

Towards Optimal Deterministic Coding for Interactive Communication

Ran Gelles* Bernhard Haeupler† Gillat Kol‡ Noga Ron-Zewi §
Avi Wigderson¶

Abstract

We show an *efficient, deterministic* interactive coding scheme that simulates any interactive protocol under random errors with nearly optimal parameters. Specifically, our main result gives an efficient deterministic interactive coding scheme that achieves a communication rate of $1 - O(\sqrt{H(\epsilon)})$ and a failure probability of $\exp(-n/\log n)$, where n is the protocol length and each bit is flipped independently with constant probability ϵ . Prior to our work, all nontrivial deterministic schemes (either efficient or not) had a rate bounded away from 1. Furthermore, the best failure probability achievable by an efficient deterministic scheme with constant rate was only quasi-polynomial, i.e., of the form $\exp(-\log^{O(1)} n)$ (Braverman, ITCS 2012). A rate of $1 - \Theta(\sqrt{H(\epsilon)})$ is essentially optimal (up to poly-logarithmic factors) by a result of Kol and Raz (STOC 2013).

A central contribution in deriving our coding scheme is a novel code-concatenation scheme, a notion standard in coding theory which we adapt for the interactive setting. Essential to our concatenation approach is an explicit, efficiently encodable and decodable *linear tree code* of length n that has relative distance $\Omega(1/\log n)$ and rate approaching 1, defined over an $O(\log n)$ -bit alphabet. The fact that our tree code is linear, and in particular can be made *systematic*, turns out to play an important role in our concatenation scheme.

We use the above tree code as the “outer code” in the concatenation scheme. The necessary deterministic “inner code” is achieved by a nontrivial derandomization of the randomized interactive coding scheme of (Haeupler, STOC 2014). This deterministic coding scheme (with *exponential* running time, but applied here to $O(\log n)$ bit blocks) can handle an ϵ fraction of adversarial errors with a communication rate of $1 - O(\sqrt{H(\epsilon)})$.

*Department of Computer Science, Princeton University, rgelles@cs.princeton.edu. Supported in part by NSF grant CCF-1149888.

†Carnegie Mellon University, haeupler@cs.cmu.edu

‡Institute for Advanced Study, Princeton, NJ. Research at the IAS supported by The Fund For Math and the Weizmann Institute of Science National Postdoctoral Award Program for Advancing Women in Science.

§School of Mathematics, Institute for Advanced Study, Princeton, NJ, USA and DIMACS, Rutgers University, Piscataway, NJ, USA. This research was partially supported by NSF grants CCF-1412958 and CCF-1445755 and the Rothschild fellowship. nogazewi@ias.edu

¶Institute for Advanced Study, Princeton, USA, avi@ias.edu. This research was partially supported by NSF grant CCF-1412958.

1 Introduction

1.1 Background

Coding for interactive communication, the subject of this paper, connects two large bodies of work, coding theory and communication complexity. Both study communication cost, but with very different settings and goals in mind. Coding Theory, born with Shannon’s and Hamming’s breakthrough papers [Sha48, Ham50], is a vast discipline which deals largely with *one-way* communication between two remote parties (Alice and Bob), each holding an input (resp. x , y , possibly from some joint distribution). Major focus is on a *single* communication task: Alice wants to convey x to Bob, and the challenge is doing so reliably when the communication channel between them is *unreliable*, e.g. some of the communicated bits are flipped randomly or adversarially. Alice’s messages are encoded by longer codewords to overcome this “noise”, and one attempts to minimize the communication cost needed to achieve high reliability in each noise model. Communication Complexity, an important research area introduced by Yao [Yao79] 30 years later, also strives to minimize communication cost, but has an opposite focus: it assumes a *perfect* communication channel between Alice and Bob, who now want to perform an *arbitrary* communication task (e.g. computing an arbitrary function $f(x, y)$) using a *two-way* interactive communication protocol.

The seminal work of Schulman [Sch92, Sch93, Sch96] merged these two important subjects, and studied coding schemes for arbitrary two-way interactive communication protocols. Given the interaction and adaptive nature of two-way protocols, this significantly extends the challenge of coding theory. For example, while trade-offs between coding parameters, like the fraction of correctable errors and redundancy for one-way communication have been well understood at least in principle already in Shannon’s paper, and efficient codes matching them were found for many channels, these questions are still far from understood in the two-way case.

In the above papers Schulman set up the basic models, proved the *existence* of nontrivial coding schemes for any interactive protocol, in both the random and the adversarial noise models, and gave an efficient *randomized* scheme for random noise. Progress on finding trade-offs between parameters, and approaching them using *efficient* coding schemes has been slow for a while, but the past few years have seen an impressive flood of new techniques and results on the many challenging questions raised by this general setting, see e.g. [BR14, GMS14, Bra12, KR13, AGS13, BKN14, BE14, GHS14, GH14, Hae14a, FGOS15]; see [Gel15] for a survey. To (informally) cite but one central recent result, Kol and Raz [KR13] (see also [Hae14a]) proved that for the binary symmetric channel BSC_ϵ (in which every communicated bit is flipped independently with probability ϵ), the communication rate is $1 - \tilde{\Theta}(\sqrt{H(\epsilon)})$, where H is the binary entropy function (this should be contrasted with the one-way setting in which the communication rate of BSC_ϵ is known to be $1 - H(\epsilon)$).

Let us describe the basic protocol structure and coding problem a bit more precisely, as in this young field models vary, and some choices actually affect the basic parameters. Many different variants are discussed in many different works but for our focus the assumptions we make are quite standard. Assume that Alice and Bob communicate to perform a distributed task, e.g. compute some function $f(x, y)$ on their respective private inputs x and y . We fix a communication protocol π for this task, in which we assume that the parties alternate: Alice sends one bit in odd steps, and Bob sends one bit in even steps. We further assume that they communicate the same number of bits on every input (the length $n = |\pi|$ of π will be our main complexity parameter). Finally, we assume that each party outputs $\pi(x, y)$, the entire transcript of their conversation (this “universal” output captures all possible tasks, e.g. computing a function). If there is no noise on the channel, this is the standard communication complexity setting.

Now assume that the communication channel is noisy. Our main result is concerned with the probabilistic noise model but throughout we shall consider also the adversarial noise model. In the probabilistic BSC_ϵ model each communicated bit is flipped independently with a constant probability ϵ . In the adversarial model an adversary can flip an ϵ fraction of the communication bits. To cope with the errors, Alice and Bob run a *coding scheme* Π that should *simulate* the noiseless protocol π over the noisy channel. That is, for any inputs x, y the parties hold and for any noiseless protocol π , after running the coding scheme Π over the noisy channel, each party should output $\pi(x, y)$ (if the coding scheme and/or the channel are probabilistic, this should happen with high probability over the randomness). We assume that the parties alternate also in the execution of Π , and as we assumed π has a fixed length (communication cost) $|\pi|$ for every input, we can assume the same for Π , denoting its length by $|\Pi|$.

One basic parameter of a coding scheme Π is the *rate*, defined as $|\pi|/|\Pi|$, which captures the redundancy of Π relative to the noiseless protocol π (this definition parallels the standard one of rate as the ratio between message length to codeword length in one-way communication). Ideally, the rate should approach the capacity of the channel. Two other important goals are computational efficiency and determinism. Ideally the computational complexity of the parties using the scheme Π (given π and their inputs) should be at most polynomial in $|\pi|$, and the coding scheme should be deterministic.¹ Furthermore, in the probabilistic noise model, we desire the coding scheme to fail with exponentially small probability in $|\pi|$ over the randomness of the channel.

1.2 Main result

In this work we focus on the probabilistic noise model BSC_ϵ which is the original model studied in [Sch92]. Specifically, in this paper Schulman gave an efficient *randomized* coding scheme of constant rate (bounded away from 1) and failure probability $\exp(-n/\log n)$ assuming the parties share a common random string, and this latter assumption was recently eliminated in [GMS14]. In the follow-up work [Sch93], Schulmann gave also a *deterministic* scheme of constant rate and exponentially small failure probability. However, this latter scheme is based on tree codes for which no efficient construction is known and therefore it is non-efficient.

In recent years there have been quite a few advancements on obtaining efficient *randomized* coding schemes with optimal parameters. Specifically, in a recent breakthrough, [BKN14] obtained the first efficient randomized scheme of constant rate and exponentially small failure probability, while [KR13] determined the capacity $1 - \tilde{\Theta}(\sqrt{H(\epsilon)})$ of the channel and obtained an efficient randomized scheme that achieves this capacity with failure probability $\exp(-n^\gamma)$ for some constant $\gamma > 0$. Finally, [Hae14a] obtained an optimal randomized scheme that achieves capacity and has an exponentially small failure probability.²

As to efficient *deterministic* coding schemes, dividing the original protocol into chunks of size $O(\log n)$ and simulating each chunk separately via the exponential time scheme of [Sch93] gives an efficient deterministic scheme with constant rate and failure probability $1/\text{poly}(n)$. Recently Braverman [Bra12] showed how to construct tree codes in sub-exponential time, and these in turn when run on chunks of size $\text{polylog}(n)$ give an efficient deterministic scheme with constant rate (bounded away from 1) and lower quasi-polynomial failure probability of the form $\exp(-\text{polylog}(n))$.

¹This presents a subtle issue in the presence of random noise. To prevent a deterministic coding scheme from using the random noise as a source of randomness, one actually uses the so-called *arbitrarily varying channel* (AVC) that extends BSC_ϵ , with which an adversary may determine the *probability* $\epsilon_i \in [0, \epsilon]$ in which the i -th bit is flipped, see e.g., [FK00]. In particular, Π must be correct also if there is no noise at all.

²The schemes of [Sch93, BKN14, Hae14a] are resilient also to adversarial errors.

Our main result provides an interactive coding scheme for BSC_ϵ channel which is both efficient, deterministic and has nearly optimal rate and failure probability.

Theorem 1.1 (Main). *For every $\epsilon > 0$ there exists an efficient deterministic interactive coding scheme Π that simulates any noiseless protocol π of length n over BSC_ϵ with rate $1 - O(\sqrt{H(\epsilon)})$ and failure probability $\exp(-\Omega(\epsilon^4 n / \log n))$.*

As mentioned above, a rate of $1 - \tilde{\Theta}(\sqrt{H(\epsilon)})$ is essentially optimal (up to poly-logarithmic factors) due to [KR13]. Achieving exponentially small failure probability (which can be done by either randomized or inefficient schemes) remains a challenging problem for which this may be a stepping stone.

1.3 Overview of techniques

Our coding scheme exploits the idea of *code concatenation*, which is a very common (simple yet powerful) technique in the one-way coding theory literature [For65]. Concatenation usually consists of two separate coding layers: an *inner code* which is defined over binary alphabet and may be inefficient, and an *outer code* which must be efficient and is defined over large alphabet. In the standard concatenation one first splits the binary message into small blocks, say of size $O(\log n)$, views each block as a single symbol in a large alphabet, and encodes the message using the outer code where each block is considered as a single input symbol of the outer code. Then one switches view again and thinks of each large output symbol of the outer code as a binary string, and encodes each such string separately via the inner code. Such concatenation results in an efficient binary code, and by choosing the right parameters one can also guarantee that this code has good rate and failure probability.

More concretely, in the typical concatenation setting one chooses the outer code to be a code of nearly optimal rate $\rho_{\text{out}} \approx 1$ that is resilient to some δ_{out} fraction of adversarial errors. The inner code on the other hand is chosen to be a code of some rate ρ_{in} that has an optimal exponentially small failure probability over BSC_ϵ . It can be verified that the rate of the final code is the product $\rho_{\text{out}} \cdot \rho_{\text{in}}$ of the rates of the outer and inner codes, and since we chose the rate ρ_{out} of the outer code to be close to 1 then the rate of the final code would be roughly the rate of the inner code ρ_{in} . Furthermore, the final code fails only if more than δ_{out} fraction of the inner codes fail, which for sufficiently large δ_{out} happens with probability at most $(2^{-\Omega(\log n)})^{\delta_{\text{out}}(n/\log n)} = 2^{-\Omega(\delta_{\text{out}} \cdot n)}$ over BSC_ϵ . To summarize, in the above setting of parameters the final code inherits on one hand the running time and the resilience of the outer code, and on the other hand the alphabet size and the rate of the inner code.

In this work we show how the above code concatenation scheme could be implemented in the interactive setting, and our inspiration came from the recent work [BKN14] which gave the first efficient *randomized* interactive coding scheme with constant rate and exponentially small failure probability. Our main observation is that one can interpret the scheme of [BKN14] as an instance of interactive concatenation, and specifically this scheme could be decomposed into an outer code which is an efficient randomized interactive scheme over $O(\log n)$ -bit alphabet that has high rate and is resilient to a constant fraction of adversarial errors, and an inner code which is an exponential-time (deterministic) interactive scheme over binary alphabet that has constant rate and exponentially small failure probability over BSC_ϵ . The inner code is the tree code based scheme of [Sch93], while the outer code is a protocol that proceeds in rounds where at each round the parties exchange the next symbol to be transmitted in the original protocol given the partial transcripts simulated so far, as well as *random* hashes and the lengths of the partial transcripts. The parties then compare the

hashes and the lengths and progress to the next round or rewind to a previous round accordingly. The point here is that if one chooses the alphabet size to be of sufficiently large $\Omega(\log n)$ size then the hash and length exchange adds very little redundancy to the overall communication and consequently a high communication rate is maintained.

Towards a *deterministic* version of interactive concatenation, we take a careful examination of the properties that the outer and inner codes should satisfy in order to enable interactive concatenation. As it turns out, assuming that both the outer and inner codes are interactive coding schemes, the only other property that is required for interactive concatenation is that the outer code could be encoded online when viewing both its input and output symbols as *binary* strings. This property is necessary since the original protocol is an interactive protocol over binary alphabet and therefore one cannot encode a large chunk of it a-priori before it was communicated. One way to ensure the online encoding property is to insist that the outer code would be *systematic*, which means that for every output symbol y_i it holds that $y_i = (x_i, r_i)$ where x_i is the i -th input symbol (the 'systematic' part) and r_i is the 'redundant' part that may depend on all previous input symbols (Indeed, in [BKN14] one can think of r_i as consisting of the hash and length of the partial transcript simulated so far). As linear codes can always be made systematic it in fact suffices that the outer code would be *linear*.

More concretely, as the outer code in our construction we use a novel construction of an efficiently encodable and decodable linear tree code of relative distance $\Omega(1/\log n)$ and rate ≈ 1 , defined over an $O(\log n)$ -bit alphabet. As the inner code we use an exponential time deterministic interactive coding scheme that can correct ϵ fraction of adversarial errors (and so has exponentially small failure probability over BSC_ϵ) with rate $1 - O(\sqrt{H(\epsilon)})$, applied to chunks of length $O(\log n)$. We obtain this latter scheme via a derandomization of a randomized scheme due to [Hae14a] with similar guarantees. The resulting coding scheme is an efficient coding scheme over binary alphabet that has rate $1 - O(\sqrt{H(\epsilon)})$ and failure probability $\exp(-n/\log n)$ over BSC_ϵ .

A by-product of our approach is an interesting connection between the adversarial and random noise settings. Our concatenation lemma (Lemma 3.1) shows that an efficient deterministic (and linear) tree code that resists δ fraction of adversarial noise (along with the inner code described above), leads to an efficient deterministic coding scheme for BSC_ϵ that succeeds with probability $\approx 1 - 2^{-\Omega(n\delta)}$ over the randomness of the channel. In particular, the failure probability of $2^{-\Omega(n/\log n)}$ over a BSC_ϵ exhibited in Theorem 1.1 follows from the fact that the tree code we use as an outer code has a relative distance of $\delta = O(1/\log n)$. Obtaining an efficient deterministic scheme whose failure probability over BSC_ϵ is *exponentially small* remains an interesting open question. Our concatenation lemma suggests that this question is closely related to the 20-year open question of finding deterministic and efficient tree codes with constant rate and relative distance [Sch93], and obtaining deterministic and efficient coding schemes with constant rate that resist a constant fraction of noise in the adversarial setting.

In what follows we describe the outer and inner codes we use in more detail and the techniques used to obtain them.

Outer code. The high-level idea of the tree code construction is as follows. For every integer t such that $\Omega(\log \log n) \leq t \leq \log n$ we partition the message to blocks of size 2^t . Each such block is separately encoded via a standard (one-way) systematic error-correcting code with relative distance $\Omega(1/\log n)$ and rate $1 - O(1/\log n)$. This yields a redundant part $R^{(t)}$ of 2^t bits which are layered across the next block, i.e., across the encodings of the next 2^t levels, so that every level gets 1 bit. This layering amortizes the redundancy across the tree, which helps keeping the rate approaching 1

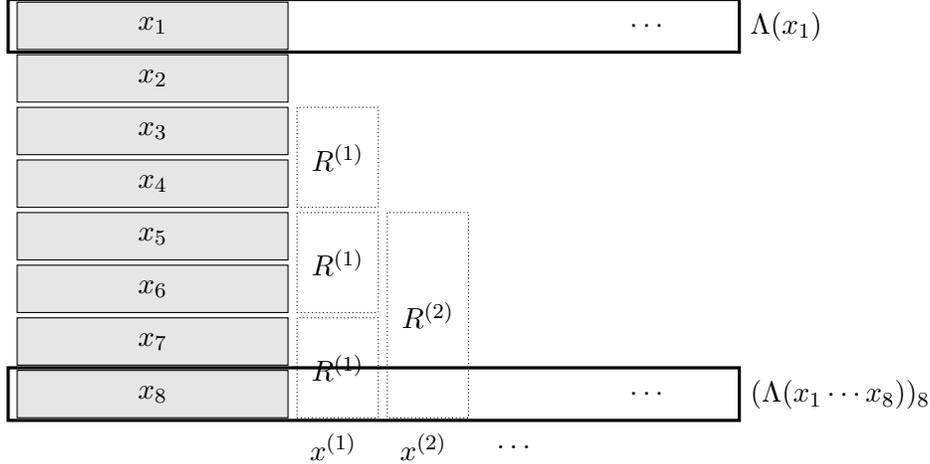


Figure 1: An illustration of the tree code’s encoding; see Section 5 for details.

while still providing the required relative distance guarantee of $\Omega(1/\log n)$, yet only over the next 2^t levels. See Figure 1 for an illustration of the construction.

A somewhat related tree code construction is outlined in an unpublished memo by Schulman [Sch03]: that construction uses the same idea of encoding prefixes of increasing lengths 2^t , using an asymptotically good error-correcting code with constant rate and relative distance, then layering the output codeword across the next 2^t levels (a similar high-level idea was used also in [Bra12]). However, since a non-systematic code with rate bounded away from 1 is used this results with rate $O(1/\log n)$; on the other hand, the obtained relative distance is constant $\Omega(1)$. One can view these two constructions as tradeoffing complementary goals: while [Sch03] optimizes the distance at the expense of a low rate, we optimize the rate at the expense of a low distance.

To turn the above tree code into an interactive coding scheme we use a scheme similar to the tree code based scheme of [Sch93]. However, the original analysis of [Sch93] only achieved a constant rate bounded away from 1, regardless of the rate of the tree code, and we provide an improved analysis of this scheme in the spirit of [Hae14a] that preserves the rate of the tree code. We also observe that the coding scheme preserves the linearity of the tree code, and consequently this scheme could be used as the outer code in our concatenation scheme.

Inner Code. The starting point for the inner code construction is an efficient randomized coding scheme by Haeupler [Hae14a]. That scheme achieves rate of $1 - O(\sqrt{\epsilon})$ over a BSC_ϵ , however it is randomized. We obtain our inner code by derandomizing this scheme, and the main technical issue here is to derandomize the scheme without damaging its rate. On the other hand, we do not need the derandomization to be efficient, since we will run the inner code only on protocols of length $O(\log n)$.

Our derandomization approach generalizes the derandomization technique of [BKN14]. More specifically, in [BKN14] the authors show an upper bound of $2^{O(n)}$ on the number of different sequences of partial transcripts that may occur during a run of their coding scheme and then show that for each sequence of transcripts at least $1 - 2^{-\Omega(\epsilon \cdot \ell \cdot n)}$ of the random strings are good for this sequence, where ℓ is the output length of the hash functions. Choosing $\ell = \Omega(1/\epsilon)$ to be a sufficiently large constant, a union bound shows the existence of a single random string that works for all sequences.

In our case, since we want to maintain a high rate we cannot afford the output length of the hash functions to be a too large constant, and so the upper bound of $2^{O(n)}$ is too large for our purposes. To deal with this we show how to relate simulated transcripts to the noise patterns that may have caused them. A fine counting argument reduces the effective number of sequences of transcripts to the number of *typical* noise patterns, i.e., the $2^{O(H(\epsilon)n)}$ patterns with at most ϵ fraction of bit flips. Then a union bound argument on all the possible noise patterns proves the existence of a single random string that works for any typical noise pattern. For this union bound to work we need to use slightly longer hashes than in the original scheme of [Hae14a] of length $\ell = \Omega(\log(1/\epsilon))$ which results in a slightly reduced rate of $1 - O(\sqrt{H(\epsilon)})$.

1.4 Organization of the paper

We begin (Section 2) by recalling several building blocks and setting up notations we will use throughout. Our main concatenation lemma is provided in Section 3, along with a formal statement of the inner and outer codes we use. These prove our main Theorem 1.1. The detailed proof of the concatenation lemma appears in Section 4. In Sections 5 and 6 we present our inner and outer codes constructions, respectively.

2 Preliminaries

All logarithms in this paper are taken to base 2. We denote by $H : [0, 1] \rightarrow [0, 1]$ the binary entropy function given by $H(p) = p \log(1/p) + (1 - p) \log(1/(1 - p))$ for $p \notin \{0, 1\}$ and $H(0) = H(1) = 0$. Let \mathbb{F}_2 denote the finite field of two elements and let \mathbb{N} denote the set of positive integers. For an integer $n \in \mathbb{N}$ let $[n] := \{1, \dots, n\}$ and for a pair of integers $m, n \in \mathbb{N}$ such that $m \leq n$ let $[m, n] := \{m, m + 1, \dots, n\}$. For a vector $x \in \Sigma^n$ and integers $1 \leq i \leq j \leq n$ we denote by $x[i, j]$ the projection of x on the coordinates in the interval $[i, j]$, and we let $|x| = n$ denote the length of x . Finally, the *relative distance* between a pair of strings $x, y \in \Sigma^n$ is the fraction of coordinates on which x and y differ, and is denoted by $\text{dist}(x, y) := |\{i \in [n] : x_i \neq y_i\}| / n$.

2.1 Error-correcting codes

A code is a mapping $C : \Sigma_{\text{in}}^k \rightarrow \Sigma_{\text{out}}^n$. We call k the *message length* of the code and n the *block length* of the code. The elements in the image of C are called *codewords*. The *rate* of C is the ratio $\frac{k \log |\Sigma_{\text{in}}|}{n \log |\Sigma_{\text{out}}|}$. We say that C has *relative distance* at least δ if for every pair of distinct vectors $x, y \in \Sigma_{\text{in}}^k$ it holds that $\text{dist}(C(x), C(y)) \geq \delta$.

Let \mathbb{F} be a finite field. We say that C is \mathbb{F} -*linear* if $\Sigma_{\text{in}}, \Sigma_{\text{out}}$ are vector spaces over \mathbb{F} and the map C is linear over \mathbb{F} . If $\Sigma_{\text{in}} = \Sigma_{\text{out}} = \mathbb{F}$ and C is \mathbb{F} -linear then we simply say that C is *linear*. Finally, if $k = n$ then we say that a code $C : \Sigma_{\text{in}}^n \rightarrow \Sigma_{\text{out}}^n$ is *systematic* if $\Sigma_{\text{in}} = \Gamma^s$ and $\Sigma_{\text{out}} = \Gamma^{s+r}$ for some alphabet Γ and integers $s, r \in \mathbb{N}$, and there exists a string $R(x) \in (\Gamma^r)^n$ such that $(C(x))_i = (x_i, (R(x))_i)$ for every $x \in \Sigma_{\text{in}}^n$ and $i \in [n]$ (that is, the projection of $(C(x))_i$ on the first s coordinates equals x_i). We call x and $R(x)$ the *systematic part* and the *redundant part* of $C(x)$, respectively.

Specific families of codes. We now mention some known constructions of error-correcting codes that we shall use as building blocks in our tree code construction, and state their relevant properties. We start with the following fact that states the existence of Reed-Solomon codes which achieve the best possible trade-off between rate and distance over large alphabets.

Fact 2.1 (Reed-Solomon codes [RS60]). *For every $k, n \in \mathbb{N}$ such that $k \leq n$, and for every finite field \mathbb{F} such that $|\mathbb{F}| \geq n$ there exists a linear code $\text{RS} : \mathbb{F}^k \rightarrow \mathbb{F}^n$ with rate k/n and relative distance at least $1 - \frac{k}{n}$. Furthermore, RS can be encoded and decoded from up to $(1 - \frac{k}{n})/2$ fraction of errors in time $\text{poly}(n, \log \mathbb{F})$.*

The next fact states the existence of asymptotically good binary codes. Such codes can be obtained for example by concatenating the Reed-Solomon codes from above with binary linear Gilbert-Varshamov codes [Gil52, Var57].

Fact 2.2 (Asymptotically good binary codes). *For every $0 < \rho < 1$ there exist $\delta > 0$ and integer $k_0 \in \mathbb{N}$ such that the following holds for any integer $k \geq k_0$. There exists a binary linear code $B : \{0, 1\}^k \rightarrow \{0, 1\}^n$ with rate at least ρ and relative distance at least δ . Furthermore, B can be encoded and decoded from up to $\delta/2$ fraction of errors in time $\text{poly}(n)$.*

2.2 Tree codes

A tree code [Sch96] is an error-correcting code $\Lambda : \Sigma_{\text{in}}^n \rightarrow \Sigma_{\text{out}}^n$ which is a *prefix-code*: for any $i \in [n]$ and $x \in \Sigma_{\text{in}}^n$ the first i symbols of $\Lambda(x)$ depend only on x_1, \dots, x_i . For simplicity we shall sometimes abuse notation and denote by Λ also the map $\Lambda : \Sigma_{\text{in}}^j \rightarrow \Sigma_{\text{out}}^j$ which satisfies that $(\Lambda(x_1, \dots, x_j))_i = (\Lambda(x_1, \dots, x_n))_i$ for every $i \in [j]$ and $x \in \Sigma_{\text{in}}^n$. Observe that this latter map is well defined as $(\Lambda(x_1, \dots, x_n))_i$ depends only on x_1, \dots, x_i .

We say that Λ has *relative tree distance* at least δ if for every pair of distinct vectors $x, y \in \Sigma_{\text{in}}^n$ such that $i \in [n]$ is the first coordinate on which x and y differ (i.e., $(x_1, \dots, x_{i-1}) = (y_1, \dots, y_{i-1})$ but $x_i \neq y_i$), and for every j such that $i \leq j \leq n$ it holds that

$$\text{dist}\left(\Lambda(x)[i, j], \Lambda(y)[i, j]\right) \geq \delta.$$

Alternatively, the relative tree distance of a tree code can be defined via the notion of *suffix distance* [FGOS15] (see also [BE14, BGMO15]). The *suffix distance* between a pair of strings $x, y \in \Sigma^n$ is

$$\text{dist}_{\text{sfx}}(x, y) := \max_{i \in [n]} \left\{ \text{dist}\left(x[i, n], y[i, n]\right) \right\}.$$

It can be shown that a tree code has relative tree distance at least δ if and only if for every pair of distinct vectors $x, y \in \Sigma_{\text{in}}^n$ it holds that $\text{dist}_{\text{sfx}}(\Lambda(x), \Lambda(y)) \geq \delta$.

Finally, we say that Λ can be *decoded from α fraction of errors* if there exists an algorithm that is given as input a vector $w \in \Sigma_{\text{out}}^j$ for some $j \in [n]$ and outputs a vector $y \in \Sigma_{\text{in}}^j$ that satisfies the following guarantee: If there exists $x \in \Sigma_{\text{in}}^j$ such that $\text{dist}_{\text{sfx}}(\Lambda(x), w) \leq \alpha$, then $y = x$.

3 Efficient Deterministic Coding Scheme

In this section we prove our main Theorem 1.1. This theorem is an immediate implication of our concatenation lemma below. The concatenation lemma proves that given an efficient deterministic systematic tree code (used as an outer code) and a possibly inefficient deterministic coding scheme (used as an inner code), one can construct an efficient deterministic coding scheme, and states the parameters of the resulting coding scheme as a function of the parameters of the outer and inner codes. We now give the concatenation lemma (whose proof appears in Section 4). Then, in Lemmas 3.2 and 3.3 below, we state the existence of outer and inner codes with good parameters, whose concatenation proves our main Theorem 1.1.

Lemma 3.1 (Concatenation). *Suppose that the following hold:*

1. (Outer code) *There exists a systematic tree code $\Lambda : \Sigma_{\text{in}}^{n_\Lambda} \rightarrow \Sigma_{\text{out}}^{n_\Lambda}$ with $\Sigma_{\text{in}} = \{0, 1\}^s$, $\Sigma_{\text{out}} = \{0, 1\}^{s+r}$ and rate ρ_Λ that can be encoded and decoded from up to δ_Λ fraction of errors in time \mathcal{T}_Λ .*
2. (Inner code) *There exists a deterministic interactive coding scheme Π that simulates any noiseless protocol π of length $s + 2(r + 1)$ with rate ρ_Π in the presence of up to δ_Π fraction of adversarial errors, and with running time \mathcal{T}_Π .*

Then for every $\gamma > 0$ there exists a deterministic interactive coding scheme Π' that simulates any noiseless protocol π' over $\text{BSC}_{\delta_\Pi/2}$ of length $n_\Lambda \cdot (s - 2) \cdot (1 - \gamma)$ with rate

$$\frac{\rho_\Lambda}{2 - \rho_\Lambda + 4/(s - 2)} \cdot \rho_\Pi \cdot (1 - \gamma),$$

and failure probability

$$\exp \left[-n_\Lambda \left(\frac{\delta_\Lambda}{36} \cdot \frac{s + 2(r + 1)}{\rho_\Pi} \cdot \frac{\delta_\Pi^2}{4} \cdot \gamma - H \left(\frac{\delta_\Lambda}{36} \cdot \gamma \right) \right) \right].$$

Furthermore, the coding scheme Π' has running time $O(n_\Lambda \cdot (\mathcal{T}_\Lambda + \mathcal{T}_\Pi))$.

The following lemmas give the tree code that will be used as the ‘outer code’ in the concatenation step, and an exponential time deterministic coding scheme that will be used as the ‘inner code’ in the concatenation step.

Lemma 3.2 (Outer code). *There exists an absolute constant $\delta_0 > 0$ such that the following holds for every $\epsilon > 0$ and $n \in \mathbb{N}$. There exists a systematic \mathbb{F}_2 -linear tree code $\Lambda : \Sigma_{\text{in}}^n \rightarrow \Sigma_{\text{out}}^n$ with $\Sigma_{\text{in}} = \{0, 1\}^{(\log n)/\epsilon}$, $\Sigma_{\text{out}} = \{0, 1\}^{(\log n)/\epsilon + \log n}$, rate $\frac{1}{1+\epsilon}$ and relative tree distance at least $\frac{\delta_0 \cdot \epsilon}{\log n}$. Furthermore, Λ can be encoded and decoded from up to a fraction $\frac{\delta_0 \cdot \epsilon}{2 \log n}$ of errors in time $\text{poly}(n)$.*

We prove the above lemma in Section 5.

Lemma 3.3 (Inner code). *For every $\epsilon > 0$ there exists a deterministic interactive coding scheme Π that simulates any noiseless protocol π of length n with rate $1 - O(\sqrt{H(\epsilon)})$ in the presence of up to a fraction ϵ of adversarial errors. Furthermore, Π has running time $\text{poly}(n)$ and can be constructed in time $2^{O(n)}$.*

We prove the above lemma in Section 6. We can now prove our main Theorem 1.1 based on the above Lemmas 3.1, 3.2 and 3.3.

Proof of Theorem 1.1. Follows from Lemma 3.1 by letting Λ be the tree code guaranteed by Lemma 3.2 for constant ϵ and integer n_Λ such that $n = n_\Lambda((\log n_\Lambda)/\epsilon - 2)(1 - \epsilon)$ (so $n_\Lambda = \Omega(\epsilon n / \log n)$), letting Π be the coding scheme guaranteed by Lemma 3.3 for constant 2ϵ , and letting $\gamma = \epsilon$. \square

4 The Concatenation Lemma: Proof of Lemma 3.1

We start with a high-level description of the coding scheme Π' . We describe below the coding scheme Π' for Alice; the coding scheme for Bob is symmetric.

Throughout the execution of the coding scheme Alice (respectively, Bob) maintains a string T^A that represents Alice's current guess for the transcript of the simulated protocol π' communicated so far. Alice also maintains a string \hat{T}^B that represents Alice's current guess for the corresponding string T^B of Bob. When the execution of the coding scheme Π' is completed the outputs of Alice and Bob are T^A and T^B , respectively.

The coding scheme Π' is executed for n_Λ iterations, where at iteration i Alice and Bob use the inner coding scheme Π to communicate the next block X_i of length $s - 2$ of π' (assuming that the transcript of π' communicated so far is T^A and T^B , respectively), as well as a pair of *position strings* $p_{i-1}^A, p_{i-1}^B \in \{0, 1\}^2$, and a pair of *hash strings* $h_{i-1}^A, h_{i-1}^B \in \{0, 1\}^r$.

Alice (respectively, Bob) then performs, one of three actions according to the output of the simulation via the inner coding scheme Π : she either appends her noisy version X_i^A of X_i to T^A , or she leaves T^A unchanged, or she erases the last block of length $s - 2$ from T^A . These actions correspond to the case where a seemingly correct simulation of X_i has occurred, a seemingly incorrect simulation has occurred, or it seems that the prefix T^A is incorrect, respectively. She then records her action in the i -th position string p_i^A (since there are only three possible actions those could be recorded using 2 bits).

Lastly, Alice views the string $(\sigma_{\text{in}})_i^A := (p_i^A, X_i^A) \in \{0, 1\}^s$ as the systematic part of the i -th output symbol of the tree code Λ and lets $(\sigma_{\text{out}})_i^A$ be the corresponding i -th output symbol of the tree code. The i -th hash string $h_i^A \in \{0, 1\}^r$ is set to be the redundant part of $(\sigma_{\text{out}})_i^A$. As described above, both the strings p_i^A and h_i^A will be communicated by Alice in iteration $i + 1$. Note that for every i , the string $((\sigma_{\text{in}})_1^A, \dots, (\sigma_{\text{in}})_i^A)$ records all the actions of Alice on T^A till iteration i and so, if decoded correctly by Bob, then Bob can extract the value of T^A at iteration i from this string (same goes for Alice). The formal definition of the coding scheme Π' appears below.

4.1 The coding scheme Π'

Coding scheme $(\Pi')^A$ for Alice:

Initialize: $T^A := \emptyset$, $\hat{T}^B := \emptyset$.

For $i = 1, \dots, n_\Lambda$ iterations:

1. Recall that p_{i-1}^A denotes the first 2 bits of $(\sigma_{\text{out}})_{i-1}^A$ and let h_{i-1}^A denote the last r bits of $(\sigma_{\text{out}})_{i-1}^A$ (for $i = 1$ let $(\sigma_{\text{out}})_0^A := 0^{s+r}$).
2. Simulate the protocol $\pi^A(|T^A|, (T^A, 0^{n_\Lambda(s-2)-|T^A|}), p_{i-1}^A, h_{i-1}^A)$ below using the inner coding scheme Π . Let the sequence $(p_{i-1}^A, \hat{p}_{i-1}^B, h_{i-1}^A, \hat{h}_{i-1}^B, X_i^A)$ denote the output of the simulation where $p_{i-1}^A, \hat{p}_{i-1}^B \in \{0, 1\}^2$, $h_{i-1}^A, \hat{h}_{i-1}^B \in \{0, 1\}^r$ and $X_i^A \in \{0, 1\}^{s-2}$.
3. Let $(\hat{\sigma}_{\text{out}})_{i-1}^B := (\hat{p}_{i-1}^B, X_{i-1}^A, \hat{h}_{i-1}^B)$. Decode the sequence $((\hat{\sigma}_{\text{out}})_1^B, \dots, (\hat{\sigma}_{\text{out}})_{i-1}^B)$ using the decoding algorithm for Λ . Let $((\hat{\sigma}_{\text{in}})_1^B, \dots, (\hat{\sigma}_{\text{in}})_{i-1}^B)$ be the decoded message and let \hat{T}^B be the transcript represented by this string (if $i = 1$ then set $\hat{T}_B = \emptyset$).

4. If $T^A = \hat{T}^B$ append X_i^A to T^A and set $p_i^A := 01$.
5. Otherwise, if $T^A \neq \hat{T}^B$ and $|T^A| < |\hat{T}^B|$ set $p_i^A := 00$.
6. Otherwise, if $T^A \neq \hat{T}^B$ and $|T^A| \geq |\hat{T}^B|$ erase the last $s-2$ bits from T^A and set $p_i^A := 10$.
7. Let $(\sigma_{\text{in}})_i^A := (p_i^A, X_i^A)$ and let $(\sigma_{\text{out}})_i^A$ be the i -th symbol of $\Lambda((\sigma_{\text{in}})_1^A, \dots, (\sigma_{\text{in}})_i^A)$. Note that since Λ is systematic it holds that $(\sigma_{\text{in}})_i^A$ is a prefix of $(\sigma_{\text{out}})_i^A$.

The output of the coding scheme is the prefix of T^A of length $n_\Lambda \cdot (s-2) \cdot (1-\gamma)$.

Next we describe the protocol π . This protocol is simulated by the inner coding scheme Π at Step 2 of the coding scheme Π' . The protocol π receives as input an integer $1 \leq t \leq n_\Lambda(s-2)$, a transcript string $T \in \{0, 1\}^{n_\Lambda(s-2)}$, a position string $p \in \{0, 1\}^2$ and a hash string $h \in \{0, 1\}^r$. The description of π for Alice's side, denoted π^A , is the following.

Protocol $\pi^A(t, T, p, h)$ for Alice:

1. Send p, h and receive \hat{p}, \hat{h} (this is done bit by bit).
2. Communicate bits $[t+1, \dots, t+(s-2)]$ of the protocol π' assuming that the first t bits of π' communicated so far are the first t bits of T .

4.2 Analysis

4.2.1 Rate and running time

The coding scheme Π' runs for n_Λ iterations and at each iteration the number of bits communicated is $((s-2) + 2(r+2))/\rho_\Pi$. Recall that $\rho_\Lambda = \frac{s}{s+r}$. Consequently, the rate of the coding scheme Π' is

$$\begin{aligned}
\frac{|\pi'|}{|\Pi'|} &= \frac{n_\Lambda \cdot (s-2) \cdot (1-\gamma)}{n_\Lambda \cdot ((s-2) + 2(r+2)) / \rho_\Pi} \\
&= \frac{s-2}{2(s+r) - (s-2)} \cdot \rho_\Pi \cdot (1-\gamma) \\
&= \frac{s-2}{2(s-2)/\rho_\Lambda + 4/\rho_\Lambda - (s-2)} \cdot \rho_\Pi \cdot (1-\gamma) \\
&= \frac{\rho_\Lambda}{2 + 4/(s-2) - \rho_\Lambda} \cdot \rho_\Pi \cdot (1-\gamma). \tag{1}
\end{aligned}$$

To analyze the running time note that the running time of each iteration is $O(\mathcal{T}_\Lambda + \mathcal{T}_\Pi)$ and therefore the total running time is $O(n_\Lambda \cdot (\mathcal{T}_\Lambda + \mathcal{T}_\Pi))$.

4.2.2 Decoding guarantees

To analyze the decoding guarantees we define a potential function Φ as follows. Let t^+ be the number of blocks of length $s-2$ contained in the longest prefix on which T^A and T^B agree, and

let $t^- = \frac{|T^A|+|T^B|}{s-2} - 2t^+$. Let $\Phi = t^+ - t^-$. Note that if at the end of the simulation it holds that $\Phi \geq n_\Lambda \cdot (1 - \gamma)$, then the simulation must be successful. The reason is that in this case $t^+ \geq n_\Lambda \cdot (1 - \gamma)$ and so a prefix of length at least $n_\Lambda \cdot (s - 2) \cdot (1 - \gamma)$ is correct in both T^A and T^B , which means the entire transcript π' was correctly simulated.

To bound the potential we shall use the notion of a *good iteration*.

Definition 4.1. We say that an iteration i is *good* if the following pair of conditions hold:

1. At Step 2 of iteration i , the simulation of π via the inner coding scheme Π is successful.
2. At Step 3 of iteration i , it holds that $T^A = \hat{T}^A$ and $T^B = \hat{T}^B$.

Claim 4.2. *The potential Φ decreases by at most 3 after any iteration. Furthermore, after any good iteration the potential increases by at least 1.*

Proof. At any single iteration, a party either leaves its transcript unchanged, erases the last block of its transcript, or adds a new block to its transcript. Therefore t^+ can change by at most 1 and t^- can change by at most 2 at each iteration, and so the total potential change at each iteration is at most 3.

Next observe that if iteration i is good, then both parties know the transcript of the other side at the beginning of iteration; they also learn the correct value of the block X_i . Therefore, if $T^A = T^B$ at the beginning of iteration i , then both parties add X_i to their transcript, t^+ increases by 1 and t^- remains zero. Otherwise, if $T^A \neq T^B$ and $|T^A| = |T^B|$, then both parties erase the last block of their transcript, thus t^+ does not change and t^- decreases by 2. Finally, if $|T^A| \neq |T^B|$, then the party with the longer transcript erases the last block of its transcript and so t^+ does not change while t^- decreases by 1. We conclude that the total potential change at a good iteration is at least 1. \square

Claim 4.2 above implies that the simulation of π' via Π' succeeds as long as the number of bad iterations throughout the execution of Π' is at most $n_\Lambda \gamma / 4$. The following claim shows that we can bound the number of bad iterations by bounding the number of iterations in which the first condition in Definition 4.1 does not hold.

Claim 4.3. *If the first condition of Definition 4.1 does not hold in at most m iterations, then the number of bad iterations is at most $9m/\delta_\Lambda$.*

Proof. By symmetry, it suffices to show that in addition to the m iteration where the first condition doesn't hold, there are at most $4m/\delta_\Lambda$ iterations in which $T_B \neq \hat{T}_B$ at Step 3.

Fix an iteration $i + 1$ in which $T_B \neq \hat{T}_B$ at Step 3 and let $((\hat{\sigma}_{\text{in}})_1^B, \dots, (\hat{\sigma}_{\text{in}})_i^B)$ be the decoded message at this step. By the decoding guarantee of Λ there exists $t(i) \in [i]$ such that in at least δ_Λ fraction of the iterations $j \in [t(i), i]$ the simulation at Step 2 failed in either iteration j or iteration $j + 1$ (since X_j is transmitted on iteration j but p_j and h_j are transmitted only on iteration $j + 1$). This implies in turn that in at least $\delta_\Lambda/2$ fraction of the iterations $j \in [t(i), i + 1]$ the simulation at Step 2 failed in iteration j . In particular, if the simulation fails at Step 2 in the segment $[t(i), i + 1]$ at most m times, then $|[t(i), i + 1]| < 2m/\delta_\Lambda$. However, we must take care of overlapping segments $[t(i), i + 1]$.

Let

$$\mathcal{I} = \left\{ [t(i), i + 1] \mid T_B \neq \hat{T}_B \text{ at Step 3 of iteration } i + 1 \right\},$$

and define $\bigcup \mathcal{I} = \bigcup_{I \in \mathcal{I}} I$. Since for each iteration $i+1$ in which $T^B \neq \hat{T}^B$ it holds that $i+1 \in \bigcup \mathcal{I}$, it suffices to show that $|\bigcup \mathcal{I}| \leq 4m/\delta_\Lambda$. Lemma 7 in [Sch96] says that there exists a subset $\mathcal{I}' \subseteq \mathcal{I}$ of disjoint intervals such that $|\bigcup \mathcal{I}'| \geq |\bigcup \mathcal{I}|/2$. The proof is completed by noting that if the simulation at Step 2 failed in at most m iterations, then it must be that $|\bigcup \mathcal{I}'| \leq 2m/\delta_\Lambda$, and so $|\bigcup \mathcal{I}| \leq 4m/\delta_\Lambda$. \square

Using the above Claim 4.3, the simulation of Π' is successful as long as the number of iterations in which the simulation at Step 2 failed is at most $\delta_\Lambda n_\Lambda \gamma / 36$. Over $\text{BSC}_{\delta_\Pi/2}$, since the inner coding scheme Π can handle δ_Π fraction of adversarial errors, the probability that the simulation at Step 2 fails is at most

$$\exp\left(-\left(\frac{\delta_\Pi}{2}\right)^2 \cdot \frac{s+2(r+1)}{\rho_\Pi}\right),$$

independently for each iteration. Therefore the probability of having more than $\delta_\Lambda n_\Lambda \gamma / 36$ iterations in which the simulation at Step 2 fails is at most

$$\begin{aligned} & \sum_{k=\delta_\Lambda n_\Lambda \gamma / 36}^{n_\Lambda} \binom{n_\Lambda}{k} \exp\left(-\frac{\delta_\Pi^2}{4} \cdot \frac{s+2(r+1)}{\rho_\Pi} \cdot k\right) \\ &= \exp\left[-n_\Lambda \left(\frac{\delta_\Lambda}{36} \cdot \frac{s+2(r+1)}{\rho_\Pi} \cdot \frac{\delta_\Pi^2}{4} \cdot \gamma - H\left(\frac{\delta_\Lambda}{36} \cdot \gamma\right)\right)\right]. \end{aligned}$$

5 The Outer Code: Proof of Lemma 3.2

5.1 The tree code construction

A main ingredient in our tree code construction is the following lemma showing the existence of a *systematic* error-correcting code $C : \Sigma_{\text{in}}^k \rightarrow \Sigma_{\text{out}}^k$ with appropriate parameters. Specifically, this lemma shows that for any integers k, n that satisfy $\Omega((\log n)/\epsilon) \leq k \leq n$, there exists a systematic code $C : \Sigma_{\text{in}}^k \rightarrow \Sigma_{\text{out}}^k$ with $|\Sigma_{\text{in}}| = \text{poly}(n)$, $|\Sigma_{\text{out}}| = \text{poly}(n)$, rate $1 - O(\frac{\epsilon}{\log n})$ and relative distance $\Omega(\frac{\epsilon}{\log n})$. The lemma follows by an application of Facts 2.1 and 2.2, and we defer its proof to Section 5.4.

Lemma 5.1. *There exists an absolute constant $k_0 \in \mathbb{N}$ such that the following holds for every $\epsilon > 0$ and integers $k, n \in \mathbb{N}$ such that $k_0 \cdot (\log n)/\epsilon \leq k \leq n$. There exists a systematic \mathbb{F}_2 -linear code $C : \Sigma_{\text{in}}^k \rightarrow \Sigma_{\text{out}}^k$ with $\Sigma_{\text{in}} = \{0, 1\}^{(\log n)/\epsilon}$, $\Sigma_{\text{out}} = \{0, 1\}^{(\log n)/\epsilon+1}$, rate $\rho' := \frac{1}{1+\epsilon/\log n}$ and relative distance at least $\delta' := \frac{1}{2(\log n)/\epsilon+1}$. Furthermore, C can be encoded and decoded from up to a fraction $\delta'/2$ of errors in time $\text{poly}(n)$.*

The construction of the tree code Λ is as follows. Let $m := k_0 \cdot (\log n)/\epsilon$, for simplicity assume that both m and n are powers of 2. The encoding $\Lambda(x)$ of a message $x \in \Sigma_{\text{in}}^n$ is the pointwise concatenation of the message string x with $\log n - \log m + 1$ binary strings $x^{(\log m)}, \dots, x^{(\log n)} \in \{0, 1\}^n$, where for $\log m \leq t \leq \log n$ the string $x^{(t)} \in \{0, 1\}^n$ is defined as follows. Let $C^{(t)} : \Sigma_{\text{in}}^{2^t} \rightarrow \Sigma_{\text{out}}^{2^t}$ be the systematic code given by Lemma 5.1 for a constant ϵ and message length $k = 2^t$, and let $R^{(t)} : \Sigma_{\text{in}}^{2^t} \rightarrow \{0, 1\}^{2^t}$ be the redundant part of $C^{(t)}$. Divide the string x into $n/2^t$ blocks $z_1, \dots, z_{n/2^t}$ of length 2^t each, and let $x^{(t)} = (0^{2^t}, R^{(t)}(z_1), \dots, R^{(t)}(z_{n/2^t-1}))$. See Figure 2.

We clearly have that Λ can be encoded in time $\text{poly}(n)$. Note furthermore that Λ is systematic and \mathbb{F}_2 -linear and that the input alphabet size of Λ is $2^{\log n/\epsilon}$ and the output alphabet size of Λ is

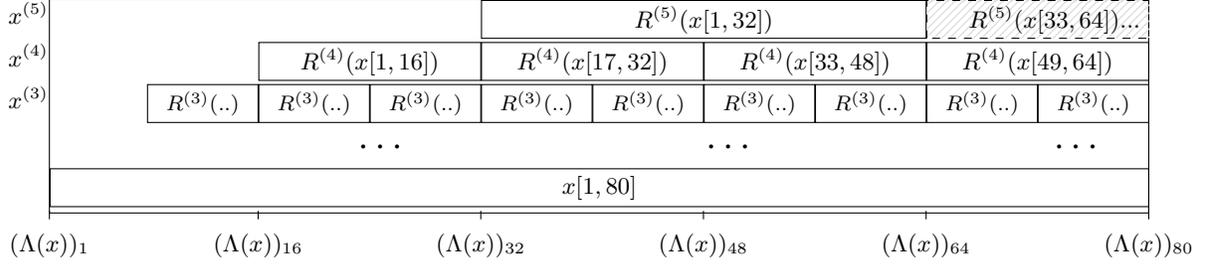


Figure 2: An illustration of the first 80 indices of $\Lambda(x)$, the encoding of $x \in \Sigma_{\text{in}}^n$ using our tree code.

$2^{\log n/\epsilon} \cdot 2^{\log n - \log m + 1} \leq 2^{\log n/\epsilon + \log n}$. The rate of Λ is then at least

$$\frac{(\log n)/\epsilon}{(\log n)/\epsilon + \log n} = \frac{1}{1 + \epsilon}.$$

It remains to analyze the distance and decoding guarantee of Λ .

5.2 Distance

The distance guarantee of the tree code stems from the fact that as long as we look at two different messages x, y that differ in their suffixes of length $\geq 2m$, then the encoding at these suffixes completely includes a pair of codewords $C^{(t)}(x') \neq C^{(t)}(y')$ for some $\log m \leq t \leq \log n$. Below, we show that either the suffix is shorter than $2m$ and then the required distance trivially holds, or we find the maximal value of t for which the above holds and then the required distance follows from the distance guarantee of the code $C^{(t)}$.

Claim 5.2. *Let $x, y \in \Sigma_{\text{in}}^n$ be a pair of distinct messages and let $i \in [n]$ be the first coordinate on which x and y differ. For any $j \in [i, n]$ it holds that*

$$\text{dist}\left(\Lambda(x)[i, j], \Lambda(y)[i, j]\right) \geq \min\left\{\frac{\epsilon}{2k_0 \log n}, \frac{1}{16(\log n)/\epsilon + 8}\right\}.$$

Lemma 3.2 then holds as a corollary of the above claim by setting $\delta_0 := 1/(32k_0)$.

Proof. If $j - i < 2m$ then

$$\text{dist}\left(\Lambda(x)[i, j], \Lambda(y)[i, j]\right) \geq \frac{1}{j - i + 1} \geq \frac{1}{2m} = \frac{\epsilon}{2k_0 \log n},$$

where the first inequality follows since $(\Lambda(x))_i \neq (\Lambda(y))_i$ due to our assumption that $x_i \neq y_i$ and the tree code being systematic.

Next assume that $j - i \geq 2m$. Let t be the maximal integer such that $2 \cdot 2^t \leq j - i$ and let $i_0 := \lfloor \frac{i-1}{2^t} \rfloor \cdot 2^t$ be $i - 1$ rounded down to the nearest multiple of 2^t . Note that

$$i_0 + 1 \leq i < i_0 + 1 + 2^t < i_0 + 2 \cdot 2^t \leq j$$

and

$$j - i < 4 \cdot 2^t,$$

due to the maximality of t .

Note that $\Lambda(x)[i_0 + 1, i_0 + 2^t]$ contains $x[i_0 + 1, i_0 + 2^t]$ as the systematic parts of $C^{(t)}(x[i_0 + 1, i_0 + 2^t])$. Also note that by our construction, $\Lambda(x)[i_0 + 1 + 2^t, i_0 + 2 \cdot 2^t]$ contains the redundant part $R^{(t)}(x[i_0 + 1, i_0 + 2^t])$ of $C^{(t)}(x[i_0 + 1, i_0 + 2^t])$. In a symmetric way, the same holds for $\Lambda(y)$ and y .

Furthermore, the assumption that $x_i \neq y_i$ implies that $x[i_0 + 1, i_0 + 2^t] \neq y[i_0 + 1, i_0 + 2^t]$ and so by the distance guarantee of $C^{(t)}$ (as given by Lemma 5.1) we have that

$$\text{dist}\left(C^{(t)}(x[i_0 + 1, i_0 + 2^t]), C^{(t)}(y[i_0 + 1, i_0 + 2^t])\right) \geq \delta'. \quad (2)$$

Equation (2) implies that either

$$\text{dist}\left(x[i_0 + 1, i_0 + 2^t], y[i_0 + 1, i_0 + 2^t]\right) \geq \frac{\delta'}{2}$$

or

$$\text{dist}\left(R^{(t)}(x[i_0 + 1, i_0 + 2^t]), R^{(t)}(y[i_0 + 1, i_0 + 2^t])\right) \geq \frac{\delta'}{2}.$$

Finally, note that in either case we get that

$$\text{dist}\left(\Lambda(x)[i, j], \Lambda(y)[i, j]\right) \geq \frac{(\delta'/2) \cdot 2^t}{j - i + 1} \geq \frac{(\delta'/2) \cdot 2^t}{4 \cdot 2^t} = \frac{\delta'}{8},$$

where the first inequality is due to the fact that $i_0 + 1 \leq i < i_0 + 1 + 2^t < i_0 + 2 \cdot 2^t \leq j$ and i is the first coordinate on which x and y differ, and the second inequality is due to the fact that $j - i < 4 \cdot 2^t$. Recall that $\delta' = \frac{1}{2(\log n)/\epsilon + 1}$ to complete the proof. \square

5.3 Decoding

Recall that the decoding procedure is given a word $w \in \Sigma_{\text{out}}^j$ for some $1 \leq j \leq n$ and is required to output a vector $y \in \Sigma_{\text{in}}^j$ such that $y = x$ whenever $x \in \Sigma_{\text{in}}^j$ is such that $\text{dist}_{\text{sfx}}(\Lambda(x), w) \leq \frac{\delta_0 \cdot \epsilon}{2 \log n}$.

For a given word $w \in \Sigma_{\text{out}}^j$, the decoded word $y \in \Sigma_{\text{in}}^j$ is obtained as follows. We decode w in parts according to its partitioning into blocks corresponding to the codes $C^{(t)}$. Specifically, we start from the largest t for which a codeword $C^{(t)}$ is fully contained in the prefix of w . We then move on to decode the remaining suffix in an iterative manner. We proceed this way until the interval at hand is shorter than $2m$, in which case we simply set y in this interval as the systematic part of w in the corresponding interval.

The formal description of the decoding procedure follows.

Decoding procedure on input $w \in \Sigma_{\text{out}}^j$:

0. Initialize: $\ell := 1$ // Left index of current interval
1. If $j - \ell < 2m$, set $y[\ell, j]$ to be the systematic part of $w[\ell, j]$ and output y .
2. Otherwise, let t be the maximal integer such that $2 \cdot 2^t \leq j - \ell$.
3. Decode the part of $w[\ell, \ell - 1 + 2 \cdot 2^t]$ that corresponds to the encoding of the code $C^{(t)}$ using the decoding procedure for $C^{(t)}$, and set $y[\ell, \ell - 1 + 2^t]$ to be the result of the decoding.
4. Set $\ell := \ell + 2^t$ and return to Step 1.

Let us give an example of the decoding process of $w \in \Sigma_{\text{out}}^{75}$. (Recall Figure 2.) For this example, let us assume that $m = 8 = 2^3$. We begin by decoding $y[1, 32]$; this is done by decoding the code $C^{(5)}$ whose systematic part lies in $w[1, 32]$ and redundant part $R^{(5)}(x[1, 32])$ lies in $w[33, 64]$. Note that we could not use the code $C^{(6)}$ since its redundant part would be in the interval $[65, 128]$ which is beyond the range of w . After we set $y[1, 32]$, we move on to the next interval. We cannot decode $y[33, 64]$ using the next $C^{(5)}$ since its redundant part lies beyond the range of w , and we need to reduce the scale to $t = 4$. Hence, the next part we decode is $y[33, 48]$, which is obtained using the code $C^{(4)}$ whose systematic part lies in $w[33, 48]$ and redundant part $R^{(4)}(x[33, 48])$ lies in $w[49, 64]$. The next $C^{(4)}$ is again beyond the currently decoded w and we reduce the scale to $t = 3$. Using the code $C^{(3)}$ we decode $y[49, 56]$, and also $y[57, 64]$. Finally, we are left with the interval $[65, 75]$ whose length is $11 < 2m$; we assume that there are no errors in this interval and simply set $y[65, 75]$ to be the systematic part of $w[65, 75]$.

We clearly have that the decoding procedure runs in time $\text{poly}(n)$. To show that the decoding procedure satisfies the required decoding guarantee we observe that our assumption—that the distance of w from $\Lambda(x)$ is small on every suffix—implies that at each iteration the part of $w[\ell, \ell - 1 + 2 \cdot 2^t]$ that corresponds to the encoding of $C^{(t)}$ is close to $C^{(t)}(x[\ell, \ell - 1 + 2^t])$. Consequently, the decoding guarantee of $C^{(t)}$ implies that $y[\ell, \ell - 1 + 2^t] = x[\ell, \ell - 1 + 2^t]$ for every iteration in which $j - \ell \geq 2m$.

In more detail, suppose that $x \in \Sigma_{\text{in}}^j$ is such that $\text{dist}_{\text{sfx}}(\Lambda(x), w) \leq \frac{\delta_0 \cdot \epsilon}{2 \log n}$. We shall show that at each iteration the coordinates of y are set to the corresponding coordinates of x and so $y = x$.

If $j - \ell < 2m$ at some iteration then we have that

$$\text{dist}\left(\Lambda(x)[\ell, j], w[\ell, j]\right) \leq \frac{\delta_0 \cdot \epsilon}{2 \log n} = \frac{1}{64m} < \frac{1}{j - \ell + 1},$$

where the equality follows due to our choice of $m = k_0(\log n)/\epsilon$ and $\delta_0 = 1/(32k_0)$. This implies in turn that $w[\ell, j] = \Lambda(x)[\ell, j]$ and so the systematic part of $w[\ell, j]$ equals $x[\ell, j]$ and consequently $y[\ell, j] = x[\ell, j]$.

Next assume that $j - \ell \geq 2m$. To show the required decoding guarantee in this case note that our assumption implies that

$$\text{dist}\left(\Lambda(x)[\ell, j], w[\ell, j]\right) \leq \frac{\delta_0 \cdot \epsilon}{2 \log n}.$$

Furthermore, due to maximality of t we have that $j - \ell < 4 \cdot 2^t$, and consequently it holds that

$$\text{dist}\left(\Lambda(x)[\ell, \ell - 1 + 2^t], w[\ell, \ell - 1 + 2^t]\right) \leq \frac{4 \cdot 2^t \cdot (\delta_0 \cdot \epsilon)/(2 \log n)}{2^t} = \frac{2 \cdot \delta_0 \cdot \epsilon}{\log n} \leq \frac{\delta'}{4},$$

and similarly

$$\text{dist}\left(\Lambda(x)[\ell + 2^t, \ell - 1 + 2 \cdot 2^t], w[\ell + 2^t, \ell - 1 + 2 \cdot 2^t]\right) \leq \frac{\delta'}{4}.$$

This implies in turn that the part of $w[\ell, \ell - 1 + 2 \cdot 2^t]$ that corresponds to the encoding of $C^{(t)}$ is of relative distance at most $\delta'/2$ from $C^{(t)}(x[\ell, \ell - 1 + 2^t])$, and so by the decoding guarantee of $C^{(t)}$ it holds that $y[\ell, \ell - 1 + 2^t] = x[\ell, \ell - 1 + 2^t]$.

5.4 Proof of Lemma 5.1

We now complete the proof of Lemma 3.2 by proving Lemma 5.1. Lemma 5.1 follows by substituting $\rho = 1/2$, $s = (\log n)/\epsilon$ and $r = 1$ in the following lemma which shows the existence of a systematic error-correcting code with good rate and distance.

Lemma 5.3. *For every $0 < \rho < 1$ there exist $\delta > 0$ and integer $k_0 \in \mathbb{N}$ such that the following holds for any integers $k, s, r \in \mathbb{N}$ that satisfy $k \cdot \frac{\rho r}{s} \geq k_0$ and $s \geq \log(k(1 + \frac{\rho r}{s}))$. There exists a systematic \mathbb{F}_2 -linear code $C : \Sigma_{\text{in}}^k \rightarrow \Sigma_{\text{out}}^k$ with $\Sigma_{\text{in}} = \{0, 1\}^s$, $\Sigma_{\text{out}} = \{0, 1\}^{s+r}$, rate $\frac{s}{s+r}$ and relative distance at least $\delta' := \min \left\{ \delta, 1 - \frac{s/\rho}{s/\rho+r} \right\}$. Furthermore, C can be encoded and decoded from up to $\delta'/2$ fraction of errors in time $\text{poly}(k, s, r)$.*

Proof. Since C is systematic it suffices to define the redundant part R of C . Roughly speaking, $R(x)$ is obtained by first encoding the message x via a systematic Reed-Solomon code, then encoding the redundant part of the resulting codeword with an asymptotically good binary code, and finally spreading the resulting bits evenly between the k coordinates of $R(x)$.

Formally, let δ and k_0 be the constants guaranteed by Fact 2.2 for rate ρ , and let B be the asymptotically good binary code guaranteed by this fact for rate ρ and message length $k \cdot \frac{\rho r}{s}$ (recall that we assume $k \cdot \frac{\rho r}{s} \geq k_0$). Let RS be the Reed-Solomon code guaranteed by Fact 2.1 for message length k and block length $k(1 + \frac{\rho r}{s})$ over a field \mathbb{F} of size 2^s , and note that our assumptions imply that $2^s \geq k(1 + \frac{\rho r}{s})$. By performing Gaussian elimination, we may assume without loss of generality that the code RS is systematic, that is, for every $x \in \mathbb{F}^k$ it holds that $\text{RS}(x) = (x, R'(x))$ for some string $R'(x) \in \mathbb{F}^{k\rho r/s}$.

Next we define the redundant part R of C . To this end, fix a string $x \in \Sigma_{\text{in}}^k = \mathbb{F}^k$ and let $R'(x) \in \mathbb{F}^{k\rho r/s}$ be the redundant part of the encoding of x via the Reed-Solomon code RS. Next view $R'(x)$ as a binary string in $\{0, 1\}^{k\rho r}$ via the usual \mathbb{F}_2 -linear isomorphism and encode this binary string via the asymptotically good binary code B , let $z_x \in \{0, 1\}^{kr}$ denote the resulting string. Finally, divide the string z_x into k blocks of size r , and for every $1 \leq i \leq k$ let $(R(x))_i \in \{0, 1\}^r$ be the i -th block of z_x .

Next we analyze the properties of C . It can be verified that C has the required rate $\frac{s}{s+r}$. To see that the relative distance of C is at least δ' , let $x \neq y \in \Sigma_{\text{in}}^k$ be a pair of strings. If

$$\text{dist}(x, y) \geq 1 - \frac{k}{k(1 + \rho r/s)} = 1 - \frac{s/\rho}{s/\rho + r}$$

then we are done due to C being systematic. Otherwise, due to the distance guarantee of the code RS we must have that $R'(x) \neq R'(y)$, and consequently the distance guarantee of the code B implies that $\text{dist}(z_x, z_y) \geq \delta$. Finally, note that grouping the coordinates of z_x and z_y cannot decrease the relative distance between the pair of strings, and so we must have that $\text{dist}(R(x), R(y)) \geq \delta$ as well. The decoding guarantees of C follow from similar considerations, based on the decoding guarantees of the codes RS and B . \square

6 The Inner Code: Proof of Lemma 3.3

The inner code is obtained via a derandomization of a randomized interactive coding scheme due to Haeupler [Hae14a, Algorithm 3]. The main use of randomness in [Hae14a] is to allow the parties to check, with high probability, whether or not they are synchronized (e.g., hold the same partial transcript). To this end, each party chooses a random hash function and communicates a short hash of its own state. Note that due to this shrinkage in length, it may happen that although the parties are unsynchronized, the hash values they exchange are the same, leading the parties to falsely believe they are synchronized. Such an event is called a *hash collision*. We show how to devise a *deterministic* variant of the coding scheme of [Hae14a], in which we fix the randomness,

and show that there exists a fixing that is “good” for all possible runs, namely, the amount of hash collisions that can occur for that fixing is low enough to complete the simulation correctly.

To this end, we observe that when the adversary is limited to corrupting at most a fraction ϵ of the transmissions, then there are only $2^{O(H(\epsilon)n)} = 2^{O(\log(1/\epsilon)\epsilon n)}$ different noise patterns that should be considered; denote these as *typical noise patterns*. We then carefully modify the way the coding scheme of [Hae14a] compares the states the parties hold, using linear hash functions. The linearity of the hash functions along with the specific way in which we perform the comparisons make hash collisions depend (roughly) only on the specific noise pattern and the randomness string, and most importantly, (almost) independent of the specific noiseless protocol π that is simulated by the coding scheme and the inputs (x, y) of the parties (The fact that hash collisions do not entirely depend only on the noise pattern and the randomness string creates further complications in our proof which we ignore in the discussion below).

Finally, we show that if we increase the output length of the hash functions from a constant (as used in [Hae14a]) to $c' \log(1/\epsilon)$ for some constant c' , then for each typical noise pattern, at least a $1 - 2^{-\Omega(c' \log(1/\epsilon)\epsilon n)}$ fraction of the randomness strings lead to at most ϵn hash collisions which is a small enough number of hash collisions that allows the simulation to be completed correctly. By setting c' to be a large enough constant, a union bound proves that there must exist a single randomness string that is “good” for *all* possible typical noise patterns. A by-product of the increase in the output length of the hash functions is that the rate of the coding scheme slightly reduces from $1 - O(\sqrt{\epsilon})$ to $1 - O(\sqrt{H(\epsilon)})$.

Concretely, we prove Lemma 3.3 in two steps. In the first step we slightly modify the original scheme of [Hae14a], specifically, by carefully modifying the way the hash comparisons are performed and slightly increasing the output length of the hash functions, as outlined above. In the second step we derandomize this latter modified coding scheme. The two steps are given in Sections 6.1 and 6.2 below, respectively. The proof and detailed analysis below build on the analysis of [Hae14a] and are not self-contained. We refer the reader to the longer version of that paper [Hae14b], and in the following all the references (lemma numbers, line numbers, variable names, etc.) correspond to that version. For a better readability, in Appendix A we re-iterate Algorithm 3 of [Hae14b].

6.1 Modified scheme

In this section we slightly modify the randomized coding scheme given by Algorithm 3 in [Hae14b] to obtain a randomized coding scheme $\tilde{\Pi}$ that is more suitable for derandomization, and state some properties of the modified scheme $\tilde{\Pi}$ that we shall use for the derandomization step. We start by describing the modified coding scheme $\tilde{\Pi}$.

6.1.1 Modified scheme

Let $\tilde{\Pi}$ be the coding scheme that is obtained from Algorithm 3 in [Hae14b] via the following modifications.

1. (Output length of hash functions) The output length o of the hash functions is increased from $\Theta(1)$ to $c' \cdot \log(1/\epsilon)$ for sufficiently large constant c' to be determined later on. Consequently, the number r_c of check bits per iteration is also increased from $\Theta(1)$ to $\Theta(\log(1/\epsilon))$, the length of the blocks $r = \sqrt{r_c/\epsilon}$ increases from $\Theta(\sqrt{1/\epsilon})$ to $\Theta(\sqrt{\log(1/\epsilon)/\epsilon})$, and the number of iterations R_{total} decreases to

$$\lceil n/r + 65\epsilon n \rceil = \Theta(\sqrt{\epsilon/\log(1/\epsilon)})n + 65\epsilon n.$$

2. (Seed length) Our modified hash comparisons described below apply hash functions to strings of length $\Theta(n \log n)$, as opposed to length $\Theta(n)$ as is done in Algorithm 3 in [Hae14b]. To this end, we increase the seed length s of the hash functions per iteration from $\Theta(n)$ to $\Theta(n \log n)$. Note that in this case the random string R at Line 5 can still be obtained by exchanging $\Theta(\sqrt{\epsilon \log(1/\epsilon)})n$ random bits sampled from a δ -biased distribution with bias $\delta = 2^{-\Theta(n\epsilon/r)} = 2^{-\Theta(\sqrt{\epsilon \log(1/\epsilon)})n}$.
3. (Position string $N(T)$) To make hash collisions depend (roughly) only on the noise pattern and the randomness string, the parties maintain throughout the execution of the coding scheme $\widetilde{\Pi}$ a *position string* $N(T) \in [R_{\text{total}}]^{R_{\text{total}} \cdot r}$ whose i -th coordinate equals the iteration in which the i -th bit of T was added to T , or is empty in the case where the i -th bit of T is empty. We denote by $N'(T) \in \{0, 1\}^{R_{\text{total}} \cdot r \cdot \log(R_{\text{total}})}$ the binary string obtained from $N(T)$ by replacing each coordinate of $N(T)$ with its binary representation of length $\log(R_{\text{total}})$ (we pad with zeros if the length is shorter than $\log(R_{\text{total}})$).
4. (Hash comparisons) Roughly speaking, our new hash comparisons will apply an \mathbb{F}_2 -linear hash function (specifically, the inner product function) to both the transcript T and the vector $N'(T)$ that encodes the iterations in which each of the bits in T were added to T . Specifically, in Line 9 we replace

$$\text{hash}_S(k), \text{hash}_S(T), \text{hash}_S(T[1, \text{MP1}]), \text{hash}_S(T[1, \text{MP2}])$$

with

$$\text{hash}_S(k), \widetilde{\text{hash}}_S(T), \widetilde{\text{hash}}_S(T[1, \text{MP1}]), \widetilde{\text{hash}}_S(T[1, \text{MP2}]),$$

where the function hash_S is as defined in [Hae14b] and the function $\widetilde{\text{hash}}_S$ is defined as follows. For integers m, o and a seed $S \in \{0, 1\}^{m \cdot o}$ let $h_S : \{0, 1\}^{\leq m} \rightarrow \{0, 1\}^o$ be the \mathbb{F}_2 -linear hash function that satisfies, for every $x \in \{0, 1\}^{\leq m}$ and $i \in [o]$, that

$$(h_S(x))_i = \langle x, S[(i-1) \cdot m + 1, im] \rangle,$$

where $\langle a, b \rangle = \sum_{i=1}^m a_i \cdot b_i \pmod{2}$ is the inner product mod 2 of $a, b \in \{0, 1\}^m$ (if $|x| < m$ then we assume that x is padded with zeroes to the right up to length m). We note that for every seed S the function h_S is \mathbb{F}_2 -linear, i.e., for any two strings $x, y \in \{0, 1\}^m$ it holds that $h_S(x \oplus y) = h_S(x) \oplus h_S(y)$.

Finally, for $m = R_{\text{total}} \cdot r \cdot \log(R_{\text{total}}) = \Theta(n \log n)$, $o = c' \log(1/\epsilon)$ and a seed $S \in \{0, 1\}^{m \cdot o}$ we let

$$\widetilde{\text{hash}}_S : \{0, 1\}^{\leq R_{\text{total}} \cdot r} \rightarrow \{0, 1\}^{3o}$$

be the hash function that satisfies

$$\widetilde{\text{hash}}_S(x) = \left(h_S(x), h_S(|x|), h_S(N'(x)) \right)$$

for every string $x \in \{0, 1\}^{\leq R_{\text{total}} \cdot r}$.

6.1.2 Properties of the modified scheme

Next we state some properties of the modified scheme $\widetilde{\Pi}$ that we use later. We start with the following lemma which says that the simulation is successful as long as at most ϵn iterations suffer from a hash collision. This lemma can be proved along the lines of the proof of Theorem 7.1 in [Hae14b].

Lemma 6.1. *Let $R \in \{0, 1\}^{R_{\text{total}} \cdot s}$ be an arbitrary string (not necessarily coming from a δ -biased distribution), and let Γ be a run of $\tilde{\Pi}$ that uses the string R as the random string sampled at Line 5, and simulates a noiseless protocol π on inputs (x, y) in the presence of up to ϵ fraction of adversarial errors. Suppose furthermore that at most ϵn iterations in Γ suffer from a hash collision. Then the output of Γ is $\pi(x, y)$ (that is, the simulation performed by Γ is successful).*

We shall also use the following claim which follows from the above lemma and says that if at most ϵn iterations suffer from a hash collision up to some iteration $t \in [R_{\text{total}}]$, then in most iterations $i \in [t]$ a new block was added to both T_A and T_B .

Claim 6.2. *Let $R \in \{0, 1\}^{R_{\text{total}} \cdot s}$ be an arbitrary string (not necessarily coming from a δ -biased distribution), and let Γ be a run of $\tilde{\Pi}$ that uses the string R as the random string sampled at Line 5, and simulates a noiseless protocol π on inputs (x, y) in the presence of up to ϵ fraction of adversarial errors. Let $t \in [R_{\text{total}}]$ be some iteration and suppose that at most ϵn iterations $i \in [t]$ in Γ suffer from a hash collision. Then there are at most $65\epsilon n$ iterations $i \in [t]$ in which no block was added to T_A and at most $65\epsilon n$ iterations $i \in [t]$ in which no block was added to T_B .*

Proof. Suppose in contradiction that there are more than $65\epsilon n$ iterations $i \in [t]$ in which no block was added to T_A or more than $65\epsilon n$ iterations $i \in [t]$ in which no block was added to T_B . By symmetry we may assume that there are more than $65\epsilon n$ iterations $i \in [t]$ in which no block was added to T_A . To arrive at a contradiction we shall modify the string R to obtain a string R' such that when the string R in the run Γ is replaced with the string R' then on one hand, at most ϵn iterations in Γ suffer from a hash collision and on the other hand, the simulation performed by Γ is unsuccessful which contradicts Lemma 6.1 above.

Specifically, let $R' \in \{0, 1\}^{R_{\text{total}} \cdot s}$ be the string which agrees with R on the first $t \cdot s$ bits and the last $(R_{\text{total}} - t) \cdot s$ bits are chosen such that no hash collision occurs after iteration t when the string R in the run Γ is replaced with the string R' . Such a choice exists since the output length of the hash functions is $o = c' \log(1/\epsilon)$ and the coding scheme is performing only a constant number of hash comparisons per iteration. Consequently, the probability that a uniform random seed in $\{0, 1\}^s$ causes a hash collision at some iteration is at most $\exp(-\Omega(c' \log(1/\epsilon)))$, and in particular there exists a seed in $\{0, 1\}^s$ that does not cause a hash collision at this iteration.

Let Γ' be the run of the coding scheme $\tilde{\Pi}$ obtained from Γ by replacing the string R with the string R' . On one hand, we have that at most ϵn iterations in Γ' suffer from a hash collision. On the other hand, since Γ' and Γ behave the same on the first t iterations there are more than $65\epsilon n$ iterations in Γ' in which no block was added to T_A . But since $\tilde{\Pi}$ is run for $n/r + 65\epsilon n$ iterations and since in each iteration at most one block is added to T_A , we have that at the end of the run Γ' less than n/r blocks of length r are present in T_A , and so the simulation is unsuccessful. This contradicts Lemma 6.1 above. \square

6.2 Derandomization

In order to derandomize the coding scheme $\tilde{\Pi}$ defined above we proceed according to the program outlined at the beginning of this section. Specifically, we observe that as long as each block in T_A was added at the same iteration in which the corresponding block in T_B was added (that is, $N(T_A) = N(T_B)$) then T_A and T_B differ only by the noise pattern corresponding to the iterations in which the blocks in T_A and T_B were added. Since the hash function h_S we use is \mathbb{F}_2 -linear, in this case we have that hash collisions, when comparing T_A and T_B , depend only on the noise pattern and the seed S used in these iterations. However, when $N(T_A) \neq N(T_B)$, hash collisions may not

depend entirely on the noise pattern and the random seed, and this creates further complications in our proof.

To cope with the above situation we replace in our analysis *noise patterns* with *behavior patterns* which include the noise pattern as well as some extra information on some of the transitions made during the execution of $\tilde{\Pi}$. We also replace *hash collisions* with *hash mismatches* which are a notion of inconsistency of hash functions that includes hash collisions as a special case. The advantage of these notions is that now hash mismatches depend *entirely* on the behavior pattern and the randomness string.

We focus on a certain subset of behavior patterns we name *typical behavior patterns*; those are a subset of the behavior patterns that can occur when the adversary is limited to ϵ fraction of errors. We then show that there are at most $2^{O(H(\epsilon)n)} = 2^{O(\log(1/\epsilon)\epsilon n)}$ different typical behavior patterns, and that for each typical behavior pattern, at least a $1 - 2^{-\Omega(c' \log(1/\epsilon)\epsilon n)}$ fraction of the randomness strings lead to at most ϵn hash mismatches. This implies in turn that for a large enough constant c' there must exist a *single* good randomness string that leads to at most ϵn hash mismatches (and thus, at most ϵn hash collisions) for *all* typical behavior patterns. So this good randomness string leads to a successful simulation whenever the adversary is limited to flipping at most a fraction ϵ of the bits. Details follow.

6.2.1 Behavior patterns and hash mismatches

We start by formally defining the notions of behavior patterns and hash mismatches and proving that hash mismatches depend only on the behavior pattern and the randomness string.

Definition 6.3 (Behavior pattern). Let Γ be a (possibly partial) run of the coding scheme $\tilde{\Pi}$ (determined by the randomness string, the simulated noiseless protocol π , the inputs (x, y) of the parties and the noise pattern). The behavior pattern \mathcal{P} of Γ consists of the following information:

1. The number of iterations in Γ .
2. The noise pattern in Γ (that is, the communication rounds in Γ in which the channel flipped a bit).
3. The iterations in Γ in which no block was added to T_A and the iterations in Γ in which no block was added to T_B .
4. For each of the iterations in Γ in which no block was added to T_A , a bit saying whether Alice made a transition on Line 25, a bit saying whether Alice returned to MP1 on Line 27 and a bit saying whether Alice returned to MP2 on Line 30. Similarly, for each of the iterations in Γ in which no block was added to T_B , a bit saying whether Bob made a transition on Line 25, a bit saying whether Bob returned to MP1 on Line 27 and a bit saying whether Bob returned to MP2 on Line 30.

Definition 6.4 (Hash mismatch). Let $i \in [R_{\text{total}}]$ be some iteration, let S be the seed used at iteration i , and let $k_A, |T_A|, N'(T_A), \text{MP1}_A$ and MP2_A (respectively, $k_B, |T_B|, N'(T_B), \text{MP1}_B$ and MP2_B) be the values of the variables of Alice (respectively, Bob) at the beginning of iteration i . Let $e \in \{0, 1\}^{|T_A|}$ be the vector that indicates the locations of the adversarial errors in the communication rounds in which the bits of T_A were transmitted. We say that a hash mismatch occurred at iteration i if at least one of the following occurred at iteration i .

1. $k_A \neq k_B$ but $\text{hash}_S(k_A) = \text{hash}_S(k_B)$.

2. $e \neq 0$ but $h_S(e) = 0$.
3. $|T_A| \neq |T_B|$ but $h_S(|T_A|) = h_S(|T_B|)$.
4. $N'(T_A) \neq N'(T_B)$ but $h_S(N'(T_A)) = h_S(N'(T_B))$.
5. There exists $b \in \{1, 2\}$ such that $e[1, \text{MP}b_A] \neq 0$ but $h_S(e[1, \text{MP}b_A]) = 0$.
6. There exist $b, b' \in \{1, 2\}$ such that $\text{MP}b_A \neq \text{MP}b'_B$ but $h_S(\text{MP}b_A) = h_S(\text{MP}b'_B)$.
7. There exist $b, b' \in \{1, 2\}$ such that $N'(T_A[1, \text{MP}b_A]) \neq N'(T_B[1, \text{MP}b'_B])$ but $h_S(N'(T_A[1, \text{MP}b_A])) = h_S(N'(T_B[1, \text{MP}b'_B]))$.

The following claim says that if some iteration does not suffer from a hash mismatch then it does not suffer from a hash collision either.

Claim 6.5. *If an iteration of $\tilde{\Pi}$ does not suffer from a hash mismatch then it does not suffer from a hash collision.*

Proof. By Condition 1 of Definition 6.4 we readily have that if $k_A \neq k_B$ then $\text{hash}_S(k_A) \neq \text{hash}_S(k_B)$. Next we show that if $T_A \neq T_B$ then $\widetilde{\text{hash}}_S(T_A) \neq \widetilde{\text{hash}}_S(T_B)$. If $|T_A| \neq |T_B|$ or $N'(T_A) \neq N'(T_B)$ then by Conditions 3 and 4 of Definition 6.4 we have that $\widetilde{\text{hash}}_S(T_A) \neq \widetilde{\text{hash}}_S(T_B)$. Otherwise, if $|T_A| = |T_B|$ and $N'(T_A) = N'(T_B)$, then we have that $T_A \oplus T_B = e$. Due to the linearity of h_S Condition 2 of Definition 6.4 implies that in this case $\widetilde{\text{hash}}_S(T_A) \neq \widetilde{\text{hash}}_S(T_B)$. A similar argument using Conditions 5, 6 and 7 of Definition 6.4 shows that if $T_A[1, \text{MP}b_A] \neq T_B[1, \text{MP}b'_B]$, for some $b, b' \in \{1, 2\}$, then $\widetilde{\text{hash}}_S(T_A[1, \text{MP}b_A]) \neq \widetilde{\text{hash}}_S(T_B[1, \text{MP}b'_B])$. \square

Finally, we show that hash mismatches depend only on the behavior pattern and the randomness string.

Claim 6.6. *Given a string $R \in \{0, 1\}^{R_{\text{total}} \cdot s}$ and a behavior pattern \mathcal{P} of a (possibly partial) run Γ that uses the string R as the random string sampled at Line 5, one can efficiently determine the iterations in Γ in which a hash mismatch occurred. In particular, whether a hash mismatch occurred at some iteration in Γ depends entirely on the string R and the behavior pattern \mathcal{P} .*

Proof. By definition it holds that whether a hash mismatch occurred at some iteration in Γ depends only on the string R , the noise pattern in Γ and the values of the variables k , $|T|$, $N'(T)$, MP1 and MP2 for both parties at the beginning of this iteration. The noise pattern is included in the description of \mathcal{P} , and it can be verified by induction on the number of iterations that the values of the variables k , $|T|$, $N'(T)$, MP1 and MP2 for both parties depend only on the behavior pattern \mathcal{P} and can be efficiently computed given \mathcal{P} . \square

6.2.2 Existence of good randomness string

In this section we show the existence of a good random string R^* that can be used to derandomize the coding scheme $\tilde{\Pi}$. For this we shall use the notion of a typical behavior pattern defined as follows.

Definition 6.7 (Typical behavior pattern). We say that a behavior pattern \mathcal{P} is *typical* if the number of bit flips in the noise pattern of \mathcal{P} is at most $2\epsilon n$, the number of iterations in \mathcal{P} in which no block was added to T_A is at most $100\epsilon n$, and the number of iterations in \mathcal{P} in which no block was added T_B is at most $100\epsilon n$.

The following claim bounds the number of typical behavior patterns.

Claim 6.8. *There are at most $2^{900H(\epsilon)n}$ different typical behavior patterns.*

Proof. First note that there are at most $R_{\text{total}} \leq n$ possible values for the number of iterations in \mathcal{P} , and that there are at most $R_{\text{total}} \cdot (r + r_c) \leq 2n$ communication rounds in \mathcal{P} . Next observe that since the noise pattern has at most $2\epsilon n$ bit flips, then the number of different noise patterns is at most

$$\sum_{i=0}^{2\epsilon n} \binom{2n}{i} \leq 2\epsilon n \cdot \binom{2n}{2\epsilon n} \leq 2^{2H(\epsilon)n}.$$

Furthermore, since there are at most $100\epsilon n$ iterations in which no block was added to T_A , the number of different sets of such iterations is at most

$$\sum_{i=0}^{100\epsilon n} \binom{R_{\text{total}}}{i} \leq 100\epsilon n \cdot \binom{n}{100\epsilon n} \leq 2^{100H(\epsilon)n}.$$

Finally, for each iteration in which no block was added to T_A we keep 3 bits of information and so the number of different possibilities for the values of these bits is at most $2^{300\epsilon n}$.

Concluding, we have that the number of different typical behavior patterns is at most

$$n \cdot 2^{2H(\epsilon)n} \cdot \left(2^{100H(\epsilon)n}\right)^2 \cdot (2^{300\epsilon n})^2 \leq 2^{900H(\epsilon)n}.$$

□

Next we show that for every behavior pattern most randomness strings lead to at most ϵn hash mismatches.

Claim 6.9. *Let \mathcal{P} be a behavior pattern and let R be a random string sampled as in Line 5. Then with probability at least $1 - 2^{-\Omega(c' \log(1/\epsilon)\epsilon n)}$ the number of iterations suffering from hash mismatches determined by \mathcal{P} and R is at most ϵn .*

Proof. Suppose first that R is a uniform random binary string in $\{0, 1\}^{R_{\text{total}} \cdot s}$. In this case, since the output length of the hash functions is $o = c' \log(1/\epsilon)$ and since there are only constant number of conditions in Definition 6.4, the probability that a hash mismatch occurs at some iteration i is at most $2^{-\Omega(c' \log(1/\epsilon))}$. Consequently, the probability that more than ϵn iterations suffer from a hash mismatch is at most

$$\binom{n}{\epsilon n} \cdot 2^{-\Omega(c' \log(1/\epsilon)\epsilon n)} \leq 2^{-\Omega(c' \log(1/\epsilon)\epsilon n)},$$

where the inequality holds for sufficiently large constant c' .

In our case R is sampled from a δ -biased distribution for $\delta = 2^{-\Theta(n\epsilon/r)}$ and consequently the probability that more than ϵn iterations suffer from a hash mismatch is at most

$$2^{-\Omega(c'\epsilon \log(1/\epsilon)n)} + 2^{-\Theta(n\epsilon/r)} = 2^{-\Omega(c' \log(1/\epsilon)\epsilon n)} + 2^{-\Theta(\sqrt{\log(1/\epsilon)\epsilon n})} = 2^{-\Omega(c'\epsilon \log(1/\epsilon)n)}.$$

□

Claims 6.8 and 6.9 above imply the existence of a single random string R^* that leads to at most ϵn hash mismatches for all typical behavior patterns.

Corollary 6.10. *For sufficiently large constant c' , there is a string $R^* \in \{0, 1\}^{R_{\text{total}} \cdot s}$ such that for every typical behavior pattern \mathcal{P} the number of iterations suffering from hash mismatches determined by \mathcal{P} and R^* is at most ϵn .*

Finally, we show that when the coding scheme $\tilde{\Pi}$ is run with the random string R^* guaranteed by the above corollary then the number of iterations suffering from *hash collisions* is at most ϵn .

Claim 6.11. *Let $R^* \in \{0, 1\}^{R_{\text{total}} \cdot s}$ be a string such that for every typical behavior pattern \mathcal{P} the number of iterations suffering from hash mismatches determined by \mathcal{P} and R^* is at most ϵn . Let Γ be a run of $\tilde{\Pi}$ that uses the string R^* as the random string sampled at Line 5 and has at most ϵ fraction of adversarial errors. Then at most ϵn iterations in Γ suffer from a hash collision.*

Proof. If Γ has a typical behavior pattern then by our assumption we have that R^* leads to at most ϵn iterations in Γ suffering from hash mismatches. By Claim 6.5 this implies in turn that at most ϵn iterations in Γ suffer from a hash collision. Therefore it suffices to show that Γ has a typical behavior pattern.

Suppose in contradiction that Γ has a non-typical behavior pattern \mathcal{P} . Let $t \in [R_{\text{total}}]$ be the first iteration in Γ such that the number of iterations $i \in [t]$ in which no block was added to T_A is more than $65\epsilon n$ or the number of iterations $i \in [t]$ in which no block was added to T_B is more than $65\epsilon n$. Let \mathcal{P}' be the (partial) behavior pattern obtained by restricting \mathcal{P} to the first t iterations. Then \mathcal{P}' is a typical behavior pattern and consequently by our assumption we have that the number of iterations suffering from hash mismatches determined by \mathcal{P}' and R^* is at most ϵn . Furthermore, since \mathcal{P} and \mathcal{P}' agree on the first t iterations we have that the number of iterations suffering from hash mismatches determined by \mathcal{P} and R^* among the first t iterations is at most ϵn . By Claim 6.5 this implies in turn that at most ϵn iterations $i \in [t]$ in Γ suffer from a hash collision which contradicts Claim 6.2. \square

6.2.3 Completing the proof of Lemma 3.3

We are now ready to complete the proof of the main result in this section.

Proof of Lemma 3.3. Corollary 6.10 and Claim 6.11 guarantee the existence of a string $R^* \in \{0, 1\}^{R_{\text{total}} \cdot s}$ such that in any run of the coding scheme $\tilde{\Pi}$ that uses the string R^* as the random string sampled at Line 5 and has at most ϵ fraction of adversarial errors, the number of iterations suffering from a hash collision is at most ϵn . By Lemma 6.1 this implies in turn that any run of the coding scheme $\tilde{\Pi}$ that uses R^* as the random string sampled at Line 5 and has at most ϵ fraction of adversarial errors successfully simulates the noiseless protocol π .

To show that $\tilde{\Pi}$ has the required rate note that the total number of bits communicated during the execution of $\tilde{\Pi}$ is

$$\begin{aligned} R_{\text{total}} \cdot (r + r_c) &= \left(\frac{n}{r} + \Theta(n\epsilon) \right) \cdot r \cdot \left(1 + \frac{r_c}{r} \right) \\ &= n \cdot \left(1 + \Theta(r\epsilon) \right) \cdot \left(1 + \frac{r_c}{r} \right) \\ &= n \cdot \left(1 + \Theta\left(r\epsilon + \frac{r_c}{r} \right) \right). \end{aligned}$$

Due to our choice of $r = \Theta(\sqrt{\log(1/\epsilon)}/\epsilon)$ and $r_c = \Theta(\log(1/\epsilon))$ the above implies in turn that the number of bits communicated in the coding scheme $\tilde{\Pi}$ is $n \cdot (1 + \Theta(\sqrt{\epsilon \log(1/\epsilon)}))$. So the rate of $\tilde{\Pi}$ is

$$1 - O(\sqrt{\epsilon \log(1/\epsilon)}) = 1 - O(\sqrt{H(\epsilon)}).$$

Finally, observe that one can find the string R^* by going over all pairs (\mathcal{P}, R) where \mathcal{P} is a typical behavior pattern and $R \in \{0, 1\}^{R_{\text{total}} \cdot s}$ is in the support of the δ -biased distribution for $\delta = 2^{-\Theta(n\alpha/r)}$ which requires $\Theta(\sqrt{\epsilon \log(1/\epsilon)})n$ random bits. Therefore, the number of possible strings R is at most $2^{O(n)}$. Furthermore, by Claim 6.8 there are at most $2^{O(n)}$ different typical behavior patterns \mathcal{P} . Therefore the total number of pairs (\mathcal{P}, R) one needs to check is at most $2^{O(n)}$. Finally, Claim 6.6 shows that for each such pair it takes $\text{poly}(n)$ time to verify whether the number of iterations suffering from a hash mismatch determined by this pair is at most ϵn , and we conclude that the total time this process takes is at most $2^{O(n)}$. \square

References

- [AGS13] S. Agrawal, R. Gelles, and A. Sahai. Adaptive protocols for interactive communication. Manuscript, arXiv:1312.4182 (cs.DS), 2013.
- [AGHP92] N. Alon, O. Goldreich, J. Håstad, and R. Peralta. Simple constructions of almost k -wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992.
- [BKN14] Z. Brakerski, Y. T. Kalai, and M. Naor. Fast interactive coding against adversarial noise. *J. ACM*, 61(6):35:1–35:30, 2014.
- [Bra12] M. Braverman. Towards deterministic tree code constructions. *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCS '12*, pp. 161–167. ACM, 2012.
- [BE14] M. Braverman and K. Efremenko. List and unique coding for interactive communication in the presence of adversarial noise. *Proceedings of the IEEE Symposium on Foundations of Computer Science, FOCS '14*, pp. 236–245. 2014.
- [BGMO15] M. Braverman, R. Gelles, J. Mao, and R. Ostrovsky. Coding for interactive communication correcting insertions and deletions. Preprint arXiv:1508.00514, 2015.
- [BR14] M. Braverman and A. Rao. Toward coding for maximum errors in interactive communication. *Information Theory, IEEE Transactions on*, 60(11):7248–7255, 2014.
- [FK00] U. Feige and J. Kilian. Finding OR in a noisy broadcast network. *Information Processing Letters*, 73(1):69–75, 2000.
- [For65] G. D. Forney. Concatenated codes. Tech. Rep. 440, Massachusetts Institute of Technology. Research Laboratory of Electronics, 1965.
- [FGOS15] M. Franklin, R. Gelles, R. Ostrovsky, and L. J. Schulman. Optimal coding for streaming authentication and interactive communication. *Information Theory, IEEE Transactions on*, 61(1):133–145, 2015.
- [Gel15] R. Gelles. Coding for interactive communication: A survey, 2015.

- [GMS14] R. Gelles, A. Moitra, and A. Sahai. Efficient coding for interactive communication. *Information Theory, IEEE Transactions on*, 60(3):1899–1913, 2014.
- [GH14] M. Ghaffari and B. Haeupler. Optimal Error Rates for Interactive Coding II: Efficiency and List Decoding. *Proceedings of the IEEE Symposium on Foundations of Computer Science*, FOCS '14, pp. 394–403. 2014.
- [GHS14] M. Ghaffari, B. Haeupler, and M. Sudan. Optimal error rates for interactive coding I: Adaptivity and other settings. *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC '14, pp. 794–803. ACM, 2014.
- [Gil52] E. N. Gilbert. A comparison of signalling alphabets. *Bell System Technical Journal*, 31:504–522, 1952.
- [Hae14a] B. Haeupler. Interactive Channel Capacity Revisited. *Proceedings of the IEEE Symposium on Foundations of Computer Science*, FOCS '14, pp. 226–235. 2014.
- [Hae14b] B. Haeupler. Interactive channel capacity revisited, 2014. Long version of [Hae14a], [online:] <http://arxiv.org/abs/1408.1467>.
- [Ham50] R. W. Hamming. Error detecting and error correcting codes. *Bell System technical journal*, 29(2):147–160, 1950.
- [KR13] G. Kol and R. Raz. Interactive channel capacity. *STOC '13: Proceedings of the 45th annual ACM Symposium on Theory of Computing*, pp. 715–724. ACM, 2013.
- [NN93] J. Naor and M. Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, 1993.
- [RS60] I. S. Reed and G. Solomon. Polynomial codes over certain finite fields. *SIAM Journal of the Society for Industrial and Applied Mathematics*, 8(2):300–304, 1960.
- [Sch92] L. J. Schulman. Communication on noisy channels: a coding theorem for computation. *Foundations of Computer Science, Annual IEEE Symposium on*, pp. 724–733, 1992.
- [Sch93] L. J. Schulman. Deterministic coding for interactive communication. *STOC '93: Proceedings of the twenty-fifth annual ACM Symposium on Theory of Computing*, pp. 747–756. ACM, 1993.
- [Sch96] L. J. Schulman. Coding for interactive communication. *IEEE Transactions on Information Theory*, 42(6):1745–1756, 1996.
- [Sch03] L. J. Schulman. A postscript to “coding for interactive communication”. [Online:] <http://www.cs.caltech.edu/~schulman/Papers/intercodingpostscript.txt>, 2003. Based on joint work with Will Evans and Michael Klugerman.
- [Sha48] C. E. Shannon. A mathematical theory of communication. *ACM SIGMOBILE Mobile Computing and Communications Review*, 5(1):3–55, 2001. Originally appeared in *Bell System Tech. J.* 27:379–423, 623–656, 1948.
- [Var57] R. R. Varshamov. Estimate of the number of signals in error correcting codes. *Doklady Akademii Nauk*, pp. 739–741, 1957.

- [Yao79] A. C.-C. Yao. Some complexity questions related to distributive computing (preliminary report). *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing*, STOC '79, pp. 209–213. ACM, 1979.

Appendix

A The efficient randomized interactive coding scheme of [Hae14b]

Algorithm 1 Coding Scheme for Oblivious Adversarial Channels [Hae14b]

```

1:  $\Pi \leftarrow n$ -round protocol to be simulated + final confirmation steps
2:  $hash \leftarrow$  inner product hash family with  $o = \Theta(1)$  and  $s = \Theta(n)$ 

3: Initialize Parameters:  $r_c \leftarrow \Theta(1)$ ;  $r \leftarrow \lceil \sqrt{\frac{r_c}{\epsilon}} \rceil$ ;  $R_{total} \leftarrow \lceil n/r + 65n\epsilon \rceil$ ;  $T \leftarrow \emptyset$ 
4: Reset Status:  $k, E, v1, v2 \leftarrow 0$ 

5:  $R \leftarrow$  Random string of length  $R_{total} \cdot s$  (can be constructed by exchanging  $\Theta(n\sqrt{\epsilon})$  random bits and expanding them to a  $\delta$ -bias string of the needed length using [AGHP92, NN93], with bias  $\delta = 2^{-\Theta(\frac{r}{o})}$ )

6: for  $R_{total}$  iterations do

7:    $k \leftarrow k + 1$ ;  $\tilde{k} \leftarrow 2^{\lceil \log_2 k \rceil}$ ;  $MP1 \leftarrow \tilde{k}r \lfloor \frac{T}{\tilde{k}r} \rfloor$ ;  $MP2 \leftarrow MP1 - \tilde{k}r$  ▷ Verification Phase
8:    $S \leftarrow s$  new preshared random bits from  $R$ 
9:   Send  $(hash_S(k), hash_S(T), hash_S(T[1, MP1]), hash_S(T[1, MP2]))$ 
10:  Receive  $(H'_k, H'_T, H'_{MP1}, H'_{MP2})$ ;
11:   $(H_k, H_T, H_{MP1}, H_{MP2}) \leftarrow (hash_S(k), hash_S(T), hash_S(T[1, MP1]), hash_S(T[1, MP2]))$ 

12:  if  $H_k \neq H'_k$  then
13:     $E \leftarrow E + 1$ 
14:  else
15:    if  $H_{MP1} \in \{H'_{MP1}, H'_{MP2}\}$  then
16:       $v1 \leftarrow v1 + 1$ 
17:    else if  $H_{MP2} \in \{H'_{MP1}, H'_{MP2}\}$  then
18:       $v2 \leftarrow v2 + 1$ 

19:  if  $k = 1$  and  $H_T = H'_T$  and  $E = 0$  then ▷ Computation Phase
20:    continue computation and transcript  $T$  for  $r$  steps
21:    Reset Status:  $k, E, v1, v2 \leftarrow 0$ 
22:  else
23:    do  $r$  dummy communications

24:  if  $2E \geq k$  then ▷ Transition Phase
25:    Reset Status:  $k, E, v1, v2 \leftarrow 0$ 
26:  else if  $k = \tilde{k}$  and  $v1 \geq 0.4 \cdot \tilde{k}$  then
27:    rollback computation and transcript  $T$  to position  $MP1$ 
28:    Reset Status:  $k, E, v1, v2 \leftarrow 0$ 
29:  else if  $k = \tilde{k}$  and  $v2 \geq 0.4 \cdot \tilde{k}$  then
30:    rollback computation and transcript  $T$  to position  $MP2$ 
31:    Reset Status:  $k, E, v1, v2 \leftarrow 0$ 
32:  else if  $k = \tilde{k}$  then
33:     $v1, v2 \leftarrow 0$ 

34: Output the outcome of  $\Pi$  corresponding to transcript  $T$ 

```
