

Distributed Computing with Channel Noise

Abhinav Aggarwal
abhiag@cs.unm.edu

Varsha Dani
varsha@cs.unm.edu

Thomas P. Hayes
hayes@cs.unm.edu

Jared Saia
saia@cs.unm.edu

Abstract

A group of n users want to run a distributed protocol π over a network where communication occurs via private point-to-point channels. Unfortunately, an adversary, who knows π , is able to maliciously flip bits on the channels. Can we efficiently simulate π in the presence of such an adversary?

We show that this is possible, even when L , the number of bits sent in π , and T , the number of bits flipped by the adversary are not known in advance. In particular, we show how to create a robust version of π that 1) fails with probability at most δ , for any $\delta > 0$; and 2) sends $\tilde{O}(L + T)$ bits, where the \tilde{O} notation hides a $\log(nL/\delta)$ term multiplying L .

Additionally, we show how to improve this result when the average message size α is not constant. In particular, we give an algorithm that sends $O(L(1 + (1/\alpha)\log(nL/\delta)) + T)$ bits. This algorithm is adaptive in that it does not require *a priori* knowledge of α . We note that if α is $\Omega(\log(nL/\delta))$, then this improved algorithm sends only $O(L + T)$ bits, and is therefore within a constant factor of optimal.

1 Introduction

Suppose we start with an n -party distributed computational protocol, to be executed over a network with point-to-point communication between adjacent parties. Now, imagine a nearly omniscient adversary, who can flip a subset of the bits sent along these communication channels. This adversary knows the protocol, as well as the players' inputs, and is unrestricted in terms of which bits she will flip. Can the players even hope to carry out their protocol successfully under conditions such as these?

Surprisingly, the answer is yes, in a surprisingly general sense! However, we will require some additional conditions. Firstly, the adversary cannot have an unlimited budget, or she could effectively cut off all communication forever by tossing a fair coin to decide whether or not to flip each bit. We thus restrict the adversary to flipping a total of T bits, where T is a finite number, which is unknown to our algorithm.

Secondly, we require that the channels are private and cannot be read by the adversary, and that players have access to private sources of randomness, which are also not readable by the adversary. This is a crucial assumption, but one for which there are strong justifications. If the players have access to shared secret randomness, i.e. a *one-time pad*, then the channels effectively become private, even if the adversary can see the traffic on them. More generally, if the players have access to *strong cryptography*, then the same conclusion holds in practice. Also, we may be interested in robustness

against an *oblivious adversary*, such as a bursty source of line noise. Such a source may not actually care about our protocol, but may still time its bursts in a way that happens to be bad for it. Finally, our assumption of private channels seems to be necessary, since without it, the adversary can institute a “man-in-the-middle” attack against any desired subset of the players. Since there is no *a priori* upper or lower bound on the number of bits, T , flipped by the adversary, such an attack can always fool the players into thinking they have successfully executed the protocol and $T = 0$, when in fact, all communication between the two subsets of the players was substituted by whatever the adversary wanted.

Our Problem and Result Consider a group of $n \geq 2$ users who want to run a noise-free asynchronous protocol π , whose length is unknown *a priori*. These users are connected via a network of arbitrary topology that is unknown to our algorithm. The edges of the network represent binary symmetric channels. Communication on these channels is *synchronous*.

An adversary is able to flip some finite number, T , of bits on any subset of the channels of the network at any time steps. The adversary chooses π , T , and which bits to flip on the channels. The adversary also knows our algorithm for transforming π to π' . However, similar to the previous work by Dani et. al [10, 11], we assume that the adversary neither knows the private random bits of any user, nor the bits sent over the channels, except when it is possible to infer these from knowledge of π and our algorithm.

Our results are summarized in the following theorem.

Theorem 1.1. *Let π be a (noise-free) asynchronous protocol for $n \geq 2$ users, and let $\delta > 0$. Then, Algorithm 3 compiles π into a robust protocol π' with the following properties. For any finite set of T (time,channel) pairs, with probability at least $1 - \delta$, when the adversary flips the bits on those channels at the corresponding times, the following hold:*

1. *The simulation succeeds. That is, each user stops after a finite number of time steps, and each ends in possession of his share of a valid transcript of π .*
2. *The total number of bits sent is $\mathcal{O}(L(1 + \frac{1}{\alpha} \log(\frac{nL}{\delta})) + T)$. Here L is the total number of bits and α is the average message length, for the particular transcript of π that was generated.*
3. *The total latency of π' is $\mathcal{O}\left(\max_p \left\{ \Lambda_p \left(1 + \frac{1}{\alpha} \log\left(\frac{n(L+T)}{\delta}\right) + T_p \right)\right\}\right)$, where p ranges over all communication paths in the asynchronous simulation of π , Λ_p is the latency of p , and T_p is the total number of bits flipped by the adversary on edges in p .*

We note the following. First, our algorithm π' always sends $\tilde{O}(L + T)$ bits, where the \tilde{O} notation hides a logarithmic term in n , L and δ . Second, when α is $\Omega(\log(nL/\delta))$, our algorithm sends only $O(L + T)$ bits, and is thus within a constant factor of optimal. Finally, we stress that our algorithm requires no *a priori* knowledge of T , L , α , or the network topology, but our algorithm does require knowledge of n , or at least a polynomial good estimate of n .

1.1 Related Work

Interactive Communication Our work is related to interactive communication. The problem of interactive communication asks how two parties can run a protocol π over a noisy channel. This

problem was first posed by Schulman [27, 28], who describes a deterministic method for simulating interactive protocols on noisy channels with only a constant-factor increase in the total communication complexity. This initial work spurred vigorous interest in the area (see [3] for an excellent survey).

Schulman’s scheme tolerates an adversarial noise rate of $1/240$, even if the adversary is *not* oblivious. It critically depends on the notion of a *tree code* for which an exponential-time construction was originally provided. This exponential construction time motivated work on more efficient constructions [4, 25, 22]. There were also efforts to create alternative codes [14, 23]. Recently, elegant computationally-efficient schemes that tolerate a constant adversarial noise rate have been demonstrated [1, 15]. Additionally, a large number of results have improved the tolerable adversarial noise rate [2, 6, 16, 13, 5], as well as tuning the communication costs to a known, but not necessarily constant, adversarial noise rate [17].

Our paper builds on a result on interactive communication by Dani et al [11], which in contrast to previous work, assumes private channels, but tolerates an unknown number of bit flips by the adversary. Their algorithm sends $L + O\left(\sqrt{L(T+1)\log L} + T\right)$ bits in expectation. They show that private channels are necessary in order to tolerate unknown T .

Recently, Braverman et al [7] show a strong lower bound for coding schemes for multiparty interactive communication when π is a synchronous protocol running on a star network, even when noise is just stochastic. Our result circumvents this lower bound since we require that π be an asynchronous protocol.

Recent work by Censor-Hillel, Gelles and Haeupler [8], like ours, shows how to make an arbitrary asynchronous distributed protocol robust against adversarial noise. In their case, the new protocol will succeed if up to $\Theta(1/n)$ of the messages are corrupted, and the overhead is a factor of $O(n \log^2 n)$ in the total communication. In this setting, the fraction $1/n$ of corrupted messages is optimal, because the adversary can simply cut off all communication to a particular player. Notably, their algorithm proceeds by essentially reducing to the case where the communication network is a tree.

Our work is not directly comparable to that of [8], because we do not quantify our results in terms of the fraction of messages that are corrupted, and more importantly, we assume that communication takes place over private channels. In our case, the stronger assumption of private channels means that up to a $1/\log(nL)$ fraction of message bits can be corrupted, and our algorithm will still succeed with high probability, with a cost overhead that is only $\log(nL)$. In our setting, the adversarial strategy of trying to cut off all communications to and from a single player requires corruption of much more than a $1/n$ fraction of the bits, because our protocol can detect the noise, and consequently increases the fraction of the total communication involving the beleaguered player.

Rateless Codes Rateless error correcting codes enable generation of potentially an infinite number of encoding symbols from a given set of source symbols with the property that given any subset of a sufficient number of encoding symbols, the original source symbols can be recovered. Fountain codes [21] and LT codes [24, 19, 18] are two classic examples of rateless codes. Erasure codes employ

feedback for stopping transmission [24, 19] and for error detection [18] at the receiver. Critically, the feedback channel is typically assumed to be noise free. We differ from this model in that we allow the adversary to flip bits on the feedback channel. Additionally, we tolerate bit flips, while most rateless codes tolerate only bit erasures.

1.2 Formal Model

Asynchrony of π Recall that a synchronous network is one in which there is a global clock and each message sent in the network takes one time step to be transmitted. On the other hand, an asynchronous network is one in which messages may be arbitrarily and adversarially delayed. In particular messages need not be received in the same order as they are sent. However, every message must be eventually delivered in a finite amount of time.

As an example, consider a synchronous leader election algorithm in which the users start by waiting for their ID number of clock ticks. At the end of this, if a user finds no other proclamations of a leader, he proclaims himself as the leader. This protocol cannot be simulated in our model. To see this, suppose processor 19 is trying to decide if he is the leader. If he hears meaningless noise on the channel for multiple time steps of π' , then he does not know whether the adversary is simply inserting this noise on a silent channel, or whether there is a processor 18 whose leadership proclamation is getting corrupted. Even if he does not hear noise, i.e. if he hears silence all throughout, then he knows that nobody has transmitted anything to him, but he does not know how many time-steps of π have passed in the view of the putative processor 18.

Synchrony of π' We assume that our protocol π' runs in a *synchronous* network. In particular, π' critically relies on the fact that all processors have synchronized clocks and that communication time is fixed and known. We define a *time step* as the amount of time that it takes to send one bit over a channel, assuming it to be the same for all the channels. As is standard in distributed computing, we assume that all local computation is instantaneous.

Silence on the channel Similar to [11], when neither user sends on a channel, we say that the channel is silent. In any contiguous sequence of silent channel steps, the bit received on the channel in the first step is set by the adversary for free. By default, the bit received in the subsequent steps of the sequence remains the same, unless the adversary pays for one bit flip each time it wants to change the value of the bit received in any contiguous sequence of silent steps.

Additional Assumptions We assume that all users know the desired error tolerance δ for π' , as well as the number of users in the network, n . However, the users do not know the values of L and T , and they have no knowledge of the underlying topology of the network connecting the users. We further assume that π has a mechanism for users to detect the start and end of messages they receive. We emphasize that this mechanism need not necessarily involve using end-of-word symbols. Techniques such as prefix-free codes that do not use any special symbols provide a good alternative.

1.3 Overview of Our Result

In simulation of π using a synchronous protocol π' , we assume a global clock which *ticks* every time step, and allows all the users to determine how long messages they should send at any time

step $t \geq 1$, which is only a function of n, δ and t . Once all messages of π have been exchanged in π' , we are sure that the users have terminated and the protocol has ended. Our approach makes this possible without the use of any synchronization messages across the network, yet keeping all communication as local to the users as possible.

A key technical challenge here is termination. The length of the simulated protocol is unknown, so the users will likely not terminate at the same time. After a user has terminated, it is a challenge for its neighbors to detect this fact based on the bits received over the noisy channel. Our approach makes use of authentication keys, which every user generates randomly for himself and shares with whoever wants to communicate with him. We use this key in a way that it makes unlikely for the adversary to successfully forge messages or silence on the channels, without detection by the affected users. Hence, a sequence of silent channel steps will suffice, with desired probability, to indicate the users of the termination of their neighbors.

Our main result is stated in Theorem 1.1. The proof of this theorem and details of how we deal with the challenges above are presented in the subsequent sections.

1.4 Rate

In most prior work it is assumed that the noise rate of the channel(s) is known and may be used as a parameter in the design of the algorithm. The parameter of interest then is the rate of the designed code. In our work we do not assume that the noise rate is known in advance, but only require that the adversary flip a finite number of bits. Nevertheless, in order to compare our result with other work, we compute the coding rate of our algorithm as a function of the *a posteriori* noise rate. Such a comparison is only meaningful when the adversary's total budget is less than L , the length of π , so for this section we will assume that $T < L$.

Let L' denote the length of π' . Theorem 1.1 states that our algorithm achieves

$$L' \leq CL \left(1 + \frac{1}{\alpha} \log \left(\frac{nL}{\delta} \right) \right) + CT$$

for some constant C , where δ is the permissible failure probability for the algorithm and α is the average message length in π . Furthermore, since $T < L$ this translates to the absolute upper bound

$$L' \leq CL \left(2 + \frac{1}{\alpha} \log \left(\frac{n(L+T)}{\delta} \right) \right) \tag{1}$$

$$\leq CL \left(2 + \frac{1}{\alpha} \log \left(\frac{2nL}{\delta} \right) \right) \tag{2}$$

Let $\varepsilon = T/L'$ be the *a posteriori* noise rate. Then $T = \varepsilon L'$. Making this substitution for T in (1), and using (2) to bound the L' inside the log, we have

$$L' \leq CL \left(2 + \frac{1}{\alpha} \log \left(\frac{n \left(L + \varepsilon CL \left(2 + \frac{1}{\alpha} \log \left(\frac{2nL}{\delta} \right) \right) \right)}{\delta} \right) \right)$$

Dividing by L ,

$$\begin{aligned}
\frac{L'}{L} &\leq 2C + \frac{C}{\alpha} \log \left(\frac{n(L + \varepsilon CL(2 + \frac{1}{\alpha} \log(\frac{2nL}{\delta})))}{\delta} \right) \\
&= 2C + \frac{C}{\alpha} \left(\log(nL/\delta) + \log \left(1 + \varepsilon \left(2C + \frac{C \log(2nL/\delta)}{\alpha} \right) \right) \right) \\
&\leq 2C + \frac{C \log(nL/\delta)}{\alpha} + \frac{\varepsilon C}{\alpha} \left(2C + \frac{C \log(2nL/\delta)}{\alpha} \right) \\
&\leq \left(2C + \frac{C \log(2nL/\delta)}{\alpha} \right) \left(1 + \frac{\varepsilon C}{\alpha} \right)
\end{aligned}$$

To summarize, for the worst case where $\alpha = 1$, the above shows that we achieve a coding rate that increases linearly with the noise rate ε and with the logarithm of nL/δ . In particular, the coding rate is $O((1+\varepsilon) \log(nL/\delta))$. For arbitrary α , we achieve a coding rate of $O((1+\varepsilon/\alpha) \log(nL/\delta)/\alpha)$.

1.5 Paper Organization

The rest of this paper is organized as follows. In section 2, we describe our main algorithm for the case where all messages in π are exactly 1 bit in length. We prove this algorithm is correct and analyze resource costs in Section 3. In Sections 4 and 5, we describe and analyze our algorithms for the case where messages in π are of an arbitrary size average (thereby achieving better results for some values of the average size α). Finally, we conclude and give directions for future work in Section 6.

2 Our Algorithm

We describe our algorithm in this section. For now, we assume that π consists of one-bit messages.

2.1 Notation and Definitions

Some helper functions and notation used in our algorithm are described here.

Algebraic Manipulation Detection Codes Our algorithm makes critical use of Algebraic Manipulation Detection (AMD) codes from [9]. For a given $\eta > 0$, called the strength of AMD encoding, these codes provide three functions: `amdEnc`, `amdDec` and `IsCodeword`. The function `amdEnc` (m, η) creates an AMD encoding of a message m . The function `IsCodeword` (m, η) takes a message m and returns true if and only if there exists some message m' such that `amdEnc` (m', η) = m . The function `amdDec` (m, η) takes a message m such that `IsCodeword` (m, η) and returns a message m' such that `amdEnc` (m', η) = m . We summarize the results from [9] to highlight the important properties of these functions in the following lemma.

Theorem 2.1. *There exist functions `amdEnc`, `amdDec` and `IsCodeword`, such that for any $\eta \in (0, \frac{1}{2}]$ and any bit string m of length x :*

1. `amdEnc` (m, η) is a string of length $x + 2 \log \left(\frac{1}{\eta} \right)$.

2. $\text{IsCodeword}(\text{amdEnc}(m, \eta), \eta)$ and $\text{amdDec}(\text{amdEnc}(m, \eta), \eta) = m$.
3. For any bit string $s \neq 0$ of length x , we have $\Pr(\text{IsCodeword}(\text{amdEnc}(m, \eta) \oplus s, \eta)) \leq \eta$.

Error-correcting Codes These codes enable us to encode a message so that it can be recovered even if the adversary corrupts a third of the bits. We will denote the encoding and decoding functions by ecEnc and ecDec , respectively. The following theorem, established by the results in [26], gives the properties of these functions.

Theorem 2.2. [26] *There exists a constant $C_e > 0$ such that for any message m , we have $|\text{ecEnc}(m)| \leq C_e|m|$. Moreover, if m' differs from $\text{ecEnc}(m)$ in at most one third of its bits, then $\text{ecDec}(m') = m$.*

With respect to this constant C_e , we define two more constants $C_1 = 12C_e + 76$ and $C_2 = 32C_e + 115$, which will be used in our algorithm.

Finally, we observe that the linearity of ecEnc and ecDec ensure that when the error correction is composed with the AMD code, the resulting code has the following properties:

1. If at most a third of the bits of the message are flipped, then the original message can be uniquely reconstructed by rounding to the nearest codeword in the range of ecEnc .
2. Even if an arbitrary set of bits is flipped, the probability of the change not being recognized is at most η , i.e. the same guarantee as for the plain AMD codes.

This is because the error-correcting code is linear, so when noise η is added by the adversary to the codeword x , effectively what happens is that the decoding function rounds the noise to the nearest codeword. Thus $\text{ecDec}(x + \eta) = \text{ecDec}(x) + \text{ecDec}(\eta) = m + \text{ecDec}(\eta)$, where m is the AMD-encoded message. But now $\text{ecDec}(\eta)$ is an obviously selected string added to the AMD-encoded codeword, and hence the result is very unlikely to be a valid message unless $\text{ecDec}(\eta) = 0$.

For a string s , we use the notation $s[i]$ to denote the i^{th} bit of s and $s[i, j]$ to denote the substring $(s[i], s[i + 1], \dots, s[j - 1])$. We let $|s|$ denote the length of string s , and use the conventions that $s = (s[0], s[1], \dots, s[|s| - 1])$, and for $j > |s|$, $s[i, j] = s[i, |s|]$.

Silence We define the function $\text{IsSilence}(s)$ to return true iff the string s has fewer than $|s|/3$ bit alternations. We also define $\mathcal{S}_\ell = \{s \in \{0, 1\}^\ell \mid \text{IsSilence}(s)\}$. We drop the subscript when ℓ is clear from the context.

2.2 Algorithm Overview

For our algorithm, we assume that each pair of neighboring users communicates over a dedicated channel. The algorithm proceeds in rounds, each of which consists of the following steps.

1. If u has a message for v , he initiates a message exchange by asking v for a key.
2. Upon receipt of this key, u sends the message along with the key.

3. v terminates the message exchange upon successful authentication and retrieval of the message.
4. u terminates the message exchange upon hearing silence from v .

This goes on until all the messages in π have been communicated to the intended recipients.

Rounds

For each message m in π that needs to be sent from some user u to his neighbor v , we communicate m through a sequence of exchanges between users u , referred to as *Alice*, and v , referred to as *Bob*, in π' using Algorithms 1 and 2, respectively. As mentioned above, the sequence of time steps corresponding to steps 1-4 of the algorithm overview constitute what we call a *round*. Thus, each round of π consists of exactly four *words*, one for each of the steps 1-4. The length of each word in round r is denoted w_r . This depends on the round number, r and is therefore a function of the time step, that can be computed independently by each user using the clock. We note that w_r gradually increases with r .

We first assume that π has single bit messages. We will extend our results to arbitrary message lengths in Section 4.

Now, since each message in π is just a single bit, it takes exactly one round to be communicated in π' , if no successful corruption happens, and more otherwise. If a round for some message m is corrupted, we attempt resending m in the subsequent round, possibly with the security increased due to the potentially increased word length. For technical reasons, Alice needs to distinguish between successive messages from π . This is because Alice and Bob may have different views on whether a particular round was successful and there may be times when Alice is resending a message that Bob has already received. If Bob encounters two progressive rounds with no silent round in between, which contain the same message, he needs to distinguish between whether Alice is resending the message or whether Alice's next message happens to be the same bits as the previous one. To disambiguate these cases, Alice appends a bit b to each message m where b is the parity of index of message m . We will denote the pair (m, b) by M .

For convenience, we assume that there are two bidirectional channels between each pair of (neighboring) users, one for each user to initiate a round with the other. Note that, by time-slicing, we could achieve the same effect on a single bidirectional channel between each pair of users, at the cost of a factor 2 increase in the number of time steps.

Both Alice and Bob generate their words for round r using a function $\mathcal{E}_r(x, k)$, described below, which returns the encoding of the word's content x using the key k based on the security settings for round r . Here x may be the message m from π , a special keyword **KEY?** used by Alice to request Bob's key, or Bob's key for the round. The key k is a string of length $2 \left\lceil \log \frac{4n\pi i}{\sqrt{\delta}} \right\rceil$ bits. In Alice's first call to \mathcal{E} , k is a random string. In Bob's call, k is the key he received from Alice in the previous word, but x is a random string of the appropriate length. In Alice's second call, k is the key she received from Bob in the previous word. Alice and Bob generate fresh random keys for each round in which they desire to send a message.

The security settings are arranged so that the word length in round r is given by: The word length in round r grows logarithmically, and is given by the following formula.:

$$w_r = 300 \lceil \log(nr/\delta) \rceil. \quad (3)$$

The function $\mathcal{E}_r(x, k)$ is formally defined as follows. First, it encodes the pair (x, k) using an AMD code with $\eta_r = \frac{\delta}{2n^2\pi^2r^2}$. The result is then encoded with a (1/3)-error-correcting code, such as a Reed-Solomon code. Finally, enough uniformly random bits are appended to bring the total word length up to w_r . These final random bits are added to ensure that, even if corrupted, any deliberately sent word is very unlikely to be mistaken for silence.

Similar to $\mathcal{E}_r(x, k)$, we define a function $\mathcal{D}_r(m)$ which returns either a pair (x', k') such that $\mathcal{E}_r(x', k') = m$ or returns (\perp, \perp) , if no such (x', k') exists. It begins by stripping off the padding at the end of m to obtain a shorter string m' . Then it decodes the error correction, computing $m'' = \text{ecDec}(m')$. If $\text{IsCodeword}(m'', \eta_r)$, then \mathcal{D} outputs $\text{amdDec}(m'', \eta)$, otherwise it returns (\perp, \perp) .

The flow of information and control in a round is illustrated in the flowchart in Fig. 1, and specified in detail in Algorithms 1 and 2.

At the beginning of each round, Bob must listen for a key request during the first w_r time steps. If no valid request is received, he idles until the start of the next round. Thus Bob is active in every round. For her part, Alice only participates in a round if, in π , she has a message to send to Bob.

There are certain events which may cause our algorithm to fail. More specifically, it is possible that the adversary (1) converts one AMD codeword to another, so that the decoded content is different from the content intended; (2) converts a non-silence word into silence; and (3) correctly guesses some user's key and uses it to communicate bits that are not in π . We will discuss these failure events in detail in Section 2.3 and analyze their probabilities of occurrence in Section 3.1. However, for the remainder of this section, we assume that none of these failure events happen.

In what follows we will describe various scenarios for what may happen in a round during the execution of the algorithm, for a particular bidirectional channel. There are n^2 such channels, and the same round may have different scenarios enacted on it on different channels. Moreover, the views of Alice and Bob may differ on which scenario was enacted.

Silent Rounds

Bob listens to the first word on the channel and hears silence. Since the adversary cannot manufacture silence, Alice had no message for him and nothing further happens in the round. In this case both views agree that the round was silent.

Progressive Rounds

These are rounds in which the number of bits flipped by the adversary is small enough that it is handled by the correction schemes. Such rounds proceed as follows.

Alice has a message m for Bob. She requests his key using the resend bit and the keyword `KEY?`. Bob decodes Alice's key request and obtains her key. He generates his own random key and sends it to Alice using his knowledge of her key to authenticate his message as actually coming from him. Alice in turn decodes Bob's communication, obtains his key and sends m , using Bob's key to authenticate her communication. Bob correctly receives m , records it in his transcript of π' and remains silent. Alice hears silence on the line and decides that Bob has successfully received her message. Here both Alice and Bob agree that the round was successful.

Corrupted Rounds

These are rounds when the adversary is active, and corrupts one or more of the words in the round.

Case 1 Alice is silent, but the adversary sends Bob a key request. Then Bob sends Alice his key, but she is not listening and remains silent. At this point the adversary can say whatever he wants on the channel. However since he does not know Bob's key, he cannot authenticate his message. So Bob receives an invalid communication, and responds with noise, but again Alice is not listening. Thus such a round is corrupted in Bob's view, but silent in Alice's. Note that Bob may realize that the round was silent in Alice's view at a later stage, but we still account for this as a corrupted round, since Bob has already incurred a cost for the corruption.

Case 2 The adversary corrupts Alice's key request. Then Bob does not receive a valid key request, and hence remains silent the rest of the round. Since Bob is silent, the adversary may say whatever he wants on the channel. In particular he may try to pretend he is Bob and send Alice a bogus key. However since he cannot guess Alice's key, he cannot authenticate himself as Bob. So Alice receives an invalid communication, and stays silent until the end of the round.

Case 3 Bob receives Alice's key request, but the adversary corrupts his communication containing his key. Since he cannot corrupt one AMD codeword into another, Alice cannot decode the message into a key for Bob, causing her to remain silent for the remainder of the round. Again, the adversary cannot install a bogus message from Alice because he cannot guess Bob's key, so Bob receives an invalid communication.

Case 4 The adversary is inactive for the first half of the round and Alice receives Bob's key. Then the adversary corrupts her communication of m . Then Bob receives garbage at his end and injects noise into the channel. Since the adversary cannot convert this noise into silence, Alice knows that the round has failed.

In cases 2, 3 and 4, Alice and Bob both know that the round has been corrupted. Alice will retry sending her message in the next round.

Case 5 The round succeeds all the way to the point where Bob correctly receives m , decides the round is successful and remains silent. Then the adversary injects noise onto the channel causing Alice to think the round has failed. In this case again, Alice and Bob's views differ. However this is not a problem because Alice will simply resend the message, and Bob will receive it again, in the next round that is successful in his view. However Bob can recognize the message as a repetition because of the parity bit. At that stage, of course, Bob realizes that the round was in fact, corrupted, but this makes no difference to his future actions.

2.3 Failure Events

As mentioned before, there are certain events can cause catastrophic failure from which the algorithm cannot recover.

1. *Failure of AMD codes* : The adversary's bit flips happen to convert an AMD codeword into another valid AMD codeword. In this case the decoded content differs from the intended content resulting in authentication failure, when the content was Bob's key, or incorrect simulation of π when the content was Alice's message.
2. *Conversion to silence* : The adversary's bit flips are such that the resulting word looks like silence to its recipient. If the noise sent by Bob in line 19 to request a resend is converted to silence, then Alice incorrectly assumes that Bob has received her message and stops transmitting it. This results in an incorrect transcript of π . Other words being converted to silence could result in a player being silent on the following word, which in conjunction with guessing the key (see below) could result in the adversary being able to say whatever he wants on the channel.
3. *Guessing the key* : On a round when Alice is silent in π , the adversary can install a bogus key request on her channel to Bob. Ordinarily this is not a problem because when Bob responds with his key, the adversary cannot read it, and therefore cannot send Bob an authenticated message. However, if he happens to guess Bob's key, then he can send Bob a message purporting to be from Alice, resulting in an incorrect simulation of π .

We will show in Section 3.1 that the probability of such a catastrophic failure over the entire run of the algorithm is at most δ .




Figure 1: Flowchart for Alice and Bob during each round. Here w denotes the word length in the current round.

Algorithm 1 Message exchange algorithm for the sender. **Send-Message** is only called at the beginning of a round.

```
1: procedure SEND-MESSAGE( $m$ )
2:    $\triangleright b$  is a persistent variable for the parity bit. On the first call to SEND-MESSAGE,  $b$  is set
3:    $\triangleright$  to 0. On subsequent calls it is whatever it was set to on the previous call.
4:   while true do
5:     Generate random key  $k_A$  of length  $\kappa_r$   $\triangleright r$  is the current round number..
6:     Send  $\mathcal{E}_r(\text{KEY?}, k_A)$ .
7:      $M_1 \leftarrow w_r$  bits on the channel from the receiver.
8:     if IsSilence( $M_1$ ) then  $\triangleright$  assume the receiver has already terminated.
9:       Stay silent for  $2w_r$  time steps and return
10:    else
11:       $(x, k) \leftarrow \mathcal{D}_t(M_1)$ 
12:      if  $k \neq k_A$  then
13:        Stay silent for  $2w_r$  time steps.
14:      else
15:         $k_B \leftarrow x$ 
16:        Send  $\mathcal{E}_t((m, b), k_B)$ .
17:         $M_2 \leftarrow w_r$  bits on the channel from the receiver.
18:        if IsSilence( $M_2$ ) then
19:           $b \leftarrow \neg b$  and return
```

Algorithm 2 Message exchange algorithm for the receiver. Receive-Message is only called at the beginning of a round.

```

1: procedure RECEIVE-MESSAGE( )
2:   ▷  $\hat{b}$  is a persistent variable for the parity bit. On the first call to RECEIVE-MESSAGE,  $\hat{b}$  is
3:   ▷ set to 0. On subsequent calls it is whatever it was set to on the previous call.
4:    $M'_1 \leftarrow w_r$  bits on the channel from the sender.
5:   if IsSilence ( $M'_1$ ) then
6:     Stay silent for  $3w_r$  time steps.                                ▷  $r$  is the current round number.
7:   else
8:      $(x', k') \leftarrow \mathcal{D}_t(M'_1)$ 
9:     if  $x' \neq \text{KEY?}$  then
10:      Send noise for  $w_r$  time steps.
11:      Stay silent for  $2w_r$  time steps.
12:     else
13:        $k_A \leftarrow k'$ .
14:       Generate random key  $k_B$  of length  $\kappa_t$ .
15:       Send  $\mathcal{E}_t(k_B, k_A)$ .
16:        $M'_2 \leftarrow w_r$  bits on the channel from the sender.
17:        $(x'', k'') \leftarrow \mathcal{D}_t(M'_2)$ 
18:       if  $k'' \neq k_B$  then
19:         Send noise for  $w_r$  time steps.
20:       else
21:          $(m', b') \leftarrow x''$ 
22:         if  $b' \neq \hat{b}$  different from last message then
23:           Set  $\hat{b} \leftarrow b'$  and record the message  $m'$  from the sender.
24:         Stay silent for  $w_r$  time steps.

```

2.4 The protocol π'

Our main protocol π' is presented as Algorithm 3. We make use of the definition of an I/O automaton (see [20] Chapter 14.1.1) to represent π . We assume that for each user u in the network, π provides an I/O automaton π_u with the following properties.

- π_u has a single *initial* state
- π_u has some subset of states that are *termination* states. Each termination state may have a value for u to *output*.
- There is a set of *transition relations*, each from one state to another state, each labeled with an *action*, where this action may be either an *input* action (e.g. receiving a message) or an *output* action (e.g. sending a message). These transitions satisfy the property for every state s , for every possible input action, a , there is a transition from s to some other state that is labelled with a .

Algorithm 3 Protocol π' , run at each node u

```
1:  $s \leftarrow$  initial state of  $\pi_u$ 
2: while  $s$  is not a termination state do
3:   if there is a transition relation from state  $s$  to state  $s'$  in  $\pi_u$  labeled with an output action
   to send message  $m$  to neighbor  $v$  then
4:     Schedule SEND-MESSAGE( $m$ ) to run on the channel to  $v$ .
     If no call to SEND-MESSAGE is currently running on the channel, begin immediately.
     Otherwise, set up the call to begin as soon as all currently scheduled SEND-MESSAGE
     calls on this channel have finished running.
5:     Transition to state  $s'$  in  $\pi_u$ ;  $s \leftarrow s'$ .
6:   else
7:     repeat
8:       for each neighbor  $v$ , in parallel run RECEIVE-MESSAGE() on the channel from  $v$ .
9:     until the RECEIVE-MESSAGE calls have recorded some non-empty set  $S$  of messages
10:    for each message  $m$  in  $S$  do
11:      Let  $s'$  be the target state on the transition relation from  $s$  with input action  $m$ .
12:      Transition to state  $s'$  in  $\pi_u$ ;  $s \leftarrow s'$ .
13: Set processor  $u$ 's output based on the termination state.
14: Continue executing any remaining scheduled SEND-MESSAGE calls until all have returned.
```

3 Analysis

We now analyze the main algorithm as presented in Section 2. We begin by computing the failure probability for the algorithm, by considering the three bad events as before, and then take a union bound over all the rounds. We will then prove that the algorithm is correct and terminates in finitely many time steps. Finally, we compute an upper bound on the expected number of bits the algorithm sends.

3.1 Probability of Failure

We define three bad events for a round.

1. AMD Failure : The adversary is able to flip the bits to convert the message into another valid word. In this case, the affected users end their rounds either with authentication failures or with knowledge of bits which are not in π .
2. Conversion to Silence : The adversary is able to flip bits in such a way that some user's random bits look like silence to his neighbor, resulting in the latter ending his round without the knowledge of this failure.
3. Key Installation : The adversary installs a correct key from Bob when Alice is silent (to successfully simulate line 16 of Algorithm 1) or installs a correct key from Alice when Bob is silent (to successfully simulate line 15 of Algorithm 2).

We now bound the probabilities of each of these failure events to eventually prove that our algorithm succeeds with probability $1 - \delta$.

Lemma 3.1. *In π' , AMD Failure occurs with probability at most $\delta/3$.*

Proof. Recall that in any round, four words are exchanged between any pair of users, and that at most $\binom{n}{2} < n^2$ pairs of users may be exchanging words with each other. Furthermore, in a given round r , the AMD failure probability for a single word is set to be at most $\eta_r = (n\pi r)^{-2} \delta/2$, as discussed in Section 2.2. Now, for $1 \leq k \leq 4, 1 \leq i \neq j \leq n$ and $r \geq 1$, define $\xi_{i,j,k,r}$ to be the event that in round r , AMD failure occurs in the k^{th} word exchanged between users i and j . Then, $\Pr(\xi_{i,j,k,r}) \leq (n\pi r)^{-2} \delta/2$. Hence, by a union bound, we get

$$\Pr\left(\bigcup_{i,j,k,r} \xi_{i,j,k,r}\right) \leq \sum_{k=1}^4 \sum_{\substack{i,j=1 \\ i \neq j}}^n \sum_{r \geq 1} \frac{\delta}{2(n\pi r)^2} \leq \frac{2\delta}{\pi^2} \sum_{r \geq 1} \frac{1}{r^2} = \frac{\delta}{3}.$$

□

Lemma 3.2. *For $b \geq 95$, the probability that a b -bit string sampled uniformly at random from $\{0,1\}^b$ has fewer than $b/3$ bit alternations is at most $e^{-b/19}$.*

Proof. Let s be a string sampled uniformly at random from $\{0,1\}^b$, where $b \geq 95$. Denote by $s[i]$ the i^{th} bit of s . Let X_i be the indicator random variable for the event that $s[i] \neq s[i+1]$, for $1 \leq i < b$. Note that all X_i 's are mutually independent. Let X be the number of bit alternations in s . Clearly, $X = \sum_{i=1}^{b-1} X_i$, which gives $\mathbb{E}(X) = \sum_{i=1}^{b-1} \mathbb{E}(X_i)$, using the linearity of expectation. Since $\mathbb{E}(X_i) = 1/2$ for all $1 \leq i < b$, we get $\mathbb{E}(X) = (b-1)/2$. Using a multiplicative version of Chernoff's bound (see [12]), we have that for $0 \leq t \leq \sqrt{b-1}$,

$$\Pr\left(X < \frac{b-1}{2} - \frac{t\sqrt{b-1}}{2}\right) \leq e^{-t^2/2}.$$

To obtain $\Pr(X < b/3)$, we set $t = \frac{b-3}{3\sqrt{b-1}}$ to get

$$\Pr(X < b/3) \leq e^{-\frac{(b-3)^2}{18(b-1)}} \leq e^{-b/19} \quad \text{for } b \geq 95,$$

where the condition $b \geq 95$ comes from rounding the solution to a quadratic equation. □

Lemma 3.3. *In π' , Conversion to Silence occurs with probability at most $\delta/3$.*

Proof. Recall that in round r , each encoded message includes $38 \left\lceil \log \left(\frac{2n\pi r}{\sqrt{\delta}} \right) \right\rceil$ random bits at the end. Then, Lemma 3.2 tells us that the probability that adversary is able flip these random bits of a word to forge silence is at most $2^{-\frac{38 \left\lceil \log \left(\frac{2n\pi r}{\sqrt{\delta}} \right) \right\rceil}{19}} \leq (n\pi r)^{-2} \delta/4$. Thus, similar to Lemma 3.1, for $1 \leq k \leq 4, 1 \leq i \neq j \leq n$ and $r \geq 1$, define $\xi_{i,j,k,r}$ to be the event that in round r , the k^{th} word exchanged between users i and j is converted to silence, so that $\Pr(\xi_{i,j,k,r}) \leq (n\pi r)^{-2} \delta/4$. Hence, by using a similar union bound as Lemma 3.1, we get the desired bound. □

Lemma 3.4. *In π' , the adversary is able to guess the key of some user with probability at most $\delta/3$.*

Proof. Recall that in round r , keys are of length $2 \left\lceil \log \left(\frac{4n\pi r}{\sqrt{\delta}} \right) \right\rceil$. Each such key is generated uniformly at random from the set of all binary strings of this length. Thus, the probability of guessing this key is at most $2^{-2 \left\lceil \log \left(\frac{4n\pi r}{\sqrt{\delta}} \right) \right\rceil} \leq (n\pi r)^{-2} \delta/16$. Again, similar to Lemma 3.1, for $1 \leq k \leq 4, 1 \leq i \neq j \leq n$ and $r \geq 1$, define $\xi_{i,j,k,r}$ to be the event that in round r , the adversary is able to guess the key in the k^{th} word exchanged between users i and j , so that $\Pr(\xi_{i,j,k,r}) \leq (n\pi r)^{-2} \delta/16$. Hence, by using a similar union bound as Lemma 3.1, we get the desired bound. \square

Lemma 3.5. *With probability at least $1 - \delta$, none of the failure events happen during a run of π' .*

Proof. A run of protocol π' fails if any of the three failure events described above happen. From Lemmas 3.1, 3.3 and 3.4, the total probability of failure is computed by using union bound over the three failure events, which gives $\delta/3 + \delta/3 + \delta/3 = \delta$. Hence, the run succeeds with probability at least $1 - \delta$. \square

3.2 Correctness

All lemmas in this section assume that none of the failure events occur. Without loss of generality, we also assume that Alice and Bob generate their keys in every round. We use the phrase *terminated in π* to mean finished executing line 13 of Algorithm 3.

Lemma 3.6. *Fix a round r of π' . Let W be any of the four sequences of w_r bits on the channel from Alice to Bob or the channel from Bob to Alice in this round and $k_A^{(r)}, k_B^{(r)}$ be the keys that Alice and Bob generate in this round. Then exactly one of the following hold.*

1. *IsSilence(W) is true, in which case the sender on the channel was silent as well.*
2. *$\mathcal{D}_t(W) = \mathcal{C} \neq \perp$, where t is the time step at which the round r began. Then,*
 - (a) *If $\mathcal{C} = (\text{KEY?}, k)$, then either Alice sent W with key $k = k_A^{(r)}$, or Alice is silent and the adversary sent W .*
 - (b) *If $\mathcal{C} = (k', k_A^{(r)})$, then Bob has not terminated in π yet, and he sent W with key $k' = k_B^{(r)}$.*
 - (c) *If $\mathcal{C} = (x, k_B^{(r)})$, then Alice sent W with content $x = m^{(r)}$, where $m^{(r)}$ is Alice's message for Bob in round r .*
 - (d) *Otherwise, the adversary sent W on a silent channel.*
3. *Bob executed line 19 of Algorithm 2.*
4. *W is the outcome of adversarial tampering on the channel.*

Proof. In any given round r , whenever Alice or Bob send a word to the other, this word does not convert to silence because we assume that the failure event *Conversion to Silence* does not occur. Hence, if silence is received on the channel, then the sender on the channel must have been silent at that time. This proves part (1) of our lemma.

If, however, a valid AMD codeword is received, which decodes into \mathcal{C} , the following cases are possible. Case 2(a) : If $\mathcal{C} = (\text{KEY?}, k)$, then either Alice issued this key request to Bob with her key $k_A^{(r)}$

(line 6 of Algorithm 1), or she was silent. In the former case, the adversary is unable to put a key $k \neq k_A^{(r)}$ in the codeword since the failure event *AMD Failure* does not occur, and hence, $k = k_A^{(r)}$. In the latter case, since the channel is silent, the adversary must have issued the key request with some key k . Case 2(b) : If $\mathcal{C} = (k', k_A^{(r)})$, then since we assume that the failure event *Key Installation* does not occur, it must be the case that Bob sent W and not the adversary. Thus, Bob must have included his own key $k_B^{(r)}$ to the codeword along with a copy of Alice's key (line 15 of Algorithm 2). This is only possible when Bob has not *terminated in π* yet. Case 2(c) : If $\mathcal{C} = (x, k_B^{(r)})$, then Alice must have sent W since the the failure event *Key Installation* does not occur. Also, since we assume that the failure event *AMD Failure* does not occur, the adversary would have not been able to convert $m^{(r)}$ into another message successfully, and hence, W must contain the message that Alice has for Bob in this round. Case 2(d) : If the AMD codeword is neither of the above three cases, then it must be the case that the sender was on the channel at the time W was on the channel, since *AMD Failure* does not occur. Hence, the adversary must have sent W on a silent channel.

If neither silence nor an AMD codeword is received, then W is noise. The only step in π' where noise is intentionally put on the channel is when Bob has to inform Alice that he did not receive her message correctly (line 19 of Algorithm 2). Thus, if Bob sent this noise, Case 3 of our lemma holds, else Case 4 must hold where the adversary has tampered with the bits on the channel so that W becomes noise. \square

Lemma 3.7. *Assume Alice calls SEND-MESSAGE(m) in some round r_1 for some message m and bit b . Then the following hold: (1) Alice returns from SEND-MESSAGE(m) in some round $r_3 \geq r_1$; and (2) either Bob records the message m (line 23 of Algorithm 2) in exactly one round r_2 where $r_1 \leq r_2 \leq r_3$, or Bob does not record the message m between rounds r_1 and r_3 because Bob terminated in π .*

Proof. We first show that if Alice calls SEND-MESSAGE(m) in round r_1 , then there must exist some round $r_3 \geq r_1$ in which this call returns. Since the adversary's budget is finite, there must be some round, r_4 , after which no bits are ever flipped. If the call returns before round r_4 , then part (1) is proven, so we only consider the cases where the call extends past round r_4 . In round $r_4 + 1$, after Alice has sent a key request she either hears silence, indicating that Bob has *terminated in π* , and the call returns, or she correctly receives Bob's key, and uses it to send m which Bob correctly receives, and records (since the bit b is different than the bit in the last message recorded). Next, Alice hears silence from Bob. This ends the call. Either way the call returns at the end of round $r_4 + 1$. Thus, in every case there is some round $r_3 \geq r_1$ in which the call to SEND-MESSAGE(m) returns.

We now prove part (2) of our lemma. We first show that the message m is recorded at most once by Bob in the rounds r_1 to r_3 . By Lemma 3.6 (2(c)), the bit b' received by Bob in line 21 of Algorithm 2 must be the same as the bit b sent by Alice from rounds r_1 to r_3 . Since this bit never changes, Bob will record a message at most once in rounds r_1 to r_3 .

If Bob *terminated in π* before round r_3 , then part (2) of our lemma statement does not require Bob to record the message m , and so that part of our lemma is trivially true. Thus, for the remainder

of the proof, we assume that Bob has not *terminated in π* before round r_3 .

Consider round r_3 in which the call to SEND-MESSAGE(m) returns. Let M_1 be the string read by Alice on line 7 of Algorithm 1. Since Bob has not terminated, Lemma 3.6 (2(b)) guarantees that Bob sent M_1 , and therefore the call to Algorithm 2 does not return on line 9. Hence, it returns on line 19. Since this is the last round, M_1 must correctly decode to (k_B, k_A) in Algorithm 1. Thus, Alice sends Bob m using k_B and hears silence subsequently. Thus, Bob must have actually been silent at this time (Lemma 3.6 (1)), which only happens if he has either now or previously recorded m . Hence, Bob must have recorded the message m in some round $r_2 \leq r_3$. \square

Lemma 3.8. *The following holds for π' . For any message m that Bob records (line 23 of Algorithm 2) in some round r_2 , Alice started a call to SEND-MESSAGE with the message m , in a round $r_1 \leq r_2$ and returned from that call in some round $r_3 \geq r_2$.*

Proof. Consider the round r_2 in which the Bob records a message m (line 23 of Algorithm 2). In round r_2 , in line 21 of Algorithm 2, let $\mathcal{D}_t(M'_1) = ((m, b), k_B)$. Thus, by Lemma 3.6 (2(c)), Alice must have sent $\mathcal{E}_t((m, b), k_B)$, during the transmission of M'_1 . Hence, Alice must be in a call to SEND-MESSAGE with the message m , and this call must have begun in some round $r_1 \leq r_2$.

Finally, we know by Lemma 3.7(1) that every call to SEND-MESSAGE ends in some round r_3 . Since, by the above, Alice is in the call during round r_2 , it must be the case that Alice returns from the call in some round $r_3 \geq r_2$. \square

Lemma 3.9. *Algorithm π' terminates with a correct simulation of an asynchronous run of π .*

Proof. We show that for every pair of users u and v , protocol π' correctly simulates a FIFO message channel(see [20] Chapter 14.1.2) from u to v during the simulation of π . Then a direct induction shows that for each processor u , π'_u simulates π_u correctly.

Fix an arbitrary channel from u to v in the network. Let $Q_{u,v}$ be the queue of SEND-MESSAGE procedures that are maintained in π' by u of messages to send to v .

We require the following facts:

1. In any time step, there is a transition in π_u across a transition relation with an output action to send a message m to user v , if and only if the procedure SEND-MESSAGE for message m to user v is pushed on the queue $Q_{u,v}$ in that time step.
2. Every SEND-MESSAGE procedure for message m on $Q_{u,v}$ will eventually start at some round r_1 and end in some round r_3 . Moreover, user v transitions across an input transition relation for message m from user u at most once in some round r_2 , $r_1 \leq r_2 \leq r_3$. Moreover, if there is no such input transition relation, than π_u has entered a termination state before round r_3 .
3. For any transition along a transition relation in π_v with an input action, in some round r_2 , user u started a call to SEND-MESSAGE with the message m , in a round $r_1 \leq r_2$ and returned from that call in some round $r_3 \geq r_2$.

Fact (1) follows directly from Algorithm 3 steps 3-6. The first sentence of Fact (2) follows by induction and Lemma 3.7, and the remainder of the fact follows directly from Lemma 3.7. Fact (3)

follows from Lemma 3.8.

Together, the facts show that no matter what the actions of the adversary, the protocol π' correctly simulates a FIFO message channel from user u to user v . In particular, we have: 1) when a transition is taken in π_u with output action to send message m to user v , this message is put on a queue; 2) all transitions in π_v with an input action to receive a message m from user u are associated with the removal of message m from the queue; and 3) all messages are eventually removed from the queue, triggering transitions across transition relations with input actions in π_v , unless π_v is already in a termination state. \square

3.3 Resource Costs

We now compute the expected number of bits sent and the latency of π' . Let $\tau(r)$ be the time step at which the r^{th} -round of our algorithm begins. Clearly, $\tau(1) = 1$.

Lemma 3.10. *In π' , round $r \geq 1$ begins at time step $\tau(r) = \Theta\left(r \log\left(\frac{nr}{\delta}\right)\right)$.*

Proof. For all $r \geq 1$, note that round $r+1$ begins as soon as the number of time steps corresponding to four words of the round r have passed. Hence, we can compute $\tau(r)$ using the recurrence $\tau(r) = \tau(r-1) + 4w_{r-1}$, where $\tau(1) = 1$. This gives $\tau(r) = \tau(1) + 4 \sum_{i=1}^{r-1} w_i$. Now, from equation (3), using $w_i \leq C_1 \log\left(\frac{ni}{\delta}\right) + C_2$, we get $\tau(r) \leq 1 + 4 \sum_{i=1}^{r-1} (C_1 \log\left(\frac{ni}{\delta}\right) + C_2)$, which gives $\tau(r) = \Theta\left(r \log\left(\frac{nr}{\delta}\right)\right)$. \square

Lemma 3.11. *If $\tau(x) \leq z$ for some $x \geq 1$ and $z \geq 1$, then $x = \mathcal{O}(z / \log\left(\frac{nz}{\delta}\right))$.*

Proof. We prove the bound on x for the case that $\tau(x) = z$. By Lemma 3.10, we know that $\tau(x) = C \left(x \log\left(\frac{nx}{\delta}\right)\right)$, for some constant C . Thus, $z = Cx \log\left(\frac{nx}{\delta}\right)$, and we get (*) $x \leq z / (C \log\left(\frac{nx}{\delta}\right))$. Note that $z = Cx \log\left(\frac{nx}{\delta}\right) \leq Cx \log x$. Taking logs of both sides, we get that $\log z \leq C' \log x$ for some constant C' , which implies that $\log\left(\frac{nz}{\delta}\right) \leq C' \log\left(\frac{nx}{\delta}\right)$. Now plugging this back into (*), we get that $x \leq C'' z / \log\left(\frac{nz}{\delta}\right)$ for some constant C'' . \square

Lemma 3.12. *If π' succeeds, then it has the following resource costs.*

- *The number of bits sent is $\mathcal{O}\left(L \log\left(\frac{nL}{\delta}\right) + T\right)$.*
- *The latency is $\mathcal{O}\left(\max_p \left\{ \Lambda_p \log\left(\frac{n(L+T)}{\delta}\right) + T_p \right\}\right)$, where the maximum is taken over all communication paths, p , in the asynchronous simulation of π , Λ_p is the latency of p , and T_p is the total number of bits flipped by the adversary on edges in p .*

Proof. To bound the number of bits sent, we assume pessimistically that in every round of π' , there is an attempt to send exactly one message. This maximizes the number of bits sent since word sizes increase with time.

Since each word in π' is ECC encoded, the adversary must flip a constant fraction of the bits to successfully corrupt the word, and thereby compromise the round. Let x be the number of rounds in which some word was successfully corrupted by the adversary. Since the length of the words increases with successive rounds, T must at least be a constant, C , times the number of bits in the first x rounds of π' , and hence, we must have $\tau(x) \leq (T+1)/C$.

By Lemma 3.11, we know that $x = \mathcal{O}\left(\frac{T}{\log(n(T+1)/\delta)}\right)$. Thus, since π' requires L progressive rounds, the total number of rounds is $r = L + \mathcal{O}\left(\frac{T}{\log(n(T+1)/\delta)}\right)$.

Hence the total number of bits sent is at most $\tau(r)$. By Lemma 3.10,

$$\begin{aligned} \tau(r) &= \mathcal{O}\left(\left(L + \frac{T}{\log(n(T+1)/\delta)}\right) \log\left(\frac{n}{\delta} \left(L + \frac{T}{\log(n(T+1)/\delta)}\right)\right)\right) \\ &\leq \mathcal{O}\left(\left(L + \frac{T}{\log(n(T+1)/\delta)}\right) \log\frac{n(L+T)}{\delta}\right) \\ &= \mathcal{O}\left(L \log\left(\frac{n(L+T)}{\delta}\right) + T\right) \\ &= \mathcal{O}\left(L \log\left(\frac{nL}{\delta}\right) + T\right). \end{aligned}$$

The second line above follows from the fact that $\frac{T}{\log(n(T+1)/\delta)} \leq T$, and the third line from the fact that if $L = O(T)$, $\log(L+T) = O(\log T)$. The final line above follows from the fact that $\log(L+T) = \log(L) + \log(1+T/L) \leq \log(L) + T/L$, and hence $L \log(L+T) \leq L \log(L) + T$. This bounds the number of bits sent.

To bound the latency, we note that the argument above holds for any communication path p in the asynchronous simulation of π . For the p which achieves the maximum, it follows by induction that all other required messages will have already been received by the time they are needed, and so p determines the overall latency. \square

Now we are ready to complete the proof of Theorem 1.1, in the case when the messages of π are all single bits.

Proof of Theorem 1.1, special case. By Lemma 3.5, no failure event happens with probability at least $1 - \delta$, and by Lemma 3.9, in such a situation, π' terminates with an asynchronous simulation of π . Upon correct termination, the resource cost bounds hold by Lemma 3.12. \square

4 Simulating Protocols with Messages of Arbitrary Length

So far, we have assumed that all messages sent in the protocol π are single bit messages. In fact π could send messages of a fixed constant size, or even of size $\log n$ and our results would still hold, provided the initial word length was chosen long enough to be able to encode the messages in π , the limiting factor here being that the encoding function $\mathcal{E}_t(\cdot)$ can only encode strings up to some length. In this section, we show how to go beyond this restriction and modify our algorithm to be able to simulate protocols π of arbitrary and variable message length, and we show that the cost of the simulation scales well relative to the average message length in π .

Before we proceed, we need to discuss what it means for π to have variable message length. After all, even in the noise-free setting, if Alice is to send Bob a previously unspecified number of bits, then how is Bob to know when to stop reading bits?

Here we assume that the messages from Alice to Bob in π come from a fixed prefix-free language $\mathcal{L} = \mathcal{L}_{A,B}$. Recall that a language \mathcal{L} is prefix-free if for any pair of strings s and s' in \mathcal{L} neither

is a prefix of the other. Note that different prefix-free languages may be used between each pair of players.

4.1 Algorithm

The main simulating algorithm π' (Algorithm 3) remains unchanged in this setting. Only the sending and receiving algorithms need to change to reflect the fact that the message to be sent may be longer than a single call to $\mathcal{E}_t(\cdot)$ can support. In Algorithms 4 and 5 below, we highlight the necessary changes in red.

Algorithm 4 Message exchange algorithm for the sender.

```

1: procedure SEND-MESSAGE( $m$ )
2:    $\triangleright$   $b$  is a persistent variable for the parity bit. On the first call to SEND-MESSAGE,  $b$  is set
3:    $\triangleright$  to 0. On subsequent calls it is whatever it was set to on the previous call.
4:    $j \leftarrow 0$ 
5:   while  $j < |m|$ 
6:      $(w_r, \kappa_t) \leftarrow \text{WORD-PARAMS}(t)$   $\triangleright t$  is the current time step.
7:     Generate random key  $k_A$  of length  $\kappa_t$ .
8:     Send  $\mathcal{E}_t(\text{KEY?}, k_A)$ .
9:      $M_1 \leftarrow w_r$  bits on the channel from the receiver.
10:    if  $\text{IsSilence}(M_1)$  then  $\triangleright$  Assume the receiver has already terminated.
11:       $b \leftarrow \neg b$ 
12:      Stay silent for  $2w_r$  time steps and return
13:    else
14:       $(x, k) \leftarrow \mathcal{D}_t(M_1)$ 
15:      if  $k \neq k_A$  then
16:        Stay silent for  $2w_r$  time steps.
17:      else
18:         $k_B \leftarrow x$ 
19:         $M \leftarrow m[j, j + \kappa_t]$   $\triangleright$  Next  $\kappa_t$  bits of  $m$ 
20:        Send  $\mathcal{E}_t((M, b), k_B)$ .
21:         $M_2 \leftarrow w_r$  bits on the channel from the receiver.
22:        if  $\text{IsSilence}(M_2)$  then
23:           $j \leftarrow j + \kappa_t$ 
24:           $b \leftarrow \neg b$ 
25:    return

```

Algorithm 5 Message exchange algorithm for the receiver.

```
1: procedure RECEIVE-MESSAGE( )
2:   ▷  $\hat{b}, \mu, \lambda$  are persistent variables for the parity bit, the partially received message, and length
3:   ▷ of the recorded message, respectively. On the first call to RECEIVE-MESSAGE,  $\hat{b} \leftarrow 0$ ,
4:   ▷  $\mu \leftarrow \emptyset, \lambda \leftarrow 0$ . On subsequent calls these variables are whatever they were set to on the
5:   ▷ previous call.
6:    $(w_r, \kappa_t) \leftarrow \text{WORD-PARAMS}(t)$  ▷  $t$  is the current time step.
7:    $M'_1 \leftarrow w_r$  bits on the channel from the sender.
8:   if IsSilence( $M'_1$ ) then
9:     Stay silent for  $3w_r$  time steps.
10:  else
11:     $(x', k') \leftarrow \mathcal{D}_t(M'_1)$ 
12:    if  $x' \neq \text{KEY?}$  then
13:      Send noise for  $w_r$  time steps.
14:      Stay silent for  $2w_r$  time steps.
15:    else
16:       $k_A \leftarrow k'$ .
17:      Generate random key  $k_B$  of length  $\kappa_t$ .
18:      Send  $\mathcal{E}_t(k_B, k_A)$ .
19:       $M'_2 \leftarrow w_r$  bits on the channel from the sender.
20:       $(x'', k'') \leftarrow \mathcal{D}_t(M'_2)$ 
21:      if  $k'' \neq k_B$  then
22:        Send noise for  $w_r$  time steps.
23:      else
24:         $(M', b') \leftarrow x''$ 
25:        if  $\mu = \emptyset$  then
26:          if  $b' \neq \hat{b}$  then
27:             $\lambda \leftarrow |M'|$ 
28:             $\mu \leftarrow M'$ 
29:          else
30:            if  $b' \neq \hat{b}$  then
31:               $\lambda \leftarrow |M'|$ 
32:              Append  $M'$  to  $\mu$ .
33:            else
34:              Replace last  $\lambda$  bits of  $\mu$  with  $M'$ .
35:               $\lambda \leftarrow |M'|$  do
36:            if  $\mu$  is a completed message of  $\mathcal{L}$  then
37:              Record the message  $\mu$  from the sender.
38:               $\mu \leftarrow \emptyset$ 
39:               $\hat{b} \leftarrow b'$ 
40:            Stay silent for  $w_r$  time steps.
```

5 Analysis

5.1 Correctness

Lemma 5.1. *Assume Alice calls $\text{SEND-MESSAGE}(m)$ in some round r_1 for some message m and bit b . Then the following hold: (1) Alice returns from $\text{SEND-MESSAGE}(m)$ in some round $r_3 \geq r_1$; and (2) either Bob records the message m (line 37 of Algorithm 5) in exactly one round r_2 where $r_1 \leq r_2 \leq r_3$, or Bob does not record the message m between rounds r_1 and r_3 because he terminated in π .*

Proof. We first show that if Alice calls $\text{SEND-MESSAGE}(m)$ in round r_1 , then there must exist some round $r_3 \geq r_1$ in which this call returns. Since the adversary's budget is finite, there must be some round, r_4 , after which no bits are ever flipped. If the call returns before or during round r_4 , then part (1) is proven, so we only consider the cases where the call extends past round r_4 . Let m' be the part of the message that remains to be sent. For $i \geq 1$ consider round $r_4 + i$. There are two possibilities:

- (a) After Alice has sent a key request in round $r_4 + i$, she hears silence and the call returns, or;
- (b) Alice correctly receives Bob's key, and uses it to send the next piece of m which Bob correctly receives. Next, Alice hears silence from Bob. If this was the last piece of the message, this ends the call. If not, the call continues into round $r_4 + i + 1$ with a shorter remaining message.

Since the message has finite length, eventually (a) occurs. Thus, in every case there is some round $r_3 \geq r_1$ in which the call to $\text{SEND-MESSAGE}(m)$ returns.

We now prove part (2) of our lemma. We first show that the message m is recorded at most once by Bob in the rounds r_1 to r_3 . By Lemma 3.6 (2(c)), the bit b' corresponding to the last partial message for m received by Bob in line 24 of Algorithm 5 must be the same as the bit b sent by Alice for this partial message from rounds r_1 to r_3 . Since this bit never changes, Bob will record m (line 37 of Algorithm 5) at most once in rounds r_1 to r_3 .

If Bob *terminated in π* before round r_3 , then part (2) of our lemma statement does not require Bob to record the message m , and so that part of our lemma is trivially true. Thus, for the remainder of the proof, we assume that Bob has not *terminated in π* before round r_3 and we must show that he records the message m exactly once.

We do this by induction on the length of m . Note that since m is a message from π , it belonged to \mathcal{L} , so it is actually possible for Bob to record m . If m is short enough to be sent in one piece, then consider round r_3 in which the call to $\text{SEND-MESSAGE}(m)$ returns. Let M_1 be the string read by Alice on line 9 of Algorithm 4. Since Bob has not *terminated in π* , Lemma 3.6 (2(b)) guarantees that Bob sent M_1 , and therefore the call to Algorithm 4 does not return on line 9. Hence, it returns on line 25. Since this is the last round, M_1 must correctly decode to (k_B, k_A) in Algorithm 4. Thus, Alice sends Bob m using k_B and hears silence subsequently. Thus, Bob must have actually been silent at this time (Lemma 3.6 (1)), which only happens if he has either now or previously recorded m . Hence, Bob must have recorded the message m in some round $r_2 \leq r_3$.

Now as an induction hypothesis suppose the conclusion about Bob recording a message exactly once is true for all strings s shorter than m . Suppose m requires more than one piece to be sent. Then since Alice continues to resend the first piece m_0 until she has received confirmation that Bob has received it, there is some round $r'_1 < r_3$ in which Alice receives this confirmation. Since the same parity bit b for this piece has been sent in all rounds $r \leq r'_1$, and received bit b' agrees with b in each of these rounds, Bob stores the partial message m_0 exactly once. Let m_1 be the remaining portion of m . Clearly it is shorter than m . Now let us examine the control flow throughout the algorithm from round r'_1 onwards. This looks exactly like a call to SEND-MESSAGE(m_1) with the persistent variable b now set to $-b$, together with a RECEIVE-MESSAGE call in which the persistent variables \hat{b} set to $-\hat{b}$, μ set to m_0 , and λ set to $|m_0|$. By induction hypothesis, there is exactly one round r_2 with $r'_1 \leq r_2 \leq r_3$ during which Bob records $m_0 \circ m_1$, which is in \mathcal{L} . But since $m = m_0 \circ m_1$, this concludes the proof. \square

Lemma 5.2. *The following holds for protocol π' . For any message m that Bob records (line 37 of Algorithm 5) in some round r_2 , Alice started a call to SEND-MESSAGE(m) in a round $r_1 \leq r_2$ and returned from that call in some round $r_3 \geq r_2$.*

Proof. Consider the round r_2 in which the Bob records a message m (line 37 of Algorithm 5). Suppose $\mu = \emptyset$ when RECEIVE-MESSAGE was called in round r_2 . Then m arose in the first component of $\mathcal{D}_t(M'_1)$ in line 24 of Algorithm 5, where the second component correctly matched k_B . Specifically the first component must have been (m, b) for some b . Thus, by Lemma 3.6 (2(c)), Alice must have sent $\mathcal{E}_t((m, b), k_B)$, during the transmission of M'_1 . Hence, Alice must be in a call to SEND-MESSAGE with some message \hat{m} from \mathcal{L} , with m as a contiguous substring, and this call must have begun in some round $r_1 \leq r_2$. We must show that $\hat{m} = m$. Suppose m is a proper contiguous substring of \hat{m} , i.e., $\hat{m} = m_0 \circ m \circ m_1$ where at least one of m_0 and m_1 is not \emptyset . Then m_0 is not in \mathcal{L} (since it is a prefix of \hat{m}). Since Alice would not proceed to sending m until she had confirmed that Bob had received m_0 , she must have previously received that confirmation, meaning that Bob previously received m_0 . But in that case, Bob would have that stored as a partial message, contradicting the assumption that $\mu = \emptyset$ when RECEIVE-MESSAGE was called in round r_2 . Thus $m_0 = \emptyset$. But that means that m is a prefix of \hat{m} . Since both m and \hat{m} are in \mathcal{L} , it follows that $m_1 = \emptyset$ and $\hat{m} = m$. Thus, Alice called SEND-MESSAGE(m).

Next suppose $\mu = s$ when RECEIVE-MESSAGE was called in round r_2 , for some string $s \neq \emptyset$. Then $m = s \circ x$, where x is the string Bob decodes in round r_2 . We first induct on the length of s to prove that Alice's call to SEND-MESSAGE had as input, some message \hat{m} from \mathcal{L} , such that \hat{m} contains $s \circ x = m$ as a contiguous substring, i.e., $\hat{m} = m_0 \circ m \circ m_1$ for some strings m_0, m_1 . The case of $s = \emptyset$ was shown above. Thus, assume there exist rounds $\rho_0 < \rho_1 \leq r_2$ and non-empty strings s_0 and s_1 , with $s = s_0 \circ s_1$ such that $\mu = s_0$ when RECEIVE-MESSAGE was called in round ρ_0 and $\mu = s_1$ when RECEIVE-MESSAGE was called in round ρ_1 . Here, ρ_0 and ρ_1 are the first rounds with this property. Then Bob decoded s_0 in round $\rho_0 - 1$ and s_1 in round $\rho_1 - 1$. Furthermore, Bob has not decoded any other string between decoding s_0 and s_1 . Then, it follows from the induction hypothesis that s_0 and s_1 came from Alice, so that Alice must be in call to SEND-MESSAGE during the course of which she sends encoded strings s_0, s_1, x in this order. Moreover, she cannot have sent any other encoded strings in between, because had she done so, she would not have moved on from it without acknowledgement that Bob decoded it. Since we stipulate that Bob did not decode any other strings, it follows that the input to Alice's call to SEND-MESSAGE had $m = s_0 \circ s_1 \circ x$ as a

contiguous substring so that $\hat{m} = m_0 \circ m \circ m_1$ for some strings m_0, m_1 . Note that m_0 is not in \mathcal{L} . Now, when Bob decoded s_0 , the call to RECEIVE-MESSAGE was passed \emptyset . Thus, from the discussion in the beginning of the proof, it follows that $m_0 = \emptyset$ and therefore $m_1 = \emptyset$ and $\hat{m} = m$. Thus, in all cases, Alice did actually have a call to SEND-MESSAGE(m) that began in some round $r_1 \leq r_2$.

Finally, we know by Lemma 5.1(1) that every call to SEND-MESSAGE ends in some round r_3 . Since, by the above, Alice is in the call during round r_2 , it must be the case that Alice returns from the call in some round $r_3 \geq r_2$. \square

5.2 Resource Costs

We now analyze the total number of bits sent by π' when the messages in π may be longer than a single bit. Let α be the average message length in π .

Lemma 5.3. *If π' succeeds, then it has the following resource costs.*

- The number of bits sent is $\mathcal{O}\left(L\left(1 + \frac{1}{\alpha} \log\left(\frac{nL}{\delta}\right) + T\right)\right)$.
- The latency is $\mathcal{O}\left(\max_p \left\{ \Lambda_p \left(1 + \frac{1}{\alpha} \log\left(\frac{n(L+T)}{\delta}\right) + T_p \right) \right\}\right)$, where p is any communication path in the asynchronous simulation of π , Λ_p is the latency of p , and T_p is the total number of bits flipped by the adversary on edges in p .

Proof. Call a string M encoded in Step 20 of Algorithm 4 a *submessage*. To upper-bound the number of bits sent, we assume pessimistically that in every round of π' , there is an attempt to send exactly one submessage. This maximizes the number of bits sent since word sizes increase with time.

Note that each message consists of at most 1 submessage that is not of length some constant times the word length in that round. Thus, the number of progressive rounds is no more than the number of messages sent, L/α , plus the largest integer x such that the number of bits sent in x rounds equals $c_1 L$ for some constant $c_1 > 1$. By Lemma 3.11, $x = O(L/\log \frac{nL}{\delta})$.

The number of non-progressive rounds is $\mathcal{O}\left(\frac{T}{\log(n(T+1)/\delta)}\right)$, by the same argument as from the proof of Lemma 3.12.

Let the number of rounds $r = L/\alpha + O(L/\log \frac{nL}{\delta}) + \mathcal{O}\left(\frac{T}{\log(n(T+1)/\delta)}\right)$. Then by Lemma 3.10, we can bound $\tau(r)$ as follows.

$$\begin{aligned} \tau(r) &\leq \left(L/\alpha + O\left(\frac{L}{\log(nL/\delta)}\right) + \mathcal{O}\left(\frac{T}{\log(n(T+1)/\delta)}\right) \right) \left(\log \frac{nr}{\delta} \right); \\ &= \mathcal{O}\left(\left(L/\alpha + \frac{L}{\log(nL/\delta)} + \frac{T}{\log(n(T+1)/\delta)} \right) \left(\log \frac{n(L+T)}{\delta} \right) \right); \\ &= \mathcal{O}\left(\frac{1}{\alpha} \log\left(\frac{n(L+T)}{\delta}\right) + L + T \right). \end{aligned}$$

Finally, to bound the latency, we note that the argument above holds for any communication path p in the asynchronous simulation of π . \square

Theorem 5.1. *The protocol π' terminates with a correct simulation of π with probability at least $1 - \delta$. If π' terminates correctly, it has the following resource costs.*

- *The expected number of bits sent is $\mathcal{O}\left(L\left(1 + \frac{1}{\alpha} \log\left(\frac{n(L+T)}{\delta}\right) + T\right)\right)$.*
- *The expected latency is $\mathcal{O}\left(\max_p \left\{\Lambda_p\left(1 + \frac{1}{\alpha} \log\left(\frac{n(L+T)}{\delta}\right) + T_p\right)\right\}\right)$, where p is any communication path in the asynchronous simulation of π , Λ_p is the latency of p , and T_p is the total number of bits flipped by the adversary on edges in p .*

Proof. By Lemma 3.5, no failure event happens with probability at least $1 - \delta$, and by Lemma 3.9, in such a situation, π' terminates with an asynchronous simulation of π . Upon correct termination, the resource cost bounds hold by Lemma 5.3. \square

6 Conclusion and Future Work

We have described the first algorithm in interactive communication for n users that deals with the case of unknown number of bits sent by the protocol, while tolerating an unbounded but finite amount of noise. Against an adversary that flips T bits, given an $\delta \in (0, 1)$, our algorithm compiles a noise free protocol π that sends L bits into a robust protocol π' that succeeds with probability $1 - \delta$, and upon successful termination, sends $\mathcal{O}\left(L\left(1 + \frac{1}{\alpha} \log\left(\frac{n(L+T)}{\delta}\right) + T\right)\right)$ bits, where α is the average message length in π . The blowup in the number of bits is constant for long messages in π and within logarithmic factors of the optimal, otherwise.

Several open problems remain including the following. First, can we adapt our results to interactive communication where π is a synchronous protocol? Second, can we handle an unknown amount of stochastic noise more efficiently, while making no assumption on the value of L or T ? Finally, for any algorithm, what is the minimum number of private random bits required to be hidden from the adversary to achieve robustness?

References

- [1] Zvika Brakerski and Yael Tauman Kalai. “Efficient Interactive Coding against Adversarial Noise”. In: *53rd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*. 2012, pp. 160–166.
- [2] Zvika Brakerski and Moni Naor. “Fast Algorithms for Interactive Coding”. In: *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2013, pp. 443–456.
- [3] Mark Braverman. “Coding for Interactive Computation: Progress and Challenges”. In: *50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. 2012, pp. 1914–1921.
- [4] Mark Braverman. “Towards Deterministic Tree Code Constructions”. In: *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS)*. 2012, pp. 161–167.

- [5] Mark Braverman and Klim Efremenko. “List and Unique Coding for Interactive Communication in the Presence of Adversarial Noise”. In: *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*. 2014, pp. 236–245.
- [6] Mark Braverman and Anup Rao. “Towards Coding for Maximum Errors in Interactive Communication”. In: *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing (STOC)*. 2011, pp. 159–166.
- [7] Mark Braverman, Klim Efremenko, Ran Gelles, and Bernhard Haeupler. “Constant-rate Coding for Multiparty Interactive Communication is Impossible”. In: *Electronic Colloquium on Computational Complexity (ECCC), TR15-197*. 2015.
- [8] Keren Censor-Hillel, Ran Gelles, and Bernhard Haeupler. “Making Asynchronous Distributed Computations Robust to Noise”. In: *arXiv preprint arXiv:1702.07403* (2017).
- [9] Ronald Cramer, Yevgeniy Dodis, Serge Fehr, Carles Padró, and Daniel Wichs. “Detection of Algebraic Manipulation with Applications to Robust Secret Sharing and Fuzzy Extractors”. In: *Advances in Cryptology–EUROCRYPT 2008*. Springer, 2008, pp. 471–488.
- [10] Varsha Dani, Thomas P. Hayes, Mahnush Movahedi, Jared Saia, and Maxwell Young. “Interactive Communication with Unknown Noise Rate”. In: *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP)*. 2015.
- [11] Varsha Dani, Thomas P. Hayes, Mahnush Movahedi, Jared Saia, and Maxwell Young. “Interactive Communication with Unknown Noise Rate”. In: *Information and Computation* (2017), special issue devoted to selected papers from ICALP 2015.
- [12] Devdatt P Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.
- [13] Matthew Franklin, Ran Gelles, Rafail Ostrovsky, and Leonard Schulman. “Optimal Coding for Streaming Authentication and Interactive Communication”. In: *IEEE Transactions on Information Theory* 61.1 (2015), pp. 133–145.
- [14] Ran Gelles, Ankur Moitra, and Amit Sahai. “Efficient and Explicit Coding for Interactive Communication”. In: *Foundations of Computer Science (FOCS)*. 2011, pp. 768–777.
- [15] Mohsen Ghaffari and Bernhard Haeupler. “Optimal Error Rates for Interactive Coding II: Efficiency and List Decoding”. In: *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*. IEEE. 2014, pp. 394–403.
- [16] Mohsen Ghaffari, Bernhard Haeupler, and Madhu Sudan. “Optimal Error Rates for Interactive Coding I: Adaptivity and Other Settings”. In: *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*. 2014, pp. 794–803.
- [17] Bernhard Haeupler. “Interactive channel capacity revisited”. In: *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*. IEEE. 2014, pp. 226–235.
- [18] Morteza Hashemi and Ari Trachtenberg. “Near Real-time Rateless Coding with a Constrained Feedback Budget”. In: *Communication, Control, and Computing (Allerton), 2014 52nd Annual Allerton Conference on*. IEEE. 2014, pp. 529–536.
- [19] Michael Luby. “LT Codes”. In: *Foundations of Computer Science (FOCS)*, 2002.
- [20] Nancy A Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.

- [21] David JC MacKay. “Fountain Codes”. In: *Communications, IEE Proceedings of Communication, Speech and Vision*. Vol. 152. 6. IET. 2005, pp. 1062–1068.
- [22] Cristopher Moore and Leonard J. Schulman. “Tree Codes and a Conjecture on Exponential Sums”. In: *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science (ITCS)*. 2014, pp. 145–154.
- [23] Rafail Ostrovsky, Yuval Rabani, and Leonard J. Schulman. “Error-Correcting Codes for Automatic Control”. In: *Information Theory, IEEE Transactions on* 55.7 (2009), pp. 2931–2941.
- [24] Ravi Palanki and Jonathan S Yedidia. “Rateless codes on noisy channels”. In: *IEEE International Symposium on Information Theory*. Citeseer. 2004, pp. 37–37.
- [25] Marcin Peczarski. “An Improvement of the Tree Code Construction”. In: *Information Processing Letters* 99.3 (Aug. 2006), pp. 92–95.
- [26] Irving S Reed and Gustave Solomon. “Polynomial Codes over Certain Finite Fields”. In: *Journal of the society for industrial and applied mathematics* 8.2 (1960), pp. 300–304.
- [27] Leonard J. Schulman. “Deterministic Coding for Interactive Communication”. In: *Proceedings of the 25th Annual ACM Symposium on Theory of Computing (STOC)*. 1993, pp. 747–756.
- [28] L.J. Schulman. “Communication on Noisy Channels: A Coding Theorem for Computation”. In: *Foundations of Computer Science, 1992. Proceedings., 33rd Annual Symposium on*. 1992, pp. 724–733.