

No More Attacks on Proof-of-Stake Ethereum?

Francesco D’Amato
francesco.damato@ethereum.org

Ertem Nusret Tas
nusret@stanford.edu

Joachim Neu
jneue@stanford.edu

David Tse
dntse@stanford.edu

ABSTRACT

The latest message driven (LMD) greedy heaviest observed sub-tree (GHOST) consensus protocol is a critical component of proof-of-stake (PoS) Ethereum. In its current form, the protocol is brittle, and intricate to reason about, as evidenced by recent attacks and patching attempts. We report on Goldfish, a considerably simplified variant of the current protocol, and a candidate under consideration for a future Ethereum protocol upgrade. We prove that Goldfish is secure in synchronous networks under dynamic participation, assuming a majority of the nodes (called *validators*) follows the protocol. Goldfish improves over Nakamoto’s longest-chain consensus in that it is *reorg resilient* (i.e., honestly produced blocks are guaranteed inclusion in the ledger) and supports *fast confirmation* (i.e., the expected confirmation latency is independent of the desired security level). We show that subsampling validators can improve the communication efficiency of Goldfish, and that Goldfish is composable with finality gadgets and accountability gadgets, which improves state-of-the-art ebb-and-flow protocols. Akin to traditional propose-and-vote-style consensus protocols, Goldfish is organized in slots, at the beginning of which a leader proposes a block containing new transactions, and subsequently members of a committee take a vote towards block confirmation. But instead of using quorums, Goldfish is powered by a new mechanism to carefully synchronize the inclusion and exclusion of blocks and votes in honest validators’ views.

1 INTRODUCTION

The latest message driven (LMD) greedy heaviest observed sub-tree (GHOST) [57, 31] consensus protocol is a key component of the Gasper protocol [9] that powers proof-of-stake (PoS) Ethereum’s beacon chain since ‘the Merge’. The initial version specified with Gasper [9] was shown to be broken using the *balancing attack*, first in synchronous and partially synchronous networks with adversarial message delay [47, 43], and later in networks with non-adversarial but merely random network delay [49, 44, 54]. In response, a patch called *proposer boosting* was added to the protocol [7]. It was subsequently shown that the LMD functionality alone can be exploited to conduct a balancing-type attack *despite* proposer boosting [50, 46], and that Gasper’s LMD GHOST component without LMD would suffer from a so called *avalanche attack* [50, 45]. Again in response, a patch called *equivocation discounting* was added to the protocol. Not least because of its complexity, the protocol has so far defied security analysis—both in terms of giving a formal security proof as well as further attacks. This leaves room for an uncomfortable amount of doubt about the security of Ethereum’s ecosystem worth hundreds of billions of US dollars.

We present a protocol, nicknamed Goldfish, with the following

key properties motivated by the application (the importance of these properties as well as their precise definition is subsequently detailed): (a) The protocol can be viewed as a small variation of the currently specified and deployed LMD GHOST protocol of the PoS Ethereum beacon chain. (b) It is *provably secure*, assuming honest majority of *validators* (i.e., nodes with stake), and network synchrony (i.e., adversarial network delay, up to a known delay upper bound Δ). (c) It can tolerate *dynamic participation* [51], i.e., a large fraction and fluctuating set of simultaneous temporary crash faults among validators. (d) It is *reorg resilient*, i.e., block proposals by honest validators are guaranteed to make their way into the output ledger, with a prefix known to the block producer at the time of block production. (e) It supports *subsampling* of validators, to improve communication efficiency and resilience to *adaptive corruption* (cf. *player-replaceability* [22, 14]). (f) It is *simple*. (g) It is composable with *finality gadgets* and *accountability gadgets* such as [8, 47, 53, 49]. The composite can achieve the *ebb-and-flow consensus formulation* [47] desired of PoS Ethereum’s beacon chain.

As a result, Goldfish can serve the following purposes: (a) The protocol can serve as a drop-in replacement for LMD GHOST in the PoS Ethereum beacon chain protocol. Due to its similarity to LMD GHOST, it is a credible candidate for a future upgrade of PoS Ethereum consensus, requiring relatively small implementation changes, and thus presents an option for the Ethereum ecosystem, should problems with the current protocol aggravate. (b) Unlike earlier negative results (attacks) on variants of LMD GHOST as is part of the PoS Ethereum beacon chain, Goldfish is the first positive result (security proof) for a close variant, slightly strengthening confidence in this family of protocols. (c) The protocol is a good pedagogical example for a simple yet feature-rich consensus protocol for synchronous networks under dynamic participation.

Akin to traditional propose-and-vote-style consensus protocols, Goldfish is organized into slots, at the beginning of which a (pseudo-randomly elected) leader proposes a block containing new transactions, and subsequently members of a (pseudo-randomly elected) committee take a vote towards block confirmation. But instead of using fix-sized quorums, Goldfish is based on two key techniques, *message buffering* and *vote expiry*, to carefully synchronize honest validators’ views, and which might be of independent interest:

- (a) *Message buffering* (also known as *view merge* [3]) first appeared in [27]. In short, buffering of votes received from the network together with carefully timed inclusion of these votes in each validator’s local view leads to the property that in periods with honest leader, all honest validators vote in favor of the leader’s proposal. This leads to reorg resilience, i.e., honest proposals are guaranteed to remain in the canonical chain. Since honest proposals contain fresh transactions and stabilize their prefix, and long streaks of only adversarial proposals are exponentially

The authors are listed alphabetically.

unlikely, safety and liveness of the protocol follow readily under high participation and without subsampling.

- (b) *Vote expiry* (also known as *ephemeral votes*) means that during each time slot only votes from the immediately preceding time slot influence the protocol’s behavior.¹ This allows the protocol to support dynamic participation, and to subsample small committees of voters per slot from the full set of validators (for improved communication efficiency). Furthermore, vote expiry keeps the set of votes small that might affect short-term future actions of honest validators. Thus, only few protocol messages need to be buffered and merged among honest validators’ views at any point in time. Vote expiry is thus a prerequisite for the feasibility and efficiency of message buffering.

Inspired by the application requirements for a drop-in replacement of LMD GHOST in the PoS Ethereum beacon chain, Goldfish was designed to achieve the following goals:

Secure consensus in synchronous networks under dynamic participation [51] and honest majority: The protocol is parametric in a security parameter κ and outputs a single ledger at each validator at any point in time. The ledger is *safe* (meaning that ledgers output by two validators at two points in time are one a prefix of the other), except with probability decaying exponentially in κ . The ledger is *live* (meaning that transactions enter the ledgers output by honest validators ‘soon’ after they were first input to an honest validator), with a confirmation delay determined by the analysis (a linear function of κ), except with probability decaying exponentially in κ . *Safety and liveness constitute security of the consensus protocol.*

Guaranteeing security under dynamic participation, *i.e.*, tolerating a large number of temporary crash faults, is key in the semi-permissionless model of public PoS blockchains, where it makes the protocol resilient to unforeseen dropouts due to, for instance, regulatory requirements or software/hardware updates. At the time of writing, approximately 70% of Ethereum validators² follow U.S. Office of Foreign Assets Control (OFAC) regulations, and ignore certain transactions. It is conceivable, that under similar future circumstances these 70% of validators selectively abstain from voting, and thus behave like temporary crash faults.

Composability with finality gadgets and accountability gadgets: Goldfish is composable with *finality gadgets and accountability gadgets* such as [8, 47, 53, 49]. The resulting composite protocol (cf. Fig. 2) can achieve the *ebb-and-flow consensus formulation* [47, 49] desired of PoS Ethereum’s beacon chain (cf. Def. 4).

Reorg resilience: As part of the Goldfish protocol, honest validators every now and then get to be the *leader* and get to propose blocks (bundles of transactions) for inclusion. The protocol is resilient to reorgs, meaning that whenever there is an honest leader, its proposal will eventually make it into the protocol’s output ledger, with a prefix ledger that can be determined at the time of block production. *This property is broadly important for incentive alignment, e.g., it reduces the risk of undercutting [34, 11], time-bandit [16], or selfish mining [19] attacks.*

Subsampling: The protocol supports subsampling, meaning that at each slot the protocol can pseudo-randomly select a small group of validators to run the protocol on behalf of the total validator

set. *The results in a considerably lower communication overhead.* Furthermore, the selected validators send only a single protocol message. *Thus, the protocol satisfies player-replaceability [22, 14] and is secure against adaptive adversaries (which can corrupt validators during protocol execution).*

Optimistic fast confirmation: Under *optimistic* conditions when participation happens to be high and a supermajority of $\frac{3}{4}$ fraction of validators is honest, Goldfish confirms with constant expected latency independent of κ .

Similarity to LMD GHOST: Goldfish is intentionally simple, and similar to LMD GHOST as currently deployed, offering a credible path to adoption of Goldfish in the short to medium term. For the two key ingredients, vote expiry can be realized entirely with minor changes to the vote accounting logic. Message buffering becomes practical due to vote expiry. While message buffering requires slight changes to the temporal structure and validator behavior of the current protocol, Goldfish and the current LMD GHOST are similar ‘in spirit’ and share their fundamental structure.

1.1 Related Works

For Goldfish, we build on the sleepy model [51] of a synchronous network where the number and identity of actively participating (*awake*) validators can change over time (*dynamic participation*). The first secure consensus protocol for the sleepy model was Nakamoto’s seminal *longest chain* (LC) protocol, first for proof-of-work (PoW) with Bitcoin [39, 21], and subsequently for PoS with protocols such as Ouroboros [32, 17, 1] and Sleepy Consensus/Snow White [51, 15]. A drawback of these protocols is that the (expected) confirmation latency scales linearly with the security parameter κ (same for Goldfish’s ‘standard’ confirmation rule). Parallel composition of LC protocol instances was suggested in [2, 20] to overcome the κ -dependence of the confirmation latency. Goldfish has an optimistic fast confirmation rule providing κ -independent latency under high participation. Unlike Goldfish, LC protocols are not reorg resilient: with selfish mining [19], every block produced by the adversary can be used to displace one honestly produced block.

In contrast, many ‘classical’ propose-and-vote-style BFT consensus protocols [12, 58, 13] have constant (expected) confirmation latency and are (or can be modified to be) reorg resilient, but do not tolerate dynamic participation. An early consensus protocol of ‘classical’ flavor for a model with unknown (but *static* rather than *dynamic*) participation is due to Khanchandani and Wattenhofer [30, 29]. A subsequent protocol of the ‘classical’ variety [24] supports dynamic participation, but with confirmation latency linear in the security parameter κ . Like Goldfish (and unlike LC), the latency of this protocol is independent of the participation level. Probabilistic security is also overcome in the permissionless PoW setting with omission faults by [52].

A recent work by Momose and Ren [38] presents the first propose-and-vote-style permissioned/PoS protocol that supports dynamic participation with confirmation latency independent of security parameter and level of participation. In the contemporary but independent work [36, 35], the prerequisites for liveness were relaxed, at the expense of reduced adversarial resilience (from $\frac{1}{2}$ down to $\frac{1}{3}$). Thus, Goldfish improves over [36, 35] in resilience.

A key challenge for ‘classical’ consensus protocols in the sleepy

¹From the alleged forgetfulness of its animal namesake stems Goldfish’s name.

²<https://www.mevwatch.info/>

setting is that quorum certificates are no longer transferable between awake honest validators. The works of Momose, Ren, and Malkhi aim to (partially) restore/replace transferability with graded agreement, but otherwise retain the structure of a ‘classical’ consensus protocol: Multiple stages of voting, with the aim of reaching quorums to progress towards confirmation. Validators keep individual state across the stages using locks, and express discontent with a proposal by abstaining from the vote. In contrast, Goldfish is closer in spirit to LC: A simple main loop in which validators repeatedly express support for the leading tip in their view (by producing blocks in LC, casting votes in Goldfish). Eventually, honest validators converge on a leading tip, which then accumulates endorsements quickly from the majority of honest validators, so that no competing tip will ever be able to supersede it.

As a result, liveness of [38] requires steady participation for a sustained period, whereas Goldfish supports fast fluctuating participation. Unlike the works of Momose, Ren, and Malkhi, Goldfish achieves constant expected confirmation latency only under optimistic conditions of high participation and less than $\frac{1}{4}$ fraction adversarial validators (the fraction can be tuned to $\frac{1}{3}$ for both optimistic-fast and ‘standard’ confirmation in Goldfish).

Unlike Goldfish, the aforementioned protocols differ substantially from the LMD GHOST component of the current PoS Ethereum protocol. Furthermore, they are considerably more involved (e.g., in number of stages) than LMD GHOST or Goldfish. Thus, adoption of these protocols for PoS Ethereum would require substantial design and engineering effort, and is unlikely.

Highway [27] employs some of the techniques also found in Goldfish and in the PoS Ethereum beacon chain. For instance, message buffering first appeared in [27]. Furthermore, Highway aims to achieve flexible finality on a gradual scale using a monolithic protocol, and does not consider dynamic participation. In contrast, we follow the ebb-and-flow formulation [47] of the beacon chain requirements (and with it adopt the extension of the sleepy model to allow for periods of asynchrony [47]) with a gradual notion of (probabilistic) confirmation for the available full ledger (which is powered by Goldfish with the help of message buffering), and a binary notion of finality for the accountable final prefix (which is provided by a separate finality/accountability gadget). In particular, we adopt the modular approach described in [8, 47, 53, 49] to designing protocols that satisfy the ebb-and-flow property using finality gadgets and accountability gadgets (Fig. 2).

1.2 Outline

We recapitulate the model of synchronous networks with dynamic participation and asynchronous periods in Sec. 2, before describing our basic Goldfish protocol in Sec. 3, and an optimistic fast confirmation rule in Sec. 4. We analyze the protocol and prove the desired security properties in Sec. 5, before concluding in Sec. 6 with a case-study discussing implementation aspects of Goldfish in the context of Ethereum.

2 MODEL & PRELIMINARIES

We review cryptographic primitives, how to model environment and adversary, and the consensus security desiderata.

2.1 Preliminaries

2.1.1 Security parameters. We denote by λ and κ the security parameters associated with the cryptographic primitives employed by the Goldfish protocol, and with the Goldfish protocol itself, respectively. We say that an event happens with probability *negligible* in a security parameter μ , denoted by $\text{negl}(\mu)$, if its probability is $o(1/\mu^d)$ for all $d > 0$. Overall, we say that an event happens *with overwhelming probability (w.o.p.)* if it happens except with probability (w.p.) $\text{negl}(\kappa) + \text{negl}(\lambda)$.

2.1.2 Digital signatures.

Definition 1 (Informal, cf. [28, 5]). A *signature* scheme $\text{Sig} = (\text{Gen}, \text{Sign}, \text{Verify})$ consists of probabilistic poly-time (PPT) algorithms so that:

- $(\text{ssk}, \text{spk}) \leftarrow \text{Sig.Gen}(1^\lambda)$ creates a secret/public key pair.
- $\sigma \leftarrow \text{Sig.Sign}(\text{ssk}, m)$ creates a signature on a message.
- $\{0, 1\} \leftarrow \text{Sig.Verify}(\text{spk}, m, \sigma)$ verifies a signature.
- *Correctness:* With overwhelming probability, for all messages, $\text{Sig.Verify}(\text{spk}, m, \text{Sig.Sign}(\text{ssk}, m)) = 1$.
- *Security (existential unforgeability):* An adversary with access to spk and to a signing oracle $\text{Sig.Sign}(\text{ssk}, \cdot)$ cannot produce a valid (m, σ) other than via the oracle.

2.1.3 Verifiable random functions. A verifiable random function (VRF) [37] is used for leader election and subsampling of the validators within the Goldfish protocol.

Definition 2 (Informal, cf. [17, Sec. 3.2, Fig. 2], [18, 14]). A *verifiable random function* (VRF) scheme $\text{Vrf} = (\text{Gen}, \text{Prove}, \text{Verify})$ consists of PPT algorithms so that:

- $(\text{vsk}, \text{vpk}) \leftarrow \text{Vrf.Gen}(1^\lambda)$ creates a secret/public key pair.
- $(y, \pi) \leftarrow \text{Vrf.Prove}(\text{vsk}, x)$ obtains the output y of the VRF at input x , and the evaluation proof π .
- $\{0, 1\} \leftarrow \text{Vrf.Verify}(\text{vpk}, x, (y, \pi))$ verifies an evaluation.
- *Correctness:* With overwhelming probability, for all inputs, $\text{Vrf.Verify}(\text{vpk}, x, \text{Vrf.Prove}(\text{vsk}, x)) = 1$.
- *Uniqueness:* Per input x , there is only one output y : if $\text{Vrf.Verify}(\text{vpk}, x, (y, \pi)) = 1$ for $(y, \pi) = (y_1, \pi_1)$ and $(y, \pi) = (y_2, \pi_2)$, then $y_1 = y_2$.
- *‘Pseudorandomness’:* Conceptually, the VRF behaves like a random oracle that is *unpredictable* (i.e., without knowledge of vsk , the VRF output cannot be distinguished from a random string) and *verifiable* (i.e., given vpk , an alleged output of the VRF can be verified). For a formal definition, see [17, Sec. 3.2, Fig. 2].

2.2 Model

2.2.1 Validators. Goldfish is run among n validators, with identities $\text{id} \in [n] \triangleq \{1, \dots, n\}$. Each validator id generates a secret/public key pair $(\text{ssk}_{\text{id}}, \text{spk}_{\text{id}})$ and $(\text{vsk}_{\text{id}}, \text{vpk}_{\text{id}})$ for the signature and the VRF scheme, respectively. The public keys are known to all validators (*public-key infrastructure*, PKI).

2.2.2 Environment and adversary. Time is divided into discrete *rounds* and the validators have synchronized clocks.³ Validators receive transactions (txs) from the environment, and continuously output transaction ledgers to it (ch_r^{id} for validator id and round r).

³Bounded clock offsets can be lumped into the subsequently discussed network delay.

The environment allows validators to *broadcast* messages to each other. The adversary is a probabilistic poly-time (PPT) algorithm that can leverage three aspects of the model (*corruption*, *sleepiness*, and *network delay*) in its attempt to undermine consensus. We first discuss these three aspects, and then the limits of the adversary.

2.2.3 Corruption. The adversary chooses f validators to corrupt, hereafter called *adversarial* validators (non-corrupt validators are *honest*). The internal state of corrupted validators is handed over to the adversary, which can subsequently make them deviate from the protocol in an arbitrary and coordinated fashion (*Byzantine faults*). We define the adversarial fraction $\beta \triangleq f/n$.

2.2.4 Sleepiness. The adversary decides for each round and each honest validator whether it is *asleep* or not. Asleep validators do not execute the protocol (*temporary crash faults*). Messages delivered to an asleep validator get picked up by it only once the validator is no longer asleep. When a validator stops being asleep, it becomes *dreamy*. During this phase, it *joins* the protocol, usually over multiple rounds, using a special *joining procedure* specified by the protocol. Upon completion of this procedure, the honest validator becomes *awake* and then follows the ‘standard path’ of the protocol. Adversarial validators are always awake. The number of awake validators at any round is bounded below by a constant n_0 .

2.2.5 Network delay. Messages sent between validators are delivered with an adversarially determined delay that can differ for each recipient. Upon picking up a message (*i.e.*, as soon as no longer asleep after delivery), an honest validator re-broadcasts it.

2.2.6 Adversary limits. A *partially synchronous network in the sleepy model* [47] has a global stabilization time (GST), a global awake time (GAT), and a delay upper-bound Δ . GST and GAT are constants unknown to the honest validators chosen *adaptively* by the adversary, *i.e.*, as causal functions of the execution, whereas Δ is a constant known to the validators. Before GST, message delays are arbitrarily adversarial (*asynchronous*). After GST, message delays are subject to the delay upper bound Δ (*synchronous*). Similarly, before GAT, the adversary can set the sleep schedule for honest validators. After GAT, all honest validators are awake.

Message delays and sleeping schedule are chosen adaptively. For corruption, Goldfish supports two assumptions. Either, we require *mildly* adaptive corruption, where it takes 3Δ rounds for corruption to take effect, together with the constraint that for every round r , the number of adversarial validators at round r must be less than the number of honest awake validators at round $r - 3\Delta$. Or, analogously to earlier works [17, 1, 14], through the use of key evolving signature and VRF schemes, we allow for fully adaptive corruption, together with the constraint that for every round r , the number of adversarial validators at round r must be less than the number of honest awake validators at round r .

2.3 Consensus Security Desiderata

2.3.1 Security. We next formalize the notion of security *after a certain time*. Security is parameterized by κ , which, in the context of longest-chain protocols and Goldfish, represents the confirmation delay for transactions. In our analysis, we consider a finite time horizon T_{hor} that is polynomial in κ . We denote a consensus

protocol’s output ledger, *e.g.*, the Goldfish ledger, in the view of a validator i at round r by ch_r^i . We write $\text{ch}_1 \leq \text{ch}_2$ to express that the ledger ch_1 is a prefix of (or the same as) ledger ch_2 .

Definition 3 (Security). Let T_{conf} be a polynomial function of the security parameter κ . We say that a state machine replication protocol that outputs a ledger ch is *secure after time* T_{sec} , and has transaction confirmation time T_{conf} , iff:

- **Safety:** For any two rounds $r, r' \geq T_{\text{sec}}$, and any two honest validators i, j awake at rounds r and r' , respectively, either $\text{ch}_r^i \leq \text{ch}_{r'}^j$ or $\text{ch}_{r'}^j \leq \text{ch}_r^i$.
- **Liveness:** If a transaction has been received by some awake honest validator by some round $r \geq T_{\text{sec}}$, then for any round $r' \geq r + T_{\text{conf}}$ and any honest validator i awake at round r' , the transaction will be included in $\text{ch}_{r'}^i$.

The protocol satisfies \bar{f} -*safety* (\bar{f} -*liveness*) if it satisfies safety (liveness) as long as the number of adversarial validators f stays below \bar{f} for all rounds. Similarly, the protocol satisfies $1/2$ -*safety* ($1/2$ -*liveness*) if it satisfies safety (liveness) if the fraction of adversarial validators β is bounded above away from $1/2$ for all rounds.

2.3.2 Accountable safety. Accountable safety provides a *trust-minimizing* strengthening of safety, with the aim to hold validators accountable for their actions. In case of a safety violation in a protocol with accountable safety resilience $\bar{f} > 0$, one can, after collecting evidence from sufficiently many honest validators, generate cryptographic proof that identifies \bar{f} adversarial validators as protocol violators [55, 49]. By definition, the proof does not falsely accuse any honest validator, except with negligible probability.

2.3.3 The ebb-and-flow formulation. As Goldfish outputs a *dynamically available* ledger (*i.e.*, live under dynamic participation), by the availability-accountability dilemma [49], its output ledger cannot satisfy *accountable safety*. Similarly, it cannot satisfy safety under a partially synchronous network (*i.e.*, *finality*), by an analogue of the CAP theorem [23, 33]. However, Goldfish can be composed with an accountability gadget in order to obtain a separate prefix ledger that attains accountable safety under partial synchrony while staying consistent with the output of Goldfish [49]. Denoting the output of Goldfish as the available ledger ch_{ava} and that of the accountability gadget as the accountable final prefix ledger ch_{acc} , the desiderata are captured in the *ebb-and-flow formulation* [47]:

Definition 4 (Ebb-and-flow formulation [47, 49]).

- (1) **(P1: Accountability and finality)** Under a partially synchronous network in the sleepy model, the accountable final prefix ledger ch_{acc} has accountable safety resilience $n/3$ at all times, (except w.p. $\text{negl}(\lambda)$), and there exists a constant C such that ch_{acc} provides $n/3$ -liveness with confirmation time T_{conf} after round $\max(\text{GST}, \text{GAT}) + C\kappa$ (w.o.p.).
- (2) **(P2: Dynamic availability)** Under a synchronous network in the sleepy model (*i.e.*, for $\text{GST} = 0$), the available ledger ch_{ava} provides $1/2$ -safety and $1/2$ -liveness at all times (w.o.p.).
- (3) **(Prefix)** For each honest id and round r , $\text{ch}_{\text{acc},r}^{\text{id}} \leq \text{ch}_{\text{ava},r}^{\text{id}}$.

The accountable final prefix ledger ch_{acc} can experience liveness violations before GST or GAT, due to lack of timely communication among sufficiently many honest validators, but ch_{acc} remains

accountably safe throughout. The available ledger ch_{ava} can experience safety violations before GST, but remains live throughout. When conditions improve, ch_{acc} catches up with ch_{ava} . This ebb-and-flow behavior lends the formulation its name. Providing the irreconcilable properties in two separate but *consistent* ledgers provides a user-dependent resolution to the CAP theorem [23, 33].

3 PROTOCOL

We first describe the Goldfish protocol that is being proposed as a drop-in replacement for LMD GHOST in PoS Ethereum’s beacon chain. We then describe how Goldfish can be securely integrated with accountability and finality gadgets.

3.1 The Goldfish Protocol

The protocol (cf. Alg. 1) proceeds in *slots* of 3Δ rounds.

3.1.1 VRF-based lotteries. The VRF PKI enables cryptographic lotteries. A *lottery* (tag, thr) is defined by a fixed tag and threshold $\text{thr} \in [0, 1]$. Each validator id receives for each time slot t a lottery ticket (id, t) . To *open* the ticket, id computes

$$\varrho \triangleq (y, \pi) \leftarrow \text{Open}_{\text{id}}^{(\text{tag}, \text{thr})}(t) \triangleq \text{Vrf.Prove}(\text{vsk}_{\text{id}}, \text{tag} \parallel t). \quad (1)$$

The *opened ticket* (with *opening* ϱ) is *winning* for (tag, thr) iff:

$$\begin{aligned} & \text{IsWinning}^{(\text{tag}, \text{thr})}((\text{id}, t), \varrho) \\ & \triangleq (\varrho.y \leq \text{thr} 2^\lambda) \wedge \text{Vrf.Verify}(\text{vpk}_{\text{id}}, \text{tag} \parallel t, (\varrho.y, \varrho.\pi)). \end{aligned} \quad (2)$$

Finally, winning opened tickets are totally ordered by increasing *precedence*, $\text{Prio}(\varrho) \triangleq \frac{\varrho.y}{2^\lambda} \in [0, 1]$.

3.1.2 Data structures. *Blocks* and *votes* are central to Goldfish. A block $B \triangleq (\text{block}, (\text{id}, t), \varrho, h, \text{txs}, \sigma)$ consists of tag ‘block’, ticket (id, t) and opening ϱ to the $(\text{block}, \text{thr}_b)$ block production lottery, hash h committing to the new block’s parent block and transactions txs (as block ‘content’), and signature σ binding together block production opportunity and the block’s content. A special *genesis block* $B_0 \triangleq (\text{block}, (\perp, 0), \perp, \perp, \emptyset, \perp)$ is known to all validators.

A block B is *valid* iff:

$$\begin{aligned} \text{IsValid}(B_0) & \triangleq 1 \\ \text{IsValid}(B) & \triangleq \text{IsWinning}^{(\text{block}, \text{thr}_b)}((B.\text{id}, B.t), B.\varrho) \\ & \wedge \text{Sig.Verify}(\text{spk}_{B.\text{id}}, \text{block} \parallel B.h \parallel B.\text{txs}, B.\sigma) \\ & \wedge \text{IsValid}(*[B.h]) \wedge (B.t > *[B.h].t). \end{aligned} \quad (3)$$

Here, $*[B.h]$ means the parent block that $B.h$ commits to. The context within which these references get resolved is detailed with the different network message types below.

A vote $v \triangleq (\text{vote}, (\text{id}, t), \varrho, h, \sigma)$ consists of tag ‘vote’, ticket (id, t) and opening ϱ to the $(\text{vote}, \text{thr}_v)$ voting lottery, hash h committing to the block voted for (as vote ‘content’), and signature σ binding together voting opportunity and the vote’s content.

A vote v is *valid* iff:

$$\begin{aligned} \text{IsValid}(v) & \triangleq \text{IsWinning}^{(\text{vote}, \text{thr}_v)}((v.\text{id}, v.t), v.\varrho) \\ & \wedge \text{Sig.Verify}(\text{spk}_{v.\text{id}}, \text{vote} \parallel v.h, v.\sigma) \\ & \wedge \text{IsValid}(*[v.h]) \wedge (v.t \geq *[v.h].t). \end{aligned} \quad (4)$$

We call *block-vote-set* (short *bvset*) a set of blocks and votes. Commitments to blocks for the purpose of the references $v.h$ or $B.h$ are computed using $H(\cdot)$. For a *bvset* \mathcal{T} we denote by $\mathcal{T}[h]$ the

block $B \in \mathcal{T}$ with $H(B) = h$, and \perp if non-existent.

In Goldfish, votes and blocks are encapsulated and exchanged in two network message types, *pieces* and *proposals*. A piece $M \triangleq (\text{piece}, x)$ consists of tag ‘piece’ and for payload x either a vote or a block, and is *valid* iff:

$$\text{IsValid}(M) \triangleq \text{IsValid}(M.x). \quad (6)$$

Pieces are used to propagate blocks and votes and abstract Ethereum’s peer-to-peer broadcast object propagation. In determining a piece’s validity, block references $*[.]$ are resolved with respect to the *bvset* \mathcal{T} each validator maintains as part of its state, see Sec. 3.1.3. If a validator does not have any matching block in \mathcal{T} , it cannot currently determine the piece’s validity. It then keeps the piece ‘in limbo’ for re-examination until its (in-)validity is established.⁴

A proposal $P \triangleq (\text{propose}, \mathcal{T}, B, \sigma)$ consists of tag ‘propose’, *bvset* \mathcal{T} and block B (as proposal content), and signature σ tying the proposal to the block production opportunity of B .

Thus, a proposal P is *valid* iff:

$$\begin{aligned} \text{IsValid}(P) & \triangleq \text{IsValid}(P.B) \wedge \text{IsConsistent}(P.\mathcal{T} \cup \{P.B\}) \\ & \wedge \text{Sig.Verify}(\text{spk}_{P.B.\text{id}}, \text{propose} \parallel P.\mathcal{T} \parallel P.B, P.\sigma) \\ & \wedge (\forall x \in P.\mathcal{T}: \text{IsValid}(x) \wedge (x.t < P.B.t)) \end{aligned} \quad (7)$$

where $\text{IsConsistent}(\mathcal{T})$ is a predicate that is satisfied on a *bvset* \mathcal{T} iff $B_0 \in \mathcal{T}$ and for every vote and block in \mathcal{T} the referenced target/parent block is also in \mathcal{T} . We call a *bvset* \mathcal{T} with $\text{IsConsistent}(\mathcal{T})$ a *block-vote-tree* (short *bvtree*). In determining the validity of proposal P , block references $*[.]$ are resolved with respect to $P.\mathcal{T}$.

3.1.3 Validator state. Each validator keeps track of the current time slot t . It also maintains a *bvtree* \mathcal{T} based on which it takes consensus decisions and actions. Finally, each validator maintains a *buffer* \mathcal{B} of network messages (*i.e.*, pieces and proposals) that ‘sits between’ network and consensus protocol.

3.1.4 Message handling. Recall that messages are delivered to validators irrespective of their sleep status. However, validators pick up delivered messages only once awake. Invalid messages are discarded. If a piece’s validity cannot be determined due to missing references, it is held in limbo until its (in-)validity is determined. Pieces and proposals ‘from the future’ (*i.e.*, in time slot t , pieces M with $M.x.t > t$ and proposals P with $P.B.t > t$) are also held in limbo. Upon picking up a *valid non-in-limbo* message from the network, the validator re-broadcasts it, and adds it to \mathcal{B} . If the message is a proposal P , the validator also re-broadcasts the blocks and votes in $P.\mathcal{T} \cup \{P.B\}$ as pieces, and adds those pieces to its own \mathcal{B} .

3.1.5 Message buffering. The validator unpacks messages from \mathcal{B} and merges them into \mathcal{T} in a way that preserves $\text{IsConsistent}(\mathcal{T})$. For this purpose, $\text{MERGE}(\mathcal{T}, \mathcal{B})$ outputs the largest *bvtree* \mathcal{T}' that is a subset of the union of \mathcal{T} and the pieces in \mathcal{B} . Merging of \mathcal{B} into \mathcal{T} takes place only at carefully chosen points in time as explicitly instructed (Alg. 1, ll. 19, 27). This *message buffering* is a key ingredient of Goldfish. First, Δ rounds into a slot, each awake validator identifies a slot leader and merges the *bvtree* proposed by

⁴Vote expiry (Sec. 3.1.6) and reorg resilience (Thm. 3) enable timely garbage collection of pieces with missing referenced blocks.

Algorithm 1 Goldfish executed by validator id with signature secret/public key $(\text{ssk}_{\text{id}}, \text{spk}_{\text{id}})$, VRF secret/public key $(\text{vsk}_{\text{id}}, \text{vpk}_{\text{id}})$, bvtree \mathcal{T} and buffer \mathcal{B} . Here, notation ‘at’ means executing the code block at the specified round, ch^{id} denotes the Goldfish chain momentarily confirmed at id . For $\text{GHOST-Eph}(\mathcal{T}, t)$, see Alg. 2.

```

1:  $(\mathcal{B}, \mathcal{T}, t) \leftarrow (\emptyset, \emptyset, 0)$   $\triangleright$  Initialize buffer  $\mathcal{B}$  and bvtree  $\mathcal{T}$ 
2:  $\triangleright$  At all rounds, only valid messages not from later than  $t$  are picked up
   from the network, re-broadcast, and put into  $\mathcal{B}$  as specified in Sec. 3.1.4
3: for  $t = 1, 2, \dots$  do  $\triangleright$  Slots
4:   at  $3\Delta t$  do  $\triangleright$  PROPOSE phase
5:      $\varrho \leftarrow \text{Open}_{\text{id}}^{(\text{block}, \text{thr}_b)}(t)$   $\triangleright$  Check if eligible to propose
6:     if  $\text{IsWinning}^{(\text{block}, \text{thr}_b)}((\text{id}, t), \varrho)$  then
7:        $\mathcal{T}' \leftarrow \text{MERGE}(\mathcal{T}, \mathcal{B})$   $\triangleright$  Bvtree to propose
8:        $B \leftarrow \text{GHOST-Eph}(\mathcal{T}', t-1)$   $\triangleright$  Parent block
9:        $\sigma \leftarrow \text{Sig.Sig}(\text{ssk}_{\text{id}}, \text{block} \parallel H(B) \parallel \text{txs})$ 
10:       $B \leftarrow (\text{block}, (\text{id}, t), \varrho, H(B), \text{txs}, \sigma)$   $\triangleright$  New block
11:       $\sigma \leftarrow \text{Sig.Sig}(\text{ssk}_{\text{id}}, \text{propose} \parallel \mathcal{T}' \parallel B)$ 
12:      Broadcast (propose,  $\mathcal{T}'$ ,  $B$ ,  $\sigma$ )  $\triangleright$  Propose
13:   at  $3\Delta t + \Delta$  do  $\triangleright$  VOTE phase
14:      $\triangleright$  Filter for proposals from slot  $t$ 
15:      $\mathcal{B}' \leftarrow \{(\mathcal{T}', B) \mid (\text{propose}, \mathcal{T}', B, \cdot) \in \mathcal{B} \wedge B.t = t\}$ 
16:      $\triangleright$  Identify the leader of slot  $t$  and its proposal
17:      $(\mathcal{T}^*, B^*) \leftarrow \arg \min_{(\mathcal{T}', B) \in \mathcal{B}'}$   $\text{Prio}(B, \varrho)$ 
18:      $\triangleright$  Merge own buffer and that of the leader into own bvtree
19:      $\mathcal{T} \leftarrow \text{MERGE}(\mathcal{T}, \mathcal{T}^* \cup \{B^*\})$ 
20:      $\varrho \leftarrow \text{Open}_{\text{id}}^{(\text{vote}, \text{thr}_v)}(t)$   $\triangleright$  Check if eligible to vote
21:     if  $\text{IsWinning}^{(\text{vote}, \text{thr}_v)}((\text{id}, t), \varrho)$  then
22:        $B \leftarrow \text{GHOST-Eph}(\mathcal{T}, t-1)$   $\triangleright$  Target block
23:        $\sigma \leftarrow \text{Sig.Sig}(\text{ssk}_{\text{id}}, \text{vote} \parallel H(B))$ 
24:        $v \leftarrow (\text{vote}, (\text{id}, t), \varrho, H(B), \sigma)$   $\triangleright$  New vote
25:       Broadcast (piece,  $v$ )  $\triangleright$  Vote
26:   at  $3\Delta t + 2\Delta$  do  $\triangleright$  CONFIRM phase
27:      $\mathcal{T} \leftarrow \text{MERGE}(\mathcal{T}, \mathcal{B})$   $\triangleright$  Merge buffer and bvtree
28:      $B \leftarrow \text{GHOST-Eph}(\mathcal{T}, t)$   $\triangleright$  Canonical GHOST-Eph chain
29:      $\text{ch}^{\text{id}} \leftarrow B^{\lceil \kappa}$   $\triangleright$  Output ledger:  $B$ 's  $\kappa$ -deep prefix in terms of slots
    
```

Algorithm 2 GHOST-Eph fork-choice rule.

```

1:  $\text{CHLDRN}(\mathcal{T}, B) \triangleq \{B' \in \mathcal{T} \mid B'.h = H(B)\}$ 
2:  $\text{VTS}(\mathcal{T}, B, t) \triangleq \{|\text{id}' \mid (\text{vote}, (\text{id}', t), \cdot, h, \cdot) \in \mathcal{T} \wedge B \leq \mathcal{T}[h]\}$ 
3: function  $\text{GHOST-Eph}(\mathcal{T}, t)$ 
4:    $B \leftarrow B_0$   $\triangleright$  Start fork-choice at genesis block
5:   forever do
6:      $\triangleright$  Choose the heaviest subtree rooted at one of the children
       blocks  $B'$  of  $B$ , by number of validators that have cast a vote for slot  $t$ 
       into the subtree rooted at  $B$ ;  $B' = \perp$  if  $\text{CHLDRN}(\mathcal{T}, B) = \emptyset$ 
7:      $B' \leftarrow \arg \max_{B' \in \text{CHLDRN}(\mathcal{T}, B)} \text{VTS}(\mathcal{T}, B, t)$ 
8:     if  $B' = \perp$  then return  $B$ 
9:      $B \leftarrow B'$ 
    
```

the leader into its bvtree (Alg. 1, l. 19). Second, 2Δ rounds into a slot, each awake validator merges its buffer into its bvtree (Alg. 1, l. 27).

3.1.6 Vote expiry. To determine the canonical chain, validators use the GHOST-Eph fork-choice function with ephemeral votes (Alg. 2). The function takes a bvtree \mathcal{T} and slot t as input, and finds the canonical GHOST-Eph chain determined by the votes within \mathcal{T} that were cast for slot t . More specifically, starting at the genesis block, the function iterates over a sequence of blocks from the bvtree,

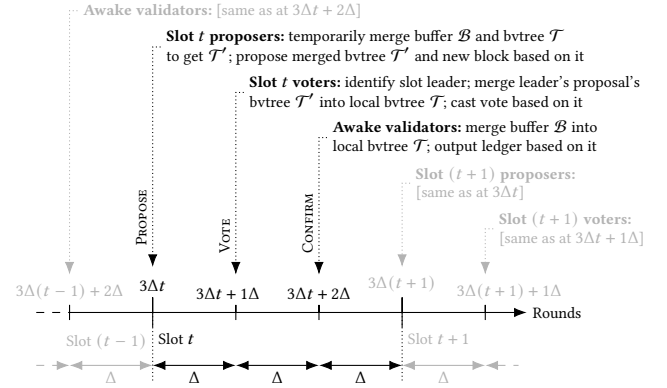


Figure 1: Throughout the execution, validators buffer received proposals and pieces, and merge the blocks and votes contained therein into their bvtrees only as explicitly instructed. Goldfish has time slots of three phases of Δ rounds each. Each time slot has proposers (one of which will later be recognized as the slot’s leader) and a committee of voters. **PROPOSE:** At the start of a slot, *proposers* temporarily merge their buffers into their local bvtrees, and propose the merger and a new block based on it. **VOTE:** One-thirds into a slot, *voters* identify the slot’s leader’s proposal, merge the proposed bvtree into their local bvtrees, and cast a vote based on their local bvtrees. **CONFIRM:** Two-thirds into a slot, *all awake validators* merge their buffers into their local bvtrees, and confirm a ledger based on their local bvtrees.

selecting as the next block the child of the current block with the maximum number of validators that have cast a slot t vote for a block within the child’s subtree. This continues until it reaches a leaf of the bvtree, and outputs a complete chain from leaf to root. The fork-choice rule ignores votes from other than slot t in its decision (votes are *ephemeral*), lending GHOST-Eph its name.

3.1.7 The complete Goldfish protocol. The three phases (PROPOSE, VOTE, CONFIRM) of each slot t are shown in Fig. 1. We describe them from the perspective of an awake honest validator id .

PROPOSE: At round $3\Delta t$, id checks if its lottery ticket (id, t) is winning for $(\text{block}, \text{thr}_b)$ (Alg. 1, l. 6). If so, id temporarily merges its bvtree with its buffer (Alg. 1, l. 7), identifies the GHOST-Eph chain tip using only slot $t-1$ votes (Alg. 1, l. 8), and proposes its temporary bvtree and a new block based on it (Alg. 1, l. 12).

VOTE: At round $3\Delta t + \Delta$, id identifies as *leader* for slot t any one of the proposals with smallest precedence (Alg. 1, l. 17). It merges the leading proposal’s bvtree into its bvtree \mathcal{T} (Alg. 1, l. 19). Validator id then checks if its lottery ticket (id, t) is winning for $(\text{vote}, \text{thr}_v)$ (Alg. 1, l. 21). If so, id identifies the GHOST-Eph chain tip using only slot $t-1$ votes (Alg. 1, l. 22), and votes for it (Alg. 1, l. 25).

CONFIRM: At round $3\Delta t + 2\Delta$, id merges its buffer \mathcal{B} into its bvtree \mathcal{T} (Alg. 1, l. 27). It then identifies the GHOST-Eph chain tip using only slot t votes (Alg. 1, l. 28), and outputs as confirmed ledger ch^{id} the transactions of those blocks in the GHOST-Eph chain that are from slots $\leq t - \kappa$ (κ -deep in time’, Alg. 1, l. 29). Since the Goldfish ledger in view of an awake honest validator id is only updated at




Figure 2: An accountability/finality gadget (AG/FG; a.k.a. *overlay*) checkpoints decisions of the dynamically available protocol Goldfish (a.k.a. *underlay*). A feedback loop ensures that Goldfish respects earlier checkpoints. This construction satisfies the *ebb-and-flow* design objective of PoS Ethereum, to produce an available full ledger that is live under dynamic participation of validators, and a prefix ledger that is accountably safe under network partition [47, 49].

this point, we may view the ledger as indexed by time slot t : ch_t^{id} .

Joining procedure: At each round, each honest validator is either asleep, dreamy or awake (Sec. 2.2.4). Once an honest validator is no longer asleep, it remains *dreamy* until the round of the next CONFIRM phase.⁵ While being dreamy, the validator does not follow Alg. 1, except for relaying messages. With the CONFIRM phase, the validator returns to being *awake* and fully resumes Alg. 1.

3.1.8 Key mechanism of Goldfish. *Message buffering* ensures that if in slot t the leading proposal is honest, then all honest voters in t will vote for the proposed block. This is because in PROPOSE, the leader’s temporary bvtree is a superset of all honest validators’ bvtrees, and thus in VOTE all honest validators adopt that leader’s bvtree. *Vote expiry* (together with majority honest validators) ensures that if in slot t all honest voters have voted into the subtree rooted at some block B , then all honest voters in slot $t + 1$ will also vote into the subtree rooted at B . An inductive argument immediately yields reorg resilience of Goldfish. Furthermore, w.o.p., every interval of κ slots has at least one honest leading proposer. The prefix of that proposal stabilizes (by reorg resilience), and the proposal includes unconfirmed transactions, leading to safety and liveness of the κ -deep confirmation rule.

3.1.9 Validator replaceability. Due to subsampling, once a validator takes an action in Goldfish, it does not play any further role, at least for a long time. As a result, Goldfish supports player replaceability [22, 14, 56] and can withstand a mildly adaptive adversary (Sec. 2.2.6). Analogously to earlier works [17, 1, 14, 25], fully adaptive corruption can be allowed through the use of key evolving signature and VRF schemes. In both cases, the adversary cannot corrupt an honest validator and make it send conflicting protocol messages ‘fast enough’ to harm the protocol execution.

3.2 Goldfish with Accountability Gadgets

For the composition of Goldfish with accountability gadgets and finality gadgets, we follow the construction of [49, 53] (Fig. 2, Alg. 4). In this construction, a partially synchronous accountably-safe consensus protocol such as Streamlet, Tendermint, or HotStuff [13, 6,

⁵We assume that messages arrive at validators while asleep (Sec. 3.1.4). To allow for extra time to download messages missed during sleep, dreaminess can be extended accordingly, but should always end at a CONFIRM phase.

Algorithm 3 GHOST-Eph (cf. Alg. 2) modified (green) to respect the latest checkpoint B . See Alg. 2 for CHLDRN and VTS.

```

1: function GHOST-EPH( $\mathcal{T}, t, B$ )
2:   ▶ Start fork-choice from latest checkpoint  $B$ 
3:   forever do
4:     ▶ Choose the heaviest subtree rooted at one of the children
       blocks  $B'$  of  $B$ , by number of validators that have cast a vote for slot  $t$ 
       into the subtree rooted at  $B$ ;  $B' = \perp$  if  $\text{CHLDRN}(\mathcal{T}, B) = \emptyset$ 
5:      $B' \leftarrow \arg \max_{B' \in \text{CHLDRN}(\mathcal{T}, B)} \text{VTS}(\mathcal{T}, B, t)$ 
6:     if  $B' = \perp$  then return  $B$ 
7:      $B \leftarrow B'$ 
    
```

Algorithm 4 Composition of Goldfish and accountability gadget (cf. Fig. 2, [49, Alg. 1]), executed by validator id . Here, Goldfish (cf. Alg. 1) uses a modified GHOST-Eph rule (Alg. 3), starting the recursion from the latest checkpoint, *i.e.*, the last block of $ch_{\text{acc}}^{\text{id}}$. Throughout, Goldfish maintains the available chain $ch_{\text{ava}}^{\text{id}}$. RUNACCOUNTABILITYGADGET attempts the next iteration of the gadget, where valid checkpoint candidates are determined using $ch_{\text{ava}}^{\text{id}}$. Iterations may fail (\perp), *e.g.*, if the gadget invokes a malicious leader.

```

1:  $ch_{\text{acc}}^{\text{id}} \leftarrow B_0$  ▶ ‘Zero-th’ checkpoint: Goldfish’s genesis block
2: for  $c = 1, 2, \dots$  do ▶ Checkpoint iterations
3:    $\text{checkpoint} \leftarrow \text{RUNACCOUNTABILITYGADGET}(ch_{\text{ava}}^{\text{id}})$ 
4:   if  $\text{checkpoint} \neq \perp$  then
5:      $ch_{\text{acc}}^{\text{id}} \leftarrow \text{checkpoint}$  ▶ Update latest checkpoint
6:     Sleep for  $T_{\text{chkpt}}$  rounds
    
```

58, 48], with accountable safety resilience of $n/3$ out of n validators, is used to determine checkpoints of Goldfish’s output ledger. To ensure that Goldfish respects earlier checkpoints, its fork-choice rule is modified to respect earlier checkpoint decisions (cf. Alg. 3). The most recent checkpoint forms the accountably-safe finalized prefix ledger ch_{acc} , while Goldfish’s output forms the dynamically available full ledger ch_{ava} (cf. *ebb-and-flow*, Def. 4). Since Goldfish now respects checkpoints, $ch_{\text{acc}} \leq ch_{\text{ava}}$, as required.

The full protocol proceeds in checkpointing iterations (cf. Alg. 4). Iterations may fail, *e.g.*, when the consensus protocol of the gadget invokes a malicious leader, or during asynchrony before GST, or while many validators are asleep before GAT. Successful checkpoint iterations are separated by at least T_{chkpt} rounds of inactivity of the gadget. In App. A.2, we apply the techniques of earlier analyses [49, 53] to the combination of Goldfish and the accountability gadget, to show how to tune T_{chkpt} as a function of the network delay Δ and the confirmation parameter κ , and to formally prove that the combination satisfies the *ebb-and-flow* desiderata.

4 OPTIMISTIC FAST CONFIRMATIONS

The Goldfish protocol described in Sec. 3.1 has reorg resilience as an advantage over protocols which use blocks as votes (*e.g.*, longest chain [39, 51, 32], GHOST [57]). On the other hand, Goldfish’s κ -slots deep confirmation rule, which leads to $\Theta(\kappa)$ latency in both the worst and the expected case, falls behind many propose-and-vote style protocols that achieve *constant* expected latency (*e.g.*, PBFT [12], Tendermint [6], HotStuff [58], Streamlet [13]). By introducing a fast confirmation rule and adding a FAST-CONFIRM phase

Algorithm 5 Goldfish executed by validator id , using both (optimistic) fast confirmation and standard confirmation (cf. Alg. 1). See Alg. 2 for VTs.

```

1:  $(\mathcal{B}, \mathcal{T}, t) \leftarrow (\emptyset, \emptyset, 0)$  ▷ Initialize buffer  $\mathcal{B}$  and btree  $\mathcal{T}$ 
2: ▷ At all rounds, only valid messages not from later than  $t$  are picked up from the network, re-broadcast, and put into  $\mathcal{B}$  as specified in Sec. 3.1.4
3: for  $t = 1, 2, \dots$  do ▷ Slots
4:   at  $4\Delta t$  do ▷ PROPOSE phase
5:     Same as PROPOSE phase in Alg. 1
6:   at  $4\Delta t + \Delta$  do ▷ VOTE phase
7:     Same as VOTE phase in Alg. 1
8:   at  $4\Delta t + 2\Delta$  do ▷ FAST-CONFIRM phase
9:      $\mathcal{T} \leftarrow \text{MERGE}(\mathcal{T}, \mathcal{B})$  ▷ Merge buffer and btree
10:     $ch_{\text{fast}}^{\text{id}} \leftarrow \arg \max_{B \in \mathcal{T}: |\text{VTs}(\mathcal{T}, B, t)| \geq n(\frac{3}{4} + \frac{\epsilon}{2})\text{thr}_v} |B|$ 
11:   at  $4\Delta t + 3\Delta$  do ▷ CONFIRM phase
12:      $\mathcal{T} \leftarrow \text{MERGE}(\mathcal{T}, \mathcal{B})$  ▷ Merge buffer and btree
13:      $B \leftarrow \text{GHOST-EPH}(\mathcal{T}, t)$  ▷ Canonical GHOST-Eph chain
14:      $ch^{\text{id}} \leftarrow \arg \max_{ch \in \{ch_{\text{fast}}^{\text{id}}, B^{\lceil \kappa} \}} |ch|$  ▷ Output Goldfish ledger
    
```

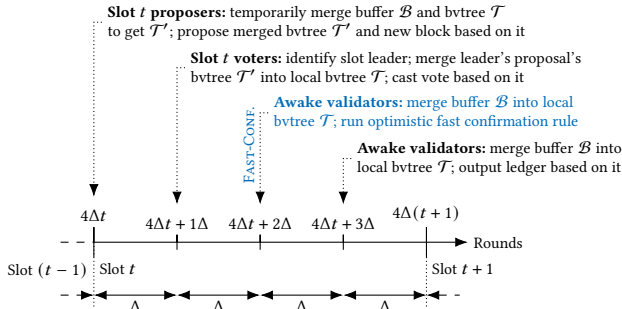


Figure 3: To enable optimistic fast confirmations, a **FAST-CONFIRM phase** (blue) of Δ rounds is inserted between **VOTE** and **CONFIRM** phase (cf. Fig. 1). **FAST-CONFIRM: Two-fourth** into a slot, **all awake validators merge their buffers into their local bvtrees, and run the optimistic fast confirmation rule based on their local bvtrees.**

to the Goldfish slot structure, we can achieve constant expected confirmation latency *under optimistic conditions*, i.e., under high participation and honest supermajority (Fig. 3, Alg. 5). In particular, validators can now confirm blocks proposed by honest leaders immediately, in the FAST-CONFIRM phase of the slot, under optimistic conditions. The κ -deep confirmation rule (Alg. 4, l. 29), to which we from now on refer as *standard* confirmation rule, still applies and guarantees liveness when optimistic conditions do not hold.

Fast confirmation phase. Slots now consist of 4Δ rounds and four phases (PROPOSE, VOTE, FAST-CONFIRM, CONFIRM), with the addition of phase FAST-CONFIRM at round $4\Delta t + 2\Delta$ (Fig. 3, Alg. 5).

In FAST-CONFIRM, a validator id first merges its buffer into its btree \mathcal{T} (Alg. 5, l. 9). It then marks a block B as fast confirmed if $|\text{VTs}(\mathcal{T}, B, t)| \geq n(\frac{3}{4} + \frac{\epsilon}{2})\text{thr}_v$ for some $\epsilon > 0$,⁶ and updates $ch_{\text{fast}}^{\text{id}}$ to the highest fast confirmed block (Alg. 5, l. 10). The other three phases are unchanged, other than for how the Goldfish ledger is

⁶The parameter $\epsilon > 0$ can be made arbitrarily small in the limit $n \rightarrow \infty$.

output in CONFIRM (Alg. 5, l. 14). Validator id outputs the highest of $ch_{\text{fast}}^{\text{id}}$ and the κ -deep prefix $B^{\lceil \kappa}$, where B is the tip of its canonical chain GHOST-Eph (cf. Alg. 4, l. 29, where $B^{\lceil \kappa}$ is output instead). For simplicity, we have omitted in Alg. 5 the mechanism to avoid temporary ledger ‘roll back’ (to ensure $\forall id, t' \geq t: ch_t^{\text{id}} \leq ch_{t'}^{\text{id}}$).

The reason for the extra Δ rounds, as opposed to just running the fast confirmation rule in the CONFIRM phase, is to guarantee that, whenever an honest validator fast confirms a block, all honest awake validators see the votes responsible for fast confirmation by the time their bvtrees are updated for the last time in the given slot, at round $4\Delta t + 3\Delta$. This ensures that the fast confirmed block eventually enters the Goldfish ledger in the view of all honest awake validators (Thm. 5), which in turn implies that fast confirmations are safe (Thm. 6). The security of the protocol with the fast confirmation rule is proven in Sec. 5.3.

Joining procedure. The joining protocol is conceptually unchanged. Once a validator stops being asleep, it remains dreamy until the next CONFIRM phase, at which point it turns fully awake and resumes protocol execution by merging its buffer with its bvtrees 3Δ rounds into a slot (phase CONFIRM, cf. Alg. 5, l. 12).

Composition with accountability and finality gadgets. When composing accountability gadgets and Goldfish with the fast confirmation rule, we stipulate that the validators input to the gadget only those blocks confirmed via the standard confirmation rule ($\text{GHOST-EPH}(\mathcal{T}, t)^{\lceil \kappa}$) in their view. This is necessary to ensure that all honest validators promptly agree on the confirmation status of the blocks input to the gadget for checkpointing, which in turn is a prerequisite for the liveness of the accountable final prefix ledger ch_{acc} . Otherwise, it is possible that a block fast confirmed by one honest validator might not become confirmed in the view of another honest validator until after κ slots, stalling the checkpointing process of the accountability gadget for that block. Thus, the fast confirmation rule is primarily for reducing the latency of the available ledger ch_{ava} , and does not affect the time for a block to enter the accountable final prefix ledger ch_{acc} .

Trading-off safety and liveness resiliences. With fast confirmation, Goldfish has two ‘parallel’ confirmation rules, ‘fast’ and ‘standard’. The overall protocol is safe only when both rules are safe, and live when one of the rules is live. To match the $\frac{1}{2}$ -safety of standard confirmation, fast confirmation’s ‘quorum’ was chosen as $n(\frac{3}{4} + \frac{\epsilon}{2})\text{thr}_v$ votes. With this parameterization, however, Goldfish cannot guarantee any fast confirmation in the presence of $\frac{n}{4} + \frac{\epsilon}{2}$ adversarial validators. It is possible to vary the quorum to trade-off safety and liveness of the fast path (and thereby of the overall protocol). With a quorum of $n(\frac{2}{3} + \frac{\epsilon}{2})\text{thr}_v$ for fast confirmation, Goldfish satisfies safety and liveness with $T_{\text{conf}} = \Theta(\kappa)$ if $\beta < \frac{1}{3} - \frac{3}{2}\epsilon$, and liveness with constant expected confirmation time if all validators are awake.

5 ANALYSIS

5.1 Goldfish

In the subsequent analysis, a valid proposal P (cf. Sec. 3.1.2) is for slot t iff $t = P.B.t$, and it has precedence p iff $p = \text{Prio}(P.B.\rho)$. A validator id is *eligible to propose at slot* t if its ticket (id, t) is winning for the lottery $(\text{block}, \text{thr}_b)$. Similarly, a validator id is *eligible to vote at*

slot t if its ticket (id, t) is winning for the lottery $(\text{vote}, \text{thr}_v)$. Recall that awake honest validators consider the proposal with lowest precedence received by $3\Delta t + \Delta$ the *leader* of slot t (Alg. 1, l. 16). We hereafter use blocks and the sequences of blocks they induce via the parent-block chain relation interchangeably. A block B_1 is a *descendant* (resp., *ancestor*) of block B_2 iff the underlying chains satisfy $B_2 \leq B_1$ (resp., $B_1 \leq B_2$). Two blocks B_1, B_2 are *conflicting* if B_1 is neither an ancestor nor a descendant of B_2 .

Let A_r and H_r denote the number of adversarial and honest validators awake at round r , respectively. Our security theorems hold for *compliant executions* that satisfy the relations on A_r and H_r laid out in Sec. 2.2.6:

Definition 5. In the absence of key-evolving cryptographic primitives (signatures and VRFs), an execution is (γ, τ) -*compliant* iff:

- $\forall r: \frac{A_r}{A_r + H_{r-\tau}} \leq \beta < \gamma - \epsilon$.
- The corruption is mildly adaptive: If the adversary decides to corrupt an honest validator at round r , then the validator becomes adversarial no earlier than at round $r + \tau$.

With key-evolving primitives, an execution is *compliant* iff:

- $\forall r: \frac{A_r}{A_r + H_r} \leq \beta < \gamma - \epsilon$.

Moreover, in both cases, $H_r > \gamma n_0 = \Theta(\kappa)$ for all rounds r , and the time horizon T_{hor} of the protocol execution satisfies $T_{\text{hor}} = \text{poly}(\kappa)$.

Intuitively, in compliant executions, honest voters outnumber adversarial voters (as long as votes have not yet expired); and every long interval of slots contains at least one slot in which all honest validators recognize the same honest validator as the slot leader.

Lemma 1. *Suppose the Goldfish execution is $(\frac{1}{2}, 3\Delta)$ -compliant.*

Then, w.o.p., for every slot t , adversarial validators at round $3\Delta(t + 1) + \Delta$ eligible to vote at slot t are less than honest validators awake at round $3\Delta t + \Delta$ and eligible to vote at slot t .⁷

Also w.o.p., all slot intervals of length κ have at least one slot t where an honest validator is recognized as the slot t leader by all awake honest validators at round $3\Delta t + \Delta$.

Lem. 1's proof uses correctness, uniqueness and pseudorandomness of VRF-based lotteries, and is given in App. A.1.

The main security results are as follows:

Theorem 1. *Suppose a $(\frac{1}{2}, 3\Delta)$ -compliant execution of Goldfish in the synchronous sleepy network model of Sec. 2.2, and validator id with proposal P^* is recognized as the leader of a slot t by all awake honest validators at round $3\Delta t + \Delta$ (Alg. 1, l. 16).*

Then, w.o.p., $P^.B \leq B$ for any B identified in Alg. 1, ll. 8, 22, 28 by any awake honest validator in any round $r \geq 3\Delta t + 2\Delta$.*

Theorem 2 (Security). *Suppose a $(\frac{1}{2}, 3\Delta)$ -compliant execution of Goldfish in the synchronous sleepy network model. Then, w.o.p., Goldfish is secure with transaction confirmation time $T_{\text{conf}} = 2\kappa + 2$ slots.*

Theorem 3 (Reorg resilience). *Suppose a $(\frac{1}{2}, 3\Delta)$ -compliant execution of Goldfish in the synchronous sleepy network model, and*

⁷For concreteness, the Ethereum validator set has over 400,000 validators as of 5-Sept-2022. Suppose we subsample with $\text{thr}_b = \frac{1}{32}$, i.e., with committee size unchanged in expectation, and that $\epsilon = 0.05$, i.e., that 55% of validators are assumed to be honest. Then, the probability of an adversarial majority at a single slot (assuming perfect randomness) is roughly $4 \cdot 10^{-15}$. There are 2628000 slots in a year, so the expected number of years before seeing an adversarial majority at a slot is $\frac{4 \cdot 10^{15}}{2628000} \approx 10^7$ years.

validator id with proposal P^ is recognized as the leader of a slot t by all awake honest validators at round $3\Delta t + \Delta$ (Alg. 1, l. 16).*

Then, w.o.p.,

$$\exists r': \forall r \geq r': \forall \text{id}: P^*.B \leq \text{ch}_r^{\text{id}}, \quad (8)$$

where ch_r^{id} denotes Goldfish's ledger at validator id and round r . In particular, $r' = 3\Delta(t + \kappa) + 2\Delta$ satisfies the above.

We first prove Thms. 2 and 3 from Thm. 1 and Lem. 1. Then, we prove Thm. 1 from Lems. 1, 2 and 3 in the remainder of the section.

PROOF OF THM. 2. By Lem. 1, w.o.p., all slot intervals of length κ have at least one slot t , where an honest validator with proposal P^* is recognized as the slot leader by all awake honest validators at round $3\Delta t + \Delta$, and, by Thm. 1, $P^*.B \leq B$ for any B identified in Alg. 1, ll. 8, 22, 28 by any awake honest validator in any $r \geq 3\Delta t + 2\Delta$.

Liveness: A transaction tx is input to an honest validator at some round r . At most 6Δ rounds (i.e., 2 slots) later the transaction is propagated to all honest validators and we have reached the beginning of a slot t_0 . For the next κ slots all honest proposers will include tx if they extend a tip whose chain does not include tx yet. By the earlier argument, one of these proposals will be an ancestor of any B identified in Alg. 1, ll. 8, 22, 28 by any awake honest validator in any $r' \geq 3\Delta(t_0 + \kappa) + 2\Delta$. From κ slots later onwards, all awake honest validators include the transaction in their ledger (Alg. 1, l. 29). Thus, Goldfish is live with $T_{\text{conf}} = 2\kappa + 2$ slots.

Safety: Pick any two honest validators id_1 and id_2 , and two slots t_1 and $t_2 \geq t_1$. By the earlier argument, there exists a block B' proposed (by an honest validator) at some slot $t' \in [t_1 - \kappa, t_1]$ such that $B' \leq B$ for any B identified in Alg. 1, ll. 8, 22, 28 by any awake honest validator in any $r' \geq 3\Delta t' + 2\Delta$. As $t' \geq t_1 - \kappa$ but by Goldfish's confirmation rule blocks in $\text{ch}_{t_1}^{\text{id}_1}$ are from no later than $t_1 - \kappa$, $\text{ch}_{t_1}^{\text{id}_1} \leq B$. Similarly, if $t' \geq t_2 - \kappa$, then $\text{ch}_{t_2}^{\text{id}_2} \leq B$; otherwise, $B \leq \text{ch}_{t_2}^{\text{id}_2}$. In both cases, either $\text{ch}_{t_1}^{\text{id}_1} \leq \text{ch}_{t_2}^{\text{id}_2}$ or $\text{ch}_{t_2}^{\text{id}_2} \leq \text{ch}_{t_1}^{\text{id}_1}$. \square

PROOF OF THM. 3. By Thm. 1, $P^*.B \leq B$ for any B identified in Alg. 1, ll. 8, 22, 28 by any awake honest validator in any $r \geq 3\Delta t + 2\Delta$. From κ slots later onwards, all awake honest validators include the transaction in their ledger (Alg. 1, l. 29). \square

Proof of Thm. 1 follows from Lems. 1, 2 and 3, and is provided at the end of this section. The structure of the argument is inductive: Lem. 2 shows that in a slot t with honest leader, all honest voters vote for the leader's proposal. Lem. 3 shows that if in slot t all honest voters have voted for a descendant of a certain block, then in slot $t + 1$ all honest voters will vote for a descendant of that block.

Lemma 2. *Suppose an execution of Goldfish in the synchronous sleepy network model. Suppose validator id^* with proposal P^* is recognized as the leader of a slot t by all awake honest validators at round $3\Delta t + \Delta$ (Alg. 1, l. 16). Then, all honest validators awake at round $3\Delta t + \Delta$ and eligible to vote at slot t , vote for $P^*.B$ at slot t .*

PROOF. Let $\mathcal{T}' = P^*. \mathcal{T}$, and \mathcal{B}^* and \mathcal{T}^* denote the buffer and bvtree of id^* at round $3\Delta t$. Since id^* is honest, it must have broadcast P^* at round $3\Delta t$ with bvtree $\mathcal{T}' = \text{MERGE}(\mathcal{T}^*, \mathcal{B}^*)$ and a new block $P^*.B$ with parent $\text{GHOST-EPH}(\mathcal{T}', t - 1)$ (Alg. 1, ll. 7, 8, 12).

By synchrony, any message that a non-asleep honest validator id could have added to its bvtree \mathcal{T}_{id} by $3\Delta(t - 1) + 2\Delta$, is received

by id^* by $3\Delta t$, and thus in \mathcal{T}' . As awake honest validators do not update their bvtrees and no honest validators turn awake in the interval $(3\Delta(t-1) + 2\Delta, 3\Delta t + \Delta)$, for any honest validator id awake at round $3\Delta t + \Delta$, $\mathcal{T}'_{\text{id}} \subseteq \mathcal{T}'$ prior to Alg. 1, l. 19.

Since id^* is recognized as the leader of slot t by all awake honest validators at round $3\Delta t + \Delta$, at that round, each awake honest validator id merges its bvtree with $\mathcal{T}' \cup \{P^*.B\}$ (Alg. 1, l. 19) and reaches $\mathcal{T}'_{\text{id}} = \mathcal{T}' \cup \{P^*.B\}$. Consequently, each honest validator id awake at round $3\Delta t + \Delta$ and eligible to vote at slot t votes for $P^*.B$ due to the recursive structure of the GHOST-Eph rule (Alg. 2). \square

Lemma 3. *Suppose a $(\frac{1}{2}, 3\Delta)$ -compliant execution of Goldfish in the synchronous sleepy network model. Consider a slot t where all honest validators awake at round $3\Delta t + \Delta$ and eligible to vote at slot t , vote for a descendant of B . Then, w.o.p., all honest validators awake at round $3\Delta(t+1) + \Delta$ and eligible to vote at slot $t+1$, vote for a descendant of B .*

PROOF. By Lem. 1, w.o.p., for every slot t , the number of adversarial validators at round $3\Delta(t+1) + \Delta$ and eligible to vote at slot t is less than the number of honest validators awake at round $3\Delta t + \Delta$ and eligible to vote at slot t .

Let t be a slot such that all honest validators awake at round $3\Delta t + \Delta$ and eligible to vote at t voted for a descendant of B . Pick any honest validator id awake at round $3\Delta(t+1) + \Delta$ and eligible to vote at slot $t+1$. Since id must have been awake at least since round $3\Delta t + 2\Delta$, its bvtree at round $3\Delta t + 2\Delta$ contains all votes broadcast by honest validators awake at round $3\Delta t + \Delta$ and eligible to vote at slot t (Alg. 1, l. 19). The same is true for its bvtree at round $3\Delta(t+1) + \Delta$, even after id merges its bvtree with that of any proposal (Alg. 1, l. 7). Moreover, the number of honest validators awake at round $3\Delta t + \Delta$ and eligible to vote at slot t is greater than the number of adversarial validators at round $3\Delta(t+1) + \Delta$ that are eligible to vote at slot t .

Consequently, upon invoking the GHOST-Eph fork-choice rule at round $3\Delta(t+1) + \Delta$ (Alg. 1, l. 22), id observes that at every iteration of the fork choice (Alg. 2, l. 7), blocks consistent with B have more votes than blocks conflicting with B . Thus, at round $3\Delta(t+1) + \Delta$, fork choice returns a descendant of B , and id votes for it. \square

PROOF OF THM. 1. From Lems. 1, 2 and 3, it follows by induction that w.o.p., for all $t' \geq t$, all honest validators awake at round $3\Delta t' + \Delta$ and eligible to vote at slot t' , vote for a descendant of $P^*.B$.

By synchrony, the honest votes of slot t' reach all honest validators awake at $3\Delta t' + 2\Delta$ by then, when they also merge the votes into their bvtrees. The number of honest validators awake at round $3\Delta t' + \Delta$ and eligible to vote at slot t' is greater than the number of adversarial validators by round $3\Delta(t'+1) + \Delta$ that are eligible to vote at slot t' (by Lem. 1). Upon invoking the GHOST-Eph rule of Alg. 1, ll. 8, 22, 28 at $3\Delta t' + 2\Delta$, $3\Delta(t'+1)$ and $3\Delta(t'+1) + \Delta$, respectively, an awake honest validator id (who must have been awake since at least $3\Delta t' + 2\Delta$, due to the joining procedure) observes that at every iteration of the fork choice (Alg. 2, l. 7), blocks consistent with $P^*.B$ have more votes than blocks conflicting with $P^*.B$. Thus, id 's fork choice reaches a descendant of $P^*.B$. \square

5.2 Goldfish with Accountability Gadget

We next provide a formal statement and proof sketch for Def. 4:

Theorem 4 (Ebb-and-flow property). *Goldfish combined with accountability gadgets (cf. Sec. 3.2) satisfies the ebb-and-flow property:*

- (1) (**P1: Accountability and finality**) *Under a partially synchronous network in the sleepy model, the accountable final prefix ledger ch_{acc} has accountable safety resilience $n/3$ at all times, (except w.p. $\text{negl}(\lambda)$), and there exists a constant C such that if the execution is $(\frac{1}{3}, 3\Delta)$ -compliant, ch_{acc} provides liveness with transaction confirmation time $T_{\text{conf}} = \Theta(\kappa^2)$ after round $\max(\text{GST}, \text{GAT}) + C\kappa$ (w.o.p.).*
- (2) (**P2: Dynamic availability**) *Under a synchronous network in the sleepy model (i.e., for $\text{GST} = 0$), if the execution is $(\frac{1}{2}, 3\Delta)$ -compliant, the available ledger ch_{ava} is secure at all times (w.o.p.).*
- (3) (**Prefix**) *For each honest id and round r , $\text{ch}_{\text{acc},r}^{\text{id}} \leq \text{ch}_{\text{ava},r}^{\text{id}}$.*

In Goldfish with accountability gadgets, a partially synchronous accountably-safe consensus protocol is used to determine checkpoints. Security of this protocol ensures the *safety and accountability of the prefix ledger ch_{acc}* in the partially synchronous sleepy network model. To ensure the *prefix property*, the fork-choice rule of Goldfish is modified to respect earlier checkpoint decisions (Alg. 3). This modification requires adjustments of the analysis of Goldfish, because it opens up the possibility that for a proposal P^* by an honest leader, $P^*.B \leq B$ no longer holds for all blocks B identified in Alg. 1, ll. 8, 22, 28 by awake honest validators at future rounds, due to a new checkpoint conflicting with $P^*.B$.

To prevent checkpoints from undermining the security in this manner, and rigorously argue security of the combination despite the modified fork-choice rule, the framework of accountability gadgets [49, 53] relies on two principles:

- **Gap Property:** After a (successful) checkpointing iteration with a new checkpoint, honest validators wait for $T_{\text{chkpt}} = \Theta(\kappa)$ rounds before participating in the next iteration.
- **Recency Property:** For checkpointing, honest validators suggest and approve only the blocks that were *recently* confirmed as part of ch_{ava} .

We prove that once the network heals and honest validators become awake at round $\max(\text{GST}, \text{GAT})$, ch_{ava} regains its security with the help of these properties. Key is the following *healing property*.

Lemma 4 (Healing property (sketch)). *Suppose the number of adversarial validators is less than $n/3$ at all rounds.*

Then, under partial synchrony in the sleepy model, the available ledger ch_{ava} is secure with transaction confirmation time $\Theta(T_{\text{chkpt}})$ after round $\max(\text{GAT}, \text{GST}) + \Theta(\kappa)$.

Formal statements for Lem. 4, its proof, and the full proof for **P1** are given in App. A.2.

PROOF SKETCH FOR LEM. 4. Since the number of adversarial validators is less than $n/3$ at all rounds, by the security of the consensus protocol used by the accountability gadgets, all checkpoints are consistent with each other. After round $\max(\text{GAT}, \text{GST})$, all awake honest validators agree on the rounds they enter and complete subsequent checkpoint iterations (up to a difference of Δ rounds). By the gap property, honest validators wait for $T_{\text{chkpt}} = \Theta(\kappa)$ rounds before participating in the next checkpointing iteration after a successful one. This ensures that no new checkpoints appear for T_{chkpt}

rounds, during which the honest validators can treat the last checkpoint as the ‘new’ genesis block. Then, via the security analysis in Sec. 5.1, there exists a slot with an honest leader such that for $P^*.B \leq B$ for all blocks B identified in Alg. 1, ll. 8, 22, 28 by awake honest validators during future rounds, until a new checkpoint is determined. By the recency property, for checkpointing, honest validators suggest and approve only the blocks that were *recently* confirmed as part of ch_{ava} . Since $P^*.B \leq B$ for all recently confirmed blocks B at the start of the next checkpointing iteration, if a new checkpoint appears in the next iteration, the above prefix relation continues to hold for all future slots after the iteration. Thus, it is possible to state an analogue of Thm. 1 for honest leaders during T_{chkpt} rounds following successful checkpoint iterations and prove security with transaction confirmation time $T_{\text{conf}} = \Theta(T_{\text{chkpt}})$ via a similar reasoning to Sec. 5.1. \square

Liveness of ch_{ava} together with the liveness of the accountability gadget’s consensus protocol imply the *liveness* of ch_{acc} in the partially synchronous sleepy network model.

In the synchronous sleepy network model, Sec. 5.1 implies that ch_{ava} remains secure until the first checkpoint is determined. However, checkpoints cannot undermine its security since only confirmed blocks in ch_{ava} are approved for checkpointing by honest validators. Proof of **P2** is given in App. A.2.

5.3 Goldfish with Fast Confirmation

In the following analysis, we consider a synchronous network in the sleepy model as described in Sec. 2. Recall that the total number of validators is n (cf. Sec. 2). Since Goldfish slots consist of 4Δ rounds in the case of fast confirmation, we hereafter assume that the Goldfish execution is $(\frac{1}{2}, 4\Delta)$ -compliant. Similarly, we state an analogue of Lem. 1, namely Lem. 6, to match the new slot structure in App. A.1. We show that Thm. 2 holds for Goldfish with fast confirmations (w.o.p.) in compliant executions. To do so, we first prove Thm. 5, an analogue of Thm. 1 for fast confirmations, showing that fast confirmed blocks are always in the canonical chain of awake validators at later rounds.

Proposition 1. Suppose $T_{\text{hor}} = \text{poly}(\kappa)$. Then, w.o.p., there can be at most $(1 + \epsilon)n \text{thr}_v$ validators that are eligible to vote at any given slot. If the Goldfish execution is $(\frac{1}{2}, 4\Delta)$ -compliant, then, w.o.p., for all slots t , the number of adversarial validators at round $4\Delta(t+1) + \Delta$, eligible to vote at slot t , is less than $\frac{1}{2}n \text{thr}_v$.

Proof follows from a Chernoff bound.

Lemma 5. Suppose the Goldfish execution is $(\frac{1}{2}, 4\Delta)$ -compliant in the synchronous sleepy network model, and an honest validator id^* fast confirms a block B at slot t . Then, w.o.p., all honest validators awake at round $4\Delta(t+1) + \Delta$ and eligible to vote at slot $t+1$, vote for a descendant of B at slot $t+1$.

PROOF. By Prop. 1, w.o.p., the number of adversarial validators at round $4\Delta(t+1) + \Delta$, eligible to vote at slot t , is less than $\frac{1}{2}n \text{thr}_v$. An eligible awake honest validator sends a single slot t vote at round $4\Delta t + \Delta$, implying that over $(\frac{3}{4} + \frac{\epsilon}{2})n \text{thr}_v - \frac{1}{2}n \text{thr}_v = (\frac{1}{4} + \frac{\epsilon}{2})n \text{thr}_v$ validators broadcast a single slot t vote by round $4\Delta(t+1) + \Delta$, and that is for a descendant of B . By Prop. 1, w.o.p., for all slots t , there can be at most $(1 + \epsilon)n \text{thr}_v$ validators that are eligible to vote at

t . Hence, the number of valid slot t votes for the descendants of any block B' conflicting with B must be less than $(1 + \epsilon)n \text{thr}_v - (\frac{1}{4} + \frac{\epsilon}{2})n \text{thr}_v = (\frac{3}{4} + \frac{\epsilon}{2})n \text{thr}_v$ at any given round. The validator id^* broadcasts B and over $(\frac{3}{4} + \frac{\epsilon}{2})n \text{thr}_v$ valid votes for it (in pieces) at round $4\Delta t + 2\Delta$. Each honest validator, awake at round $4\Delta(t+1) + \Delta$ and eligible to vote at slot $t+1$, observes these votes in its bvtree at the round of voting (Alg. 5, l. 12). Upon invoking the GHOST-Eph fork-choice rule at any of the rounds $4\Delta t + 3\Delta$, $4\Delta(t+1)$ or $4\Delta(t+1) + \Delta$ (Alg. 1, ll. 8, 22, 28), for any awake honest validator id with bvtree \mathcal{T}' , $V_{\text{TS}}(\mathcal{T}', B, t) > V_{\text{TS}}(\mathcal{T}', B', t)$ for any block B' conflicting with B . This implies that all honest validators, awake at round $4\Delta(t+1) + \Delta$ and eligible to vote at slot $t+1$ all vote for B or one of its descendants at slot $t+1$. \square

Theorem 5. Suppose the Goldfish execution is $(\frac{1}{2}, 4\Delta)$ -compliant in the synchronous sleepy network model, and an honest validator id^* fast confirms a block B at slot t . Then, w.o.p., $B \leq B$ for any B identified in Alg. 1, ll. 8, 22, 28 by any awake honest validator in any round $r \geq 4\Delta(t+1) + \Delta$.

PROOF. Follows by Lems. 6, 5 and 3, by the same inductive argument used in the proof of Thm. 1, in that case following from Lems. 1, 2 and 3. Here, Lem. 6 is the analogue of Lem. 1 with the new slot structure, and Lem. 5 provides the base case, substituting Lem. 2. \square

Theorem 6. Suppose the Goldfish execution is $(\frac{1}{2}, 4\Delta)$ -compliant. Then, Goldfish with fast confirmations satisfies safety (w.o.p.).

PROOF. If an honest validator fast confirms a block B at slot t , then B is in the canonical GHOST-Eph chain of every awake honest validator at all slots larger than t by Thm. 5. Therefore, B is in the κ -slots-deep prefix of the canonical GHOST-Eph chains of all awake honest validators at slot $t + \kappa$, and thus confirmed by them with the standard confirmation rule. Therefore, Thm. 2 implies the safety of the protocol. \square

In $(\frac{1}{2}, 4\Delta)$ -compliant executions, we automatically get liveness of Goldfish with fast confirmations from the liveness of the standard confirmation rule, since fast confirmation is not needed for a block to be confirmed. Under optimistic conditions, liveness of fast confirmations holds as well. We prove that a block within an honest, valid proposal is immediately fast confirmed within the same slot by the awake honest validators, if there are over $(\frac{3}{4} + \frac{3}{2}\epsilon)n$ awake, honest validators at the voting time of the given slot, implying the liveness of fast confirmations under optimistic conditions.

Theorem 7. Suppose the Goldfish execution is $(\frac{1}{2}, 4\Delta)$ -compliant. Then, Goldfish with fast confirmations satisfies liveness with $T_{\text{conf}} = \Theta(\kappa)$ (w.o.p.).

Consider a slot t , such that there are $(\frac{3}{4} + \frac{3}{2}\epsilon)n \text{thr}_v$ honest validators eligible to vote at slot t and awake at round $4\Delta t + \Delta$. Suppose an honest validator id with proposal P^* is recognized as the leader of a slot t by all awake honest validators at round $4\Delta t + \Delta$ (Alg. 1, l. 16).

Then all honest validators awake at round $4\Delta t + 2\Delta$ fast confirm $P^*.B$ in Alg. 5, l. 10.

PROOF. Proof of liveness follows from Thm. 2.

For the second part of the proof, by Lem. 2, all of eligible and awake honest validators vote for $P^*.B$ at slot t . Then, the buffer

of any honest validator awake at round $4\Delta t + 2\Delta$ contains at least $(\frac{3}{4} + \frac{\epsilon}{2})n$ thr_v votes (by Chernoff bound) for the block $P^*.B$, implying that all honest validators awake at rounds $4\Delta t + 2\Delta$ fast confirm $P^*.B$ at the respective slots. \square

6 PRACTICAL CONSIDERATIONS

Goldfish is designed to be a provably secure alternative to the Gasper [9] protocol used in Ethereum, but due to its qualities, it is more broadly practically applicable to decentralized, permissionless blockchains, as we illustrate using Ethereum as a case study.

Comparison of security properties with current PoS Ethereum. In its standalone form without an accountability gadget, Goldfish is reorg resilient and dynamically available, while still supporting subsampling of validators per slot. Reorg resilience is a broadly desirable property in real world blockchain systems, because rationality considerations in practice complement honesty assumptions, and thus incentive compatibility is paramount. Subsampling of validators allows for a large validator set, as is the case in Ethereum, without requiring a large load on validators, altogether facilitating the decentralization of the system. Finally, dynamic availability is a robust notion of security, which allows for unforeseeable disruptions to the set of consensus participants, such as node software updates. It is therefore well suited for a permissionless system, in which there is little control over the participants, and limited ability to react to system-wide disruptions in real time.

On the contrary, Gasper’s LMD GHOST satisfies none of these properties, as we extensively discuss in App. C. In the original protocol, ex-ante reorgs and balancing attacks [43, 44, 54] prevent security even in the full participation setting and without subsampling. The proposer boost technique [7] mitigates these issues, but is itself not compatible with dynamic participation, and it entails a lower adversarial tolerance ($\frac{1}{3}$) than what is obtained with message buffering ($\frac{1}{2}$). Moreover, ex-ante reorgs [54] are still possible with subsampling, compromising reorg resilience, and the latest message rule (LMD) itself is not compatible with dynamic participation, both issues which Goldfish solves through vote expiry. In its combination with accountability gadgets, Goldfish satisfies the Ebb-and-flow properties [47], which formalize the requirements of a solution to the availability-accountability dilemma [49]. On the other hand, the interaction of LMD GHOST and Casper FFG in Gasper is known to be susceptible to bouncing attacks [10, 40, 41], which exploit the checkpointing process to jeopardize security of the available ledger, and consequently also liveness of the accountable ledger.

Wire format, message size, and spamming vectors. Another practical consideration in designing a consensus protocol for a decentralized blockchain are resource constraints. A challenge which is specific to PoS protocols is dealing with equivocations, which are for instance not possible in PoW. Care has to be taken in order to prevent spamming of equivocations being a relatively cheap Denial-of-Service (DoS) vector, while preserving security properties [42]. In fact, multiple attacks related to the handling of equivocations were discovered for LMD GHOST [46, 45]. We address this issue by introducing *equivocation discounting*, i.e., not counting votes in fork-choice from validators who have equivocated. We discuss this

at length, and show to not compromise security, in App. B.

Together with vote expiry, equivocation discounting alleviates related concerns about the practical feasibility of the message buffering technique, in particular concerning the size of proposal messages. Due to vote expiry, the only votes which have to be included in the proposed bvtree are those from the previous slot, and equivocation discounting ensures that at most two votes per validator are sufficient, even in the presence of multiple equivocations. Notice also that, while we have for ease of exposition talked about proposal messages as including a bvset, it is in practice not necessary to *fully* include any messages in a proposal, as references (hashes) are sufficient. The proposer’s role in the message buffering technique is in fact only to point validators to messages which *they already have in their buffer*, notifying them that they can safely be merged into their bvtree. For the same reason, only leaf blocks with nonzero weight need to be at all referenced, because each leaf block is a reference to its entire chain, and ones without weight cannot possibly favorably influence the fork choice towards the proposer’s block. Since leaf blocks with nonzero weight correspond to a subset of votes, only votes need to be referenced. As a concrete example, were Goldfish to be applied to Ethereum today, the maximum number of votes which would need to be referenced is roughly 2000, in the worst case in which all attestation aggregators (validators which aggregate votes) equivocate at a slot, for a total of 64 KBs if only hashes are included in the proposed bvtree (missing objects can be exchanged by reference to the hash). Moreover, the κ -deep confirmation rule allows for garbage collection of blocks ‘stuck in limbo’ (i.e., with dangling references preventing assessment of validity) that are κ slots old (then they cannot enter the chain permanently anymore), and vote expiry allows to garbage collect votes within 2 slots (then they have no more effect on fork-choice).

From LMD GHOST to Goldfish. For Goldfish to be used as a drop-in replacement for LMD GHOST in Ethereum, only a few adjustments are required. The proposer selection mechanism, RANDAO, would have to be replaced with the VRF lottery, in order to preserve adaptive security and a confirmation time which is independent of participation. The former can also be preserved by using a single secret leader election [4], which has already been researched for use in Ethereum [26]. Moreover, vote expiry and message buffering would have to be introduced, with the latter replacing proposer boost. Finally, in order to benefit from the security guarantees of Goldfish in its combination with an accountability gadget, the interaction with Casper FFG would have to be modified to fit the construction from [49], which we have also employed in this work. Casper FFG is itself not a fully specified consensus protocol, as it lacks specification of a proposal mechanism and slot structure, but it is very similar to *streamlined*, accountably safe consensus protocols like Streamlet [13] or Chained Hotstuff [58], which are all feasible for use in combination with Goldfish.

ACKNOWLEDGMENTS

We thank Aditya Asgaonkar, Carl Beekhuizen, Vitalik Buterin, Justin Drake, Dankrad Feist, Sreeram Kannan, Georgios Konstantopoulos, Barnabé Monnot, Dan Robinson, Danny Ryan, Caspar Schwarz-Schilling, and Fan Zhang for fruitful discussions. The work of JN was conducted in part while at Paradigm. JN, ENT and DT

are supported by a gift from the Ethereum Foundation. JN is supported by the Protocol Labs PhD Fellowship and the Reed-Hodgson Stanford Graduate Fellowship. ENT is supported by the Stanford Center for Blockchain Research.

REFERENCES

- [1] Christian Badertscher, Peter Gazi, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. 2018. Ouroboros genesis: composable proof-of-stake blockchains with dynamic availability. In *CCS*. ACM, 913–930.
- [2] Vivek Kumar Bagaria, Sreeram Kannan, David Tse, Giulia Fanti, and Pramod Viswanath. 2019. Prism: deconstructing the blockchain to approach physical limits. In *CCS*. ACM, 585–602.
- [3] Carl Beekhuizen, Caspar Schwarz-Schilling, and Francesco D’Amato. 2021. Change fork choice rule to mitigate balancing and reorging attacks. Ethereum Research. (Oct. 28, 2021). Retrieved June 28, 2022 from <https://ethresear.ch/t/change-fork-choice-rule-to-mitigate-balancing-and-reorging-attacks/11127>.
- [4] Dan Boneh, Saba Eskandarian, Lucjan Hanzlik, and Nicola Greco. 2020. Single secret leader election. In *AFT*. ACM, 12–24.
- [5] Dan Boneh and Victor Shoup. 2015. *A Graduate Course in Applied Cryptography*. <http://toc.cryptobook.us/book.pdf>.
- [6] Ethan Buchman, Jae Kwon, and Zarko Milosevic. 2018. The latest gossip on bft consensus. (2018). arXiv: 1807.04938v3 [cs. DC].
- [7] Vitalik Buterin. 2020. Proposal for mitigation against balancing attacks to LMD GHOST. Retrieved Apr. 20, 2021 from https://notes.ethereum.org/@vbuterin/lmd_ghost_mitigation.
- [8] Vitalik Buterin and Virgil Griffith. 2017. Casper the friendly finality gadget. (2017). arXiv: 1710.09437v4 [cs. CR].
- [9] Vitalik Buterin, Diego Hernandez, Thor Kampefner, Khiem Pham, Zhi Qiao, Danny Ryan, Juhyeok Sin, Ying Wang, and Yan X Zhang. 2020. Combining ghost and casper. (2020). arXiv: 2003.03052v3 [cs. CR].
- [10] Vitalik Buterin and Alistair Stewart. 2018. Beacon chain casper mini-spec (comments #17, #19). Ethereum Research. (Sept. 25, 2018). Retrieved Aug. 18, 2020 from <https://ethresear.ch/t/beacon-chain-casper-mini-spec/2760/17>.
- [11] Miles Carlsten, Harry A. Kalodner, S. Matthew Weinberg, and Arvind Narayanan. 2016. On the instability of bitcoin without the block reward. In *CCS*. ACM, 154–167.
- [12] Miguel Castro and Barbara Liskov. 1999. Practical byzantine fault tolerance. In *OSDI*. USENIX Association, 173–186.
- [13] Benjamin Y. Chan and Elaine Shi. 2020. Streamlet: textbook streamlined blockchains. In *AFT*. ACM, 1–11.
- [14] Jing Chen and Silvio Micali. 2019. Algorand: A secure and efficient distributed ledger. *Theor. Comput. Sci.*, 777, 155–183.
- [15] Phil Daian, Rafael Pass, and Elaine Shi. 2019. Snow white: robustly reconfigurable consensus and applications to provably secure proof of stake. In *Financial Cryptography* (LNCS). Vol. 11598. Springer, 23–41.
- [16] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. 2020. Flash boys 2.0: frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *IEEE Symposium on Security and Privacy*. IEEE, 910–927.
- [17] Bernardo David, Peter Gazi, Aggelos Kiayias, and Alexander Russell. 2018. Ouroboros praos: an adaptively-secure, semi-synchronous proof-of-stake blockchain. In *EUROCRYPT* (2) (LNCS). Vol. 10821. Springer, 66–98.
- [18] Yevgeniy Dodis and Aleksandr Yampolskiy. 2005. A verifiable random function with short proofs and keys. In *Public Key Cryptography* (LNCS). Vol. 3386. Springer, 416–431.
- [19] Ittay Eyal and Emin Gün Sirer. 2018. Majority is not enough: bitcoin mining is vulnerable. *Commun. ACM*, 61, 7, 95–102.
- [20] Matthias Fitz, Peter Gazi, Aggelos Kiayias, and Alexander Russell. 2018. Parallel chains: improving throughput and latency of blockchain protocols via parallel composition. Cryptology ePrint Archive, Paper 2018/1119. (2018). <https://eprint.iacr.org/2018/1119>.
- [21] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. 2015. The bitcoin backbone protocol: analysis and applications. In *EUROCRYPT* (2) (LNCS). Vol. 9057. Springer, 281–310.
- [22] Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. 2017. Algorand: scaling byzantine agreements for cryptocurrencies. In *SOSP*. ACM, 51–68.
- [23] Seth Gilbert and Nancy A. Lynch. 2002. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, 33, 2, 51–59.
- [24] Vipul Goyal, Hanjun Li, and Justin Raizes. 2021. Instant block confirmation in the sleepy model. In *Financial Cryptography* (2) (LNCS). Vol. 12675. Springer, 65–83.
- [25] Gene Itkis and Leonid Reyzin. 2001. Forward-secure signatures with optimal signing and verifying. In *CRYPTO* (LNCS). Vol. 2139. Springer, 332–354.
- [26] George Kadianakis. 2022. Whisk: a practical shuffle-based SSLE protocol for Ethereum. Ethereum Research. (Jan. 13, 2022). Retrieved Sept. 5, 2022 from <https://ethresear.ch/t/whisk-a-practical-shuffle-based-ssle-protocol-for-ethereum/11763>.
- [27] Daniel Kane, Andreas Fackler, Adam Gagol, and Damian Straszak. 2021. High-way: efficient consensus with flexible finality. (2021). arXiv: 2101.02159v2 [cs. DC].
- [28] Jonathan Katz and Yehuda Lindell. 2014. *Introduction to Modern Cryptography, Second Edition*. CRC Press.
- [29] Pankaj Khanchandani and Roger Wattenhofer. 2020. Brief announcement: byzantine agreement with unknown participants and failures. In *PODC*. ACM, 178–180.
- [30] Pankaj Khanchandani and Roger Wattenhofer. 2021. Byzantine agreement with unknown participants and failures. In *IPDPS*. IEEE, 952–961.
- [31] Aggelos Kiayias and Giorgos Panagiotakos. 2017. On trees, chains and fast transactions in the blockchain. In *LATINCRYPT* (LNCS). Vol. 11368. Springer, 327–351.
- [32] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. 2017. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *CRYPTO* (1) (LNCS). Vol. 10401. Springer, 357–388.
- [33] Andrew Lewis-Pye and Tim Roughgarden. 2020. Resource pools and the CAP theorem. (2020). arXiv: 2006.10698v1 [cs. DC].
- [34] Kevin Liao and Jonathan Katz. 2017. Incentivizing blockchain forks via whale transactions. In *Financial Cryptography Workshops* (LNCS). Vol. 10323. Springer, 264–279.
- [35] Dahlia Malkhi, Atsuki Momose, and Ling Ren. 2022. Byzantine consensus under fully fluctuating participation. Cryptology ePrint Archive, Paper 2022/1448. (2022). <https://eprint.iacr.org/2022/1448>.
- [36] Dahlia Malkhi, Atsuki Momose, and Ling Ren. 2022. Instant finality in byzantine generals with unknown and dynamic participation. Chainlink Labs Research. (Aug. 28, 2022). Retrieved Sept. 5, 2022 from <https://blog.chain.link/instant-finality-in-byzantine-generals-with-unknown-and-dynamic-participation/>.
- [37] Silvio Micali, Michael O. Rabin, and Salil P. Vadhan. 1999. Verifiable random functions. In *FOCS*. IEEE Computer Society, 120–130.
- [38] Atsuki Momose and Ling Ren. 2022. Constant latency in sleepy consensus. In *CCS*. ACM, 2295–2308.
- [39] Satoshi Nakamoto. 2008. Bitcoin: a peer-to-peer electronic cash system. <https://bitcoin.org/bitcoin.pdf>. (2008).
- [40] Ryuya Nakamura. 2019. Analysis of bouncing attack on FFG. Ethereum Research. (Sept. 8, 2019). Retrieved Aug. 18, 2020 from <https://ethresear.ch/t/analysis-of-bouncing-attack-on-ffg/6113>.
- [41] Ryuya Nakamura. 2019. Prevention of bouncing attack on FFG. Ethereum Research. (Sept. 8, 2019). Retrieved Aug. 18, 2020 from <https://ethresear.ch/t/prevention-of-bouncing-attack-on-ffg/6114>.
- [42] Joachim Neu, Srivatsan Sridhar, Lei Yang, David Tse, and Mohammad Alizadeh. 2022. Longest chain consensus under bandwidth constraint. In *4th ACM Conference on Advances in Financial Technologies* (AFT ’22). ACM. <https://arxiv.org/abs/2111.12332>.
- [43] Joachim Neu, Ertem Nusret Tas, and David Tse. 2020. A balancing attack on Casper, the current candidate for Eth2’s beacon chain. Ethereum Research. (Oct. 20, 2020). Retrieved Apr. 22, 2021 from <https://ethresear.ch/t/a-balancing-attack-on-gasper-the-current-candidate-for-eth2s-beacon-chain/8079>.
- [44] Joachim Neu, Ertem Nusret Tas, and David Tse. 2021. Attacking Casper without adversarial network delay. Ethereum Research. (July 25, 2021). Retrieved Sept. 2, 2021 from <https://ethresear.ch/t/attacking-gasper-without-adversarial-network-delay/10187>.
- [45] Joachim Neu, Ertem Nusret Tas, and David Tse. 2022. Avalanche attack on proof-of-stake GHOST. Ethereum Research. (Jan. 24, 2022). Retrieved Jan. 24, 2022 from <https://ethresear.ch/t/avalanche-attack-on-proof-of-stake-ghost/11854>.
- [46] Joachim Neu, Ertem Nusret Tas, and David Tse. 2022. Balancing attack: LMD edition. Ethereum Research. (Jan. 24, 2022). Retrieved Jan. 24, 2022 from <https://ethresear.ch/t/balancing-attack-lmd-edition/11853>.
- [47] Joachim Neu, Ertem Nusret Tas, and David Tse. 2021. Ebb-and-flow protocols: A resolution of the availability-finality dilemma. In *IEEE Symposium on Security and Privacy*. IEEE, 446–465.
- [48] Joachim Neu, Ertem Nusret Tas, and David Tse. 2020. Snap-and-chat protocols: system aspects. (2020). arXiv: 2010.10447v1 [cs. CR].
- [49] Joachim Neu, Ertem Nusret Tas, and David Tse. 2022. The availability-accountability dilemma and its resolution via accountability gadgets. In *Financial Cryptography* (LNCS). Vol. 13411. Springer, 541–559.
- [50] Joachim Neu, Ertem Nusret Tas, and David Tse. 2022. Two more attacks on proof-of-stake GHOST/Ethereum. In *Proceedings of the 2022 ACM Workshop on Developments in Consensus* (ConsensusDay ’22). ACM, (Nov. 2022). doi: 10.1145/3560829.3563560.
- [51] Rafael Pass and Elaine Shi. 2017. The sleepy model of consensus. In *ASIACRYPT* (2) (LNCS). Vol. 10625. Springer, 380–409.

- [52] Youer Pu, Lorenzo Alvisi, and Ittay Eyal. 2022. Safe permissionless consensus. In *DISC (LIPICs)*. Vol. 246. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 33:1–33:15.
- [53] Suryanarayana Sankagiri, Xuechao Wang, Sreeram Kannan, and Pramod Viswanath. 2021. Blockchain CAP theorem allows user-dependent adaptivity and finality. In *Financial Cryptography (2) (LNCS)*. Vol. 12675. Springer, 84–103.
- [54] Caspar Schwarz-Schilling, Joachim Neu, Barnabé Monnot, Aditya Asgaonkar, Ertem Nusret Tas, and David Tse. 2022. Three attacks on proof-of-stake Ethereum. In *Financial Cryptography (LNCS)*. Vol. 13411. Springer, 560–576.
- [55] Peiyao Sheng, Gerui Wang, Kartik Nayak, Sreeram Kannan, and Pramod Viswanath. 2021. BFT protocol forensics. In *CCS. ACM*, 1722–1743.
- [56] Peiyao Sheng, Gerui Wang, Kartik Nayak, Sreeram Kannan, and Pramod Viswanath. 2022. Player-replaceability and forensic support are two sides of the same (crypto) coin. *Cryptology ePrint Archive*, Paper 2022/1513. (2022). <https://eprint.iacr.org/2022/1513>.
- [57] Yonatan Sompolinsky and Aviv Zohar. 2015. Secure high-rate transaction processing in bitcoin. In *Financial Cryptography (LNCS)*. Vol. 8975. Springer, 507–527.
- [58] Maofan Yin, Dahlia Malkhi, Michael K. Reiter, Guy Golan-Gueta, and Ittai Abraham. 2019. Hotstuff: BFT consensus with linearity and responsiveness. In *PODC. ACM*, 347–356.

A ANALYSIS

A.1 Proof of Lem. 1

PROOF OF LEM. 1. By the *pseudorandomness* property of the VRF-based lottery (Secs. 2.1.3 and 3.1.1), for any given slot t and validators id_1 and id_2 , $id_1 \neq id_2$,

$$\Pr \left[\text{IsWinning}^{(\text{vote}, \text{thr}_v)}((id, t), \text{Open}_{id_1}^{(\text{vote}, \text{thr}_v)}(t)) \right] = \text{thr}_v \quad (9)$$

$$\Pr \left[\text{IsWinning}^{(\text{block}, \text{thr}_b)}((id, t), \text{Open}_{id_1}^{(\text{block}, \text{thr}_b)}(t)) \right] = \text{thr}_b \quad (10)$$

$$\Pr \left[\text{Prio}(\text{Open}_{id_1}^{(\text{block}, \text{thr}_b)}(t)) < \text{Prio}(\text{Open}_{id_2}^{(\text{block}, \text{thr}_b)}(t)) \right] = \frac{1}{2}, \quad (11)$$

and $\text{Open}_{id_1}^{(\text{vote}, \text{thr}_v)}(t)$, $\text{Open}_{id_1}^{(\text{block}, \text{thr}_b)}(t)$, $\text{Open}_{id_2}^{(\text{block}, \text{thr}_b)}(t)$, and $\text{Open}_{id_2}^{(\text{vote}, \text{thr}_v)}(t)$ are independent random variables.

We first consider the protocol *without key-evolving primitives*. By the *uniqueness* property of the lottery (Sec. 2), w.o.p., for all validators id and slots t , the ticket (id, t) can be opened at most one unique opening (Alg. 1, l. 20). Let \tilde{H}_t denote the number of honest validators awake at round $3\Delta t + \Delta$ and eligible to vote at slot t . Let \tilde{A}_t denote the number of adversarial validators at round $3\Delta(t+1) + \Delta$ that are eligible to vote at slot t . Recall that A_r and H_r denote the number of adversarial and honest validators awake at round r respectively (note that the honest validators have been awake since the closest round $3\Delta t + 2\Delta$ same as or preceding r). Let $n_t = H_{3\Delta t + \Delta} + A_{3\Delta(t+1) + \Delta} \geq n_0 = \Theta(\kappa)$.

By the pseudorandomness property, the adversary cannot predict in advance which honest validators will become eligible to vote or propose at a given slot. Moreover, if the adversary decides to corrupt the honest validators eligible to vote at a slot t after learning their identities at round $3\Delta t + \Delta$, it takes over 3Δ rounds for the corruption to take effect, implying that these validators cannot be counted as part of \tilde{A}_t . Hence, as $\frac{A_r}{A_r + H_r - 3\Delta} \leq \beta < \frac{1}{2} - \epsilon$ for all rounds r , w.o.p.,

$$\mathbb{E}[\tilde{H}_t] = H_{3\Delta t + \Delta} \text{thr}_v \geq \left(\frac{1}{2} + \epsilon\right) n_t \text{thr}_v$$

$$\mathbb{E}[\tilde{A}_t] = A_{3\Delta(t+1) + \Delta} \text{thr}_v \leq \left(\frac{1}{2} - \epsilon\right) n_t \text{thr}_v$$

By a Chernoff bound,

$$\Pr \left[\tilde{H}_t < \frac{1}{2} n_t \text{thr}_v \right] \leq e^{-\frac{\epsilon^2}{1+2\epsilon} n_t \text{thr}_v}$$

$$\Pr \left[\tilde{A}_t > \frac{1}{2} n_t \text{thr}_v \right] \leq e^{-\frac{\epsilon^2}{1+3\epsilon} n_t \text{thr}_v}.$$

Thus, at any given slot t , $\tilde{H}_t > \tilde{A}_t$, except with probability

$$2 \exp\left(-\frac{\epsilon^2}{1+3\epsilon} n_0 \text{thr}_v\right).$$

By a union bound, every slot t has more honest validators awake at round $3\Delta t + \Delta$ and eligible to vote at slot t than adversarial validators at round $3\Delta(t+1) + \Delta$, eligible to vote at slot t (and more than $\frac{1}{2} n_0 \text{thr}_v$ such honest validators), except with probability

$$2T_{\text{hor}} \exp\left(-\frac{\epsilon^2}{1+3\epsilon} n_0 \text{thr}_v\right) + \text{negl}(\lambda) = \text{negl}(\kappa) + \text{negl}(\lambda),$$

since $n_0 = \Theta(\kappa)$ and $T_{\text{hor}} = \Theta(\kappa)$. By the same reasoning, w.o.p., every slot t has more honest validators awake and eligible to propose for slot t at round $3\Delta t$ than adversarial validators at round $3\Delta t + \Delta$, eligible to propose for slot t .

Finally, for any given slot t , each valid slot t proposal broadcast within rounds $[3\Delta t, 3\Delta t + \Delta]$ has the same probability of achieving the minimum precedence up to terms negligible in λ^8 . Now, at a slot t , if an honest validator's proposal achieves the minimum precedence among the valid slot t proposals broadcast by Δ rounds into the slot, then that validator is identified as the slot leader by all honest validators awake at round $3\Delta t + \Delta$. Taking a fixed $t \geq \kappa$, the probability that no awake honest validator's proposal has the minimum precedence among the valid slot s proposals broadcast by Δ rounds into the slot, during the slots $s \in [t - \kappa, t]$, is upper bounded by $2^{-\kappa} + \text{negl}(\kappa) + \text{negl}(\lambda)$. Union bounding over all T_{hor} many such intervals, we find that w.o.p., all slot intervals of length κ have at least one slot t , where an honest validator is identified as the slot leader by all awake honest validators at round $3\Delta t + \Delta$.

Now with *key-evolving primitives*, we define $\tilde{H}_t = H_{3\Delta t + \Delta}$ and $\tilde{A}_t = A_{3\Delta t + \Delta}$. Similarly, we define $n_t = H_{3\Delta t + \Delta} + A_{3\Delta t + \Delta} \geq n_0 = \Theta(\kappa)$. In this case, $\frac{A_r}{A_r + H_r} \leq \beta < \frac{1}{2} - \epsilon$ for all rounds r . Note that the adversary cannot predict in advance which honest validators will become eligible to vote or propose at a given slot due to the pseudorandomness property of the lottery. Moreover, if the adversary corrupts the honest validators eligible to vote at a slot t after learning their identities at round $3\Delta t + \Delta$, it cannot make these validators broadcast new valid votes for slot t since the keys for slot t would have been evolved prior to adversarial corruption (i.e., these corrupted validators cannot be counted as part of \tilde{A}_t). Hence, the number of valid slot t votes adversarial validators can broadcast by round $3\Delta(t+1) + \Delta$ is upper bounded by the number of adversarial validators at round $3\Delta t + \Delta$ that are eligible to vote at slot t . Finally, by the same calculations as above, every slot t has more honest validators eligible to vote and awake at round $3\Delta t + \Delta$ than the adversarial validators at round $3\Delta(t+1) + \Delta$ eligible to vote at slot t (and more than $\frac{1}{2} n_0 \text{thr}_v$ such honest validators), except with probability

$$2T_{\text{hor}} \exp\left(-\frac{\epsilon^2}{1+3\epsilon} n_0 \text{thr}_v\right) + \text{negl}(\lambda) = \text{negl}(\kappa) + \text{negl}(\lambda).$$

Similarly, w.o.p., every slot t has more honest validators awake and eligible to propose for slot t at round $3\Delta t$ than adversarial validators at round $3\Delta t + \Delta$ eligible to propose for slot t . Thus, via the same

⁸We assume that $\text{poly}(\kappa) \text{negl}(\lambda) = \text{negl}(\lambda)$.

argument, w.o.p., all slot intervals of length κ have at least one slot t , where an honest validator is identified as the slot leader by all awake honest validators at round $3\Delta t + \Delta$. \square

Since Goldfish slots consist of 4Δ rounds in the case of fast confirmation, we state an analogue of Lem. 1 to match the new slot structure:

Lemma 6. *Suppose the Goldfish execution is $(\frac{1}{2}, 4\Delta)$ -compliant.*

Then, w.o.p., for every slot t , the number of adversarial validators at round $4\Delta(t+1) + \Delta$, eligible to vote at slot t , is less than the number of honest validators, awake at round $4\Delta t + \Delta$ and eligible to vote at slot t .

Also w.o.p., all slot intervals of length κ have at least one slot t , where an honest validator is identified as the slot leader by all awake honest validators at round $4\Delta t + \Delta$.

Proof of Lem. 6 is very similar to the proof of Lem. 1, and follows from the same arguments using $(\frac{1}{2}, 4\Delta)$ -compliant executions.

A.2 Goldfish with Accountability Gadgets

We now prove the ebb-and-flow property for Goldfish combined with accountability gadgets (Fig. 2). The following analysis extensively refers to the details of the accountability gadgets described in [49, Section 4]. These gadgets can be *instantiated* with any BFT protocol that satisfies accountable safety (e.g., PBFT [12], HotStuff [58]).

To distinguish the votes cast by validators as part of the accountability gadget iterations from those broadcast within Goldfish, we will refer to the former as *gadget votes*. Similarly, to distinguish the leaders of accountability gadget iterations from the leaders of Goldfish slots, we will refer to the former as the *iteration leaders*. We refer the reader to [49] for the accountability gadget specific definitions of the timeout parameter T_{tmout} and the confirmation delay T_{bft} of the BFT protocol. We highlight that honest iteration leaders propose only the blocks B^* that are *confirmed* in their view of ch_{ava} , i.e., $B^* \leq B^{\kappa}$ for B identified in Alg. 1, ll. 8, 22, 28 run using ch_{ava} . Similarly, honest validators send accepting gadget votes only for the checkpointing proposals that are *confirmed* in their view of ch_{ava} . We set T_{chkpt} , the time gap between the accountability gadget iterations, to be at least $6\Delta(\kappa+1) + T_{\text{tmout}} + T_{\text{bft}}$ (this is necessary for proving the ebb-and-flow property as will be evident in the following proofs). This makes the upper bound T_{upper} on the total duration of an iteration $T_{\text{chkpt}} + T_{\text{tmout}} + T_{\text{bft}} = 6\Delta(\kappa+1) + 2(T_{\text{tmout}} + T_{\text{bft}}) = \Theta(\kappa)$.

We first show that ch_{ava} remains secure under synchrony in the sleepy network model, despite the added gadget.

Proposition 2. *Suppose a $(\frac{1}{2}, 3\Delta)$ -compliant execution of Goldfish in the synchronous sleepy network model of Sec. 2.2. If a block B is observed to be checkpointed by an honest validator for the first time at some round r , then B is in the common prefix of the chains identified in Alg. 1, ll. 8, 22, 28 right before round r by all awake honest validators.*

PROOF. Since the execution is $(\frac{1}{2}, 3\Delta)$ -compliant, for a block to become checkpointed, at least one honest validator must have sent an accepting gadget vote for that block. Let B_i denote the sequence of checkpointed blocks listed in the order of the rounds r_i at which, an awake honest validator observed B_i to be checkpointed for the first time. Proof is by induction on the indices of these blocks.

Induction Hypothesis: B_i is in the common prefix of the chains identified in Alg. 1, ll. 8, 22, 28 right before round r_i by all awake honest validators, and stays so until at least round r_{i+1} .

Base Case: Since an honest validator sends an accepting gadget vote only for a confirmed block (i.e., κ slots deep), B_1 must have been confirmed by an honest validator at some slot t_1 before round r_1 . As all honest validators start the fork-choice at the genesis block prior to r_1 and B_1 is confirmed in an honest view, it is in the prefix of a block proposed by an honest leader by Lem. 1 and Thm. 1. Hence, B_1 is in the common prefix of the chains identified in Alg. 1, ll. 8, 22, 28 right before round r_1 by all awake honest validators. It also stays in the common prefix until at least round r_2 .

Inductive Step: By the induction hypothesis, checkpointing of the blocks B_1, \dots, B_{i-1} does not alter the fork-choice rule at Alg. 3, l. 2 for any awake honest validator. Hence, by the same reasoning above, B_i is in the common prefix of the chains identified in Alg. 1, ll. 8, 22, 28 right before round r_i by all awake honest validators, and stays so until at least round r_{i+2} . \square

Lemma 7 (Safety and liveness of ch_{ava} under synchrony). *Suppose a $(\frac{1}{2}, 3\Delta)$ -compliant execution of Goldfish in the synchronous sleepy network model of Sec. 2.2. Then, w.o.p., the available ledger ch_{ava} satisfies 1/2-safety and 1/2-liveness (at all times).*

PROOF. By Prop. 2, checkpointing of blocks does not alter the fork-choice rule at Alg. 3, l. 2 for any awake honest validator. Concretely, if the honest validators started the fork-choice rule from the genesis block at all rounds instead of the latest checkpoint in view, then they would end up with the same execution. Thus, the security of ch_{ava} follows from Thm. 2. \square

We next demonstrate the liveness of ch_{acc} after $\max(\text{GST}, \text{GAT})$. In the subsequent analysis, the total number of validators is denoted by n (cf. Sec. 2). The accountability gadget is instantiated with a BFT protocol that has an accountable safety resilience of $n/3$.

Proposition 3 (Prop. 2 of [49]). *The BFT protocol satisfies $n/3$ -liveness after $\max(\text{GST}, \text{GAT})$ with transaction confirmation time $T_{\text{bft}} < \infty$.*

Proposition 4 (Prop. 3 of [49]). *Consider a $(\frac{1}{3}, 3\Delta)$ -compliant execution of Goldfish in the partially synchronous sleepy network model of Sec. 2.2. Suppose a block from iteration c was checkpointed in the view of an honest validator at round r . Then, every honest validator enters iteration $c+1$ by round $\max(\text{GST}, \text{GAT}, r) + \Delta$.*

Let c' be the largest iteration such that a block B was checkpointed in the view of some honest validator before $\max(\text{GAT}, \text{GST})$. (Let $c' = 0$ and B be the genesis block if there does not exist such an iteration.) If an honest validator enters an iteration $c'' > c'$ at some round $r \geq \max(\text{GAT}, \text{GST}) + \Delta + T_{\text{chkpt}}$, every honest validator enters iteration c'' by round $r + \Delta$.

PROOF. Suppose a block B from iteration c was checkpointed in the view of an honest validator id at round r . Then, there are over $2n/3$ accepting gadget votes for B from iteration c on $\text{LOG}_{\text{bft}, \text{id}}^r$, the output ledger of the BFT protocol in id 's view at round r . All gadget votes and BFT protocol messages observed by id by round r are delivered to all other honest validators by round $\max(\text{GST}, \text{GAT}, r) + \Delta$. Hence, by the safety of the BFT protocol when $f < n/3$, for any

honest validator id' , the ledger $\text{LOG}_{\text{bft},id}^r$ is the same as or a prefix of the ledger observed by id' at round $\max(\text{GST}, \text{GAT}, r) + \Delta$. Thus, for any honest validator id' , there are over $2n/3$ accepting gadget votes for B from iteration c on LOG_{bft} at round $\max(\text{GST}, \text{GAT}, r) + \Delta$. This implies every honest validator enters iteration $c + 1$ by round $\max(\text{GST}, \text{GAT}, r) + \Delta$.

Finally, by the reasoning above, all honest validators enter iteration $c' + 1$ by round $\max(\text{GAT}, \text{GST}) + \Delta$. Thus, entrance time of the honest validators to subsequent iterations have become synchronized by round $\max(\text{GAT}, \text{GST}) + \Delta + T_{\text{chkpt}}$: If an honest validator enters an iteration $c'' > c'$ at some round $r \geq \max(\text{GAT}, \text{GST}) + \Delta + T_{\text{chkpt}}$, every honest validator enters iteration c'' by round $r + \Delta$. Similarly, if a block from iteration c'' is first checkpointed in the view of an honest validator at some round after $\max(\text{GAT}, \text{GST}) + \Delta + T_{\text{chkpt}}$, then it is checkpointed in the view of all honest validators within Δ rounds. \square

Lemma 8 (Liveness of ch_{acc} , analogue of Thm. 4 of [49]). *Consider a $(\frac{1}{3}, 3\Delta)$ -compliant execution of Goldfish in the partially synchronous sleepy network model of Sec. 2.2. Suppose ch_{ava} is secure (safe and live) after some round $T_{\text{heal}} \geq \max(\text{GST}, \text{GAT}) + \Delta + T_{\text{chkpt}}$. Then, w.o.p., ch_{acc} satisfies $n/3$ -liveness after round T_{heal} with transaction confirmation time $T_{\text{conf}} = \Theta(\kappa^2)$.*

PROOF. By Prop. 3, LOG_{bft} is live with transaction confirmation time T_{bft} after $\max(\text{GST}, \text{GAT})$, a fact we will use subsequently.

Let c' be the largest iteration such that a block B was checkpointed in the view of some honest validator before $\max(\text{GAT}, \text{GST})$ (Let $c' = 0$ and B be the genesis block if there does not exist such an iteration). Then, by Prop. 4, entrance times of the honest validators to subsequent iterations become synchronized by round $\max(\text{GAT}, \text{GST}) + \Delta + T_{\text{chkpt}}$: If an honest validator enters an iteration $c > c'$ at some round $r \geq \max(\text{GAT}, \text{GST}) + \Delta + T_{\text{chkpt}}$, every honest validator enters iteration c by round $r + \Delta$.

Suppose an iteration $c > c'$ has an honest iteration leader $L^{(c)}$, which sends a checkpoint proposal, denoted by \hat{b}_c , at some round $r > T_{\text{heal}} + T_{\text{chkpt}}$. The proposal \hat{b}_c is received by every honest validator by round $r + \Delta$. Since the entrance times of the validators are synchronized by $T_{\text{heal}} \geq \max(\text{GST}, \text{GAT}) + \Delta + T_{\text{chkpt}}$, every honest validator sends a gadget vote by round $r + \Delta$. By Lem. 10, $\hat{b}_c \leq B^{\lceil \kappa}$ for any B identified in Alg. 1, ll. 8, 22, 28 by any awake honest validator after r . Moreover, \hat{b}_c is a descendant all of the checkpoints seen by the honest validators until then. Consequently, at iteration c , every honest validator sends a gadget vote accepting \hat{b}_c by round $r + \Delta$, all of which appear within LOG_{bft} in the view of every honest validator by round $r + \Delta + T_{\text{bft}}$. Thus, \hat{b}_c becomes checkpointed in the view of every honest validator by round $r + \Delta + T_{\text{bft}}$. (Here, we assume that T_{tmout} was chosen large enough for $T_{\text{tmout}} > \Delta + T_{\text{bft}}$ to hold.)

Since $r > T_{\text{heal}} + T_{\text{chkpt}}$, by Lem. 10, \hat{b}_c contains at least one honest block since an earlier checkpointed block in its prefix from before iteration c . This implies that the prefix of \hat{b}_c contains at least one fresh honest block that enters ch_{acc} by round $r + \Delta + T_{\text{bft}}$.

Next, we show that an adversarial leader cannot make an iteration last longer than $\Delta + T_{\text{tmout}} + T_{\text{bft}}$ for any honest validator after the initial T_{chkpt} period elapsed. Indeed, if an honest validator id

enters an iteration c at round $r - T_{\text{chkpt}}$, by round $r + T_{\text{tmout}}$, either it sees a block (potentially \perp) become checkpointed for iteration c , or it sends a reject vote for iteration c . In the first case, every honest validator sees a block checkpointed for iteration c by round at most $r + T_{\text{tmout}} + \Delta$. In the second case, rejecting gadget votes from over $2n/3 > n/3$ validators appear in LOG_{bft} in the view of every honest validator by round at most $r + T_{\text{tmout}} + \Delta + T_{\text{bft}}$. Hence, a new checkpoint, potentially \perp , is output in the view of every honest validator by round $r + T_{\text{tmout}} + \Delta + T_{\text{bft}}$.

Finally, we observe that except with probability $(1/3)^\kappa$, there exists a checkpoint iteration with an honest leader within κ consecutive iterations. Since an iteration lasts at most $\max(\Delta + T_{\text{tmout}} + T_{\text{bft}}, \Delta + T_{\text{chkpt}} + T_{\text{bft}}) \leq \Delta + T_{\text{chkpt}} + T_{\text{tmout}} + T_{\text{bft}} = \Theta(\kappa)$ rounds, and a new checkpoint containing a fresh honest block in its prefix appears when an iteration has an honest leader (Lem. 10), w.o.p., any transaction received by an honest validator at round t appears within ch_{acc} in the view of every honest validator by round at most $t + \kappa(\Delta + T_{\text{tmout}} + T_{\text{bft}} + T_{\text{chkpt}})$. Hence, via a union bound over the total number of iterations (which is a polynomial in κ), we observe that if ch_{ava} satisfies security after some round T_{heal} , then w.o.p., ch_{acc} satisfies liveness after T_{heal} with a transaction confirmation time $T_{\text{conf}} = \Theta(\kappa^2)$. \square

The latency expression $T_{\text{conf}} = \Theta(\kappa^2)$ stated in Lem. 8 is a *worst-case* latency to guarantee that an honest block enters the accountable, final prefix ledger ch_{acc} with overwhelming probability. In the expression, the first κ term comes from the requirement to have $T_{\text{chkpt}} = \Theta(\kappa)$ slots in between the accountability gadget iterations, and the second κ term comes from the fact that it takes $\Theta(\kappa)$ iterations for the accountability gadget to have an honest iteration leader except with probability $\text{negl}(\kappa)$. The accountability gadget protocol asks honest validators to wait for $T_{\text{chkpt}} = \Theta(\kappa)$ slots in between iterations to ensure the security of the protocol, reasons for which will be evident in the proof of Lem. 10.

Unlike the worst-case latency, the expected latency for an honest block to enter ch_{acc} after ch_{ava} regains its security would be $\Theta(\kappa)$ as each checkpointing iteration has an honest leader with probability at least $2/3$. In this context, the latency of $\Theta(\kappa)$ is purely due to the requirement to have $T_{\text{chkpt}} = \Theta(\kappa)$ slots in between the accountability gadget iterations. Here, waiting for T_{chkpt} slots in between iterations guarantees the inclusion of a new honest block in ch_{ava} , which in turn appears in the prefix of the next checkpoint, implying a liveness event whenever there is an honest iteration leader.

Lem. 8 requires the available ledger ch_{ava} to eventually regain security under partial synchrony when there are less than $n/3$ adversarial validators. Towards this goal, we first analyze the gap and recency properties, the core properties that must be satisfied by the accountability gadget for recovery of security of ch_{ava} . The gap property states that the blocks are checkpointed sufficiently apart in time, controlled by the parameter T_{chkpt} :

Proposition 5 (Gap property, analogue of Prop. 4 of [49]). *Consider a $(\frac{1}{3}, 3\Delta)$ -compliant execution of Goldfish in the partially synchronous sleepy network model of Sec. 2.2. Given any round interval of size T_{chkpt} , no more than a single block can be checkpointed in the interval in the view of any honest validator.*

Proof of Prop. 5 follows from the fact that upon observing a new checkpoint that is not \perp for an iteration, honest validators wait for T_{chkpt} rounds before sending gadget votes for the checkpoint proposal of the next iteration, and there cannot be two conflicting checkpoints for the same iteration in the view of any honest validator.

As in [49] and [53], we state that a block B^* checkpointed at iteration c and round $r > \max(\text{GST}, \text{GAT})$ in the view of an honest validator id is T_{rcnt} -recent if $B^* \leq B^{\lceil \kappa}$ for B identified in Alg. 1, l. 28 by id' at some round within $[r - T_{\text{rcnt}}, r]$. Then, we can express the recency property as follows:

Lemma 9 (Recency property, analogue of Lem. 1 of [49]). *Consider a $(\frac{1}{3}, 3\Delta)$ -compliant execution of Goldfish in the partially synchronous sleepy network model of Sec. 2.2. Every checkpointed block proposed after $\max(\text{GST}, \text{GAT})$ is T_{rcnt} -recent for $T_{\text{rcnt}} = \Delta + T_{\text{tmout}} + T_{\text{bft}}$.*

PROOF. By the proof of Lem. 8, if a block B proposed after $\max(\text{GST}, \text{GAT})$ is checkpointed in the view of an honest validator at some round r , it should have been proposed after round $r - (\Delta + T_{\text{tmout}} + T_{\text{bft}})$. Moreover, over $2n/3$ validators must have sent accepting gadget votes for B by round r . Let id denote such an honest validator. It would vote for B only after it sees the checkpoint proposal for iteration c , i.e., after round $r - T_{\text{rcnt}} = r - (\Delta + T_{\text{tmout}} + T_{\text{bft}})$, and only if the proposal is confirmed in its view. Hence, B must be κ slots deep in the chain returned at Alg. 1, l. 28 by validator id at some round within $[r - T_{\text{rcnt}}, r]$. This concludes the proof that every checkpointed block proposed after $\max(\text{GST}, \text{GAT})$ is T_{rcnt} -recent. \square

Lemma 10 (Healing property, analogue of Thm. 5 of [49]). *Consider a $(\frac{1}{3}, 3\Delta)$ -compliant execution of Goldfish in the partially synchronous sleepy network model of Sec. 2.2. Then, ch_{ava} is secure with transaction confirmation time $T_{\text{chkpt}} + T_{\text{tmout}} + T_{\text{bft}} = \Theta(\kappa)$ after round $\max(\text{GAT}, \text{GST}) + \Delta + 2T_{\text{chkpt}}$.*

Moreover, for the iteration proposal \hat{b}_c of an honest iteration leader broadcast at round r , it holds that $\hat{b}_c \leq B^{\lceil \kappa}$ for any B identified in Alg. 1, ll. 8, 22, 28 by any awake honest validator after r , and \hat{b}_c contains a fresh honest block that is not in the prefix of any checkpoint from before iteration c .

PROOF. By [49, Thm. 3], ch_{acc} provides accountable safety with resilience $n/3$ except with probability $\text{negl}(\lambda)$ in the partially synchronous sleepy network model. As the execution is $(\frac{1}{3}, 3\Delta)$ -compliant, w.o.p., no two checkpoints observed by awake honest validators conflict.

Let c be the largest iteration such that a block B was checkpointed in the view of some honest validator before $\max(\text{GAT}, \text{GST})$. (Let $c = 0$ and B be the genesis block if there does not exist such an iteration.) Then, by Prop. 4, if an honest validator enters an iteration $c' > c$ at some round $r \geq \max(\text{GAT}, \text{GST}) + \Delta + T_{\text{chkpt}}$, every honest validator enters iteration c by round $r + \Delta$. Let c' be the first iteration such that the first honest validator to enter c' enters it after round $\max(\text{GAT}, \text{GST}) + \Delta + T_{\text{chkpt}}$ (e.g., at some round r such that $\max(\text{GAT}, \text{GST}) + \Delta + T_{\text{chkpt}} < r < \max(\text{GAT}, \text{GST}) + \Delta + 2T_{\text{chkpt}}$). Then, all honest validators enter iteration c' and agree on the last checkpointed block within Δ rounds. Subsequently, the honest validators wait for T_{chkpt} rounds before casting any gadget vote for

a checkpoint proposal of iteration c' , during which no block can be checkpointed (Prop. 5, gap property).

By Lem. 1, w.o.p., the slot interval of length κ starting after round $r + \Delta$ contains a slot t with an honest leader and proposal P^* . After round $r \geq \text{GST}$, all messages broadcast by honest validators are received by all honest validators within Δ rounds. As honest validators agree on the last checkpointed block during the interval $[r + \Delta, r + T_{\text{chkpt}}]$, by the absence of new checkpoints, the GHOST-Eph fork-choice rule starts at the same last checkpointed block for all honest validators during the interval (Alg. 2, l. 2). Then, by Lem. 1, w.o.p., $P^*.B \leq B$ for any B identified in Alg. 1, ll. 8, 22, 28 by any awake honest validator in any round after $3\Delta t + 2\Delta$, until at least a new block is checkpointed in the view of an honest validator.

By Lem. 9 (recency property), the next block checkpointed in the view of an honest validator (which happens earliest at some iteration $c'' \geq c'$ and round $r' \geq r + T_{\text{chkpt}}$ by Prop. 5, the gap property) must have been confirmed by some honest validator id at some round within $[r' - T_{\text{rcnt}}, r']$, where $r' - T_{\text{rcnt}} \geq r + 6\Delta\kappa + 4\Delta$. Hence, the new checkpointed block is κ slots deep in the chains identified in Alg. 1, ll. 8, 22, 28 by id , and is a descendant of $P^*.B$. This implies $P^*.B \leq B$ for any B identified in Alg. 1, ll. 8, 22, 28 by any awake honest validator in any round after $3\Delta t + 2\Delta$ indefinitely.

Note that if the iteration leader was honest, for its proposal \hat{b}_c broadcast at some round r'' , it holds that $\hat{b}_c \leq B^{\lceil \kappa}$ for any B identified in Alg. 1, ll. 8, 22, 28 by any awake honest validator after round r . Moreover, $P^*.B \leq \hat{b}_c$, implying that honest checkpoint proposals contain fresh honest blocks in their prefixes.

Finally, we extend the above argument to future checkpoints by induction. Let B_n denote the sequence of checkpointed blocks, ordered by their iteration numbers $c_n \geq c'$, $c_1 = c''$. The rounds r_n , at which the blocks B_n are first checkpointed in the view of an honest validator satisfy the relation $r_{n+1} \geq r_n + T_{\text{chkpt}}$ and $r_1 = r''$. Via the inductive assumption and the reasoning above, w.o.p., in each interval $[r_n + \Delta, r_{n+1} - T_{\text{rcnt}}]$, there exists a slot t_n with an honest leader and proposal P_n such that $P_n.B \leq B$ for any B identified in Alg. 1, ll. 8, 22, 28 by any awake honest validator in any round after $3\Delta t_n + 2\Delta$ indefinitely. Hence, for a sufficiently large confirmation time exceeding the maximum possible length of an iteration (i.e., $T_{\text{conf}} \geq T_{\text{chkpt}} + T_{\text{tmout}} + T_{\text{bft}}$), these honest blocks imply the security of the Goldfish protocol after round $\max(\text{GAT}, \text{GST}) + \Delta + 2T_{\text{chkpt}}$. \square

Note that Thm. 1 holds for the honest blocks proposed during the intervals $[r_n + \Delta, r_{n+1} - T_{\text{rcnt}}]$ as all honest validators agree on the latest checkpoint during these intervals.

PROOF OF THM. 4. We first show the property **P1**, namely, the accountable safety and liveness of the accountable, final prefix ledger ch_{acc} under partial synchrony in the sleepy model. By [49, Thm. 3], ch_{acc} provides accountable safety with resilience $n/3$ except with probability $\text{negl}(\lambda)$ under partial synchrony in the sleepy model. By Lem. 10, under the same model, the available ledger ch_{ava} is secure after round $\max(\text{GAT}, \text{GST}) + \Delta + 2T_{\text{chkpt}}$. Using this fact and Lem. 8, we can state that, w.o.p., ch_{acc} satisfies liveness after round $\max(\text{GAT}, \text{GST}) + \Delta + 2T_{\text{chkpt}}$ with transaction confirmation time $T_{\text{conf}} = \Theta(\kappa^2)$.

Finally, the property **P2** follows from Lem. 7, and **Prefix** follows

by construction of the ledgers ch_{acc} and ch_{ava} . This concludes the proof of the ebb-and-flow property. \square

B EQUIVOCATION DISCOUNTING TO MITIGATE SPAMMING

Goldfish deals with equivocating votes simply by accepting all of them, but counting at most one per subtree (Alg. 2, l. 7). This does not give any additional fork-choice power to an equivocating validator, and it does not allow for irreconcilable splits of honest validators’ views, which would be the case with a naive first-seen (or last-seen) approach. Instead, it guarantees that honest validators can always end up with the same view, in particular through the vote buffering mechanism, and that their fork-choice outputs agree when they do. Nonetheless, this approach is vulnerable to spamming attacks, because it requires validators to accept all the votes they receive. Even a single adversarially controlled validator can be used to create an arbitrarily large number of equivocating votes at a slot, with the goal of creating network congestion and making it impossible for honest validators to download all of the other votes in time.

Equivocations are attributable faults, punishable by slashing *a posteriori*, but this does not prevent the attack vector *a priori* given that only one validator is required for it, and that there is no immediate recovery, because the same validator can continue producing equivocating attestations in subsequent slots as well. It is perhaps possible to mitigate this attack vector without breaking the strong properties of vote buffering with approaches similar to those of [42], *i.e.*, by introducing more refined rules for downloading and forwarding consensus messages. The approach we take is instead to introduce *equivocation discounting*. This general technique is already present in the current implementation of PoS Ethereum, but the ephemerality of votes in Goldfish allows for a simpler rule, with clear bounds on the number of messages required for honest views to converge. This is particularly important in order to have guarantees about the functioning of the vote buffering technique, and in turn about the security of the whole protocol, which heavily relies on reorg resilience. We formalize the simple equivocation discounting rule here, as a combination of a modification to the GHOST-Eph fork-choice, a download rule, and a validity condition for proposals.

Equivocation discounting.

- (a) *Fork-choice discounting:* When running the GHOST-Eph fork-choice rule at slot t , only count the valid slot $t - 1$ votes from those validators for which your btree contains a single valid slot $t - 1$ vote, *i.e.*, those which are not viewed to have equivocated at slot $t - 1$.
- (b) *Download rule:* Only download (or forward as part of the peer-to-peer gossip layer) votes from the current and prior slots, and at most two votes per *eligible* validator (*i.e.*, the opened ticket (id, t) for the validator id is winning for the tag $(\text{vote}, \text{thr}_v)$, cf. Sec. 5).
- (c) *Validity condition for proposals:* A proposal whose btree contains more than two valid votes for the same slot from some validator is invalid, and so is one which contains any invalid vote.

Discounting equivocations from the fork-choice preserves the property that there cannot be irreconcilable splits of validator views, because all that is needed for convergence is agreement on the list of equivocators from the previous slot, which in turn only needs all views to have compatible equivocation evidence, *i.e.*, pairs of equivocating votes for the same list of equivocating validators. The download rule and validity condition ensure that a validator only ever needs to download at most two votes per subsampled validator of the current and previous slot. Setting the subsampling parameters so that this is manageable, we can ensure that equivocations cannot succeed at creating network congestion sufficient to prevent the functioning of vote buffering. Previously, this meant guaranteeing that an honest proposer’s btree be a superset of honest validators’ btrees. Instead, the success of vote buffering now only requires that a leader’s view of votes from voters which have not equivocated in the last slot is a superset of the validators’ views of such votes, and so is its view of *the list of equivocators from the previous slot*. Agreement on these two is sufficient for agreement on the fork-choice output, *i.e.*, Lem. 2 still holds. Note that the leader still only needs to include its btree in the proposal message, because following the download rule guarantees that it will contain exactly all valid votes from validators which have not equivocated in the previous slot, together with a pair of votes, *i.e.*, equivocation evidence, for validators which have.

The security analysis for Goldfish with equivocation discounting is then the same as that for vanilla Goldfish. Vote buffering implies that all honest validators vote together when the proposal with the minimum precedence is honest, as in Lem. 2, and all honest validators voting together implies that the proposal is never reorged, as in Lem. 3. The latter is not affected by equivocation discounting, because it relies on the valid votes of honest validators, which do not equivocate. From these two properties, we obtain reorg resilience as in Thm. 3, and from reorg resilience, we eventually obtain safety and liveness.

Optimistic fast confirmations are also compatible with equivocation discounting, without any loss of resilience. Liveness and fast confirmation of honest proposals follow from Thm. 7, since equivocation discounting plays no role in it. For safety, the key ingredient is Lem. 5, from which Thm. 6 follows unchanged. We thus prove Lem. 5 here for Goldfish with equivocation discounting, by making a very small modification to the argument:

PROOF OF LEM. 5 WITH EQUIVOCATION DISCOUNTING. By Prop. 1, w.o.p., the number of adversarial validators at round $4\Delta(t + 1) + \Delta$, eligible to vote at slot t , is less than $\frac{1}{2}n \text{thr}_v$. An eligible awake honest validator sends a single slot t vote at round $4\Delta t + \Delta$, implying that over $(\frac{3}{4} + \frac{\epsilon}{2})n \text{thr}_v - \frac{1}{2}n \text{thr}_v = (\frac{1}{4} + \frac{\epsilon}{2})n \text{thr}_v$ validators broadcast a single slot t vote by round $4\Delta(t + 1) + \Delta$, and that is for a descendant of B . By Prop. 1, w.o.p., for all slots t , there can be at most $(1 + \epsilon)n \text{thr}_v$ validators that are eligible to vote at t . Hence, the number of valid slot t votes for the descendants of any block B' conflicting with B , and which are from validators which have not also cast one of the $(\frac{3}{4} + \frac{\epsilon}{2})n \text{thr}_v$ votes for B , must be less than $(1 + \epsilon)n \text{thr}_v - (\frac{3}{4} + \frac{\epsilon}{2})n \text{thr}_v = (\frac{1}{4} + \frac{\epsilon}{2})n \text{thr}_v$ at any given round. The validator id^* broadcasts B and over $(\frac{3}{4} + \frac{\epsilon}{2})n \text{thr}_v$ valid votes for it (in pieces) at round $4\Delta t + 2\Delta$. Each honest validator, awake at round $4\Delta(t + 1) + \Delta$ and eligible to vote at slot $t + 1$, observes

these votes in its bvtree at the round of voting (Alg. 5, l. 12). Upon invoking the GHOST-Eph fork-choice rule at any of the rounds $4\Delta t + 3\Delta$, $4\Delta(t + 1)$ or $4\Delta(t + 1) + \Delta$, using only the votes from validators which are not seen to be equivocating at slot $t - 1$, the votes for the descendants of any block B' conflicting with B are then less than $(\frac{1}{4} + \frac{\epsilon}{2})n \text{ thr}_v$, and the votes for descendants of B are over $(\frac{1}{4} + \frac{\epsilon}{2})n \text{ thr}_v$. This implies that all honest validators, awake at round $4\Delta(t + 1) + \Delta$ and eligible to vote at slot $t + 1$, all vote for B or one of its descendants at slot $t + 1$. \square

Invalid proposals and pieces are not merged into the bvtrees by the honest validators. Proposals and pieces whose validity is in limbo can be purged from the buffer in certain cases. A slot t block in limbo can be purged after slot $t + T_{\text{conf}}$ since if the block has not been validated and included as part of the canonical GHOST-Eph chain by that slot, it will never appear as part of the confirmed Goldfish chain. A slot t vote in limbo can be purged at slot $t + 2$ since it will not have any effect on the fork-choice rule slot $t + 2$ onwards due to vote expiry. Purging proposals and pieces in limbo further reduces the storage requirement expected of the Goldfish validators.

C COMPARISON WITH CURRENT POS ETHEREUM

Goldfish is a simple, provably secure, dynamically available and reorg resilient protocol which minimally differs from the LMD GHOST component of the current implementation of the Gasper protocol, responsible for the consensus of Ethereum’s beacon chain. Gasper has so far defied formal security analysis even in the simpler, full participation setting, not least because of its complexity. Moreover, it is not reorg resilient even in that setting, and it is not dynamically available. We first analyze these shortcomings and their origins in various components of the protocol, then discuss incorporating Goldfish into Gasper.

C.1 Limitations of Gasper

Interaction of LMD GHOST and Casper FFG. The combination of Goldfish with the accountability gadget in Sec. 3 follows the generic construction of [49], which is proven to be secure for any appropriately secure dynamically available protocol and accountable BFT protocol. On the other hand, the combination of LMD GHOST and Casper FFG in HLMD GHOST, the hybrid fork-choice rule of [9], is ad-hoc and complicated to reason about. Firstly, it is known to be susceptible to a *bouncing attack* [40]. Instead of LMD GHOST starting its fork-choice iteration from the last block *finalized* by Casper FFG, it starts from the last *justified* block in the terminology of Casper FFG, *i.e.*, the last block that has been the target of FFG votes by a supermajority of *all* n validators. This is sufficient to ensure accountable safety of the finalized checkpoints; however, it hinders safety of the available ledger ch_{ava} (after $\max(\text{GST}, \text{GAT})$) under partial synchrony in the sleepy model, in particular negating the healing property (Lem. 10) of ch_{ava} , preventing us from proving the ebb-and-flow property. The current mitigation for the bouncing attack causes other problems such as the splitting attack [41], akin to the balancing attacks [47]. It is perhaps possible to resolve these through the use of techniques akin to vote buffering, to avoid

the adversary being able to consistently split honest views. Even then, it is not at all clear that the bouncing attack cannot still be executed by exploiting other aspects of the complex interaction of LMD GHOST and Casper FFG. One such aspect is the fact that the FFG votes at any *Ethereum epoch* point at the last epoch boundary block of that epoch, regardless of its confirmation status by the underlying LMD GHOST rule. (In fact, there is no confirmation rule specified for LMD GHOST.) Another one is that the FFG voting schedule is staggered throughout an epoch, as FFG votes are cast together with LMD GHOST votes, so it is not clear how to ensure that the views of honest validators when casting FFG votes are consistent.

On-chain inclusion of consensus messages. Another peculiarity of Gasper is that the inclusion of consensus messages (FFG votes) into blocks is crucial to the consensus process itself. In particular, its hybrid fork-choice rule filters out all branches whose state (at the tip) has not justified the latest justified checkpoint, meaning that either not enough FFG votes have been included, or that a state transition processing them has not yet occurred. This rule makes the protocol even harder to reason about and formally analyze, and also introduces attack vectors similar to those already mentioned for the bouncing attack, *i.e.*, related to the adversary obtaining private justifications. Future work is required to carefully analyze its role in the protocol, and whether it can be removed from it.

Stale votes in LMD GHOST. Without vote expiry, the votes of honest asleep validators can be weaponized by an adversary controlling a small fraction of the validator set to execute an arbitrarily long reorg. This implies that the protocol is not dynamically available with *any* confirmation rule with finite confirmation time T_{conf} . Consider for example a validator set of size $n = 2m + 1$, and a partition of the validator set into three sets, V_1, V_2, V_3 , with $|V_1| = |V_2| = m$ and $|V_3| = 1$. The validators in V_1, V_2 are all honest, while the one in V_3 is adversarial. Suppose that the adversarial validator in V_3 is the leader of slots t , and that it broadcasts two proposals, with conflicting blocks B_1 and B_2 . It does so in such a way that validators in V_1 see only B_1 before voting, and validators in V_2 only B_2 . Validators in V_1 then vote for B_1 , and so does the adversarial validator, while validators in V_2 vote for B_2 . B_1 becomes canonical, since it has received $m + 1$ votes. The adversary then puts all validators in V_2 to sleep, and they do not become awake for the remainder of the protocol. The adversarial validator does not cast any more votes for a while. Meanwhile, validators in V_1 , keep voting for descendants of B_1 . After waiting for $> T_{\text{conf}}$ slots, the adversarial validator votes for B_2 . Since the m latest votes of the validators in V_2 are still for B_2 , it now has $m + 1$ votes and becomes canonical, resulting in all awake honest validators experiencing a reorg of all blocks confirmed after slot t . If there are no such blocks, liveness is violated, and otherwise safety is violated.

Proposer boost. Proposer boost is not compatible with dynamic availability, because the artificial fork-choice weight it temporarily provides to proposals is independent of participation: the lower the participation, the more powerful the boost is relative to the weight of real attestations from awake validators, and thus the more it can be exploited by the adversary. When the weight of awake honest validators is less than the boost, the adversary has complete

control of the fork-choice during the slots in which it is elected as the leader.

Reorg resilience. Even in the setting of full participation, where the adversary cannot take advantage of votes of asleep validators, LMD GHOST lacks reorg resilience. This is firstly due to subsampling without vote expiry, because it allows the adversary to accumulate fork-choice weight by withholding blocks and attestations, *i.e.*, to execute *ex ante* reorgs [54]. Without subsampling, LMD GHOST is indeed reorg resilient in the full participation setting, *if proposer boost is replaced by vote buffering*. In fact, Thm. 3 obtains reorg resilience as a consequence of two properties, Lems. 2 and 3, respectively the property that all honest awake validators vote for an honest proposal, and the property that all honest validators voting together guarantee the inclusion of honest blocks in the canonical GHOST-Eph chain, both of which also hold for LMD GHOST with vote buffering.

With proposer boost, LMD GHOST is not reorg resilient for $\beta \geq \frac{1}{3}$, even in the full participation setting and without subsampling, because those two properties are in conflict for such β , for any boost value W_p . The first property only holds if $W_p > \beta$, *i.e.*, the amount of adversarial votes which might be withheld in an *ex ante* reorg attempt. On the other hand, the second property only holds if $W_p + \beta < 1 - \beta$, because otherwise an adversarial proposer can make use of boost to conclude an *ex post* reorg. Therefore, we can only have reorg resilience for $\beta < \frac{1}{3}$, by setting $W_p = \frac{1}{3}$.

C.2 Bringing Goldfish to Gasper

Replacing LMD GHOST with Goldfish. LMD GHOST could be replaced by Goldfish in Gasper [9], with only minor changes needed. Firstly, multiple potential leaders are elected through VRFs instead of a single leader through RANDAO. This results in additional bandwidth consumption due to multiple proposals being propagated, but helps maintain the confirmation time as a constant as participation drops. Moreover, the election process via VRFs is not biasable by the participants, which is not the case with RANDAO [26], and it automatically provides privacy to the leader before they reveal themselves, protecting them from targeted denial-of-service (DoS) attacks. It is not clear whether VRFs can also be utilized for subsampling, because the functioning of the beacon chain heavily relies on the aggregation of signatures, in order to support a large validator set. Augmenting the votes with VRF proofs is not compatible with aggregation, since the latter requires all messages to be the same (compatibility would require an aggregation scheme for VRF outputs). Nonetheless, RANDAO could still be used for subsampling, though its biasability would affect the tolerable adversarial fraction of validators. Some care has to also be taken to make attestation aggregation compatible with vote buffering.

Combination of Goldfish and FFG. The protocol resulting from replacing LMD GHOST with Goldfish in Gasper still does not satisfy all the properties we want, and which we have proved for the combination of Goldfish with an accountability gadget, when following the construction of [49]. In particular, it does not escape the negative result from App. C.1, due to the bouncing attack under low participation. Since Casper FFG is not a complete protocol, as it lacks a message schedule for proposals and votes, more work is

needed to understand if it is possible to use it as the BFT protocol in the aforementioned construction.

An alternative approach, and perhaps easier to integrate in the protocol if successful, is to try to achieve security by adapting the construction to the protocol, rather than trying to use its black box approach. A simple modification is to stipulate that the honest validators cast their FFG votes only for the blocks that are confirmed by Goldfish, *i.e.*, the blocks that are part of their available ledgers ch_{ava} . As already mentioned, this is different from the current PoS Ethereum specification for Casper FFG, where the FFG votes at any *Ethereum epoch* point at the last epoch boundary block of that epoch, regardless of its confirmation status by the underlying LMD GHOST rule. Unfortunately, this is not sufficient to guarantee the security of the accountable, final prefix ledger ch_{acc} outputted as the prefix of the finalized Casper FFG checkpoints, again due to the bouncing attack. To avoid the latter, and inspired by the aforementioned construction, a second modification is to start the iteration of the hybrid fork-choice from the latest *finalized* checkpoint, rather than the latest justified. The question of whether this is sufficient to ensure security is left for future work.