

Towards Optimal Parallel Broadcast under a Dishonest Majority

Daniel Collins¹, Sisi Duan², Julian Loss³, Charalampos Papamanthou⁴,
Giorgos Tsimos⁵, and Haochen Wang⁶

¹ Texas A&M University, danielpatcollins@gmail.com

² Tsinghua University and State Key Laboratory of Cryptography and Digital
Economy Security, duansisi@tsinghua.edu.cn

³ CISA Helmholtz Center for Information Security, lossjulian@gmail.com

⁴ Yale University, charalampos.papamanthou@yale.edu

⁵ University of Maryland, tsimos@umd.edu

⁶ Tsinghua University, whc20@mails.tsinghua.edu.cn

Abstract. The parallel broadcast (PBC) problem generalizes the classic Byzantine broadcast problem to the setting where all n nodes broadcast a message and deliver $O(n)$ messages. PBC arises naturally in many settings including multi-party computation. The state-of-the-art PBC protocol, TRUSTEDPBC, is due to Tsimos, Loss, and Papamanthou (CRYPTO 2022), which is secure under an adaptive adversary assuming $f < (1-\epsilon)n$, where f is the number of Byzantine failures and $\epsilon \in (0, 1)$. TRUSTEDPBC focuses on single-bit inputs and achieves $\tilde{O}(n^2\kappa^4)$ communication and $O(\kappa \log n)$ rounds.

In this work, we propose three PBC protocols for L -bit messages, for any size L , that significantly improve TRUSTEDPBC. First, we propose a new extension protocol that uses a κ -bit PBC as a black box and achieves i) communication complexity of $O(Ln^2 + n^3\kappa + \mathcal{P}(\kappa))$, where $\mathcal{P}(\kappa)$ is the communication complexity of the κ -bit PBC, and ii) round complexity same as the κ -bit PBC. By comparison, the state-of-the-art extension protocol for regular broadcast (Nayak et al., DISC 2020) incurs $O(n)$ additional rounds of communication. Next, we propose a protocol that is secure against a static adversary, for κ -bit messages with $O(n^2\kappa^{1+K} + n\kappa^3 + \kappa^4)$ communication and $O(\kappa)$ round complexity, where K is an arbitrarily small constant such that $0 < K < 1$. Finally, we propose an adaptively-secure protocol for κ -bit messages with $\tilde{O}(n^2\kappa^2 + n\kappa^3)$ communication overhead and $O(\kappa \log n)$ round complexity. Notably, our latter two protocols are $\tilde{O}(\kappa^{2-K})$ and $O(\kappa^2)$ times more communication-efficient, respectively, than the state-of-the-art protocols while achieving the same round complexity.

1 Introduction

Byzantine broadcast (BC) is a fundamental primitive for many cryptographic protocols and distributed systems. The goal of BC is to allow a designated sender to distribute its input value such that all honest nodes output the same value,

$ m $	Protocol	Model	Adv.	$f <$	Communication	Rounds
1	BULLETINBC [‡] [20]	bulletin	static	$(1-\epsilon)n$	$\tilde{O}(n^3\kappa^2) (= \tilde{O}(\mathbf{C}^5))$	$O(n)$
	FLOODBC [‡] [6]	trusted	static	$(1-\epsilon)n$	$\tilde{O}(n^2\kappa^3) (= \tilde{O}(\mathbf{C}^5))$	$O(\kappa)$
	BULLETINPBC [20]	bulletin	adaptive	$(1-\epsilon)n$	$\tilde{O}(n^3\kappa^2) (= \tilde{O}(\mathbf{C}^8))$	$O(n \log n)$
	TRUSTEDPBC [20]	trusted	adaptive	$(1-\epsilon)n$	$\tilde{O}(n^2\kappa^4) (= \tilde{O}(\mathbf{C}^6))$	$O(\kappa \log n)$
	PBC ₁ ^{static} (§4)	trusted	static	$(1-\epsilon)n$	$O(n^2\kappa^{1+K} + n\kappa^3 + \kappa^4) = O(\mathbf{C}^4)$	$O(\kappa)$
	PBC ₁ ^{adaptive} (§5)	trusted	adaptive	$(1-\epsilon)n$	$\tilde{O}(n^2\kappa^2 + n\kappa^3) = \tilde{O}(\mathbf{C}^4)$	$O(\kappa \log n)$
L	ANS [3]	trusted	adaptive	$n/2$	$O(n^2L + n^3\kappa) = O(\mathbf{C}^2L + \mathbf{C}^4)$	$O(1)$
	NRSVX [19]	*	*	$(1-\epsilon)n$	$O(n^2L + \mathcal{P}(\kappa) + n^3\kappa + n^4)$ $(= O(\mathbf{C}^2L + \mathcal{P}(\mathbf{C}) + \mathbf{C}^4))$	$O(n)$
	TLP [20]	trusted	adaptive	$(1-\epsilon)n$	$\tilde{O}(n^2\kappa^4L) (= \tilde{O}(\mathbf{C}^6L))$	$O(\kappa \log n)$
	AC [5]	trusted	static	$n/3$	$O(n^2L + n^3 \log^2 n)$ $(= O(\mathbf{C}^2L) + \tilde{O}(\mathbf{C}^3))$	$O(1)$
	PBC _L [*] (§3)	SRS+*	*	$(1-\epsilon)n$	$O(n^2L + n^3\kappa + \mathcal{P}(\kappa))$ $(= O(\mathbf{C}^2L + \mathcal{P}(\mathbf{C}) + \mathbf{C}^4))$	$O(\mathcal{T}(\kappa))$
	PBC _L ^{static} (§3 & §4)	trusted	static	$(1-\epsilon)n$	$O(n^2L + n^3\kappa + n^2\kappa^{1+K} + n\kappa^3 + \kappa^4)$ $(= O(\mathbf{C}^2L + \mathbf{C}^4))$	$O(\kappa)$
	PBC _L ^{adaptive} (§3 & §5)	trusted	adaptive	$(1-\epsilon)n$	$O(n^2L + n^3\kappa) + \tilde{O}(n^2\kappa^2 + n\kappa^3)$ $(= O(\mathbf{C}^2L) + \tilde{O}(\mathbf{C}^4))$	$O(\kappa \log n)$

Table 1: Comparison of the PBC protocols where honest nodes broadcast messages length $\leq |m|$. †PBC that runs n parallel instances. *The assumptions (*bulletin* board PKI, *trusted* PKI and/or structured reference string (*SRS*)) and the adversarial model (*static* or *adaptive*) depend on the underlying κ -bit PBC oracle. $\mathcal{P}(x)$ is the communication complexity of x -bit PBC, and $\mathcal{T}(\kappa)$ is the round complexity of κ -bit PBC. K is an arbitrarily small constant such that $0 < K < 1$. $\tilde{O}(f(n))$ indicates that the complexity of an algorithm is $O(f(n) \cdot \text{poly}(\log n))$ for some polynomial poly . \mathbf{C} captures practical settings where $\mathbf{C} = O(n) \approx O(\kappa)$.

even if a fraction of Byzantine nodes (including, potentially, the sender) fail arbitrarily. In spite of a large body of work studying broadcast with a single sender, in many applications such as multi-party computation (MPC) and verifiable secret sharing (VSS) broadcast is most commonly required *in parallel*, i.e., with every sender broadcasting simultaneously.

Motivated by this observation, Tsimos, Loss, and Papamanthou [20] gave an efficient designated parallel broadcast (PBC) protocol under dishonest majority, TRUSTEDPBC. Denoting n as the number of nodes and κ as the length of a signature, TRUSTEDPBC achieves $\tilde{O}(n^2\kappa^4)$ communication against up to $f < (1-\epsilon)n$ adaptive and malicious corruptions (for some $0 < \epsilon < 1$) under the assumption of a trusted PKI. Compared to naively running n parallel BC instances, TRUSTEDPBC improved substantially the communication with respect to n . While TRUSTEDPBC already achieves improved communication, its communication is still high, especially when n and κ are close. Additionally, TRUSTEDPBC is limited to single-bit inputs.

Our contributions. In this work, we study PBC with L -bit inputs in the synchronous setting assuming $f < (1-\epsilon)n$ where $0 < \epsilon < 1$. The single-bit variants of our PBC protocols simply follow, which also enjoy improved communication. We consider both the static and weakly adaptive adversarial models (adaptive for short). As summarized in Table 1 and Figure 1, we provide three protocols with improved communication: a new extension protocol PBC_L^{*} that achieves




Fig. 1: Overview of our results.

both improved communication and round; a κ -bit PBC $\text{PBC}_{\kappa}^{\text{static}}$ in the static adversary model that achieves improved communication and round; a κ -bit PBC $\text{PBC}_{\kappa}^{\text{adaptive}}$ with improved communication in the adaptive adversary model. Our solutions do not trade factors of κ for factors in n and solely decrease factors in κ .

We begin with a new extension protocol for PBC, PBC_L^* , that reduces the L -bit PBC problem to a κ -bit PBC oracle. Compared to prior extension protocols, e.g., running n BC instances of the protocol of Nayak, Ren, Shi, Vaidya, and Xiang (NRSVX) [19], PBC_L^* achieves both improved communication and round complexity. The adversarial assumption of PBC_L^* depends on the underlying κ -bit PBC oracle. In particular, if the κ -bit PBC is adaptively secure, then so is PBC_L^* .

We then present $\text{PBC}_{\kappa}^{\text{static}}$, a κ -bit PBC protocol in the static adversarial setting. $\text{PBC}_{\kappa}^{\text{static}}$ can be generalized to L -bit PBC. However, using it as a κ -bit PBC in our extension protocol results in a more communication-efficient PBC. Compared to the state-of-the-art protocols BULLETINBC [20] and FLOODBC [6], $\text{PBC}_{\kappa}^{\text{static}}$ enjoys substantially improved communication complexity and the same or better round complexity. The core idea is to reduce the problem of PBC among n nodes to L -bit PBC among a small committee of κ nodes. Based on the most optimal constructions known so far for L -bit PBC, $\text{PBC}_{\kappa}^{\text{static}}$ achieves $O(n^2\kappa^{1+K} + n\kappa^3 + \kappa^4)$ communication and $O(\kappa)$ rounds for κ -bit broadcast, where K is an arbitrarily small constant such that $0 < K < 1$.

Finally, we present $\text{PBC}_{\kappa}^{\text{adaptive}}$, a κ -bit PBC protocol secure under an adaptive adversary. Our starting point for building PBC under an adaptive adversary is TRUSTEDPBC of Tsimos et al. [20], the most efficient 1-bit PBC protocol known so far that achieves $\tilde{O}(n^2\kappa^4)$ communication. We first construct a κ -bit PBC with $O((n^2\kappa^2 + n\kappa^3) \cdot \log^2 n)$ communication, a $O(\kappa^3)$ improvement over that of TRUSTEDPBC for κ -sized messages. Similarly to $\text{PBC}_{\kappa}^{\text{static}}$, we can use $\text{PBC}_{\kappa}^{\text{adaptive}}$ as a κ -bit PBC oracle in our extension protocol to obtain a more communication-efficient L -bit PBC.

2 Preliminaries

Model. We consider a system with n nodes $\{P_1, \dots, P_n\}$, running over authenticated channels. Among the n nodes, f of them may become Byzantine and fail arbitrarily. We assume $f < (1 - \epsilon)n$, where ϵ is a constant and $0 < \epsilon < 1$. Nodes that are not Byzantine are called *honest*. We consider a synchronous network, where there exists an upper bound on the network and message processing delay.

We consider both the static and the adaptive adversary models. In the static model, the adversary corrupts nodes prior to the start of the protocol. In the adaptive model, the adversary can choose the set of corrupted nodes at any moment during the execution of the protocol based on its current state. In this work, we focus on the weakly adaptive adversary model, where the adversary cannot perform “after-the-fact-removal” and retroactively erase the messages the node sent before they become corrupted. Additionally, we restrict the adversary by assuming *atomic sends* [7] where an honest node P_i can send to multiple nodes simultaneously, without the adversary being able to corrupt P_i in between sending to two nodes.

We assume a trusted setup unless otherwise specified, where a trusted party generates and distributes keys to the nodes prior to the protocol execution.

Normalizing security parameters. Let κ denote the cryptographic security parameter, i.e., the length of hashes or digital signatures. In this work, we also use λ as the statistical parameter. We may consider $\lambda = O(\kappa)$, as typically $\lambda < \kappa$. We can also say that the security parameter of the system is the maximum of λ and the cryptographic security parameter (e.g., length of digital signatures). When we discuss the concrete complexities in the main body of the paper, we differentiate λ and κ . Also note that κ and λ are independent of n and we usually assume $n \gg \kappa$ and $n \gg \lambda$. In Table 1, we provide \mathbf{C} assuming $O(n) \approx O(\kappa) \approx O(\lambda)$ for the ease of understanding.

2.1 Definitions

Parallel broadcast (PBC). In a system with n nodes $\{P_1, \dots, P_n\}$, PBC executes n parallel BC, where each node P_i provides an input v_i and outputs an n -value vector \mathbf{v}_i . Each slot s in \mathbf{v}_i is dedicated for the value broadcast by P_s , the output of which is denoted as $\mathbf{v}_i[s]$. In this work, we study PBC with both 1-bit inputs and L -bit inputs where $L > 1$.

Definition 1 (*f*-Secure Parallel Broadcast). Let Π be a protocol executed by nodes $\{P_1, \dots, P_n\}$, where each node P_i holds an input v_i and each node outputs a n -size vector \mathbf{v}_i . Π should achieve the following properties with probability $1 - \text{negl}(\kappa)$ whenever at most f nodes are corrupted.

- ***f*-Validity:** If P_s is honest, the output \mathbf{v}_i at any honest node P_i satisfies $\mathbf{v}_i[s] = v_s$.
- ***f*-Consistency:** All honest nodes output the same vector \mathbf{v}' .

We will need an *external validity* property for some of our constructions defined as follows. There exists a predicate $Q(\cdot)$ known by all nodes. Given a value v , every node can query $Q(v)$ to validate v . In the literature, the value v can be validated by via some additional data such as digital signatures [9]. Alternatively, v can be validated according to the local state of some nodes [1, 13]. In this case, we may call the predicate a *locally validated predicate*. We use the state-based predicate in this paper.

<p>Initialization:</p> <ul style="list-style-type: none"> - Mining probability p_{mine}. - Let $call_i \leftarrow \perp$ for any $i \in [n]$ <hr/> <p>On input $\mathcal{F}_{mine}(\text{type}, val, i)$ from node P_i:</p> <ul style="list-style-type: none"> - If $call_i = \perp$, output $b = 1$ with probability p_{mine}, or $b = 0$ with probability $1 - p_{mine}$ and set $call_i = b$. - Else output $call_i$. <p>On input $\mathcal{F}_{mine}.\text{verify}(\text{type}, val, j)$ from node P_i:</p> <ul style="list-style-type: none"> - If $call_j = 1$, output 1, otherwise output 0.
--

Fig. 2: Functionality \mathcal{F}_{mine} . val can be \perp or consists of multiple values.

- ***f-External validity:** Given a predicate Q , any honest node P_i that terminates outputs a value \mathbf{v}_i such that for each $\mathbf{v}_i[s] \neq \perp$, $Q(\mathbf{v}_i[s])$ holds by at least one honest node.*

Protocol naming convention PBC_x^y . To differentiate the protocols we study in this paper, we use the PBC_x^y to denote a PBC protocol where each node provides an x -size input that is secure under y model. For example, $\text{PBC}_1^{\text{static}}$ denotes a 1-bit PBC assuming a static adversary.

2.2 Building Blocks

We review the building blocks. Due to space limitations, we provide detailed definitions and descriptions where relevant in our full paper.

Aggregate signatures. We assume recursively combinable aggregate signatures of size $O(\kappa + S \log n)$ when signed by S nodes or $O(\kappa + n)$ using a bitmask, e.g., using BLS signatures based on pairings in the random oracle model [8] or a signature scheme and generic zero-knowledge proofs.

The \mathcal{F}_{mine} oracle. We follow prior work [2, 11, 20] and define the \mathcal{F}_{mine} ideal functionality that we use for random committee selection. \mathcal{F}_{mine} is parameterised by the total number of nodes and a *mining* probability p_{mine} . \mathcal{F}_{mine} provides two interfaces: \mathcal{F}_{mine} and $\mathcal{F}_{mine}.\text{verify}()$, as illustrated in Figure 2. For our static PBC, this can be implemented by nodes multicasting $O(\kappa)$ -sized proofs using an SRS [2, 14]. For our adaptive PBC, we assume generic zero-knowledge proofs for composing signature aggregation and \mathcal{F}_{mine} proofs (or for our protocols, proving committee membership), which also can be instantiated using an SRS [14].

Erasure codes. An (m, n) erasure coding scheme over a data block M is specified by two algorithms (`encode`, `decode`). The `encode` algorithm takes as input m data fragments of M , and outputs $n > m$ coded fragments. The `decode` algorithm takes as input any m -size subset coded fragments and outputs the original data block containing m data fragments. Namely, if $\mathbf{d} \leftarrow \text{encode}(M)$ and $\mathbf{d} = [d_1, \dots, d_n]$, then $\text{decode}(d_{i_1}, \dots, d_{i_m}) = M$ for any distinct $i_1, \dots, i_m \in [1..n]$.

Erasure coding proof (ECP) system. The idea of ECP [4] is to allow the encoder to prove succinctly and non-interactively that an erasure-coded fragment is consistent with a commitment to the original data block. Consider an (m, n)

erasure code that encodes a message M into a set of n fragments d_1, d_2, \dots, d_n . An ECP system is designed to allow for efficient dispersal of these fragments. A proof contains two parts: a constant-sized commitment ϕ plus a per-node witness π_i that is around size $O(|M|/m + \kappa)$. Together, ϕ and π_i convince node P_i that d_i is the correct data fragment for the message committed to by ϕ .

An ECP system consists of three algorithms:

- **setup**. The **setup** algorithm receives a security parameter κ and sets up the system parameters pp .
- **prove_{pp}**. The **prove_{pp}** algorithm takes as input a block of data M and outputs $(\phi, \mathbf{d}, \boldsymbol{\pi})$ where $|\mathbf{d}| = |\boldsymbol{\pi}| = n$. Here, ϕ is a (computationally) binding commitment to all erasure-coded fragments $\mathbf{d} \leftarrow \text{encode}(M)$, and each π_i is intended to serve as a proof that the corresponding d_i is the i -th data fragment with respect to the commitment ϕ .
- **verify_{pp}**. The **verify_{pp}** algorithm takes as input (ϕ, d_i, π_i) and outputs a bit. If **verify_{pp}** $(\phi, d_i, \pi_i) = 1$, then we say d_i is a valid fragment w.r.t. ϕ .

A secure ECP system achieves *EC-correctness* and *EC-consistency*. We use ECP-1 in this work, one of the two constructions provided in the paper. Under the trusted setup assumption (relying on a trusted setup to generate a powers-of-tau structured reference string [15] or SRS hereafter) and when the data block is of at least length $O(\kappa)$, the size of the witness has the same length as each data fragment.

Forward-secure public-key encryption (FS-PKE). A forward-secure public-key encryption scheme [10], or FS-PKE, is a probabilistic public-key cryptosystem that additionally allows the secret key to be updated such that previous keys and encrypted plaintexts cannot be derived from an updated key. It consists of algorithms **gen**, **enc**, **dec** and **upd**, the first three as in standard PKE, and the last updating the secret key into a new *epoch*. We require an appropriate IND-CCA security notion where a challenge is made in epoch j , and the adversary has access to the secret key in any epoch $i > j$. FS-PKE can be implemented using pairings with $O(\kappa \log E)$ -sized keys and constant-sized ciphertexts to support E epochs [10].

3 PBC_L^* : An Extension Protocol for L -bit PBC

3.1 Technical Overview

We present a new extension protocol for L -bit PBC. PBC_L^* reduces L -bit PBC to a κ -bit PBC PBC_κ^* and uses an ECP system. We only require that the κ -bit PBC protocol is transformed into a *validated* PBC by adding a locally validated predicate to PBC_κ^* . We show that such a validated PBC can be easily achieved in our protocol (cf. Lemma 1). As described above, ECP works like an accumulator scheme for erasure coding and can be used to determine if a given fragment corresponds to the original data block. Our extension protocol achieves improved communication compared to prior extension protocols. Additionally, the round complexity remains essentially the same as the κ -bit PBC, incurring three extra

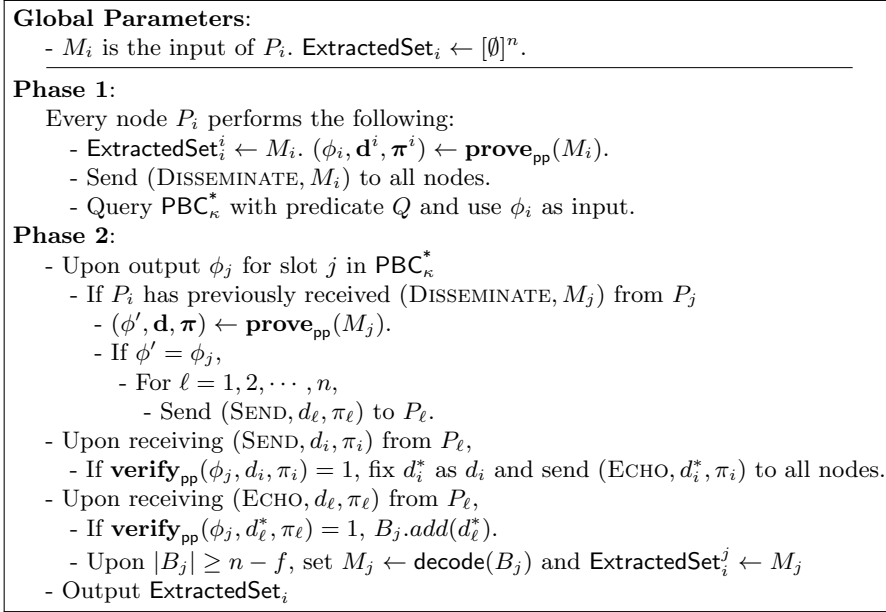


Fig. 3: The PBC_L^* protocol. Q is the locally validated-predicate evaluated within PBC_κ^* defined as follows: $Q(\phi_i)$ is valid for a node P_j if, during execution, P_j has previously received M_i from P_i such that for $(\phi, \mathbf{d}, \boldsymbol{\pi}) \leftarrow \mathbf{prove}(M_i)$ it holds that $\phi_i = \phi$.

rounds of communication. In contrast, the most communication-efficient extension protocol known so far (for BC rather than PBC) [19] incurs $O(n)$ rounds on top of the underlying κ -bit BC oracle.

Our extension protocol is secure under an adaptive adversary, as long as PBC_κ^* is adaptively secure. The same paradigm can also be extended to obtain a communication-efficient L -bit extension protocol for (non-parallel) BC.

3.2 The Extension Protocol

The pseudocode of our extension protocol is shown in Figure 3. Briefly speaking, each node first disseminates its input to all nodes. Then, it queries PBC_κ^* and uses the ECP commitment ϕ_i as the input. The predicate we add to PBC_κ^* ensures that if PBC_κ^* completes, at least one honest node holds M_i . Finally, after PBC_κ^* outputs some value, nodes reconstruct the data via two communication rounds.

It is worth mentioning that in our construction, we view ECP as a tailored and computation-efficient *proof* that proves the correctness of the encoding function of erasure coding. There exists more generic approach that achieves the same communication complexity as our approach, e.g., using zero knowledge proofs [12] to prove the correctness of the encoding function.

Our protocol has the following properties. Note that proofs of claims made hereafter are available in the full version of this work.

Lemma 1. *The PBC_κ^* protocol with predicate Q satisfies f -external validity.*

Theorem 1. *Assuming an SRS, the PBC_L^* protocol presented in Figure 3 satisfies f -validity and f -consistency with probability $1 - \text{negl}(\lambda)$.*

Theorem 2. *The PBC_L^* protocol achieves $O(n^2L + n^3\kappa + \mathcal{P}(\kappa))$ communication and the round complexity is asymptotically the same as PBC_κ^* , where $\mathcal{P}(\kappa)$ is the communication complexity of PBC_κ^* .*

4 $\text{PBC}_\kappa^{\text{static}}$: Efficient PBC under a Static Adversary

4.1 Technical Overview

We present $\text{PBC}_\kappa^{\text{static}}$, a two-layer protocol that reduces the PBC problem to best effort broadcast and a $\mathcal{C}()$ protocol among λ committee members. Although $\text{PBC}_\kappa^{\text{static}}$ itself can clearly be generalized to an L -bit PBC, we obtain a more efficient L -bit PBC integrating $\text{PBC}_\kappa^{\text{static}}$ with our extension protocol.

Our motivation is a tempting solution for committee sampling based protocols: the committee members can *reach an agreement* on some value and then convey the results to all nodes. While this is feasible for a system in the honest majority setting [2,16–18], there is no straightforward way to properly convey the results under a corrupt majority, an observation also made in prior works [11,21].

Our $\text{PBC}_\kappa^{\text{static}}$ protocol makes the above tempting solution work under a corrupt majority. We use a reduction from PBC to a so-called $\mathcal{C}()$ protocol, among $\lambda = O(\kappa)$ committee members. By carefully defining the security properties of $\mathcal{C}()$ and building an efficient construction from L -bit PBC among λ committee members, we ensure that if an honest committee member sees some value output by $\mathcal{C}()$ for the first time, so does any other honest committee member. Accordingly, the maximum number of interactions between an honest node and committee members is bounded by a constant. In particular, $\mathcal{C}()$ takes as input a vector of sets of ‘valid’ (defined below) messages, and outputs a vector of sets corresponding to the union of valid messages that were input.

Briefly speaking, our protocol roughly works as follows. Each node first disseminates its input to all nodes. Then they proceed in a constant number of rounds. In each round, every node first sends its current received values to the committees. Then, the committees query the $\mathcal{C}()$ protocol. After $\mathcal{C}()$ terminates, committee members create signatures for the values they have seen and send them to all nodes. Finally, nodes merge the signatures they receive. At the end of the protocols, for each P_s , every honest node P_i delivers some value from P_s only after P_i has received a sufficiently large number of signatures from the committee.

4.2 The $\mathcal{C}()$ Protocol

To achieve the goals mentioned above, the input of $\mathcal{C}()$ needs to be *validated* and the output needs to be *verifiable* [9]. For the protocol to be validated, the input

message \mathbf{M} must satisfy a global predicate. In our case, for \mathbf{M} to be validated, it must consist of n vectors of valid $(r - 1)$ -s batches, as defined below. For the protocol to be verifiable, the output message, once sent to an honest node, should also be a valid $(r - 1)$ -s batches.

In our PBC protocol, each committee member P_i receives M_j from each node P_j , where $j \in [n] = \{1, 2, \dots, n\}$ and M_j is an n -value vector in the form of $[M_j^1, \dots, M_j^n]$. Each M_j^k is either \perp or consists of up to two valid $(r - 1)$ -s batches. To facilitate the exposition of our protocol, we provide some definitions.

Definition 2. (Valid r -s batch). A valid r -s batch on a message/slot pair (u, s) for (some round) $r \geq 0$ is in the form of $u||s||\text{SIG}_r$, where $u \in \{0, 1\}^L$, $s \in [n]$, and SIG_r is a set of signatures that contains one signature from P_s and $\frac{3r(\epsilon-\mu)(1-\epsilon)}{\mu^2} \log \frac{1}{\delta}$ signatures on $[u, s]$ from members in the committee, where μ is a small constant such that $0 < \mu < \epsilon$ and δ is the desired failure probability.

Definition 3 (Valid tuple for round r). A valid tuple M_i for round $r \geq 1$ is in the form of $[M_i^1, \dots, M_i^n]$ where each M_i^j is either \perp , or consists of at most two valid $(r - 1)$ -s batches, one for a pair (u, j) and one for a pair (u', j) .

We now specify the input and output of the $\mathcal{C}()$ protocol as follows. In some round r , the input of each node P_i for the $\mathcal{C}()$ protocol is \mathbf{M} which consists of up to n vectors $\{M_1, \dots, M_n\}$. Any $M_j \in \mathbf{M}$ is sent by node P_j . Each M_j is validated if it is a valid tuple for round r . After running the $\mathcal{C}()$ protocol, each honest committee member P_i outputs an n -value vector Merged_i . Merged_i is verified if it is a valid tuple for round r . An interesting finding is that we can build $\mathcal{C}()$ from a κ -bit PBC among λ committee members.

Definition 4 (t -Secure $\mathcal{C}()$). Let $\mathcal{C}()$ be a protocol executed by c nodes $\{P_1, \dots, P_c\}$, as specified above. $\mathcal{C}()$ should satisfy the following properties for some round r with probability $1 - \text{negl}(\kappa)$ whenever at most t nodes are corrupted.

- **t -Validity:** If an honest node P_i provides \mathbf{M} as input, any valid tuple $M_j \in \mathbf{M}$ for some round r is part of Merged_k for any honest node P_k in this round.
- **t -Consistency:** For each slot $s \in [n]$, if an honest node P_i outputs Merged_i^s , another honest node P_j outputs Merged_j^s , $\text{Merged}_i^s = \text{Merged}_j^s$.

4.3 The $\text{PBC}_{\kappa}^{\text{static}}$ Protocol

We present the workflow of $\text{PBC}_{\kappa}^{\text{static}}$ in Figure 5. The protocol is round-based, starting from round 0 to round R where $R = \lceil \frac{(1-\epsilon+\mu)c+1}{(\epsilon-\mu)c} \rceil = O(\frac{1}{\epsilon-\mu})$, i.e., a constant number. Each round consists of three *mini-rounds*, the first and the third using a constant number of message delays, and the second being $\mathcal{C}()$ executed by the committee ($O(\lambda)$ round complexity).

We optimize the 1st mini-round and reduce the communication from $O(n^2\kappa\lambda)$ to $O(n^2\kappa\lambda^K)$, where $0 < K < 1$ via a new sampling protocol. Additionally, using ECP as a building block, we can reduce the communication of the 3rd mini-round from $O(n^2\kappa\lambda)$ to $O(n^2\kappa+n\kappa\lambda)$. We provide a detailed description of the protocol and the optimizations in the full version of this work.

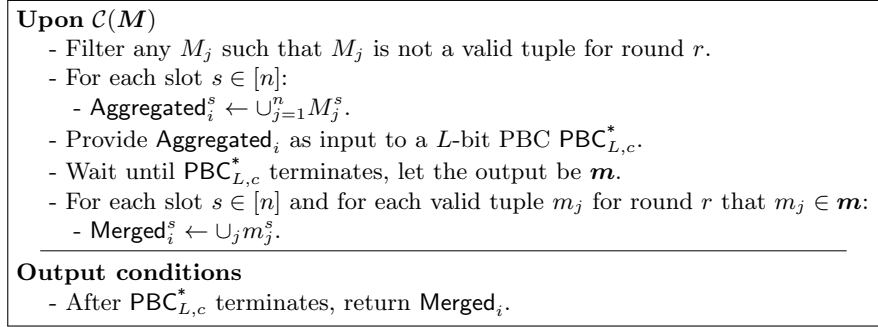


Fig. 4: The $\mathcal{C}()$ protocol.

Theorem 3. *Assuming a trusted PKI and SRS, $\text{PBC}_{\kappa}^{\text{static}}$ is an f -Secure Parallel Broadcast protocol with probability $1 - \text{negl}(\lambda)$.*

Theorem 4. *Assuming a trusted PKI and SRS, the $\text{PBC}_{\kappa}^{\text{static}}$ protocol has $O(\lambda)$ round complexity and $O(n^2\kappa\lambda^K + n\kappa\lambda^2 + \kappa\lambda^3 + \lambda^4)$ communication complexity.*

5 $\text{PBC}_{\kappa}^{\text{adaptive}}$: Efficient PBC under an Adaptive Adversary

We present our adaptively-secure PBC protocol that, for κ -sized messages, has a communication complexity of $\tilde{O}(n^2\kappa^2 + n\kappa^3)$ given $\kappa = O(\lambda)$, or in general $\tilde{O}(n^2\kappa\lambda + n^2\lambda^2 + n\kappa\lambda^2 + n\lambda^3)$. By direct application of our adaptive extension protocol PBC_L^* from §3, we obtain an L -bit protocol with a communication complexity of $O(n^2L + n^3\kappa) + \tilde{O}(n^2\kappa^2 + n\kappa^3)$ bits.

5.1 Review of TrustedPBC

We briefly review TRUSTEDPBC [20], our starting point of $\text{PBC}_{\kappa}^{\text{adaptive}}$. Every node P_i holds a bit b_i as input and outputs a vector of values \mathbf{v}_i . For each bit b_j and slot j , r signatures on $[b_j, j]$ (including one from P_j) from the committee members are collectively called a valid r -batch on $[b_j, j]$. The protocol is round-based and proceeds as follows.

- In round $r = 1$, every node P_i creates a signature for $[b_i, i]$ and sends to all nodes. Each node now holds a vector of valid 1-batches, denoted as M_i , which continues to be updated throughout the protocol.
- In rounds $r = 2, \dots, 2\kappa/\epsilon$, there are two mini-rounds in each round.
 - In the first mini-round, each node P_i executes a \mathcal{M} -DistinctConverge protocol, the idea being for honest nodes to disseminate their received valid r -batches to all other nodes. Their implementation of \mathcal{M} -DistinctConverge uses $\tilde{O}(n^2\kappa^3)$ total communication and $\lceil \log(\epsilon \cdot n) \rceil$ rounds. It is an iterative gossiping protocol, where each step a node sends each message it propagates to $O(\kappa)$ other nodes, where the number of nodes that has seen a given message doubles each round (thus converging in a logarithmic number of rounds).

Global Parameters:

- Let u_i be the input of P_i . Set $\text{ExtractedSet}_i \leftarrow [\emptyset]^n$ and $\text{VotedSet}_i \leftarrow [\emptyset]^n$.

Round 0:

- $\text{ExtractedSet}_i^i \leftarrow u_i$. Send $(\text{SIGN}, u_i, \sigma_i)$ to all where σ_i is a signature on $[u_i, i]$.
- Upon receiving a $(\text{SIGN}, u_j, \sigma_j)$ from P_j , add σ_j to Received_i^j .
- Query $b \leftarrow \mathcal{F}_{\text{mine}}(\text{static}, i)$, if $b = 1$, broadcast (COM, i) to all nodes.
- Upon receiving (COM, j) s.t. $\mathcal{F}_{\text{mine}}.\text{verify}(\text{static}, j) = 1$, add P_j to committee.

Round $r = 1, \dots, R$: each round has three mini-rounds.

1st mini-round: Every node P_i performs the following:

- Send $(\text{ECHO}, \text{Received}_i)$ to all committee members
- Upon receiving $(\text{ECHO}, \text{Received}_j)$ from P_j , $M[j] \leftarrow \text{Received}_j$.

2nd mini-round: Every committee member queries $\mathcal{C}(M)$ and obtains Merged_i .

3rd mini-round: Set $\text{Received}_i \leftarrow [\perp]^n$. For each $s \in [n]$:

Every committee member P_i performs the following:

- Send $(\text{SEND}, \text{Merged}_i)$ to all.
- If a valid $(r-1)$ -s batch on $[u_i^s, s]$ is included in Merged_i^s but $u_i^s \notin \text{VotedSet}_i^s$:
- Set $\text{VotedSet}_i^s \leftarrow \text{VotedSet}_i^s \cup u_i^s$, create a signature for $[u_i^s, s]$ and send to all.
- If at least two valid $(r-1)$ -s batches on different pairs $[u_i^s, s]$ and $[v_i^s, s]$ are included in Merged_i^s and $u_i^s, v_i^s \notin \text{VotedSet}_i^s$:
- Extract the first two of valid $(r-1)$ -s batches and send to all.

For every node P_i , upon receiving valid $(\text{SEND}, \text{Merged}_j)$

- Merge the $(r-1)$ -s batches from the (SEND) messages into an n -value vector Received_i s.t. each Received_i^s contains at most two valid r -s batches.
- If there exists u_j^s s.t. a valid r -s batch for $[u_j^s, s]$ is included in Received_i^s and $u_j^s \notin \text{ExtractedSet}_i^s$ and $|\text{ExtractedSet}_i^s| < 2$, $\text{ExtractedSet}_i^s \leftarrow \text{ExtractedSet}_i^s \cup u_j^s$.

Output conditions. At the end of round R , for each slot $s \in [n]$:

- (Event 1) If $|\text{ExtractedSet}_i^s| = 1$ and $\text{ExtractedSet}_i^s = \{u\}$, $v_i[s] \leftarrow u$.
- (Event 2) If $|\text{ExtractedSet}_i^s| = 0$ or 2 , $v_i[s] \leftarrow \perp$.

- Output v_i .

Fig. 5: The $\text{PBC}_{\kappa}^{\text{static}}$ protocol. $R = \lceil \frac{(1-\epsilon+\mu)c+1}{(\epsilon-\mu)c} \rceil$.

- In the second mini-round, whenever an honest committee member P_i observes a valid r -batch in M_i for the first time for slot j , it creates a digital signature for $[b_j, j]$, appends the signature to the valid r -batch, and sends to all nodes. Upon receiving a valid $(r+1)$ -batch on $[b_j, j]$, each node P_i updates its M_i and adds b_j to ExtractedSet_i^s .
- At the end of the protocol, for each slot j , if there is only one value b_j in ExtractedSet_i^j , P_i sets $v_i[j]$ as b_j . Otherwise, P_i sets $v_i[j]$ as a canonical bit \perp . Finally, P_i outputs the vector v_i .

5.2 An Improved \mathcal{M} -DistinctConverge Protocol

Towards describing our protocol, we first formally define the aforementioned \mathcal{M} -DistinctConverge problem below.

Definition 5 (distinct_k function). For any set M , $\text{distinct}_k(M)$ is a subset of M that contains all messages in M with distinct k -bit prefixes.¹

Definition 6 (t -secure \mathcal{M} -DistinctConverge protocol). Let $\mathcal{M} \subseteq \{0, 1\}^*$ be an efficiently recognizable set. A protocol Π executed by n nodes, where every honest node P_i initially holds input set $M_i \subseteq \mathcal{M}$ and constraint set $C_i \subseteq \mathcal{M}$, is a t -secure \mathcal{M} -DistinctConverge protocol if all remaining honest nodes upon termination, with probability $1 - \text{negl}(\kappa)$, output a set

$$S_i \supseteq \text{distinct}_k \left(\bigcup_{P_i \in \mathcal{H}} M_i - \bigcup_{P_i \in \mathcal{H}} C_i \right),$$

when at most t nodes are corrupted and where \mathcal{H} is the set of honest nodes at the beginning of the protocol.

We build a new \mathcal{M} -DistinctConverge protocol (that improves the same function in TRUSTEDPBC) to reduce the overall communication from $\tilde{O}(n^2\kappa^4)$ to $\tilde{O}(n^2\kappa^3)$. The novelty is to provide a new way for nodes to disseminate the messages in each round via a more efficient *sampling* approach. We first present a Propagate sub-protocol, and define its ideal functionality in Figure 6.

Let n be the number of nodes and $m = 10/\epsilon + \kappa$. For every node $i \in [n]$, $\mathcal{F}_{\text{prop}}$ keeps a set O_i which is initialized to \emptyset . Let M_i be node i 's input messages' set.

- On input (**SendRandom**, M_i) by honest node i :
 - If $|M_i| < \log n$, then for all $j \in [n]$ set $O_j = M_i$.
 - Else, for all $x \in M_i$ and for all $j \in [n]$ add (i, x) to O_j with probability m/n ;
 - **return** M_i to adversary \mathcal{A} ;
 - **return** O_i to node i .
- On input (**SendDirect**, \mathbf{x}, J) by adversary \mathcal{A} (for a corrupted node i):
 - Add $(i, x[j])$ to O_j for all $j \in J$;
 - **return** O_i to adversary \mathcal{A} .

Fig. 6: Functionality $\mathcal{F}_{\text{prop}}$.

We show the pseudocode of our Propagate protocol in Figure 7. In Propagate, nodes send in a given round a set of messages to $O(\lambda)$ nodes. Like [20], to achieve adaptive security, message lists are encrypted and padded to the same size, preventing the adversary from learning who sent what and, e.g., blocking a single message from being propagated. [20] implement Propagate in two rounds, where in the first round nodes sample fresh public keys that are then used in the second round. To improve concrete efficiency, our protocol reduces this to one round using *forward-secure public-key encryption* [10], since the public key is fixed at initialisation and only secret keys are evolved in a one-way fashion each time Propagate is called, which suffices for security.

¹ For example, for $M = \{01001, 01111, 11000, 10000\}$ we have that $\text{distinct}_2(M) = \{01001, 11000, 10000\}$. Note that distinct_k is an one-to-many function, e.g., $\text{distinct}_2(M)$ is also $\{01111, 11000, 10000\}$.

<p>Input: A set of messages M_i.</p> <p>Output: A set of messages O_i.</p> <p>Global Parameters:</p> <ul style="list-style-type: none"> - PK is a set of FS-PKE public keys inherited from the higher-level protocol. - sk_i is a FS-PKE secret key inherited by the higher-level protocol. <hr/> <p>Upon Propagate(M_i)</p> <ul style="list-style-type: none"> - If $M_i \leq \log n$: - For $\ell \in [n]$: Send (NOENC, M_i) to P_ℓ. - Else: // {let $A_p = 2m M_i /n$} - For $j \in [n]$: - For $x \in M_i$: - Add x to list \mathcal{L}_j with probability m/n. - If $M_p \leq n \log n/\lambda$ and $\mathcal{L}_j > A_p$, set $\mathcal{L}_j \leftarrow \perp$ and $j \leftarrow j - 1$ // {resample} - For $j \in [n]$: - Pad \mathcal{L}_j to size A_p. - If $(\text{pk}, j) \in \text{PK}$: $\text{ct}_j \leftarrow \text{enc}(\text{pk}, e, \mathcal{L}_j)$. - Erase \mathcal{L}_j from memory. - If $(\text{pk}, j) \in \text{PK}$: Send ct_j to P_j. <p>Upon Δ time after invoking Propagate(M_i):</p> <ul style="list-style-type: none"> - For all ct_j received from P_j: - $\mathcal{L}_j \leftarrow \text{dec}(\text{pk}, e, \text{sk}_e, \text{ct}_j)$. - If $\mathcal{L}_j \neq \perp$: Add \mathcal{L}_j to O_i. - $\text{sk}_{e+1} \leftarrow \text{upd}(\text{pk}_i, e + 1, \text{sk}_e)$. and erase sk_e from memory. - For all (NOENC, M_j) received from P_j s.t. ct_j was not received from P_j: - Add M_j to O_i. - Output O_i.
--

Fig. 7: Our new Propagate protocol, an instantiation of the propagation process.

The challenge is then to ensure that honest nodes receive all messages with overwhelming probability while reducing communication. We use a *resampling* technique to address the issue. In particular, if any list exceeds a predetermined size (which is approximately $\kappa \times |\text{size of the set of messages to be propagated}|/n$), then the node resamples the list until it does not exceed that size. In our case, we allow nodes to resample their lists $O(n)$ times in the worst case to keep the padding size to a minimum. By contrast, [20] pads lists to a predetermined maximum size in all cases, thereby incurring $O(\lambda)$ more communication than us.

Lemma 2. *If M is the input set of node P in Propagate, the size of each list \mathcal{L}_j is:*

$$\begin{cases} O(\log n), & \text{if } |M| < \log n \\ \leq 2m|M|/n, & \text{else,} \end{cases}$$

with probability $1 - \text{negl}(\lambda)$ if $m = \Theta(\lambda)$.

Input: Sets of messages $M_i \subseteq \mathcal{M}$ and $C_i \setminus \mathcal{M}$, and integer k .

Output: A set of messages S_i .

Round $r = 1, \dots, \lceil \log(\epsilon n) \rceil$:

- $O_i \leftarrow \mathcal{F}_{\text{prop}}(\text{SendRandom}, \text{distinct}_k(M_i - C_i))$.

- For $\ell \in [n]$: Send (NOENC, M_i) to P_ℓ .

- $\text{Local}_i \leftarrow \text{Local}_i \cup O_i$.

- $C_i \leftarrow C_i \cup M_i$.

- $M_i \leftarrow \text{Local}_i \cap \mathcal{M}$.

Output M_p .

Fig. 8: \mathcal{M} -DistinctCV protocol [20, Fig. 7] implementing \mathcal{M} -DistinctConverge.

Lemma 3. *Let s be the length in bits of each message in M for node P . Then, the communication complexity of P during Propagate is with prob. $1 - \text{negl}(\lambda)$:*

$$\begin{cases} O(n \log n \cdot s), & \text{if } |M| < \log n \\ O(m \cdot |M| \cdot s), & \text{else.} \end{cases}$$

Lemma 4. *Assuming a FS-PKE scheme, Propagate is a secure instantiation of the F_{prop} functionality.*

Implementing \mathcal{M} -DistinctConverge. We use the \mathcal{M} -DistinctCV [20, Fig. 7] protocol in our work, besides that we now use our new Propagate protocol, which we provide pseudocode for in Figure 8. Here, nodes input a set of messages M_i and a constraint set C_i . Running for $O(\log n)$ rounds, in each round, nodes first call Propagate with $M_i \setminus C_i$ as input, which outputs a set O_i . The output is appended to a set Local_i . M_i is then added to the constraint set, since there is no need for the caller to propagate them again, and then the set of messages M_i is set to $\text{Local}_i \cap \mathcal{M}$. Ultimately, set M_i is returned.

Lemma 5. *DistinctCV is an adaptively f -secure \mathcal{M} -DistinctConverge protocol, for $f < (1 - \epsilon)n$. The number of bits sent over all nodes is*

$$O\left(\sum_{l=1}^{\log \epsilon n} \cdot \sum_{i \in [n]} CC(\text{Propagate}(M_i^l \setminus C_i^l))\right),$$

where $CC(\text{Propagate}(M_i^l \setminus C_i^l))$ denotes the communication cost of node P_i calling Propagate($(M_i^l \setminus C_i^l)$). Moreover, DistinctCV has $O(\log n)$ round complexity.

5.3 The $\text{PBC}_\kappa^{\text{adaptive}}$ Protocol

We are finally ready to present our $\text{PBC}_\kappa^{\text{adaptive}}$ protocol. Note that we have previously defined the notion of a valid r -s batch for our static PBC protocol in §4. We need an appropriate notion of a valid r -batch as defined below.

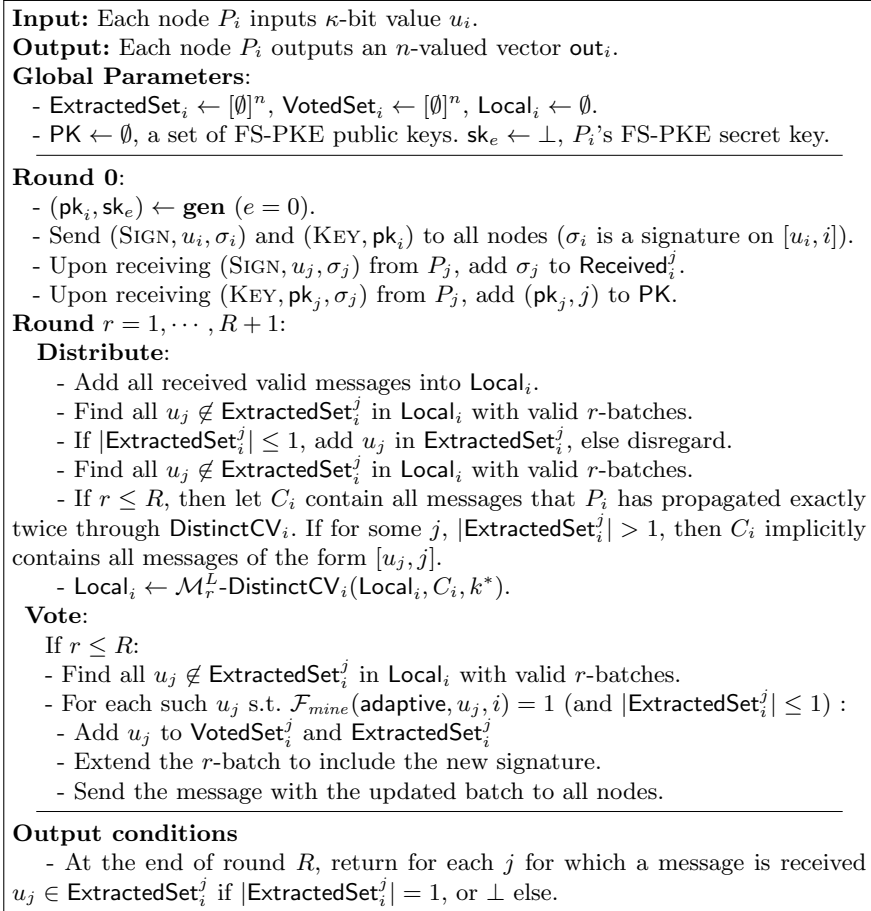


Fig. 9: The $\text{PBC}_{\kappa}^{\text{adaptive}}$ protocol.

Definition 7 ((u, j)-committee). For each message/slot pair (u, j) , the (u, j) -committee is a subset of nodes such that for each node P_i in the (u, j) -committee, whenever $\mathcal{F}_{\text{mine}}$ is queried on input $\mathcal{F}_{\text{mine}}.\text{verify}(\text{adaptive}, u, j)$, $\mathcal{F}_{\text{mine}}$ outputs 1.

Definition 8. Let \mathcal{M}_r denote the set of all possible valid r -batches for all $m \in \{0, 1\}^{\kappa}$ and for all $s \in [n]$.

Lemma 6. Let k^* be the number of bits required to describe $s||m$, where $s \in [n]$ and $m \in \{0, 1\}^{\kappa-1}$ is such that distinguishes between exactly 2 messages in \mathcal{M}_r and where distinct_{k^*} is defined in Definition 5. Then $|\text{distinct}_{k^*}(\mathcal{M}_r)| = 2 \cdot n$.

Definition 9 (Valid r -batch). A valid r -batch on pair (u, j) is the element

$$u||j||\text{SIG}_r,$$

where SIG_r is a set of at least r signatures (or aggregate signature) on $[u, j]$ consisting of one signature from node P_j and at least $r - 1$ signatures from nodes

in the (u, j) -committee (resp. or an aggregate signature with the contributions of P_j and at least $r - 1$ other nodes in the (u, j) -committee).

Our protocol (Figure 9) follows the template of TRUSTEDPBC of [20], which itself follows the template of the broadcast protocol of Chan et al. [11], save for the following notable changes.

First, TRUSTEDPBC is defined only for single-bit PBC. Therefore, we generalize it for multiple nodes, the main difference coming from our use of \mathcal{F}_{mine} . Abstractly, there are an exponential number of possible committees, one per message/slot pair (this number is quadratic in n for TRUSTEDPBC), but since \mathcal{F}_{mine} can be evaluated on-demand for a given input, this is not an issue for complexity. Also, in our protocol, we guarantee that each node will forward at most two messages from the same sender, since they are sufficient to show that the sender is dishonest. Therefore, the size of the message space does not affect the total communication, except for the message length.

Note that at the beginning of protocol execution, nodes send to all nodes their input value u_i and a signature σ_i . Recall in **Propagate** that we use FS-PKE instead of regular public-key encryption. To bootstrap keys, each node therefore sample a FS-PKE key pair and send to all nodes their public key. Recall that we use the DISTINCTCONVERGE protocol in a single round so that each honest node P_i disassembles its message and all honest nodes receive it at the end of this round; TRUSTEDPBC does not use constraint sets to this end.

Then, as shown in Figure 9, each round r is divided into two phases. In the *distribution* phase, nodes propagate r -batches of messages associated with a given node P_j that they have not previously propagated (using DISTINCTCONVERGE), and for any such r -batches, they add the corresponding message to a set ExtractedSet_i^j . In the *voting* phase, nodes check, for each r -batch that they have received in the distribution phase, whether they are in the committee or not for the corresponding message/slot pair using \mathcal{F}_{mine} . If so, and they have not previously added their signature to the r -batch, they do so. Finally, at the end of $R = O(\lambda)$ rounds, nodes output a vector of values for each node P_j , which is \perp if $|\text{ExtractedSet}_i^j| \neq 1$ and the message in ExtractedSet_i^j otherwise.

Theorem 5. *Assuming a trusted PKI and SRS, protocol $\text{PBC}_\kappa^{\text{adaptive}}$ satisfies f -consistency with probability $1 - \text{negl}(\kappa) - \text{negl}(\lambda)$.*

Theorem 6. *Assuming a trusted PKI and SRS, protocol $\text{PBC}_\kappa^{\text{adaptive}}$ satisfies f -validity with probability $1 - \text{negl}(\kappa)$.*

Theorem 7. *Protocol $\text{PBC}_\kappa^{\text{adaptive}}$ has $O(\kappa \log n)$ round complexity and a communication complexity of $O(\log \epsilon n(n^2 \lambda \kappa + n \lambda^2 \kappa + n \lambda^3 \log n + n^2 \lambda^2 \log n))$.*

Acknowledgments

Sisi and Haochen were supported in part by the Beijing Natural Science Foundation under M23015 and National Natural Science Foundation of China under

92267203. Daniel and Giorgos completed part of this work while visiting CISPA, and Daniel while at EPFL, Purdue University and Georgia Institute of Technology, and was supported in part by Sunday Group, Inc. and AnalytiXIN.

References

1. Abraham, I., Asharov, G., Patra, A., Stern, G.: Perfectly secure asynchronous agreement on a core set in constant expected time. In: TCC (2024)
2. Abraham, I., Chan, T.H., Dolev, D., Nayak, K., Pass, R., Ren, L., Shi, E.: Communication complexity of byzantine agreement, revisited. In: PODC (2019)
3. Abraham, I., Nayak, K., Shrestha, N.: Communication and round efficient parallel broadcast protocols. Cryptology ePrint Archive (2023)
4. Alhaddad, N., Duan, S., Varia, M., Zhang, H.: Succinct erasure coding proof systems. Cryptology ePrint Archive (2021)
5. Asharov, G., Chandramouli, A.: Perfect (parallel) broadcast in constant expected rounds via statistical vss. In: EUROCRYPT. pp. 310–339. Springer (2024)
6. Blum, E., Boyle, E., Cohen, R., Liu-Zhang, C.D.: Communication Lower Bounds for Cryptographic Broadcast Protocols. In: DISC (2023)
7. Blum, E., Katz, J., Liu-Zhang, C.D., Loss, J.: Asynchronous byzantine agreement with subquadratic communication. In: TCC. Springer (2020)
8. Boneh, D., Lynn, B., Shacham, H.: Short signatures from the weil pairing. In: International conference on the theory and application of cryptology and information security. pp. 514–532. Springer (2001)
9. Cachin, C., Kursawe, K., Petzold, F., Shoup, V.: Secure and efficient asynchronous broadcast protocols. In: CRYPTO. pp. 524–541. Springer (2001)
10. Canetti, R., Halevi, S., Katz, J.: A forward-secure public-key encryption scheme. Journal of Cryptology **20**, 265–294 (2007)
11. Chan, T.H.H., Pass, R., Shi, E.: Sublinear-round byzantine agreement under corrupt majority. In: PKC. pp. 246–265. Springer (2020)
12. Civit, P., Gilbert, S., Guerraoui, R., Komatovic, J., Monti, M., Vidigueira, M.: Every bit counts in consensus. DISC (2023)
13. Duan, S., Wang, X., Zhang, H.: Fin: Practical signature-free asynchronous common subset in constant time. In: ACM CCS (2023)
14. Groth, J.: On the size of pairing-based non-interactive arguments. In: Advances in Cryptology–EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8–12, 2016, Proceedings, Part II 35. pp. 305–326. Springer (2016)
15. Kate, A., Zaverucha, G.M., Goldberg, I.: Constant-size commitments to polynomials and their applications. In: ASIACRYPT (2010)
16. Katz, J., Koo, C.Y.: On expected constant-round protocols for byzantine agreement. In: CRYPTO. pp. 445–462 (2006)
17. King, V., Saia, J.: Breaking the $o(n^2)$ bit barrier: scalable Byzantine agreement with an adaptive adversary. JACM **58**(4), 18 (2011)
18. King, V., Saia, J., Sanwalani, V., Vee, E.: Scalable leader election. In: SODA (2006)
19. Nayak, K., Ren, L., Shi, E., Vaidya, N.H., Xiang, Z.: Improved extension protocols for byzantine broadcast and agreement. DISC (2020)
20. Tsimos, G., Loss, J., Papamanthou, C.: Gossiping for communication-efficient broadcast. In: CRYPTO (2022)
21. Wan, J., Xiao, H., Shi, E., Devadas, S.: Expected constant round byzantine broadcast under dishonest majority. In: TCC. pp. 381–411. Springer (2020)