

# Byzantine Reliable Broadcast with Low Communication and Time Complexity

THOMAS LOCHER, DFINITY, Switzerland

Byzantine reliable broadcast is a fundamental problem in distributed computing, which has been studied extensively over the past decades. State-of-the-art algorithms are predominantly based on the approach to share encoded fragments of the broadcast message, yielding an asymptotically optimal communication complexity when the message size exceeds the network size, a condition frequently encountered in practice. However, algorithms following the standard coding approach incur an overhead factor of at least 3, which can already be a burden for bandwidth-constrained applications. Minimizing this overhead is an important objective with immediate benefits to protocols that use a reliable broadcast routine as a building block.

This paper introduces a novel mechanism to lower the communication and computational complexity. Two algorithms are presented that employ this mechanism to reliably broadcast messages in an asynchronous network where less than a third of all nodes are Byzantine. The first algorithm reduces the overhead factor to 2 and has a time complexity of 3 if the sender is honest, whereas the second algorithm attains an optimal time complexity of 2 with the same overhead factor in the absence of equivocation. Moreover, an optimization for real-world implementations is proposed, reducing the overhead factor to  $3/2$  under normal operation. Lastly, a lower bound is proved that an overhead factor lower than  $3/2$  cannot be achieved for a relevant class of reliable broadcast algorithms.

CCS Concepts: • **Computer systems organization** → **Fault-tolerant network topologies**; • **Theory of computation** → **Distributed algorithms**.

Additional Key Words and Phrases: asynchronous networks, reliable broadcast, communication complexity

## 1 INTRODUCTION

The goal of Byzantine reliable broadcast is to disseminate a message efficiently and reliably despite the presence of Byzantine nodes that may interfere with the protocol execution in arbitrary ways. A reliable broadcast routine is a powerful primitive with a broad range of applications including asynchronous atomic broadcast [15, 18, 19, 21, 27], distributed key generation [2, 12, 20], secure data replication [8], and secret sharing [26]. Moreover, Byzantine reliable broadcast plays a pivotal role in Byzantine fault-tolerant consensus protocols where message dissemination is treated separately from message ordering to improve throughput [10].

The first Byzantine reliable broadcast algorithm is due to Bracha [6]. While it is elegant in its simplicity, its main drawback is that the message  $m$  is broadcast by every node during the execution of the algorithm, i.e.,  $O(|m|n^2)$  bits are sent overall, where  $|m|$  denotes the size of  $m$  in bits and  $n$  is the number of nodes in the network. In the seminal paper by Cachin and Tessaro [9], this bound was improved to  $O(|m|n + \kappa n^2 \log(n))$  bits using *erasure coding*, where  $\kappa$  is the output size of a collision-resistant hash function. Since each node must receive the message in a successful broadcast, this result is asymptotically optimal if  $|m| \in \Omega(\kappa n \log(n))$ . Subsequent work has primarily focused on getting closer to the lower bound of  $\Omega(|m|n + n^2)$  [1, 4, 11, 22, 23]. These pieces of work follow the blueprint laid out by Cachin and Tessaro and augment it with error-correction and cryptographic primitives to achieve better asymptotic bounds.

Curiously, there is little work on minimizing the constant in the  $O(|m|n)$  term, although it likely dominates the actual bandwidth consumption in practice and is therefore crucial for real-world applications. The design by Cachin and Tessaro is based on a  $(n, t+1)$ -erasure code, where  $t < n/3$  is the largest number of nodes that may exhibit Byzantine behavior. In this design, nodes broadcast encoded *fragments* of size  $|m|/(t+1) \approx 3|m|/n$  instead of broadcasting  $m$ , sending at least  $3|m|/n$




Fig. 1. Given an input of size 1MiB, encoding and decoding times for erasure codes with parameters  $(n, t+1)$  and  $(n, 2t+1)$  are shown for different network sizes. The circles and bars mark the averages and the 5th and 95th percentiles, respectively, of the measured encoding and decoding times over 1000 runs.

bits altogether. An overhead factor of 3—compared to the ideal scenario where each node receives exactly  $|m|$  bits—may already prove too costly for applications with strict bandwidth constraints. All schemes that follow this approach naturally inherit the overhead inherent in this design.

In this paper, a mechanism is introduced that departs from this blueprint in order to reduce the communication complexity, i.e., the number of bits that need to be exchanged in the worst case, for large messages. The core idea is quite simple: a  $(n, 2t+1)$ -erasure code is used instead, reducing the fragment size to  $|m|/(2t+1) \approx \frac{3}{2}|m|/n$ . This modification obviously reduces the cost of broadcasting fragments; however, correctness is no longer given without additional changes because some nodes may receive  $2t+1$  fragments whereas others merely obtain  $t+1$  fragment, which is not enough to recover the message. This problem is addressed by introducing the following step. If a node with  $2t+1$  fragments manages to reconstruct the message, it disseminates fragments again but only to the  $t$  nodes from which no fragment was received. As we will see, this step is sufficient to ensure correctness while keeping the communication complexity low.

As an added benefit, the transition from a  $(n, t+1)$ -erasure code to a  $(n, 2t+1)$ -erasure code results in substantial performance improvements, both with respect to the encoding and decoding of fragments. Figure 1 shows performance numbers when encoding/decoding a message of size 1MiB for various network sizes and the two different parameterizations. The experiments were conducted on standard hardware using an optimized library<sup>1</sup> for Reed-Solomon codes [24]. The figure confirms the practicality of erasure coding due to its efficiency, exhibiting encoding and decoding times in the millisecond range, and scalability with respect to the network size. Moreover, the figure reveals that the encoding (decoding) time is lower by a factor of 2.2-2.4 (1.8-2.0) for all network sizes when using an  $(n, 2t+1)$ -erasure code. Another interesting property is that the variance is also significantly lower, leading to more predictable encoding and decoding times.

The paper is organized as follows. The model is formally introduced in §2. A concrete algorithm that makes use of the novel mechanism is presented in §3. Its communication complexity is  $2|m|n$  plus a term independent of  $|m|$ . Given an *honest* (i.e., non-faulty) sender, each honest node obtains the message after 3 rounds of communication. A natural question is whether there is an algorithm with an overhead factor lower than 3 and that terminates in an optimal 2 rounds given an honest sender. In §4, this question is answered in the affirmative by presenting an algorithm

<sup>1</sup>See <https://github.com/AndersTrier/reed-solomon-simd>.

with this property and the same overhead factor if the sender does not equivocate. Otherwise, the worst-case overhead factor is  $\frac{5}{2}$  with respect to the largest fragment size transmitted by honest nodes. While the first algorithm only requires collision-resistant hash functions, the second algorithm makes use of threshold signatures. Another aspect of practical relevance is the (worst-case) storage requirement at each node. Both algorithms bound the required space to a small multiple of a predetermined maximum message size. Practical considerations with the goal of reducing the communication complexity further in real-world deployments and simplifying the implementation are discussed in §5. A key observation is that the overhead can be reduced to  $\frac{3}{2}$  during periods of synchrony and in the absence of failures. In §6, a lower bound is presented, showing that an overhead factor of  $\frac{3}{2}$  is optimal for a particular class of algorithms. Related work is summarized in §7 before the paper concludes in §8.

## 2 MODEL

The considered network comprises  $n = 3t + 1$  nodes, where  $t$  denotes the maximum number of *Byzantine nodes*, which may deviate arbitrarily from any given protocol. The other  $2t + 1$  nodes faithfully execute the given protocol and are called *honest*. Any two nodes can communicate directly over an authenticated channel by exchanging messages. Communication is assumed to be *asynchronous* in the sense that the delivery of messages, while guaranteed, may be delayed indefinitely. In this communication model, nodes cannot make any assumption about message delays, neither about the time required until sent messages arrive at their destinations nor about the time that elapses until certain (expected) messages are received.

A node may send a message to a single node or to multiple nodes. When a node  $v$  sends a message to all nodes, we say that  $v$  *broadcasts* this message. A faulty node may fail to send messages as specified in the protocol. The goal of *reliable broadcast* is to ensure, under some conditions, that all honest nodes eventually *deliver* a certain message, i.e., mark it as the accepted outcome of the reliable broadcast. The required properties of a reliable broadcast protocol are stated formally in the following definition.

*Definition 2.1 (Reliable broadcast).* A *reliable broadcast protocol* is a distributed protocol to send a message  $m$  from a specific node called the *sender* to all nodes with the following properties.

- **Validity:** If the sender is honest and broadcasts message  $m$ , then every honest node eventually delivers  $m$ .
- **Agreement:** If two honest nodes deliver messages  $m$  and  $m'$ , then  $m = m'$ .
- **Integrity:** Every honest node delivers at most one message  $m$ .
- **Totality:** If an honest node delivers message  $m$ , then all honest nodes eventually deliver  $m$ .

In general, nodes may engage in multiple reliable broadcasts in parallel. In this case, the integrity condition effectively requires that messages are attributable to a uniquely identifiable execution context. To this end, we define that every reliable broadcast execution has a unique identifier. In the following, we assume that every message implicitly contains an identifier and that each node runs a separate instance of the protocol for each identifier. The sender for a particular execution context is assumed to be globally known.

As mentioned above, the objective is to reliably broadcast some message  $m$ , where the notation  $|m|$  is used to denote the size of the message in bits. Instead of sending  $m$ , the algorithms introduced in §3 and §4 send *fragments* of  $m$ , which can be either chunks of  $m$  itself or encoded data derived from  $m$ . Without loss of generality, we assume that there is some upper bound  $\ell_{max}$  on the message size, which in turn limits the fragment size. State-of-the-art algorithms use either *erasure codes*, which can handle missing data, or *error-correcting codes*, which support the correction of erroneous

data. The algorithms in this paper exclusively use an  $(n, k)$ -erasure codes with optimal reception efficiency, i.e., exactly  $k$  out of  $n$  symbols are sufficient to reconstruct a message of  $k$  symbols. Throughout this paper, the parameters  $n$  and  $k$  correspond to the total number of nodes and the number of honest nodes, i.e.,  $k := 2t + 1$ . We assume that each node has access to the routine `get_fragments` that takes a message  $m$  as input and returns a list of  $n$  fragments of size  $\frac{|m|}{k} = \frac{|m|}{2f+1} < \frac{3}{2} \frac{|m|}{n}$  each. Furthermore, the nodes use the routine `recover_message` to recover  $m$  given any subset  $F$  of the  $n$  fragments of size  $|F| \geq k = 2t + 1$ . This routine is assumed to always return *some* message, which may simply be a bit string of zeroes in case of an error, e.g., when the input contains fragments of different sizes.

In addition to fragments, the nodes also send hashes, identifiers, and, in the case of the algorithm presented in §4, also signatures. The hashes are assumed to be *cryptographically strong* in the sense that it is infeasible to find hash collisions (except with negligible probability). Naturally, it is also assumed that it is computationally infeasible to spoof signatures. Since the cryptographic properties only hold for hashes and signatures of a certain minimum size, we introduce a security parameter  $\kappa$  and define that the size of all data types other than fragments is bounded by  $O(\kappa)$ .

Reliable broadcast algorithms are evaluated against multiple complexity measures. For a specific number  $n$  of nodes, number  $t$  of Byzantine nodes, and message size  $\ell$ , the *communication complexity*  $C(n, t, \ell)$  of an algorithm is the total number of bits sent by all honest nodes in the worst case. In the case of a Byzantine sender that sends fragments of different sizes, we define that  $\ell$  is the message size corresponding to the largest fragment sent by an honest node. Given that at least  $n\ell$  bits need to be transferred and we are concerned with the overhead for large messages, the primary goal is to minimize  $\mathcal{L}(n, t) := \lim_{\ell \rightarrow \infty} \frac{C(n, t, \ell)}{n\ell}$ . Additionally, the *time complexity* of an algorithm is of great practical importance, measuring the duration of an execution when normalizing the maximum message delay to 1 time unit. According to Definition 2.1, an execution with a Byzantine sender may never terminate. Therefore, we restrict our attention to the case where the sender is honest. Due to the validity condition, this *good-case* time complexity must be bounded. Lastly, the *space complexity* is considered as well, which is defined as the number of bits that any honest node stores during the execution of the algorithm in the worst case.

### 3 ALGORITHM

#### 3.1 Overview

The algorithm, referred to as  $\mathcal{A}_{bit}$ , effectively works as outlined in §1 with a few important additions. In the first step, the sender disseminates the fragments. The validity of fragments is verified using Merkle proofs (as in [9]). After validating the received fragment, each honest node broadcasts a *proposal* to accept the message with the root hash in the received Merkle proof. The proposals serve to keep the communication complexity low and ensure that the integrity property holds. Honest nodes broadcast their fragments only if they receive at least  $2t + 1$  proposals for the corresponding root hash. Lastly, if an honest node receives at least  $2t + 1$  fragments and manages to reconstruct the correct message, it first carries out the crucial step of sending the  $j^{th}$  fragment to node  $v_j$  if it did not receive any fragment from  $v_j$  to guarantee that the totality property holds in case of a Byzantine sender before delivering the message.

#### 3.2 Description

Algorithm  $\mathcal{A}_{bit}$  makes use of the routines `get_fragments` and `recover_message` introduced in §2 to generate fragments of a given message and recover the message given at least  $2t + 1$  valid fragments, respectively. Furthermore, it requires routines for Merkle tree operations, specifically, `get_merkle_root` yields the Merkle root hash for a given set of fragments, `get_merkle_proof`

---

**Algorithm 1** Algorithm  $\mathcal{A}_{bit}$ : Triggered actions at node  $v_i$ . Initially,  $F = R = H = P = \{\}$ .

---

```

if reliable_broadcast( $m$ ) invoked and  $v_i = \text{sender}$  then
  ( $f_1, \dots, f_n$ ) := get_fragments( $m$ )
   $h$  := get_merkle_root( $f_1, \dots, f_n$ )
  for  $v_j \in V$  do
     $\pi_j$  := get_merkle_proof( $(f_1, \dots, f_n), j$ )
    send fragment( $h, j, f_j, \pi_j$ ) to  $v_j$ 

if received fragment( $h, j, f_j, \pi_j$ ) from  $v_k$  and ( $j = i$  or  $j = k$ ) then
  if ( $|H(v_k)| < 2$  or  $h \in H(v_k)$ ) and valid_merkle_proof( $h, f_j, j, \pi_j$ ) then
     $H(v_k) := H(v_k) \cup \{h\}$ ,  $R(h) := R(h) \cup \{v_k\}$ ,  $F(h, v_j) := (f_j, \pi_j)$ 
    if  $i = j$  and first fragment from  $v_k = \text{sender}$  then
      broadcast_proposal( $h$ )

if received proposal( $h$ ) from  $v_k$  then
  if  $|H(v_k)| < 2$  or  $h \in H(v_k)$  then
     $H(v_k) := H(v_k) \cup \{h\}$ ,  $P(h) := P(h) \cup \{v_k\}$ 

```

---

returns the Merkle proof for a specific fragment, and `valid_merkle_proof` indicates whether a given Merkle proof is valid.

Each node executing  $\mathcal{A}_{bit}$  locally maintains the data structures  $F$ ,  $R$ ,  $H$ , and  $P$  per execution context. Let  $\mathcal{F}$ ,  $\Pi$ , and  $\mathcal{H}$  denote the set of all possible fragments, Merkle proofs, and hashes, respectively. The hash map  $F : \mathcal{H} \times V \rightarrow \mathcal{F} \times \Pi$  stores the fragment  $f_j \in \mathcal{F}$  and Merkle proof  $\pi_j \in \Pi$  of node  $v_j$  for the message with the root hash  $h \in \mathcal{H}$ . If this fragment is locally available, then  $F(h, v_j) = (f_j, \pi_j)$ , and  $F(h, v_j) = \perp$  otherwise. This hash map is used to collect fragments with the goal of eventually recovering the corresponding message. The hash map  $R : \mathcal{H} \rightarrow 2^V$  stores the nodes from which a fragment for root hash  $h$  has been obtained. This hash map is needed to determine which nodes may still be missing their fragments once a node is able to recover the message. Conversely, for any node  $v$ , the hash map  $H : V \rightarrow 2^{\mathcal{H}}$  contains the set of hashes for which a proposal or fragment was received from node  $v$ . The purpose of this data structure is to bound both the communication and space complexity. Lastly,  $P : \mathcal{H} \rightarrow 2^V$  indicates which nodes have proposed the delivery of the message associated with the root hash  $h$ . The collected proposals are used to ensure that only fragments with sufficient support are broadcast. All hash maps are initially empty. In addition to the hash maps, each node further uses the Boolean variable *done*, initially *false*, to capture the information whether the execution has terminated, either with or without the delivery of a message. In the latter case, the totality condition implies that no honest node will deliver a message for this execution. Let further  $\mathbf{H} := \cup_{v \in V} H(v)$  and  $h_{max} := \arg \max_{h \in \mathbf{H}} |P(h)|$ , breaking ties arbitrarily if no single hash received the most proposals.

We distinguish between “triggered actions”, which are executed when a function is invoked or a message is received, and “state-based actions”, which occur when some conditions hold for the local state. This separation facilitates the description and analysis of the algorithm’s properties. The triggered actions are formally stated in Algorithm 1 and discussed next.

Given a message  $m$ , a reliable broadcast is triggered by invoking the routine `reliable_broadcast` at node  $v_i = \text{sender}$ , which performs the standard steps of generating  $n$  fragments  $f_1, \dots, f_n$ , computing the Merkle root hash  $h$ , and then transmitting the *fragment message* `fragment( $h, j, f_j, \pi_j$ )`, which contains the corresponding Merkle proof  $\pi_j$ , to  $v_j$  for all  $j \in \{1, \dots, n\}$ . A node  $v_i$  only accepts a received message `fragment( $h, j, f_j, \pi_j$ )` from some node  $v_k$

---

**Algorithm 2** Algorithm  $\mathcal{A}_{bit}$ : State-based actions at node  $v_i$ . Initially,  $F = R = P = \{\}$ ,  $done = false$ . Let  $h_{max} := \arg \max_{h \in H} |P(h)|$ .

---

**if**  $|P(h_{max})| \geq 2t + 1$  **and not** broadcast  $(f_i, \pi_i) := F(h_{max}, v_i) \neq \perp$  **before then**  
**broadcast**  $fragment(h_{max}, f_i, i, \pi_i)$

**if**  $(|F(h_{max})| \geq t + 1$  **and not** broadcast  $proposal(h_{max})$  **before then**  
**broadcast**  $proposal(h_{max})$

**if**  $|P(h_{max})| \geq 2t + 1$  **and**  $|F(h_{max})| \geq 2t + 1$  **and not**  $done$  **then**  
 $m := recover\_message(F(h_{max}))$   
 $(f_1, \dots, f_n) := get\_fragments(m)$   
 $h := get\_merkle\_root(f_1, \dots, f_n)$   
**if**  $h = h_{max}$  **then**  
**for**  $v_j \in V \setminus R(h_{max})$  **do**  
 $\pi_j := get\_merkle\_proof((f_1, \dots, f_n), j)$   
**send**  $fragment(h_{max}, j, f_j, \pi_j)$  to  $v_j$   
 $deliver(m)$   
 $done := true$

} Execute delivery

---

under the following conditions. First, the sender  $v_k$  must have sent the fragment of the recipient  $v_i$  (i.e.,  $j = i$ ) or its own fragment (i.e.,  $j = k$ ). Second, the sender has not sent messages for two Merkle root hashes other than  $h$  before (formally,  $|H(v_k)| < 2$  or  $h \in H(v_k)$ ) and, lastly, the Merkle proof  $\pi_j$  in the received message is valid. If all conditions are met, the hash maps  $H$  and  $R$  are updated by adding  $h$  and  $v_k$ , respectively, and the fragment and Merkle proof are stored ( $F(h, v_j) := (f_j, \pi_j)$ ). In the final step, if the recipient received its fragment (i.e.,  $j = i$ ) and it is the first fragment received from the dedicated sender (i.e.,  $v_k = sender$ ), then  $v_i$  broadcasts  $proposal(h)$ . Whenever a proposal of the form  $proposal(h)$  is received from node  $v_k$ , it is again only accepted if  $v_k$  did not send messages associated with two other hashes before, in which case  $h$  is added to  $H(v_k)$  and  $v_k$  is added to  $P(v_k)$ .

Algorithm 2 describes the state-based actions of algorithm  $\mathcal{A}_{bit}$ . An honest node  $v_i$  broadcasts its fragment  $f_i$ , at most once, after collecting  $2t + 1$  proposals for the corresponding Merkle root hash  $h$ . By contrast, a node  $v_i$  broadcasts a proposal for a specific root hash  $h$  not only when it receives its fragment from  $sender$  but also when at least  $t + 1$  fragments for the root hash  $h = h_{max}$  have been received and  $v_i$  has not broadcast the proposal before. Lastly, if a node  $v_i$  receives at least  $2t + 1$  fragments and proposals for  $h_{max}$ , it recovers the message  $m$  and, additionally, recomputes the fragments and the corresponding root hash. If the computed hash matches  $h_{max}$ , node  $v_i$  concludes that the fragments are valid and the message  $m$  can be delivered. In this case,  $v_i$  first recomputes the Merkle proof  $\pi_j$  and sends  $fragment(h_{max}, j, f_j, \pi_j)$  to each  $v_j \in V \setminus R(h_{max})$  before calling  $deliver(m)$  to deliver the message. Note that  $done$  is set to  $true$  after executing these steps even if  $h \neq h_{max}$  because it is computationally infeasible to find fragments such that equality holds for some subsets of fragments of size  $2t + 1$ .

### 3.3 Analysis

In this section, we prove the correctness of  $\mathcal{A}_{bit}$  as well as its communication, time, and space complexity. A series of lemmas is used to simplify the proof structure. The first lemma is concerned with the fragments that honest nodes broadcast, showing that the proposal mechanism successfully ensures that honest nodes only broadcast fragments for one specific root hash.

LEMMA 3.1. *If honest nodes  $v$  and  $v'$  broadcast fragments for root hashes  $h$  and  $h'$ , then  $h = h'$ .*

PROOF. Without loss of generality, let  $h$  be the first hash for which  $2t + 1$  proposals are received at some node  $v$ . Thus, there are at least  $t + 1$  honest nodes that received their fragments from *sender* and then broadcast their first proposal for root hash  $h$ .

Let  $v'$  be the first node that broadcasts a fragment for a hash  $h' \neq h$ . Since it must hold that  $|P(h')| \geq 2t + 1$  at node  $v'$ , there must be at least  $t + 1$  honest nodes that broadcast the proposal for root hash  $h'$ . Consequently, there must be an honest node  $v^*$  that first proposed  $h$  and then  $h'$ . However,  $v^*$  only sends a proposal for  $h'$  if  $|F(h')| \geq t + 1$  according to Algorithm 2, which implies that it must have received a fragment for root hash  $h'$  from an honest node before. This is a contradiction to the assumption that  $v'$  is the first node that broadcasts such a fragment.  $\square$

Since algorithm  $\mathcal{A}_{bit}$  imposes restrictive rules for the acceptance of received messages, we must show that honest nodes always accept messages from other honest nodes regardless of the messages of Byzantine nodes.

LEMMA 3.2. *If an honest node  $v$  sends a proposal or fragment message to an honest node  $v'$ , then  $v'$  accepts and processes the received message.*

PROOF. According to Algorithm 1,  $v'$  accepts messages associated with at most two different hashes for any sender  $v$ . After an initial proposal for some root hash  $h$ ,  $v$  may send a second proposal for  $h'$  if  $|F(h')| \geq t + 1$ , i.e., there is at least one honest node that has broadcast a fragment for root hash  $h'$ . Assume that  $v$  sends another proposal for a different root hash  $h''$ , which again implies that at least one honest node must have broadcast a fragment for this root hash, a contradiction to Lemma 3.1.

Regarding the transmission of fragments, Lemma 3.1 also implies that an honest node only sends fragments for one root hash. Assume that  $v$  sends fragments for a root hash  $h''$  that differs from the hashes  $h$  and  $h'$  for which it sends proposals. In this case,  $v$  must have received at least  $t + 1$  fragments for one of the root hashes  $h$  or  $h'$ . Hence it follows that an honest node broadcast a fragment for a root hash other than  $h''$ , again contradicting Lemma 3.1.  $\square$

An important criterion for the totality property is that all honest eventually manage to get sufficiently many fragments of a message  $m$  and proposals for the corresponding root hash if there is an honest node that delivers  $m$ . The last lemma states that this is indeed the case.

LEMMA 3.3. *If an honest node  $v$  delivers  $m$  with root hash  $h$ , all honest nodes will eventually store at least  $2t + 1$  proposals for  $h$  and  $2t + 1$  fragments of  $m$ .*

PROOF. Assume that  $v$  delivers  $m$  with root hash  $h$ . Since  $|F(h)| \geq 2t + 1$ ,  $v$  received fragments of  $m$  from at least  $t + 1$  honest nodes, i.e., every honest node eventually receives at least  $t + 1$  fragments of  $m$ . As a result, every honest node broadcasts *proposal*( $h$ ) at some point and thus  $|P(h)| \geq 2t + 1$  eventually holds at all honest nodes.

Node  $v$  adds all nodes from which it received fragments for root hash  $h$  to the set  $R(h)$ . For any node  $v' \in R(h)$  it holds that  $v'$  either sent its own fragment or  $v$ 's fragment. However, the latter case also implies that  $v'$  must have its own fragment. According to Algorithm 1,  $v$  sends fragment  $f_j$  to  $v_j$  for all  $v_j \in V \setminus R(h)$ , guaranteeing that the remaining nodes eventually receive and, due to Lemma 3.2, store their fragments as well. Thus, it eventually holds that  $|P(h)| \geq 2t + 1$  and  $F(h) \neq \perp$  at all honest nodes, causing them to broadcast their fragments. Consequently, every honest node will eventually receive fragments of  $m$  from at least  $2t + 1$  nodes.  $\square$

We are now in the position to prove the following main result.

**THEOREM 3.4.** *Algorithm  $\mathcal{A}_{bit}$  implements reliable broadcast in the asynchronous communication model with  $t < n/3$  Byzantine nodes.*

**PROOF.** The four conditions of reliable broadcast are proved separately.

**Validity:** If the sender  $v$  of a message  $m$  with root hash  $h$  is honest, it sends the fragment  $f_j$  to  $v_j$  for all  $v_j \in V$ . Subsequently, every honest nodes broadcasts  $proposal(h)$ . Since all honest nodes eventually receive at least  $2t + 1$  proposals for hash  $h$  (and only for hash  $h$ ), they all broadcast their fragments. Thus, eventually  $|F(h)| \geq 2t + 1$  and  $|P(h)| \geq 2t + 1$  holds at all honest nodes, triggering the delivery of  $m$ .

**Agreement:** Assume for the sake of contradiction that honest nodes  $v$  and  $v'$  deliver distinct messages  $m$  and  $m'$  with root hashes  $h$  and  $h'$ , respectively. Since  $|F(h)| \geq 2t + 1$  at  $v$  and  $|F(h')| \geq 2t + 1$  at  $v'$ , there must be honest nodes that have broadcast fragments for  $h$  and  $h'$ , which contradicts Lemma 3.1.

**Integrity:** Lemma 3.1 implies that there can only be sufficiently many fragments for at most one hash and hence at most one message can be delivered.

**Totality:** Let  $v$  be a node that delivers a message  $m$  with root hash  $h$ . Due to Lemma 3.3, it holds eventually that  $|F(h)| \geq 2t + 1$  and  $|P(h')| \geq 2t + 1$  at all honest nodes. Since  $v$  managed to verify the correctness of the fragments and the corresponding root hash, the verification also succeeds at all other nodes. Thus, all honest nodes eventually deliver  $m$ .  $\square$

The following theorem states the communication complexity of algorithm  $\mathcal{A}_{bit}$ .

**THEOREM 3.5 (COMMUNICATION COMPLEXITY).** *It holds for algorithm  $\mathcal{A}_{bit}$  that  $\mathcal{L}(n, t) = 2$ .*

**PROOF.** Let  $|f| := \frac{3}{2}|m|/n + O(\kappa \log(n))$  denote the size of the largest fragment message sent by an honest node. If the sender is honest,  $|f|$  corresponds to the size of a fragment message containing a valid fragment of message  $m$ .

The sender may send fragment messages to all other nodes for a total of at most  $(n - 1)|f|$  bits. Each honest node may broadcast one fragment message containing its fragment and at most two proposals of size  $O(\kappa)$ . After reconstructing the message  $m$ , an honest node further sends a fragment message containing  $f_j$  to all  $v_j \in V \setminus R(h)$ , where  $|V \setminus R(h)| \leq t < n/3$ . Not counting messages that nodes send to themselves, the communication complexity  $\mathcal{C}(n, t, |m|)$  is therefore upper bounded by

$$(n - 1)|f| + n \cdot (n - 1 + t)|f| + O(\kappa n^2) < 2n|m| + O(\kappa n^2 \log(n)),$$

and therefore  $\mathcal{L}(n, t) = \lim_{|m| \rightarrow \infty} \frac{1}{n|m|} \mathcal{C}(n, t, |m|) = 2$ .  $\square$

As stated in §2, we consider the good-case time complexity with an honest sender. The following theorem states the time complexity of  $\mathcal{A}_{bit}$  for this case.

**THEOREM 3.6 (TIME COMPLEXITY).** *If the sender is honest,  $\mathcal{A}_{bit}$  has a time complexity of 3.*

**PROOF.** If the sender is honest, all honest nodes receive their fragments and send the proposal for the corresponding root hash after at most 1 time unit. Thus, after at most 2 time units, all honest nodes get at least  $2t + 1$  proposals, triggering the broadcasting of fragments. After at most 3 time units, the honest nodes receive the fragments from all honest nodes and deliver  $m$ .  $\square$

If the sender is Byzantine, an honest node may still deliver  $m$  eventually. It is easy to see that it takes at most 3 time units for every honest node to deliver after the first honest node delivers  $m$ .

As far as the space complexity is concerned, recall that  $\ell_{max}$  denotes the largest permissible message size, which is a lower bound on the space complexity. Given this bound, the following result shows that honest nodes require little additional storage space.



**THEOREM 3.7 (SPACE COMPLEXITY).** *Algorithm  $\mathcal{A}_{bit}$  has a space complexity of  $2\ell_{max} + O(n\kappa)$ .*

**PROOF.** While honest nodes send fragments for at most one root hash, Byzantine nodes cannot send fragments for more than two hashes, resulting in a total of at most  $t \cdot 2 + (2t + 1) \cdot 1 < \frac{4}{3}n$  fragments of size  $\frac{3}{2}\ell_{max}/n$  each. Thus, the space required to store received fragments is upper bounded by  $2\ell_{max}$ . The other data structures merely contain identifiers and hashes of size  $O(\kappa)$ . Since they may contain entries for at most 2 hashes per node, it follows that the total size of these data structures is bounded by  $O(n\kappa)$ .  $\square$

It is worth pointing out that the space complexity can be lowered to  $\frac{3}{2}\ell_{max} + O(n\kappa)$  by introducing the rule that received fragments for different hashes from the same node are rejected. This rule is omitted from Algorithm 1 for ease of exposition.

## 4 VARIANT WITH THRESHOLD SIGNATURES

### 4.1 Overview

In this section, a variant of  $\mathcal{A}_{bit}$ , called  $\mathcal{A}_{sig}$ , is introduced with an optimal time complexity of 2 while still guaranteeing that  $\mathcal{L}(n, t) < 3$ . The challenge is that this time complexity cannot be reached when sending proposals to agree on a root hash first before sending fragments. However, if fragments are broadcast right away, a Byzantine sender may send fragments for different root hashes to different nodes. Naturally, honest nodes cannot simply broadcast any received fragments without an adverse effect on the communication complexity.

Algorithm  $\mathcal{A}_{sig}$  addresses this challenge by using *threshold signatures*, i.e., we assume that each node has the keying material to *threshold-sign* messages and that every node can verify the validity of signatures of all nodes. At least  $2t+1$  threshold signatures are required to create a valid signature. The proposals of  $\mathcal{A}_{bit}$  are replaced by messages containing a root hash and a threshold signature or a full signature. The crucial advantage is that a (full) signature serves as proof that the sender has received threshold signatures from  $2t+1$  nodes, ensuring that the execution can make progress whenever an honest nodes meets the conditions to deliver  $m$ .

### 4.2 Description

Algorithm  $\mathcal{A}_{sig}$  uses *signed proposal messages* of the form  $proposal(h, \sigma)$ , where  $\sigma$  is either the threshold signature of  $h$  of a particular node or it is a signature of  $h$  derived from at least  $2t+1$  threshold signatures. Note that it is a signature of  $h$  and the identifier in practice in order to tie the signature to a particular execution. The algorithm still uses the hash maps  $F$ ,  $H$ , and  $R$ . Threshold signatures are collected in the hash map  $S : \mathcal{H} \times V \rightarrow \mathcal{S}$ , where  $\mathcal{S}$  denotes the set of all possible signatures. We further define that  $S(h) := \cup_{v \in V} S(h, v)$ . In this section,  $h_{max}$  is defined as  $h_{max} := \arg \max_{h \in \mathcal{H}} |S(h)|$ . In addition to *done*, each node further stores  $h^*$ , the root hash of the message that will be delivered if set (initially,  $h^* = \perp$ ).

The steps of algorithm  $\mathcal{A}_{sig}$  for triggered actions are shown in Algorithm 3. When `reliable_broadcast` is invoked, exactly the same steps as in Algorithm 1 are executed. A message  $fragment(h, j, f_j, \pi_j)$  is handled the same way as in  $\mathcal{A}_{bit}$  with the exception that  $h$  is signed and both  $proposal(h, \sigma_i)$  and  $fragment(h, f_i, i, \pi_i)$  are broadcast if it is the first (valid) fragment received from *sender*. When receiving  $proposal(h, \sigma)$  from some node  $v_k$ , the action depends on  $\sigma$ . If it is a valid signature,  $h^*$  is set to  $h$ . Otherwise, if it is a valid threshold signature and  $H(v_k)$  does not contain two hashes different from  $h$ , then  $h$  is added to  $H(v_k)$  and  $S(h, v_k)$  is set to  $\sigma$ .

The state-based actions at node  $v_i$  are shown in Algorithm 4. If at least  $2t+1$  threshold signatures have been collected for some root hash  $h_{max}$ , then  $h^*$  is set to  $h_{max}$ . Node  $v_i$  broadcasts its fragment  $f_i$  (with the Merkle proof  $\pi_i$ ) if  $h^*$  is set,  $f_i$  is locally available, and it has not been broadcast before.

---

**Algorithm 3** Algorithm  $\mathcal{A}_{sig}$ : Triggered actions at node  $v_i$ . Initially,  $F = R = H = S = \{\}$  and  $h^* = \perp$ .

---

**if** reliable\_broadcast( $m$ ) invoked **and**  $v_i = sender$  **then**  
     Execute reliable\_broadcast( $m$ ) of Algorithm 1

**if** received fragment( $h, j, f_j, \pi_j$ ) from  $v_k$  **and** ( $j = i$  **or**  $j = k$ ) **then**  
     **if** ( $|H(v_k)| < 2$  **or**  $h \in H(v_k)$ ) **and** valid\_merkle\_proof( $h, f_j, j, \pi_j$ ) **then**  
          $H(v_k) := H(v_k) \cup \{h\}$ ,  $R(h) := R(h) \cup \{v_k\}$ ,  $F(h, v_j) := (f_j, \pi_j)$   
         **if**  $i = j$  **and** first fragment from  $v_k = sender$  **then**  
              $\sigma_i := \text{threshold\_sign}(h)$   
             **broadcast** [ $proposal(h, \sigma_i)$ , fragment( $h, f_i, i, \pi_i$ )]

**if** received proposal( $h, \sigma$ ) from  $v_k$  **then**  
     **if** valid\_multi\_sig( $h, \sigma$ ) **then**  
          $h^* := h$   
     **else if** ( $|H(v_k)| < 2$  **or**  $h \in H(v_k)$ ) **and** valid\_sig( $h, v_k, \sigma$ ) **then**  
          $H(v_k) := H(v_k) \cup \{h\}$ ,  $S(h, v_k) := \sigma$

---

Lastly, if  $h^* \neq \perp$  and  $|F(h^*)| \geq 2t + 1$ ,  $v_i$  computes the signature and broadcasts the corresponding proposal before executing the same delivery steps as in Algorithm 2.

### 4.3 Analysis

It is evident from the description of algorithm  $\mathcal{A}_{sig}$  that honest nodes must set  $h^*$  to the same hash as otherwise the integrity property may be violated. The following lemma states that honest nodes indeed cannot set  $h^*$  to different hashes.

LEMMA 4.1. *If honest nodes set  $h^*$ , they set it to the same hash.*

PROOF. According to Algorithm 3, only the first message from *sender* is threshold-signed, i.e., every honest node provides a threshold signature for at most one hash. However, setting  $h^*$  to some hash  $h$  requires  $2t + 1$  threshold signatures of  $h$ . If we assume that there are two different hashes for which  $2t + 1$  threshold signatures have been collected, there must be at least one honest node that threshold-signed two different hashes, which contradicts the rule that honest nodes only threshold-sign at most one hash.  $\square$

A Byzantine sender may send fragments for different root hashes to the honest nodes, causing wasteful transmissions of fragments. However, algorithm  $\mathcal{A}_{sig}$  ensures that the number of transmissions is bounded as the next lemma shows.

LEMMA 4.2. *Every honest node sends fragments for at most two different root hashes.*

PROOF. An honest node  $v_i$  broadcasts its fragment  $f_i$  for some root hash  $h$  when receiving  $f_i$  from *sender*. Apart from this transmission, according to Algorithm 3,  $v_i$  only sends fragments associated with  $h^*$ . Since  $h^*$  never changes once set, the claim follows.  $\square$

Since algorithm  $\mathcal{A}_{sig}$  enforces similar restrictions for the acceptance of messages as algorithm  $\mathcal{A}_{bit}$ , Lemma 3.2 holds for  $\mathcal{A}_{sig}$  as well.

LEMMA 4.3. *Lemma 3.2 ("If an honest node  $v$  sends a proposal or fragment message to an honest node  $v'$ , then  $v'$  accepts and processes the received message.") holds for  $\mathcal{A}_{sig}$ .*

---

**Algorithm 4** Algorithm  $\mathcal{A}_{sig}$ : State-based actions at node  $v_i$ . Initially,  $F = R = S = \{\}$ ,  $h^* = \perp$ , and  $done = false$ . Let  $h_{max} := \arg \max_{h \in \mathcal{H}} |S(h)|$ .

---

**if**  $|S(h_{max})| \geq 2t + 1$  **and**  $h^* = \perp$  **then**  
 $h^* := h_{max}$

**if**  $h^* \neq \perp$  **and not** broadcast  $(f_i, \pi_i) := F(h^*, i) \neq \perp$  **before then**  
**broadcast**  $fragment(h^*, f_i, i, \pi_i)$

**if**  $h^* \neq \perp$  **and**  $|F(h^*)| \geq 2t + 1$  **and not**  $done$  **then**  
 $\sigma := \text{compute\_signature}(S(h^*))$   
**broadcast**  $proposal(h^*, \sigma)$   
Execute delivery as in Algorithm 2

---

PROOF. Algorithm  $\mathcal{A}_{sig}$  and  $\mathcal{A}_{bit}$  share the property that messages for two different hashes are accepted from any node. According to Algorithm 3, an honest node broadcasts a signature and fragment message for the same hash  $h$  when receiving its fragment from *sender*. Algorithm 4 states that any further transmission of a signature or fragment message must be for hash  $h^* \neq \perp$ . Since  $h^*$  never changes when set, honest nodes send messages for at most 2 hashes.  $\square$

Algorithm  $\mathcal{A}_{sig}$  depends on the same mechanism to ensure totality as  $\mathcal{A}_{bit}$  and therefore the following variant of Lemma 3.3 holds.

LEMMA 4.4. *If an honest node delivers  $m$  with root hash  $h$ , all honest nodes will eventually store at least  $2t + 1$  fragments of  $m$  and set  $h^*$  to  $h$ .*

PROOF. The same argument as for algorithm  $\mathcal{A}_{bit}$  applies, which proves that every honest node gets its fragment eventually after an honest node  $v$  has delivered a message  $m$  for some hash  $h$ , which it accepts and stores due to Lemma 4.3. Since  $v$  must have computed and broadcast a valid signature  $\sigma$ , it eventually holds at all honest nodes that  $h^* = h$ . According to Lemma 4.1, honest nodes all set  $h^*$  to the same value. Consequently, every honest node will eventually broadcast its fragment. Thus, every honest node will eventually have at least  $2t + 1$  fragments of  $m$ .  $\square$

As in §3, the first main result is that algorithm  $\mathcal{A}_{sig}$  is a correct reliable broadcast algorithm.

THEOREM 4.5. *Algorithm  $\mathcal{A}_{sig}$  implements reliable broadcast in the asynchronous communication model with  $t < n/3$  Byzantine nodes.*

PROOF. The four conditions of reliable broadcast are again proved in sequence.

**Validity:** If the sender  $v$  of a message  $m$  with root hash  $h$  is honest, it sends  $f_j$  to  $v_j$  for all  $v_j \in V$ . Every honest node  $v_j$  broadcasts its fragment and threshold signature and, consequently, receives at least  $2t + 1$  signatures and fragments eventually. Subsequently, every honest node computes the signature  $\sigma$ , sets  $h^* := h$ , and proceeds to deliver  $m$ .

**Agreement:** If two honest nodes  $v$  and  $v'$  deliver different messages  $m$  and  $m'$ , they must have set  $h^*$  to different hashes, a contradiction to Lemma 4.1.

**Integrity:** Due to Lemma 4.1, every honest node delivers at most one message.

**Totality:** If an honest node  $v$  delivers a message  $m$  with root hash  $h$ , it eventually holds at all honest nodes that  $h^* = h$  and  $|F(h^*)| \geq 2t + 1$  due to Lemma 4.4. Since  $v$  managed to verify the correctness of the fragments and the root hash, it holds again that the verification must succeed at other nodes as well. Every honest node delivers  $m$  after the successful verification.  $\square$

Compared to  $\mathcal{A}_{bit}$ , the communication complexity of algorithm  $\mathcal{A}_{sig}$  is worse, which appears to be inevitable given that the nodes must broadcast fragments without delay to achieve an optimal time complexity. However, the following lemma states that even a Byzantine sender cannot induce many (wasteful) transmissions of fragments.

LEMMA 4.6. *At most  $t$  honest nodes send fragments for 2 different root hashes.*

PROOF. Let  $S$  denote the set of honest nodes that send fragments for more than one root hash. Algorithm 3 dictates that an honest node only sends another fragment when  $h^* \neq \perp$ . If there are  $t' \leq t$  Byzantine nodes in the execution, at least  $2t + 1 - t'$  honest nodes must receive their fragments for hash  $h^*$  in order to obtain a signature  $\sigma$  for this hash. Since none of these nodes sends fragments for other root hashes, we get that  $|S| \leq 3t + 1 - (2t + 1 - t') - t' = t$ .  $\square$

As the following theorem shows, algorithm  $\mathcal{A}_{sig}$  achieves an overhead factor that tends to  $5/2$  for large messages.

THEOREM 4.7 (COMMUNICATION COMPLEXITY). *It holds for algorithm  $\mathcal{A}_{sig}$  that  $\mathcal{L}(n, t) = \frac{5}{2}$ .*

PROOF. The dissemination of fragments by the sender requires at most  $(n - 1)|f|$  bits, where  $|f|$  again denotes the size of the largest fragment message sent by an honest node. Due to Lemma 4.6, at most  $t$  honest nodes broadcast fragments twice, and the other other honest nodes broadcast at most one fragment. Additionally, each node sends  $O(\kappa n)$  bits for the signed proposal messages and  $f_j$  to all  $v_j \in V \setminus R(h)$ , where  $|V \setminus R(h)| \leq t < n/3$  as before. Hence it follows that  $C(n, t, |m|)$  is at most

$$\begin{aligned} & (n - 1)|f| + (2t + 1)(n - 1)|f| + 2t(n - 1)|f| + t(n - 1)|f| + O(\kappa n^2) \\ & < (n - 1)|f| + \frac{5}{3}n(n - 1)|f| + O(\kappa n^2) < \frac{5}{2}n|m| + O(\kappa n^2 \log(n)), \end{aligned}$$

where the first inequality holds because  $(5t + 1) < \frac{5}{3}(3t + 1) = \frac{5}{3}n$ . Thus,  $\mathcal{L}(n, t) = \frac{5}{2}$ .  $\square$

It is worth noting that the communication complexity is only worse compared to  $\mathcal{A}_{bit}$  if the sender deliberately transmits conflicting fragments, i.e.,  $\mathcal{L}(n, t) = 2$  still holds otherwise. The main advantage of  $\mathcal{A}_{sig}$  over  $\mathcal{A}_{bit}$  is its superior time complexity.

THEOREM 4.8 (TIME COMPLEXITY). *If the sender is honest,  $\mathcal{A}_{sig}$  has a time complexity of 2.*

PROOF. If the sender is honest, all honest nodes receive their fragments within 1 time unit and broadcast their threshold signatures and fragments. After another time unit, all honest nodes must have received at least  $2t + 1$  threshold signatures and fragments for the same root hash. As a result, each honest node computes the signature and sets  $h^*$ , which triggers the delivery of  $m$ .  $\square$

If an honest node  $v$  delivers  $m$  despite a Byzantine sender, it is straightforward to show that each honest node delivers  $m$  at most 2 time units later. While the time complexity of  $\mathcal{A}_{sig}$  is lower, it has a slightly higher space complexity.

THEOREM 4.9 (SPACE COMPLEXITY). *Algorithm  $\mathcal{A}_{sig}$  has a space complexity of  $\frac{5}{2}\ell_{max} + O(n\kappa)$ .*

PROOF. According to Lemma 4.6, there are at most  $t$  honest nodes that send their fragments for 2 root hashes, i.e., there are at least  $t + 1$  honest nodes that only send their fragment once. Thus, an honest node may store  $2t \cdot 2 + (t + 1) \cdot 1 \leq \frac{5}{3}n$  fragments of size  $\frac{3}{2}\ell_{max}/n$  for a total of  $\frac{5}{3}\ell_{max}$ .

Algorithm  $\mathcal{A}_{sig}$  also utilizes  $R$ , containing at most 2 node identifiers of size  $O(\kappa)$  per node, and  $S$ , storing at most 2 signatures of size  $O(\kappa)$  per node. Thus, the space complexity of these data structures is upper bounded by  $O(n\kappa)$ .  $\square$

## 5 PRACTICAL CONSIDERATIONS

In this section, optimizations that may be relevant for practical applications are discussed.

**Partially synchronous communication.** Byzantine behavior and unpredictable message latencies are often exceptional situations rather than the norm in practice. While both algorithm  $\mathcal{A}_{sig}$  and  $\mathcal{A}_{bit}$  are tailored to the asynchronous communication model, they can easily be adapted to the *partially synchronous model* where periods of synchrony are assumed. Given an upper bound  $d$  on the message delay that holds for some periods of time, both algorithms can be modified to improve the communication complexity during these times in the absence of faults. Specifically, the additional constraint can be introduced that at least  $\delta$  time must have passed since the first fragment was received before delivering  $m$ . The following theorem states the effect of this modification on the communication complexity of both algorithms for an algorithm-specific  $\delta$ .

**THEOREM 5.1.** *If communication is synchronous from the start of the execution for at least  $\delta = 3d$  ( $\delta = 2d$ ) time, then  $\mathcal{L}(n, t) = \frac{3}{2}$  for the adapted version of  $\mathcal{A}_{bit}$  ( $\mathcal{A}_{sig}$ ) in the absence of faults.*

**PROOF.** As there are no faults by assumption, Theorem 3.6 and Theorem 4.8 imply that all nodes receive  $n$  fragments after at most  $3d$  ( $2d$ ) time in case of algorithm  $\mathcal{A}_{bit}$  ( $\mathcal{A}_{sig}$ ). Thus,  $R(h) = V$  for the root hash  $h$  of  $m$ , which entails that the nodes do not disseminate any additional fragments before delivering, resulting in a communication complexity of  $(n-1)|f| + n(n-1)|f| + O(\kappa n) < n^2|f| + O(\kappa n) \leq \frac{3}{2}|m|n + O(\kappa n^2 \log(n))$  and thus  $\mathcal{L}(n, t) = \frac{3}{2}$ .  $\square$

Naturally, this modification does not improve any theoretical (worst-case) bound because it still holds that  $\mathcal{L}(n, t) = 2$  even for  $t$  crash failures but it may prove to be beneficial for real-world applications where failures are rare nonetheless.

**Simplified structure.** The second optimization only concerns  $\mathcal{A}_{sig}$ . Instead of handling fragments and proposals in separate messages, signatures can be appended to the fragment message. As a result, nodes only send and process messages of a single type. Fragments and proposals are kept separate in algorithm  $\mathcal{A}_{sig}$  to retain as much similarity to  $\mathcal{A}_{bit}$  as possible in order to emphasize the key differences and simplify the presentation. This modification requires the addition of a simple rule: a node appends its threshold signature as long as  $h^* = \perp$  and the signature otherwise, which can be computed if  $h^* \neq \perp$  according to Algorithm 4. The signature in the message is still processed as shown in Algorithm 3.

Algorithm 4 dictates that the signature is broadcast before executing the message delivery. This step can be omitted entirely when adding the rule above without violating any of the reliable broadcast conditions. Since the signature is broadcast to ensure totality, we briefly argue why this condition continues to hold. Let  $v$  be an honest node that has delivered message  $m$  with root hash  $h$ . If all honest nodes broadcast fragments for root hash  $h$  at some point, then all honest nodes will eventually deliver  $m$  even if they never receive a signature. On the other hand, if there is at least one honest node  $v'$  that never sends a fragment for the right hash, then  $v' \notin R(h)$  at node  $v$ , and thus  $v$  must send  $v'$ 's fragment and the signature to  $v'$ . Upon receipt of this message,  $v'$  will broadcast its fragment and the signature, ensuring that all honest nodes eventually receive the signature, i.e., the signature is broadcast anyway in this case.

## 6 LOWER BOUND

The preceding section showed that an overhead factor of  $\frac{3}{2}$  for large messages can be attained in practice under normal network conditions and in the absence of faults. In this section, we shed light on the question whether we can hope to further improve upon this bound. The question is answered in the negative, at least for a specific but relevant class of algorithms.

The lower bound holds in the *synchronous communication model* where all nodes perform some computation, send messages, and receive the sent messages within the same *round* of a bounded and known duration. Furthermore, the lower bound holds in the *crash failure model* where a faulty node stops executing at any point during the execution but it never deviates from the correct protocol execution until it fails. In order to utilize the available bandwidth efficiently, it is beneficial to minimize the maximum bandwidth consumption over all nodes. If each node sends the same amount of data up to a constant factor, the algorithm is called *balanced* [4]. Any imbalance beyond a constant factor is usually due to the sender transmitting more data, typically in the first round. We focus on a broader class of “weakly balanced” algorithms where the sender can send arbitrarily sized messages to other nodes but at most  $o(|m|n)$  bits overall in the first round. Note that each node sends  $O(|m| + \kappa n \log(n))$  bits in both  $\mathcal{A}_{bit}$  and  $\mathcal{A}_{sig}$ .

We further restrict our attention to reliable broadcast algorithms that, given an honest sender, guarantee that all honest nodes deliver  $m$  after at most 3 (synchronous) rounds. This class of algorithms is interesting because it covers  $\mathcal{A}_{bit}$ ,  $\mathcal{A}_{sig}$ , as well as other algorithms in the literature including the algorithm by Cachin and Tessaro.

Formally, a  $(b, r)$ -reliable broadcast algorithm is defined as a reliable broadcast algorithm that sends messages of size at most  $b$  bits in the first round and delivers the message of an honest sender at all honest nodes within at most  $r$  rounds.

Let  $b_j$  denote the bits (of entropy) that the sender sends to  $v_j$  in the first round according to the given algorithm in some execution *without any failures*. As mentioned before, we assume that

$$\sum_{j=1}^n b_j \in o(|m|n). \quad (1)$$

Let  $b_{ij}$  further denote the bits that  $v_i$  sends to  $v_j$  in rounds  $2, \dots, r$  in the same execution. Due to the validity condition and the fact that we consider algorithms with a good-case time complexity of at most  $r$ , all honest nodes must receive  $m$  by the end of round  $r$ . Let  $S_j$  denote the set  $\{v_i \in V \mid b_{ij} > 0\}$ , i.e., the set of nodes that send some bits to  $v_j$ . Moreover, let  $R := \{v_j \in V \mid b_j \geq |m|\}$  denote the set of nodes that receive at least  $|m|$  bits in round 1. Since  $\sum_{j=1}^n b_j \in o(|m|n)$ , it must hold that  $|R| \in o(n)$ . Let  $\bar{b}_j := \frac{1}{|S_j|} \sum_{i=1}^n b_{ij}$  denote the average number of bits sent to  $v_j$  from the nodes in  $S_j$  in rounds  $2, \dots, r$ .

In the following, we consider the case  $r = 3$ , which permits a simple strategy to derive a lower bound due to the fact that a node failing to receive expected messages in the second round can only inform other nodes about the failure in the third round. Requiring a fourth round to send the missing bits violates the requirement that all honest nodes deliver the message by round 3 in case of a non-faulty sender. Therefore, the strategy is simply to mark those nodes as faulty that send many bits in rounds 2 and 3.

Given this strategy, it is easy to see that  $|S_j| > t$  and  $\bar{b}_j > 0$  must hold for any  $v \in V \setminus R$  as otherwise  $v_j$  may not receive  $m$  by the end of round 3. The following lemma provides a lower bound on  $\bar{b}_j$  for the case  $|S_j| > t$ .

LEMMA 6.1. *For all  $j \in \{1, \dots, n\}$ ,  $|S_j| > t$ , it must hold that*

$$\bar{b}_j \geq (|m| - b_j) / (|S_j| - t). \quad (2)$$

PROOF. Note that  $b_j \geq |m|$  implies that  $\bar{b}_j \geq 0$ , which holds trivially. Thus, we consider the case  $b_j < |m|$ . Assume that  $\bar{b}_j < (|m| - b_j) / (|S_j| - t)$  and that the  $t$  nodes in  $S_j$  that send the largest number of bits are faulty and never send anything. Let  $S'_j$  denote the remaining nodes that send

bits in rounds 2 and 3. In this case, node  $v_j$  receives

$$b_j + \sum_{v_i \in S'_j} b_{ij} \leq b_j + (|S_j| - t)\bar{b}_j < b_j + (|S_j| - t)(|m| - b_j)/(|S_j| - t) = |m|$$

bits by the end of round 3, which implies that  $v_j$  does not receive the (whole) message.  $\square$

The following theorem states the main result that the overhead factor must be at least  $\frac{3}{2}$  for the considered class of algorithms as  $n$  tends to infinity.

**THEOREM 6.2.** *If communication is synchronous and there are  $t < n/3$  faulty nodes, it holds that  $\mathcal{L}(n, t) \geq \frac{3}{2}$  for every  $(o(|m|n), 3)$ -reliable broadcast algorithm as  $n \rightarrow \infty$ .*

**PROOF.** The communication complexity is at least

$$\begin{aligned} \sum_{j=1}^n b_j + \sum_{j=1}^n \sum_{i=1}^n b_{ij} &\geq \sum_{v_j \in V \setminus R} \sum_{i=1}^n b_{ij} = \sum_{v_j \in V \setminus R} |S_j| \bar{b}_j \stackrel{(2)}{\geq} \sum_{v_j \in V \setminus R} |S_j| \frac{|m| - b_j}{|S_j| - t} \geq \sum_{v_j \in V \setminus R} \frac{n}{n-t} (|m| - b_j) \\ &= (n - |R|) \frac{n}{n-t} |m| - \sum_{v_j \in V \setminus R} \frac{n}{n-t} b_j \stackrel{(1)}{\geq} (n - |R|) \frac{n}{n-t} |m| - o(|m|n) \stackrel{n \rightarrow \infty}{=} n \frac{n}{n-n/3} |m| = \frac{3}{2} |m|n. \end{aligned}$$

$\square$

## 7 RELATED WORK

As mentioned in §1, the first reliable broadcast algorithm, achieving a communication complexity of  $O(|m|n^2)$ , was presented by Bracha [6]. An algorithm with a much improved bound of  $O(|m|n + \kappa n^2 \log(n))$ , based on the use of erasure codes, was published 18 years later by Cachin and Tessaro [9], making use of collision-resistant hash functions of size  $\kappa$ . By contrast, Bracha's algorithm is *error-free*, i.e., it does not rely on any cryptographic assumptions.

Subsequently, error-free reliable broadcast algorithms with lower communication complexity were proposed, guaranteeing bounds of  $O(n|m| + n^4 \log(n))$  [23] and  $O(n|m| + n^3 \log(n))$  [22]. While the latter algorithm achieves a lower asymptotic communication complexity, it is not *balanced* in that a single node, the sender, must transmit more bits than the other nodes. A reliable broadcast algorithm is said to be *balanced* if all nodes send the same number of bits up to a constant factor [4]. Note that the algorithms presented in §3 and §4 are both balanced. The best known bound for error-free algorithms is  $O(n|m| + n^2 \log(n))$  [4]. Abraham and Asharov proposed a probabilistic algorithm with a similar bound of  $O(n|m| + n^2 \log(n^3/\epsilon))$ , guaranteeing validity but the agreement and totality properties only hold with probability  $1 - \epsilon$  [1].

It has also been shown how to achieve a communication complexity of  $O(n|m| + \kappa n^2)$  using only collision-resistant hash functions [11], improving upon the algorithm by Cachin and Tessaro. The downside of this algorithm is that it is not balanced and it has a higher computational cost. Algorithm  $\mathcal{A}_{bit}$  of §3 falls into the category of reliable broadcast algorithms that rely on collision-resistant hash functions as well.

There are reliable broadcast algorithms that require other cryptographic primitives. For example, an algorithm has been proposed that achieves a communication complexity of  $O(n|m| + \kappa n^2)$  but requires a trusted setup for a public key infrastructure and cryptographic accumulators [22]. This bound has been improved to  $O(n|m| + \kappa n + n^2)$  using threshold signatures [4]. The algorithm  $\mathcal{A}_{sig}$  defined in §4 uses threshold signatures to achieve an optimal (good-case) time complexity.

The best known upper bounds, with and without cryptographic assumptions, are almost asymptotically tight considering the lower bound of  $\Omega(n|m| + n^2)$  [14]. The algorithms in this paper improve the *effective* communication complexity for large messages by reducing the constant of the

first term. Moreover, most error-free algorithms, excluding Bracha’s algorithm, and also the algorithms that use more cryptographic tooling than hash functions, have a higher time complexity than  $\mathcal{A}_{bit}$  and  $\mathcal{A}_{sig}$ , the latter of which having an optimal time complexity [3].

Byzantine reliable broadcast has further been studied in a variety of models that differ from the model used in most related work. There is a probabilistic algorithm based on stochastic samples that allows each property to be violated with a fixed and small probability [17]. An algorithm for a model with dynamic membership has also been proposed [16]. There is further work on *consistent broadcast*, a variant of reliable broadcast without the totality property, and its applications [7, 25]. Lastly, in an effort to minimize the actual latency and bandwidth consumption, simulations have been used to show the efficacy when combining Bracha’s algorithm with Dolev’s reliable communication protocol [13] to disseminate messages reliably in partially connected networks [5].

## 8 CONCLUSION

The mechanism introduced in this paper lowers the communication complexity of Byzantine reliable broadcast for large messages. The presented algorithms, which utilize this mechanism, further guarantee near-optimal and optimal bounds on the time complexity while also keeping the space complexity close to optimal. As mentioned in §1, numerous applications make use of reliable broadcast as a subroutine. Therefore, it may be worthwhile to revisit selected applications to determine if the techniques introduced in this paper can be used to obtain stronger results.

A fundamental open question is whether a lower communication complexity for large messages, i.e., a lower overhead factor, can yet be attained, ideally without increasing the time complexity substantially. We conjecture that the lower bound of  $3/2$ , which has been shown for a specific class of algorithms, holds more generally, i.e., with one or both of the imposed restrictions lifted. Deriving a tight bound for the overhead factor is another important avenue for future research.

## REFERENCES

- [1] Ittai Abraham and Gilad Asharov. 2022. Gradecast in Synchrony and Reliable Broadcast in Asynchrony with Optimal Resilience, Efficiency, and Unconditional Security. In *Proc. 43rd ACM Symposium on Principles of Distributed Computing (PODC)*. 392–398.
- [2] Ittai Abraham, Philipp Jovanovic, Mary Maller, Sarah Meiklejohn, Gilad Stern, and Alin Tomescu. 2021. Reaching Consensus for Asynchronous Key Generation. In *Proc. 42nd ACM Symposium on Principles of Distributed Computing (PODC)*. 363–373.
- [3] Ittai Abraham, Kartik Nayak, Ling Ren, and Zhuolun Xiang. 2021. Good-case Latency of Byzantine Broadcast: A Complete Categorization. In *Proc. 42nd ACM Symposium on Principles of Distributed Computing (PODC)*. 331–341.
- [4] Nicolas Alhaddad, Sourav Das, Sisi Duan, Ling Ren, Mayank Varia, Zhuolun Xiang, and Haibin Zhang. 2022. Balanced Byzantine Reliable Broadcast with Near-Optimal Communication and Improved Computation. In *Proc. 43rd ACM Symposium on Principles of Distributed Computing (PODC)*. 399–417.
- [5] Silvia Bonomi, Jérémie Decouchant, Giovanni Farina, Vincent Rahli, and Sébastien Tixeuil. 2021. Practical Byzantine Reliable Broadcast on Partially Connected Networks. In *Proc. 41st International Conference on Distributed Computing Systems (ICDCS)*. 506–516.
- [6] Gabriel Bracha. 1987. Asynchronous Byzantine Agreement Protocols. *Information and Computation* 75, 2 (1987), 130–143.
- [7] Christian Cachin, Klaus Kursawe, Frank Petzold, and Victor Shoup. 2001. Secure and Efficient Asynchronous Broadcast Protocols. In *Annual International Cryptology Conference*. 524–541.
- [8] Christian Cachin and Jonathan A Poritz. 2002. Secure Intrusion-tolerant Replication on the Internet. In *Proc. International Conference on Dependable Systems and Networks (DSN)*. 167–176.
- [9] Christian Cachin and Stefano Tessaro. 2005. Asynchronous Verifiable Information Dispersal. In *Proc. 24th IEEE Symposium on Reliable Distributed Systems (SRDS)*. 191–201.
- [10] George Danezis, Lefteris Kokoris-Kogias, Alberto Sonnino, and Alexander Spiegelman. 2022. Narwhal and Tusk: A DAG-based Mempool and Efficient BFT Consensus. In *Proc. 17th European Conference on Computer Systems (EuroSys)*. 34–50.



- [11] Sourav Das, Zhuolun Xiang, and Ling Ren. 2021. Asynchronous Data Dissemination and its Applications. In *Proc. ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 2705–2721.
- [12] Sourav Das, Thomas Yurek, Zhuolun Xiang, Andrew Miller, Lefteris Kokoris-Kogias, and Ling Ren. 2022. Practical Asynchronous Distributed Key Generation. In *Proc. IEEE Symposium on Security and Privacy (S&P)*. 2518–2534.
- [13] Danny Dolev. 1981. Unanimity in an Unknown and Unreliable Environment. In *Proc. 22nd Annual Symposium on Foundations of Computer Science (FOCS)*. 159–168.
- [14] Danny Dolev and Rüdiger Reischuk. 1985. Bounds on Information Exchange for Byzantine Agreement. *Journal of the ACM (JACM)* 32, 1 (1985), 191–204.
- [15] Sisi Duan, Michael K. Reiter, and Haibin Zhang. 2018. BEAT: Asynchronous BFT Made Practical. In *Proc. ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 2028–2041.
- [16] Rachid Guerraoui, Jovan Komatovic, Petr Kuznetsov, Yvonne-Anne Pignolet, Dragos-Adrian Seredinschi, and Andrei Tonkikh. 2020. Dynamic Byzantine Reliable Broadcast. In *Proc. 24th International Conference on Principles of Distributed Systems (OPODIS)*. 23:1–23:18.
- [17] Rachid Guerraoui, Petr Kuznetsov, Matteo Monti, Matej Pavlovic, Dragos-Adrian Seredinschi, and Yann Vonlanthen. 2019. Scalable Byzantine Reliable Broadcast. In *Proc. 33rd International Symposium on Distributed Computing (DISC)*.
- [18] Bingyong Guo, Zhenliang Lu, Qiang Tang, Jing Xu, and Zhenfeng Zhang. 2020. Dumbo: Faster Asynchronous BFT Protocols. In *Proc. ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 803–818.
- [19] Idit Keidar, Eleftherios Kokoris-Kogias, Oded Naor, and Alexander Spiegelman. 2021. All You Need is DAG. In *Proc. 42nd ACM Symposium on Principles of Distributed Computing (PODC)*. 165–175.
- [20] Eleftherios Kokoris Kogias, Dahlia Malkhi, and Alexander Spiegelman. 2020. Asynchronous Distributed Key Generation for Computationally-Secure Randomness, Consensus, and Threshold Signatures. In *Proc. ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 1751–1767.
- [21] Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. 2016. The Honey Badger of BFT Protocols. In *Proc. ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 31–42.
- [22] Kartik Nayak, Ling Ren, Elaine Shi, Nitin H Vaidya, and Zhuolun Xiang. 2020. Improved Extension Protocols for Byzantine Broadcast and Agreement. In *Proc. 34th International Symposium on Distributed Computing (DISC)*.
- [23] Arpita Patra. 2011. Error-free Multi-valued Broadcast and Byzantine Agreement with Optimal Communication Complexity. In *Proc. 15th International Conference On Principles Of Distributed Systems (OPODIS)*. 34–49.
- [24] Irving S Reed and Gustave Solomon. 1960. Polynomial Codes over Certain Finite Fields. *J. Soc. Indust. Appl. Math.* 8, 2 (1960), 300–304.
- [25] Michael K. Reiter. 1994. Secure Agreement Protocols: Reliable and Atomic Group Multicast in Rampart. In *Proc. 2nd ACM Conference on Computer and Communications Security (CCS)*. 68–80.
- [26] Thomas Yurek, Licheng Luo, Jaiden Fairoze, Aniket Kate, and Andrew Miller. 2022. hbACSS: How to Robustly Share Many Secrets. In *Proc. 29th Annual Network and Distributed System Security Symposium (NDSS)*.
- [27] Haibin Zhang and Sisi Duan. 2022. PACE: Fully Parallelizable BFT from Reproducible Byzantine Agreement. In *Proc. ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 3151–3164.