



Near-Optimal Communication Byzantine Reliable Broadcast Under a Message Adversary

Downloaded from: <https://research.chalmers.se>, 2026-04-21 07:04 UTC

Citation for the original published paper (version of record):

Albouy, T., Frey, D., Gelles, R. et al (2025). Near-Optimal Communication Byzantine Reliable Broadcast Under a Message Adversary. Leibniz International Proceedings in Informatics, LIPIcs, 324. <http://dx.doi.org/10.4230/LIPIcs.OPODIS.2024.14>

N.B. When citing this work, cite the original published paper.

Near-Optimal Communication Byzantine Reliable Broadcast Under a Message Adversary

Timothé Albouy  



Univ Rennes, Inria, CNRS, IRISA,
35042 Rennes-cedex, France

Ran Gelles  



Bar-Ilan University, Ramat Gan, Israel

Michel Raynal  

Univ Rennes, Inria, CNRS, IRISA,
35042 Rennes-cedex, France

François Taïani  

Univ Rennes, Inria, CNRS, IRISA,
35042 Rennes-cedex, France

Davide Frey  


Univ Rennes, Inria, CNRS, IRISA,
35042 Rennes-cedex, France

Carmit Hazay  

Bar-Ilan University, Ramat Gan, Israel

Elad Michael Schiller  

Chalmers University of Technology,
Gothenburg, Sweden

Vassilis Zikas  

Georgia Institute of Technology,
Atlanta, GA, USA

Abstract

We address the problem of Reliable Broadcast in asynchronous message-passing systems with n nodes, of which up to t are malicious (faulty), in addition to a *message adversary* that can drop some of the messages sent by correct (non-faulty) nodes. We present a Message-Adversary-Tolerant Byzantine Reliable Broadcast (MBRB) algorithm that communicates $\mathcal{O}(|m| + n\kappa)$ bits per node, where $|m|$ represents the length of the application message and $\kappa = \Omega(\log n)$ is a security parameter. This communication complexity is optimal up to the parameter κ . This significantly improves upon the state-of-the-art MBRB solution (Albouy, Frey, Raynal, and Taïani, TCS 2023), which incurs communication of $\mathcal{O}(n|m| + n^2\kappa)$ bits per node. Our solution sends at most $4n^2$ messages overall, which is asymptotically optimal. Reduced communication is achieved by employing coding techniques that replace the need for all nodes to (re-)broadcast the entire application message m . Instead, nodes forward authenticated fragments of the encoding of m using an erasure-correcting code. Under the cryptographic assumptions of threshold signatures and vector commitments, and assuming $n > 3t + 2d$, where the adversary drops at most d messages per broadcast, our algorithm allows at least $\ell = n - t - (1 + \epsilon)d$ (for any arbitrarily low $\epsilon > 0$) correct nodes to reconstruct m , despite missing fragments caused by the malicious nodes and the message adversary.

2012 ACM Subject Classification Theory of computation → Distributed algorithms

Keywords and phrases Asynchronous message-passing, Byzantine fault-tolerance, Message adversary, Reliable broadcast, Erasure-correction codes, Threshold signatures, Vector commitments

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2024.14

Funding This work was partially supported by the French ANR projects ByBloS (ANR-20-CE25-0002-01) and PriCLESS (ANR-10-LABX-07-81), devoted to the modular design of building blocks for large-scale Byzantine-tolerant multi-users applications. Research supported in part by the United States-Israel Binational Science Foundation (BSF) through Grant No. 2020277.

Acknowledgements R. Gelles would like to thank Paderborn University and CISPA – Helmholtz Center for Information Security for hosting him while part of this research was done.

1 Introduction

Reliable Broadcast allows n asynchronous nodes to agree eventually on a message sent by a designated node, the *sender*, despite the possible malicious behavior by some nodes and the transmission network. Reliable broadcast plays a crucial role in key applications,



© Timothé Albouy, Davide Frey, Ran Gelles, Carmit Hazay, Michel Raynal, Elad Michael Schiller, François Taïani, and Vassilis Zikas;

licensed under Creative Commons License CC-BY 4.0

28th International Conference on Principles of Distributed Systems (OPODIS 2024).

Editors: Silvia Bonomi, Letterio Galletta, Etienne Rivière, and Valerio Schiavoni; Article No. 14; pp. 14:1–14:29



Leibniz International Proceedings in Informatics

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

including consensus algorithms, replication, event notification, and distributed file systems. These systems sometimes require broadcasting large messages or files (*e.g.*, permissioned blockchains), and thus, reducing the communication overhead to a minimum is an important aspect of achieving scalability. In that vein, this work aims at providing *communication efficient* solutions for the task of reliable broadcast in the presence of node and link faults.

Byzantine nodes [20, 29] are faulty nodes that are assumed to act cooperatively in an arbitrary manner to hinder the non-faulty nodes (also known as correct nodes) from reaching an agreement on the value of a sent message. These faulty nodes can manifest in various ways, such as delivering fake messages that were never sent, altering the payload of messages sent by any faulty nodes, delaying message delivery, or even omitting messages altogether. Also, a Byzantine failure can present itself differently to different nodes.

Solving reliable broadcast in the presence of Byzantine nodes (a problem known as BRB for Byzantine Reliable Broadcast [10]) has been extensively studied for at least four decades. Bracha [11, 12] in particular showed that BRB could be implemented in asynchronous networks as soon as the number t of Byzantine nodes is limited to be less than a third of the nodes. This seminal result has since been followed by hundreds of works, with a various range of optimizations and tradeoffs between different parameters such as resilience, efficiency, and communication; see [38] for an excellent book on this topic.

A significant challenge to reliable broadcast algorithms arises when the message-passing system is unreliable and possibly cooperates with the Byzantine nodes. Link faults [41, 43] give Byzantine nodes (potentially limited) control over certain network links, enabling them to omit or corrupt messages (an ability captured under the umbrella term *message adversary* [37]). This work focuses on a specific type of *message adversary* [37] that can only omit messages sent by correct nodes, but that cannot alter their content. This message adversary abstracts cases related to *silent churn*, where nodes may voluntarily or involuntarily disconnect from the network without explicitly notifying other nodes. During disconnection, the adversary causes correct nodes to pause the execution of their algorithm temporarily and resume upon reconnecting. In the message adversary model, correct nodes may miss messages sent over reliable communication channels by other nodes while they are disconnected, as there is no explicit notification about the message omission.

Problem overview. We assume n nodes over an asynchronous network, where a message can be delayed for an arbitrary yet finite amount of time (unless omitted by the message adversary). We assume the existence of t Byzantine nodes and a message adversary capable of omitting up to d messages per node's broadcast. To be more precise, a node communicates through a `comm` primitive (or a similar multicast/unicast primitive that targets a dynamically defined subset of nodes), which results in the transmission of n messages, with each node being sent one message, including the sender. The message adversary can omit messages in transit to a subset of at most d correct nodes. The adversary is only limited by the size of that subset. For instance, between different `comm` invocations, the adversary can modify the set of correct nodes to which messages are omitted. Furthermore, a designated sender node holds a message m that it wishes to broadcast to all the nodes.

An algorithm that satisfies the requirements of reliable broadcast despite Byzantine nodes and a message adversary is called a *Message-adversary Byzantine Reliable Broadcast* (MBRB) algorithm. The detailed version of MBRB's requirements was formulated in [4], see Section 2. We informally summarize them here as follows. **(1)** For any sender invoking the broadcast algorithm, no two correct nodes deliver m and m' , such that $m' \neq m$. **(2)** For any sender invoking the broadcast algorithm, either zero or at least ℓ correct nodes will deliver m .

The quantity ℓ might depend on the adversary's power, *i.e.*, on t and d . **(3)** If a correct node delivers some message m from a *correct sender*, this correct sender has broadcast m previously and at least ℓ correct nodes will deliver it.

Background. Albouy, Frey, Raynal, and Taïani [4] recently proposed a Message-adversary Byzantine Reliable Broadcast algorithm (denoted AFRT for short) for asynchronous networks that withstands the presence of t Byzantine nodes and a message adversary capable of omitting up to d messages per node's broadcast. AFRT guarantees the reliable delivery of any message when $n > 3t+2d$. Moreover, they demonstrate the necessity of this bound on the number of Byzantine nodes and omitted messages, as no reliable broadcast algorithm exists otherwise.

One caveat of AFRT regards its communication efficiency. While it achieves an optimal number of $\mathcal{O}(n^2)$ messages, and an optimal delivery power $\ell = n - t - d$, each node's communication requires $\mathcal{O}(n \cdot (|m| + n\kappa))$ bits, where $|m|$ represents the number of bits in the broadcast message and κ is the length of the digital signatures used in their algorithm. In the current work, we design an algorithm that significantly reduces the communication cost per node while preserving the total number of messages communicated. Our solution features at most $4n$ messages per correct node (corresponding to $4n^2$ messages overall), and only $\mathcal{O}(|m| + n\kappa)$ bits per correct node. Overall, $\mathcal{O}(n|m| + n^2\kappa)$ bits are communicated by correct nodes. This bound is tight (up to the size of the signature κ) for deterministic algorithms using signatures [18, 33], as every correct node must receive the message m , and as the reliable broadcast of a single bit necessitates at least $\Omega(n^2)$ messages [21].

Contributions and techniques. This paper is the first to present an MBRB algorithm able to tolerate a hybrid adversary combining t Byzantine nodes and a Message Adversary of power d , while providing optimal Byzantine resilience, near-optimal communication, and near-optimal delivery power ℓ .

► **Theorem 1.1** (Main, informal). *For any $\varepsilon > 0$, there exists an efficient MBRB algorithm, such that every message m broadcast via this scheme is delivered by at least $\ell = n - t - (1 + \varepsilon)d$ correct nodes under the assumption $n > 3t + 2d$. Each correct node communicates no more than $4n$ messages and $\mathcal{O}(|m| + n\kappa)$ bits overall, where $|m|$ is the length of the message m .*

The above asymptotic communication complexity holds assuming a sufficiently long message m . Further, $n - t - d$ is a natural upper bound on the delivery power ℓ of any MBRB algorithm. This bound arises from the power of the message adversary to isolate a subset of the correct parties of size d , and omit all messages sent to this subset. Our solution obtains a delivery power ℓ that is as close to the limit as desired, at the cost of increasing communication (through the hidden constants in the asymptotic $\mathcal{O}(\cdot)$ term, which depends on ε). Finally, $n > 3t + 2d$ is a necessary condition to implement MBRB under asynchrony [3], thus making our solution optimal in terms of Byzantine resilience.

The starting point of our algorithm is the AFRT algorithm [4]. This algorithm achieves all the desired MBRB properties (Definition 2.3), albeit with a large communication cost of at least $n^2|m|$ bits overall. This communication cost stems from the re-emission strategy used by AFRT. In the AFRT algorithm, the sender first disseminates the message m to all nodes. To counter a possibly faulty sender, each node that receives m signs it and forwards it to the entire network, along with its own signature and any other signature observed so far for that message. This re-broadcasting step leads to $n^2|m|$ bits of communication.

14:4 Near-Optimal Communication BRB Under a Message Adversary

In order to reduce the communication costs, we apply a coding technique, inspired by an approach by Alhaddad *et al.* [5]. Instead of communicating the message m directly, the sender first encodes the message using an error-correction code and “splits” the resulting codeword between the nodes, so that each node receives one fragment of size $\mathcal{O}(|m|/n)$ bits. Now, each node needs to broadcast only its fragment of the message rather than the entire message. This reduced per-node communication effectively reduces the overall communication for disseminating the message itself to $n|m|$ bits.

Due to the message adversary and the actions of the Byzantine nodes, some of the fragments might not arrive at their destination. Error-correction codes have the property that the message m can be reconstructed from any sufficiently large subset of the fragments. But Byzantine nodes can do even worse, namely, they can propagate an incorrect fragment. Correct nodes cannot distinguish correct fragments from incorrect ones (at least, not until enough fragments are collected, and the message is reconstructed). Without this knowledge, correct nodes might assist the Byzantine nodes in propagating incorrect fragments, possibly harming the correctness and/or performance of the algorithm. To prevent this, the sender could sign each fragment that it sends. A node receiving a fragment could then verify that it is correctly signed by the sender and ignore it otherwise. The drawback of this solution is that only the sender can generate signatures for fragments.

In our MBRB algorithm, we rely on correct nodes that have already reconstructed the correct message to disseminate its fragments to the nodes that have not received any (say, due to the message adversary). In principle, when a node reconstructs the correct message, it can generate the codeword and obtain all the fragments, even if it did not receive some of them beforehand. However, the node cannot generate the sender’s signature for the fragments it generated by itself. Because of this, the node cannot relay these fragments to the other nodes, potentially leading to a reduced delivery power ℓ .

We avert this issue by exploiting vector commitments [15]. This cryptographic primitive generates a unique short digest C for any input vector of elements V . Additionally, it generates succinct proofs of inclusion for each element in V . In our system, the fragments of the (coded) message m form the vector V , and the inclusion proofs replace the need to sign each fragment separately. In more detail, every fragment of the codeword communicated by some node is accompanied by two pieces of information: the commitment C for the vector V containing all fragments of m , and a proof of inclusion showing that the specific fragment indeed belongs to V (see Section 2 for a formal definition of these properties). The sender signs only the commitment C . This means that Byzantine nodes cannot generate an incorrect fragment and a proof that will pass the verification, since they cannot forge the sender’s signature on C . Yet, given the message m , generating a proof of inclusion for any specific fragment can be done by any node. The vector commitment on the message m creates the same commitment C and the same proofs of inclusion generated by the sender. These could then be propagated to any other node along with the sender’s signature on C .

To complete the description of our MBRB algorithm, we mention that, similar to AFRT, our algorithm tries to form a quorum of signatures on some specific vector commitment C . In parallel, nodes collect fragments they verify as part of the message whose vector commitment is C . Once a node collects enough signatures (for some C) and at the same time obtains enough message fragments that are proven to belong to the same C , the node can reconstruct m and deliver (accept) it. At this point, the node also disseminates the quorum of signatures along with (some of) the fragments. This allows other correct nodes to reconstruct the message and verify a quorum has been reached. In fact, the dissemination of fragments, including fragments that this node did not have before reconstructing the message, is a crucial step in amplifying the number of nodes that deliver m to our stated level of $\ell = n - t - (1 + \varepsilon)d$. See the full description of the MBRB algorithm in Section 3.

Although our algorithm builds quorums on commitments, it departs substantially from the BRB algorithm proposed by Das, Xiang, and Ren [18], which avoids signatures and relies on hashes only. Their solution provides an overall communication complexity in $\mathcal{O}(n|m| + n^2\kappa)$ that is optimal up to the κ parameter. Following the sender’s initial dissemination of message m , their proposal runs Bracha’s algorithm on the hash value of the broadcast message to ensure agreement. Unfortunately, when used with a message adversary, Bracha’s algorithm loses the sub-optimal Byzantine resilience $n > 3t + 2d$ that AFRT and our solution provide, which is why the solution presented in this paper avoids it. (See Section C for a more detailed discussion of why this is so.) Due to the page limit, our algorithm’s complete analysis and proof details appear in the Appendix.

Related work. *Byzantine reliable broadcast* (BRB) can be traced back to Pease, Shostak, and Lamport [34], which considered the particular case of synchronous message-passing systems. Since then, solutions for reliable broadcast, together with the related *consensus* problem, have been considered for many distributed systems [38]. In asynchronous systems, BRB solutions can be traced back to Bracha and Toueg [11, 12]. Recent advances [30, 27, 9, 23] in Byzantine fault-tolerant (BFT) solutions to the problem of BRB include the above-mentioned AFRT [4], which safeguards also against the message adversary. Our solution features substantially lower communication than AFRT, without harming the other properties, *e.g.*, the number of messages or the delivery power.

Computation in networks with link faults, namely, with Byzantine links was considered by Santoro and Widmayer [41, 42], who discussed various types of Byzantine actions, *e.g.*, omitting messages, corrupting messages, injecting messages, and combinations thereof. The works of [6, 43] focus on the case of reliable broadcast with such faults. In [35], Pelc proved that robust communication is feasible over graphs whose edge-connectivity is more than $2f$, assuming the number of Byzantine links is bounded by f . This is also implied by the work of Dolev [20]. Censor-Hillel *et al.* [16] and Frei *et al.* [24] show that any computation can be performed when all links suffer arbitrary substitution faults (but no crashes), given that the network is 2-edge connected. When all links suffer corruption, but their overall amount is restricted, any computation can be reliably performed by Hoza and Schulman [28], for synchronous networks where the topology is known, or by Censor-Hillel, Gelles, and Haeupler [17], for asynchronous networks with unknown topology, see also [25].

Settings that allow Byzantine nodes in addition to faulty links were considered [7, 8, 19, 26, 36, 46]. Building on [4], the algorithm was also extended to signature-free asynchronous systems [3], albeit with lower delivery guarantees, and a weaker resilience with respect to d and t . We consider fully connected asynchronous networks, Byzantine nodes, and omission only link failures. But, our MBRB is the first, to the best of our knowledge, to offer a near-optimal communication (up to the length of the signature, κ) and delivery power ℓ .

Coding techniques are implemented to minimize the dissemination costs associated with message transmission across the network, ensuring the ability to reconstruct data in the event of node failures or adversarial compromises. In the context of Blockchains, significant contributions have been made by Cachin and Tessaro [14] as well as Cachin and Poritz in SINTRA [13], followed by its successors such as HoneyBadgersBFT [32], BEAT [22], DispersedLedger [47] and Alhaddad *et al.* [5]. These solutions leverage digital signatures and coding techniques to provide a balanced and reliable broadcast. Our work contributes to the advancement of the state of the art in the field of coded reliable broadcast by offering improved fault-tolerance guarantees that are stronger than the aforementioned solutions.

2 Preliminaries

General notations and conventions. For a positive integer n , let $[n]$ denote the set $\{1, 2, \dots, n\}$. A sequence of elements (x_1, \dots, x_n) is shorthand as $(x_i)_{i \in [n]}$. We use the symbol ‘-’ to indicate any possible value. That is, $(h, -)$ means a tuple where the second index includes any arbitrary value which we do not care about. All logarithms are base 2.

Nodes and Network. We focus on asynchronous message-passing systems that have no guarantees of communication delay. The system consists of a set, $\mathcal{P} = \{p_1, \dots, p_n\}$, of n fail-prone nodes that cannot access a clock or use timeouts. We identify node i with p_i .

Communication means. Any ordered pair of nodes $p_i, p_j \in \mathcal{P}$ has access to a communication channel, $channel_{i,j}$. Each node can send messages to all nodes (possibly by sending a different message to each node). That is, any node, $p_i \in \mathcal{P}$, can invoke the transmission macro, $\text{comm}(m_1, \dots, m_n)$, that communicates the message m_j to p_j over $channel_{i,j}$. The message m_j can also be empty, in which case nothing will be sent to p_j . However, in our algorithms, all messages sent in a single comm activation will have the same length. Furthermore, when a node sends the same message m to all nodes, we write $\text{broadcast}(m) = \text{comm}(m, m, \dots, m)$ for shorthand. We call each message m_j transmitted by the protocol an *implementation message* (or simply, a *message*) to distinguish such messages from the *application-level* messages, *i.e.*, the one the sender wishes to broadcast.

Byzantine nodes. Faulty nodes are called *Byzantine*, and their adversarial behavior can deviate from the proposed algorithm in any manner. For example, they may crash or send fake messages. Their ability to communicate and collude is unlimited. They might perform any arbitrary computation, and we assume their computing power is at least as strong as that of non-faulty nodes, yet not as strong as to undermine the security of the cryptographic signatures we use. We assume that at most t nodes are faulty, where t is a value known to the nodes. Non-faulty nodes are called *correct nodes*. The set of correct nodes contains c nodes where $n - t \leq c \leq n$. The value of c is unknown.

Faulty nodes may deviate arbitrarily from the correct implementation of $\text{comm}(\cdot)$. For instance, they may unicast messages to only a subset of the nodes in \mathcal{P} . As mentioned, each pair of nodes can communicate using $channel_{i,j}$. While $channel_{i,j}$ is assumed to be a reliable channel that is not prone to message corruption, duplication, or the creation of fake messages that were never sent by nodes; the *message adversary* [41, 43, 37], which we specify next, has a limited ability to cause message loss.

Message adversary. This entity can remove implementation messages from the communication channels used by correct nodes when they invoke $\text{comm}(\cdot)$. More precisely, during each activation of $\text{comm}(m_1, \dots, m_n)$, the adversary has the discretion to choose up to d messages from the set $\{m_i\}$ and eliminate them from the corresponding communication channels where they were queued. We assume that the adversary has full knowledge of the contents of all messages $\{m_i\}$, and thus it makes a worst-case decision as to which messages to eliminate.

The failures injected by a message adversary differ from those of classical sender and/or receiver omissions in that they are *mobile*. *I.e.*, they are not pinned to a set of particular nodes but may move between any correct nodes during the same execution since they are defined per invocation of $\text{comm}(\cdot)$. In particular, no node is immune to the message adversary.

Note that for the case of $d = 0$, the adversary is as weak as the common settings in which all communication channels are reliable (since no message is ever lost) and t nodes are Byzantine. Assumption 2.1 limits the adversary’s power to avoid network partitions. As mentioned above, this is necessary for any MBRB algorithm [4].

► **Assumption 2.1** (Adversary-power-assumption). $n > 3t + 2d$.

Since the message adversary can omit all implementation messages that are sent to a given set $D \subseteq \mathcal{P} : |D| = d$, we know that ℓ , the number of correct nodes that are guaranteed to output the broadcast m correctly, must satisfy $\ell \leq c - d$.

Such a hybrid fault model has been studied in the past in synchronous networks [8], but has remained little studied in an asynchronous setting, except for the work of Schmid and Fetzer [44] that limits itself to round-based algorithms (unlike us) and does not cover full disconnections of correct nodes (which we do). Modeling full disconnections is relevant as this captures correct nodes that remain disconnected for long periods. In addition to bounding the maximal number of outgoing messages that a correct sender might lose (as we do), the model proposed by Schmid and Fetzer [44] also bounds the maximal number of incoming messages that any correct node might miss. This adds an elegant symmetry to the model but poses significant challenges when considering asynchronous networks, as there is no obvious scope on which to limit the number of incoming messages missed by a given node. Schmid and Fetzer [44] therefore restrict the fault model to round-based algorithms. By contrast, our model allows for algorithms that do not follow this structure.

Error Correction Codes. A central tool used in our algorithm is an error-correction code (ECC) [40]. Intuitively speaking, an ECC takes a message as input and adds redundancy to create a codeword from which the original message can be recovered even when parts of the codeword are corrupted. In this work, we focus on *erasures*, a corruption that replaces a symbol of the codeword with a special erasure mark \perp .

Let \mathbb{F} denote a finite field whose size we set later, and let \perp be a special symbol $\perp \notin \mathbb{F}$. Given two strings of the same length, $x, y \in \mathbb{F}^n$, their *Hamming distance* is the number of indices where they differ, $\Delta(x, y) = |\{i \in [n] \mid x_i \neq y_i\}|$. Given a subset $I \subseteq [n]$, we denote by $x_I \in \mathbb{F}^{|I|}$ the string x restricted to the indices in I .

To avoid confusion with global parameters, we denote the ECC-specific parameters by using a bar (e.g., \bar{x}). An *error-correction code* is a function $\text{ECC} : \mathbb{F}^{\bar{k}} \rightarrow \mathbb{F}^{\bar{n}}$, with *rate* $\bar{r} = \bar{k}/\bar{n}$, and *distance* $\bar{d} = \min_{x, y \in \mathbb{F}^{\bar{k}}, x \neq y} \Delta(\text{ECC}(x), \text{ECC}(y))$. The Singleton bound determines that $\bar{d} \leq \bar{n} - \bar{k} + 1$, and when the equality holds, the code is said to be maximum distance separable (MDS). A prominent example of MDS codes is Reed-Solomon (RS) codes [39], which exist for any \bar{k}, \bar{n} , and $|\mathbb{F}| \geq \bar{n}$. Such codes can be efficiently encoded and decoded [40]. The erasure correction capabilities of a code depend on its distance, as given by the following fact.

► **Fact 2.2** (Erasure Correction Capability). *Any error-correction code of distance \bar{d} can recover up to $\bar{d} - 1$ erasures. That is, for any $y \in (\mathbb{F} \cup \{\perp\})^{\bar{n}}$, let $E = \{i \mid y_i = \perp\}$ the set of erased indices. Then, if $|E| < \bar{d}$, there is at most a single $x \in \mathbb{F}^{\bar{k}}$ such that $y_{[\bar{n}] \setminus E} = \text{ECC}(x)_{[\bar{n}] \setminus E}$.*

Cryptographic Primitives. Our algorithm relies on cryptographic assumptions. We assume that the Byzantine nodes are computationally bounded with respect to the security parameter, denoted by κ . That is, all cryptographic algorithms are polynomially bounded in the input 1^κ . We assume that $\kappa = \Omega(\log n)$. We further assume the availability of Public Key Infrastructure (PKI), which is the setting that assumes each node is assigned a pair of public/private keys generated according to some standard key-generation algorithm. Further, at the start of the computation, each node holds its own private key and the public key of all other parties. This setting implies private and authenticated channels. In particular, each node has public and private keys to support the following cryptographic primitives.

Threshold signatures. In a (τ, n) threshold signature scheme [45], at least τ out of all n nodes (the threshold) produce individual *signatures shares* σ for the same message m , which are then aggregated into a fixed-size *threshold signature* Σ . Verifying Σ does not require the public keys of the signers; one needs to use a single system-wide public key, the same for all threshold signatures produced by the scheme. This system public key, known to everyone, is generated during the system setup phase and distributed through the PKI.

Formally, we define a (τ, n) threshold signature scheme as a tuple of (possibly randomized) algorithms $\text{TSIG} = (\text{ts_sign_share}, \text{ts_verify_share}, \text{ts_combine}, \text{ts_verify})$. The signing algorithm executed by node p_i (denoted, ts_sign_share_i) takes a message m (and implicitly a private key) and produces a signature $\sigma = \text{ts_sign_share}_i(m)$. The share verification algorithm takes a message m , a signature share σ_i , and the identity i of its signer p_i (and implicitly p_i 's public key), and outputs a single bit, $b = \text{ts_verify_share}(m, \sigma_i, i) \in \{\text{valid}, \text{invalid}\}$, which indicates whether the signature share is valid or not. The combination algorithm takes a set sigs of τ valid signature shares produced by τ out of n nodes and the associated message m (and implicitly the system public key) and outputs a threshold signature $\Sigma = \text{ts_combine}(\text{sigs})$. The threshold signature verification algorithm takes a message m and a threshold signature Σ (and implicitly the system public key) and outputs a single bit $b = \text{ts_verify}(m, \Sigma) \in \{\text{valid}, \text{invalid}\}$, indicating if the threshold signature is valid or not.

We require the conventional robustness and unforgeability properties for threshold signatures. This scheme is parameterized by a security parameter κ , and the size of signature shares and threshold signatures, $|\sigma| = |\Sigma| = \mathcal{O}(\kappa)$, is independent of the size of the signed message, m . In our algorithm, we take $\tau = \lfloor \frac{n+t}{2} \rfloor + 1$ (i.e., the integer right above $\frac{n+t}{2}$).

Vector commitments (VC). A *vector commitment* (VC) is a short digest C for a vector of elements V , upon which a user can then generate a *proof of inclusion* π (sometimes called *partial opening*) of some element in V without disclosing the other elements of V to the verifier: the verifier only needs C , π , the element, and its index in the vector to verify its inclusion in V . A Merkle tree [31] is a notable example of vector commitment, although with several sub-optimal properties. For example, for a hash size of κ , a Merkle proof of inclusion is of $\mathcal{O}(\kappa \log |V|)$ bits, which is significantly larger than modern schemes such as Catalano-Fiore vector commitments [15], which produce proofs of inclusion with an optimal size of $\mathcal{O}(\kappa)$ bits. In our construction, we use these optimal VC schemes (such as Catalano-Fiore's), which provide commitments and proofs in $\mathcal{O}(\kappa)$ bits. The VC scheme provides two operations, parameterized by the security parameter κ : $\text{vc_commit}(\cdot)$ and $\text{vc_verify}(\cdot)$, that work as follows. For any vector of strings $V = (x_1, \dots, x_n) = (x_i)_{i \in [n]}$, the function $\text{vc_commit}(V) \rightarrow (C, \pi_1, \dots, \pi_n)$ returns C , the commitment, and every π_i , the proof of inclusion for x_i . The following hold.

1. **Proof of inclusion (Correctness):** Let $(C, (\pi_i)_{i \in [n]}) = \text{vc_commit}((x_1, \dots, x_n))$. Then for any $i \in [n]$, it holds that $\text{vc_verify}(C, \pi_i, x_i, i) = \text{valid}$.
2. **Collision-resistance (Binding):** For any $j \in [n]$ and any randomized algorithm A taking $(x_i)_{i \in [n]}$ and $(C, (\pi_i)_{i \in [n]}) = \text{vc_commit}(x_1, \dots, x_n)$ as input, $\Pr[A \text{ outputs } (x'_j, \pi'_j, j) \wedge \text{vc_verify}(C, \pi'_j, x'_j, j) = \text{valid}] < 2^{-\Omega(\kappa)}$. Namely, it is difficult to generate $x'_j \neq x_j$ and a valid proof π'_j for the same commitment C .

We omit the traditional Hiding property of VC schemes (a commitment does not leak information on the vector [2]) as it is unneeded in our algorithm. We also implicitly assume that the vc_commit operation is deterministic: it always returns the same commitment C given the same input vector V , no matter the calling node p_i . This is the case for Catalano-Fiore's scheme [15], which does not use random salt.

■ **Table 1** Notations used by Algorithm 2.

n	total number of nodes
t	upper bound on the number of Byzantine nodes
d	removal power of the message adversary
k	reconstruction threshold of the erasure code, k out of n
p_s	the designated sending node (with identity s)
σ_i	signature share by node p_i
sig_i	the pair (σ_i, i)
$sigs, sigs_i$	sets of (signature share, id) pairs
Σ	threshold signature (TS)
τ	threshold of the TS scheme (set to $\tau = \lfloor \frac{n+t}{2} \rfloor + 1$ in our algorithm)
m, m', m_i	application messages
C, C', C_m	vector commitments
\tilde{m}_i	i^{th} message fragment of application message m
$\tilde{\pi}_i$	proof of inclusion of fragment \tilde{m}_i

Specification of the MBRB primitive. MBRB's objective is to guarantee a reliable broadcast, meaning it aims to ensure that a bounded minimum number of correct nodes ultimately deliver the broadcast messages to the application while upholding specific safety and liveness criteria. This assurance holds even when confronted with Byzantine faults and a message adversary capable of selectively suppressing messages.

An MBRB algorithm provides the MBRB-broadcast and MBRB-deliver operations. The following specification is presented in a single-shot, single-sender version, where p_s is the sending node. A multi-sender, multi-shot version can be derived by adding the sender's identifier and a running sequence number to messages and signatures. Nodes invoke the MBRB-deliver operation to deliver (to the application layer) messages broadcast by p_s .

Definition 2.3 specifies the safety and liveness properties of MBRB. Safety ensures that messages are delivered correctly without spurious messages, duplication, or duplicity. Liveness guarantees that if a correct node broadcasts a message, it will eventually be delivered by at least one correct node (MBRB-Local-delivery). If a correct node delivers a message from any specific sender, that message will eventually be delivered by a sufficient number, ℓ , of correct nodes (MBRB-Global-delivery), where ℓ is a measure of the *delivery power* of the MBRB object. The parameter ℓ might depend on the adversary's power, *i.e.*, on t and d . Since the message adversary can omit all implementation messages that are sent to an unknown set $D \subseteq \mathcal{P} : |D| = d$, we know that $\ell \leq c - d$.

► **Definition 2.3.** *An MBRB is an algorithm that satisfies the following properties.*

- **MBRB-Validity.** *If p_s is correct and a correct node, p_i , MBRB-delivers an application message m , then, node p_s has MBRB-broadcast m (before that MBRB-delivery).*
- **MBRB-No-duplication.** *A correct node p_i MBRB-delivers at most one application message m .*
- **MBRB-No-duplicity.** *No two different correct nodes MBRB-deliver different application messages from node p_s .*
- **MBRB-Local-delivery.** *Suppose p_s is correct and MBRB-broadcasts an application message m . At least one correct node, p_j , eventually MBRB-delivers m from node p_s .*
- **MBRB-Global-delivery.** *Suppose a correct node, p_i , MBRB-delivers an application message m from p_s . Then, at least ℓ correct nodes MBRB-deliver m from p_s .*

3 The Coded-MBRB algorithm

The proposed solution, named Coded MBRB (Algorithm 2), allows a distinguished sender p_s (known to everyone) to disseminate one specific application message m . In Section A, we discuss how to extend this algorithm so that it implements a general MBRB algorithm, allowing any node to be the sender, as well as allowing multiple instances of the MBRB, either with the same or different senders, to run concurrently. In the algorithm, we instantiate the threshold signature scheme with the threshold value set to $\tau = \lfloor \frac{n+t}{2} \rfloor + 1$ (see Section 2).

Algorithm 2 introduces the MBRB-BROADCAST(m) operation, which takes message m and disseminates it reliably to a minimum bound of correct nodes, denoted ℓ . That is, after executing Algorithm 2, and assuming a correct sender, at least ℓ correct nodes will have invoked the MBRB-DELIVERY(m) procedure, while no correct node will have invoked MBRB-DELIVERY with $m' \neq m$. Table 1 summarizes the notations used by Algorithm 2.

Algorithm description. MBRB-BROADCAST(m) allows the sender p_s to start disseminating the application message, m (line 2). The sender (line 2) starts by invoking COMPUTEFRAGVECCOMMIT(m) (Algorithm 1). This function encodes the message m using an error-correction code, divides it into n fragments and constructs a vector commitment with an inclusion proof for each fragment. The function returns several essential values: the commitment C , and the fragment details (\tilde{m}_j, π_j, j) , which contain the fragment data itself \tilde{m}_j (the j -th part of the codeword ECC(m); see below for detail), a proof of inclusion π_j for that part, and each fragment's respective index j . For ease of reference, let Commitment(m)

■ **Algorithm 1** Helper functions of the Coded MBRB Algorithm (code for p_i).

```

1 Function COMPUTEFRAGVECCOMMIT( $m$ ) is ▷ Computes ECC fragments and VC
2    $\tilde{m} \leftarrow \text{ECC}(m)$  ▷ See “The error-correction code in use” paragraph
3   let  $\tilde{m}_1, \dots, \tilde{m}_n$  be  $n$  equal size fragments of  $\tilde{m}$ 
4    $(C, \pi_1, \dots, \pi_n) \leftarrow \text{vc\_commit}(\tilde{m}_1, \dots, \tilde{m}_n)$  ;
5   return  $(C, (\tilde{m}_j, \pi_j, j)_{j \in [n]})$ 

6 Function ISVALID( $C, fragsSet, sigs, isThreshSig$ ) is ▷ Checks the validity of received msgs
7   ▷ If  $sigs$  is a set of signature shares
8   if  $\neg isThreshSig$  then
9     ▷ Each signature in  $sigs$  must be valid
10    if  $\exists (\sigma_x, x) \in sigs : \neg \text{ts\_verify\_share}(C, \sigma_x, x)$  then return False
11    ▷  $sigs$  must contain the sender's signature
12    if  $(-, s) \notin sigs$  then return False
13    ▷ If  $sigs$  is a threshold signature, check if it aggregates at least  $\tau = \lfloor \frac{n+t}{2} \rfloor + 1$  valid shares
14    else if  $\neg \text{ts\_verify}(C, sigs)$  then return False
15    ▷ Each proof of inclusion in  $fragsSet$  must be valid
16     $fragsSet \leftarrow fragsSet \setminus \{\perp\}$  ▷ Ignoring  $\perp$  values
17    if  $\exists (\tilde{m}_x, \pi_x, x) \in fragsSet : \neg \text{vc\_verify}(C, \pi_x, \tilde{m}_x, x)$  then return False
18    return True

19 Function GETTHRESHSIG( $C$ ) is ▷ Get the TS for  $C$  if it exists,  $\perp$  otherwise
20    $sigs_C \leftarrow \{\text{signatures shares stored by } p_i \text{ for commitment } C\}$ 
21   return  $\Sigma_C \leftarrow \begin{cases} \text{the threshold signature saved for } C & \text{if it exists} \\ \text{else, } \text{ts\_combine}(sigs_C) & \text{if } |sigs_C| > \frac{n+t}{2} \\ \perp & \text{otherwise} \end{cases}$ 

```

represent the commitment C obtained from $\text{COMPUTEFRAGVECCOMMIT}(m)$. This commitment serves as a compact representation of the entire message m . The sender node p_s is responsible for signing the computed commitment C and generating a signature share sig_s (line 2) which includes p_s 's identifier. The sender then initiates m 's propagation by employing the operation comm (line 2), which sends to each node, p_j , an individual message, v_j . The message v_j includes several components: the message type (SEND), the commitment C , the j -th fragment details (\tilde{m}_j, π_j, j) , and the signature share sig_s (line 2) for C .

The rest of the algorithm progresses in two phases, which we describe in turn. The first phase is responsible for message dissemination, which forwards message fragments received by the sender. The other role of this phase is reaching a quorum of nodes that vouch for the same message. A node vouches for a single message by signing its commitment. Nodes collect and store signature shares until it is evident that sufficiently many nodes agree on the same message. The subsequent phase focuses on disseminating the quorum of signature shares so that it is observed by at least ℓ correct nodes, and on successfully terminating while ensuring the delivery of the reconstructed message.

Validating message integrity. The validity of the signatures and inclusion proofs are checked each time a new message is received (at line 2) using the function ISVALID (Algorithm 1). All message types (SEND , FORWARD , and BUNDLE) carry a vector commitment (C or C') and up to two message fragments with their inclusion proofs. Moreover, the SEND and FORWARD types contain up to two signature shares for the provided commitment, and the BUNDLE type contains a threshold signature for the provided commitment. The validation hinges on three key criteria. Every enclosed signature share or threshold signature must be valid and correspond to the accompanying commitment. For SEND or FORWARD messages, the signature share from the designated sending node p_s must be present. All message fragments must contain valid inclusion proofs for the provided commitment. Note that π_i , the proof of inclusion of \tilde{m}_i , does not need to be signed by p_s , as the commitment already is.

Phase I: Message dissemination. This phase facilitates the widespread distribution of the message fragments \tilde{m}_j . Recall that the sender has sent each node a different (encoded) fragment of the message m , however, no node currently holds enough information to retrieve m . The phase also sets the ground for forming a quorum on m . When a node receives a SEND message from the sender, it begins by validating the fragment's authenticity (line 2), and it forwards this fragment to all other nodes by broadcasting a FORWARD message (line 2).

Upon receiving a $\langle \text{SEND}, C', (\tilde{m}_i, \pi_i, i), sig_s \rangle$ message from p_s , the recipient p_i validates the incoming message (line 2). p_i then determines whether it had previously broadcast a FORWARD message at line 2 or signed a commitment C'' from p_s distinct from the currently received C' , in which case the incoming message is discarded (line 2). Otherwise, p_i proceeds to store the received information (line 2), encompassing the fragment \tilde{m}_i and the associated signature share sig_s , linked to the specific commitment C' . We clarify that p_i never stores multiple copies of the same information, *i.e.*, all store operations are to be read as adding an item to a set. Subsequently, p_i generates its own signature share sig_i for the commitment C' , storing it for later utilization (line 2). Node p_i then disseminates all the relevant information, by broadcasting the message $\langle \text{FORWARD}, C', (\tilde{m}_i, \pi_i, i), \{sig_s, sig_i\} \rangle$.

The broadcast of a FORWARD message is instrumental in disseminating information for several reasons. First, up to d nodes might not receive the sender's SEND message. Second, this is the node's way to disseminate its own fragment and signature share for that specific C' .

Upon the arrival of a $\langle \text{FORWARD}, C', \text{fragtuple}_j, sigs_j \rangle$ message from p_j (line 2), the recipient p_i validates the incoming message using the ISVALID function (Algorithm 1), discarding invalid messages (line 2). As for SEND messages, p_i checks if it already signed

■ **Algorithm 2** Phases of the Coded MBRB Algorithm (code for p_i , single-shot, single-sender, threshold for the TS scheme $\tau = \lfloor \frac{n+t}{2} \rfloor + 1$).

```

17 Function MBRBBROADCAST( $m$ ) is  $\triangleright$  only executed by the sender,  $p_s$ 
18    $(C, (\tilde{m}_j, \pi_j, j)_j) \leftarrow \text{COMPUTEFRAgVECCOMMIT}(m)$ 
19    $sig_s \leftarrow (\text{ts\_sign\_share}_s(C), s)$ 
20    $\text{comm}(v_1, \dots, v_n)$  where  $v_j = \langle \text{SEND}, C, (\tilde{m}_j, \pi_j), sig_s \rangle$ 

```

Phase I: Message dissemination

```

21 Upon  $\langle \text{SEND}, C', (\tilde{m}_i, \pi_i, i), sig_s \rangle$  arrival from  $p_s$  do
22   if  $\neg \text{ISVALID}(C', \{(\tilde{m}_i, \pi_i, i)\}, \{sig_s\}, isThreshSig=False)$  then return
23   if  $p_i$  already executed line 2 or signed some commitment  $C'' \neq C'$  then return
24    $sig_i \leftarrow (\text{ts\_sign\_share}_i(C'), i)$  ; store  $\tilde{m}_i, sig_s$ , and  $sig_i$  for  $C'$ 
25   broadcast  $\langle \text{FORWARD}, C', (\tilde{m}_i, \pi_i, i), \{sig_s, sig_i\} \rangle$ 

26 Upon  $\langle \text{FORWARD}, C', fragtuple_j, sigs_j = \{sig_s, sig_i\} \rangle$  arrival from  $p_j$  do
27   if  $\neg \text{ISVALID}(C', \{fragtuple_j\}, sigs_j, isThreshSig=False)$  then return
28   if  $p_i$  already signed some commitment  $C'' \neq C'$  then return
29   store  $sigs_j$  for  $C'$ 
30   if  $fragtuple_j \neq \perp$  then
31      $(\tilde{m}_j, \pi_j, j) \leftarrow fragtuple_j$  ; store  $\tilde{m}_j$  for  $C'$ 
32   if no FORWARD message sent yet then
33      $sig_i \leftarrow (\text{ts\_sign\_share}_i(C'), i)$  ; store  $sig_i$  for  $C'$ 
34     broadcast  $\langle \text{FORWARD}, C', \perp, \{sig_s, sig_i\} \rangle$ 

```

Phase II: Reaching Quorum and Termination

```

35 When  $\left\{ \begin{array}{l} \exists C' : \text{GETTHRESHSIG}(C') \neq \perp \wedge |\{\text{stored } \tilde{m}_j \text{ for } C'\}| \geq k \\ \wedge \text{no message has been MBRB-delivered yet} \end{array} \right\}$  do
36    $m_i \leftarrow \text{ECC}^{-1}(\tilde{m}_1, \dots, \tilde{m}_n)$ ,  $\left\{ \begin{array}{l} \text{where } \tilde{m}_j \text{ are taken from line 35;} \\ \text{when a fragment is missing use } \perp. \end{array} \right.$ 
37    $(C, (\tilde{m}'_j, \pi'_j, j)_j) \leftarrow \text{COMPUTEFRAgVECCOMMIT}(m_i)$ 
38   if  $C' \neq C$  then return
39    $\Sigma_C \leftarrow \text{GETTHRESHSIG}(C)$ 
40    $\text{comm}(v_1, \dots, v_n)$  where  $v_j = \langle \text{BUNDLE}, C, (\tilde{m}'_i, \pi'_i, i), (\tilde{m}'_j, \pi'_j, j), \Sigma_C \rangle$ 
41    $\text{MBRBDELIVER}(m_i)$ 

42 Upon  $\langle \text{BUNDLE}, C', (\tilde{m}'_j, \pi'_j, j), fragtuple'_i, \Sigma \rangle$  arrival from  $p_j$  do
43   if  $\neg \text{ISVALID}(C', \{(\tilde{m}'_j, \pi'_j, j), fragtuple'_i\}, \Sigma, isThreshSig=True)$  then return
44   store  $(\tilde{m}'_j, \pi'_j, j)$  and  $\Sigma$  for  $C'$ 
45   if no BUNDLE message has been sent yet  $\wedge fragtuple'_i \neq \perp$  then
46      $(\tilde{m}'_i, \pi'_i, i) \leftarrow fragtuple'_i$ 
47     store  $(\tilde{m}'_i, \pi'_i, i)$  for  $C'$ 
48     broadcast  $\langle \text{BUNDLE}, C', (\tilde{m}'_i, \pi'_i, i), \perp, \Sigma \rangle$ 

```

a message from p_s with a different commitment C'' , in which case it discards the message (line 2). Subsequently, p_i stores the set of signature shares sig_s linked to the specific commitment C' (line 2) and $fragtuple_j$ contained in this message, if any (line 2). Also, p_i assesses whether a FORWARD message has been previously dispatched. If it has already done so, there is no reason to re-send it, and the processing ends here. Otherwise, similar to above, p_i generates its own signature share sig_i for the commitment C' , and broadcasts the message $\langle \text{FORWARD}, C', \perp, sig_s, sig_i \rangle$. Note that, in this case, p_i is unaware of his own fragment (*i.e.*, it has not received a SEND message, or otherwise it would have already sent a FORWARD message in line 2); therefore it sends the sentinel value \perp instead.

Phase II: Reaching quorum and termination. This phase relies on the GETTHRESHSIG function described in Algorithm 1, which, given a commitment C , either returns a threshold signature for C (received beforehand or aggregating $\tau = \lfloor \frac{n+t}{2} \rfloor + 1$ signature shares stored for C) if it exists, or \perp otherwise. This phase focuses on ensuring that, once a Byzantine quorum (represented by the threshold signature returned by GETTHRESHSIG) and enough message fragments for reconstructing the original message m are gathered, at least ℓ correct nodes deliver m and terminate. Node p_i enters Phase II only when there is a commitment C' for which GETTHRESHSIG returns a valid threshold signature, and p_i stores at least k message fragments. As long as no application message from p_s was delivered (line 2), p_i reconstructs the application message m_i (line 2) using the stored message fragments, and use this message as an input to COMPUTEFRAGVECCOMMIT (line 2), which outputs its commitment $C = \text{Commitment}(m_i)$ along with coded message fragments and proofs of inclusion, $(\tilde{m}'_j, \pi'_j, j)$. Node p_i then ensures that the computed commitment C matches the stored commitment C' (line 2). If this condition holds true, then $m_i = m$ is the message sent by the sender¹, and in particular, p_i now holds *all* the possible fragments for m along with their valid proof of inclusion, including fragments it has never received before! Node p_i then retrieves the threshold signature Σ_C of C using the GETTHRESHSIG function (line 2), and disseminates it along with the message fragments to the rest of the network. In particular, to each p_j in the network, p_i sends a BUNDLE message (line 2) that includes the commitment C , fragment details $(\tilde{m}'_i, \pi'_i, i)$ and $(\tilde{m}'_j, \pi'_j, j)$, and the associated threshold signature Σ_C . After these transmissions, p_i can MBRB-deliver the reconstructed message m_i (line 2).

The parameter k used at line 2 is the number of (valid) fragments sufficient to reconstruct the application message m by the error-correction code ECC. This parameter should be practically selected by the desired ℓ given in Theorem A.12. That is, one needs to set $\varepsilon > 0$ for $\ell = n - t - (1 + \varepsilon)d$ and then choose $k \leq 1 + \frac{\varepsilon}{1+\varepsilon}(n - t - d)$. See details in Section A.

Upon the arrival of a BUNDLE message, *i.e.*, $\langle \text{BUNDLE}, C', (\tilde{m}'_j, \pi'_j, j), fragtuple'_i, \Sigma \rangle$ arriving from p_j (line 2), the recipient p_i validates the received message using the ISVALID function (with the *isThreshSig* parameter set to True to indicate that we verify a threshold signature) and discards invalid messages (line 2). Node p_i proceeds to store the arriving information regarding the message segments $(\tilde{m}'_j, \pi'_j, j)$ and threshold signature Σ for the specific commitment C' (line 2). In the case that no BUNDLE message was sent by p_j and the received $fragtuple'_i$ is nonempty (so p_i learns its fragment, which it stores at line 2, unless already known), p_i broadcasts a $\langle \text{BUNDLE}, C', (\tilde{m}'_i, \pi'_i, i), \perp, \Sigma \rangle$ message (line 2).

The use of \perp values appears also in BUNDLE messages (line 2). A BUNDLE message might contain up to two fragments: the sender's fragment $(\tilde{m}'_i$ in the pseudo-code), which is always included, and the receiver's fragment (\tilde{m}'_j) , which is included only when the sender was

¹ To see why this is needed, consider a Byzantine sender that disseminates fragments $\tilde{m}_1, \dots, \tilde{m}_n$, that do not form a proper codeword.

able to reconstruct the message m (at line 2). The sender's fragments are collected by the receivers and allow reconstruction of the message once enough BUNDLE messages are received. The receiver's fragment allows the receiver to send BUNDLE messages (with its fragment), facilitating the dissemination of both threshold signatures and fragments.

The error-correction code in use. COMPUTEFRAGVECCOMMIT (Algorithm 1) uses an error-correction code at line 1 to encode the application message m , before it is split into n fragments that will be disseminated by the parties. The code uses a fixed parameter k , that can be set later. Our algorithm requires that the ECC will be able to decode the message m from any subset of k fragments out of the n fragments generated at line 1. That is, we need an ECC that can deal with erasures, where the erased symbols are those contained in the $n - k$ missing fragments. To that end, we use a Reed-Solomon code $\text{ECC} : \mathbb{F}^{\bar{k}} \rightarrow \mathbb{F}^{\bar{n}}$ with $\bar{k} > |m|/\log |\mathbb{F}|$. Each fragment contains \bar{n}/n symbols of the codeword, and to be able to recover from $(n - k) \cdot \bar{n}/n$ erased symbols by Theorem 2.2, we can set the code's distance to be $\bar{d} > (n - k) \cdot \bar{n}/n$. Since a Reed-Solomon code is MDS (see Section 2), $\bar{d} \leq \bar{n} - \bar{k} + 1$, and we can set $\bar{n} > \frac{n}{k}(\bar{k} - 1)$. The code will have a constant rate, *i.e.*, $|\text{ECC}(m)| = \mathcal{O}(|m|)$, as long as m is sufficiently long, *i.e.*, $|m| = \Omega(n \log |\mathbb{F}|)$, which implies that $\bar{k} = \Omega(n)$, and as long as $k = \Omega(n)$. Recall also that $|\mathbb{F}| \geq \bar{n}$ is in a Reed-Solomon code.

Analysis. The following main theorem states that our algorithm is correct.

► **Theorem 3.1 (main).** *Assume $n > 3t + 2d$, $k \leq (n - t - 2d)$ and $\varepsilon > 0$. Algorithm 2 implements MBRB with $\ell > n - t - (1 + \varepsilon)d$. Any algorithm activation on the input message m communicates $4n^2$ messages, where each node communicates $\mathcal{O}(|m| + nk)$ bits.*

Due to the page limit, the complete proof appear in Sections A and B. Here, we sketch the proof of the MBRB-Global-delivery property in Theorem 3.2 (assuming the other properties hold) and the communication analysis in Theorem 3.3.

► **Lemma 3.2 (MBRB-Global-delivery).** *If a correct node p_i MBRB-delivers an application message m , then at least $\ell = c - d/(1 - ((k - 1)/(c - d)))$ correct nodes MBRB-deliver m .*

Proof Sketch of Lemma 3.2. Let $C_m = \text{Commitment}(m)$. The proof counts the BUNDLE messages disseminated by correct nodes. If a correct node disseminates a BUNDLE message both at line 2, we only consider the one from line 2. Let B_{send} be the set of correct nodes that disseminate at least one BUNDLE message, let B_{recv} be the set of correct nodes that receive at least one valid BUNDLE message from a correct node during their execution, and let $B_{k,\text{recv}}$ be the set of correct nodes that receive BUNDLE messages from at least k distinct correct nodes. The following holds.

► **Observation 3.2.1.** $c \geq |B_{\text{recv}}| \geq c - d$ and $c \geq |B_{\text{send}}| \geq c - d$.

Proof of Observation 3.2.1. Since B_{send} and B_{recv} contain only correct nodes, trivially $c \geq |B_{\text{send}}|$ and $c \geq |B_{\text{recv}}|$. $|B_{\text{recv}}| \geq c - d$ and $|B_{\text{send}}| \geq c - d$ follow from the definition of the message adversary, and the way the algorithm chains BUNDLE messages. ◀

► **Observation 3.2.2.** $|B_{k,\text{recv}}| \times |B_{\text{send}}| + (k - 1)(|B_{\text{recv}}| - |B_{k,\text{recv}}|) \geq |B_{\text{send}}|(c - d)$.

Proof of Observation 3.2.2. The inequality follows from a counting argument on the overall number of valid BUNDLE messages received by correct nodes from distinct correct senders. In particular, we use the fact that nodes in $B_{\text{recv}} \setminus B_{k,\text{recv}}$ each receives at most $k - 1$ valid BUNDLE messages from distinct correct senders. ◀

► **Observation 3.2.3.** $|B_{k,\text{recv}}| \geq c - d/(1 - ((k-1)/(c-d)))$.

Proof of Observation 3.2.3. The observation is obtained by isolating $B_{k,\text{recv}}$ in Observation 3.2.2, and minimizing the right-hand side to remove the dependence on $|B_{\text{send}}|$. ◀

► **Observation 3.2.4.** *All nodes in $B_{k,\text{recv}}$ MBRB-deliver m .*

Proof of Observation 3.2.4. The observation results from the properties of the vector commitment scheme, the unforgeability of signatures, and the ability of the ECC scheme to reconstruct m from k distinct valid fragments for m . ◀

As all nodes in $B_{k,\text{recv}}$ are correct, the above observations yield the lemma. (The full proof of this lemma can be found in Section A, page 24.) ◀

► **Lemma 3.3.** *Correct nodes collectively communicate at most $4n^2$ messages. Each correct node sends at most $\mathcal{O}(|m| + n\kappa)$ bits. Overall, the system sends at most $\mathcal{O}(n|m| + n^2\kappa)$ bits.*

Proof Sketch of Lemma 3.3. For concision, we use in the following the verb *disseminate* for referring to a correct node that sends a message to all n nodes of the system, whether it be using the broadcast operation or the comm operation.

Let us first consider message complexity. If the sender p_s is correct, it disseminates one SEND message to all nodes at line 2, receives it immediately at line 2, and disseminates one FORWARD message at line 2 (as p_s cannot not pass line 2 afterward). Every other correct node disseminates at most two FORWARD messages at line 2. Moreover, each correct node disseminates at most two BUNDLE messages at line 2. This amounts to at most $4n^2$ messages sent by correct nodes.

Let us now analyze the bit complexity of the algorithm. If the sender p_s is correct, it disseminates a SEND message (line 2) including a fragment of m ($\mathcal{O}(|m|/n)$ bits) with its proof of inclusion ($\mathcal{O}(\kappa)$ bits), a commitment ($\mathcal{O}(\kappa)$ bits), and a signature share ($\mathcal{O}(\kappa)$ bits). Thus, p_s communicates at most $\mathcal{O}(|m| + n\kappa)$ bits in SEND messages. Moreover, a FORWARD message contains a commitment ($\mathcal{O}(\kappa)$ bits), at most one message fragment with its proof of inclusion ($\mathcal{O}(|m|/n + \kappa)$ bits), and two signature shares ($\mathcal{O}(\kappa)$ bits). Hence, every correct node communicates at most $\mathcal{O}(|m| + \kappa n)$ bits in FORWARD messages (line 2 or 34). Additionally, a BUNDLE message contains a commitment ($\mathcal{O}(\kappa)$ bits), at most two message fragments with their proof of inclusion ($\mathcal{O}(|m|/n + \kappa)$ bits), and one threshold signature ($\mathcal{O}(\kappa)$ bits). Therefore, every correct node communicates at most $\mathcal{O}(|m| + \kappa n)$ bits in BUNDLE messages (line 2 or 34). This amounts to a total communication of $\mathcal{O}(|m| + n\kappa)$ bits sent per correct node, and $\mathcal{O}(n|m| + n^2\kappa)$ bits sent overall.

Let us remark that the above analysis also holds in the presence of Byzantine nodes, even if p_s is Byzantine. (The full proof of this lemma can be found in Section B, page 27.) ◀

4 Conclusion

We introduced a Coded MBRB algorithm that significantly improves the state-of-the-art solution. It achieves optimal communication (up to the size of cryptographic parameter κ) while maintaining a high delivery power, *i.e.*, it ensures that messages are delivered by at least $\ell = n - t - (1 + \varepsilon)d$ correct nodes, where $\varepsilon > 0$ is a tunable parameter. The proposed solution is deterministic up to its use of cryptography (threshold signatures and vector commitments). Each correct node sends no more than $4n$ messages and communicates at most $\mathcal{O}(|m| + n\kappa)$ bits, where $|m|$ represents the length of the input message and κ is a security parameter. We note that the algorithm's communication efficiency holds for sufficiently long messages and approaches the natural upper bound on delivery power, $n - t - d$, which accounts for

the message adversary’s ability to isolate a subset of correct nodes. The proposed approach achieves a delivery power ℓ that can be made arbitrarily close to this limit, albeit with a marginal increase in communication costs, which depends on the chosen ε . This work represents a significant advancement in Byzantine Reliable Broadcast, offering a practical solution for robust communication in asynchronous message-passing systems with malicious nodes and message adversaries. One intriguing question is whether it is possible to devise an (M)BRB algorithm without the κ parameter in its communication complexity or the ϵ parameter in its delivery power ℓ , *e.g.*, by leveraging randomization [1] or error-freedom [5].

References

- 1 Ittai Abraham, T.-H. Hubert Chan, Danny Dolev, Kartik Nayak, Rafael Pass, Ling Ren, and Elaine Shi. Communication complexity of byzantine agreement, revisited. *Distributed Comput.*, 36(1):3–28, 2023. doi:10.1007/S00446-022-00428-8.
- 2 Timoth e Albouy, Emmanuelle Anceaume, Davide Frey, Mathieu Gustin, Arthur Rauch, Michel Raynal, and Fran ois Taiani. Asynchronous BFT asset transfer: quasi-anonymous, light, and consensus-free. *CoRR*, arXiv:2405.18072, 2024. doi:10.48550/arXiv.2405.18072.
- 3 Timoth e Albouy, Davide Frey, Michel Raynal, and Fran ois Taiani. A modular approach to construct signature-free BRB algorithms under a message adversary. In *26th International Conference on Principles of Distributed Systems (OPODIS 2022)*, volume 253 of *LIPICs*, pages 26:1–26:23, 2023. doi:10.4230/LIPICs.OPODIS.2022.26.
- 4 Timoth e Albouy, Davide Frey, Michel Raynal, and Fran ois Taiani. Asynchronous Byzantine reliable broadcast with a message adversary. *Theor. Comput. Sci.*, 978:114110, 2023. doi:10.1016/J.TCS.2023.114110.
- 5 Nicolas Alhaddad, Sourav Das, Sisi Duan, Ling Ren, Mayank Varia, Zhuolun Xiang, and Haibin Zhang. Balanced Byzantine reliable broadcast with near-optimal communication and improved computation. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, PODC’22, pages 399–417, 2022. doi:10.1145/3519270.3538475.
- 6 Piotr Berman, Krzysztof Diks, and Andrzej Pelc. Reliable broadcasting in logarithmic time with Byzantine link failures. *Journal of Algorithms*, 22(2):199–211, 1997. doi:10.1006/jagm.1996.0810.
- 7 Martin Biely. An optimal Byzantine agreement algorithm with arbitrary node and link failures. In *IASTED Parallel and Distributed Computing and Systems*, volume 1, pages 146–151, 2003. URL: <https://www.auto.tuwien.ac.at/Projects/W2F/documents/pdcs03.ps.gz>.
- 8 Martin Biely, Ulrich Schmid, and Bettina Weiss. Synchronous consensus under hybrid process and link failures. *Theoretical Computer Science*, 412(40):5602–5630, 2011. doi:10.1016/j.tcs.2010.09.032.
- 9 Silvia Bonomi, J er mie Decouchant, Giovanni Farina, Vincent Rahli, and S ebastien Tixeuil. Practical Byzantine reliable broadcast on partially connected networks. In *ICDCS*, pages 506–516. IEEE, 2021. doi:10.1109/ICDCS51616.2021.00055.
- 10 Gabriel Bracha. Asynchronous byzantine agreement protocols. *Inf. Comput.*, 75(2):130–143, 1987. doi:10.1016/0890-5401(87)90054-X.
- 11 Gabriel Bracha and Sam Toueg. Resilient consensus protocols. In *Proceedings of the Second Annual ACM Symposium on Principles of Distributed Computing*, pages 12–26. ACM, August 1983. doi:10.1145/800221.806706.
- 12 Gabriel Bracha and Sam Toueg. Asynchronous consensus and broadcast protocols. *J. ACM*, 32(4):824–840, 1985. doi:10.1145/4221.214134.
- 13 C. Cachin and J.A. Poritz. Secure INtrusion-tolerant replication on the internet. In *Proceedings International Conference on Dependable Systems and Networks (DSN)*, pages 167–176, 2002. doi:10.1109/DSN.2002.1028897.

- 14 C. Cachin and S. Tessaro. Asynchronous verifiable information dispersal. In *24th IEEE Symposium on Reliable Distributed Systems (SRDS'05)*, pages 191–201, 2005. doi:10.1109/RELDIS.2005.9.
- 15 Dario Catalano and Dario Fiore. Vector commitments and their applications. In *Proc. 16th Int'l Conference on Practice and Theory in Public-Key Cryptography (PKC'13)*, volume 7778 of *Lecture Notes in Computer Science*, pages 55–72. Springer, 2013. doi:10.1007/978-3-642-36362-7_5.
- 16 Keren Censor-Hillel, Shir Cohen, Ran Gelles, and Gal Sela. Distributed computations in fully-defective networks. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, PODC'22, pages 141–150, 2022. doi:10.1145/3519270.3538432.
- 17 Keren Censor-Hillel, Ran Gelles, and Bernhard Haeupler. Making asynchronous distributed computations robust to noise. *Distributed Computing*, 32(5):405–421, 2019. doi:10.1007/s00446-018-0343-5.
- 18 Sourav Das, Zhuolun Xiang, and Ling Ren. Asynchronous data dissemination and its applications. In *CCS*, pages 2705–2721. ACM, 2021. doi:10.1145/3460120.3484808.
- 19 Pallab Dasgupta. Agreement under faulty interfaces. *Information Processing Letters*, 65(3):125–129, 1998. doi:10.1016/S0020-0190(97)00202-0.
- 20 Danny Dolev. The Byzantine generals strike again. *Journal of Algorithms*, 3(1):14–30, 1982. doi:10.1016/0196-6774(82)90004-9.
- 21 Danny Dolev and Rüdiger Reischuk. Bounds on information exchange for Byzantine agreement. *J. ACM*, 32(1):191–204, January 1985. doi:10.1145/2455.214112.
- 22 Sisi Duan, Michael K. Reiter, and Haibin Zhang. BEAT: asynchronous BFT made practical. In *CCS*, pages 2028–2041. ACM, 2018. doi:10.1145/3243734.3243812.
- 23 Romaric Duvignau, Michel Raynal, and Elad M. Schiller. Self-stabilizing Byzantine fault-tolerant repeated reliable broadcast. In *Stabilization, Safety, and Security of Distributed Systems*, pages 206–221. Springer International Publishing, 2022. doi:10.1007/978-3-031-21017-4_14.
- 24 Fabian Frei, Ran Gelles, Ahmed Ghazy, and Alexandre Nolin. Content-oblivious leader election on rings. In *38th International Symposium on Distributed Computing (DISC 2024)*, volume 319, pages 24:1–24:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.DISC.2024.26.
- 25 Ran Gelles. Coding for interactive communication: A survey. *Foundations and Trends® in Theoretical Computer Science*, 13(1–2):1–157, 2017. doi:10.1561/04000000079.
- 26 Li Gong, Patrick Lincoln, and John Rushby. Byzantine agreement with authentication: Observations and applications in tolerating hybrid and link faults. In *Dependable Computing and Fault Tolerant Systems*, volume 10, pages 139–158. IEEE, 1995.
- 27 Rachid Guerraoui, Jovan Komatovic, Petr Kuznetsov, Yvonne-Anne Pignolet, Dragos-Adrian Seredinschi, and Andrei Tonkikh. Dynamic Byzantine reliable broadcast. In *24th International Conference on Principles of Distributed Systems (OPODIS 2020)*, volume 184, pages 23:1–23:18, Dagstuhl, Germany, 2021. doi:10.4230/LIPICS.OPODIS.2020.23.
- 28 William M. Hoza and Leonard J. Schulman. The adversarial noise threshold for distributed protocols. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 240–258, 2016. doi:10.1137/1.9781611974331.ch18.
- 29 Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982. doi:10.1145/357172.357176.
- 30 Alexandre Maurer and Sébastien Tixeuil. Self-stabilizing Byzantine broadcast. In *2014 IEEE 33rd International Symposium on Reliable Distributed Systems*, pages 152–160, 2014. doi:10.1109/SRDS.2014.10.
- 31 Ralph C. Merkle. A certified digital signature. In Gilles Brassard, editor, *Advances in Cryptology — CRYPTO' 89 Proceedings*, pages 218–238, New York, NY, 1990. Springer New York. doi:10.1007/0-387-34805-0_21.

- 32 Andrew Miller, Yu Xia, Kyle Croman, Elaine Shi, and Dawn Song. The honey badger of BFT protocols. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, pages 31–42, 2016. doi:10.1145/2976749.2978399.
- 33 Kartik Nayak, Ling Ren, Elaine Shi, Nitin H. Vaidya, and Zhuolun Xiang. Improved extension protocols for Byzantine broadcast and agreement. In *34th International Symposium on Distributed Computing (DISC 2020)*, volume 179, pages 28:1–28:17, 2020. doi:10.4230/LIPIcs.DISC.2020.28.
- 34 M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *J. ACM*, 27(2):228–234, April 1980. doi:10.1145/322186.322188.
- 35 Andrzej Pelc. Reliable communication in networks with Byzantine link failures. *Networks*, 22(5):441–459, 1992. doi:10.1002/net.3230220503.
- 36 Kenneth J. Perry and Sam Toueg. Distributed agreement in the presence of processor and communication faults. *IEEE Trans. Softw.*, SE-12(3):477–482, 1986. doi:10.1109/TSE.1986.6312888.
- 37 Michel Raynal. Message adversaries. In *Encyclopedia of Algorithms*, pages 1272–1276. Springer, 2016. doi:10.1007/978-1-4939-2864-4_609.
- 38 Michel Raynal. *Fault-Tolerant Message-Passing Distributed Systems - An Algorithmic Approach*. Springer, 2018. doi:10.1007/978-3-319-94141-7.
- 39 Irving S. Reed and Gustave Solomon. Polynomial codes over certain finite fields. *Journal of the Society for Industrial and Applied Mathematics*, 8(2):300–304, 1960.
- 40 Ron M. Roth. *Introduction to coding theory*. Cambridge University Press, 2006.
- 41 Nicola Santoro and Peter Widmayer. Time is not a healer. In *STACS 89*, pages 304–313, Berlin, Heidelberg, 1989. Springer Berlin Heidelberg. doi:10.1007/BFb0028994.
- 42 Nicola Santoro and Peter Widmayer. Distributed function evaluation in the presence of transmission faults. In *International Symposium on Algorithms*, pages 358–367. Springer, 1990. doi:10.1007/3-540-52921-7_85.
- 43 Nicola Santoro and Peter Widmayer. Agreement in synchronous networks with ubiquitous faults. *Theor. Comput. Sci.*, 384(2-3):232–249, 2007. doi:10.1016/J.TCS.2007.04.036.
- 44 Ulrich Schmid and Christof Fetzer. Randomized asynchronous consensus with imperfect communications. In *SRDS*, pages 361–370. IEEE Computer Society, 2003. doi:10.1109/RELDIS.2003.1238089.
- 45 Victor Shoup. Practical threshold signatures. In Bart Preneel, editor, *Proc. Int'l Conference on the Theory and Application of Cryptographic Techniques (EUROCRYPT'2000)*, volume 1807 of *Lecture Notes in Computer Science*, pages 207–220. Springer, 2000. doi:10.1007/3-540-45539-6_15.
- 46 Hin-Sing Siu, Yeh-Hao Chin, and Wei-Pang Yang. Byzantine agreement in the presence of mixed faults on processors and links. *IEEE Transactions on Parallel and Distributed Systems*, 9(4):335–345, April 1998. doi:10.1109/71.667895.
- 47 Lei Yang, Seo Jin Park, Mohammad Alizadeh, Sreeram Kannan, and David Tse. Dispersed-ledger: High-throughput byzantine consensus on variable bandwidth networks. In *NSDI*, pages 493–512. USENIX Association, 2022. URL: <https://www.usenix.org/conference/nsdi22/presentation/yang>.

A Correctness analysis

We now analyze Algorithm 2 and show it satisfies the MBRB properties specified in Theorem 2.3. This analysis, along with the communication analysis in Theorem B.1 prove our main Theorem 3.1.

► **Assumption A.1** (coded-MBRB-assumption). $n > 3t + 2d$ and $k \leq (n - t) - 2d$.

► **Theorem A.2.** For any network that satisfies Theorem A.1 and for any $\varepsilon > 0$, Algorithm 2 implements an MBRB algorithm (Theorem 2.3) with $\ell > n - t - (1 + \varepsilon)d$.

Recall that any MBRB algorithm requires $n > 3t + 2d$ [4]. Our coded MBRB algorithm uses an error-correction code that can reconstruct any encoded message from k fragments of the codeword. While we have some flexibility in selecting the value of k , which affects the parameters of the ECC and thus the communication complexity, our proof requires that k will not be too large. We begin with a few technical lemmas.

► **Lemma A.3.** *If a correct node p_u stores a message fragment \tilde{m}_j associated to a proof of inclusion π_j for some commitment C' and node identity j , then π_j is valid with respect to C' , that is $\text{vc_verify}(C', \pi_j, \tilde{m}_j, j) = \text{valid}$.*

Proof of Lemma A.3. A correct node stores fragments for a commitment C' at line 2, when receiving SEND, FORWARD, or BUNDLE messages, respectively. The fragments stored at these lines and their proof have been received through the corresponding message, whose content is verified by a call to ISVALID (at line 2). ISVALID (described in Algorithm 1) checks that proofs of inclusion are valid for the corresponding commitment. ◀

The following notion of *valid messages* will be used throughout the analysis to indicate messages containing only valid information, as the algorithm dictates.

► **Definition A.4** (Valid messages). *We say that a message of type SEND, FORWARD, or BUNDLE is valid if and only if ISVALID returns True at line 22, 27, or 43, respectively, upon the receipt of that message.*

Operatively, valid messages satisfy the following which is immediate from the definition of the ISVALID function (Algorithm 1).

► **Corollary A.5.** *To be valid, a message must meet the following criteria: (i) all the signatures shares or threshold signatures it contains must be valid and correspond to the commitment included in the message; (ii) if it is of type SEND or FORWARD, it must contain a signature by the designated sending node p_s ; and (iii) all inclusion proofs must be valid with respect to the commitment included in the message.*

We now show that the correct parties always send valid messages. However, they might receive invalid messages sent by Byzantine nodes.

► **Lemma A.6.** *All SEND, FORWARD, or BUNDLE messages sent by a correct node p_u , are valid.*

Proof of Lemma A.6. The only correct node that sends SEND messages is p_s at line 2. Indeed, when p_s is correct, this message will contain a valid signature share by p_s and all proofs of inclusion are valid, by line 2.

Now consider a FORWARD message sent either at lines 25 or 34. To reach there, p_u must have passed line 2 or line 2, which guarantees p_u received a valid signature for C' made by p_s (where C' is the commitment in the received message triggering this code). Then, at line 2 or at line 2, p_u stores a signature of p_s for this C' , and at line 2, p_u signs the same C' . Thus, conditions (i) and (ii) of Theorem A.5 hold, and if the FORWARD is sent at line 2, then condition (iii) vacuously holds as well. If the FORWARD message is sent at line 2, it contains a fragment that was stored by p_u for the same C' , and by Theorem A.3, its associated proof of inclusion is valid; thus condition (iii) holds in this case as well.

Finally, consider a BUNDLE message. First off, this type of message is not concerned by condition (ii) of Theorem A.5. For the transmission at line 2, condition (i) follows from the construction of the threshold signature Σ_C at line 2. Σ_C is guaranteed to be non- \perp by the

condition at line 2 of Algorithm 2, and is provided by the helper function $\text{GETTHRESHSIG}(\cdot)$ (line 1 of Algorithm 1). When executing $\text{GETTHRESHSIG}(\cdot)$, the first possibility is that Σ_C is already known by p_u because it was received by p_u at line 2 and stored at line 2. In this case, the validity of Σ_C is ensured by the check at line 2. The second possibility is that Σ_C aggregates $\tau = \lfloor \frac{n+t}{2} \rfloor + 1$ signature shares received by p_u at line 2 or line 2, and stored at line 2 or line 2, respectively. In this case, the validity of all these signature shares is ensured by the checks at line 2 and line 2, respectively, and thus the aggregated threshold signature Σ_C is also valid. Condition (iii) follows since the proofs of inclusion were computed at line 2 by p_u and match the same commitment C' used in that BUNDLE message, as enforced by line 2. For the broadcast at line 2, conditions (i) and (iii) follow since the threshold signature Σ and the fragment tuple $(\tilde{m}'_j, \pi'_j, j)$ come from the incoming BUNDLE message at line 2, whose validity (w.r.t. C') has been verified at line 2. \blacktriangleleft

► **Lemma A.7.** *A correct node p signs a most a single commitment C .*

Proof of Lemma A.7. p signs a commitment either at line 2 (for p_s), 24 or 33. We consider two cases, depending on whether p is p_s or not.

- *Case 1:* Assume $p \neq p_s$. p can sign some commitment only at line 2 or 33. By the conditions at line 2, line 2 are executed only if either p has not signed any commitment yet, or has already signed the exact same commitment C' .
- *Case 2:* If $p = p_s$, because valid messages must contain p_s 's signature share (due to calls to $\text{ISVALID}()$ at line 2), and because we have assumed that signatures cannot be forged, line 2 is always executed before line 2. By the same reasoning as Case 1, p_s therefore never signs a different commitment at line 2 or 33. \blacktriangleleft

We recall that the above lemmas, and as a consequence, the upcoming theorems, hold with high probability, assuming a computationally-bounded adversary that forges signature shares/threshold signatures or finds commitment collisions with only negligible probability. We can now prove the properties required for an MBRB algorithm, as depicted in Theorem 2.3.

► **Lemma A.8 (MBRB-Validity).** *Suppose p_s is correct and a correct node, p_i , MBRB-delivers an application message, m . Then, node p_s has previously MBRB-broadcast m .*

Proof of Lemma A.8. Suppose p_i MBRB-delivers m at line 2. Consider C' the commitment that renders true the condition at line 2, and C the commitment that is computed at line 2. It holds that $C' = C$ by line 2, or otherwise p_i could not have reached line 2.

Consider the threshold signature Σ_C returned by the GETTHRESHSIG function at line 2. Using the same reasoning as in the proof of Theorem A.6, Σ_C is valid, and must, therefore, aggregate at least $\tau = \lfloor \frac{n+t}{2} \rfloor + 1$ valid signature shares for C . Let us remark that, out of all these valid signature shares, at least $\lfloor \frac{n+t}{2} \rfloor + 1 - t = \lfloor \frac{n-t}{2} \rfloor + 1 \geq 1$ are generated by correct nodes². Thus, at least one correct node p_j must have produced a signature share for C , whether it be at line 2 or line 2 if $p_j \neq p_s$, or at line 2 if $p_j = p_s$. However, in all these cases, the sender p_s must have necessarily produced a signature share for C : the case $p_j = p_s$ is trivial, and in the case $p_j \neq p_s$, p_j must have verified the presence of a valid signature share from p_s in the message it received, at line 2 or line 2, respectively.

Under the assumption that the adversary cannot forge signature shares/threshold signatures (Section 2), and recalling that p_s is correct, the only way in which p_s could have signed C' is by executing line 2 when MBRB-broadcasting some message m' at line 2; see also the proof of Theorem A.7. Furthermore, recall that the commitment is collision-resistant (or binding, see Section 2), meaning that except with negligible probability, the message m' that p_s uses in line 2 satisfies $m' = m$, since it holds that $C' = \text{Commitment}(m') = \text{Commitment}(m) = C$. \blacktriangleleft

² Remind that, $\forall x \in \mathbb{R}, i \in \mathbb{Z} : \lfloor x \rfloor + i = \lfloor x + i \rfloor$.

► **Lemma A.9** (MBRB-No-duplication). *A correct node p_i MBRB-delivers at most one application message, m .*

Proof of Lemma A.9. The condition at line 2 directly implies the proof. ◀

► **Lemma A.10** (MBRB-No-duplicity). *No two different correct nodes MBRB-deliver different application messages.*

Proof of Lemma A.10. Suppose, towards a contradiction, that p_i MBRB-delivers m and p_j MBRB-delivers $m' \neq m$, where p_i and p_j are both correct nodes. Let us denote by C , resp. C' , the commitment returned by COMPUTEFRAGVECCOMMIT() for m , resp. for m' . As commitments are assumed to be collision-resistant (Section 2), $m \neq m'$ implies $C \neq C'$.

By the condition at line 2, p_i gets a threshold signature $\Sigma_i \neq \perp$ from the GETTHRESHSIG function that aggregates a set Q_i containing $\tau = \lfloor \frac{n+t}{2} \rfloor + 1$ valid signature shares for C . Similarly, p_j gets a threshold signature Σ_j aggregating a set Q_j of signature shares for C' . We know that $|Q_i \cup Q_j| = |Q_i| + |Q_j| - |Q_i \cap Q_j|$. Moreover, we know that, $\forall x \in \mathbb{R}, k \in \mathbb{Z} : k = \lfloor x \rfloor + 1 \implies k > x$, and hence we have $Q_i > \frac{n+t}{2} < Q_j$. Thus, $|Q_i \cap Q_j| \geq |Q_i| + |Q_j| - n > 2 \frac{n+t}{2} - n = t$. In other words, Q_i and Q_j have at least one correct node, p_u , in common that has signed both C and C' . Theorem A.7, and the fact that p_u has signed both C and C' leads the proof to the needed contradiction. Thus, $m = m'$, and the lemma holds. ◀

► **Lemma A.11** (MBRB-Local-delivery). *Suppose p_s is correct and MBRB-broadcasts m . At least one correct node, say, p_j , MBRB-delivers m .*

Proof of Lemma A.11. Let us denote by C_m the commitment computed at line 1 when executing COMPUTEFRAGVECCOMMIT(m). The proof of the lemma will follow from Observations A.11.1–A.11.8 stated and proven below.

► **Observation A.11.1.** *All valid SEND, FORWARD, or BUNDLE messages received by some correct node p_u contain C_m .*

Proof of Observation A.11.1. Recall that p_s MBRB-broadcasts m , thus we know that p_s has included its own signature share, $sig_s = (\text{ts_sign_share}_s(C_m), s)$, when it propagates $\langle \text{SEND}, C_m, (\tilde{m}_j, \pi_j, j), sig_s \rangle$ (lines 2–2). Consider a correct node p_u that receives a valid SEND, FORWARD, or BUNDLE message containing a commitment C_u at lines 21, 26, or 42. If the message is of type SEND or FORWARD, then, as it is valid, it must contain p_s 's signature on C_u . If the message is of type BUNDLE, then, similarly to Theorem A.8, its valid threshold signature for C_u aggregates a set of valid signature shares for C_u that contains at least one share produced by a correct node p_x . But for p_x to produce this share, p_s must also have produced a valid signature share for C_u , either because p_x must have checked its existence at line 2 or line 2 (prior to signing, at line 2 or line 2, respectively), or because p_x is the sender. Hence, in any case, p_s produces a signature share for C_u . Since p_s is correct, by Theorem A.7, it does not sign another commitment $C' \neq C_m$. Under the assumption that signatures cannot be forged, the above implies that $C_u = C_m$. ◀

► **Observation A.11.2.** *A correct node p_u only signs valid signature shares for C_m .*

Proof of Observation A.11.2. If $p_u = p_s$, it MBRB-broadcasts a single message and executes line 2 only once, signing C_m . Besides line 2, a correct node p_u only signs signature shares after receiving a valid SEND or FORWARD message (at line 2), and when it does, p_u only ever signs the commitment received in the message. By Observation A.11.1, this implies p_u never signs any $C' \neq C_m$. ◀

► **Observation A.11.3.** *If a correct node p_u broadcasts a FORWARD message, this message is of the form $\langle \text{FORWARD}, C_m, -, \{sig_s, sig_u\} \rangle$, where sig_s, sig_u are p_s 's and p_u 's valid signature shares for C_m .*

Proof of Observation A.11.3. Consider a correct node p_u that broadcasts a message $\langle \text{FORWARD}, C', -, sigs \rangle$ either at lines 25 or 34. By Observation A.11.1, $C' = C_m$. The observation then follows from Theorem A.6. ◀

We now define F_{recv} to be the set of correct nodes that receive a valid message $\langle \text{FORWARD}, C_m, -, sigs \rangle$ at line 2, where $sigs$ contains p_s 's valid signature for C_m . We analyze its size and the behavior of such nodes in the following observations.

► **Observation A.11.4.** F_{recv} contains at least one correct node, i.e., $F_{\text{recv}} \neq \emptyset$.

Proof of Observation A.11.4. If p_s is correct and MBRB-broadcasts m , it executes line 2 and disseminates messages of the form $\langle \text{SEND}, C_m, (\tilde{m}_j, \pi_j), sig_s \rangle$ to all nodes, where sig_s is p_s 's signature share of C_m . By definition of the message adversary, these SEND messages are received by at least $c - d$ correct nodes.

By Theorem A.1, $n > 3t + 2d$, and therefore $c - d \geq n - t - d > 0$. At least one correct node p_x , therefore, receives one of the SEND messages disseminated by p_s at line 2. As p_s is correct, by Theorem A.6, this message is valid, and is handled by p_x at lines 2–2.

By Observation A.11.2, p_x only signs signature shares for C_m , and thus passes the test at line 2, and reaches line 2, where it disseminates a FORWARD message. By Observation A.11.3, this message is of the form $\langle \text{FORWARD}, C_m, -, \{sig_s, sig_x\} \rangle$, and is valid. As above, by definition of the message adversary, this FORWARD message is received by at least $c - d > 0$ correct nodes. By definition these nodes belong to F_{recv} , which yield $|F_{\text{recv}}| > 0$ and $F_{\text{recv}} \neq \emptyset$. ◀

► **Observation A.11.5.** *Any $p_u \in F_{\text{recv}}$ broadcasts a $\langle \text{FORWARD}, C_m, -, \{sig_s, sig_u\} \rangle$ message, where sig_s and sig_u are p_s and p_u 's valid signature shares for C_m , respectively.*

Proof of Observation A.11.5. Let $p_u \in F_{\text{recv}}$ upon receiving a valid FORWARD message $\langle \text{FORWARD}, C_m, -, sigs \rangle$ at line 2. By the condition of line 2, p_u has either previously sent a FORWARD message at line 2 or it will send such a message at line 2. In both cases, Observation A.11.3 applies and guarantees that this message contains C_m and both p_u 's and p_s 's valid signature shares. ◀

Note that F_{recv} is defined over an entire execution of Algorithm 2. Observation A.11.5 therefore states that any correct node p_u that receives a valid FORWARD message *at some point of its execution* also broadcasts a matching FORWARD message *at some point of its execution*. The two events (receiving and sending a FORWARD message) might, however, occur in any order. For instance, p_u might first receive a SEND message from p_s at line 2, disseminate a FORWARD message as a result at line 2, and later on possibly receive a FORWARD message from some other node at line 2. Alternatively, p_u might first receive a FORWARD message at line 2, and disseminate its own FORWARD message at line 2 as a result. In this second case, p_u might also eventually receive a SEND message from p_s (at line 2). If this happens, p_u will disseminate a second FORWARD message at line 2. A correct node, however, never disseminates more than two FORWARD messages (at line 2).

► **Observation A.11.6.** *Any broadcast of $\langle \text{FORWARD}, C_m, \text{fragtuple}, \{sig_s, sig_u\} \rangle$ by a correct $p_u \in F_{\text{recv}}$ arrives to at least $c - d$ correct nodes that are each, eventually, in F_{recv} .*

Proof of Observation A.11.6. Each broadcast by a correct p_u of a FORWARD message is eventually received by at least $c - d$ correct nodes by definition of the message adversary. By Observation A.11.3 the FORWARD message contains C_m , by Theorem A.6 it is valid. Thus, each of its at least $c - d$ correct recipients belong in F_{recv} , by definition. \blacktriangleleft

Because FORWARD messages are disseminated at line 2, the reception and sending of FORWARD messages by correct nodes will induce a “chain reaction” until a correct node is reached that has already disseminated a FORWARD message. This “chain reaction” mechanism is the intuitive reason why some correct node will eventually receive enough distinct FORWARD messages to trigger an MBRB-delivery, as captured by the following observation.

► **Observation A.11.7.** *There exists a correct node p_j that receives messages $\langle \text{FORWARD}, C_m, -, \text{sig}_s = \{\text{sig}_s, \text{sig}_u\} \rangle$ from at least $(c - d)$ distinct correct nodes p_u , where $\text{sig}_s = (\text{ts_sign_share}_s(C_m), s)$ and $\text{sig}_u = (\text{ts_sign_share}_u(C_m), u)$ are p_s and p_u 's valid signature shares for C_m , respectively, and the FORWARD message is the last message sent by p_u .*

Proof of Observation A.11.7. By Observation A.11.5, any nodes $p_u \in F_{\text{recv}}$ broadcasts at least one message $\langle \text{FORWARD}, C_m, -, \text{sig}_s = \{\text{sig}_s, \text{sig}_u\} \rangle$, that includes p_u 's valid signature share for C_m , $\text{sig}_u = (\text{ts_sign_share}_u(C_m), u)$. Consider all the FORWARD messages sent by nodes in F_{recv} during the *last time* they perform such a broadcast. By Observation A.11.6, there are $|F_{\text{recv}}|$ senders, $p_u \in F_{\text{recv}}$, such that each of p_u 's last broadcast of a FORWARD message is guaranteed to be delivered to at least $c - d$ correct nodes p_x , such that eventually $p_x \in F_{\text{recv}}$. Thus, at least $|F_{\text{recv}}|(c - d)$ such messages are received by nodes in F_{recv} , overall. By Observation A.11.4, F_{recv} contains at least one node. We can, therefore, apply the pigeonhole principle, where F_{recv} are the holes and the above $|F_{\text{recv}}|(c - d)$ messages are the pigeons, and observe that there exists a node $p_j \in F_{\text{recv}}$ that will receive at least $|F_{\text{recv}}|(c - d) / |F_{\text{recv}}|$ such messages. Since we limit the discussion to a *single*, *i.e.*, the last, broadcast performed by each node in F_{recv} , no node in F_{recv} receives two of the above messages that were originated by the same node in F_{recv} . Therefore, we deduce that p_j has received messages of the form $\langle \text{FORWARD}, C_m, -, \text{sig}_s \rangle$ from at least $(c - d)$ distinct correct nodes p_u and the FORWARD message is the *last* message sent by p_u . \blacktriangleleft

► **Observation A.11.8.** *At least one correct node MBRB-delivers m from p_s .*

Proof of Observation A.11.8. By Observation A.11.7, there is a correct node p_j that receives messages of the form $\langle \text{FORWARD}, C_m, -, \text{sig}_s \rangle$ from at least $(c - d)$ distinct correct nodes p_u , such that these FORWARD messages are the *last* message sent by each p_u . Let us denote by U the set of such nodes p_u , hence, $|U| \geq c - d$.

Still by Observation A.11.7, p_j receives a valid signature share $\text{sig}_u = (\text{ts_sign_share}_u(C_m), u)$ from each node $p_u \in U$. It thus receives at least $(c - d)$ distinct signature shares for C_m . Theorem A.1 says $3t + 2d < n$, and thus, $n + 3t + 2d < 2n$ and $n + t < 2n - 2t - 2d$. Since $n - t \leq c$, we have $(n + t)/2 < n - t - d \leq c - d$. Thus, p_j receives more than $(n + t)/2$ valid distinct signature shares for C_m .

Let us now consider the set of correct nodes S that receive the initial SEND messages disseminated by p_s at line 2. Any node $p_x \in S$ receives through the SEND message its allocated fragment (\tilde{m}_x, π_x, x) from p_s . By definition of the message adversary, the SEND messages disseminated at line 2 are received by at least $c - d$ correct nodes, therefore $|S| \geq c - d$. Furthermore, all nodes in S broadcast a FORWARD message at line 2, and this will be their *last* FORWARD message, due to the condition of line 2. By the above reasoning, this FORWARD message will contain their message fragment, that is, it will be of the form $\langle \text{FORWARD}, C_m, (\tilde{m}_x, \pi_x, x), \text{sig}_s \rangle$. By Theorem A.6, they are all valid.

By definition of S and U , both these sets contain only correct nodes, thus, $|S \cup U| \leq c$. As a result, $|S \cap U| = |S| + |U| - |S \cup U| \geq 2 \times (c - d) - c = c - 2d$. The last FORWARD messages broadcast by nodes in $S \cap U$ are received by p_j by the definition of U . As argued above about nodes in S (and thus applying to nodes in $S \cap U$), FORWARD messages sent by a node in $S \cap U$ contain their valid message fragment and proof of inclusion (\tilde{m}_x, π_x, x) . It follows that p_j accumulates at least $c - 2d$ distinct such message fragments with their (valid) proof of inclusion. By Theorem A.1, $c - 2d \geq k$.

To conclude the proof, note that we have shown that p_j eventually receives more than $(n+t)/2$ valid distinct signature shares for C_m , and additionally, that p_j accumulates at least k valid message fragments with their proof of inclusion. At this point, the condition of line 2 becomes True for p_j . Because the commitment is collision-resistant (Section 2), once C_m is fixed, we can assume that, except with negligible probability, all the message fragments that p_j has received correspond to the fragments computed by p_s at line 2. By the parameters of the ECC we use, it can recover the message m from any k or more (correct) fragments generated by p_s , where missing fragments are considered as erasures. Therefore, the message m_j reconstructed at line 2 by p_j is the message initially MBRB-broadcast by p_s . As a result $m_j = m$, and p_j MBRB-delivers m at line 2. ◀

Theorem A.12 is the detailed version of Theorem 3.2.

► **Lemma A.12** (MBRB-Global-delivery). *Suppose a correct node, p_i , MBRB-delivers an application message m . At least $\ell = c - d \left(\frac{1}{1 - \left(\frac{k-1}{c-d} \right)} \right)$ correct nodes MBRB-deliver m .*

Proof of Lemma A.12. Suppose a correct node p_i MBRB-delivers m (line 2). Let us denote by C_m the commitment returned by COMPUTEFRAgVECCOMMIT(m). The proof follows a counting argument on the BUNDLE messages disseminated by correct nodes at line 2. In the following, if a correct node disseminates a BUNDLE message both at line 2, we only consider the one from line 2.

► **Observation A.12.1.** *All valid BUNDLE messages exchanged during the execution of Algorithm 2 contain C_m , the commitment of the message m , where m is the message MBRB-delivered by p_i .*

Proof of Observation A.12.1. Consider a valid message $\langle \text{BUNDLE}, C', \text{fragtuple}'_j, \text{fragtuple}'_i, \Sigma' \rangle$. By definition of a valid BUNDLE message, Σ' aggregates a set sigs' of $\tau = \lfloor \frac{n+t}{2} \rfloor + 1$ valid signature shares for C' . Similarly, when p_i MBRB-delivers m at line 2, it has a threshold signature Σ_m which aggregates a set sigs_m of $\tau = \lfloor \frac{n+t}{2} \rfloor + 1$ valid signature shares for C_m . By a reasoning identical to that of Theorem A.10, these two inequalities imply that $\text{sigs}' \cap \text{sigs}_m$ contains the signature shares from at least one common correct node, p_u . As signature shares cannot be forged, p_u has issued signature shares for both C' and C_m , and by Theorem A.7, $C' = C_m$. To complete the proof, note that by the definition of a valid BUNDLE message, the threshold signature it contains is valid with respect to the commitment it carries. Hence, all valid BUNDLE messages must contain the commitment C_m of the application message m that matches the threshold signature Σ' they contain. ◀

Let B_{send} be the set of correct nodes that disseminate at least one BUNDLE message during their execution. Similarly, let B_{recv} be the set of correct nodes that receive at least one valid BUNDLE message from a correct node during their execution. The following holds.

Observation A.12.2 is a detailed version of Observation 3.2.1 in Theorem 3.2.

► **Observation A.12.2.** $c \geq |B_{\text{recv}}| \geq c - d$ and $c \geq |B_{\text{send}}| \geq c - d$.


Proof of Observation A.12.2. Since B_{send} and B_{recv} contain only correct nodes, trivially $c \geq |B_{\text{send}}|$ and $c \geq |B_{\text{recv}}|$. Since p_i MBRB-delivers m at line 2, it must have disseminated BUNDLE messages of the form $\langle \text{BUNDLE}, C_m, (\tilde{m}'_i, \pi'_i, i), (\tilde{m}'_j, \pi'_j, j), \Sigma_C \rangle$ at line 2. The BUNDLE messages sent by p_i eventually reach at least $c - d$ correct nodes, as the message adversary can remove at most d of these BUNDLE messages. By Theorem A.6, these BUNDLE messages are valid. Hence, $|B_{\text{recv}}| \geq c - d > 0$ proves the lemma's first part.

The nodes in B_{recv} (which are correct) execute line 2, and reach line 2. Because p_i has included a non- \perp second fragment in all its BUNDLE message, any of the $(c - d)$ nodes of B_{recv} that receive one of p_i 's BUNDLE messages and has not already sent a BUNDLE message passes the condition at line 2. Each such node then disseminates a (valid) BUNDLE message at line 2. This behavior yields $|B_{\text{send}}| \geq c - d$. ◀

Let $B_{k,\text{recv}}$ be the set of correct nodes that receive BUNDLE messages from at least k distinct correct nodes.

Observation A.12.3 is a detailed version of Observation 3.2.2 in Theorem 3.2.

► **Observation A.12.3.** $|B_{k,\text{recv}}| \times |B_{\text{send}}| + (k - 1)(|B_{\text{recv}}| - |B_{k,\text{recv}}|) \geq |B_{\text{send}}|(c - d)$.



■ **Figure 1** Distribution of distinct BUNDLE messages received by correct nodes. The proof of Observation A.12.4 shows that $|B_{\text{send}}| > k - 1$.

Proof of Observation A.12.3. Let us denote by $\#\text{BUNDLE}$ the overall number of valid BUNDLE messages received by correct nodes from distinct correct senders. More specifically, in the case when a correct node disseminates BUNDLE messages both at line 2, we only consider the *last* BUNDLE message, *i.e.*, the one of line 2. We know that each $p \in B_{\text{send}}$ sends a BUNDLE message, which by Theorem A.6 is valid. As the message adversary may drop up to d out of the n messages of this comm, we are guaranteed that at least $c - d$ correct nodes receive p 's BUNDLE message. This immediately implies that

$$\#\text{BUNDLE} \geq |B_{\text{send}}|(c - d). \quad (1)$$

As illustrated in Figure 1, the nodes in $B_{k,\text{recv}}$ may receive up to $|B_{\text{send}}|$ valid BUNDLE messages from distinct correct senders (one from each sender in B_{send}), for a maximum of $|B_{k,\text{recv}}| \times |B_{\text{send}}|$ BUNDLE messages overall. The remaining nodes of $B_{\text{recv}} \setminus B_{k,\text{recv}}$ may each receive up to $k - 1$ valid BUNDLE messages from distinct correct senders, by definition of $B_{k,\text{recv}}$. As $B_{k,\text{recv}} \subseteq B_{\text{recv}}$ by definition, $|B_{\text{recv}} \setminus B_{k,\text{recv}}| = |B_{\text{recv}}| - |B_{k,\text{recv}}|$, and the nodes of $B_{\text{recv}} \setminus B_{k,\text{recv}}$ accounts for up to $(k - 1)(|B_{\text{recv}}| - |B_{k,\text{recv}}|)$ valid BUNDLE messages overall. As the BUNDLE messages counted by $\#\text{BUNDLE}$ are received either by correct nodes in $B_{k,\text{recv}}$ or in $B_{k,\text{recv}} \setminus B_{\text{recv}}$, these observations lead to

$$|B_{k,\text{recv}}| \times |B_{\text{send}}| + (k-1)(|B_{\text{recv}}| - |B_{k,\text{recv}}|) \geq \#\text{BUNDLE}. \quad (2)$$

Combining Equations (1) and (2) yields the desired inequality. ◀

Observation A.12.4 is a detailed version of Observation 3.2.3 in Theorem 3.2.

► **Observation A.12.4.** $|B_{k,\text{recv}}| \geq c - d \left(\frac{1}{1 - \left(\frac{k-1}{c-d} \right)} \right).$

Proof of Observation A.12.4. Rearranging the terms of Observation A.12.3, and recalling that $|B_{\text{recv}}| \leq c$ and $k \geq 1$, we get

$$|B_{k,\text{recv}}| \times (|B_{\text{send}}| - k + 1) \geq |B_{\text{send}}|(c - d) - |B_{\text{recv}}|(k - 1) \geq |B_{\text{send}}|(c - d) - c(k - 1). \quad (3)$$

By Observation A.12.2 and Theorem A.1, $|B_{\text{send}}| \geq c - d \geq c - 2d \geq k$, therefore $|B_{\text{send}}| - k + 1 > 0$, and the previous equation can be transformed in

$$|B_{k,\text{recv}}| \geq \frac{|B_{\text{send}}|(c - d) - c(k - 1)}{|B_{\text{send}}| - k + 1}. \quad (4)$$

Note that the right-hand side of Equation (4) is a monotone increasing function in $|B_{\text{send}}|$ when $|B_{\text{send}}| > k - 1$, as its derivative, $\frac{d(k+1)}{(|B_{\text{send}}| - k + 1)^2}$, is positive. By Observation A.12.2, $|B_{\text{send}}| \in [c - d, c] \subseteq [k, c]$. The minimum of the right-hand side of Equation (4) is therefore obtained for $B_{\text{send}} = c - d$, yielding

$$|B_{k,\text{recv}}| \geq \frac{(c - d)^2 - c(k - 1)}{(c - d) - k + 1} = c - d \left(\frac{1}{1 - \left(\frac{k-1}{c-d} \right)} \right). \quad (5)$$

Observation A.12.5 is a detailed version of Observation 3.2.4 in Theorem 3.2.

► **Observation A.12.5.** *All nodes in $B_{k,\text{recv}}$ MBRB-deliver m .*

Proof of Observation A.12.5. Consider $p_u \in B_{k,\text{recv}}$. By the definition of $B_{k,\text{recv}}$, the node p_u receives k valid BUNDLE messages from k distinct correct nodes. Let us denote by $\langle \text{BUNDLE}, C_x, (\tilde{m}_x, \pi_x, x), -, \Sigma_x \rangle$ these k messages with $x \in [k]$. By Observation A.12.1, for all $x \in [k]$, $C_x = C_m$. In addition, p_u stores each received threshold signature Σ_x , which is valid for C_m .

Because the messages are valid, so are the proofs of inclusions π_x , and as we have assumed that the commitments are collision-resistant, $C_x = C_m$ implies that the received fragments \tilde{m}_x all belong to the set of fragments computed by p_s at line 2 for m . As the BUNDLE messages were received from k distinct correct nodes, the node p_u receives at least k distinct valid fragments for m during its execution. If p_u has not MBRB-delivered any message yet, the condition at line 2 eventually becomes true for C_m , and p_u reconstructs m at line 2, since it possesses at least k (correct) message fragments, which are sufficient for the correct recovery of m by our choice of ECC. Then, p_u MBRB-delivers m at line 2. On the other hand, if p_u has already MBRB-delivered some message m' , then Theorem A.10 (MBRB-No-duplicity) implies $m' = m$, since p_i is known to have MBRB-delivered m . Therefore, in all possible cases, p_u MBRB-delivers m . ◀

Theorem A.12 follows from Observations A.12.4 and A.12.5 and the fact that all nodes in $B_{k,\text{recv}}$ are correct. ◀

Discussion: Selection of k . In the above analysis, we set k to be a parameter that controls the number of fragments that allow decoding the ECC. To obtain the communication depicted in Section 3, we assumed $k = \Omega(n)$. Furthermore, this parameter affects the delivery power of the MBRB algorithm, as seen in Theorem A.12, namely $\ell = c - d \left(1 - \left(\frac{k-1}{c-d}\right)^{-1}\right)$.

Let us assume that we wish to design an MBRB algorithm with a specified delivery power of $\ell = c - (1 + \varepsilon)d$, for some $\varepsilon > 0$. Plugging in Theorem A.12, we need the delivery power ℓ provided by the Algorithm 2 to surpass $c - (1 + \varepsilon)d$, thus

$$c - (1 + \varepsilon)d \leq c - d \left(\frac{1}{1 - \left(\frac{k-1}{c-d}\right)} \right)$$

leading to $k \leq \frac{\varepsilon}{1+\varepsilon}(c-d) + 1$. That is, choosing any integer $k \leq \frac{\varepsilon}{1+\varepsilon}(n-t-d) + 1$, satisfies the above. Recall that the blowup of the ECC is given by $\bar{n}/k \approx n/k$ (Section 3), which implies that for any application message m , we have $|\text{ECC}(m)| \approx \frac{n}{k}|m| = \frac{1+\varepsilon}{\varepsilon} \cdot \frac{n}{n-t-d}|m|$.

Together with Theorem A.1, we conclude that the constraints on k that support delivery power of $\ell \geq n - t - (1 + \varepsilon)d$, are

$$k \leq \min \left\{ n - t - 2d, \frac{\varepsilon}{1 + \varepsilon}(n - t - d) + 1 \right\}.$$

Supporting multiple instances and multiple senders. We remark that the above analysis fits the single-shot broadcast algorithm with a fixed sender. As mentioned above, a multi-shot multi-sender algorithm can be achieved by communicating the identity of the sender and a sequence number along with any piece of information communicated or processed during this algorithm. This added information uniquely identifies any piece of information with the respective instance. Additionally, signature shares, threshold signatures, commitments, and proofs of inclusion should be performed on the application message m augmented with the sender's identifier and the sequence number. This will prevent Byzantine nodes from using valid signature shares/threshold signatures from one instance in a different instance. As a result, an additive factor of $\mathcal{O}(\log n)$ bits has to be added to each communicated message, which yields additive communication of $\mathcal{O}(n^2 \log n)$ and has no effect on the asymptotic communication, as we explained in the proof of Theorem B.1. Other changes, such as augmenting the application message m with the sender's identifier and sequence number do not affect the length of signature shares, threshold signatures, commitments, and proofs of inclusion.

B Communication analysis

This appendix section analyzes the communication cost of Algorithm 2. Theorem B.1 is the detailed version of Theorem 3.3.

► **Lemma B.1.** *Correct nodes collectively communicate at most $4n^2$ messages. Each correct node sends at most $\mathcal{O}(|m| + n\kappa)$ bits. Overall, the system sends at most $\mathcal{O}(n|m| + n^2\kappa)$ bits.*

Proof of Lemma B.1. Let us count the messages communicated by counting `comm` and `broadcast` invocations. The sender, p_s , sends `SEND` messages at line 2. In Phase I, each correct node that has received a `SEND` message broadcasts a `FORWARD` message once (line 2). However, if it receives a `FORWARD` before the `SEND` arrives, it performs one additional `FORWARD` broadcast (line 2). This yields at most 2 `comm` and `broadcast` invocations per

correct node until the end of Phase I. We can safely assume that a correct sender always sends a single FORWARD (*i.e.*, it immediately and internally receives the SEND message sent to self). Thus, p_s is also limited to at most 2 invocations up to this point. In Phase II, each correct node that MBRB-delivers a message at line 2 transmits BUNDLE messages at line 2. This can only happen once due to the condition at line 2. Additionally, it may transmit BUNDLE messages also at line 2, upon the reception of a BUNDLE. However, this second BUNDLE transmission can happen at most once, due to the if-statement at line 2. This leads to at most 2 additional comm and broadcast invocations per correct node. Thus, as the number of correct nodes is bounded by n , the two phases incur in total at most $4n$ invocations of comm and broadcast overall. Since each invocation communicates exactly n messages, the total message cost for correct nodes when executing one instance of Algorithm 2 is upper bounded by $4n^2$. Note that the above analysis holds for correct nodes also in the presence of Byzantine participants, including when p_s is dishonest.

We now bound the number of bits communicated by correct nodes throughout a single instance of Algorithm 2. Consider COMPUTEFragVECCommit. Let m be a specific application message. We have $|\tilde{m}| = \mathcal{O}(|m|)$ since we use a code with a constant rate. Thus, any specific fragment \tilde{m}_i has length $|\tilde{m}_i| = \mathcal{O}(|m|/n)$. Recall that the sizes of a signature share σ , a threshold signature Σ , a commitment C , and an inclusion proof π all have $\mathcal{O}(\kappa)$ bits (Section 2). Along a signature share pair sig , the identifier of the signing node is included, which takes additional $\mathcal{O}(\log n)$ bits. However, since $\kappa = \Omega(\log n)$, the inclusion of this field does not affect asymptotic communication costs.

We now trace all the comm and broadcast instances in Algorithm 2 and analyze the number of bits communicated in each. The SEND comm (line 2) communicates n messages, where each message includes a fragment of m ($\mathcal{O}(|m|/n)$ bits) with its proof of inclusion ($\mathcal{O}(\kappa)$ bits), a commitment ($\mathcal{O}(\kappa)$ bits), and a signature share ($\mathcal{O}(\kappa)$ bits). Thus, this operation allows the sender to communicate at most $\mathcal{O}(|m| + n\kappa)$ bits. Each FORWARD broadcast in lines 25 and 34 sends n copies of a message containing a commitment ($\mathcal{O}(\kappa)$ bits), at most one message fragment with its proof of inclusion ($\mathcal{O}(|m|/n + \kappa)$ bits), and two signature shares ($\mathcal{O}(\kappa)$ bits). Hence, each one of lines 25 and 34 communicates a total of $\mathcal{O}(|m| + \kappa n)$ bits. The BUNDLE communication (lines 40 or 48) sends n messages, where each contains a commitment ($\mathcal{O}(\kappa)$ bits), at most two message fragments with their proof of inclusion ($\mathcal{O}(|m|/n + \kappa)$ bits), and one threshold signature ($\mathcal{O}(\kappa)$ bits). Hence, each line communicates at most $\mathcal{O}(|m| + n\kappa)$ bits. As analyzed above, the sending node (p_s , when correct) performs at most one comm of SEND messages, while each correct node performs at most two broadcast of FORWARD messages, and at most two comm/broadcast of BUNDLE messages. Thus, each node communicates at most $\mathcal{O}(|m| + n\kappa)$ bits. Overall, the total bit communication by correct nodes during Algorithm 2's execution is $\mathcal{O}(n|m| + n^2\kappa)$. As mentioned above, the analysis holds in the presence of Byzantine nodes, even if p_s is dishonest. ◀

C Using Bracha's BRB on hash values under a message adversary

Das, Xiang, and Ren [18] have proposed a communication optimal BRB algorithm that does not use signatures and relies on Bracha instead to reliably broadcast a hash value of the initial sender's message. One might legitimately ask whether this approach could not be easily adapted to withstand a message adversary, possibly resulting in an MBRB algorithm exhibiting optimal communication complexity (up to the size of hashes κ), optimal Byzantine resilience ($n > 3t + 2d$), and optimal delivery power $n - t - d$ (or at least some close-to-optimal delivery power ℓ , up to some factor ϵ).

Unfortunately, under a message adversary, Bracha's BRB leads to a sub-optimal Byzantine resilience, and degraded delivery power ℓ . In particular, Albouy, Frey, Raynal and Taiani [3] have shown that Bracha can be used to implement a MBRB algorithm, but their solution requires a sub-optimal resilience bound ($n > 3t + 2d + 2\sqrt{td}$) and yields a reduced delivery power $\ell_{BRB} = \left\lceil n - t - \left(\frac{n-t}{n-3t-d}\right)d \right\rceil$. Disappointingly, these less-than-optimal properties would in turn be passed on to any MBRB algorithm using Bracha's BRB along the lines of Das, Xiang, and Ren's solution.³ By contrast, the algorithm we propose is optimal in terms of communication cost (up to κ) and Byzantine resilience, and close to optimal in terms of delivery power (up to some parameter ϵ that can be chosen arbitrarily small).

To provide a hint of why Bracha's BRB leads to degraded resilience and delivery power when confronted with a message adversary (MA), consider the classical ECHO phase of Bracha's BRB [10]. At least one correct node must receive $(n+t)/2$ ECHO messages to ensure the first READY message by a correct node can be emitted. To ensure Local-delivery, the threshold $(n+t)/2$ must remain lower than the worst number of ECHO messages a correct node can expect to receive when the sender is correct. Without an MA this constraint leads to $(n+t)/2 < n-t$, which is verified by assuming $n > 3t$. With an MA, the analysis is more complex. Applying a similar argument to that of the proof of Theorem 3.2, one can show that in the worst case the adversary can ensure that no correct node receives more than $(n-t-d)^2/(n-t)$ ECHO messages. Ensuring that at least one correct node reaches the Byzantine quorum threshold $(n+t)/2$ requires therefore that

$$(n+t)/2 < (n-t-d)^2/(n-t).$$

This leads to a quadratic inequality involving n , t and d , which results in the following constraint on n :

$$n > 2t + 2d + \sqrt{(d+t)^2 + d^2} \geq 3t + 3d.$$

In their analysis [3], Albouy, Frey, Raynal, and Taiani improve on this resilience bound by systematically optimizing the various retransmission and phase thresholds used by Bracha's BRB algorithm, but their solution still falls short of the optimal lower bound $n > 3t + 2d$, which the solution presented in this paper provides.

³ Taking into account the initial dissemination of the message m by the broadcaster, which is also impacted by the message adversary, such an algorithm could in fact at most reach a delivery power of $\max(0, (n-t-d) + \ell_{BRB} - (n-t)) = \max(0, \ell_{BRB} - d) = \max(0, \lceil n - t - \left(\frac{n-t}{n-3t-d} + 1\right)d \rceil)$.