

Make Every Word Count: Adaptive Byzantine Agreement with Fewer Words

Shir Cohen ✉

Technion, Haifa, Israel

Idit Keidar ✉

Technion, Haifa, Israel

Alexander Spiegelman ✉

Aptos, San Francisco, CA, USA

Abstract

Byzantine Agreement (BA) is a key component in many distributed systems. While Dolev and Reischuk have proven a long time ago that quadratic communication complexity is necessary for worst-case runs, the question of what can be done in practically common runs with fewer failures remained open. In this paper we present the first Byzantine Broadcast algorithm with $O(n(f+1))$ communication complexity in a model with resilience of $n = 2t + 1$, where $0 \leq f \leq t$ is the actual number of process failures in a run. And for BA with strong unanimity, we present the first optimal-resilience algorithm that has linear communication complexity in the failure-free case and a quadratic cost otherwise.

2012 ACM Subject Classification Theory of computation → Distributed algorithms

Keywords and phrases Byzantine Agreement, Byzantine Broadcast, Adaptive communication

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2022.18

Funding *Shir Cohen*: Supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

1 Introduction

Byzantine Agreement (BA) is a key component in many distributed systems. As these systems are being used at larger scales, there is an increased need to find efficient solutions for BA. Arguably, the most important aspect of an efficient BA solution is its communication costs. That is, how much information needs to be transferred in the network to solve the BA problem. Indeed, improving the communication complexity, often measured as word complexity, was the focus of many recent works and deployed systems [1, 11, 16, 2, 13, 7].

In the BA problem, a set of n processes attempt to agree on a *decision* value despite the presence of Byzantine processes. One of the properties of a BA algorithm is a threshold t on how many Byzantine processes it can withstand. Namely, the algorithm is correct as long as up to t processes are corrupted in the course of a run. In this paper we focus on $n = 2t + 1$ and we assume a trusted setup of a public-key infrastructure (PKI) that enables us to use a threshold signature scheme [15, 4, 6].

A large and growing body of literature has investigated how to reduce the word complexity of BA algorithms. Recently, Momose and Ren [13] have presented a synchronous protocol with $O(n^2)$ words, which meets Dolev and Reischuk's long-standing lower bound [9]. Spiegelman [16] considered the more common case, where the number of actual failures, denoted by f , is smaller than t with resilience of $n = 3t + 1$. In this paper we consider better resilience and ask:

Can we design a BA protocol with $O(n(f+1))$ communication complexity in runs with $f \leq t$ failures, where $n = 2t + 1$?



© Shir Cohen, Idit Keidar, and Alexander Spiegelman;
licensed under Creative Commons License CC-BY 4.0


26th International Conference on Principles of Distributed Systems (OPODIS 2022).

Editors: Eshcar Hillel, Roberto Palmieri, and Etienne Rivière; Article No. 18; pp. 18:1–18:21



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Relation between various Byzantine Agreement solutions. Each box uses the primitives within it.

Whereas Dolev and Reischuk’s better-known lower bound applies to worst-case runs, they further proved a lower bound of $\Omega(nt)$ signatures in failure-free runs ($f = 0$) in a model with a PKI. At the time, one could have thought that this bound extends to the communication complexity, rendering it $\Omega(nt)$ even with small f values. However, the introduction of threshold signature schemes [8, 15, 4, 6] exposed the possibility to compact many signatures into one word, potentially saving many words.

In this paper, we first revisit the original problem as stated in Dolev and Reischuk’s work. In this problem there is a single sender who proposes a value and we refer to this problem as Byzantine Broadcast (BB). We prove that although $O(nt)$ signatures are inevitable, $O(nt)$ messages are not necessary with $f \in o(t)$ failures by presenting an adaptive BB solution with $O(n(f + 1))$ words.

The idea behind our algorithm is to reduce this problem to another BA variant. There is a simple reduction from BB to BA with the strong unanimity validity property (from hereon: *strong BA*), which states that if all correct processes propose the same value, this is the only allowed decision. In this reduction, the sender initially sends its value to all other processes who then run a BA solution. Unfortunately, no *adaptive strong BA* is known to date. I.e., a strong BA solution where communication complexity depends on f , rather than on t . Instead, in Section 5 we reduce the problem to a new *weak BA* problem with a weaker validity property, *unique validity*, which we define in this paper.

Intuitively, the validity condition of weak BA is somewhere between weak unanimity, where if all processes are correct and propose the same value this is the only allowed decision, and external validity [5], where a decision value must satisfy some external predicate. In weak BA, one can define its desired predicate and the requirement is that if all correct processes propose the same value and Byzantine processes cannot devise a value that satisfies the chosen predicate, then the decision must be valid. Otherwise, \perp is allowed.

While the unique validity condition seems to be weak, it is surprisingly powerful when provided the “right” external predicate. For example, we can determine that a value is valid if it has at least $t + 1$ unique signatures, assuring that some correct process in the system knows this value. Unique validity may be of independent interest as a tool for designing algorithms. We present our adaptive weak BA in Section 6. The weak BA, in turn, exploits the quadratic solution by Momose and Ren [13]. Figure 1 describes the relation between the various solutions.

Finally, we consider strong BA. In Section 7, we present the first optimally resilient strong binary BA protocol with $O(n)$ communication complexity in the failure-free case. This leaves open the question whether a fully adaptive (to any f) strong BA protocol exists. We summarize the results in Table 1.

■ **Table 1** Bounds on communication complexity of deterministic synchronous Byzantine Agreement algorithms with resilience $n = 2t + 1$.

	Upper Bound	Lower Bound
Byzantine Broadcast	$O(n(f + 1))$ Section 5 + Section 6	$\Omega(nf)$ ($\Omega(n^2)$ signatures) [9]
Strong BA	$O(n^2)$ multi-valued Momose-Ren [13] $O(n)$ with $f = 0$, binary Section 7	$\Omega(nf)$ binary ($\Omega(n^2)$ signatures) [9]
Weak BA	$O(n(f + 1))$ multi-valued Section 6	$\Omega(n)$

2 Model and Preliminaries

We consider a distributed system consisting of a well-known static set Π of n processes and an adaptive adversary. The adversary may adaptively corrupt up to $t < n$, $n = 2t + 1$ processes in the course of a run. A corrupted process is *Byzantine*; it may deviate arbitrarily from the protocol. In particular, it may crash, fail to send or receive messages, and send arbitrary messages. As long as a process is not corrupted by the adversary, it is *correct* and follows the protocol. We denote by $0 \leq f \leq t$ the actual number of corrupted processes in a run.

Cryptographic tools. We assume a trusted public-key infrastructure (PKI) and a computationally bounded adversary. Hence, we can construct and use a threshold signature scheme [15, 4, 6]. We denote by $\langle m \rangle_p$ the message m signed by process p . Using a (k, n) -threshold signature scheme, k unique signatures on the same message m can be batched into a threshold signature for m with the same length as an individual signature. For simplicity we abstract away the details of cryptography and assume the threshold signature schemes are ideal. In practice, our results hold except with arbitrarily small probability, depending on the security parameters.

Communication. Every pair of processes is connected via a reliable link. If a correct process p_i receives a message m indicating that m was sent by a correct process p_j , then m was indeed generated by p_j and sent to p_i . The network is synchronous. Namely, there is a known bound δ on message delays, allowing us to design protocols that proceed in rounds. Specifically, if a correct process sends a message to any other correct process at the beginning of some round, it is received by the end of the same round.

Complexity. We use the following standard complexity notions [2, 16, 13]. While measuring complexity, we say that a *word* contains a constant number of signatures and values from a finite domain, and each message contains at least 1 word. The communication complexity of a protocol is the maximum number of words sent by all correct processes, across all runs. The *adaptive* complexity is a complexity that depends on f .

3 Problem Definitions

We consider a family of agreement problems all satisfy agreement and termination defined as follows:

Agreement No two correct processes decide differently.

Termination Every correct process eventually decides.

In addition, each variant of the problem satisfies some validity property. In the Byzantine Broadcast (BB) problem, a designated sender has an input to broadcast to all n processes. The goal is that all correct processes decide upon the sender's value. If the sender is Byzantine, however, it is enough that all correct processes decide upon some common value. Formally,

► **Definition 1** (Byzantine Broadcast). *In Byzantine Broadcast, a designated sender has an input value v_{sender} to broadcast to all processes, and each correct process decides on an output value decision_i . BB solution must satisfy agreement, termination and the following validity property:*

Validity *If sender is correct, then all correct processes decide v_{sender} .*

Byzantine Agreement (BA) is a closely related problem to BB. In this problem, a set Π of n processes each propose an initial value and they all attempt to reach a common decision. In addition, the decided value must be “valid” in some sense that makes the problem non-trivial. The classic notion of validity states that if all correct processes in Π share the same initial value, then the decision must be on this value. This property is known as *strong unanimity*, and it entails a limitation on the resilience of a protocol, requiring that $n \geq 2t + 1$. For hereon we refer to BA with strong unanimity validity condition as *strong BA*. Formally,

► **Definition 2** (Strong Byzantine Agreement). *In Byzantine Agreement, each correct process $p_i \in \Pi$ proposes an input value v_i and decides on an output value decision_i . Any strong BA solution must satisfy agreement, termination and the following validity property:*

Strong unanimity *If all correct processes propose the same value v , then the output is v .*

A different validity property requires that a decision satisfies some external boolean predicate (we call such value a *valid* value). It is used under the assumption that all correct processes propose valid values. This is known as *external validity* [5] and only requires $n > t$. External validity by itself is trivial in case there is a well-known predefined value that satisfies the predicate. However, it is commonly used in settings with signatures, where valid values can be verified by all but generated only by specific users or sets thereof. For instance, consider a predicate that verifies that v is signed by $n - t$ processes – no process can unilaterally generate a default valid value.

Our notion of unique validity adopts external validity to allow default values to be decided in cases when there is no unanimous valid value. We say that a value v *exists* in a run of a BA protocol if v is either the input value of a correct process or can be generated by a Byzantine process. E.g., any value signed by a non-Byzantine process cannot be generated locally by a Byzantine process. Unique validity stipulates that there is a default value if and only if there exists more than one valid value in a BA run. Formally,

► **Definition 3** (Weak Byzantine Agreement). *In weak Byzantine Agreement, each correct process $p_i \in \Pi$ proposes an input value v_i and decides on an output value decision_i . Any weak BA solution must satisfy agreement, termination and the following validity property:*

Unique Validity *Assume an arbitrary predicate $\text{validate}(v) \in \{\text{true}, \text{false}\}$ that can be computed locally. If a correct process decides v then either $v = \perp$ or $\text{validate}(v) = \text{true}$, and if $v = \perp$ then more than one valid value exists in the run.*

As the definition suggests, unique validity is satisfied in weak BA with respect to any chosen external predicate. This allows for the application level to determine the desired properties, and choose the relevant external predicate accordingly. As a simple example, one can think of a predicate that specifies that a value is valid if it is signed by at least $t + 1$ processes stating that this value was their initial value. In this scenario, unique validity yields exactly the common strong unanimity property on the underlying signed values.

In fact, unique validity is a useful tool when designing distributed algorithms as it allows to use BA as a framework. Different applications may require different validity conditions, yet still unique validity prevents the system from having a trivial solution in the presence of Byzantine processes. Note, in addition, that every solution to BA with external validity property immediately solves weak BA.

4 Related Work

The starting point of this work goes back to 1985 when Dolev and Reischuk proved two significant lower bounds for the Byzantine Broadcast problem. Specifically, they have studied the worst-case message complexity over all runs and proved it to be $\Omega(nt)$. Moreover, in the authenticated model, which was somewhat undeveloped at the time, they proved a lower bound of $\Omega(nt)$ signatures – even in a failure-free run.

Since the publication of their fundamental results, the paradigm of complexity measurement has shifted. The number of messages is of little importance nowadays, compared to the number of words it entails. The total number of words (the communication complexity) better reflects the load on the system and is commonly used today when analyzing distributed algorithms. For example, Dolev and Reischuk presented in their paper a BB algorithm that matches their messages' lower bound. It requires $O(nt)$ messages, but as a single message can be composed of many different signatures it requires a cubic number of words. It was not until recently that a solution with quadratic communication complexity was presented for synchronous BA with optimal resilience [13].

Dolev and Reischuk's complementary lower bound on signatures does not translate to a bound on the communication complexity of an algorithm. Only a few years after Dolev and Reischuk's work, the threshold signature scheme was introduced [8]. This scheme allows multiple signatures to be compacted into a single combined signature of the same size. That is, a single word can carry multiple signatures. In this work, we focus on the communication complexity of the BB and BA problems while taking advantage of such schemes.

To make our algorithms efficient in real-world systems, we adjust the complexity to match the actual number of faults. Moreover, we do so without compromising the worst-case complexity. If all t possibly Byzantine processes crash, the complexity of our algorithms is $O(nt)$. However, in most runs, where systems do not exhibit the worst crash patterns, the complexity is much lower. In fact, it is linear in the number of faults times n .

While consensus algorithms were designed to be adaptive in the number of failures over 30 years ago [10], these works focus on the number of rounds that it takes to reach a decision rather than on communication complexity. A special case of adaptivity is focusing on failure-free runs. This problem was addressed both by Amdur et al. [3] (only for crash failures) and by Hadzilacos and Halpern [12]. However, both works measure the number of messages rather than words and have sub-optimal communication complexity.

A recent work by Spiegelman [16] tackled the problem of adaptive communication complexity in the asynchronous model. It presents a protocol that achieves correctness in asynchronous runs and requires $O(ft + t)$ communication in synchronous runs. However, due to the need to tolerate asynchrony, its resilience is only $n \geq 3t + 1$. This solution relies on threshold signatures schemes, as we do.

As noted also by Momose and Ren [13], designing optimally-resilient protocols for the synchronous model limits the use of threshold signatures. While this primitive has been used in various eventually synchronous and asynchronous works over the last few years [1, 16, 14, 6], usually with a threshold of $n - t$. Using this threshold in settings with resilience $n = 3t + 1$, we get certificates signed by at least $t + 1$ correct processes. However, for a resilience of $n = 2t + 1$, this is no longer the case. The threshold signatures “lose” their power as $n - t = t + 1$ for which no intersection properties between correct processes signing two distinct certificates can be derived. In this work, we exploit threshold signatures with this improved resilience by carefully choosing a better threshold for our needs, as we discuss in Section 6. We mention that although not using threshold schemes, Xiang et al. [17] also benefit from collecting more than $n - t$ signatures in some scenarios.

5 From Weak BA to Adaptive Byzantine Broadcast

In this section we study the BB problem, and optimize its adaptive communication complexity over all runs. We present a new BB protocol with resilience $n = 2t + 1$ and adaptive communication complexity of $O(n(f + 1))$.

Recall that in the BB problem there is only one sender who aims to broadcast its initial value and have all correct processes agree on it. If the sender is Byzantine, it may attempt to cause disagreement across correct processes. There is a known simple and efficient reduction from BB to strong BA. Given a strong BA solution, the designated sender starts by sending its value to all processes, and then they all execute the BA solution and decide on its output. It is easy to see that if the sender is correct, all correct processes begin the strong BA algorithm with the same input, and by strong unanimity they then decide upon the sender's value.

However, trying to apply the same reduction from BB to weak BA no longer works. If the sender is Byzantine, the correct processes do not have a valid initial value for the BA. Nonetheless, in this section we present a reduction from BB to weak BA¹, which incurs a cost of $O(n(f + 1))$ words. Thus, together with an adaptive weak BA with the same complexity, we obtain a synchronous adaptive BB algorithm with a total of $O(n(f + 1))$ words and resilience $n = 2t + 1$. At this point we assume that such adaptive weak BA is given as a black box. An implementation for this primitive is presented in Section 6.

■ **Algorithm 1** BB algorithm: code for process p_i , sender's input is v_{sender} .

Initially $v_i, val, decision, ba_decision = \perp$

Round 1:

```

1: if  $sender = p_i$  then
2:   send  $\langle v_{sender} \rangle_{sender}$  to all
3: if received message  $\langle v \rangle_{sender}$  from  $sender$  then
4:    $v_i \leftarrow \langle v \rangle_{sender}$ 
5: for  $j = 1$  to  $n$  do
6:    $val \leftarrow invokePhase(j, v_i)$ 
7:   if  $val \neq \perp$  then
8:      $v_i \leftarrow val$ 
9:  $ba\_decision \leftarrow$  weak BA with  $BB\_valid$  predicate and initial value  $v_i$ 
10: if  $ba\_decision$  is of the form  $\langle v \rangle_{sender}$  then
11:    $decision \leftarrow v$ 
12: else
13:    $decision \leftarrow \perp$ 

```

Our algorithm, presented in Algorithms 1 and 2, is composed of three parts. The first part (lines 1 – 4 in Algorithm 1) is the first round in which the leader disseminates its value. Processes that receive that value adopt it as their BA initial value (line 4). The second part (lines 5 – 8 in Algorithm 1 and Algorithm 2) is a “vetting” part. It consists of n phases, with a rotating leader. Leaders initiate phases to learn about the first part's initial value. Finally, the third part (lines 9 – 13 in Algorithm 1) is a weak BA execution.

¹ This reduction only works if $n \geq 2t + 1$.

Deciding upon the weak BA output takes care of the agreement and termination properties. It is left to (1) satisfy the BB validity property and (2) make sure that the preconditions for the weak BA hold, that is, each correct process has a valid input to propose. To achieve these properties, we define the $BB_valid(v)$ predicate in the following way. $BB_valid(v) = true$ if and only if v is signed by either the sender or by $t + 1$ processes.

Note that if the sender happens to be Byzantine, it is acceptable to decide on any value. However, it is important to make sure that if the sender is correct, then the only valid value is its initial BB input. Simply setting a value to be valid only if it is signed by the sender would not work, as it allows a faulty sender to cause a scenario in which there are no valid values to agree upon by not sending its value to any process. Note that we cannot simply fix this by introducing some default valid value: If we were to do so, it would be valid to agree on that value also in the case of a correct sender, violating the BB validity condition.

Our algorithm makes sure that if the sender is correct, the second condition in the BB_valid definition cannot be satisfied, and hence there is only one possible outcome to the BA algorithm. However, if the sender is Byzantine, it is guaranteed that there is some value to decide upon. That is, all correct processes start the weak BA with an initial value that satisfies the predicate.

In the vetting part of the algorithm, we ensure that the above-mentioned conditions hold. Moreover, we do so with a communication complexity that is adaptive to the number of actual process failures. The core idea is to work in leader-based *phases*. Every phase has a unique leader and is composed of a constant number of leader-to-all and all-to-leader synchronous rounds. Every phase is initiated by a leader-to-all message. If the leader decides not to send the initial message then no messages will be sent by correct processes in this phase and we say that this phase is *silent*, and otherwise, it is *non-silent*. In our algorithm, a phase is non-silent if the phase's leader did not choose an initial value for the BA prior to that phase.

In every phase, each process p_i starts the phase with some initial value v_i and if the phase is non-silent it returns some value. The requirements from the phase are: (1) If the phase's leader is correct and the phase is non-silent, then all correct processes return a valid value. (2) All correct processes return either \perp or a valid v . And (3) if the sender is correct, then no correct process returns a value signed by $t + 1$ processes.

Upon a non-silent phase, the leader starts by asking all processes for help by sending a `help_req` message (line 16). A correct process that receives a help request message answers the leader. If it has set a BA initial value, it sends it to the leader at line 19, and otherwise, it sends a signed `idk` (i don't know) message at line 21. If the leader receives a value signed by the designated sender it broadcasts it (line 24). Otherwise, if it receives $t + 1$ `idk` messages, it uses a threshold signature scheme to create an `idk` quorum certificate and broadcasts it (line 27). A process that receives from the leader a value signed by either the sender or any $t + 1$ processes returns it. Otherwise, it returns \perp .

At the end of each non-silent phase, a correct process that returns a $v \neq \perp$ from the phase, updates its local v_i accordingly at line 8. This value at the end of the n^{th} phase is the input for the weak BA algorithm. Since we execute n phases, all correct processes set valid values by the end of all phases. This is because once there is a correct process that did not set a value it initiates its phase and then all correct processes return with a valid value. At this point, all processes execute the weak BA and decide upon its output (line 9).

A formal correctness proof of Algorithms 1 and 2 appears in Appendix A, proving the following theorem:

► **Theorem 4.** *Algorithm 1 solves BB.*

■ **Algorithm 2** *invokePhase*(j, v_i): code for process p_i .

```

14:  $leader \leftarrow p_{j \bmod n}$ 
    Round 1:
15: if  $leader = p_i$  and  $v_i = \perp$  then
16:   broadcast the message  $\langle \text{help\_req}, j \rangle_{leader}$ 
    Round 2:
17: if received  $\langle \text{help\_req}, j \rangle_{leader}$  then
18:   if  $v_i \neq \perp$  then
19:     send  $\langle v_i, j \rangle$  to  $leader$ 
20:   else
21:     send  $\langle \text{idk}, j \rangle_{p_i}$  to  $leader$ 
    Round 3:
22: if  $leader = p_i$  then
23:   if received  $\langle v', j \rangle$  s.t.  $v' = \langle v \rangle_{sender}$  then
24:     broadcast the message  $\langle \langle v \rangle_{sender}, j \rangle$ 
25:   else if received  $t + 1$  unique signatures  $\langle \text{idk}, j \rangle_{p'}$  then
26:     batch these messages into  $QC_{\text{idk}}$  using a  $(t + 1, n)$ -threshold signature scheme
27:     broadcast the message  $\langle QC_{\text{idk}}, j \rangle$ 
28: if received  $\langle v, j \rangle$  from  $leader$  and  $BB\_valid(v) = true$  then
29:   return  $v$ 
30: else
31:   return  $\perp$ 

```

Complexity

We prove that the complexity of Algorithms 1 and 2 is $O(n(f + 1))$.

Each non-silent phase is composed of a constant number of all-to-leader and leader-to-all rounds and thanks to the use of threshold signatures, all messages sent have a size of one word. Thus, each phase incurs $O(n)$ words. In total, there are potentially n phases. However, we prove in Appendix A that after the first non-silent phase by a correct leader, all following phases with correct leaders are silent. Thus, the number of non-silent phases is linear in f . We conclude that all phases in lines 5 – 8 use $O(n(f + 1))$ words. The complexity of the weak BA black box is also $O(n(f + 1))$ (as we will show in the next section), resulting in a total of $O(n(f + 1))$ words.

6 Adaptive Weak BA

In this section, we present a synchronous adaptive weak BA algorithm with resilience $n = 2t + 1$. This algorithm is the missing link for the adaptive BB presented in the previous section. Once again, we use the concept of phases and exploit the pattern of possible *silent* phases. In this algorithm, the phases are slightly different and the decision to start a phase as a leader depends on whether or not the leader has reached a decision in previous phases.

Unlike the BB problem, in BA every process begins the algorithm with its own input value. Communication-efficient solutions to this problem usually employ threshold signatures schemes [1, 16]. This technique is widely used in asynchronous and eventually synchronous protocols, with resilience $n = 3t + 1$. In these contexts, one can use a scheme of $(n - t)$ -out-of- n signatures and benefit from the fact that any two such quorum certificates intersect by at least $t + 1$ processes, and therefore at least one correct process.

Unfortunately, when trying to apply the same technique to a system with resilience $n = 2t + 1$, it fails. A correct process might be unable to obtain $2t + 1$ unique signatures on any value as Byzantine processes might not sign it. On the other hand, a quorum certificate with only $t + 1$ unique signatures is not very useful as it does not guarantee the desired intersection property.

Our first key observation is that the intersection property can be achieved as long as we have $\lceil \frac{n+t+1}{2} \rceil$ unique signatures. If we obtain this number of signatures out of $n = 2t + 1$, safety is preserved in the sense that conflicting certificates cannot be formed by a malicious adversary. Of course, there are runs in which we cannot reach that threshold since $\lceil \frac{n+t+1}{2} \rceil > n - t$ (e.g., if t processes crash immediately as the run begins). But in this case, $f \geq \frac{t}{2}$, and $O(f)$ becomes asymptotically $O(t)$. Hence, we can use a fallback algorithm with $O(nt)$ communication complexity.

As we assume that $t \in \Theta(n)$, we can use Momose and Ren's synchronous algorithm that has $O(n^2)$ communication complexity [13] for the fallback. We denote that algorithm $\mathcal{A}_{fallback}$. Note that their algorithm is "stronger" than our proposed algorithm as it provides strong unanimity for validity (i.e., it solves strong BA). We can use their solution by checking the validity of $\mathcal{A}_{fallback}$'s output according to the predicate. If it is valid, this is the decision value, and otherwise a default valid value is decided. Equipped with these insights, we next present our algorithm.

During the phases part of the protocol, a correct process must commit a value before reaching a decision. When it has certainty about a value it updates that value in a *commit* variable, along a *commit_proof* of this commitment (a quorum certificate, signed by sufficiently many processes). Once a correct process commits to a certain value it does not commit to any other value during the run. It may, however, decide on another value eventually. For example, if it reaches the fallback and no correct process has decided. Once a correct process reaches a decision it updates it in its local *decision* variable as well as a matching quorum certificate in *decide_proof* variable.

A single phase. The code for a single phase is given in Algorithm 4. Each process p_i starts a phase with its initial value v_i and information about possible previous commits (*commit*, *commit_proof*) and decisions (*decision*, *decide_proof*). Correct processes return with updated information about commits and decisions that were made in that phase (or prior to that). The guarantees of the phases are: (1) Every *decision* updated during a phase is valid; (2) All decisions updated by correct processes are the same and there exists at most one valid *decide_proof* in the system; and (3) If the phase's leader is correct, the phase is non-silent, and $n - f > \lceil \frac{n+t+1}{2} \rceil$, then all correct processes return with the same valid decision.

Every non-silent phase starts with the leader broadcasting a **propose** message with its value in line 32. Upon receiving this message, correct processes either vote for this value by signing it (line 34) or answer with a value that was previously committed as well as its commit quorum certificate (line 36). If the leader receives a committed value it simply broadcasts it. Otherwise, if it manages to achieve the required $\lceil \frac{n+t+1}{2} \rceil$ threshold of signatures, it can form a quorum certificate committing its proposed value (line 40).

Note that at this point the committed value is not "safe enough" to be decided by correct processes. Byzantine leaders may cause correct processes to participate in forming a commit certificate for more than one value. As correct processes that have decided do not initiate phases, they might never communicate without going through Byzantine leaders. Thus,

18:10 Make Every Word Count: Adaptive Byzantine Agreement with Fewer Words

■ **Algorithm 3** weak BA algorithm: code for process p_i with initial value v_i .

Initially $decision = undecided$, $bu_decision = v_i$, $fallback_start \leftarrow \infty$
 $decide_proof$, $commit$, $commit_proof$, bu_proof , $fallback_val$, $phase_decision = \perp$

- 1: **for** $j = 1$ to $t + 1$ **do**
- 2: $phase_decision$, $decide_proof$, $commit$, $commit_proof \leftarrow$
 $invokePhase(j, v_i, decision, commit, commit_proof)$
- 3: **if** $decision = undecided$ and $phase_decision \neq undecided$ **then**
- 4: $decision \leftarrow phase_decision$
- Round 1:**
- 5: **if** $decision = undecided$ **then**
- 6: broadcast $\langle help_req \rangle_{p_i}$
- Round 2:**
- 7: **if** received $\langle help_req \rangle_{p'}$ message and $decision \neq undecided$ **then**
- 8: send $\langle help, decision, decide_proof \rangle_{p_i}$ to p'
- 9: **if** received $t + 1$ messages of $\langle help_req \rangle_{p'}$ from different processes **then**
- 10: batch these messages into $QC_{fallback}(v)$ using a $(t + 1, n)$ -threshold signature scheme
- 11: broadcast the message $\langle fallback, QC_{fallback}, decision, proof \rangle_{p_i}$
- 12: $fallback_start \leftarrow now + 2\delta$
- Round 3:**
- 13: **if** received $\langle help, v, decide_proof \rangle_{p'}$ with valid v and $decide_proof$ for v and $decision =$
 $undecided$ **then**
- 14: $decision \leftarrow v$
- 15: $bu_decision \leftarrow decision$
- 16: **while** $fallback_start > now$ **do**
- 17: **if** received valid $\langle fallback, QC_{fallback}, v, proof_{p'} \rangle_{p'}$ **then**
- 18: **if** $decision = undecided$ and $proof_{p'} \neq \perp$ is a valid proof for a valid v **then**
- 19: $bu_decision \leftarrow v$
- 20: $bu_proof \leftarrow proof_{p'}$
- 21: **if** $fallback_start = \infty$ **then**
- 22: broadcast the message $\langle fallback, QC_{fallback}, bu_decision, bu_proof \rangle_{p_i}$
- 23: $fallback_start \leftarrow now + 2\delta$
- 24: $fallback_val \leftarrow A_{fallback}$ with $\delta' = 2\delta$ and initial value $bu_decision$
- 25: **if** $decision = undecided$ **then**
- 26: **if** $fallback_val$ is valid **then**
- 27: $decision \leftarrow fallback_val$
- 28: **else**
- 29: $decision \leftarrow \perp$

we need another level of certainty, in the form of the finalize certificate (to be stored in $decide_proof$). We maintain the invariant that if a correct process receives a valid finalize certificate, then no finalize certificate on another value can be formed.

Thus, after a correct process learns about a committed certificate it sends a matching **decide** message to the leader (line 44). In addition, if this is the first commit certificate it receives, it commits to it. If the leader receives the necessary threshold of **decide** messages, it forms a finalize quorum certificate. Every process that receives such a certificate can safely return the certificate's value as its decision.

■ **Algorithm 4** *invokePhase*($j, v_i, decision, decide_proof, commit, commit_proof$): code for process p_i .

```

30:  $leader \leftarrow p_{j \bmod n}$ 
    Round 1:
31: if  $leader = p_i$  and  $decision = \perp$  then
32:   broadcast the message  $\langle \text{propose}, v_i, j \rangle_{leader}$ 
    Round 2:
33: if received  $\langle \text{propose}, v, j \rangle_{leader}$  with a valid  $v$  for the first time and  $commit = \perp$  then
34:   send  $\langle \text{vote}, v, j \rangle_{p_i}$  to  $leader$ 
35: else if received  $\langle \text{propose}, v, j \rangle_{leader}$  and  $commit \neq \perp$  then
36:   send  $\langle \text{commit}, commit, commit\_proof, j \rangle_{p_i}$  to  $leader$ 
    Round 3:
37: if  $leader = p_i$  then
38:   if received  $\langle \text{commit}, w, QC_{\text{commit}}(w), j \rangle'_p$  then
39:     broadcast the message  $\langle \text{commit}, w, QC_{\text{commit}}(w), j \rangle_{leader}$ 
40:   else if received  $\lceil \frac{n+t+1}{2} \rceil$  messages of  $\langle \text{vote}, v, j \rangle_{p'}$  then
41:     batch these messages into  $QC_{\text{commit}}(v)$  using a  $(\lceil \frac{n+t+1}{2} \rceil, n)$ -threshold signature
    scheme
42:     broadcast the message  $\langle \text{commit}, v, QC_{\text{commit}}(v), j \rangle_{leader}$ 
    Round 4:
43: if received  $\langle \text{commit}, v, QC_{\text{commit}}(v), j \rangle_{leader}$  then
44:   send  $\langle \text{decide}, v, j \rangle_{p_i}$  to  $leader$ 
45:   if  $commit = \perp$  then
46:      $commit \leftarrow v$ 
47:      $commit\_proof \leftarrow QC_{\text{commit}}(v)$ 
    Round 5:
48: if  $leader = p_i$  then
49:   if received  $\lceil \frac{n+t+1}{2} \rceil$  messages of  $\langle \text{decide}, v, j \rangle_{p'}$  then
50:     batch these messages into  $QC_{\text{finalized}}(v)$  using a  $(\lceil \frac{n+t+1}{2} \rceil, n)$ -threshold signature
    scheme
51:     broadcast the message  $\langle \text{finalized}, v, QC_{\text{finalized}}(v), j \rangle_{leader}$ 
52: if received  $\langle \text{finalized}, v, QC_{\text{finalized}}(v), j \rangle_{leader}$  then
53:    $decision \leftarrow v$ 
54:    $decide\_proof \leftarrow QC_{\text{finalized}}(v)$ 
55: return  $(decision, decide\_proof, commit, commit\_proof)$ 

```

Main algorithm. The BA algorithm is given in Algorithm 3, using the phase algorithm as a building block. In our algorithm, all correct processes eventually decide by updating their *decision* variable. However, they do not halt. In our BA algorithm, we start by executing n phases with a rotating leader, ensuring that every correct process has a chance to reach a decision before executing the fallback algorithm. After the phases end there are several possibilities. First, if there are at most $\frac{n-t-1}{2}$ Byzantine processes, all correct processes must have decided. If there are more Byzantine processes, it may be the case that some correct processes decided and others did not. This could happen, for example, if a Byzantine leader causes the single correct leader to decide and not initiate its phase. By the phase guarantees, we know that all correct processes that decide by this point, decide the same valid value.

18:12 Make Every Word Count: Adaptive Byzantine Agreement with Fewer Words

To address the case where not all correct processes decided, we have processes that have not decided ask for help from all other processes (line 6). If a correct process has decided and receives a `help_req` message, it answers with a help message including the decision value along with its proof at line 8. Note that in this round, the number of messages sent by correct processes is linear in the number of help requests. Specifically, if only Byzantine processes send `help_req` messages, the number of answers is $O(nf)$ and independent of t .

We note that if $t + 1$ help requests are sent, then at least one of them is sent by a correct process that did not manage to form quorum certificates when it served as leader. Thus, in this case, $f \in \Theta(t)$, and we can execute the fallback algorithm. To make sure that all correct processes participate in the fallback algorithm, a `fallback` certificate with $t + 1$ signature is formed.

We now encounter a new challenge. We must have all correct processes start a synchronous fallback algorithm at the same time. However, an adversary can form the `fallback` certificate and deal it to only some correct processes. This scenario can happen, for example, if less than $t + 1$ `help_req` messages are sent, and the adversary adds t `help_req` signatures of its own. We thus require a correct process that receives a `fallback` certificate to broadcast it (line 22). This ensures that whenever one correct process runs the fallback algorithm, all of them do, but may still cause different correct processes to start the fallback at different times. Nevertheless, we know that the starting time difference is at most the δ it takes the message to arrive. We therefore run the fallback algorithm with $\delta' = 2\delta$, ensuring that all correct processes enter a fallback round before any of them exits from it.

Another subtle point is making sure that the fallback algorithm does not output a decision value that contradicts previous decisions made by correct processes. For that reason we add another 2δ safety window between getting notified about a fallback and initiating it. Correct processes that broadcast the `fallback` certificate attach their decision value and a proof (if exists). In the 2δ safety window, processes that learn about a decision value in the system adopt it as the initial value for the fallback algorithm (line 17). Recall that $\mathcal{A}_{fallback}$ is a strong BA protocol. If a correct process decides v prior to the fallback algorithm, all other correct processes learn about v during the safety window. Then, by strong unanimity, they all decide v .

Note that if the decision returned from $\mathcal{A}_{fallback}$ is not valid then it must be that strong unanimity preconditions are not satisfied (since correct processes always have valid inputs) and a default value is returned. Furthermore, whenever the strong unanimity precondition is not satisfied, it follows that not all correct processes propose the same value. As a result, there must exist more than one valid value in the run (the different correct proposals). And the \perp default value is a valid weak BA output.

A formal correctness proof of Algorithms 3 and 4 appears in Appendix B, proving the following theorem:

► **Theorem 5.** *Algorithm 3 solves weak BA.*

Complexity

We show that if $f < \frac{n-t-1}{2}$, correct processes never perform the fallback algorithm.

► **Lemma 6.** *If $f < \frac{n-t-1}{2}$, correct processes never perform the fallback algorithm.*

Proof. In Appendix B we prove that if a correct process is the leader of a non-silent phase and $f < \frac{n-t-1}{2}$, then all correct processes return the same valid decision. Since Algorithm 3 is composed of n phases, every correct process has a chance to invoke its phase and all correct processes decide by line 4. Assume by way of contradiction that there exists a correct

process that invokes the fallback algorithm. By the code, it has received a **fallback** certificate. However, such certificate can only be formed by $t + 1$ unique **help_req** signatures, meaning that at least one correct process sent a **help_req** message. But this is impossible if all correct processes decide by line 4. ◀

Each phase is composed of a constant number of all-to-leader and leader-to-all rounds. Thus, it incurs $O(n)$ words. Potentially, there are n phases. However, a lemma in Appendix B proves that once a correct leader invokes *invokePhase()* and the number of actual failures is $f < \frac{n-t-1}{2}$, all correct processes decide by the end of that phase. Since correct leaders that had already decided do not invoke their phases (their phases are silent), the number of invoked phases depends on f itself. Thus, all phases combined send $O(n(f + 1))$ words.

After n *invokePhase* invocations end, help request messages are sent only by correct processes that did not decide. By the above-mentioned lemma, it happens only if $f > \frac{n-t-1}{2}$. In this case, $f = \Theta(n)$ and since $t = \Theta(n)$ it holds that $O(nf) = O(n^2)$. Correct processes that decide by this point answer directly to whoever sent them help requests, without affecting the asymptotic complexity. If some correct process receives a **fallback** certificate, another all-to-all round is added, keeping the complexity $O(n^2)$. All other communication costs are incurred in the fallback algorithm, whose complexity is also $O(n^2)$.

7 strong BA: the failure free case

Recall that the optimal resilience for strong BA is $n = 2t + 1$. In this section we present a binary strong BA protocol that has communication complexity of $O(n)$ in the failure free case. Otherwise, it has complexity $O(n^2)$. The question of whether an adaptive protocol with $O(n(f + 1))$ complexity can be designed for strong BA with optimal resilience remains open.

In the algorithm, presented in Algorithm 5, a single leader first collects all initial values. Since we solve binary agreement, in the failure-free case there must be a value proposed by $t + 1$ different processes. Thus, the leader can use a threshold signature scheme to aggregate a quorum certificate on this proposed value.

As a second step, the leader sends this certificate to all processes and attempts to collect n different signatures on the value. If it succeeds, it broadcasts it. Every process that receives a signed-by-all certificate can safely decide upon its value. If a correct process does not decide, it broadcasts a fallback message. Every process that hears such a message, echoes it at most once, and execute $\mathcal{A}_{fallback}$ after 2δ time with 2δ -long rounds, as in Section 6. In Appendix C we prove the correctness of Algorithm 5. That is, we prove the following theorem:

▶ **Theorem 7.** *Algorithm 5 solves binary strong BA.*

Complexity

We show that if the run is failure-free, correct processes never perform the fallback algorithm.

▶ **Lemma 8.** *If $f = 0$, correct processes never perform the fallback algorithm.*

Proof. If all processes are correct then they all send their initial values to the leader at line 2. Since values are binary, and there are $n = 2t + 1$ processes, there must be a value v such that the leader receives $t + 1$ unique signatures on v . Then, the leader broadcasts a **propose** certificate on v (line 6). Every correct process that receives this certificate replies with a signed **decide** message at line 8. Since all processes are correct, the leader then receives n signatures and then broadcasts a **decide** certificate on v (line 12). All processes then receive this certificate and decide v at line 14. None of them sends a fallback message. ◀

18:14 Make Every Word Count: Adaptive Byzantine Agreement with Fewer Words

■ **Algorithm 5** strong BA algorithm: code for process p_i with initial value v_i .

Initially $decision, proof, bu_decision, bu_proof, fallback_val = \perp$
 $fallback_start \leftarrow \infty$

1: $leader \leftarrow p_1$
Round 1:
2: send $\langle v_i \rangle_{p_i}$ to $leader$
Round 2:
3: **if** $leader = p_i$ **then**
4: **if** received $t + 1$ messages of $\langle v \rangle_{p'}$ for some v **then**
5: batch these messages into $QC_{\text{propose}}(v)$ using a $(t + 1, n)$ -threshold signature scheme
6: broadcast the message $\langle \text{propose}, v, QC_{\text{propose}}(v) \rangle_{leader}$
Round 3:
7: **if** received valid $\langle \text{propose}, v, QC_{\text{propose}}(v) \rangle_{leader}$ **then**
8: send $\langle \text{decide}, v \rangle_{p_i}$ to $leader$
Round 4:
9: **if** $leader = p_i$ **then**
10: **if** received n messages of $\langle \text{decide}, v \rangle_{p'}$ **then**
11: batch these messages into $QC_{\text{decide}}(v)$ using a (n, n) -threshold signature scheme
12: broadcast the message $\langle \text{decide}, v, QC_{\text{decide}}(v) \rangle_{leader}$
Round 5:
13: **if** received valid $\langle \text{decide}, v, QC_{\text{decide}}(v) \rangle_{leader}$ and $decision = \perp$ **then**
14: $decision \leftarrow v$
15: $proof \leftarrow QC_{\text{decide}}(v)$
16: **else**
17: broadcast the message $\langle \text{fallback}, \perp, \perp \rangle_{p_i}$
18: $fallback_start \leftarrow now + 2\delta$
19: $bu_decision \leftarrow decision$
20: **while** $fallback_start > now$ **do**
21: **if** received $\langle \text{fallback}, v, proof_{p'} \rangle_{p'}$ **then**
22: **if** $decision = \perp$ and $proof_{p'} \neq \perp$ is a valid proof for a valid v **then**
23: $bu_decision \leftarrow v$
24: $bu_proof \leftarrow proof_{p'}$
25: **if** $fallback_start = \infty$ **then**
26: broadcast the message $\langle \text{fallback}, bu_decision, bu_proof \rangle_{p_i}$
27: $fallback_start \leftarrow now + 2\delta$
28: $fallback_val \leftarrow \mathcal{A}_{\text{fallback}}$ with $\delta' = 2\delta$ and initial value $bu_decision$
29: **if** $decision = \perp$ **then**
30: $decision \leftarrow fallback_val$

By Lemma 8, if all processes are correct then they never perform the fallback algorithm, and there are 4 all-to-leader and leader-to-all rounds, with a total of $O(n)$ words. Otherwise, the complexity is the complexity of the fallback algorithm, which is $O(n^2)$.

8 Conclusions and Future Directions

We have presented solutions for both Byzantine Broadcast and weak Byzantine Agreement with adaptive communication complexity of $O(n(f + 1))$ and resilience $n = 2t + 1$. To construct the weak BA algorithm, we utilized a threshold on the number of signatures such that on one hand, this number is sufficient to ensure a safe algorithm with adaptive communication in case there are not “many” Byzantine processes. On the other hand, failing to achieve this threshold indicates that there is a high number of failures, which allows the use of a quadratic fallback algorithm.

This weak BA algorithm is taken as a black box to construct our adaptive BB algorithm. Here, we carefully choose the predicate for the validity property, to allow us to reduce one problem to the other. Finally, for strong BA we propose a binary solution with optimal resilience. Our solution is linear in n in the practically common failure-free case, and quadratic in any other case. The question of whether a fully adaptive strong BA with optimal resilience exists or not remains open.

While $n = 2t + 1$ is optimal for strong BA, this is not the case for BB and weak BA, where any $t < n$ can be tolerated². Thus, another possible future direction is improving the resilience of an adaptive BB or adaptive weak BA to support any $t < n$. Our weak BA algorithm relies on the current resilience to satisfy that if $f > n - \lceil \frac{n+t+1}{2} \rceil$ then f is linear in t . Note that this remains true for any resilience of $n = \alpha t + \beta$, for $\alpha > 1, \beta > 0$ without compromising the intersection property required for safety. Should a quadratic solution for weak BA be developed, it could be used to improve the total resilience of our adaptive algorithm (instead of Momose and Ren’s algorithm [13]).

References

- 1 Ittai Abraham, Guy Golan-Gueta, and Dahlia Malkhi. Hot-stuff the linear, optimal-resilience, one-message bft devil. *CoRR*, abs/1803.05069, 2018. [arXiv:1803.05069](https://arxiv.org/abs/1803.05069).
- 2 Ittai Abraham, Dahlia Malkhi, and Alexander Spiegelman. Asymptotically optimal validated asynchronous byzantine agreement. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 337–346, 2019.
- 3 Eugene S Amdur, Samuel M Weber, and Vassos Hadzilacos. On the message complexity of binary byzantine agreement under crash failures. *Distributed Computing*, 5(4):175–186, 1992.
- 4 Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the weil pairing. In *International conference on the theory and application of cryptography and information security*, pages 514–532. Springer, 2001.
- 5 Christian Cachin, Klaus Kursawe, Frank Petzold, and Victor Shoup. Secure and efficient asynchronous broadcast protocols. In *Annual International Cryptology Conference*, pages 524–541. Springer, 2001.
- 6 Christian Cachin, Klaus Kursawe, and Victor Shoup. Random oracles in constantinople: Practical asynchronous byzantine agreement using cryptography. *Journal of Cryptology*, 18(3):219–246, 2005.
- 7 Shir Cohen, Idit Keidar, and Alexander Spiegelman. Not a coincidence: Sub-quadratic asynchronous byzantine agreement whp. In *34th International Symposium on Distributed Computing (DISC 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- 8 Yvo Desmedt. Society and group oriented cryptography: A new concept. In *Conference on the Theory and Application of Cryptographic Techniques*, pages 120–127. Springer, 1987.

² For weak BA, this stems from the resilience for external validity.

18:16 Make Every Word Count: Adaptive Byzantine Agreement with Fewer Words

- 9 Danny Dolev and Rüdiger Reischuk. Bounds on information exchange for byzantine agreement. *Journal of the ACM (JACM)*, 32(1):191–204, 1985.
- 10 Danny Dolev, Ruediger Reischuk, and H Raymond Strong. Early stopping in byzantine agreement. *Journal of the ACM (JACM)*, 37(4):720–741, 1990.
- 11 Yossi Gilad, Rotem Hemo, Silvio Micali, Georgios Vlachos, and Nickolai Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. In *Proceedings of the 26th Symposium on Operating Systems Principles, SOSP '17*, New York, NY, USA, 2017. ACM. doi:10.1145/3132747.3132757.
- 12 Vassos Hadzilacos and Joseph Y Halpern. Message-optimal protocols for byzantine agreement. *Mathematical systems theory*, 26(1):41–102, 1993.
- 13 Atsuki Momose and Ling Ren. Optimal communication complexity of authenticated byzantine agreement. In *35th International Symposium on Distributed Computing (DISC 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.
- 14 Oded Naor and Idit Keidar. Expected linear round synchronization: The missing link for linear byzantine smr. In *34th International Symposium on Distributed Computing (DISC 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- 15 Victor Shoup. Practical threshold signatures. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 207–220. Springer, 2000.
- 16 Alexander Spiegelman. In search for an optimal authenticated byzantine agreement. In *35th International Symposium on Distributed Computing (DISC 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.
- 17 Zhuolun Xiang, Dahlia Malkhi, Kartik Nayak, and Ling Ren. Strengthened fault tolerance in byzantine fault tolerant replication. In *2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS)*, pages 205–215. IEEE, 2021.

A Adaptive Byzantine Broadcast: Correctness

A.1 Correctness

We start by proving the phase’s requirements. First, immediately from lines 29 – 31 we get that all correct processes return either \perp or a valid v . Next, the following lemma shows that in non-silent phases with correct leaders all correct processes return a valid value.

► **Lemma 9.** *If a phase is non-silent and its leader is correct, then all correct processes return a valid value.*

Proof. If the leader is correct it broadcasts a `help_req` message at line 16. All correct processes then answer at round 2. If the leader receives a value signed by the sender at line 23, it broadcasts it at line 24. Otherwise, no correct processes received a value signed by the sender and sends an `idk` message at line 21. Since $n = 2t + 1$, the leader receives at least $t + 1$ `idk` messages (from the correct processes) and forms an `idk` certificate. It broadcasts this value at line 27. In both cases, all correct processes return a valid value at line 29. ◀

The next lemma proves that if all correct processes invoke a phase with a value other than \perp , then they can return only one type of a valid value – a value signed by the sender.

► **Lemma 10.** *If all correct processes invoke a phase with value $v \neq \perp$, there does not exist a value signed by $t + 1$ processes in the system.*

Proof. If all correct processes invoke a phase with value $v \neq \perp$, they reply to the `help_req` messages at line 19 and never send an `idk` message. Since there are at most t Byzantine processes, the leader cannot receive $t + 1$ `idk` messages and form an `idk` certificate signed by $t + 1$ different processes. ◀

We now prove the correctness of the BB algorithm. First, to be able to use the weak BA, all correct processes must execute it with valid initial values.

► **Lemma 11.** *All correct processes execute line 9 with a valid initial value.*

Proof. Let p_i be a correct process. In n phases, there is one phase with p_i as leader. If p_i has updated v_i prior to that phase, it happened either line 4 or at line 8. Immediately from the code we get that in both cases p_i updates a valid value. If p_i did not update a value, it initiates a non-silent phase, and by Lemma 9 returns a valid value. ◀

Note that agreement and termination stem immediately from the code and the correctness of the weak BA. The following lemma proves validity.

► **Lemma 12.** *If sender is correct, then all correct processes decide v_{sender} .*

Proof. If *sender* is correct then all correct process learn v_{sender} by the end of round 1 and update their values at line 4. By Lemma 10, in no phase can any process create a value signed by $t + 1$ processes. Thus, when executing the weak BA v_{sender} signed by the sender is the only valid value that exists in the run. By unique validity and since the *sender* does not sign more than one initial value, v_{sender} is the only possible BA output. It follows that all correct processes execute line 11 and return the sender's value. ◀

We conclude the following theorem:

► **Theorem 13.** *Algorithm 1 solves BB.*

B Weak BA: Correctness

We start by proving some lemmas about the phase's guarantees. First, we prove that if the *decision* is updated in a given phase, then its new value is valid.

► **Lemma 14.** *If a correct process updates decision during `invokePhase`, then v is a valid decision value.*

Proof. If a correct process updates its *decision* value at line 53 of `invokePhase` then it must have received a finalized certificate signed by $\lceil \frac{n+t+1}{2} \rceil$ processes. Hence, at least one correct process p' signed the decide message for v at line 44. By the code, p' signed the decide message for v if it received a commit certificate signed by $\lceil \frac{n+t+1}{2} \rceil$ processes. Hence, at least one correct process p'' signed the vote message for v at line 34. By the code, this is possible only if v is a valid value (line 33). ◀

Next, we prove that all correct processes that update their *decision* variable do so the same value. Moreover, at most one valid *decide_proof* can exist in the system. That is, a Byzantine process cannot devise a *decide_proof* that conflicts with any other *decide_proof* known by correct processes.

► **Lemma 15.** *All correct processes that update decision during `invokePhase` return the same decision. In addition, at most one finalize certificate can be formed in all phases.*

Proof. Note that every correct process that updates *decision* sets its at line 53. Assume that a correct process p_i sets its decision value to v in phase l and a correct process p_j sets its decision value to w in phase $k \geq l$.

18:18 Make Every Word Count: Adaptive Byzantine Agreement with Fewer Words

If $k = l$, then p_i and p_j set their decision value in the same round and they both receive a finalize certificate signed by $\lceil \frac{n+t+1}{2} \rceil$ different processes. At least one correct process signed both certificates and since correct processes sign at most one finalize message per phase, $v = w$.

For the case where $k > l$: in phase l , p_i receives a finalize certificate signed by $\lceil \frac{n+t+1}{2} \rceil$ different processes. Thus, at least $\lceil \frac{n+t+1}{2} \rceil - t \geq \frac{n-t+1}{2}$ correct processes updated their commit to v in that phase. These processes do not vote for any value in following rounds. Thus at most $n - t - \frac{n-t+1}{2} = \frac{n-t-1}{2}$ correct processes can sign a conflicting value. Since $\frac{n-t-1}{2} + t < \lceil \frac{n+t+1}{2} \rceil$, in any phase greater than l , no process can collect $\lceil \frac{n+t+1}{2} \rceil$ signatures on any value other than v . Hence, no process can send a commit message on $w \neq v$ with a valid commit certificate. For the same reason, no process can form and send a valid finalize certificate and decide upon any other value. Thus, $v = w$. ◀

We prove next that once a correct process is the leader of a non-silent phase, all correct processes return the same valid decision value by the end of that phase.

► **Lemma 16.** *If a correct leader invokes `invokePhase` in phase k and $f < \frac{n-t-1}{2}$, then all correct processes return the same valid decision by the end of the phase and this decision is a proposal of a correct process.*

Proof. The leader broadcasts its value v to all processes. If there is a correct process p for which `commit` $\neq \perp$, it sends the message $\langle \text{commit}, w, \text{proof}, j \rangle_p$ to the leader. If the leader receives $\langle \text{commit}, w, \text{proof}, j \rangle_{p'}$ (from any process), it broadcasts in round 3 a commit certificate for w . Otherwise, since $f < \frac{n-t-1}{2}$, leader receives $\lceil \frac{n+t+1}{2} \rceil$ messages voting for v and broadcasts a commit certificate for v . Then, all correct processes send the leader a finalize messages on v or w . Again, the leader receives $\lceil \frac{n+t+1}{2} \rceil$ messages finalizing v and broadcasts a finalize certificate for v . Correct processes receive this message and update their `decision` and `decide_proof` accordingly. Then, by the code they all return v . ◀

We now prove the correctness of the main BA algorithm. The following two lemmas prove that although some processes may start executing $\mathcal{A}_{\text{fallback}}$ at different times, they all successfully execute the fallback algorithm.

► **Lemma 17.** *If some correct process executes the fallback algorithm in Algorithm 3, all correct process do so and they all start at most δ time apart.*

Proof. Let p be the first correct process that executes the fallback algorithm at line 24 of Algorithm 3 at time t . This means that at time $t - 2\delta$, p broadcasts the fallback certificate to all other processes (line 22). By synchrony, this certificate is guaranteed to arrive at all correct processes by $t - \delta$, causing them to execute the fallback algorithm by $t + \delta$ if they have not done so earlier. ◀

► **Lemma 18.** *Consider a synchronous algorithm \mathcal{A} . Let σ be a synchronized run of \mathcal{A} defined as follows. Let t be the time that the first correct process starts executing \mathcal{A} in σ . All correct processes start executing \mathcal{A} by $t + \delta$. The round duration is 2δ . In round r that begins (locally) in t_r , round r messages are processed if they are received in the time window $[t_r - \delta, t_r + 2\delta]$. Then σ is a correct run of \mathcal{A} .*

Proof. Consider a process p that starts round r at time t_r^p . Let p' be another correct process that starts round r at time $t_r^{p'}$, and sends a message to p in round r . By assumption, $t_r^{p'} = t_r^p + \epsilon$ where $-\delta \leq \epsilon \leq \delta$, and a message sent by p' at $t_r^{p'}$ arrives at time t_a where $t_r^{p'} \leq t_a \leq t_r^{p'} + \delta$. Note that round r ends at p at time $t_{r+1}^p = t_r^p + 2\delta$. Hence, $t_r^p - \delta \leq t_a \leq t_r^p + 2\delta$, as needed. ◀

Next, the following lemma states that if a correct process manages to reach a decision prior to the fallback algorithm, then this is the only possible decision. Moreover, this decision value must be a valid one.

► **Lemma 19.** *If some correct process decides v before executing the fallback algorithm, then all correct processes decide v and v is valid.*

Proof. If there exists a correct process p that decides at line 4, then by Lemma 15 and the code all processes that decide at line 4 decide v as well. Moreover, all other correct processes that have not decided by line 5, send `help_req` messages. Process p answers them and they all decide at line 14. Otherwise, no correct process decides at line 4 and they all send `help_req` messages at line 6. Then, they all receive $t + 1$ `help` messages and by the code perform the fallback algorithm. In addition, by the lemma assumption, it must be that p decides v at line 14.

If correct processes execute the fallback algorithm, then by the code they all wait a time period of 2δ before the execution, during which they receive all decisions made by other correct processes and update `bu_decision` accordingly (line 24). Specifically, they receive v from p . It follows from Lemma 15 that `bu_decision` is updated with the same value at all correct processes. Thus, all correct processes execute $\mathcal{A}_{fallback}$ with the same input, and by strong unanimity they set `fallback_val` to v at line 24.

We now prove that v is valid. If p decides v at line 4, then it must have updated `decision` in the scope of the relevant phase. By Lemma 14 this value is valid. Otherwise, if p decides v at line 14, then the validity follows from the code. Hence, since v is valid, all correct processes decide it by line 27. ◀

Finally, we are ready to prove the required BA properties.

► **Lemma 20 (Agreement).** *In Algorithm 3 all correct processes decide on the same value.*

Proof. First, by Lemma 15, all correct processes that decide in line 4 decide the same value v . In addition, it follows from the same lemma that every correct process that decides at line 14 after receiving a valid `finalize` certificate decides v , as at most one `finalize` certificate can be formed.

It is left to show that if not all correct processes decide before the fallback algorithm at line 24, they still decide upon the same value. If at least one correct process p receives a fallback certificate it follows from Lemma 17 that all correct processes receive the certificate within at most δ time of p . Then, by the code, all correct processes execute the fallback algorithm at line 24 and by Lemma 18 and the fallback algorithm solves strong BA, providing agreement. By Lemma 19, we get that processes that decide before running the fallback decide on the same value. ◀

► **Lemma 21 (Termination).** *In Algorithm 3 all correct processes decide.*

Proof. If not all correct processes decide before line 5 and no correct process receives a fallback certificate, it follows that less than $t + 1$ correct processes broadcast `help` messages at line 6. Hence, at least one correct process p has decided by line 5. Process p receives all of the correct `help` messages at line 7 and answers them at line 8. All correct processes that asked for `help` then decide at line 14.

It remains to examine the case that at least one correct process p receives a fallback certificate. It follows from Lemma 17 that all correct processes receive the certificate within at most δ time of p . Then, by the code, all correct processes execute the fallback algorithm at line 24 and by Lemma 18 and the fallback algorithm solves BA, providing termination. ◀

18:20 Make Every Word Count: Adaptive Byzantine Agreement with Fewer Words

► **Lemma 22** (Unique Validity). *In Algorithm 3 if a correct process decides v then either $v = \perp$ or $\text{validate}(v) = \text{true}$, and if $v = \perp$ then more than one valid value exists in the run.*

Proof. Let v be the decision value of a correct process in Algorithm 3. First, by lines 27 – 29 $\text{validate}(v) = \text{true}$ or $v = \perp$. We prove that if $v = \perp$, then at least two valid values exist in the run.

By the code, all processes execute the fallback algorithm with valid inputs (either their initial valid values, or a valid value they adopt at line 19). By strong unanimity of $\mathcal{A}_{\text{fallback}}$, if all correct processes start with the same valid value v' , then v' must be the returned decision value. This contradicts the fact that \perp is returned at line 29. Therefore, not all correct processes execute $\mathcal{A}_{\text{fallback}}$ with the same value. As they all execute the fallback algorithm with valid inputs, it follows that at least two valid values exist in the run. ◀

In addition, we need to prove that every correct process updates its *decision* at most once.

► **Lemma 23.** *In Algorithm 3 all correct process decide at most once.*

Proof. Any correct process updates *decision* at line 4, line 14 or lines 27 – 29. In all cases, it only does so if *decision* = *undecided*. Since by the code it does not update *decision* to the value *undecided*, it follows that *decision* is updated at most once. ◀

From Lemmas 20, 21, 22, and 23 we conclude:

► **Theorem 24.** *Algorithm 3 solves weak BA.*

C Strong BA: Correctness

► **Lemma 25.** *If some correct process executes the fallback algorithm in Algorithm 5, all correct process do so and they all start at at most δ time apart.*

Proof is similar to Lemma 17 in Section 6.

► **Lemma 26** (Agreement). *In Algorithm 5 all correct process decide on the same value.*

Proof. First, as correct processes only sign one *decide* message, every process that receives $QC_{\text{decide}}(v)$ receives the same quorum certificate. Thus, all correct processes that decide at line 14 decide the same v . If at least one correct process receives a fallback message then by Lemma 25, they all execute the fallback algorithm at most δ time apart. Thus, if at least one correct process decides at line 14, then all correct processes that have not yet decided learn about v in the 2δ safety window, and adopt it as their initial value for the fallback (line 23). It follows that all correct processes decide with the same input value v and by strong unanimity this is the only possible decision. ◀

► **Lemma 27** (Termination). *In Algorithm 5 all correct process decide.*

Proof. If not all correct processes decide by line 14, then a correct process broadcasts a fallback message at line 17. It follows from Lemma 25 that all correct processes receive the certificate within at most δ time of p . Then, by the code, all correct process execute the fallback algorithm at line 28 and by Lemma 18 and the fallback algorithm solves strong BA, providing termination. ◀

► **Lemma 28** (Validity). *In Algorithm 5 if all correct processes propose the same value v , then the output is v .*

Proof. Correct processes only send **decide** messages on values with valid **propose** quorum certificates. Note that such a quorum certificate can only be formed with $t + 1$ unique signatures. Hence, if all correct processes propose the same value v , then the only possible **propose** quorum certificate is with v . As a result, the only possible **decide** quorum certificate is with v as well.

The fallback algorithm is executed with either the original initial values or with a value that has a corresponding **decide** quorum certificate. Thus, if correct processes execute the fallback algorithm, they all start with v and by strong unanimity of $\mathcal{A}_{fallback}$, the decision is v . ◀

Finally, we prove that every correct process updates its *decision* at most once.

▶ **Lemma 29.** *In Algorithm 5 all correct process decide at most once.*

Proof. Any correct process updates *decision* either at line 14 or at line 30. In both cases, it only does so if *decision* = \perp . Since it does not update *decision* to the value \perp at any step of the algorithm, it follows that *decision* is updated at most once. ◀

From Lemmas 26, 27, 28, and 29 we conclude:

▶ **Theorem 30.** *Algorithm 5 solves binary strong BA.*