

Binary Codes with Resilience Beyond 1/4 via Interaction

Klim Efremenko
Ben-Gurion University
Be'er Sheva, Israel
klimefrem@gmail.com

Gillat Kol
Princeton University
Princeton, USA
gillat.kol@gmail.com

Raghuvansh R. Saxena
Microsoft Research
Cambridge, USA
raghuvansh.saxena@gmail.com

Zhijun Zhang
Princeton University
Princeton, USA
zhijunz@princeton.edu

Abstract—In the *reliable transmission problem*, a sender, Alice, wishes to transmit a bit-string x to a remote receiver, Bob, over a binary channel with adversarial noise. The solution to this problem is to encode x using an *error correcting code*. As it is long known that the distance of binary codes is at most $1/2$, reliable transmission is possible only if the channel corrupts (flips) at most a $1/4$ -fraction of the communicated bits.

We revisit the reliable transmission problem in the two-way setting, where both Alice and Bob can send bits to each other. Our main result is the construction of *two-way error correcting codes* that are resilient to a constant fraction of corruptions strictly larger than $1/4$. Moreover, our code has constant rate and requires Bob to only send one short message. We mention that our result resolves an open problem by Haeupler, Kamath, and Velingker [APPROX-RANDOM, 2015] and by Gupta, Kalai, and Zhang [STOC, 2022].

Curiously, our new two-way code requires a fresh perspective on classical error correcting codes: While classical codes have only one distance guarantee for all pairs of codewords (*i.e.*, the minimum distance), we construct codes where the distance between a pair of codewords depends on the “compatibility” of the messages they encode. We also prove that such codes are necessary for our result.

Index Terms—error correcting code, interactive communication, noise resilience

I. INTRODUCTION

As errors are everywhere, essentially any telecommunication system crucially uses *error correcting codes*. Classical “one-way” error correcting codes date back to the 40’s [1] and are designed to solve the *reliable transmission problem*, where a sender, Alice, wishes to send a message x to a remote receiver, Bob, but she can only communicate with him over a noisy one-way channel that corrupts some of her communication.

As the price of interaction goes down, systems are becoming more interactive. In this paper we study *two-way error correcting codes*, that are designed to solve the same problem, but work assuming a two-way channel instead of a one-way channel, allowing the parties to interact. Specifically, we consider the reliable transmission problem, where Alice and Bob are connected by a pair of *binary channels* with *adversarial corruption noise* (bit flips), one in each direction.

Klim Efremenko is supported by the Israel Science Foundation (ISF) through grant No. 1456/18 and European Research Council Grant number: 949707. Gillat Kol is supported by a National Science Foundation CAREER award CCF-1750443 and by a BSF grant No. 2018325.

The two most important parameters in the study of error correcting codes are the (relative) *distance* and the *rate* of the code. Plotkin showed in the 60’s [2] that the minimum (or even average) relative distance of a binary error correcting code is at most $\frac{1}{2}$, which implies that binary codes can be resilient to up to $\frac{1}{4}$ fraction of adversarial errors. *Can interaction improve the error resilience of binary codes?*

A. Our Result

1) *Main Result: Binary Two-Way Codes with Error Resilience $> 1/4$* : The main result of this paper is Theorem I.1 (see Theorem V.1 for a formal statement), that gives a positive answer to the above question, resolving an open problem by Haeupler, Kamath, and Velingker (Section 6 in [3]) and by Gupta, Kalai, and Zhang (Section 1.1 in [4]).

Theorem I.1 (Main, Informal). *There exists a constant rate, deterministic, binary two-way error correcting code with error resilience $\frac{1}{4} + 10^{-5}$, where Bob sends a single message.*

We mention that in the two-way code we construct, Bob sends a single message whose length is less than 2% of the total communication, thus showing that even a minimal amount of interaction can already improve the noise resilience of binary codes, while keeping their rate constant. Finding the maximal noise tolerance of (constant rate or even zero-rate) binary two-way codes is an intriguing question we leave open.

2) *Impossibility for Equally-Spaced Codes*: To construct the binary two-way code promised by Theorem I.1, we design a new binary one-way code where the guaranteed distance between certain carefully chosen pairs of “compatible” codewords is strictly greater than $\frac{1}{2}$ (at a high level, the adversary is more likely to want to confuse between these pairs of codewords). Recall, however, that by the *Plotkin bound*, the average distance of a binary code is at most $\frac{1}{2}$, thus the distance between some of the other pairs of codewords is strictly smaller than $\frac{1}{2}$.

Theorem I.2 below (see Theorem VI.1 for a formal statement) shows that these types of codes are required to obtain our result, as solely using *equally-spaced codes*, where the distance between every pair of codewords is roughly equal, is insufficient for breaking the $\frac{1}{4}$ error resilience barrier. This is in contrast to the one-way setting where equally-spaced codes attain the maximum resilience (*e.g.*, random binary codes are

equally-spaced and have resilience of $\frac{1}{4}$). We also mention that all prior work concerned with the noise resilience of two-way channels and channels with feedback, surveyed in Section I-B, essentially only uses equally-spaced codes (see Section II-C).

Theorem I.2 (Informal). *The maximum error resilience of a binary two-way error correcting code (of any rate) that uses equally-spaced codes and has Bob sending a single message is $\frac{1}{4}$.*

a) Why non-equally-spaced: Intriguingly, when we started this project over a year ago, we believed that the answer to the above question should be negative, that is, that binary two-way codes cannot break the $\frac{1}{4}$ resilience barrier. In fact, we had a sketch of an impossibility result designed to show that any two-way code (say, with three messages, Alice, Bob, Alice) is essentially of the following form, and that codes of this form cannot have resilience better than $\frac{1}{4}$: On input x , Alice sends x encoded by a list-decodable code (for simplicity, assume lists are of size 2). Bob decodes to obtain two candidates x_1 and x_2 with the promise $x \in \{x_1, x_2\}$. Bob sends $\{x_1, x_2\}$ encoded by a list-decodable code. Alice decodes to obtain two candidates of the form $\{x, x_3\}$ and $\{x, x_4\}$. Observe that after this second message, the parties are left with the following communication task F : Bob knows $\{x_1, x_2\}$ such that $x \in \{x_1, x_2\}$, Alice knows $x, \{x_3, x_4\}$ such that $\{x_1, x_2\} \subseteq \{x, x_3, x_4\}$. Bob wishes to learn x . The two-way code then solves F with a single message from Alice to Bob.

Our strategy for proving an impossibility result was to show the following two lemmas: (1) A round elimination lemma saying that a reliable transmission protocol with resilience γ implies a one-message protocol for the task F with resilience γ . (2) A lemma showing that F cannot be solved with a single message of constant size over the *noiseless* channel.¹ We observed that under the assumption that Alice uses an equally-spaced code C to solve F ,² the second lemma implies that the distance between any two messages she may send is at most $\frac{1}{2}$.³ In this case F cannot have resilience better than $\frac{1}{4}$, and using the first lemma, we get our impossibility result. However, as should be expected, despite our best effort, we were unsuccessful in showing that the assumption is without loss of generality... Nevertheless, under this assumption, we were able to formalize this impossibility sketch, yielding Theorem I.2 (see additional discussion in Section II-B).

We remark that most proofs are omitted in this version. Interested readers are referred to the full version [7].

¹The task F is interesting on its own right. In Section VI we show that if $x, x_1, x_2, x_3, x_4 \in [N]$, then the one-way communication complexity of F is $\Theta(\log \log N)$. We also mention that F is very related to compression efforts by [5], [6].

²That is, the set C of all possible messages by Alice for all possible inputs, forms an equally-spaced code.

³In more detail, for equally-spaced codes, the distance between every pair of codewords is roughly the minimum distance, and the minimum distance of codes of super-constant size is $\frac{1}{2}$.

B. Related Work

1) *Two-Way Erasure Codes:* In a recent work, Gupta, Kalai, and Zhang [4] study two-way error correcting codes over the adversarial binary *erasure channel*, where the adversary may replace some of the sent bits by ‘?’. Their main result is a code that is resilient to a $\frac{3}{5}$ fraction of adversarial errors, improving on the noise tolerance of the one-way binary erasure channel that is known to be $\frac{1}{2}$. Gupta and Zhang [8] give a two-way code over the same channel that is also of constant rate. We mention that the two-way coding schemes of [4], [8] exchange (almost) linear number of messages and are generally very different than ours.

[4] also give an upper bound of $\frac{2}{3}$ on the maximum tolerance of the two-way adversarial binary erasure channel, and an upper bound of $\frac{2}{7}$ on the maximum tolerance of the two-way adversarial binary channel (the model assumed by our work). As mentioned above, bridging the gap between our lower bound on the noise tolerance and their upper bound is a great problem.

2) Reliable Exchange and Interactive Coding:

a) Reliable exchange: In the *reliable exchange problem*, two parties, Alice, holding a private input x , and Bob, holding a private input y , communicate over the two-way binary adversarial channel with the goal of learning each other’s input. Observe that the reliable transmission problem is at least as hard as the reliable exchange problem, in the sense that a transmission protocol with resilience θ implies an exchange protocol with resilience $\frac{\theta}{2}$: Alice sends x using the transmission protocol, then Bob sends y using the transmission protocol. Now, if an adversary corrupts at most $\frac{\theta}{2}$ fraction of the total communication, it also corrupts at most θ fraction of each transmission and both transmissions go through. Since one-way codes solve the transmission problem with error resilience $\frac{1}{4}$, the exchange problem is easily solvable with error resilience $\frac{1}{8}$.

Efremenko, Kol, and Saxena [9] show how to go beyond $\frac{1}{8}$ and obtain an exchange protocol that is resilient to a $\frac{5}{39}$ fraction of adversarial errors with a constant number of rounds and constant overhead. As explained in Section II, we use [9] as a stepping stone towards our two-way code. The resilience constant was later improved by [10] from $\frac{5}{39}$ to $\frac{1}{6}$, which was known to be optimal [11], [12]. Note however, that the [10] protocol has linear overhead and many communication rounds.

b) Interactive coding: The reliable exchange problem (and therefore also reliable transmission) are special cases of the interactive coding problem: Given a two-party communication protocol Π that works assuming the noiseless channel, simulate Π by a protocol Π' that works over a noisy channel. The study of interactive coding was first suggested in seminal works by Schulman [13]–[15], and is now an active research area, see [16] for an excellent survey.

Observe that the transmission problem corresponds to the noiseless protocol Π where Alice sends her input x , and the exchange problem corresponds to Π where Alice sends her input x and Bob sends his input y . The exchange problem is also “complete” in the sense that after exchanging x and y

the parties can run any other protocol without communication. Thus, an exchange protocol resilient to θ fraction of errors implies an interactive coding scheme with the same resilience. Note however that this scheme may have a huge overhead.

Braverman and Rao [17], building on [15], gave an interactive code with constant overhead and optimal resilience of $\frac{1}{4}$ for the case where the alphabet set is large, and showed that it implies a binary interactive code with resilience $\frac{1}{8}$; see also [18] on asymmetric corruptions. [9] gave a binary interactive code with constant overhead that is resilient to a $\frac{5}{39}$ fraction of errors (that is, they showed that their reliable exchange scheme can be generalized to an interactive coding scheme).

The maximum resilience of interactive coding schemes was also studied for other channels, such as the erasure channel, the channel with feedback, and the insertion-deletion channel [12], [19]–[24]. Another channel that received quite a bit of attention in this context is the adaptive channel, where several parties may speak at the same round and collisions may occur⁴ [25]–[28].

3) *Reliable Transmission with Feedback*: The works surveyed so far consider two-way channels, but are inspired by classical results from the 60’s showing that the maximum resilience of one-way error correcting codes can be improved assuming the channel provides *feedback*. At a very high level, these two-way results work by implementing feedback (over channels with no build-in feedback) using interaction.

In more detail, the *feedback channel* allows Alice to communicate symbols to Bob, but upon receiving each sent symbol, Bob sends the received symbol back to Alice as feedback [29]. Alice can then use it when deciding what to send next. Note that it is typically assumed that Bob’s feedback is not corrupted by the channel.

Observe that any protocol that can be run over a two-way channel can also be run over the feedback version of the same channel: Given a communication protocol over the two-way channel, Alice can simulate the messages sent by Bob as she knows everything he knows (which, as Bob has no input, is just his received transcript). The two main differences between two-way channels and feedback channels are: (1) The noise in a feedback channel is one-way: while the communication from Alice to Bob may be noisy, the communication from Bob to Alice is noiseless. (2) For a feedback channel, the length of the communication is defined as the number of rounds where Alice communicates (Bob’s feedback rounds do not count towards the length of the protocol). In particular, the noise tolerance is measured as a fraction of Alice’s rounds.

While Shannon showed that feedback does not increase the maximum noise tolerance of stochastic channels [30], feedback can, in fact, increase the noise tolerance of adversarial channels [11], [31]–[33]. Specifically, it was shown by Berlekamp that the noise tolerance of the binary adversarial channel increases from $\frac{1}{4}$ to $\frac{1}{3}$ given feedback [11]. Haeupler, Kamath, and Velingker [3] considered the setting where the

⁴More formally, the communication order is not predetermined. In every round, each party may decide whether to send or listen depending on his input and received transcript.

feedback is partial, and showed that even if Alice receives feedback bits from Bob for an arbitrarily small constant fraction of her transmissions, resilience close to $\frac{1}{3}$ is possible.

Partial *noisy* feedback was considered by Wang, Qin, and Chang [34], who constructed a binary two-way code that is resilient to any constant fraction strictly smaller than 1 of adversarial erasures from Bob to Alice, but only up to $\frac{1}{2}$ fraction of adversarial erasures from Alice to Bob (*cf.* [4], where the *total* noise tolerance is strictly greater than $\frac{1}{2}$).

II. PROOF SKETCH

We now give a detailed overview of our main result, a 3-message (or a 3-step) protocol for the reliable transmission problem with error resilience strictly larger than $\frac{1}{4}$. Recall that the maximum distance possible using a binary code depends on the number of codewords it has. Precisely, if one wants an even⁵ number C of codewords, then the maximum distance⁶ one can get equals $\frac{C}{2(C-1)}$. Put another way, this means that if Alice wants to send one of C strings to Bob over a *one-way* channel that corrupts (flips) a γ fraction of the symbols sent, she can do this if and only if $\gamma < \frac{C}{4(C-1)}$. As C goes to infinity, this bound becomes $\gamma < \frac{1}{4}$, which is the error resilience of the one-way channel. However, for smaller values of C , the error resilience is much higher: It is $\frac{1}{2}$ when $C = 2$, as demonstrated by the codewords $00 \cdots 0$ and $11 \cdots 1$, and $\frac{1}{3}$ when $C = 4$, *etc.*

As a prelude to our main result, we first overview the result of [9]. We mention that [9] were concerned with the reliable exchange problem and the interactive coding question (see Section I-B2), and that we are presenting a simplified and slightly modified version of their scheme for the restricted case where Bob has no input and the channel has one-way error, meaning that the adversary can only corrupt symbols sent from Alice to Bob, but cannot corrupt the ones sent from Bob to Alice⁷.

A. The [9] Result: Beating $\frac{1}{4}$ Given One-Way Error

At an extremely high level, the [9] protocol works by using two-way communication (with one-way error) to reduce the effective number of codewords that Alice and Bob need to consider, and then uses the higher resilience guarantees achievable by codes with small C to get an improved error resilience.

In more detail, [9] describe a 3-message protocol where in the first step, Alice sends to Bob an encoding of her input x using a *list-decodable* error correcting code. Bob’s goal in this step is not to recover x exactly, which would limit the error

⁵The bound is slightly different for odd numbers, but exhibits the same phenomenon.

⁶Recall that the distance of a code is defined as the minimum distance between any two codewords. However, the bound stated here even holds for the average distance between pairs of codewords.

⁷The one-way error setting differs from the noiseless feedback setting (see Section I-B3) in two respects: (1) It allows Bob to send any value it wishes as feedback. (2) The length of the protocol is the total number of bits communicated by *both* parties.

resilience to $\frac{1}{4}$, but instead to compute a set of size 2,⁸ say $S = \{x_1, x_2\}$ such that $x \in S$. This weaker guarantee, known as list-decoding, allows the parties to tolerate strictly higher than $\frac{1}{4}$ errors in this step.

Then, in the second step, Bob sends the set S to Alice, and the one-way error guarantee implies that Alice will always receive the set S correctly. Additionally, the one-way error guarantee also means that the second step can be arbitrarily short.

Overall, these two steps guarantee that before the third step begins, both Alice and Bob agree on a set S of size 2 that contains x , and in the third step Alice only needs to tell Bob which of the two elements in S is her correct input. As S is of size 2, this can be done with a high error resilience using the codes with *small* C discussed above.

B. Challenges in Going from One-Way to Two-Way Error

The [9] protocol described above breaks down once the adversary is allowed to corrupt the bits sent from Bob to Alice in the second step, in addition to those sent from Alice to Bob in the first and the third steps. This poses the following challenges:

a) *No unique-decoding the second step:* First and foremost, the guarantee that Alice and Bob agree on a small set S containing x that [9] crucially relied on no longer holds in the two-way error case. As the number of possible sets Bob could send in this step is large, by corrupting just a $\frac{1}{4}$ fraction of this step, the adversary can make sure that the set Alice decodes is different from the set Bob sent. This is fatal to the protocol, as without the common knowledge of S , Bob has no way to interpret the message sent by Alice in the third step, which means he cannot output x .

b) *List-decoding the second step seems insufficient:* A possible remedy that one might consider for the foregoing problem is to have Bob send the set S to Alice using a list-decodable error correcting code similarly to the way Alice sent him x in the first step. This would allow Alice to have two sets T_1, T_2 , both containing x such that one of the two sets equals S . One may then hope to suitably adapt the third step of the protocol to work with this weaker guarantee while still using codes with small C .

It turns out that such an adaptation is impossible. As we show in Section VI, specifically, in Lemma VI.4, the third step in any such adaptation must use a super-constant number C of codewords (also see discussion after Theorem I.2). However, a super-constant C means that the error resilience guarantee reduces to $\frac{1}{4}$, and we get no improvement.

c) *List-decoding the second step is impossible:* Not only does it seem hard to get a protocol that works given a list-decoding guarantee for the second step (*i.e.*, assuming Alice obtains two sets T_1, T_2 such that one of them is S), but it is actually *impossible* for Alice to obtain such sets T_1, T_2 . The reason is that any protocol for the transmission problem

⁸We mention that the bounds one gets for list-decoding implies that the set S needs to be of size at least 3 to get guarantees better than $\frac{1}{4}$. Nonetheless, for the sake of this sketch, we shall stick to sets of size 2.

with error resilience larger than $\frac{1}{4}$ must have Bob speak in at most a $\frac{1}{4}$ fraction of the communication rounds. This is because of a classical result by [11], [29] showing that even when Alice knows everything Bob knows, the maximum error resilience possible (as a fraction of Alice’s rounds) is $\frac{1}{3}$ (see Section I-B3). Therefore, if Bob speaks in more than a $\frac{1}{4}$ fraction of the rounds, Alice speaks in at most $\frac{3}{4}$ fraction of the rounds. Now, even if the adversary does not corrupt Bob at all and Alice knows everything Bob knows, the maximum error resilience is at most $\frac{1}{3} \cdot \frac{3}{4} = \frac{1}{4}$.⁹

However, the facts that Bob speaks in at most a $\frac{1}{4}$ fraction of the rounds, and that the adversary can corrupt strictly more than a $\frac{1}{4}$ fraction of the rounds, mean that the adversary can, if he wants, corrupt Bob’s transmission entirely (that is, corrupt the communication from Bob to Alice in all 3 steps). Since Bob is effectively shut off with this attack, Alice will not be able to obtain sets T_1, T_2 such that one of them is S . Thus, even if one could adapt the third step of the [9] protocol to work when Alice has a list of 2 sets (which has its own challenges), one cannot guarantee that Alice will have this list.

C. Our Main Idea: Non-Equally-Spaced Codes

Although the challenges mentioned above make a pretty solid case for an impossibility result, we were able to construct a two-way scheme with resilience better than $\frac{1}{4}$ using new ideas. Our main new idea, and where our work differs most significantly from all prior work, is the use of codes where not all codewords are *equally-spaced*. Note that the distance between any two codewords sent by Alice roughly captures the amount of corruptions that the adversary needs to invest to confuse Bob between those two codewords. We observe that, in our protocol, some pairs of codewords are more prone to be corrupted by the adversary than others, and ensure that such pairs have a high distance to begin with, thereby implying that the adversary needs a high number of corruptions to confuse Bob.

Implementing this idea is not simple, as whatever code we come up with still needs to obey the aforementioned bound of $\frac{C}{2(C-1)}$ on the average distance between codewords. Thus, if we want to have a higher distance between some pairs of codewords, we must also have a lower distance between some other pair of codewords, and we need to ensure that these low distance pairs are carefully chosen to not affect the overall error resilience of our protocol.

Before explaining which pairs of codewords can have a lower than average distance and which pairs need to be farther apart, we mention that such a careful analysis of the distance is a novelty of our paper. Most prior work in the area of error resilience, and the area of binary codes in general, only uses one measure, the minimum distance between two codewords, when understanding the distance properties of a given code. Moreover, many constructions have a lot of “symmetry”, *e.g.*, picking codewords at random, that implies the distance

⁹We mention that this idea can also be used to get an upper bound of $\frac{2}{7}$ on the maximum error resilience, without any restriction on the number of rounds Bob speaks in, see [4].

between any two pairs of codewords is more or less the same. This also holds for the codes with small C that were used in the [9] protocol and followup work (see Section I-B), with [10] being a minor exception as the codes in [10], albeit not equally-spaced overall, can be seen as a union of a small number (four) of codes that are equally-spaced, and still have the issues described above.

D. Distance Requirements for Pairs of Codewords

We now explain why our protocol benefits if certain pairs of codewords are farther apart than other pairs of codewords. Recall that our protocol has 3 steps, and let L_1, L_2, L_3 be the lengths of these steps, and let $T = L_1 + L_2 + L_3$. Our protocol sticks to the framework in Sections II-A and II-B for the first two steps, using equally-spaced codes for these two steps.

In this section we are going to make the following two simplifying assumptions: first is that at the end of the first step Bob has a list of size two instead of three, second one is that at the second round adversary can either completely corrupt message to another codeword by investing $L_2/2$ errors or not to corrupt this message at all. Removing this assumptions is not trivial and we will discuss it in details in Section II-F.

As a result, if x is the input that Alice starts the protocol with, then at the end of the first step Bob has a set S of size 2 that is guaranteed to contain x . Moreover, Bob knows that Alice's encodings for the two elements of S are at least $L_1/2$ apart in Hamming distance (as the list-decoding radius is $\frac{1}{2}$).

For the second step, Alice gets a set T of size 2 that contains x , as she knows that Bob encoded a set that contains x , but is otherwise arbitrary. As Alice knows both T and x , this is equivalent to her knowing an *ordered* pair (x, a) , where x is her input and a is the element in T that is different from her input. Moreover, Bob knows that if Alice indeed got a set $T \neq S$, then the adversary must have invested at least $L_2/2$ corruptions in the second step.

For the third step, Alice encodes an ordered pair (x, a) , and sends the encoding to Bob. As Bob knows a set S of size 2 that contains x , say $S = \{x, x'\}$, where $x \neq x'$, he knows that Alice either encoded a pair of the form (x, a) or a pair of the form (x', b) , for some inputs $a \neq x, b \neq x'$, and his goal is to use the messages he received in the first and third step to figure out whether the pair sent was of the form (x, a) or the form (x', b) . Note that he does not need to know what the pair was exactly as his goal is to output Alice's input (and not the pair).

If the target error resilience is $\frac{1}{4} + \theta$, Bob can achieve this goal only if it is impossible for two pairs, one of the form (x, a) and the other of the form (x', b) , to give rise to the same messages (m_1, m_3) in the first and the third steps, with at most $(\frac{1}{4} + \theta) \cdot T$ corruptions. With this in mind, let us now look at the set of messages in the first and third steps that the adversary can generate using $(\frac{1}{4} + \theta) \cdot T$ corruptions from pairs of the form (x, a) .

Let ECC_1 and ECC_3 denote the error correcting codes used by Alice in the first and third steps respectively, so that the messages Alice sends are $ECC_1(x)$ and $ECC_3(x, a)$. As

explained above, we either have $a = x'$ or the adversary spent at least $L_2/2$ corruptions in the second step. Let Δ denote the Hamming distance, and let

$$f(x, x', m_1, m_3) = \Delta(ECC_1(x), m_1) + \Delta(ECC_3(x, x'), m_3).$$

The above discussion implies that all the pairs of messages (m_1, m_3) that the adversary can generate using $(\frac{1}{4} + \theta) \cdot T$ corruptions from the pair (x, a) must either satisfy $a = x'$ (that is, Alice's second element is one of the elements in Bob's list) and

$$f(x, x', m_1, m_3) \leq \left(\frac{1}{4} + \theta\right) \cdot T, \quad (1)$$

or satisfy $a \notin \{x, x'\}$ (that is, Alice's second element is not one of the elements in Bob's list) and

$$f(x, a, m_1, m_3) \leq \left(\frac{1}{4} + \theta\right) \cdot T - \frac{L_2}{2}, \quad (2)$$

Using these inequalities and the fact that Alice's encodings for x and x' in the first step are $L_1/2$ apart in Hamming distance, one gets that the following guarantees on the distances between codewords of ECC_3 are both necessary and sufficient for a protocol to have error resilience $\frac{1}{4} + \theta$:

- 1) For Bob to not be confused between messages for the pairs (x, x') and (x', x) we need there not to be a pair (m_1, m_3) that satisfies both $f(x, x', m_1, m_3) \leq (\frac{1}{4} + \theta) \cdot T$ and $f(x', x, m_1, m_3) \leq (\frac{1}{4} + \theta) \cdot T$. The reason is that if Alice has (x, x') then her input is x and her second element is in Bob's list. Similarly, if Alice has (x', x) then her input is x' and her second element is in Bob's list. Thus, on both pairs we can use (1).

Take m_1 to be the middle point between $ECC_1(x)$ and $ECC_1(x')$ (a point with equal Hamming distance from both) and m_3 be the middle point between $ECC_3(x, x')$ and $ECC_3(x', x)$. Very roughly, we claim that these are the "worst" m_1 and m_3 .

By summing the two f inequalities for this m_1 and m_3 we get

$$\begin{aligned} & \Delta(ECC_1(x), ECC_1(x')) + \Delta(ECC_3(x, x'), ECC_3(x', x)) \\ & \leq 2 \left(\frac{1}{4} + \theta\right) \cdot T. \end{aligned}$$

Recall that $\Delta(ECC_1(x), ECC_1(x')) = L_1/2$ and that $T = L_1 + L_2 + L_3$, and get that we must have:

$$\Delta(ECC_3(x, x'), ECC_3(x', x)) \geq \frac{L_3}{2} + 2\theta T + \frac{L_2}{2}.$$

- 2) For Bob to not be confused between messages for the pairs (x, a) and (x', x) (and similarly for pairs (x, x') and (x', b)), we must have:

$$\Delta(ECC_3(x, a), ECC_3(x', x)) \geq \frac{L_3}{2} + 2\theta T.$$

This follows from a similar argument to Item 1 where we use (1) on (x', x) and (2) on (x, a) .

- 3) For Bob to not be confused between messages for the pairs (x, a) and (x', b) (note that it is possible that $a = b$

but both are always different from x and x'), we must have:

$$\Delta(\text{ECC}_3(x, a), \text{ECC}_3(x', b)) \geq \frac{L_3}{2} + 2\theta T - \frac{L_2}{2}.$$

This follows from similar argument to Item 1 when we use (2) on both pairs.

Importantly, as the number of pairs that require the weakest possible guarantee in Item 3 above is much larger than those requiring the stronger guarantees, if we choose $\theta > 0$ small enough so that L_2 is significantly larger than θT , say, $L_2 = 16\theta T$, these distance requirements do not violate the $L_3/2$ bound on the average distance implied by the Plotkin bound. Thus, such an ECC_3 code is theoretically possible, and we provide a construction below.

E. Location-Sensitive Codes: The Construction

The upshot of the discussion above is that the code Alice uses in the third step to encode pairs must ensure that the distance between the encodings of pairs that have one or two common elements *at different locations* must be large (Items 1 and 2), while the distance between the encodings of pairs that have no common elements or the common elements are *at the same location* (Item 3) can be smaller than the distance obtained by a random code.

We call such a code a *location-sensitive code* and note that its distance guarantees are very different from standard equally-spaced codes. To capture the fact that the distances need to be larger between the encodings of two pairs where the same element appears in different locations, we employ the following approach while encoding a pair (s, t) : Let C be an equally-spaced code that encodes a single input s and has relative distance $\frac{1}{2}$ between pairs of codewords. We construct our location-sensitive code $\text{LSC}(s, t)$ by setting each coordinate i to be coordinate i of $C(s)$ with probability p and coordinate i of $C(t)$ with probability $1 - p$, for some carefully selected $p > 0$. In other words, the encoding $\text{LSC}(s, t)$ is *positively correlated* with $C(s)$ and *negatively correlated* with $C(t)$.

Roughly speaking, the above approach has the property that for pairs (x, a) and (x', x) where the same element appears at different locations, the encodings $\text{LSC}(x, a)$ and $\text{LSC}(x', x)$ will be positively and negatively correlated with $C(x)$ respectively, and therefore should have a high distance. In contrast, for pairs (x, a) and (x', a) where the same element appears at the same location, the encodings $\text{LSC}(x, a)$ and $\text{LSC}(x', a)$ will both be negatively correlated with $C(a)$ and will suffer from a lower distance. By choosing the parameters carefully, we are able to meet the requirements in Section II-D.

F. Finalizing the Details

We finish the overview with a discussion of some remaining issues that we have not covered well so far.

a) Other attacks in the second step: One assumption that we made in our discussion above is that Alice always gets the encoding of one pair in the second step. This loses generality as nothing stops the adversary from giving Alice a combination of the encoding of various pairs. To get around this, we use list-decoding in the second step to give Alice a pair of pairs that contains the right pair unless there were too many corruptions.

When Alice tries to encode this pair of pairs in the third step using a location-sensitive code, she actually just encodes both the pairs separately, and simply sends a message that is positively correlated with both the encodings. The actual correlations intricately depend on the exact message received by her in the second step, with larger correlations to one of the pairs if the encoding of that pair was not too far from what Alice received in the second step.

b) Sets arising from list-decoding: Another assumption we made throughout the above sketch was that using list-decodable codes allows the parties to compute sets of size 2 that are guaranteed to contain the correct codeword even if there are strictly more than $\frac{1}{4}$ errors. This assumption is not correct for binary codes, and one needs to have sets of size at least 3. Correspondingly, in the actual proof, we make the entire argument above work with list-decodable codes that yield lists of size 3.

One crucial change this entails is that our location-sensitive codes must also now take triples instead of pairs. While we believe that such location-sensitive codes still exist, we found it easier to convert triple to pairs in the following way: When Alice wants to encode a triple (x, a, b) , where x is her true input and $a, b \neq x$, she instead encodes the pairs (x, a) and (x, b) and sends a message positively correlated with both these encodings. We then are able to make the analysis work by carefully controlling the correlations, which forms a lot of the technical work in this paper.

III. MODEL AND PRELIMINARIES

A. Notation

All logarithms are base 2. We use $\log^{(k)} n$ to denote the k -times iterated logarithms of n . For $n \in \mathbb{N}$, $x, y \in \{0, 1\}^n$, we denote by $\Delta(x, y) = |\{i \in [n] \mid x_i \neq y_i\}|$ the Hamming distance between x and y . For a set S and an integer $k \geq 0$, the notation $\binom{S}{k}$ denotes the set of all subsets of S that have exactly k elements. Also define, for a set S , the set $\mathcal{D}(S)$ to be the set of all distributions over S .

B. The Binary Two-Way Communication Channel

We now define (deterministic) protocols over the binary two-way communication channel. Such a protocol is defined by a tuple:

$$\Pi = (\mathcal{X}^A, \mathcal{X}^B, \mathcal{Y}, T, p, f^A, f^B, \text{out}),$$

where (1) \mathcal{X}^A is the set of all possible inputs for Alice, (2) \mathcal{X}^B is the set of all possible inputs for Bob, (3) \mathcal{Y} is the set of all possible outputs (for Bob), (4) T is the length of the protocol (the number of rounds), (5) $p \in \{A, B\}^T$ is the order of turns,

(6) $f^A : \mathcal{X}^A \times \{0,1\}^{<T} \rightarrow \{0,1\}$ is the message function for Alice, (7) $f^B : \mathcal{X}^B \times \{0,1\}^{<T} \rightarrow \{0,1\}$ is the message function for Bob, (8) $\text{out} : \mathcal{X}^B \times \{0,1\}^T \rightarrow \mathcal{Y}$ is the output function (for Bob).

a) *Execution of a protocol:* An adversary for such a protocol is defined by a function $\text{Adv} : \mathcal{X}^A \times \mathcal{X}^B \rightarrow \{0,1\}^T$. For $i \in [T]$, we shall use $\text{Adv}_i(\cdot)$ to denote the function that outputs the i^{th} bit of $\text{Adv}(\cdot)$. We next define an execution of Π in the presence of an adversary Adv for Π : At the beginning of the execution, Alice starts with an input $x^A \in \mathcal{X}^A$ and Bob starts with an input $x^B \in \mathcal{X}^B$. The execution consists of T rounds and before the i^{th} round, for $i \in [T]$, Alice and Bob have transcripts $\pi^A, \pi^B \in \{0,1\}^{i-1}$ respectively. In round i , if $p_i = A$, then Alice transmits the symbol $f^A(x^A, \pi^A)$ while Bob receives the symbol $\text{Adv}_i(x^A, x^B)$. Both the parties add these symbols to π^A and π^B respectively. Similarly, if $p_i = B$, then Bob transmits the symbol $f^B(x^B, \pi^B)$ while Alice receives the symbol $\text{Adv}_i(x^A, x^B)$. Both the parties add these symbols to π^A and π^B respectively.

After T such rounds, Bob outputs $\text{out}(x^B, \pi^B)$. Observe that this execution, and therefore π^A and π^B , are completely determined by x^A, x^B, Π , and Adv . We denote the output of Π on inputs $x^A \in \mathcal{X}^A$ and $x^B \in \mathcal{X}^B$ in the presence of adversary Adv by $\text{out}_{\Pi, \text{Adv}}(x^A, x^B)$.

b) *Phases:* We say that $i \in [T-1]$ is an *alternation round* of Π if $p_{i+1} \neq p_i$. Assume that Π has P alternations and let $i_1 < i_2 < i_3 < \dots < i_P$ be all alternation rounds. Define $i_0 = 0$ and $i_{P+1} = T$. For $t \in [P+1]$, we define *Phase t* of Π as the set of rounds $(i_{t-1}, i_t]$. Informally, Phase t is the t^{th} message by one of the parties.

c) *Corruptions:* Consider an execution of Π in the presence of the adversary Adv . For $R \subseteq [T]$, $x^A \in \mathcal{X}^A$, and $x^B \in \mathcal{X}^B$, we define the number of corruptions in the rounds in R to be

$$\text{corr}_{\Pi, \text{Adv}, R}(x^A, x^B) = \sum_{i \in R} \mathbb{1}(\pi_i^A \neq \pi_i^B).$$

Recall that π^A, π^B are completely determined by x^A, x^B, Π , and Adv and therefore corr is well defined. We omit the subscript R when $R = [T]$.

d) *Computing a function:* Let $\mathcal{I} \subseteq \mathcal{X}^A \times \mathcal{X}^B$ and let $\gamma \in [0,1]$. Let $F : \mathcal{I} \rightarrow \mathcal{Y}$ be a (possibly partial) function. Let $\Pi = (\mathcal{X}^A, \mathcal{X}^B, \mathcal{Y}, T, p, f^A, f^B, \text{out})$ be a protocol. We say that Π *computes F* against a γ fraction of corruptions if for all adversaries Adv and all inputs $(x^A, x^B) \in \mathcal{I}$, it holds that $\text{out}_{\Pi, \text{Adv}}(x^A, x^B) = F(x^A, x^B)$ as long as $\text{corr}_{\Pi, \text{Adv}}(x^A, x^B) \leq \lceil \gamma T \rceil$. Similarly, we say that Π *computes F* against a γ fraction of corruptions *per phase* if for all adversaries Adv and all inputs $(x^A, x^B) \in \mathcal{I}$, it holds that $\text{out}_{\Pi, \text{Adv}}(x^A, x^B) = F(x^A, x^B)$ as long as $\text{corr}_{\Pi, \text{Adv}, (i_{t-1}, i_t]}(x^A, x^B) \leq \lceil \gamma(i_t - i_{t-1}) \rceil$ for all $t \in [P+1]$ (recall that rounds $(i_{t-1}, i_t]$ constitute Phase t).

1) *Protocols with EQUALLY SPACED CODE:* Let $n, k \in \mathbb{N}$ and $C \subseteq \{0,1\}^n$ be such that $|C| = k$. Let $\delta > 0$. We say

that C is a δ -EQUALLY SPACED CODE if for all $c_1, c_2 \in C$ it holds that

$$\Delta(c_1, c_2) \leq \begin{cases} \left(\frac{k}{2(k-1)} + \delta \right) n, & \text{if } k \text{ is even} \\ \left(\frac{k+1}{2k} + \delta \right) n, & \text{if } k \text{ is odd} \end{cases}.$$

We mention that the Plotkin bound implies that for *any* code C the average distance, $\Delta(c_1, c_2)$, between two codewords $c_1, c_2 \in C$ satisfies the last inequality, even with $\delta = 0$. Thus, a code C is an EQUALLY SPACED CODE if all the distances are at most the best possible average distance of a binary code.

Let Π be a protocol of length T over the binary two-way communication channel. For every Phase $t \in [P+1]$ of Π , let $C_{t, \mathcal{I}, \gamma}$ be the set of all possible messages sent in Phase t for inputs in \mathcal{I} and adversaries with at most γ fraction of corruptions per phase. More formally, if Alice is the sender in Phase t , then $C_{t, \mathcal{I}, \gamma}$ is $\pi_{(i_{t-1}, i_t]}^A$ for all possible transcripts π^A obtained by taking all possible input pairs $(x^A, x^B) \in \mathcal{I}$ and adversaries Adv with $\forall t' \in [t-1] : \text{corr}_{\Pi, \text{Adv}, (i_{t'-1}, i_{t'})}(x^A, x^B) \leq \lceil \gamma(i_{t'} - i_{t'-1}) \rceil$. (Recall that π^A is completely determined by x^A, x^B, Π , and Adv .) Similarly, for the case that Bob is the sender in Phase t .

Let $\delta > 0$. We say that Π uses $(\mathcal{I}, \gamma, \delta)$ -EQUALLY SPACED CODE if for all $t \in [P+1]$ the set $C_{t, \mathcal{I}, \gamma}$ is a δ -EQUALLY SPACED CODE.

2) *The Message Transfer Function:* We now define the message transfer function MsgTrans that is the focus of this paper. Let $n \in \mathbb{N}$. The *message transfer function* $\text{MsgTrans}_n : \{0,1\}^n \times \{\perp\} \rightarrow \{0,1\}^n$ is given by $\text{MsgTrans}_n(x, \perp) = x$ for all $x \in \{0,1\}^n$.

IV. LOCATION-SENSITIVE CODES

This section is devoted to the construction of our *location-sensitive code* C , formally stated in Theorem IV.1 below. This code encodes a pair of binary strings $(x_1, x_2) \in \{0,1\}^n \times \{0,1\}^n$ such that the distance between the encodings of two different pairs, $\Delta = \Delta(C(x_{1,1}, x_{1,2}), C(x_{2,1}, x_{2,2}))$, satisfies:

- 1) If the two pairs do not share an element, *i.e.*, $|\{x_{1,1}, x_{1,2}\} \cap \{x_{2,1}, x_{2,2}\}| = 0$, then $\Delta = \frac{1}{2}$.
- 2) If the two pairs share a single element and the element is in the same location, *i.e.*, $|\{x_{1,1}, x_{1,2}\} \cap \{x_{2,1}, x_{2,2}\}| = 1$ and $\exists j \in [2] : x_{1,j} = x_{2,j}$, then Δ is a constant smaller than $1/2$.
- 3) If the two pairs share a single element and the element is in different locations, *i.e.*, $|\{x_{1,1}, x_{1,2}\} \cap \{x_{2,1}, x_{2,2}\}| = 1$ and $\exists j \in [2] : x_{1,j} = x_{2,3-j}$, then Δ is a constant greater than $1/2$.
- 4) If the two pairs share both elements, *i.e.*, $|\{x_{1,1}, x_{1,2}\} \cap \{x_{2,1}, x_{2,2}\}| = 2$, then Δ is a constant strictly greater than 1 minus the constant in Item 2 (and, in particular, greater than $1/2$).

Theorem IV.1. *For all $\epsilon > 0$, there exists a constant K_5 such that for all $K' \geq K_5$ and $n > 0$, there exists a code $C : \{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}^{K'n}$ such that the following holds*

for all $(x_{1,1}, x_{1,2}), (x_{2,1}, x_{2,2}) \in \{0, 1\}^n \times \{0, 1\}^n$ satisfying $x_{i,1} \neq x_{i,2}$ for all $i \in [2]$:

$$\left| \frac{1}{K'n} \cdot \Delta\left(C(x_{1,1}, x_{1,2}), C(x_{2,1}, x_{2,2})\right) - \left(\frac{1}{2} + \eta\right) \right| \leq \epsilon,$$

where:

$$\eta = -\frac{1}{128} \cdot \sum_{j, j' \in [2]} (-7)^{4-j-j'} \cdot \mathbb{1}(x_{1,j} = x_{2,j'}).$$

(Note that η may be positive.)

To prove Theorem IV.1, we will need some (fairly standard) results about random codes (given in Section IV-A), as well as new ideas (given in Sections IV-B and IV-C).

A. Random Coding

Define the function $\text{reg} : [4] \rightarrow \mathbb{R}$ as follows:

$$\text{reg}(z) = \begin{cases} 0, & z = 1 \\ \frac{1}{2}, & z = 2 \\ \frac{3}{4}, & z = 3 \\ \frac{5}{4}, & z = 4 \end{cases}.$$

Roughly speaking, $\text{reg}(z)$ for $z \in [4]$ captures the expected sum of the fractional Hamming distance of any z codewords to the bitwise majority of the z codewords when the code is chosen uniformly at random. Using the probabilistic method, we show a code for which these fractional Hamming distances are always attained.

Lemma IV.2. *For all $\epsilon > 0$, there exists a constant K_1 such that for all $K' \geq K_1$ and $n > 0$, there exists a code $C : \{0, 1\}^n \rightarrow \{0, 1\}^{K'n}$ such that the following holds for all $z \in [4]$, distinct $x_1, \dots, x_z \in \{0, 1\}^n$, and (not necessarily distinct) bits $b_1, \dots, b_z \in \{0, 1\}$:*

$$\left| \frac{1}{K'n} \cdot \sum_{j \in [K'n]} \mathbb{1}(\forall i \in [z] : C_j(x_i) = b_i) - \frac{1}{2^z} \right| \leq \epsilon.$$

Corollary IV.3. *For all $\epsilon > 0$, there exists a constant K_2 such that for all $K' \geq K_2$ and $n > 0$, there exists a code $C : \{0, 1\}^n \rightarrow \{0, 1\}^{K'n}$ such that the following holds for all $z \in [4]$ and distinct $x_1, \dots, x_z \in \{0, 1\}^n$:*

$$\left| \frac{1}{K'n} \cdot \sum_{j \in [K'n]} \min_{b \in \{0, 1\}} \sum_{i \in [z]} \Delta\left(C_j(x_i), b\right) - \text{reg}(z) \right| \leq \epsilon.$$

B. The Code Merging Operation

Lemma IV.4. *For all $\epsilon > 0, k, m \in \mathbb{N}$, and a set $D \in \binom{\mathcal{D}([k])}{m}$, there exists a constant K_3 such that for all $K' \geq K_3, n > 0$, a set S with $|S| \leq 2^n$, and codes $C_j : S \rightarrow \{0, 1\}^n$ for all $j \in [k]$, there exists a code $C : S^k \times D \rightarrow \{0, 1\}^{K'n}$ such that the following holds for all distinct $\left((s_{i,j})_{j \in [k]}, d_i\right)_{i \in [2]} \in (S^k \times D)^2$:*

$$\left| \Delta\left(C\left((s_{1,j})_{j \in [k]}, d_1\right), C\left((s_{2,j})_{j \in [k]}, d_2\right)\right) - \nabla \right| \leq \epsilon K'n,$$

where:

$$\nabla = K' \cdot \sum_{j, j' \in [k]} d_1(j) d_2(j') \cdot \Delta\left(C_j(s_{1,j}), C_{j'}(s_{2,j'})\right).$$

C. Proof of Theorem IV.1

Lemma IV.5 below constructs an error correcting code that encodes a binary string and a bit $(x, \ell) \in \{0, 1\}^n \times \{0, 1\}$ such that the distance between the encodings of two different messages, $\Delta\left(C(x_1, \ell_1), C(x_2, \ell_2)\right)$, is $1/2$ if $x_1 \neq x_2$ and is 1 if $x_1 = x_2$. Note that in the latter case, we must have $\ell_1 \neq \ell_2$.

Lemma IV.5. *For all $\epsilon > 0$, there exists a constant K_4 such that for all $K' \geq K_4$ and $n > 0$, there exists a code $C : \{0, 1\}^n \times \{0, 1\} \rightarrow \{0, 1\}^{K'n}$ such that the following holds for all $x_1, x_2 \in \{0, 1\}^n$ and $\ell_1, \ell_2 \in \{0, 1\}$:*

$$\left| \frac{1}{K'n} \cdot \Delta\left(C(x_1, \ell_1), C(x_2, \ell_2)\right) - \left(\frac{1}{2} + \kappa\right) \right| \leq \epsilon,$$

where:

$$\kappa = \frac{1}{2} \cdot (-1)^{\ell_1 + \ell_2 + 1} \cdot \mathbb{1}(x_1 = x_2).$$

(Note that κ may be negative.)

We are now ready to prove the main result of this section, Theorem IV.1.

Proof of Theorem IV.1: Let K_4 be the constant from Lemma IV.5 for $\epsilon' = \frac{\epsilon}{2}$, and $C' : \{0, 1\}^n \times \{0, 1\} \rightarrow \{0, 1\}^{K_4 n}$ the code guaranteed by Lemma IV.5 for K_4, n . Define $C_1(\cdot) := C'(\cdot, 0)$ and $C_2(\cdot) := C'(\cdot, 1)$. Let K_3 be the constant from Lemma IV.4 for $\epsilon', k = 2, m = 1$ and D being the singleton set containing the distribution μ over $[2]$ that gives a probability of $\frac{7}{8}$ to 1 and a probability of $\frac{1}{8}$ to 2. Set $K_5 = K_4 K_3$. For all $K' \geq K_5$, let $C_0 : (\{0, 1\}^n)^2 \times D \rightarrow \{0, 1\}^{K'n}$ be the code guaranteed by Lemma IV.4 for $K'/K_4, K_4 n$ and codes C_1, C_2 . We show the following code C satisfies the lemma:

$$C(x_1, x_2) := C_0(x_1, x_2, \mu).$$

By Lemma IV.4, we get that for all $(x_{1,1}, x_{1,2}), (x_{2,1}, x_{2,2}) \in \{0, 1\}^n \times \{0, 1\}^n$ as in the lemma statement, we have:

$$\nabla - \epsilon' K'n \leq \Delta\left(C(x_{1,1}, x_{1,2}), C(x_{2,1}, x_{2,2})\right) \leq \nabla + \epsilon' K'n,$$

where, using Lemma IV.5, we have:

$$(\kappa^* - \epsilon') \cdot K'n \leq \nabla \leq (\kappa^* + \epsilon') \cdot K'n,$$

and:

$$\kappa^* = \sum_{j, j' \in [2]} \frac{7^{4-j-j'}}{64} \cdot \left(\frac{1}{2} + \frac{(-1)^{j+j'+1}}{2} \cdot \mathbb{1}(x_{1,j} = x_{2,j'}) \right).$$

To finish the proof, note that $\eta = \kappa^* - \frac{1}{2}$ and $\epsilon = 2\epsilon'$. ■

V. OUR PROTOCOL

We formalize Theorem I.1 as Theorem V.1:

Theorem V.1. *Define $\theta = 10^{-5}$. There exists a constant K such that for all $n \in \mathbb{N}$ there exists a two party protocol $\Pi = \Pi_n$ with length $T = Kn$ that computes the message transfer function MsgTrans_n against a $(\frac{1}{4} + \theta)$ fraction of corruptions.*

In this section we give the protocol that proves Theorem V.1. We first give the notation and definitions used by the protocol (Sections V-A to V-C). The protocol itself is in Algorithm 1 (Section V-D). The analysis of Algorithm 1 (and the proof of Theorem V.1) is deferred to the full version [7].

For the rest of this section, since Bob has only one possible input $x^B = \perp$ in a protocol Π for MsgTrans_n , we omit Bob's input from our notation (e.g., we write $\text{out}_{\Pi, \text{Adv}}(x)$ to mean $\text{out}_{\Pi, \text{Adv}}(x^A = x, x^B = \perp)$).

A. Stories

We heavily rely on the following definition of a story.

Definition V.2. *Let $n, M \in \mathbb{N}$. We define an (n, M) -story to be a tuple $\mathcal{Z} = (x, U, V, w) \in \{0, 1\}^n \times \binom{\{0, 1\}^n}{2} \times \binom{\{0, 1\}^n}{4} \times [M]$ such that $U \cap V = \emptyset$ and $x \notin U \cup V$. Denote by $\text{Stories}_{n, M}$ the set of all (n, M) -stories.*

a) *The function $\text{story}(\cdot)$:* Recall that our algorithm starts with Alice sending her input x using a list-decodable code. This allows Bob to compute a set S of size 3 that contains x . Bob then sends this set S back to Alice using another list-decodable code, which allows Alice to compute three sets $T_1, T_2, T_3 \in \binom{\{0, 1\}^n}{3}$ such that all of them contain x and (if not too many errors were introduced), one of them is S . Additionally, Alice can compute the distance $w \in [M]$ of the message she receives to the closest correct codeword (see Line 4). Thus, before Phase 3, Alice can compute a tuple $\mathcal{T} = (T_1, T_2, T_3, w)$, which is guaranteed to be an element of the set Data defined next.

Definition V.3. *For $x \in \{0, 1\}^n$, we define the set $\text{Data}(x)$ to be the set containing all tuples $\mathcal{T} = (T_1, T_2, T_3, w) \in \binom{\{0, 1\}^n}{3} \times [M]$ such that T_1, T_2, T_3 are distinct and $x \in \bigcap_{\ell \in [3]} T_\ell$. We also define the set $\text{Data} = \bigcup_{x \in \{0, 1\}^n} \text{Data}(x)$.*

Next, we define a function $\text{story}(\cdot)$ that Alice can use to take a tuple $\mathcal{T} \in \text{Data}(x)$ and her input x to output a story $\text{story}(x, \mathcal{T}) \in \text{Stories}_{n, M}$ she can encode in Phase 3. More precisely, we define $\text{story}(x, \mathcal{T}) = (x, U, V, w)$, where $U = T_1 \setminus \{x\}$ and V is the set $(T_2 \cup T_3) \setminus T_1$ padded using dummy elements to have size 4.

B. Location-Sensitive Codes for Stories

Our protocol uses a location-sensitive code that encodes stories, given in Lemma V.4, and is based on the location-sensitive code constructed in Section IV. This code uses the following functions $d_1, d_2, d_3 : [0, 1] \rightarrow [0, 1]$:

$$\begin{aligned} d_1(z) &= z, \\ d_2(z) &= \max\left(z, \frac{1}{2} - z\right), \\ d_3(z) &= \max\left(z, \frac{1}{2} - z, \frac{5}{12} - \frac{z}{3}\right). \end{aligned}$$

Lemma V.4. *For all $\epsilon > 0$ and $M \in \mathbb{N}$, there exists a constant K_6 such that for all $K' \geq K_6$ and $n > 0$, there exists a code $C : \text{Stories}_{n, M} \rightarrow \{0, 1\}^{K'n}$ such that the following holds for all $\mathcal{Z}_1 = (x_1, U_1, V_1, w_1), \mathcal{Z}_2 = (x_2, U_2, V_2, w_2) \in \text{Stories}_{n, M}$ satisfying $x_1 \neq x_2$:*

$$\begin{aligned} &\Delta\left(C(\mathcal{Z}_1), C(\mathcal{Z}_2)\right) \\ &\geq \left(0.5511 - \epsilon - 0.1 \cdot \sum_{i \in [2]} d(x_{3-i}, U_i, V_i, w_i)\right) \cdot K'n, \end{aligned}$$

where:

$$d(y, U, V, w) = \begin{cases} d_1\left(\frac{w}{M}\right), & \text{if } y \in U \\ d_2\left(\frac{w}{M}\right), & \text{if } y \in V \\ d_3\left(\frac{w}{M}\right), & \text{if } y \notin U \cup V \end{cases}.$$

We defer the proof of Lemma V.4 to the full version [7].

C. Protocol Definitions

a) *Constants:* We shall assume that there are at least 20 ‘‘dummy’’ strings (strings that cannot be inputs for Alice) in $\{0, 1\}^n$. This is without loss of generality as one can simply increase n by 10 and have enough dummy strings. We define $\epsilon = \theta^{10}$ and $M = \frac{480}{\epsilon}$. Let K_2 be the constant from Corollary IV.3 for this value of ϵ . Similarly, let K_6 be the constant from Lemma V.4 for this value of ϵ and M . Define $K = 90^{10} \cdot \max(K_2, K_6)$ and the parameters:

$$L_1 = \frac{401}{500} \cdot K, \quad L_2 = \frac{9}{500} \cdot K, \quad L_3 = \frac{90}{500} \cdot K.$$

Note that all these parameters are integers divisible by 3 and larger than $\max(K_2, K_6)$.

b) *Error correcting codes:* Fix $n > 0$ for the rest of this paper. The protocol Π that we define shall use several different types of codes, which we define next. Let ECC_1 and ECC_2 be the codes promised by Corollary IV.3 for L_1, n and $L_2/3, 3n$ respectively. Also, let ECC_3 be the code promised by Lemma V.4 for L_3, n .

c) *The function $\text{corr-lb}^{(2)}(\cdot)$:* As explained above, in Phase 3, Alice sends the encoding of a story to Bob. To compute its output, Bob decodes this message together with the message he received in Phase 1. During this decoding, he will also need to estimate (actually, lower bound) the number of corruptions in Phase 2.

This is done using a function $\text{corr-lb}^{(2)}(\cdot)$, parameterized by the set $S \in \binom{\{0, 1\}^n}{3}$ that Bob computed in Phase 1, and

Algorithm 1 The MsgTrans protocol Π .

Input: Alice has input $x \in \{0, 1\}^n$.

Output: Bob outputs $y \in \{0, 1\}^n$.

Phase 1:

- 1: Alice sends $\text{ECC}_1(x)$ bit by bit over $L_1 n$ rounds.
- 2: Bob receives $\rho \in \{0, 1\}^{L_1 n}$ and computes

$$S = \arg \min_{S' \in \binom{\{0, 1\}^n}{3}} \sum_{x' \in S'} \Delta(\text{ECC}_1(x'), \rho).$$

Phase 2:

- 3: Bob sends $\text{ECC}_2(S)$ bit by bit over $L_2 n$ rounds.
- 4: Alice receives $\sigma \in \{0, 1\}^{L_2 n}$. She orders all sets $T \in \binom{\{0, 1\}^n}{3}$ containing x in increasing order of the value $\Delta(\text{ECC}_2(T), \sigma)$, and denotes by T_1, T_2, T_3 the first three sets in this ordering, with T_1 being the first. Let $\mathcal{T} = (T_1, T_2, T_3, w)$ where $w = \left\lceil M \cdot \frac{\Delta(\text{ECC}_2(T_1), \sigma)}{L_2 n} \right\rceil$.

Phase 3:

- 5: Alice sends $\text{ECC}_3(\text{story}(x, \mathcal{T}))$ bit by bit over $L_3 n$ rounds.
- 6: Bob receives $\tau \in \{0, 1\}^{L_3 n}$ and outputs

$$y = \arg \min_{x' \in S} \text{corr-lb}(x'),$$

where

$$\begin{aligned} \text{corr-lb}(x') &= \Delta(\text{ECC}_1(x'), \rho) + \min_{\mathcal{T}' \in \text{Data}(x')} \left(\text{corr-lb}_S^{(2)}(\mathcal{T}') \right) \\ &\quad + \Delta(\text{ECC}_3(\text{story}(x', \mathcal{T}')), \tau). \end{aligned}$$

takes as input $\mathcal{T} \in \text{Data}$, which is Bob's candidate for what Alice may have encoded in Phase 2. Formally,

$$\text{corr-lb}_S^{(2)}(\mathcal{T}) = L_2 n \cdot \begin{cases} d_1\left(\frac{w}{M}\right), & \text{if } S = T_1 \\ d_2\left(\frac{w}{M}\right), & \text{if } S = T_2 \text{ or } S = T_3 \\ d_3\left(\frac{w}{M}\right), & \text{otherwise} \end{cases}.$$

At a high level, $\text{corr-lb}_S^{(2)}(\mathcal{T})$ is a lower bound on the number of corruptions in Phase 2 when S is the set Bob computed in Phase 1 and \mathcal{T} is the data Alice computed in Phase 2. Note however that none of the parties can actually compute $\text{corr-lb}_S^{(2)}(\mathcal{T})$ as they do not know both S and \mathcal{T} .

D. The Protocol

We are now ready to define the protocol Π in Algorithm 1. We note that ties in all $\arg \min$ are broken arbitrarily.

VI. IMPOSSIBILITY RESULT FOR EQUALLY SPACED CODE

In this section we prove Theorem VI.1 below, which is the formal version of Theorem I.2.

Theorem VI.1. *For any $\theta > 0$, there exists $n_0 > 0$ such that for every $n \geq n_0$, there does not exist a 3-phase protocol using*

$(\{0, 1\}^n, \frac{1}{4} + \theta, \frac{\theta}{2})$ -EQUALLY SPACED CODE that computes MsgTrans_n against a $(\frac{1}{4} + \theta)$ fraction of corruptions.

A. Proof of Theorem VI.1

The proof of Theorem VI.1 uses the following functions:

- 1) $F_{1,N}$: Let $\mathcal{I}_{1,N} = [N] \times \{\perp\}$. Define $F_{1,N} : \mathcal{I}_{1,N} \rightarrow [N]$ by $F_{1,N}(x, \perp) = x$.
- 2) $F_{2,N}$: Let $\mathcal{I}_{2,N} = \{(x, B) \in [N] \times \binom{[N]}{2} \mid x \in B\}$. Define $F_{2,N} : \mathcal{I}_{2,N} \rightarrow [N]$ by $F_{2,N}(x, B) = x$.
- 3) $F_{3,N}$: Let $\mathcal{I}_{3,N} = \left\{ ((x, A), B) \in \left([N] \times \binom{[N]}{2} \right) \times \binom{[N]}{2} \mid x \in B, x \notin A, B \subseteq A \cup \{x\} \right\}$. Define $F_{3,N} : \mathcal{I}_{3,N} \rightarrow [N]$ by $F_{3,N}((x, A), B) = x$.

For $i \in [3]$, we are interested in $(4 - i)$ -phase protocols using EQUALLY SPACED CODE that computes $F_{i,N}$ against a $(\frac{1}{4} + \theta)$ fraction of corruptions per phase for $\theta > 0$. Without loss of generality, we assume protocols alternate between Alice and Bob across phases with Alice always sending in the last phase.

To prove Theorem VI.1, we first note that for $N = 2^n$, $F_{1,N}$ is MsgTrans_n , thus we need to prove a lower bound against 3-phase protocols for $F_{1,N}$. To this end, we show that a 3-phase protocol using EQUALLY SPACED CODE that computes $F_{1,N}$ against a $(\frac{1}{4} + \theta)$ fraction of corruptions per phase for $\theta > 0$, implies a similar 2-phase protocol that computes $F_{2,N'}$, for some smaller N' (Lemma VI.2). We then show that a 2-phase protocol that computes $F_{2,N'}$ implies a 1-phase protocol that computes $F_{3,N''}$, for some even smaller N'' (Lemma VI.3). Finally, we give a lower bound against any 1-phase protocol that computes $F_{3,N''}$ (Lemma VI.4).

We mention that the (noiseless) communication complexity of $F_{2,N}$ was studied by [5], who showed a tight bound of $\Theta(\log N)$. Furthermore, his upper bound protocol consists of only two phases. Observe that a 2-phase protocol for $F_{2,N}$ implies a 1-phase protocol for $F_{3,N}$. Lemma VI.4 gives an $\Omega(\log \log N)$ lower bound on the communication complexity of a 1-phase protocol for $F_{3,N}$. While our result in this paper does not require an upper bound for this problem, in Section VI-B we include a matching $O(\log \log N)$ upper bound as we believe it may be of independent interest. We also remark that the 1-phase complexity of $F_{3,N}$ is closely related to the question of deterministic compression with uncertain priors studied in [6].

We will use the following definition: For $n > 0$ and $x, y \in \{0, 1\}^n$ we define the string $\text{mid}(x, y) \in \{0, 1\}^n$. Informally, $\text{mid}(x, y)$ is a string whose Hamming distance from x and from y is equal (up-to ± 1). Formally, if $i_1 < i_2 < \dots < i_{\Delta(x,y)}$ are the elements of $\{i \in [n] \mid x_i \neq y_i\}$, then, for all $i \in [n]$, coordinate i of the string $\text{mid}(x, y)$ is defined as:

$$\text{mid}(x, y) = \begin{cases} x_i, & \text{if } \forall j \in [\Delta(x, y)] : i \neq i_j \\ x_i, & \text{if } \exists j \in \left[\left\lceil \frac{\Delta(x, y)}{2} \right\rceil \right] : i = i_j \\ y_i, & \text{if } \exists j \in [\Delta(x, y)] \setminus \left[\left\lceil \frac{\Delta(x, y)}{2} \right\rceil \right] : i = i_j \end{cases}.$$

1) Going from 3 Phases to 2:

Lemma VI.2. *If there exists a 3-phase protocol Π using $(\mathcal{I}_{1,N}, \frac{1}{4} + \theta, \frac{\theta}{2})$ -EQUALLY SPACED CODE that computes $F_{1,N}$ against a $(\frac{1}{4} + \frac{\theta}{2})$ fraction of corruptions per phase, where $N, \theta > 0$, then there also exists a 2-phase protocol Π' using $(\mathcal{I}_{2,\theta N}, \frac{1}{4} + \theta, \frac{\theta}{2})$ -EQUALLY SPACED CODE that computes $F_{2,\theta N}$ against a $(\frac{1}{4} + \frac{\theta}{2})$ fraction of corruptions per phase.*

2) Going from 2 Phases to 1:

Lemma VI.3. *If there exists a 2-phase protocol Π using $(\mathcal{I}_{2,N}, \frac{1}{4} + \theta, \frac{\theta}{2})$ -EQUALLY SPACED CODE that computes $F_{2,N}$ against a $(\frac{1}{4} + \frac{\theta}{2})$ fraction of corruptions per phase, where $N, \theta > 0$, then there also exists a 1-phase protocol Π' using $(\mathcal{I}_{3,\log^\theta N}, \frac{1}{4} + \theta, \frac{\theta}{2})$ -EQUALLY SPACED CODE that computes $F_{3,\log^\theta N}$ against a $(\frac{1}{4} + \frac{\theta}{2})$ fraction of corruptions (per phase).*

3) Lower Bound for 1 Phase:

Lemma VI.4. *For any 1-phase protocol Π computing $F_{3,N}$, Alice uses at least $\log \log N$ distinct codewords across all possible inputs.*

Finally, we are ready to prove Theorem VI.1.

Proof of Theorem VI.1: We first claim that a 3-phase protocol that computes MsgTrans_n (even against a 0 fraction of corruptions) must have at least $\frac{\log n}{10}$ rounds. Suppose there exists a 3-phase protocol that computes MsgTrans_n with $T < \frac{\log n}{10}$ rounds, then there also exists a 1-phase protocol that computes MsgTrans_n with 4^T rounds. But a 1-phase protocol that computes MsgTrans_n (one message from Alice to Bob) must communicate at least n bits, a contradiction.

Fix $\theta > 0$ and let n_0 be an arbitrary integer such that $\log \log \left(\log \frac{\theta}{100} \left(\frac{\theta}{100} \cdot n_0 \right) \right) \geq \frac{100}{\theta}$. Let $n \geq n_0$. We next claim that any 3-phase protocol Π that computes MsgTrans_n against a $(\frac{1}{4} + \theta)$ fraction of corruptions, also computes MsgTrans_n against a $(\frac{1}{4} + \theta')$ fraction of corruptions *per phase*, where $\theta' = \frac{\theta}{2}$. To see this, for $t \in [3]$, let L_t be the length of Phase t of Π , and let $T = L_1 + L_2 + L_3$. Let Adv be an adversary for Π that corrupts at most $\lceil (\frac{1}{4} + \theta') L_t \rceil$ rounds in Phase t . The total number of rounds that are corrupted by Adv is at most $(\frac{1}{4} + \theta')T + 3$, which is no more than $\lceil (\frac{1}{4} + \theta)T \rceil$ as $T \geq \frac{\log n}{10} \geq \frac{10}{\theta}$.

Let $N = 2^n$. Observe that $F_{1,N}$ is exactly MsgTrans_n . Thus, it remains to show that there is no 3-phase protocol Π using $(\mathcal{I}_{1,N}, \frac{1}{4} + \theta, \theta')$ -EQUALLY SPACED CODE that computes $F_{1,N}$ against a $(\frac{1}{4} + \theta')$ fraction of corruptions per phase.

We prove the last claim by contradiction. Assume that there exists a 3-phase protocol Π using $(\mathcal{I}_{1,N}, \frac{1}{4} + \theta, \theta')$ -EQUALLY SPACED CODE that computes $F_{1,N}$ against a $(\frac{1}{4} + \theta')$ fraction of corruptions per phase. We can apply Lemmas VI.2 and VI.3 in sequence to get a 1-phase protocol Π' using $(\mathcal{I}_{3,N'}, \frac{1}{4} + \theta, \theta')$ -EQUALLY SPACED CODE that computes $F_{3,N'}$ against a $(\frac{1}{4} + \theta')$ fraction of corruptions, where $N' = \log^{\theta'}(\theta'N)$. Furthermore, Lemma VI.4 shows Alice must use

Algorithm 2 The protocol computing $F_{3,N}$.

Input: Alice has input $A = (x_1^A, \{x_2^A, x_3^A\}) \in [N] \times \binom{[N]}{2}$, and Bob has input $B = \{x_1^B, x_2^B\}$.

Output: Bob outputs $y \in [N]$.

Alice:

- 1: Alice finds the smallest index $i_1 \in [\log N]$ such that $x_{1,i_1}^A \neq x_{2,i_1}^A$ and the smallest index $i_2 \in [\log N]$ such that $x_{1,i_2}^A \neq x_{3,i_2}^A$. Alice also tries to find an index $j \in [\log \log N]$ such that $i_{1,j} \neq i_{2,j}$.
- 2: Alice sends a codeword to Bob as follows (the first satisfying case applies):

$$\begin{cases} (0, 0, 0), & x_2^A, x_3^A \geq x_1^A \\ (0, 1, 1), & x_2^A, x_3^A \leq x_1^A \\ \left(j, \mathbb{1}(x_1^A > x_2^A), \mathbb{1}(x_1^A > x_3^A) \right), & i_{1,j} < i_{2,j} \\ \left(j, \mathbb{1}(x_1^A > x_3^A), \mathbb{1}(x_1^A > x_2^A) \right), & i_{1,j} > i_{2,j} \end{cases} \quad (3)$$

Bob:

- 3: Bob finds the smallest index $i' \in [\log N]$ such that $x_{1,i'}^B \neq x_{2,i'}^B$.
- 4: Bob receives (j', b_0, b_1) from Alice and outputs

$$y = \begin{cases} \min(x_1^B, x_2^B), & (j', b_0, b_1) = (0, 0, 0) \text{ or } b_{i'} = 0 \\ \max(x_1^B, x_2^B), & (j', b_0, b_1) = (0, 1, 1) \text{ or } b_{i'} = 1 \end{cases}$$

at least $N'' = \log \log N'$ distinct codewords across all possible inputs in Π' . However, this contradicts the assumption that Π' is using $(\mathcal{I}_{3,N'}, \frac{1}{4} + \theta, \theta')$ -EQUALLY SPACED CODE since $\frac{1}{2} + \frac{1}{2(N''-1)} + \theta'$ is less than $2 \cdot (\frac{1}{4} + \theta')$ under the above assumptions. \blacksquare

B. Tightness of Lemma VI.4

In this section, we present a simple 1-phase protocol computing $F_{3,N}$. It shows the lower bound of Lemma VI.4 is essentially tight up to constant factors.

Theorem VI.5. *There exists a 1-phase protocol computing $F_{3,N}$ in which Alice uses $O(\log \log N)$ distinct codewords across all possible inputs.*

The protocol proving Theorem VI.5 is presented in Algorithm 2. In Algorithm 2, we identify an integer $x \in [N]$ with a $(\log N)$ -bit string corresponding to its binary representation, from its most significant bit to its least significant bit. Similarly, we also identify an integer $i \in [\log N]$ with a $(\log \log N)$ -bit string corresponding to its binary representation, from its most significant bit to its least significant bit. The length of the string can always be inferred from context.

It is not hard to see Alice uses $O(\log \log N)$ distinct codewords across all possible inputs. Now, we first show that the protocol is well-defined as (3) covers all possible cases. In particular, if there exists no such index j (equivalently, $i_1 = i_2$), one of the first two cases of (3) must hold.

Claim VI.6. If $i_1 = i_2$, then it holds that either $x_1^A = \min(x_1^A, x_2^A, x_3^A)$ or $x_1^A = \max(x_1^A, x_2^A, x_3^A)$.

Then, we finish the proof of Theorem VI.5 by showing the correctness of Algorithm 2.

Proof of Theorem VI.5: When $j' = 0$, since $x_1^A \in \{x_1^B, x_2^B\} \subseteq \{x_1^A, x_2^A, x_3^A\}$, it is clearly that either $y = \min(x_1^B, x_2^B) = \min(x_1^A, x_2^A, x_3^A) = x_1^A$ or $y = \max(x_1^B, x_2^B) = \max(x_1^A, x_2^A, x_3^A) = x_1^A$. Now consider the case where $j' = j \in [\log N]$. Assume without loss of generality that $i_{1,j} < i_{2,j}$, implying that $i_{1,j} = 0$ and $i_{2,j} = 1$. (The other case where $i_{1,j} > i_{2,j}$ is symmetric.) Observe that Bob must have $i' = i_1$ if $\{x_1^B, x_2^B\} = \{x_1^A, x_2^A\}$ and $i' = i_2$ if $\{x_1^B, x_2^B\} = \{x_1^A, x_3^A\}$. Moreover, Bob can distinguish between the cases by simply examining the value of $i'_{j'}$: $i'_{j'} = 0$ implies the former case while $i'_{j'} = 1$ implies the latter one. By the construction of Alice's message in (3), $b_{i'_{j'}}$ always conveys the critical information about the relative order of the correct answer x_1^A in the two-candidate set $\{x_1^B, x_2^B\}$. Therefore, the correctness of y follows as Bob already figures out which one of b_0, b_1 conveys the correct information. ■

REFERENCES

- [1] C. E. Shannon, "A mathematical theory of communication," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 5, no. 1, pp. 3–55, 2001. Originally appeared in *Bell System Tech. J.* 27:379–423, 623–656, 1948.
- [2] M. Plotkin, "Binary codes with specified minimum distance," *IRE Transactions on Information Theory*, vol. 6, no. 4, pp. 445–450, 1960.
- [3] B. Haeupler, P. Kamath, and A. Velingker, "Communication with partial noiseless feedback," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, ser. LIPIcs, vol. 40, 2015, pp. 881–897.
- [4] M. Gupta, Y. T. Kalai, and R. Y. Zhang, "Interactive error correcting codes over binary erasure channels resilient to 1/2 adversarial corruption," in *Symposium on Theory of Computing (STOC)*, 2022.
- [5] A. Orlitsky, "Worst-case interactive communication. i. two messages are almost optimal," *IEEE Transactions on Information Theory*, vol. 36, no. 5, 1990.
- [6] E. Haramaty and M. Sudan, "Deterministic compression with uncertain priors," *Algorithmica*, vol. 76, no. 3, pp. 630–653, 2016.
- [7] K. Efremenko, G. Kol, R. Saxena, and Z. Zhang, "Binary codes with resilience beyond 1/4 via interaction," *Electron. Colloquium Comput. Complex.*, p. 129, 2022.
- [8] M. Gupta and R. Y. Zhang, "Positive rate binary interactive error correcting codes resilient to 1/2 adversarial erasures," *CoRR*, vol. abs/2201.11929, 2022. [Online]. Available: <https://arxiv.org/abs/2201.11929>
- [9] K. Efremenko, G. Kol, and R. R. Saxena, "Binary interactive error resilience beyond 1/8," in *Foundations of Computer Science (FOCS)*, 2020, pp. 470–481.
- [10] M. Gupta and R. Y. Zhang, "The optimal error resilience of interactive communication over binary channels," in *Symposium on Theory of Computing (STOC)*, 2022.
- [11] E. R. Berlekamp, "Block coding for the binary symmetric channel with noiseless, delayless feedback," *Error-correcting codes*, pp. 61–68, 1968.
- [12] K. Efremenko, R. Gelles, and B. Haeupler, "Maximal noise in interactive communication over erasure channels and channels with feedback," *IEEE Trans. Inf. Theory*, vol. 62, no. 8, pp. 4575–4588, 2016.
- [13] L. J. Schulman, "Communication on noisy channels: A coding theorem for computation," in *Foundations of Computer Science (FOCS)*, 1992, pp. 724–733.
- [14] —, "Deterministic coding for interactive communication," in *Symposium on Theory of Computing (STOC)*, 1993, pp. 747–756.
- [15] —, "Coding for interactive communication," *IEEE Transactions on Information Theory*, vol. 42, no. 6, pp. 1745–1756, 1996.
- [16] R. Gelles, "Coding for interactive communication: A survey," *Foundations and Trends® in Theoretical Computer Science*, vol. 13, no. 1–2, pp. 1–157, 2017.
- [17] M. Braverman and A. Rao, "Towards coding for maximum errors in interactive communication," in *Symposium on Theory of Computing (STOC)*, 2011, pp. 159–166.
- [18] M. Braverman and K. Efremenko, "List and unique coding for interactive communication in the presence of adversarial noise," *SIAM J. Comput.*, vol. 46, no. 1, pp. 388–428, 2017.
- [19] M. Franklin, R. Gelles, R. Ostrovsky, and L. J. Schulman, "Optimal coding for streaming authentication and interactive communication," *IEEE Transactions on Information Theory*, vol. 61, no. 1, pp. 133–145, 2015.
- [20] R. Gelles and B. Haeupler, "Capacity of interactive communication over erasure channels and channels with feedback," *SIAM Journal on Computing*, vol. 46, no. 4, pp. 1449–1472, 2017.
- [21] D. Pankratov, "On the power of feedback in interactive channels," *Manuscript*, 2013.
- [22] M. Braverman, R. Gelles, J. Mao, and R. Ostrovsky, "Coding for interactive communication correcting insertions and deletions," *IEEE Transactions on Information Theory*, vol. 63, no. 10, pp. 6256–6270, 2017.
- [23] A. A. Sherstov and P. Wu, "Optimal interactive coding for insertions, deletions, and substitutions," in *Foundations of Computer Science (FOCS)*, 2017, pp. 240–251.
- [24] B. Haeupler, A. Shahrasbi, and E. Vitercik, "Synchronization strings: Channel simulations and interactive coding for insertions and deletions," in *International Colloquium on Automata, Languages, and Programming (ICALP)*, vol. 107, 2018, pp. 75:1–75:14.
- [25] M. Ghaffari, B. Haeupler, and M. Sudan, "Optimal error rates for interactive coding i: Adaptivity and other settings," in *Symposium on Theory of Computing (STOC)*, 2014, pp. 794–803.
- [26] M. Ghaffari and B. Haeupler, "Optimal Error Rates for Interactive Coding II: Efficiency and List Decoding," in *Foundations of Computer Science (FOCS)*, ser. FOCS, 2014, pp. 394–403.
- [27] K. Efremenko, G. Kol, and R. Saxena, "Interactive error resilience beyond 2/7," in *Symposium on Theory of Computing (STOC)*. ACM, 2020.
- [28] K. Efremenko, G. Kol, and R. R. Saxena, "Optimal error resilience of adaptive message exchange," in *Symposium on Theory of Computing (STOC)*, 2021, pp. 1235–1247.
- [29] E. R. Berlekamp, "Block coding with noiseless feedback," Ph.D. dissertation, Massachusetts Institute of Technology (MIT), 1964.
- [30] C. E. Shannon, "The zero error capacity of a noisy channel," *IRE Transactions on Information Theory*, vol. 2, no. 3, pp. 8–19, 1956.
- [31] K. Zigangirov, "On the number of correctable errors for transmission over a binary symmetrical channel with feedback," *Problems of Information Transmission*, vol. 12, pp. 85–97, 1976.
- [32] J. Spencer, P. Winkler, and S. St., "Three thresholds for a liar," *Combinatorics, Probability and Computing*, vol. 1, pp. 81–93, 1992.
- [33] R. Ahlswede, C. Deppe, and V. S. Lebedev, "Non-binary error correcting codes with noiseless feedback, localized errors, or both," in *International Symposium on Information Theory (ISIT)*, 2006, pp. 2486–2487.
- [34] G. Wang, Y. Qin, and C. Chang, "Communication with partial noisy feedback," in *IEEE Symposium on Computers and Communications (ISCC)*, 2017, pp. 602–607.