

Polygraph: Accountable Byzantine Agreement

Pierre Civot
University of Sydney

Seth Gilbert
National University of Singapore

Vincent Gramoli
University of Sydney
Data61-CSIRO

Abstract

In this paper, we introduce *Polygraph*, the first accountable Byzantine consensus algorithm for partially synchronous systems. If among n users $t < \frac{n}{3}$ are malicious then it ensures consensus, otherwise it eventually detects malicious users that cause disagreement. Polygraph is appealing for blockchain applications as it allows them, if $t < \frac{n}{3}$, to totally order blocks in a chain, hence avoiding forks and double spending and, otherwise, to punish malicious users when a fork occurs. We first show that stronger forms of the problem including a fixed time detection are impossible to solve before proving our solution. We also show that Polygraph has a bounded justification size so that each of its asynchronous rounds exchanges only $O(n^2)$ messages. Finally, we use a blockchain application to evaluate Polygraph on a geodistributed system of 80 machines and show that accountability reduces the performance of the non-accountable baseline by about a third, still committing thousands of transactions per second.

1 Introduction

Over the last several years we have seen a boom in the development of new Byzantine agreement protocols, in large part driven by the excitement over blockchains and cryptocurrencies. The (virtual) gold rush is on, and as in the Wild West of yore, the outlaws are ever present. Byzantine agreement protocols act as the locks on the bank doors, preventing the gangs from making off with the loot.

Limitations. Unfortunately, Byzantine agreement protocols have some inherent limitations: it is impossible to ensure correct operation when more than $\frac{1}{3}$ of the processing power in the system is controlled by a single malicious party, unless the network can guarantee per-

fect synchrony in communication.¹ And yet, as mining pools become larger and more centralized, we are in danger of a single group’s power surpassing the $\frac{1}{3}$ threshold. (For example, the largest Bitcoin mining pool today controls approximately 19% of the hashing power.) This problem is even more severe in consortium and private blockchains where a small number of predetermined parties maintain control over the chain, and hence are each more likely to control a significant fraction of the compute power of the network.

At first, one might hope to relax the liveness guarantees, while always ensuring safety. For example, perhaps we could ensure that on termination, the honest users would always agree; if the dishonest users have too much power, then the protocol never terminates. Alas, in a partially synchronous network, this type of guarantee is impossible (see Theorem 1). If the adversary controls more than $\frac{1}{3}$ of the computing power, it can always force disagreement.

Accountability. Since our Byzantine agreement “locks” are not good enough by themselves to protect the banks; we need a new sheriff in town to bring the guilty parties to justice. What if, instead of *preventing* bad behavior by a party that controls too much of the network power, we guarantee *accountability*, i.e., we can provide irrefutable evidence of the bad behavior and the perpetrator of those illegal actions? Much in the way we prevent crime in the real world, we can prevent bad blockchain behavior via defense-in-depth: the basic Byzantine agreement protocol disallows illegal actions if the attacker has $< \frac{1}{3}$ of the network under their control, or if the network infrastructure is working properly and ensuring timely message delivery; when these guarantees fail, we record sufficient information to catch the criminal and take remedial actions. For example, the offending transactions

¹Bitcoin is typically assumed to be correct if at least 50% of the network is honest. However, this holds only under the assumption that communication is synchronous and timely. Eclipse attacks [12], for example, can violate correctness even with a majority honest.

could be voided or reparations provided to those that were stolen from; a mining pool that behaves badly could be banned from the blockchain; or in the case of a consortium blockchain, the offline legal system could be used to pursue damages. Alternatively, the threat of being caught might function as an incentive for good behavior (e.g., if a certain deposit/stake were at risk).

The idea of accountability in distributed systems was pioneered by Haeberlen, Kuznetsov, and Druschel [11], and over the last several years, accountability has been increasingly discussed as an important goal in blockchains, especially in proof-of-stake systems; see, e.g., [5], discussed below, which provides accountability for a synchronous blockchain protocol.² There is an increasing consensus in the community that blockchain systems *should* provide accountability.

Intuition. Technically, there is nothing preventing this sort of accountability, even in a partially synchronous system. From an intuitive perspective, the impossibility of Byzantine agreement depends on a lack of timely communication. For example, the adversary controlling $> \frac{1}{3}$ of the network might divide the remaining honest users into two groups whose communication between groups is delayed; this would allow the adversary to convince each of the groups of a different outcome. However, when communication is eventually restored between the groups, they can discover the trickery and generate a proof that can bring the guilty parties to justice. Thus, in a partially synchronous network in which communication can be sometimes delayed (whether by simple network failures or by adversarial attacks), it remains possible to ensure strong forms of accountability as long as communication can eventually be restored among the honest parties.

At the same time, however, we must protect the honest users from false accusation. It is not enough for a small number of users to observe the bad behavior and report it to the authorities. How do we know whether those users were themselves honest? Instead, we need enough participants to observe the bad behavior to prove irrefutably that the criminal really did perpetrate the crime.

One may think that any Byzantine consensus can be extended with an independent detection algorithm executed right before the decision is taken and without changes to the core of the Byzantine consensus protocol itself. Such a generic solution is actually not trivial, as the Byzantine processes may produce a disagreement and refuse to participate in the detection algorithm when asked to do so, making it impossible (without synchrony) for honest nodes to prove their guilt.

²An interesting suggestion in Casper is that when a user is caught cheating, their stake is slashed, incentivizing good behavior.

Results. In this paper, we develop a new Byzantine agreement algorithm among n nodes out of which t can be Byzantine, which we refer to as the *Polygraph Protocol*, with the following guarantees: (i) if $t < \frac{n}{3}$, then consensus is guaranteed, i.e., every honest device produces the same valid output; (ii) no matter the number of Byzantine nodes, if a disagreement occurs between two honest nodes, every honest node eventually produces an irrefutable proof as to the identity of some malicious users.³

Polygraph message and communication complexities are $O(n^3)$ and $O(n^4)$, respectively, thanks to a bounded size justification. Note that the naive detection that exchanges signed messages that were broadcast, would be unpractical. Specifically, consider that every correct process simply forwards every signed message m from p to all other processes so that a correct node delivers m only when he has received $\lfloor \frac{2n}{3} \rfloor + 1$ occurrences of the same correctly signed messages, hence discarding possible lies. This would multiply the message complexity of the original consensus protocol by $\Omega(n)$ instead of $O(1)$ in Polygraph.

We also show that stronger forms of accountability are impossible. For example, we cannot guarantee agreement when $t > \frac{n}{3}$, even if we are willing to tolerate a failure of liveness; and processes cannot detect even one guilty party by a fixed time limit (e.g., prior to decision), since (intuitively) that would enable processes to determine guilt before deciding in a way that leads to disagreement. Nor can we guarantee detection of more than $\frac{n}{3}$ malicious users, since it takes only $\frac{n}{3}$ malicious users to cause disagreement and additional malicious users could simply stay mute to not be detected.

Finally, we describe how Polygraph can be used to hold participants accountable for their misbehavior in a blockchain application. First, we explain how to transform the accountable binary Byzantine consensus algorithm to accept arbitrary values, like the blocks of a blockchain. Second, we explain how consensus nodes can also be held accountable to client nodes that do not run the consensus as long as $t < 2\frac{n}{3}$. To evaluate Polygraph we implement it in the Red Belly Blockchain [9], and deploy it on 80 geodistributed machines. In particular, we compare the performance of the Red Belly Blockchain to this new “accountable” Red Belly Blockchain and observe that the cost associated with accountability, although largely visible, remains manageable.

Basic idea. The key observation is as follows: we can design Byzantine agreement protocols that can be defeated only by *disassembling*, i.e., by a malicious user send-

³While it might seem counterintuitive that Polygraph works regardless of the number of Byzantine nodes, note that when $t = n - 1$ or $t = n$, no disagreement can occur by definition.

ing different messages to different honest users, when (according to the protocol) it was supposed to send the same message to every honest user. And if there is dissembling, then the honest users have evidence of the bad behavior. Unfortunately, the malicious users might lie and accuse honest nodes of dissembling, and so extra care is needed before we can prove bad behavior.

The basic algorithm is based on the Byzantine agreement protocol in [8]. The protocol proceeds in (asynchronous) rounds, where processes try to converge on a single estimate, and try to decide if enough processes have already adopted that estimate. First, a reliable broadcast (similar to that described in [3]) is used to distribute the proposal values. Then, a second phase of communication is used to determine whether enough processes have converged on a single value. Finally the processes decide, if they can; and if not, they update their estimate in an attempt to converge on a single value.

The key to accountability is requiring that each of the messages in the second phase of communication are signed, and collecting *ledgers* of signed signatures that justify the estimate adopted. The exact requirement of the ledger depends on the situation; the important aspect is that a ledger guarantees that in a certain round, at least $2\frac{n}{3}$ processes broadcast a proposal for a single value. If we have two ledgers showing that in the same round, $2\frac{n}{3}$ processes broadcast a proposal for ‘1’ and $2\frac{n}{3}$ processes broadcast a proposal for ‘0’, then we can identify at least $\frac{n}{3}$ processes that dissembled, proposing both ‘0’ and ‘1’ in the same round.

We can show that, by propagating the ledgers properly, we can ensure that any process that decides has sufficient information to justify their decision. Honest processes that eventually disagree can exchange sufficient information to ensure that every process can provably identify at least $\frac{n}{3}$ users that acted maliciously.

Roadmap. The background is given in Section 2. The model and the accountable Byzantine consensus problem are presented in Section 3. In Section 4, we show that it is impossible to detect Byzantine failures in fixed time in a consensus algorithm. Section 5 depicts Polygraph an algorithm that solves the accountable binary Byzantine consensus problem. Section 6 analyses empirically Polygraph in a geodistributed blockchain and Section 7 concludes. Appendices A, B, C and D present the proof of the impossibility result, the proof of correctness of Polygraph, the multivalued extension and the application of Polygraph to blockchains, respectively.

2 Background and Related Work

In this section, we review existing work on accountability in distributed systems.

PeerReview. Haerberlen, Kuznetsov, and Druschel [11] pioneered the idea of accountability in distributed systems. They developed a system called *PeerReview* that implemented accountability as an add-on feature for any distributed system. Each process in the system records messages in tamper-evident logs; an authenticator can challenge a process, retrieve its logs, and simulate the original protocol to ensure that the process behaved correctly. They show that in doing so, you can always identify at least one malicious process (if some process acts in a detectably malicious way). Their technique is quite powerful, given its general applicability which can be used in any (deterministic) distributed system!

A natural alternative to The Polygraph Protocol, then, would be to take a standard existing consensus protocol like Paxos [22] and apply the PeerReview technique. The first issue is that, unlike in PeerReview, we want to guarantee correct outcomes when $t < \frac{n}{3}$, and only rely on accountability when t is large. Thus our base protocol must already be a (deterministic) Byzantine Agreement protocol, such as PBFT [6] or DBFT [8].

A second issue has to do with (partial) synchrony. The PeerReview approach is challenge-based: to prove misbehavior, an auditor must receive a response from the malicious process. If no response is received, the auditor cannot determine whether the process is malicious, or whether the network has not yet stabilized. It follows that the malicious coalition will only be suspected forever but not proved guilty. There is no fixed point at which the auditor can be *completely certain* that the sender is malicious; the auditor may never have definitive *proof* that the process is malicious; it always might just be poor network performance.⁴ The Polygraph Protocol, by contrast, produces a concrete proof of malicious behavior that is completely under the control of the honest processes.

A third issue has to do with message (and communication) complexity. The Polygraph Protocol has no (asymptotic) increase in message complexity over the base Byzantine Agreement Protocol on which it is based, and it increases the communication complexity by a factor of $\Theta(n)$. By contrast, PeerReview increases the message complexity by $\Theta(n^2)$ per audit interval (or $\Theta(n \log n)$ for probabilistic guarantees), while also increasing the communication complexity by a factor of

⁴To be precise, they refer to this situation as one where a malicious process is permanently *suspected* but never *exposed*; our goal is to guarantee that at least $\frac{n}{3}$ malicious processes are exposed, in their terminology.

$\Theta(n)$. Thus The Polygraph Protocol is more message efficient, while still preserving deterministic accountability guarantees.

As such, The Polygraph Protocol is more suitable for use in blockchain protocols. The key disadvantage of The Polygraph Protocol, of course, is that it is not general: it is solving a specific problem (Byzantine Agreement), rather than giving a solution for any distributed system. Moreover, The Polygraph Protocol detects misbehavior only when there is disagreement. The Byzantine processes can violate the protocol as much as they like, without detection, as long as the resulting consensus has no disagreement!

Accountable Blockchains. Recently, accountability has been an important goal in “proof-of-stake” blockchains, where users that violate the protocol can be punished by confiscating their deposited stake. Buterin and Grifith [5] have proposed a blockchain protocol, Casper, that provides this type of accountability guarantee. Validators try to agree on (or “finalize”) a branch of k hundreds of consecutive blocks, by gathering signatures for this branch or “link” from validators jointly owning at least $2\frac{2}{3}$ of the deposited stake. If a validator signs multiple links at the same height, Casper uses its signatures as proofs to slash its deposited stake. This is very similar in intent to The Polygraph Protocol’s notion of identifying $\frac{n}{3}$ malicious users when there is disagreement.

Like most blockchain protocols, however, Casper implicitly assumes some synchronous underlying (overlay) network and allows the blockchain to fork into a tree until some branch is finalized. To guarantee “plausible liveness” or that Casper does not block when not enough signatures are collected to finalize a link, validators are always allowed to sign links that overlap but extend links they already signed. However, this does not guarantee that consensus terminates. Consider, for example, that Byzantine validators always create the blocks at indices multiple of 100 and sends the block to less $2\frac{2}{3}$ honest validators.⁵ Repeating this process guarantees that $2\frac{2}{3}$ honest processes can never vote for the same link, and consensus will never be reached.

Even so, as a result of Casper, there has been much discussion in and around the blockchain community about accountability. It is certainly a popular idea today that accountability would be a desirable property of blockchains.

Accountability in the context of blockchain fairness was raised by Herlihy and Moir in a keynote address [13], and the idea of “accountable Byzantine fault tolerance” has been discussed [4]. (The goal in the latter case was to identify all the malicious nodes, which is

⁵Note that this requires Byzantine users have a large mining power if proof-of-work is needed for block creation.

unfortunately impossible.)

Earlier work on accountability. Even before PeerReview, others had suggested the idea of accountability in distributed systems as an alternate approach to security (see, e.g., [17,24,25]). Yumerefendi and Chase [25] developed an accountable system for network storage, and Repeat and Compare [21] developed an accountable peer-to-peer content distribution network.

The idea of accountability appeared less explicit in many earlier systems. For example, Aiyer et al. [1] proposed the BAR model for distributed systems, which relied on incentives to ensure good behavior; one key idea was in detecting and punishing bad behaviors. And Intrusion Detection Systems (e.g., [10,14,19] provided heuristics and techniques for detecting malicious behaviors in a variety of different systems.

Failure Detectors. There is a connection between accountability and failure detectors. A failure detector is designed to provide each process in the system with some advice, typically a list of processes that are faulty in some manner. However, failure detectors tend to have a different set of goals. They are used during an execution to help make progress, while accountability is usually about what can be determined *post hoc* after a problem occurs. They provide advice to a process, rather than proofs of culpability that can be shared. They tend to be designed for use in fully asynchronous systems (i.e., to capture synchrony assumptions), and protocols that rely on them tend to assume that $t < \frac{n}{3}$.

Most of the work in this area has focused on detecting crash failures (see, e.g., [7]). There has been some interesting work extending this idea to detecting Byzantine failures [11,15,19]. Malkhi and Reiter [19] introduced the concept of an unreliable Byzantine failure detector that could detect *quiet* processes, i.e., those that did not send a message when they were supposed to. They showed that this was sufficient to solve Byzantine Agreement.

Kihlstrom, Moser, and Melliar-Smith [15] continue this direction, considering failures of both omission and commission. Of note, they define the idea of a *mutant message*, i.e., a message that was received by multiple processes and claimed to be identical (e.g., had the same header), but in fact was not. The Polygraph Protocol is designed so that only malicious users sending a mutant message can cause disagreement. In fact, the main task of accountability in this paper is identifying processes that were supposed to broadcast a single message to everyone and instead sent different messages to different processes.

Mazières and Shasha propose SUNDR [20] that detects Byzantine behaviors in a network file system if all clients are honest and can communicate directly. Polygraph clients request multiple signatures from servers

so that they do not need to be honest. Li and Mazières [18] improves on SUNDR with BFT2F, a weakly consistent protocol when the number of failures is $\frac{n}{3} \leq t < 2\frac{n}{3}$ and its BFTx variant that copes with more than $2\frac{n}{3}$ failures but does not guarantee liveness even with less than t failures.

3 Model and Problem

We first define the problem in the context of a traditional distributed computing setting. (We later discuss applications to blockchains.)

System. We consider n processes. A subset C of the processes are honest, i.e., always follow the protocol; the remaining $t < n$ are Byzantine, i.e., may maliciously violate the protocol. We define $t_0 = \max(t \in \mathbb{N}_0 : t < \frac{n}{3})$, i.e., $t_0 = \lceil \frac{n}{3} \rceil - 1$, a useful threshold for Byzantine behavior.

Processes execute one step at a time and are asynchronous, proceeding at their own arbitrary, unknown speed. We assume local computation time is zero, as it is negligible with respect to message delays.

We assume that there is an idealized PKI (public-key infrastructure) so that each process has a public/private key pair that it can use to sign messages and to verify signatures.

Partial synchrony. We consider a partially synchronous network. During some intervals of time, messages are delivered in a reliable and timely fashion, while in other intervals of time messages may be arbitrarily delayed. More specifically, we assume that there is some time t_{GST} known as the *global stabilization time*, unknown to the processes, such that any message sent after time t_{GST} will be delivered with latency at most d . We say that an event occurs *eventually* if there exists an unknown but finite time when the event occurs.

Verification Algorithm. A verification algorithm V takes as input the state of a process and returns a set G of undeniable guilty processes, that is, every process-id of G is tagged with an unforgeable proof of culpability. (More formally, this means that for every computationally bounded adversary, for every execution in which a process p_j is honest, for every state s generated during the execution or constructed by Byzantine users, the probability that the verification algorithm returns a set containing p_j is negligible. In practice, this will reduce to the non-forgeability of signatures.)

Accountable Byzantine Agreement. The problem of Byzantine Agreement, first introduced by Pease, Shostak, and Lamport [16], assumes that each process begins with a binary *input*, i.e., either a 0 or a 1, outputs a *decision*, and requires three properties: agreement, validity, and termination.

We define the Accountable Byzantine Agreement problem in a similar way, with the additional requirement that there exists a verification algorithm that can identify at least $t_0 + 1$ Byzantine users whenever there is disagreement. (Recall that $t_0 = \lceil \frac{n}{3} \rceil - 1$.) More precisely:

Definition 1 (Accountable Byzantine Agreement). *We say that an algorithm solves Accountable Byzantine Agreement if each process takes an input value, possibly produces a decision, and satisfies the following properties:*

- **Agreement:** *If $t \leq t_0$, then every honest process that decides outputs the same decision value.*
- **Validity:** *If all processes are honest and begin with the same value, then that is the only decision value.*
- **Termination:** *If $t \leq t_0$, every honest process eventually outputs a decision value.*
- **Accountability:** *There exists a verification algorithm V such that: if two processes output disagreeing decision values, then eventually for every process p_j , for every state s_j reached by p_j from that point onwards, the verification $V(s_j)$ outputs a guilty set of size at least $t_0 + 1$.*

Our validity definition is sometimes called weak validity [23], but Lemma 8 shows that our accountable binary Byzantine consensus protocol ensures even a stronger validity property.

4 Impossibility Results

A couple of natural questions arise regarding accountable algorithms:

- Can we design an algorithm that always guarantees agreement, and simply fails to terminate if there are too many Byzantine users? If so, we would trivially get accountability!
- Can we design an algorithm that provides earlier evidence of Byzantine behavior, even before the decision is possible? If so, we could provide stronger guarantees than are provided in this paper.

Alas, neither is possible. The details of the following theorems are deferred to Appendix A (and follow from standard partitioning arguments).

Theorem 1. *In a partially synchronous system, no algorithm solves both the Byzantine consensus problem when $t < \frac{n}{3}$ and the agreement and validity of the Byzantine consensus problem when $t_0 < t \leq T$.*

We say that a verification algorithm V is *swift* if it provides the following guarantee: assume p_i has already decided some value v , and that p_j is in a state s wherein it will decide $w \neq v$ in its next step; then $V(s) \neq \emptyset$. Notice that a swift verification algorithm may only detect *one* Byzantine process (i.e., it is not sufficient evidence for p_j to decide never to decide).

Theorem 2. *For a consensus solved while $t < \frac{n}{3}$, there does not exist a swift verification algorithm for $t_0 < t \leq T$.*

The result follows directly from the impossibility of devising the needed swift verification algorithm.

5 Polygraph, an Accountable Byzantine Consensus Algorithm

In this section, we introduce *Polygraph*, a Byzantine agreement protocol that is accountable. We begin by giving the basic outline of the protocol for ensuring agreement when $t < \frac{n}{3}$ (which is derived from the Byzantine agreement protocol in [8]). Then, we focus on the key aspects that lead to accountability, specifically, the “ledgers” and “certificates.” In this section, we focus on binary agreement; in Section C, we show how to support arbitrary values. In Appendix B, we prove that the algorithm is correct.

As a notational issue, we indicate that a process p_i sends a message to every other process by: $\text{broadcast}(TAG, m) \rightarrow \text{messages}$, where TAG is the type of the message, m is the message content, and messages is the location to store any messages received.

Throughout we assume that every message is signed by the sender so the receiver can authenticate who sent it. (Any improperly signed message is discarded.) Thus we can identify messages sent by distinct processes. Similarly, the protocol will at times include cryptographically signed “ledgers” in messages; again, any message that is missing a required ledger or has an improperly formed ledger is discarded. (See the discussion below regarding ledgers.)

Protocol Overview. The basic protocol operates in two phases, after which a possible decision is taken. Each process maintains an estimate. In the first phase, each process broadcasts its estimate to all other users using a reliable broadcast service, *bv-broadcast* (discussed below). The protocol also uses a rotating coordinator; whoever is the assigned coordinator for the round broadcasts its estimate with a special designation.

All processes then wait until they receive at least one message, and until a timer expires. (The timeout is increased with each iteration, so that eventually once the network stabilizes it is long enough.) If a process receives a message from the coordinator, then it chooses the coordinator’s value to “echo”, i.e., to rebroadcast to everyone in the second phase. Otherwise, it simply echoes all the messages received in the first phase.

At this point, each process p_i waits until it receives enough *compatible* ECHO messages. Specifically, it waits to receive at least $(n - t_0)$ messages sent by distinct processes where every value in those messages was also received by p_i in the first phase. In this case, it adopts

the collection of values in those $(n - t_0)$ messages as its candidate set. In fact, if a process p_i receives a set of $(n - t_0)$ messages that *all* contain exactly the coordinator’s value, then it chooses only that value as the candidate value.

Finally, the processes try to come to a decision. If process p_i has only one candidate value v , then p_i adopts that value v as its estimate. In that case, it can decide v if it matches the parity of the round, i.e., if $v = r_i \pmod 2$. Otherwise, if p_i has more than one candidate value, then it adopts as its estimate $r_i \pmod 2$, the parity of the round.

To see that this ensures agreement (when $t < \frac{n}{3}$), consider a round in which some process p_i decides value $v = r_i \pmod 2$. Since p_i receives $(n - t_0)$ echo messages containing *only* the value v , we know that every honest process must have value v in every possible set of $(n - t_0)$ echo messages, and hence every honest process included v in its candidate set. Every honest process that *only* had v as a candidate also decided v . The remaining honest processes must have adopted $v = r_i \pmod 2$ as their estimate when they adopted the parity bit of the round. And if all the honest processes begin a round with estimate v , then that is the only possible decision due to the reliable broadcast *bv-broadcast* in Phase 1 (see below).

Processes always continue to make progress, if $t < \frac{n}{3}$. Termination is a consequence of the coordinator: eventually, after GST when the network stabilizes, there is a round where the coordinator is honest and the timeout is larger than the message delay. At this point, every honest process receives the coordinator’s Phase 1 message and echoes the coordinator’s value. In that round, every honest process adopts the coordinator’s estimate, and the decision follows either in that round or the next one (if $t < \frac{n}{3}$).

BV-Broadcast. The protocol relies in Phase 1 on a reliable broadcast routine *bv-broadcast*, which is used to ensure validity, i.e., any estimate adopted (and later decided) must have been proposed by some honest process. Moreover, it guarantees that if every honest process begins a round with the same value, then that is the only possible estimate for the remainder of the execution (if $t < \frac{n}{3}$). Specifically, *bv-broadcast* guarantees the following critical properties while $t < \frac{n}{3}$: (i) every message broadcast by $t_0 + 1$ honest processes is eventually delivered to every honest process (see Lemma 2); (ii) every message delivered to an honest process was broadcast by at least $t + 1$ processes (see Lemma 1).

These properties are ensured by a simple echo procedure. When a process first tries to *bv-broadcast* a message, it broadcasts it to everyone. When a process receives $t_0 + 1$ copies of a message, then it echoes it. When a process receives $n - t_0$ copies of a message, then it de-

Algorithm 1 The Polygraph Protocol

```

1: bin-propose( $v_i$ ):
2:    $est_i = v_i$ 
3:    $r_i = 0$ 
4:    $timeout_i = 0$ 
5:    $ledger_i[0] = \emptyset$ 
6:   repeat
7:      $r_i \leftarrow r_i + 1$ ; ▷ Increment the round number and the timeout.
8:      $timeout_i \leftarrow timeout_i + 1$ 
9:      $coord_i \leftarrow ((r_i - 1) \bmod n) + 1$  ▷ Rotate the coordinator.
  ▷ Phase 1:
10:  bv-broadcast( $EST[r_i], est_i, ledger_i[r_i - 1], i, bin\_values_i$ ) ▷ Binary value broadcast the current estimate.
11:  if  $i = coord_i$  then ▷ Coordinator rebroadcasts first value received.
12:    wait until ( $bin\_values_i[r_i] = \{w\}$ ) ▷  $bin\_values$  stores messages received by binary value broadcast.
13:    broadcast( $COORD[r_i], w$ )  $\rightarrow messages_i$ 
14:    StartTimer( $timeout_i$ )
15:    wait until ( $bin\_values_i[r_i] \neq \emptyset \wedge timer_i$  expired)
  ▷ Phase 2:
16:   $timer_i \leftarrow timeout_i$  ▷ Reset the timer.
17:  if ( $(COORD[r_i], w) \in messages_i$  from  $p_{coord_i} \wedge w \in bin\_values_i[r_i]$ ) then ▷ Prioritize the coordinator.
18:     $aux_i \leftarrow \{w\}$ 
19:  else  $aux_i \leftarrow bin\_values_i[r_i]$  ▷ Otherwise, use any value received.
20:   $signature_i = \text{sign}(aux_i, r_i, i)$  ▷ Sign the messages.
21:  broadcast( $ECHO[r_i], aux_i[r_i], signature_i$ )  $\rightarrow messages_i$  ▷ Broadcast second phase message.
22:  wait until  $values_i = \text{ComputeValues}(messages_i, bin\_values_i, aux_i) \neq \emptyset$ 
  ▷ Decision phase:
23:  if  $values_i = \{v\}$  then ▷ If there is only one value, then adopt it.
24:     $est_i \leftarrow v$ 
25:    if  $v = (r_i \bmod 2)$  then ▷ Decide if value matches parity.
26:      if no previous decision by  $p_i$  then decide( $v$ )
27:    else
28:       $est_i \leftarrow (r_i \bmod 2)$  ▷ Otherwise, adopt the current parity bit.
29:       $ledger_i[r_i] = \text{ComputeJustification}(values_i, est_i, r_i, bin\_values_i, messages_i)$  ▷ Broadcast certificate if decision.
Rules:
  1. Every message that is not properly signed by the sender is discarded.
  2. Every message that is sent by bv-broadcast without a valid ledger after Round 1, except for messages containing value 1 in Round 2, are discarded.
  3. On first discovering a ledger  $\ell$  that conflicts with a certificate, send ledger  $\ell$  to all processes.

```

livers it. Notice that if a message is not bv-broadcast by at least $t_0 + 1$ processes, then it is never echoed and hence never delivered. And if a message is bv-broadcast by $t_0 + 1$ (honest) processes is echoed by every honest process and hence delivered to every honest process.

This reliable broadcast routine ensures validity, since a Phase 1 message that is echoed in Phase 2 must have been delivered by bv-broadcast, and hence must have been bv-broadcast by at least one honest process.

Ledgers and Certificates. In order to ensure accountability, we need to record enough information during the execution to justify any decision that is made, and hence to allow processes to determine accountability. For this purpose, we define two types of justifications: ledgers and certificates. A ledger is designed to justify adopting a specific value. A certificate justifies a decision. We will attach ledgers to certain messages; any message containing an invalid or malformed ledger is discarded.

We define a ledger for round r and value v as follows. If $v \neq r \bmod 2$, then the ledger consists of the $(n - t_0)$ ECHO messages, each properly signed, received in Phase 2 of round r that contain only value v (and no other value). If $v = r \bmod 2$, then the ledger is simply a copy of *any* other ledger from the previous round $r - 1$ justifying value v . (The asymmetry may seem strange, but is useful in finding the guilty parties!)

We define a certificate for a decision of value v in round r to consist of $(n - t_0)$ echo messages, each properly signed, received in Phase 2 of round r that contain only value v (and no other value).

Accountability. We now explain how the ledgers and certificates are used. In every round, when a process uses bv-broadcast to send a message containing a value, it attaches a ledger from the previous round justifying why that value was adopted. (There is one exception: in Round 1, no ledger will be available to justify value 1, and so no ledger is generated in that case.)

Algorithm 2 Helper Components

```

1: bv-broadcast(MSG, val, ledger, i, bin_values):
2:   broadcast(BVAL, ⟨val, ledger, i⟩) → msgs ▷ Broadcast message.
3:   After round 2, and in round 1 if val = 0, discard all messages received without a proper ledger.
4:   upon receipt of (BVAL, ⟨v, ·, j⟩)
5:     if (BVAL, ⟨v, ·, ·⟩) received from (t0 + 1) distinct processes and (BVAL, ⟨v, ·, ·⟩) not yet broadcast then
6:       Let ℓ be any non-empty ledger received in these messages. ▷ Picking one of the received ledgers is enough.
7:       broadcast(BVAL, ⟨v, ℓ, j⟩) ▷ Echo after receiving (t0 + 1) copies.
8:     if (BVAL, ⟨v, ·, ·⟩) received from (2t0 + 1) distinct processes then
9:       Let ℓ be any non-empty ledger received in these messages. ▷ Picking one of the received ledgers is enough.
10:      bin_values ← bin_values ∪ {⟨v, ℓ, j⟩} ▷ Deliver after receiving (2t0 + 1) copies.

11: ComputeValues(messages, b_set, aux_set): ▷ Check if there are n - t0 compatible messages.
12:   if ∃S ⊆ messages where the following conditions hold:
13:     (i) |S| contains (n - t0) distinct ECHO[ri] messages
14:     (ii) aux_set is equal to the set of values in S.
15:     then return(aux_set)
16:   if ∃S ⊆ messages where the following conditions hold:
17:     (i) |S| contains (n - t0) distinct ECHO[ri] messages
18:     (ii) Every value in S is in b_set.
19:     then return(V = the set of values in S)
20:   else return(∅)

21: ComputeJustification(valuesi, esti, ri, bin_valuesi, messagesi): ▷ Compute ledger and broadcast certificate if decision.
22:   if esti = (ri mod 2) then
23:     if ri > 1 then
24:       return ledger[ri]i = ledger ℓ where (EST[ri], ⟨v, ℓ, ·⟩) ∈ bin_valuesi
25:     else return ledger[ri]i = ∅
26:   else return ledger[ri]i = (n - t0) signed messages from messagesi containing only value esti
27:   if valuesi = {(ri mod 2)} ∧ no previous decision by pi in previous round then
28:     certificatei = (n - t0) signed messages from messagesi containing only value esti
29:     broadcast(esti, ri, i, certificatei) ▷ Transmit certificate to everyone.

```

The bv-broadcast ignores the ledger for the purpose of deciding when to echo a message. When it echoes a message m , it chooses any arbitrary non-empty ledger that was attached to a message containing m (if any such ledgers are available). However, every message that does not contain a valid ledger justifying its value is discarded, with the following exception: in Round 2, messages containing the value 1 can be delivered without a ledger (since no justification is available for adopting the value 1 in Round 1).

Whenever there is only one candidate value received in Phase 2, a process adopts that value and either: (i) decides and constructs a certificate, or (ii) does not decide and constructs a ledger. In both cases, this construction simply relies on the signed messages received in Phase 2 of that round (and hence is always feasible).

If a process decides a value v in round $r > 1$, or adopts v because it is the parity bit for round $r > 1$, then it also constructs a ledger justifying why it adopted that value v . It accomplishes this by examining all the bv-broadcast messages received for value v and copying a round $r - 1$ ledger. Again, this is always possible since any message that is not accompanied by a valid ledger is ignored. (The only possible problem occurs in Round 2 where messages for value 1 are not accompanied by a ledger;

however ledgers for value 1 in round 2 do not require copying old ledgers.)

Proving Culpability. How do disagreeing processes decide which processes were malicious? When a process decides in round r , it sends its certificate to all the other processes. Any process that decides a different value in a round $> r$ can prove the culpability of at least $\lceil \frac{n}{3} \rceil$ Byzantine processes by comparing this certificate to its logged ledgers. (It can then broadcast the proper logged ledgers to ensure that everyone can identify the malicious processes.)

We will say that a certificate (e.g., from p_1) and a ledger (e.g., from p_2) *conflict* if they are constructed in the same round r , but for different values v and w . That is, both the certificate and the ledger attest to $(n - t_0)$ ECHO messages from round r sent to p_1 and p_2 (respectively) that contain only value v and only value w , respectively. Since every two sets of size $(n - t_0)$ intersect in at least $(t_0 + 1)$ locations, this conflict identifies $(n - t_0)$ processes that sent different Phase 2 messages in round r to p_1 and p_2 and hence they are malicious.

We now discuss how to find conflicting certificates and ledgers. Assume that process p_i decides value v in round r , and that process p_j decides a different value w in a round $> r$. (Recall that v is the only possible value

that can be decided in round r .) There are two cases to consider, depending on whether p_j decides in round $r + 1$ or later.

- *Round $r + 1$:* If p_j decides in round $r + 1$, then value w was the only candidate value after Phase 2. This implies that w was received by some bv-broadcast message. Since $r > 1$, we know that the message must have contained a valid ledger ℓ from round r for value $w \neq v$. This ledger ℓ conflicts with the decision certificate of p_i .
- *Round $\geq r + 2$:* Since p_j decides $w \neq v$, it does not decide v in round $r + 2$. This means that p_j has w as a candidate value, which implies that p_j received w in a bv-broadcast. Since $r > 1$, we know that the message must have contained a valid ledger ℓ from round $r + 1$ for value $w \neq v$. This ledger ℓ consists of a copy of a ledger from round r for value w which conflicts with the decision certificate of p_i .

In either case, if p_j does not decide v , then, by looking at the messages received in round $r + 1$ and $r + 2$, it can identify a ledger that conflicts with the decision certificate of p_i and hence can prove the culpability of at least $t_0 + 1$ malicious processes.

Analysis. In Appendix B.1, we show that the the BV-broadcast routine provides the requisite properties. This then allows us to prove the main correctness theorem, which follows immediately from Lemma 7, Corollary 1, and Lemma 9 in Appendix B.2:

Theorem 3. *The Polygraph Protocol is a correct Byzantine agreement protocol guaranteeing agreement, validity, and termination.*

Accountability follows from Lemma 10, which shows that disagreement leads to every honest process eventually receiving a certificate and a ledger that conflict:

Theorem 4. *The Polygraph Protocol is accountable.*

Finally, we show in Lemma 11 that each round has a message complexity of $O(n^2)$ and a communication complexity of $O(n^3)$. If all the processes are honest, or if the Byzantine corruptions are oblivious to the processor identities, then the protocol terminates in $O(1)$ rounds after GST. Otherwise, it may take $t + 1$ rounds after GST to terminate.

In Appendix C, we show that the multivalued generalization of The Polygraph Protocol is also correct and accountable.

6 Experimenting Polygraph with a Blockchain Application

To understand the overhead of Polygraph over a non-accountable consensus, we compare the throughput of

the original Red Belly Blockchain [9] based on DBFT [8] and the “Accountable Red Belly Blockchain” based on Polygraph as described in Appendices C and D. We deploy both blockchains on up to $n = 80$ AWS virtual machines located in 5 availability zones on two continents: California, Oregon, Ohio, Germany and Ireland. All machines issue transactions, insert transactions in their “mempool” or memory pool, propose blocks of 10,000 transactions, verify transaction signatures and account integrity, and run their respective consensus algorithm with $t = t_0 = \lceil \frac{n}{3} \rceil - 1$, before storing decided blocks to non-volatile storage.

Figure 1 represents the throughput while increasing the number of consensus participants from 10 (2 machines per zone) to 80




Figure 1: Accountability cost on a geodistributed system.

(16 machines per zone). We observe that Red Belly Blockchain is about a third faster than the Accountable Red Belly Blockchain. We also observe that performance decreases for both blockchains while the Red Belly Blockchain is known to be scalable [9] but we attribute this difference to the size of the blocks, ours contain $10 \times$ more transactions which can create network congestion. Finally, the Accountable Red Belly Blockchain commits several thousands of transactions per second at 80 nodes, which indicates that the cost of accountability remains practical.

7 Conclusion

We introduced Polygraph, the first accountable Byzantine consensus algorithm for partially synchronous systems. If $t < \frac{n}{3}$, it ensures consensus, otherwise it eventually detects malicious users that cause disagreement. Polygraph is practical thanks to its bounded justification size that does not increase with the number of rounds. As future work, we could extend Polygraph to detect malicious users even when there is agreement and detect all malicious “active” users [15] instead of simply up to $\frac{n}{3}$.

Acknowledgements

This research is in part supported under Australian Research Council Discovery Projects funding scheme (project number 180104030) entitled “Taipan: A Blockchain with Democratic Consensus and Validated Contracts”. Vincent Gramoli is a Future Fellow of the Australian Research Council.

References

- [1] Amitanand S. Aiyer, Lorenzo Alvisi, Allen Clement, Michael Dahlin, Jean-Philippe Martin, and Carl Porth. BAR fault tolerance for cooperative services. *SOSP*, 2005.
- [2] Emmanuelle Anceaume, Antonella Del Pozzo, Romaric Ludinard, Maria Potop-Butucaru, and Sara Tucci Piergiovanni. Blockchain abstract data type. *IACR Cryptology ePrint Archive*, 2018:561, 2018.
- [3] Gabriel Bracha. Asynchronous Byzantine agreement protocols. *Information and Computation* 75:130-143, 1985.
- [4] Ethan Buchman. On the design and accountability of Byzantine fault tolerant protocols, 2017. Personal communication.
- [5] Vitalik Buterin and Virgil Griffith. Casper the friendly finality gadget. Technical Report 1710.09437v4, arXiv, Jan 2019.
- [6] Miguel Castro and Barbara Liskov. Practical Byzantine fault tolerance and proactive recovery. *ACM Transactions on Computer Systems*, 20(4), 2002.
- [7] Tushar Deepak Chandra and Sam Toueg. Unreliable failure detectors for reliable distributed systems. *Journal of the ACM*, Volume 43 Issue 2, Pages 225-267, 1996.
- [8] Tyler Crain, Vincent Gramoli, Mikel Larrea, and Michel Raynal. DBFT: Efficient leaderless Byzantine consensus and its applications to blockchains. In *Proceedings of the 17th IEEE International Symposium on Network Computing and Applications (NCA'18)*. IEEE, 2018.
- [9] Tyler Crain, Chris Natoli, and Vincent Gramoli. Evaluating the Red Belly Blockchain. Technical Report 1812.11747, arXiv, 2018.
- [10] Dorothy E. Denning. An intrusion-detection model. *IEEE Trans. Software Eng.*, 13(2), 1987.
- [11] Andreas Haeberlen, Petr Kouznetsov, and Peter Druschel. PeerReview: Practical accountability for distributed systems. *SOSP'07*, 2007.
- [12] Ethan Heilman, Alison Kendler, Aviv Zohar, and Sharon Goldberg. Eclipse attacks on bitcoin’s peer-to-peer network. In *24th USENIX Security Symp. (USENIX Security 15)*, pages 129–144. USENIX Association, 2015.
- [13] Maurice Herlihy and Mark Moir. Blockchains and the logic of accountability: Keynote address. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016*, pages 27–30, 2016.
- [14] Márk Jelasity, Alberto Montresor, and Ozalp Babaoglu. Detection and removal of malicious peers in gossip-based protocols. In *In Proceedings of FuDiCo*, June 2004.
- [15] Kim P. Kihlstrom, Louise E. Moser, and Peter M. Melliari-Smith. Byzantine fault detectors for solving consensus. *British Computer Society*, 2003.
- [16] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.
- [17] Butler W. Lampson. Computer security in the real world. In *In Proc. Annual Computer Security Applications Conference*, December 2000.
- [18] Jinyuan Li and David Mazières. Beyond one-third faulty replicas in Byzantine fault tolerant systems. In *Proceedings of the 4th USENIX Conference on Networked Systems Design & Implementation*, NSDI'07, pages 10–10, 2007.
- [19] Dahlia Malkhi and Michael K. Reiter. Unreliable intrusion detection in distributed computations. In *10th Computer Security Foundations Workshop (CSFW '97)*, June 10-12, 1997, Rockport, Massachusetts, USA, pages 116–125, 1997.
- [20] David Mazières and Dennis Shasha. Building secure file systems out of Byzantine storage. In *Proceedings of the Twenty-first Annual Symposium on Principles of Distributed Computing*, PODC '02, pages 108–117, 2002.
- [21] Nikolaos Michalakis, Robert Soulé, and Robert Grimm. Ensuring content integrity for untrusted peer-to-peer content distribution networks. In *4th Symposium on Networked Systems Design and Implementation (NSDI 2007)*, April 11-13, 2007, Cambridge, Massachusetts, USA, Proceedings., 2007.

- [22] Roberto De Prisco, Butler W. Lampson, and Nancy A. Lynch. Revisiting the PAXOS algorithm. *Theor. Comput. Sci.*, 243(1-2):35–91, 2000.
- [23] Roberto De Prisco, Dahlia Malkhi, and Michael K. Reiter. On k -set consensus problems in asynchronous systems. In *Proceedings of the Eighteenth Annual ACM Symposium on Principles of Distributed Computing, PODC, '99Atlanta, Georgia, USA, May 3-6, 1999*, pages 257–265, 1999.
- [24] Aydan R. Yumerefendi and Jeffrey S. Chase. Trust but verify: accountability for network services. In *Proceedings of the 11st ACM SIGOPS European Workshop, Leuven, Belgium, September 19-22, 2004*, page 37, 2004.
- [25] Aydan R. Yumerefendi and Jeffrey S. Chase. Strong accountability for network storage. *TOS*, 3(3):11:1–11:33, 2007.

A Proofs of the Impossibility Result

Process. More formally, a process $p_i, i \in [1, n]$ is seen as a Mealy machine $(S, S_0, \Sigma, \Lambda, \delta, \omega)$ where S is a set of states, S_0 a set of initial states, Σ an input alphabet of receivable message, Λ an output alphabet of messages to send, $\delta: \Sigma^n \times S \rightarrow S$ a transition function and $\omega: \Sigma^n \times S \rightarrow \Lambda^n$ an output function. The absence of message is denoted by $\{\epsilon\}$, $\{\epsilon\} \in \Sigma$ and $\{\epsilon\} \in \Lambda$. The process is fitted with a clock which allows it to take a transition after a certain timeout with $\{\epsilon\}$ as argument for δ .

The agreement according to a local view. A node processing an agreement protocol follows an ordered succession of states $p(s_k) = s_0, \dots, s_k$ called a path. On the path, the node has received messages which constitutes an history $h(p(s_k))$. The set S is a partition of the set of states before decision S_{bd} and the set of states after decision S_{ad} . The image of the restriction of δ on S_{ad} is S_{ad} , which means a decision is irrevocable. Among the attribute of $s \in S$, we note $s.r$ the associated round number and $s.dec$ the decided value.

To show that it is impossible to devise a swift T -Accountability algorithm, we show that no consensus algorithm can be safe when $t \geq \frac{n}{3}$ and that it is impossible given a consensus algorithm to design a swift verification algorithm for it when $t_0 < t \leq T$. Let $n \in \{3t_0 + 1, 3t_0 + 2, 3t_0 + 3\}$ and $T > t_0$. Let P, Q and R be three sets of n processes such that $|P| \leq t_0, |R| \leq t_0$ and $|Q| = n - |P| - |R|$. The proofs rely on the indistinguishability between pairs of distinct scenarios A, B and C such that processes in P cannot distinguish A from C and R cannot distinguish B from C.

- **Scenario A:** All initial values are 0 and the processes in R are inactive. The messages sent from $P \cup Q$ to $P \cup Q$ are delivered in time 1. By t -resiliency, processes in P reach a decision (0 by validity) within a certain time, noted T_A . To take this decision, every node i of P followed a path $s_0^i, \dots, s_{bd}^i, s_{ad}^i$ where $s_{bd}^i \in S_{bd}$ and $s_{ad}^i.dec = 0$.
- **Scenario B:** All initial values are 1 and the processes in P are inactive. The messages sent from $R \cup Q$ to $R \cup Q$ are delivered in time 1. By t -resiliency, processes in R reach a decision (1 by validity) within a certain time, noted T_B . To take this decision, every node j of R followed a path $s_0^j, \dots, s_{bd}^j, s_{ad}^j$ where $s_{bd}^j \in S_{bd}$ and $s_{ad}^j.dec = 1$.
- **Scenario C:** All initial values in P are 0, all initial values in R are 1 and processes of Q are Byzantine. These Byzantine processes behave with respect to those in P exactly as they do in Scenario A and with respect to those in R exactly as they do in Scenario B. Messages sent from $P \cup Q$ to $P \cup Q$ are delivered in time 1, as well as the ones from $R \cup Q$ to $R \cup Q$, while the ones from $P \cup R$ to $P \cup R$ are delivered in a time greater than $\max(T_A, T_B)$.

Theorem 5 (Theorem 1). *In a partially synchronous system, no algorithm solves both the Byzantine consensus problem when $t < \frac{n}{3}$ and the agreement and validity of the Byzantine consensus problem when $t_0 < t \leq T$.*

Proof. Assume for the sake of contradiction that it exists an algorithm preserving the agreement for $t_0 < t \leq T$. Because scenarios A and B are indistinguishable from P 's standpoint while scenarios B and C are indistinguishable for R 's standpoint, P must decide 0 while R must decide 1 in scenario C. This yields a contradiction. \square

Theorem 6 (Theorem 2). *For a consensus solved while $t < \frac{n}{3}$, it does not exist a swift verification algorithm for $t_0 < t \leq T$.*

Proof. Assume for the sake of contradiction that such a swift verification algorithm exists. The paths followed by the honest processes are the same in the scenario C, than in the scenarios A and B. Because no comission message has been sent in A and B, $\forall i \in P, V(s_{bd}^i) = \emptyset$ and $\forall j \in R, V(s_{bd}^j) = \emptyset$. But in scenario C, we should have $\exists i \in PV(s_{bd}^i) \neq \emptyset$ or $\exists j \in R, V(s_{bd}^j) \neq \emptyset$ which yields a contradiction. \square

B Proof of Correctness of Polygraph

In this section, we provide the complete proofs that the binary consensus protocol presented in Section 5 is a correct, accountable Byzantine agreement protocol. First, in

Section B.1, we focus on the properties satisfied by the accountable BV-Broadcast algorithm. Then we return to show the correctness of the consensus algorithm, as well as prove that it ensures accountability.

B.1 Accountable BV-Broadcast

The BV-broadcast algorithm is presented in lines 1–9 of Algorithm 2. The protocol and proof presented here are quite similar to that in [8], with small changes to accommodate ledgers.

First, we prove that it satisfies certain useful properties when $t \leq t_0$. (These properties will not necessarily hold when there are more than t_0 Byzantine processes.) We begin with a simple structural property (that is later useful for ensuring validity of consensus): a value is only delivered if at least one honest process sends it.

Lemma 1 (BV-Justification). *If $t \leq t_0$ and if p_i is honest and $v \in \text{bin_values}$, then v has been BV-broadcast by some honest process.*

Proof. By contraposition, assume v has been BV-broadcast only by faulty processes, i.e., by at most t_0 processes. Then no process ever receives $t_0 + 1$ distinct BVAL messages for v , and hence the conditions on lines 4 and 7 are never met. Thus, no process ever adds v to bin_values . \square

We next show that BV-broadcast satisfies some typical properties of reliable broadcast, i.e., if at least $t_0 + 1$ process BV-broadcast the same value, then every honest process delivers it; and if p_i is honest and delivers a value, then every honest process also delivers it.

Lemma 2 (BV-Obligation). *If $t \leq t_0$ and at least $(t_0 + 1)$ honest processes BV-Broadcast the same value v , then v is eventually added to the set bin_values of each non-faulty process p_i .*

Proof. Let v be a value such that $(t_0 + 1)$ honest processes invoke $\text{BV-broadcast}(\text{MSG}, v, \cdot, \cdot, \cdot)$. Each of these processes then sends a BVAL message with value v and a valid ledger, and consequently each honest process receives at least $t_0 + 1$ BVAL messages for value v along with valid ledgers for v . Therefore each honest process (i.e., at least $2t_0 + 1 \leq n - t_0$) broadcasts a BVAL message for v with a valid ledger, and consequently eventually every honest process receives at least $2t_0 + 1$ BVAL messages for v . Thus every honest process adds value v to bin_values . \square

Lemma 3 (BV-Uniformity). *If $t \leq t_0$ and a value v is added to the set bin_values of an honest process p_i , eventually $v \in \text{bin_values}$ at every honest process p_j .*

Proof. If a value v is added to the set bin_values of an honest process p_i , then this process has received at least $(2t_0 + 1)$ distinct BVAL messages for v with valid ledgers (line 7). This implies that at least $(t_0 + 1)$ different honest processes sent BVAL messages with valid ledgers. So every non-faulty process receives at least $(t_0 + 1)$ BVAL messages with the value v . Therefore every honest process eventually broadcasts a BVAL message for v , and so eventually every honest process receives at least $2t_0 + 1 \leq n - t_0$ BVAL messages for v with valid ledgers. Thus every honest process adds value v to bin_values . \square

Finally, we show termination:

Lemma 4 (BV-Termination). *If $t \leq t_0$ and if every honest process BV-broadcasts some value, then eventually, every honest process has at least one value in bin_values .*

Proof. As there are at least $(n - t_0)$ honest processes, each of them BV-broadcasts some value, and ‘0’ and ‘1’ are the only possible values, it follows that there is a value $v \in \{0, 1\}$ that is BV-broadcast by at least $(n - t_0)/2 \geq t_0 + 1$ processes (since one of the two values must be BV-broadcast by at least half the honest processes). The claim then follows by Lemma 2. \square

Finally, we observe the straightforward fact that messages are only delivered with valid ledgers (with the exception of messages in Round 1, and message containing value 1 in Round 2):

Lemma 5 (BV-Accountability). *If a value v is added to the set bin_values of a non-faulty process p_i in round r (with the exception of messages in Round 1, and message containing value 1 in Round 2), then associated with the value v is a valid ledger from round $r - 1$.*

Proof. Since every BVAL message without a valid ledger is discarded (aside from above mentioned exceptions), it follows immediately that when the conditions on lines 4 and 7 are met, then the process has access to a valid ledger, which is then included when the value is added to bin_values . \square

B.2 Accountable Byzantine Agreement

Here, we prove that The Polygraph Protocol is a correct accountable Byzantine agreement protocol. We begin with the standard properties of consensus, which hold when $t \leq t_0$, and then continue to discuss accountability. First, we observe that if every honest process begins a round r with the same estimate, then that value is decided either in round r or round $r + 1$. This follows immediately from the fact that if every honest process BV-broadcasts the same value, then that is the only value

delivered, and so it is the only value that remains in the system.

Lemma 6. *Assume that each honest process begins round r with the estimate v . Then every honest process decides v either at the end of round r or round $r + 1$.*

Proof. Since every honest process BV-broadcasts the value v , we know from Lemma 2 that v is eventually delivered to every honest process, and from Lemma 1 that v is the only value delivered to each honest process. Since v is the only value in bin_values_i for each honest p_i , it is also the only value echoed via aux_i messages, and eventually it is the only value in values_i .

There are now two cases. If $v = r \bmod 2$, every honest process decides v . Otherwise, every honest process continues to the next round with its estimate equal to v . In the next round, $v = (r + 1) \bmod 2$, and by the same argument, every process will then decide v in round $r + 1$. \square

We can now observe that if $t \leq t_0$, we get agreement, because after the first round in which some process decides v , then every process adopts value v as its estimate.

Lemma 7 (Agreement). *If $t \leq t_0$ and some honest processes p_i and p_j decide v and w , respectively, then $v = w$.*

Proof. Without loss of generality, assume that p_i decides no later than p_j . Assume for the sake of contradiction that $v \neq w$. Assume that p_i decides v in round r . If process p_j also decides in round r , then $v = w = r \bmod 2$. Thus, we conclude that p_j decides in some round $> r$.

In round r , we know that $\text{values}_i = \{v\}$. This implies that i received at least $n - t_0$ distinct ECHO messages containing only value v . Consider now the ECHO messages received by some other honest process p_k . If $\text{values}_k = \{v\}$ or $\text{values}_k = \{v, w\}$, then process p_k adopts estimate v . Otherwise, process p_k has $\text{values}_k = \{w\}$, which implies that it received at least $n - t_0$ distinct ECHO messages containing only value w .

Thus there are at least $t_0 + 1$ processes that sent an ECHO message to p_i containing only v and an ECHO message to p_k containing only w . That, however, is illegal (see line 21), as a process must send the same ECHO message to all. Since only $t \leq t_0$ processes are Byzantine, this is impossible, so we conclude that every honest process adopts estimate v by the end of round r .

We then conclude, by Lemma 6, that every honest process decides value v in either round $r + 1$ or round $r + 2$. (In this case, of course, it will be round $r + 2$.) \square

It is immediate from BV-broadcast that Polygraph guarantees a stronger form of validity (Lemma 8) than the one required (Corollary 1). (The General Polygraph

Protocol stated in Appendix C that applies to arbitrary values only ensures the required validity.)

Lemma 8 (Strong validity). *If $t \leq t_0$ and an honest process decides v , then some honest process proposed v .*

Proof. Let round r be the first round where a process p_i adopts a value v that was not initially proposed by an honest process at the beginning of round 1. There are two possibilities depending on whether p_i adopts v in line 24 or line 28.

Assume that honest process p_i in round r adopts value v in line 24. Then $\text{values}_i = \{v\}$ in round r . This can only happen if $\text{bin_values}_i = \{v\}$, since ComputeValues only returns values that are in bin_values_i or aux_i , and aux_i only includes values in bin_values_i . However, bin_values_i only includes values delivered by BV-broadcast, and Lemma 1 implies that every value delivered value was BV-broadcast by an honest process. So value v was the estimate of an honest process p_j at the beginning of round r . If $r = 1$ then we are done. If $r > 1$, then we conclude that v was adopted as an estimate in round $r - 1$ by process p_j , and hence by induction, we conclude that v was initially proposed by an honest process at the beginning of round 1.

Assume that honest process p_i executes line 28, adopting the parity of the round number as its estimate. This implies that $\text{values}_i = \{0, 1\}$ in round r . This can only happen if $\text{bin_values}_i = \{0, 1\}$ also, since ComputeValues only returns values that are in bin_values_i or aux_i , and aux_i only includes values in bin_values_i . However, bin_values_i only includes values delivered by BV-broadcast, and Lemma 1 implies that every value delivered value was BV-broadcast by an honest process. So at least one honest process p_j began round r with value '0' and at least one honest process p_k began round r with '1'. If $r = 1$, then we are done. Otherwise, we conclude that p_j adopted '0' in round $r - 1$ and p_k adopted '1' in round $r - 1$, and hence by induction we conclude that both '0' and '1' were initially proposed by honest processes at the beginning of round 0. \square

Corollary 1 (Validity). *If all processes are honest and begin with the same value, then that is the only decision value.*

Finally, we argue that the protocol terminates:

Lemma 9 (Termination). *If $t \leq t_0$, every honest process decides.*

Proof. First, we observe that processes continue executing increasing rounds (i.e., no process gets stuck in some round). Assume, for the sake of contradiction, that r is the first round where some process gets stuck forever, and p_i is the process that gets stuck.

A process cannot get stuck in line 12 or line 15 waiting for a BV-broadcast, since every honest process performs

a BV-broadcast and so by Lemma 4, every honest process eventually delivers a value. (And a process cannot get stuck waiting on a timer, since the timer will always eventually expire.)

A process also cannot get stuck waiting on line 22: Eventually process p_i will receive ECHO messages from each of the $n - t_0$ honest processes. And by Lemma 3, every value that is delivered by BV-broadcast to an honest process will eventually be delivered to p_i . Specifically, every value that is included in an ECHO message from an honest process is eventually delivered to bin_values_i . Thus, eventually a set of $n - t_0$ messages is identified, and the waiting condition on line 22 is satisfied.

The last issue that might prevent progress is if process p_i , or some other process, cannot transmit a proper message due to missing ledgers. This can only be a problem with messages being BV-broadcast. If a process completed round $r - 1$ and its estimate was $\neq r - 1 \pmod 2$, then it could always construct a proper ledger in round $r - 1$ from the estimates received. If a process completed round $r - 1$ and its estimate was $= r - 1 \pmod 2$, then it must have received a valid ledger for the value in round $r - 1$ as part of the BV-broadcast; otherwise, it could not have completed round $r - 1$.

Thus we conclude that every process executes an infinite number of rounds. The remaining question is whether processes ever decide. Consider the first round r after GST where the timer is sufficiently large that (i) every BV-broadcast message is delivered, and (ii) the coordinators message is delivered before the timer expires. In this case, every honest process will prioritize the coordinator's value, adopting it as their *aux* message in line 18, echoing it in line 21, and adding only that value to *values* in line 22. Thus at the end of round r , every process adopts the same value, and hence decides either in round r or round $r + 1$ by Lemma 6. \square

Thus we conclude (see Theorem 3) that when $t \leq t_0$, the binary agreement protocol is correct. We now consider the case where $t > t_0$ and show that it still provides accountability.

Lemma 10 (Accountability). *If $t > t_0$ and two honest processes p_i and p_j decide different values v and w , then eventually every honest process receives a ledger and a certificate that conflict (providing irrefutable proof that a specific collection of $t_0 + 1$ processes are Byzantine).*

Proof. Assume that p_i decided v in round r and p_j decided w in the round r' where $w = \text{not}(v) = 1 - v$ and $r \leq r'$. It is undeniable (by construction, line 25–26) that $v = r \pmod 2$. There are only four possible cases to consider:

1. Case 1: $values_j^r \neq \{0,1\}$
2. Case 2: $values_j^r = \{0,1\}$ and $values_j^{r+1} \neq \{v\}$

3. Case 3: $values_j^r = \{0,1\}$ and $values_j^{r+1} = \{v\}$ and $values_j^{r+2} \neq \{v\}$
4. Case 4: $values_j^r = \{0,1\}$ and $values_j^{r+1} = \{v\}$ and $values_j^{r+2} = \{v\}$

We now consider each of the cases in turn.

Case 1. If $values_j^r = \{v\}$, then process p_j would have decided $v \neq w$ in round r . So we conclude that $values_j^r = \{w\}$. In that case, $ledger[r]_j$ and $certificate[r]_i$ conflict.

Case 2. Assume $values_j^r = \{0,1\}$ and $w \in values_j^{r+1}$. This implies that $w \in bin_values_j$ in round $r + 1$. Notice that round $r + 1$ cannot be round 1, and if it is round 2, then the value $v = 1$ and so $w \neq 1$. Thus, BV-Accountability (Lemma 5 indicates that p_j receives a valid ledger for w from round r . That ledger conflicts with the certificate for v from r (and consists of $n - t_0$ distinct ECHO messages containing only value w in round r).

Case 3. Assume $values_j^r = \{0,1\}$, $values_j^{r+1} = \{v\}$, and $w \in values_j^{r+2}$. This implies that $w \in bin_values_j$ in round $r + 2$. Since round $r + 2 > 2$, BV-Accountability (Lemma 5 indicates that p_j receives a valid ledger for w from round $r + 1$. Since $w = r + 1 \pmod 2$, the ledger for w from round $r + 1$ is a copy of a ledger for w from round r , which therefore conflicts with the certificate for v for round r (and consists of $n - t_0$ distinct ECHO messages containing only value w in round r).

Case 4. Assume $values_j^r = \{0,1\}$, $values_j^{r+1} = \{v\}$, and $values_j^{r+2} = \{v\}$. Then process p_j decides v in round $r + 2$ and there is agreement.

All the cases have been examined, and in each case, process p_j has a ledger constructed in round r conflicting with the certificate delivered from process p_i . The conflicting ledger/certificate each contain $n - t_0$ signed, distinct ECHO messages containing only value v and only value w respectively. Since any two sets of size $n - t_0$ have an intersection of size $t_0 + 1$, the signatures in the conflicting ledgers prove the existence of a set G of $t_0 + 1$ Byzantine processes. \square

Lastly, we bound the message and communication complexity of the protocol. The number of rounds depends on when the network stabilizes (i.e., we cannot guarantee a decision for any consensus protocol prior to GST). However, we can bound the communication costs of each round:

Lemma 11 (Complexity). *Each round of The Polygraph Protocol has message complexity $O(n^2)$ and communication complexity $O(n^3)$.*

Proof. The BV-broadcast has message complexity $O(n^3)$, as each of the n processes echoes up to 2 (binary) BVAL messages, each of which is re-transmitted to all n processes. The communication complexity is at most $O(n^3)$, since each message may contain a ledger which contains $O(n)$ signatures. The remainder of the protocol involves only $O(n^2)$ messages and only $O(n^3)$ communication complexity (e.g., for the coordinator to broadcast its message, and for processes to send their ECHO messages). \square

C Multivalued Consensus

In this section, we discuss how to generalize the binary consensus to ensure accountable Byzantine agreement for arbitrary values. We follow the approach from [8]: First, all n processes use a reliable broadcast service to send their proposed value to all the other n processes. Then, all the process participate in parallel in n binary agreement instances, where each instance is associated with one of the processes. Lastly, if j is the smallest binary consensus instance to decide 1, then all the processes decide the value received from process p_j .

The key to making this work is that we need the reliable broadcast service to be accountable, that is, if it violates the reliable delivery guarantees, then each honest process has irrefutable proof of the culpability of $t + 1$ processes. Specifically, we want a single-use reliable broadcast service that allows each process to send one message, delivers at most one message from each process, and guarantees the following properties:

- *RB-Validity:* If an honest process RB-delivers a message m from an honest process p_j , then p_j RB-broadcasts m .
- *RB-Send:* If $t \leq t_0$ and p_j is honest and RB-broadcasts a message m , then all honest processes eventually RB-deliver m from p_j .
- *RB-Receive:* If $t \leq t_0$ and an honest process RB-delivers a message m from p_j (possibly faulty) then all honest processes eventually RB-deliver the same message m from p_j .
- *RB-Accountability:* If an honest process p_i RB-delivers a message m from p_j and some other honest process p_j RB-delivers m' from p_j , and if $m \neq m'$, then eventually every process has irrefutable proof of the culpability of $t_0 + 1$ processes.

The resulting algorithm provides a weaker notion of validity: if all processes are honest, then the decision value is one of the values proposed. (A stronger version

of validity could be achieved with a little more care, but is not needed for blockchain applications, which depend on an external validity condition.)

We first, in Section C.1, present the general multivalued algorithm, and prove that it is correct—assuming the existing of a reliable broadcast service satisfying the above properties. Then, in Section C.2, we describe the reliable broadcast service and prove that it is correct.

C.1 Accountable Byzantine Agreement

We now present the algorithm in more detail. The general algorithm has three phases.

- First, in lines 1–8, each process uses reliable broadcast to transmit its value to all the others. Then, whenever a process receives a reliable broadcast message from a process p_k , it proposes ‘1’ in binary consensus instance k . The first phase ends when there is at least one decision of ‘1’.
- Second, in lines 10–12, each process proposes ‘0’ in every remaining binary consensus instance for which it has not yet proposed a value. The second phase ends when every consensus instance decides.
- Third, in lines 14–16, each process identifies the smallest consensus instance j that has decided ‘1’. (If there is no such consensus instance, then it does not decide at all.) It then waits until it has received the reliable broadcast message from p_j and outputs that value.

First, we argue that it solves the consensus problem as long as $t \leq t_0$:

Lemma 12 (Agreement). *If $t \leq t_0$, then every honest process eventually decides the same value. If all processes are honest and propose the same value, that is the only possible decision.*

Proof. First we focus on termination. Since $t \leq t_0$, by the RB-Send property, we know that every honest process eventually delivers every value that was proposed by an honest process. Assume for the sake of contradiction that no binary consensus instance ever decides ‘1’. Then eventually, every honest process proposes ‘1’ for every binary consensus instance associated with an honest process. By the validity and termination properties of binary consensus (and since $t \leq t_0$), we conclude that these instances all decide ‘1’, which is a contradiction. Thus eventually every honest process executes line 10.

Since every honest process eventually proposes a value to every binary consensus instance, we conclude (since $t \leq t_0$) that eventually every binary consensus instance decides and every honest process reaches line 15. Let j be the minimum binary consensus instance that decides ‘0’. Assume (for the sake of contradiction) that

Algorithm 3 The General Polygraph Protocol

```

1: gen-propose( $v_i$ ):
2:   RB-broadcast( $\text{EST}, \langle v_i, i \rangle$ )  $\rightarrow$   $\text{messages}_i$  ▷ Reliable broadcast value to all.
3:
4:   repeat ▷ When you receive a value from  $p_k$ , begin consensus instance  $k$  with a proposal of 1.
5:     if  $\exists v, k : (\text{EST}, \langle v, k \rangle) \in \text{messages}_i$  then
6:       if  $\text{BIN-CONSENSUS}[k]$  not yet invoked then
7:          $\text{BIN-CONSENSUS}[k].\text{bin-propose}(1) \rightarrow \text{bin-decisions}[k]_i$ 
8:       until  $\exists k : \text{bin-decisions}[k] = 1$  ▷ Wait until the first decision.
9:
10:    for all  $k$  such that  $\text{BIN-CONSENSUS}[k]$  not yet invoked do ▷ Now begin consensus on the remaining instances.
11:       $\text{BIN-CONSENSUS}[k].\text{bin-propose}(0) \rightarrow \text{bin-decisions}[k]_i$  ▷ For these, propose 0.
12:    wait until for all  $k, \text{bin-decisions}[k] \neq \perp$  ▷ Wait until all the instances decide.
13:
14:     $j = \min\{k : \text{bin-decisions}[k] = 1\}$  ▷ Choose the smallest instance that decides 1.
15:    wait until  $\exists v : (\text{EST}, \langle v, j \rangle) \in \text{messages}_i$  ▷ Wait until you receive that value.
16:    decide  $v$  ▷ Return that value.

```

no honest process received the value from p_j . Then every honest process proposed ‘0’ to the binary consensus instance for j , and hence by the validity property, the decision would have been ‘0’, i.e., a contradiction. Thus we know that at least one honest process received the value that was reliably broadcast by p_j .

Finally, by the RB-Receive property, we know that every honest process must eventually deliver the value that was reliably broadcast by p_j , and hence every process eventually returns a value, i.e., we satisfy the termination property.

Next, we argue agreement. Since $t \leq t_0$, by the guarantees of the binary consensus instances, every honest process decides the same thing for each of the instances. Therefore, all honest processes will choose the same j that is the minimum binary consensus instance that decides 1. As we have already argued, every honest process must eventually deliver the value reliably broadcast by p_j , and that must be the same value. This guarantees agreement.

Finally, we argue validity: if all the processes are honest, then every value received by reliable broadcast is from an honest value, and hence validity is immediate. \square

Next, we prove that if there is any disagreement, then the algorithm guarantees accountability:

Lemma 13 (Accountability). *If two honest process p_i and p_j decide different values, then honest processes eventually receive irrefutable proof of at least $t_0 + 1$ Byzantine processes.*

Proof. First, assume that p_i and p_j decide different values for some binary consensus instance. Then, by the accountability of binary consensus, we know that every process eventually receives the desired irrefutable proof. Alternatively, if p_i and p_j agree for every instance of binary consensus, then they choose the same value

k that is the minimum binary consensus instance to decide ‘1’, and they output the value delivered by the reliable broadcast service from p_k . (Recall that the reliable broadcast service delivers only one value from each process, and p_i and p_j do not decide until they receive that value.) However, by the RB-accountability property, we conclude that eventually every process receives irrefutable proof of at least $t_0 + 1$ Byzantine processes. \square

C.2 Reliable Broadcast

We now describe the reliable broadcast service, which is a straightforward extension of the broadcast protocol proposed by Bracha [3]. A process begins by broadcasting its message to everyone. Every process that receives the message directly, echoes it, along with a signature. Every process that receives $n - t_0$ distinct ECHO messages, sends a READY message. And if a process receives $t_0 + 1$ distinct READY messages, it also sends a READY message. Finally, if a process receives $n - t_0$ distinct READY messages, then it delivers it.

The key difference from [3] is that, as in the binary value consensus protocol, we construct ledgers to justify the messages we send. Specifically, when a process sends a READY message, if it has received $n - t_0$ distinct ECHO messages, each of which is signed, it packages them into a ledger, and forwards that with its READY message. Alternatively, if a process sends a READY message because it received $t_0 + 1$ distinct READY messages, then it simply copies an existing (valid) ledger. Either way, if a process p_i sends a READY message for value v which was sent by process p_j , then it has stored a ledger containing $n - t_0$ signed ECHO messages for v , and it has sent that ledger to everyone.

As before, two ledgers conflict if they justify two different values v and v' , both supposedly sent by the same

process p_j . In that case, one ledger contains $n - t_0$ signed ECHO message for v and the other contains $n - t_0$ signed ECHO message for v' . Since any two sets of size $n - t_0$ have an intersection of size $t_0 + 1$, this immediately identifies at least $t_0 + 1$ processes that illegally sent ECHO messages for both v and v' . These process can by ir-refutably proved to be Byzantine.

We now prove that the reliable broadcast protocol satisfies the desired properties. First, we show that it delivers only one value from each process, and that if $t \leq t_0$, then it only delivers a value if it was previously RB-broadcast by that process:

Lemma 14 (RB-Unicity). *At most one value (v, j) is delivered from process p_j .*

Proof. Follows immediately by inspection: only one message from p_j is every delivered to each process.. \square

Lemma 15. (RB-Validity) *If $t \leq t_0$, and if an honest process RB-delivers a value v from an honest process p_j , then p_j RB-broadcasts m .*

Proof. A process only delivers a value v for p_j if it received $(\text{READY}, v, \cdot, j)$ messages from $n - t_0$ processes, implying that at least one honest process sent a READY message for v . Let p_i be the first honest process to send a $(\text{READY}, v, \cdot, j)$. In that case, we know that p_i must have received $n - t_0$ distinct (ECHO, v, j) , implying that at least one honest process sent an ECHO message for v . An honest process only sends an (ECHO, v, j) message if it received v directly from p_j . And if p_j is honest, it only sends v if the value was RB-broadcast. \square

Next, we prove a key lemma, showing that either all honest process send READY messages only for one value, or two conflicting ledgers are eventually received by all honest processes.

Lemma 16. *Assume that p_i is an honest process that sends a READY message for value v and that p_j is an honest process that sends a READY message for value v' . Then either $v = v'$ or the ledgers ℓ_i and ℓ_j constructed by p_i and p_j , respectively, conflict.*

Proof. Process p_i sends the READY message for v either because (i) it has received $n - t_0$ distinct messages (ECHO, v, \cdot) and has constructed a ledger ℓ_i containing those signed messages, or (ii) it has received $t_0 + 1$ distinct READY messages $(\text{ECHO}, v, \cdot, j)$, each containing a valid ledger for v , one of which it copies as ℓ_i . Either way, process p_i has a valid ledger ℓ_i containing $n - t_0$ signed echo messages for v . Similarly, by the same logic, process p_j has a valid ledger ℓ_j containing $n - t_0$ signed echo message for v' .

Since any two sets of size $n - t_0$ must have an intersection of size at least $t_0 + 1$, we conclude that if $v \neq v'$,

then the ledgers ℓ_i and ℓ_j conflict, i.e., prove that at least $t_0 + 1$ processes illegally sent ECHO messages for both v and v' . \square

We can now show that if an honest process performs a RB-broadcast, then as long as $t \leq t_0$, every honest process delivers its message:

Lemma 17. (RB-Send) *If $t \leq t_0$, and if p_j is honest and RB-broadcasts a value v , then all honest processes eventually RB-deliver v from p_j .*

Proof. If p_j is honest, then it broadcasts its value v , properly signed, to all processes. All honest processes receive it directly from p_j and immediately broadcast an ECHO message. (Moreover, there is no other message they could echo, because there is no other message they could have received directly from p_j .)

Since $t \leq t_0$, we know that there are at least $n - t_0$ honest processes that perform the ECHO, and hence every honest process receives at least $n - t_0$ ECHO messages, and hence every honest process broadcasts a READY message. (Of course there is no other message that an honest process could send a READY message for, since the first honest process to send a READY message for some other value must have received at least $n - t_0$ distinct ECHO messages for that value; at least one of those ECHO messages must have been sent by an honest process which received it directly from p_j , which—being honest—only sent value v .)

Since $t \leq t_0$, there are at least $n - t_0$ honest processes that send READY messages, and so every honest process receives $n - t_0$ READY messages and delivers the value v from p_j . (Of course p_j cannot have delivered any other values v' from p_j earlier, since p_j is honest there is no other value v' that it RB-broadcast.) \square

Next, we can show that if any honest process delivers a value v , then every honest process also delivers value v —as long as $t \leq t_0$.

Lemma 18. (RB-Receive) *If $t \leq t_0$ and an honest process p_k RB-delivers a value v from p_j (possibly faulty), then all honest processes eventually RB-deliver the same message v from p_j .*

Proof. Assume p_k delivers v from p_j . In this case, p_k must have received at least $n - t_0$ valid READY messages for (v, j) . Therefore, there must have been at least $t_0 + 1$ honest processes that broadcast valid READY messages. This implies that every honest process receives at least $t_0 + 1$ valid READY messages for (v, j) , and hence also sends a READY message for (v, j) . (Since $t \leq t_0$, we know that an honest process cannot send a READY message for any other value $v' \neq v$ for process j , by Lemma 16.) Therefore everyone honest process receives

Algorithm 4 Reliable Broadcast

```

1: RB-broadcast( $v_i$ ):
2:   broadcast(INITIAL,  $v_i$ )
3: upon receiving a message (INITIAL,  $v$ ) from  $p_j$ :
4:   broadcast(ECHO,  $v, j$ )
5: upon receiving  $n - t_0$  distinct messages (ECHO,  $v, j$ ) and not having sent a READY message:
6:   Construct a ledger  $\ell_i$  containing the  $n - t_0$  signed messages (ECHO,  $v, j$ ).
7:   broadcast(READY,  $v, \ell_i, j$ )
8: upon receiving  $t_0 + 1$  distinct messages (READY,  $v, \cdot, j$ ) and not having sent a READY message:
9:   Set  $\ell_i$  to be one of the (valid) ledgers received (READY,  $v, \ell, j$ ).
10:  broadcast(READY,  $v, \ell, j$ )
11: upon receiving  $n - t_0$  distinct messages (READY,  $v, \cdot, j$ ) and not having delivered a message from  $j$ :
12:   Let  $\ell$  be one of the (valid) ledgers received (READY,  $v, \ell, j$ ).
13:   deliver( $v, j$ )

```

▷ Only executed by the source
▷ Broadcast value v_i to all.

▷ Echo value v to all.

▷ Send READY message and ledger for v to all.

▷ Send READY message for v to all.

▷ Send READY message for v to all.

at least $2t_0 + 1$ valid READY message for value v for process j , and hence delivers value v from p_j . \square

Finally, we show the accountability property: either all honest processes deliver the same value, or two conflicting ledgers are received by every honest process:

Lemma 19. RB-Accountability: *If an honest process p_i RB-delivers value v from p_i and some other honest process p_j RB-delivers v' from p_j , and if $v \neq v'$, then eventually every honest process receives two ledgers ℓ and ℓ' that conflict.*

Proof. If process p_i delivers v from p_j , then it received at least $n - t_0$ READY messages for value v , which implies that at least one honest process p_u sent a READY message for v . Similarly, since p_j delivers v' from p_j , we know that at least one honest process p_w sent a READY message for v' . By Lemma 16, we know that if $v \neq v'$, then the ledgers ℓ_u and ℓ_w conflict. Moreover, since both p_u and p_w are honest, they sent their respective READY messages to all processes, and hence every honest process receives the conflicting ledgers. \square

D Application to Blockchain

In this section we explain how the General Polygraph Protocol can hold blockchain service providers accountable to blockchain client nodes that do not run the consensus as long as $t < \frac{n}{3}$. Note that a blockchain service can be implemented with a replicated state machine to which separate clients send requests. A predetermined set of n nodes, called a *consortium*, can propose and decide valid blocks that they append to their local view of the blockchain through the *General Polygraph Protocol* that accepts arbitrary values. A client can send a get requests to the members of the consortium and if it receives the same view of the blockchain from a certain number m of members ($m = (n - t_0)$ by default), it considers the transactions of this common view as valid.

In case of disagreement ($t > t_0$), the blockchain forks in that multiple blocks get appended to the same index of the chain, which could lead to *double spending* if the resulting branches have conflicting transactions.

Preliminaries. We now restate the existing blockchain formalism by Anceaume et al. [2]. A *blockchain* is a chain of blocks whose `score()` function takes as input a blockchain and returns its score s as a natural number, which can be its height, its weight, etc. A *blocktree* is a Mealy's machine whose states are countable, with an input alphabet comprising operation `append(block)` to append a block to the blocktree and operation `read()` that returns a blockchain, and an oracle Θ distributing permission tokens to processes for them to include a new block. The *blocktree strong (resp. eventual) consistency* is the conjunction of the following properties:

- *block validity*: each block in a blockchain returned by a `read()` operation is valid and has been inserted in the blocktree with the `append()` operation.
- *local monotonic read*: given a sequence of `read()` operations at the same process, the score of the returned blockchains never decreases.
- *ever growing tree*: given an infinite sequence of `append()` and `read()` operations, the score of the returned blockchains eventually grows.
- *strong (resp. eventual) prefix property*: for each blockchain returned by a `read()` operation with score s , then (resp. eventually) all the `read()` operations return blockchains sharing the same maximum common prefix of at least s blocks.

Now we consider three cases depending on the number of Byzantine processes, namely the nominal case, the degrading case and the zombie case when respectively $t \leq t_0$, $t_0 < t < (n - t_0)$ and $(n - t_0) \leq t \leq n$.

Nominal case ($t \leq t_0$). In this case the strong consistency is preserved. To read the state of a blockchain, a client asks the n members of the consortium (`read()` invocation) and waits for $m = n - t_0$ identical answers (`read()` response events). Indeed, if the assumption $t \leq t_0$ holds,

the consensus will finish (the ever growing tree property is ensured) and the client will eventually receive at least $n - t_0$ identical answers. But nothing prevents the Byzantine processes to stay mute forever or give false answers, that is why a client does not expect (a priori) more than $n - t_0$ answers. While $t \leq t_0$, the frugal $\Theta_{F,k=1}$ oracle manages tokens in a controlled way to guarantee that no more than $k = 1$ forks can occur on a given block (k -fork coherence), thus the consortium blockchain implements the blocktree strong consistency.

Degraded case ($t_0 < t < (n - t_0)$). In this case the ever-growing tree is violated but not the eventual prefix property. Moreover, if the threat of punishment (allowed by the accountability) disincentivizes a malicious coalition from attacking, then the strong prefix property is ensured. Let $t_0 < t < n - t_0$. A malicious coalition can do either one of these actions:

- Follow the protocol.
- Stay mute to violate the liveness property of the consensus and so the ever growing tree property of the blockchain (and so the (even eventual) consistency of the blockchain).
- Attempt an attack to create a disagreement among the consortium. Whatever the result of the bid, a proof of guilt will eventually be spread among all the honest processes (consortium members and clients). Then, regardless the sentence applied to the malicious nodes, a special fork can be created to drop the illegitimate forks (labelled in conse-

quence) due to the attack. The selection function returns then the unique blockchain which does not pass by a fork labelled illegitimate. The eventual prefix property is then satisfied. Moreover, if the potential punishment discourages any attempt of disagreement-attack (for example with a negative utility function in game theory approach), then the strong prefix property is ensured.

So, as long $t < (n - t_0)$, the eventual prefix property is ensured, but the ever-growing tree property is violated if $t > t_0$.

Zombie case ($(n - t_0) \leq t \leq n$). In this case a super coalition of $t > 2\frac{n}{3}$ Byzantine nodes can override the General Polygraph Protocol by proposing directly two conflicting views to two different clients to then perform a double-spending attack. The coalition does not participate to the consensus in order to violate the liveness property. It follows that the ever growing tree property is violated. Note that safety is also violated: When a client invokes the `read()` primitive, the coalition can answer arbitrary values, despite the non-termination of the legitimate consensus. The client is supposed to trust the coalition, like all the other clients who can forever receive a different output for the `read()` primitive. Hence, for $t \geq n - t_0$, the eventual prefix property is violated. This makes the blockchain vulnerable to a double-spending attack.