

A Noise Resilient Transformation for Streaming Algorithms

Meghal Gupta *
UC Berkeley

Rachel Yun Zhang†
MIT

July 14, 2023

Abstract

In a streaming algorithm, Bob receives an input $x \in \{0, 1\}^n$ via a stream and must compute a function f in low space. However, this function may be fragile to errors in the input stream. In this work, we investigate what happens when the input stream is corrupted. Our main result is an encoding of the incoming stream so that Bob is still able to compute any such function f in low space when a constant fraction of the stream is corrupted.

More precisely, we describe an encoding function $\text{enc}(x)$ of length $\text{poly}(n)$ so that for any streaming algorithm A that on input x computes $f(x)$ in space s , there is an explicit streaming algorithm B that computes $f(x)$ in space $s \cdot \text{polylog}(n)$ as long as there were not more than $\frac{1}{4} - \varepsilon$ fraction of (adversarial) errors in the input stream $\text{enc}(x)$.

*E-mail: meghal@berkeley.edu. This author was supported by an UC Berkeley Chancellor's fellowship.

†E-mail: rachelyz@mit.edu. This research was supported in part by DARPA under Agreement No. HR00112020023, an NSF grant CNS-2154149, and NSF Graduate Research Fellowship 2141064.

Contents

1	Introduction	1
1.1	Our Results	1
1.2	Open Questions	3
1.3	Related Works	4
2	Overview of Techniques	5
3	Preliminaries	8
3.1	Error-Correcting Codes	8
4	Locally Decoding with Confidences	8
5	Noise Resilient Streaming for Linear Algorithms	11
5.1	Statement of Algorithm	11
5.2	Proof of Theorem 1.3	13
5.2.1	Algorithmic Complexity	13
5.2.2	Correctness	15
5.2.3	Proof of Lemma 5.6	16
6	Noise Resilient Streaming for General Algorithms	18
6.1	Statement of Algorithm	18
6.2	Proof of Theorem 1.1	20
6.2.1	Algorithmic Complexity	21
6.2.2	Correctness	22
6.2.3	Proof of Lemma 6.6	23

1 Introduction

In a streaming algorithm, Bob witnesses a stream of bits x_1, \dots, x_n , and he wishes to compute a function $f(x_1, \dots, x_n)$. The goal is to do this in space significantly less than that necessary to store the entire stream.

However, streaming algorithms are very rigid when it comes to errors in the stream. For instance, corruption of even one bit of the stream can change the value of the parity function, an index function, or the result of a decision tree. In this work, we study what happens when the stream itself is corrupted. In particular, we ask if it's possible to convert a given streaming algorithm into one that is robust to errors in the stream.

In order to protect the output against corruption to the stream, the sender, Alice, encodes her message as $\text{enc}(x)$ and sends Bob $\text{enc}(x)$ bit-by-bit. Then, using the stream $\text{enc}(x)$, Bob has to compute $f(x)$ in low space.

Of course, if Alice knows the function f that Bob wishes to compute, this all is very simple: she can simply compute $f(x) \in \{0, 1\}$ and send Bob $f(x)$ using e.g. the repetition code.¹ What we actually desire is for the encoding $\text{enc}(x)$ to be agnostic to the function that Bob wishes to compute: the encoding $\text{enc}(x)$ must work for a broad class of functions.

As an example, consider the case where $f(x) = f_y(x) = x \cdot y \pmod 2$ for some $y \in \{0, 1\}^n$. In the noiseless setting, Alice can simply send x bit by bit to Bob, and after receiving each bit Bob adds $x_i \cdot y_i$ to a running sum. Bob only needs to track i and the running sum modulo 2, which is in total $\log n + 1$ bits of space.

However, when there are any errors in the input stream x , the value $f_y(x)$ Bob computes may be wrong. If Alice instead uses a generic large distance error-correcting code to encode x , the stream's information is now protected against error but the issue is that Bob must receive and store all of the received word in order to decode, which is far too much space. Even if Alice sends a locally decodable code, which has the property that to determine a single bit of x one only needs to track a few random bits of the codeword, it is not clear how Bob can decode all n bits without $\Omega(n)$ storage.

In this work, we investigate whether such an encoding $\text{enc}(x)$ exists that allows Bob to compute $f(x)$ in low space, regardless of which function f Bob wishes to compute.

Our main result is a scheme that protects any streaming algorithm against noise in the incoming stream. More precisely, we give an encoding $\text{enc}(x)$ of length $O(n^{4+\delta})$ such that any streaming algorithm running in s space on the noiseless stream can be converted to one running in $s \cdot \text{polylog}(n)$ space on the encoded stream. It is correct with high probability whenever at most $\frac{1}{4} - \varepsilon$ of the stream was corrupted. In the specific case where Bob's function is a linear function such as dot product, our encoding can be made to have length $O(n^{2+\delta})$.

1.1 Our Results

Our main result is a noise-resilient conversion for deterministic streaming algorithms.

Theorem 1.1. *Fix $\varepsilon, \delta > 0$. There is a function $\text{enc} : X \subseteq \{0, 1\}^n \rightarrow \{0, 1\}^m$ where $m = O_\varepsilon(n^{4+\delta})$ and an explicit transformation B such that the following holds: For any deterministic algorithm A that takes as input $x \in X \subseteq \{0, 1\}^n$ as a stream, runs in space s and time t , and outputs $A(x)$, the algorithm $B_A = B(A)$ takes as input $z \in \{0, 1\}^m$ as a stream, runs in space $s \cdot O_\varepsilon((\log n)^{O(1/\delta)})$ and*

¹In fact, she could have just sent Bob $f(x)$ instead of x in the noiseless version as well.

time $m \cdot O_{\varepsilon, \delta} \left(1 + \frac{t}{n^2}\right)$, and satisfies that whenever $\Delta(z, \text{enc}(x)) < \left(\frac{1}{4} - \varepsilon\right) \cdot m$ then $B_A(z)$ outputs $A(x)$ with probability $\geq 1 - 2^{O_\varepsilon(\log^2 n)}$.

In other words, given an algorithm A that accepts a stream of length n and uses s space, we demonstrate a noise resilient algorithm $B(A)$ computing A that uses $s \cdot \text{polylog}(n)$ space. The transformation enc of the stream x from the sender is independent of A and has length $O(n^{4+\delta})$.

A priori, it is not obvious that there is *any* low-memory noise resilient transformation of the algorithm A even with an arbitrary blow-up in communication. For example, if the sender were to encode their message using an arbitrary error-correcting code, the receiver would need to store the entire message in memory in order to decode it before applying the algorithm A . Our result shows that not only does this low-memory transformation exist, but that it can be done efficiently (both with respect to communication complexity and computational complexity).

Linear Streaming Algorithms. In the above transformation, the length of the resulting stream is $O(n^{4+\delta})$ for any $\delta > 0$. In the specific case where the streaming algorithm is *linear*, we construct a scheme where the stream length is only quadratic, namely $O(n^{2+\delta})$. The key property of linear streaming algorithms that we will be leveraging is that they can be computed in pieces in any order, and later computations do not depend on the results of previous ones.

Definition 1.2. [*Linear Streaming Algorithms*] A linear streaming algorithm $A : \{0, 1\}^n \rightarrow \mathbb{F}_q$ is described by a list of functions $g_i : \{0, 1\} \rightarrow \mathbb{F}_q$ for $i \in [n]$, and computes the value $A(x) = g_1(x_1) + \dots + g_n(x_n)$ (where addition is over \mathbb{F}_q) by tracking the partial sum $g_1(x_1) + \dots + g_i(x_i)$ upon receiving the i 'th bit.

Note, for example, that every linear function on codomain $\{0, 1\}^n$ admits a linear streaming algorithm. We note that linear streaming algorithms capture a large class of interesting functions, including linear sketching algorithms (see for example [CM05, CCFC04, AGM12, JL84]). For linear streaming algorithms, we show the following result.

Theorem 1.3. Fix $\varepsilon, \delta > 0$. There is a function $\text{enc} : X \subseteq \{0, 1\}^n \rightarrow \{0, 1\}^m$ where $m = O_\varepsilon(n^{2+\delta})$ and an explicit transformation B such that the following holds: For any linear streaming algorithm A that takes as input $x \in X \subseteq \{0, 1\}^n$ as a stream, runs in space s and time t , and outputs $A(x)$, the algorithm $B_A = B(A)$ takes as input $z \in \{0, 1\}^m$ as a stream, runs in space $s \cdot O_\varepsilon((\log n)^{O(1/\delta)})$ and time $m \cdot O_{\varepsilon, \delta} \left(1 + \frac{t}{n^2}\right)$, and satisfies that whenever $\Delta(z, \text{enc}(x)) < \left(\frac{1}{4} - \varepsilon\right) \cdot m$ then $B_A(z)$ outputs $A(x)$ with probability $\geq 1 - 2^{O_\varepsilon(\log^2 n)}$.

Randomized Algorithms. We remark that our transformations in Theorems 1.1 and 1.3 are only for deterministic algorithms. However, this easily implies the result for algorithms A that *pre-sample* all of their randomness (that is, algorithms that fix their randomness before receiving any bits of x , at which point they are deterministic for the remainder of the algorithm). Such algorithms can be viewed as sampling from a distribution over deterministic algorithms. Our transformation can then be applied to the ensuing deterministic algorithm, thus correctly computing the function A with high probability while being resilient to noise.

We remark that this notion of randomized algorithms whose only access to randomness is pre-sampled at the start of computation is quite general: randomized algorithms that sample its randomness online can often be converted to protocols where the randomness is pre-sampled, see e.g. [Nis90].

Non-Binary Input Alphabets. Our main theorems are stated for the setting where the input stream is binary. We remark that this assumption is without loss of generality: a larger alphabet stream can be converted to a binary one by assigning each alphabet symbol to a unique binary string.

1.2 Open Questions

Our work opens up many interesting questions. We list a few below.

First, in the case of linear streaming algorithms, Theorem 1.3 gives a scheme where the input stream is length $O(n^{2+\delta})$. We conjecture that even in the simple case where A is a linear function over \mathbb{F}_2 , a communication complexity of at least $\tilde{\Omega}(n^2)$ is necessary to make the computation noise resilient and still low-memory. This would imply that Theorem 1.3 is essentially tight.

Open Question 1.1. *Does there exist $\varepsilon > 0$ and a function $\text{enc} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ with $m = O(n^{2-\delta})$ for some $\delta > 0$ such that the following holds? For any linear function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, there exists an algorithm B that uses $\text{polylog}(n)$ bits of space with the guarantee that for every input $z \in \{0, 1\}^m$ with $\Delta(z, \text{enc}(x)) < \varepsilon$, B outputs $f(z)$ with probability $> \frac{2}{3}$.*

It would be interesting to disprove even the existence of a positive rate ($O(n)$ length) encoding enc that admits a noise resilient transformation for all streaming algorithms with $\text{polylog}(n)$ space overhead. More specifically, we ask what the optimal length of such an encoding is (recall that our transformation for general algorithms requires a quartic encoding length).

Open Question 1.2. *What is the optimal exponent β such that there exists an encoding $\text{enc}(x) : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and $\varepsilon > 0$ where $m = O(n^\beta)$, such that for any streaming algorithm A that runs in space s , there exists a streaming algorithm B that reads $z \in \{0, 1\}^m$ and runs in space $s \cdot \text{polylog}(n)$, with the property that whenever $\Delta(z, \text{enc}(x)) < \varepsilon$, B outputs $A(x)$ with probability $> \frac{2}{3}$?*

Finally, we present an open question about the space blow-up incurred by our algorithm. The space blow-up of the algorithms in our paper is $\text{polylog}(n)$. Most deterministic algorithms require at least logarithmic space (for example, even to just count the number of bits in the input), so a $\text{polylog}(n)$ blow-up is not that large. However, some algorithms avoid this. For example, consider the following streaming problem.

Definition 1.4 (Index Problem). *The stream is the list of pairs in order (i, y_i) for some hidden string $y \in \{0, 1\}^n$. The solver has a target index j and wants to output y_j .*

There is a deterministic algorithm that accomplishes this task in $O(\log \log(n))$ bits of space.² However, our transformation would increase the space to $\text{polylog}(n)$, which is an exponential blow-up in the space complexity. We ask whether this blow-up is necessary.

Open Question 1.3. *Does there exist $\varepsilon > 0$, a function $\text{enc} : \{0, 1\}^n \rightarrow \{0, 1\}^m$, and a list of algorithms $B_1 \dots B_n$ such that the following holds? For all $i \in [n]$ and for every input $z \in \{0, 1\}^m$ with $\Delta(z, \text{enc}(y)) < \varepsilon$, the algorithm B_i runs in space $\text{polyloglog}(n)$ and outputs y_i with probability $> \frac{2}{3}$.*

²The algorithm to do this is as follows. Upon receiving the pair (i, y_i) in a stream, Bob checks if $i = j$ in $\log \log n$ bits of space (by checking if each bit in the binary representation of i and j are equal), which can be done because both i and j are length $\log n$. Then, if $i = j$, he outputs y_i .

³We remark that in the noiseless protocol, the stream itself is not the target vector y ; it is $x(y) := (1, y_1) \dots (n, y_n)$. In the previous discussion, we had been describing enc as a function of the noiseless stream x , but here it is simpler to think of it as a function of the target vector y (although either definition is fine).

1.3 Related Works

Streaming Algorithms. The study of streaming algorithms began with the work Morris [Mor78] on approximate counting; a rigorous analysis was given later in [Fla85]. Later works introduced other classic problems in streaming, including heavy hitters [CCFC02], ℓ_p approximation [AMS96, MW10, IW05], and finding a nonzero entry in a vector (for turnstile algorithms) [MW10].

Many works, for example [GKLR21, CZ16, MWOC21], consider the problem of processing noisy data sets using streaming algorithms. [GKLR21] shows a memory lower bound for learning a parity function with noisy samples of random linear equations.

However, this type of noise is quite different from our treatment. In these problems, noise is inherent in the way samples are generated, whereas we are investigating noise that occurs in the communication process.

Error Correcting Codes. Our result can be viewed through the lens of low-space noise resilient one-way communication.

Noise resilient one-way communication, in other words, error correction [Sha48, Ham50], is one of the most well-studied problems in computer science. One specific line of works related to our result is that of *locally decodable codes*. Locally decodable codes [Y⁺12] can be viewed as a low-time and low-space version of error-correcting codes, where the goal is to learn a single bit of the original message. By contrast, for us, the decoder must be able to piece together any function of the input that is noiselessly computable in low space, with the tradeoff that the decoder accesses the entire encoding via a stream rather than via query access. Local decodable codes have been constructed in a variety of different parameter regimes, including constant query [Efr09, DGY11] and near 1 rate [KSY14]. In our work, we will use Reed-Muller codes [Mul54, Ree54] that achieve polylog(n) query complexity and slightly super-linear block length.

Another line of works related to ours is that of streaming algorithms for local testing of codes [RU10, MRU11]. However, this direction was rendered moot by the recent discovery of locally testable codes with constant rate, distance, and locality [DEL⁺22].

Coding for Interactive Communication. The analogue of noise resilient one-way communication in the interactive setting, where Alice and Bob have inputs x, y and engage in a protocol to compute some $f(x, y)$ while being resilient to error, comprises the study of *coding for interactive communication*. Interactive coding was studied starting with the seminal works of Schulman [Sch92, Sch93, Sch96] and continuing in a prolific sequence of followup works, including [BR11, Bra12, BK12, BN13, Hae14, BE14, DHM⁺15, GHK⁺16, GH17, EGH16, GH13, GI18, EKS20, GZ22a, GZ22b]. We refer the reader to an excellent survey by Gelles [Gel17] for an extensive list of related work.

The recent work of [EHK⁺23] studies the space complexity of interactive coding schemes. Their main result is an interactive coding scheme that preserves the space complexity of the original noiseless protocol up to a $\log(n)$ factor, where n is the length of the protocol. Their result can be viewed as the interactive analogue of ours, where their noiseless and noise resilient protocols are both interactive. We remark that their techniques are quite different than ours and do not apply to our setting, however, since their approach crucially relies on the communication of feedback from Bob to Alice.

2 Overview of Techniques

Consider the task of computing a linear function $f(x) = \sum_{i \in [n]} g_i(x_i)$. In a noiseless setting, one can compute $f(x)$ by tracking the partial sum $\sum_{1 \leq i \leq n'} g_i(x_i)$ and updating it with $g_{n'+1}(x_{n'+1})$ upon receiving the $(n' + 1)$ 'th bit of x .

As a first attempt to create a noise resilient version of this algorithm, one might consider the input stream $\text{enc}(x) = \text{LDC}(x)$, where LDC is a locally decodable code. In short, a locally decodable code is a noise resilient encoding of x such that in order to decode any specific bit of x , say x_i , the decoder simply has to read $\text{polylog}(n)$ (randomized) bits of $\text{LDC}(x)$. In a streaming setting, the decoder can record only these $\text{polylog}(n)$ bits into memory, then run the decoding algorithm to determine x_i after all such query bits have been seen.

However, it's not clear that this local decoding property extends to arbitrary linear functions.⁴ If we attempt to simultaneously decode all bits x_i from $\text{LDC}(x)$ using a separate query set Q_i for each index, the issue is that we will need to track $\Omega(n)$ queries simultaneously since we cannot ensure that the local decoding for any index finishes before another: the query sets Q_i are typically randomized so that any bit of the codeword is equally likely to belong to the query set for any particular index.

A second attempt is to send $\text{LDC}(x)$ n times, with the intent that Bob locally decodes x_i in the i 'th $\text{LDC}(x)$ and computes $f(x)$ by tracking the partial sum $\sum_{1 \leq i \leq n'} g_i(x_i)$ at the end of the i 'th chunk. Now, the only space requirements are the space required to store this partial sum and the space required to locally decode x_i in each $\text{LDC}(x)$, which is $\text{polylog}(n)$. However, the adversary can simply corrupt the first $\text{LDC}(x)$ so that Bob mis-decodes x_1 , thereby corrupting his final output, and ultimately, this approach is not so different from just sending the bits $x_1 \dots x_n$ in order.

What Bob would like to do is randomize which x_i he is computing from a given copy of $\text{LDC}(x)$ so that the adversary does not know. This crucially uses the property that LDC's allow Bob to decode *any* x_i not known to the encoder or adversary. Then, if he computes each x_i many times and takes the majority, he will learn x_i correctly with high probability, regardless of the adversary's corruption process. Since the adversary does not know in advance which chunks Bob is computing x_i in, she cannot target corruptions towards a specific x_i . However, this runs into an issue with the amount of space Bob requires. Since he computes each x_i many times throughout the protocol and only takes the majority when he has finished these computations, he needs to track his current guess for its value through (almost) the entire protocol. When each x_i was computed only once, Bob could just add $g_i(x_i)$ to the running sum, but now he needs to track each x_i individually between the first and last time he queries it, for a total of $\Omega(n)$ bits of space.

The crux of our construction is an organized system for aggregating decoding attempts in space much smaller than storing all n attempts simultaneously.

A recursive computation. Let $r = \text{polylog}(n)$. Suppose that we already know how to compute linear functions on inputs of size n/r in space $s_{n/r}$, that is resilient to $\varepsilon_{n/r}$ fraction of errors in the stream of $M_{n/r}$ copies of $\text{LDC}(x)$. We will use this to build a scheme for computing linear functions on inputs of size n . Note that the base case is given by computing linear functions on a single bit, which can be done by local decoding on a single $\text{LDC}(x)$.

⁴The smallest local decodable code we know of that supports local decoding to any arbitrary linear function is an extension of the Hadamard code, which has block length exponential in n . This will not be a satisfactory solution for us because (a) we would like encodings that are efficiently computable, and (b) keeping track of which bit of the stream one is receiving takes n memory.

To compute $f(x) = \sum_{i \in [n]} g_i(x_i)$, we just need to compute $f_1(x) = \sum_{i \in [n/r]} g_i(x_i)$, $f_2(x) = \sum_{i \in [n/r+1, 2n/r]} g_i(x_i)$, \dots , $f_r(x) = \sum_{i \in [n-n/r+1, n]} g_i(x_i)$. Each sub-computation can individually be done with the guarantee that if $< \varepsilon_{n/r}$ fraction of the stream is corrupted, then the function is computed correctly.

Consider a stream consisting of $\ell \cdot r \cdot M_{n/r}$ copies of $\text{LDC}(x)$ (you can think of $\ell = \log^2 n$). We split up this stream into ℓ chunks each consisting of r blocks of $M_{n/r}$ $\text{LDC}(x)$'s. In each of the ℓ chunks, we will assign the r blocks of $M_{n/r}$ LDC 's to the computation of a random permutation of $f_1(x), \dots, f_r(x)$. Throughout the algorithm, we will keep track of all the outputs from each of the ℓ sub-computations for each of $f_1(x), \dots, f_r(x)$. At the end, we can take a majority vote for each of the $f_j(x)$. We illustrate the recursion process in Figure 1.




Figure 1: In our transformation, each sub-function is computed ℓ times instead of 1, and in each of the ℓ chunks, the sub-functions f_i are computed in according to a random permutation π .

Because of the way we've randomly assigned blocks to f_1, \dots, f_r , the adversary cannot concentrate her attack on the computation of any specific f_i . More precisely, we can show that the amount of error present in the blocks corresponding to the computation of f_k is (roughly) proportional to the total error present. Then, as long as at least half the blocks corresponding to f_k have $< \varepsilon_{n/r}$ fraction of errors, we have that the majority vote for each of f_1, \dots, f_r will be correct. In total, this means that our algorithm for computing $f(x)$ is resilient to $\frac{\varepsilon_{n/r}}{2}$ fraction of errors.

Unfortunately, this factor of 2 loss in the error resilience in each iteration of the recursion will prove to be problematic, as we can no longer guarantee a positive error resilience when n gets large.

Decoding with confidences. In the above analysis, our worst case error distribution was that where just under one-half of the chunks have $\varepsilon_{n/r}$ error while the rest are error-free. In this case, the chunks with $\varepsilon_{n/r}$ error will produce a wrong output, while the error-free ones will produce a correct output.

However, we notice that in certain local decoding constructions (for us, we will use a Reed-Muller code), one can actually obtain an estimate of how much error there was in the codeword based on the inconsistency of the queries. In particular, in chunks where there were $\varepsilon_{n/r}$ errors, even though we will obtain a wrong answer, we can also see that there were a lot of errors in the codeword, and thereby be less sure of the obtained answer.

This motivates our definition of *local decoding with confidences*: in addition to outputting an answer, the decoding algorithm should also output a *confidence* $c \in [0, 1]$ indicating how confident it is that the answer is correct. A confidence of 0 indicates that it's not sure at all: the fraction of errors seen was near the threshold ε_1 ; whereas a confidence of 1 indicates that it's quite sure: either

the answer is right and no errors occurred, or at least $2\varepsilon_1$ errors were required to obtain a wrong answer with confidence 1. Another way to view the confidence is as an estimate of how much error was present in the total codeword (how far the actual error is from ε_1).

Now, when we perform the recursive step to obtain a noise resilient algorithm for f from f_1, \dots, f_r , in addition to recording the outputted answer from each chunk, we will also record the associated confidence. Suppose that the ℓ answers and confidences obtained for f_k were $(\hat{q}^{(1)}, \hat{c}^{(1)}), \dots, (\hat{q}^{(\ell)}, \hat{c}^{(\ell)})$. Then, the weighted majority vote of these answers is defined to be the \hat{q} with the largest total confidence. We can show that as long as the cumulative fractional error c_n is at most $c_{n/r}$, we're guaranteed that the correct answer has more the highest total confidence.

We remark that one does not have to store all ℓ pairs $(\hat{q}^{(j)}, \hat{c}^{(j)})$. Rather, we can keep track of a most likely q and associated confidence c . To update (q, c) with a new pair (\hat{q}, \hat{c}) , one can update c by $\pm \hat{c}$ ($+\hat{c}$ if $\hat{q} = q$, and $-\hat{c}$ otherwise). If the resulting confidence c is less than 0, then it means that the answer \hat{q} is supported by less overall error, so we set our most likely q to be \hat{q} and flip the sign of c .

Now, we analyze the space and communication complexity of the resulting algorithm. The space required is simply $s_{n/r}$, the space required to do a size n/r sub-computation, plus the space required to store a pair (q, c) for each of f_1, \dots, f_r . Overall, this means that for each layer of the recursion, we are only gaining an additive factor of $r \cdot (|q| + |c|)$ space overhead, so the total space overhead is at most polylogarithmic in n . As for the size of the resulting stream, we see that the number of LDC(x)'s is $\approx (r \cdot \ell)^{\log_r n}$,⁵ which is polynomial in n .

Computing non-linear functions. Our algorithm so far has only captured functions f that are splittable into r sub-computations whose computations do not depend on each other. However, for a general function, this may not be the case. Indeed, for many functions, the computations must be done *sequentially*, where later computations can only be done once earlier computations have finished. We may split a general function f into r sub-functions where $q_1 := f_1(x[1 : n/r], \emptyset)$ is the state of the algorithm after n/r bits of the stream, $q_2 := f_2(x[n/r + 1 : 2n/r], q_1)$ is the state of the algorithm after receiving the next n/r bits of the stream, and so on. The difference from the linear function case is that to compute f_j , we must have the correct output of f_{j-1} .

In order to handle such sequential computations, we modify our above algorithm as follows. In the recursive layer computing f from functions f_1, \dots, f_r , we always compute f_j from the starting state q_{j-1} which is our current best guess for the output of f_{j-1} . This value of q_{j-1} may not always be correct, in which case the computation done for f_j will be doomed.

As with our algorithm for linear functions, after every chunk, we update the most likely state q_j and confidence c_j for the sub-function that was computed in that chunk. However, when a state q_j changes, this also means that computations done for $f_{j+1} \dots f_n$ were based on wrong information. In the case of linear functions, the computations are independent so this does not matter, but for sequential functions, we have to discard the guesses $q_{j+1} \dots q_n$. Formally, we set these states to \emptyset and set the corresponding confidences all to 0.

Because each sub-function q_j is only computed usefully when q_{j-1} is correct and also the value of q_j is frequently discarded, it is less clear that over many iterations the q_j 's will converge to the correct values. Nevertheless, we prove in Section 6 that over enough iterations, the guesses q_j will eventually all be correct as long as there are not too many errors in the stream.

⁵In our actual protocol, we will repeat the procedure $\log^2 n$ times and take a final majority vote, so the number of LDC(x)'s used will be $(r \cdot \ell)^{\log_r n} \cdot \log^2 n$.

Outline. The rest of this paper is organized as follows. In Section 3, we start with some preliminary theorems about error-correcting codes. In Section 4, we define our notion of local decoding with confidences and prove that Reed-Muller codes satisfy this property. Then, in Sections 5 and 6, we present our noise resilient streaming algorithms for linear and general (sequential) functions, respectively.

3 Preliminaries

Notation.

- The function \log is in base 2 unless otherwise specified.
- The set $[n]$ denotes the integers $1 \dots n$.
- We use $x[i : j]$ to denote the i 'th to j 'th bits of x inclusive. We use $x(i : j]$ to refer to the $(i + 1)$ 'th to j 'th bit inclusive.
- We use e as a variable rather than the universal constant e .

3.1 Error-Correcting Codes

We begin with some results about error-correcting codes. The first is the existence of a relative distance $\frac{1}{2}$ binary code.

Theorem 3.1. [GS00] *For all $\epsilon > 0$, there exists an explicit error-correcting code $\text{ECC}_\epsilon = \{\text{ECC}_{\epsilon,n} : \{0, 1\}^n \rightarrow \{0, 1\}^m\}_{n \in \mathbb{N}}$ with relative distance $\frac{1}{2} - \epsilon$ and $m = O_\epsilon(n)$, and a $\text{poly}_\epsilon(n)$ -time decoding algorithm $\text{DEC}_\epsilon : \{0, 1\}^m \rightarrow \{0, 1\}^n$, such that for any $x \in \{0, 1\}^n$ and $w \in \{0, 1\}^m$ satisfying $\Delta(\text{ECC}_\epsilon(x), w) < (\frac{1}{4} - \frac{\epsilon}{4}) \cdot m$, it holds that $x = \text{DEC}_\epsilon(w)$.*

Our second theorem is about the efficient decoding of Reed-Solomon codes [RS60].

Theorem 3.2 (Berlekamp-Welch; see e.g. [GS92]). *Given a collection of tuples $(\alpha_1, v_1), \dots, (\alpha_n, v_n)$ where $\alpha_i, v_i \in \mathbb{F}_q$ and a parameter $d \leq n$, there is an algorithm running in $\text{poly}(n, \log q)$ time that outputs a polynomial $g(\alpha) \in \mathbb{F}_q[\alpha]$ of degree at most $d-1$ such that $\Delta((g(\alpha_i))_{i \in [n]}, (v_i)_{i \in [n]}) < \frac{n-d+1}{2}$.*

Finally, we recall generalized minimum distance decoding, which allows us to decode concatenated codes efficiently.

Theorem 3.3 (GMD Decoding [For66, Gur06]). *Given two codes, $N_{\text{outer}} : \{0, 1\}^k \rightarrow \Sigma^n$ of distance $1 - \epsilon_{\text{outer}}$ with decoding time T_{outer} , and $N_{\text{inner}} : \Sigma \rightarrow \{0, 1\}^m$ of distance $\frac{1}{2} - \epsilon_{\text{inner}}$ with decoding time D_{inner} , there is an algorithm running in time $\text{poly}(D_{\text{inner}}, T_{\text{outer}})$ that on input $w \in \{0, 1\}^{n \cdot m}$ outputs $x \in \{0, 1\}^k$ such that $\Delta(w, N_{\text{outer}} \circ N_{\text{inner}}(x)) < \frac{(1 - \epsilon_{\text{outer}})(\frac{1}{2} - \epsilon_{\text{inner}})}{2}$.*

4 Locally Decoding with Confidences

We now recall the definition of *locally decodable codes*. For us, besides just correctness of local decoding when the distance to a codeword is small, we require that the decoder output a *confidence* indicating how much error it witnessed (less error means higher confidence).

Theorem 4.1. (*Locally Decodable Code*) For any $\varepsilon > 0$ and $d = d(n) = \log^{1/\delta} n$, there is a code $\text{LDC} : \{0, 1\}^n \rightarrow \{0, 1\}^{N(n)}$ where $N(n) = O_\varepsilon(n^{1+\delta})$ that satisfies the following properties:

- **Distance:** For any $x \neq y \in \{0, 1\}^n$, it holds that $\Delta(\text{LDC}(x), \text{LDC}(y)) \geq (\frac{1}{2} - \varepsilon) \cdot N$.
- **Locally Decoding with Confidence:** For any index $i \in [n]$ and any word $w \in \{0, 1\}^N$, there exists a randomized algorithm \mathcal{D}_i that reads $Q = O_\varepsilon(\log^{2/\delta} n)$ bits of w , runs in $O_\varepsilon(n^{O(1/\delta)})$ time, and outputs a bit \hat{x}_i along with confidence $\text{conf} \in [0, 1]$, satisfying that

- For any $x \in \{0, 1\}^n$ such that $\Delta(w, \text{LDC}(x)) = (\frac{1}{4} - 2\varepsilon - e) \cdot N$ where $e \geq 0$, it holds that

$$\hat{x}_i = x_i \quad \text{and} \quad \text{conf} > e$$

with probability at least $1 - \exp(-\log^2 n)$.

- For any $x \in \{0, 1\}^n$ such that $\Delta(w, \text{LDC}(x)) = (\frac{1}{4} - 2\varepsilon + e) \cdot N$ where $e > 0$, it holds that

$$\hat{x}_i = x_i \quad \text{or} \quad \text{conf} < e$$

with probability at least $1 - \exp(-\log^2 n)$.

Furthermore, the confidence conf can be represented as a rational number with denominator $4Q$.

Proof. The locally decodable code that we will be using is the concatenation of two codes. The outer code C_{outer} is a Reed-Muller code of $m = \log_d n$ -variate polynomials over $\mathbb{F} = \mathbb{F}_q$ of degree $d - 1$ where $q \in [2d/\varepsilon, 4d/\varepsilon]$ is a power of 2. C_{outer} has dimension $d^m = n$ and block length $N_{\text{outer}} = q^m < (4d/\varepsilon)^m = n \cdot (\frac{4}{\varepsilon})^m = n^{1+\delta \cdot \log(4/\varepsilon)/\log \log n} = O_\varepsilon(n^{1+\delta/2})$. The inner code C_{inner} is a binary code of relative distance $\frac{1}{2} - \frac{\varepsilon}{2}$ (see Theorem 3.1) with block length $N_{\text{inner}} = O_\varepsilon(\log q) = o_\varepsilon(n^{\delta/2})$. Thus, the block length of the concatenated code C is $N = N_{\text{outer}} \cdot N_{\text{inner}} = O_\varepsilon(n^{1+\delta})$. We will index the bits of the concatenated code by pairs (v, j) where $v \in \mathbb{F}^m$ and $j \in [N_{\text{inner}}]$. We remark that the outer code C_{outer} is systematic.

By the Schwartz-Zippel lemma [Sch80, Zip79], the distance of the outer code is $\geq (q - d + 1) \cdot q^{m-1} \geq (1 - \frac{\varepsilon}{2}) \cdot N_{\text{outer}}$. So the relative distance of the concatenated code is $\geq (1 - \frac{\varepsilon}{2}) \cdot (\frac{1}{2} - \frac{\varepsilon}{2}) \geq \frac{1}{2} - \varepsilon$.

The decoding algorithm is as follows. Suppose we want the i 'th index of x which is located at $v_0 \in \mathbb{F}^m$ in the outer code. Then, we do the following $k = 10 \log^2 n / \varepsilon^2$ times: Pick a random degree-2 polynomial $p(\lambda) = v_0 + v_1 \lambda + v_2 \lambda^2$ by sampling $v_1, v_2 \leftarrow \mathbb{F}^m$. Then, we query all bits of w located at $(p(\lambda), j)$ for $\lambda \in \mathbb{F} \setminus \{0\}$, $j \in [N_{\text{inner}}]$ (we denote these collection of values by $w|_p$). This means that $Q = k \cdot (q - 1) = O_\varepsilon(\log^{2/\varepsilon})$. By Theorem 3.3 (combined with Theorems 3.1 and 3.2), one can in $\text{poly}(q \log q) = O_\varepsilon(n^{O(1/\delta)})$ time find the unique degree $2d - 2$ polynomial $h \in \mathbb{F}[\lambda]$ such that the distance between $w|_p$ and the C_{inner} encodings of $h(\lambda)$ is at most $(q - 2d + 1) \cdot (\frac{1}{2} - \frac{\varepsilon}{2}) \cdot N_{\text{inner}}/2 \geq (\frac{1}{2} - \varepsilon) \cdot \frac{Q}{k}$. Then, the decoding algorithm outputs \hat{x}_i set to be the majority of all the values of $h(0)$ among the k iterations (where if such a polynomial h of distance $(q - 2d + 1) \cdot (\frac{1}{2} - \frac{\varepsilon}{2}) \cdot N_{\text{inner}}/2$ doesn't exist, we say that $h(0)$ evaluates to 0).

The confidence conf is calculated as follows. For each of the k polynomials p , we let Δ_p^h be the distance of the queried bits from the C_{inner} encodings of $h(\lambda)$. If $h(0) = \hat{x}_i$, we set $\Delta_p = \Delta_p^h$, and otherwise we set $\Delta_p = (q - 2d + 1) \cdot \frac{1}{2} - \frac{\varepsilon}{2} \cdot N_{\text{inner}}/2 - \Delta_p^h$. This gives us an estimate Δ_p of the

amount of error present in w conditioned on the output guess \hat{x}_i being the actual value of $p(0)$. In particular, $\Delta_p \leq \Delta(w|_p, \text{LDC}(x)|_p)$. Then, we set $\text{conf} := \frac{1}{4} - \sum_p \Delta_p / Q$.

We now discuss why the local decoding properties hold. First, suppose $x \in \{0, 1\}^n$ satisfies that $\Delta(w, \text{LDC}(x)) \leq (\frac{1}{4} - 2\varepsilon) \cdot N$. Then,

$$\begin{aligned} & \Pr \left[\hat{x}_i \neq x_i \text{ or } \text{conf} \leq \frac{1}{4} - 2\varepsilon - \frac{\Delta(w, \text{LDC}(x))}{N} \right] \\ & \leq \Pr[\hat{x}_i \neq x_i] + \Pr \left[\sum_p \Delta_p \geq \left(\frac{\Delta(w, \text{LDC}(x))}{N} + 2\varepsilon \right) \cdot Q \mid \hat{x}_i = x_i \right] \\ & \leq e^{-\varepsilon^2 k / 10}, \end{aligned}$$

where we've used the following two inequalities:

- By Proposition 2.7 in [Yek12], the probability that for any of the k polynomials p it holds that $p(0) \neq x_i$ is

$$\leq \frac{4(\Delta - \Delta^2)}{(q-1) \cdot ((1-\varepsilon) \cdot (\frac{1}{2} - \frac{\varepsilon}{2}) - 2\Delta)^2} \leq \frac{1}{\varepsilon(q-1)} \leq \frac{1}{\log^{1/\delta} n}$$

when $\Delta := \frac{\Delta(w, \text{LDC}(x))}{N} \leq \frac{1}{4} - 2\varepsilon$. Then, by the Chernoff bound,

$$\Pr[\hat{x}_i \neq x_i] \leq \Pr \left[\# [p : p(0) \neq x_i] \geq \frac{k}{2} \right] \leq e^{-k/10}.$$

- Given that $\hat{x}_i = x_i$, we have that $\Delta_p \leq \Delta(w|_p, \text{LDC}(x)|_p)$, so that $\mathbb{E}[\Delta_p] \leq \frac{\Delta(w, \text{LDC}(x))}{N} \cdot \frac{Q}{k}$. By Hoeffding's inequality,

$$\Pr \left[\sum_p \Delta_p \geq \left(\frac{\Delta(w, \text{LDC}(x))}{N} + 2\varepsilon \right) \cdot Q \right] < e^{-8\varepsilon^2 \cdot k}.$$

Second, suppose $x \in \{0, 1\}^n$ satisfies that $\Delta(w, \text{LDC}(x)) > (\frac{1}{4} - 2\varepsilon) \cdot N$. Then,

$$\begin{aligned} & \Pr \left[\hat{x}_i \neq x_i \text{ and } \text{conf} \leq \frac{\Delta(w, \text{LDC}(x))}{N} - \left(\frac{1}{4} - 2\varepsilon \right) \right] \\ & \leq \Pr \left[\text{conf} \leq \frac{\Delta(w, \text{LDC}(x))}{N} - \left(\frac{1}{4} - 2\varepsilon \right) \mid \hat{x}_i \neq x_i \right] \\ & = \Pr \left[\sum_p \Delta_p \geq \left(\frac{1}{2} - 2\varepsilon - \frac{\Delta(w, \text{LDC}(x))}{N} \right) \cdot Q \mid \hat{x}_i \neq x_i \right]. \end{aligned}$$

Note that since $\hat{x}_i \neq x_i$, and the restriction of $\text{LDC}(x)|_p$ to any polynomial p is a code of distance $\geq (q-2d+1) \cdot (\frac{1}{2} - \frac{\varepsilon}{2}) \cdot N_{\text{inner}} \geq (\frac{1}{2} - \varepsilon) \cdot \frac{Q}{k}$. Thus, if Γ_p denotes the actual amount of corruptions to the points of $\text{LDC}(x)|_p$, then $\Gamma_p + \Delta_p \geq (\frac{1}{2} - \varepsilon) \cdot \frac{Q}{k}$. The above inequality is then

$$\leq \Pr \left[\sum_p \Gamma_p \geq \left(\frac{\Delta(w, \text{LDC}(x))}{N} + \varepsilon \right) \cdot Q \right].$$

Since $\mathbb{E}[\Gamma_p] = \frac{\Delta(w, \text{LDC}(x))}{N} \cdot \frac{Q}{k}$, by Hoeffding's Inequality,

$$\Pr \left[\sum_p \Gamma_p \geq \left(\frac{\Delta(w, \text{LDC}(x))}{N} + \varepsilon \right) \cdot Q \right] \leq e^{-2\varepsilon^2 \cdot k}.$$

□

5 Noise Resilient Streaming for Linear Algorithms

Our first result is a noise resilient conversion for linear streaming algorithms A . We begin by recalling the definition of a linear streaming algorithm.

Definition 1.2. [*Linear Streaming Algorithms*] A linear streaming algorithm $A : \{0, 1\}^n \rightarrow \mathbb{F}_q$ is described by a list of functions $g_i : \{0, 1\} \rightarrow \mathbb{F}_q$ for $i \in [n]$, and computes the value $A(x) = g_1(x_1) + \dots + g_n(x_n)$ (where addition is over \mathbb{F}_q) by tracking the partial sum $g_1(x_1) + \dots + g_i(x_i)$ upon receiving the i 'th bit.

For such algorithms, we describe a noise resilient conversion that incurs quadratic blow-up in communication complexity.

Theorem 1.3. Fix $\varepsilon, \delta > 0$. There is a function $\text{enc} : X \subseteq \{0, 1\}^n \rightarrow \{0, 1\}^m$ where $m = O_\varepsilon(n^{2+\delta})$ and an explicit transformation B such that the following holds: For any linear streaming algorithm A that takes as input $x \in X \subseteq \{0, 1\}^n$ as a stream, runs in space s and time t , and outputs $A(x)$, the algorithm $B_A = B(A)$ takes as input $z \in \{0, 1\}^m$ as a stream, runs in space $s \cdot O_\varepsilon((\log n)^{O(1/\delta)})$ and time $m \cdot O_{\varepsilon, \delta} \left(1 + \frac{t}{n^2}\right)$, and satisfies that whenever $\Delta(z, \text{enc}(x)) < \left(\frac{1}{4} - \varepsilon\right) \cdot m$ then $B_A(z)$ outputs $A(x)$ with probability $\geq 1 - 2^{O_\varepsilon(\log^2 n)}$.

5.1 Statement of Algorithm

Let $\varepsilon > 0$. Throughout the section let $r = \log^{1/\delta} n$ and define $\ell := \log^2 n$. We also assume for simplicity that $\log_r n$ is an integer. We will also assume that $n > \max\{\exp(\exp(8\delta/\varepsilon)), \exp(1/\varepsilon)\}$ is sufficiently large, so that also $\ell < n$.

We begin by specifying the encoding $\text{enc}(x)$ that Alice sends in the stream.

Definition 5.1 ($\text{enc}(x)$). Let $\text{LDC}(x)$ be a locally decodable code with $N(n) = |\text{LDC}(x)| = O_\varepsilon(n^{1+\delta})$ and query complexity $Q = O_\varepsilon(\log^{2/\delta} n)$ satisfying the guarantees of Theorem 4.1 for $\varepsilon/8$. Then Alice sends $\text{enc}(x) := \text{LDC}(x)^{\log^2 n \cdot M_n}$ where $M_n = (r \cdot \ell)^{\log_r n}$ (that is, $\text{enc}(x)$ is $\log^2 n \cdot M_n$ copies of $\text{LDC}(x)$). In particular, $m = |\text{enc}(x)| = (r \cdot \ell)^{\log_r n} \cdot \log^2 n \cdot N(n) = O_\varepsilon(n^{2+4\delta})$.

Throughout this section, we use the notation $A_{i,j}(x)$ to denote the quantity $g(x_{i+1}) + \dots + g_j(x_j)$. Then, $A(x) = A_{0,n}(x)$. We define $N_{j-i} := (r \cdot \ell)^{\log_r(j-i)} \cdot N(n)$ which, as we'll see in the description below, represent the number of bits read by the algorithm to approximate $A_{i,j}(x)$.

We describe Bob's streaming algorithm $B_A := B(A)$. Before stating it formally, we provide a high-level description.

Description: Let z be the incoming stream that Bob receives. At a high level, the algorithm B_A runs by generating many guesses for $A(x)$ along with confidences indicating how likely that guess was correct (depending on how much error is witnessed while generating the guess). At the end, the guess that has the most cumulative confidence is outputted.

For i, j where $j - i$ is an integer power of r , we define the algorithm $\text{estA}(i, j)$ that takes in two input two indices $i, j \in [n]$. It reads N_{j-i} many bits of z after which it outputs a guess for $A_{i,j}(x)$ along with a confidence. In particular, $\text{estA}(0, n)$ outputs a guess for $A(x)$ along with a confidence. By running $\text{estA}(0, n)$ many times and aggregating the guesses weighted by the confidences, we obtain a final output.

To compute $\text{estA}(i, j)$, we break down the interval $(i : j]$ into r subintervals $(i_0 : i_1], (i_1 : i_2], \dots, (i_{r-1} : i_r]$ where $i_a = i + a \cdot \frac{j-i}{r}$. We then recursively formulate guesses for each $A(x(i_{a-1} : i_a))$ over many iterations by calling $\text{estA}(i_{a-1}, i_a)$ each many times. The choice of a in each iteration must be randomized, so that the adversary cannot attack any single choice of a . More specifically, we will split up the N_{j-i} length input stream into ℓ chunks, each of length N_{j-i}/ℓ bits. Each chunk is split into r sections, each of size $N_{j-i}/(\ell \cdot r) = N_{(j-i)/r}$. In each chunk, we pick a random permutation $(\pi_1, \pi_2, \dots, \pi_r)$ of $[r]$. In the a 'th section of the chunk, we will compute on the subproblem corresponding to interval $(i_{\pi_a-1} : i_{\pi_a}]$ by calling $\text{estA}(i_{\pi_a-1}, i_{\pi_a})$.

Whenever a recursive call to $\text{estA}(i_{a-1}, i_a)$ is completed, outputting (\hat{q}, \hat{c}) , then: if $q_a = \hat{q}$ then the confidence c_a is increased by \hat{c} , and if $q_a \neq \hat{q}$ then the confidence c_a is decreased by \hat{c} . If the confidence c_a becomes negative, then we are more sure that the correct state is \hat{q} rather than q_a : then q_a is replaced by \hat{q} and the confidence is negated (so that it's positive).

Algorithm 1 Bob's Noise Resilient Algorithm B_A

```
1: input  $n \in \mathbb{N}$  and stream  $z \in \{0, 1\}^m$ .
2: function  $\text{estA}(i, j)$  ▷compute the state of  $A$  starting at state  $q$  between steps  $i$  and  $j$ 
3:   if  $j = i + 1$  then
4:     Read the next  $|\text{LDC}(x)|$  bits of the stream  $y$ .
5:     Using Theorem 4.1 compute guess  $\hat{b}$  for  $x[j]$  and confidence  $\hat{c}$  in  $O_\varepsilon((\log n)^{O(1/\delta)})$  bits of
     space.
6:     return  $g_i(\hat{b}), c$ .
7:   else
8:     Let  $i_0 = i, i_1 = i + \frac{j-i}{r}, i_2 = i + \frac{2(j-i)}{r}, \dots, i_r = j$ .
9:     Initialize a list of pairs  $(q_1, c_1), \dots, (q_r, c_r)$  each to  $(\emptyset, 0)$ . ▷cumulative best guesses and
     confidences
10:    for  $\ell$  iterations do
11:      Let  $\pi_1 \dots \pi_r$  be a random permutation of  $[r]$ .
12:      for  $a \in [r]$  do
13:        Compute  $(\hat{q}, \hat{c}) \leftarrow \text{estA}(i_{\pi_a-1}, i_{\pi_a})$ .
14:        if  $\hat{q} = q_{\pi_a}$  then ▷update the confidence on  $\text{estA}(i_{\pi_a}, i_{\pi_j+1}, q_{\pi_a-1})$ 
15:           $c_{\pi_a} \leftarrow c_{\pi_a} + \hat{c}$ 
16:        else
17:           $c_{\pi_a} \leftarrow c_{\pi_a} - \hat{c}$ 
18:          if  $c_{\pi_a} < 0$  then ▷if the guess changes, flip its confidence
19:             $q_{\pi_a} \leftarrow \hat{q}$  and  $c_{\pi_a} \leftarrow -c_{\pi_a}$ .
20:      return  $q_1 + \dots + q_r, \min(c_1 \dots c_r)/\ell$ 
21:
22: Initialize a pair  $(q, c)$  to  $(\emptyset, 0)$ .
23: for  $\log^2 n$  iterations do ▷amplification step
24:   Let  $(\hat{q}, \hat{c}) \leftarrow \text{estA}(0, n)$ 
25:   if  $\hat{q} = q$  then
26:      $c \leftarrow c + \hat{c}$ 
27:   else
28:      $c \leftarrow c - \hat{c}$ 
29:     if  $c < 0$  then
30:        $q \leftarrow \hat{q}$  and  $c \leftarrow -c$ .
31: output  $q$ . ▷output  $A(x)$ 
```

5.2 Proof of Theorem 1.3

We begin by proving the stated claims about the communication, time, and space complexities, and then proceed to prove correctness.

5.2.1 Algorithmic Complexity

We've already shown by Definition 5.1 that the length of the encoding $\text{enc}(x)$ is $m = O_\varepsilon(n^{2+4\delta})$. We begin by proving the computational complexity claim.

Lemma 5.2. *If A runs in time t , Algorithm 2 runs in time $m \cdot O_{\varepsilon, \delta} \left(1 + \frac{t}{n^2}\right)$.*

Proof. Notice that each function $g_i(x_i)$ is computed exactly $\frac{\log^2 n \cdot M_n}{n}$ times throughout the protocol. (To see this, one can show inductively that $\frac{1}{j-i}$ fraction of the LDC(x)'s in the computation of $A(i, j)$ are used towards the decoding of each x_k , $k \in (i, j]$.) Other than that, one must decode of each of the $\log^2 n \cdot M_n$ copies of LDC(x) throughout the stream, each of which takes $N(n) + \text{polylog}(n)$ time to read and decode (see Theorem 4.1). This gives a total computation time of $\log^2 n \cdot M(n) \cdot (N(n) + \text{polylog}(r) + \frac{t}{n}) = O_{\varepsilon, \delta}(m \cdot (1 + \frac{t}{n^2}))$. \square

Next, we show that Algorithm 1 satisfies the required space guarantees. We'll first show that the confidences c can be stored in $O_\delta(\log n)$ bits.

Lemma 5.3. *For any i, j, q where $j - i$ is a power of r , let $(\hat{q}, \hat{c}) := \text{estA}(i, j)$. It holds that \hat{c} can be computed and stored as a fraction $(c_1, c_2) = c_1/c_2$ where $c_1 \leq c_2 = (j - i)^{\log_r \ell} \cdot 4Q$.*

Proof. We show this by induction on $j - i$ for $j - i \leq n$. For the base case where $j - i = 1$, this statement holds by Theorem 4.1 since $c_2 = 4Q$.

Assume the statement holds for all q' and $j' - i' < j - i$ where $j' - i'$ is a power of r . In the computation of $\text{estA}(i, j)$, for $a \in r$, each value of c_a is ultimately the sum of $\leq \ell$ values outputted by the function $\text{estA}(i_{a-1}, i_a)$. Since $\frac{i_a - i_{a-1}}{r} = \frac{j-i}{r}$, we can use the induction hypothesis to write each value as a fraction with denominator $\left(\frac{j-i}{r}\right)^{\log_r \ell} \cdot 4Q$. Then, each c_a/ℓ is a fraction with denominator

$$\left(\frac{j-i}{r}\right)^{\log_r \ell} \cdot 4Q \cdot \ell = \ell^{\log_r \frac{j-i}{r}} \cdot 4Q \cdot \ell = \ell^{\log_r(j-i)} \cdot 4Q = (j-i)^{\log_r \ell} \cdot 4Q,$$

and so the output confidence $\min_{a \in [r]} \{c_a/\ell\}$ can be written with denominator $c_2 = (j-i)^{\log_r \ell} \cdot 4Q$. \square

Now, we return to showing that Algorithm 1 satisfies the required space guarantees.

Lemma 5.4. *Algorithm 1 uses $s \cdot O_\varepsilon((\log n)^{O(1/\delta)})$ bits of space.*

Proof. Define $\alpha = \alpha(n) = O_\varepsilon(\log^{O(1/\delta)} n)$ to be the runtime and thus an upper bound on the space required in the local decoding with confidence of LDC(x) as given in Theorem 4.1. We claim the computation of $\text{estA}(0, n)$ takes $\alpha + s + \log_r n \cdot r \cdot (s + O_{\varepsilon, \delta}(\log n))$ bits of storage. We show by induction on $j - i$ (where $j - i$ is a power of r that the computation of $\text{estA}(i, j)$ requires at most $\alpha + s + \log_r(j - i) \cdot r \cdot (s + O_{\varepsilon, \delta}(\log n))$ bits of space.

The base case holds because when $j - i = 1$ because then only $\alpha + s$ space is required.

For the inductive step, the computation of $\text{estA}(i, j)$ requires ℓ computations of $\text{estA}(i_{a-1}, i_a)$ for each $a \in [r]$. By the inductive hypothesis, each such computation takes $\alpha + s + \log_r \left(\frac{j-i}{r}\right) \cdot r \cdot (s + O_{\varepsilon, \delta}(\log n))$ bits of space. Each such computation is done in series, and between iterations only the pairs $\{(q_a, c_a)\}_{a \in [r]}$ are tracked. Each q_a is a state of size s . Each c_a is the sum of $\leq \ell$ confidences that are the output of computation $\text{estA}(i_{a-1}, i_a)$. By Lemma 6.3 the output confidence of $\text{estA}(i_{a-1}, i_a)$ is a fraction ≤ 1 with denominator $\left(\frac{j-i}{r}\right)^{\log_r \ell} \cdot 4Q$, so c_a can be represented in $2 \log_2 \left(\ell \cdot \left(\frac{j-i}{r}\right)^{\log_r \ell} \cdot 4Q \right) < O_{\varepsilon, \delta}(\log_2 n)$ bits. Moreover, in each chunk, we must compute a

permutation of r , which takes at most $r \log_2 r = r \cdot o_\delta(\log n)$ bits. In total, this takes

$$\begin{aligned} &\leq \left[\alpha + s + \log_r \left(\frac{j-i}{r} \right) \cdot r \cdot (s + O_{\varepsilon,\delta}(\log n)) \right] + r \cdot (s + O_{\varepsilon,\delta}(\log n)) \\ &= \alpha + s + \log_r(j-i) \cdot r \cdot (s + O_{\varepsilon,\delta}(\log n)) \end{aligned}$$

bits of space, as desired.

Finally, the amplification part of Algorithm 2 requires $\log^2 n$ iterations. Between iterations, we are only storing a pair (q, c) . This increases the number of bits required by at most $s + O_{\varepsilon,\delta}(\log n)$.

In total, the total space needed throughout the entire algorithm is

$$< s \cdot \log^{1/\delta+1} n + O_{\varepsilon,\delta}(\log^{1/\delta+2} n) + \alpha < s \cdot O_\varepsilon \left((\log n)^{O(1/\delta)} \right)$$

bits. □

5.2.2 Correctness

We now show correctness. Formally, correctness is shown by the following lemma.

Lemma 5.5. *When at most $\frac{1}{4} - \varepsilon$ fraction of the stream $\text{enc}(x)$ is corrupted, Algorithm 1 outputs $A(x)$ with probability $\geq 1 - \exp(-\varepsilon^2 \log^2 n / 32)$.*

We prove the following statement which will easily imply Lemma 5.5. For a given i, j and associated string that is read by the algorithm in the computation, let the random variable $(q(i, j), c(i, j))$ denote the state and confidence after the computation $\text{estA}(i, j)$. We define the *signed confidence*, denoted $c^\pm(i, j)$, to be defined as $+c(i, j)$ if $q(i, j) = A_{i,j}(x)$ and $-c(i, j)$ otherwise. For intuition, the more positive $c^\pm(i, j)$ is, the more correct with higher confidence the output $(q(i, j), c(i, j))$ is, whereas the more negative $c_\pm(i, j)$ is, the more $q(i, j)$ is incorrect with high confidence. So, $c^\pm(i, j)$ gives us a scaled measure of how correct the output of $\text{estA}(i, j)$ is.

Lemma 5.6. *For any $0 \leq i < j \leq n$ and $e \in \mathbb{R}$, given $\leq \frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log(j-i)}{\log n} - e$ fraction of corruptions in the bits read by the computation of $\text{estA}(i, j)$, we have $\mathbb{E}[c^\pm(i, j)] > e$.*

We defer the proof to Section 5.2.3 and return to the proof of Lemma 5.5.

Proof of Lemma 5.5. In the special case where $i = 0, j = n$ and there are $\leq \frac{1}{4} - \frac{\varepsilon}{2} - e < \frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log(j-i)}{r} - e$ errors (where the inequality follows because $\log n < \varepsilon \cdot r$), by Lemma 5.6 we have that $\mathbb{E}[c^\pm(0, n)] > e$.

Then, the final amplification step of the protocol computes the pair $(q(0, n), c(0, n))$ for $\log^2 n$ times, where in each chunk i we denote the fraction of error to be $(\frac{1}{4} - \frac{\varepsilon}{2} - e_i)$, where $\sum_i e_i \geq \frac{\varepsilon \log^2 n}{2}$ since we assumed the total error was $\leq \frac{1}{4} - \varepsilon$. Also, the protocol outputs $A(x)$ if and only if $\sum_i c_i^\pm(0, n)$ (denoting the value of $c^\pm(0, n)$ in the i 'th chunk) is positive. By Azuma's inequality,

$$\Pr \left[\sum_i c_i^\pm(0, n) > 0 \right] \geq \Pr \left[\sum_i (c_i^\pm(0, n) - e_i) > -\frac{\varepsilon \log^2 n}{2} \right] \geq 1 - \exp \left(-\frac{\varepsilon^2 \log^2 n}{32} \right).$$

□

Combining Lemmas 5.2, 5.4, and 5.5, this concludes the proof of Theorem 1.3.

5.2.3 Proof of Lemma 5.6

We prove Lemma 5.6 by induction on $j - i$. We'll show the statement assuming it holds for $j' - i' < j - i$ (where $j' - i'$ is a power of r).

We begin with the base case of $j - i = 1$. In this case, the algorithm simply reads (a corrupted version of) $\text{LDC}(x)$, and by Theorem 4.1, it holds that $\mathbb{E}[c^\pm(i, j)] > e$ since there is at most $(\frac{1}{4} - \frac{\varepsilon}{4} - e)$ error.

We start with some useful notation. In the computation of $\text{estA}(i, j)$, we say that the output of the computation of $\text{estA}(i_{a-1}, i_a)$ in the b 'th chunk is $(\hat{q}_{a,b}, \hat{c}_{a,b})$. We define the *signed confidence* $\hat{c}_{a,b}^\pm$ to be the confidence $\hat{c}_{a,b}$ signed by whether the guess $\hat{q}_{a,b}$ is correct: that is,

$$\hat{c}_{a,b}^\pm := \begin{cases} \hat{c}_{a,b} & \hat{q}_{a,b} = A_{i_{a-1}, i_a}(x) \\ -\hat{c}_{a,b} & \hat{q}_{a,b} \neq A_{i_{a-1}, i_a}(x). \end{cases}$$

The cumulative best guess and confidence for $A_{i_{a-1}, i_a}(x)$ after all ℓ chunks is denoted by (q_a, c_a) . We let $q_0 = \emptyset$. We define a *cumulative signed confidence* c_a^\pm as follows:

$$c_a^\pm := \begin{cases} c_a & q_a = A_{i_{a-1}, i_a}(x) \\ -c_a & q_a \neq A_{i_{a-1}, i_a}(x). \end{cases}$$

We show a lower bound on the cumulative signed confidence at the end of all ℓ chunks c_a^\pm that holds with high probability.

Lemma 5.7. *Let $a \in [r]$ index the r intervals $(i_{a-1}, i_a]$ in the computation of $\text{estA}(i, j)$. Let each of the ℓ chunks in the computation be indexed by $b \in [\ell]$ and suppose chunk b has $(\frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log \frac{j-i}{r}}{\log n} - e_b)$ fraction of corruption. Then, for all $a \in [r]$, with probability $\geq 1 - \exp(-\frac{\log^2 n}{8})$ it holds that $c_a^\pm > \sum_{b \leq \ell} e_b - \frac{\ell}{\log n}$.*

Proof. We first show that for all a , it holds that $c_a^\pm \geq \sum_{b \leq \ell} \hat{c}_{a,b}^\pm$. To see this, define $c_{a,b}^\pm$ to be the value of c_a at the end of chunk b if q_a is correct at the end of chunk b , and otherwise, $-c_a$, so that $c_a^\pm = c_{a,\ell}^\pm$. Then, $c_{a,b}^\pm - c_{a,b-1}^\pm \geq \hat{c}_{a,b}^\pm$, which implies that $c_a^\pm \geq \sum_{b \leq \ell} \hat{c}_{a,b}^\pm$.

Consider the sequence $\{X_b := \sum_{d \leq b} (\hat{c}_{a,d}^\pm - e_d)\}_{b \in [\ell]}$. Because $\mathbb{E}[\hat{c}_{a,b}^\pm] > e_b$ by induction on induction on $j - i$, the sequence $\{X_b\}_{b \in [\ell]}$ is a submartingale. It also holds that $|\hat{c}_{a,b}^\pm - e_b| \leq 2$.

Then, by Azuma's inequality, it holds that

$$\Pr \left[c_a^\pm > \sum_{b \leq \ell} e_b - \frac{\ell}{\log n} \right] \geq \Pr \left[X_\ell > -\frac{\ell}{\log n} \right] > 1 - \exp\left(-\frac{\log^2 n}{8}\right).$$

□

Now, we need a way to reconcile the values of e and e_b . Recall that from the statement of Lemma 5.6 that e is defined to be such that there is $(\frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log j-i}{\log n} - e)$ fraction of error in

the bits read by $\text{estA}(i, j)$. In Lemma 5.7, chunk b had $\left(\frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log \frac{j-i}{r}}{\log n} - e_b\right)$ fraction of error. This means that the total fraction of error is equal to

$$\begin{aligned} \frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log(j-i)}{\log n} - e &= \frac{1}{\ell} \sum_{b \leq \ell} \left(\frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log \frac{j-i}{r}}{\log n} - e_b \right) \\ \implies e &= \frac{1}{\ell} \cdot \sum_{b \leq \ell} e_b - \frac{\varepsilon}{4} \cdot \frac{\log r}{\log n}. \end{aligned}$$

Next, we show that the cumulative signed confidence is larger than proportional to e at the end of each of the ℓ chunks.

Lemma 5.8. *With probability at least $1 - \frac{1}{n}$, it holds that for all $a \in [r]$ that $\frac{c_a^\pm}{\ell} > e + \frac{1}{\log n}$.*

Proof. Using Lemma 5.7 with a union bound, it holds with probability at least

$$1 - \sum_{1 \leq a \leq r} \exp\left(-\frac{\log^2 n}{8}\right) \geq 1 - r \cdot \exp\left(-\frac{\log^2 n}{8}\right) \geq 1 - \frac{1}{n}$$

that

$$\frac{c_a^\pm}{\ell} > \frac{\sum_{b \leq \ell} e_b - \frac{\ell}{\log n}}{\ell} \geq e + \frac{\varepsilon}{4} \cdot \frac{\log r}{\log n} - \frac{1}{\log n} \geq e + \frac{1}{\log n}$$

since for sufficiently large n it holds that $\frac{\varepsilon}{4} \cdot \log r > 2$. \square

We return to the computation of $\mathbb{E}[c^\pm(i, j)]$. Note that

$$c^\pm(i, j) = \begin{cases} + \min_{a \in [r]} \left\{ \frac{c_a}{\ell} \right\} & q(i, j) = A_{i,j}(x) \\ - \min_{a \in [r]} \left\{ \frac{c_a}{\ell} \right\} & q(i, j) \neq A_{i,j}(x). \end{cases}$$

Define the event E to be the event where $\frac{c_a^\pm}{\ell} > e + \frac{1}{\log n}$ for all $a \in [r]$. By Lemma 5.8, E holds with probability at least $1 - \frac{1}{n}$. We split the proof of Lemma 5.6 based on whether $e + \frac{1}{\log n} \geq 0$.

If $e + \frac{1}{\log n} \geq 0$, then whenever E holds we have that for all a we have $c_a^\pm > 0$ and thus $q(i, j) = A_{i,j}(x)$. Therefore,

$$\begin{aligned} \mathbb{E}[c^\pm(i, j)] &\geq \mathbb{E} \left[\min_{a \in [r]} \left\{ \frac{c_a}{\ell} \right\} \middle| E \right] \cdot \Pr[E] + (\text{min value of } c^\pm(i, j)) \cdot (1 - \Pr[E]) \\ &> \left(e + \frac{1}{\log n} \right) \left(1 - \frac{1}{n} \right) + (-1) \cdot \frac{1}{n} \\ &= e + \frac{n - e \log n - 1 - \log n}{n \log n} \\ &> e. \end{aligned}$$

If $e + \frac{1}{\log n} < 0$, then if $q(i, j) = A_{i,j}(x)$, then $\mathbb{E}[c^\pm(i, j)] \geq 0 > e + \frac{1}{\log n} \geq e$. Otherwise, there is at least one value of a for which q_a is incorrect. Conditioned on E , it holds that $\frac{c_a^\pm}{\ell} > e + \frac{1}{\log n}$. Since q_a is incorrect it holds that $c_a^\pm = -c_a$, implying that $\frac{c_a}{\ell} < -e - \frac{1}{\log n}$. As such,

$$\begin{aligned} \mathbb{E}[c^\pm(i, j)] &\geq \mathbb{E} \left[-\min_{a \in [r]} \left\{ \frac{c_a}{\ell} \right\} \middle| E \right] \cdot \Pr[E] + (\text{min value of } c^\pm(i, j)) \cdot (1 - \Pr[E]) \\ &> \left(e + \frac{1}{\log n} \right) \left(1 - \frac{1}{n} \right) + (-1) \cdot \frac{1}{n} \\ &= e + \frac{n - e \log n - 1 - \log n}{n \log n} \\ &> e. \end{aligned}$$

This concludes the proof of Lemma 5.6.

6 Noise Resilient Streaming for General Algorithms

We now consider general streaming algorithms, whose computational may be sequential in nature. Our main result is a noise resilient conversion for deterministic streaming algorithms A . Compared to our scheme for linear algorithms from Section 5, the length of our encoding is larger: $n^{4+\delta}$ compared to $n^{2+\delta}$.

Theorem 1.1. *Fix $\varepsilon, \delta > 0$. There is a function $\text{enc} : X \subseteq \{0, 1\}^n \rightarrow \{0, 1\}^m$ where $m = O_\varepsilon(n^{4+\delta})$ and an explicit transformation B such that the following holds: For any deterministic algorithm A that takes as input $x \in X \subseteq \{0, 1\}^n$ as a stream, runs in space s and time t , and outputs $A(x)$, the algorithm $B_A = B(A)$ takes as input $z \in \{0, 1\}^m$ as a stream, runs in space $s \cdot O_\varepsilon((\log n)^{O(1/\delta)})$ and time $m \cdot O_{\varepsilon, \delta} \left(1 + \frac{t}{n^2} \right)$, and satisfies that whenever $\Delta(z, \text{enc}(x)) < \left(\frac{1}{4} - \varepsilon \right) \cdot m$ then $B_A(z)$ outputs $A(x)$ with probability $\geq 1 - 2^{O_\varepsilon(\log^2 n)}$.*

Throughout this section, we use the notation

$$A(q, \hat{x}) \in \{0, 1\}^s$$

to denote the state of algorithm A when starting with the state $q \in \{0, 1\}^s$ and executing on \hat{x} received in a stream. By definition, there is an explicit algorithm that computes $A(q, \hat{x})$ in s space. We also use the shorthand $A(\hat{x})$ to denote $A(\emptyset, \hat{x})$. Notice that $A(x) = A(\emptyset, x)$ is simply the output state of the stream.

6.1 Statement of Algorithm

Let $\varepsilon > 0$. Throughout the section let $r = \log^{1/\delta} n$ and define $\ell := r^2 \log^4 n$. We also assume for simplicity that $\log_r n$ is an integer. We will also assume that $n > \max \{ \exp(\exp(8\delta/\varepsilon)), \exp(1/\varepsilon) \}$ is sufficiently large, so that also $\ell < n$.

We begin by specifying the encoding $\text{enc}(x)$ that Alice sends in the stream.

Definition 6.1 ($\text{enc}(x)$). *Let $\text{LDC}(x)$ be a locally decodable code with $N(n) = |\text{LDC}(x)| = O_\varepsilon(n^{1+\delta})$ and query complexity $Q = O_\varepsilon(\log^{2/\delta} n)$ satisfying the guarantees of Theorem 4.1 for $\varepsilon/8$. Then Alice sends $\text{enc}(x) := \text{LDC}(x)^{\log^2 n \cdot M_n}$ where $M_n = (r \cdot \ell)^{\log_r n}$ (that is, $\text{enc}(n)$ is $\log^2 n \cdot M_n$ copies of $\text{LDC}(x)$). In particular, $m = |\text{enc}(x)| = (r \cdot \ell)^{\log_r n} \cdot \log^2 n \cdot N(n) = O_\varepsilon(n^{4+6\delta})$.*

Next, we describe Bob's streaming algorithm $B_A := B(A)$. Before stating it formally, we provide a high-level description.

Description: Let z be the incoming stream that Bob receives. At a high level, the algorithm B_A runs by generating many guesses for $A(x)$ along with confidences indicating how likely that guess was correct (depending on how much error is witnessed while generating the guess). At the end, the guess that has the most cumulative confidence is outputted.

For i, j where $j - i$ is an integer power of r , we define the algorithm $\text{estA}(i, j, q)$ that has the following syntax:

- $\text{estA}(i, j, q)$ takes as input two indices $i, j \in [n]$ along with a state $q \in \{0, 1\}^s$ which represents a guess for the state of $A(q, x(i : j))$.
- It reads $N_{j-i} := (r \cdot \ell)^{\log_r(j-i)} \cdot N(n)$ many bits of z after which it outputs a guess for $A(q, x(i : j))$ along with a confidence.

In particular, $\text{estA}(0, n, q)$ outputs a guess for $A(x)$ along with a confidence. By running $\text{estA}(0, n, q)$ many times and aggregating the guesses weighted by the confidences, we obtain a final output.

At a high level, to compute $\text{estA}(i, j, q)$, we break down the interval $(i : j]$ into r subintervals $(i_0 : i_1], (i_1 : i_2], \dots, (i_{r-1} : i_r]$ where $i_a = i + a \cdot \frac{j-i}{r}$. We then recursively formulate guesses for each $A(x(1 : i_a))$ over many iterations by calling $\text{estA}(i_{a-1}, i_a, q_{a-1})$ for some state q_{a-1} . The choice of a in each iteration must be randomized, so that the adversary cannot attack any single choice of a . More specifically, we will split up the N_{j-i} length input stream into ℓ chunks, each of length N_{j-i}/ℓ bits. Each chunk is split into r sections, each of size $N_{j-i}/(\ell \cdot r) = N_{(j-i)/r}$. In each chunk, we pick a random permutation $(\pi_1, \pi_2, \dots, \pi_r)$ of $[r]$. In the a 'th section of the chunk, we will compute on the subproblem corresponding to interval $(i_{\pi_{a-1}} : i_{\pi_a}]$ by calling $\text{estA}(i_{\pi_{a-1}}, i_{\pi_a}, q_{\pi_{a-1}})$.

Note that as the algorithm A is computed sequentially, computing a guess for $A(x[1 : i_a])$ requires having a starting state q_{a-1} for $A(x[1 : i_{a-1}])$. Thus, throughout the entire computation, we will keep track of the best guess and associated confidence for each of the r states $A(x[1 : i_a])$, denoted (q_a, c_a) and all initialized to $(\emptyset, 0)$.

These pairs (q_a, c_a) are maintained as follows. Whenever a recursive call to $\text{estA}(i_{a-1}, i_a, q_{a-1})$ is completed, outputting (\hat{q}, \hat{c}) , then: if $q_a = \hat{q}$ then the confidence c_a is increased by \hat{c} , and if $q_a \neq \hat{q}$ then the confidence c_a is decreased by \hat{c} . If the confidence c_a becomes negative, then we are more sure that the correct state is \hat{q} rather than q_a : then q_a is replaced by \hat{q} and the confidence is negated (so that it's positive). In this last case where q_a is replaced, we have no more reason to believe that further computations of $q_{a'}, a' > a$, are correct since they all depended on q_a , so we erase all such pairs $(q_{a'}, c_{a'})$ and reset them to $(\emptyset, 0)$. A key point in our analysis is to understand why this does not cause the error probability to accumulate.

Algorithm 2 Bob's Noise Resilient Algorithm B_A

```
1: input  $n \in \mathbb{N}$  and stream  $z \in \{0, 1\}^m$ .
2: function  $\text{estA}(i, j, q)$  ▷compute the state of  $A$  starting at state  $q$  between steps  $i$  and  $j$ 
3:   if  $j = i + 1$  then
4:     Read the next  $|\text{LDC}(x)|$  bits of the stream  $y$ .
5:     Using Theorem 4.1 compute guess  $\hat{b}$  for  $x[j]$  and confidence  $\hat{c}$  in  $s \cdot O_\varepsilon((\log n)^{O(1/\delta)})$  bits
     of space.
6:     return  $A(q, \hat{b}), c$ .
7:   else
8:     Let  $i_0 = i, i_1 = i + \frac{j-i}{r}, i_2 = i + \frac{2(j-i)}{r}, \dots, i_r = j$ .
9:     Initialize a list of pairs  $(q_1, c_1), \dots, (q_r, c_r)$  each to  $(\emptyset, 0)$ . ▷cumulative best guesses and
confidences
10:    for  $\ell$  iterations do
11:      Let  $\pi_1 \dots \pi_r$  be a random permutation of  $[r]$ .
12:      Set  $(q'_1, c'_1) = (q_1, c_1), \dots, (q'_r, c'_r) = (q_r, c_r)$ .
13:      for  $a \in [r]$  do
14:        Compute  $(\hat{q}, \hat{c}) \leftarrow \text{estA}(i_{\pi_a-1}, i_{\pi_a}, q'_{\pi_a-1})$  where  $q'_0 := q$ .
15:        if  $\hat{q} = q'_{\pi_a}$  then ▷update the confidence on  $\text{estA}(i_{\pi_a}, i_{\pi_j+1}, q_{\pi_a-1})$ 
16:           $c_{\pi_a} \leftarrow c'_{\pi_a} + \hat{c}$ 
17:        else
18:           $c_{\pi_a} \leftarrow c'_{\pi_a} - \hat{c}$ 
19:          if  $c_{\pi_a} < 0$  then ▷if the guess changes, reset its confidence
20:             $q_{\pi_a} \leftarrow \hat{q}$  and  $c_{\pi_a} \leftarrow -c_{\pi_a}$ .
21:          if some  $q_a$  was changed from its value at the beginning of the chunk then
22:            For all  $i > a$ , set  $(q_i, c_i) \leftarrow (\emptyset, 0)$ . ▷reset state and confidence for states with  $i > a$ 
23:        return  $q_r, \min(c_1 \dots c_r)/\ell$ 
24:
25: Initialize a pair  $(q, c)$  to  $(\emptyset, 0)$ .
26: for  $\log^2 n$  iterations do ▷amplification step
27:   Let  $(\hat{q}, \hat{c}) \leftarrow \text{estA}(0, n, \emptyset)$ 
28:   if  $\hat{q} = q$  then
29:      $c \leftarrow c + \hat{c}$ 
30:   else
31:      $c \leftarrow c - \hat{c}$ 
32:     if  $c < 0$  then
33:        $q \leftarrow \hat{q}$  and  $c \leftarrow -c$ .
34: output  $q$ . ▷output  $A(x)$ 
```

6.2 Proof of Theorem 1.1

The proofs that the algorithm satisfies the required space, computational and communication guarantees are essentially the same as in Section 5, but we reproduce them for completeness. The proof of correctness is somewhat more complicated but also follows the same general outline.

6.2.1 Algorithmic Complexity

We begin by proving the stated claims about communication, time, and space complexities. We've already shown by Definition 6.1 that the length of the encoding $\text{enc}(x)$ is $m = O_\varepsilon(n^{4+6\delta})$. We thus proceed by proving the computational complexity claim.

Lemma 6.2. *If A runs in time t , Algorithm 2 runs in time $m \cdot O_{\varepsilon,\delta} \left(1 + \frac{t}{n^2}\right)$.*

Proof. Notice that each step $A(q_{i-1}, x[i])$ is computed exactly $\frac{\log^2 n \cdot M_n}{n}$ times throughout the protocol. (To see this, one can show inductively that $\frac{1}{j-i}$ fraction of the $\text{LDC}(x)$'s in the computation of $A(i, j, q)$ are used towards the decoding of x_k for each $k \in (i, j]$.) Other than that, one must decode each of the $\log^2 n \cdot M_n$ copies of $\text{LDC}(x)$ throughout the stream, each of which takes $N(n) + \text{polylog}(n)$ time to read and decode (see Theorem 4.1). This gives a total computation time of $\log^2 n \cdot M(n) \cdot (N(n) + \text{polylog}(r) + \frac{t}{n}) = O_{\varepsilon,\delta} \left(m \cdot \left(1 + \frac{t}{n^2}\right)\right)$. \square

Next, we show that Algorithm 2 satisfies the required space guarantees. We'll first show that the confidences c can be stored in $O_\delta(\log n)$ bits.

Lemma 6.3. *For any i, j, q where $j - i$ is a power of r , let $(\hat{q}, \hat{c}) := \text{estA}(i, j, q)$. It holds that \hat{c} can be computed and stored as a fraction $(c_1, c_2) = c_1/c_2$ where $c_1 \leq c_2 = (j - i)^{\log_r \ell} \cdot 4Q$.*

Proof. We show this by induction on $j - i$ for $j - i \leq n$. For the base case where $j - i = 1$, this statement holds by Theorem 4.1 since $c_2 = 4Q$.

Assume the statement holds for all q' and $j' - i' < j - i$ where $j' - i'$ is a power of r . In the computation of $\text{estA}(i, j, q)$, for $a \in r$, each value of c_a is ultimately the sum of $\leq \ell$ values outputted by the function $\text{estA}(i_{a-1}, i_a)$. Since $\frac{i_a - i_{a-1}}{r} = \frac{j-i}{r}$, we can use the induction hypothesis to write each value as a fraction with denominator $\left(\frac{j-i}{r}\right)^{\log_r \ell} \cdot 4Q$. Then, each c_a/ℓ is a fraction with denominator

$$\left(\frac{j-i}{r}\right)^{\log_r \ell} \cdot 4Q \cdot \ell = \ell^{\log_r \frac{j-i}{r}} \cdot 4Q \cdot \ell = \ell^{\log_r(j-i)} \cdot 4Q = (j-i)^{\log_r \ell} \cdot 4Q,$$

and so the output confidence $\min_{a \in [r]} \{c_a/\ell\}$ can be written with denominator $c_2 = (j-i)^{\log_r \ell} \cdot 4Q$. \square

Now, we return to showing that Algorithm 2 satisfies the required space guarantees.

Lemma 6.4. *Algorithm 2 uses $s \cdot O_\varepsilon((\log n)^{O(1/\delta)})$ bits of space.*

Proof. Define $\alpha = \alpha(n) = O_\varepsilon(\log^{O(1/\delta)} n)$ to be the runtime and thus an upper bound on the space required in the local decoding with confidence of $\text{LDC}(x)$ as given in Theorem 4.1. We claim the computation of $\text{estA}(0, n, \emptyset)$ takes $\alpha + s + \log_r n \cdot r \cdot (s + O_{\varepsilon,\delta}(\log n))$ bits of storage. We show by induction on $j - i$ (where $j - i$ is a power of r that the computation of $\text{estA}(i, j, q)$ requires at most $\alpha + s + \log_r(j-i) \cdot r \cdot (s + O_{\varepsilon,\delta}(\log n))$ bits of space.

The base case holds because when $j - i = 1$ because then only $\alpha + s$ space is required.

For the inductive step, the computation of $\text{estA}(i, j, q)$ requires ℓ computations of $\text{estA}(i_{a-1}, i_a, q_{a-1})$ for each $a \in [r]$. By the inductive hypothesis, each such computation takes $\alpha + s + \log_r \left(\frac{j-i}{r}\right) \cdot r \cdot (s + O_{\varepsilon,\delta}(\log n))$ bits of space. Each such computation is done in series, and

between iterations only the pairs $\{(q_a, c_a)\}_{a \in [r]}$ are tracked. Each q_a is a state of size s . Each c_a is the sum of $\leq \ell$ confidences that are the output of computation $\text{estA}(i_{a-1}, i_a, q_{a-1})$. By Lemma 6.3 the output confidence of $\text{estA}(i_{a-1}, i_a, q_{a-1})$ is a fraction ≤ 1 with denominator $\left(\frac{j-i}{r}\right)^{\log_r \ell} \cdot 4Q$, so c_a can be represented in $2 \log_2 \left(\ell \cdot \left(\frac{j-i}{r}\right)^{\log_r \ell} \cdot 4Q \right) < O_{\varepsilon, \delta}(\log_2 n)$ bits. Moreover, in each chunk, we must compute a permutation of r , which takes at most $r \log_2 r = r \cdot o_\delta(\log n)$ bits. In total, this takes

$$\begin{aligned} &\leq \left[\alpha + s + \log_r \left(\frac{j-i}{r} \right) \cdot r \cdot (s + O_{\varepsilon, \delta}(\log n)) \right] + r \cdot (s + O_{\varepsilon, \delta}(\log n)) \\ &= \alpha + s + \log_r (j-i) \cdot r \cdot (s + O_{\varepsilon, \delta}(\log n)) \end{aligned}$$

bits of space, as desired.

Finally, the amplification part of Algorithm 2 requires $\log^2 n$ iterations. Between iterations, we are only storing a pair (q, c) . This increases the number of bits required by at most $s + O_{\varepsilon, \delta}(\log n)$.

In total, the total space needed throughout the entire algorithm is

$$< s \cdot \log^{1/\delta+1} n + O_{\varepsilon, \delta}(\log^{1/\delta+2} n) + \alpha < s \cdot O_\varepsilon \left((\log n)^{O(1/\delta)} \right)$$

bits. □

6.2.2 Correctness

Next, we show correctness. Formally, correctness is shown by the following lemma.

Lemma 6.5. *When at most $\frac{1}{4} - \varepsilon$ fraction of the stream $\text{enc}(x)$ is corrupted, Algorithm 2 outputs $A(x)$ with probability $\geq 1 - \exp(-\varepsilon^2 \log^2 n / 32)$.*

We prove the following statement which will easily imply Lemma 6.5. For a given i, j, q and associated string that is read by the algorithm in the computation, let the random variable $(q(i, j, q), c(i, j, q))$ denote the state and confidence after the computation $\text{estA}(i, j, q)$. We define the *signed confidence*, denoted $c^\pm(i, j, q)$, to be defined as $+c(i, j, q)$ if $q(i, j, q) = A(q, x(i : j))$ and $-c(i, j, q)$ otherwise. For intuition, the more positive $c^\pm(i, j, q)$ is, the more correct with higher confidence the output $(q(i, j, q), c(i, j, q))$ is, whereas the more negative $c_\pm(i, j, q)$ is, the more $q(i, j, q)$ is incorrect with high confidence. So, $c^\pm(i, j, q)$ gives us a scaled measure of how correct the output of $\text{estA}(i, j, q)$ is.

Lemma 6.6. *For any $0 \leq i < j \leq n$, $q \in \{0, 1\}^s$, and $e \in \mathbb{R}$, given $\leq \frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log(j-i)}{\log n} - e$ fraction of corruptions in the bits read by the computation of $\text{estA}(i, j, q)$, we have $\mathbb{E}[c^\pm(i, j, q)] > e$.*

We defer the proof to Section 6.2.3 and return to the proof of Lemma 6.5.

Proof of Lemma 6.5. In the special case where $i = 0, j = n, q = \emptyset$ and there are $\leq \frac{1}{4} - \frac{\varepsilon}{2} - e < \frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log(j-i)}{r} - e$ errors (where the inequality follows because $\log n < \varepsilon \cdot r$), by Lemma 6.6 we have that $\mathbb{E}[c^\pm(0, n, \emptyset)] > e$.

Then, the final amplification step of the protocol computes the pair $(q(0, n, \emptyset), c(0, n, \emptyset))$ for $\log^2 n$ times, where in each chunk i we denote the fraction of error to be $(\frac{1}{4} - \frac{\varepsilon}{2} - e_i)$, where

$\sum_i e_i \geq \frac{\varepsilon \log^2 n}{2}$ since we assumed the total error was $\leq \frac{1}{4} - \varepsilon$. Also, the protocol outputs $A(x)$ if and only if $\sum_i c_i^\pm(0, n, \emptyset)$ (denoting the value of $c^\pm(0, n, \emptyset)$ in the i 'th chunk) is positive. By Azuma's inequality,

$$Pr \left[\sum_i c_i^\pm(0, n, \emptyset) > 0 \right] \geq Pr \left[\sum_i (c_i^\pm(0, n, \emptyset) - e_i) > -\frac{\varepsilon \log^2 n}{2} \right] \geq 1 - \exp \left(-\frac{\varepsilon^2 \log^2 n}{32} \right).$$

□

Combining Lemmas 6.2, 6.4, and 6.5, this concludes the proof of Theorem 1.1.

6.2.3 Proof of Lemma 6.6

We remark that this section constitutes the main difference in the analysis of Algorithm 1 and Algorithm 2. In Algorithm 1, the recursive calls to $\text{estA}(i_{a-1}, i_a)$ can be computed in any order, and so $c^\pm(i, j)$ is easy to analyze and converges quickly. In Algorithm 2, the recursive call to $\text{estA}(i_{a-1}, i_a, q_{a-1})$ has no hope of providing any useful information until q_{a-1} is correct, and so in this algorithm, the quantity $c^\pm(i, j, q)$ requires a more careful analysis, and ultimately converges slower.

We prove Lemma 6.6 by induction on $j - i$. We'll show the statement for a specific value of (i, j, q) where $j - i$ is a power of r assuming it holds for any (i', j', q') where $j' - i' < j - i$ (and $j' - i'$ is a power of r).

We begin with the base case of $j - i = 1$. In this case, the algorithm simply reads (a corrupted version of) $\text{LDC}(x)$, and by Theorem 4.1, it holds that $\mathbb{E}[c^\pm(i, j, q)] > e$ since there is at most $(\frac{1}{4} - \frac{\varepsilon}{4} - e)$ error.

Now, suppose we are in the computation of $\text{estA}(i, j, q)$. We start with some useful notation.

We say that a state q'_a is *correct* when it is equal to the state produced by the algorithm A starting at q and executing steps i to i_a (as defined in Line 8 of the algorithm). That is, it is equal to $A(q, x(i : i_a))$. We note that this is different from requiring only $\text{estA}(i_{a-1}, i_a, q_{a-1})$ to be done correctly from the starting state q_{a-1} .

In the computation of $\text{estA}(i, j, q)$, we say that the output of the computation of $\text{estA}(i_{a-1}, i_a, q_{a-1})$ in the b 'th chunk is $(\hat{q}_{a,b}, \hat{c}_{a,b})$. We define the *signed confidence* $\hat{c}_{a,b}^\pm$ to be the confidence $\hat{c}_{a,b}$ signed by whether the guess $\hat{q}_{a,b}$ is correct: that is,

$$\hat{c}_{a,b}^\pm := \begin{cases} \hat{c}_{a,b} & \hat{q}_{a,b} \text{ is correct} \\ -\hat{c}_{a,b} & \hat{q}_{a,b} \text{ is incorrect.} \end{cases}$$

The cumulative best guess and confidence for $A(q, x(i : i_a))$ after chunk b , as initialized in Line 9, are denoted by $(q_{a,b}, c_{a,b})$. We let $q_{0,b} = \emptyset$. We can define a *cumulative signed confidence* $c_{a,b}^\pm$ as follows:

$$c_{a,b}^\pm := \begin{cases} c_{a,b} & q_{a,b} \text{ is correct} \\ -c_{a,b} & q_{a,b} \text{ is incorrect.} \end{cases}$$

We first show that $c_{a,b}^\pm$ is correct proportional to $\sum_{d \leq b} e_d$ assuming that $q_{a-1,b}$ is correct.

Lemma 6.7. *Let $a \in [r]$ index the r intervals $(i_{a-1}, i_a]$ in the computation of $\text{estA}(i, j, q)$. Let each of the ℓ chunks in the computation be indexed by $b \in [\ell]$ and suppose chunk b has*

$\left(\frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log \frac{j-i}{r}}{\log n} - e_b\right)$ fraction of corruption. Then, for all $a \in [r]$, with probability $\geq 1 - a\ell^2 \exp\left(-\frac{\log^2 n}{32}\right)$ it holds for all $b \in [\ell]$ that either $c_{a,b}^\pm > \sum_{d \leq b} e_d - \frac{a\ell}{r \log n}$ or $q_{a-1,b}$ is incorrect.

Proof. We use induction on a . Let us prove the statement for a assuming it is true for $a-1$. Our base case is $a=0$, and it holds that $q_{0,b}$ is correct (meaning equal to q) for all $b \in [\ell]$.

Fix some chunk b . For $d \leq b$, we define the *altered confidence* $\hat{c}_{a,d}^{\text{alt}}$ to be the signed confidence, with a caveat for when $q_{a-1,d-1}$ is incorrect, as follows:

$$\hat{c}_{a,d}^{\text{alt}} := \begin{cases} 1 & q_{a-1,d-1} \text{ is incorrect} \\ c_{a,b}^\pm & q_{a-1,d-1} \text{ is correct.} \end{cases}$$

Then, for any $0 \leq b' < b$, the sequence $\left\{X_{b',d} := \sum_{b' < d' \leq d} (\hat{c}_{a,d'}^{\text{alt}} - e_{d'})\right\}_{b' < d \leq b}$ is a submartingale.

To see this, we have that $\mathbb{E}[\hat{c}_{a,d}^{\text{alt}}] \geq \mathbb{E}[c_{a,d}^\pm] > e_d$, where the last inequality was by induction on induction on $j-i$. Furthermore, $|c_{a,d}^\pm - e_d| \leq 2$. Then, by Azuma's inequality,

$$\Pr\left[X_{b',b} \leq -\frac{\ell}{2r \log n}\right] \leq \exp\left(-\frac{\ell^2}{32r^2 \log^2 n(b-b')}\right) \leq \exp\left(-\frac{\ell}{32r^2 \log^2 n}\right) = \exp\left(-\frac{\log^2 n}{32}\right).$$

recalling that $\ell = r^2 \log^4 n$. By a union bound,

$$\Pr\left[X_{b',b} > -\frac{\ell}{2r \log n} \quad \forall 0 \leq b' < b\right] \geq 1 - \ell \cdot \exp\left(-\frac{\log^2 n}{32}\right). \quad (1)$$

Now, we return to proving the lemma statement. We have to prove that the statement holds simultaneously for all $b \in [\ell]$ with high probability.

By the inductive hypothesis, Equation (1), and a union bound, we have that with probability $\geq 1 - (a-1)\ell^2 \cdot \exp\left(-\frac{\log^2 n}{32}\right) - \ell \cdot \ell \cdot \exp\left(-\frac{\log^2 n}{32}\right) = 1 - a\ell^2 \cdot \exp\left(-\frac{\log^2 n}{32}\right)$, the following two conditions simultaneously hold:

1. for all d for which $\sum_{d' \leq d} e_{d'} \geq \frac{(a-1)\ell}{r}$ it holds that $c_{a-1,d}^\pm > 0$, or in other words $q_{a-1,d}$ is correct, and
2. for all $b \in [\ell]$, it holds that $X_{b',b} > -\frac{\ell}{2r \log n} \quad \forall 0 \leq b' < b$.

Assume for now that we are not within this failure probability. For any chunk b , let b^* be the last chunk $\leq b$ in which $q_{a-1,d}$ changed values. If q_{a-1,b^*} is incorrect, then the lemma is proven, because $q_{a-1,b} = q_{a-1,b^*}$ is also incorrect. So assume that q_{a-1,b^*} is correct.

If $b^* = b$ then $c_{a,b}^\pm = 0$ since $(q_{a,b}, c_{a,b})$ was reset to $(\emptyset, 0)$. Moreover, because $q_{a-1,b-1}$ has to be incorrect, we have that $\sum_{d < b} e_d < \frac{(a-1)\ell}{r}$ by assumption 1. Then, we have that

$$c_{a,b}^\pm = 0 > \frac{(a-1)\ell}{r \log n} + 1 - \frac{a\ell}{r \log n} \geq \sum_{d < b} e_d + e_b - \frac{a\ell}{r \log n} = \sum_{d \leq b} e_d - \frac{a\ell}{r \log n},$$

as desired, where the first inequality follows because $\ell = r^2 \log^4 n > r \log n$.

Otherwise, $b^* < b$. Then, since q_{a-1,b^*} is correct, the value of q_{a-1,b^*-1} in chunk $b^* - 1$ was incorrect, and so $\sum_{d < b^*} e_d \leq \frac{(a-1)\ell}{r \log n}$ by assumption 1. Since $q_{a-1,d}$ is correct for all $b^* \leq d < b$, we have that $\hat{c}_{a,d}^{\text{alt}} = \hat{c}_{a,d}^{\pm}$ for all $b^* < d \leq b$. Moreover, $c_{a,b}^{\pm} \geq \sum_{b^* < d \leq b} \hat{c}_{a,d}^{\pm}$. To see this, notice that for all b , it holds that $c_{a,b}^{\pm} - c_{a,b-1}^{\pm} \geq \hat{c}_{a,b}^{\pm}$. Then,

$$c_{a,b}^{\pm} \geq \sum_{b^* < d \leq b} \hat{c}_{a,d}^{\pm} = \sum_{b^* < d \leq b} \hat{c}_{a,d}^{\text{alt}} = X_{b^*,b} + \sum_{b^* < d \leq b} e_d.$$

Therefore, by Equation (1),

$$\begin{aligned} c_{a,b}^{\pm} &= X_{b^*,b} + \sum_{b^* < d \leq b} e_d \\ &> -\frac{\ell}{2r \log n} + \sum_{d \leq b} e_d - e_{b^*} - \sum_{d < b^*} e_d \\ &> -\frac{\ell}{2r \log n} + \sum_{d \leq b} e_d - 1 - \frac{(a-1)\ell}{r \log n} \\ &> \sum_{d \leq b} e_d - \frac{a\ell}{r \log n}. \end{aligned}$$

This holds for all $b \in [\ell]$, as desired. \square

We now take $b = \ell$ in Lemma 6.7. First, we need a way to reconcile the values of e and e_b . Recall that from the statement of Lemma 6.6 that e is defined to be such that there is $\left(\frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log(j-i)}{\log n} - e\right)$ fraction of error in the bits read by $\text{estA}(i, j, q)$. In Lemma 6.7, chunk b had $\left(\frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log \frac{j-i}{r}}{\log n} - e_b\right)$ fraction of error. This means that the total fraction of error is equal to

$$\begin{aligned} \frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log(j-i)}{\log n} - e &= \frac{1}{\ell} \sum_{b \leq \ell} \left(\frac{1}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} \cdot \frac{\log \frac{j-i}{r}}{\log n} - e_b \right) \\ \implies e &= \frac{1}{\ell} \cdot \sum_{b \leq \ell} e_b - \frac{\varepsilon}{4} \cdot \frac{\log r}{\log n}. \end{aligned}$$

Lemma 6.8. *With probability at least $1 - \frac{1}{n}$, it holds that for all $a \in [r]$ where $q_{a-1,\ell}$ is correct that $\frac{c_{a,\ell}^{\pm}}{\ell} > e + \frac{1}{\log n}$.*

Proof. By Lemma 6.7, using a union bound, it holds with probability at least

$$\begin{aligned} 1 - \sum_{1 \leq a \leq r} a\ell^2 \exp\left(-\frac{\log^2 n}{32}\right) &\geq 1 - r^2 \ell^2 \exp\left(-\frac{\log^2 n}{32}\right) \\ &\geq 1 - \ell^3 \exp\left(-\frac{\log^2 n}{32}\right) \\ &\geq 1 - \frac{1}{n} \end{aligned}$$

(because $\ell < n$) that for all values of a where $q_{a-1,\ell}$ is correct, we have

$$\begin{aligned} \frac{c_{a,\ell}^\pm}{\ell} &> \frac{\sum_{d \leq \ell} e_d - \frac{a\ell}{r \log n}}{\ell} \\ &\geq e + \frac{\varepsilon}{4} \cdot \frac{\log r}{\log n} - \frac{1}{\log n} \\ &\geq e + \frac{1}{\log n} \end{aligned}$$

since for sufficiently large n , we have that $\frac{\varepsilon}{4} \cdot \log r > 2$. □

We return to the computation of $\mathbb{E}[c^\pm(i, j, q)]$. Note that

$$c^\pm(i, j, q) = \begin{cases} + \min_{a \in [r]} \left\{ \frac{c_{a,\ell}}{\ell} \right\} & q_r \text{ is correct} \\ - \min_{a \in [r]} \left\{ \frac{c_{a,\ell}}{\ell} \right\} & q_r \text{ is incorrect.} \end{cases}$$

We split the proof of Lemma 6.6 based on whether $e + \frac{1}{\log n} \geq 0$.

If $e + \frac{1}{\log n} \geq 0$, define the event E to be the event where $\frac{c_{a,\ell}^\pm}{\ell} > e + \frac{1}{\log n}$ for all $a \in [r]$. We show that E holds with probability at least $1 - \frac{1}{n}$. This is due to Lemma 6.8, which says that with probability $\geq 1 - \frac{1}{n}$, simultaneously for all $a \in [r]$ if $q_{a-1,\ell}$ is correct then $\frac{c_{a,\ell}^\pm}{\ell} > e + \frac{1}{\log n} \geq 0$. Let E be the event that this holds. Then, since $q_{0,\ell} = q$ is correct by definition, it follows that $q_{a,\ell}$ is correct for all $a \in [r]$. Since $q_{a-1,\ell}$ is correct for all $a \in [r]$, it holds that $\frac{c_{a,\ell}^\pm}{\ell} > e + \frac{1}{\log n}$ also. Therefore,

$$\begin{aligned} \mathbb{E}[c^\pm(i, j, q)] &\geq \mathbb{E} \left[\min_{a \in [r]} \left\{ \frac{c_{a,\ell}}{\ell} \right\} \middle| E \right] \cdot \Pr[E] + (\text{min value of } c^\pm(i, j, q)) \cdot (1 - \Pr[E]) \\ &> \left(e + \frac{1}{\log n} \right) \left(1 - \frac{1}{n} \right) + (-1) \cdot \frac{1}{n} \\ &= e + \frac{n - e \log n - 1 - \log n}{n \log n} \\ &> e. \end{aligned}$$

If $e + \frac{1}{\log n} < 0$, then if $q_{r,\ell}$ is correct, then $\mathbb{E}[c^\pm(i, j, q)] \geq 0 > e + \frac{1}{\log n} \geq e$. Otherwise, $q_{r,\ell}$ is incorrect. By Lemma 6.8, with probability $\geq 1 - \frac{1}{n}$, it holds that $\frac{c_{a,\ell}^\pm}{\ell} > e + \frac{1}{\log n}$ for all a as long as $q_{a-1,\ell}$ is correct. Let the event this occurs be E . If $q_{r,\ell}$ is incorrect, there is a smallest $a > 0$ such that $q_{a,\ell}$ is incorrect (note that $q_{0,\ell} = 0$ is always correct), for which it holds that $\frac{c_{a,\ell}^\pm}{\ell} > e + \frac{1}{\log n}$ given E . Since $q_{a,\ell}$ is incorrect it holds that $c_{a,\ell}^\pm = -c_{a,\ell}$, implying that $\frac{c_{a,\ell}}{\ell} < -e - \frac{1}{\log n}$. As such,

$$\begin{aligned} \mathbb{E}[c^\pm(i, j, q)] &\geq \mathbb{E} \left[- \min_{a \in [r]} \left\{ \frac{c_{a,\ell}}{\ell} \right\} \middle| E \right] \cdot \Pr[E] + (\text{min value of } c^\pm(i, j, q)) \cdot (1 - \Pr[E]) \\ &> \left(e + \frac{1}{\log n} \right) \left(1 - \frac{1}{n} \right) + (-1) \cdot \frac{1}{n} \\ &= e + \frac{n - e \log n - 1 - \log n}{n \log n} \\ &> e. \end{aligned}$$

This concludes the proof of Lemma 6.6.

References

- [AGM12] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. *Analyzing Graph Structure via Linear Measurements*, pages 459–467. 2012. [2](#)
- [AMS96] Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 20–29, 1996. [4](#)
- [BE14] Mark Braverman and Klim Efremenko. List and Unique Coding for Interactive Communication in the Presence of Adversarial Noise. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 236–245, Los Alamitos, CA, USA, oct 2014. IEEE Computer Society. [4](#)
- [BK12] Zvika Brakerski and Yael Tauman Kalai. Efficient Interactive Coding against Adversarial Noise. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 160–166, 2012. [4](#)
- [BN13] Zvika Brakerski and Moni Naor. Fast Algorithms for Interactive Coding. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '13*, page 443–456, USA, 2013. Society for Industrial and Applied Mathematics. [4](#)
- [BR11] Mark Braverman and Anup Rao. Towards Coding for Maximum Errors in Interactive Communication. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing, STOC '11*, page 159–166, New York, NY, USA, 2011. Association for Computing Machinery. [4](#)
- [Bra12] Mark Braverman. Towards Deterministic Tree Code Constructions. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCS '12*, page 161–167, New York, NY, USA, 2012. Association for Computing Machinery. [4](#)
- [CCFC02] Moses Charikar, Kevin Chen, and Martin Farach-Colton. Finding frequent items in data streams. In *International Colloquium on Automata, Languages, and Programming*, pages 693–703. Springer, 2002. [4](#)
- [CCFC04] Moses Charikar, Kevin Chen, and Martin Farach-Colton. Finding frequent items in data streams. *Theoretical Computer Science*, 312(1):3–15, 2004. Automata, Languages and Programming. [2](#)
- [CM05] Graham Cormode and S. Muthukrishnan. An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms*, 55(1):58–75, 2005. [2](#)
- [CZ16] Jiecao Chen and Qin Zhang. Bias-aware sketches. *arXiv preprint arXiv:1610.07718*, 2016. [4](#)
- [DEL⁺22] Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes. Locally testable codes with constant rate, distance, and locality. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 357–374, 2022. [4](#)

- [DGY11] Zeev Dvir, Parikshit Gopalan, and Sergey Yekhanin. Matching vector codes. *SIAM Journal on Computing*, 40(4):1154–1178, 2011. 4
- [DHM⁺15] Varsha Dani, Thomas P. Hayes, Mahnush Movahedi, Jared Saia, and Maxwell Young. Interactive Communication with Unknown Noise Rate, 2015. 4
- [Efr09] Klim Efremenko. 3-query locally decodable codes of subexponential length. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 39–44, 2009. 4
- [EGH16] Klim Efremenko, Ran Gelles, and Bernhard Haeupler. Maximal Noise in Interactive Communication Over Erasure Channels and Channels With Feedback. *IEEE Trans. Inf. Theory*, 62(8):4575–4588, 2016. 4
- [EHK⁺23] Klim Efremenko, Bernhard Haeupler, Yael Tauman Kalai, Gillat Kol, Nicolas Resch, and Raghuvansh R Saxena. Interactive coding with small memory. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3587–3613. SIAM, 2023. 4
- [EKS20] Klim Efremenko, Gillat Kol, and Raghuvansh R. Saxena. Binary Interactive Error Resilience Beyond $1/8$ (or why $(1/2)^3 > 1/8$). In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 470–481, 2020. 4
- [Fla85] Philippe Flajolet. Approximate counting: a detailed analysis. *BIT Numerical Mathematics*, 25(1):113–134, 1985. 4
- [For66] G. Forney. Generalized minimum distance decoding. *IEEE Transactions on Information Theory*, 12(2):125–131, 1966. 8
- [Gel17] Ran Gelles. Coding for Interactive Communication: A Survey. *Foundations and Trends® in Theoretical Computer Science*, 13:1–161, 01 2017. 4
- [GH13] Mohsen Ghaffari and Bernhard Haeupler. Optimal Error Rates for Interactive Coding II: Efficiency and List Decoding. *Proceedings - Annual IEEE Symposium on Foundations of Computer Science, FOCS*, 12 2013. 4
- [GH17] Ran Gelles and Bernhard Haeupler. Capacity of Interactive Communication over Erasure Channels and Channels with Feedback. *SIAM Journal on Computing*, 46:1449–1472, 01 2017. 4
- [GHK⁺16] Ran Gelles, Bernhard Haeupler, Gillat Kol, Noga Ron-Zewi, and Avi Wigderson. *Towards Optimal Deterministic Coding for Interactive Communication*, pages 1922–1936. 2016. 4
- [GI18] Ran Gelles and Siddharth Iyer. Interactive coding resilient to an unknown number of erasures. *arXiv preprint arXiv:1811.02527*, 2018. 4
- [GKLR21] Sumegha Garg, Pravesh K Kothari, Pengda Liu, and Ran Raz. Memory-sample lower bounds for learning parity with noise. *arXiv preprint arXiv:2107.02320*, 2021. 4

- [GS92] Peter Gemmell and Madhu Sudan. Highly resilient correctors for polynomials. *Information Processing Letters*, 43(4):169–174, 1992. 8
- [GS00] Venkatesan Guruswami and Madhu Sudan. List decoding algorithms for certain concatenated codes. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, STOC '00, page 181–190, New York, NY, USA, 2000. Association for Computing Machinery. 8
- [Gur06] Venkat Guruswami. Error-correcting codes: Constructions and algorithms, lecture no. 11, 2006. 8
- [GZ22a] Meghal Gupta and Rachel Yun Zhang. Efficient interactive coding achieving optimal error resilience over the binary channel. *arXiv preprint arXiv:2207.01144*, 2022. 4
- [GZ22b] Meghal Gupta and Rachel Yun Zhang. The Optimal Error Resilience of Interactive Communication Over Binary Channels. In *Symposium on Theory of Computing, STOC 2012, New York, NY, USA, June 20 - June 24, 2022*, STOC '22. ACM, 2022. 4
- [Hae14] Bernhard Haeupler. Interactive Channel Capacity Revisited. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 226–235, 2014. 4
- [Ham50] Richard W Hamming. Error detecting and error correcting codes. *The Bell system technical journal*, 29(2):147–160, 1950. 4
- [IW05] Piotr Indyk and David Woodruff. Optimal approximations of the frequency moments of data streams. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 202–208, 2005. 4
- [JL84] William Johnson and Joram Lindenstrauss. Extensions of lipschitz maps into a hilbert space. *Contemporary Mathematics*, 26:189–206, 01 1984. 2
- [KSY14] Swastik Kopparty, Shubhangi Saraf, and Sergey Yekhanin. High-rate codes with sublinear-time decoding. *J. ACM*, 61(5), sep 2014. 4
- [Mor78] Robert Morris. Counting large numbers of events in small registers. *Communications of the ACM*, 21(10):840–842, 1978. 4
- [MRU11] Andrew McGregor, Atri Rudra, and Steve Uurtamo. Polynomial fitting of data streams with applications to codeword testing. *Symposium on Theoretical Aspects of Computer Science (STACS2011)*, 9, 03 2011. 4
- [Mul54] David E Muller. Application of boolean algebra to switching circuit design and to error detection. *Transactions of the IRE professional group on electronic computers*, (3):6–12, 1954. 4
- [MW10] Morteza Monemizadeh and David P Woodruff. 1-pass relative-error lp-sampling with applications. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 1143–1160. SIAM, 2010. 4

- [MWOC21] Chaoyi Ma, Haibo Wang, Olufemi Odegbile, and Shigang Chen. Noise measurement and removal for data streaming algorithms with network applications. In *2021 IFIP Networking Conference (IFIP Networking)*, pages 1–9. IEEE, 2021. 4
- [Nis90] N. Nisan. Pseudorandom generators for space-bounded computations. In *Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing*, STOC '90, page 204–212, New York, NY, USA, 1990. Association for Computing Machinery. 2
- [Ree54] Irving S Reed. A class of multiple-error-correcting codes and the decoding scheme. *IEEE Transactions on Information Theory*, 4(4):38–49, 1954. 4
- [RS60] I. S. Reed and G. Solomon. Polynomial codes over certain finite fields. *Journal of the Society for Industrial and Applied Mathematics*, 8(2):300–304, 1960. 8
- [RU10] Atri Rudra and Steve Uurtamo. Data stream algorithms for codeword testing. In *International Colloquium on Automata, Languages, and Programming*, pages 629–640. Springer, 2010. 4
- [Sch80] J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, oct 1980. 9
- [Sch92] Leonard J. Schulman. Communication on noisy channels: a coding theorem for computation. In *Proceedings., 33rd Annual Symposium on Foundations of Computer Science*, pages 724–733, 1992. 4
- [Sch93] Leonard J. Schulman. Deterministic Coding for Interactive Communication. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '93, page 747–756, New York, NY, USA, 1993. Association for Computing Machinery. 4
- [Sch96] Leonard J. Schulman. Coding for interactive communication. *IEEE Transactions on Information Theory*, 42(6):1745–1756, 1996. 4
- [Sha48] Claude E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, 1948. 4
- [Y⁺12] Sergey Yekhanin et al. Locally decodable codes. *Foundations and Trends[®] in Theoretical Computer Science*, 6(3):139–255, 2012. 4
- [Yek12] Sergey Yekhanin. Locally decodable codes. *Foundations and Trends[®] in Theoretical Computer Science*, 6(3):139–255, 2012. 10
- [Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *Symbolic and Algebraic Computation*, 1979. 9