

Communication-Optimal Convex Agreement

Diana Ghinea

ETH Zürich
ghinead@ethz.ch

Chen-Da Liu-Zhang

Lucerne University
of Applied Sciences and Arts
& Web3 Foundation
chendaliu@gmail.com

Roger Wattenhofer

ETH Zürich
wattenhofer@ethz.ch

Abstract

Byzantine Agreement (BA) allows a set of n parties to agree on a value even when up to t of the parties involved are corrupted. While previous works have shown that, for ℓ -bit inputs, BA can be achieved with the optimal communication complexity $O(\ell n)$ for sufficiently large ℓ , BA only ensures that honest parties agree on a meaningful output when they hold the same input, rendering the primitive inadequate for many real-world applications.

This gave rise to the notion of Convex Agreement (CA), introduced by Vaidya and Garg [PODC'13], which requires the honest parties' outputs to be in the convex hull of the honest inputs. Unfortunately, all existing CA protocols incur a communication complexity of at least $O(\ell n^2)$. In this work, we introduce the first CA protocol with the optimal communication of $O(\ell n)$ bits for inputs in \mathbb{Z} of size $\ell = \Omega(\kappa \cdot n \log^2 n)$, where κ is the security parameter.

1 Introduction

Reaching collaborative decisions becomes tricky in decentralized systems, especially when participants might be unreliable or even malicious. This is where agreement protocols come in, acting as crucial tools for finding common ground. One such primitive is Byzantine Agreement (BA), enabling n parties to agree on a value even if t of the parties are byzantine.

The standard BA definition comes with certain limitations when applied to real-world scenarios. Consider, for instance, a network of sensors deployed within a cooling room, responsible for reporting the room's temperature. One can expect minor errors in the measurements, such as correct sensors obtaining temperatures between $-10.05^\circ C$ and $-10.03^\circ C$. Standard BA would then allow the parties to agree on a value proposed by the byzantine parties, such as $+100^\circ C$, instead of requiring the output to reflect the correct sensors' measurements. A stronger variant of BA, known as Convex Agreement (CA), addresses this issue, as it requires the honest parties to agree on a value within the convex hull of their inputs (or within the range of their inputs, if the input space is uni-dimensional). CA-related problems have gained significant attention in recent years. The interest is driven by both theoretical curiosity [15, 43, 50] and practical concerns. Real-world applications of CA-related problems span diverse fields, including aviation control systems [36, 47], robotic coordination [44], blockchain oracles [5], transaction ordering in blockchain [14], and distributed machine learning [4, 18, 19, 48].

The synchronous model, where parties have synchronized clocks and messages get delivered within a publicly known amount of time, facilitates a straightforward approach for achieving CA through Synchronous Broadcast (BC, also known as *Byzantine Broadcast*). Essentially, each party sends its input value via BC, which provides the parties with an identical view of the inputs. Afterwards, the parties decide on a common output by applying a deterministic function to the values received. While this approach yields optimal solutions in terms of resilience and round complexity, there is still a gap in terms of communication: if the honest parties hold inputs of at most ℓ bits, a lower bound on the communication complexity is $\Omega(\ell n + n^2)$ bits [41], and this approach incurs a sub-optimal cost of at least $O(\ell n^2)$ bits ($O(\ell n^2 + n^3)$ for deterministic protocols). In fact, the communication complexity of existing CA protocols [15, 41, 50] is adversarially chosen, as they involve steps where honest parties forward messages sent by corrupted parties. In real-world distributed systems, excessive communication can be detrimental: it may lead to network congestion and hence cause messages to be delayed or lost, compromising the system’s reliability.

For regular BA and BC, this issue regarding communication complexity was solved in a line of works [8, 23, 24, 34, 41] via so-called *extension protocols*, that achieve a communication cost of $O(\ell n + \text{poly}(n, \kappa))$ bits, where κ is a security parameter (these protocols make use of cryptographic primitives that compress the input values to $O(\kappa)$ bits). Note that this communication cost is optimal for large enough ℓ . While these results are adequate for real-world scenarios such as fault-tolerant distributed storage systems handling large files, they fall short in scenarios where CA is more suitable than BA. Our work will then focus on closing the communication complexity gap in the synchronous model for CA. Concretely, we ask the following question:

*Can we achieve CA with communication complexity $O(\ell n + \text{poly}(n, \kappa))$,
which is optimal for large enough ℓ ?*

We answer this question in the affirmative. More concretely, we introduce a deterministic protocol in the plain model (i.e. unauthenticated setting) that achieves the optimal resilience $t < n/3$, communication complexity $O(\ell n + n^2 \kappa \log^2 n)$, and round complexity $O(n \log n)$. The protocol makes use of collision-resistant hash functions and takes as inputs bitstrings interpreted as integer values. This is without loss of generality and only used to establish an ordering between the inputs (one could alternatively interpret the inputs being rational numbers with some arbitrary pre-defined precision).

1.1 Related work

Convex Agreement. The requirement of obtaining outputs within the honest inputs’ range has been first introduced in [16] for Approximate Agreement (AA). AA relaxes the agreement requirement, where parties’ outputs may deviate by a predefined error $\varepsilon > 0$. This relaxation allows for deterministic asynchronous protocols, circumventing the FLP result [22]. AA has been a subject of an extensive line of works, focusing on optimal convergence rates [6, 20, 21], higher resilience [1, 26, 33], and different input spaces, such as multidimensional inputs [27, 37, 50], or graphs and abstract convexity spaces [3, 15, 32, 43]. CA was formally defined by Vaidya and Garg in [38, 50] for multidimensional input values. Feasibility with optimal resilience has been considered for abstract convexity spaces as well [15, 43]. Another line of works has investigated the feasibility of an even stronger requirement for inputs in \mathbb{R} , i.e. that the output is *close* to the median of the honest inputs [14, 47], or, more generally, to the k -th lowest honest input [36].

Extension protocols. The problem of reducing the communication complexity of BA on multi-valued inputs was first addressed by Turpin and Coan [49], where the authors assume $t < n/3$ and give a reduction from long-messages BA to short-messages BA with a communication cost

of $O(\ell n^2)$ bits. Fitzi and Hirt [23] later achieve BA in the honest majority setting with the asymptotically optimal communication complexity $O(\ell n + \text{poly}(n, \kappa))$ bits, assuming a universal hash function. Further works have provided error-free solutions focusing on reducing the additional $\text{poly}(n, \kappa)$ factor in the communication complexity both in the $t < n/3$ [24, 35, 41] setting and in the honest-majority setting [8, 24, 41]. Extension protocols have also been a topic of interest for problems related to BA, such as BC in the $t < n$ setting [11, 28], or asynchronous Reliable Broadcast [10, 41].

Protocols for short messages. Reducing the communication complexity is not only a topic of interest for long inputs, but also for short inputs (i.e., one bit or a constant number of bits). Dolev and Reischuk [17] showed that deterministic BA protocols (and hence also deterministic CA protocols) incur communication complexity $\Omega(t^2)$ if t of the parties involved are byzantine, therefore $\Omega(n^2)$ if $t = \Theta(n)$. This lower bound is tight [12, 40]. However, randomized protocols have offered a path to subquadratic communication [2, 9, 13, 25, 29–31] under different assumptions.

To the best of our knowledge, our work introduces the first CA protocol with asymptotically optimal-communication for sufficiently long messages, at least $\ell = \Omega(\kappa \cdot n \log^2 n)$. Our protocol is also deterministic. We leave the question of achieving communication-optimal CA protocols for shorter messages as an interesting open question.

1.2 Comparison to previous works

In terms of techniques, our solution differs significantly from both prior works on BA extension protocols and prior works on CA or AA. In comparison to BA, the honest-range requirement of CA adds a new level of challenges when it comes to reducing the communication. Roughly, in prior works on communication-optimal BA, each party first computes a short κ -bit encoding of its long ℓ -bit input value (using e.g. a hash function). Afterwards, the parties agree on an encoding z^* using a BA protocol for short messages. Finally, parties holding the (unique) input value v^* matching the encoding z^* non-trivially distribute v^* to all the parties. The main issue when trying to adapt this approach to CA is that the short κ -bit encodings lost information about the ordering of the original values, and in particular cannot reflect the honest inputs' range. On the other hand, existing protocols satisfying this validity requirement, regardless of whether they achieve CA or AA, involve some step where all parties send their ℓ -bit values to all other parties. It might seem intuitive that the parties need a possibly consistent or identical view over their actual values to decide on a valid output. However, we show that this intuition is not true.

Our protocol relies on a byzantine variant of the *longest common prefix* problem, and makes use of a BA protocol for short messages as a building block. The central insight behind our approach is that the longest common prefix of the honest parties' inputs represented as bitstrings reveals a subset of the honest inputs' range. While the byzantine parties prevent us from finding the exact longest common prefix of the honest inputs, the longest common prefix of any values in the honest inputs' range will suffice to obtain an output.

2 Preliminaries

We consider a setting with n parties P_1, P_2, \dots, P_n in a fully connected network, where each pair of parties is connected by an authenticated channel. We assume that the network is synchronous: the parties' clocks are synchronized, all messages get delivered within Δ time, and Δ is publicly known. We consider an adaptive adversary that can corrupt up to $t < n/3$ parties at any point in the protocol's execution, causing them to become byzantine: corrupted parties may deviate arbitrarily from the protocol. We consider a security parameter κ , a collision-resistant hash function

$H_\kappa : \{0, 1\}^* \rightarrow \{0, 1\}^\kappa$, and we assume that the adversary is computationally bounded. Informally, $H_\kappa : \{0, 1\}^* \rightarrow \{0, 1\}^\kappa$ is collision-resistant if, for any computationally-bounded adversary \mathcal{A} , the probability that $\mathcal{A}(1^\kappa)$ outputs two values $x \neq y$ such that $H_\kappa(x) = H_\kappa(y)$ is negligible in κ .¹ For simplicity of presentation, our proofs will assume that H_κ is *collision-free*; our protocols are secure conditioned on the event that no collision occurs.

Definitions. We recall the definitions of CA and BA. We mention that, throughout the paper, we will use *valid value* to refer to a value satisfying Convex Validity, as defined below.

Definition 1 (Convex Agreement). *Let Π be an n -party protocol where each party holds a value v_{IN} as input, and parties terminate upon generating an output v_{OUT} . Π achieves Convex Agreement if the following properties hold even when t of the n parties are corrupted: (Termination) All honest parties terminate; (Convex Validity) Honest parties' outputs lie in the honest inputs' convex hull; (Agreement) All honest parties output the same value.*

Definition 2 (Byzantine Agreement). *Let Π be an n -party protocol where each party holds a value v_{IN} as input, and parties terminate upon generating an output v_{OUT} . Π achieves Byzantine Agreement if the following properties hold even when t of the n parties are corrupted: (Termination) All honest parties terminate; (Validity) If all honest parties hold the same input value v , they output $v_{\text{OUT}} = v$; (Agreement) All honest parties output the same value.*

We denote the communication complexity for ℓ -bit inputs of a protocol Π by $\text{BITS}_\ell(\Pi)$: this is the worst-case total number of bits sent by honest parties if they all hold inputs of at most ℓ bits. In addition, $\text{ROUNDS}_\ell(\Pi)$ denotes the worst-case round complexity of Π .

Binary representations. We establish a few notations that will be used throughout the paper. For a value $v \in \mathbb{N}$, we define its binary representation $\text{BITS}(v) := B_1 B_2 \dots B_k$ such that $2^{k-1} \leq v < 2^k$, $B_i \in \{0, 1\}$ for every $1 \leq i \leq k$, and $\sum_{i=1}^k B_i \cdot 2^{k-i} = v$. For $\ell \geq k$, we also define v 's ℓ -bit representation $\text{BITS}_\ell(v)$ as the ℓ -bit string obtained by prepending $\ell - k$ zeroes to $\text{BITS}(v)$. In addition, for $1 \leq i \leq \ell$, $B_\ell^i(v)$ denotes the i -th leftmost bit in the ℓ -bit representation of v $\text{BITS}_\ell(v)$.

The reverse operation of $\text{BITS}(\cdot)$ will be $\text{VAL}(\text{BITS})$: given a bitstring $\text{BITS} := B_1 B_2 \dots B_k$ (where every $B_i \in \{0, 1\}$), $\text{VAL}(\text{BITS}) := \sum_{i=1}^k B_i \cdot 2^{k-i} = v$. We denote the length of a bitstring BITS by $|\text{BITS}|$, and \parallel is the concatenation operator.

We will also need to define $\text{MAX}_\ell(\text{BITS})$ and $\text{MIN}_\ell(\text{BITS})$. $\text{MAX}_\ell(\text{BITS})$ is the highest ℓ -bit value having prefix BITS , obtained by appending $\ell - |\text{BITS}|$ ones to BITS . Similarly, $\text{MIN}_\ell(\text{BITS})$ is the lowest ℓ -bit value having prefix BITS , obtained by appending $\ell - |\text{BITS}|$ zeroes to BITS .

3 CA for Long Inputs in \mathbb{N} of Fixed Length

Building towards our CA protocol for \mathbb{Z} , we first focus on \mathbb{N} . For now, we assume that the inputs' length ℓ is fixed: there is a publicly known ℓ such that every honest input v_{IN} satisfies $v_{\text{IN}} < 2^\ell$. In this section, we present a protocol FIXEDLENGTHCA achieving CA under these assumptions. When $\ell \in \text{poly}(n)$, FIXEDLENGTHCA has communication complexity $O(\ell \cdot n + \text{poly}(n, \kappa))$ and round complexity $O(n \log n)$.

FIXEDLENGTHCA searches for a valid value by only working with values' prefixes. We first use the honest parties' inputs to identify a bitstring PREFIX^* that is the prefix of an ℓ -bit valid value. If PREFIX^* consists of ℓ bits, the parties may output $\text{VAL}(\text{PREFIX}^*)$. Otherwise, we ensure that PREFIX^* satisfies a few special properties enabling the parties to efficiently find an output. Concretely, we ensure that sufficiently many honest parties *know* valid values that do not have

¹See [46] for a formal definition.

PREFIX^{*} as a prefix: such values are either lower than any value with prefix PREFIX^{*}, meaning that $\text{MIN}_\ell(\text{PREFIX}^*)$ is valid, or higher than any value with prefix PREFIX^{*}, meaning that $\text{MAX}_\ell(\text{PREFIX}^*)$ is valid. These parties will announce which of these two options they believe to valid, enabling all parties to decide on the final output.

We split the implementation of FIXEDLENGTHCA into three subprotocols. The first one is FINDPREFIX, where the parties agree on a bitstring PREFIX^{*}, and each party obtains two ℓ -bit valid values v and v_\perp such that: (i) the values v have PREFIX^{*} as a prefix, and (ii) for any bitstring BITS of $|\text{PREFIX}^*| + 1$ bits, there are $t + 1$ honest parties whose values v_\perp do not have BITS as a prefix. If $|\text{PREFIX}^*| = \ell$, then the parties hold the same valid value v , and may simply output v . Otherwise, in our second subprotocol, ADDLASTBIT, the parties append one bit to PREFIX^{*} using their values v , ensuring that the extended PREFIX^{*} is still some valid values' prefix. Now there are $t + 1$ honest parties holding values v_\perp that do not have prefix PREFIX^{*}, and these differences are announced in the third subprotocol, GETOUTPUT, where the parties agree on the final output.

FIXEDLENGTHCA(ℓ, v)

Code for party P

- 1: PREFIX^{*}, v , $v_\perp := \text{FINDPREFIX}(v, \ell)$. If $|\text{PREFIX}^*| = \ell$, output v and terminate.
- 2: PREFIX^{*} := ADDLASTBIT(PREFIX^{*}, v, ℓ).
- 3: Return $v := \text{GETOUTPUT}(\text{PREFIX}^*, v_\perp, \ell)$.

In the following, we present each of these subprotocols in detail.

Subprotocol FindPrefix. Roughly, FINDPREFIX aims to identify the longest common prefix of the *honest parties' inputs* using binary search. Identifying the precise longest common prefix is impossible, as the byzantine parties can act as honest parties with inputs of their own choice. However, we can identify a valid value's prefix that is at least as long as the honest inputs' longest common prefix. Roughly, the parties will be looking for some index i^* such that running BA on their values' i^* -bit prefixes would return an honest prefix, but running BA on their $(i^* + 1)$ -bit prefixes would not offer the same guarantee. We need a BA protocol *for long messages* that satisfies two additional properties, defined below. The first one is *Intrusion Tolerance*: recall that, unless the honest parties hold identical inputs, BA's Validity condition does not impose any restrictions. Intrusion Tolerance requires that the output is either an honest party's input or a special symbol \perp . The second property, *Bounded Pre-Agreement*, prevents agreement on \perp when an honest input can be easily identified.

Definition 3. *Intrusion Tolerance: Honest parties output an honest party's input or \perp .*

Definition 4. *Bounded Pre-Agreement: If the parties agree on \perp , then there are fewer than $n - 2t$ honest parties holding the same input value.*

In Section 7, we describe a protocol $\Pi_{\ell\text{BA}+}$ achieving these guarantees with communication complexity $O(\ell n + \text{poly}(n, \kappa))$, as described in Theorem 1. The main technical challenge behind $\Pi_{\ell\text{BA}+}$ is to design a protocol for κ -bit messages with communication complexity $O(\kappa \cdot n^2)$.

Theorem 1. *Given a BA protocol Π_{BA} resilient against $t < n/3$ corruptions, there is a BA protocol $\Pi_{\ell\text{BA}+}$ resilient against $t < n/3$ corruptions that achieves Intrusion Tolerance and Bounded Pre-Agreement, with communication complexity $\text{BITS}_\ell(\Pi_{\ell\text{BA}+}) = O(\ell n + \kappa \cdot n^2 \log n) + \text{BITS}_\kappa(\Pi_{\text{BA}})$, and round complexity $\text{ROUNDS}_\ell(\Pi_{\ell\text{BA}+}) = O(1) + \text{ROUNDS}_\kappa(\Pi_{\text{BA}})$.*

We proceed in $O(\log \ell)$ iterations. In the first iteration, the parties check whether $\Pi_{\ell\text{BA}+}$ returns \perp on the first half of their values' bitstrings $B_1 \parallel \dots \parallel B_{\text{MID}}$. If $\Pi_{\ell\text{BA}+}$ returns \perp , $\text{MID} - 1$ is

an upper bound for index i^* . The Bounded Pre-agreement property ensures that at most $n - 2t$ honest parties hold values v with the same prefix of MID bits. This means that, for any bitstring BITS of MID bits, there are at least $(n - t) - (n - 2t) \geq t + 1$ honest parties holding values v that do not have BITS as a prefix. The parties set $v_\perp := v$ and then continue the search for i^* within bits $B_1, \dots, B_{\text{MID}-1}$ in the next iteration, using an identical approach.

Otherwise, if $\Pi_{\ell\text{BA}+}$ returns a bitstring of MID bits $\text{PREFIX}_1^* \parallel \dots \parallel \text{PREFIX}_{\text{MID}}^*$, Intrusion Tolerance ensures that this is the prefix of an honest party's (valid) value v^* . The parties holding values v with a different prefix update v to match this prefix: if $\text{VAL}(B_1 \parallel \dots \parallel B_{\text{MID}}) < \text{VAL}(\text{PREFIX}_1^* \parallel \dots \parallel \text{PREFIX}_{\text{MID}}^*)$, meaning that $v < v^*$, then $\text{MIN}_\ell(\text{PREFIX}_1^* \parallel \dots \parallel \text{PREFIX}_{\text{MID}}^*)$ is in $[v, v^*]$ and therefore is valid. Otherwise, if $\text{VAL}(B_1 \parallel \dots \parallel B_{\text{MID}}) > \text{VAL}(\text{PREFIX}_1^* \parallel \dots \parallel \text{PREFIX}_{\text{MID}}^*)$, meaning that $v > v^*$, then $\text{MAX}_\ell(\text{PREFIX}_1^* \parallel \dots \parallel \text{PREFIX}_{\text{MID}}^*)$ is in $[v^*, v]$ and therefore is valid. The parties then proceed to the next iteration, where they continue the search for i^* on the second half of their (updated) values' ℓ -bit representation. After $O(\log \ell)$ iterations, either $\Pi_{\ell\text{BA}+}$ never returned \perp and the parties hold identical values v , or i^* is found. We present the code below.

FINDPREFIX(ℓ, v)

Code for party P

```

1: LEFT := 1, RIGHT :=  $\ell + 1$ ;  $v := v_{\text{IN}}$ ;  $v_\perp := v_{\text{IN}}$ ; PREFIX* := empty string.
2: loop
3:   If LEFT = RIGHT, exit the loop.
4:   ( $B_1, B_2, \dots, B_\ell$ ) := BITS $_\ell(v)$ .
5:   Join  $\Pi_{\ell\text{BA}+}$  with input  $B_{\text{LEFT}} \parallel \dots \parallel B_{\text{MID}}$ , where  $\text{MID} := \lfloor (\text{LEFT} + \text{RIGHT})/2 \rfloor$ .
6:   If  $\Pi_{\ell\text{BA}+}$  has returned  $\perp$ , set  $v_\perp := v$  and  $\text{RIGHT} := \text{MID}$ .
7:   Otherwise, if  $\Pi_{\ell\text{BA}+}$  has returned MID - LEFT + 1 bits  $\text{PREFIX}_{\text{LEFT}}^* \parallel \dots \parallel \text{PREFIX}_{\text{MID}}^*$ :
8:     PREFIX* := PREFIX*  $\parallel$  PREFIX $_{\text{LEFT}}^* \parallel \dots \parallel$  PREFIX $_{\text{MID}}^*$ .
9:     If  $\text{VAL}(B_1 \parallel \dots \parallel B_{\text{MID}}) < \text{VAL}(\text{PREFIX}^*)$ :  $v := \text{MIN}_\ell(\text{PREFIX}^*)$ .
10:    If  $\text{VAL}(B_1 \parallel \dots \parallel B_{\text{MID}}) > \text{VAL}(\text{PREFIX}^*)$ :  $v := \text{MAX}_\ell(\text{PREFIX}^*)$ .
11:    Set LEFT := MID + 1.
12: end loop
13: Return PREFIX*,  $v$ ,  $v_\perp$ .

```

The formal proof of the lemma below is included in Appendix A.2.

Lemma 1. *Assume a BA protocol Π_{BA} , and that honest parties join FINDPREFIX with the same ℓ , and with valid ℓ -bit values v . Then, the honest parties obtain the same bitstring PREFIX*, and each honest party obtains two valid ℓ -bit values v, v_\perp such that: (i) PREFIX* is a prefix of BITS $_\ell(v)$; (ii) for any bitstring BITS of $|\text{PREFIX}^*| + 1$ bits, at least $t + 1$ honest parties hold values v_\perp such that BITS $_\ell(v_\perp)$ does not have prefix BITS.*

FINDPREFIX achieves communication complexity $\text{BITS}_\ell(\text{FINDPREFIX}) = O(\ell \cdot n + \kappa \cdot n^2 \log n \log \ell) + O(\log \ell) \cdot \text{BITS}_\kappa(\Pi_{\text{BA}})$, and round complexity $\text{ROUNDS}_\ell(\text{FINDPREFIX}) = O(\log \ell) \cdot \text{ROUNDS}_\kappa(\Pi_{\text{BA}})$.

Proof Sketch. We consider the properties below. If these properties are satisfied at the beginning of iteration $i \geq 1$ of the loop, then either the stopping condition $\text{LEFT} = \text{RIGHT}$ is met in iteration i , or the properties hold at the beginning of iteration $i + 1$ as well. We also note that these properties hold at the beginning of iteration 1 due to the variables' initialization.

- (A) All honest parties hold the same indices $1 \leq \text{LEFT} \leq \text{RIGHT} \leq \ell + 1$, and the same bitstring PREFIX* consisting of $\text{LEFT} - 1$ bits.
- (B) $0 \leq \text{RIGHT} - \text{LEFT} \leq 2^{\lceil \log_2 \ell \rceil - (i-1)}$.
- (C) Honest parties hold valid ℓ -bit values v such that BITS $_\ell(v)$ has PREFIX* as a prefix.

(D) Honest parties hold valid ℓ -bit values v_\perp , and, for any bitstring BITS of RIGHT bits, there are $t + 1$ honest parties holding values v_\perp such that $\text{BITS}_\ell(v_\perp)$ does not have prefix BITS.

Property (B) implies that the condition $\text{LEFT} = \text{RIGHT}$ is met by iteration $i = \lceil \log_2 \ell \rceil + 2$. Once this condition is met, Property (A) ensures that parties hold the same bitstring PREFIX^* of $\text{LEFT} - 1$ bits. Property (C) ensures that parties hold valid ℓ -bit values v with prefix PREFIX^* . Finally, property (D) ensures that, honest parties hold valid ℓ -bit values v_\perp such that for any bitstring BITS of $\text{RIGHT} = \text{LEFT} = |\text{PREFIX}^*| + 1$ bits, there are $t + 1$ honest parties whose values v_\perp do not have PREFIX^* as a prefix.

Since each iteration invokes $\Pi_{\ell\text{BA}+}$ once, $\text{ROUNDS}_\ell(\text{FINDPREFIX}) = O(\log \ell) \cdot \text{ROUNDS}_\ell(\Pi_{\ell\text{BA}+})$, and applying Theorem 1 gives our claimed round complexity. In each iteration $i < \lceil \log_2 n \rceil + 2$, Property (B) ensures that FINDPREFIX runs $\Pi_{\ell\text{BA}+}$ on bitstrings of at most $2^{\lceil \log_2 \ell \rceil - i} \leq \ell/2^{i-1}$ bits. Therefore, $\text{BITS}_\ell(\text{FINDPREFIX}) = \sum_{i=1}^{\lceil \log_2 \ell \rceil + 1} \text{BITS}_{\ell/2^{i-1}}(\Pi_{\ell\text{BA}+})$. Using Theorem 1 and that $\sum_{i=0}^{\infty} 1/2^i \leq 2$, we obtain our claimed communication cost. \square

Extending the prefix agreed upon. We now describe the subprotocol ADDLASTBIT , where parties extend the bitstring PREFIX^* agreed upon in FINDPREFIX with one bit (assuming $|\text{PREFIX}^*| < \ell$). The resulting bitstring should still be a valid values' prefix. As each party holds a valid value v with prefix $\text{PREFIX}^*||0$ or $\text{PREFIX}^*||1$, we extend PREFIX^* by using a bit BA protocol Π_{BA} .

$\text{ADDLASTBIT}(\ell, v, \text{PREFIX}^*)$

Code for party P .

- 1: Join Π_{BA} with input $B_\ell^{i^*+1}(v)$, where $i^* = |\text{PREFIX}^*|$, and obtain output B^* . Return $\text{PREFIX}^* || B^*$.

The proof of the lemma below is included in Appendix A.2.

Lemma 2. Assume a BA protocol Π_{BA} , and that honest parties join ADDLASTBIT with the same value ℓ , the same bitstring PREFIX^* of $i^* < \ell$ bits, and with valid ℓ -bit values v such that $\text{BITS}_\ell(v)$ has prefix PREFIX^* . Then, the honest parties agree on a bitstring of $i^* + 1$ bits that is the prefix of a valid value's ℓ -bit representation. ADDLASTBIT has communication complexity $\text{BITS}_\ell(\text{ADDLASTBIT}) = \text{BITS}_1(\Pi_{\text{BA}})$ and round complexity $\text{ROUNDS}_\ell(\text{ADDLASTBIT}) = \text{ROUNDS}_1(\Pi_{\text{BA}})$.

Obtaining the final output. After running ADDLASTBIT , the parties hold a bitstring PREFIX^* of $i^* + 1$ bits that is a valid value's prefix. Moreover, $t + 1$ honest parties hold valid values v_\perp that do not have prefix PREFIX^* . These parties' values v_\perp are either lower than $\text{MIN}_\ell(\text{PREFIX}^*)$ or higher than $\text{MAX}_\ell(\text{PREFIX}^*)$, meaning that at least one of these two options is valid. Each of these parties may announce (by sending a bit) which of the two options it believes to be valid. Then, every party becomes aware of a valid option by looking at the bit received the most (and therefore sent by at least one honest party). Afterwards, the parties use Π_{BA} to agree on a valid option.

$\text{GETOUTPUT}(\ell, v_\perp, \text{PREFIX}^*)$

Code for party P

- 1: If PREFIX^* is not a prefix of $\text{BITS}_\ell(v_\perp)$:
- 2: Set $B := 0$ if $v_\perp < \text{MIN}_\ell(\text{PREFIX}^*)$ and $B := 1$ otherwise.
- 3: Send B to all parties.
- 4: Set $m :=$ the number of bits B received, $\text{CHOICE} :=$ a bit B received from $\lceil m/2 \rceil$ parties.
- 5: Join Π_{BA} with input CHOICE . If the bit agreed upon is 0, return $\text{MIN}_\ell(\text{PREFIX}^*)$. Otherwise, return $\text{MAX}_\ell(\text{PREFIX}^*)$.

The proof of the lemma below is included in Appendix A.2.

Lemma 3. Assume a BA protocol Π_{BA} , and that honest parties join GETOUTPUT with the same value ℓ , and with the same bitstring PREFIX^* representing the prefix of some valid value's ℓ -bit representation. In addition, assume that each party joins with some valid ℓ -bit input v_\perp such that the ℓ -bit representations of $t + 1$ honest parties' values v_\perp do not have PREFIX^* as a prefix. Then, the honest parties obtain the same valid value v_{OUT} .

GETOUTPUT has communication complexity $\text{BITS}_\ell(\text{GETOUTPUT}) = O(n^2) + \text{BITS}_1(\Pi_{\text{BA}})$ and round complexity $\text{ROUNDS}_\ell(\text{GETOUTPUT}) = O(1) + \text{ROUNDS}_1(\text{BA})$.

Protocol Analysis. We have now reached the final theorem of this section, which presents the guarantees of FIXEDLENGTHCA. We highlight that, when $\ell \in \text{poly}(n)$ and therefore $\log \ell \in O(\log n)$, we achieve communication complexity $\text{BITS}_\ell(\text{FIXEDLENGTHCA}) = O(\ell n + \kappa \cdot n^2 \log^2 n) + O(\log n) \cdot \text{BITS}_\kappa(\Pi_{\text{BA}})$, and round complexity $\text{ROUNDS}_\ell(\text{FIXEDLENGTHCA}) = O(\log n) \cdot \text{ROUNDS}_\kappa(\Pi_{\text{BA}})$.

Theorem 2. Assume a BA protocol Π_{BA} resilient against $t < n/3$ corruptions. Then, if the honest parties hold ℓ -bit inputs $v_{\text{IN}} \in \mathbb{N}$ and ℓ is publicly known, FIXEDLENGTHCA is a CA protocol resilient against $t < n/3$ corruptions, with communication complexity $\text{BITS}_\ell(\text{FIXEDLENGTHCA}) = O(\ell n + \kappa \cdot n^2 \log n \log \ell) + O(\log \ell) \cdot \text{BITS}_\kappa(\Pi_{\text{BA}})$, and round complexity $\text{ROUNDS}_\ell(\text{FIXEDLENGTHCA}) = O(\log \ell) \cdot \text{ROUNDS}_\kappa(\Pi_{\text{BA}})$.

Proof. Lemma 1 ensures that FINDPREFIX enables the parties to agree on a bitstring PREFIX^* , and provides them with valid ℓ -bit values v, v_\perp such that the values v have prefix PREFIX^* . If $|\text{PREFIX}^*| = \ell$, the honest parties hold the same valid value v , and therefore CA is achieved. Otherwise, Lemma 2 ensures that parties obtain the same bitstring PREFIX^* such that there are $t + 1$ honest parties whose values v_\perp do not have PREFIX^* as a prefix. Then, GETOUTPUT's preconditions are met, and Lemma 3 ensures that CA is achieved. The communication complexity and the round complexity follow from summing up the complexities of each subprotocol. \square

4 CA for Very Long Inputs in \mathbb{N} of Fixed length

The protocol FIXEDLENGTHCA presented in Section 3 achieves our communication complexity goal when $\ell \in \text{poly}(n)$. We now consider values $\ell \geq n^2$ that may not satisfy this condition, and we build a round-efficient CA protocol with communication complexity $O(\ell n + \text{poly}(n, \kappa))$ by making small adjustments to the protocol of Section 3. We maintain the same assumptions: parties hold ℓ -bit inputs in \mathbb{N} , and ℓ is publicly known.

We first describe the adjustments for FINDPREFIX. Instead of comparing substrings of *bits* in each iteration, we compare substrings of *blocks*. For simplicity, we assume that ℓ is a multiple of n^2 . Each party will split its ℓ -bit value into n^2 blocks $\text{BLOCK}_1, \text{BLOCK}_2, \dots, \text{BLOCK}_{n^2}$ of ℓ/n^2 bits each. For an ℓ -bit value $v \in \mathbb{N}$, we define $\text{BLOCKS}(v) := (\text{BLOCK}_1, \text{BLOCK}_2, \dots, \text{BLOCK}_{n^2})$ such that $\text{BITS}_\ell(v) = \text{BLOCK}_1 \parallel \text{BLOCK}_2 \parallel \dots \parallel \text{BLOCK}_{n^2}$, and, for any $1 \leq i \leq n^2$, $|\text{BLOCK}_i| = \ell/n^2$. We denote BLOCK_i by $\text{BLOCK}_i(v)$, and we use the term *block* to refer to such sequences of ℓ/n^2 bits.

Then, in each iteration, instead of comparing via $\Pi_{\ell\text{BA}+}$ the sequences of bits $\text{B}_{\text{LEFT}} \parallel \dots \parallel \text{B}_{\text{MID}}$ of honest parties' values v , we compare sequences of blocks $\text{BLOCK}_{\text{LEFT}} \parallel \dots \parallel \text{BLOCK}_{\text{MID}}$. This change reduces the number of iterations from $O(\log \ell)$ to $O(\log n)$: after $O(\log n)$ iterations, parties agree on a bitstring PREFIX^* of i^* blocks, and each party obtains two ℓ -bit valid values v, v_\perp . The values v have prefix PREFIX^* , and, for any bitstring of $i^* + 1$ blocks, there are $t + 1$ honest parties whose values v_\perp do not have that bitstring as prefix. We present the modified subprotocol below.

FINDPREFIXBLOCKS(ℓ_{EST}, v)

Code for party P

```

1: LEFT := 1, RIGHT :=  $n + 1$ ;  $v := v_{\text{IN}}, v_{\perp} := v_{\text{IN}}, \text{PREFIX}^* := \text{empty string}$ .
2: loop
3:   If LEFT = RIGHT, set  $i^* := \text{LEFT}$  and exit the loop.
4:   ( $\text{BLOCK}_1, \text{BLOCK}_2, \dots, \text{BLOCK}_n$ ) := BLOCKS( $v$ ).
5:   Join  $\Pi_{\ell\text{BA}+}$  with input  $\text{BLOCK}_{\text{LEFT}} \parallel \dots \parallel \text{BLOCK}_{\text{MID}}$ , where  $\text{MID} := \lfloor (\text{LEFT} + \text{RIGHT})/2 \rfloor$ .
6:   If  $\Pi_{\ell\text{BA}+}$  has returned  $\perp$ , set  $v_{\perp} := v$  and  $\text{RIGHT} := \text{MID}$ .
7:   Otherwise, if  $\Pi_{\ell\text{BA}+}$  has returned  $\text{MID} - \text{LEFT} + 1$  blocks  $\text{PREFIX}_{\text{LEFT}}^* \parallel \dots \parallel \text{PREFIX}_{\text{MID}}^*$ :
8:      $\text{PREFIX}^* := \text{PREFIX}^* \parallel \text{PREFIX}_{\text{LEFT}}^* \parallel \dots \parallel \text{PREFIX}_{\text{MID}}^*$ .
9:     If  $\text{VAL}(\text{BLOCK}_1 \parallel \dots \parallel \text{BLOCK}_{\text{MID}}) < \text{VAL}(\text{PREFIX}^*)$ :  $v := \text{MIN}_{\ell_{\text{EST}}}(\text{PREFIX}^*)$ .
10:    If  $\text{VAL}(\text{BLOCK}_1 \parallel \dots \parallel \text{BLOCK}_{\text{MID}}) > \text{VAL}(\text{PREFIX}^*)$ :  $v := \text{MAX}_{\ell_{\text{EST}}}(\text{PREFIX}^*)$ .
11:    Set  $\text{LEFT} := \text{MID} + 1$ .
12: end loop
13: Return  $\text{PREFIX}^*, v, v_{\perp}$ .

```

The lemma below is the *block version* of Lemma 1. The proof is included in Appendix A.3.

Lemma 4. *Assume a BA protocol Π_{BA} , and that honest parties join FINDPREFIXBLOCKS with the same (multiple of n^2) ℓ , and with valid ℓ -bit values v . Then, the honest parties obtain the same bitstring PREFIX^* of i^* blocks, and each honest party obtains two valid ℓ -bit values v, v_{\perp} such that: (i) the ℓ -bit representations of the values v have prefix PREFIX^* ; (ii) for any bitstring BITS of $i^* + 1$ blocks, at least $t + 1$ honest parties hold values v_{\perp} such that $\text{BITS}_{\ell}(v_{\perp})$ does not have prefix BITS .*

FINDPREFIXBLOCKS achieves communication complexity $\text{BITS}_{\ell}(\text{FINDPREFIXBLOCKS}) = O(\ell \cdot n + \kappa \cdot n^2 \log^2 n) + O(\log n) \cdot \text{BITS}_{\kappa}(\Pi_{\text{BA}})$ and round complexity $\text{ROUNDS}_{\ell}(\text{FINDPREFIXBLOCKS}) = O(\log n) \cdot \text{ROUNDS}_{\kappa}(\Pi_{\text{BA}})$.

ADDLASTBIT becomes ADDLASTBLOCK: in order to decide on a final output using the values v_{\perp} , we need to append one block to PREFIX^* , so that the extended PREFIX^* is a valid values' prefix. As the honest parties values' v have PREFIX^* as a common prefix of i^* blocks, any block within the range of values $\text{VAL}(\text{BLOCK}_{i^*+1}(v))$ of the honest parties' values v suffices. Finding such a block comes down to solving CA on inputs of ℓ/n^2 bits. Since we only run this step once, and on inputs of ℓ/n^2 bits, we may use a high communication complexity approach. For instance, we may use the protocol of [47], with minor adjustments. We describe this protocol in Appendix A.4.

Theorem 3 (Theorem 4 of [47]). *There is a CA protocol HIGHCOSTCA for \mathbb{N} resilient against $t < n/3$ corruptions, with communication complexity $\text{BITS}_{\ell}(\text{HIGHCOSTCA}) = O(\ell \cdot n^3)$, and round complexity $\text{ROUNDS}_{\ell}(\text{HIGHCOSTCA}) = O(n)$.*

We present ADDLASTBLOCK below. The proof of Lemma 5 is deferred to Appendix A.3.

ADDLASTBLOCK(ℓ, v, PREFIX^*)

Code for party P

```

1: Set  $i^* := \lfloor \text{PREFIX}^* \rfloor / (\ell/n^2)$ , i.e., the number of blocks in  $\text{PREFIX}^*$ .
2: Set  $\text{BLOCK}'_{i^*+1} := \text{HIGHCOSTCA}(\text{BLOCK}_{i^*+1}(v))$ . Return  $\text{PREFIX}^* \parallel \text{BLOCK}'_{i^*+1}$ .

```

Lemma 5. *Assume that the honest parties join ADDLASTBLOCK with the same value ℓ (that is a multiple of n^2), with the same bitstring prefix PREFIX^* of $i^* < n^2$ blocks, and with valid ℓ -bit values*

v that have PREFIX^* as a prefix. Then, the honest parties agree on a bitstring of $i^* + 1$ blocks that is the prefix of a valid value's ℓ -bit representation. ADDLASTBLOCK has communication complexity $\text{BITS}_\ell(\text{ADDLASTBLOCK}) = O(\ell \cdot n)$ and round complexity $\text{ROUNDS}_\ell(\text{ADDLASTBLOCK}) = O(n)$.

Afterwards, as in FIXEDLENGTHCA , the parties obtain their output using the subprotocol GETOUTPUT presented in Section 3. We present the code of $\text{FIXEDLENGTHCABLOCKS}$ below.

$\text{FIXEDLENGTHCABLOCKS}(\ell, v)$

Code for party P

- 1: $\text{PREFIX}^*, v, v_\perp := \text{FINDPREFIXBLOCKS}(v, \ell)$. If $|\text{PREFIX}^*| := \ell$, output v and terminate.
- 2: $\text{PREFIX}^* := \text{ADDLASTBLOCK}(\text{PREFIX}^*, v, \ell)$.
- 3: Return $v := \text{GETOUTPUT}(\text{PREFIX}^*, v_\perp, \ell)$.

Similarly to the proof of Theorem 2, we can prove the following theorem by showing that the preconditions of each subprotocol are met. We have included the proof in Appendix A.3.

Theorem 4. Assume a BA protocol Π_{BA} resilient against $t < n/3$ corruptions. If the honest parties hold ℓ -bit inputs $v_{\text{IN}} \in \mathbb{N}$, where ℓ is a publicly known multiple of n^2 , $\text{FIXEDLENGTHCABLOCKS}$ is a CA protocol resilient against $t < n/3$ corruptions.

The protocol achieves communication complexity $\text{BITS}_\ell(\text{FIXEDLENGTHCABLOCKS}) = O(\ell n + \kappa \cdot n^2 \log^2 n) + O(\log n) \cdot \text{BITS}_\kappa(\Pi_{\text{BA}})$, and round complexity $\text{ROUNDS}_\ell(\text{FIXEDLENGTHCABLOCKS}) = O(n) + O(\log n) \cdot \text{ROUNDS}_\kappa(\Pi_{\text{BA}})$.

5 Final CA Protocol for \mathbb{N}

In the previous sections, we have presented two protocols achieving CA given that the honest parties hold ℓ -bit input values in \mathbb{N} and ℓ is publicly known. FIXEDLENGTHCA matches our communication complexity goal when $\ell \in \text{poly}(n)$, while $\text{FIXEDLENGTHCABLOCKS}$ does not impose an upper bound on ℓ , but instead implicitly requires that $\ell \geq n^2$. Our final protocol combines these two, and removes the assumption that ℓ is publicly known.

The parties decide which protocol to run using a bit BA protocol Π_{BA} : each party joins Π_{BA} with input 0 if $|\text{BITS}(v_{\text{IN}})| \leq n^2$, and input 1 otherwise. If Π_{BA} returns 0, the parties obtain the estimation ℓ_{EST} in $O(\log n) \cdot \text{ROUNDS}_1(\Pi_{\text{BA}})$ rounds by comparing their inputs' length with powers of two, and afterwards they run FIXEDLENGTHCA . Otherwise, if Π_{BA} returns 1, the parties obtain the estimation ℓ_{EST} by agreeing on a block size using the high-communication-cost protocol HIGHCOSTCA , and afterwards they run $\text{FIXEDLENGTHCABLOCKS}$.

Protocol $\Pi_{\mathbb{N}}$

Code for party P with input v_{in}

- 1: Join Π_{BA} with input 0 if $|\text{BITS}(v_{\text{IN}})| \leq n^2$ and otherwise with input 1.
- 2: If the bit agreed upon is 0:
- 3: If $|\text{BITS}(v_{\text{IN}})| > n^2$, set $v_{\text{IN}} := 2^{n^2} - 1$.
- 4: For $i = 0 \dots \lceil \log_2 n^2 \rceil$:
- 5: Join Π_{BA} with input 0 if $|\text{BITS}(v_{\text{IN}})| \leq 2^i$ and with input 1 otherwise. If Π_{BA} returns 0:
- 6: Let $\ell_{\text{EST}} := 2^i$. If $|\text{BITS}(v_{\text{IN}})| > \ell_{\text{EST}}$, set $v_{\text{IN}} := 2^{\ell_{\text{EST}}} - 1$.
- 7: Output $v_{\text{OUT}} := \text{FIXEDLENGTHCA}(\ell_{\text{EST}}, v)$ and terminate.
- 8: If the bit agreed upon is 1:
- 9: Set $\text{BLOCKSIZE} := \lceil |\text{BITS}(v_{\text{IN}})| / n^2 \rceil$, and $\text{BLOCKSIZE}' := \text{HIGHCOSTCA}(\text{BLOCKSIZE})$.

- 10: Set $\ell_{\text{EST}} := \text{BLOCKSIZE}' \cdot n^2$. If $|\text{BITS}(v_{\text{IN}})| \geq \ell_{\text{EST}}$, $v_{\text{IN}} := 2^{\ell_{\text{EST}}} - 1$.
 11: Output $v_{\text{OUT}} := \text{FIXEDLENGTHCABLOCKS}(\ell_{\text{EST}}, v)$ and terminate.

Theorem 5. Assume a BA protocol Π_{BA} resilient against $t < n/3$ corruptions. Then, if the honest parties hold ℓ -bit inputs $v_{\text{IN}} \in \mathbb{N}$, $\Pi_{\mathbb{N}}$ is a CA protocol resilient against $t < n/3$ corruptions, with communication complexity $\text{BITS}_{\ell}(\Pi_{\mathbb{N}}) = O(\ell n + \kappa \cdot n^2 \log^2 n) + O(\log n) \cdot \text{BITS}_{\kappa}(\Pi_{\text{BA}})$, and round complexity $\text{ROUNDS}_{\ell}(\Pi_{\mathbb{N}}) = O(n) + O(\log n) \cdot \text{ROUNDS}_{\kappa}(\Pi_{\text{BA}})$.

Proof. Let ℓ_{\min} and ℓ_{\max} denote the lengths of the lowest and the highest honest inputs respectively.

We first assume that the Π_{BA} invocation in line 1 returns 0. Then, at least one honest party has joined with input 0 due to Π_{BA} 's Validity condition and therefore $\ell_{\min} \leq n^2$. If any honest party holds an input value longer than n^2 bits, $2^{n^2} - 1$ is within the honest inputs' range. Hence, the parties join the loop in line 4 with valid values of at most n^2 bits. If Π_{BA} returns 0 in some iteration i , then at least one honest party has joined with input 0 and therefore $\ell_{\min} \leq 2^i$. Then, if an honest party holds an input value longer than i bits, $2^i - 1$ is in the honest inputs range. Note that Π_{BA} returns 0 by iteration $i = \lceil \log_2 \min(\ell_{\max}, n^2) \rceil$ the latest, since all honest parties hold values of at most $\min(\ell_{\max}, n^2)$ bits. This ensures that, the parties agree on an estimation $\ell_{\text{EST}} \leq 2 \cdot \min(\ell_{\max}, n^2) \leq 2n^2$ with $O(\log n)$ iterations. Finally, parties join FIXEDLENGTHCA with the same value $\ell_{\text{EST}} \leq n^2$ and valid ℓ_{EST} -bit values. The parties agree on a valid output v_{OUT} , which ensures that CA is achieved. The communication complexity in this case is $O(\log n) \cdot \text{BITS}_1(\Pi_{\text{BA}}) + \text{BITS}_{2 \cdot \min(\ell, n^2)}(\text{FIXEDLENGTHCA}) = O(\ell n + \kappa n^2 \log^2 n) + O(\log n) \cdot \text{BITS}_{\kappa}(\Pi_{\text{BA}})$, while the round complexity is $O(\log n) \cdot \text{ROUNDS}_1(\Pi_{\text{BA}}) + \text{ROUNDS}_{2 \cdot \min(\ell, n^2)}(\text{FIXEDLENGTHCA}) = O(\log n) \cdot \text{ROUNDS}_{\kappa}(\Pi_{\text{BA}})$.

Otherwise, if the Π_{BA} invocation in line 1 returns 1, then there is an honest party holding an input value longer than n^2 bits: $n^2 < \ell_{\max} \leq \ell$. Parties join HIGHCOSTCA with inputs $\text{BLOCKSIZE} := \lceil |\text{BITS}(v_{\text{IN}})|/n^2 \rceil$ and obtain a value $\text{BLOCKSIZE}'$ satisfying $\lceil \ell_{\min}/n^2 \rceil \leq \text{BLOCKSIZE}' \leq \lceil \ell_{\max}/n^2 \rceil$. Note that the values BLOCKSIZE can be represented via $O(\log(\ell/n^2))$ bits, hence $O(\ell/n^2)$ bits, and therefore this step has communication cost $O(\ell n)$. Therefore, $\ell_{\text{EST}} := \text{BLOCKSIZE}' \cdot n^2$ satisfies $\ell_{\min} \leq \ell_{\text{EST}} \leq \ell_{\max} + n^2 \leq 2 \cdot \ell$. Since $\ell_{\text{EST}} \geq \ell_{\min}$, if an honest party's input value is longer than ℓ_{EST} bits, $2^{\ell_{\text{EST}}} - 1$ is guaranteed to be in the honest inputs' range. It follows that the parties join $\text{FIXEDLENGTHCABLOCKS}$ with the same value $\ell_{\text{EST}} \leq 2 \cdot \ell$ (that is a multiple of n^2) and valid ℓ_{EST} -bit values. The parties agree on a valid value v_{OUT} and therefore CA is achieved. The total bit complexity is $\text{BITS}_1(\Pi_{\text{BA}}) + \text{BITS}_{O(\ell/n^2)}(\text{HIGHCOSTCA}) + \text{BITS}_{2\ell}(\text{FIXEDLENGTHCABLOCKS}) = O(\ell n + \kappa n^2 \log^2 n)$, and the round complexity is $\text{ROUNDS}_1(\Pi_{\text{BA}}) + \text{ROUNDS}_{O(\ell/n^2)}(\text{HIGHCOSTCA}) + \text{ROUNDS}(\text{FIXEDLENGTHCABLOCKS}) = O(n) + O(\log n) \cdot \text{ROUNDS}_{\kappa}(\Pi_{\text{BA}})$. \square

6 CA Protocol for \mathbb{Z}

To extend the input space to \mathbb{Z} , we assume that the parties' inputs v_{IN} are represented as $(-1)^{\text{SIGN}_{\text{IN}}} \cdot v_{\text{IN}}^{\mathbb{N}}$, where $\text{SIGN}_{\text{IN}} \in \{0, 1\}$ and $v_{\text{IN}}^{\mathbb{N}} \in \mathbb{N}$. Then, in order to cover negative numbers using $\Pi_{\mathbb{N}}$, the parties make use of the assumed BA protocol Π_{BA} to agree on their values' sign. If the sign agreed upon, denoted by SIGN_{OUT} , differs from a party P 's SIGN_{IN} , P updates $v_{\text{IN}}^{\mathbb{N}}$ to 0, since it is guaranteed to be valid. Afterwards, the parties join $\Pi_{\mathbb{N}}$ with their possibly updated inputs $v_{\text{IN}}^{\mathbb{N}}$ and agree on $v_{\text{OUT}}^{\mathbb{N}}$ such that $v_{\text{OUT}} := (-1)^{\text{SIGN}_{\text{OUT}}} \cdot v_{\text{IN}}^{\mathbb{N}}$ is valid. We present the code and the guarantees of $\Pi_{\mathbb{Z}}$ below. The formal proof is included in Appendix A.5.

Protocol $\Pi_{\mathbb{Z}}$

Code for party P with input $v_{\text{in}} = (-1)^{\text{sign}_{\text{in}}} \cdot v_{\text{in}}^{\mathbb{N}}$

- 1: Join Π_{BA} with input SIGN_{IN} and obtain output SIGN_{OUT} .
- 2: If $\text{SIGN}_{\text{OUT}} \neq \text{SIGN}_{\text{IN}}$, set $v_{\text{IN}}^{\mathbb{N}} := 0$. Join $\Pi_{\mathbb{N}}$ with input $v_{\text{IN}}^{\mathbb{N}}$ and obtain output $v_{\text{OUT}}^{\mathbb{N}}$.
- 3: Output $v_{\text{OUT}} := (-1)^{\text{SIGN}_{\text{OUT}}} \cdot v_{\text{OUT}}^{\mathbb{N}}$.

Corollary 1. *Assume a BA protocol Π_{BA} resilient against $t < n/3$ corruptions. Then, if the honest parties hold inputs $(-1)^{\text{SIGN}_{\text{IN}}} \cdot v_{\text{IN}}^{\mathbb{N}} \in \mathbb{Z}$, such that $v_{\text{IN}}^{\mathbb{N}} \in \mathbb{N}$ with $|\text{BITS}(v_{\text{IN}}^{\mathbb{N}})| \leq \ell$, $\Pi_{\mathbb{Z}}$ is a CA protocol resilient against $t < n/3$ corruptions, with communication complexity $\text{BITS}_{\ell}(\Pi_{\mathbb{Z}}) = O(\ell n + \kappa \cdot n^2 \log^2 n) + O(\log n) \cdot \text{BITS}_{\kappa}(\Pi_{\text{BA}})$, and round complexity $\text{ROUNDS}_{\ell}(\Pi_{\mathbb{Z}}) = O(n) + O(\log n) \cdot \text{ROUNDS}_{\kappa}(\Pi_{\text{BA}})$.*

We state the final corollary, where we instantiate the assumed BA protocol Π_{BA} with a deterministic BA protocol with quadratic communication (e.g. [12]).

Corollary 2. *If the honest parties hold inputs $(-1)^{\text{SIGN}_{\text{IN}}} \cdot v_{\text{IN}}^{\mathbb{N}} \in \mathbb{Z}$, such that $v_{\text{IN}}^{\mathbb{N}} \in \mathbb{N}$ with $|\text{BITS}(v_{\text{IN}}^{\mathbb{N}})| \leq \ell$, $\Pi_{\mathbb{Z}}$ is a CA protocol resilient against $t < n/3$ corruptions, with communication complexity $\text{BITS}_{\ell}(\Pi_{\mathbb{Z}}) = O(\ell n + \kappa \cdot n^2 \log^2 n)$, and round complexity $\text{ROUNDS}_{\ell}(\Pi_{\mathbb{Z}}) = O(n \log n)$.*

7 BA for Long Messages with Additional Properties

We recall that our CA protocol relies on a BA protocol $\Pi_{\ell\text{BA}+}$ with communication complexity $O(\ell n + \text{poly}(n, \kappa))$ that satisfies two additional properties: *Intrusion Tolerance* and *Bounded Pre-Agreement*, introduced in Section 3. We restate the theorem describing this protocol below.

Theorem 1. *Given a BA protocol Π_{BA} resilient against $t < n/3$ corruptions, there is a BA protocol $\Pi_{\ell\text{BA}+}$ resilient against $t < n/3$ corruptions that achieves *Intrusion Tolerance* and *Bounded Pre-Agreement*, with communication complexity $\text{BITS}_{\ell}(\Pi_{\ell\text{BA}+}) = O(\ell n + \kappa \cdot n^2 \log n) + \text{BITS}_{\kappa}(\Pi_{\text{BA}})$, and round complexity $\text{ROUNDS}_{\ell}(\Pi_{\ell\text{BA}+}) = O(1) + \text{ROUNDS}_{\kappa}(\Pi_{\text{BA}})$.*

In this section, we describe the construction behind this theorem. The main technical challenge lies in building a communication-efficient BA protocol *for short messages* (κ bits) that achieves the two additional properties, denoted by $\Pi_{\ell\text{BA}+}$. Afterwards, $\Pi_{\ell\text{BA}+}$ is constructed using the outline of prior works [8, 41].

Protocol $\Pi_{\ell\text{BA}+}$. In our protocol $\Pi_{\ell\text{BA}+}$, the parties first distribute their input values. Each party P receives $n - t$ values v from honest parties, plus at most t values from corrupted parties, and checks whether there is any value received from $n - 2t$ parties. Note that P sees at most two values v_1, v_2 with this property. Moreover, if there is some value v held as input by $n - 2t$ honest parties, all honest parties observe this value.

The parties then *vote* for the values they have seen $n - 2t$ times by sending either $\text{VOTE}(\cdot)$, $\text{VOTE}(v_1)$, or $\text{VOTE}(v_1, v_2)$, depending on how many such values they have seen. Each party P looks at the values that have received $n - t$ votes: there will be at most two such values. If there are two, P denotes them by a and b such that $a \leq b$. If there is only one such value v , P sets $a := v$ and $b := v$. If there are none, P simply sets $a := \perp$ and $b := \perp$.

The key observation will be that, if there is a value v held as input by $n - 2t$ honest parties, then the honest parties hold the same $a = v$ or the same $b = v$. The parties then first try to agree on a : they run a BA protocol Π_{BA} on their values a and obtain an output a' . Afterwards they

check whether they are happy with a' by joining Π_{BA} once again: with input 1 if $a = a'$ and 0 otherwise. If Π_{BA} returns 1, the parties output a' . If Π_{BA} returns 0, the parties check whether they hold the same value b with the same strategy: they join Π_{BA} with input b , and obtain output b' . Afterwards, they join Π_{BA} again with input 1 if $b = b'$ and 0 otherwise. If the output is 1, they output b' , and otherwise they output \perp .

This way, if $n - 2t$ honest parties hold the same input, the output is guaranteed to be non- \perp , which ensures that Bounded Pre-Agreement holds. In addition, if the parties output a non- \perp value, we are able to show that some honest party has received it from $n - 2t > t$ parties in the first step, and therefore it is an honest input, which ensures that Intrusion Tolerance holds. We present the code below.

Protocol $\Pi_{\text{BA}+}$

Code for party P with input v_{in}

- 1: Send v_{in} to all parties.
- 2: Check if there is any value received from $n - 2t$ parties. If there is none, send $\text{VOTE}(\cdot)$ to all parties. If there is only one, let this value be v_1 and send $\text{VOTE}(v_1)$ to all parties. If there are two, let these values be v_1 and v_2 and send $\text{VOTE}(v_1, v_2)$ to all parties.
- 3: Let $a := \perp, b := \perp$. If there is a single value v voted by $n - t$ parties, set $a := v$ and $b := v$. If there are two values $v \leq v'$ voted by $n - t$ parties, set $a := v$ and $b := v'$.
- 4: Join Π_{BA} with input a and obtain output a' . If $a' = a \neq \perp$, join Π_{BA} with input 1, and otherwise with input 0. If Π_{BA} returns 1, output a and terminate.
- 5: Join Π_{BA} with input b and obtain output b' . If $b' = b \neq \perp$, join Π_{BA} with input 1, and otherwise with input 0. If Π_{BA} returns 1, output b and terminate. Otherwise, output \perp and terminate.

Theorem 6. *Assume a BA protocol Π_{BA} resilient against $t < n/3$ corruptions. Then, there is a BA protocol $\Pi_{\text{BA}+}$ resilient against $t < n/3$ that additionally achieves Intrusion Tolerance and Bounded Pre-Agreement. $\Pi_{\text{BA}+}$ has communication complexity $\text{BITS}_{\kappa}(\Pi_{\text{BA}+}) = O(\kappa n^2) + \text{BITS}_{\kappa}(\Pi_{\text{BA}})$ and round complexity $\text{ROUNDS}(\Pi_{\text{BA}+}) = O(1) \cdot \text{ROUNDS}_{\kappa}(\Pi_{\text{BA}})$.*

Proof. We first show that $\Pi_{\text{BA}+}$ is indeed a BA protocol. Agreement and Termination follow from the fact that Π_{BA} satisfies these properties. If all honest parties hold the same input value v , then no honest party receives $v' \neq v$ from $n - 2t > t$ parties, and therefore all honest parties send (VOTE, v) . Then, the honest parties see at most t votes for any value $v' \neq v$, and therefore all honest parties set $a = b = v$. They join the Π_{BA} invocation in line 4 with input $a = v$, and since Π_{BA} achieves Validity, they agree on $a' = v$. All honest parties join the next Π_{BA} invocation with input 1, and therefore they agree on 1 and output v , hence Validity holds.

We now focus on Intrusion Tolerance. If the honest parties output some value $v \neq \perp$, then this is the value a or b obtained by an honest party P (due to Π_{BA} 's Validity). P has received $n - t$ votes for v , hence at least one vote from some honest party P' . Then, P' has received v from $n - 2t > t$ parties, hence from at least one honest party, and therefore Intrusion Tolerance holds.

For Bounded Pre-Agreement, we show that, if there is a value v held as input by $n - 2t$ honest parties, then the value agreed upon is not \perp .

We establish that, in line 2, every party sees at most two input values from $n - 2t$ parties each. Assuming that a party has received at least three input values sent by $n - 2t$ parties each, we obtain that $3 \cdot (n - 2t) \leq n$, which contradicts $n > 3t$. Then, each party sees at most two values from $n - 2t$ parties each, and one of these is v . Hence, each honest party sends a vote for v , and possibly for a second value.

Then, each party receives the $n - t$ votes for v from the honest parties. We add that each party receives $n - t$ votes for two values at most. Otherwise, since each party votes for at most

two values, there are at most $2n$ votes in total. Assuming that a party receives $n - t$ votes for at least three values implies $3 \cdot (n - t) \leq 2n$, contradicting $n > 3t$.

If every honest party sees v as the only value with $n - t$ votes, then every honest party sets $a = b = v$, and therefore all honest parties output v in line 4. Otherwise, let P and P' be two honest parties. We assume that P sees two values with $n - t$ votes each: one of these we know to be v , and the other is $v' \neq v$. We have showed that P' also sees v with $n - t$ (honest) votes, and we assume that P' receives $n - t$ votes for $v'' \notin \{v, v'\}$. This leads to a contradiction as P' has received $n - t$ honest votes for v , and at least $n - 2t$ honest votes for v' . Since every party may vote for at most two different values, there are at most t honest votes and t votes from byzantine parties left for v'' : these are $2t < n - t$ in total.

Hence, every honest party obtains a, b such that $v \in \{a, b\} \subseteq \{v, v'\}$. The parties then try to agree on a in line 4. If the second Π_{BA} invocation of line 4 returns 0, then the parties hold different values a : $a = v$ for some honest parties, and $a = v'$ for the others. Since every honest party has set a and b such that $a \leq b$, this means that the honest parties hold the same value $b = v$ and output v in line 5.

For the round complexity, note that $\Pi_{\text{BA}+}$ incurs two rounds of communication and afterwards runs Π_{BA} at most four times. For the communication complexity, note that each party sends at most three values to all parties, and each of these values is an honest party's input, and therefore consists of κ bits. Afterwards, they run Π_{BA} on κ -bit inputs at most twice and on bits at most twice. \square

From $\Pi_{\text{BA}+}$ to $\Pi_{\ell\text{BA}+}$. We may now describe our protocol $\Pi_{\ell\text{BA}+}$ for long messages, relying on $\Pi_{\text{BA}+}$ and on the outline of prior works [8, 41].

$\Pi_{\ell\text{BA}+}$ makes use of Reed-Solomon (RS) codes [45], which allow each party to split its value into n codewords so that reconstructing the original value only requires $n - t$ of these n codewords. To enable the parties to detect corrupted codewords, and also to compress values, prior works make use of collision-free cryptographic accumulators [42]. Essentially, accumulators convert a set (in our case, the n codewords) into a κ -bit value and provide witnesses confirming the accumulated set's contents. For this task, we use Merkle Trees (MT) [39], which do not require a trusted dealer. We briefly describe RS codes and MT below.

$\Pi_{\ell\text{BA}+}$ assumes standard RS codes with parameters $(n, n - t)$. This provides us with a deterministic algorithm $\text{RS.ENCODE}(v)$, which takes a value v as input and converts it into n codewords (s_1, \dots, s_n) of $O(|\text{BITS}(v)|/n)$ bits each. The codewords s_i are elements of a Galois Field $\mathbb{F} = GF(2^a)$ with $n \leq 2^a - 1$. To reconstruct the original value, RS codes provide a decoding algorithm, RS.DECODE , which takes as input $n - t$ of the n codewords and returns the original value v . Any $n - t$ of the n codewords uniquely determine the original value v .

An MT is a balanced binary tree that enables us to compress a multiset of values into a κ -bit encoding, and to efficiently verify (with high probability) that some value belongs to the compressed multiset. Given a multiset $S = \{s_1, \dots, s_n\}$, the MT is built bottom-up, using the collision-resistant hash function H_κ : starting with n leaves, where the i -th leaf stores $H_\kappa(s_i)$. Each non-leaf node stores $H_\kappa(h_{\text{LEFT}} \parallel h_{\text{RIGHT}})$, where h_{LEFT} and h_{RIGHT} are the hashes stored by the node's left and resp. right child. This way, the hash stored by the root represents the encoding of S . Given the root's hash z , one can prove that s_i belongs to the compressed multiset using a witness w_i of $O(\kappa \cdot \log n)$ bits. The witness w_i contains the hashes needed to verify the path from the i -th leaf to the root. Note that the collision-resistance assumption leads to different encodings for different multisets, and prevents the adversary from producing witnesses for values of its own choice. We will use $\text{MT.BUILD}(S)$ to denote the (deterministic) algorithm that creates the MT for the given multiset S and returns the hash stored by the root z and the witnesses w_1, w_2, \dots, w_n .

Afterwards, $\text{MT.VERIFY}(z, i, s_i, w_i)$ returns **true** if w_i proves that $H_\kappa(s_i)$ is indeed stored on the i -th leaf of the MT with root hash z and **false** otherwise.

Then, $\Pi_{\ell\text{BA}+}$ consists of three steps. In the first step, every party computes $s_1, \dots, s_n := \text{RS.ENCODE}(v_{\text{IN}})$ and $z, w_1, \dots, w_n := \text{MT.BUILD}(\{s_1, \dots, s_n\})$. In the second step, the parties agree on an encoding z^* with the help of $\Pi_{\text{BA}+}$. In the third step, the parties obtain the final output. If $\Pi_{\text{BA}+}$ returns \perp , the parties output \perp . Otherwise, if $\Pi_{\text{BA}+}$ returns z^* , every party P^* holding $z = z^*$ distributes $v^* := v_{\text{IN}}$ to all the parties. To achieve this using only $O(\ell n + \text{poly}(n, \kappa))$ bits, P^* sends s_i and its MT witness w_i to each party P_i . The MT witnesses allow the parties to detect and discard any corrupted codewords. In addition, RS codes are deterministic, so each party P_i obtains a unique codeword s_i from $\text{RS.ENCODE}(v^*)$. Every party P_i then sends (s_i, w_i) to all parties, which allows the parties to reconstruct v^* .

Protocol $\Pi_{\ell\text{BA}+}$

Code for party P_i with input v_{in}

- 1: Let $s_1, s_2, \dots, s_n := \text{RS.ENCODE}(v_{\text{IN}})$; $z, w_1, w_2, \dots, w_n := \text{MT.BUILD}(\{s_1, s_2, \dots, s_n\})$.
- 2: Join $\Pi_{\text{BA}+}$ with input z . If $\Pi_{\text{BA}+}$ has returned \perp , output \perp . Otherwise, if $\Pi_{\text{BA}+}$ has returned $z^* \neq \perp$, run the **distributing step**:
- 3: If $z^* = z$: for every $1 \leq j \leq n$, send (j, s_j, w_j) to P_j .
- 4: If you have received a tuple (i, s_i, w_i) such that $\text{MT.VERIFY}(i, z^*, s_i, w_i) = \text{true}$:
- 5: Send (i, s_i, w_i) to all parties.
- 6: Discard any tuples (j, s_j, w_j) where $\text{MT.VERIFY}(i, z^*, s_i, w_i) = \text{false}$.
- 7: Let $S :=$ the set of correct tuples received. Output $v^* := \text{RS.DECODE}(S)$.

We may now sketch the proof of Theorem 1. The formal proof is included in Section A.1.

Proof sketch of Theorem 1. As $\Pi_{\text{BA}+}$ achieves Termination, $\Pi_{\ell\text{BA}+}$ achieves Termination as well. Then, note that $\Pi_{\ell\text{BA}+}$ returns \perp whenever $\Pi_{\text{BA}+}$ returns \perp , and the parties obtain non- \perp in $\Pi_{\ell\text{BA}+}$ whenever $\Pi_{\text{BA}+}$ returns a non-bot value. Moreover, the Intrusion Tolerance property of $\Pi_{\text{BA}+}$ ensures that, whenever $\Pi_{\text{BA}+}$ returns non- \perp , the parties agree on the encoding z^* of an honest party's input, which means that the parties successfully decode a value that is an honest party's input. Hence, both Agreement and Intrusion Tolerance hold. For Bounded Pre-Agreement, since $\Pi_{\text{BA}+}$ only returns \perp when fewer $n - 2t$ parties hold the same input value, $\Pi_{\ell\text{BA}+}$ also only returns \perp when fewer $n - 2t$ parties hold the same input value. Finally, for Validity, if all honest parties hold the same input value v^* , $\Pi_{\text{BA}+}$'s Validity ensures that the parties agree on the encoding z^* of v^* , and therefore the parties successfully decode z^* .

The round complexity follows immediately from the round complexity of $\Pi_{\text{BA}+}$. For the communication complexity, note that, in Step 3, each party sends at most two shares and two MT witnesses to each party. This leads $O(\ell n + \kappa \cdot n^2 \log n) + \text{BITS}_\kappa(\Pi_{\text{BA}+})$ bits of communication. \square

8 Conclusions

Achieving solutions with optimal-communication has been the subject of an extensive line of works [8, 23, 24, 34, 41]. These works have primarily focused on the fundamental primitives BA and BC, where $\Omega(\ell n)$ bits of communication are necessary, and have presented protocols with communication complexity $O(\ell n + \text{poly}(n, \kappa))$, proving the lower bound tight for large enough ℓ . Our work shows that the lower bound $\Omega(\ell n)$ is also tight for synchronous CA on integers given that ℓ is large enough, namely $\ell = \Omega(\kappa \cdot n \log^2 n)$. We have presented a protocol that relies on

finding some valid values' longest common prefix, achieving CA with optimal resilience, asymptotically optimal communication complexity, and efficient round complexity. Our protocol is also deterministic and operates without trusted setup.

We leave a number of exciting open problems. While we expect that our techniques can be easily extended to the asynchronous setting for a lower number of corruptions $t < n/5$, it would be interesting to see whether achieving asymptotically optimal communication complexity for $t < n/3$ corruptions in the asynchronous model is possible. The same question applies to the synchronous model with $t < n/2$ corruptions assuming cryptographic setup. A different direction could investigate whether the round complexity can be reduced from $O(n \log n)$ to the optimal $O(n)$ while maintaining the communication complexity. Further works could also consider reducing the $\text{poly}(n, \kappa)$ factor, or extending our question to input spaces beyond \mathbb{Z} .

References

- [1] Ittai Abraham, Yonatan Amit, and Danny Dolev. Optimal resilience asynchronous approximate agreement. In Teruo Higashino, editor, *Principles of Distributed Systems*, pages 229–239, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- [2] Ittai Abraham, TH Hubert Chan, Danny Dolev, Kartik Nayak, Rafael Pass, Ling Ren, and Elaine Shi. Communication complexity of byzantine agreement, revisited. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 317–326, 2019.
- [3] Dan Alistarh, Faith Ellen, and Joel Rybicki. Wait-free approximate agreement on graphs. In Tomasz Jurdziński and Stefan Schmid, editors, *Structural Information and Communication Complexity*, pages 87–105, Cham, 2021. Springer International Publishing. doi:10.1007/978-3-030-79527-6_6.
- [4] Dan Alistarh, Jerry Ho, Sashank Kannan, Aditya Sarwate, and Leonidas Zhang. Collaborative learning in the jungle (decentralized, byzantine, heterogeneous, asynchronous and nonconvex learning). In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 11990–12001, 2021.
- [5] Akhil Bandarupalli, Adithya Bhat, Saurabh Bagchi, Aniket Kate, Chen-Da Liu-Zhang, and Michael K. Reiter. Delphi: Efficient asynchronous approximate agreement for distributed oracles. In *54th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2024.
- [6] Michael Ben-Or, Danny Dolev, and Ezra N. Hoch. Brief announcement: Simple gradecast based algorithms. In Nancy A. Lynch and Alexander A. Shvartsman, editors, *Distributed Computing*, pages 194–197, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- [7] Piotr Berman, Juan A. Garay, and Kenneth J. Perry. Towards optimal distributed consensus (extended abstract). In *Proceedings of the 30th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 410–415. IEEE Computer Society, 1989.
- [8] Amey Bhangale, Chen-Da Liu-Zhang, Julian Loss, and Kartik Nayak. Efficient adaptively-secure byzantine agreement for long messages. In *Advances in Cryptology–ASIACRYPT 2022: 28th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, December 5–9, 2022, Proceedings, Part I*, pages 504–525, Berlin, Heidelberg, 2023. Springer, Springer-Verlag.
- [9] Erica Blum, Jonathan Katz, Chen-Da Liu-Zhang, and Julian Loss. Asynchronous byzantine agreement with subquadratic communication. In *TCC 2020, Part I*, LNCS, pages 353–380. Springer, Heidelberg, March 2020. doi:10.1007/978-3-030-64375-1_13.
- [10] Christian Cachin and Stefano Tessaro. Asynchronous verifiable information dispersal. In *24th IEEE Symposium on Reliable Distributed Systems (SRDS’05)*, pages 191–201, Orlando, FL, USA, 2005. IEEE, IEEE Computer Society. doi:10.1109/RELDIS.2005.9.
- [11] Wutichai Chongchitmate and Rafail Ostrovsky. Information-theoretic broadcast with dishonest majority for long messages. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part I*, volume 11239 of LNCS, pages 370–388, Cham, November 2018. Springer, Heidelberg. doi:10.1007/978-3-030-03807-6_14.

- [12] Brian A Coan and Jennifer L Welch. Modular construction of a byzantine agreement protocol with optimal message bit complexity. *Information and Computation*, 97(1):61–85, 1992.
- [13] Shir Cohen, Idit Keidar, and Alexander Spiegelman. Not a coincidence: Sub-quadratic asynchronous byzantine agreement whp. In *34th International Symposium on Distributed Computing*, 2020.
- [14] Andrei Constantinescu, Diana Ghinea, Lioba Heimbach, Zilin Wang, and Roger Wattenhofer. A Fair and Resilient Decentralized Clock Network for Transaction Ordering. In Alysson Bessani, Xavier Défago, Junya Nakamura, Koichi Wada, and Yukiko Yamauchi, editors, *27th International Conference on Principles of Distributed Systems (OPODIS 2023)*, volume 286 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 8:1–8:20, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. URL: <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.OPODIS.2023.8>, doi: 10.4230/LIPIcs.OPODIS.2023.8.
- [15] Andrei Constantinescu, Diana Ghinea, Roger Wattenhofer, and Floris Westermann. Convex Consensus with Asynchronous Fallback. In *38th International Symposium on Distributed Computing (DISC)*, Madrid, Spain, October 2024.
- [16] Danny Dolev, Nancy A. Lynch, Shlomit S. Pinter, Eugene W. Stark, and William E. Weihl. Reaching approximate agreement in the presence of faults. *J. ACM*, 33(3):499–516, May 1986. doi:10.1145/5925.5931.
- [17] Danny Dolev and Rüdiger Reischuk. Bounds on information exchange for byzantine agreement. In Robert L. Probert, Michael J. Fischer, and Nicola Santoro, editors, *1st ACM PODC*, pages 132–140. ACM, August 1982. doi:10.1145/800220.806690.
- [18] El-Mahdi El-Mhamdi, Rachid Guerraoui, Arsany Guirguis, Lê Nguyễn Hoàng, and Sébastien Rouault. Genuinely distributed byzantine machine learning. In *Proceedings of the 39th Symposium on Principles of Distributed Computing*, pages 355–364, 2020.
- [19] El-Mahdi El-Mhamdi, Rachid Guerraoui, Arsany Guirguis, Lê Nguyễn Hoàng, and Sébastien Rouault. Genuinely distributed byzantine machine learning. In *Proceedings of the 39th Symposium on Principles of Distributed Computing*, PODC ’20, page 355–364, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3382734.3405695.
- [20] A. D. Fekete. Asynchronous approximate agreement. In *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing*, PODC ’87, page 64–76, New York, NY, USA, 1987. Association for Computing Machinery. doi:10.1145/41840.41846.
- [21] Alan David Fekete. Asymptotically optimal algorithms for approximate agreement. *Distributed Computing*, 4(1):9–29, 1990.
- [22] Michael J Fischer, Nancy A Lynch, and Michael S Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM (JACM)*, 32(2):374–382, 1985.
- [23] Matthias Fitzi and Martin Hirt. Optimally efficient multi-valued Byzantine agreement. In Eric Ruppert and Dahlia Malkhi, editors, *25th ACM PODC*, pages 163–168, New York, NY, USA, July 2006. ACM. doi:10.1145/1146381.1146407.

- [24] Chaya Ganesh and Arpita Patra. Broadcast extensions with optimal communication and round complexity. In George Giakkoupis, editor, *35th ACM PODC*, pages 371–380, New York, NY, USA, July 2016. ACM. doi:10.1145/2933057.2933082.
- [25] Yuval Gelles and Ilan Komargodski. Optimal load-balanced scalable distributed agreement. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, page 411–422, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3618260.3649736.
- [26] Diana Ghinea, Chen-Da Liu-Zhang, and Roger Wattenhofer. Optimal synchronous approximate agreement with asynchronous fallback. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, PODC’22, page 70–80, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3519270.3538442.
- [27] Diana Ghinea, Chen-Da Liu-Zhang, and Roger Wattenhofer. Multidimensional approximate agreement with asynchronous fallback. In *Proceedings of the 35th ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA ’23, page 141–151, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3558481.3591105.
- [28] Martin Hirt and Pavel Raykov. Multi-valued byzantine broadcast: The $t < n$ case. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 448–465, Berlin, Heidelberg, December 2014. Springer, Heidelberg. doi:10.1007/978-3-662-45608-8_24.
- [29] Valerie King and Jared Saia. From almost everywhere to everywhere: Byzantine agreement with $\tilde{O}(n^{\frac{3}{2}})$ bits. In Idit Keidar, editor, *Distributed Computing*, pages 464–478, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.
- [30] Valerie King and Jared Saia. Breaking the $O(n^2)$ bit barrier: scalable byzantine agreement with an adaptive adversary. In Andréa W. Richa and Rachid Guerraoui, editors, *29th ACM PODC*, pages 420–429. ACM, July 2010. doi:10.1145/1835698.1835798.
- [31] Valerie King, Jared Saia, Vishal Sanwalani, and Erik Vee. Scalable leader election. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm*, SODA ’06, page 990–999, USA, 2006. Society for Industrial and Applied Mathematics.
- [32] Jérémy Ledent. Brief announcement: Variants of approximate agreement on graphs and simplicial complexes. In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, PODC’21, page 427–430, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3465084.3467946.
- [33] Christoph Lenzen and Julian Loss. Optimal clock synchronization with signatures. In *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, PODC’22, page 440–449, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3519270.3538444.
- [34] Guanfeng Liang and Nitin Vaidya. Error-free multi-valued consensus with byzantine failures. In *Proceedings of the 30th annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, pages 11–20, New York, NY, USA, 2011. Association for Computing Machinery.

- [35] Guanfeng Liang and Nitin H. Vaidya. Error-free multi-valued consensus with byzantine failures. In Cyril Gavoille and Pierre Fraigniaud, editors, *30th ACM PODC*, pages 11–20. ACM, June 2011. doi:10.1145/1993806.1993809.
- [36] Darya Melnyk and Roger Wattenhofer. Byzantine agreement with interval validity. In *2018 IEEE 37th Symposium on Reliable Distributed Systems (SRDS)*, pages 251–260, Salvador, Brazil, 2018. IEEE Computer Society. doi:10.1109/SRDS.2018.00036.
- [37] Hammurabi Mendes and Maurice Herlihy. Multidimensional approximate agreement in byzantine asynchronous systems. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 391–400, Palo Alto, CA, USA, June 2013. ACM Press. doi:10.1145/2488608.2488657.
- [38] Hammurabi Mendes, Maurice Herlihy, Nitin Vaidya, and Vijay K Garg. Multidimensional agreement in byzantine systems. *Distributed Computing*, 28(6):423–441, 2015.
- [39] Ralph C Merkle. A digital signature based on a conventional encryption function. In *Conference on the theory and application of cryptographic techniques*, pages 369–378. Springer, 1987.
- [40] Atsuki Momose and Ling Ren. Optimal communication complexity of authenticated byzantine agreement. In *35th International Symposium on Distributed Computing (DISC 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.
- [41] Kartik Nayak, Ling Ren, Elaine Shi, Nitin H. Vaidya, and Zhuolun Xiang. Improved Extension Protocols for Byzantine Broadcast and Agreement. In Hagit Attiya, editor, *34th International Symposium on Distributed Computing (DISC 2020)*, volume 179 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 28:1–28:17, Dagstuhl, Germany, 2020. Schloss Dagstuhl-Leibniz-Zentrum für Informatik. URL: <https://drops.dagstuhl.de/opus/volltexte/2020/13106>, doi:10.4230/LIPIcs.DISC.2020.28.
- [42] Lan Nguyen. Accumulators from bilinear pairings and applications. In *Topics in Cryptology—CT-RSA 2005: The Cryptographers’ Track at the RSA Conference 2005, San Francisco, CA, USA, February 14–18, 2005. Proceedings*, pages 275–292. Springer, 2005.
- [43] Thomas Nowak and Joel Rybicki. Byzantine Approximate Agreement on Graphs. In Jukka Suomela, editor, *33rd International Symposium on Distributed Computing (DISC 2019)*, volume 146 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 29:1–29:17, Dagstuhl, Germany, 2019. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2019/11336>, doi:10.4230/LIPIcs.DISC.2019.29.
- [44] Maria Potop-Butucaru, Michel Raynal, and Sebastien Tixeuil. Distributed computing with mobile robots: An introductory survey. pages 318 – 324, 10 2011. doi:10.1109/NBiS.2011.55.
- [45] 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.
- [46] Phillip Rogaway. Formalizing human ignorance. In Phong Q. Nguyen, editor, *Progress in Cryptology - VIETCRYPT 2006*, pages 211–228, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

- [47] David Stolz and Roger Wattenhofer. Byzantine Agreement with Median Validity. In Emmanuelle Anceaume, Christian Cachin, and Maria Potop-Butucaru, editors, *19th International Conference on Principles of Distributed Systems (OPODIS 2015)*, volume 46 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 1–14, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2016/6591>, doi:10.4230/LIPIcs.OPODIS.2015.22.
- [48] Lili Su and Nitin H Vaidya. Fault-tolerant multi-agent optimization: optimal iterative distributed algorithms. In *Proceedings of the 2016 ACM symposium on principles of distributed computing*, pages 425–434, 2016.
- [49] Russell Turpin and Brian A Coan. Extending binary byzantine agreement to multivalued byzantine agreement. *Information Processing Letters*, 18(2):73–76, 1984.
- [50] Nitin H. Vaidya and Vijay K. Garg. Byzantine vector consensus in complete graphs. In Panagiota Fatourou and Gadi Taubenfeld, editors, *32nd ACM PODC*, pages 65–73, Montreal, QC, July 2013. ACM. doi:10.1145/2484239.2484256.

A Appendix

A.1 BA for long messages with additional properties

This section provides the analysis of our protocol $\Pi_{\ell\text{BA}+}$. We formally prove Theorem 1, restated below.

Theorem 1. *Given a BA protocol Π_{BA} resilient against $t < n/3$ corruptions, there is a BA protocol $\Pi_{\ell\text{BA}+}$ resilient against $t < n/3$ corruptions that achieves Intrusion Tolerance and Bounded Pre-Agreement, with communication complexity $\text{BITS}_{\ell}(\Pi_{\ell\text{BA}+}) = O(\ell n + \kappa \cdot n^2 \log n) + \text{BITS}_{\kappa}(\Pi_{\text{BA}})$, and round complexity $\text{ROUNDS}_{\ell}(\Pi_{\ell\text{BA}+}) = O(1) + \text{ROUNDS}_{\kappa}(\Pi_{\text{BA}})$.*

The lemma below ensures that, if the $\Pi_{\text{BA}+}$ invocation in line 2 of $\Pi_{\ell\text{BA}+}$ returns a non- \perp value, the honest parties agree on an honest input. The argument is identical to [41, Lemma 6].

Lemma 6. *Assume that the parties join the distributing step and that at least one honest party has proposed $z = z^*$. Then, the honest parties agree on a value v^* that is an honest party's input. In addition, this step has communication complexity $O(\ell n + \kappa \cdot n^2 \log n)$ and round complexity $O(1)$.*

Proof. Since at least one honest party P_i holds $z := z^*$, P_i holds an input value v^* whose RS encoding s_1, \dots, s_n leads to an MT tree with root z^* . P_i sends to each party P_j a tuple (j, s_j, w_j) such that $\text{MT.VERIFY}(z^*, j, s_j, w_j) = \text{true}$.

Note that party P_j ignores any tuples (j, s'_j, w'_j) with $s'_j \neq s_j$: a different RS encoding $(s'_1, \dots, s'_n) \neq (s_1, \dots, s_n)$ leads to an MT with root $z \neq z^*$. Hence, such a tuple is sent by a corrupted party. We note that finding a witness w'_j with $\text{MT.VERIFY}(z^*, j, s'_j, w'_j) = \text{true}$ requires the adversary to find collisions for H_{κ} , which we assumed to be impossible. Therefore, $\text{MT.VERIFY}(z^*, j, s'_j, w'_j) = \text{false}$, and P_j discards this tuple.

Then, every party P_i holds a unique correct tuple (i, s_i, w_i) (possibly received from multiple parties), and forwards this tuple to all parties. Each party P_i receives $n - t$ correct tuples from honest parties, plus at most t tuples from corrupted parties. Once again, if an honest party P_j receives (j, s'_j, w'_j) with an incorrect codeword s'_j , P_j discards this tuple: $\text{MT.VERIFY}(z^*, j, s'_j, w'_j) = \text{false}$. Hence, all (at least $n - t$) tuples remaining are correct, which allows the parties to reconstruct v^* correctly. Therefore, the parties agree on an honest party's input value.

It remains to discuss the communication complexity and the round complexity. There are two communication rounds, where every party sends to all parties at most two tuples. Each such tuple contains an index of $O(\log n)$ bits, a RS codeword of $O(\ell/n)$ bits, and a MT witness of $O(\kappa \cdot \log n)$ bits. Therefore, this step has a total communication cost of $O(\ell n + \kappa \cdot n^2 \log n)$ bits. \square

Below we provide the analysis of $\Pi_{\ell\text{BA}+}$. This is also similar to the analysis of [41], with the exception that we also verify the additional properties Intrusion Tolerance and Bounded Pre-Agreement. Theorem 6 and the lemma presented below directly imply Theorem 1.

Lemma 7. *Assume a BA protocol $\Pi_{\text{BA}+}$ secure against $t < n/3$ corruptions that additionally achieves Intrusion Tolerance and Bounded Pre-Agreement. Then, $\Pi_{\ell\text{BA}+}$ achieves the same guarantees, with communication complexity $\text{BITS}_{\ell}(\Pi_{\ell\text{BA}+}) = O(\ell n + \kappa \cdot n^2 \log n) + \text{BITS}_{\kappa}(\Pi_{\text{BA}+})$, and round complexity $\text{ROUNDS}_{\ell}(\Pi_{\ell\text{BA}+}) = O(1) + \text{ROUNDS}_{\kappa}(\Pi_{\text{BA}+})$.*

Proof. In the following, we first show that $\Pi_{\ell\text{BA}+}$ achieves the standard BA properties.

The parties obtain the same output in $\Pi_{\text{BA}+}$: either z^* or \perp . If the output returned by $\Pi_{\text{BA}+}$ is \perp , the parties output \perp , hence Agreement holds in this case. Otherwise, there is an honest

party who proposed $z = z^*$ since $\Pi_{\text{BA}+}$ achieves Intrusion Tolerance, and Lemma 6 ensures that the parties output the same value.

Honest parties holding the same value v_{IN} obtain the same encoding z since the algorithms for computing the RS encoding and the MT are deterministic. This implies that, if all honest parties hold the same input v , then all honest parties obtain the same value z , and $\Pi_{\text{BA}+}$ returns $z^* = z$. Lemma 6 ensures that the parties output an honest party's input, therefore they output v . Therefore, Validity also holds and BA is achieved.

If the honest parties obtain a non- \perp output, they have obtained this value via the distributing step. Since this step is only executed if there is an honest party holding $z = z^*$, Lemma 6 ensures that the Intrusion Tolerance property holds.

If there are $n - 2t$ honest parties holding the same input value v , then these parties join $\Pi_{\text{BA}+}$ with the same encoding z . The Bounded Pre-Agreement property of $\Pi_{\text{BA}+}$ ensures that the parties agree on $z^* \neq z$. Afterwards, Lemma 6 ensures that the honest parties agree on a non- \perp value in $\Pi_{\ell\text{BA}+}$, and therefore Bounded Pre-Agreement holds.

We have obtained that $\Pi_{\ell\text{BA}+}$ indeed maintains the properties of $\Pi_{\text{BA}+}$. Running $\Pi_{\text{BA}+}$ with inputs z requires $\text{BITS}_{\kappa}(\Pi_{\text{BA}+})$ bits and $\text{ROUNDS}_{\kappa}(\Pi_{\text{BA}+})$ rounds. If the output is \perp , there is no further communication. Otherwise, the parties run the distributing step, and Lemma 6 shows that this step has an additional cost of $O(\ell n + \kappa \cdot n^2 \log n)$ bits and $O(1)$ rounds. Then, the total bit complexity of $\Pi_{\ell\text{BA}+}$ is $O(\ell n + \kappa \cdot n^2 \log n) + \text{BITS}_{\kappa}(\Pi_{\text{BA}+})$, and the round complexity is $O(1) + \text{ROUNDS}_{\kappa}(\Pi_{\text{BA}+})$. \square

A.2 Missing proofs for FixedLengthCA

In this section, we focus on analyzing each of the subprotocols of FIXEDLENGTHCA. We first include two remarks which will enable us to show that values computed by comparing prefixes of valid values are valid.

Remark 1. Consider two values $v, v' \in \mathbb{N}$ satisfying $v \leq v' < 2^{\ell}$, and let COMMON_PREFIX denote the **longest** common prefix of $\text{BITS}_{\ell}(v)$ and $\text{BITS}_{\ell}(v')$. If $|\text{COMMON_PREFIX}| < \ell$, then $\text{MAX}_{\ell}(\text{COMMON_PREFIX} \parallel 0), \text{MIN}_{\ell}(\text{COMMON_PREFIX} \parallel 1) \in [v, v']$.

Proof. We show that $v \leq \text{MAX}_{\ell}(\text{COMMON_PREFIX} \parallel 0) \leq \text{MIN}_{\ell}(\text{COMMON_PREFIX} \parallel 1) \leq v'$.

We first note that, since $v \leq v'$, $\text{BITS}_{\ell}(v)$ has prefix $\text{COMMON_PREFIX} \parallel 0$, while $\text{BITS}_{\ell}(v')$ has prefix $\text{COMMON_PREFIX} \parallel 1$. Secondly, since $\text{MAX}_{\ell}(\text{COMMON_PREFIX} \parallel 0)$ is the highest ℓ -bit value having prefix $\text{COMMON_PREFIX} \parallel 0$, and v is an ℓ -bit value with the same prefix, $v \leq \text{MAX}_{\ell}(\text{COMMON_PREFIX} \parallel 0)$.

In addition, note that $\text{MAX}_{\ell}(\text{COMMON_PREFIX} \parallel 0) + 1 = \text{MIN}_{\ell}(\text{COMMON_PREFIX} \parallel 1)$.

We use a similar argument to show that $v' \geq \text{MIN}_{\ell}(\text{COMMON_PREFIX} \parallel 1)$: v' is an ℓ -bit value with prefix $\text{COMMON_PREFIX} \parallel 1$, while $\text{MIN}_{\ell}(\text{COMMON_PREFIX} \parallel 1)$ is the lowest ℓ -bit value having prefix $\text{COMMON_PREFIX} \parallel 1$. \square

Remark 2. Consider two values $v, v' \in \mathbb{N}$ such that $v, v' < 2^{\ell}$, and let COMMON_PREFIX denote the **longest** common prefix of $\text{BITS}_{\ell}(v)$ and $\text{BITS}_{\ell}(v')$. Let NEXT_BITS and $\text{NEXT_BITS}'$ denote two non-empty bitstrings of equal length such that $\text{COMMON_PREFIX} \parallel \text{NEXT_BITS}$ is a prefix of $\text{BITS}_{\ell}(v)$, and $\text{COMMON_PREFIX} \parallel \text{NEXT_BITS}'$ is a prefix of $\text{BITS}_{\ell}(v')$.

If $\text{VAL}(\text{NEXT_BITS}) < \text{VAL}(\text{NEXT_BITS}')$, then:

$$\text{MIN}_{\ell}(\text{COMMON_PREFIX} \parallel \text{NEXT_BITS}'), \text{MAX}_{\ell}(\text{COMMON_PREFIX} \parallel \text{NEXT_BITS}) \in [v, v'].$$

Proof. As $\text{BITS}_\ell(v)$ has prefix $\text{COMMON_PREFIX} \parallel \text{NEXT_BITS}$, v is at most the highest ℓ -bit value having prefix $\text{COMMON_PREFIX} \parallel \text{NEXT_BITS}$. Similarly, as $\text{BITS}_\ell(v')$ has prefix $\text{COMMON_PREFIX} \parallel \text{NEXT_BITS}'$, v is at least the lowest ℓ -bit value having this prefix $\text{COMMON_PREFIX} \parallel \text{NEXT_BITS}'$. Since $\text{VAL}(\text{NEXT_BITS}) < \text{VAL}(\text{NEXT_BITS}')$, we have that $\text{MAX}(\text{COMMON_PREFIX} \parallel \text{NEXT_BITS}) \leq \text{MIN}_\ell(\text{COMMON_PREFIX} \parallel \text{NEXT_BITS}')$. Therefore, we have obtained the following inequality: $v \leq \text{MAX}_\ell(\text{COMMON_PREFIX} \parallel \text{NEXT_BITS}) \leq \text{MIN}_\ell(\text{COMMON_PREFIX} \parallel \text{NEXT_BITS}') \leq v'$. \square

Missing proofs for FindPrefix. We first prove the invariants of each iteration, as described in the proof sketch of Lemma 1.

Lemma 8. *Assume that the following properties hold at the beginning of iteration i .*

- (A) *All honest parties hold the same indices $1 \leq \text{LEFT} \leq \text{RIGHT} \leq \ell + 1$, and the same bitstring PREFIX^* consisting of $\text{LEFT} - 1$ bits.*
- (B) $0 \leq \text{RIGHT} - \text{LEFT} \leq 2^{\lceil \log_2 \ell \rceil - (i-1)}$.
- (C) *Honest parties hold valid ℓ -bit values v such that $\text{BITS}_\ell(v)$ has PREFIX^* as a prefix.*
- (D) *Honest parties hold valid ℓ -bit values v_\perp , and, for any bitstring BITS of RIGHT bits, there are $t + 1$ honest parties holding values v_\perp such that $\text{BITS}_\ell(v_\perp)$ does not have prefix BITS .*

Then, either the condition $\text{LEFT} = \text{RIGHT}$ is met in iteration i , or the properties hold at the beginning of iteration $i + 1$.

Proof. We assume that the condition $\text{LEFT} = \text{RIGHT}$ is not yet met in iteration i (otherwise, the statement trivially holds). Then, Property (B) ensures that $\text{LEFT} < \text{RIGHT}$, and we may prove that the properties hold at the beginning of iteration $i + 1$ as well. The honest parties obtain the same output in the $\Pi_{\ell\text{BA}+}$ invocation of iteration i : either \perp or a sequence of bits, and we split the analysis into these two cases. In the following, we make the iteration number explicit to differentiate between variables' values at the beginning of iteration i and at the beginning of iteration $i + 1$ (i.e. $\text{PREFIX}^*(i)$ is the value held at the beginning of iteration i , and $\text{PREFIX}^*(i + 1)$ is the value computed during iteration i and held at the beginning of iteration $i + 1$).

We first assume that $\Pi_{\ell\text{BA}+}$ returns \perp :

- (A) Honest parties compute the $\text{RIGHT}(i + 1)$ index identically, while all other values remain unchanged. Note that $\text{LEFT}(i) \leq \text{MID} < \text{RIGHT}(i)$ and $\text{RIGHT}(i + 1) := \text{MID}$ still satisfies $1 \leq \text{RIGHT}(i + 1) \leq \ell + 1$. Therefore, Property (A) holds at the beginning of iteration $i + 1$.
- (B) All honest parties compute $\text{RIGHT}(i + 1) := \text{MID} \geq \text{LEFT}(i)$, while the LEFT index remains unchanged: $\text{LEFT}(i + 1) := \text{LEFT}(i)$. We obtain the inequality below, which ensures that Property (B) holds at the beginning of iteration $i + 1$.

$$\begin{aligned} 0 \leq \text{RIGHT}(i + 1) - \text{LEFT}(i + 1) &= \lfloor (\text{LEFT}(i) + \text{RIGHT}(i))/2 \rfloor - \text{LEFT}(i) \\ &\leq (\text{LEFT}(i) + \text{RIGHT}(i))/2 - \text{LEFT}(i) \\ &= (\text{RIGHT}(i) - \text{LEFT}(i))/2 \leq 2^{\lceil \log_2 \ell \rceil - ((i+1)-1)}. \end{aligned}$$

- (C) Since $v(i + 1) := v(i)$, $\text{LEFT}(i + 1) := \text{LEFT}(i)$ and $\text{PREFIX}^*(i + 1) := \text{PREFIX}^*(i)$, Property (C) holds at the beginning of iteration $i + 1$.
- (D) Note that $v_\perp(i + 1) := v(i)$ is a valid ℓ -bit value according to Property (C). We also need to show that, given an arbitrary bitstring $\text{B}_1 \parallel \dots \parallel \text{B}_{\text{MID}}$ of $\text{RIGHT}(i + 1) = \text{MID}$ bits, there are $t + 1$ honest parties holding values $v(i)$ such that $\text{BITS}_\ell(v(i))$ does not have BITS as a prefix. This is ensured by the *Bounded Pre-Agreement* property of $\Pi_{\ell\text{BA}+}$: fewer than $n - 2t$ honest

parties have proposed $B_{\text{LEFT}(i)} \parallel \dots \parallel B_{\text{MID}}$. Therefore, at least $(n-t) - (n-2t-1) \geq t+1$ honest parties hold values $v(i)$ satisfying $B_{\text{LEFT}(i)}(v(i)) \parallel \dots \parallel B_{\text{MID}}(v(i)) \neq B_{\text{LEFT}(i)} \parallel \dots \parallel B_{\text{MID}}$, which implies that the bit representations $\text{BITS}_\ell(v(i))$ do not have prefix $B_1 \parallel \dots \parallel B_{\text{MID}}$. Therefore, Property (D) holds at the beginning of iteration $i+1$.

We now assume that $\Pi_{\ell\text{BA}+}$ returns $\text{PREFIX}_{\text{LEFT}(i)}^* \parallel \dots \parallel \text{PREFIX}_{\text{MID}}^*$:

- (A) The parties compute their $\text{LEFT}(i+1)$ index and the sequence of bits $\text{PREFIX}^*(i+1)$ identically, while the RIGHT index remains unchanged ($\text{RIGHT}(i+1) := \text{RIGHT}(i)$). Note that $\text{PREFIX}^*(i+1)$ is obtained by adding $\text{LEFT}(i+1) - \text{LEFT}(i)$ bits to $\text{PREFIX}^*(i)$, therefore $\text{PREFIX}^*(i+1)$ consists of $\text{LEFT}(i+1) - 1$ bits. In addition, $\text{LEFT}(i) \leq \text{MID} < \text{RIGHT}(i)$ and $\text{LEFT}(i+1) := \text{MID} + 1$ satisfies $1 \leq \text{LEFT}(i+1) \leq \ell + 1$. Therefore, Property (A) holds at the beginning of iteration $i+1$.
- (B) Since $\text{LEFT}(i+1) := \text{MID} + 1 \leq \text{RIGHT}(i)$, while $\text{RIGHT}(i+1) := \text{RIGHT}(i)$, we obtain the inequality below, which ensures that Property (B) holds at the beginning of iteration $i+1$ as well.

$$\begin{aligned} 0 &\leq \text{RIGHT}(i+1) - \text{LEFT}(i+1) = \text{RIGHT}(i) - (\lfloor (\text{LEFT}(i) + \text{RIGHT}(i))/2 \rfloor + 1) \\ &\leq \text{RIGHT}(i) - (\text{LEFT}(i) + \text{RIGHT}(i))/2 \\ &\leq (\text{RIGHT}(i) - \text{LEFT}(i))/2 \leq 2^{\lceil \log_2 \ell \rceil - ((i+1)-1)}. \end{aligned}$$

- (C) Honest parties either hold values $v(i)$ having $\text{PREFIX}^*(i+1)$ as a prefix, or they set $v(i+1)$ to some ℓ -bit value having prefix $\text{PREFIX}^*(i+1)$. This implies that, at the beginning of iteration $i+1$, all honest parties hold ℓ -bit values $v(i)$ with prefix $\text{PREFIX}^*(i+1)$. We still need to prove that values $v(i+1)$ are valid. If $v(i+1) = v(i)$, this follows from Property (C) holding for values $v(i)$.

Otherwise, let P denote an honest party holding $v(i+1) \neq v(i)$. The *Intrusion Tolerance* property of $\Pi_{\ell\text{BA}+}$ ensures that parties agree on a sequence of bits $\text{PREFIX}_{\text{LEFT}(i)}^* \parallel \dots \parallel \text{PREFIX}_{\text{MID}}^*$ that was proposed by an honest party holding value v^* . Then, Property (C) ensures that v^* is a valid ℓ -bit value such that $\text{BITS}_\ell(v^*)$ has prefix $\text{PREFIX}^*(i+1)$. In addition, $v(i)$ is a valid ℓ -bit value, such that $\text{BITS}_\ell(v(i))$ does not have prefix $\text{PREFIX}^*(i+1)$. Then, Remark 2 guarantees that P 's updated value $v(i+1)$ is in $[\min(v(i), v^*), \max(v(i), v^*)]$, and therefore it is an ℓ -bit value within the honest inputs' range.

- (D) Since $\text{RIGHT}(i+1) := \text{RIGHT}(i)$ and $v_\perp(i+1) := v_\perp(i)$, Property (D) is maintained.

□

We may now focus on the proof of Lemma 1.

Lemma 1. *Assume a BA protocol Π_{BA} , and that honest parties join FINDPREFIX with the same ℓ , and with valid ℓ -bit values v . Then, the honest parties obtain the same bitstring PREFIX^* , and each honest party obtains two valid ℓ -bit values v, v_\perp such that: (i) PREFIX^* is a prefix of $\text{BITS}_\ell(v)$; (ii) for any bitstring BITS of $|\text{PREFIX}^*| + 1$ bits, at least $t+1$ honest parties hold values v_\perp such that $\text{BITS}_\ell(v_\perp)$ does not have prefix BITS .*

FINDPREFIX achieves communication complexity $\text{BITS}_\ell(\text{FINDPREFIX}) = O(\ell \cdot n + \kappa \cdot n^2 \log n \log \ell) + O(\log \ell) \cdot \text{BITS}_\kappa(\Pi_{\text{BA}})$, and round complexity $\text{ROUNDS}_\ell(\text{FINDPREFIX}) = O(\log \ell) \cdot \text{ROUNDS}_\kappa(\Pi_{\text{BA}})$.

Proof. The properties listed in Lemma 8 hold in iteration 1 due to the variables' initialization. Hence, these properties hold for every iteration of the loop.

Property (A) ensures that honest parties hold the same indices LEFT and RIGHT in every iteration of the loop. Once the condition LEFT = RIGHT is met, Property (C) guarantees that honest parties hold valid ℓ -bit values v having the bitstring PREFIX^{*} as a common prefix. According to Property (A), this common prefix consists of $i^* := \text{LEFT} - 1$ bits. From Property (D), it follows that the honest parties hold valid ℓ -bit values v_\perp . The same property implies that, for any bitstring BITS of RIGHT = $i^* + 1$ bits, the ℓ -bit representations of $t + 1$ honest parties do not have BITS as a prefix. Hence, once the stopping condition holds, honest parties hold values i^* , v , and v_\perp satisfying the guarantees in the lemma's statement. It remains to show that the stopping condition indeed holds eventually.

Note that the condition LEFT = RIGHT is met (for all honest parties simultaneously, due to Property (A)) by iteration $i := \lceil \log_2 \ell \rceil + 2$. Property (B) ensures that, at the beginning of iteration i , $0 \leq \text{RIGHT} - \text{LEFT} \leq 2^{\lceil \log_2 \ell \rceil - (i-1)}$. Then, if this condition was not met by iteration $i := \lceil \log_2 \ell \rceil + 2$, the indices LEFT and RIGHT obtained by the honest parties in iteration i satisfy $0 \leq \text{RIGHT} - \text{LEFT} \leq 2^{\lceil \log_2 \ell \rceil - (\lceil \log_2 \ell \rceil + 1)} \leq 2^{-1}$. Since the indices LEFT and RIGHT are natural numbers, we may conclude that RIGHT - LEFT = 0.

We may then discuss the round complexity of FINDPREFIX: since $O(\log \ell)$ iterations are sufficient and each iteration invokes $\Pi_{\ell\text{BA}+}$ once, we obtain that $\text{ROUNDS}_\ell(\text{FINDPREFIX}) = O(\log \ell) \cdot \text{ROUNDS}_\ell(\Pi_{\ell\text{BA}+})$. Then, Theorem 1 leads to the result claimed in the lemma's statement.

For the communication complexity, Property (B) of Lemma 8 ensures that, in each iteration $i < \lceil \log_2 \ell \rceil + 2$, FINDPREFIX runs $\Pi_{\ell\text{BA}+}$ on inputs of at most $2^{\lceil \log_2 \ell \rceil - i}$ bits, hence of at most $\ell/2^{i-1}$ bits. Therefore, $\text{BITS}_\ell(\text{FINDPREFIX}) = \sum_{i=1}^{\lceil \log_2 \ell \rceil + 1} \text{BITS}_{\ell/2^{i-1}}(\Pi_{\ell\text{BA}+})$. Using Theorem 1, and the fact that $\sum_{i=0}^{\infty} 1/2^i \leq 2$, we obtain that $\text{BITS}_\ell(\text{FINDPREFIX}) = O(\ell \cdot n + \kappa \cdot n^2 \log n \log \ell) + O(\log \ell) \cdot \text{BITS}_\kappa(\Pi_{\text{BA}})$. \square

Missing proofs for AddLastBit. We present the proof of Lemma 2. This provides the guarantees of ADDLASTBIT, which enables the honest parties to extend the prefix obtained in FINDPREFIX with one bit.

Lemma 2. *Assume a BA protocol Π_{BA} , and that honest parties join ADDLASTBIT with the same value ℓ , the same bitstring PREFIX^{*} of $i^* < \ell$ bits, and with valid ℓ -bit values v such that $\text{BITS}_\ell(v)$ has prefix PREFIX^{*}. Then, the honest parties agree on a bitstring of $i^* + 1$ bits that is the prefix of a valid value's ℓ -bit representation. ADDLASTBIT has communication complexity $\text{BITS}_\ell(\text{ADDLASTBIT}) = \text{BITS}_1(\Pi_{\text{BA}})$ and round complexity $\text{ROUNDS}_\ell(\text{ADDLASTBIT}) = \text{ROUNDS}_1(\Pi_{\text{BA}})$.*

Proof. Since the honest parties hold the same bitstring PREFIX^{*} when joining the subprotocol, Π_{BA} ensures that they obtain the same bitstring of PREFIX^{*} \parallel B^{*} of $|\text{PREFIX}^*| + 1$ bits. Moreover, the Validity property of Π_{BA} ensures that the bit agreed upon B^{*} was proposed by an honest party. Hence, there is an honest party holding a valid value v whose ℓ -bit representation has PREFIX^{*} \parallel B^{*} as prefix. The communication complexity and the round complexity follow from the fact that ADDLASTBIT only invokes Π_{BA} once on one-bit inputs. \square

Missing proofs for GetOutput. We present the proof of Lemma 3. This describes the subprotocol GETOUTPUT, which enables the honest parties to obtain the final output.

Lemma 3. *Assume a BA protocol Π_{BA} , and that honest parties join GETOUTPUT with the same value ℓ , and with the same bitstring PREFIX^{*} representing the prefix of some valid value's ℓ -bit representation. In addition, assume that each party joins with some valid ℓ -bit input v_\perp such that*

the ℓ -bit representations of $t + 1$ honest parties' values v_\perp do not have PREFIX^* as a prefix. Then, the honest parties obtain the same valid value v_{OUT} .

GETOUTPUT has communication complexity $\text{BITS}_\ell(\text{GETOUTPUT}) = O(n^2) + \text{BITS}_1(\Pi_{\text{BA}})$ and round complexity $\text{ROUNDS}_\ell(\text{GETOUTPUT}) = O(1) + \text{ROUNDS}_1(\text{BA})$.

Proof. There are $t + 1$ honest parties holding values v_\perp such that $\text{BITS}_\ell(v_\perp)$ does not have PREFIX^* as prefix. For each of these parties, v_\perp is either lower than $\text{MIN}_\ell(\text{PREFIX}^*)$ or higher than $\text{MAX}_\ell(\text{PREFIX}^*)$. This implies that there are at least $t + 1$ honest parties sending bits B . Hence, each party receives $m \geq t + 1$ bits B .

We need to show that each honest party's CHOICE is a bit B sent by an honest party. Let P denote an honest party that obtained some value CHOICE , and assume that no honest party has sent $B = \text{CHOICE}$. Hence, P has received at most t bits CHOICE , and at least $t + 1$ honest bits $1 - \text{CHOICE}$. We obtain a contradiction: P has received $m \geq 2t + 1$ bits B , and $\lceil m/2 \rceil > t$. This means that P did not receive CHOICE from $\lceil m/2 \rceil$ parties.

Then, each honest party joins Π_{BA} with an honest party's bit B as input, and therefore they agree on an honest party's bit B due to Π_{BA} 's Validity condition.

If the bit agreed upon is 0, some honest party holds $v_\perp < \text{MIN}_\ell(\text{PREFIX}^*)$. Since PREFIX^* is some valid value's prefix, $\text{MIN}_\ell(\text{PREFIX}^*)$ is valid. Similarly, if the bit agreed upon is 1, some honest party holds $v_\perp > \text{MAX}_\ell(\text{PREFIX}^*)$. Since PREFIX^* is some valid value's prefix, $\text{MAX}_\ell(\text{PREFIX}^*)$ is a valid value.

For the communication complexity and round complexity, note that GETOUTPUT makes use of one round of communication where the parties sent at most a bit to all parties, and afterwards the parties run Π_{BA} on one-bit inputs. \square

A.3 Missing proofs for FixedLengthCABlocks

Missing proofs for FindPrefixBlocks. We recall that the main difference between FINDPREFIX and FINDPREFIXBLOCKS is that the first implements binary search on *bits*, while the latter implements binary search on *blocks*. This will be the main difference in the analysis as well.

We start by analyzing the invariants of each iteration. The lemma below is a variant of Lemma 8 on blocks, and the proof is also reflects this.

Lemma 9. *Assume that the following properties hold at the beginning of iteration i .*

- (A) *All honest parties hold the same indices $1 \leq \text{LEFT} \leq \text{RIGHT} \leq \ell + 1$, and the same bitstring PREFIX^* consisting of $\text{LEFT} - 1$ blocks.*
- (B) *$0 \leq \text{RIGHT} - \text{LEFT} \leq 2^{\lceil \log_2 n^2 \rceil - (i-1)}$.*
- (C) *Honest parties hold valid ℓ -bit values v such that $\text{BITS}_\ell(v)$ has PREFIX^* as a prefix.*
- (D) *Honest parties hold valid ℓ -bit values v_\perp , and, for any bitstring BITS of RIGHT blocks, the ℓ -bit representations of the values v_\perp of $t + 1$ honest parties do not have prefix BITS .*

Then, either the condition $\text{LEFT} = \text{RIGHT}$ is met in iteration i , or the properties hold at the beginning of iteration $i + 1$.

Proof. We assume that the condition $\text{LEFT} = \text{RIGHT}$ is not yet met in iteration i (otherwise, the statement trivially holds). Then, Property (B) ensures that $\text{LEFT} < \text{RIGHT}$, and we may prove that the properties hold at the beginning of iteration $i + 1$ as well. The honest parties obtain the same output in the $\Pi_{\ell\text{BA}+}$ invocation of iteration i : either \perp or a sequence of blocks, and we split the analysis into these two cases. In the following, we make the iteration number explicit to differentiate between variables' values at the beginning of iteration i and at the beginning of iteration $i + 1$ (i.e. $\text{PREFIX}^*(i)$ is the value held at the beginning of iteration i , and $\text{PREFIX}^*(i + 1)$ is the value computed during iteration i and held at the beginning of iteration $i + 1$).

We first assume that $\Pi_{\ell\text{BA}+}$ returns \perp :

- (A) Honest parties compute the $\text{RIGHT}(i+1)$ index identically, while all other values remain unchanged. Note that $\text{LEFT}(i) \leq \text{MID} < \text{RIGHT}(i)$ and therefore $\text{RIGHT}(i+1) := \text{MID}$ still satisfies $1 \leq \text{RIGHT}(i+1) \leq n^2 + 1$. Therefore, Property (A) holds at the beginning of iteration $i+1$ as well.
- (B) All honest parties compute $\text{RIGHT}(i+1) := \text{MID} \geq \text{LEFT}(i)$, while the LEFT index remains unchanged: $\text{LEFT}(i+1) := \text{LEFT}(i)$. We obtain the inequality below, which ensures that Property (B) holds at the beginning of iteration $i+1$.

$$\begin{aligned} 0 \leq \text{RIGHT}(i+1) - \text{LEFT}(i+1) &= \lfloor (\text{LEFT}(i) + \text{RIGHT}(i))/2 \rfloor - \text{LEFT}(i) \\ &\leq (\text{LEFT}(i) + \text{RIGHT}(i))/2 - \text{LEFT}(i) \\ &= (\text{RIGHT}(i) - \text{LEFT}(i))/2 \leq 2^{\lceil \log_2 n^2 \rceil - ((i+1)-1)}. \end{aligned}$$

- (C) Since $v(i+1) := v(i)$, $\text{LEFT}(i+1) := \text{LEFT}(i)$ and $\text{PREFIX}^*(i+1) := \text{PREFIX}^*(i)$, Property (C) holds at the beginning of iteration $i+1$.
- (D) Note that $v_{\perp}(i+1) := v(i)$ is a valid ℓ -bit value according to Property (C). We also need to show that, given an arbitrary bitstring $\text{BLOCK}_1 \parallel \dots \parallel \text{BLOCK}_{\text{MID}}$ of $\text{RIGHT}(i+1) = \text{MID}$ blocks, there are $t+1$ honest parties holding values $v(i)$ such that $\text{BITS}_{\ell}(v(i))$ does not have BITS as a prefix. This is ensured by the *Bounded Pre-Agreement* property of $\Pi_{\ell\text{BA}+}$: $t+1$ honest parties hold values $v(i)$ satisfying $\text{BLOCK}_{\text{LEFT}(i)}(v(i)) \parallel \dots \parallel \text{BLOCK}_{\text{MID}}(v(i)) \neq \text{BLOCK}_{\text{LEFT}(i)} \parallel \dots \parallel \text{BLOCK}_{\text{MID}}$, which implies that the bit representations $\text{BITS}_{\ell}(v(i))$ do not have prefix $\text{BLOCK}_1 \parallel \dots \parallel \text{BLOCK}_{\text{MID}}$. Therefore, Property (D) holds at the beginning of iteration $i+1$.

We now assume that $\Pi_{\ell\text{BA}+}$ returns $\text{PREFIX}_{\text{LEFT}(i)}^* \parallel \dots \parallel \text{PREFIX}_{\text{MID}}^*$:

- (A) The parties compute their $\text{LEFT}(i+1)$ index and the sequence of blocks $\text{PREFIX}^*(i+1)$ identically, while the RIGHT index remains unchanged ($\text{RIGHT}(i+1) := \text{RIGHT}(i)$). Note that $\text{PREFIX}^*(i+1)$ is obtained by adding $\text{LEFT}(i+1) - \text{LEFT}(i)$ blocks to $\text{PREFIX}^*(i)$, therefore $\text{PREFIX}^*(i+1)$ consists of $\text{LEFT}(i+1) - 1$ blocks. In addition, $\text{LEFT}(i) \leq \text{MID} < \text{RIGHT}(i)$ and therefore $\text{LEFT}(i+1) := \text{MID} + 1$ still satisfies $1 \leq \text{LEFT}(i+1) \leq n^2 + 1$. Therefore, Property (A) holds at the beginning of iteration $i+1$.
- (B) Since $\text{LEFT}(i+1) := \text{MID} + 1 \leq \text{RIGHT}(i)$, while $\text{RIGHT}(i+1) := \text{RIGHT}(i)$, we obtain the inequality below, which ensures that Property (B) holds at the beginning of iteration $i+1$ as well.

$$\begin{aligned} 0 \leq \text{RIGHT}(i+1) - \text{LEFT}(i+1) &= \text{RIGHT}(i) - (\lfloor (\text{LEFT}(i) + \text{RIGHT}(i))/2 \rfloor + 1) \\ &\leq \text{RIGHT}(i) - (\text{LEFT}(i) + \text{RIGHT}(i))/2 \\ &\leq (\text{RIGHT}(i) - \text{LEFT}(i))/2 \leq 2^{\lceil \log_2 n^2 \rceil - ((i+1)-1)}. \end{aligned}$$

- (C) Honest parties either hold values $v(i)$ having $\text{PREFIX}^*(i+1)$ as a prefix, or they set $v(i+1)$ to some ℓ -bit value having prefix $\text{PREFIX}^*(i+1)$. This implies that, at the beginning of iteration $i+1$, all honest parties hold ℓ -bit values $v(i)$ with prefix $\text{PREFIX}^*(i+1)$. We still need to prove that values $v(i+1)$ are valid. If $v(i+1) = v(i)$, this follows from Property (C) holding for values $v(i)$. Otherwise, let P denote an honest party holding $v(i+1) \neq v(i)$.

The *Intrusion Tolerance* property of $\Pi_{\ell\text{BA}+}$ ensures that parties agree on a sequence of blocks $\text{PREFIX}_{\text{LEFT}(i)}^* \parallel \dots \parallel \text{PREFIX}_{\text{MID}}^*$ that was proposed by an honest party holding value v^* . Then, Property (C) ensures that v^* is a valid ℓ -bit value such that $\text{BITS}_\ell(v^*)$ has prefix $\text{PREFIX}^*(i+1)$. On the other hand, $v(i)$ is a valid ℓ -bit value such that $\text{BITS}_\ell(v(i))$ has $\text{PREFIX}^*(i)$ as a prefix, but not $\text{PREFIX}^*(i+1)$. Remark 2 guarantees that P 's updated value $v(i+1)$ is in $[\min(v(i), v^*), \max(v(i), v^*)]$, and therefore it is an ℓ -bit value within the honest inputs' range.

(D) Since $\text{RIGHT}(i+1) := \text{RIGHT}(i)$ and $v_\perp(i+1) := v_\perp(i)$, Property (D) is maintained. \square

Then, the proof of Lemma 4 will be similar to that of Lemma 1: the main difference is, once again, that we consider blocks instead of bits.

Lemma 4. *Assume a BA protocol Π_{BA} , and that honest parties join FINDPREFIXBLOCKS with the same (multiple of n^2) ℓ , and with valid ℓ -bit values v . Then, the honest parties obtain the same bitstring PREFIX^* of i^* blocks, and each honest party obtains two valid ℓ -bit values v, v_\perp such that: (i) the ℓ -bit representations of the values v have prefix PREFIX^* ; (ii) for any bitstring BITS of $i^* + 1$ blocks, at least $t + 1$ honest parties hold values v_\perp such that $\text{BITS}_\ell(v_\perp)$ does not have prefix BITS .*

FINDPREFIXBLOCKS achieves communication complexity $\text{BITS}_\ell(\text{FINDPREFIXBLOCKS}) = O(\ell \cdot n + \kappa \cdot n^2 \log^2 n) + O(\log n) \cdot \text{BITS}_\kappa(\Pi_{\text{BA}})$ and round complexity $\text{ROUNDS}_\ell(\text{FINDPREFIXBLOCKS}) = O(\log n) \cdot \text{ROUNDS}_\kappa(\Pi_{\text{BA}})$.

Proof. The properties listed in Lemma 9 hold in iteration 1 due to the variables' initialization. Hence, these properties hold for every iteration of the loop.

Property (A) ensures that honest parties hold the same indices LEFT and RIGHT in every iteration of the loop. Once the condition $\text{LEFT} = \text{RIGHT}$ is met, Property (C) guarantees that honest parties hold valid ℓ -bit values v having the bitstring PREFIX^* as a common prefix. According to Property (A), this common prefix consists of $i^* := \text{LEFT} - 1$ blocks. From Property (D), it follows that the honest parties hold valid ℓ -bit values v_\perp . The same property implies that, for any bitstring BITS of $\text{RIGHT} = i^* + 1$ blocks, the ℓ -bit representations of $t + 1$ honest parties' values v_\perp do not have BITS as a prefix. Hence, once the stopping condition is met, honest parties hold values i^*, v , and v_\perp satisfying the guarantees in the lemma's statement. It remains to show that the stopping condition indeed holds eventually.

Note that the condition $\text{LEFT} = \text{RIGHT}$ is met (for all honest parties simultaneously, due to Property (A)) by iteration $i := \lceil \log_2 n^2 \rceil + 2$. Property (B) ensures that, at the beginning of iteration i , $0 \leq \text{RIGHT} - \text{LEFT} \leq 2^{\lceil \log_2 n^2 \rceil - (i-1)}$. Then, if this condition was not met by iteration $i := \lceil \log_2 n^2 \rceil + 2$, the indices LEFT and RIGHT obtained by the honest parties in iteration i satisfy $0 \leq \text{RIGHT} - \text{LEFT} \leq 2^{\lceil \log_2 n^2 \rceil - (\lceil \log_2 n^2 \rceil + 1)} \leq 2^{-1}$. Since the indices LEFT and RIGHT are natural numbers, we may conclude that $\text{RIGHT} - \text{LEFT} = 0$.

We may then discuss the round complexity of FINDPREFIXBLOCKS : since $O(\log n)$ iterations are sufficient and each iteration invokes $\Pi_{\ell\text{BA}+}$ once, we obtain that $\text{ROUNDS}_\ell(\text{FINDPREFIXBLOCKS}) = O(\log n) \cdot \text{ROUNDS}_\ell(\Pi_{\ell\text{BA}+})$. Then, Theorem 1 leads to the result claimed in the lemma's statement.

For the communication complexity, Property (B) of Lemma 9 ensures that, in each iteration $i < \lceil \log_2 n^2 \rceil + 2$, FINDPREFIXBLOCKS runs $\Pi_{\ell\text{BA}+}$ on inputs of at most $2^{\lceil \log_2 n^2 \rceil - i}$ blocks, hence of at most $2^{\lceil \log_2 n^2 \rceil - i} \cdot \ell / n^2 \leq \ell / 2^{i-1}$ bits. Therefore, $\text{BITS}_\ell(\text{FINDPREFIXBLOCKS}) = \sum_{i=1}^{\lceil \log_2 n^2 \rceil + 1} \text{BITS}_{\ell/2^{i-1}}(\Pi_{\ell\text{BA}+})$. Using Theorem 1, and the fact that $\sum_{i=0}^{\infty} 1/2^i \leq 2$, we obtain that $\text{BITS}_\ell(\text{FINDPREFIXBLOCKS}) = O(\ell \cdot n + \kappa \cdot n^2 \log^2 n) + O(n \log n) \cdot \text{BITS}_\kappa(\Pi_{\text{BA}})$. \square

Missing proofs for AddLastBlock. We include the proof of Lemma 5 below. This describes the subprotocol ADDLASTBLOCK, which enables the honest parties to extend the prefix obtained in FINDPREFIXBLOCKS with one block.

Lemma 5. *Assume that the honest parties join ADDLASTBLOCK with the same value ℓ (that is a multiple of n^2), with the same bitstring prefix PREFIX^* of $i^* < n^2$ blocks, and with valid ℓ -bit values v that have PREFIX^* as a prefix. Then, the honest parties agree on a bitstring of $i^* + 1$ blocks that is the prefix of a valid value's ℓ -bit representation. ADDLASTBLOCK has communication complexity $\text{BITS}_\ell(\text{ADDLASTBLOCK}) = O(\ell \cdot n)$ and round complexity $\text{ROUNDS}_\ell(\text{ADDLASTBLOCK}) = O(n)$.*

Proof. HIGHCOSTCA ensures that the honest parties obtain the same bitstring BLOCK'_{i^*+1} , that is within the honest range of blocks $\text{BLOCK}_{i^*+1}(v)$. That is, some honest parties P_1 and P_2 have joined with block $\text{BLOCK}^1, \text{BLOCK}^2$ such that $\text{VAL}(\text{BLOCK}^1) \leq \text{VAL}(\text{BLOCK}'_{i^*+1}) \leq \text{VAL}(\text{BLOCK}^2)$.

Then, since all honest parties have joined HIGHCOSTCA with bitstrings of ℓ/n^2 bits, BLOCK'_{i^*+1} is also a block. Then, since the honest parties hold the same bitstring PREFIX^* of i^* blocks, they obtain the same bitstring $\text{PREFIX}^* \parallel \text{BLOCK}'_{i^*+1}$ of $i^* + 1$ blocks.

It remains to show that $\text{PREFIX}^* \parallel \text{BLOCK}'_{i^*+1}$ is some valid values' prefix. The honest parties P_1 and P_2 hold valid values v^1 and v^2 such that $\text{BITS}_\ell(v^1)$ has prefix $\text{PREFIX}^* \parallel \text{BLOCK}^1$, and $\text{BITS}_\ell(v^2)$ has prefix $\text{PREFIX}^* \parallel \text{BLOCK}^2$. Then, since $\text{VAL}(\text{BLOCK}^1) \leq \text{VAL}(\text{BLOCK}'_{i^*+1}) \leq \text{VAL}(\text{BLOCK}^2)$, there is a valid value v^* such that $\text{BITS}_\ell(v^*)$ has $\text{PREFIX}^* \parallel \text{BLOCK}'_{i^*+1}$ as prefix.

The communication and round complexities follow from those of HIGHCOSTCA presented in Theorem 3, as ADDLASTBLOCK invokes HIGHCOSTCA once on inputs of ℓ/n^2 bits. \square

Missing proofs for FixedLengthCABlocks. We include the proof of Theorem 4, describing the guarantees of FIXEDLENGTHCABLOCKS, restated below.

Theorem 4. *Assume a BA protocol Π_{BA} resilient against $t < n/3$ corruptions. If the honest parties hold ℓ -bit inputs $v_{\text{IN}} \in \mathbb{N}$, where ℓ is a publicly known multiple of n^2 , FIXEDLENGTHCABLOCKS is a CA protocol resilient against $t < n/3$ corruptions.*

The protocol achieves communication complexity $\text{BITS}_\ell(\text{FIXEDLENGTHCABLOCKS}) = O(\ell n + \kappa \cdot n^2 \log^2 n) + O(\log n) \cdot \text{BITS}_\kappa(\Pi_{\text{BA}})$, and round complexity $\text{ROUNDS}_\ell(\text{FIXEDLENGTHCABLOCKS}) = O(n) + O(\log n) \cdot \text{ROUNDS}_\kappa(\Pi_{\text{BA}})$.

Proof. Lemma 4 enables the parties to agree on a bitstring PREFIX^* , and provides them with valid ℓ -bit values v, v_\perp such that the values v have prefix PREFIX^* . If $|\text{PREFIX}^*| = \ell$, the honest parties hold the same valid value v , and therefore CA is achieved. Otherwise, Lemma 5 ensures that parties obtain the same bitstring PREFIX^* such that there are $t + 1$ honest parties whose values v_\perp do not have PREFIX^* as a prefix. Then, GETOUTPUT's preconditions are met, and Lemma 3 ensures that CA is achieved. The communication complexity and the round complexity follow by summing up the complexities of each subprotocol. \square

A.4 High-Communication-Cost CA

In subprotocol ADDLASTBLOCK of FIXEDLENGTHCABLOCKS, we have used a high-communication-cost CA protocol, described in Theorem 3, restated below.

Theorem 3 (Theorem 4 of [47]). *There is a CA protocol HIGHCOSTCA for \mathbb{N} resilient against $t < n/3$ corruptions, with communication complexity $\text{BITS}_\ell(\text{HIGHCOSTCA}) = O(\ell \cdot n^3)$, and round complexity $\text{ROUNDS}_\ell(\text{HIGHCOSTCA}) = O(n)$.*

In this section, we present the protocol HIGHCOSTCA, obtained by making minor adjustments to the Median Validity protocol of [47]. This is a variant of the well-known King BA protocol [7].

This protocol involves a *setup stage*, where the parties distribute their inputs and each party P estimates a *trusted* interval. In the protocol of [47], this is an interval containing values *close* to the honest median. For us, any interval included in the honest inputs' range suffices, allowing us to slightly simplify the protocol. If a party receives $n - t + k$ values, then at most k out of these values were sent by the byzantine parties. Hence, the interval between the $(k + 1)$ -th lowest and the $(k + 1)$ -th highest values received is included in the honest inputs' range. Each party afterwards sends its trusted interval to all parties. We need to be careful about a small technical detail here – byzantine parties may send non-integer values that fit into the trusted intervals, and honest parties forward forwarding such values would increase the communication cost. To prevent this, we take into account that we only run HIGHCOSTCA on values in \mathbb{N} , and therefore the honest parties may ignore any values outside \mathbb{N} in each step of the protocol.

At the end of the setup stage, each party chooses a SUGGESTION value: this is a value that appears in $n - t$ of the intervals received, hence in at least $t + 1$ honest intervals, which is roughly *likely* to get support in the *search stage*.

The parties then start the *search stage*. The only adjustment we make here is that parties ignore all values outside \mathbb{N} . The parties run $t + 1$ sequential phases. In each phase i , P_i is *the king*. Roughly, the king distributes the value it believes the parties should agree on. If the king is honest (and $t + 1$ phases means at least one honest king), all honest parties accept the king's suggestion, and agreement is achieved. In addition, the agreement obtained is maintained in all further phases. We present the code below.

HIGHCOSTCA

Code for party P with input v_{in}

- 1: *Setup stage*
- 2: Send v_{in} to all parties.
- 3: Out of the $n - t + k$ values in \mathbb{N} received, set INTERVALMIN := the $(k + 1)$ -th lowest value received, and INTERVALMAX := the $(k + 1)$ -th highest value received.
- 4: Send INTERVALMIN, INTERVALMAX to all parties. (ℓn^2).
- 5: Let SUGGESTION := some value in \mathbb{N} that appears in $n - t$ of the intervals received.
- 6: Let CURRENT := SUGGESTION.
- 7:
- 8: *Search stage*
- 9: **for** $i = 1 \dots t + 1$ **do**
- 10: Send CURRENT to all parties.
- 11: If you have received the same value $v \in \mathbb{N}$ from $n - t$ parties, send (PROPOSE, v) to all parties.
- 12: If you have received the same (PROPOSE, v) with $v \in \mathbb{N}$ from $t + 1$ parties, set CURRENT = v .
- 13: **King P_i only:**
- 14: If you have received the same (PROPOSE, v) with $v \in \mathbb{N}$ from $t + 1$ parties: KINGVALUE := v .
- 15: Otherwise, set KINGVALUE := SUGGESTION.
- 16: Send KINGVALUE to all parties.
- 17: If KINGVALUE = CURRENT or KINGVALUE \in [INTERVALMIN, INTERVALMAX] \cap \mathbb{N} :
- 18: Send (VOTE, KINGVALUE) to all parties.
- 19: If you have not received $n - t$ messages (PROPOSE, v') for any value $v' \in \mathbb{N}$:
- 20: If you have received $t + 1$ messages (VOTE, KINGVALUE) for some value KINGVALUE \in \mathbb{N} :
- 21: Set CURRENT := KINGVALUE.
- 22: **end for**
- 23: Output CURRENT.

We now present the analysis of HIGHCOSTCA. Most of the lemmas below follow the analysis of [47].

The lemma below ensures that the interval obtained by each honest party is indeed in the honest inputs' range.

Lemma 10. *Let v_1, v_2, \dots, v_{n-t} be the $n - t$ honest inputs arranged in increasing order. Then, for every honest party P , INTERVALMIN and INTERVALMAX are well-defined and satisfy: $v_1 \leq \text{INTERVALMIN} \leq v_{t+1} \leq \text{INTERVALMAX} \leq v_{n-t}$.*

Proof. P receives $n - t + k$ values v_{IN} , where $0 \leq k \leq t$. Since the $n - t$ honest inputs are received, only k of these $n - t + k$ values are sent by byzantine parties, and hence may be outside the honest inputs' range. Hence, there are at most k values lower than v_1 , and at most k values higher than v_{n-t} .

Note that both INTERVALMIN and INTERVALMAX are well defined, since $k + 1 \leq (n - t + k)$: there is a $(k + 1)$ -th lowest value received, and a $(k + 1)$ -th highest value received.

Since there are at most k values lower than v_1 , INTERVALMIN := the $(k + 1)$ -th lowest value received is at least v_1 . Moreover, since all honest inputs are received, INTERVALMIN $\leq v_{k+1} \leq v_{t+1}$.

Similarly, since there are at most k values higher than v_{n-t} , INTERVALMAX := the $(k + 1)$ -th highest value received is at most v_{max} . Moreover, since all honest inputs are received, INTERVALMAX $\geq v_{(n-t)-k} \geq v_{n-2t}$. Since $n > 3t$, we obtain that $n - 2t \geq t + 1$, and therefore INTERVALMAX $\geq v_{t+1}$. \square

The following two properties are immediate corollaries of Lemma 10 and ensure the success of the setup stage.

Corollary 3. *For any honest party P , $[\text{INTERVALMIN}, \text{INTERVALMAX}]$ is non-empty and it is a subset of the honest inputs' range.*

Corollary 4. *The intervals $[\text{INTERVALMIN}, \text{INTERVALMAX}]$ obtained by the honest parties have a non-empty intersection. Moreover, the intersection contains some natural number.*

The next lemma ensures that, at all times, every honest party holds a value CURRENT that is in some honest party's trusted interval. This also implies that honest parties hold valid values at all times.

Lemma 11. *Assume that at the beginning of iteration i , every honest party P holds a value CURRENT that is in some honest party's interval $[\text{INTERVALMIN}, \text{INTERVALMAX}]$.*

Then, at the end of iteration i , the same property holds: every honest party P holds a value CURRENT that is in some honest party's interval $[\text{INTERVALMIN}, \text{INTERVALMAX}]$.

Proof. P may update its value CURRENT to some value v if it receives $t + 1$ messages (PROPOSE, v). If this is the case, at least one of these PROPOSE messages was sent by an honest party P' , and v is the value CURRENT held by P' at the beginning of iteration i . This means that v is also a value in within some honest party's interval $[\text{INTERVALMIN}, \text{INTERVALMAX}]$. Hence, up to line 12, all honest parties hold values CURRENT satisfying this property.

Afterwards, if P did not receive $n - t$ messages (PROPOSE, v) for some value v , it may try to update its value CURRENT to the king's suggestion. P first checks if it has received $t + 1$ messages (VOTE, KINGVALUE). If this is the case, then at least one honest party P' has sent a (VOTE, KINGVALUE) message. Then, there are two cases: KINGVALUE satisfies KINGVALUE $\in [\text{INTERVALMIN}, \text{INTERVALMAX}] \cap \mathbb{N}$ for P' , or KINGVALUE is the value CURRENT held by P' , which we have proved to be in some honest party's interval $[\text{INTERVALMIN}, \text{INTERVALMAX}]$. \square

We now ensure that, if two honest parties update their value `CURRENT` at the beginning of a phase, then they update `CURRENT` to the same value.

Lemma 12. *If, in iteration i , an honest party P receives $t + 1$ messages $(\text{PROPOSE}, v)$ for some value v , no honest party receives $t + 1$ messages $(\text{PROPOSE}, v')$ for some value $v' \neq v$.*

Proof. Assume that an honest party P' receives $t + 1$ messages $(\text{PROPOSE}, v')$ for some value $v' \neq v$.

P has received $t + 1$ `PROPOSE` messages for v , hence at least one honest party P'' has sent $(\text{PROPOSE}, v)$. This implies that P'' has received v from $n - t$ parties in line 11, hence from at least $n - 2t$ honest parties.

Similarly, P' has received $t + 1$ `PROPOSE` messages for v' , hence at one honest party P''' has sent $(\text{PROPOSE}, v')$. This implies that P''' has received v' from $n - t$ parties in line 10, hence from at least $n - 2t$ honest parties. We obtain a contradiction: at least $n - 2t$ out honest parties have sent $v \neq v'$ and there are only $(n - t) - (n - 2t) = t < n - 2t$ honest parties that could have sent v' . \square

The lemma below ensures that, once agreement is reached in some phase, it is maintained in all further phases.

Lemma 13. *If all honest parties hold the same value `CURRENT` at the beginning of iteration i , no honest party changes its value `CURRENT` during iteration i .*

Proof. All honest parties send `CURRENT`, and therefore all honest parties receive $n - t$ `PROPOSE` messages for this value `CURRENT`. No honest party sends $(\text{PROPOSE}, v)$ with $v \neq \text{CURRENT}$, and therefore no honest party receives $t + 1$ `PROPOSE` messages for $v \neq \text{CURRENT}$. Regardless of whether the king of this iteration is honest or not, parties have received $n - t$ messages `PROPOSE` for `CURRENT` and therefore no honest party updates its value to the king's suggestion. Therefore, the honest parties maintain their value `CURRENT`. \square

The lemma below ensures that, in the first phase with an honest king, agreement is reached.

Lemma 14. *If, in iteration i , the king P_i is honest, then the honest parties hold the same value `CURRENT` at the end of iteration i .*

Proof. If the honest parties have started iteration i with the same value `CURRENT`, Lemma 13 ensures that the honest parties complete iteration i with the same value `CURRENT`.

We may then assume that honest parties held different values `CURRENT` at the beginning of iteration i . The remainder of the proof will be split into two cases, depending on how the honest king P_i defines its `KINGVALUE`.

Case 1: P_i has received the same $(\text{propose}, v)$ from $t + 1$ parties. Then, P_i sets `KINGVALUE` $:= v$.

We first show that at least $t + 1$ honest parties send (VOTE, v) in line 18. Since P_i has received $t + 1$ `PROPOSE` messages for v , at least one honest party has sent $(\text{PROPOSE}, v)$ and therefore it has received v from $n - t$ parties. This implies that $t + 1$ honest parties held `CURRENT` $:= v$ at the beginning of the iteration. Lemma 12 ensures that every honest party receives at most t `PROPOSE` messages for any value $v' \neq v$, and therefore these honest parties still hold `CURRENT` $:= v$ in line 17. Therefore, the condition `KINGVALUE` = `CURRENT` in line 17 holds for at least $t + 1$ honest parties, and $t + 1$ honest parties send (VOTE, v) . Moreover, since P_i is honest, no honest party sends (VOTE, v') for $v' \neq v$.

We now show that every honest party P holds `CURRENT` = v by the end of the iteration. P receives $t + 1$ messages (VOTE, v) and at most t votes for any other values. Then, if P has not received $n - t$ `PROPOSE` messages for some value v' , it sets `CURRENT` $:= v$ in line 21. Otherwise,

if P has received $n - t$ PROPOSE messages for some value v' , these messages are for $v' = v$, as guaranteed by Lemma 12. This implies that P has set $\text{CURRENT} := v$ in line 12.

Case 2: P_i has not received the same (propose, v) from $t + 1$ parties. In this case, P_i sets $\text{KINGVALUE} := \text{SUGGESTION}$. P_i has set SUGGESTION in line 5 to some natural number that appears in $n - t$ of the intervals $[\text{INTERVALMIN}, \text{INTERVALMAX}]$ it has received, therefore in at least $n - 2t \geq t + 1$ of the honest parties' intervals.

Then, the condition $\text{KINGVALUE} \in [\text{INTERVALMIN}, \text{INTERVALMAX}] \cap \mathbb{N}$ holds for at least $t + 1$ honest parties, and these honest parties send (VOTE, KINGVALUE). In addition, since P_i is honest, every honest party receives at most t messages (VOTE, v') for $v' \neq \text{KINGVALUE}$.

Note that no honest party has received $n - t$ messages (PROPOSE, v) for some value v : otherwise, P_i would have received at least the $n - 2t \geq t + 1$ messages (PROPOSE, v) sent by honest parties, which contradicts our assumption for this case. Hence, all honest parties reach line 12, and all honest parties have received the $t + 1$ messages (VOTE, KINGVALUE) sent by honest parties, and at most t messages (VOTE, v') for $v' \neq \text{KINGVALUE}$. This means that every honest party sets $\text{CURRENT} := \text{KINGVALUE}$. \square

We may now present the proof of Theorem 3.

Proof of Theorem 3. Termination follows from the protocol's construction.

For Agreement, note that at least one of the $t + 1$ kings of the $t + 1$ iterations is honest. Lemma 14 ensures that, in the first iteration i where the king P_i is honest, the honest parties obtain the same value CURRENT . Afterwards, Lemma 13 ensures that the honest parties do not change their CURRENT value in any further iteration, and therefore they output the same value CURRENT .

For Convex Validity, Lemma 6 guarantees that the honest parties enter the loop with valid values. Then, Lemma 11 implies that, in each of the iterations, each honest party holds a value CURRENT that is in some honest party's interval $[\text{INTERVALMIN}, \text{INTERVALMAX}]$, which is a subset of the honest inputs' range according to Lemma 11. Therefore, at the end of the $t + 1$ iterations, the honest parties output a valid value.

The round complexity follows from the protocol's construction: the setup stage consists of $O(1)$ communication rounds, and the search stage of $t + 1$ iterations, and each iteration consists of $O(1)$ communication rounds.

It remains to discuss the communication complexity. In the setup phase, each honest party sends three ℓ -bit values to all parties, with a communication cost of $O(\ell n^2)$ bits. Afterwards, in each iteration in the search stage, each honest party sends at most four ℓ -bit values to all parties, implying a communication cost of $O(\ell n^2)$ bits per iteration. Since there are $t + 1$ iterations, we obtain a communication cost of $O(\ell n^3)$ bits in total. \square

A.5 Protocol for \mathbb{Z} : missing proofs

We include the proof of Corollary 1, describing protocol $\Pi_{\mathbb{Z}}$.

Corollary 1. *Assume a BA protocol Π_{BA} resilient against $t < n/3$ corruptions. Then, if the honest parties hold inputs $(-1)^{\text{SIGN}_{\text{IN}}} \cdot v_{\text{IN}}^{\mathbb{N}} \in \mathbb{Z}$, such that $v_{\text{IN}}^{\mathbb{N}} \in \mathbb{N}$ with $|\text{BITS}(v_{\text{IN}}^{\mathbb{N}})| \leq \ell$, $\Pi_{\mathbb{Z}}$ is a CA protocol resilient against $t < n/3$ corruptions, with communication complexity $\text{BITS}_{\ell}(\Pi_{\mathbb{Z}}) = O(\ell n + \kappa \cdot n^2 \log^2 n) + O(\log n) \cdot \text{BITS}_{\kappa}(\Pi_{\text{BA}})$, and round complexity $\text{ROUNDS}_{\ell}(\Pi_{\mathbb{Z}}) = O(n) + O(\log n) \cdot \text{ROUNDS}_{\kappa}(\Pi_{\text{BA}})$.*

Proof. We first show that $\Pi_{\mathbb{Z}}$ achieves CA. In $\Pi_{\mathbb{Z}}$, parties first agree on their values' sign with the help of Π_{BA} . If Π_{BA} returns $\text{SIGN}_{\text{OUT}} = 0$, then there is an honest party holding a non-negative

input. If a party holds $v_{\text{IN}} < 0$, then $v_{\text{IN}}^{\mathbb{N}} := 0$ is a valid value. Parties then join $\Pi_{\mathbb{N}}$ with valid values $v_{\text{IN}}^{\mathbb{N}}$ and therefore agree on a valid output according to Theorem 5. Otherwise, if Π_{BA} returns $\text{SIGN}_{\text{OUT}} = 1$, there is an honest party holding a non-positive input. If a party holds $v_{\text{IN}} > 0$, then $v_{\text{IN}}^{\mathbb{N}} := 0$ is a valid value. Therefore, all honest parties hold valid values $(-1) \cdot v_{\text{IN}}^{\mathbb{N}}$. Parties then join $\Pi_{\mathbb{N}}$ with inputs $v_{\text{IN}}^{\mathbb{N}}$ and, according to Theorem 5, they agree on a value $v_{\text{OUT}}^{\mathbb{N}}$ such that $v_{\text{OUT}} := (-1) \cdot v_{\text{OUT}}^{\mathbb{N}}$ is valid.

$\Pi_{\mathbb{Z}}$ first runs Π_{BA} once with bits as inputs, and afterwards it runs $\Pi_{\mathbb{N}}$ on inputs of at most ℓ bits. Then, we obtain that $\text{BITS}_{\ell}(\Pi_{\mathbb{N}}) = \text{BITS}_1(\Pi_{\text{BA}}) + \text{BITS}_{\ell}(\Pi_{\mathbb{N}})$, and $\text{ROUNDS}_{\ell}(\Pi_{\mathbb{Z}}) = \text{ROUNDS}_1(\Pi_{\text{BA}}) + \text{ROUNDS}_{\ell}(\Pi_{\mathbb{N}})$. Theorem 5 leads to the results claimed in the corollary's statement. \square