Byzantine Fault Tolerance For Distributed Ledgers Revisited

YONGGE WANG, UNC Charlotte, USA

The problem of Byzantine Fault Tolerance (BFT) has received a lot of attention in the last 30 years. Due to the popularity of Proof of Stake (PoS) blockchains in recent years, several BFT protocols have been deployed in the large scale of Internet environment. We analyze several popular BFT protocols such as Capser FFG / CBC-FBC for Ethereum 2.0 and GRANDPA for Polkadot. Our analysis shows that the security models for these BFT protocols are slightly different from the models commonly accepted in the academic literature. For example, we show that, if the adversary has a full control of the message delivery order in the underlying network, then none of the BFT protocols for Ethereum blockchain 2.0 and Polkadot blockchain could achieve liveness even in a synchronized network. Though it is not clear whether a practical adversary could *actually* control and re-order the underlying message delivery system (at Internet scale) to mount these attacks, it raises an interesting question on security model gaps between academic BFT protocols and deployed BFT protocols in the Internet scale. With these analysis, this paper proposes a Casper CBC-FBC style binary BFT protocol and shows its security in the traditional academic security model with complete asynchronous networks. For partial synchronous networks, we propose a multi-value BFT protocol BDLS based on the seminal DLS protocol and show that it is one of the most efficient practical BFT protocols at large scale networks in the traditional academic BFT security model. The implementation of BDLS is available at https://github.com/yonggewang/bdls. Finally, we propose a multi-value BFT protocol XP for complete asynchronous networks and show its security in the traditional academic BFT security model.

CCS Concepts: • **Theory of computation** → **Cryptographic protocols**; • **Security and privacy** → **Distributed systems security**; • **Applied computing** → **Electronic commerce**; • **Computing methodologies** → *Distributed algorithms.*

Additional Key Words and Phrases: Byzantine Fault Tolerance; distributed computing; partial synchronous networks; security models; blockchain

ACM Reference Format:
Yongge Wang. 2022. Byzantine Fault Tolerance For Distributed Ledgers Revisited. *Distrib. Ledger Technol.* 1, 1, Article 1 (April 2022), 28 pages. https://doi.org/10.1xxx/1xxxx.1xxxxx

1 INTRODUCTION

Consensus is hard to achieve in open networks such as partial synchronous networks or complete asynchronous networks. Several practical protocols such as Paxos [11] and Raft [13] have been designed to tolerate $\lfloor \frac{n-1}{2} \rfloor$ non-Byzantine faults. For example, Google, Microsoft, IBM, and Amazon have used Paxos in their storage or cluster management systems. Lamport, Shostak, and Pease [12] and Pease, Shostak, and Lamport [14] initiated the study of reaching consensus in face of Byzantine failures and designed the first synchronous solution for Byzantine agreement. Dolev and Strong [7] proposed an improved protocol in a synchronous network with $O(n^3)$ communication complexity. By assuming the existence of digital signature schemes and a public-key infrastructure, Katz and

Author's address: Yongge Wang, yonwang@uncc.edu, UNC Charlotte, 9201 University City Blvd., Charlotte, NC, USA, 28223-0001.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
© 2022 Association for Computing Machinery.
2769-6472/2022/4-ART1 $15.00
https://doi.org/10.1xxx/1xxxx.1xxxxx

https://mc.manuscriptcentral.com/dlt

Koo [10] proposed an expected constant-round BFT protocol in a synchronous network setting against $\lfloor \frac{n-1}{2} \rfloor$ Byzantine faults.

Fischer, Lynch, and Paterson [9] showed that there is no deterministic protocol for the BFT problem in face of a single failure. Several researchers have tried to design BFT consensus protocols to circumvent the impossibility. The first category of efforts is to use a probabilistic approach to design BFT consensus protocols in completely asynchronous networks. This kind of work was initiated by Ben-Or [2] and Rabin [15] and extended by others such as Cachin, Kursawe, and Shoup [5]. The second category of efforts was to design BFT consensus protocols in partial synchronous networks which was initiated by Dwork, Lynch, and Stockmeyer [8]. Though the network communication model could be different for these protocols, the assumption on the adversary capability is generally same. That is, there is a threshold t such that the adversary could coordinate the activities of the malicious t participating nodes. Furthermore, it is also assumed that the adversary could re-order messages on communication networks.

In recent years, many practical BFT protocols have been designed and deployed at the Internet scale. For example, Ethereum foundation has designed a BFT finality gadget for their Proof of Stake (PoS) blockchain. The current Ethereum 2.0 beacon network uses Casper Friendly Finality Gadget (Casper FFG) [4] and Ethereum foundation has been advocating the "Correct-by-Construction" (CBC) family consensus protocols [23, 24] for their future release of Ethereum blockchain. Similarly, the Polkadot blockchain deployed their home-brew BFT protocol GRANDPA [18]. The analysis in this paper shows that these protocols have an assumption that the adversary cannot control the message delivery order in the underlying networks. Our examples (as noted in the extended abstract [22]) show that if the adversary could control the the message delivery order, then these blockchains could not achieve liveness property. This brings up an interesting question to the research community: what kind of models are appropriate for the Internet scale BFT protocols? Does an adversary have the capability to co-ordinate/control one-third of the participating nodes and to reschedule message delivery order for a blockchain at Internet scale?

Before we have a complete understanding about the impact of the new security assumptions for these blockchain BFT protocols (i.e., the adversary cannot control the message delivery order on the underlying networks), we should still design practical large-scale BFT protocols that are robust in the traditional academic security model. For complete asynchronous networks, we present an Casper CBC-FBC style binary BFT protocol and a multi-value BFT protocol XP and prove their security in the traditional security model. For partial synchronous networks, we re-investigate the seminal DLS BFT protocol by Dwork, Lynch, and Stockmeyer [8]. In the DLS protocol, the protocol continues until every node decides on the same message. This is obviously inefficient. In a practical deployment, after one honest node decides on a message, it is more efficient for the decided participant to reliably broadcast his decision together with the proof. Each node that receives this decision can decide on the message instead of continuing the DLS protocol negotiation process. Since the traditional academic security model for BFT protocols assumes that even though the adversary controls the entire network, the message by an honest participant should be delivered to other honest participants eventually. Thus the revised DLS protocol does not break its security assumption and achieves liveness and safety in the same security model with both reduced round complexity and reduced communication complexity. We call this revised DLS protocol BDLS. When threshold digital signature schemes are used, BDLS achieves linear authenticator complexity with 4 rounds. It is noted that Facebook's HotStuff BFT/LibraBFT protocol achieves linear authenticator complexity with 7 rounds using threshold digital signature schemes.

The structure of the paper is as follows. Section 2 introduces system models and Byzantine agreement. Section 3 briefly discusses reliable broadcast communication channels. Section 4 shows that Ethereum blockchain 2.0's BFT protocol Casper FFG could not achieve liveness if the adversary

can re-order messages in the network. Section 5 shows that Ethereum blockchain's candidate BFT protocol Casper FBC for future deployment could not achieve liveness if the adversary can re-order messages in the network. Section 5 also proposes a Casper FBC style binary BFT protocol that achieves both safety and liveness in the traditional academic security model for complete asynchronous networks. Section 6 reviews the Polkadot's GRANDPA BFT protocol and shows that it cannot achieve liveness if the adversary is allowed to reschedule the message delivery order in the underlying networks. Section 7 presents the multi-value BDLS-BFT protocol design for partial synchronous networks and proves its security in the traditional academic security model. Section 8 discusses BLDS implementation details and presents the evaluation results. Section 9 proposes a multi-value BFT protocol XP for complete asynchronous networks and proves its security.

2 SYSTEM MODEL AND BYZANTINE AGREEMENT

For the Byzantine general problem, there are n participants and an adversary that is allowed to corrupt up to t of them. The adversary model is a static one wherein the adversary must decide whom to corrupt at the start of the protocol execution. For the network setting, we consider three kinds of networks: synchronous networks, partial synchronous networks by Dwork, Lynch, and Stockmeyer [8], and complete asynchronous networks by Fischer, Lynch, and Paterson [9].

(1) In a synchronous network, the time is divided into discrete units called slots $T_0, T_1, T_2, \cdots$ where the length of the time slots are equal. Furthermore, we assume that: (1) the current time slot is determined by a publicly-known and monotonically increasing function of current time; and (2) each participant has access to the current time. In a synchronous network, if an honest participant P_1 sends a message m to a participant P_2 at the start of time slot T_i, the message m is guaranteed to arrive at P_2 at the end of time slot T_i.

(2) In partial synchronous networks, the time is divided into discrete units as in synchronous networks. The adversary can selectively delay or re-order any messages sent by honest parties. In other words, if an honest participant P_1 sends a message m to an honest participant P_2 at the start of time slot T_{i_1}, P_2 will receive the message m eventually at time T_{i_2} where $i_2 = i_1 + \Delta$. Based on the property of Δ, we can distinguish the following two scenarios:
   - Type I partial synchronous network: $\Delta < \infty$ is unknown. That is, there exists a Δ but participants do not know the exact (or even approximate) value of Δ.
   - Type II partial synchronous network: $\Delta < \infty$ holds eventually. That is, the participant knows the value of Δ. But this Δ only holds after an unknown time slot $T = T_i$. Such a time T is called the Global Stabilization Time (GST).

   For Type I partial synchronous networks, the protocol designer supplies the consensus protocol first, then the adversary chooses her Δ. For Type II partial synchronous networks, the adversary picks the Δ and the protocol designer (knowing Δ) supplies the consensus protocol, then the adversary chooses the GST.

(3) In a complete asynchronous network, we make no assumptions about the relative speeds of processes or about the delay time in delivering a message. We also assume that processes do not have access to synchronized clocks. Thus algorithms based on time-outs cannot be used.

In all of the network models, we assume that the adversary has *complete control of the network*. That is, the adversary may schedule/reorder the delivery of messages as he wishes, and may insert messages as he wishes. The honest participants are completely passive: they simply follow the protocol steps and maintain their internal state between protocol steps.

The computations made by the honest participants and the adversary are modeled as polynomial-time computations. We assume that public key cryptography is used for message authentications. In particular, each participant should have authentic public keys of all other participants. This

https://mc.manuscriptcentral.com/dlt

means that if two participants P_i and P_j are honest and P_j receives a message from P_i over the network, then this message must have been generated by P_i at some prior point in time. A Byzantine agreement protocol must satisfy the following properties:

- **Safety**: If an honest participant decides on a value, then all other honest participants decides on the same value. That is, it is computationally infeasible for an adversary to make two honest participants to decide on different values.
- **Liveness (termination)**: There exists a function $B(\cdot)$ such that all honest participants should decide on a value after the protocol runs at most $B(n)$ steps. It should be noted that $B(n)$ could be exponential in n. In this case, we should further assume that 2^n is significantly smaller than 2^κ where κ is the security parameter for the underlying authentication scheme. In other words, one should not be able to break the underlying authentication scheme within $O(B(n))$ steps.
- **Non-triviality (Validity)**: If all honest participants start the protocol with the same initial value, then all honest participants that decide must decide on this value.

3 RELIABLE AND STRONGLY RELIABLE BROADCAST COMMUNICATION PRIMITIVES

For the BFT protocol design, it is important to understand what kind of communication channels are required. If a BFT protocol assume a reliable broadcast channel then the implementation must use a reliable broadcast primitive to achieve this channel since our Internet does not provide a robust broadcast channel. If the broadcast is not reliable, then Wang [21] showed that some other popular BFT protocols (such as Tendermint-BFT) cannot achieve liveness and the blockchain would go to deadlock. The similar attack also holds for Facebook's HotStuff BFT/LibraBFT (see Wang [21]).

The difference between point-to-point communication channels and broadcast communication channels has been extensively studied in the literature. A reliable broadcast channel requires that the following two properties be satisfied.

1. Correctness: If an honest participant broadcasts a message m, then every honest participant accepts m.
2. Unforgeability: If an honest participant does not broadcast a message m, then no honest participant accepts m.

By the above definition, a broadcast channel is unreliable if an honest participant broadcasts a message m to all participants and only a proper subset of honest participants receives this message m. That is, some honest participants receive the message m while other honest participants receive nothing at all (this could happen if the time is before GST in Type II networks). Thus we need to assume that the broadcast channel is unreliable before GST in Type II partial synchronous networks.

The above definition does not say anything about dishonest participants. In practice, a dishonest participant may send different messages to different participants or send the message only to a proper subset of honest participants even after GST in Type II networks. In order to defeat dishonest participants from carrying out these attacks, Bracha [3] designed a *strongly reliable broadcast primitive* (see Appendix for details) with the following additional requirement by assuming that all messages are delivered eventually in the network:

- If a dishonest participant P_i broadcasts a message, then either all honest participants accept the identical message or no honest participant accepts any value from P_i.

One should take precautions for using reliable or strongly reliable broadcast primitives since these primitives generally have assumptions about the underlying network topology. For example, Bracha's primitive [3] assumes that the underlying network is a complete network. Indeed, it is

https://mc.manuscriptcentral.com/dlt

sufficient to assume that there is a reliable point-to-point communication channel for each pair of participants in Bracha's primitive. For a given integer k, a network is called k-connected if there exist k-node disjoint paths between any two nodes within the network. In non-complete networks, it is well known that $(2t+1)$-connectivity is necessary for reliable communication against t Byzantine faults (see, e.g., Wang and Desmedt [20] and Desmedt-Wang-Burmester [6]). On the other hand, for broadcast communication channels, Wang and Desmedt [19] showed that there exists an efficient protocol to achieve probabilistically reliable and perfectly private communication against t Byzantine faults when the underlying communication network is $(t+1)$-connected. The crucial point to achieve these results is that: in a point-to-point channel, a malicious participant P_1 can send a message m_1 to participant P_2 and send a different message m_2 to participant P_3 though, in a broadcast channel, the malicious participant P_1 has to send the same message m to multiple participants including P_2 and P_3. If a malicious P_1 sends different messages to different participants in a reliable broadcast channel, it will be observed by its neighbors.

Though broadcast channels at physical layers are commonly used in local area networks, it is not trivial to design reliable broadcast channels over the Internet infrastructure since the Internet connectivity is not a complete graph and some direct communication paths between participants are missing (see, e.g., [12, 20]). In addition to Bracha's strongly reliable broadcast primitive [3], quite a few alternative broadcast primitives have been proposed in the literature using message relays (see, e.g., Srikanth and Toueg [17] and Dwork-Lynch-Stockmeyer [8]). In the message relay based broadcast protocol, if an honest participant accepts a message signed by another participant, it relays the signed message to other participants.

4 CASPER THE FRIENDLY FINALITY GADGET (FFG)

Buterin and Griffith [4] proposed the BFT protocol Casper the Friendly Finality Gadget (Casper FFG) as an overlay atop a block proposal mechanism. Casper FFG has been deployed in the Proof of Stake Based Ethereum 2.0. In Casper FFG, weighted participants validate and finalize blocks that are proposed by an existing proof of work chain or other mechanisms. To simplify our discussion, we assume that there are $n = 3t+1$ validators of equal weight. The Casper FFG works on the checkpoint tree that only contains blocks of height $100 * k$ in the underlying block tree. Each validator P_i can broadcast a signed vote $\langle P_i : s, t\rangle$ where s and t are two checkpoints and s is an ancestor of t on the checkpoint tree. For two checkpoints a and b, we say that $a \rightarrow b$ is a supermajority link if there are at least $2t+1$ votes for the pair. A checkpoint a is justified if there are supermajority links $a_0 \rightarrow a_1 \rightarrow \cdots \rightarrow a$ where a_0 is the root. A checkpoint a is finalized if there are supermajority links $a_0 \rightarrow a_1 \rightarrow \cdots \rightarrow a_i \rightarrow a$ where a_0 is the root and a is the direct son of a_i. In Casper FFG, an honest validator P_i should not publish two distinct votes

$$\langle P_i : s_1, t_1\rangle \quad \text{AND} \quad \langle P_i : s_2, t_2\rangle$$

such that either

$$h(t_1) = h(t_2) \quad \text{OR} \quad h(s_1) < h(s_2) < h(t_2) < h(t_1)$$

where $h(\cdot)$ denotes the height of the node on the checkpoint tree. In other words, *an honest validator should neither publish two distinct votes for the same target height nor publish a vote strictly within the span of its other votes.* Otherwise, the validator's deposit will be slashed. The authors [4] claimed that Casper FFG achieves accountable safety and plausible liveness where

(1) achieve accountable safety means that two conflicting checkpoints cannot both be finalized (assuming that there are at most t malicious validators), and
(2) plausible liveness means that supermajority links can always be added to produce new finalized checkpoints, provided there exist children extending the finalized chain.

https://mc.manuscriptcentral.com/dlt

In order to achieve the liveness property, [4] proposed to use the "correct by construction" fork choice rule: the underlying block proposal mechanism should "*follow the chain containing the justified checkpoint of the greatest height*".

The authors in [4] proposed to defeat the long-range revision attacks by a fork choice rule to never revert a finalized block, as well as an expectation that each client will "log on" and gain a complete up-to-date view of the chain at some regular frequency (e.g., once per month). In order to defeat the catastrophic crashes where more than t validators crash-fail at the same time (i.e., they are no longer connected to the network due to a network partition, computer failure, or the validators themselves are malicious), the authors in [4] proposed to slowly drains the deposit of any validator that does not vote for checkpoints, until eventually its deposit sizes decrease low enough that the validators who are voting are a supermajority. Related mechanism to recover from related scenarios such as network partition is considered an open problem in [4].

No specific network model is provided in [4]. In the implementation of the Casper FFG (see GO-Ethereum implementation), a participating node broadcasts his message as soon as he receives a sufficient number of messages to move forward. In other words, even if the network is a synchronized network, a participant may just make his decision on the first $2t + 1$ messages and ignore the remaining messages if these first $2t + 1$ messages are sufficient for him to move forward. This is reasonable and necessary since the remaining t nodes could be malicious ones and will never send any message at all. Based on this observation, we show that if the adversary could reschedule the message delivery order on the underlying networks, Casper FFG cannot achieve liveness property even in synchronized networks.

Fig. 1. Casper FFG cannot achieve liveness

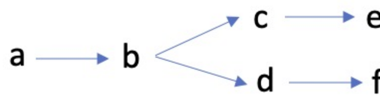

As an example, assume that, at time T, the checkpoint a is finalized where there is a supermajority link from a to its direct child b (that is, b is justified) and no vote for b's descendant checkpoint has been broadcast by any validator yet (see Figure 1). Now assume that the underlying block production mechanism produces a fork starting from b. That is, b has two descendant checkpoints c and d. The adversary who controls the network can arrange t honest validators to receive c first and $t + 1$ honest validators to receive d first where $h(c) = h(d)$. Thus t honest validators vote for $b \rightarrow c$, $t+1$ honest validators vote for $b \rightarrow d$, and t malicious validators vote randomly so that both $b \rightarrow c$ and $b \rightarrow d$ receives same number of votes. This means that c and d could not be finalized since neither the link $b \rightarrow c$ nor the link $b \rightarrow d$ could get a supermajority vote. It should be noted that by the two slashing rules in Casper FFG, an honest validator who voted for $b \rightarrow c$ is allowed to vote for $b \rightarrow f$ later since the two votes on $b \rightarrow c$ and $b \rightarrow f$ are not slashable. Next assume that the adversary schedules the message delivery order so that t honest validators receive e first and $t + 1$ honest validators receive f first (without loss of generality, we may assume that $h(e) = h(f)$). Thus t honest validators vote for $b \rightarrow e$, $t + 1$ honest validators vote for $b \rightarrow f$, and t malicious validators vote randomly so that both $b \rightarrow e$ and $b \rightarrow f$ receives same number of votes. Thus e and f could not be finalized since neither the link $b \rightarrow e$ nor the link $b \rightarrow f$ could get a supermajority vote. This process continues forever and no checkpoint after a could be finalized. That is, Casper FFG could not achieve liveness with this kind of message delivery schedule by the adversary.

5 CBC CASPER THE FRIENDLY BINARY CONSENSUS (FBC)

The network model for Casper FFG is not clearly defined. In order to make Ethereum blockchain robust in complete asynchronous networks, Ethereum foundation has been advocating the "Correct-by-Construction" (CBC) family of Casper blockchain consensus protocols [23, 24] for their future release of Ethereum blockchain. The CBC Casper the Friendly Ghost emphasizes the safety property. But it does not try to address the liveness requirement for the consensus process. Indeed, it explicitly says that [23] "*liveness considerations are considered largely out of scope, and should be treated in future work*". Thus in order for CBC Casper to be deployable, a lot of work needs to be done since the Byzantine Agreement Problem becomes challenging only when both safety and liveness properties are required to be satisfied at the same time. It is simple to design BFT protocols that only satisfy one of the two requirements (safety or liveness). The Ethereum foundation community has made several efforts to design safety oracles for CBC Casper to help participants to make a decision when an agreement is reached (see, e.g., [16]). However, this problem is at least as hard as coNP-complete problems. So no satisfactory solution has been proposed yet.

CBC Casper has received several critiques from the community. For example, Ali et al [1] concluded that "*the definitions and proofs provided in [24] result in neither a theoretically sound nor practically useful treatment of Byzantine fault-tolerance. We believe that considering correctness without liveness is a fundamentally wrong approach. Importantly, it remains unclear if the definition of the Casper protocol family provides any meaningful safety guarantees for blockchains*". Though CBC Casper is not a deployable solution yet and it has several fundamental issues yet to be addressed, we think these critiques as in [1] may not be fair enough. Indeed, CBC Casper provides an interesting framework for consensus protocol design. In particular, the algebraic approach proposed by CBC Casper has certain advantages for describing Byzantine Fault Tolerance (BFT) protocols. The analysis in this section shows that the current formulation of CBC Casper could not achieve liveness property. However, if one revises the CBC Casper's algebraic approach to include the concept of "waiting" and to enhance participant's capability to identify more malicious activities (that is, to consider general malicious activities in addition to equivocating activities), then one can design efficiently constructive liveness concepts for CBC Casper even in complete asynchronous networks.

5.1 Casper FBC protocol description

CBC Casper contains a binary version and an integer version. In this paper, we only consider Casper the Friendly Binary Consensus (FBC). Our discussion can be easily extended to general cases. For the Casper FBC protocol, each participant repeatedly sends and receives messages to/from other participants. Based on the received messages, a participant can infer whether a consensus has been achieved. Assume that there are n participants $P_1, \cdots, P_n$ and let $t < n$ be the Byzantine-fault-tolerance threshold. The protocol proceeds from step to step (starting from step 0) until a consensus is reached. Specifically the step s proceeds as follows:

- Let $\mathcal{M}_{i,s}$ be the collection of valid messages that P_i has received from all participants (including himself) from steps $0, \cdots, s-1$. P_i determines whether a consensus has been achieved. If a consensus has not been achieved yet, P_i sends the message

$$m_{i,s} = \langle P_i, e_{i,s}, \mathcal{M}_{i,s} \rangle \tag{1}$$

  to all participants where $e_{i,s}$ is P_i's estimated consensus value based on the received message set $\mathcal{M}_{i,s}$.

In the following, we describe how a participant P_i determines whether a consensus has been achieved and how a participant P_i calculates the value $e_{i,s}$ from $\mathcal{M}_{i,s}$.

For a message $m = \langle P_i, e_{i,s}, \mathcal{M}_{i,s}\rangle$, let $J(m) = \mathcal{M}_{i,s}$. For two messages m_1, m_2, we write $m_1 \prec m_2$ if m_2 depends on m_1. That is, there is a sequence of messages $m'_1, \cdots, m'_v$ such that

$$\begin{array}{c} m_1 \in J(m'_1) \\ m'_1 \in J(m'_2) \\ \cdots \\ m'_v \in J(m_2) \end{array}$$

For a message m and a message set $\mathcal{M} = \{m_1, \cdots, m_v\}$, we say that $m \prec \mathcal{M}$ if $m \in \mathcal{M}$ or $m \prec m_j$ for some $j = 1, \cdots, v$. The *latest message* $m = L(P_i, \mathcal{M})$ by a participant P_i in a message set $\mathcal{M}$ is a message $m \prec \mathcal{M}$ satisfying the following condition:

- There does not exist another message $m' \prec \mathcal{M}$ sent by participant P_i with $m \prec m'$.

It should be noted that the "latest message" concept is well defined for a participant P_i if P_i has not equivocated, where a participant P_i equivocates if P_i has sent two messages $m_1 \neq m_2$ with the properties that "$m_1 \not\prec m_2$ and $m_2 \not\prec m_1$".

For a binary value $b \in \{0, 1\}$ and a message set $\mathcal{M}$, the score of a binary estimate for b is defined as the number of non-equivocating participants P_i whose latest message voted for b. That is,

$$\texttt{score}(b, \mathcal{M}) = \sum_{L(P_i,\mathcal{M})=(P_i,b,*)} \lambda(P_i, \mathcal{M}) \tag{2}$$

where

$$\lambda(P_i, \mathcal{M}) = \begin{cases} 0 & \text{if } P_i \text{ equivocates in } \mathcal{M}, \\ 1 & \text{otherwise.} \end{cases}$$

To estimate consensus value: Now we are ready to define P_i's estimated consensus value $e_{i,s}$ based on the received message set $\mathcal{M}_{i,s}$ as follows:

$$e_{i,s} = \begin{cases} 0 & \text{if } \texttt{score}(0, \mathcal{M}_{i,s}) > \texttt{score}(1, \mathcal{M}_{i,s}) \\ 1 & \text{if } \texttt{score}(1, \mathcal{M}_{i,s}) > \texttt{score}(0, \mathcal{M}_{i,s}) \\ b & \text{otherwise, where } b \text{ is coin-flip output} \end{cases} \tag{3}$$

To infer consensus achievement: For a protocol execution, it is required that for all i, s, the number of equivocating participants in $\mathcal{M}_{i,s}$ is at most t. A participant P_i determines that a consensus has been achieved at step s with the received message set $\mathcal{M}_{i,s}$ if there exists $b \in \{0, 1\}$ such that

$$\forall s' > s : \texttt{score}(b, \mathcal{M}_{i,s'}) > \texttt{score}(1 - b, \mathcal{M}_{i,s'}). \tag{4}$$

5.2 Efforts to achieve liveness for CBC Casper FBC

From CBC Casper protocol description, it is clear that CBC Casper is guaranteed to be correct against equivocating participants. However, the "inference rule for consensus achievement" requires a mathematical proof that is based on infinitely many message sets $\mathcal{M}_{i,s'}$ for $s' > s$. This requires each participant to verify that for each potential set of t Byzantine participants, their malicious activities will not overturn the inequality in (4). This problem is at least co-NP hard. Thus even if the system reaches a consensus, the participants may not realize this fact. In order to address this challenge, Ethereum community provides three "safety oracles" (see [16]) to help participants to determine whether a consensus is obtained. The first "adversary oracle" simulates some protocol execution to see whether the current estimate will change under some Byzantine attacks. As mentioned previously, this kind of problem is co-NP hard and the simulation cannot be exhaustive generally. The second "clique oracle" searches for the biggest clique of participant graph to see whether there exist more than 50% participants who agree on current estimate and all acknowledge the agreement. That is, for each message, the oracle checks to see if, and for how long, participants

have seen each other agreeing on the value of that message. This kind of problem is equivalent to the complete bipartite graph problem which is NP-complete. The third "Turan oracle" uses Turan's Theorem to find the minimum size of a clique that must exist in the participant edge graph. In a summary, currently there is no satisfactory approach for CBC Casper participants to determine whether finality has achieved. Thus no liveness is guaranteed for CBC Casper. Indeed, we can show that it is impossible to achieve liveness in CBC Casper.

5.3 Impossibility of achieving liveness in CBC Casper

In this section, we use a simple example to show that without a protocol revision, no liveness could be achieved in CBC Casper. Assume that there are $3t+1$ participants. Among these participants, $t-1$ of them are malicious and never vote. Furthermore, assume that $t+1$ of them hold value 0 and $t+1$ of them hold value 1. Since the message delivery system is controlled by the adversary, the adversary can let the first $t+1$ participants to receive $t+1$ voted 0 and t voted 1. On the other hand, the adversary can let the next $t+1$ participants to receive $t+1$ voted 1 and t voted 0. That is, at the end of this step, we still have that $t+1$ of them hold value 0 and $t+1$ of them hold value 1. This process can continue forever and never stop.

In CBC Casper FBC [23, 24], a participant is identified as malicious only if he equivocates. This is not sufficient to guarantee liveness (or even safety) of the protocol. For example, if no participant equivocates and no participant follows the equation (3) for consensus value estimation, then the protocol may never make a decision (that is, the protocol cannot achieve liveness property). However, the protocol execution satisfies the valid protocol execution condition of [23, 24] since there is zero equivocating participant.

5.4 Revising CBC Casper FBC

CBC Casper does not have an in-protocol fault tolerance threshold and does not have any timing assumptions. Thus the protocol works well in complete asynchronous settings. Furthermore, it does not specify when a participant P_i should broadcast his step s protocol message to other participants. That is, it does not specify when P_i should stop waiting for more messages to be included $\mathcal{M}_{i,s}$. We believe that CBC Casper authors do not specify the time for a participant to send its step s protocol messages because they try to avoid any timing assumptions. In fact, there is a simple algebraic approach to specify this without timing assumptions. First, we revise the message set $\mathcal{M}_{i,s}$ as the collection of messages that P_i receives from all participants (including himself) during step $s-1$. That is, the message set $\mathcal{M}_{i,s}$ is a subset of E_s where E_s is defined recursively as follows:

$$
\begin{aligned}
&E_0 = \emptyset \\
&E_1 = \{\langle P_j, b, \emptyset\rangle : j = 1, \cdots, n; b = 0, 1\} \\
&E_2 = \{\langle P_j, b, \mathcal{M}_{j,1}\rangle : j = 1, \cdots, n; b = 0, 1; \mathcal{M}_{j,1} \subset E_1\} \\
&\cdots \\
&E_s = \{\langle P_j, b, \mathcal{M}_{j,s-1}\rangle : j = 1, \cdots, n; b = 0, 1; \mathcal{M}_{j,s-1} \subset E_{s-1}\} \\
&\cdots
\end{aligned}
$$

Then we need to revise the latest message definition $L(P_j, \mathcal{M}_{i,s})$ accordingly:

$$
L(P_j, \mathcal{M}_{i,s}) = \begin{cases} m & \text{if } \langle P_j, b, m\rangle \in \mathcal{M}_{i,s} \\ \emptyset & \text{otherwise} \end{cases} \tag{5}
$$

As we have mentioned in the preceding section, CBC Casper FBC [23, 24] only considers equivocating as malicious activities. This is not sufficient to guarantee protocol liveness against Byzantine faults. In our following revised CBC Casper model, we consider any participant that does not follow the protocol as malicious and exclude their messages:

https://mc.manuscriptcentral.com/dlt

- For a message set $\mathcal{M}_{i,s}$, let $I(\mathcal{M}_{i,s})$ be the set of identified malicious participants from $\mathcal{M}_{i,s}$. Specifically, let

$$I(\mathcal{M}_{i,s}) = E(\mathcal{M}_{i,s}) \cup F(\mathcal{M}_{i,s})$$

where $E(\mathcal{M}_{i,s})$ is the set of equivocating participants within $\mathcal{M}_{i,s}$ and $F(\mathcal{M}_{i,s})$ is the set of participants that does not follow the protocols within $\mathcal{M}_{i,s}$. For example, $F(\mathcal{M}_{i,s})$ includes participants that do not follow the consensus value estimation process properly or do not wait for enough messages before posting his own protocol messages.

With the definition of $I(\mathcal{M}_{i,s})$, we should also redefine the score function (2) by revising the definition of $\lambda(P_i, \mathcal{M})$ accordingly:

$$\lambda(P_i, \mathcal{M}) = \begin{cases} 0 & \text{if } P_i \in I(\mathcal{M}), \\ 1 & \text{otherwise.} \end{cases}$$

5.5 Secure BFT protocol in the revised CBC Casper

With the revised CBC Casper, we are ready to introduce the "waiting" concept and specify when a participant P_i should send his step s protocol message:

- A participant P_i should wait for at least $n - t + |I(\mathcal{M}_{i,s})|$ valid messages $m_{j,s-1}$ from other participants before he can broadcast his step s message $m_{i,s}$. That is, P_i should wait until $|\mathcal{M}_{i,s}| \geq n - t + |I(\mathcal{M}_{i,s})|$ to broadcast his step s protocol message.
- In case that a participant P_i receives $n - t + |I(\mathcal{M}_{i,s})|$ valid messages $m_{j,s-1}$ from other participants (that is, he is ready to send step s protocol message) before he could post his step $s - 1$ message, he should wait until he finishes sending his step $s - 1$ message.
- After a participant P_i posts his step s protocol message, it should discard all messages from steps $s - 1$ or early except decision messages that we will describe later.

It is clear that these specifications does not have any restriction on the timings. Thus the protocol works in complete asynchronous networks.

In Ben-Or's BFT protocol [2], if consensus is not achieved yet, the participants autonomously toss a coin until more than $\frac{n+t}{2}$ participant outcomes coincide. For Ben-Or's maximal Byzantine fault tolerance threshold $t \leq \lfloor \frac{n}{5} \rfloor$, it takes exponential steps of coin-flipping to converge. It is noted that, for $t = O(\sqrt{n})$, Ben-Or's protocol takes constant rounds to converge. Bracha [3] improved Ben-Or's protocol to defeat $t < \frac{n}{3}$ Byzantine faults. Bracha first designed a reliable broadcast protocol with the following properties (Bracha's reliable broadcast protocol is briefly reviewed in the Appendix): If an honest participant broadcasts a message, then all honest participants will receive the same message in the end. If a dishonest participants P_i broadcasts a message, then either all honest participants accept the identical message or no honest participant accepts any value from P_i. By using the reliable broadcast primitive and other validation primitives, Byzantine participants can be transformed to fail-stop participants. In the following, we assume that a reliable broadcast primitive such as the one by Bracha is used in our protocol execution and present Bracha's style BFT protocol in the CBC Casper framework. At the start of the protocol, each participant P_i holds an initial value in his variable $x_i \in \{0, 1\}$. The protocol proceeds from step to step. The step s consists of the following sub-steps.

(1) Each participant P_i reliably broadcasts $\langle P_i, x_i, \mathcal{M}_{i,s,0} \rangle$ to all participants where $\mathcal{M}_{i,s,0}$ is the message set that P_i has received during step $s - 1$. Then P_i waits until it receives $n - t$ valid messages in $\mathcal{M}_{i,s,1}$ and computes the estimate $e_{i,s}$ using the value estimation function (3).
(2) Each participant P_i reliably broadcasts $\langle P_i, e_{i,s}, \mathcal{M}_{i,s,1} \rangle$ to all participants and waits until it receives $n - t$ valid messages in $\mathcal{M}_{i,s,2}$. If there is a b such that $\texttt{score}(b, \mathcal{M}_{i,s,2}) > \frac{n}{2}$, then let $e'_{i,s} = b$ otherwise, let $e'_{i,s} = \perp$.

https://mc.manuscriptcentral.com/dlt

(3) Each participant P_i reliably broadcasts $\langle P_i, e'_{i,s}, \mathcal{M}_{i,s,2}\rangle$ to all participants and waits until it receives $n-t$ valid messages in $\mathcal{M}_{i,s,3}$. P_i distinguishes the following three cases:
- If $\texttt{score}(b, \mathcal{M}_{i,s,2}) > 2t+1$ for some $b \in \{0,1\}$, then P_i decides on b and broadcasts his decision together with justification to all participants.
- If $\texttt{score}(b, \mathcal{M}_{i,s,2}) > t+1$ for some $b \in \{0,1\}$, then P_i lets $x_i = b$ and moves to step $s+1$.
- Otherwise, P_i flips a coin and let x_i to be coin-flip outcome. P_i moves to step $s+1$.

Assume that $n = 3t+1$. The security of the above protocol can be proved be establishing a sequence of lemmas.

LEMMA 5.1. *If all honest participants hold the same initial value b at the start of the protocol, then every participant decides on b at the end of step $s = 0$.*

Proof. At sub-step 1, each honest participant receives at least $t+1$ value b among the $2t+1$ received values. Thus all honest participants broadcast b at sub-step 2. If a malicious participant P_j broadcasts $1-b$ during sub-step 2, then it cannot be justified since P_j could not receive $t+1$ messages for $1-b$ during sub-step 1. Thus P_j will be included in $I(\mathcal{M})$. That is, each honest participant receives $2t+1$ messages for b at the end of sub-step 2 and broadcasts b during sub-step 3. Based on the same argument, all honest participants decide on b at the end of sub-step 3. □

LEMMA 5.2. *If an honest participant P_i decides on a value b at the end of step s, then all honest participants either decide on b at the end of step s or at the end of step $s+1$.*

Proof. If an honest participant P_i decides on a value b at the end of sub-step 3, then P_i receives $2t+1$ valid messages for the value b. Since the underlying broadcast protocol is reliable, each honest participant receives at least $t+1$ these valid messages for the value b. Thus if a participant P_i does not decide on the value b at the end of sub-step 3, it would set $x_i = b$. That is, all honest participants will decide during step $s+1$. □

The above two Lemmas show that the protocol is a secure Byzantine Fault Tolerance protocol against $\lfloor \frac{n-1}{3} \rfloor$ Byzantine faults in complete asynchronous networks. The above BFT protocol may take exponentially many steps to converge. However, if a common coin such as the one in Rabin [15] is used, then the above protocol converges in constant steps. It should be noted that Ethereum 2.0 provides a random beacon which could be used as a common coin for the above BFT protocol. Thus the above BFT protocol could be implemented with constant steps on Ethereum 2.0.

6 POLKADOT'S BFT PROTOCOL GRANDPA

The project Polkadot (https://github.com/w3f) proposed an algebraic approach based BFT finality gadget protocol GRANDPA which is similar to Casper FBC in some sense. By May 2021, Polkadot has a global market cap of 28.821 billion US dollar and ranks No. 8 among the entire cryptography currency market (after: Bitcoin, Ethereum, BNB, USDT, XRP, ADA, DOGE). There are different versions of GRANDPA protocol. In this paper, we refer to the most recent one [18] dated on June 19, 2020. Specifically, Polkadot implements a nominated proof-of-stake (NPoS) system. At certain time period, the system elects a group of validators to serve for block production and the finality gadget. Nominators also stake their tokens as a guarantee of good behavior, and this stake gets slashed whenever their nominated validators deviate from their protocol. On the other hand, nominators also get paid when their nominated validators play by the rules. Elected validators get equal voting power in the consensus protocol. Polkadot uses BABE as its block production mechanism and GRANDPA as its BFT finality gadget. Here we are interested in the finality gadget GRANDPA (GHOST-based Recursive ANcestor Deriving Prefix Agreement) that is implemented for the Polkadot relay chain. GRANDPA contains two protocols. The first protocol works in partially synchronous

networks and tolerates 1/3 Byzantine participants. The second protocol works in full asynchronous networks (requiring a common random coin) and tolerates 1/5 Byzantine participants. The first GRANDPA protocol assumes that the underlying network is a Type I partial synchronous network. In the following paragraphs, we will show that GRANDPA cannot achieve liveness property in partial synchronous networks if the adversary is allowed to reschedule the message delivery order.

Assume that there are $n = 3t + 1$ participants $P_0, \cdots, P_{n-1}$ and at most t of them are malicious. Each participant stores a tree of blocks produced by the block production mechanism with the genesis block as the root. A participant can vote for a block on the tree by digitally signing it. For a set S of votes, a participant P_i equivocates in S if P_i has more than one vote in S. A set S of votes is called safe if the number of participants who equivocate in S is at most t. A vote set S has supermajority for a block B if

$$|\{P_i : P_i \text{ votes for } B*\} \cup \{P_i : P_i \text{ equivocates}\}| \geq 2t + 1$$

where P_i votes for $B*$ mean that P_i votes for B or a descendant of B.

In GRANDPA, the 2/3-GHOST function $g(S)$ returns the block B of the maximal height such that S has a supermajority for B or a “nil” if no such block exists. If a safe vote set S has a supermajority for a block B, then there are at least $t + 1$ voters who do vote for B or its descendant but do not equivocate. Based on this observation, it is easy to check that if $S \subseteq T$ and T is safe, then $g(S)$ is an ancestor of $g(T)$.

The authors in [18] defined the following concept of *possibility* for a vote set to have a supermajority for a block: “We say that it is *impossible* for a set S to have a supermajority for a block B if at least $2t + 1$ voters either equivocate or vote for blocks who are not descendant of B. Otherwise it is *possible* for S to have a supermajority for B”. Then they claimed (the second paragraph above Lemma 2.6 in [18]) that “a vote set S is possible to have a supermajority for a block B if and only if there exists a safe vote set $T \supseteq S$ such that T has a supermajority for B”. Unfortunately, this claim is not true in practice if the adversary selects a non-equivocating strategy which may introduce deadlock to the system (on the other hand, this claim is true if all t malicious voters MUST equivocate).

Example 6.1. Assume that the underlying block production mechanism BABE has produced blocks C and D that are inconsistent (that is, C is not an ancestor of D and D is not an ancestor of C). The adversary who controls the message delivery system may arrange the block C to reach the adversary-selected $t + 1$ voters first and let the block D to reach these $t + 1$ voters after these voters have voted for C. Similarly, the adversary can arrange the block D to reach the remaining $2t + 1$ voters first and let the block C to reach them after they have voted for D. In a summary, the vote set S will now contain the following votes:

(1) $t + 1$ voters vote for C.
(2) $2t$ voters vote for D.
(3) no voter equivocates.

By the fact that $2t$ votes in S do “either equivocate or vote for blocks who are not descendant of C” and by the following definition from GRANDPA

> “We say that it is *impossible* for a set S to have a supermajority for a block B if at least $2t + 1$ voters either equivocate or vote for blocks who are not descendant of B. Otherwise it is *possible* for S to have a supermajority for B”

The GRANDPA system concludes that S is *possible* to have a supermajority for the block C. Furthermore, the malicious voters (controlled by the adversary) decide not to equivocate at all during the protocol run. Since all $3t + 1$ votes have been cast already in S and no one equivocates in potential supersets of S, for all semantically valid safe vote set $T \supseteq S$, we have $S = T$. In other words, no semantically valid safe vote set $T \supseteq S$ could have a supermajority for C. Similarly, by the above

https://mc.manuscriptcentral.com/dlt

definition, S is *possible* to have a supermajority for a block D and no semantically valid safe vote set $T \supseteq S$ could have a supermajority for D. In a summary, in the GRANDPA protocol, "*possible*" does not really means it is possible.

In the following sections, we will use Example 6.1 to show that the GRANDPA protocol will enter deadlock and cannot achieve the liveness property if the adversary is allowed to reschedule the message delivery order.

6.1 GRANDPA protocol

The GRANDPA protocol starts from round 1. For each round, one participant is designated as the primary and all participants know who is the primary. Each round consists of two phases: *prevote* and *precommit*. Let $V_{r,i}$ and $C_{r,i}$ be the sets of prevotes and precommits received by P_i during round r respectively. Let $E_{0,i}$ be the genesis block and $E_{r,i}$ be the last ancestor block of $g(V_{r,i})$ that is possible for $C_{r,i}$ to have a supermajority. If either $E_{r,i} < g(V_{r,i})$ or it is impossible for $C_{r,i}$ to have a supermajority for any children of $g(V_{r,i})$, then we say that P_i sees that round r is *completable*. Let Δ be a time bound such that it suffices to send messages and gossip them to everyone. The protocol proceeds as follows.

(1) P_i starts round $r > 1$ if round $r-1$ is completable and P_i has cast votes in all previous rounds. Let $t_{r,i}$ be the time P_i starts round r.
(2) The primary voter P_i of round r broadcasts $E_{r-1,i}$.
(3) **prevote**: P_i waits until either it is at least time $t_{r,i} + 2\Delta$ or round r is completable. P_i *prevotes* for the head of the best chain containing $E_{r-1,i}$ unless P_i receives a block B from the primary with $g(V_{r-1}, i) \geq B > E_{r-1,i}$. In this case, P_i uses the best chain containing B.
(4) **precommit**: P_i waits until $g(V_{r,i}) \geq E_{r-1,i}$ and one of the following holds
  (a) it is at least time $t_{r,i} + 4\Delta$
  (b) round r is completable
  Then P_i broadcasts a precommit for $g(V_{r,i})$

At any time after the precommit step of round r, if P_i sees that $B = g(C_{r,i})$ is descendant of the last finalized block and $V_{r,i}$ has a supermajority, then P_i finalizes B.

6.2 GRANDPA cannot achieve liveness in partial synchronous networks

In this section, we show that GRANDPA BFT protocol cannot achieve liveness property in partial synchronous networks. Assume that $E_{r-1,0} = \cdots = E_{r-1,n-1} = A$ and all $3t+1$ voters prevote and precommit to A during round $r-1$ and A is finalized by all voters during round $r-1$. Also assume that no voter will ever equivocate. During round r, the block production mechanisms produces a fork of of A. That is, we get two children blocks C and D of A.
Counter-example 1: By adjusting the message delivery schedule (this could happen before GST in partial synchronous networks), $t+1$ voters only receive the block C before time $t_{r,i} + 2\Delta$ and $2t$ voters only receive the block D before time $t_{r,i} + 2\Delta$. However, all voters will receive both blocks C and D before time $t_{r,i} + 3\Delta$.

At step 2 of round r, the primary voter broadcasts $A = E_{r-1,i}$. At step 3, both $V_{r,i}$ and $C_{r,i}$ are empty initially, the round r cannot be completable until time $t_{r,i} + 2\Delta$. Thus voter P_i waits until time $t_{r,i} + 2\Delta$ to submit its prevote. The $t+1$ voters that received block C would prevote for C and the other $2t$ voters that received block D would prevote for D. The adversary allows all prevotes of Step 3 to be delivered to all voters synchronously before time $t_{r,i} + 4\Delta$. During Step 4, each voter P_i receives $t+1$ prevotes for C and $2t$ prevotes for D. Since $C_{r,i}$ is empty until it receives any precommit, round r is not completable until time $t_{r,i} + 4\Delta$. That is, each voter P_i waits until $t_{r,i} + 4\Delta$ to precommit $g(V_{r,i}) = A$. The adversary allows all voters to receive all precommit votes for A. Now

each voter P_i estimates $E_{r,i} = g(V_{r,i}) = A$. By the fact that $C_{r,i} = \{3t+1$ precommit votes for $A\}$, we have $g(C_{r,i}) = A$. Since A has already been finalized, P_i will not finalize any block during round r.

In order for the round r to be completable, we need "either $E_{r,i} < g(V_{r,i})$ or it is impossible for $C_{r,i}$ to have a supermajority for any children of $g(V_{r,i})$". However, we have $E_{r,i} = g(V_{r,i}) = A$ and $C_{r,i} = \{3t+1$ precommit votes for $A\}$. That is, by definition of "*possibility*", it is "**possible**" for $C_{r,i}$ to have a supermajority for both children C and D of $g(V_{r,i}) = A$. In other words, the round r is NOT "completable" and GRANDPA cannot start Step 1 of round $r+1$.

Fig. 2. Counter-example 2 for GRANDPA

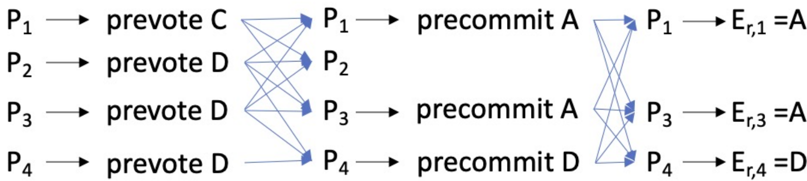

Counter-example 2: This example is more involved than counter-example 1 and an example with $t = 1$ is shown in Figure 2. By adjusting the message delivery schedule (this could happen before GST in partial synchronous networks), by time $t_{r,i} + 2\Delta$, we have t voters received block C and $2t+1$ voters received block D. Furthermore, all voters will receive both blocks C and D before time $t_{r,i} + 3\Delta$.

At step 2 of round r, the primary voter broadcasts $A = E_{r-1,i}$. At step 3, both $V_{r,i}$ and $C_{r,i}$ are empty initially, the round r cannot be completable until time $t_{r,i} + 2\Delta$. Thus voter P_i waits until time $t_{r,i} + 2\Delta$ to submit its prevote. The t voters that received block C would prevote for C and the other $2t+1$ voters that received block D would prevote for D. Durng Step 4, the adversary schedules the message delivery in such a way that, by time $t_{r,i} + 4\Delta$, t voters receive "$2t+1$ prevotes for D" and $2t+1$ voters receive "t prevotes for C and $t+1$ prevotes for D". Since $C_{r,i}$ is empty until it receives any precommit, round r is not completable until time $t_{r,i} + 4\Delta$. That is, each voter P_i waits until $t_{r,i} + 4\Delta$ to precommit $g(V_{r,i})$. At time $t_{r,i} + 4\Delta$, t voters precommit for $D = g(V_{r,i})$, $2t$ voters precommit for $A = g(V_{r,i})$, and one malicious voter does not precommit. The adversary let all precommit messages to be delivered to all voters synchronously.

Now t voters estimates $E_{r,i} = g(V_{r,i}) = D$ and $2t+1$ voters P_i estimates $E_{r,i} = g(V_{r,i}) = A$. By the fact that

$$C_{r,i} = \{t \text{ precommit votes for } D \text{ and } 2t \text{ precommit votes for } A\},$$

we have $g(C_{r,i}) = A$. Since A has already been finalized, P_i will not finalize any block during round r. In order for the round r to be completable, we need "either $E_{r,i} < g(V_{r,i})$ or it is impossible for $C_{r,i}$ to have a supermajority for any children of $g(V_{r,i})$". However, we have $E_{r,i} = g(V_{r,i})$ for all voters and, by Example 6.1, it is "**possible**" for $C_{r,i}$ to have a supermajority for all children of $g(V_{r,i})$. In other words, the round r is NOT "completable" and GRANDPA cannot start Step 1 of round $r+1$.

Paper [18, page 7] mentions that "*$C_{r,i}$ and $V_{r,i}$ may change with time and also that $E_{r-1,i}$, which is a function of $V_{r-1,i}$ and $C_{r-1,i}$, can also change with time if P_i sees more votes from the previous round*". However, this has no impact on our preceding examples since after an honest voter prevotes/precommits, the honest voter cannot change his prevote/prevommit votes anymore (otherwise, it will be counted as equivocation).

7 A SECURE BFT PROTOCOL IN PARTIAL SYNCHRONOUS NETWORKS

We have showed that BFT protocols for Ethereum and Polkadot blockchains cannot achieve liveness property in partial synchronous networks if the adversary is allowed to reschedule the message delivery order. Section 5.5 proposes a secure BFT protocol for complete asynchronous networks that requires a common random beacon for efficiency purpose. In this section, we propose a secure and efficient BFT protocol that achieves safety and liveness properties in partial synchronous networks that does not require a common random beacon. Though our protocol could be used in other scenarios such as State Machine Replication (SMR), we present the protocol as a finality gadget for blockchains. Assume that there is a separate block proposal mechanism that produces children blocks for finalized blocks by our BFT finality gadget. Let $B^0, \cdots, B^{h-1}$ be the blockchain where B^0 is the genesis block and B^{h-1} is the most recently finalized head block. The block proposal mechanism may produce several child blocks $B^h_0, B^h_1, \cdots, B^h_{n_0-1}$ of the current head block B^{h-1}. These child blocks are strictly ordered. For example, in proof of stake blockchain applications, each participant has a stake value for the chain height h and these child blocks may be ordered using proposer's stake values. However, it is beyond the scope of this paper to specify how these child blocks are ordered for general blockchains. It is the task for the BFT finality gadget to select the maximal block among these ordered candidate child blocks as the next block B^h. Though the goal of the BFT protocol is to select the maximal child block as the final version of block B^h, this may not be true in certain scenarios. For example, if $t+1$ honest participants have seen the child block $B^h_{n_0-2}$ and have not seen the maximal block $B^h_{n_0-1}$ at the start of the protocol (at the same time, we may assume that the other t honest participants have seen the maximal block $B^h_{n_0-1}$), then our BFT protocol BDLS will finalize $B^h_{n_0-2}$ instead of $B^h_{n_0-1}$ (assuming that the t malicious participants submit the block $B^h_{n_0-2}$ to the leader). Secondly, our BFT protocol leverages the fact that a candidate block is self-certified. That is, the validity of a candidate child block can be verified by using the information contained in the candidate block itself against the currently finalized blockchain.

7.1 The BFT protocol BDLS

Our BFT protocol is based on the original DLS protocol in Dwork, Lynch, and Stockmeyer [8] and we call it a Blockchain version of DLS (BDLS). The DLS protocol is inefficient since the protocol proceeds until all honest participants decides on a value. BDLS employs a reliable broadcast channel so it stops as soon as one honest participants decides on a message. This reduces the complexity by a factor of $O(n)$.

For each blockchain height h, BDLS protocol runs from round to round until it reaches an agreement for the height h. Then the protocol moves to the next blockchain height $h+1$. Let $P_0, \cdots, P_{n-1}$ be the $n = 3t+1$ participants of the protocol. Assume that there are n_0 valid candidate proposals $B^h_0 \prec B^h_1 \prec \cdots \prec B^h_{n_0-1}$ for the block B^h. During the protocol run, each participant P_i maintains a local variable $\text{BLOCK}_i \subseteq \{B^h_0, B^h_1, \cdots, B^h_{n_0-1}\}$ that contains the candidate blocks that it has learned so far. Participant P_i prefers the maximal block in BLOCK_i to be selected as the final block for B^h. The goal of the BDLS protocol is for participants $P_0, \cdots, P_{n-1}$ to reach a consensus on the finalized block B^h.

Generally, we can use a robust threshold signature scheme to achieve linear authenticator complexity. For simplicity, the following protocol description is based on a standard digital signature scheme. It could be easily revised to use a threshold signature scheme. Following Dwork, Lynch, and Stockmeyer [8], we assume that all messages after the unknown GST (Global Stabilization Time) is delivered in the same round with any order chosen by the adversary and messages before round GST could get lost. Furthermore, though all participants have a common numbering for the

round, they do not know when the round GST occurs. A candidate block B' is acceptable to P_i if P_i does not have a lock on any value except possibly B'. There is a public function $leader(h, r)$ that returns the round leader for a given round r of the height h. For each height h, the BDLS protocol proceeds from round to round (starting from round 0) until the participant decides on a value. The round r of the height h starts when at least $2t + 1$ participants submit a round-change message to the leader participant. The round r proceeds as follows where $P_i = leader(h, r)$ is the leader for round r:

(1) Each participant P_j (including P_i) sends the signed message $(\langle h, r\rangle_j, \langle h, r, B'_j\rangle_j)$ to the leader P_i where $B'_j \in \text{BLOCK}_j$ is the maximal acceptable candidate block for P_j. The message $\langle h, r\rangle_j$ is considered as a round-change message. After sending the round-change message, P_j will not accept messages for round $r' < r$ anymore except a "decide" message from a previous round.

(2) If P_i receives at least $2t + 1$ round-change messages (including himself), it enters round r. In these round-change messages, if there are at least $2t + 1$ signed messages from $2t + 1$ participants with the same candidate block $B' \neq NULL$, then P_i broadcasts the following signed message (6) to all participants

$$\langle \text{lock}, h, r, B', \text{proof}\rangle_i \quad (6)$$

where proof is a list of at least $2t + 1$ signed messages showing that B' is the candidate blocks for at least $2t + 1$ participants (the proof also shows that round-change request has been authorized by at least $2t+1$ participants). If P_i does not receive such a block B', then P_i adds all received candidate blocks to its local variable BLOCK_i and broadcasts $\langle \text{select}, h, r, B'', \text{proof}\rangle$ where B'' is the candidate block $B'' = \max\{B : B \in \text{BLOCK}_i\}$ and proof is a list of at least $2t+1$ round-change messages. It should be noted that in order to achieve linear communication complexity when a threshold signature scheme employed, the "proof" in the `lock`-message and `select`-message are different: In the `lock`-message, the "proof" contains an assembled digital signature on the message $\langle h, r, B'\rangle$ while, in the `select`-message, the "proof" contains an assembled digital signature on the message $\langle h, r\rangle$. See Remark 3 for details.

(3) If a participant P_j (including P_i) receives a valid $\langle \text{select}, h, r, B'', \text{proof}\rangle$ from P_i during Step 2, then it adds B'' to its BLOCK_j. If a participant P_j (including P_i) receives a valid message $\langle \text{lock}, h, r, B', \text{proof}\rangle_i$ from P_i in Step 2, then it does the following:
   (a) releases any potential lock on B' from previous round, but does not release locks on any other potential candidate blocks
   (b) locks the candidate block B' by recording the valid lock (6)
   (c) sends the following signed commit message to the leader P_i.

$$\langle \text{commit}, h, r, B'\rangle_j. \quad (7)$$

(4) If P_i receives at least $2t + 1$ commit messages (7), then P_i *decides* on the value B' and *strongly reliable broadcast* (e.g., using Bracha's strongly reliable broadcast primitive in Section 3) the following decide message to all participants

$$\langle \text{decide}, h, r, B', \text{proof}\rangle_i. \quad (8)$$

where proof is a list of at least $2t + 1$ commit messages (7).

(5) If a participant P_j (including P_i) receives a decide message (8) from Step 4 or from its neighbor, it decides on the block B' for B^h and moves to the next height $h + 1$ (that is, run the Step 1 of height $h + 1$ by sending the round-change message). Otherwise, it goes to the following lock-release step:

https://mc.manuscriptcentral.com/dlt

- (*lock-release*) If a participant P_j (including P_i) has some locked values, it broadcasts all of its locked values with proofs. A participant releases its lock on a value $\langle \text{lock}, h, r'', B'', \text{proof} \rangle_{i''}$ if it receives a lock $\langle \text{lock}, h, r', B', \text{proof} \rangle_{i'}$ with $r' \geq r''$ and $B' \neq B''$.
- Move to the next round $r + 1$ (i.e., run the Step 1 of height h with $r + 1$).

(6) *height synchronization*: At any time during the protocol, if P_j receives a finalized bock of height h (e.g., a decide message (8)), P_j decides for height h and moves to height $h + 1$.
(7) *round synchronization*: At any time during the protocol, if P_j receives a valid "lock" or "select" or "decide" message for a round $r' > r$, P_j moves to round r' and processes the "lock" or "select" or "decide" message.
(8) *timeout*: For each step, P_j should set an appropriate timeout counter. If P_j does not receive enough messages to move forward before timeout counter expires, it moves to the next step.

Remark 1: In the BDLS protocol, the lock-release step is a mesh network broadcast. In some applications, one may prefer a star network to reduce the total number of messages from n^2 to n (achieving linear communication complexity). One may achieve this kind of needs by replacing the "lock-release" step with the following additions to the protocol. At the Step 1 of round r, each participant P_j sends the message

$$\text{all-locked-values}, \langle h, r, B'_j \rangle_j$$

instead of only sending the message $\langle h, r, B'_j \rangle_j$ to P_i, where "all-locked-values" is the set of candidate blocks that P_j has locks on. During Step 2, if P_i cannot lock a candidate block during round r, then it broadcasts the candidate block $B'' = \max\{B : B \in \text{BLOCK}_i\}$ together with a best locked candidate block from all received locks (the "best lock" is defined according to the lock-release process in Step 5). It is straightforward to check that our security analysis in the next section remains unchanged for this protocol revision.
Remark 2: During Step 5, BDLS strongly reliable broadcast (see, e.g., Section 3) the decide message. Alternatively, one may use the regular broadcast primitive and each participants who receives a decide message keeps propagating/broadcasting the `decide` message to its neighbors regularly until it receives at least $2t$ broadcasts of the `decide` message for height h from other $2t$ participants.
Remark 3: To achieve linear communication/authenticator complexity with threshold digital signature schemes, participant P_j sends the signed message $(\langle h, r \rangle_j, \langle h, r, B'_j \rangle_j)$ to the leader P_i during step 1. It should be noted that if there are $2t + 1$ participants that send the same B'_j to the leader, then the leader P_i can assembly a signature for $\langle h, r, B'_j \rangle$. If there is no such value B'_j, then the leader can only assembly a digital signature for $\langle h, r \rangle$ which can be used for the `select` message. In the security proof for BDLS in the next section, the leader does not need to assemble a digital signature for B'_j if it only broadcasts a `select` message.

7.2 Liveness and Safety

The security of BDLS protocol is proved by establishing a series of Lemmas. The proofs for Lemmas 7.1, 7.2, 7.3 and Theorem 7.4 follow from straightforward modifications of the corresponding Lemmas/Theorem in [8]. For completeness, we include these proofs here also.

LEMMA 7.1. *It is impossible for two candidate blocks B' and B'' to get locked in the same round r of height h.*

Proof. In order for two blocks B' and B'' to get locked in one round r of height h, the leader $P_i = leader(h, r)$ must send two conflict lock messages (6) with different proofs. This can only happen if there exist at least $t + 1$ participants P_j each of whom equivocates two messages $\langle h, r, B' \rangle_j$ and $\langle h, r, B'' \rangle_j$ to P_i. This is impossible since there are at most t malicious participants. □

https://mc.manuscriptcentral.com/dlt

Lemma 7.2. *If the leader P_i decides a block value B' at round r of height h and r is the smallest round at which a decision is made. Then at least $t+1$ honest participants lock the candidate block B' at round r. Furthermore, each of the honest participants that locks B' at round r will always have a lock on B' for round $r' \geq r$.*

Proof. In order for P_i to decide on B', at least $2t+1$ participants send commit messages (7) to P_i at round r of height h. Thus at least $t+1$ honest participants have locks on B' at round r. Assume that the second conclusion is false. Let $r' > r$ be the first round that the lock on B' is released. In this case, the lock is released during the lock-release step of round r' if some participant has a lock on another block $B'' \neq B'$ with associated round r'' where $r' \geq r'' \geq r$. Lemma 7.1 shows that it is impossible for a participant to have a lock on B'' ar round r. Thus the participant acquired the lock on B'' in round r'' with $r' \geq r'' > r$. This implies that, at the step 1 of round r'', more than $2t+1$ participants send signed messages $\langle h, r'', B'' \rangle$ to the leader participant. That is, at least $2t+1$ participants have not locked B' at the step 1 of round r''. This contradicts the fact that at least $t+1$ participants have locked B' at the start of round r''. □

Lemma 7.3. *Immediately after any lock-release step at or after the round GST, the set of candidate blocks locked by honest participants contains at most one value.*

Proof. This follows from the lock-release step. □

Theorem 7.4. *(Safety) Assume that there are at most t malicious participants. It is impossible for two participants to decide on different block values.*

Proof. Suppose that an honest participant P_i decides on B at round r and this is the smallest round at which the decision is made. Lemma 7.2 implies that at least $t+1$ participants will lock B' in all future rounds. Consequently, no other block values other than B' will be acceptable to $2t+1$ participants. Thus no participants will decide on any other values than B'. □

Theorem 7.5. *(Liveness) Assume that there are at most t malicious participants and valid candidate child blocks for B^h are always produced by the block proposal mechanism before the start of first round for height h for all h. Then BDLS protocol will finalize blocks for each height h. That is, the BDLS protocol will not reach a deadlock.*

Proof. We consider two cases. For the first case, assume that no decision has been made by any honest participants and no honest participant locks a candidate block at round r where $r \geq$ GST is the first round after GST that the leader participant is honest. In this case, if P_i receives $2t+1$ signed messages for a candidate block B' in step 1 of round r, then all honest participants will decides on B' by the end of round r. Otherwise, P_i broadcasts the maximal candidate block B'' during step 2 of round r. Thus all honest participants will receive this maximum block and this candidate block becomes the maximum acceptable candidate block for all honest participants. Then, in round $r' > r$ where r' is the smallest round after r that the leader participant is honest, all honest participants decide on a maximal block.

For the second case, assume that no candidate block is locked at the start of round GST and some participants hold a lock on a candidate block B'. By Lemma 7.3, there is at most one value locked by honest participants at the end of round GST. Furthermore, at the end of round GST, all the honest participants either decide on B' or obtain a lock on B'. Thus if no decision is made during round GST, the decision will be made during round GST+1. □

7.3 Complexity/Efficiency analysis

In this section, we compare the performance of BDLS against commonly deployed BFT protocols such as PBFT, Tendermint BFT, and LibraBFT (HotStuff BFT). Three kinds of primitives are used

https://mc.manuscriptcentral.com/dlt

Table 1. Comparison of BFT protocols with honest leader after GST

Steps	PBFT	Tendermint BFT	LibraBFT	BDLS
1	📡	📡	📡	👂
2	👋	👋	👂	📡
3	👋	👋	📡	👂
4			👂	📡
5			📡	
6			👂	
7			📡	
message complexity	$2n^2 + n$	$2n^2 + n$	$7n$	$4n$
authenticator complexity	$O(n^2)$	$O(n)$	$O(n)$	$O(n)$

in these protocol design: (1) broadcast from the leader to all participants; (2) all participants send messages to the leader; and (3) all participants broadcast. We use the following symbols to denote these primitives.

- 📡: leader broadcasts
- 👂: all participants send messages to the leader
- 👋: all participants broadcast

For the comparison, we focus on the performance of these protocols after the network is synchronized (that is, after GST) and when the round has an honest leader. For all of these protocols, they will reach agreement within one run of the protocol assuming all participants have all the necessary input values at the start of the protocol and the leader is honest. Table 1 lists the steps of one run of these protocols. Furthermore, for BDLS, we use the approaches discussed in the Remarks after the BDLS protocol description to embed the lock-release step into Steps 1 and 2. For each 📡 or 👂 step, there is a total of n messages communicated in the network. For each 👋 step, there is a total of n^2 messages communicated in the network. The row "message complexity" of Table 1 lists the total number of messages communicated in the network for each run of the protocol. That is, in the ideal synchronized network, this is the total number of messages that are needed to achieve a consensus. These numbers show that BDLS has the smallest number of messages for a consensus in the synchronized network. Another way to compare the performance of BFT protocols is to compare the number of authenticator operations (signing and verifying) that are needed to achieve a consensus. Assume that all these schemes (except PBFT) use threshold digital signature schemes, then the row "authenticator complexity" of Table 1 lists the total number authenticator operations needed for each run of the protocol.

8 IMPLEMENTATION AND PERFORMANCE EVALUATION

8.1 Chained BDLS and other implementation related issues

In order to improve efficiency, several blockchain BFT protocols (e.g., Ethereum Casper FFG, and LibraBFT/HotStuff BFT) adopt the chaining paradigm where the BFT protocol phases for commitment are spread across rounds. That is, every phase is carried out in a round and contains a new proposal. The same techniques could be used to construct a chained BDLS. For chained BFT protocol implementation, the BFT protocol participants for various rounds/heights should be relatively static. If the BFT protocol participants change from rounds to rounds or from heights to heights, it is not realistic to implement chained BFT protocols. Thus chained BFT protocol

https://mc.manuscriptcentral.com/dlt

implementation is suitable for permissioned blockchains such as Libra blockchain while it is not suitable for permissionless blockchains where BFT protocol participants change frequently.

In most distributed BFT protocols, when the participants could not reach an agreement in one round, participants move to a new round by submitting round-change request. Thus BFT participants may be in different status and receive different messages. It is important to maximize the period of time when at least $2t + 1$ honest participants are in the same round. PBFT protocol achieves round synchronization by exponentially increasing the timeout length for each round. That is, if the round 0 of height h has a timeout length of Δ, then the round r of height h will have a timeout length of $2^r\Delta$. On the other hand, Tendermint BFT achieves round synchronization by linearly increasing the timeout length for each round. That is, the round r has a timeout length of $r\Delta$ where Δ is the timeout length for round 0 of height h. HotStuff proposes a functionality called PaceMaker to achieve round synchronization without details on how to implement the PaceMaker. LibraBFT implemented the PaceMaker functionality as follows: When a participant gives up on a certain round r, it broadcasts a timeout message carrying a certificate for entering the round. This brings all honest participants to r within the transmission delay bound. When timeout messages are collected from a quorum of participants, they form a timeout certificate.

8.2 BDLS with Pacemaker

Though BDLS may use the PBFT mechanism to keep round synchronization (that is, the timeout period for round r is $2^r\Delta$), it seems to be more efficient to use Pacemaker for BDLS round synchronization. Similar to LibraBFT, the advancement of rounds in BDLS is governed by a module called Pacemaker. The Pacemaker keeps track of votes and of time. We revise BDLS slightly so that a Pacemaker could be seamlessly integrated into the protocol without extra workload. The major change is Step 1 where Pacemaker timeout messages are combined with round-change requests for efficiency. The round r of the height h for a participant P_j starts when (1) its Pacemaker receives round-change messages from at least $2t+1$ participants, or (2) if its timeout for round $r-1$, or (3) if it receives a valid `Resync` message for round r while it is still in round less than r, or (4) if it receives a "lock" or a "select" or a "decide" message for round r. At the beginning of the protocol, each participant sets `Resync(0) = NULL`. In the protocol, each time when a timeout counter expires, the participant will broadcast its current Resync variable to all participants and we will not explicitly mention this. The round r proceeds as follows where $P_i = leader(h, r)$ is the leader for round r:

(1) (If $r > 0$, this is done at the end of round $r-1$ of height h. If $r = 0$, this is done after a decision for height $h-1$ is made) The Pacemaker of each participant P_j (including P_i) broadcasts the signed message $(\langle h, r\rangle_j, \langle h, r, B'_j\rangle_j)$ and `Resync` variable where $B'_j \in \text{BLOCK}_j$ is the maximal acceptable candidate block for P_j of height h. The message $\langle h, r\rangle_j$ is considered as a round-change message for round r. After P_j broadcasts the round-change message for round r, it will set a timeout message Δ_0 and enters `roundchanging` status. During `roundchanging` status, a participant will not accept any messages except the following messages: (1) round-change messages of rounds $\geq r$. (2) "decide" messages for the height h of any round, and (3) valid `Resync` message. Furthermore, if $r > 0$, then each participant P_j (including P_i) initializes all of its variables except the locked block variable. If $r = 0$, then each participant P_j (including P_i) initializes all of its variables including the locked block variable. For any participant P_j who is in `roundchanging` status, if it does not enter the `lock` status of Step 2 before Δ_0 expires, it resends the round-change message, its `Resync` value, and resets its Δ_0.

(2) During any time of the protocol, if P_j (including P_i) receives a Resync(r) or at least $2t+1$ round-change messages (including round-change message from himself) for round r (which is larger than its current round status), it enters `lock` status of round r and sets Resync=Resync(r)

https://mc.manuscriptcentral.com/dlt

where Resync(r) is a set of $2t+1$ valid round-change messages for round r. If P_j has not broadcast the round-change message yet, it broadcasts now. Then P_j sets the timeout counter Δ_1 for `lock` status[1]. Furthermore, as soon as the leader P_i enters `lock` status of round r, it starts a timeout counter $\Delta'_1 < \Delta_1$ concurrently[2]. The leader P_i stops the time counter Δ'_1 as soon as he receives n round-change requests or as soon as he receives $2t+1$ round-change requests with an identical proposed block. As soon as the time counter Δ'_1 expires or the leader P_i stops the time counter Δ'_1, P_i distinguishes the following two cases:

(a) Among all round-change messages that P_i has received, if there are at least $2t+1$ signed messages from $2t+1$ participants with the same candidate block $B' \neq NULL$, then P_i broadcasts the following signed message (6) to all participants

$$\langle \text{lock}, h, r, B', \text{proof} \rangle_i \tag{9}$$

where proof shows that at least $2t+1$ participants singed B' (the proof also shows that round-change request has been authorized by at least $2t+1$ participants).

(b) If P_i does not receive such a block B', then P_i adds all received candidate blocks to its local variable BLOCK_i and broadcasts

$$\langle \text{select}, h, r, B'', \text{proof} \rangle \tag{10}$$

where B'' is the candidate block $B'' = \max\{B : B \in \text{BLOCK}_i\}$ and proof shows that round-change request has been authorized by at least $2t+1$ participants from Step 1.

(3) If a participant P_j (including P_i) does not receive a valid message from the leader P_i during Step 2 and the timeout counter Δ_1 expires, P_j enters `commit` status of round r and sets the timeout counter Δ_2 for `commit` status[3]. Otherwise, if a participant P_j (including P_i) receives a valid message (9) or (10) from P_i before Δ_1 expires, P_j stops the time counter Δ_1 and distinguishes the following two cases:

- If P_j receives a valid $\langle \text{select}, h, r, B'', \text{proof} \rangle$ from P_i during Step 2, then it adds B'' to its BLOCK_j and enters `lock-release` status of round r and sets the timeout counter Δ_3 for `lock-release` status.
- If P_j (including P_i) receives a valid message $\langle \text{lock}, h, r, B', \text{proof} \rangle_i$ from P_i in Step 2, then it does the following and enters `commit` status by setting the timeout counter Δ_2:

(a) releases any potential lock on B' from previous round, but does not release locks on any other potential candidate blocks

(b) locks the candidate block B' by recording the valid lock (9)

(c) sends the following signed commit message to the leader P_i.

$$\langle \text{commit}, h, r, B' \rangle_j. \tag{11}$$

[1]The `lock` status timeout counters could be set as follows: For round $r=0$, the timeout counter $\Delta_1 = \Delta_{1,0}$ should be at least 4 network transmission delays plus some time for each participant to process the messages. For round $r>0$, the timeout counter could be defined as $r\Delta_{1,0}$.

[2]Though it is sufficient for a non-leader participant to collect only $2t+1$ round-change requests, the leader should collect as many round-change messages as possible. In particular, the leader should try to collect all round-change messages from all participants. It is recommended that after the leader P_i collects $2t+1$ round-change requests and starts the `lock` status timeout counter Δ_1, it initiates another timeout counter $\Delta'_1 < \Delta$ to collect as many as possible round-change requests if more round-change requests still arrive. Generally, we can set Δ_1 as two network transmission delays. This mechanism is used to avoid the following attack: the malicious t participants may send random round-change messages to the leader. If the leader only checks the first $2t+1$ messages (among them, t could be malicious), then the system may never reach an agreement. However, the leader should not wait forever since the t malicious participants may choose not to send round-change request at all.

[3]The `commit` status timeout counters could be set as follows: For round $r=0$, the timeout counter $\Delta_2 = \Delta_{2,0}$ should be at least 2 network transmission delays plus some time for each participant to process the messages. For round $r>0$, the timeout counter could be defined as $r\Delta_{2,0}$.

Table 2. Deciding time for various number of participants

No. nodes	20	30	50	80	100
deciding time	1.48s	1.54s	1.87s	3.56s	4.7s

(4) If P_i receives at least $2t+1$ commit messages (11) for the round r of height h with the locked value B' of (9) before Δ_2 expires, then P_i *decides* on the value B' and *strongly reliable broadcasts* (e.g., using Bracha's strongly reliable broadcast primitive) the following decide message to all participants

$$\langle \text{decide}, h, r, B', \text{proof}\rangle_i. \tag{12}$$

where proof is a list of at least $2t+1$ commit messages (11).

(5) If a participant P_j (including P_i) receives a decide message (12) from Step 4 or from its neighbor before the timeout counter Δ_2 expires, it decides on the block B' for B^h and the Pacemaker of P_j goes to Step 1 of height $h+1$. At the same time, the participant P_j propagates (broadcasts) the decide message (12) to all of its neighbors if it has not done so yet. Otherwise, if P_j (including P_i) does not receive a `decide` message from the leader P_i or its neighbors before the timeout counter Δ_2 expires, P_j enters `lock-release` status of round r and sets the timeout counter Δ_3 for `lock-release` status[4].

(6) (*lock-release*) If a participant P_j (including P_i) has some locked values, then P_j calculates

$$r_1 = \max\{r' : P_j \text{ holds a lock } \langle \text{lock}, h, r', B', \text{proof}\rangle_{i'}\}.$$

P_j releases all locks $\langle \text{lock}, h, r'', B'', \text{proof}\rangle_{i''}$ with $r'' \neq r_1$. P_j then broadcasts the following lock-release message

$$\langle \text{lock-release}, h, r, \langle \text{lock}, h, r_1, B', \text{proof}\rangle_{i_1}\rangle. \tag{13}$$

If P_j receives a lock-release message $\langle \text{lock-release}, h, r, \langle \text{lock}, h, r_1', B''', \text{proof}\rangle_{i_1'}\rangle$ with $r_1' > r_1$ from another participant before the timeout Δ_3 expires, then P_j releases its lock $\langle \text{lock}, h, r_1, B', \text{proof}\rangle_{i_1}$ and records the lock $\langle \text{lock}, h, r_1', B''', \text{proof}\rangle_{i_1'}$. After the timeout Δ_3 expires, the Pacemaker of P_j goes to Step 1 for round $r+1$ of height h.

(7) *height synchronization*: At any time of the protocol run, if P_j receives a finalized bock of height h (e.g., a decide message (12)), P_j decides for height h and moves to height $h+1$.

(8) *round synchronization*: At any time of the protocol run, if P_j receives a valid "lock" or "select" or "decide" or "`Resync`" message for a round $r' > r$, P_j moves to round r' and processes the "lock" or "select" or "decide" or "`Resync`" message. Furthermore, at any time, if P_j receives from more than $t+1$ participants valid messages for round $r' > r$ (including round-change messages for round r'), P_j goes to Step 1 for round r' of height h.

8.3 Performance evaluation

The BDLS consensus algorithm with pacemaker in Section 8.2 has been implemented using Go Programming Language. The implementation is based on the flowchart in Figure 3 and the source codes are available at https://github.com/yonggewang/bdls. We tested the BDLS on AWS EC2 with globally distributed nodes. Table 2 shows the average deciding (reaching agreement) time for 20, 30, 50, 80, and 100 nodes respectively.

[4]The `lock-release` status timeout counters could be set as follows: For round $r=0$, the timeout counter $\Delta_3 = \Delta_{3,0}$ should be at least 2 network transmission delays plus some time for each participant to process the messages. For round $r>0$, the timeout counter could be defined as $r\Delta_{3,0}$.

9 MULTI-VALUE BFT PROTOCOLS FOR ASYNCHRONOUS NETWORKS

Section 5.5 proposed a binary BFT finality gadget in complete asynchronous networks and Section 7 proposed a multi-value BFT finality gadget for partial synchronous networks. Furthermore, the BFT protocol in Section 5.5 requires a strongly reliable broadcast channel. In this section, we present a constant round multi-value BFT protocol XP for complete asynchronous networks that does not require strongly reliable broadcast channels. The XP protocol is motivated by the probabilistic binary BFT protocol in Cachin, Kursawe, and Shoup [5]. For a probabilistic BFT protocol, the randomness could be local coin flips or commonly shared random bits. Similar to the protocol in [5], our BFT XP protocol uses a commonly shared random sequence. This commonly shared random sequence is generally modelled as a random beacon in blockchains. For example, the Ethereum 2.0 provides a random beacon protocol.

Similar to Section 7, we assume that there is a partial order on the list of candidate blocks to be finalized: $\mathcal{B} = \{B_j : 1 \le j \le \tau\}$ where $B_1 \prec B_2 \prec \cdots \prec B_\tau$. During the protocol run, each participant P_i maintains a list of known candidate blocks in its local variable $X_i \subseteq \mathcal{B}$. At the start of the protocol run, X_i contains the list of candidate blocks that the participant P_i has learned and could be empty. During the protocol run, we assume that there is a random coin shared by all participants. For example, for the Ethereum 2.0, one may use the existing random beacon protocol as a common coin. Let σ be the random string shared by all participants for step s. Then participant P_i sets the "common" block X_i^σ as a block $B_j \in X_i$ such that $H(B_j, s)$ and $H(\sigma, s)$ has the maximal common prefix within X_i, where $H(\cdot)$ is a hash function. If there are two candidate blocks $B_{j_1} \prec B_{j_2}$ such that

$$\text{commonPrefix}(H(B_{j_1}, s), H(\sigma, s)) = \text{commonPrefix}(H(B_{j_2}, s), H(\sigma, s)),$$

then P_i sets $X_i^\sigma = B_{j_2}$. It is easy to observe that if $X_{i_1} = X_{i_2}$, then $X_{i_1}^\sigma = X_{i_2}^\sigma$. However, if $X_{i_1} \ne X_{i_2}$, then $X_{i_1}^\sigma$ and $X_{i_2}^\sigma$ may be different.

The protocol proceeds from step to step until an agreement is achieved and the protocol does not have any assumption on the time setting. Each participant waits for at least $n-t$ justified messages from participants (including himself) to proceed to the next sub-step. The step $s \ge 0$ for a participant P_i consists of the following sub-steps:

- **lock**: If $s = 0$, then let B be the maximal element in X_i. If $s > 0$ then wait for $n-t$ justified commit-votes from step $s-1$ and let

$$B = \begin{cases} B' & P_i \text{ receives a commit-vote for } B' \text{ in step } s-1 \\ X_i^\sigma & P_i \text{ receives } 2t+1 \text{ commit-votes for } \perp \text{ and } \sigma \text{ is common coin} \end{cases} \tag{14}$$

  Then P_i sends the following message to all participants.

$$\langle P_i, \texttt{lock}, s, B, \texttt{justification} \rangle \tag{15}$$

  where justification consists of messages to justify the selection of the value B.
- **commit**: P_i collects $n-t$ justified lock messages (15) and lets

$$\bar{B} = \begin{cases} B & \text{if there are } n-t \text{ locks for } B \\ \perp & \text{otherwise} \end{cases} \tag{16}$$

  Then P_i sends the following message to all participants

$$\langle P_i, \texttt{commit}, s, \bar{B}, X_i, \texttt{justification} \rangle \tag{17}$$

  where justification consists of messages to justify the selection of the value $\bar{B}$.
- **check-for-decision**: Collect $n-t$ properly justified commit votes (17) and lets $X_i = X_i \cup (\cup_j X_j)$ where X_j are from messages (17). Furthermore, if these are $n-t$ commit-votes for a

block $\bar{B}$, then P_i decides the block $\bar{B}$ and continues for one more step (up to `commit` sub-step). Otherwise, simply proceed.

Assume that $n = 3t + 1$. The security of the above protocol can be proved by establishing a sequence of lemmas.

LEMMA 9.1. *If an honest participant P_i decides on the value $\bar{B}$ at the end of step s (but no honest participant has ever decided before step s), then all honest participants either decide on $\bar{B}$ at the end of step s or at the end of step $s + 1$.*

Proof. If an honest participant P_i decides on the value $\bar{B}$ at the end of step s, then at least $t + 1$ honest participants commit-vote for $\bar{B}$. Thus each participant (including malicious participant) receives at least one commit-vote for $\bar{B}$ at the end of step s. This means that a malicious participant cannot create a justification that she has received a commit-vote for another block $B \neq \bar{B}$ or has received $2t + 1$ commit-votes for $\perp$ during step s. In other words, if a participant broadcasts a lock message for a block $B \neq \bar{B}$ during step $s + 1$, it cannot be justified and will be discarded by honest participants. This means that, all honest participants will commit-vote for the block $\bar{B}$ during step $s + 1$ and any commit-vote for other blocks cannot be justified. Thus, all honest participants will collect $n - t$ justified commit-vote for the block $\bar{B}$ and decide on block $\bar{B}$ at the end of step $s + 1$. □

LEMMA 9.2. *Block B in equation (14) is uniquely defined for each honest participant.*

Proof. It is sufficient to show that each participant P_i (including both honest and dishonest participants) can not receive commit-votes for two different blocks $\bar{B}_1$ and $\bar{B}_2$ during step s. For a contradiction, assume that P_i receives commit-vote for both $\bar{B}_1$ and $\bar{B}_2$ during step s. Then there are $2t + 1$ participants who submit `lock` messages for $\bar{B}_1$ and $2t + 1$ participants who submit `lock` messages for $\bar{B}_2$. This means that at least $t + 1$ participants (thus at least one honest participant) submit `lock` messages for both $\bar{B}_1$ and $\bar{B}_2$ which is impossible. □

LEMMA 9.3. *During step s, if participants P_i and P_j receive commit votes for $\bar{B}_1$ and $\bar{B}_2$ respectively, then $\bar{B}_1 = \bar{B}_2$.*

Proof. For a contradiction, assume that $\bar{B}_1 \neq \bar{B}_2$. Then there are $2t + 1$ lock messages for $\bar{B}_1$ and $2t + 1$ lock messages for $\bar{B}_2$ during step s. This means that at least $t + 1$ participants (thus at least one honest participant) submit `lock` messages for both $\bar{B}_1$ and $\bar{B}_2$ which is impossible. □

LEMMA 9.4. *If all honest participants hold the the same local value $X_i = \mathcal{B}$ at the start of step s, then with high probability, every participant decides by the end of step $s + \tau$.*

Proof. The Lemma is proved by distinguishing the following two cases:

(1) $s = 0$: At step 0, each honest participant broadcasts the lock for B_τ though dishonest participant may broadcast a lock for another block. At the commit phase, each honest participant P_i broadcasts $\mathcal{B}$ and a commit message for $\perp$ or B_τ depending on what he receives. If some participant decides at the end of Step 0, by Lemma 9.1, all honest participants decide by the end of Step 1. Assume that no participant decides by the end of Step 0. During Step 1, when a participant broadcasts a lock for a block, he needs to include $2t + 1$ commit messages from Step 0 as the justification. Among these $2t + 1$ commit messages, at least $t + 1$ come from honest participants which contain $\mathcal{B}$. Thus from now on, each participant must include its local variable $X_i = \mathcal{B}$ in its justification message. In other words, if a participant broadcasts a lock for a block based on the common coin, this locked block must be identical for all participants who use the common coin. Therefore, from Step 1 on, a participant can only broadcast a lock for a block committed in the immediate previous step (cf. Lemma 9.3) or a lock for a block determined by the common coin. With probability $\frac{1}{\tau}$, the block determined

by the common coin is identical to the committed block from the previous step. Thus all honest participants are expected to decide by Step τ.

(2) $s > 0$: for this case, we distinguish the following three cases:
  (a) By the end of step $s-1$, at least one participant (including dishonest participant) can legally decide on a block (this means at least one honest participant receives a commit-vote for a block $B \neq \perp$ during step $s-1$): By Lemma 9.1, all honest participants decides by the end of step s.
  (b) By the end of step $s-1$, no participant (including dishonest participant) can legally decide on a block: From Step s and on, each honest participant broadcasts a lock message for the unique block X_i^σ determined by the common coin or a unique block that was committed in the immediate previous Step (cf. Lemma 9.3). With probability $\frac{1}{\tau}$, the block determined by the common coin is identical to the committed block from the immediate previous Step. Thus all honest participants are expected to decide by Step $s+\tau$.

This completes the proof of the Lemma. □

Lemma 9.5. *All honest participant decides in constant steps.*

Proof. If no participant decides by the end of Step $s+\tau$, then, by Lemma 9.4, with high probability, at least one honest participant P_i revises its local variable X_i to include at least one more element during the Steps from s to $s+\tau$. Since there are at most τ candidate blocks, this process continues until no honest participant revises its local variable X_i. Then, by Lemma 9.4, all honest participants hold the same candidate block and the consensus will be reached. □

The above five Lemmas show that the protocol XP is a secure Byzantine Fault Tolerance protocol against $\lfloor \frac{n-1}{3} \rfloor$ Byzantine faults in complete asynchronous networks.

REFERENCES

[1] M. Ali, J. Nelson, and A. Blankstein. Peer review: CBC Casper. available at: https://medium.com/@muneeb/peer-review-cbc-casper-30840a98c89a, December 6, 2018.
[2] M. Ben-Or. Another advantage of free choice: Completely asynchronous agreement protocols (extended abstract). In *Proc. 2nd ACM PODC*, pages 27–30, 1983.
[3] G. Bracha. An asynchronous $[(n-1)/3]$-resilient consensus protocol. In *Proc. 3rd ACM PODC*, pages 154–162. ACM, 1984.
[4] V. Buterin and V. Griffith. Casper the friendly finality gadget. *arXiv preprint arXiv:1710.09437v4*, 2019.
[5] C. Cachin, K. Kursawe, and V. Shoup. Random oracles in constantinople: Practical asynchronous byzantine agreement using cryptography. *Journal of Cryptology*, 18(3):219–246, 2005.
[6] Yvo Desmedt, Yongge Wang, and Mike Burmester. A complete characterization of tolerable adversary structures for secure point-to-point transmissions without feedback. In *International Symposium on Algorithms and Computation*, pages 277–287. Springer, 2005.
[7] D. Dolev and H.R. Strong. Polynomial algorithms for multiple processor agreement. In *Proc. 14th ACM STOC*, pages 401–407. ACM, 1982.
[8] C. Dwork, N. Lynch, and L. Stockmeyer. Consensus in the presence of partial synchrony. *JACM*, 35(2):288–323, 1988.
[9] M.J. Fischer, N. A Lynch, and M.S. Paterson. Impossibility of distributed consensus with one faulty process. *JACM*, 32(2):374–382, 1985.
[10] J. Katz and C.-Y. Koo. On expected constant-round protocols for byzantine agreement. *Journal of Computer and System Sciences*, 75(2):91–112, 2009.
[11] L. Lamport. The part-time parliament. *ACM Transactions on Computer Systems (TOCS)*, 16(2):133–169, 1998.
[12] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 4(3):382–401, 1982.
[13] D. Ongaro and J. Ousterhout. In search of an understandable consensus algorithm. In *2014 USENIX Annual Technical Conference*, pages 305–319, 2014.
[14] M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *JACM*, 27(2):228–234, 1980.
[15] M.O. Rabin. Randomized byzantine generals. In *24th IEEE FOCS*, pages 403–409. IEEE, 1983.

[16] Ethereum Research. CBC Casper FAQ. available at: https://github.com/ethereum/cbc-casper/wiki/FAQ, November 27, 2018.
[17] TK Srikanth and S. Toueg. Simulating authenticated broadcasts to derive simple fault-tolerant algorithms. *Distributed Computing*, 2(2):80–94, 1987.
[18] A. Stewart and E. Kokoris-Kogia. GRANDPA: a byzantine finality gadge https://github.com/w3f/consensus/blob/master/pdf/grandpa.pdf, June 19, 2020.
[19] Y. Wang and Y. Desmedt. Secure communication in multicast channels: the answer to Franklin and Wright's question. *Journal of Cryptology*, 14(2):121–135, 2001.
[20] Y. Wang and Y. Desmedt. Perfectly secure message transmission revisited. *Information Theory, IEEE Tran.*, 54(6):2582–2595, 2008.
[21] Yongge Wang. Byzantine fault tolerance in partially connected asynchronous networks. *http://eprint.iacr.org/2019/1460*, 2019.
[22] Yongge Wang. The adversary capabilities in practical byzantine fault tolerance. In *Proc. 17th International Workshop on Security and Trust Management, STM 2021, LNCS 13075*, pages 1–20, 2021.
[23] V. Zamfir. Casper the friendly ghost: A correct by construction blockchain consensus protocol. *Whitepaper: https://github.com/ethereum/research/tree/master/papers*, 2017.
[24] V. Zamfir, N. Rush, A. Asgaonkar, and G. Piliouras. Introducing the minimal cbc casper family of consensus protocols. *DRAFT v1.0: https://github.com/cbc-casper/*, 2018.

A BRACHA'S STRONGLY RELIABLE BROADCAST PRIMITIVE

Assume $n > 3t$. Bracha [3] designed a broadcast protocol for asynchronous networks with the following properties:

- If an honest participant broadcasts a message, then all honest participants accept the message.
- If a dishonest participant P_i broadcasts a message, then either all honest participants accept the same message or no honest participant accepts any value from P_i.

Bracha's broadcast primitive runs as follows:

(1) The transmitter P_i sends the value $\langle P_i, initial, v\rangle$ to all participants.
(2) If a participant P_j receives a value v with one of the following messages
   - $\langle P_i, \texttt{initial}, v\rangle$
   - $\frac{n+t}{2}$ messages of the type $\langle \texttt{echo}, P_i, v\rangle$
   - $t+1$ message of the type $\langle \texttt{ready}, P_i, v\rangle$

   then P_j sends the message $\langle \texttt{echo}, P_i, v\rangle$ to all participants.
(3) If a participant P_j receives a value v with one of the following messages
   - $\frac{n+t}{2}$ messages of the type $\langle \texttt{echo}, P_i, v\rangle$
   - $t+1$ message of the type $\langle \texttt{ready}, P_i, v\rangle$

   then P_j sends the message $\langle \texttt{ready}, P_i, v\rangle$ to all participants.
(4) If a participant P_j receives $2t+1$ messages of the type $\langle \texttt{ready}, P_i, v\rangle$, then P_j accepts the message v from P_i.

Assume that $n = 3t+1$. The intuition for the security of Bracha's broadcast primitive is as follows. First, if an honest participant P_i sends the value $\langle P_i, initial, v\rangle$, then all honest participant will receive this message and echo the message v. Then all honest participants send the ready message for v and all honest participants accept the message v.

Secondly, if honest participants P_{j_1} and P_{j_2} send ready messages for u and v respectively, then we must have $u = v$. This is due to the following fact. A participant P_j sends a $\langle \texttt{ready}, P_j, u\rangle$ message only if it receives $t+1$ ready messages or $2t+1$ echo messages. That is, there must be an honest participant who received $2t+1$ echo messages for u. Since an honest participant can only send one message of each type, this means that all honest participants will only sends ready message for the value u.

In order for an honest participant P_j to accept a message u, it must receive $2t+1$ ready messages. Among these messages, at least $t+1$ ready messages are from honest participants. An honest

https://mc.manuscriptcentral.com/dlt

participant can only send one message of each type. Thus if honest participants P_{j_1} and P_{j_2} accept messages u and v respectively, then we must have $u = v$. Furthermore, if a participant P_j accepts a message u, we just showed that at least $t + 1$ honest participants have sent the ready message for u. In other words, all honest participants will receive and send at least $t + 1$ ready message for u. By the argument from the preceding paragraph, each honest participant sends one ready message for u. That is, all honest participants will accept the message u.

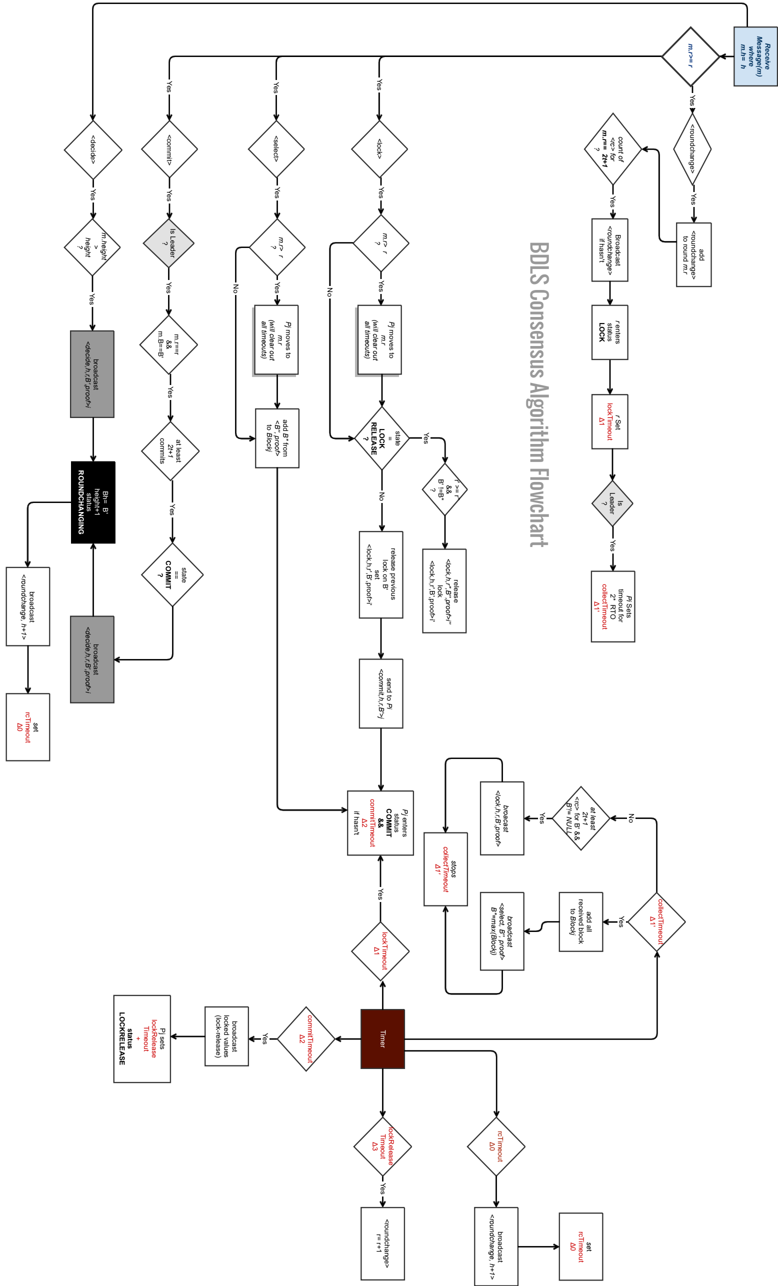

Fig. 3. BDLS Consensus Algorithm Flowchart

https://mc.manuscriptcentral.com/dlt