive in slot $t-1$, could potentially aid the adversary.

Note that our approach diverges from the *generalized sleepy model* proposed by D’Amato and Zanolini [7]. Specifically, we require that Equation 1 only holds after GST and we refer to this model as the *generalized partially synchronous τ -sleepy model* (or wlog, when the context is clear, as the *τ -sleepy model* for short). Finally, we say that an execution in the generalized partially synchronous sleepy model is *τ -compliant* if it satisfies τ -sleepiness (Equation 1).

3.2 Validator internals

View. A *view* (at a given round r), denoted by \mathcal{V} , is a subset of all the messages that a validator has received until r . The notion of view is *local* for the validators. For this reason, when we want to focus the attention on a specific view of a validator v_i , we denote with \mathcal{V}_i the view of v_i (at a round r).

Blocks and chains. Let’s consider two chains, ch_1 and ch_2 . We denote $\text{ch}_1 \prec \text{ch}_2$ if ch_1 acts as a prefix to ch_2 . When block B is at the end of chain ch , we refer to it as the *head of ch* , and we equate the entire chain with B . Therefore, if $\text{ch}' \prec \text{ch}$ and A is the head of ch' , we also express this as $\text{ch}' \prec B$ and $A \prec B$.

Fork-choice functions. A *fork-choice function* is a deterministic function, denoted as FC. This function, when given a view \mathcal{V} and a slot t as inputs, produces a block B . If B is a block extending $\text{FC}(\mathcal{V}, t)$, then $\text{FC}(\mathcal{V} \cup B, t)$ equals B . The result of FC is referred to as the *head of the canonical chain* in \mathcal{V} , and the chain with B as its head is referred to as the *canonical chain* in \mathcal{V} . Every validator keeps track of its canonical chain and updates it using FC, according to its local view. The canonical chain for validator v_i at round r is represented as ch_i^r . In this work we will focus our attention on a specific class of fork-choice functions based on GHOST [17]. D’Amato and Zanolini [7] characterize a GHOST-based fork-choice function by a view filter FIL, which takes as input a view \mathcal{V} and a slot t , and outputs (\mathcal{V}', t) , where \mathcal{V}' is another view such that $\mathcal{V}' \subseteq \mathcal{V}$. Then, $\text{FC}(\mathcal{V}, t) := \text{GHOST}(\text{FIL}(\mathcal{V}, t))$, i.e., $\text{FC} := \text{GHOST} \circ \text{FIL}$.

3.3 Security

Security Parameters. In this work we treat λ and κ as the security parameters related to the cryptographic components utilized by the protocol and the protocol’s own security parameter, respectively. We also account for a finite time horizon, represented as T_{hor} , which is polynomial in relation to κ . An event is said to occur with *overwhelming probability* if it happens except with probability which is $\text{negl}(\kappa) + \text{negl}(\lambda)$. The properties of cryptographic primitives hold true with a probability of $\text{negl}(\lambda)$, signifying an overwhelming probability, although we will not explicitly mention this in the subsequent sections of this work.

Confirmed chain. The protocols we consider always specify a *confirmation rule*, with whom validators can identify a *confirmed prefix* of the canonical chain. Alongside the canonical chain, validators then also keep track of a *confirmed chain*. We refer to the confirmed chain of validator v_i at round r as Ch_i^r (cf. ch_i^r for the canonical chain).

Definition 1 (Secure protocol [6]). We say that a protocol outputting a confirmed chain Ch is *secure* after time T_{sec} , and has confirmation time T_{conf}^1 , if Ch satisfies:

¹If the protocol satisfies liveness, then at least one honest proposal is added to the confirmed chain of all active validators every T_{conf} slots. Since honest validators include all transactions they see, this ensures that transactions are confirmed within time $T_{\text{conf}} + \Delta$ (assuming infinite block sizes or manageable transaction volume).

Safety For any two rounds $r, r' \geq T_{\text{sec}}$, and any two honest validators v_i and v_j (possibly $i = j$) at rounds r and r' respectively, either $\text{Ch}_i^r \prec \text{Ch}_j^{r'}$ or $\text{Ch}_j^{r'} \prec \text{Ch}_i^r$.

Liveness For any rounds $r \geq T_{\text{sec}}$ and $r' \geq r + T_{\text{conf}}$, and any honest validator v_i active at round r' , $\text{Ch}_i^{r'}$ contains a block proposed by an honest validator at a round $> r$.

A protocol satisfies τ -*safety* and τ -*liveness* if it satisfies safety and liveness, respectively, in the τ -*sleepy model*, i.e., in τ -compliant executions. A protocol satisfies τ -*security* if it satisfies τ -safety and τ -liveness.

We now recall the definitions of *dynamic availability* and *reorg resilience* from [7]. We consider them only under network synchrony, i.e., for $\text{GST} = 0$, as this is the only setting in which we utilize them. Note that it is customary to only analyze dynamic availability with $\text{GST} = 0$, when analyzing the behavior of ebb-and-flow protocols.

Definition 2 (Dynamic availability). We say that a protocol is τ -*dynamically-available* if and only if it satisfies τ -security with confirmation time $T_{\text{conf}} = O(\kappa)$ when $\text{GST} = 0$. Moreover, we say that a protocol is *dynamically available* if it is 1-dynamically-available, as this corresponds to the usual notion of dynamic availability.

Definition 3 (Reorg resilience). An execution with $\text{GST} = 0$ satisfies *reorg resilience* if any honest proposal B from a slot t is always in the canonical chain of all active validators at rounds $\geq 4\Delta t + \Delta$. A protocol is τ -*reorg-resilient* if all τ -compliant executions with $\text{GST} = 0$ satisfy reorg resilience.

Definition 4 (Accountable safety). We say that a protocol has *accountable safety* with resilience $f > 0$ if, upon a safety violation, it is possible to identify at least f responsible participants. In particular, it is possible to collect evidence from sufficiently many honest participants and generate a cryptographic proof that identifies f adversarial participants as protocol violators. Such proof cannot falsely accuse any honest participant that followed the protocol correctly. Finally, we also say that a chain is f -*accountable* if the protocol outputting it has accountable safety with resilience f . If a protocol Π outputs multiple chains $\text{Ch}_1, \dots, \text{Ch}_k$, we say that Ch_i is f -accountable if Π_i is, where Π_i is the protocol which runs Π and outputs only Ch_i .

Ebb-and-flow protocols. Neu *et al.* [15] propose a protocol with two confirmation rules that outputs two chains, one that provides liveness under dynamic participation (and synchrony), and one that provides accountable safety even under network partitions. This protocol is called *ebb-and-flow* protocol. We present a generalization of it, in the τ -sleepy model.

Definition 5 (τ -secure ebb-and-flow protocol). A τ -secure *ebb-and-flow protocol* outputs an available chain chAva that is τ -dynamically-available if $\text{GST} = 0$, and a finalized (and accountable) chain chFin that, if $f < \frac{n}{3}$, is always safe and is live after $\max\{\text{GST}, \text{GAT}\}$. Moreover, for each honest validator v_i and for every round r , chFin_i^r is a prefix of chAva_i^r .

4 Propose-vote-merge protocols

The aim of this work is to present a secure ebb-and-flow [15] protocol that can finalize (at most) one block per slot and, in particular, that can finalize within slot t the block proposed in t . This is achieved by revisiting the *propose-vote-merge* protocol RLMD-GHOST introduced by D'Amato and Zanolini [7] as the basis for our protocol implementation. Propose-vote-merge protocols proceed in *slots* consisting of k rounds², each having a proposer v_p , chosen through a proposer selection mechanism among the set of validators. In particular, at the beginning of each slot t , the proposer v_p proposes a block B . Then, all active validators (also referred as *voters*) vote after Δ rounds. Every validator v_i has a buffer \mathcal{B}_i , a collection of messages received from other

²D'Amato and Zanolini [7] implement RLMD-GHOST with fast confirmation with $k = 3\Delta$ (Appendix B [7]). However, we will consider $k = 4\Delta$, following the approach taken by D'Amato *et al.* [6] when presenting *Goldfish* with *fast confirmation*. We will show how RLMD-GHOST with fast confirmation can be changed into its variant with $k = 4\Delta$ in Section 5 while presenting our protocol.

validators, and a view \mathcal{V}_i , used to make consensus decisions, which admits messages from the buffer only at specific points in time.

Propose-vote-merge protocols are defined through a deterministic fork-choice function FC , which is used by honest proposers and voters to decide how to propose and vote, respectively, based on their view at the round in which they are performing those actions. It is moreover used as the basis of a *confirmation rule* (Section 5.2), which defines the output of the protocol, and thus with respect to which the security of the protocol is defined. In the case of RLMD-GHOST, its fork-choice function RLMD-GHOST considers the last (non equivocating) messages sent by validators that are not older than $t - \eta$ slots (for an expiration period η), in order to make protocol's decisions. In particular, the filter function $\text{FIL}_{\text{rlmd}}(\mathcal{V}, t)$ removes *all but the latest messages within the expiry period $[t - \eta, t)$ for slot t , from non-equivocating validators*, i.e., $\text{FIL}_{\text{rlmd}} = \text{FIL}_{\text{lmd}} \circ \text{FIL}_{\eta\text{-exp}} \circ \text{FIL}_{\text{eq}}$. Here, $\text{FIL}_{\text{lmd}}(\mathcal{V}, t)$ removes all but the latest votes of every validator (possibly more than one) from \mathcal{V} and outputs the resulting view, i.e., it implements the *latest message* (LMD) rule, $\text{FIL}_{\eta\text{-exp}}(\mathcal{V}, t)$ removes all votes from slots $< t - \eta$ from \mathcal{V} and outputs the resulting view, and $\text{FIL}_{\text{eq}}(\mathcal{V}, t)$ removes all votes by *equivocating validators in \mathcal{V} [1]*, i.e., validators for which \mathcal{V} contains multiple, equivocating, votes for some slot t .

A propose-vote-merge protocol proceeds in three phases:

PROPOSE: In this phase, which starts at the beginning of a slot, the proposer v_p merges its view \mathcal{V}_p with its buffer \mathcal{B}_p , i.e., $\mathcal{V}_p \leftarrow \mathcal{V}_p \cup \mathcal{B}_p$, and sets $\mathcal{B}_p \leftarrow \emptyset$. Then, v_p runs the fork-choice function FC with inputs its view \mathcal{V}_p and slot t , obtaining the head of the chain $B' = \text{FC}(\mathcal{V}_p, t)$. Proposer v_p extends B' with a new block B , and updates its canonical chain accordingly, setting $\text{ch}_p \leftarrow B$. Finally, it broadcasts the message $[\text{PROPOSE}, B, \mathcal{V}_p \cup \{B\}, t, v_p]$.

VOTE: Here, every validator v_i that receives a proposal message $[\text{PROPOSE}, B, \mathcal{V}, t, v_p]$ from v_p merges its view with the proposed view \mathcal{V} , by setting $\mathcal{V}_i \leftarrow \mathcal{V}_i \cup \mathcal{V}$. Then, it broadcasts votes for some blocks based on its view. We omit, for the moment, for which blocks a validator v_i votes: it will become clear once we present the full protocol.

MERGE: In this phase, every validator v_i merges its view with its buffer, i.e., $\mathcal{V}_i \leftarrow \mathcal{V}_i \cup \mathcal{B}_i$, and sets $\mathcal{B}_i \leftarrow \emptyset$.

The MERGE phase, along with all other operations involving views and buffers discussed in the previous section, are implemented using the *view-merge* technique [6][7][10]. The idea behind the view-merge technique involves synchronizing the views of all honest validators with the view \mathcal{V}_p of the proposer for a specific slot *before* the validators broadcast their votes in that slot.

D'Amato *et al.* [6] introduce the notion of *active* validators³: awake validators that have terminated a *joining protocol* at a round r , described as it follows. Assuming a propose-vote-merge protocol proceeding in slots of $k = 4\Delta$ rounds, when an honest validator v_i wakes up at some round $r \in (4\Delta(t - 1) + 3\Delta, 4\Delta t + 3\Delta]$, it immediately receives all the messages that were sent while it was asleep, and it adds them into its buffer \mathcal{B}_i , without actively participating in the protocol yet. All new messages which are received are also added to the buffer \mathcal{B}_i . Validator v_i then waits for the *next view-merge opportunity*, at round $4\Delta t + 3\Delta$, in order to merge its buffer \mathcal{B}_i into its view \mathcal{V}_i . At this point, v_i starts executing the protocol. From this point on, validator v_i becomes *active*, until either corrupted or put to sleep by the adversary. We consider such a joining protocol when describing our propose-vote-merge protocol.

5 Protocol specification

5.1 Data structures

We consider five message types: PROPOSE, BLOCK, CHECKPOINT, HEAD-VOTE, and FFG-VOTE. We make no distinctions between network-level representation of blocks and votes, and their representation in a validator's view, i.e., there is no difference between BLOCK and *-VOTE messages and blocks and votes, and we usually just refer to the latter. We describe the objects as tuples (DATA-TYPE, ...) with their data type as a tag, but in practice mostly refer to them without the tag. We use dot notation to refer to the fields. For the tag, we do so simply with .tag, for the other fields we use the generic names specified in the object descriptions

³Observe that D'Amato *et al.* [6] actually refer to *awake* validators to indicate what we call active, and to *dreamy* validators to indicate what we call awake (but not active).

below, to access the different fields, *e.g.*, $B.t$ is the slot of block B . In the following, t is a slot and v_i a validator.

Blocks and checkpoints. A block is a tuple $B = (\text{BLOCK}, b, t, v_i)$, where b is a *block body*, *i.e.*, the protocol-specific content of the block⁴. A checkpoint is a tuple $\mathcal{C} = (\text{CHECKPOINT}, B, t)$, where B is a block and $\mathcal{C}.t \geq B.t$.

Votes. A head vote is a tuple $[\text{HEAD-VOTE}, B, t, v_i]$, where B is a block. An FFG vote is a tuple $[\text{FFG-VOTE}, \mathcal{C}_1, \mathcal{C}_2, v_i]$, where $\mathcal{C}_1, \mathcal{C}_2$ are checkpoints, $\mathcal{C}_1.t < \mathcal{C}_2.t$, and $\mathcal{C}_1.B \prec \mathcal{C}_2.B$. We refer to the two checkpoints as *source* and *target*, respectively, and to FFG votes as *links* between source and target. When v_i is clear from context, we also write $\mathcal{C}_1 \rightarrow \mathcal{C}_2$ for the whole vote, *e.g.*, we say that v_i *casts* a $\mathcal{C}_1 \rightarrow \mathcal{C}_2$ vote.

Proposals. A proposal is a tuple $[\text{PROPOSE}, B, \mathcal{V}, t, v_i]$ where B is a block and \mathcal{V} a view. We refer to \mathcal{V} as a *proposed view*.

Gossip behavior. Votes and blocks are gossiped at any time, regardless of whether they are received directly or as part of another message. For example, a validator receiving a vote also gossips the block that it contains, and a validator receiving a proposal also gossips the blocks and votes contained in the proposed view. Finally, a proposal from slot t is gossiped only during the first Δ rounds of slot t .

5.2 Confirmation rule

A confirmation rule allows validators to identify a *confirmed prefix* of the canonical chain, for which safety properties hold, and which is therefore used to define the output of the protocol. Since the protocol we are going to present outputs two chains, the available chain chAva and the finalized chain chFin , we have two confirmation rules. One is *finality*, which we introduce in Section 5.3, and defines chFin . The other confirmation rule, defining chAva , is the one adopted by RLMD-GHOST, in its variant supporting fast confirmation⁵. It is itself essentially split in two rules, a *slow* κ -deep confirmation rule, which is live also under dynamic participation, and a *fast optimistic rule*, requiring $\frac{2}{3}n$ honest validators to be awake, *i.e.*, a stronger assumption than just τ -compliance. Both rules are employed at round $4\Delta t + 2\Delta$, and chAva is updated to the highest block confirmed by either one, so that liveness of chAva only necessitates liveness of one of the two rules. In particular, τ -compliance is sufficient for liveness. On the other end, safety of chAva requires both rules to be safe.

5.3 FFG component

As mentioned above, a propose-vote-merge protocol is characterized by a fork-choice function that identifies for every slot the current head of the canonical chain for a given validator. Moreover, we described two kind of votes that a validator v_i executes in the VOTE phase: a HEAD-VOTE, used to vote for the head of the canonical chain, *i.e.*, the output of the fork-choice function evaluated at the current slot, and an FFG-VOTE, used by the *FFG-component* of our protocol⁶.

The FFG component of our protocol aims at finalizing one block per slot by counting FFG-VOTES cast at a given slot.

Justification. We say that a set of $\frac{2}{3}n$ distinct FFG votes $\mathcal{C}_1 \rightarrow \mathcal{C}_2$ is a *supermajority link* between \mathcal{C}_1 and \mathcal{C}_2 . We say that a checkpoint \mathcal{C} is *justified* if there is a chain of $k \geq 0$ supermajority links $(B_{\text{genesis}}, 0) \rightarrow \mathcal{C}_1 \cdots \rightarrow \mathcal{C}_{k-1} \rightarrow \mathcal{C}$. In particular, $(B_{\text{genesis}}, 0)$ is justified. Finally, we say that a block B is *justified* if there exists a justified checkpoint \mathcal{C} with $\mathcal{C}.B = B$.

⁴For simplicity, we omit a reference to the parent block.

⁵With some minor changes, as RLMD-GHOST still has 3Δ rounds per slots, by requiring an optimistic assumption on network latency in order for fast confirmations to be live.

⁶The component of our protocol that outputs chFin is almost identical to Casper [3], the *friendly* finality gadget (FFG) adopted by the Ethereum consensus protocol Gasper [4]. This is the reason why we decided to use the *FFG* terminology already accepted within the Ethereum ecosystem.

Slashing. The slashing rules are the same as in Casper FFG. Validator v_i is slashable (see Section 3) for two *distinct* FFG votes $(\mathcal{C}_1, \mathcal{C}_2, v_i)$ and $(\mathcal{C}_3, \mathcal{C}_4, v_i)$ if either: **E₁** (Equivocation) $\mathcal{C}_2.t = \mathcal{C}_4.t$ or **E₂** (Surround voting) $\mathcal{C}_3.t < \mathcal{C}_1.t < \mathcal{C}_2.t < \mathcal{C}_4.t$.

Latest justified checkpoint and block. A checkpoint is justified in a view \mathcal{V} if \mathcal{V} contains the chain of supermajority links justifying it. We refer to the justified checkpoint \mathcal{C} of highest slot $\mathcal{C}.t$ in \mathcal{V} as the *latest justified checkpoint* in \mathcal{V} , or $\mathcal{LJ}(\mathcal{V})$, and to $\mathcal{LJ}(\mathcal{V}).B$ as the *latest justified block* in \mathcal{V} , or $LJ(\mathcal{V})$. Ties are broken arbitrarily (the occurrence of a tie implies that $\frac{n}{3}$ validators are slashable for equivocation). For brevity, we also use \mathcal{LJ}_i to refer to $\mathcal{LJ}(\mathcal{V}_i)$, the latest justified checkpoint in the view \mathcal{V}_i of validator v_i .

Finality. A checkpoint \mathcal{C} is *finalized* if it is justified and there exists a supermajority link $\mathcal{C} \rightarrow \mathcal{C}'$ with $\mathcal{C}'.t = \mathcal{C}.t + 1$. A block B is finalized if there exists a finalized checkpoint \mathcal{C} with $B = \mathcal{C}.B$.

5.4 Voting

Fork-choice. Similarly to Gasper [4], we adopt an hybrid *justification-respecting* fork-choice, namely HFC, building upon RLMD-GHOST [7] fork-choice function. $\text{HFC}(\mathcal{V}, t)$ starts from $LJ(\mathcal{V})$, the *latest justified block* in \mathcal{V} , instead of B_{genesis} , and then proceeds as RLMD-GHOST, *i.e.*, it runs GHOST using the view filtered by FIL_{rlmd} . Formally, we can define it by using another view filter, FIL_{ffg} , *i.e.*, $\text{HFC} = \text{RLMD-GHOST} \circ \text{FIL}_{\text{ffg}}$. $\text{FIL}_{\text{ffg}}(\mathcal{V}, t)$ outputs (\mathcal{V}', t) , where \mathcal{V}' filters out blocks in \mathcal{V} that conflict with $LJ(\mathcal{V})$. In other words, it filters out *branches which do not contain $LJ(\mathcal{V})$* , so $LJ(\mathcal{V})$ is guaranteed to be canonical.

Algorithm 1 HFC, the justification-respecting fork-choice function

```

1: function HFC( $\mathcal{V}, t$ )
2:   return RLMD-GHOST( $\text{FIL}_{\text{ffg}}(\mathcal{V}, t)$ )
3: function  $\text{FIL}_{\text{ffg}}(\mathcal{V}, t)$ 
4:    $\mathcal{V}' \leftarrow \mathcal{V} \setminus \{B \in \mathcal{V}, B.\text{tag} = \text{BLOCK} : LJ(\mathcal{V}) \not\prec B \wedge B \not\prec LJ(\mathcal{V})\}$ 
5:   return  $(\mathcal{V}', t)$ 

```

Voting rules. Consider a validator v_i voting at slot t . Head votes work exactly as in RLMD-GHOST, or any propose-vote-merge protocol, *i.e.*, validators vote for the output of their fork-choice: when it is time to vote, validator v_i casts vote $[\text{HEAD-VOTE}, \text{HFC}(\mathcal{V}_i, t), t, v_i]$. FFG votes always use the *latest justified checkpoint as source*. The target block is the *highest confirmed descendant of the latest justified block, or the latest justified block itself if there is none*. The target checkpoint is then $\mathcal{C}_{\text{target}} = (\arg \max_{B \in \{LJ_i, \text{chAva}\}} |B|, t)$, with $|B|$ being the height of block B , and the FFG vote of v_i is $[\text{FFG-VOTE}, \mathcal{LJ}_i, \mathcal{C}_{\text{target}}, v_i]$, voting for the link $\mathcal{LJ}_i \rightarrow \mathcal{C}_{\text{target}}$.

5.5 Protocol execution

Our protocol is implemented in Algorithm 2 and it works as it follows. Note that the PROPOSE and HEAD-VOTE phases are *exactly* as in a generic propose-vote-merge protocol (see Section 4). Moreover, a slot t in our protocol begins at round $4\Delta t$. At any time, the finalized chain chFin_i of validator v_i just consists of the finalized blocks according to its view \mathcal{V}_i , so we omit explicit updates to chFin in the following.

PROPOSE: At round $4\Delta t$, proposer v_p merges its view \mathcal{V}_p with its buffer \mathcal{B}_p , *i.e.*, $\mathcal{V}_p \leftarrow \mathcal{V}_p \cup \mathcal{B}_p$, and sets $\mathcal{B}_p \leftarrow \emptyset$. Then, v_p runs the fork-choice function HFC with inputs its view \mathcal{V}_p and slot t , obtaining the head of the chain $B' = \text{HFC}(\mathcal{V}_p, t)$. Proposer v_p extends B' with a new block B , and updates its canonical chain accordingly, by setting $\text{ch}_p \leftarrow B$. Finally, it broadcasts the proposal $[\text{PROPOSE}, B, \mathcal{V}_p \cup \{B\}, t, v_p]$.

HEAD-VOTE: In rounds $[4\Delta t, 4\Delta t + \Delta]$, a validator v_i , upon receiving a proposal message (PROPOSE, B , \mathcal{V}, t, v_p) from v_p , merges its view with the proposed view \mathcal{V} by setting $\mathcal{V}_i \leftarrow \mathcal{V}_i \cup \mathcal{V}$. At round $4\Delta t + \Delta$, even if no proposal is received, validator v_i updates its canonical chain by setting $\text{ch}_i \leftarrow \text{HFC}(\mathcal{V}_i, t)$, and casts the head vote (HEAD-VOTE, $\text{HFC}(\mathcal{V}_i, t), t, v_i$).

Algorithm 2 Single slot finality protocol – code for validator v_i

```

1: State
2:    $\mathcal{V}_i \leftarrow \{\mathcal{B}_{\text{genesis}}\}$ : view of validator  $v_i$ 
3:    $\mathcal{B}_i \leftarrow \emptyset$ : buffer of validator  $v_i$ 
4:    $\text{ch}_i \leftarrow \mathcal{B}_{\text{genesis}}$ : canonical chain of validator  $v_i$ 
5:    $t \leftarrow 0$ : the current slot
6:    $r \leftarrow 0$ : the current round

  PROPOSE
7: at  $r = 4\Delta t$  do
8:   if  $v_i = v_p^t$  then
9:      $\mathcal{V}_i \leftarrow \mathcal{V}_i \cup \mathcal{B}_i$ ,  $\mathcal{B}_i \leftarrow \emptyset$ ,  $B' \leftarrow \text{HFC}(\mathcal{V}_i, t)$ 
10:     $B \leftarrow \text{NewBlock}(B')$ ,  $\text{ch}_i \leftarrow B$ 
11:    send message [PROPOSE,  $B$ ,  $\mathcal{V}_i \cup \{B\}$ ,  $t$ ,  $v_i$ ] through gossip

  HEAD-VOTE
12: at  $r = 4\Delta t + \Delta$  do
13:    $\text{ch}_i \leftarrow \text{HFC}(\mathcal{V}_i, t)$ 
14:   send message [HEAD-VOTE,  $\text{HFC}(\mathcal{V}_i, t)$ ,  $t$ ,  $v_i$ ] through gossip

  CONFIRM AND FFG-VOTE
15: at  $r = 4\Delta t + 2\Delta$  do
16:    $B_{\text{fast}} \leftarrow B_{\text{genesis}}$ 
17:    $S_{\text{fast}} \leftarrow \{B \prec \text{ch}_i : |\{v_j : \exists B' \succ B : [\text{HEAD-VOTE}, B', t, v_j] \in \mathcal{B}_i\}| \geq \frac{2}{3}n\}$ 
18:   if  $S_{\text{fast}} \neq \emptyset$  then:
19:      $B_{\text{fast}} \leftarrow \arg \max_{S_{\text{fast}}} |B|$ 
20:   if  $\neg(B_{\text{fast}} \prec \text{chAva}_i \wedge \text{ch}_i^{\lceil \kappa} \prec \text{chAva}_i)$  then:
21:      $\text{chAva}_i \leftarrow \arg \max_{\text{ch} \in \{\text{ch}_i^{\lceil \kappa}, B_{\text{fast}}\}} |\text{ch}|$ 
22:      $\mathcal{C}_{\text{target}} \leftarrow \left( \arg \max_{B \in \{LJ_i, \text{chAva}_i\}} |B|, t \right)$ 
23:     send message [FFG-VOTE,  $\mathcal{LJ}_i$ ,  $\mathcal{C}_{\text{target}}$ ,  $v_i$ ] through gossip

  MERGE
24: at  $r = 4\Delta t + 3\Delta$  do
25:    $\mathcal{V}_i \leftarrow \mathcal{V}_i \cup \mathcal{B}_i$ 
26:    $\mathcal{B}_i \leftarrow \emptyset$ 
27: upon receiving a gossiped message [PROPOSE,  $B$ ,  $\mathcal{V}$ ,  $t$ ,  $v_p^t$ ] do
28:    $\mathcal{B}_i \leftarrow \mathcal{B}_i \cup \{B\}$ 
29:   if  $r \in [4\Delta t, 4\Delta t + \Delta]$  then
30:      $\mathcal{V}_i \leftarrow \mathcal{V}_i \cup \mathcal{V}$ 
31: upon receiving a gossiped BLOCK  $B$  or a gossiped *-VOTE  $V$  from  $v_j$  do
32:    $\mathcal{B}_i \leftarrow \mathcal{B}_i \cup \{B\}$  or  $\mathcal{B}_i \leftarrow \mathcal{B}_i \cup \{V\}$ 

```

CONFIRM: At round $4\Delta t + 2\Delta$, a validator v_i selects for fast confirmation the highest *canonical* block $B_{\text{fast}} \prec \text{ch}_i$ such that \mathcal{B}_i contains $\geq \frac{2}{3}n$ votes from slot t for descendants of B_{fast} , from distinct validators. It then updates its confirmed chain chAva_i to the highest between B_{fast} and $\text{ch}_i^{\lceil \kappa}$, the κ -deep prefix of its canonical chain, *as long as this does not result in updating chAva_i to some prefix of it* (we do not needlessly revert confirmations).

FFG-VOTE: At round $4\Delta t + 2\Delta$, after updating chAva_i , a validator v_i casts the FFG vote (FFG-VOTE, \mathcal{LJ}_i , $\mathcal{C}_{\text{target}}$, v_i), where $\mathcal{C}_{\text{target}} = \left(\arg \max_{B \in \{LJ_i, \text{chAva}_i\}} |B|, t \right)$

MERGE: At round $4\Delta t + 3\Delta$, every validator v_i merges its view with its buffer, *i.e.*, $\mathcal{V}_i \leftarrow \mathcal{V}_i \cup \mathcal{B}_i$, and sets $\mathcal{B}_i \leftarrow \emptyset$.

6 Analysis

Algorithm 2 works in the generalized partially synchronous sleepy model, and is in particular a τ -secure ebb-and-flow protocol, *if we strengthen τ -compliance to require that less than $\frac{n}{3}$ validators are ever slashable*

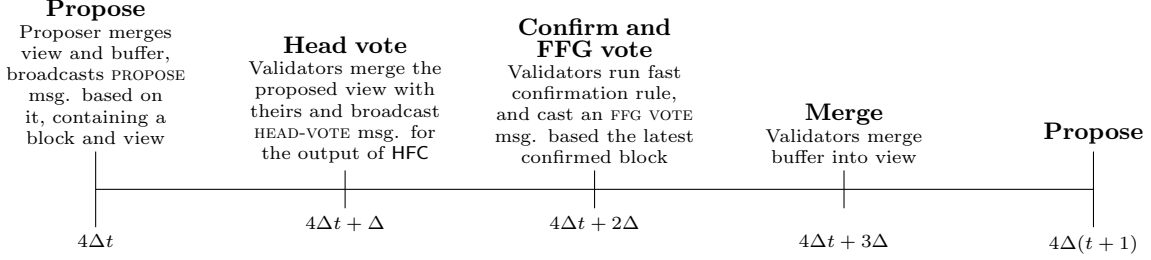


Figure 1: Slot t of our protocol, with its four phases.

for equivocation, for reasons that will be explained shortly. For $\text{GST} = 0$, we show in Section 6.1 that, if the execution is τ -compliant in this stronger sense, then all the properties of RLMD-GHOST [7] keep holding. In Section 6.2 we show that the finalized chain chFin is $\frac{n}{3}$ -accountable, and thus always safe if $f < \frac{n}{3}$. Moreover, if $f < \frac{n}{3}$, chFin is live after $\max\{\text{GST}, \text{GAT}\}$.

Before proceeding with the analysis under synchrony and partial synchrony, we state without proof the *view-merge property*, which follows from the usage of the view-merge technique, since it enables proposers to synchronize the view of honest voters with theirs. It corresponds to Lemma 2 as presented by D’Amato and Zanolini [7], with an addition regarding synchronization of the latest justified checkpoint.

Lemma 1. *Suppose that t is a slot with an (honest) active proposer and that network synchrony holds in rounds $[4\Delta t - \Delta, 4\Delta t + \Delta]$. Say the proposed block is B , and the latest justified checkpoint in the view of the proposer is \mathcal{LJ}_p . Then, at round $4\Delta t + \Delta$, all active validators broadcast a HEAD-VOTE for the honest proposal B of slot t . Moreover, $\mathcal{LJ}_i = \mathcal{LJ}_p$ for any such active validator v_i .*

6.1 Synchrony

Throughout this part of the analysis, we assume that $\text{GST} = 0$, and that $< \frac{n}{3}$ validators are ever slashable for equivocation, by which here we mean signing multiple HEAD-VOTES for a single slot, rather than violating \mathbf{E}_1 . In other words, we are not concerned about equivocation with FFG-VOTES, but rather with HEAD-VOTES, which can similarly be declared a slashable offense. Observe that, in RLMD-GHOST with fast confirmations (Appendix B [7]), this assumption is strictly needed for safety (and only for clients which use fast confirmations), but for example not for reorg resilience or liveness, because fast confirmations do not affect the canonical chain. On the other hand, the protocol we present here utilizes confirmations as a prerequisite for justification, and justification does affect the canonical chain, since HFC filters out branches conflicting with the latest justified block. Therefore, we require that $< \frac{n}{3}$ validators are ever slashable for equivocation for all of the properties which we are going to prove. As already mentioned, to avoid stating it repeatedly, we further restrict η -compliant executions to those executions in which the assumption holds.

Our single slot finality protocol implemented in Algorithm 2 uses the HFC fork-choice function, dealing with checkpoints and justifications. However, one could implement it using also different fork-choice functions. In particular, we show that by substituting HFC with RLMD-GHOST (with equal expiration period η), *i.e.*, if we ignore justifications and consider the *vanilla* fork-choice function introduced by D’Amato and Zanolini [7], then the resulting protocol is equivalent to the RLMD-GHOST protocol with fast confirmation (Appendix B [7]). This because FFG votes have no effect at all, and as such it is η -reorg-resilient, and η -dynamically-available. Moreover, the following two results about fast confirmations (Appendix B [7]) also apply.

Theorem 1 (Reorg resilience of fast confirmations). *Let us consider an η -compliant execution with $\text{GST} = 0$. A block fast confirmed by an honest validator at a slot t is always in the canonical chain of all active validators at rounds $\geq 4\Delta(t+1) + \Delta$.*

Theorem 2 (Liveness of fast confirmations). *An honest proposal B from a slot t in which $|H_t| \geq \frac{2}{3}n$ is fast confirmed by all active validators at round $4\Delta t + \Delta$.*

We show that, under synchrony, *i.e.*, with $\text{GST} = 0$, these properties are preserved by our justification-respecting protocol, which uses HFC instead. To do so, we show that for every η -compliant execution,

Algorithm 2 using FC = RLMD-GHOST and Algorithm 2 using FC = HFC are *equivalent*, i.e., the sequence of outputs of Algorithm 2 is the same regardless of which fork-choice function is used. All properties of Algorithm 2 with FC = RLMD-GHOST in such η -compliant executions then also apply to Algorithm 2 with FC = HFC. In particular, it is also η -reorg-resilient and η -dynamically-available, and it also satisfies reorg resilience and liveness of fast confirmations, i.e., Theorems 1 and Theorem 2 hold.

Theorem 3 (Execution equivalence). *Let us consider an η -compliant execution with GST = 0 and with Algorithm 2 using FC = HFC. Furthermore, let us consider the same execution, with the same adversarial strategy and randomness, with Algorithm 2 using FC = RLMD-GHOST. The sequence of outputs of the two algorithms correspond exactly.*

Proof. Since the only difference between the two protocols is the fork-choice function FC, the sequences of outputs correspond as long as the outputs of HFC and RLMD-GHOST obtained by active validators are always the same in the two executions. FC is used only twice in Algorithm 2, in Line 9 for proposing, and in Lines 13-14, with the same input, for broadcasting a HEAD-VOTE. We are going to prove by induction that the canonical chain of an active validator *at any voting round* is the same in both executions. Since Line 13 sets $\text{ch}_i \leftarrow \text{FC}(\mathcal{V}_i, t)$, and this value is the same as in Line 14, we only need to show that the fork-choice output in Line 9 coincides in the two executions as well. In Line 14, an honest validator uses the fork-choice output to construct their HEAD-VOTES, so HEAD-VOTES correspond in both executions. Moreover, the view-merge property applies in both executions, so honestly proposed blocks correspond to the honest HEAD-VOTES from their slot. Therefore, HEAD-VOTES coinciding in the two executions implies that honestly proposed blocks coincide as well. Since honestly proposed blocks extend the output of the fork-choice at Line 9, this output is then also the same in both executions, completing the proof. We now carry out the induction.

Induction hypothesis: At any slot $t' \leq t$ and for $r = 4\Delta t' + \Delta$, ch_i^r coincides in both executions, for any active validator i .

Base case: In rounds $[0, \Delta]$, the two executions are exactly the same, because the only justified checkpoint is B_{genesis} , so HFC = RLMD-GHOST. Therefore, the statement holds for $t = 0$.

Inductive step: Suppose now that the statement holds for t , and consider round $r = 4\Delta(t + 1) + \Delta$. Consider an active validator v_i with view \mathcal{V}_i at round r , and latest justified block $B = LJ(\mathcal{V}_i)$. Let t' be minimal such that there exists a justified checkpoint $\mathcal{C} = (B, t')$, i.e., slot t' is the first slot in which block B was justified. The supermajority link with target \mathcal{C} contains at least one FFG vote from an honest validator v_k . By minimality of t' , B could not have been already justified in the view of v_k when broadcasting an FFG vote at slot t' . Therefore, by Lines 22-23 of Algorithm 2, it must be the case that $B \prec \text{chAva}_k$ at round $4\Delta t' + 2\Delta$, i.e., that it had been confirmed by v_k . If it was fast confirmed at a slot $\leq t'$, then, in the execution with FC = RLMD-GHOST, Theorem 1 implies that $B \prec \text{ch}_j^{r'}$ for all active validators v_j at any round $r' \geq 4\Delta(t' + 1) + \Delta$, and so in particular that $B \prec \text{ch}_i^r$, since $t > t'$. If instead $B \prec \text{ch}_k^{\kappa}$ at round $4\Delta t' + 2\Delta$, i.e., B is confirmed by v_k due to being κ -deep in its canonical chain, then with overwhelming probability there exists a pivot slot $t'' \in [t' - \kappa, t')$ (Lemma 3 [7]), with proposed block B' . In the execution with FC = RLMD-GHOST, η -reorg-resilience then implies that $B' \prec \text{ch}_j^{r'}$ for all active validators v_j at any round $r' \geq 4\Delta t'' + \Delta$. In particular, $B' \prec \text{ch}_k^{r'}$ at round $r' = 4\Delta t' + 2\Delta$, and $B' \prec \text{ch}_i^r$. The former implies $B \prec B'$, since $B.t \leq t' - \kappa \leq B'.t$, and we then have $B \prec B' \prec \text{ch}_i^r$.

Anyway, regardless of how B has been confirmed by v_k , we have $B \prec \text{ch}_i^r$. Therefore, $LJ(\mathcal{V}_i) = B \prec \text{RLMD-GHOST}(\mathcal{V}_i, t+1)$, which in turn implies $\text{RLMD-GHOST}(\mathcal{V}_i, t+1) = \text{RLMD-GHOST} \circ \text{FIL}_{\text{FFG}}(\mathcal{V}_i, t+1) = \text{HFC}(\mathcal{V}_i, t+1)$. Therefore, after v_i updates its canonical chain ch_i at round r by setting $\text{ch}_i \leftarrow \text{FC}(\mathcal{V}, t+1)$, with FC dependent on the execution, ch_i is the same in both executions. \square

6.2 Partial synchrony

Throughout this section we assume that $f < \frac{\eta}{3}$. First, we prove that the finalized chain is accountably safe, exactly as done in Casper [3]. Then, we show that honest proposals made after $\max(\text{GST}, \text{GAT}) + \Delta$ are justified within their proposal slot, which implies liveness of the finalized chain.

Theorem 4 (Accountable safety). *The finalized chain chFin is accountably safe, i.e., two conflicting finalized blocks imply that at least $\frac{\eta}{3}$ adversarial validators can be detected to have violated either \mathbf{E}_1 or \mathbf{E}_2 .*

Proof. We assume throughout that there are no double justifications, *i.e.*, there are no checkpoints $\mathcal{C} \neq \mathcal{C}'$ with $\mathcal{C}.t = \mathcal{C}'.t$, and we refer to this as the non-equivocation assumption. If that's not the case, clearly $\geq \frac{n}{3}$ validators are slashable for violating \mathbf{E}_1 . Consider two conflicting finalized blocks B and B' . By definition, there are also finalized checkpoints \mathcal{C} and \mathcal{C}' with $B = \mathcal{C}.B$, $B' = \mathcal{C}'.B$. Say \mathcal{C} is finalized by the chain of supermajority links $(B_{\text{genesis}}, 0) \rightarrow \mathcal{C}_1 \cdots \rightarrow \mathcal{C}_k = \mathcal{C} \rightarrow \mathcal{C}_{k+1}$, with $\mathcal{C}_{k+1}.t = \mathcal{C}.t + 1$, and \mathcal{C}' by the chain $(B_{\text{genesis}}, 0) \rightarrow \mathcal{C}'_1 \cdots \rightarrow \mathcal{C}'_{k'} = \mathcal{C}' \rightarrow \mathcal{C}'_{k'+1}$, with $\mathcal{C}'_{k'+1}.t = \mathcal{C}'.t + 1$. Let $t_i = \mathcal{C}_i.t$, and $t'_i = \mathcal{C}'_i.t$. By the non-equivocation assumption, $t_k \neq t'_k$, and without loss of generality we take $t_k < t'_k$. Let $j = \min\{i \leq k' : t_k < t'_i\}$, so $t_k < t'_j \leq t'_{k'}$, and $t'_{j-1} \leq t_k$ by minimality of t'_j . By the non-equivocation assumption, $t'_j = t_{k+1}$ implies that $\mathcal{C}_{k+1} = \mathcal{C}'_j$. We then have $B = \mathcal{C}.B \prec \mathcal{C}_{k+1}.B = \mathcal{C}'_j.B \prec \mathcal{C}'.B = B'$, contradicting that B and B' are conflicting. Therefore, $t'_j > t_k + 1 = t_{k+1}$ as well. Similarly, $t'_{j-1} < t_k$. Therefore, we have $t'_{j-1} < t_k < t_{k+1} < t'_k$, *i.e.*, $\mathcal{C}'_{j-1}.t < \mathcal{C}_k.t < \mathcal{C}_{k+1}.t < \mathcal{C}'_j.t$. The intersection of the two sets of voters of the supermajority links $\mathcal{C}_k \rightarrow \mathcal{C}_{k+1}$ and $\mathcal{C}'_{j-1} \rightarrow \mathcal{C}'_j$ contains at least $\frac{n}{3}$ validators, which are then slashable for violating \mathbf{E}_2 . \square

Lemma 2. *If an honest proposer v_p proposes a block B at a slot t after $\max(\text{GST}, \text{GAT}) + \Delta$, and the latest justified checkpoint in the view of the proposer is \mathcal{LJ}_p , then the checkpoint (B, t) is justified in all honest views at round $4\Delta t + 3\Delta$, by supermajority link $\mathcal{LJ}_p \rightarrow (B, t)$.*

Proof. Since t is after $\text{GAT} + \Delta$, all $> \frac{2}{3}n$ honest validators are awake since at least round $4\Delta t - \Delta$, so at slot t they have completed the joining protocol and are active. Moreover, the view-merge property (Lemma 1) applies to all of them. Consider now an honest validator v_i . By the view-merge property, validator v_i broadcasts a HEAD-VOTE for B at round $4\Delta t + \Delta$. Also by the view-merge property, $\mathcal{LJ}_i = \mathcal{LJ}_p$ at round $4\Delta t + \Delta$, but \mathcal{LJ}_i does not change until round $4\Delta t + 3\Delta$, since \mathcal{V}_i does not. Therefore, $\mathcal{LJ}_i = \mathcal{LJ}_p$ at round $4\Delta t + 2\Delta$. By that round, all $\geq \frac{2}{3}$ honest HEAD-VOTES for B are received by all honest validators, including v_i . Since also $B \prec \text{ch}_i$, v_i fast confirms B , and thus broadcasts an FFG vote $\mathcal{LJ}_i \rightarrow (B, t) = \mathcal{LJ}_p \rightarrow (B, t)$. All honest validators receive such votes by round $4\Delta t + 3\Delta$, and merge them into their view then. Therefore, checkpoint (B, t) is justified in all honest views at that round. \square

Theorem 5 (Liveness). *Consider two consecutive slots t and $t+1$ with honest proposers after $\max(\text{GST}, \text{GAT}) + 4\Delta$. The block B proposed at slot t is finalized at the end of slot $t + 1$.*

Proof. By Lemma 2, checkpoint (B, t) is justified in all honest views at round $4\Delta t + 3\Delta$. Since at the beginning of slot $t + 1$ there cannot be any justified checkpoint with slot $> t$, and there cannot be any other justified checkpoint with slot t , (B, t) is therefore the latest justified block in the view of the proposer of slot $t + 1$. B is then canonical in its view, and it proposes a block B' which extends B . Again by Lemma 2, $(B', t+1)$ is justified in all honest views at round $4\Delta(t+1) + \Delta$, by the supermajority link $(B, t) \rightarrow (B', t+1)$. Therefore, B is finalized in all honest views. \square

7 Single slot finality

The protocol implemented in Algorithm 2 is a an η -secure ebb-and-flow protocol which (at best) finalizes a block in every slot, but it does not achieve *single slot finality*, *i.e.*, it cannot finalize a proposal *within its proposal slot*. At best, it lags behind by one slot, finalizing a proposal from slot t at the end of slot $t + 1$. A straightforward extension of our protocol which achieves single slot finality is one with 5Δ rounds per slot, allowing for an additional FFG voting phase. This would be very costly in Ethereum, for two reasons. First, it would in practice significantly increase the slot time, because each voting round requires aggregating hundreds of thousands (if not millions) of BLS signatures, likely requiring a lengthier multi-step aggregation process. Moreover, it would be expensive in terms of bandwidth consumption and computation, because such votes would have to all be gossiped and verified by each validator, costly even if already aggregated. For these reasons, we describe here an alternative way to enhance to protocol for the purpose of achieving single slot finality, without suffering from the drawbacks just described. We introduce a new type of message, *acknowledgment*, and a new slashing condition alongside it. Acknowledgments do not influence the protocol in any way, except in case of slashing, and are mainly intended to be consumed by external observers which want to have the earliest possible finality guarantees. Therefore, they do not need to be gossiped to and verified by all validators. They can then simply be gossiped in smaller sub-networks (similar to the *attestation subnets*

which Ethereum employs today), requiring limited bandwidth and verification resources. If an observer wants to have faster finality guarantees than they could have by simply following the chain or listening to the global gossip, they can opt to participate in all such sub-networks, and collect all acknowledgments. As doing so is permissionless, it can also be expected that aggregate acknowledgments, or equivalent proofs, might become available through some other channels.

Acknowledgment. An *acknowledgment* is a tuple $[\text{ACK}, \mathcal{C}, t, v]$, where \mathcal{C} is a checkpoint with $\mathcal{C}.t = t$. We also refer to this as an acknowledgment of \mathcal{C} . A *supermajority acknowledgment* of \mathcal{C} is a set of $\geq \frac{2}{3}n$ distinct acknowledgments of \mathcal{C} . At round $4\Delta t + 3\Delta$, after merging the buffer \mathcal{B}_i , validator v_i broadcasts the acknowledgment $[\text{ACK}, \mathcal{L}\mathcal{J}_i, t, v_i]$ if $\mathcal{L}\mathcal{J}_i.t = t$, i.e., if $\mathcal{L}\mathcal{J}_i$ has been justified in the current slot. An observer which receives a supermajority acknowledgment of a *justified* checkpoint \mathcal{C} may consider \mathcal{C} to be finalized.

Slashing rule for finality voting. When validator v_i broadcasts an acknowledgment of (\mathcal{C}, t) , it *acknowledges* that, at the end of slot t , it knows about \mathcal{C} being justified. Since the FFG voting rules prescribe that the source of an FFG vote should be the latest known justified checkpoint, subsequent FFG votes with a source whose slot is $< t$ constitute a provable violation, which is analogous to surround voting. Accordingly, we formulate a third slashing rule, which ensures that finality via a supermajority acknowledgment is accountably safe. In particular, validator v_i is slashable for an FFG vote $(\mathcal{C}_1, \mathcal{C}_2)$ and an acknowledgment (\mathcal{C}, t) , if they satisfy **E₃**, i.e., $\mathcal{C}_1.t < \mathcal{C}.t < \mathcal{C}_2.t$. In other words, the link $\mathcal{C}_1 \rightarrow \mathcal{C}_2$ *surrounds* the acknowledged checkpoint.

Theorem 6 (Accountable safety with acknowledgments). *The finalized chain is accountably safe even when it is updated via acknowledgments as well, i.e., two conflicting finalized checkpoints imply that more than $\frac{n}{3}$ adversarial validators can be detected to have violated **E₁**, **E₂**, or **E₃**.*

Proof. The proof largely follows that of Theorem 4. We again consider two conflicting finalized blocks B and B' , and corresponding finalized checkpoints \mathcal{C} and \mathcal{C}' . Regardless of whether finalization is through a supermajority link or a supermajority acknowledgment, \mathcal{C} and \mathcal{C}' have to be justified, by chains of supermajority links $(B_{\text{genesis}}, 0) \rightarrow \mathcal{C}_1 \cdots \rightarrow \mathcal{C}_k = \mathcal{C}$ and $(B_{\text{genesis}}, 0) \rightarrow \mathcal{C}'_1 \cdots \rightarrow \mathcal{C}'_{k'} = \mathcal{C}'$. Let $t_i = \mathcal{C}_i.t$, and $t'_i = \mathcal{C}'_i.t$. By the non-equivocation assumption considered in Theorem 4, we again have $t_k \neq t'_k$, and without loss of generality we take $t_k < t'_k$. As before, we let $j = \min\{i \leq k' : t_k < t'_i\}$, so $t_k < t'_j \leq t'_{k'}$, and $t'_{j-1} \leq t_k$ by minimality of t'_j . Moreover, also by the non-equivocation assumption, $t'_{j-1} < t_k$. If \mathcal{C} is finalized through a supermajority link, the proof of Theorem 4 already shows that at least $\frac{n}{3}$ validators must have violated **E₂**, and it is still applicable here because it does not use the last supermajority link in the chain finalizing \mathcal{C}' (which may or may not exist here). If instead \mathcal{C} is finalized through a supermajority acknowledgment, i.e., there are $\frac{2}{3}n$ acknowledgments of \mathcal{C} , then at least $\frac{n}{3}$ validators have violated **E₃**, because $\mathcal{C}'_{j-1}.t < \mathcal{C}.t < \mathcal{C}'_j.t$. \square

Theorem 7 (Single Slot Finality). *An honest proposal from a slot t after $\max(\text{GST}, \text{GAT}) + 4\Delta$ is finalized in round $4\Delta(t + 1)$ by a supermajority acknowledgment.*

Proof. Say the honestly proposed block is B . By Lemma 2, checkpoint $\mathcal{C} = (B, t)$ is justified in all honest views at round $4\Delta t + 3\Delta$. Therefore, all honest validators broadcast an acknowledgment of \mathcal{C} . Any observer which listens for acknowledgments would receive all such messages by rounds $4\Delta(t + 1)$, and thus possesses a supermajority acknowledgment of \mathcal{C} . Such observer may then consider \mathcal{C} , and thus also B , to be finalized. \square

8 Conclusions

In this work, we have made significant strides towards realizing a secure and reorg-resilient ebb-and-flow protocol that has the potential to replace Ethereum’s current consensus protocol, Gasper. We have provided a comprehensive analysis and modifications to D’Amato and Zanolini’s RLMD-GHOST protocol, integrating it with a partially synchronous finality gadget. In particular, our protocol introduces a novel approach for achieving single slot finality.

Another significant contribution of our work lies in the expansion of the generalized sleepy model introduced by D’Amato and Zanolini. Our generalized partially synchronous sleepy model introduces stronger constraints related to the adversary’s corruption and sleepiness power and incorporates the concept of partial

synchrony. This extension not only enhances the original model but also provides a generalized framework suitable for a wider array of practical scenarios.

However, despite the security guarantees of our protocol, we acknowledge that it is not (yet) practical for real-world implementation. This challenge is due to the current structure of Ethereum, which employs a large pool of validators. Requiring every validator to vote at each slot would necessitate extensive message exchanges – a process that is far from optimal given the scale of Ethereum’s network. Therefore, while our current findings represent a crucial stride towards an improved consensus protocol, they also highlight the need for additional research. Specifically, we need to focus on how we can refine the voting mechanism to better manage and aggregate the messages involved in this process.

References

- [1] Aditya Asgaonkar. Remove equivocating validators from fork choice consideration. URL: <https://github.com/ethereum/consensus-specs/pull/2845>.
- [2] Vitalik Buterin. Paths toward single-slot finality, 2023. URL: https://notes.ethereum.org/@vbuterin/single_slot_finality.
- [3] Vitalik Buterin and Virgil Griffith. Casper the friendly finality gadget. *CoRR*, abs/1710.09437, 2017. URL: <http://arxiv.org/abs/1710.09437>, arXiv:1710.09437.
- [4] Vitalik Buterin, Diego Hernandez, Thor Kampefner, Khiem Pham, Zhi Qiao, Danny Ryan, Juhyeok Sin, Ying Wang, and Yan X Zhang. Combining GHOST and Casper. 2020.
- [5] Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In Margo I. Seltzer and Paul J. Leach, editors, *Proceedings of the Third USENIX Symposium on Operating Systems Design and Implementation (OSDI), New Orleans, Louisiana, USA, February 22-25, 1999*, pages 173–186. USENIX Association, 1999.
- [6] Francesco D’Amato, Joachim Neu, Ertem Nusret Tas, and David Tse. No more attacks on proof-of-stake ethereum? *CoRR*, abs/2209.03255, 2022. URL: <https://doi.org/10.48550/arXiv.2209.03255>.
- [7] Francesco D’Amato and Luca Zanolini. Recent latest message driven ghost: Balancing dynamic availability with asynchrony resilience, 2023. URL: <https://arxiv.org/abs/2302.11326>.
- [8] Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, 1988.
- [9] Seth Gilbert and Nancy A. Lynch. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, 33(2):51–59, 2002.
- [10] Daniel Kane, Andreas Fackler, Adam Gagol, and Damian Straszak. Highway: Efficient consensus with flexible finality. *CoRR*, abs/2101.02159, 2021. URL: <https://arxiv.org/abs/2101.02159>, arXiv:2101.02159.
- [11] Andrew Lewis-Pye and Tim Roughgarden. Resource pools and the CAP theorem. *CoRR*, abs/2006.10698, 2020. URL: <https://arxiv.org/abs/2006.10698>.
- [12] Dahlia Malkhi, Atsuki Momose, and Ling Ren. Byzantine consensus under fully fluctuating participation. *IACR Cryptol. ePrint Arch.*, page 1448, 2022.
- [13] Atsuki Momose and Ling Ren. Constant latency in sleepy consensus. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS 2022, Los Angeles, CA, USA, November 7-11, 2022*, pages 2295–2308. ACM, 2022.
- [14] Joachim Neu, Ertem Nusret Tas, and David Tse. A balancing attack on Gasper, the current candidate for Eth2’s beacon chain. URL: <https://ethresear.ch/t/a-balancing-attack-on-gasper-the-current-candidate-for-eth2s-beacon-chain/8079>.

- [15] Joachim Neu, Ertem Nusret Tas, and David Tse. Ebb-and-flow protocols: A resolution of the availability-finality dilemma. In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*, pages 446–465. IEEE, 2021.
- [16] Rafael Pass and Elaine Shi. The sleepy model of consensus. In *ASIACRYPT (2)*, volume 10625 of *Lecture Notes in Computer Science*, pages 380–409. Springer, 2017.
- [17] Yonatan Sompolinsky and Aviv Zohar. Secure high-rate transaction processing in Bitcoin. In *International Conference on Financial Cryptography and Data Security*, pages 507–527. Springer, 2015.
- [18] Vlad Zamfir. Casper the friendly ghost. a correct-by-construction blockchain consensus protocol. URL: <https://github.com/ethereum/research/blob/master/papers/cbc-consensus/AbstractCBC.pdf>.