2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)

Quasilinear time list-decodable codes for space bounded channels

Swastik Kopparty*, Ronen Shaltiel†, Jad Silbak‡
*Department of Mathematics and Department of Computer Science, Rutgers University. swastik.kopparty@gmail.com
†Department of Computer Science, University of Haifa. ronen@cs.haifa.ac.il
‡School of Computer Science, Tel Aviv University. jadsilbak@mail.tau.ac.il

***Abstract*—We consider codes for space bounded channels. This is a model for communication under noise that was studied by Guruswami and Smith (J. ACM 2016) and lies between the Shannon (random) and Hamming (adversarial) models. In this model, a channel is a space bounded procedure that reads the codeword in one pass, and modifies at most a p fraction of the bits of the codeword.**

Guruswami and Smith, and later work by Shaltiel and Silbak (RANDOM 2016), gave constructions of list-decodable codes with rate approaching $1 - H(p)$ against channels with space $s = c \log n$, with encoding/decoding time $\mathrm{poly}(2^s) = \mathrm{poly}(n^c)$.

In this paper we show that for every constant $0 \leq p < \frac{1}{2}$, and every sufficiently small constant $\epsilon > 0$, there are codes with rate $R \geq 1 - H(p) - \epsilon$, list size $\mathrm{poly}(1/\epsilon)$, and furthermore:

- **Our codes can handle channels with space $s = n^{\Omega(1)}$, which is much larger than $O(\log n)$ achieved by previous work.**
- **We give encoding and decoding algorithms that run in time $n \cdot \mathrm{polylog}(n)$. Previous work achieved large and unspecified $\mathrm{poly}(n)$ time (even for space $s = 1 \cdot \log n$ channels).**
- **We can handle space bounded channels that read the codeword in any order, whereas previous work considered channels that read the codeword in the standard order.**

Our construction builds on the machinery of Guruswami and Smith (with some key modifications) replacing some nonconstructive codes and pseudorandom objects (that are found in exponential time by brute force) with efficient explicit constructions. For this purpose we exploit recent results of Haramaty, Lee and Viola (SICOMP 2018) on pseudorandom properties of "t-wise independence + low weight noise" which we quantitatively improve using techniques by Forbes and Kelly (FOCS 2018).

To make use of such distributions, we give new explicit constructions of binary linear codes that have dual distance of $n^{\Omega(1)}$, and are also polynomial time list-decodable from relative distance $\frac{1}{2} - \epsilon$, with list size $\mathrm{poly}(1/\epsilon)$. To the best of our knowledge, no such construction was previously known.

Somewhat surprisingly, we show that Reed-Solomon codes with dimension $k < \sqrt{n}$, have this property if interpreted as binary codes (in some specific interpretation) which we term: "Raw Reed-Solomon Codes". A key idea is viewing Reed-Solomon codes as "bundles" of certain dual-BCH codewords.

***Index Terms*—error-correcting codes; online-space channels; pseudorandomness;**

I. INTRODUCTION

A longstanding open problem in coding theory is to construct binary list-decodable codes that achieve list-decoding capacity, with efficient encoding and list-decoding algorithms. We start with a definition of list-decodable codes. Thinking ahead, the definition of codes below is stated in terms of algorithmic properties of encoding and decoding (rather than combinatorial properties using distance and Hamming balls).

Definition I.1 (Codes). *For $z \in \{0,1\}^n$, let weight(z) denote the Hamming weight of z. Namely, weight$(z) = |\{i \in [n] : z_i \neq 0\}|$. We say that $\mathrm{Enc} : \{0,1\}^k \to \{0,1\}^n$ is an encoding function for a code that is:*

- ***decodable from t errors**, if there exists a function $\mathrm{Dec} : \{0,1\}^n \to \{0,1\}^k$ such that for every $m \in \{0,1\}^k$ and every $e \in \{0,1\}^n$ with weight$(e) \leq t$, $\mathrm{Dec}(\mathrm{Enc}(m) \oplus e) = m$.*
- ***L-list-decodable from t errors**, if the function Dec is allowed to output a list of size at most L, and for every $m \in \{0,1\}^k$ and every $e \in \{0,1\}^n$ with weight$(e) \leq t$, $\mathrm{Dec}(\mathrm{Enc}(m) \oplus e) \ni m$.*

The rate of a code is $R = \frac{k}{n}$.

We will be interested in codes for $t = pn$ errors, where $0 \leq p < \frac{1}{2}$ is a constant, and n is sufficiently large. The "list-decoding capacity" in this setup is $R = 1 - H(p)$, meaning that for every constant $\epsilon > 0$, and sufficiently large n, there exist L-list decodable codes for pn errors, with rate $R \geq 1 - H(p) - \epsilon$, and list size $L = \mathrm{poly}(1/\epsilon)$. Despite substantial effort, it is not known how to construct such codes with poly-time encoding algorithms (even if one does not insist on poly-time list-decoding).

2575-8454/19/$31.00 ©2019 IEEE
DOI 10.1109/FOCS.2019.00028

It is known that codes with rate $R < 1 - H(p)$ must have exponential size lists. The best known uniquely decodable codes have rate $R \leq 1 - H(2p)$, and (unlike the case of list-decoding) the precise capacity of unique decoding is not completely understood.

a) Hamming versus Shannon scenarios: The list-decoding task of Definition I.1 is in the "Hamming scenario" in which the codeword $z = \mathrm{Enc}(m)$ is corrupted by an "unbounded channel" $C(\cdot)$ which given z produces an arbitrary "error pattern" $e = C(z) \in \{0,1\}^n$ with *weight*$(e) \leq pn$, and the decoding algorithm is required to decode (or list-decode) given the "corrupted received word" $z \oplus C(z) = z \oplus e$.

The "Shannon scenario" considers a "restricted channel" C which prepares the "error pattern" $e \in \{0,1\}^n$ *without looking at the codeword* $z = \mathrm{Enc}(m)$. The most well known example is a *binary symmetric channel* (BSC), in which the error pattern $e \in \{0,1\}^n$ is sampled from a distribution which we denote by BSC_p^n, in which the bits $e_1, \ldots, e_n$ are independent, and each e_i is one with probability p. The capacity of such a channel is $R = 1 - H(p)$, and a long line of works give explicit codes matching capacity with efficient (and in fact linear time) encoding and (unique) decoding algorithms [GI05].

Many other "channel distributions" are considered, and in some of them (like "bursts of errors") the individual bits of e are not chosen independently, but rather by a process with "small space".

b) Computationally bounded channels: Note that in Shannon's scenario, channels produce an error pattern that *does not* depend on the codeword $z = \mathrm{Enc}(m)$, whereas in Hamming's scenario there is no restriction, and channels may choose the error pattern as an arbitrary function of the codeword $z = \mathrm{Enc}(m)$. A natural intermediate scenario (considered by Lipton [Lip94]) is to allow the channel $C(z) = e$ to choose the error pattern as a function of z (while insisting that *weight*$(e) \leq pn$), but restrict our attention to channels C from some complexity class.

In this paper (following [GS16, SS16]) we will consider *space bounded channels* which read $z = Enc(m)$ in one pass, using limited space. In this scenario it is helpful to consider *stochastic codes* in which encoding and decoding procedures are randomized.

We now explain why randomization helps. If we insist on the standard notion of *deterministic* encoding algorithms for codes, then the notion of codes for bounded channels coincides with the standard (combinatorial) notion of codes (and so we don't gain by restricting to bounded channels). More specifically, Let $\mathrm{Enc} : \{0,1\}^k \to \{0,1\}^n$ be some function, and let m, e be the "worst message and error pattern" for Enc in terms of "combinatorial list decoding". That is, that for some number L, a Hamming ball of radius pn around $\mathrm{Enc}(m) \oplus e$ contains more than L codewords. This in particular means that the code is not L-list decodable from pn errors. The constant function $C(z) = e$ is a channel from which L-list decoding is impossible. This channel has low complexity in any nonuniform complexity class (as e depends on n, but not on z). In other words, if encoding is deterministic, then the worse combinatorial attack has low complexity. This argument does not apply if encoding algorithms are randomized, and decoding is only guaranteed to succeed with high probability (as is the case in stochastic codes that are defined below).

A. Stochastic codes for space bounded channels

Guruswami and Smith [GS16] considered a notion of *stochastic codes* in which encoding is randomized.[1] In this framework, the encoding algorithm Enc also receives d random bits, and the encoding of a message m, is a random variable $X = \mathrm{Enc}(m, U_d)$ (where U_d denote the uniform distribution on d bits). The channel C receives the "codeword" X as input, and produces an error pattern $e = C(X)$. The decoding algorithm Dec receives the corrupted received word $X \oplus e = X \oplus C(X)$, and needs to decode (or list decode) with high probability (over the choice of the random coins of encoding and decoding algorithms). *We stress that the decoding algorithm **does not** need to receive the random coins of the encoder*. A formal definition follows:

Definition I.2 (Stochastic codes for bounded channels [GS16])**.** *Let $\mathcal{C}$ be a class of functions from n bits to n bits. We say that* $\mathrm{Enc} : \{0,1\}^k \times \{0,1\}^d \to \{0,1\}^n$ *is an encoding function for a stochastic code that is:*

- ***decodable** for "channel class" $\mathcal{C}$, with success probability $1-\nu$, if there exists a (possibly randomized) procedure* $\mathrm{Dec} : \{0,1\}^n \to \{0,1\}^k$ *such that for every $m \in \{0,1\}^k$ and every $C \in \mathcal{C}$, setting* $X = \mathrm{Enc}(m, U_d)$*, we have that* $\Pr[\mathrm{Dec}(X \oplus C(X)) = m] \geq 1-\nu$*, where the probability is over coin tosses of the encoding and decoding procedures.*
- *L-**list-decodable** for "channel class" $\mathcal{C}$, with success probability $1 - \nu$, if the procedure* Dec *is*

---

[1] In the definition below we also allow the decoding algorithm to be randomized. This allows us to speed up decoding algorithms in some cases. Previous work of [GS16, SS16] did not use randomized decoding algorithms, partly because the running time of decoding was inherently large, and there was no gain in using randomness to speed it up.

allowed to output a list of size at most L, and $\Pr[\mathrm{Dec}(X \oplus C(X)) \ni m] \geq 1-\nu$*, where the probability is over coin tosses of the encoding and decoding procedures.*

The rate of a stochastic code is $R = \frac{k}{n}$.

Following [GS16, SS16] we will be interested in the class $\mathcal{C}$ of functions that are computable in one pass using small space. This is captured by the model of oblivious read once branching programs (ROBP). Loosely speaking, a space s ROBP, $C : \{0,1\}^n \to \{0,1\}^n$ is a model of computation that on input $x \in \{0,1\}^n$ performs the following: The ROBP has an internal state q of s bits (initiated to zeros). At step i, the ROBP reads x_i and uses a "transition function" $\delta_i : \{0,1\}^s \times \{0,1\} \to \{0,1\}^s \times \{0,1\}$ to update its "internal state", and output a bit. Overall, on input $x \in \{0,1\}^n$, C produces an output $C(x)$ of n bits. (We sometimes need to consider ROBPs that output a single bit, and in this case the ROBP "decides" on its output bits as a function of its final state). A (more general) precise definition of ROBPs is given in Section III-B.

We will consider channels that induce pn errors and are computable by space s ROBPs that may choose to read the bits of the codeword z in any order. This is captured by allowing the ROBP to permute the order of the bits using some permutation $\sigma : [n] \to [n]$ prior to reading it. We also "un-permute" the output bits, so that the output is presented in the standard order. The definition below is restated more formally in Definition III.2.

Definition I.3 (Space bounded channels, informal)**.** *We say that a function* $C : \{0,1\}^n \to \{0,1\}^n$ *induces* t *errors if for every* $z \in \{0,1\}^n$*, weight*$(C(z)) \leq t$.

The class of ***space*** s ***channels*** *is the class of functions* $C : \{0,1\}^n \to \{0,1\}^n$ *computed by space* s *ROBPs.*

Let $\sigma : [n] \to [n]$ *be a permutation, and for* $z \in \{0,1\}^n$*, let* $\sigma(z)$ *denote the* n *bit string* z'*, in which* $z'_i = z_{\sigma(i)}$*. The class of* ***any-order space*** s ***channels*** *is the class of functions from* n *bits to* n *bits, of the form* $e_C^{\sigma} = \sigma^{-1} \circ C \circ \sigma$ *where* $\sigma : [n] \to [n]$ *is a permutation, and* $C : \{0,1\}^n \to \{0,1\}^n$ *is a space* s *ROBP.*

B. Our Results

1) New constructions of stochastic codes for space bounded channels: Guruswami and Smith [GS16] (and later work by Shaltiel and Silbak [SS16]) gave constructions of stochastic codes for space $s = O(\log n)$ channels with rate approaching $1 - H(p)$. However, a significant drawback of these works is that when set up against channels with space $s = c \log n$ for some constant c, the running time of encoding and decoding in [GS16, SS16] is polynomial in n, for a polynomial that is significantly larger than $2^s = n^c$. This means that one has to pay severely in efficiency, even when considering channels with moderate space.[2]

Guruswami and Smith [GS16] posed the open problem of removing this dependence, and coming up with a code for space $s = c \log n$ channels that has encoding and decoding that run in time n^{c_0} where c_0 is a universal constant, and does not grow with c. In this paper we *solve* this open problem, and in fact, go much farther. Our techniques give explicit constructions of stochastic codes with rate approaching $1-H(p)$, and the following additional improvements:

- Our codes can handle channels with space $s = n^{\Omega(1)}$, which is much larger than $O(\log n)$ achieved by previous work.
- We give encoding and decoding algorithms that run in time $n \cdot \mathrm{polylog}(n)$. Previous work achieved large and unspecified $\mathrm{poly}(n)$ time (even for space $s = 1 \cdot \log n$ channels).
- Our success probability is $1 - \nu$ for $\nu = 2^{-\log^{O(1)} n} = n^{-\omega(1)}$, whereas previous works could only achieve $\nu = n^{-O(1)}$.
- We can handle any-order channels, whereas previous work considered channels that read the codeword in the standard order.

These improvements are summarized in the following theorem (which is our main result). A more general version with more precise description of the dependencies between parameters is stated in Theorem VI.1.

Theorem I.4 (quasilinear time codes for space $n^{\Omega(1)}$ channels)**.** *For every constant* $0 \leq p < \frac{1}{2}$ *and sufficiently small constant* $\epsilon > 0$*, there is an infinite family of stochastic codes with rate* $R \geq 1 - H(p) - \epsilon$*, that are* $(L = \mathrm{poly}(1/\epsilon))$*-list decodable for any-order space* $s = n^{\Omega(1)}$ *channels that induce* pn *errors, with success*

---

[2] We remark that the construction of Guruswami and Smith [GS16] is a "Monte-Carlo construction", meaning that it requires a preprocessing stage, in which a random string of length $\mathrm{poly}(n^c)$ is shared between the encoding and decoding algorithm. The correctness of encoding and decoding algorithms is guaranteed w.h.p. over the choice of this string. (This string need not be kept secret from the channel, but note that this string is longer than the "description length" of the channel). In the final version of [GS16] it is observed that this Monte-Carlo approach can be extended to any class of channels where all channels have description length $\mathrm{poly}(n^c)$, like for example size n^c circuits. Shaltiel and Silbak removed the need for a "Monte-Carlo" construction, and gave a construction that does not require this preprocessing step. However, their construction still suffers from running time of encoding and decoding that is exponential in $s = c \log n$. More details are given in Section I-C.

probability $1 - 2^{-\log^2 n}$. *Furthermore, encoding and decoding run in time* $n \cdot$ *polylog*(n).

Our approach builds on the approach of [GS16] (and refinements of [SS16]) with some modifications and simplifications. The key to our improvements is a better explicit construction of some component that we call "control code" (for which [GS16, SS16] gave constructions based on showing existence by a non-constructive argument, and then finding the object by brute force search). We replace these inefficient arguments by an explicit construction. We give a detailed high level overview of the proof of Theorem I.4 in Section II.

We can also handle channels with space $s = n/\text{polylog}(n)$, but we only know how to do this for small values of p, and then, encoding and decoding run in polynomial time (rather than quasilinear time). A more general version of the next theorem is stated in Theorem VI.3.

Theorem I.5 (polynomial time codes for space $n/\text{polylog}(n)$ channels)**.** *There exist constants* $p_0 > 0$ *and* $c_0 \geq 1$ *such that for every constant* $0 \leq p < p_0$ *and sufficiently small constant* $\epsilon > 0$, *there is an infinite family of stochastic codes with rate* $R \geq 1 - H(p) - \epsilon$ *that are* $(L = \text{poly}(1/\epsilon))$*-list decodable for any-order space* $s = \frac{n}{(\log n)^{c_0}}$ *channels that induce* pn *errors, with success probability* $1 - 2^{-\log^2 n}$. *Furthermore, encoding and decoding run in time* $\text{poly}(n)$.

a) Perspective: Our results clearly extend to any channel that is a convex combination of any-order space s channels. Furthermore, with an additional $\log n$ space, a channel can count the number of error that it induces, and avoid inducing more than pn errors. This means that our theorems handle any distribution over any-order space s channels in which the probability of inducing significantly more than pn errors is small.

It was pointed out by Guruswami and Smith [GS16] that all the "stochastic channels" studied in Shannon's scenario are captured by this framework. Consequently, Theorem I.4 can be seen as providing a *unified* solution that handles all such channels with rate approaching $1 - H(p)$ and quasilinear time encoding and decoding.

On a more philosophical level, one may postulate that the behavior of most conceivable channels that are not "fully adversarial" is captured by this framework of Guruswami and Smith, which can now be implemented in quasilinear time (without the severe penalty of the dependence of running time on the space of the channel).

2) Raw Reed-Solomon Codes: One of the tools that we require in order to prove our main theorem, is a *binary linear code* $\text{Enc} : \mathbb{F}_2^k \to \mathbb{F}_2^n$ with the following properties.[3]

- Distance $(1/2 - o(1)) \cdot n$.
- Large dual distance of at least $n^{\Omega(1)}$. (In fact, we need a slightly stronger property to be explained later).
- Polynomial time list-decoding with list size $\text{poly}(1/\epsilon)$ from $(\frac{1}{2} - \epsilon) \cdot n$ errors, for every sufficiently small constant $\epsilon > 0$. (This in particular implies poly-time unique decoding up $(\frac{1}{4} - \epsilon) \cdot n$ errors, by pruning the list and keeping only the unique codeword that is closest to the received word).

To the best of our knowledge, no construction with these properties is known. In this paper we exhibit such codes. Interestingly, we show that for some settings of parameters, Reed-Solomon codes have these properties if interpreted as binary codes suitably. We call these resulting binary codes *Raw Reed-Solomon codes*.

More specifically, let m be an integer, and $n_{\text{RS}} = 2^m - 1$. We consider the field $\mathbb{F}_{2^m}$, and Reed-Solomon codes of degree $\leq d$ with n_{RS} evaluation points given by $D = \mathbb{F}_{2^m} \setminus \{0\}$. That is, the encoding of a message $w \in \mathbb{F}_{2^m}^{d+1}$, is $\text{Enc}(w) = (\sum_{0 \leq i \leq d} w_i \cdot t^i)_{(t \in S)}$. This code has distance $n_{\text{RS}} - d$, and alphabet size $q_{\text{RS}} = 2^m$. It can be interpreted as a binary linear code by choosing some $\mathbb{F}_2$-linear bijection $\Phi : \mathbb{F}_{2^m} \to \mathbb{F}_2^m$ (which is used to interpret field elements as m bit vectors), and applying Φ on each of the n_{RS} symbols of the codeword. This gives a binary linear code $\mathsf{VerySimpleRawRS}$ with dimension $k = (d+1) \cdot m$ and block length $n = n_{\text{RS}} \cdot m$. It immediately follows that this standard construction has distance at least $n_{\text{RS}} - d$, but note that in terms of *relative distance*, this quite general argument does quite poorly, since it can never show that the relative distance is more than:

$$\frac{n_{\text{RS}}}{n} = \frac{1}{m} = \Theta\left(\frac{1}{\log n}\right) = o(1).$$

In fact, this is the truth, and $\mathsf{VerySimpleRawRS}$ truly does have $o(1)$ relative distance.

Nevertheless, we show that a slight modification of this code has extremely good distance (and keeps the dual distance). Let $\mathsf{SimpleRawRS}$ be the subcode of $\mathsf{VerySimpleRawRS}$ which only includes codewords that

---

[3] In this section we use a more standard notation of coding theory. With our notation, a binary linear code is a code that has a linear encoding function $\text{Enc} : \mathbb{F}_2^k \to \mathbb{F}_2^n$. The image of this function is a subspace C of the vector space $\mathbb{F}_2^n$. The **dual code** is the dual subspace $C^{\perp} = \left\{v \in \mathbb{F}_2^n : \forall c \in C, < v, c >= 0\right\}$. The **dual distance** is the minimum distance of $C^{\perp}$. The precise standard definitions are given in Section III-E2.

come from polynomials which have 0 constant term. Using deep algebraic tools (very specific to the algebraic situation at hand) we show that if the degree bound $d < n_{\mathrm{RS}}^{o(1)}$ (so that the dimension k satisfies $k = n^{o(1)}$), then $\mathsf{SimpleRawRS}$ has relative distance $\frac{1}{2} - o(1)$. We also define another variant, $\mathsf{OddRawRS}$, which has the same relative distance but which can achieve any dimension that is $o(n^{1/2})$. Finally, using the powerful algorithmic decoding algorithms known for Reed-Solomon codes, we show that these codes are also list-decodable.

A more detailed description of Raw Reed Solomon codes appears in Section IV. In particular we prove the following theorem.

Theorem I.6 (Codes with large distance and dual distance)**.** *For every constant $0 < \alpha < 1/2$, and every sufficiently large m, setting $n = (2^m - 1) \cdot m$, and $k = n^\alpha$, there is a binary linear $[n, k]_2$-code C that satisfies:*

- *C has distance $(\frac{1}{2} - O((\frac{\log n}{n})^{\frac{1}{2}-\alpha}))n = (\frac{1}{2} - o(1)) \cdot n$.*
- *C has dual distance $\Omega(\frac{n^\alpha}{\log n})$.*
- *C has a linear encoding map $\mathrm{Enc} : \mathbb{F}_2^k \to \mathbb{F}_2^n$ that runs in time $\mathrm{poly}(n)$.*
- *There exists a universal constant b, such that for every $\epsilon \geq b\sqrt{\alpha}$, Enc is $O(\frac{1}{\epsilon^2})$-list-decodable from $(\frac{1}{2} - \epsilon)n$ errors in time $\mathrm{poly}(n)$.*

A key property that we use in the analysis is that (modulo some caveats) a codeword of length $n_{\mathrm{RS}} \cdot m$ of the $\mathsf{SimpleRawRS}$ and $\mathsf{OddRawRS}$ codes can be viewed as the juxtaposition of m binary strings of length n_{RS} in a natural way. These m binary strings turn out to be (correlated) codewords of a dual-BCH code, and our analysis of the distance exploits this. This is explained more precisely is Section II.

Dual-BCH codes themselves satisfy the first three requirements above, but they are not known to have efficient decoding. Curiously, our result shows that a code of correlated tuples of dual-BCH codewords *can* be decoded efficiently, while retaining the other good properties of dual-BCH codes.

Reed Solomon codes have a lot of structure and many useful properties (in addition to their distance properties) and so, we believe that the fact that Raw Reed-Solomon codes have the additional properties listed above (when viewed as binary codes) is of independent interest, and may prove useful in other applications.

C. More related work on codes for bounded channels

1) Stochastic codes for other classes of channels:

a) Additive channels: Guruswami and Smith [GS16] gave constructions of stochastic codes with rate approaching $1 - H(p)$ that are uniquely decodable for additive channels that induce pn errors, with success probability $1 - 2^{-\Omega(n/\log n)}$. In our notation these are the constant functions $C(\cdot) = e$ where e is a constant string with weight at most pn. The encoding and decoding algorithms in [GS16] run in polynomial time. Our approach can be used in this setup, and can speed up the encoding and decoding algorithms to run in quasilinear time, if the success probability is reduced to $1 - 2^{-\mathrm{polylog}(n)}$.

b) poly-size circuits and bounded space channels: In the same paper, Guruswami and Smith also gave constructions of stochastic codes with rate approaching $1 - H(p)$ that are list-decodable for space $s = c \log n$ channels (or size n^c circuits) that induce pn errors, with success probability n^{-c}. As explained earlier in Section I-B1, a significant drawback of these results is that the running time of the encoding algorithm was polynomial in n^c, for a large and unspecified polynomial (meaning that efficiency quickly deteriorates even for conservative estimates on channel complexity). The construction of [GS16] is "Monte-Carlo". Meaning that it requires a preprocessing stage, in which a random string of length $\mathrm{poly}(n^c)$ is shared between the encoding and decoding algorithm. The correctness of encoding and decoding algorithms is guaranteed w.h.p. over the choice of this string. (This string need not be kept secret from the channel).

Shaltiel and Silbak [SS16] removed the need for a preprocessing stage by slightly modifying the construction of Guruswami and Smith, and providing explicit constructions for the modified components. They give results for space $s = c \log n$ channels, and size n^c circuits (here a complexity assumption that there are functions in *DTIME*$(2^{O(n)})$ that are hard for small circuits is used, and is necessary). Shaltiel and Silbak also consider channels that are implementable by constant depth circuits, and provide constructions of stochastic codes for this setup.

2) Other coding scenarios with randomized encoding/bounded channels:

The notion of computationally bounded channels was also studied in other setups. We mention some of these works below.

a) Shared private randomness: We start with the notion of codes with "shared private randomness". While this setup was considered before the notion of stochastic codes, in this paper, it is natural to view it as a version of stochastic codes in which the decoding algorithm *does* receive the randomness S chosen by the encoding algorithm. This corresponds to a standard symmetric cryptography setup in which honest parties (the encoder and decoder) share a uniform private key S, and the bad party (the channel) does not get the key. Lipton [Lip94] and following work (see [Smi07] for more details) gave explicit constructions of uniquely decodable codes against computationally bounded channels, in this setup, with rate approaching $1 - H(p)$, under cryptographic assumptions.

Note that the setup of stochastic codes is lighter. The encoder and decoder do not need to share a private random key. Moreover, a fresh key can be chosen on the spot every time the encoder encodes a message.

A related notion of "private codes" was studied by Langberg [Lan04]. This is also in the setup of shared private randomness. Here channels are computationally unbounded, codes are existential rather than explicit, and have rate approaching $1 - H(p)$. The focus is on minimizing the length of the shared key. Langberg provides asymptotically matching upper and lower bounds of $\Theta(\log n + \log(1/\nu))$, on the amount of randomness that needs to be shared for unique decoding in this setup, where ν is the error parameter.

b) Non malleable codes: Non-malleable codes (introduce by Dziembowski, Pietrzak, and Wichs [DPW18]) consider channels that are not restricted in the number of errors that they induce. Instead, channels are assumed to come from some limited class of functions (or complexity class). Codes are stochastic (meaning that the encoding procedure is randomized) and it is required that following the corruption by the channel, the decoder either reproduces the encoded message, or an "unrelated" message. The definition of "unrelated" is given using the simulation paradigm from cryptography. Several classes have been considered, and some of the constructions rely on cryptographic assumptions. The reader is referred to [DPW18] and the references therein for precise definition and a survey of results in non-mallable codes.

c) Encoding a uniform message: Haviv and Langberg [HL11] consider a model where encoding and decoding is deterministic, but the message to be encoded is chosen uniformly at random. They show the existence of codes in this setup, that have unique decoding, and beat the Gilbert-Varshamov bound.

d) Public key setup: Micali et al. [MPSW10] considered computationally bounded channels, and a cryptographic public key setup. Their focus is to use this setup to convert a given (standard) explicit list-decodable code into an explicit uniquely decodable codes (in this specific public key setup).

D. Organization of the paper

In Section II we give a high level overview of the ideas and techniques in this paper. In Section III we give definitions and past work on the tools and ingredients that are used in our construction. In Section IV we state and prove our results on raw Reed-Solomon codes. In Section V we use raw Reed-Solomon codes to construct stochastic control codes. In Section VI we give our main construction of stochastic codes for space bounded channels (which relies on stochastic control codes). The reader is referred to the full version for the proof of correctness of the main construction.

II. OVERVIEW OF THE TECHNIQUE

In this section we give a high level overview of the ideas and techniques that we use. We allow ourselves to be informal and imprecise (in order to highlight the main ideas). Complete definitions, theorem statements and proofs, appear in later sections (which do not rely on the informal description given in this section).

A. Stochastic control codes

The construction of codes for bounded channels of Guruswami and Smith [GS16], as well as later modification by Shaltiel and Silbak [SS16] use a component which we call a "stochastic control code".

Definition II.1 (Stochastic control code, informal)**.** *A function* $\mathrm{Enc}_{\mathrm{ctrl}} : \{0,1\}^k \times \{0,1\}^d \to \{0,1\}^n$ *is a stochastic control code that is:*

- *pseudorandom, if for every* $x \in \{0,1\}^k$, $\mathrm{Enc}_{\mathrm{ctrl}}(x, U_d)$ *is pseudorandom for small space ROBPs.*
- *List decodable, if for every sufficiently small constant* $\epsilon > 0$, *there is an explicit list-decoding algorithm* $\mathrm{Dec}_{\mathrm{ctrl}}$ *with constant size lists, such that for every* $x \in \{0,1\}^k$, $y \in \{0,1\}^d$, *and* $e \in \{0,1\}^n$ *with weight*$(e) \leq (\frac{1}{2} - \epsilon) \cdot n$, *list-decoding succeeds, that is,* $x \in \mathrm{Dec}_{\mathrm{ctrl}}(\mathrm{Enc}_{\mathrm{ctrl}}(x, y) \oplus e)$.

A more precise definition (with precise quantities) appears in Section V. The construction of capacity achieving stochastic codes for space bounded channels (Theorem I.4) will rely on stochastic control codes that are pseudorandom and list-decodable. Definition II.1 requires recovery from *adversarial* errors (that may be induced by an *unbounded* channel) while also requiring the additional pseudorandomness property. This makes the requirements *stronger* than the codes for bounded space channels that we aim to construct.

A key idea is that in the final construction, the control code will be used to encode a short "control string", and so, the rate of this code does not need to approach $1 - H(p)$, and we will be able to use control codes in which $k = n^{\Omega(1)}$ in our final stochastic codes for bounded channels, and still have rate approaching $1 - H(p)$.

Theorem II.2 (Control code, informal)**.** *There is an explicit control code* $\mathrm{Enc}_{\mathrm{ctrl}} : \{0,1\}^{n^{\Omega(1)}} \times \{0,1\}^{O(n)} \to \{0,1\}^n$ *which satisfies Definition II.1.*

We start by explaining this construction, and later show how to use it to obtain stochastic codes for small space channels with rate approaching $1 - H(p)$.

1) Raw Reed-Solomon codes: The first step in our construction of control codes, are explicit (standard) binary linear codes with large dual distance and poly-time list decoding. These are the codes stated in Theorem I.6. For concreteness, let m be an integer, and consider the field $\mathbb{F}_{2^m}$. Let $D = \mathbb{F}_{2^m} \setminus \{0\}$. We will consider evaluations of polynomials of degree at most d with $\mathbb{F}_{2^m}$ coefficients at the points of D, and then convert these evaluations to binary vectors of length m using an $\mathbb{F}_2$-linear bijection $\Phi : \mathbb{F}_{2^m} \to \mathbb{F}_2^m$. More generally, this conversion to binary vectors can also be done using a different $\mathbb{F}_2$-linear bijection Φ_x at each point $x \in D$. Overall, this gives us binary codewords of length $n = m \cdot (2^m - 1)$. We call the codes obtained this way *Raw Reed-Solomon Codes*, RawRS.

In this high level overview we will explain the analysis of a particular Raw Reed-Solomon code which we call *Odd Raw Reed-Solomon codes*[4], OddRawRS. This is an instance of the above RawRS family of codes (for a particular choice of Φ_x), but it has a more direct description which we give next. Let $k = m \cdot (d+1)$ and $n = m \cdot (2^m - 1)$. Given a message $w \in \mathbb{F}_2^k$, we break it into $d+1$ blocks of m bits each, use these m-bit blocks to specify elements $\gamma_0, \gamma_1, \ldots, \gamma_d \in \mathbb{F}_{2^m}$, and consider the polynomial $P(X) = \sum_{j=0}^{d} \gamma_j X^{2j+1}$ which has *only odd degree monomials*. The codeword $c : D \times [m] \to \mathbb{F}_2$ corresponding to this message w is then the n bit long string obtained by taking all the evaluations of P and writing them in bits using Φ:

$$c(x, i) = \Phi(P(x))_i.$$

A key observation is that this interpretation is closely related to the dual-BCH code [MS77]. Dual-BCH codes are known to satisfy the first three properties in Theorem IV.1. That is they have poly-time encoding, large distance, and large dual distance. However, they are not known to have poly-time decoding or list-decoding.

We show that OddRawRS codes satisfy all four requirements. First we elaborate on the connection to dual-BCH codes. For each $i \in [m]$, consider the function $c_i : D \to \mathbb{F}_2$ given by $c_i(x) = c(x, i)$ (this is just a subset of the bits of c). Then each c_i is a codeword of the dual-BCH code. More specifically, the dual BCH codeword $c^{\mathsf{dual-BCH}} : D \to \mathbb{F}_2$ that corresponds to $w = (\gamma_0, \ldots, \gamma_d)$, can be defined as $c^{\mathsf{dual-BCH}}(x) = \mathrm{Tr}(\sum_{j=0}^{d} \gamma_j \cdot x^{2j+1})$, where $\mathrm{Tr} : \mathbb{F}_{2^m} \to \mathbb{F}_2$ is the $\mathbb{F}_2$-linear map that is the *field trace*. Furthermore, any bijective linear map $\Phi : \mathbb{F}_{2^m} \to \mathbb{F}_2^m$ can be expressed as m $\mathbb{F}_2$-linear maps $\Phi_i : \mathbb{F}_{2^m} \to \mathbb{F}_2$, where each Φ_i is defined by $\Phi_i(x) = \mathrm{Tr}(b \cdot x)$, for some nonzero $b \in \mathbb{F}_{2^m}$. It follows that $c_i(x) = \mathrm{Tr}(\sum_{j=0}^{d} (b \cdot \gamma_j) \cdot x^{2j+1})$, which can be viewed as the dual-BCH encoding of the nonzero word $(b \cdot \gamma_0, \ldots, b \cdot \gamma_d)$.

Thus the codewords of OddRawRS are just a sequence of correlated nonzero dual-BCH codewords. Using this connection to dual-BCH codes we get that OddRawRS codes have large distance.[5] This part of the argument does not work for general RawRS codes.

The remaining three properties hold for general RawRS codes. Efficient encoding is clear. The dual distance of OddRawRS codes follows from the dual distance of a related Reed-Solomon code, and the fact that Φ is a bijection.

Finally we come to the decodability. This is where we go beyond what is known for dual-BCH codes. The crucial point here is the connection to Reed-Solomon codes, for which amazing decoding algorithms are known [Sud97, GS99]. We show that the natural 2-stage list-decoding algorithm for OddRawRS (which is naturally viewed as a concatenated code) indeed decodes from $(1/2 - \epsilon)$-fraction errors. The first stage is list-decoding of the inner blocks, which leads to a huge list of candidate symbols for each coordinate (since the inner blocks are all codes with minimum distance only 1). Then the efficient list-recovery algorithms known for Reed-Solomon codes, which can handle huge lists and $(1 - o(1))$-fraction error, enables us find a large list that contains all nearby codewords. This implies the unique decodability from $(1/4 - \epsilon)$-fraction errors. Finally, for list-decodability, using the fact that OddRawRS has relative distance $1/2 - o(1)$, the Johnson bound [Joh62] implies that the list of $(1/2 - \epsilon)$-fraction close codewords is in fact $\mathrm{poly}(1/\epsilon)$, and we get the desired list-decoding algorithm.

---

[4] A similar but more involved analysis applies to the more natural code SimpleRawRS.

[5] In the OddRawRS case it is almost immediate, in the SimpleRawRS case we need to understand the correlations between the component dual-BCH codewords.

2) From Raw Reed-Solomon codes to stochastic control codes.: We now explain how to prove Theorem II.2 using $\mathsf{OddRawRS}$ codes (and specifically, the code stated in Theorem I.6). This approach is inspired by a related argument that was used by Shaltiel and Silbak [SS16] to construct control codes against AC^0 circuits.

We will construct the control code $\mathrm{Enc}_{\mathrm{ctrl}} : \{0,1\}^{k/2} \times \{0,1\}^{d=k/2+n\cdot\log(1/\eta)} \to \{0,1\}^n$ as follows: Given $x \in \{0,1\}^{k/2}$, $r \in \{0,1\}^{k/2}$ and $v \in \{0,1\}^{n\cdot\log(1/\eta)}$. We use v as random coins to sample an element from BSC_η^n (that is n i.i.d. coins that evaluate to one with probability η) and define:

$$\mathrm{Enc}_{\mathrm{ctrl}}(x;(r,v)) = \mathrm{Enc}_{\mathsf{OddRawRS}}(r \circ x) \oplus \mathrm{BSC}_\eta^n.$$

That is, we use the linear code from the previous section to encode $r \circ x$ and xor the output with BSC_η^n. By Theorem I.6 the $\mathsf{OddRawRS}$ code has dual distance $t = n^{\Omega(1)}$, and this can be used to show that for every $x \in \{0,1\}^{k/2}$, $\mathrm{Enc}_{\mathsf{OddRawRS}}(U_{k/2} \circ x)$ is $(t-1)$-wise independent.

Actually, for this to hold we need a stronger property, namely that if we consider the code $\mathrm{Enc}^{\mathrm{trunc}} : \mathbb{F}_2^{k/2} \to \mathbb{F}_2^n$ defined by $\mathrm{Enc}^{\mathrm{trunc}}(r) = \mathrm{Enc}_{\mathsf{OddRawRS}}(r \circ 0^{k/2})$, then this linear code has dual distance t. (This stronger property also holds for $\mathsf{OddRawRS}$ codes). Once we have that, by linearity, $\mathrm{Enc}_{\mathsf{OddRawRS}}(U_{k/2} \circ x) = \mathrm{Enc}^{\mathrm{trunc}}(U_{k/2}) \oplus L(x)$ where L is some linear function. Thus, it is enough to show that $\mathrm{Enc}^{\mathrm{trunc}}(U_{k/2})$ is $(t-1)$-wise independent. Note that encoding by $\mathrm{Enc}^{\mathrm{trunc}}$ is done by multiplying the message by the generator matrix of $\mathrm{Enc}^{\mathrm{trunc}}$ (which is the parity check matrix of the dual code). As the dual distance of $\mathrm{Enc}^{\mathrm{trunc}}$ is at least t, every $t-1$ columns of the latter matrix are linearly independent, and so $\mathrm{Enc}^{\mathrm{trunc}}(U_{k/2})$ is $(t-1)$-wise independent. It follows that the output distribution $\mathrm{Enc}_{\mathrm{ctrl}}(x, U_d)$ is a "t-wise independent distribution plus low weight BSC noise".

3) t-wise independence + low weight BSC noise.: Haramaty, Lee and Viola [HLV18] studied the pseudorandomness of such distributions, and showed that distributions of this form are pseudorandom for any-order small space ROBPs if t is sufficiently larger than $n^{2/3}$, and η is not too small (any constant $\eta > 0$ will do). Note that the list decoding property of Theorem II.2 immediately follows for our construction (as $\mathsf{OddRawRS}$ codes have list decoding up to $(\frac{1}{2}-\epsilon)\cdot n$ errors, and so, if η is sufficiently small, the additional relative error of η can be "swallowed" in ϵ.

We now turn our attention to the pseudorandomness property of Theorem II.2. Unfortunately, the dual distance t of $\mathsf{OddRawRS}$ codes cannot be larger than $\sqrt{n}$. This means, that at best, $\mathrm{Enc}_{\mathsf{OddRawRS}}(U_{k/2} \circ x)$ is t-wise independent for $t < \sqrt{n}$, whereas the results of [HLV18] give nothing unless $t \gg n^{2/3}$.

Recent work by Lee and Viola [LV17], and Forbes and Kelly [FK18], showed that "t-wise independence + *large weight* noise" *is* pseudorandom for small space any-order ROBPs even for small t (e.g., $t = O(s + \log n)$). However, these results use noise that is a conjunction of a t-wise independent distribution with a uniform distribution, and such noise has relative weight roughly $\frac{1}{4}$. This will not do for our list-decoding argument.

Fortunately, we can use the technique of Forbes and Kelly [FK18] to give a better analysis than Haramaty, Lee and Viola [HLV18] and show that $\mathrm{Enc}(x, U_d)$ is pseudorandom for any-order space $s = \Omega(t)$ ROBPs even for $t = n^{\Omega(1)}$. The precise statement appears in Theorem V.6. This shows that $\mathrm{Enc}_{\mathrm{ctrl}}$ satisfies the properties in Definition II.1 and proves Theorem II.2.

B. The construction of stochastic codes for space bounded channels

In this section we give a sketch of the construction of stochastic codes for any-order space bounded channels. Our construction heavily builds on the machinery developed by Guruswami and Smith [GS16] (which in turn relies on previous ideas by Lipton [Lip94] and Smith [Smi07]). We also use the refinements of Shaltiel and Silbak [SS16], as well as several new modifications and simplification.

Recall that given a constant $\epsilon > 0$, our goal is to design a stochastic code $\mathrm{Enc} : \{0,1\}^{RN} \times \{0,1\}^d \to \{0,1\}^N$ that has rate $R \geq 1 - H(p) - \epsilon$, and is list-decodable for small space channels that induce pN errors. Furthermore, we aim for quasilinear time encoding and list-decoding with constant sized lists.

1) The encoding algorithm: To encode a message $m \in \{0,1\}^{RN}$ the encoding algorithm will encode m by a code $\mathrm{Enc}_{\mathrm{BSC}}$ for binary symmetric channels BSC_p. There are explicit constructions of such codes with rate approaching $R = 1 - H(p)$ and linear time encoding and decoding [GI05]. Thinking ahead, we set the block length of $\mathrm{Enc}_{\mathrm{BSC}}$ to be $N_{\mathrm{data}} = (1-\epsilon)\cdot N$, which we can do by choosing a smaller constant $\epsilon_{\mathrm{BSC}} = \epsilon/10$ for the BSC code. The encoding algorithm will also select random "seeds" to activate several "pseudorandom components". The seeds of all these components will be of length N^α for some small constant $\alpha > 0$, and so their length is negligible compared to N.

The first seed s_π will be used to generate an "almost t-wise independent permutation" $\pi : [N_{\mathrm{data}}] \to [N_{\mathrm{data}}]$. In this high level overview we will pretend that π is a random permutation. The encoding algorithm computes $x = \mathrm{Enc}_{\mathrm{BSC}}(m)$, and $y = \pi^{-1}(x)$ (recall that this means that the bits of x are "reordered" according to π^{-1}). Loosely speaking, this is done so that if the channel is an *additive channel* (namely one that has a fixed error pattern $e \in \{0,1\}^{N_{\mathrm{data}}}$ of weight pN) and if the decoding algorithm has a copy of s_π, then the decoding algorithm can apply π on the received word $y \oplus e$ and obtain $\mathrm{Enc}(m) \oplus \pi(e)$. For a random permutation π, the distribution $\pi(e)$ is very similar to BSC_p, and so the decoding algorithm can decode by applying $\mathrm{Dec}_{\mathrm{BSC}}$.

The argument above (which was suggested by Lipton [Lip94]) crucially requires that the decoding algorithm receives the seed s_π. The approach of Guruswami and Smith is to encode the seed s_π (as well as other seeds that we introduce soon) by a "control code" and "merge" y and this "control encoding" together, in the hope that the channel is not able to "wipe out" the control information, and furthermore, that the decoding algorithm is able to identify and correctly decode the control information.

For this purpose, the encoding algorithm also chooses a random seed s_{PRG} for a pseudorandom generator G that fools any-order small space ROBPs (we use the PRG of Forbes and Kelly [FK18]). When preparing the data part, the encoding xors the string $y = \pi^{-1}(\mathrm{Enc}_{\mathrm{BSC}}(m))$ with $G(s_{\mathrm{PRG}})$ to obtain the "data codeword" $c_{\mathrm{data}} = \pi^{-1}(\mathrm{Enc}_{\mathrm{BSC}}(m)) \oplus G(s_{\mathrm{PRG}})$. Loosely speaking, this means that c_{data} looks random to the channel.

The encoding algorithm now prepares the control codeword. For this purpose, the encoding algorithm divides the N output bits into $n = N^{1-\lambda}$ blocks of length $b = N^\lambda$, where $\lambda > 0$ is some small constant. It chooses an additional random seed s_{samp} for an "averaging sampler". This seed is used to specify $\epsilon \cdot n$ distinct indices $i_1, \ldots, i_{\epsilon \cdot n} \in [n]$. These blocks are called "control blocks", and the remaining blocks are called "data blocks". In this high level overview we pretend that the indices of control blocks are uniformly distributed in $[n]$. (Loosely speaking, the definition of averaging samplers allows us to make this assumption).

The final codeword $c \in \{0,1\}^N$ is prepared as follows: Note that the total length of data blocks is N_{data}, and the encoding algorithm "places" the data codeword c_{data} in these blocks. The remaining ϵn blocks are used to encode the "control information" $s = (s_\pi, s_{\mathrm{PRG}}, s_{\mathrm{samp}})$. This is done as follows: for each control block i, the encoding algorithm sets $c_i = \mathrm{Enc}_{\mathrm{ctrl}}(s, U_d)$ with fresh randomness for each block (where $\mathrm{Enc}_{\mathrm{ctrl}}$ is the stochastic control code of the previous section).[6] Note that this indeed gives a codeword c of length N.

2) The list decoding algorithm: In order to decode, the decoding algorithm first applies the list-decoding algorithm $\mathrm{Dec}_{\mathrm{ctrl}}$ on all the n blocks of the received word. The decoding algorithm obtains $\ell = \frac{n}{\epsilon^c}$ outcomes (where $c = 2$ is the exponent of the list size of $\mathrm{Dec}_{\mathrm{ctrl}}$), and it "passes on" each outcome that appears at least say $\epsilon^2 n$ times. (Note that there are at most $\mathrm{poly}(1/\epsilon)$ such outcomes). For each such candidate s' for control information, the decoding produces a message. Namely, it identifies the partition of blocks into control blocks and data blocks according to s'_{samp}. It then xors the data part with $G(s'_{\mathrm{PRG}})$, and applies the permutation π (defined by seed s'_π). Finally, it performs decoding by $\mathrm{Dec}_{\mathrm{BSC}}$. This process indeed produces a list of size $\mathrm{poly}(1/\epsilon)$ of candidate messages.

The analysis will show that for any small space channel, w.h.p. the "correct control information" s is one of the candidates s' considered by the decoding algorithm, and that with $s' = s$, the correct message m is decoded w.h.p.

3) Analyzing the construction: The analysis of the construction is quite involved and is presented in detail in the full version. On a high level, the key observation is that from the point of view of a space bounded channel, the data part looks random (as it is xored with the output of a pseudorandom generator) and each control block looks random (by the pseudorandom property of $\mathrm{Enc}_{\mathrm{ctrl}}$). This intuitively means that the channel cannot distinguish data blocks from control blocks, and therefore, from its point of view, the position of control blocks is random (as they were chosen by the sampler). It intuitively follows that the channel cannot hope to "wipe out" the short control part. At best, it can place a p fraction of errors on the control part, and it is likely that an ϵ fraction of the control blocks will be decoded correctly by $\mathrm{Dec}_{\mathrm{ctrl}}$, meaning that the correct control information s is one of the candidates s' that is considered by the decoding algorithm.

The channel C chooses the error pattern e as a function of the codeword c. However, as c looks random to the channel, the channel intuitively chooses e in a way that is independent of the seed s_π. Therefore, the analysis used earlier for additive channels (in which e is fixed and π is random) can be applied, and the correct

---

[6] We mention that here we simplify previous work by [GS16, SS16] that also used an "outer control code" that was chosen to be list-recoverable. This simplification allows us to speed up the encoding and decoding as explained later.

message appears in the list.[7]

4) Achieving quasilinear time encoding and decoding.: In order to achieve quasilinear time encoding and decoding, we need to first verify that none of the components we use, runs in larger polynomial time. In some cases (e.g., the PRG of Forbes and Kelly [FK18]) we need to delve into the construction and analysis in order to implement it in a more efficient manner.

A key observation is that we don't have to optimize the exponent of the polynomial in the time of encoding and decoding $\mathrm{Enc}_{\mathrm{ctrl}}$. This is crucial as for example, the time of the decoding algorithm that we give for OddRawRS codes is inherently at least quadratic (just for writing all the large lists in the list recovering stage).

We make the following observation: When encoding, we run $\mathrm{Enc}_{\mathrm{ctrl}}$ (which in turn runs the linear code $\mathrm{Enc}_{\mathsf{OddRawRS}}$) many times. Therefore, we only care about amortized encoding time (rather than worst case running time). Any linear function $L : \mathbb{F}_2^{n^{0.1}} \to \mathbb{F}_2^n$ can be computed in amortized time $O(n \cdot (\log^2 n))$ (following a pre-processing step that prepares the matrix of L). This is because making n such computations, can be reduced to matrix multiplication of an $n \times n^{0.1}$ matrix by an $n^{0.1} \times n$ matrix (which can be done in time $O(n^2 \cdot \log^2 n)$ by Coppersmith [Cop82]. Recall that we have less than $N^{1-\lambda}$ applications of $\mathrm{Enc}_{\mathrm{ctrl}}$ where each one is over block length N^{λ}, and so overall, all these applications take time $O(N \cdot \log^2 N)$

A second observation is that when list-decoding, the decoding algorithm doesn't need to try all $n = N^{1-\lambda}$ blocks. It can instead sample a polylogarithmic number of blocks in $[n]$, and only try to decode the control code on the sampled blocks. Each such block is of length $b = N^{\lambda}$, and so, even if applying $\mathrm{Dec}_{\mathrm{ctrl}}$ takes time b^c for a large constant c, by choosing the constant $\lambda > 0$ to be sufficiently small, this step takes time $\mathrm{polylog}(N) \cdot N^{\lambda \cdot c} \leq N$.

Finally, we mention that an obvious bottleneck that prevents our encoding and decoding to run in linear time, is that computing a permutation $\pi : [N] \to [N]$ on all N inputs, requires time $\Omega(N \cdot \log N)$ just to read the inputs and write the outputs.

III. PRELIMINARIES, AND INGREDIENTS USED IN THE CONSTRUCTION

In this section we give formal definitions of the notions and ingredients used in the construction. We also cite previous results from coding theory and pseudorandomness that are used in the construction.

a) General notation: We use U_n to define the uniform distribution on n bits. The statistical distance between two distributions P, Q over Ω is $\max_{A \subseteq \Omega} |P(A) - Q(A)|$. Random variables $R_1, \dots, R_n$ are t-wise independent if for every $i_1, \dots, i_t \in [n]$, $R_{i_1}, \dots, R_{i_t}$ are uniform and independent.

The Hamming weight of $x \in [q]^n$ is *weight*$(x) = |\{i : x_i \neq 0\}|$. The Hamming distance between $x, y \in [q]^n$ is $|\{i : x_i \neq y_i\}|$ and the relative Hamming distance is the Hamming distance divided by n.

A. Permuting strings

We will use a permutation $\pi : [n] \to [n]$ to "reorder" the bits of a string $x \in \{0,1\}^n$: The i'th bit in the rearranged string will be $\pi(i)$'th bit in x. This is captured in the definition below.

Definition III.1 (Permuting strings)**.** *Given a string $x \in \{0,1\}^n$ and a permutation $\pi : [n] \to [n]$. Let $\pi(x)$ denote the string $x' \in \{0,1\}^n$ with $x'_i = x_{\pi(i)}$.*

B. Read once branching program (in any order)

1) Formal definition of ROBPs and bounded space channels: We give a more formal definition of bounded space computation and channels, restating Definition I.3 in a more formal notation. The model that we consider is that of oblivious read once branching programs (ROBP). In the definition below, we will consider several variants depending on whether the ROBP outputs a single bit, or one bit per any input bit (which is the case for channels that are function $C : \{0,1\}^n \to \{0,1\}^n$). We will also consider ROBPs that are allowed to choose the order in which they read the input bits.

Definition III.2 (Read Once Branching Programs (ROBP))**.** *A space s **ROBP** C which receives input in $\{0,1\}^n$ is defined by picking n transition functions*

---

[7] This high level argument is an oversimplification and the actual proof is quite involved. We need to show that if the channel is able to prevent the decoding algorithm from decoding, then it can be used to break one of the pseudorandom components. A significant difficulty, is that the channel cannot run the decoding algorithm (which cannot be run by a small space ROBP) and therefore, the channel "does not know" whether it succeeded in preventing the decoding algorithm from decoding correctly. This is a problem as the distinguisher (for the PRG) that we aim to construct, will intuitively want to distinguish the output of the PRG from random, by distinguishing between the case that decoding succeeded from the one where it doesn't (and in particular, the distinguisher will want to run decoding algorithms). The argument used to construct this distinguisher relies on additional specific properties of the BSC code. Our approach to handling this issue, builds heavily on the previous arguments of [GS16, SS16] with some modifications.

We also remark that a possible behavior of a channel is to inject "false control strings" in order to make the decoding algorithm decode to incorrect values. Indeed, there are bounded space channels that can cause the decoding algorithm to have incorrect messages in the list (in addition to the correct one).

$\delta_1, \dots, \delta_n$ where for each i, $\delta_i : \{0,1\}^s \times \{0,1\} \to \{0,1\}^s$. On input $x \in \{0,1\}^n$, the computation path of C is the sequence $r_0, \dots, r_n$ of states defined by $r_0 = 0^s$ and for $i \geq 1$, $r_i = \delta_i(r_{i-1}, x_i)$. We distinguish between two types of ROBPs:

- *If $C : \{0,1\}^n \to \{0,1\}$ is an ROBP that outputs a single bit, then C also has an output function $o : \{0,1\}^s \to \{0,1\}$ and $C(x)$ is defined to be $o(r_n)$.*
- *If $C : \{0,1\}^n \to \{0,1\}^n$ is an ROBP that outputs n bits, then C also has n output functions $o_1, \dots, o_n$ where for each i, $o_i : \{0,1\}^s \to \{0,1\}$ and $C(x)$ is defined by the n bit string $o_1(r_1), \dots, o_n(r_n)$.*

We are stating this definition in the terminology of "transition functions" and "output functions" which is more convenient when discussing ROBPs that output more than one bit. However, we stress that this definition is equivalent to the more common definition of width $w = 2^s$ ROBPs in terms of a layered graph with $n+1$ layers, where the i'th transition function specifics the edges from the $(i-1)$'th level to the i'th level.

Another remark is that the definition above forces an ROBP that outputs many bits to output its i'th bit *before* seeing the $(i+1)$'th bit. This is done in order to have a simple definition of ROBPs that output many bits. However, all our results hold for a more general model in which the ROBP can delay outputting the i'th bit to a later stage and look ahead at the next input bits.[8]

We now define *any-order* ROBPs that are allowed to reorder their input bits using a permutation $\sigma : [n] \to [n]$ prior to reading the input. The next claim immediately follows from the definition.

Definition III.3 (any-order ROBPs)**.** *Given an ROBP C over n bits, and a permutation $\sigma : [n] \to [n]$ we define C^σ to be the function $C^\sigma(x) = C(\sigma(x))$. (Here $\sigma(x)$ is the function from Definition III.1). The class of* ***any-order space** s **ROBPs*** *is the class of all functions C^σ where C is a space s ROBP and $\sigma : [n] \to [n]$ is a permutation.*

We now observe that if we restrict the input of an any-order space s ROBP, then we obtain an any-order space s ROBP.

Claim III.4 (Restrictions of any-order ROBPs)**.** *Given a space s ROBP $C : \{0,1\}^n \to \{0,1\}$, a permutation $\sigma : [n] \to [n]$, $T \subseteq [n]$ of size t, and $v \in \{0,1\}^t$, the function $f : \{0,1\}^{n-t} \to \{0,1\}$ defined by $f(y) = C^\sigma(x)$ where $x_T = v$ and $x_{[n]\setminus T} = y$, is computable by an any-order space s ROBP. That is, there exists a permutation $\tau : [n-t] \to [n-t]$ and a space s ROBP $D : \{0,1\}^{n-t} \to \{0,1\}$ such that $D^\tau(y) = f(y)$.*

Definition III.3 applies also to ROBPs that output n bits. Note that in that case, the output of the ROBP is also ordered by σ. When considering channels, it is more convenient to reorder the bits back to the natural order. This is done in the next definition.

Definition III.5 (any-order bounded space channels)**.** *The class of* ***any-order space** s **channels*** *is the class of all functions $e^\sigma_C : \{0,1\}^n \to \{0,1\}^n$, where $C : \{0,1\}^n \to \{0,1\}^n$ is a space s ROBP, $\sigma : [n] \to [n]$ is a permutation, and $e^\sigma_C(x) = \sigma^{-1}(C(\sigma(x)))$.*

It should be noted that the "reordered function" e^σ_C is not necessarily computable by a small space ROBP. However, when applying channels e^σ_C on a codeword, it is more natural to reorder the bits in the order used by the codeword.

Using this notation, the bounded space channel model considered in [GS16, SS16] (which was called "online space s channels") corresponds to space s channels with the identity permutation (namely, channels that read their input bits in the standard order).

2) PRGs for any-order ROBPs: We need the following standard definition of pseudorandom distributions and generators.

Definition III.6 (Pseudorandom generators)**.** *A distribution X on n bits is ϵ-**pseudorandom** for a class $\mathcal{C}$ of functions from n bits to one bit, if for every $C \in \mathcal{C}$, $|\Pr[C(X) = 1] - \Pr[C(U_n)] = 1| \leq \epsilon$. A function $G : \{0,1\}^d \to \{0,1\}^n$ is an ϵ-**PRG** for $\mathcal{C}$ if $G(U_d)$ is ϵ-pseudorandom for $\mathcal{C}$.*

We will use the following PRG by Forbes and Kelly [FK18]

Theorem III.7. *[FK18] For every $\log n \leq s \leq n$, there exists an ϵ-PRG $G : \{0,1\}^d \to \{0,1\}^n$ for any-order space s ROBPs that output one bit, with $d = O((s + \log \frac{1}{\epsilon}) \cdot \log^2 n)$. Furthermore, G can be computed in time $O(n \cdot polylog(n))$.*

---

[8] To make this statement more concrete, a space s ROBP that wants to look ahead and read t additional bits before outputing the i'th output bit, can store these additional t bits, and the number of bits it outputted so far, in its memory, and this can be done in space $s + t + \log n$. Our results on space s channels also apply to these kind of channels, namely channels that use space $\Omega(s)$ and look ahead at the next $\Omega(s)$ input bits. More generally, our results apply to any "reasonable" model of ROBPs that output many bits, in which if $C_1, C_2 : \{0,1\}^n \to \{0,1\}^n$ are space s ROBPs, then the composition $C_1 \circ C_2$ can be computed by an ROBP with space, say $O(s)$.

Forbes and Kelly [FK18] do not carefully estimate the running time of their pseudorandom generator, and only claim that it runs in polynomial time. The proof below, will prove the "furthermore" clause in the theorem above.

Proof. (of the "furthermore" clause in Theorem III.7) The construction of Forbes and Kelly works as follows: Let k be a parameter to be chosen later. The generator G is constructed iteratively, by setting G_0 to be a $320k$-wise independent distribution, and $G_{i+1} = D_i \oplus (T_i \wedge G_i)$ where D_i is a $2k$-wise independent distribution, and T_i is a k-wise independent distribution. (Different copies of D_i's and T_i's are sampled independently). The final distribution is $G = G_r$ where r is a parameter chosen by the construction.

Note that sampling a k-wise independent distribution X on n bits can be done by a deterministic procedure that receives a seed of length $k \log n$, in time $n \cdot \text{polylog}(n)$. This can be done by the standard Reed-Solomon based construction. Namely, encoding the $k \log n$ bit seed as k elements $a_0, \dots, a_{k-1} \in \mathbb{F}_n$ (here we assume w.l.o.g. that n is a power of 2) and for $\alpha \in [n]$ (which can be interpreted as $\alpha \in \mathbb{F}_n$) setting X_α to be (the first bit of) $\sum_{0 \le i < k} a_i \alpha^i$. This gives a k-wise independent distribution and can be computed in time $n \cdot \text{polylog}(n)$, using univariate multipoint evaluation, that can be done with $O(n \log n)$ field operations.

Forbes and Kelly show that taking $r = O(\log n)$ and $k = O(s + \log n + \log(1/\epsilon))$ gives an ϵ-PRG for any-order space s ROBPs. This gives total running time of $r \cdot n \cdot \text{polylog}(n) = n \cdot \text{polylog}(n)$.

This follows by the proof of [FK18, Lemma 4.2] which proves the correctness of the PRG with a slightly different construction, but also applies to the construction described above. ☐

C. Averaging Samplers

The reader is referred to Goldreich's survey [Gol97] on averaging samplers.

Definition III.8 (Averaging Samplers). *A function* $\text{Samp} : \{0,1\}^n \to (\{0,1\}^m)^t$ *is an* (ϵ, δ)-***Sampler*** *if for every* $f : \{0,1\}^m \to [0,1]$, $\Pr_{(z_1,\dots,z_t) \leftarrow \text{Samp}(U_n)}[|\frac{1}{t}\sum_{i \in [t]} f(z_i) - \frac{1}{2^m}\sum_{x \in \{0,1\}^m} f(x)| > \epsilon] \le \delta$. *A sampler has* ***distinct samples*** *if for every* $x \in \{0,1\}^n$, *the* t *elements in* $\text{Samp}(x)$ *are distinct.*

The next theorem follows from the "expander sampler". This particular form can be found (for example) in [Vad04].

Theorem III.9. *For every sufficiently large* m *and every* $\epsilon \ge \delta > 0$ *such that* $m \le \log(1/\delta)$ *there is an* (ϵ, δ)*-sampler with distinct samples,* $\text{Samp} : \{0,1\}^{O(\log(1/\delta) \cdot \text{poly}(1/\epsilon))} \to (\{0,1\}^m)^t$ *for any* $t \le 2^m$ *such that* $t \ge \text{poly}(1/\epsilon) \cdot \log(1/\delta)$*. Furthermore,* Samp *is computable in time* $t \cdot \text{poly}(1/\epsilon, \log(1/\delta))$ *and has distinct samples.*

D. Almost t-wise permutations

We also need the following notion of almost t-wise permutations.

Definition III.10 (Almost t-wise independent permutations). *A function* $\pi : \{0,1\}^d \times [n] \to [n]$ *is an* (ϵ, t)*-wise independent permutation if:*

- *For every* $s \in \{0,1\}^d$ *the function* $\pi_s(i) = \pi(s, i)$ *is a permutation over* $[n]$.
- *For every distinct* $i_1, \dots, i_t \in [n]$, *the random variable* $R = (R_1, \dots, R_t)$ *defined by* $R_j = \pi(s, i_j)$: $s \leftarrow U_d$, *is* ϵ*-close to* t *uniform samples without repetition from* $[n]$.

Theorem III.11. *[KNR09] For every* t *and every sufficiently large* n, *there exists an* (ϵ, t)*-wise independent permutation with* $d = O(t \cdot \log n + \log(1/\epsilon))$*. Furthermore, computing* $\pi(s, i)$ *on inputs* $s \in \{0,1\}^d$ *and* $i \in [n]$ *can be done in time* $\text{poly}(d, \log n)$.[9]

We will use (ϵ, t)-wise independent permutations to permute strings. Consider the following example: Let $e \in \{0,1\}^n$ be a string with Hamming weight pn, and let $\pi : \{0,1\}^d \times [n] \to [n]$ be an (ϵ, t)-wise independent permutation. We will be interested in the distribution $X = \pi_{U_d}(e)$ (here, $\pi(e)$ is the "permuted string" defined in Definition III.1). We would like to apply "Chernoff style bounds for t-wise independence" [BR94, SSS95] on $X_1, \dots, X_n$. A technical issue is that it is not true that $X_1, \dots, X_n$ are t-wise independent (even in the case that $\epsilon = 0$). What is true is that for every t-tuple of distinct indices $i_1, \dots, i_t \in [n]$, $\Pr[X_{i_1} = \ldots = X_{i_t} = 1] \le p^t + \epsilon$. The latter condition is sufficient to obtain Chernoff style behavior (at least when ϵ is sufficiently small compared to p^t) by the following lemma.

Lemma III.12 (tail bounds for almost t-wise independent permuted strings). *Let* $X_1, \dots, X_n$ *be binary random variables, such that for every set of distinct* t *indices*

---

[9] We will be interested in the time it takes to compute the permutation on all $i \in [n]$ (namely given s, we want to compute $(\pi(s,i))_{i \in [n]}$) and will use $n \cdot \text{poly}(d)$ as a bound on the time for this task. Note that this also gives that computing $(\pi^{-1}(s,i))_{i \in [n]}$ can be done within the same time bound.

$i_1, \cdots, i_t \in [n]$, $\Pr[X_{i_1} = \ldots = X_{i_t} = 1] \leq \mu^t$. *If* $0 < \delta \leq 1$ *and* $t \leq \frac{\delta \cdot \mu \cdot n}{2}$ *then*

$$Pr[\sum_{j=1}^{n} X_j \geq (1+\delta) \cdot \mu \cdot n] \leq e^{-\Omega(\delta t)}$$

This type of lemma follows by using the approach of [SSS95]. It was applied for t-wise independent permutations (in a related setup) in [DHRS07]. For completeness we provide a proof.

Proof. (of Lemma III.12) We use X to denote the n bit long random variable composed of $(X_1, \ldots, X_n)$. For every t-tuple $y = (i_1, \ldots, i_t)$ of indices in $[n]$, let A_y be the event $A_y = \{X_{i_1} = \ldots, X_{i_t} = 1\}$. Let Y be an independent, uniformly chosen t-tuple of indices in $[n]$. We have that $\Pr[A_Y] \leq \mu^t$. Let $\ell = (1+\delta) \cdot \mu \cdot n$. Note that for every $x \in \{0,1\}^n$ such that $\sum_{j \in [n]} x_j \geq \ell$, we have that:

$$\Pr[A_Y | X = x] \geq \frac{\binom{\ell}{t}}{\binom{n}{t}} = \frac{\ell \cdot \ldots \cdot (\ell - t + 1)}{n \cdot \ldots \cdot (n - t + 1)}$$
$$\geq \frac{(\ell - t + 1)^t}{n^t} \geq \mu^t \cdot (1 + \frac{\delta}{2})^t.$$

This gives that,

$$\Pr[A_Y] \geq \Pr[A_Y \cap \{X \geq \ell\}]$$
$$\geq \Pr[X \geq \ell] \cdot \Pr[A_Y | X \geq \ell]$$
$$\geq \Pr[X \geq \ell] \cdot \mu^t \cdot (1 + \frac{\delta}{2})^t.$$

Rearranging, we get that:

$$\Pr[X \geq \ell] \leq \frac{\Pr[A_Y]}{\mu^t \cdot (1+\frac{\delta}{2})^t} \leq \frac{\mu^t}{\mu^t \cdot (1+\frac{\delta}{2})^t} = e^{-\Omega(\delta \cdot t)}.$$

□

E. Error-Correcting Codes

In this section we give definitions of the various notions of error correcting codes used in this paper. We also state some previous constructions that will be used in this paper.

1) The standard notion of error correcting codes: We give a more general version of Definition I.1 that discusses codes over non-binary alphabets, as well as codes in Shannon's scenario. For our purposes it is more natural to define codes in terms of a pair (Enc, Dec) of encoding and decoding algorithms. Different variants are obtained by considering different properties required by the encoding and decoding algorithms and different types of error patterns.

Definition III.13 (Codes)**.** *Let* k, n, q *be parameters and let* $\text{Enc} : \{0,1\}^k \to (\{0,1\}^{\log q})^n$ *be a function. We say that* Enc *is an encoding function for a code that is:*

- ***decodable from*** t ***errors****, if* $t \in [n]$*, and there exists a function* $\text{Dec} : (\{0,1\}^{\log q})^n \to \{0,1\}^k$ *such that for every* $m \in \{0,1\}^k$ *and every* $e \in (\{0,1\}^{\log q})^n$ *with Hamming weight at most* t*,* $\text{Dec}(\text{Enc}(m) \oplus e) = m$*.*
- L***-list-decodable from*** t ***errors****, if the function* Dec *is allowed to output a list of size at most* L*, and for every* $m \in \{0,1\}^k$ *and every* $e \in (\{0,1\}^{\log q})^n$ *with Hamming weight at most* t*,* $\text{Dec}(\text{Enc}(m) \oplus e) \ni m$*.*
- ***decodable from*** P*, with success probability* $1-\nu$*, if* P *is a distribution over* $(\{0,1\}^{\log q})^n$*,* $0 \leq \nu \leq 1$*, and there exists a function* $\text{Dec} : (\{0,1\}^{\log q})^n \to \{0,1\}^k$ *such that for every* $m \in \{0,1\}^k$*,* $\Pr_{e \leftarrow P}[\text{Dec}(\text{Enc}(m) \oplus e) = m] \geq 1 - \nu$*.*

A code has ***encoding time*** *[resp.* ***decoding time****]* $T(\cdot)$*, if* Enc *[resp.* Dec*] can be computed in time* $T(n \log q)$*. The code is* ***explicit*** *if both encoding and decoding run in polynomial time. (Naturally, this makes sense only for a family of encoding and decoding functions with varying block length* n*, message length* $k(n)$*, and alphabet size* $q(n)$*).*

The rate of the code is the ratio of the message length and output length of Enc*, where both lengths are measured in bits. That is the rate* $R = \frac{k}{n \cdot \log q}$*.*

2) Linear Codes and Dual Distance: We also define the standard notion of linear codes.

Definition III.14 (Linear codes and dual codes)**.** *Let* q *be a prime power, and let* $\mathbb{F}_q$ *denote the field with* q *element. An* $[n,k]_q$ ***linear code*** *is a linear subspace of* $C \subseteq \mathbb{F}_q^n$ *of dimension* k*. We say that* C *has distance* d*, if the Hamming weight of every nonzero vector in* C *is at least* d*. Such codes are called* $[n,k,d]_q$ *codes. A linear map* $\text{Enc} : \mathbb{F}_q^k \to \mathbb{F}_q^n$ *is an* ***encoding function*** *for* C*, if* $\text{Enc}(\mathbb{F}_q^k) = C$*. For a code* C*, we use* $C^\perp$ *to denote the dual vector space. We say that* C *has* ***dual distance*** d *if* $C^\perp$ *has distance* d*.*

It is standard that C is an $[n,k,2t+1]_q$ code iff C has a linear encoding function $\text{Enc} : \mathbb{F}_q^k \to \mathbb{F}_q^n$ that is decodable from t errors. We will use the standard fact that encoding functions for codes with dual distance r yield $(r-1)$-wise dual independent distributions.

Lemma III.15 ($(t-1)$-wise independence from linear codes with dual distance t)**.** *Let* $\text{Enc} : \mathbb{F}_q^k \to \mathbb{F}_q^n$ *be an encoding function for a linear* $[n,k]_q$*-code* C *with dual distance* t*. Applying* Enc *on a uniformly chosen message*

$m \leftarrow \mathbb{F}_q^k$ yields a distribution $(Z_1, \ldots, Z_n)$ over $\mathbb{F}_q^n$ that is $(t-1)$-wise independent, and every Z_i is uniformly distributed over $\mathbb{F}_q$.

Proof. Applying Enc on some $v \in \mathbb{F}^k$, can be seen as multiplying v by a generator matrix of C (which is a transposed parity check matrix of $C^\perp$). As $C^\perp$ has distance t, every $t-1$ columns of the generator matrix of C are linearly independent, and the lemma follows. □

3) Stochastic Codes: We restate Definition I.2 using slightly more precise notation.

Definition III.16 (Stochastic codes for channels). *Let k, n, d be parameters and let $\mathrm{Enc} : \{0,1\}^k \times \{0,1\}^d \to \{0,1\}^n$ be a function. Let $\mathcal{C}$ be a class of functions from n bits to n bits. We say that Enc is an encoding function for a stochastic code that is:*

- ***decodable** for "channel class" $\mathcal{C}$, with success probability $1-\nu$, if there exists a (possibly randomized) procedure $\mathrm{Dec} : \{0,1\}^n \to \{0,1\}^k$ such that for every $m \in \{0,1\}^k$ and every $C \in \mathcal{C}$, setting $X = \mathrm{Enc}(m, U_d)$, we have that $\Pr[\mathrm{Dec}(X \oplus C(X)) = m] \geq 1-\nu$, where the probability is over coin tosses of the encoding and decoding procedures.*
- *L-**list-decodable** for "channel class" $\mathcal{C}$, with success probability $1-\nu$, if the procedure Dec is allowed to output a list of size at most L, and $\Pr[\mathrm{Dec}(X \oplus C(X)) \ni m] \geq 1-\nu$, where the probability is over coin tosses of the encoding and decoding procedures.*

*A code has **encoding time** [resp. **decoding time**] $T(\cdot)$, if Enc [resp. Dec] can be computed in time $T(k+n+d)$. The code is **explicit** if both encoding and decoding run in polynomial time. (Naturally, this makes sense only for a family of encoding and decoding functions with varying block length n, message length $k(n)$ and seed length $d(n)$).*

The rate of the code is the ratio of the message length and output length of Enc, where both lengths are measured in bits. That is the rate $R = \frac{k}{n}$.

4) Codes for binary-symmetric channels and related variants: We will make use of known constructions of codes for binary symmetric channels.

Definition III.17 (Binary symmetric channel). *Let BSC_p^n denote the distribution over n bit strings in which individual bits are i.i.d. and each is one with probability p.*

There are constructions of codes with rate approaching $1-H(p)$ that are decodable from BSC_p^n with very high suvcess probability, and have linear time encoding and decoding [GI05].

We are interested in codes for an intuitively similar (but more general) scenario in which the error distribution is obtained by taking a string $e \in \{0,1\}^n$ of weight pn, an (ϵ, t)-wise independent permutation $\pi : \{0,1\}^d \times [n] \to [n]$ and considering the error distribution $e' = \pi_{U_d}(e)$.

This distribution is somewhat similar to BSC_p in the sense that if we project both distributions to a "not too large" tuple of indices, the distributions are statistically close. More precisely, for every choice of t distinct indices $I = (i_1, \ldots, i_t)$, the distribution $(\mathrm{BSC}_p)_I$ and $(\pi_{U_d}(e))_I$ are $(\epsilon + t^2/n)$-close in statistical distance. This can be used to argue that current constructions for BSC_p also work for $\pi_{U_d}(e)$ (for certain parameters).

However, for technical reasons, this isn't sufficient for our purposes, and we will require that the code has some additional structure (which we will use in our construction). We now explain the additional structure that we need: The known codes for BSC_p are constructed by code concatenation, and for technical reasons, we will be interested in some properties of the inner and outer codes (and not just properties of the concatenated code). We first give the following standard definition of code concatenation.

Definition III.18 (Concatenated code). *Given functions:*

- $\mathrm{Enc}_{\mathrm{out}} : \{0,1\}^{k_{\mathrm{out}}} \to (\{0,1\}^{\log q_{\mathrm{out}}})^{n_{\mathrm{out}}}$*, and*
- $\mathrm{Enc}_{\mathrm{in}} : \{0,1\}^{k_{\mathrm{in}}} \to (\{0,1\}^{\log q_{\mathrm{in}}})^{n_{\mathrm{in}}}$,

*such that $\log q_{\mathrm{out}} = k_{\mathrm{in}}$ we define the **concatenated encoding function** $\mathrm{Enc} : \{0,1\}^{k_{\mathrm{out}}} \to (\{0,1\}^{\log q_{\mathrm{in}}})^{n_{\mathrm{out}} \cdot n_{\mathrm{in}}}$ denoted by $\mathrm{Enc}_{\mathrm{out}} \circ \mathrm{Enc}_{\mathrm{in}}$ as follows: For $i_{\mathrm{out}} \in [n_{\mathrm{out}}]$, $i_{\mathrm{in}} \in [n_{\mathrm{in}}]$, and $i = (i_{\mathrm{out}} - 1) \cdot n_{\mathrm{in}} + i_{\mathrm{in}}$ we define $\mathrm{Enc}(m)_i = \mathrm{Enc}_{\mathrm{in}}(\mathrm{Enc}_{\mathrm{out}}(m)_{i_{\mathrm{out}}})_{i_{\mathrm{in}}}$.*

Concatenated codes can be decoded by "concatenated decoding".

Definition III.19 (Concatenated decoding). *Let $\mathrm{Enc} = \mathrm{Enc}_{\mathrm{out}} \circ \mathrm{Enc}_{\mathrm{in}}$ be a concatenated code, and let $\mathrm{Dec}_{\mathrm{out}} : (\{0,1\}^{\log q_{\mathrm{out}}})^{n_{\mathrm{out}}} \to \{0,1\}^{k_{\mathrm{out}}}$, $\mathrm{Dec}_{\mathrm{in}} : (\{0,1\}^{\log q_{\mathrm{in}}})^{n_{\mathrm{in}}} \to \{0,1\}^{k_{\mathrm{in}}}$ be functions. For $i \in [n_{\mathrm{out}}]$ we define $\mathrm{Dec}_{\mathrm{in}}^i : (\{0,1\}^{\log q_{\mathrm{in}}})^{n_{\mathrm{out}} \cdot n_{\mathrm{in}}} \to \{0,1\}^{k_{\mathrm{in}}}$ by:*

$$\mathrm{Dec}_{\mathrm{in}}^i(z) = (z_{(i-1)\cdot n_{\mathrm{in}}+1}, \ldots, z_{i \cdot n_{\mathrm{in}}}).$$

*The **concatenated decoding** function $\mathrm{Dec} : (\{0,1\}^{\log q_{\mathrm{in}}})^{n_{\mathrm{out}} \cdot n_{\mathrm{in}}} \to \{0,1\}^{k_{\mathrm{out}}}$ is defined by:*

$$\mathrm{Dec}(z) = \mathrm{Dec}_{\mathrm{out}}(\mathrm{Dec}_{\mathrm{in}}^1(z), \ldots, \mathrm{Dec}_{\mathrm{in}}^{n_{\mathrm{out}}}(z)).$$

In the following theorem we revisit the code construction of [GI05] for BSC_p^n, and observe that the constructed concatenated code has some properties that we will use later on.

Theorem III.20. *For every constant $0 < p < 1/2$, and every sufficiently small constant $\epsilon > 0$, there exist integer constants $k_{\text{in}}, n_{\text{in}}, q_{\text{out}}$ and real constants $\lambda_1, \lambda_2, \lambda_3 > 0$ such that $k_{\text{in}} = \log q_{\text{out}} \le \frac{1}{\epsilon^7}$, and for infinitely many choices of n_{out} there exist functions:*

- $\text{Enc}_{\text{out}} : \{0,1\}^{k_{\text{out}}} \to (\{0,1\}^{\log q_{\text{out}}})^{n_{\text{out}}}$,
- $\text{Enc}_{\text{in}} : \{0,1\}^{k_{\text{in}}} \to \{0,1\}^{n_{\text{in}}}$,

such that:

- $R_{\text{out}} = \frac{k_{\text{out}}}{n_{\text{out}} \cdot \log q_{\text{out}}} \ge 1 - \frac{\epsilon}{10}$*, and* Enc_{out} *is decodable from* $w = \lambda_1 \cdot n_{\text{out}}$ *errors with linear time encoding and decoding.*
- $R_{\text{in}} = \frac{k_{\text{in}}}{n_{\text{in}}} \ge 1 - H(p) - \epsilon/10$*, and* Enc_{in} *is decodable from* $\text{BSC}_p^{n_{\text{in}}}$ *with probability* $1 - 2^{-\lambda_2 \cdot n_{\text{in}}}$*. This decoding is achieved by a function* Dec_{in} *that implements "maximum likelihood decoding".*
- *Consequently, setting* $n = n_{\text{out}} \cdot n_{\text{in}}$*, and* $q_{\text{in}} = 2$*, the concatenated code* $\text{Enc} = \text{Enc}_{\text{out}} \circ \text{Enc}_{\text{in}} : \{0,1\}^{k_{\text{out}}} \to \{0,1\}^n$ *is well defined, has rate* $R = \frac{k_{\text{out}}}{n} \ge 1 - H(p) - \epsilon$*, and is encodable in time* $O(n)$ *(where the constant* c *hidden in the* $O(\cdot)$ *depends on* ϵ*, and* $c = c(\epsilon) = 2^{\text{poly}(1/\epsilon)}$*).*
- *Let* $t \le n^{0.1}$*, and let* $\pi : \{0,1\}^d \times [n] \to [n]$ *be a* $(2^{-10 \cdot t}, t)$*-wise independent permutation. Let* $m \in \{0,1\}^{k_{\text{out}}}$*, and let* $A_m : \{0,1\}^n \to \{0,1\}$ *be the function that on input* $e' \in \{0,1\}^n$*, outputs one iff*

  $$\left| \left\{ i \in [n_{\text{out}}] : \text{Dec}_{\text{in}}^i(\text{Enc}(m) \oplus e') \ne \text{Enc}_{\text{out}}(m)_i \right\} \right| \le \frac{w}{10}.$$

  (Note that $A_m(e') = 1$ *implies that concatenated decoding that is applied on* $\text{Enc}(m) \oplus e'$ *indeed recovers* m*).*
  For every $e \in \{0,1\}^n$ *of Hamming weight at most* pn*,*

  $$\Pr[A_m(\pi_{U_d}(e)) = 1] \ge 1 - 2^{-\lambda_3 \cdot t}.$$

- *Consequently, for every* $e \in \{0,1\}^n$ *of Hamming weight at most* pn*, the code* Enc *is decodable from* $\pi_{U_d}(e)$ *with probability* $1 - 2^{-\lambda_3 \cdot t}$*. Furthermore, the concatenated decoding algorithm runs in time* $O(n)$ *(where the constant* c *hidden in the* $O(\cdot)$ *depends on* ϵ*, and* $c = c(\epsilon) = 2^{\text{poly}(1/\epsilon)}$*).*

The final item in Theorem III.20 follows from the penultimate item. However, for our purposes, the final item will not be sufficiently strong, and we will need to use the penultimate item (as well as the previous items). The advantage of the penultimate item is that we get that for every m, there exists a space $O(\log n)$ ROBP which implements the function A_m, in contrast to the entire concatenated decoding algorithm that does not seem to be implemented by small space ROBPs.

Theorem III.20 follows by noticing that the proofs of known construction of codes for binary symmetric channels (see e.g., [For65]) are achieved by code concatenation of codes with the properties listed above. The fourth item follows by using Lemma III.12 to analyze the behavior of this concatenated code on errors from the distribution $\pi_{U_d}(e)$. The proof appears in the full version.

Very similar arguments to the proof of Theorem III.20 were made by Smith [Smi07] and in an early version of [GS16].

IV. RAW REED-SOLOMON CODES

For our intended application, we need linear binary $[n,k]_2$ codes with:

- Relative distance $(1/2 - o(1))$.
- Large dual distance of at least $n^{\Omega(1)}$. (In fact, we need a slightly stronger property to be explained below).
- Polynomial time encoding.
- Polynomial time unique decoding from p-fraction errors for every $p < \frac{1}{4}$.
- Polynomial time list-decoding with list size $\text{poly}(1/\epsilon)$ from $(\frac{1}{2} - \epsilon)$-fraction errors, for every constant $\epsilon > 0$.

In this paper, we construct binary codes with the properties above. To the best of our knowledge, this is the first construction of such codes.

First a slight abuse of notation: for this section only, we will use the word distance to denote *relative distance*, as opposed to *absolute distance*. This helps with the exposition.

Reed Solomon codes exhibit all the properties above (in addition to constant rate, and larger dual distance) but only for large alphabets. As far as we are aware, there are only two known families of codes over the binary alphabet which have $\Omega(1)$ distance and $n^{\Omega(1)}$ absolute dual distance. The first family is dual-BCH codes, but we do not know decoding algorithms for these codes from $\Omega(1)$-fraction errors for this setting of parameters (it is known [KS07, KS13] how to decode from $\Omega(1)$-fraction errors only when the absolute dual distance is $O(\log n)$). The second family is based on Algebraic-Geometric codes (see the appendix to [Shp09] for a detailed exposition). AG codes are generalizations of Reed-Solomon codes, and retain many of the good features of Reed-Solomon codes while having the advantage of being realizable over constant size alphabets. An AG code with suitable parameters over a constant size alphabet $\mathbb{F}_{2^t}$ has $\Omega(1)$ distance and $\Omega(1)$ dual distance. To bring the alphabet down to binary, one can do code concatenation. However, typically *concatenation destroys dual distance*. But not always! If we concatenate with a *trivial code*,[10] that maps $\mathbb{F}_{2^t}$ to t-bit strings, the absolute dual distance is preserved under concatenation. On the other hand, using the trivial code makes the distance shrink by a factor t. This yields codes with $\Omega(1)$ distance and dual distance with efficient decoding algorithms (these are the codes that we use to prove Theorem I.5). However the lower bound on the distance that follows is nowhere near[11] $1/2$.

The binary codes that we construct here are obtained by concatenating Reed-Solomon codes (over a large alphabet) with a different trivial code for each coordinate of the Reed-Solomon code. We call the general class of such codes *Raw Reed-Solomon codes*. In the positive direction, concatenating with trivial codes preserves the absolute dual distance, and we get the required dual-distance property. On the other hand, since the outer Reed-Solomon code is over a large alphabet, the trivial codes must have superconstant block-length, and thus $o(1)$ distance. By default, concatenating with inner codes of $o(1)$ distance leads to the final codes having $o(1)$ distance ($O(1/\log n)$ to be precise). However, for *special choices* of the trivial codes, we use some deep algebraic tools[12] to give a direct analysis of the distance of these codes, which miraculously turns out to be $1/2 - o(1)$. Finally, using the powerful list-decoding machinery available for Reed-Solomon codes, we show that Raw Reed-Solomon codes can be list-decoded from nearly $(1/2 - \epsilon)$-fraction errors.

Reed Solomon codes have a lot of structure and many useful properties (in addition to their distance properties) and so, we believe that the fact that Reed-Solomon codes have the additional properties listed above (when viewed as binary codes appropriately) is of independent interest and may prove useful in other applications.

In the theorem below, we focus on codes with the parameters that we require for our application. The theorem below will follow from a more general result on Raw Reed-Solomon codes stated later.

Theorem IV.1 (Codes with large distance and dual distance)**.** *For every constant $0 < \alpha < 1/2$, and every sufficiently large m, setting $n = (2^m - 1) \cdot m$, and $k = n^{\alpha}$, there is a binary linear $[n, k]_2$-code C that satisfies:*

- *C has a linear encoding map $\mathrm{Enc} : \mathbb{F}_2^k \to \mathbb{F}_2^n$ that runs in time $\mathrm{poly}(n)$.*
- *C has relative distance $(\frac{1}{2} - O((\frac{\log n}{n})^{\frac{1}{2}-\alpha})) = (\frac{1}{2} - o(1))$.*
- *For every constant $p < 1/4$, Enc is decodable from p-fraction errors in time $\mathrm{poly}(n)$.*
- *There exists a universal constant b such that for every $\epsilon \geq b\sqrt{\alpha}$, Enc is $O(\frac{1}{\epsilon^2})$-list-decodable from $(\frac{1}{2} - \epsilon)$-fraction errors in time $\mathrm{poly}(n)$.*[13]
- *C has absolute dual distance $\Omega(\frac{n^{\alpha}}{\log n})$.*
- *Moreover, define $\mathrm{Enc}^{\mathrm{trunc}} : \mathbb{F}_2^{k/2} \to \mathbb{F}_2^n$ by $\mathrm{Enc}^{\mathrm{trunc}}(x) = \mathrm{Enc}(x \circ 0^{k/2})$, and consider the linear code $C' = \mathrm{Enc}^{\mathrm{trunc}}(\mathbb{F}_2^{k/2})$. It holds that C' has absolute dual distance $\Omega(\frac{n^{\alpha}}{\log n})$.*

In the remainder of the section we introduce Raw Reed-Solomon codes, study their properties, and use them to prove Theorem IV.1.

A. General Raw Reed-Solomon codes

Let $q = 2^m$. We will discuss a family of binary codes that are derived from Reed-Solomon codes over $\mathbb{F}_q$.

Start with an evaluation domain $D \subseteq \mathbb{F}_q$ and a degree bound d, and consider the Reed-Solomon code of evaluations on D of polynomials of degree at most d over $\mathbb{F}_q$. In order to convert this code to a binary code, we also choose a sequence $\mathbf{\Phi} = (\Phi_x)_{x \in D}$, where each Φ_x is an $\mathbb{F}_2$-linear bijection between $\mathbb{F}_q$ and $\mathbb{F}_2^m$.

In terms of this data, we define the Raw Reed-Solomon code $\mathsf{RawRS}[\mathbb{F}_q, d, D, \mathbf{\Phi}]$ as follows. The coordinates of the code are indexed by pairs $(x, i) \in D \times [m]$, and the codewords are indexed by polynomials $P(X) \in \mathbb{F}_q[X]$ of degree at most d. The codeword corresponding to $P(X)$ is $c : D \times [m] \to \mathbb{F}_2$ given by:

$$c(x, i) = \Phi_x(P(x))_i.$$

This can also be expressed as the Reed-Solomon code concatenated with a different trivial code $\Phi_x : \mathbb{F}_q \to \mathbb{F}_2^m$

---

[10] In this paper, "trivial code" will always refer to bijective $\mathbb{F}_2$-linear maps $\Phi : \mathbb{F}_{2^m} \to \mathbb{F}_2^m$ for some m. They are trivial because their absolute minimum distance equals 1. Note that for any given m, there are many different choices of trivial codes (corresponding to invertible $m \times m$ matrices over $\mathbb{F}_2$).

[11] The lower bound obtained on the distance of the resulting codes is always at most $\frac{5}{84} < 0.06$.

[12] The Weil bounds, which are also used to analyze the distance of dual-BCH codes.

[13] We remark that the constant hidden in the notation $\mathrm{poly}(n)$ here (and in the previous item) is universal and does not depend on α. However, the choice of which m is sufficiently large, does depend on α.

in each coordinate (in the spirit of Justesen [Jus72] and Thommesen [Tho83]).

A lot, but not all, of requirements for the code we desire are already satisfied by arbitrary Raw Reed-Solomon codes. We now pick out two special codes in this family, SimpleRawRS and OddRawRS which do satisfy all the requirements (and whose analysis will be more specialized).

- SimpleRawRS: Let Φ be an $\mathbb{F}_2$-linear bijection from $\mathbb{F}_q$ to $\mathbb{F}_2^m$. Let $D = \mathbb{F}_q \setminus \{0\}$. For each $x \in D$, define[14] $\Phi_x(y) = \Phi(xy)$ for all y, and take $\mathbf{\Phi} = (\Phi_x)_{x \in D}$. The Simple Raw Reed-Solomon code $\mathsf{SimpleRawRS}[\mathbb{F}_q, d, \Phi]$ is defined to be $\mathsf{RawRS}[\mathbb{F}_q, d, D, \mathbf{\Phi}]$.

  In this code, the codeword corresponding to polynomial $P(X)$ is obtained by writing down, for each $x \in D$, the m bits of $\Phi(xP(x))$. Observe that $XP(X)$ is a polynomial of degree at most $d+1$ with 0 constant term. Thus, the codewords are obtained by taking a polynomial of degree at most $d+1$ with 0 constant term and writing down all its values using Φ. This is the way these codes are described in the introduction.
- OddRawRS: Let Φ be an $\mathbb{F}_2$-linear bijection from $\mathbb{F}_q$ to $\mathbb{F}_2^m$. Let $D = \mathbb{F}_q \setminus \{0\}$. For each $x \in D$, define[15] $\Phi_x(y) = \Phi(xy^2)$ for all y, and take $\mathbf{\Phi} = (\Phi_x)_{x \in D}$. The Odd Raw Reed-Solomon code $\mathsf{OddRawRS}[\mathbb{F}_q, d, \Phi]$ is defined to be $\mathsf{RawRS}[\mathbb{F}_q, d, D, \mathbf{\Phi}]$.

  In this code, the codeword corresponding to polynomial $P(X)$ is obtained by writing down, for each $x \in D$, the m bits of $\Phi(xP(x)^2)$. Observe that $XP(X)^2$ is a polynomial of degree at most $2d+1$ with only odd degree monomials. Thus, the codewords are obtained by taking a polynomial of degree at most $2d+1$ with only odd degree monomials and writing down all its values using Φ. This is the way these codes are described in the introduction.

Our results for OddRawRS are technically simpler and quantitatively stronger, but SimpleRawRS is arguably a more natural code whose parameters are not far behind, so we feel it is interesting to see that too.

For contrast, it is also worth keeping in mind the following example:

- VerySimpleRawRS: Let Φ be an $\mathbb{F}_2$-linear bijection from $\mathbb{F}_q$ to $\mathbb{F}_2^m$. Let $D = \mathbb{F}_q \setminus \{0\}$. For each $x \in D$, define $\Phi_x = \Phi$, and take $\mathbf{\Phi} = (\Phi_x)_{x \in D}$. The Very Simple Raw Reed-Solomon code $\mathsf{VerySimpleRawRS}[\mathbb{F}_q, d, \Phi]$ is defined to be $\mathsf{RawRS}[\mathbb{F}_q, d, D, \mathbf{\Phi}]$.

  In this code, the codeword corresponding to polynomial $P(X)$ is obtained by writing down, for each $x \in D$, the m bits of $\Phi(P(x))$. Thus, the codewords are obtained by taking a polynomial of degree at most d and writing down all its values using Φ.

  Note that this code is the usual concatenation of Reed-Solomon codes with the trivial code given by the map Φ. Also note that this code is very closely related to SimpleRawRS: it is obtained by adding the constant functions to a suitable SimpleRawRS.

Our plan now is as follows. First we study some properties of all Raw Reed-Solomon codes, including the rate and dual-distance. Next we prove the list-decodability of all Raw Reed-Solomon codes: this is more sophisticated, but still works in full generality.

Finally, we give a specialized analysis to show that SimpleRawRS and OddRawRS have good distance (nearly $1/2$ for the setting of interest). This is in contrast to VerySimpleRawRS which has distance $O(1/\log n)$.

Lemma IV.2 (Easy Properties of Raw Reed-Solomon codes)**.** *Let* $\mathbb{F}_q, d, D, m, \mathbf{\Phi}$ *be as above, and let* $C = \mathsf{RawRS}[\mathbb{F}_q, d, D, \mathbf{\Phi}]$. *Then:*

1) *The block-length of* C *is* $m \cdot |D|$.
2) *The dimension of* C *is* $m \cdot (d+1)$.
3) C *has absolute dual distance at least* $(d+2)$.

Proof. The first two items are trivial.

To show the final item, we use an alternate characterization of the dual distance: a linear code has absolute dual distance at least b if and only if the uniform distribution on the codewords of the code is $(b-1)$-wise independent.

Let $P(X) \in \mathbb{F}_q[X]$ be a uniformly random polynomial of degree at most d. We need to show that for any set $S \subseteq D \times [m]$, the random variables $(\Phi_x(P(x))_i)_{(x,i) \in S}$ are independent.

To see this, first note that for

$$A = \{x \in D \mid \exists i \in [m] \text{ with } (x,i) \in S\},$$

we have that the random variables $(P(x))_{x \in A}$ are uniform and independent. This is because $|A| \leq |S| \leq d+1$, and the evaluations of uniformly random degree d polynomials are uniform and $(d+1)$-wise independent.

Next, we observe that for any x, setting $S_x = \{i \in [m] \mid (x,i) \in S\}$, the random variables $(\Phi(P(x))_i)_{i \in S_x}$ are independent. This is because $P(x)$ is uniformly distributed over $\mathbb{F}_q$, and since Φ is a bijection, the image

---

[14] Note that for $x \in D$, $y \mapsto xy$ is a linear bijection of $\mathbb{F}_q$.

[15] Note that $y \mapsto y^2$ is a linear bijection of $\mathbb{F}_q$.

under Φ of a uniformly random element of $\mathbb{F}_q$ is uniform on $\mathbb{F}_2^m$.

Combining these facts, we get the desired $(d+1)$-wise independence. $\square$

B. List-decoding algorithm

We now give a list-decoding algorithm for (general) Raw Reed-Solomon codes, which is interesting in the setting of polynomially small rate and where $|D| = \Omega(q)$.

Lemma IV.3 (Decodability of Raw Reed-Solomon codes)**.** *Let $\mathbb{F}_q, d, m, D, \mathbf{\Phi}$ be as above, and let $C =$ $\mathsf{SimpleRawRS}[\mathbb{F}_q, d, \mathbf{\Phi}]$. Let $\eta > 0$, and suppose*

$$d \leq \eta^2 \cdot \frac{q^{O(\eta^2)}}{q} \cdot |D|.$$

(If $|D| = \Omega(q)$, this is roughly the same as $d \leq \eta^2 \cdot |D|^{O(\eta^2)}$.) Let $n = m \cdot |D|$ be the block-length of C.

Then C is list-decodable from $1/2 - \eta$ fraction errors in time $\mathrm{poly}(n)$ with list size $O(n^2)$.

Proof. We use the natural 2-stage list-decoding strategy for concatenated codes. This will reduce our problem to list-recovery of Reed-Solomon codes, for which we have the following fundamental result.

Theorem IV.4 (List-recovery of Reed-Solomon codes [Sud97, GS99])**.** *Suppose we are given, for each $x \in D$, an "input list" $L_x \subseteq \mathbb{F}_q$ with $|L_x| \leq \ell$. Then we can find, in $\mathrm{poly}(q)$ time, the list of all polynomials $P(X)$ of degree at most d such that:*

$$\Pr_{x \in D}[P(x) \in L_x] \geq \alpha,$$

provided:

$$\alpha \geq \sqrt{\frac{d\ell}{|D|}}.$$

Furthermore, the output list size is at most $O(q^2)$.

Let $w : D \times m \to \mathbb{F}_2$ be a given received word. For $x \in D$, let $w(x) \in \mathbb{F}_2^m$ denote the vector whose ith coordinate is $w(x,i)$.

For each $x \in D$, we define the input-list $L_x \subseteq \mathbb{F}_q$ as follows:

$$L_x = \{u \in \mathbb{F}_q \mid \Delta(\Phi_x(u), w(x)) \leq 1/2 - \eta/2\}.$$

Then for any $c \in C$ with $\Delta(w,c) < 1/2 - \eta$, we have that:

$$\Pr_{x \in D}[c(x) \in L_x] \geq \eta/2.$$

This is because $\mathbf{E}_{x \in D}[\Delta(w(x), c(x))] = \Delta(w,c) < 1/2 - \eta$, and so by Markov's inequality,

$$\Pr_{x \in D}[\Delta(w(x), c(x)) < 1/2 - \eta/2] > \eta/2. \quad (1)$$

We have that each L_x has size

$$\begin{aligned}\ell &= Vol(\text{Ball of radius } (1/2 - \eta/2) \text{ in } \mathbb{F}_2^m) \\ &\leq 2^{(1-\Omega(\eta^2))m} = q^{1-\Omega(\eta^2)}.\end{aligned}$$

Thus we have that

$$\sqrt{\frac{d\ell}{|D|}} \leq \sqrt{\frac{dq^{1-\Omega(\eta^2)}}{|D|}} \leq \eta/2.$$

Thus the list-recovery algorithm of Theorem IV.4 will find all $P(X) \in \mathbb{F}_q[X]$ of degree at most d such that

$$\Pr_{x \in D}[P(x) \in L_x] \geq \eta/2.$$

By Equation (1), all the codewords we are interested in will be recovered by this procedure.

We summarize the algorithm below:

- Create, for each $x \in D$, an input list $L_x \subseteq \mathbb{F}_q$.
- Use the Reed-Solomon list-recovery algorithm to find all polynomials $P(X)$ for which $P(x) \in L_x$ for a noticeable fraction of $x \in D$.
- For each such polynomial $P(X)$, include the corresponding codeword $c : \mathbb{F}_q \times [m] \to \mathbb{F}_2$ in the output list.

$\square$

C. *Explicit* RawRS *codes with good distance*

Now we come to the most delicate part: the minimum distance.

In general, a Raw Reed-Solomon code could have minimum distance as small as $1/m = O(1/\log n)$. Indeed VerySimpleRawRS does have small distance. If we take some $\alpha \in \mathbb{F}_q$ for which $\Phi(\alpha)$ has absolute weight equal to 1, then the codeword of a VerySimpleRawRS code which corresponds to the constant polynomial α has minimum distance $O(1/\log n)$.

Nevertheless, the following results shows that OddRawRS and SimpleRawRS have good minimum distance. Our first result shows that OddRawRS has distance $1/2 - o(1)$ for $d = o(q^{1/2})$. Our second result shows that SimpleRawRS has distance about $\frac{1-\epsilon}{2}$ when $d < q^\epsilon$ for any $\epsilon < 1/2$ (and in particular the distance is $1/2 - o(1)$ for $d = q^{o(1)}$).

Lemma IV.5 (Distance of OddRawRS)**.** *Let $\mathbb{F}_q, d, m, \Phi$ be as above, and let $C =$ $\mathsf{OddRawRS}[\mathbb{F}_q, d, \Phi]$.*

Then C has minimum distance at least

$$\left(\frac{1}{2} - \frac{2d}{\sqrt{q}}\right).$$

Proof. The key ingredient in the proof is the Weil bound on additive character sums.

First we recall the field trace function $\mathrm{Tr} : \mathbb{F}_q \to \mathbb{F}_2$. This is an $\mathbb{F}_2$-linear function given by:

$$\mathrm{Tr}(x) = x + x^2 + x^4 + \ldots + x^{2^i} + \ldots + x^{2^{m-1}}.$$

Theorem IV.6 ([Wei48])**.** *Let* $\mathrm{Tr} : \mathbb{F}_q \to \mathbb{F}_2$ *denote the finite field trace. Let* $R(X) \in \mathbb{F}_q[X]$ *be a* nonzero *polynomial of degree at most d with only* odd degree *monomials. Then:*

$$\left|\sum_{x\in\mathbb{F}_q} (-1)^{\mathrm{Tr}(R(x))}\right| \leq (d-1)\sqrt{q}.$$

It says that for low degree polynomials R with only odd degree monomials, $\mathrm{Tr}(R(x))$ is approximately uniformly distributed over $\mathbb{F}_2$. The hypothesis about odd degree is needed to avoid pathological situations where $\mathrm{Tr}(R(x))$ is constant (for example, $\mathrm{Tr}(x + x^2) = 0$ for all x). The statement above can be found in [Sch06, Chapter II.2, Theorem 2E]. Elementary proofs were given by Stepanov, Schmidt and Bombieri (see [Sch06, Mor93, Kop10] for expositions).

We also need some simple facts about Tr.

- Every $\mathbb{F}_2$-linear function $g : \mathbb{F}_q \to \mathbb{F}_2$ is of the form $g(x) = \mathrm{Tr}(\beta x)$ for some $\beta \in \mathbb{F}_q$.
- $\mathrm{Tr}(y) = \mathrm{Tr}(y^2)$ for all $y \in \mathbb{F}_q$.

By the first fact above, there are $\beta_1, \ldots, \beta_m \in \mathbb{F}_q$ such that $\Phi : \mathbb{F}_q \to \mathbb{F}_2^m$ is given by:

$$\Phi(y) = (\mathrm{Tr}(\beta_1 y), \mathrm{Tr}(\beta_2 y), \ldots, \mathrm{Tr}(\beta_m y)).$$

Since Φ is injective, we get that $\beta_1, \ldots, \beta_m$ are linearly independent over $\mathbb{F}_2$, and thus are a basis for $\mathbb{F}_q$ over $\mathbb{F}_2$.

Let $c : \mathbb{F}_q \times [m] \to \mathbb{F}_2$ be a nonzero codeword. We break it into m functions $c_1, c_2, \ldots, c_m : \mathbb{F}_q \to \mathbb{F}_2$ given by:

$$c_i(x) = c(x, i).$$

It will turn out that each c_i is a nonzero codeword of a dual-BCH code.

Let $P(X)$ be the polynomial underlying c. We have $\deg(P(X)) \leq d$. Let $P(X) = \sum_{j=0}^{d} \gamma_j X^j$.

By definition of $\mathsf{OddRawRS}$, we have $\Phi_x(y)_i = \mathrm{Tr}(\beta_i x y^2) = \mathrm{Tr}(\beta_i x P(x)^2)$. The crucial point is that

$$\begin{aligned} c_i(x) = \Phi_x(P(x))_i &= \mathrm{Tr}(\beta_i x P(x)^2) \\ &= \mathrm{Tr}\left(\beta_i x \left(\sum_{j\leq d} \gamma_j^2 x^{2j}\right)\right) \\ &= \mathrm{Tr}\left(\beta_i \sum_{\ell\leq 2d+1, \ell \text{ odd}} \gamma_{(\ell-1)/2}^2 x^{\ell}\right) \\ &= \mathrm{Tr}(R_i(x)), \end{aligned}$$

where:

$$R_i(X) = \sum_{\ell\leq 2d+1, \ell \text{ odd}} \beta_i \gamma_{(\ell-1)/2}^2 X^{\ell}$$

is a *nonzero* polynomial of degree at most $2d+1$ with only odd degree monomials. This allows us to apply the Weil bound (Theorem IV.6) directly. It tells us that for *all* $i \in [m]$,

$$\left|\Pr_x[\mathrm{Tr}(R_i(x)) = 0] - \Pr_x[\mathrm{Tr}(R_i(x)) = 1]\right| \leq \frac{2d}{\sqrt{q}} + \frac{1}{q}.$$

So, using wt to denote the relative weight,

$$\mathsf{wt}(c_i) = \Pr_{x\in D}[\mathrm{Tr}(R_i(x)) \neq 0] \geq \left(1/2 - \frac{d}{\sqrt{q}} - \frac{1}{2q}\right).$$

Averaging over all i, we get that

$$\mathsf{wt}(c) = \mathbf{E}_{i\in[m]}[\mathsf{wt}(c_i)] \geq \left(1/2 - \frac{d}{\sqrt{q}} - \frac{1}{2q}\right).$$

Thus the minimum distance of C is at least that quantity, as desired. ☐

Lemma IV.7 (Distance of $\mathsf{SimpleRawRS}$)**.** *Let $\mathbb{F}_q, d, m, \Phi$ be as above, and let* $C = \mathsf{SimpleRawRS}[\mathbb{F}_q, d, \Phi]$.

Then C has minimum distance at least

$$\left(1 - \frac{\log(d+1)}{\log q}\right)\left(\frac{1}{2} - \frac{d}{\sqrt{q}}\right).$$

Proof. The proof is very similar to the previous one.

Again we have a basis $\beta_1, \ldots, \beta_m$ of $\mathbb{F}_q$ over $\mathbb{F}_2$ such that

$$\Phi(y) = (\mathrm{Tr}(\beta_1 y), \mathrm{Tr}(\beta_2 y), \ldots, \mathrm{Tr}(\beta_m y)).$$

Let c be a nonzero codeword. We define $c_1, \ldots, c_m : D \to \mathbb{F}_2$ as before:

$$c_i(x) = c(x, i).$$

Here again we will get that the c_i are codewords of the dual-BCH code. However, unlike the previous proof, here it might be the case that some c_i is identically 0. We will show that at most $\log_2(d+1)$ of these c_i are identically 0, and the remaining c_i have weight at least $\left(\frac{1}{2} - \frac{d}{\sqrt{q}}\right)$. This will imply that:

$$\begin{aligned} \mathsf{wt}(c) &\geq \mathbf{E}_{i\in[m]}[\mathsf{wt}(c_i)] \\ &\geq \left(1 - \frac{\log(d+1)}{\log q}\right)\left(\frac{1}{2} - \frac{d}{\sqrt{q}}\right), \end{aligned}$$

and thus the minimum distance of C is at least that quantity, completing the proof.

Let $P(X)$ be the polynomial underlying c. We have $\deg(P(X)) \leq d$. Let $P(X) = \sum_{j=0}^{d} \gamma_j X^j$. By construction, $c_i(x) = \mathrm{Tr}(\beta_i x P(x))$. The polynomial $\beta_i X P(X)$ may have monomials of even degree, and so we cannot directly apply Theorem IV.6 to it. Instead we will reduce the even degree monomials using the identity $\mathrm{Tr}(y^2) = \mathrm{Tr}(y)$, and then hope that the reduction does not leave us with the zero polynomial.

$$\begin{aligned} &\beta_i X P(X) = \beta_i \left(\sum_{j=1}^{d+1} \gamma_{j-1} X^j \right) \\ &= \sum_{j=1}^{d+1} \beta_i \gamma_{j-1} X^j \\ &= \sum_{\ell \leq d+1, \ell \text{ odd}} \left(\sum_{r\geq 0, 2^r\ell \leq d+1} \beta_i \gamma_{(\ell\cdot 2^r - 1)} X^{\ell\cdot 2^r} \right), \end{aligned}$$

where in the last equality, we grouped all the powers of X according to the largest odd factor of the exponent (for example, $X^3, X^6, X^{12}, X^{24}, \ldots$ are all in the same group).

Thus for every $x \in \mathbb{F}_q$, we have:

$$\begin{aligned} &\mathrm{Tr}(\beta_i x P(x)) \\ &= \sum_{\ell \leq d+1, \ell \text{ odd}} \left(\sum_{r\geq 0, 2^r\ell \leq d+1} \mathrm{Tr}(\beta_i \gamma_{(\ell\cdot 2^r - 1)} x^{\ell\cdot 2^r}) \right) \\ &= \sum_{\ell \leq d+1, \ell \text{ odd}} \left(\sum_{r\geq 0, 2^r\ell \leq d+1} \mathrm{Tr}\left(\left(\beta_i^{1/2^r} \gamma_{(\ell\cdot 2^r - 1)}^{1/2^r} x^{\ell} \right)^{2^r} \right) \right) \\ &= \sum_{\ell \leq d+1, \ell \text{ odd}} \left(\sum_{r\geq 0, 2^r\ell \leq d+1} \mathrm{Tr}\left(\beta_i^{1/2^r} \gamma_{(\ell\cdot 2^r - 1)}^{1/2^r} x^{\ell} \right) \right) \\ &\quad \text{Since } \mathrm{Tr}(y^{2^r}) = \mathrm{Tr}(y) \\ &= \mathrm{Tr}\left(\sum_{\ell \leq d+1, \ell \text{ odd}} \left(\sum_{r\geq 0, 2^r\ell \leq d+1} \beta_i^{1/2^r} \gamma_{(\ell\cdot 2^r - 1)}^{1/2^r} \right) x^{\ell} \right) \\ &= \mathrm{Tr}\left(\sum_{\ell \leq d+1, \ell \text{ odd}} E_\ell(\beta_i) x^{\ell} \right), \end{aligned}$$

where $E_\ell : \mathbb{F}_q \to \mathbb{F}_q$ is the function

$$E_\ell(y) = \sum_{r\geq 0, 2^r\ell \leq d+1\cdot} y^{1/2^r} \gamma_{(\ell\cdot 2^r - 1)}^{1/2^r}.$$

Let $R_i(X) \in \mathbb{F}_q[X]$ be given by:

$$R_i(X) = \sum_{\ell \leq d+1, \ell \text{ odd}} E_\ell(\beta_i) X^{\ell}.$$

Summarizing, we have

$$c_i(x) = \mathrm{Tr}(R_i(x)).$$

Since $P(X)$ is a nonzero polynomial, some coefficient $\gamma_{j_0} \neq 0$. Let $j_0 = \ell_0 \cdot 2^{r_0} - 1$, where ℓ_0 is odd. Observe that E_{ℓ_0} satisfies:

1) E_{ℓ_0} is $\mathbb{F}_2$-linear,
2) $E_{\ell_0}(y) = A(y^{1/2^\rho})$ for some nonnegative integer ρ and some polynomial $A(Z) \in \mathbb{F}_q[Z]$ of degree at most $d+1$,
3) $A(Z)$ has some power of γ_{j_0} as a coefficient of some monomial, and is thus a nonzero polynomial.

These three facts imply that E_{ℓ_0} can vanish on at most $\log_2(d+1)$ of the β_i. Indeed, since the β_i are linearly independent, if E_{ℓ_0} vanishes on t of them, then by linearity we get that E_{ℓ_0} vanishes on their span, which has 2^t points. However E_{ℓ_0} cannot have more than $d+1$ roots (since A has degree at most $d+1$). Thus $t \leq \log_2(d+1)$.

In particular, this means that at most $\log_2(d+1)$ of the $R_i(X)$ are identically zero.

Fix an i where $R_i(X)$ is not identically 0. We have that $R_i(X)$ is a nonzero polynomial of degree at most $d+1$ with only monomials of odd degree. Thus Theorem IV.6 applies. It tells us that

$$\left|\Pr_{x\in D}[\mathrm{Tr}(R_i(x))=0]-\Pr_{x\in D}[\mathrm{Tr}(R_i(x))=1]\right|$$
$$\leq \frac{d+1}{\sqrt{q}}+\frac{1}{q}.$$

Then we get:

$$\mathsf{wt}(c_i)=\Pr_{x\in D}[\mathrm{Tr}(R_i(x))\neq 0]\geq\left(\frac{1}{2}-\frac{d+1}{2\sqrt{q}}-\frac{1}{q}\right).$$

Since there are at least $\log_2 q-\log_2(d+1)$ such i, we get the desired claim about the weight of c. This completes the proof of the minimum distance of SimpleRawRS codes. □

D. Discussion

1) Consider (general) Raw Reed-Solomon codes with $D=\mathbb{F}_q$ and $d\leq q^{0.01}$. Somewhat surprisingly, even though these codes need not have $\Omega(1)$ distance, they all have polynomial list-size for list-decoding upto radius almost $1/2$.
   Indeed, since the linear bijections Φ_x are completely arbitrary, we can choose them so that some particular polynomial $P(X)$ has the property that $\Phi_x(P(x))$ has Hamming weight ≤ 1 for all $x\in D$. The codeword of C corresponding to $P(X)$ will have relative Hamming weight $1/m=\Theta(1/\log n)$. However, as the list-decodability implies, the underlying algebraic structure somehow forces that one cannot choose the $(\Phi_x)_{x\in\mathbb{F}_q}$ so that this happens for many other $P(X)$.
2) Thommesen [Tho83] showed that if we choose the entries of $\mathbf{\Phi}=(\Phi_x)_{x\in\mathbb{F}_q}$ independently and uniformly at random (i.e., each Φ_x is an independently chosen uniformly random $\mathbb{F}_2$-linear bijection from $\mathbb{F}_q$ to $\mathbb{F}_2^m$), then for all d the resulting Raw Reed-Solomon code $C=\mathsf{RawRS}(\mathbb{F}_q,d,D,\mathbf{\Phi})$, for arbitrary D, meets the Gilbert-Varshamov bound[16] with high probability. In particular, even for $d=\Omega(|D|)$ (when the rate is $\Omega(1)$), there are Raw Reed-Solomon codes that have distance $\Omega(1)$. It is easy to see that no Simple Raw Reed-Solomon code has this property.
   Finding an explicit such code seems like a deep and very interesting open question.
3) VerySimpleRawRS and SimpleRawRS are closely related, yet have very different minimum distances. The results about SimpleRawRS explain the structure of VerySimpleRawRS. VerySimpleRawRS is the space spanned by SimpleRawRS along with codewords corresponding to the constant polynomials. SimpleRawRS has good minimum distance, it is only the small dimensional space of constant polynomials that spoils the minimum distance. This also explains why VerySimpleRawRS has good list-decodability despite having bad distance.
4) An important fact underlying our analysis of the distance of SimpleRawRS and OddRawRS is that Raw Reed-Solomon codes are just a bunch of (correlated) dual-BCH codewords written together. We don't know how to decode dual-BCH codes efficiently, but if we take about $\log n$ dual-BCH codewords together, then this resulting code magically can be decoded, while still retaining the good distance and dual distance of dual-BCH codes.

E. Proof of Theorem IV.1

We can now put everything together and prove Theorem IV.1.

Proof. Take $q=2^m$, and let $d=q^\alpha$ be an even number. Let $\Phi:\mathbb{F}_q\to\mathbb{F}_2^m$ be an arbitrary $\mathbb{F}_2$-linear bijection. Take $C=\mathsf{OddRawRS}[\mathbb{F}_q,d,\Phi]$. Recall that this code is the Raw Reed-Solomon code with evaluation domain $D=\mathbb{F}_q\setminus\{0\}$ with certain special $\Phi_x:\mathbb{F}_q\to\mathbb{F}_2^m$.

We now specify a linear encoding map for C. We take the encoding map $\mathrm{Enc}:\mathbb{F}_2^k\to\mathbb{F}_2^n$ to be the one which partitions the k input bits into blocks of size m, interprets the ith block as specifying (in an $\mathbb{F}_2$-linear way) the coefficient of the X^i monomial in a polynomial $P(X)$, and outputs the codeword c of C corresponding to the polynomial $P(X)$. Then clearly $\mathrm{Enc}^{\mathrm{trunc}}$ is simply the encoding map of $C'=\mathsf{OddRawRS}[\mathbb{F}_q,d/2,\Phi]$.

Since OddRawRS is an instance of RawRS, we can apply Lemma IV.2. It gives us the following basic properties of C:

1) The blocklength n of C equals $m\cdot(2^m-1)$.
2) The dimension k of C equals $m\cdot(d+1)\geq m\cdot q^\alpha\geq m^{1-\alpha}\cdot n^\alpha$.
3) The absolute dual distance of C is at least $d+2\geq q^\alpha\geq\left(\frac{n}{\log n}\right)^\alpha$.

Apply the same lemma to C' tells us that the absolute dual distance of C' is at least $d/2+2\geq\Omega\left(\left(\frac{n}{\log n}\right)^\alpha\right)$.

---

[16] The Gilbert-Varshamov bound $R=1-H(\delta)$ is the best *known* rate for codes with relative distance δ. This result is not constructive - deterministically constructing codes that meet this bound is a central open question.

Next we invoke Lemma IV.5. This is the place where we use the specific structure of *Odd* Raw Reed-Solomon codes. We get that C has distance at least:

$$\delta = \frac{1}{2} - O\left(\frac{q^{\alpha}}{q}\right) \geq \frac{1}{2} - n^{-\Omega(1)}.$$

Next we invoke Lemma IV.3. Set $\eta = b\sqrt{\alpha}$ for some absolute constant b. Since $|D| = q - 1$, we get that

$$d = q^{\alpha} \leq \eta^2 q^{O(\eta^2)} \leq \eta^2 q^{O(\eta^2)} \cdot \frac{|D|}{q}.$$

Thus C can be list-decoded from $1/2 - \eta$ fraction errors in time $\mathrm{poly}(q) \leq \mathrm{poly}(n)$. As an immediate consequence, since $1/2 - \eta > \delta/2$, we get that C can be unique decoded from $\delta/2 > 1/4 - o(1) > p$ fraction errors in polynomial time: we simply run the list-decoder and find the unique (if any) element of the output list which is within distance p from the received word.

The list-size guaranteed by Lemma IV.3 only implies that the list-size is at most $\mathrm{poly}(n)$, which is weaker than what we want. However now we only seek a combinatorial bound on the list-size, and this follows immediately from the Johnson bound [Joh62], which states that binary codes with minimum distance $\geq 1/2 - o(1)$ have list-size at most $O(1/\epsilon^2)$ for list-decoding from $(1/2-\epsilon)$-fraction errors.

This completes the proof of the theorem. ☐

V. STOCHASTIC CONTROL CODES

In this section we consider a more stringent notion of stochastic codes that are decodable from t errors. We will also require that such codes have an additional "pseudorandom property", namely that for every message $m \in \{0,1\}^k$, $\mathrm{Enc}(m, U_d)$ is pseudorandom for small space ROBPs.

Definition V.1 (Pseudorandom stochastic Codes decodable from errors)**.** *Let k, n, d be parameters and let* $\mathrm{Enc} : \{0,1\}^k \times \{0,1\}^d \to \{0,1\}^n$ *be a function. We say that* Enc *is an encoding function for a stochastic code that is:*

- ***ϵ-pseudorandom** for a class $\mathcal{C}$ of functions from n bits to one bit, if for every $m \in \{0,1\}^k$,* $\mathrm{Enc}(m, U_d)$ *is ϵ-pseudorandom for $\mathcal{C}$.*
- ***decodable from** t **errors**, if $t \in [n]$, and there exists a function* $\mathrm{Dec} : \{0,1\}^n \to \{0,1\}^k$ *such that for every $m \in \{0,1\}^k$, $s \in \{0,1\}^d$, and $e \in \{0,1\}^n$ with Hamming weight at most t,* $\mathrm{Dec}(\mathrm{Enc}(m, s) \oplus e) = m$.
- ***L-list-decodable from t-errors**, if the function* Dec *is allowed to output a list of size at most L, and for every $m \in \{0,1\}^k$, $s \in \{0,1\}^d$, and $e \in \{0,1\}^n$ with Hamming weight at most t,* $\mathrm{Dec}(\mathrm{Enc}(m, s) \oplus e) \ni m$.

*A code has **encoding time** [resp. **decoding time**] $T(\cdot)$, if* Enc *[resp.* Dec*] can be computed in time $T(k + n + d)$. (Naturally, this makes sense only for a family of encoding and decoding functions with varying block length n, message length $k(n)$ and seed length $d(n)$).*

Remark V.2 (This notion is only interesting for pseudorandom codes)**.** *We remark that the notion of stochastic codes decodable (or list-decodable) from t-errors is not interesting by itself. This is because for any such code* $\mathrm{Enc}(m, s)$*, we can define a standard (not stochastic) code* $\mathrm{Enc}'(m) = \mathrm{Enc}(m, s')$ *for some fixed seed s', and this code will be decodable (or list decodable) from t-errors.*

This means that designing stochastic codes that are decodable from t errors is a harder task *than designing standard codes that are decodable from errors, and we don't gain (and in fact make our task more difficult) by allowing* Enc *to receive a seed.*

This notion of codes decodable from errors becomes interesting when it is coupled with the pseudorandomness requirement. Loosely speaking, one can think of such codes as "standard codes" with an additional pseudorandomness property.

Remark V.3 (The use of this notion in past work)**.** *Similar notions appear in [GS16, SS16]. Specifically, Guruswami and Smith [GS16] considered a notion similar to "list-decodable from errors" with the stronger requirement that the decoding function needs to produce the randomness s, in addition to the message m.*

Shaltiel and Silbak [SS16] referred to this stronger requirement as "strongly list-decodable" and to the weaker notion defined here as "weakly list-decodable". The fact that the weaker notion (in which decoding does not need to produce the randomness) suffices for the intended application of stochastic codes for bounded channels, was key in [SS16] (as this weaker codes are easier to construct). The same also holds for this paper, as the list-decoding algorithms that we construct will not be able to reproduce the randomness s used by the encoding.

The main result of this section (that is stated in the theorem below) is a construction of stochastic codes that are pseudorandom for small space ROBPs. We plan to use these codes to encode very short strings, and so, their rate is not that important to us. The construction uses the $\mathsf{OddRawRS}$ of Section IV.

The theorem below gives a construction of a stochastic code that will be used as a "control code" in the construction of Section VI. We will use this "control code" in the proof of Theorem I.4, which is our main construction for bounded channels.

Theorem V.4 (Stochastic control codes for space $n^{\Omega(1)}$, with list-decoding up to $\frac{1}{2}$)**.** *For every constant* $\beta > 0$ *there exists a constant* $0 < \alpha \leq 0.1$ *such that for every sufficiently large* m*, setting* $n = (2^m - 1) \cdot m$*,* $k = n^{\alpha}$*,* $d = n \log n$*, and* $s = \frac{n^{\alpha}}{\log^3 n}$*, there is a stochastic code* $\mathrm{Enc} : \{0,1\}^k \times \{0,1\}^d \to \{0,1\}^n$ *that is:*

- 2^{-s}*-pseudorandom for any-order space* s *ROBPs.*
- *For every constant* $p < 1/4$*,* Enc *is decodable from* pn *errors in time* $\mathrm{poly}(n)$*.*
- *For every constant* $\epsilon > \beta$*,* Enc *is* $O(1/\epsilon^2)$*-list decodable from* $(\frac{1}{2} - \epsilon) \cdot n$ *errors in time* $\mathrm{poly}(n)$*.*
- *There exists a constant* c*, such that* Enc *can be computed in time* n^c*. Furthermore, encoding* n^c *inputs takes "amortized time"* $O(n \cdot \log^2 n)$*, namely, for every* $(m_1, s_1), \ldots, (m_{n^c}, s_{n^c}) \in \{0,1\}^k \times \{0,1\}^d$*, computing* $(\mathrm{Enc}(m_i, S_i))_{(i \in [n^c])}$ *takes time* $n^c \cdot O(n \cdot \log^2 n)$*.*

In the construction proving Theorem I.4 we will apply the stochastic control code many times, and this is why we care about amortized encoding time (that can be made quasilinear) rather than the time of encoding one message.

We can also get a different tradeoff that gives pseudorandomness for larger space. However, this comes with a cost of decoding only from $p_0 \cdot n$ errors for some small constant $p_0 > 0$ (rather than a number of errors that approaches $\frac{1}{2} \cdot n$). The encoding algorithm for this code is also less efficient than the one in Theorem V.4, and we don't get encoding in amortized linear time. The theorem below will be used in the proof of Theorem I.5.

Theorem V.5 (Stochastic control codes for space $n/\mathrm{polylog}(n)$, that decode from few errors)**.** *There exist constants* $p_{\max} > 0$*, and* $R > 0$ *such that for every sufficiently large* m*, setting* $n = 8^m - 8$*,* $k = Rn$*,* $d = n \log n$*, and* $s = n/\log^2 n$ *there is a stochastic code* $\mathrm{Enc} : \{0,1\}^k \times \{0,1\}^d \to \{0,1\}^n$ *that is:*

- 2^{-s}*-pseudorandom for any-order space* s *ROBPs.*
- Enc *is decodable from* $p_{\max} \cdot n$ *errors in time* $\mathrm{poly}(n)$*.*
- Enc *is encodable in time* $\mathrm{poly}(n)$*.*

In the remainder of the section we prove Theorem V.4 and Theorem V.5. In Section V-A we revisit the idea of "bounded independence plus noise" of [HLV18, LV17, FK18]. We state and prove a quantitative variant of this approach that will be used in our proof. In Section V-B we use the method of "bounded independence plus noise" to transform linear codes with certain additional properties into stochastic control codes that are pseudorandom for ROBPs. Finally, in Section V-C we prove Theorem V.4 (by relying on the OddRawRS of Section IV) and Theorem V.5 (by relying on algebraic geometric codes of Garcia and Stichtenoth [GS96]).

A. Bounded independence plus low weight noise fools ROBPs

Haramaty, Lee and Viola [HLV18] showed that "k-wise independence plus low weight noise" is pseudorandom for certain classes of distinguishers. Specifically, they consider xoring a k-wise independent distribution D_k^n on n bits, with low weight noise, chosen according to BSC_η^n for some small $\eta > 0$. They show that this distribution, namely $D_k^n \oplus \mathrm{BSC}_\eta^n$ (where the two distributions are independent) is pseudorandom for small space ROBPs if k is sufficiently larger than $n^{2/3}$, and $\eta > 0$ is a positive constant (that can be arbitrarily small).

This result is specifically appealing as this gives a distribution that is pseudorandom for ROBPs, regardless of the order in which they read the n bits. However, it requires a very large seed length, of at least $n^{2/3}$, (even for generating the distribution D_k^n).

We are interested in constructing stochastic control codes (which combine requirements from coding theory and pseudorandomness) and will make use of the particular structure of the distribution of [HLV18] (in addition to their pseudorandomness properties). More specifically, the fact that w.h.p. BSC_η has low hamming weight, will be important in our intended application.

Unfortunately, for our intended application, taking $k > n^{2/3}$ is too large, and we need k to be smaller. Subsequent work [LV17, FK18], gives pseudorandom distribution in which k is much smaller, but they use a different distribution in which the "noise distribution" BSC_η is replaced by distributions which have large hamming weight of $\approx \frac{n}{4}$). Such large weight is not useful for our application.[17]

---

[17] Lee and Viola [LV17] and Forbes and Kelly [FK18] consider the noise distribution $T_k^n \wedge U_n$ (where T_k^n is a k-wise independent distribution). Their motivation is that using this approach, one can think of T_k^n as selecting approximately $n/2$ of the n indices, and then placing uniform bits on these $n/2$ indices. This view enables a recursive construction in which the $n/2$ uniform bits are replaced with pseudorandom bits, and this approach can yield pseudorandom generators with polylogarithmic seed [FK18].

In this paper we observe that the ideas and technique use by Forbes and Kelly [FK18] to reduce the amount of independence for "high weight noise", can also be applied in the case of noise with low hamming weight. This is stated precisely in the following theorem.

Theorem V.6 (Improved analysis for [HLV18])**.** *For every integers n, s and $\epsilon, \eta > 0$, the distribution $D_k^n \oplus \mathrm{BSC}_\eta^n$ where D_k^n is independent of BSC_η^n, and is k-wise independent over $\{0,1\}^n$, for $k = O(\frac{s+\log n+\log(1/\epsilon)}{\eta})$, is ϵ-pseudorandom for any-order space s ROBPs.*

Note that Theorem V.6 kicks in, whenever $k > \log n$ whereas the previous analysis by [HLV18] only kicked in if $k \geq n^{2/3}$.

Proof. Forbes and Kelly [FK18, Lemma 6.3] show that the distribution $D_{2k}^n \oplus (T_k^n \wedge U_n)$ where the three distributions are independent, and:

- D_{2k}^n is $2k$-wise independent over $\{0,1\}^n$.
- T_k^n is k-wise independent over $\{0,1\}^n$.

is ϵ-pseudorandom for width $w = 2^s$ ROBPs, with $\epsilon = \frac{nw}{2^{k/2}}$. We first note that $\mathrm{BSC}_\eta^n = \mathrm{BSC}_{2\eta}^n \wedge U_n$. Thus, if we change D_k^n to D_{2k}^n (which we can do because of the $O(\cdot)$ notation in our statement) we can think of our target distribution $D_{2k}^n \oplus \mathrm{BSC}_\eta^n$ as $D_{2k}^n \oplus (\mathrm{BSC}_{2\cdot\eta}^n \wedge U_n)$. Consequently, in order to prove our result, we need to show that the analysis of [FK18] can be carried out in case T_k^n is replaced by $\mathrm{BSC}_{2\eta}$. This is indeed the case, and the analysis gives $\epsilon = nw \cdot (1-2\eta)^{k/2}$ under this modification. This is sufficient to derive our result.

On an intuitive level, this follows because the distribution T_k^n is only used to argue that if $\alpha \in \{0,1\}^n$ is a "Fourier coefficient" with weight k, then $\Pr[T_k^n \wedge \alpha = 0] = (\frac{1}{2})^k$. The distribution $\mathrm{BSC}_{2\eta}^n$ (which plays the role of T_k^n in our case) gives the similar (though slightly weaker) bound of $\Pr[\mathrm{BSC}_{2\eta}^n \wedge \alpha = 0] = (1-2\eta)^k$, and this suffices for the argument.

More precisely, inspecting the proof of [FK18, Lemma 6.3], one can observe that the only place where a specific property of the distribution T_k^n is used is at the final equality in the proof of Lemma 6.2, and that replacing T_k^n with $\mathrm{BSC}_{2\eta}^n$ yields a version of Lemma 6.2. in which the term $(\frac{1}{2})^k$. is replaced by $(1-2\eta)^k$. Finally, Lemma 6.2 is used to derive the last inequality in Lemma 6.3. and substituting the modified quantity gives the final result. (The rest of the proof goes through unchanged). ☐

B. Linear codes with large dual distance yield pseudorandom stochastic codes

In this section we show that linear codes with large dual distance can be used to construct stochastic control codes.

In the definition below we define a function $\mathrm{BSC}_\eta^n(\cdot)$ which when given uniform input, generates the distribution BSC_η^n. We then consider a truncated version $\mathrm{BSC}_\eta^{n,\mathrm{trunc}}(\cdot)$ which evaluates to 0^n if the generated string has hamming weight larger than $2\eta \cdot n$. This is done to guarantee that with probability one, over a uniform input, $\mathrm{BSC}_\eta^{n,\mathrm{trunc}}$ produces a string with hamming weight $\leq 2\eta \cdot n$.

Definition V.7 (Generating and truncating BSC)**.** *For an integer k, $\eta = \frac{1}{2^k}$, and an integer n, we define the function $\mathrm{BSC}_\eta^n : \{0,1\}^{n\cdot\log(1/\eta)} \to \{0,1\}^n$ by $\mathrm{BSC}(s)_i = 1$ iff $s_{(i-1)\cdot\log(1/\eta)+1}, \ldots, s_{i\cdot\log(1/\eta)} = 1$, so that $\mathrm{BSC}_\eta^n(U_{n\cdot\log(1/\eta)})$ is the distribution BSC_η^n.*

The truncated version $\mathrm{BSC}_\eta^{n,\mathrm{trunc}} : \{0,1\}^{n\cdot\log(1/\eta)} \to \{0,1\}^n$ is defined as follows: Given $s \in \{0,1\}^{n\cdot\log(1/\eta)}$, if $\mathrm{BSC}_\eta^n(s)$ has hamming weight larger than $2\eta\cdot n$, then $\mathrm{BSC}_\eta^{n,\mathrm{trunc}}(s)$ is set to 0^n, and otherwise, it is set to $\mathrm{BSC}_\eta^n(s)$.

Note that by a multiplicative Chernoff bound, the statistical distance between $\mathrm{BSC}_\eta^n(U_{n\cdot\log(1/\eta)})$ and $\mathrm{BSC}_\eta^{n,\mathrm{trunc}}(U_{n\cdot\log(1/\eta)})$ is $2^{-\Omega(n)}$. (This will allow us to replace the former by the latter).

The next definition shows how to convert a linear code Enc into a stochastic code Enc_η (and we soon show that Enc_η is pseudorandom for any-order small space ROBPs).

Definition V.8 (stochastic control codes from linear codes)**.** *Given a function $\mathrm{Enc} : \{0,1\}^k \to \{0,1\}^n$, we define $\mathrm{Enc}^{\mathrm{trunc}} : \{0,1\}^{k/2} \to \{0,1\}^n$ by $\mathrm{Enc}^{\mathrm{trunc}}(x) = \mathrm{Enc}(x \circ 0^{k/2})$.*

Given a function $\mathrm{Enc} : \{0,1\}^k \to \{0,1\}^n$ and $\eta > 0$, we define $d = k/2 + n \cdot \log(1/\eta)$, and the function $\mathrm{Enc}_\eta : \{0,1\}^{k/2} \times \{0,1\}^d \to \{0,1\}^n$ as follows: Given inputs $m \in \{0,1\}^{k/2}$ and $s \in \{0,1\}^d$, we interpret s as a pair $s = (s_1, s_2)$ where $s_1 \in \{0,1\}^{k/2}$ and $s_2 \in \{0,1\}^{n\cdot\log(1/\eta)}$, and define:

$$\mathrm{Enc}_\eta(m,s) = \mathrm{Enc}(s_1 \circ m) \oplus \mathrm{BSC}_\eta^{n,\mathrm{trunc}}(s_2)$$

The following lemma shows that if Enc is a linear encoding map for a code with large dual distance, then Enc_η is a stochastic code which is pseudorandom, and inherits the decoding capabilities of Enc.

We plan to apply the stochastic code on many inputs, and are therefore interested in the *amortized* encoding time. The last item of the following lemma says that if k is sufficiently small compared to n, encoding Enc_η takes amortized *quasilinear time* even if one evaluation of Enc takes polynomial (but not necessarily quasilinear) time.

Lemma V.9. *Let* $\mathrm{Enc} : \mathbb{F}_2^k \to \mathbb{F}_2^n$ *be a linear function. Let* $\eta > 0$*, and* $d = k/2 + n \cdot \log(1/\eta)$.

- *If the linear code* $C' = \mathrm{Enc}^{\mathrm{trunc}}(\mathbb{F}_2^{k/2})$ *has dual distance* $r \geq \frac{10 \log n}{\eta}$*, then* $\mathrm{Enc}_\eta : \{0,1\}^{k/2} \times \{0,1\}^d$ *is* 2^{-s}*-pseudorandom for any-order, space* s *ROBPs, for* $s = \Omega(r \cdot \eta)$.
- *If* Enc *is decodable [resp. L-list decodable] from* pn *errors, then* Enc_η *is decodable [resp. L-list decodable] from* $(p - 2\eta) \cdot n$ *errors. Furthermore, the decoding time for* Enc_η *is the same as that of* Enc.
- *Computing* Enc_η *takes time* $O(n \cdot \log(1/\eta))$ *plus the time it takes to compute* Enc.[18]
- *If* $k \leq n^{0.1}$ *and* Enc *can be computed in time* n^c *for some constant c, then computing* Enc_η *on* n^c *pairs* $(m_1, s_1), \ldots, (m_{n^c}, s_{n^c}) \in \{0,1\}^{k/2} \times \{0,1\}^d$*, takes time* $n^c \cdot O(n \cdot (\log^2 n + \log(1/\eta)))$ *(that is* amortized time $O(n \cdot (\log^2 n + \log(1/\eta)))$*).*

Proof. We start with the first item. By Lemma III.15 we have that $\mathrm{Enc}^{trunc}(U_{k/2})$ is $(r-1)$-wise independent. Let $\mathrm{Enc}^{\mathrm{suf}} : \mathbb{F}_2^{k/2} \to \mathbb{F}_2^n$ be the function $\mathrm{Enc}^{\mathrm{suf}}(x) = \mathrm{Enc}(0^{k/2} \circ x)$. We first observe that using the linearity of Enc:

$$\mathrm{Enc}(s_1 \circ m) = \mathrm{Enc}^{\mathrm{trunc}}(s_1) \oplus \mathrm{Enc}^{\mathrm{suf}}(m)$$

It follows that for every $m \in \{0,1\}^{k/2}$, $\mathrm{Enc}(U_{k/2} \circ m) = \mathrm{Enc}^{\mathrm{trunc}}(U_{k/2}) \oplus \mathrm{Enc}^{\mathrm{suf}}(m)$ is also $(r-1)$-wise independent. Therefore,

$$\mathrm{Enc}_\eta(m, U_d) = \mathrm{Enc}(U_{k/2} \circ m) \oplus \mathrm{BSC}_\eta^{n,\mathrm{trunc}}(U_{n \cdot \log(1/\eta)}),$$

is $2^{-\Omega(n)}$-close to a distribution of the form $D_{r-1}^n \oplus \mathrm{BSC}_\eta^n$, for some $(r-1)$-wise independent distribution D_{r-1}^n. Therefore, $\mathrm{Enc}_\eta(m, U_d)$ is a distribution that is very close to "$(r-1)$-wise independent plus low weight noise". By Theorem V.6 with $\epsilon = 2^{-2s}$, $D_{r-1}^n \oplus \mathrm{BSC}_\eta^n$ is 2^{-2s}-pseudorandom for any-order ROBPs with space $\Omega(r \cdot \eta) - \log n - 2s \geq s$ for our choice of parameters.

The distribution $\mathrm{Enc}_\eta(m, U_d)$ is therefore ϵ'-pseudorrandom for any-order ROBPs with space s, for $\epsilon' = 2^{-2s} + 2^{-\Omega(n)} \leq 2^{-s}$ by noting that $s \leq r \leq n$ and choosing the constant hidden in the definition of s to be sufficiently small.

For the second item, note that if we encode a message m, by $\mathrm{Enc}_\eta(m, s) = \mathrm{Enc}(s_1 \circ m) \oplus \mathrm{BSC}_\eta^{n,\mathrm{trunc}}(s_2)$ and xor it with an error vector $e \in \{0,1\}^n$ is of hamming weight $(p - 2\eta) \cdot n$, then (as the hamming weight of $\mathrm{BSC}_\eta^{n,\mathrm{trunc}}(s_2)$ is at most $2\eta \cdot n$) we obtain a string that is within hamming distance $(p-2\eta) \cdot n + 2\eta \cdot n = pn$ from $\mathrm{Enc}(s_1 \circ m)$. Consequently, decoding (or list decoding) is guaranteed to decode (or list decode) the message $s_1 \circ m$ from which we can recover m.

The third item follows directly by construction.

For the fourth item, we note that by the previous item, we only need to show how given $x_1, \ldots, x_{n^c} \in \{0,1\}^k$, we can compute $\mathrm{Enc}(x_1), \ldots, \mathrm{Enc}(x_{n^c})$ in time $n^c \cdot O(n \cdot \log^2 n)$. Using the fact that Enc can be computed in time n^c, gives that in time n^{c+1} we can compute the $k \times n$ generator matrix G of Enc.

We will use fast matrix multiplication. Specifically, that multiplying an $n \times n^{0.1}$ matrix by an $n^{0.1} \times n$ matrix can be performed in time $O(n^2 \cdot \log^2 n)$ [Cop82] (See [Wil14] for more details on this algorithm). For $i \in [n^{c-1}]$, let $A^{(i)}$ be the $n \times k$ matrix in which the j'th row is $x_{(i-1)n+j}$. Note that the outputs $\mathrm{Enc}(x_1), \ldots, \mathrm{Enc}(x_{n^c})$ that we want to compute are the rows of the matrices $A^{(i)} \cdot G$ where i ranges over $[n^{c-1}]$. Each multiplication of $A^{(i)} \cdot G$ takes time $O(n^2 \cdot \log^2 n)$ and therefore n^{c-1} such computations take time $O(n^{c+1} \cdot \log^2 n)$. Overall, the entire computation takes time $O(n^{c+1} \cdot \log^2 n)$, which gives amortized $O(n \cdot \log^2 n)$ time. ☐

C. Proof of Theorem V.4 and Theorem V.5

In this section, we put everything together and prove Theorem V.4 and Theorem V.5.

Proof of Theorem V.4. We use the linear code of Theorem IV.1, and apply Lemma V.9, choosing $\eta = \frac{1}{2 \cdot \log n}$, so that $k/2 + n \log(1/\eta) \leq n \cdot \log n$. The properties in Theorem V.4 follow directly from Theorem IV.1 and Lemma V.9, and the fact that $\eta = o(n)$. ☐

The proof of Theorem V.5 follows in the same way, by using the following construction of error correcting codes, due to Garcia and Stichtenoth [GS96].

Theorem V.10. *[GS96] There exist constants* $p_{\max} > 0$*,* $\delta > 0$ *and* $R > 0$ *such that for every sufficiently large* m*, setting* $n = 8^m - 8$*,* $k = Rn$*, there is a binary linear* $[n, k]_2$*-code that satisfies:*

---

[18] We will use η which is only slightly smaller than constant, and so the term $\log(1/\eta)$ is immaterial. We could have been more careful and reduce the running time from $n \cdot \log(1/\eta)$ to expected running time $2n$ by observing that determining whether $\log(1/\eta)$ uniform bits are all one, can be done by querying only two of the bits (in expectation).

- *C has a linear encoding map* $\mathrm{Enc} : \mathbb{F}_2^k \to \mathbb{F}_2^n$ *that runs in time* $\mathrm{poly}(n)$*.*
- Enc *is decodable from* $p_{\max} \cdot n$ *errors in time* $\mathrm{poly}(n)$*.*
- *C has dual distance $\delta \cdot n$.*
- *Moreover, define* $\mathrm{Enc}^{\mathrm{trunc}} : \mathbb{F}_2^{k/2} \to \mathbb{F}_2^n$ *by* $\mathrm{Enc}^{\mathrm{trunc}}(x) = \mathrm{Enc}(x \circ 0^{k/2})$*, and consider the linear code* $C' = \mathrm{Enc}^{\mathrm{trunc}}(\mathbb{F}_2^{k/2})$*. It holds that C' has dual distance $\delta \cdot n$.*

The code of Garcia and Stichtenoth is not a binary code, but rather a code over constant size alphabet. The statement above is obtained by interpreting the code as a binary code. The reader is referred to the appendix of [Shp09] (which was written by Venkat Guruswami) for a precise description of this code and a proof of Theorem V.10.

VI. STOCHASTIC CODES FOR SPACE BOUNDED CHANNELS

In this section we state our main construction of stochastic codes for any-order bounded space channels. We start by restating Theorem I.4 more precisely. The statements below allow a wider range of parameters, and also give a more precise dependence of the parameters on each other.

Theorem VI.1. *There exists a universal constant c_0, such that for every constant $0 \leq p < \frac{1}{2}$, there exists a constant $\delta > 0$, such that for every constant $c_\nu \geq 1$, and every sufficiently small constant $\epsilon > 0$, there exists a constant* $L = \mathrm{poly}(1/\epsilon)$*, such that for infinitely many N, there is a stochastic code* $\mathrm{Enc} : \{0,1\}^{RN} \times \{0,1\}^{O(N \cdot \log N)} \to \{0,1\}^N$ *that satisfies the following properties:*

- Enc *has rate $R \geq 1 - H(p) - \epsilon$.*
- *There is a list-decoding algorithm* Dec *showing that* Enc *is L-list decodable for any-order space $s = N^\delta$ channels that induce at most pN errors, with probability $1 - \nu$, for $\nu = 2^{-(\log N)^{c_\nu}}$.*
- Enc *can be computed in time $N \cdot (\log N)^{c_0 \cdot c_\nu}$.*
- Dec *can be computed in time $N \cdot (\log N)^{c_0 \cdot c_\nu}$.*

Remark VI.2 (Dependence on constants)**.** *In this remark we give more details on the dependence of the parameters on the chosen constants.*

- *The list size achieved in the proof of Theorem VI.1 is $L = O(\frac{1}{\epsilon^4})$. However, if p is sufficiently smaller than $1/2$ (say $p < 0.49$) then the list size can be reduced to $L = O(\frac{1}{\epsilon^2})$ by a more careful argument, and if p is sufficiently smaller than $1/4$ (say $p < 0.24$) then it can be further reduced to $L = O(1/\epsilon)$.*
- *The running time of encoding and decoding depends on ϵ as follows: For every $\epsilon > 0$ there exists a constant $c_\epsilon = 2^{\mathrm{poly}(1/\epsilon)}$ such that the running time is $c_\epsilon \cdot N \cdot (\log N)^{c_0 \cdot c_\nu}$. This dependence is inherited from the use of explicit codes for binary symmetric channels [For65, GI05]. All other ingredients allow* $c_\epsilon = \mathrm{poly}(1/\epsilon)$*. Polar codes [GX15, HAU14] achieve running time $c_\epsilon n \log n$, where* $c_\epsilon = \mathrm{poly}(1/\epsilon)$*. It is plausible that using polar codes (and some other modifications) the dependence on ϵ can be reduced to* $\mathrm{poly}(1/\epsilon)$*.*
- *The theorem statement does not explicitly state which choices of infinitely many N are possible. Again, the reason that we don't get "for every sufficiently large N" is solely because linear time codes for binary symmetric channels [For65, GI05] are stated for infinitely many N. We remark that an inspection of these results reveals that (in the very least) there exists a universal polynomial $q(\cdot)$ such that for every $\epsilon > 0$, there is a constant c'_ϵ such that for every sufficiently large m, a suitable N can be found between $c'_\epsilon \cdot q(m)$ and $c'_\epsilon \cdot q(m+1)$. The same property is inherited by our construction. We remark that we are less picky and can allow quasi-linear time codes for binary symmetric channels, which can be constructed more easily using code concatenation and yield a denser family of N's.*

We can also achieve a different tradeoff where the channel has space $N/\mathrm{polylog}(N)$, for small p, in polynomial time. The following theorem is the more formal restatement of Theorem I.5.

Theorem VI.3. *There exist universal constants $p_{\max} > 0$ and c_1, such that for every constants $0 \leq p \leq p_{\max}$, and $c_\nu \geq 1$, and every sufficiently small constant $\epsilon > 0$, there exists a constant* $L = \mathrm{poly}(1/\epsilon)$*, such that for infinitely many N, there is a stochastic code* $\mathrm{Enc} : \{0,1\}^{RN} \times \{0,1\}^{O(N \cdot \log N)} \to \{0,1\}^N$ *that satisfies the following properties:*

- Enc *has rate $R \geq 1 - H(p) - \epsilon$.*
- *There is a list-decoding algorithm* Dec *showing that* Enc *is L-list decodable for any-order space $s = \frac{N}{(\log N)^{c_1 + c_\nu}}$ channels that induce at most pN error, with probability $1 - \nu$, for $\nu = 2^{-(\log N)^{c_\nu}}$.*
- Enc *can be computed in time* $\mathrm{poly}(N)$*.*
- Dec *can be computed in time* $\mathrm{poly}(N)$*.*

In Section VI-A we present our construction. The construction expects to receive a stochastic control codes with certain properties. In Section VI-B we plug in the specific control codes of Section V to obtain our main results. The correction of the construction is proven in the full version.

A. The construction

In this section we present our construction of stochastic codes for bounded channels. The construction is detailed in three figures. Figure 1 lists parameters and ingredients, Figure 2 which describes the encoding algorithm, and Figure 3 which describes the decoding algorithm. We start with some notation and definitions. We remark that an intuitive explanation of the construction appears in Section II-B.

a) Partitioning codewords into control blocks and data blocks: The construction will think of codewords $c \in \{0,1\}^N$ as being composed of $n = n_{\text{ctrl}} + n_{\text{data}}$ blocks of length $b = N/n$. Given a subset $I \subseteq [n]$ of n_{ctrl} distinct indices, we can decompose c into its data part $c_{\text{data}} \in \{0,1\}^{N_{\text{data}} = n_{\text{data}} \cdot b}$ and its control part $c_{\text{ctrl}} \in \{0,1\}^{N_{\text{ctrl}} = n_{\text{ctrl}} \cdot b}$. Similarly, given strings c_{data} and c_{ctrl} we can prepare the codeword c (which we denote by $(c_{\text{data}}, c_{\text{ctrl}})^I$ by the reverse operation. This is stated formally in the definition below.

Definition VI.4. *Let $I = \{i_1, \dots, i_{n_{\text{ctrl}}}\} \subseteq [n]$ be a subset of indices of size n_{ctrl}.*

- *Given strings $c_{\text{data}} \in \{0,1\}^{N_{\text{data}}}$ and $c_{\text{ctrl}} \in \{0,1\}^{N_{\text{ctrl}}}$ we define an N bit string c denoted by $(c_{\text{data}}, c_{\text{ctrl}})^I$ as follows: We think of $c_{\text{data}}, c_{\text{ctrl}}, c$ as being composed of blocks of length b (that is $c_{\text{data}} \in (\{0,1\}^b)^{n_{\text{data}}}$, $c_{\text{ctrl}} \in (\{0,1\}^b)^{n_{\text{ctrl}}}$ and $c \in (\{0,1\}^b)^n$). We enumerate the indices in $[n] \setminus I$ by $j_1, \dots, j_{n_{\text{data}}}$ and set $c_\ell = \begin{cases} (c_{\text{ctrl}})_k & \textit{if } \ell = i_k \textit{ for some } k; \\ (c_{\text{data}})_k & \textit{if } \ell = j_k \textit{ for some } k \end{cases}$*
- *Given a string $c \in \{0,1\}^N$ (which we think of as $c \in (\{0,1\}^b)^n$) we define strings $c^I_{\text{data}}, c^I_{\text{ctrl}}$ by $c^I_{\text{ctrl}} = c|_I$ and $c^I_{\text{data}} = c|_{[n]\setminus I}$, (namely the strings restricted to the indices in I, $[n] \setminus I$, respectively).*

We omit the superscript I when it is clear from the context.

Theorem VI.5 (correctness of the construction)**.** *There exists a universal constant c_0 such that for every constants $0 \le p < \frac{1}{2}$, $c_\nu \ge 1$, and every sufficiently small constant $\epsilon > 0$, for infinitely many N we have that: for every b, ℓ, s', and stochastic code* $\text{Enc}_{\text{ctrl}} : \{0,1\}^\ell \times \{0,1\}^d \to \{0,1\}^b$ *that satisfy the requirements in Figure 1. The encoding and decoding functions* $\text{Enc} : \{0,1\}^{Rn} \times \{0,1\}^{\ell + n_{\text{ctrl}} \cdot d} \to \{0,1\}^N$ *and* $\text{Dec} : \{0,1\}^N \to (\{0,1\}^{RN})^{\frac{100 \cdot L_{\text{ctrl}}}{\epsilon^2}}$ *specified in Figures 2 and 3 using the ingredients and parameter choices in Figure 1 satisfy the following properties:*

- Enc *has rate $R \ge 1 - H(p) - \epsilon$.*
- Dec *is a list-decoding algorithm showing that* Enc *is $O(\frac{L_{ctrl}}{\epsilon^2})$-list decodable for any-order space s channels, with probability $1 - \nu$, for $s = s' - s_A$, $s_A \le (\log N)^{2c_\nu + 3}$, and $\nu = 2^{-(\log N)^{c_\nu}}$.*
- Enc *can be computed in time $N \cdot (\log N)^{c_0 \cdot c_\nu} + T$, where T is a bound on the time it takes to perform the following task: Given n pairs $(m_1, s_1), \dots, (m_n, s_n) \in \{0,1\}^\ell \times \{0,1\}^d$ output* $\text{Enc}_{\text{ctrl}}(m_1, s_1), \dots, \text{Enc}_{\text{ctrl}}(m_n, s_n)$.
- Dec *can be computed in time $L_{\text{ctrl}} \cdot N \cdot (\log N)^{c_0 \cdot c_\nu} + n' \cdot T'$, where $n' = (\log N)^{c_\nu + 2}$ and T' is the running time of* Dec_{ctrl} *on input in $\{0,1\}^b$.*

We prove Theorem VI.5 in the full version. In the next section, we plug in our control codes from Section V to get specific constructions.

B. Deriving the main theorems

In this section we use specific stochastic control codes to derive our main results. We first use the control code of Theorem V.4 to prove Theorem VI.1.

Proof. (of Theorem VI.1) We want to choose parameters b, ℓ, s' and a control code to plug into Theorem VI.5. Our plan is to use Theorem V.4 as a control code. Let $\beta = \frac{1}{2} - p - 2\epsilon$ and note that β is a positive constant as $p < \frac{1}{2}$ is a constant, and $\epsilon > 0$ is sufficiently small. Given $\beta > 0$, Theorem V.4 provides a constant $0 < \alpha < 0.1$. The running time of encoding and decoding algorithms in Theorem V.4 is n^c for some universal constant c. Let $\lambda = \frac{1}{2(c+1)}$. We want to choose $b = N^\lambda$ and use it as a block length in Theorem V.4. However, we must verify that b is of the form $(2^m - 1) \cdot m$ in Theorem V.4, and so we choose a number b of this form such that $N^\lambda \le b \le N^{2\lambda}$ (and such a number exists). We apply Theorem V.4 to obtain a code $\text{Enc} : \{0,1\}^{b^\alpha} \times \{0,1\}^d \to \{0,1\}^b$ that is $2^{-\hat{s}}$-pseudorandom for any order space $\hat{s}$ ROBPs with $\hat{s} = \frac{b^\alpha}{\log^3 b}$. We also obtain that this code is $O(1/\hat{\epsilon}^2)$-list decodable from $(\frac{1}{2} - \hat{\epsilon})b$ errors in time b^c for every constant $\hat{\epsilon} > \beta$.

We choose $\ell = b^\alpha$ and $s' = \frac{\ell}{\log^3 N} \le \hat{s}$. It follows that Enc is $2^{-s'}$-pseudorandom for any-order space s' ROBPs. It also follows that Enc is L_{ctrl}-list decodable from $(p + \epsilon) \cdot b$ errors for $L_{\text{ctrl}} = O(\frac{1}{(\frac{1}{2} - (p+\epsilon))^2})$ in time b^c. This follows because $p + \epsilon < p + 2\epsilon = \frac{1}{2} - \beta$. This means that $\text{Enc}_{\text{ctrl}} : \{0,1\}^\ell \times \{0,1\}^d \to \{0,1\}^b$ satisfies the requirements from a control code in Figure 1.

Constants:

- $p > 0$ - The fraction of errors we need to recover from.
- $\epsilon > 0$ - The final code will have rate $R \geq 1 - H(p) - \epsilon$. We assume that $\epsilon > 0$ is sufficiently small compared to p.
- $c_\nu \geq 1$ - We are shooting for a code with success probability $1 - \nu$ for $\nu = 2^{-(\log N)^{c_\nu}}$.

Parameters that are allowed to vary with N:

- N - The length (in bits) of the codeword. Throughout, we assume that N is sufficiently large, and that other parameters are chosen as a function of N. Later choices will also restrict N to be a number of a special form.
- b - We divide the N output bits to $n := N/b$ blocks of length b.
- ℓ - This is the total length of a "control seed". Let $\ell' = \ell/3$. This will be the length of individual "seeds".
- s' - The code will work for any-order space s channels, for $s = s' - (\log N)^{2c_\nu+3}$.

Stochastic control code: The construction receives as a component a stochastic code $\mathrm{Enc}_{\mathrm{ctrl}} : \{0,1\}^\ell \times \{0,1\}^d \to \{0,1\}^b$ such that:

- $\mathrm{Enc}_{\mathrm{ctrl}}$ is $(2^{-s'})$-pseudorandom for any-order space s' ROBPs.
- $\mathrm{Enc}_{\mathrm{ctrl}}$ is L_{ctrl}-list-decodable from $(p+\epsilon) \cdot b$ errors (for some parameter L_{ctrl} to be chosen later) with decoding $\mathrm{Dec}_{\mathrm{ctrl}} : \{0,1\}^b \to (\{0,1\}^\ell)^{L_{\mathrm{ctrl}}}$.

Requirements: $(\log N)^{c_\nu+10} \leq b \leq \frac{N}{(\log N)^{c_\nu+10}}$, $s' \geq (\log N)^{c_\nu+3}$, and $\ell \geq s' \cdot (\log N)^3$.

Some more parameters:

- $\epsilon' = \epsilon/3$ will be the fraction of "control blocks", and set $n_{\mathrm{ctrl}} = \epsilon' \cdot n$, $n_{\mathrm{data}} = n - n_{\mathrm{ctrl}}$.
- Let $N_{\mathrm{ctrl}} = b \cdot n_{\mathrm{ctrl}}$ and $N_{\mathrm{data}} = b \cdot n_{\mathrm{data}}$. (Note that: $n = n_{\mathrm{ctrl}} + n_{\mathrm{data}}$, $N = N_{\mathrm{ctrl}} + N_{\mathrm{data}}$).

Other ingredients that are used:

- Let $\alpha = \frac{\epsilon}{10 \log \frac{1}{p}}$, $p_{\mathrm{BSC}} = p + \alpha$, and $R_{\mathrm{BSC}} = 1 - H(p_{\mathrm{BSC}}) - \epsilon/3$. We apply Theorem III.20 using $p_{\mathrm{BSC}}, \epsilon/3, N_{\mathrm{data}}$ as choices for p, ϵ, n, respectively. Theorem III.20 only guarantees the code for infinitely many block lengths, and so we require that $N_{\mathrm{data}} = (1 - \epsilon') \cdot N$ is one of these block lengths. This translates into a restriction on N (which is satisfied for infinitely many N). (See Remark VI.2 for a discussion on the "density" of the block lengths). We obtain an encoding function $\mathrm{Enc}_{\mathrm{BSC}} : \{0,1\}^{R_{\mathrm{BSC}} \cdot N_{\mathrm{data}}} \to \{0,1\}^{N_{\mathrm{data}}}$.
- Let $t = (\log N)^{c_\nu+2}$. We use the $(2^{-10 \cdot t}, t)$-wise permutation $\pi : \{0,1\}^{\ell'} \times [N_{\mathrm{data}}] \to [N_{\mathrm{data}}]$. By Theorem III.11 we have an explicit construction with seed length $O(t \cdot \log N) \leq \ell'$.
- We use Theorem III.9 to obtain an $(\frac{\alpha}{100}, \frac{\nu}{N^3})$-sampler with distinct samples $\mathrm{Samp} : \{0,1\}^{\ell'} \to [n]^{n_{\mathrm{ctrl}}}$. By Theorem III.9 we have an explicit construction with seed length $O(\log \frac{N^3}{\nu}) = O((\log N)^{c_\nu+1}) \leq \ell'$. We use n_{ctrl} samples, and indeed $n_{\mathrm{ctrl}} = \epsilon' \cdot n \geq \epsilon' \cdot (\log N)^{c_\nu+10} \gg \log \frac{N^3}{\nu}$ (as required in Theorem III.9).
- Let $G : \{0,1\}^{\ell'} \to \{0,1\}^{N_{\mathrm{data}}}$ be the $(2^{-s'})$-PRG for any-order space s' ROBPs, provided by Theorem III.7. We verify that for the constant c hidden in the statmement of Theorem III.7, for sufficiently large N, $\ell' \geq c \cdot s' \cdot (\log N_{\mathrm{data}})^2$.

Fig. 1: Parameters and ingredients for stochastic code

Furthermore, the requirements in Figure 1 are met by our choices of b, ℓ and s'. Specifically:

$$(\log N)^{c_\nu+10} \leq N^\lambda \leq b \leq N^{2\lambda} \leq \frac{N}{(\log N)^{c_\nu+10}},$$

$$s' = \frac{\ell}{\log^3 N} = \frac{b^\alpha}{\log^3 N} \geq \frac{N^{\lambda \cdot \alpha}}{\log^3 N} \geq (\log N)^{c_\nu+3},$$

and we chose $s' = \frac{\ell}{\log^3 N}$ so that the requirement $\ell \geq s' \cdot (\log N)^3$ is met. It follows that we meet all the conditions of Theorem VI.5 and obtain that:

- Enc has rate $R \geq 1 - H(p) - \epsilon$.
- Let $L = O(\frac{L_{\mathrm{ctrl}}}{\epsilon^2}) = O(\frac{1}{(\frac{1}{2} - (p+\epsilon))^2 \cdot \epsilon^2}) = O(\frac{1}{\epsilon^4})$. There is a list-decoding algorithm Dec showing that Enc is L-list decodable for any-order space $s = s' - s_A \geq N^{\frac{\alpha \cdot \lambda}{2}}$ channels, with probability $1 - \nu$. In other words, we can have $\delta = \frac{\alpha \cdot \lambda}{2}$.
- Enc can be computed in time $N \cdot (\log N)^{c_0 \cdot c_\nu} + T$ where T is the time it takes to perform $n = N/b$ encodings of Enc_{ctrl}. By Theorem V.4 there exists a constant c such that performing b^c encodings takes time $O(b^{c+1} \cdot \log^2 b)$. We can break the $n = \frac{N}{b} \geq b^c$ encodings into n/b^c groups of size b^c. Each such group takes time $O(b^{c+1} \log^2 b)$ and so,

$$\begin{aligned} T =& O(\frac{n \cdot b^{c+1} \cdot \log^2 b}{b^c}) \\ =& O(n \cdot b \cdot \log^2 b) \leq O(N \cdot \log^2 N). \end{aligned}$$

- As ϵ is constant, Dec can be computed in time $N \cdot (\log N)^{c_0 \cdot c_\nu} + n' \cdot T'$ where $n' = (\log N)^{c_\nu+2}$ and $T' \leq b^c \leq N^{2\lambda c} \leq N$.

This completes the proof of Theorem VI.1 □

Input:

- A message $m \in \{0,1\}^{R_{\mathrm{BSC}} \cdot N_{\mathrm{data}}}$. (This gives $R = \frac{R_{\mathrm{BSC}} \cdot N_{\mathrm{data}}}{N}$).
- A "random part" for the stochastic encoding that consists of: a string $s = (s_{\mathrm{samp}}, s_\pi, s_{\mathrm{PRG}})$ where $s_{\mathrm{samp}}, s_\pi, s_{\mathrm{PRG}} \in \{0,1\}^{\ell'}$ so that $s \in \{0,1\}^{\ell}$, and $r_1, \dots, r_{n_{\mathrm{ctrl}}} \in \{0,1\}^d$.

Output: A codeword $c = \mathrm{Enc}(m; (s, r_1, \dots, r_{n_{\mathrm{ctrl}}}))$ of length N.

Operation:

Determine control blocks: Apply $\mathrm{Samp}(s_{\mathrm{samp}})$ to generate $I = \{i_1, \dots, i_{n_{\mathrm{ctrl}}}\} \subseteq [n]$. These blocks will be called "control blocks", and the remaining n_{data} blocks will be called "data blocks".

Prepare data part: We prepare a string c_{data} of length N_{data} as follows:

- Encode m by $x = \mathrm{Enc}_{\mathrm{BSC}}(m)$.
- Generate an N_{data} bit string y by reordering the N_{data} bits of the encoding using the (inverse of) the permutation $\pi_{s_\pi}(\cdot) = \pi(s_\pi, \cdot)$. More precisely, $y = \pi_{s_\pi}^{-1}(x) = \pi_{s_\pi}^{-1}(\mathrm{Enc}_{\mathrm{BSC}}(m))$.
- Mask y using PRG. That is, $c_{\mathrm{data}} = y \oplus G(s_{\mathrm{PRG}}) = \pi_{s_\pi}^{-1}(\mathrm{Enc}_{\mathrm{BSC}}(m)) \oplus G(s_{\mathrm{PRG}})$.

Prepare control part: We prepare a string c_{ctrl} of length N_{ctrl} (which we view as n_{ctrl} blocks of length b) as follows:

- $(c_{\mathrm{ctrl}})_j = \mathrm{Enc}_{\mathrm{ctrl}}(s, r_j)$.

Merge data and control parts: We prepare the final output codeword $c \in \{0,1\}^N$ by merging c_{data} and c_{ctrl}. That is, $c = (c_{\mathrm{data}}, c_{\mathrm{ctrl}})^I$.

Fig. 2: Encoding algorithm for stochastic code

Input: A "received word" $c' \in \{0,1\}^N$.

Output: A list of messages $m \in \{0,1\}^{RN}$, where the list is of size at most $\frac{100 \cdot L_{\mathrm{ctrl}}}{\epsilon^2}$.

Operation:

Determine candidates for control information:

Decode control code: Generate $n' = (\log N)^{c_\nu + 2}$ lists of size L_{ctrl} as follows: choose uniformly distributed and independent n' indices $i_1, \dots, i_{n'} \in [n]$, and for every $j \in [n']$ apply the list decoding algorithm, $\mathrm{Dec}_{\mathrm{ctrl}}$ on c'_{i_j} (here, c'_i is the i'th block of c'). This gives a size L_{ctrl} list, $\mathrm{List}_i = \mathrm{Dec}_{\mathrm{ctrl}}(c'_{i_j})$.

Prune list of candidates: Let $\mathrm{List}_{\mathrm{ctrl}}$ consist of all $s \in \{0,1\}^{\ell}$ such that $s \in \mathrm{List}_i$ for at least $\frac{\epsilon^2 \cdot n'}{100}$ of $i \in [n']$. Note that $\mathrm{List}_{\mathrm{ctrl}}$ is of size at most $\frac{100 \cdot L_{\mathrm{ctrl}}}{\epsilon^2}$.

Use each control candidate to decode data: For each $s = (s_{\mathrm{samp}}, s_\pi, s_{\mathrm{PRG}}) \in \mathrm{List}_{\mathrm{ctrl}}$ we produce a candidate messages $m_s \in \{0,1\}^{RN}$.

Determine control blocks: Apply $\mathrm{Samp}(s_{\mathrm{samp}})$ to generate $I = \{i_1, \dots, i_{n_{\mathrm{ctrl}}}\}$. Compute $c'_{\mathrm{data}} = (c')^I_{\mathrm{data}}$.

Unmask PRG: Compute $y'_{\mathrm{data}} = c'_{\mathrm{data}} \oplus G(s_{\mathrm{PRG}})$.

Reverse permutation: Let x' be the N_{data} bit string obtained by "undoing" the permutation. More precisely, let $\pi_{s_\pi}(\cdot) = \pi(s_\pi, \cdot)$, and let $x' = \pi_{s_\pi}(y'_{\mathrm{data}}) = \pi_{s_\pi}(c'_{\mathrm{data}} \oplus G(s_{\mathrm{PRG}}))$.

Decode data: Compute $m_s = \mathrm{Dec}_{\mathrm{BSC}}(x')$.

Prepare output list: The final output is $\mathrm{List} = \{m_s : s \in \mathrm{List}_{\mathrm{ctrl}}\}$.

Fig. 3: List-decoding algorithm for stochastic code

We use the control code of Theorem V.5 to prove Theorem VI.3.

Proof. (of Theorem VI.3) We want to choose parameters b, ℓ, s' and a control code to plug into Theorem VI.5. Our plan is to use Theorem V.5 as a control code. We want to choose $b = \frac{N}{(\log N)^{c_1 + c_\nu}}$ and use it as a block length in Theorem V.5. However, we must verify that b is of the form $8^m - 8$ in Theorem V.4, and so we choose a number b of this form such that $\frac{N}{(\log N)^{c_1 + c_\nu}} \leq b \leq \frac{N}{(\log N)^{c_1 + c_\nu - 1}}$ (and such a number exists). We apply Theorem V.5 to obtain a code $\mathrm{Enc} : \{0,1\}^{R \cdot b} \times \{0,1\}^d \to \{0,1\}^b$ that is $2^{-\hat{s}}$-pseudorandom for any order space $\hat{s}$ ROBPs with $\hat{s} = \frac{b}{\log^2 b}$. We also obtain that this code is decodable from $p_{max} \cdot b$ errors, where $p_{\max} > 0$ is the constant from Theorem V.5.

We choose $\ell = R \cdot b$ and $s' = \frac{\ell}{\log^3 N} \leq \hat{s}$. It follows that Enc is $2^{-s'}$-pseudorandom for any-order space s' ROBPs. This means that $\mathrm{Enc}_{\mathrm{ctrl}} : \{0,1\}^{\ell} \times \{0,1\}^d \to \{0,1\}^b$ satisfies the requirements from a control code in Figure 1.

Furthermore, the requirements in Figure 1 are met by our choices of b, ℓ and s'. Specifically, for a sufficiently large constant c_1:

$$
\begin{aligned}
(\log N)^{c_\nu+10} &\leq \frac{N}{(\log N)^{c_1+c_\nu}} \leq b \\
&\leq \frac{N}{(\log N)^{c_1+c_\nu-1}} \leq \frac{N}{(\log N)^{c_\nu+10}},
\end{aligned}
$$

$$
\begin{aligned}
s' &= \frac{\ell}{\log^3 N} = \frac{R \cdot b}{\log^3 N} \\
&\geq \frac{N}{(\log N)^{c_1+c_\nu} \cdot \log^3 N} \geq (\log N)^{c_\nu+3},
\end{aligned}
$$

and we chose $s' = \frac{\ell}{\log^3 N}$ so that the requirement $\ell \geq s' \cdot (\log N)^3$ is met. It follows that we meet all the conditions of Theorem VI.5 and obtain that:

- Enc has rate $R \geq 1 - H(p) - \epsilon$.
- Let $L = O(\frac{L_{\text{ctrl}}}{\epsilon^2}) = O(\frac{1}{\epsilon^2})$. There is a list-decoding algorithm Dec showing that Enc is L-list decodable for any-order space $s = s' - s_A \geq \frac{N}{(\log N)^{c_1+c_\nu} \cdot \log^3 N}$ channels, with probability $1 - \nu$.
- Enc can be computed in time $\text{poly}(N)$.
- As ϵ is constant, Dec can be computed in time $\text{poly}(N)$.

This completes the proof of Theorem VI.3 $\square$

VII. CONCLUSION AND OPEN PROBLEMS

A natural open problem is to improve the running time of encoding and decoding to linear time. We remark that the step of applying a permutation $\pi[n] \rightarrow [n]$ on all n inputs, takes at least time $O(n \log n)$ (just to write down the inputs and outputs) and this is an obvious bottleneck for the approach used in this paper.

Guruswami and Smith [GS16] showed that we cannot expect to have uniquely decodable stochastic codes for space $\log n$ channels if $p > \frac{1}{4}$. However, it is not known whether for $p < \frac{1}{4}$, uniquely decodable stochastic codes for bounded space channels are possible with rate $R > 1 - H(2p)$ that is larger than the Gilbert-Varshamov bound (or even with a rate that matches the Gilbert-Varshamov bound, and efficient encoding and decoding).

ACKNOWLEDGEMENT

We are very grateful to Benny Applebaum, Michael Forbes, William Hoza, Shachar Lovett, Or Meir, Noga Ron-Zewi, Chris Umans and Emanuele Viola for helpful conversations and advice.

We are very grateful to the organizers of Dagstuhl seminar 18391 on "Algebraic Methods in Computational Complexity" held on September 23-28, 2018. Some progress on this project was inspired by talks and personal conversations at the seminar.

We are very grateful to anonymous referees for a detailed and very helpful comments.

Swastik Kopparty was supported in part by NSF grants CCF-1253886, CCF-1540634, CCF-1814409 and CCF-1412958, and BSF grant 2014359. Some of this research was done while visiting the Institute for Advanced Study.

Ronen Shaltiel was supported by ISF grant 1628/17.

Jad Silbak was supported by an ERC starting grant 638121 and by ISF grant 1628/17. Some of this research was done while Jad Silbak was a student at the University of Haifa.

REFERENCES

[BR94] Mihir Bellare and John Rompel. Randomness-efficient oblivious sampling. In *35th Annual Symposium on Foundations of Computer Science*, pages 276–287, 1994.

[Cop82] Don Coppersmith. Rapid multiplication of rectangular matrices. *SIAM J. Comput.*, 11(3):467–471, 1982.

[DHRS07] Y. Ding, D. Harnik, A. Rosen, and R. Shaltiel. Constant-round oblivious transfer in the bounded storage model. *J. Cryptology*, 20(2):165–202, 2007.

[DPW18] Stefan Dziembowski, Krzysztof Pietrzak, and Daniel Wichs. Non-malleable codes. *J. ACM*, 65(4):20:1–20:32, 2018.

[FK18] Michael A. Forbes and Zander Kelley. Pseudorandom generators for read-once branching programs, in any order. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 946–955, 2018.

[For65] G David Forney. *Concatenated codes.* PhD thesis, Massachusetts Institute of Technology, 1965.

[GI05] Venkatesan Guruswami and Piotr Indyk. Linear-time encodable/decodable codes with near-optimal rate. *IEEE Transactions on Information Theory*, 51(10):3393–3400, 2005.

[Gol97] Oded Goldreich. A sample of samplers - a computational perspective on sampling (survey). *Electronic Colloquium on Computational Complexity (ECCC)*, 4(20), 1997.

[GS96] A. Garcia and H. Stichtenoth. On the asymptotic behavior of some towers of function fields over finite fields. *Journal of Number Theory*, 61(2):248–273, 1996.

[GS99] V. Guruswami and M. Sudan. Improved decoding of reed-solomon and algebraic-geometry codes. *IEEE Transactions on Information Theory*, 45(6):1757–1767, 1999.

[GS16] Venkatesan Guruswami and Adam Smith. Optimal rate code constructions for computationally simple channels. *Journal of the ACM (JACM)*, 63(4):35, 2016.

[GX15] Venkatesan Guruswami and Patrick Xia. Polar codes: Speed of polarization and polynomial gap to capacity. *IEEE Transactions on Information Theory*, 61(1):3–16, 2015.

[HAU14] Seyed Hamed Hassani, Kasra Alishahi, and Rüdiger L Urbanke. Finite-length scaling for polar codes. *IEEE Transactions on Information Theory*, 60(10):5875–5898, 2014.

[HL11] Ishay Haviv and Michael Langberg. Beating the gilbert-varshamov bound for online channels. In *2011 IEEE International Symposium on Information Theory Proceedings, ISIT 2011, St. Petersburg, Russia, July 31 - August 5, 2011*, pages 1392–1396, 2011.

[HLV18] Elad Haramaty, Chin Ho Lee, and Emanuele Viola. Bounded independence plus noise fools products. *SIAM J. Comput.*, 47(2):493–523, 2018.

[Joh62] Selmer Johnson. A new upper bound for error-correcting codes. *IEEE Transactions on Information Theory*, 8(3):203–207, 1962.

[Jus72] Jørn Justesen. Class of constructive asymptotically good algebraic codes. *IEEE Transactions on Information Theory*, 18(5):652–656, 1972.

[KNR09] E. Kaplan, M. Naor, and O. Reingold. Derandomized constructions of k-wise (almost) independent permutations. *Algorithmica*, 55(1):113–133, 2009.

[Kop10] Swastik Kopparty. *Algebraic methods in randomness and pseudorandomness*. PhD thesis, Massachusetts Institute of Technology, 2010.

[KS07] Tali Kaufman and Madhu Sudan. Sparse random linear codes are locally decodable and testable. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), October 20-23, 2007, Providence, RI, USA, Proceedings*, pages 590–600. IEEE Computer Society, 2007.

[KS13] Swastik Kopparty and Shubhangi Saraf. Local list-decoding and testing of random linear codes from high error. *SIAM Journal on Computing*, 42(3):1302–1326, 2013.

[Lan04] Michael Langberg. Private codes or succinct random codes that are (almost) perfect. In *45th Symposium on Foundations of Computer Science (FOCS 2004)*, pages 325–334, 2004.

[Lip94] Richard J. Lipton. A new approach to information theory. In *11th Annual Symposium on Theoretical Aspects of Computer Science*, pages 699–708, 1994.

[LV17] Chin Ho Lee and Emanuele Viola. More on bounded independence plus noise: Pseudorandom generators for read-once polynomials. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:167, 2017.

[Mor93] Carlos Moreno. *Algebraic curves over finite fields*. Number 97 in Cambridge Tracts in Mathematics Series. Cambridge University Press, 1993.

[MPSW10] Silvio Micali, Chris Peikert, Madhu Sudan, and David A. Wilson. Optimal error correction for computationally bounded noise. *IEEE Trans. Information Theory*, 56(11):5673–5680, 2010.

[MS77] Florence Jessie MacWilliams and Neil James Alexander Sloane. *The theory of error-correcting codes*, volume 16. Elsevier, 1977.

[Sch06] Wolfgang M Schmidt. *Equations over finite fields: an elementary approach*, volume 536. Springer, 2006.

[Shp09] Amir Shpilka. Constructions of low-degree and error-correcting epsilon-biased generators. *Computational Complexity*, 18(4):495–525, 2009.

[Smi07] Adam D. Smith. Scrambling adversarial errors using few random bits, optimal information reconciliation, and better private codes. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 395–404, 2007.

[SS16] Ronen Shaltiel and Jad Silbak. Explicit list-decodable codes with optimal rate for computationally bounded channels. In *APPROX/RANDOM*, pages 45:1–45:38, 2016.

[SSS95] Jeanette P. Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff-hoeffding bounds for applications with limited independence. *SIAM J. Discrete Math.*, 8(2):223–250, 1995.

[Sud97] M. Sudan. Decoding of Reed Solomon codes beyond the error-correction bound. *Journal of Complexity*, 13, 1997.

[Tho83] Christian Thommesen. The existence of binary linear concatenated codes with reed - solomon outer codes which asymptotically meet the gilbert- varshamov bound. *IEEE Trans. Information Theory*, 29(6):850–853, 1983.

[Vad04] Salil P. Vadhan. Constructing locally computable extractors and cryptosystems in the bounded-storage model. *J. Cryptology*, 17(1):43–77, 2004.

[Wei48] André Weil. On some exponential sums. *Proceedings of the National Academy of Sciences of the United States of America*, 34(5):204, 1948.

[Wil14] Ryan Williams. Faster all-pairs shortest paths via circuit complexity. In *Symposium on Theory of Computing, STOC*, pages 664–673, 2014.