

LDPC Codes Achieve List Decoding Capacity*

Jonathan Mosheiff¹, Nicolas Resch¹, Noga Ron-Zewi², Shashwat Silas³, and Mary Wootters³

¹*Carnegie Mellon University*

²*University of Haifa*

³*Stanford University*

July 25, 2020

Abstract

We show that Gallager’s ensemble of Low-Density Parity Check (LDPC) codes achieves list-decoding capacity with high probability. These are the first graph-based codes shown to have this property. This result opens up a potential avenue towards truly linear-time list-decodable codes that achieve list-decoding capacity.

Our result on list decoding follows from a much more general result: any *local* property satisfied with high probability by a random linear code is also satisfied with high probability by a random LDPC code from Gallager’s distribution. Local properties are properties characterized by the exclusion of small sets of codewords, and include list-decoding, list-recovery and average-radius list-decoding.

In order to prove our results on LDPC codes, we establish sharp thresholds for when local properties are satisfied by a random linear code. More precisely, we show that for any local property \mathcal{P} , there is some R^* so that random linear codes of rate slightly less than R^* satisfy \mathcal{P} with high probability, while random linear codes of rate slightly more than R^* with high probability do not. We also give a characterization of the threshold rate R^* .

*JM is partially supported by NSF grants CCF-1814603 and CCF-1563742. A significant portion of this work was accomplished while JM was a postdoctoral fellow at the Weizmann Institute, partially supported by Irit Dinur’s ERC-CoG grant 772839. NRe is partially supported by NSERC grant CGSD2-502898, NSF grants CCF-1422045, CCF-1814603, CCF-1527110, CCF-1618280, CCF-1910588, NSF CAREER award CCF-1750808 and a Sloan Research Fellowship. NRo is partially supported by BSF grant 2014359 and ISF grant 735/20. SS and MW are partially supported by NSF grants CCF-1844628, CCF-1814629, and a Sloan Research Fellowship. SS is partially supported by a Google Graduate Fellowship.

1 Introduction

In this paper, we study sets $C \subset \Sigma^n$ of strings of length n , with the combinatorial property that not too many elements of C are contained in any small enough Hamming ball. In the language of coding theory, such a C is a *list-decodable code*. List-decoding is an important primitive in coding theory, with applications ranging from communication to complexity theory. However, as discussed below, most constructions of *capacity-achieving* (aka, optimal) list-decodable codes are fundamentally algebraic, despite a rich history of combinatorial—and in particular, graph-based—constructions of error correcting codes.

We show that a random ensemble of *Low-Density Parity-Check (LDPC) codes* achieves list-decoding capacity with high probability. LDPC codes are the prototypical example of graph-based codes, and are popular both in theory and in practice because of their extremely efficient algorithms. One of the motivations for this work is that we do not currently know any linear-time algorithms for list-decoding any code up to capacity; since graph-based codes offer linear-time algorithms for a variety of other coding-theoretic tasks, our result opens up the possibility of using these constructions for linear-time list-decoding algorithms.

List Decoding. Formally, a code $C \subset \Sigma^n$ is (α, L) -list-decodable if for all $z \in \Sigma^n$,

$$|\{c \in C : \text{dist}(c, z) \leq \alpha\}| \leq L.$$

Above, $\text{dist}(c, z)$ is the relative Hamming distance,

$$\text{dist}(c, z) = \frac{1}{n} |\{i : c_i \neq z_i\}|.$$

Elements $c \in C$ are called codewords, Σ is called the alphabet, and n is called the length of the code.

The fundamental trade-off in list-decoding is between the parameter α and the size $|C|$ of the code, given that the list size L is reasonably small. We would like both α and $|C|$ to be large, but these requirements are at odds: the larger the code C is, the closer together the codewords have to be, which means that α cannot be as large before some Hamming ball of radius α has many codewords in it. The size of a code C is traditionally quantified by the rate R of C , which is defined as

$$R = \frac{\log_{|\Sigma|}(|C|)}{n}.$$

The rate of C is a number between 0 and 1, and larger rates are better.

List-decoding has been studied since the work of Elias and Wozencraft in the 1950's [Eli57, Woz58], and by now we have a good understanding of what is possible and what is not. The classical *list-decoding capacity theorem* states that there exist codes over alphabets of size $|\Sigma| = q$ and of rate $R \geq 1 - h_q(\alpha) - \varepsilon$ which are $(\alpha, 1/\varepsilon)$ -list-decodable, where

$$h_q(x) := x \log_q(q-1) - x \log_q(x) - (1-x) \log_q(1-x) \tag{1}$$

is the q -ary entropy function. Conversely, any such code with rate $R \geq 1 - h_q(\alpha) + \varepsilon$ must have exponential list sizes, in the sense that there is some $z \in \Sigma^n$ so that $|\{c \in C : \text{dist}(c, z) \leq \alpha\}| = \exp_{\varepsilon, \alpha}(n)$.¹

¹Here and throughout the paper, $\exp(n)$ denotes $2^{\Theta(n)}$, and subscripts indicate that we are suppressing the dependence on those parameters.

A code of rate $R \geq 1 - h_q(\alpha) - \varepsilon$ that is (α, L) -list decodable for $L = O_{\varepsilon, \alpha}(1)$ is said to achieve list-decoding capacity, and a major question in list-decoding is which codes have this property. By now we have three classes of examples. First, it is not hard to see that completely random codes achieve list-decoding capacity with high probability. Second, a long line of work (discussed more below) has established that *random linear codes* do as well: we say that a code over the alphabet $\Sigma = \mathbb{F}_q$ is linear if it is a linear subspace of \mathbb{F}_q^n ,² and a random linear code is a random subspace. Third, there are several explicit constructions of codes which achieve list-decoding capacity; as discussed below, most of these constructions rely importantly on algebraic techniques.

LDPC Codes. Graph-based codes, such as LDPC codes, are a class of codes which is notably absent from the list of capacity-achieving codes above. Originally introduced by Gallager in the 1960's [Gal62], codes defined from graphs have become a class of central importance in the past 30 years.

Here is one way to define a code using a graph. Suppose that $G = (V, W, E)$ is a bipartite graph with $|V| = n$ and $|W| = m$ for $m \leq n$. Then G naturally defines a linear code $C \subset \mathbb{F}_q^n$ of rate at least $1 - m/n$ as follows:

$$C = \left\{ c \in \mathbb{F}_q^n : \forall j \in W, \sum_{i \in \Gamma(j)} \alpha_{i,j} c_i = 0 \right\},$$

where $\Gamma(i)$ denotes the neighbors of i in G and $\alpha_{i,j} \in \mathbb{F}_q$ are fixed coefficients. (See Figure 1). That is, each vertex in W serves as a **parity check**, and the code is defined as all possible labelings of vertices in V which obey all of the parity checks. When the right-degree of G is small, the resulting code is called a Low-Density Parity Check (LDPC) code.

LDPC codes and related constructions (in particular, Tanner codes [Tan81] and expander codes [SS94, Zém01]) are notable for their efficient algorithms for unique decoding; in fact, the only linear-time encoding/decoding algorithms we have for unique decoding (that is, list-decoding with $L = 1$) are based on such codes.

Motivating question. We currently do not know of any linear-time algorithms to list-decode any code to capacity. Since graph-based codes and LDPC codes in particular are notable for their linear-time algorithms, this state of affairs motivates the following question:

Question 1.1. *Are there (families) of LDPC codes that achieve list-decoding capacity?*

1.1 Contributions

Motivated by Question 1.1, our contributions are as follows.

- (1) We show that the answer to Question 1.1 is “yes.” More precisely, we show that random LDPC codes (the same ensemble studied by Gallager in his seminal work nearly 60 years ago [Gal62]), achieve list-decoding capacity with high probability.

²Here and throughout the paper, \mathbb{F}_q denotes the finite field with q elements.

- (2) In fact, we show a stronger result: random LDPC codes satisfy with high probability any *local* property that random linear codes satisfy with high probability. We define local properties precisely below; informally, a local property is one defined by the exclusion of certain bad sets. List-decodability is a local property—it can be defined by the exclusion of any big set of vectors that are too close together—and this answers Question 1.1.
- (3) Along the way, we develop a characterization of the local properties that are satisfied with high probability by a random linear code. We show that for any local property \mathcal{P} , there is a threshold R^* so that random linear codes of rate slightly less than R^* satisfy \mathcal{P} with high probability, while random linear codes of rate slightly greater than R^* with high probability do not. Moreover, we give a characterization of the threshold R^* .

In [GLM⁺20], the above characterization is used to compute lower bounds on the list-decoding and list-recovery parameters of random linear codes. This additional application does not directly relate to LDPC codes.

We describe each of these contributions in more detail below.

(1) Random LDPC codes achieve list-decoding capacity. We study the so-called “Gallager ensemble” of binary LDPC codes introduced by Gallager in the 1960’s [Gal62], as well as its natural generalization to larger alphabets.³

Fix a rate $R \in (0, 1)$ and a sparsity parameter s , and let $t = (1 - R)s$. We assume that t is an integer. To define the ensemble of random s -LDPC codes of rate R , we need to specify a distribution on the underlying bipartite graphs and a distribution on the coefficients $\alpha_{i,j}$. We define the distribution on graphs as follows. Let $G_i = (V, W_i, E_i)$ for $i = 1, \dots, t$ be independent uniformly random $(1, s)$ -regular bipartite graphs with a shared left vertex set V of size n and disjoint right vertex sets W_i , each of size n/s . Then let $G = (V, W, E)$ be the union of these graphs, where $W = \bigcup_{i=1}^t W_i$. Finally, we choose the coefficients $\alpha_{i,j}$ for $(i, j) \in E$ to be uniformly random in \mathbb{F}_q^* . The ensemble of random s -LDPC codes of rate R is illustrated in Figure 1.

Our main theorem about the list-decodability of random LDPC codes is a reduction from the list-decodability of random linear codes:

Theorem 1.2. *For any $R \in (0, 1)$, $\varepsilon > 0$, prime power q , $\alpha \in (0, 1 - 1/q)$ and $L \geq 1$ there exists $s_0 \geq 1$ such that the following holds for any odd $s \geq s_0$. Suppose that a random linear code of rate R over \mathbb{F}_q is (α, L) -list decodable with high probability. Then a random s -LDPC code of rate $R - \varepsilon$ over \mathbb{F}_q is (α, L) -list decodable with high probability.*

Remark 1.3 (The parity of s). *All of our results hold for even s as well as odd s . However, the proof is slightly simpler for odd s , so for clarity we state and prove the theorem in this case.*

Instantiating this with a result of [GHK11] on list decoding of random linear codes, we get the following corollary.

Corollary 1.4. *For any prime power q , $\alpha \in (0, 1 - 1/q)$, and $\varepsilon \in (0, 1 - h_q(\alpha))$ there exists $L = O_\alpha(1/\varepsilon)$ and $s \geq 1$ so that a random s -LDPC code of rate $1 - h_q(\alpha) - \varepsilon$ over \mathbb{F}_q is (α, L) -list-decodable with high probability.*

³For binary codes, our definition coincides with Gallager’s. For larger alphabets our definition is somewhat different: Gallager’s ensemble chooses the coefficients $\alpha_{i,j}$ to be all ones, while we choose them to be random elements of \mathbb{F}_q^* .




Figure 1: A random (t, s) -regular bipartite graph that gives rise to a random s -LDPC code of rate R . Here, we set $t := s(1 - R)$.

Remark 1.5 (Other parameter regimes). We state Corollary 1.4 as one example of what can be obtained by combining Theorem 1.2 with one result on random linear codes. The result of [GHK11] degrades as $\alpha \rightarrow 1 - 1/q$, and so Corollary 1.4 degrades as well. However, there has been a great deal of work on the list-decodability of random linear codes as $\alpha \rightarrow 1 - 1/q$ (summarized in Section 1.2 below), and Theorem 1.2 implies that these results carry over to random LDPC codes as well.

(2) Random LDPC codes achieve any local property that random linear codes achieve. Theorem 1.2 follows as a corollary of a much more general theorem. We show that any “local” property that is satisfied by random linear codes with high probability is also satisfied by random LDPC codes with high probability.

Informally, a local property is a property which can be defined by the exclusion of certain bad sets. For example, a code C is (α, L) -list-decodable if it does *not* contain any sets $B \subseteq \Sigma^n$ of size larger than L so that B is contained in a Hamming ball of radius α . Along with list-decodability, local properties include many related notions like *list recovery*, *average-radius list decoding*, and *erasure list decoding*. A long line of work (discussed more in Section 1.2) has established that these properties hold for random linear codes with high probability, so our reduction immediately implies that they hold with high probability for LDPC codes as well.

Formally, we define a local property as follows. Let $\pi : [n] \rightarrow [n]$ be a permutation on $[n]$. For a string $x \in \Sigma^n$, we let $\pi(x) \in \Sigma^n$ denote the string obtained by permuting the coordinates of x according to π , and for a subset $B \subseteq \Sigma^n$, we let $\pi(B) := \{\pi(x) \mid x \in B\}$. We say that a collection \mathcal{B} of subsets of Σ^n is permutation invariant if for any $B \in \mathcal{B}$ and permutation $\pi : [n] \rightarrow [n]$, we also have that $\pi(B) \in \mathcal{B}$.

Definition 1.6 (Local property). Let $\mathcal{P} = \{P_n\}_{n \in \mathbb{N}}$, where each P_n is a property of length n codes over Σ . We say that \mathcal{P} is a b -local property if for any $n \in \mathbb{N}$ there exists a permutation-invariant collection \mathcal{B}_n of subsets of Σ^n , where $|B| \leq b$ for all $B \in \mathcal{B}_n$, such that

$$C \subseteq \Sigma^n \text{ satisfies } P_n \iff B \not\subseteq C \text{ for all } B \in \mathcal{B}_n.$$

We say that a family of random codes $C = \{C_{n_i}\}_{i \in \mathbb{N}}$ (where $\{n_i\}$ is an increasing sequence) satisfies \mathcal{P} with high probability if $\lim_{i \rightarrow \infty} \Pr[C_{n_i} \text{ satisfies } P_{n_i}] = 1$. Similarly, we say that C almost surely does not satisfy \mathcal{P} if $\lim_{i \rightarrow \infty} \Pr[C_{n_i} \text{ satisfies } P_{n_i}] = 0$.

A code property is **monotone decreasing** if given a code C satisfying P , it holds that every code $C' \subseteq C$ also satisfies P . Note that every local property is monotone decreasing.

A random linear code of rate R over \mathbb{F}_q is defined⁴ as the kernel of a uniformly random matrix $H \in \mathbb{F}_q^{(1-R)n \times n}$. Notice that such a code has rate R with high probability.

For any $n \in \mathbb{N}$ and $R \in [0, 1]$ such that $R \cdot n \in \mathbb{N}$, we denote a random linear length n code of rate R by $C_{\text{RLC}}^n(R)$. Likewise, given s, n and R such that $s \mid n$ and $R \cdot s \in \mathbb{N}$, we denote a random s -LDPC code of length n and rate R by $C_{s\text{LDPC}}^n(R)$. Whenever we use these notations, it is implicitly assumed that the relevant divisibility conditions are satisfied.

Let $\mathcal{P} = \{P_n\}_{n \in \mathbb{N}}$ be a monotone decreasing property of linear codes. We define

$$R_{\text{RLC}}^n(\mathcal{P}) := \begin{cases} \sup \{R \in [0, 1] : \Pr[C_{\text{RLC}}^n(R) \text{ satisfies } P_n] \geq 1/2\} & \text{if there is such an } R \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Remark 1.7. *If \mathcal{P} is a monotone decreasing property then the function $\Pr[C_{\text{RLC}}^n(R) \text{ satisfies } P_n]$ is monotone decreasing in R . This can be proved by a standard coupling argument, akin to [Bol01, Thm. 2.1].*

With the notation out of the way, we are ready to state our more general theorem about random LDPC codes. Essentially, this theorem says that every local property that holds with high probability for a random linear code also holds with high probability for a random s -LDPC code of approximately the same rate. This approximation improves as s grows.

Theorem 1.8 (Main). *Let $\mathcal{P} = (P_n)_{n \in \mathbb{N}}$ be a b -local property with $\bar{R} := \limsup_{n \rightarrow \infty} R_{\text{RLC}}^n(\mathcal{P}) < 1$. For any $\varepsilon > 0$ and prime power q , there exists $s_0 = s_0(\varepsilon, \bar{R}, q, b) \geq 1$ such that for any odd $s \geq s_0$ and any sequence $\{R_n\}_{n \in \mathbb{N}}$, if $R_n \leq R_{\text{RLC}}^n(\mathcal{P}) - \varepsilon$ for all n , then the code ensemble $C_{s\text{LDPC}}^n(R_n)$ satisfies \mathcal{P} with high probability.*

Remark 1.9 (The dependence on $\varepsilon, \bar{R}, q, b$). *An inspection of the proof (see Remarks 4.2 and 5.3) shows that we may take*

$$s_0 = O\left(\frac{b \log(q) + \log(q/\varepsilon)}{h_q^{-1}(1 - \bar{R})}\right).$$

The existence of a reduction like the one in Theorem 1.8 is surprising, at least to the authors. There is a lot more structure in a random LDPC code than in a random linear code. For example, we know of linear-time unique decoding algorithms for random LDPC codes,⁵ but it is unlikely that any efficient unique decoding algorithm exists for random linear codes.⁶ Thus it is unexpected that this much more structured ensemble would share many properties—in a black-box way—with random linear codes.

⁴There are a few natural ways to define a random linear code: for example we could also define it as a uniformly random subspace of dimension Rn , or we could define it as the image of a uniformly random $n \times Rn$ matrix, or we could define it as we do here, as the kernel of a uniformly random $(1 - R)n \times n$ matrix. It can be shown that these distributions are quite close to each other, and in particular, any property that holds for one with high probability holds for the others.

⁵This follows, for example, from [SS94] because the underlying random graph is with high probability a good expander.

⁶Unique decoding of random linear codes is related to the problem of Learning Noisy Parities (LNP) and Learning With Errors (LWE), which are thought to be hard.

Remark 1.10 (A converse to Theorem 1.8?). *One may be tempted to conjecture that the converse of Theorem 1.8 holds as well. Namely, in the setting of Theorem 1.8, if $R_{n_i} \geq R_{\text{RLC}}^n(\mathcal{P}) + \varepsilon$ for all i , then the code ensemble $C_{\text{sLDPC}}(R_n)$ almost surely does not satisfy \mathcal{P} . However, this turns out to be false, due to the following example. Assume that $q = 2$ and consider the 1-local property $\mathcal{P} := (P_n)_{n \in \mathbb{N}}$, where P_n is the set of all length n linear codes that only contain even weight codewords. It is not hard to see (e.g., using Theorem 2.8) that $R_{\text{RLC}}^n(\mathcal{P})$ tends to 0 as $n \rightarrow \infty$. On the other hand, if $\frac{n}{s}$ is even, then every s -LDPC code (including, say, a code of rate $\frac{1}{2}$) satisfies \mathcal{P} , contradicting this conjecture.*

However, the above counter-example relies on a technicality involving divisibility criteria. It is an interesting question whether a natural converse of Theorem 1.8 holds if we additionally assume that \mathcal{P} belongs to some natural class of “nicely behaved” properties that precludes counter-examples of this sort.

(3) A characterization of local properties satisfied by random linear codes. In order to prove Theorems 1.2 and 1.8, we develop a new characterization of the local properties satisfied by a random linear code. Our formal theorem is given as Theorem 2.8. Informally, this theorem implies that for any monotone decreasing property \mathcal{P} , there is a sharp threshold R^* so that random linear codes of rate slightly less than R^* with high probability satisfy \mathcal{P} , while random linear codes of rate slightly larger than R^* with high probability do not. Moreover, we give a characterization of R^* .

Formally, we have the following definition, recalling the definition of $R_{\text{RLC}}^n(R_n)$ from (2).

Definition 1.11 (Sharpness for random linear codes). *We say that the property \mathcal{P} is sharp for random linear codes if for every $\varepsilon > 0$ there holds:*

- *If $R_n \leq R_{\text{RLC}}^n(\mathcal{P}) - \varepsilon$ for large enough n , then the code ensemble $C_{\text{RLC}}^n(R_n)$ ($n \in \mathbb{N}$) satisfies \mathcal{P} with high probability.*
- *If $R_n \geq R_{\text{RLC}}^n(\mathcal{P}) + \varepsilon$ for large enough n , then the code ensemble $C_{\text{RLC}}^n(R_n)$ ($n \in \mathbb{N}$) almost surely does not satisfy \mathcal{P} .*

If a property \mathcal{P} is sharp, we sometimes refer to $R_{\text{RLC}}^n(\mathcal{P})$ as the *threshold* for \mathcal{P} .

Theorem 2.8 has two corollaries. The first is that local properties are sharp for random linear codes:

Corollary 1.12. *Every local property is sharp for random linear codes.*

The second corollary of Theorem 2.8 is a characterization of $R_{\text{RLC}}^n(\mathcal{P})$. This characterization requires some definitions to state formally, so we defer the formal statement to Theorem 2.8. However, it has an intuitive interpretation, which we sketch here.

Recall that a local property is defined by a permutation-invariant collection \mathcal{B}_n of excluded sets. For simplicity of exposition, suppose that all of the sets $B \in \mathcal{B}_n$ have size exactly b , and moreover that they all have dimension exactly b . (This assumption is helpful for exposition but not necessary for our analysis). In this case, it is easy to compute the probability that each individual set $B \in \mathcal{B}_n$ is contained in $C_{\text{RLC}}(R)$ (see Fact 2.2):

$$\Pr [B \subseteq C_{\text{RLC}}(R)] = q^{-(1-R)nb}.$$

Thus, we have

$$\mathbb{E}|\{B \in \mathcal{B}_n : B \subseteq C_{\text{RLC}}(R)\}| = |\mathcal{B}_n| \cdot q^{-(1-R)nb}.$$

Thus, as long as

$$R < R_{\text{RLC}}^{\mathbb{E}}(\mathcal{B}_n) := 1 - \frac{\log |\mathcal{B}_n|}{nb},$$

we are guaranteed by Markov's inequality that with high probability, no elements of \mathcal{B}_n appear in $C_{\text{RLC}}(R)$. However, what if $R > R_{\text{RLC}}^{\mathbb{E}}(\mathcal{B}_n)$? It turns out that the statement above is not tight: in some cases it is likely that no elements of \mathcal{B}_n appear in $C_{\text{RLC}}(R)$ even if the rate R is significantly larger than $R_{\text{RLC}}^{\mathbb{E}}(\mathcal{B}_n)$. We give an example in Example 2.5 of when this can occur.

Our result in Theorem 2.8 pins down exactly when this can occur. Informally, it happens only because some projection \mathcal{B}'_n of the collection \mathcal{B}_n is more favorable than one might expect, in the sense that $R_{\text{RLC}}^{\mathbb{E}}(\mathcal{B}'_n)$ is larger than one might expect. In this case, the “correct” threshold is precisely $R_{\text{RLC}}^{\mathbb{E}}(\mathcal{B}'_n)$.

Thus, Theorem 2.8 also provides a characterization of which sorts of “bad” lists B (up to a permutation of the coordinates) are contained in a random linear code of a particular rate. We hope that this characterization will be useful in the study of random linear codes themselves, in addition to random LDPC codes.

The full power of Theorem 2.8 (including the characterization of $R_{\text{RLC}}^n(\mathcal{P})$ described above) is used to prove Theorem 1.8. However, given Theorem 1.8, Theorem 1.2 readily follows from Corollary 1.12 itself:

Proof of Theorem 1.2. Let \mathcal{P} denote the property of being (α, L) -list-decodable. Note that \mathcal{P} is a local property: for any $n \in \mathbb{N}$, take \mathcal{B}_n to be the collection of all sets of $L+1$ vectors in \mathbb{F}_q^n contained in some Hamming ball of radius α . Now, fix some $R \in (0, 1)$ and assume that a random linear code of rate R satisfies \mathcal{P} with high probability. Corollary 1.12 implies that $R_{\text{RLC}}^n(\mathcal{P}) \leq R + o_{n \rightarrow \infty}(1)$.

Next, it is not hard to verify that $\limsup_{n \rightarrow \infty} R_{\text{RLC}}^n(\mathcal{P}) \leq 1 - h_q(\alpha) < 1$. Indeed, it follows from the list-decoding capacity theorem (e.g. [LW18, Thm 1.1]) that for large enough n there are no (α, L) -list-decodable codes of rate $1 - h_q(\alpha) + \varepsilon$. In particular, this means that a random linear code of rate $1 - h_q(\alpha) + \varepsilon$ almost surely does not satisfy \mathcal{P} .

Theorem 1.8 now immediately yields Theorem 1.2. \square

We give a high-level overview of the proof of Theorem 1.8 in Section 2 below after a discussion of related work in Section 1.2.

1.2 Related Work

List-decodability of random ensembles of codes. As mentioned above, it is not hard to see that a completely random code $C \subset \Sigma^n$ achieves list-decoding capacity. There has also been work studying more structured random ensembles of codes, notably random linear codes. Zyablov and Pinsker [ZP81] showed that random linear codes of rate $1 - h_q(\alpha) - \varepsilon$ are (α, L) -list-decodable with high probability, where L is independent of n but depends exponentially on $1/\varepsilon$. Two decades later, [GHSZ02] showed that there exist binary linear codes with list-size $O(1/\varepsilon)$, and their techniques were recently extended to hold with high probability in [LW18]. In the meantime, [GHK11] showed that random linear codes over any constant-sized alphabet achieve capacity with $L = O(1/\varepsilon)$ when α is bounded away from $1 - 1/q$; [CGV13, Woo13, RW14, RW18] extended these results to get list sizes nearly as good even for large α , although the problem is still open in some parameter regimes.

Several variants of list-decoding have been studied for random linear codes, including *list-recovery* [RW18], *average-radius list-decoding* [Woo13, RW14, RW18], and list-recovery from erasures [Gur03].⁷ All of these properties are local, and so our main theorem implies that LDPC codes satisfy them with high probability.

List-decodability of explicit codes. Obtaining explicit constructions of codes which achieve list-decoding capacity was a major open problem until it was solved about a decade ago. The first explicit codes to provably achieve capacity were the *Folded Reed-Solomon Codes* of Guruswami and Rudra [GR08]. These codes are variants on the classic *Reed-Solomon codes* and are based on polynomials over finite fields. Since then, there have been several constructions of such codes, also based on algebraic techniques, including *Univariate Multiplicity Codes* [GW13, Kop15, KRSW18], variants of Algebraic-Geometry Codes [GX12, GX13], and manipulations of these codes [DL12, GK16, HRW17, KRRZ⁺19]. However, the state-of-the-art for explicit constructions still requires quite large (but constant) alphabet and list sizes. These codes can be efficiently list-decoded in polynomial time; the fastest algorithm is that of [HRW17, KRRZ⁺19], which runs in nearly-linear time $O(n^{1+o(1)})$.

While graph-based techniques have been used to modify the underlying algebraic constructions (for example the expander-based distance-amplification technique of [AEL95] is used in [HRW17, KRRZ⁺19] to obtain near-linear-time list-decoding), to the best of our knowledge there are no results establishing list-decodability up to capacity for purely graph-based codes such as LDPC codes or expander codes.⁸

Finally, we note that recent work [DHK⁺19] has given an algorithm to list-decode codes based on high-dimensional expanders, but these results are far from list-decoding capacity.

LDPC Codes Achieve Capacity on the Binary Symmetric Channel. LDPC Codes have been studied extensively in the context of unique decoding, especially in a model of random errors. Informally, a code is said to achieve capacity on the Binary Symmetric Channel (BSC) if there is some algorithm which can, with high probability, uniquely decode a code of rate $R = 1 - h_2(\alpha) - \varepsilon$ from an α -fraction of *random* errors. It is known that Gallager’s LDPC codes nearly achieve capacity on the BSC as n gets large, under maximum-likelihood decoding [Gal62, Gur06], and recently it was shown that certain LDPC codes achieve capacity for smaller block lengths under efficient decoding algorithms as well [KRU13]. Achieving capacity on the BSC is related to achieving list-decoding capacity (in particular, the capacities are the same, $R = 1 - h_q(\alpha)$). However, there is no formal connection along these lines, and to the best of our knowledge these results about the BSC do not imply anything about the list-decodability of LDPC codes.

⁷ List-recovery is a generalization of list-decoding where the input is a list of sets Z_1, \dots, Z_n of size at most ℓ (instead of a received word $z \in \Sigma^n$, which can be seen as the $\ell = 1$ case), and goal is to find all of the codewords $c \in C$ so that $c_i \in Z_i$ for at least a $1 - \alpha$ fraction of the $i \in [n]$. Average-radius list-decoding is a strengthening of list-decoding where instead of requiring that no set of $L + 1$ codewords are *all* close to some z , we require that no set of $L + 1$ codewords has small *average* distance to z . List-decoding from erasures is a weaker notion than list-decoding, where $z \in (\Sigma \cup \{\perp\})^n$ has some *erased* symbols, and the goal is to recover all $c \in C$ which agree with z on the observed coordinates.

⁸We note that [HW18] give capacity-achieving graph-based codes for zero-error list-recovery (with erasures), where the input is lists Z_1, \dots, Z_n so that most lists have small size, and the goal is to return all codewords $c \in C$ that satisfy $c_i \in Z_i$ for all i . It does not seem easy to adapt these techniques for general list-recovery and hence for list-decoding.

Relationship to graph properties. Finally, we note that our results providing sharp thresholds of local properties for random linear codes are reminiscent of classic results about local properties of random graphs. We discuss this connection more in Remark 2.10.

1.3 Discussion and open questions

In this work, we answer Question 1.1 with a very strong “yes.” There are LDPC codes that achieve list-decoding capacity, and moreover there are many of them, and moreover these codes also likely satisfy any local property—that is, any property which can be defined by ruling out small bad sets of codewords—which is likely satisfied by a random linear code. Our results raise several interesting questions:

1. **What other properties are local?** We have shown that random LDPC codes satisfy with high probability any local property that random linear codes satisfy with high probability. There are several natural examples of local properties, including distance, list-decoding and list-recovery. What other examples are there?
2. **What other applications of Theorem 2.8 are there?** In subsequent work [GLM⁺20], the characterization of a sharp threshold for local properties of random linear codes (Theorem 2.8) was already demonstrated to be useful beyond our work on LDPC codes. We hope to see additional applications of this result. For example, Remark 2.9 implies that to prove that $C_{\text{RLC}}(R - \varepsilon)$ satisfies a local property \mathcal{P} with probability $1 - 2^{-\Omega(n)}$, it suffices to show that $C_{\text{RLC}}(R)$ satisfies \mathcal{P} with some tiny probability (at least $2^{-o(n)}$). Are there situations where this could be useful?
3. **Derandomization?** Our results hold for a random ensemble of LDPC codes. It is natural to ask whether (or to what extent) this construction can be derandomized. In particular, it does not seem as though the underlying graph being an expander would be sufficient.
4. **Algorithms?** Our results are combinatorial, but one of our main motivations is algorithmic. At the moment we do not know of any truly linear-time list-decoding algorithms for any capacity-achieving list-decodable codes. Since essentially all known linear-time algorithms in coding theory arise from graph-based codes, such codes are a natural candidate for linear-time list-decoding. Now that we know that random LDPC codes achieve list-decoding capacity combinatorially, can we list-decode them efficiently?

1.4 Organization and main building blocks

In Section 2, we give a high-level overview of the proof of Theorem 1.8. This proof relies on three building blocks:

- First, Lemma 2.7 establishes sharp thresholds for certain local properties, and effectively characterizes the sorts of sets $B \subseteq \mathbb{F}_q^n$ that are contained in a random linear code. We prove this lemma in Section 3. Using Lemma 2.7 we prove Theorem 2.8, which pins down a sharp threshold for any local property of a random linear code.
- Second, Lemma 2.13 shows that for a set B with a certain property called δ -smoothness, the probability that B appears in a random s -LDPC code is not much larger than the probability

that it appears in a random linear code of the same rate. We prove this Lemma 2.13 in Section 4 using Fourier analysis.

Together with Lemma 2.7, Lemma 2.13 implies that any property satisfied with high probability by a random linear code is also satisfied with high probability by a random s -LDPC code of similar rate, provided that we can restrict our attention to δ -smooth sets B . It turns out that for any code with good distance,⁹ we may indeed restrict our attention to such sets, so it remains to show that random s -LDPC codes have good distance.

- Third, Theorem 2.14 shows that random s -LDPC codes do indeed have good distance with high probability. This was already shown by Gallager in the binary case; we give an alternative proof of this fact that also extends to large alphabets. We prove Theorem 2.14 in Section 5 using techniques from exponential families.

Together, these three building blocks can be used to establish Theorem 1.8, as we show next in Section 2.

2 High-level idea: proof of Theorem 1.8

In this section we prove our main theorem (Theorem 1.8) using the building blocks outlined in Section 1.4. We will establish these building blocks in later sections. The purpose of this section is to give a high-level idea of the structure of the proof, deferring the technical parts to later sections. However, we will need a few technical definitions, outlined in Section 2.1.

2.1 Notation and definitions

Because we are studying local properties, we need some notation around sets $B \subseteq \mathbb{F}_q^n$. For such a set B of size ℓ , it will be convenient to view B as a matrix $M \in \mathbb{F}_q^{n \times \ell}$ with the elements of B as the columns. (The ordering of the columns will not matter.) We say that M is contained in a code $C \subseteq \mathbb{F}_q^n$ (written “ $M \subset C$ ”) if all of the columns of M belong to C .

The notion of permutation-invariant properties leads us to think about permutations of the rows of such a matrix $M \in \mathbb{F}_q^{n \times \ell}$. Motivated by this, we define τ_M , the row distribution of M , as follows: for any $v \in \mathbb{F}_q^\ell$,

$$\tau_M(v) := \frac{\text{number of appearances of } v \text{ as a row in } M}{n}.$$

Let $\mathcal{D}_{n,\ell}$ denote the collection of possible row distributions of matrices in $\mathbb{F}_q^{n \times \ell}$, i.e., distributions τ over \mathbb{F}_q^ℓ where $\tau(v) \cdot n \in \mathbb{N}$ for any $v \in \text{supp}(\tau)$.¹⁰ The number of possible row distributions of matrices in $\mathbb{F}_q^{n \times \ell}$ is just the number of ways to partition n things into at most q^ℓ groups, so

$$|\mathcal{D}_{n,\ell}| \leq \binom{n + q^\ell - 1}{q^\ell - 1}. \tag{3}$$

⁹The distance of a code is the minimum distance between any two codewords.

¹⁰Notice that $\mathcal{D}_{n,\ell}$ depends on q as well, but we suppress this dependence in the notation for readability.

For a distribution $\tau \in \mathcal{D}_{n,\ell}$, let $\mathcal{M}_{n,\tau}$ denote the collection of matrices $M \in \mathbb{F}_q^{n \times \ell}$ with row distribution τ . We say that a code C *contains* τ to mean that $M \subset C$ for some matrix $M \in \mathcal{M}_{n,\tau}$. Let

$$\mathcal{L}_\tau = \{n \in \mathbb{N} \mid \tau(u) \cdot n \text{ is an integer for all } u \in \mathbb{F}_q^\ell\}.$$

Note that for C to contain τ , a trivial necessary condition is that the length of C belongs to \mathcal{L}_τ . Let \mathcal{P}_τ denote the ℓ -local property of not containing any matrix from the set $\mathcal{M}_{n,\tau}$. Properties of the form \mathcal{P}_τ are particularly useful to us due to the following observation:

Observation 2.1 (Local property decomposition). *Let $\mathcal{P} = (P_n)_{n \in \mathbb{N}}$ be an ℓ -local property for some $\ell \in \mathbb{N}$. Then, for every $n \in \mathbb{N}$ there exists $T_n \subseteq \mathcal{D}_{n,\ell}$ such that*

$$C \subseteq \mathbb{F}_q^n \text{ satisfies } P_n \iff C \text{ satisfies } P_\tau \text{ for all } \tau \in T_n.$$

Proof. Note that for every $\tau \in \mathcal{D}_{n,\ell}$, the set of matrices $\mathcal{M}_{n,\tau}$ is closed under row permutations. The lemma now follows immediately from the definition of a local property. \square

Finally, let $H(\tau)$ and $H_q(\tau)$ denote the entropy and base- q -entropy of a random variable distributed according to τ :

$$H(\tau) := - \sum_{x \in \text{supp}(\tau)} \tau(x) \log(\tau(x)) \quad \text{and} \quad H_q(\tau) := \frac{H(\tau)}{\log q}.$$

Let

$$d(\tau) := \dim(\text{span}(\text{supp}(\tau))).$$

We will work with the parity-check matrix view of a random s -LDPC code C . Let $H \in \mathbb{F}_q^{(1-R)n \times n}$ be the adjacency matrix of the graph G in Figure 1 where the nonzero entries are given by the coefficients $\alpha_{i,j}$ of the parity checks. Then we can define a random s -LDPC code C as

$$C = \{x \in \mathbb{F}_q^n : H \cdot x = 0\}.$$

We introduce some notation to talk about the structure of H , which we will use throughout the paper. This is illustrated in Figure 2.

Let $F \in \{0,1\}^{(n/s) \times n}$ be the matrix $F = (F_1 \mid F_2 \mid \dots \mid F_{n/s})$, where each $F_i \in \{0,1\}^{(n/s) \times s}$ has all-ones i -th row, and the rest of the rows are all-zeros. Let $\Pi \in \{0,1\}^{n \times n}$ be a random permutation matrix, and let $D \in \mathbb{F}_q^{n \times n}$ be a diagonal matrix with diagonal entries that are uniform in \mathbb{F}_q^* . Let $H_1, \dots, H_{(1-R) \cdot s}$ be sampled independently according to the distribution $F \cdot \Pi \cdot D$. Then let $H \in \mathbb{F}_q^{(1-R)n \times n}$ be the matrix obtained by stacking $H_1, \dots, H_{(1-R) \cdot s}$ on top of each other (see Figure 2). Then H is the parity-check matrix for a random s -LDPC code of rate R . We will refer to each H_i as a “layer” of H .

We will also require the following standard facts:

Fact 2.2. *A matrix $M \in \mathbb{F}_q^{n \times \ell}$ is contained in a random linear code $C \subseteq \mathbb{F}_q^n$ of rate R with probability $q^{-(1-R) \cdot \text{rank}(M) \cdot n}$.*

We include the proof of Fact 2.2 for completeness.




Figure 2: The matrices F and H . Each layer H_i of H is drawn independently according to the distribution $F \cdot \Pi \cdot D$, where $\Pi \in \{0, 1\}^{n \times n}$ is a random permutation and $D \in \mathbb{F}_q^{n \times n}$ is a diagonal matrix with diagonal entries that are uniform in \mathbb{F}_q^* .

Proof. Let $v_1, \dots, v_{\text{rank}(M)}$ be columns of M that form a basis for the column span of M . Then for each v_i , $\Pr[v_i \in C] = q^{-(1-R)n}$. Since the v_i are linearly independent, the events that they are contained in a random linear code C are stochastically independent, and so the probability that all $\text{rank}(M)$ of these vectors are contained in C is $q^{-(1-R) \cdot \text{rank}(M) \cdot n}$. \square

Fact 2.3 ([CS⁺04], Lemma 2.2). *For any distribution $\tau \in \mathcal{D}_{n,\ell}$,*

$$q^{H_q(\tau) \cdot n} \cdot \left(\frac{n + q^\ell - 1}{q^\ell - 1} \right)^{-1} \leq |\mathcal{M}_{n,\tau}| \leq q^{H_q(\tau) \cdot n}.$$

2.2 Sharp thresholds for local properties for random linear codes

The first building block is Lemma 2.7 below, which shows that for every distribution $\tau \in \mathcal{D}_{n,\ell}$, the property \mathcal{P}_τ is sharp for random linear codes. Moreover we give a simple characterization of $R_{\text{RLC}}(\mathcal{P}_\tau)$. As an easy corollary, we get Theorem 2.8, which generalizes Lemma 2.7 to any local property, not necessarily of the form \mathcal{P}_τ .

Before stating Lemma 2.7 we give some intuition. Fix some distribution τ over \mathbb{F}_q^ℓ . Let C be a random linear code of length $n \in \mathcal{L}_\tau$ and rate R . We seek a threshold rate, above which C is likely to contain τ . It is natural to attempt a first-moment approach to this problem and ask what is the expected number of matrices from $\mathcal{M}_{n,\tau}$ which are contained in C . Note that $|\mathcal{M}_{n,\tau}| = q^{n \cdot H_q(\tau)} \cdot \text{poly}(n)$. Indeed, if u_1, \dots, u_{q^ℓ} are an enumeration of \mathbb{F}_q^ℓ , then $\mathcal{M}_{n,\tau}$ is in one-to-one correspondence with partitions on $[n]$ into q^ℓ subsets of sizes $n\tau(u_1), \dots, n\tau(u_{q^\ell})$. That is, $|\mathcal{M}_{n,\tau}| = \binom{n}{n\tau(u_1), \dots, n\tau(u_{q^\ell})} = q^{n H_q(\tau)} \cdot \text{poly}(n)$, where the last estimate follows from Fact 2.3, and relies on our assumption that $n \in \mathcal{L}_\tau$.

Given $M \in \mathcal{M}_{n,\tau}$, the code C contains M with probability $q^{-n \cdot (1-R) \cdot d(\tau)}$ (see Fact 2.2). Hence, in expectation, C contains roughly $q^{n \cdot (H_q(\tau) - (1-R) \cdot d(\tau))}$ matrices from $\mathcal{M}_{n,\tau}$. In particular, this expectation grows (resp. decays) exponentially in n , when R is larger (resp. smaller) than $1 - \frac{H_q(\tau)}{d(\tau)}$. This motivates the following definition.

Definition 2.4 (Expectation threshold). *Given a distribution τ over \mathbb{F}_q^ℓ , define the expectation-threshold*

$$R_{\text{RLC}}^{\mathbb{E}}(\tau) := 1 - \frac{H_q(\tau)}{d(\tau)}.$$

It follows immediately from a first-moment argument that if $R < R_{\text{RLC}}^{\mathbb{E}}(\tau)$ then C satisfies \mathcal{P}_τ with probability $1 - e^{-\Omega(n)}$. In particular, as n grows we get the lower bound

$$R_{\text{RLC}}^n(\mathcal{P}_\tau) \geq R_{\text{RLC}}^{\mathbb{E}}(\tau) - o(1). \quad (4)$$

However, as the following example shows, this bound is not tight.

Example 2.5. Let $q = 2$, $\ell = 3$ and consider the distribution τ over \mathbb{F}_2^3 , given by the following table:

u	$\tau(u)$
(1, 0, 0)	1/4
(0, 1, 0)	1/4
(1, 0, 1)	1/4
(0, 1, 1)	1/4
Every other vector	0

It is straightforward to compute $R_{\text{RLC}}^{\mathbb{E}}(\tau) = 1 - \frac{H_2(\tau)}{d(\tau)} = 1 - \frac{2}{3} = \frac{1}{3}$.

We claim that $R_{\text{RLC}}^n(\mathcal{P}_\tau)$ is bounded away from $R_{\text{RLC}}^{\mathbb{E}}(\tau)$. Let $A := \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \in \mathbb{F}_2^{2 \times 3}$ represent the linear map which projects a vector onto its first two coordinates. Let τ' denote the distribution of Au , where u is a random vector sampled from τ . Thus, τ' is distributed as follows:

u	$\tau'(u)$
(1, 0)	1/2
(0, 1)	1/2
Every other vector	0

Note that a code C which contains a matrix M from $\mathcal{M}_{n,\tau}$ must contain the first two columns of M : that is, the matrix MA^T . Consequently, every code which satisfies $\mathcal{P}_{\tau'}$ also satisfies \mathcal{P}_τ , and so $R_{\text{RLC}}^n(\mathcal{P}_\tau) \geq R_{\text{RLC}}^n(\mathcal{P}_{\tau'})$.

Finally, (4) yields

$$R_{\text{RLC}}^n(\mathcal{P}_{\tau'}) \geq R_{\text{RLC}}^{\mathbb{E}}(\tau') - o(1) = 1 - \frac{H_2(\tau')}{d(\tau')} - o(1) = 1 - \frac{1}{2} - o(1) = \frac{1}{2} - o(1)$$

and we conclude that

$$R_{\text{RLC}}^n(\mathcal{P}_\tau) \geq \frac{1}{2} - o(1) > \frac{1}{3} = R_{\text{RLC}}^{\mathbb{E}}(\tau)$$

for large n .

In Example 2.5, the bound of $R_{\text{RLC}}^{\mathbb{E}}(\tau)$ was not tight, in that the rate can actually be much higher than we would expect from a first-moment argument. The reason was that there was some linear map A so that $\tau' = A\tau$ had a larger value of $R_{\text{RLC}}^{\mathbb{E}}(\tau')$. We will show below that this is the only reason that $R_{\text{RLC}}^{\mathbb{E}}(\tau)$ might not be the right answer. To make this precise, we introduce the following definition.

Definition 2.6 (Implied distribution). Let τ be a distribution over \mathbb{F}_q^ℓ and let $A \in \mathbb{F}_q^{m \times \ell}$ be a rank m matrix for some $m \leq \ell$. The distribution of the random vector Au , where u is randomly sampled from τ , is said to be τ -implied. We denote the set of τ -implied distributions by \mathcal{I}_τ .

Note that whenever $\tau' \in \mathcal{I}_\tau$, a linear code satisfying $\mathcal{P}_{\tau'}$ must also satisfy \mathcal{P}_τ . Indeed, in the setting of Definition 2.6 assume that C contains a matrix $M \in \mathcal{M}_{n,\tau}$. By linearity, C also contains the matrix MA^T , which belongs to $\mathcal{M}_{n,\tau'}$. Hence, not satisfying \mathcal{P}_τ implies not satisfying $\mathcal{P}_{\tau'}$. Consequently, $R_{\text{RLC}}^n(\mathcal{P}_\tau) \geq R_{\text{RLC}}^n(\mathcal{P}_{\tau'})$.

Inequality (4) now yields the stronger bound

$$R_{\text{RLC}}(\mathcal{P}_\tau)^n \geq \max_{\tau' \in \mathcal{I}_\tau} R_{\text{RLC}}^{\mathbb{E}}(\tau') - o(1). \quad (5)$$

Lemma 2.7 below essentially says that (5) is tight, and that \mathcal{P}_τ is sharp for random linear codes. We prove this Lemma in Section 3.

Lemma 2.7 (Sharp threshold for \mathcal{P}_τ for random linear codes). *Let $\ell \in \mathbb{N}$ and let τ be a distribution over \mathbb{F}_q^ℓ . Denote $R_\tau^* = \max_{\tau' \in \mathcal{I}_\tau} R_{\text{RLC}}^{\mathbb{E}}(\tau')$. Then*

$$R_{\text{RLC}}^n(\mathcal{P}_\tau) = R_\tau^* \pm o_{n \rightarrow \infty}(1).$$

for $n \in \mathcal{L}_\tau$. Moreover, \mathcal{P}_τ is sharp for random linear codes. Specifically, fix any $\varepsilon > 0$, and let C be a random linear code of rate R and length $n \in \mathcal{L}_\tau$. The following holds:

(i) If $R \leq R_\tau^* - \varepsilon$, then

$$\Pr[\exists M \in \mathcal{M}_{n,\tau}, M \subset C] \leq q^{-\varepsilon n}.$$

(ii) If $R \geq R_\tau^* + \varepsilon$, then

$$\Pr[\exists M \in \mathcal{M}_{n,\tau}, M \subset C] \geq 1 - \binom{n + q^{2\ell} - 1}{q^{2\ell} - 1}^3 \cdot q^{-\varepsilon n}.$$

We now can conclude a more general result.

Theorem 2.8 (Sharp thresholds for local properties for random linear codes). *Fix $\ell \in \mathbb{N}$. Let $\mathcal{P} = (P_n)_{n \in \mathbb{N}}$ be an ℓ -local property and let $(T_n)_{n \in \mathbb{N}}$ be as in Observation 2.1. Then \mathcal{P} is sharp for random linear codes and*

$$R_{\text{RLC}}^n(\mathcal{P}) = \min_{\tau \in T_n} \max_{\tau' \in \mathcal{I}_\tau} R_{\text{RLC}}^{\mathbb{E}}(\tau') \pm o_{n \rightarrow \infty}(1).$$

Proof of Theorem 2.8. Denote

$$R_n^* = \min_{\tau \in T_n} \max_{\tau' \in \mathcal{I}_\tau} R_{\text{RLC}}^{\mathbb{E}}(\tau')$$

and fix $\varepsilon > 0$. To prove the theorem, it suffices to show the following:

1. $\lim_{n \rightarrow \infty} \Pr[C_{\text{RLC}}^n(R_n^* - \varepsilon) \text{ satisfies } \mathcal{P}] = 1$
2. $\lim_{n \rightarrow \infty} \Pr[C_{\text{RLC}}^n(R_n^* + \varepsilon) \text{ satisfies } \mathcal{P}] = 0.$

For the first statement, let $C = C_{\text{RLC}}^n(R_n^* - \varepsilon)$. For each $\tau \in T_n$, Lemma 2.7(i) guarantees that $\Pr[C \text{ contains } \tau] \leq q^{-\varepsilon n}$. We take a union bound over all $\tau \in T_n$ noting that

$$|T_n| \leq |\mathcal{D}_{n,\ell}| \leq \binom{n + q^\ell - 1}{q^\ell - 1} \leq (n + q^\ell)^{q^\ell}$$

due to (3). This yields

$$\Pr[C \text{ satisfies } P_n] \leq (n + q^\ell)^{q^\ell} \cdot q^{-\varepsilon n} \leq o_{n \rightarrow \infty}(1).$$

We turn to the second statement. Let $C = C_{\text{RLC}}^n(R_n^* + \varepsilon)$, and let $\tau \in T_n$ such that

$$\max_{\tau' \in \mathcal{T}_\tau} R_{\text{RLC}}^{\mathbb{E}}(\mathcal{P}'_{\tau'}) = R^*.$$

By Lemma 2.7(ii), C almost surely contains τ , which is a sufficient condition for the code not to satisfy \mathcal{P} . \square

Remark 2.9 (Probability of satisfying \mathcal{P} in Theorem 2.8). *Fix $\varepsilon > 0$. An inspection of the proof of Theorem 2.8 shows that $C_{\text{RLC}}^n(R_{\text{RLC}}^n(\mathcal{P}) - \varepsilon)$ satisfies \mathcal{P} with probability $1 - 2^{-\Omega(n)}$. Likewise, $C_{\text{RLC}}^n(R_{\text{RLC}}^n(\mathcal{P}) + \varepsilon)$ satisfies \mathcal{P} with probability $2^{-\Omega(n)}$.*

Remark 2.10 (Relationship to random graphs). *Lemma 2.7 has an analog in the theory of random graphs. Fix a constant-sized graph H and let G be a random graph in the $G(n, p)$ model. A natural problem is to determine the threshold for the appearance of H as a sub-graph of G . The answer (see for example [Bol01, Sec. 4.2]) is that a copy of H is likely to occur in G whenever p is large enough so that every subgraph of H has, in expectation, $\omega(1)$ copies as subgraphs of G . To complete the analogy, equate H with τ , and a subgraph of H with a τ -implied distribution.*

We also mention the recent breakthrough result of Frankston et al., which studies this relationship between thresholds and expectations of sub-structures in a more general framework [FKNP19]. However, since the properties that they study are not necessarily local, it is impossible for that work to precisely pinpoint the thresholds, as we do in our work.

2.3 Probability that a matrix is contained in a random s -LDPC code

The second building block shows that given a matrix $M \in \mathbb{F}_q^{n \times \ell}$, the probability that M is contained in a random s -LDPC code is not much larger than that of appearing in a random linear code, provided that M is δ -smooth (defined below).

Definition 2.11 (Smooth distribution). *Let $\delta > 0$. We say that a distribution τ over \mathbb{F}_q^ℓ is δ -smooth if $\Pr_{v \sim \tau}[\langle u, v \rangle \neq 0] \geq \delta$ for all $u \in \mathbb{F}_q^\ell \setminus \{0\}$. If $M \in \mathbb{F}_q^{n \times \ell}$ is such that τ_M is δ -smooth, we also say that M is δ -smooth.*

Remark 2.12 (Relationship to distance). *In coding-theoretic terms, τ_M is δ -smooth if and only if the code $\{Mu : u \in \mathbb{F}_q^\ell\}$ has relative distance at least δ and M is full-rank. Indeed, the relative weight of any codeword Mu in this code is*

$$\frac{1}{n} \sum_{i \in [n]} \mathbf{1}_{\langle u, e_i^T M \rangle \neq 0} = \Pr_{v \sim \tau}[\langle u, v \rangle \neq 0].$$

The following lemma bounds the probability that a matrix with smooth row distribution is contained in a random LDPC code with sufficiently large sparsity parameter. We prove this lemma in Section 4.

Lemma 2.13 (Probability that a random LDPC code contains a matrix). *For any $\delta, \varepsilon > 0$, prime power q , and $\ell \geq 1$ there exists $s_0 \geq 1$ such that the following holds for any odd $s \geq s_0$, and sufficiently large n . Let $M \in \mathbb{F}_q^{n \times \ell}$ be δ -smooth. Then the probability p that M is contained in a random s -LDPC code of length n and rate R satisfies*

$$p \leq q^{-(1-\varepsilon) \cdot (1-R) \cdot \ell \cdot n}.$$

Given a smooth distribution τ , in light of Fact 2.2, Lemma 2.13 says that the expected number of matrices from $\mathcal{M}_{n,\tau}$ in a random s -LDPC code is not much larger than this number for a random linear code. If we ignore the constraint that τ must be smooth, then together with Lemma 2.7 the above would imply Theorem 1.8. Indeed, if a distribution τ is unlikely to appear in a random linear code then Lemma 2.7 shows that some τ -implied distribution τ' appears $o(1)$ times in expectation in the random linear code. By Lemma 2.13, τ' appears $o(1)$ times in the random LDPC code as well, so the LDPC code is unlikely to contain τ' . Thus, it is also unlikely to contain τ . (Of course, we cannot ignore the constraint that τ must be smooth; we will address this in our next building block discussed in Section 2.4).

The proof of Lemma 2.13 proceeds by Fourier analysis. The basic idea is as follows: since C is a random s -LDPC code, each parity-check corresponds (essentially) to an independent and uniformly random set of s coordinates in $[n]$.¹¹ Thus, the probability that a matrix $M \in \mathcal{M}_{n,\tau}$ is in C can be derived from the probability that s random vectors $v_1, \dots, v_s \sim \tau$ sum to zero. This probability is given by a convolution $\tau^{*s}(0) = \tau * \tau * \dots * \tau(0)$ of τ with itself s times. The convolution is in turn controlled by s 'th powers of the Fourier coefficients $\hat{\tau}(w)$ of τ . As we will see, the condition that τ be δ -smooth implies that the nonzero Fourier coefficients $\hat{\tau}(w)$ are bounded away from 1, and this means that if s is large enough, the contributions $\hat{\tau}(w)^s$ of the nonzero coefficients to $\tau^{*s}(0)$ will become small.

2.4 Distance of random s -LDPC codes

As noted above, the first two building blocks show that for any δ -smooth distribution $\tau \sim \mathbb{F}_q^\ell$, a random LDPC code of rate slightly below $R_{\text{RLC}}^n(\mathcal{P}_\tau)$ is unlikely to contain τ . The third and final building block shows that we may restrict our attention to δ -smooth distributions.

As noted in Remark 2.12, the condition that M be δ -smooth is the same as the condition that the code generated by M has relative distance at least δ . Thus, if $C \subset \mathbb{F}_q^n$ has relative distance at least δ , it does not contain any matrices that are not δ -smooth. Fortunately, it is well-known that *binary* random s -LDPC codes have good distance, and that in fact the distance approaches the *Gilbert-Varshamov* (GV) bound with high probability.¹² Theorem 2.14 generalizes this result to s -LDPC codes over any alphabet. Below, $h_q(x)$ is the q -ary entropy function (as in (1)).

Theorem 2.14 (Random LDPC codes achieve the GV bound). *For any $\delta \in (0, 1 - 1/q)$, $\varepsilon > 0$, and prime power q there exists $s_0 \geq 1$ such that the following holds for any $s \geq s_0$. Let $R \leq 1 - h_q(\delta) - \varepsilon$. Then a random s -LDPC code of rate R over \mathbb{F}_q has relative distance at least δ with high probability.*

Remark 2.15 (Comparison to Gallager's proof). *Gallager's proof for binary random s -LDPC codes in [Gal62] uses generating functions. We give an alternative proof using ideas from exponential*

¹¹This is not exactly true because the parity checks that belong to the same layer are not independent; however, we show that this does not significantly affect the probability of the event of interest.

¹²The GV bound refers to the rate-distance trade-off $R = 1 - h_q(\delta)$, which is approached by a random linear code.

families, which follows the approach of recent work by Linial and the first author [LM20]. Our proof extends to random s -LDPC codes over any alphabet. We note that Gallager left it as an open problem in [Gal62] to obtain a result like this for larger alphabets, but his definition was slightly different than ours: the coefficients $\alpha_{i,j}$ in his parity checks were all 1's, while ours are taken randomly from \mathbb{F}_q^* .

Despite having different frameworks, our proof and that of [Gal62] turn out to yield similar equations. In particular our proof of Lemma 5.2 is very similar to the corresponding proof in [Gal62] at a technical level. We highlight where the proofs diverge in Remark 5.9.

2.5 Proof of Theorem 1.8 from Lemma 2.7, Lemma 2.13 and Theorem 2.14

Theorem 1.8 now follows as an immediate consequence of the building blocks above. We restate Theorem 1.8 here:

Theorem 1.8 (Main). *Let $\mathcal{P} = (P_n)_{n \in \mathbb{N}}$ be a b -local property with $\bar{R} := \limsup_{n \rightarrow \infty} R_{\text{RLC}}^n(\mathcal{P}) < 1$. For any $\varepsilon > 0$ and prime power q , there exists $s_0 = s_0(\varepsilon, \bar{R}, q, b) \geq 1$ such that for any odd $s \geq s_0$ and any sequence $\{R_n\}_{n \in \mathbb{N}}$, if $R_n \leq R_{\text{RLC}}^n(\mathcal{P}) - \varepsilon$ for all n , then the code ensemble $C_{s\text{LDPC}}^n(R_n)$ satisfies \mathcal{P} with high probability.*

Proof. Fix a sufficiently large odd integer s (depending on \bar{R} , ε , q and b). For $n \in \mathbb{N}$, let $C := C_{s\text{LDPC}}^n(R_n)$ for some $R_n \leq R_{\text{RLC}}^n(\mathcal{P}) - \varepsilon$. Let T_n be as in Observation 2.1. Let

$$\delta := \frac{h_q^{-1}(1-R)}{2} > 0.$$

Let D denote the event that the relative distance of C is at most δ .

Fix some $\tau \in T_n$. Let $\tau' \in \mathcal{I}_\tau$ be a maximizer of $R_{\text{RLC}}^{\mathbb{E}}(\tau')$. We may assume that τ' is a distribution over $\mathbb{F}_q^{d(\tau')}$. Indeed, otherwise, let $A : \text{span}(\text{supp}(\tau')) \rightarrow \mathbb{F}_q^{d(\tau')}$ be a linear bijection, and take the distribution of Au (for $u \sim \tau'$) in place of τ' itself.

By Lemma 2.7, for n large enough,

$$\begin{aligned} R_n &\leq R_{\text{RLC}}^n(\mathcal{P}) - \varepsilon \\ &\leq R_{\text{RLC}}^n(\mathcal{P}_\tau) - \varepsilon \\ &\leq R_{\text{RLC}}^{\mathbb{E}}(\tau') - \frac{\varepsilon}{2} \\ &= 1 - \frac{H_q(\tau')}{d(\tau')} - \frac{\varepsilon}{2}, \end{aligned}$$

where the first line is our assumption on R_n ; the second line follows from the fact that any code satisfying \mathcal{P} must in particular satisfy \mathcal{P}_τ ; the third line is Lemma 2.7; and the fourth line is the definition of $R_{\text{RLC}}^{\mathbb{E}}(\tau')$.

Consider the case where τ' is δ -smooth. Let p denote the probability that C contains a given matrix from $M_{n,\tau'}$. By Lemma 2.13, for s large enough we have $p \leq q^{-(1-\frac{\varepsilon}{4})(1-R_n) \cdot d(\tau') \cdot n}$. Thus,

the expected number of such matrices in C is at most

$$\begin{aligned}
|M_{n,\tau'}| \cdot p &\leq q^{H(\tau') \cdot n} \cdot p \\
&\leq q^{(H(\tau') - (1 - \frac{\varepsilon}{4})(1 - R_n) \cdot d(\tau')) \cdot n} \\
&\leq q^{(H(\tau') - (1 - \frac{\varepsilon}{4})(\frac{H(\tau')}{d(\tau')} + \frac{\varepsilon}{2}) \cdot d(\tau')) \cdot n} \\
&= q^{(\frac{\varepsilon}{4} \cdot H(\tau') - (1 - \frac{\varepsilon}{4}) \frac{\varepsilon}{2} \cdot d(\tau')) \cdot n} \\
&\leq q^{(\frac{\varepsilon}{4} \cdot d(\tau') - (1 - \frac{\varepsilon}{4}) \frac{\varepsilon}{2} \cdot d(\tau')) \cdot n} \\
&= q^{-\frac{\varepsilon}{8} \cdot d(\tau') \cdot n} \\
&\leq q^{-\frac{\varepsilon}{8} n}.
\end{aligned} \tag{6}$$

Here, we used the fact that $H(\tau') \leq \log_q |\text{supp}(\tau')| = d(\tau')$.

On the other hand, assume that τ' is not δ -smooth. By Remark 2.12, if C contains τ' then the event D must hold (in that setting of that remark, our assumption that the domain of τ' is $\mathbb{F}_q^{d(\tau')}$, is equivalent to M having full-rank). Since any code containing τ must also contain τ' , Markov's inequality and (6) yield

$$\Pr(C \text{ contains } \tau, \text{ and } D \text{ does not hold}) \leq q^{-\frac{\varepsilon}{8} n}.$$

Taking a union bound over all $\tau \in T_n$ and using (3), we get

$$\Pr(C \text{ satisfies } \mathcal{P}, \text{ and } D \text{ does not hold}) \leq q^{-\frac{\varepsilon}{8} n} \cdot |T_n| \leq q^{-\frac{\varepsilon}{8} n} \cdot |\mathcal{D}_{n,\ell}| \leq q^{-\frac{\varepsilon}{8} n} \cdot \binom{n + q^\ell - 1}{q^\ell - 1} \leq o_{n \rightarrow \infty}(1).$$

Finally, Theorem 2.14 says that D almost surely does not hold, and we conclude that C satisfies \mathcal{P} with high probability. \square

3 Sharp thresholds of local properties for random linear codes: proof of Lemma 2.7

In this section we prove Lemma 2.7, which we restate below.

Lemma 2.7 (Sharp threshold for \mathcal{P}_τ for random linear codes). *Let $\ell \in \mathbb{N}$ and let τ be a distribution over \mathbb{F}_q^ℓ . Denote $R_\tau^* = \max_{\tau' \in \mathcal{I}_\tau} R_{\text{RLC}}^{\mathbb{E}}(\tau')$. Then*

$$R_{\text{RLC}}^n(\mathcal{P}_\tau) = R_\tau^* \pm o_{n \rightarrow \infty}(1).$$

for $n \in \mathcal{L}_\tau$. Moreover, \mathcal{P}_τ is sharp for random linear codes. Specifically, fix any $\varepsilon > 0$, and let C be a random linear code of rate R and length $n \in \mathcal{L}_\tau$. The following holds:

(i) *If $R \leq R_\tau^* - \varepsilon$, then*

$$\Pr[\exists M \in \mathcal{M}_{n,\tau}, M \subset C] \leq q^{-\varepsilon n}.$$

(ii) *If $R \geq R_\tau^* + \varepsilon$, then*

$$\Pr[\exists M \in \mathcal{M}_{n,\tau}, M \subset C] \geq 1 - \left(\frac{n + q^{2\ell} - 1}{q^{2\ell} - 1} \right)^3 \cdot q^{-\varepsilon n}.$$

We note that statements (i) and (ii) of Lemma 2.7 also imply the rest of the lemma. Thus, it suffices to prove them.

3.1 Proof of Statement (i)

Assume that τ is such that $R_\tau^* = \max_{\tau' \in \mathcal{I}_\tau} R_{\text{RLC}}^{\mathbb{E}}(\tau')$ satisfies

$$R \leq R_\tau^* - \varepsilon.$$

Choose $\tau' \in \mathcal{I}_\tau$ achieving $R_{\text{RLC}}^{\mathbb{E}}(\tau') = R_\tau^*$ and let $A \in \mathbb{F}_q^{m \times \ell}$ be such that τ' is given by Av for $v \sim \tau$. By Fact 2.2, a matrix $M' \in \mathcal{M}_{n, \tau'}$ is contained in $C = C_{\text{RLC}}^n(R)$ with probability $q^{-(1-R) \cdot \text{rank}(M') \cdot n} = q^{-(1-R) \cdot d(\tau') \cdot n}$, and so

$$\Pr [\exists M \in \mathcal{M}_{n, \tau'}, M \subset C] \leq |\mathcal{M}_{n, \tau'}| \cdot q^{-(1-R) \cdot d(\tau') \cdot n} \leq q^{(H_q(\tau') - (1-R) \cdot d(\tau')) \cdot n} \leq q^{-\varepsilon n},$$

where the first inequality follows by a union bound, the second applies Fact 2.3, and the final inequality uses $R_{\text{RLC}}^{\mathbb{E}}(\tau') = 1 - \frac{H_q(\tau')}{d(\tau')} \geq R + \varepsilon$.

Finally, note that if C contains some matrix $M \in \mathcal{M}_{n, \tau}$, then by linearity, $M' := MA^T \in \mathcal{M}_{n, \tau'}$ is also contained in C . So we conclude

$$\Pr [\exists M \in \mathcal{M}_{n, \tau}, M \subset C] \leq q^{-\varepsilon n}.$$

3.2 Proof of Statement (ii)

We now proceed to the second part of the theorem, which is more involved. Suppose that $\tau \in \mathcal{D}_{n, \ell}$ is such that $R_\tau^* = \max_{\tau' \in \mathcal{I}_\tau} R_{\text{RLC}}^{\mathbb{E}}(\tau')$ satisfies $R \geq R_\tau^* + \varepsilon$.

First, we will argue that we may assume without loss of generality that $d(\tau) = \ell$. By the definition of $d(\tau)$, there is some matrix $B \in \mathbb{F}_q^{d(\tau) \times \ell}$ of rank $d(\tau)$ so that the distribution $\tilde{\tau}$ given by $Bv, v \sim \tau$ has $d(\tilde{\tau}) = d(\tau)$. We claim that

$$\max_{\tau' \in \mathcal{I}_\tau} R_{\text{RLC}}^{\mathbb{E}}(\tau') \leq R - \varepsilon$$

implies that

$$\max_{\tilde{\tau}' \in \mathcal{I}_{\tilde{\tau}}} R_{\text{RLC}}^{\mathbb{E}}(\tilde{\tau}') \leq R - \varepsilon.$$

To see this, we prove the contrapositive. Suppose that there is some $\tilde{\tau}' \in \mathcal{I}_{\tilde{\tau}}$ so that $R_{\text{RLC}}^{\mathbb{E}}(\tilde{\tau}') > R - \varepsilon$. Then by the definition of $\mathcal{I}_{\tilde{\tau}}$, there is some matrix $A \in \mathbb{F}_q^{m \times d(\tilde{\tau})}$ so that $\tilde{\tau}'$ is given by $Aw, w \sim \tilde{\tau}$. But this is the same as the distribution $ABv, v \sim \tau$, using the definition of $\tilde{\tau}$. Thus, $\tilde{\tau}' \in \mathcal{I}_\tau$, and this implies that $\max_{\tau' \in \mathcal{I}_\tau} R_{\text{RLC}}^{\mathbb{E}}(\tau') > R - \varepsilon$. Finally, we observe that $\binom{n+q^{2\ell}-1}{q^{2\ell}-1}$ is increasing in ℓ , we conclude that to prove the statement (ii), we may as well work with the distribution $\tilde{\tau}$ on $\mathbb{F}_q^{d(\tilde{\tau})}$. Thus, in the following we will assume that $d(\tau) = \ell$.

For a matrix $M \in \mathbb{F}_q^{n \times \ell}$, let X_M be the indicator variable for the event that $M \subseteq C$, and let $X = \sum_{M \in \mathcal{M}_{n, \tau}} X_M$. Our goal then is to show that $X > 0$ with high probability, and we do so by showing that $\text{Var}(X) = o(\mathbb{E}^2[X])$.

We first show a lower bound on $\mathbb{E}[X]$. By Facts 2.2 and 2.3,

$$\mathbb{E}[X] = |\mathcal{M}_{n, \tau}| \cdot q^{-(1-R) \cdot \ell \cdot n} \geq q^{(H_q(\tau) - (1-R) \cdot \ell) \cdot n} \cdot \binom{n+q^\ell-1}{q^\ell-1}^{-1}. \quad (7)$$

Next we show an upper bound on $\text{Var}(X)$. Given a pair of matrices $M, M' \in \mathcal{M}_{n,\tau}$, we let $(M|M')$ denote the $(n \times (2\ell))$ -matrix consisting of a left $n \times \ell$ block equal to M , and a right $n \times \ell$ block equal to M' . Then in this notation we have

$$\begin{aligned} \text{Var}(X) &= \sum_{M, M' \in \mathcal{M}_{n,\tau}} \left(\mathbb{E}[X_M \cdot X_{M'}] - \mathbb{E}[X_M] \cdot \mathbb{E}[X_{M'}] \right) \\ &= \sum_{M, M' \in \mathcal{M}_{n,\tau}} \left(\Pr[(M|M') \subseteq C] - \Pr[M \subseteq C] \cdot \Pr[M' \subseteq C] \right) \\ &= \sum_{M, M' \in \mathcal{M}_{n,\tau}} \left(q^{-(1-R) \cdot \text{rank}(M|M') \cdot n} - q^{-2 \cdot (1-R) \cdot \ell \cdot n} \right). \end{aligned}$$

Notice that in the above sum, terms for which $\text{rank}(M|M') = 2\ell$ vanish. Let

$$\mathcal{M} := \left\{ (M|M') \mid M, M' \in \mathcal{M}_{n,\tau} \text{ and } \text{rank}(M|M') < 2\ell \right\},$$

and

$$\mathcal{D} := \{\tau_M : M \in \mathcal{M}\}. \quad (8)$$

Then we have

$$\begin{aligned} \text{Var}(X) &\leq \sum_{M \in \mathcal{M}} q^{-(1-R) \cdot \text{rank}(M) \cdot n} \\ &= \sum_{\tau' \in \mathcal{D}} \sum_{M \in \mathcal{M}_{n,\tau'}} q^{-(1-R) \cdot \text{rank}(M) \cdot n} \\ &= \sum_{\tau' \in \mathcal{D}} |\mathcal{M}_{n,\tau'}| \cdot q^{-(1-R) \cdot d(\tau') \cdot n}. \\ &\leq \sum_{\tau' \in \mathcal{D}} q^{(H_q(\tau') - (1-R) \cdot d(\tau')) \cdot n} \end{aligned}$$

where the last inequality follows by Fact 2.3. Finally, Claim 3.1 below shows that for any $\tau' \in \mathcal{D}$,

$$H_q(\tau') - (1-R) \cdot d(\tau') \leq 2(H_q(\tau) - (1-R) \cdot \ell) - \varepsilon,$$

which implies in turn that

$$\text{Var}(X) \leq |\mathcal{D}| \cdot q^{2(H_q(\tau) - (1-R) \cdot \ell) \cdot n} \cdot q^{-\varepsilon n} \leq \left(\frac{n + q^{2\ell} - 1}{q^{2\ell} - 1} \right) \cdot q^{2(H_q(\tau) - (1-R) \cdot \ell) \cdot n} \cdot q^{-\varepsilon n}. \quad (9)$$

Above, we used the fact that $\mathcal{D} \subseteq \mathcal{D}_{n,2\ell}$ and applied (3). Combining (7) and (9), by Chebyshev's inequality we conclude that

$$\Pr[X = 0] \leq \frac{\text{Var}(X)}{\mathbb{E}^2[X]} \leq \left(\frac{n + q^{2\ell} - 1}{q^{2\ell} - 1} \right)^3 \cdot q^{-\varepsilon n}.$$

To complete the proof, we prove Claim 3.1 which we used above.




Figure 3: Notation in the proof of Claim 3.1

Claim 3.1. Let \mathcal{D} be as in (8). For any $\tau' \in \mathcal{D}$,

$$H_q(\tau') - (1 - R) \cdot d(\tau') \leq 2(H_q(\tau) - (1 - R) \cdot \ell) - \varepsilon.$$

Proof. In what follows, let $d := d(\tau')$, and $V := \text{span}(\text{supp}(\tau')) \subseteq \mathbb{F}_q^{2\ell}$. Let $w_1, \dots, w_{2\ell-d} \in \mathbb{F}_q^{2\ell}$ be a basis for V^\perp . Let $\pi_1 : \mathbb{F}_q^{2\ell} \rightarrow \mathbb{F}_q^\ell$ (respectively, π_2) denote the projection of a vector $w \in \mathbb{F}_q^{2\ell}$ to the first (respectively, last) ℓ coordinates. Finally, let A be the matrix whose rows are $w_1, \dots, w_{2\ell-d}$, and let $A_1 \in \mathbb{F}_q^{(2\ell-d) \times \ell}$ (A_2 , respectively) denote the matrix whose rows are $\pi_1(w_1), \dots, \pi_1(w_{2\ell-d})$ ($\pi_2(w_1), \dots, \pi_2(w_{2\ell-d})$, respectively). See Figure 3 for a diagram of this notation.

We claim that all rows of A_1 are linearly independent, and so $\text{rank}(A_1) = 2\ell - d$. To see this suppose in contradiction that $\pi_1(w_1), \dots, \pi_1(w_{2\ell-d})$ are linearly dependent. Then there exists a non-trivial linear combination of $w_1, \dots, w_{2\ell-d}$ that sums to a non-zero vector of the form $(0, w)$. But this means that $\pi_2(\text{supp}(\tau')) = \text{supp}(\tau)$ is orthogonal to w , in contradiction to our assumption that $\text{span}(\text{supp}(\tau)) = \mathbb{F}_q^\ell$. Consequently, recalling that $d(\tau) = \ell$, the distribution τ'' given by $A_1 w$ for $w \sim \tau$ has $d(\tau'') = 2\ell - d$. As $\tau'' \in \mathcal{I}_\tau$, $R_{\text{RLC}}^{\mathbb{E}}(\tau'') \leq R - \varepsilon$.

Let $I_q(X; Y) = H_q(X) - H_q(X | Y)$ denote the base- q mutual information of X and Y . Now for $v \sim \tau'$ we have,

$$\begin{aligned} H_q(\tau') &= H_q(v) \\ &= H_q(\pi_1(v)) + H_q(\pi_2(v)) - I_q(\pi_1(v); \pi_2(v)) \\ &= 2H_q(\tau) - I_q(\pi_1(v); \pi_2(v)) \\ &\leq 2H_q(\tau) - I_q(A_1\pi_1(v); -A_2\pi_2(v)) \\ &= 2H_q(\tau) - H_q(A_1\pi_1(v)) \\ &\leq 2H_q(\tau) - (1 - R + \varepsilon) \cdot d(\tau'') \\ &= 2H_q(\tau) - (1 - R + \varepsilon) \cdot (2\ell - d), \end{aligned}$$

where the first inequality follows since $A_1\pi_1(v) + A_2\pi_2(v) = Av = 0$, and the last inequality follows because $1 - \frac{H_q(\tau'')}{d(\tau'')} = R_{\text{RLC}}^{\mathbb{E}}(\tau'') \leq R - \varepsilon$. Rearranging, and recalling the assumption that $2\ell > d$, gives the desired conclusion. \square

4 Matrices contained in a random LDPC code: proof of Lemma 2.13

In this section we prove our second building block, Lemma 2.13, which we re-state below.

Lemma 2.13 (Probability that a random LDPC code contains a matrix). *For any $\delta, \varepsilon > 0$, prime power q , and $\ell \geq 1$ there exists $s_0 \geq 1$ such that the following holds for any odd $s \geq s_0$, and sufficiently large n . Let $M \in \mathbb{F}_q^{n \times \ell}$ be δ -smooth. Then the probability p that M is contained in a random s -LDPC code of length n and rate R satisfies*

$$p \leq q^{-(1-\varepsilon) \cdot (1-R) \cdot \ell \cdot n}.$$

Remark 4.1 (The parity of s , again). *Lemma 2.13 holds for even s as well as odd s , but the proof is slightly simpler for odd s , so we state and prove it in this case for clarity. This is the only place in the proof of Theorem 1.8 where we use the parity of s , and so this remark implies Remark 1.3.*

We begin with some definitions from Fourier analysis which we will need.

4.1 Fourier-analytic facts

We give here some basic definitions and facts from Fourier analysis of functions on \mathbb{F}_q . We refer the reader to, for example, [LN94, O'D14] for more details and proofs of these facts. In what follows assume that $q = p^h$ for a prime p . The trace map of \mathbb{F}_q over \mathbb{F}_p is the function $\text{tr} : \mathbb{F}_q \rightarrow \mathbb{F}_p$ given by

$$\text{tr}(\alpha) = \alpha + \alpha^p + \alpha^{p^2} + \cdots + \alpha^{p^{h-1}}.$$

For a function $f : \mathbb{F}_q^n \rightarrow \mathbb{C}$, we define the Fourier transform $\hat{f} : \mathbb{F}_q^n \rightarrow \mathbb{C}$ of f by

$$\hat{f}(y) = \mathbb{E}_{x \in \mathbb{F}_q^n} \left[f(x) \cdot \overline{\chi_x(y)} \right],$$

where $\chi_x(y) = \omega_p^{\text{tr}(\langle x, y \rangle)}$, and $\omega_p = e^{2\pi i/p}$. Then we have the decomposition

$$f(x) = \sum_{y \in \mathbb{F}_q^n} \hat{f}(y) \cdot \chi_y(x).$$

We define an inner product on the space of \mathbb{C} -valued functions on \mathbb{F}_q^n by

$$\langle f, g \rangle = \mathbb{E}_{x \in \mathbb{F}_q^n} \left[f(x) \cdot \overline{g(x)} \right].$$

Plancherel's identity then asserts that

$$\langle f, g \rangle = \sum_{x \in \mathbb{F}_q^n} \hat{f}(x) \cdot \overline{\hat{g}(x)}.$$

An important special case is Parseval's identity:

$$\langle f, f \rangle = \sum_{x \in \mathbb{F}_q^n} |\hat{f}(x)|^2.$$

The convolution of a pair of functions $f, g : \mathbb{F}_q^n \rightarrow \mathbb{C}$ is given by

$$(f * g)(x) = \mathbb{E}_{y \in \mathbb{F}_q^n} [f(y) \cdot g(x - y)].$$

Convolution interacts nicely with the Fourier transform:

$$\widehat{f * g}(x) = \hat{f}(x) \cdot \hat{g}(x).$$

Finally, we define inductively $f^{*1} := f$, and $f^{*s} = f^{*(s-1)} * f$ for an integer $s \geq 2$.

4.2 Proof of Lemma 2.13

Let $H \in \mathbb{F}_q^{((1-R) \cdot n) \times n}$ be the parity-check matrix of C with layers $H_1, H_2, \dots, H_{(1-R) \cdot s}$, as in Section 2.1. Recall that each layer H_i is an independent sample from $FD\Pi$, where F is as in Figure 2, $\Pi \in \{0, 1\}^{n \times n}$ is a random permutation matrix, and $D \in \mathbb{F}_q^{n \times n}$ is a diagonal matrix with diagonal entries that are independent and uniformly random in \mathbb{F}_q^* . Let Λ be a random matrix sampled according to the distribution $D\Pi M$. Then by independence of the layers,

$$\begin{aligned} \Pr[M \subseteq C] &= \Pr[HM = 0] \\ &= (\Pr[H_1 M = 0])^{(1-R) \cdot s} \\ &= (\Pr[FD\Pi M = 0])^{(1-R) \cdot s} \\ &= (\Pr[F\Lambda = 0])^{(1-R) \cdot s}. \end{aligned} \tag{10}$$

So it suffices to bound the probability that $F\Lambda = 0$.

Next observe that each row in Λ has the same marginal distribution P given by λv for $v \sim \tau_M$ and uniformly random $\lambda \in \mathbb{F}_q^*$. Let $\Lambda' \in \mathbb{F}_q^{n \times \ell}$ be a random matrix in which each row is independently sampled according to P . We claim that

$$\Pr[F\Lambda = 0] \leq O\left(n^{\frac{q^\ell - 1}{2}}\right) \cdot \Pr[F\Lambda' = 0]. \tag{11}$$

Indeed,

$$\begin{aligned} \Pr[F\Lambda = 0] &= \Pr[F\Lambda' = 0 \mid \tau_{\Lambda'} = \tau_M] \\ &= \frac{\Pr[F\Lambda' = 0 \wedge \tau_{\Lambda'} = \tau_M]}{\Pr[\tau_{\Lambda'} = \tau_M]} \\ &\leq \frac{\Pr[F\Lambda' = 0]}{\Pr[\tau_{\Lambda'} = \tau_M]}. \end{aligned}$$

Now we have

$$\Pr[\tau_{\Lambda'} = \tau_M] = \binom{n}{n \cdot \tau_M(v_1), \dots, n \cdot \tau_M(v_{q^\ell})} \cdot \prod_{v \in \mathbb{F}_q^\ell} \tau_M(v)^{n \cdot \tau_M(v)}$$

where v_1, \dots, v_{q^ℓ} are the elements of \mathbb{F}_q^ℓ . Noting that $\prod_{v \in \mathbb{F}_q^\ell} \tau_M(v)^{n \cdot \tau_M(v)} = q^{nH_q(\tau_M)}$, (11) follows from Fact 2.3.

Thus, it is enough to bound the probability that $F\Lambda' = 0$. We can express this probability as

$$\Pr[F\Lambda' = 0] = \left(\Pr_{u_1, \dots, u_s \sim P} \left[\sum_{i=1}^s u_i = 0 \right] \right)^{n/s} = \left(q^{\ell \cdot (s-1)} \cdot P^{*s}(0) \right)^{n/s}. \tag{12}$$

Next we bound $P^{*s}(0)$. In terms of Fourier transform, we can write

$$P^{*s}(0) = \sum_{y \in \mathbb{F}_q^\ell} \widehat{P^{*s}}(y) \cdot \chi_y(0) = \sum_{y \in \mathbb{F}_q^\ell} \left(\widehat{P}(y) \right)^s.$$

Claim 4.3 below shows that $\hat{P}(y) \leq q^{-\ell} \cdot \left(1 - \frac{q}{q-1} \cdot \delta\right)$ for any $y \in \mathbb{F}_q^\ell \setminus \{0\}$ (in particular, it's a real number), and by the assumption that s is odd this implies in turn that

$$P^{*s}(0) = \left(\hat{P}(0)\right)^s + \sum_{y \in \mathbb{F}_q^\ell \setminus \{0\}} \left(\hat{P}(y)\right)^s \leq q^{-\ell \cdot s} + q^{-\ell \cdot (s-1)} \cdot \left(1 - \frac{q}{q-1} \cdot \delta\right)^s. \quad (13)$$

Finally, combining Equations (10), (11), (12), and (13) we conclude that

$$\Pr[M \subseteq C] \leq O\left(n^{\frac{q^\ell - 1}{2} \cdot (1-R) \cdot s}\right) \cdot \left(q^{-\ell} + \left(1 - \frac{q}{q-1} \cdot \delta\right)^s\right)^{(1-R) \cdot n} \leq q^{-(1-\varepsilon) \cdot (1-R) \cdot \ell \cdot n},$$

where the last inequality holds for large enough s depending on $\delta, \varepsilon, q, \ell$, and sufficiently large n .

Remark 4.2 (The choice of s). *An inspection of the proof shows that we may take*

$$s_0 = O\left(\frac{\ell}{\log_q\left(\frac{1}{1-\delta/(1-1/q)}\right)}\right).$$

In particular, noting that $\ell \leq b$ and that

$$\log_q\left(\frac{1}{1-\delta/(1-1/q)}\right) = \frac{1}{\ln(q)} \sum_{i=1}^{\infty} \frac{1}{i} \left(\frac{\delta}{1-1/q}\right)^i,$$

this part of the proof requires us to take

$$s_0 \geq C_0 \cdot \frac{b \log(q)}{\delta}$$

for some constant $C_0 > 0$. There is one other place in the proof of Theorem 1.8 that requires s_0 to be sufficiently large; we comment on this again in Remark 5.3.

Now, all that remains is to prove Claim 4.3 which we used above.

Claim 4.3. *For any $y \in \mathbb{F}_q^\ell \setminus \{0\}$, $\hat{P}(y) \in \mathbb{R}$ and*

$$\hat{P}(y) \leq q^{-\ell} \cdot \left(1 - \frac{q}{q-1} \cdot \delta\right).$$

Proof of Claim 4.3. We have

$$\begin{aligned} \hat{P}(y) &= q^{-\ell} \cdot \sum_{x \in \mathbb{F}_q^\ell} P(x) \cdot \overline{\omega_p^{\text{tr}(\langle y, x \rangle)}} \\ &= q^{-\ell} \cdot \mathbb{E}_{v \sim \tau_M} \mathbb{E}_{\lambda \in \mathbb{F}_q^*} [\omega_p^{-\text{tr}(\langle y, \lambda v \rangle)}] \\ &= q^{-\ell} \cdot \left(\Pr_{v \sim \tau_M} [\langle v, y \rangle \neq 0] \cdot \mathbb{E}_{\xi \in \mathbb{F}_q^*} [\omega_p^{\text{tr}(\xi)}] + \Pr_{v \sim \tau_M} [\langle v, y \rangle = 0] \cdot \mathbb{E}_{\lambda \in \mathbb{F}_q^*} [\omega_p^{\text{tr}(0)}] \right) \\ &= q^{-\ell} \cdot \left(\Pr_{v \sim \tau_M} [\langle v, y \rangle \neq 0] \cdot \frac{-1}{q-1} + \Pr_{v \sim \tau_M} [\langle v, y \rangle = 0] \cdot 1 \right) \\ &\leq q^{-\ell} \cdot \left(\frac{-\delta}{q-1} + (1-\delta) \right) = q^{-\ell} \cdot \left(1 - \frac{q}{q-1} \cdot \delta\right), \end{aligned}$$

where the last inequality follows by assumption that τ_M is δ -smooth. □

This completes the proof of Lemma 2.13.

5 Random LDPC codes achieve the GV bound: proof of Theorem 2.14

In this section we prove Theorem 2.14, which shows that an LDPC code over any alphabet approaches the Gilbert-Varshamov bound with high probability. We restate the theorem below.

Theorem 2.14 (Random LDPC codes achieve the GV bound). *For any $\delta \in (0, 1 - 1/q)$, $\varepsilon > 0$, and prime power q there exists $s_0 \geq 1$ such that the following holds for any $s \geq s_0$. Let $R \leq 1 - h_q(\delta) - \varepsilon$. Then a random s -LDPC code of rate R over \mathbb{F}_q has relative distance at least δ with high probability.*

5.1 Proof of Theorem 2.14, given a lemma

In this section we give an outline of the proof of Theorem 2.14 and prove the theorem based on Lemma 5.2 that we state below and prove in subsequent subsections.

Our goal is to show that a random s -LDPC code C has good distance, or equivalently that there are no low-weight codewords in C with high probability. To that end, we introduce the following notation.

Definition 5.1. *For $\lambda \in (0, 1)$ such that λn is an integer, let $P_\lambda = \Pr[u \in C]$, where $u \in \mathbb{F}_q^n$ has weight λ . Note that this probability is the same for every u of weight λ , so P_λ is well-defined.*

Our main challenge is to find sufficiently tight upper bounds on these terms P_λ for $0 < \lambda \leq \delta$. The proof proceeds by giving a bound on P_λ in terms of a certain function $\varphi : (0, \frac{q-1}{q}] \rightarrow \mathbb{R}_{\leq 0}$. We will prove the following lemma below in Sections 5.2 and 5.3. We will define φ below in Section 5.2, but for now we introduce its important properties in the following lemma (which we also prove below).

Lemma 5.2. *There is a function $\varphi : (0, \frac{q-1}{q}] \rightarrow \mathbb{R}_{\leq 0}$ which has the following properties.*

1. For every $\lambda \in (0, \frac{q-1}{q}]$,

$$\log_q P_\lambda \leq \varphi(\lambda)(1 - R)n.$$

2. The function φ satisfies

$$\varphi(\lambda) \leq \log_q \left(1 + (q-1) \left(1 - \frac{q}{q-1} \lambda \right)^s \right) - 1$$

for all $\lambda \in (0, \frac{q-1}{q}]$.

3. The function $\frac{\varphi(\lambda)}{h_q(\lambda)}$ is strictly increasing in the range $0 < \lambda \leq \frac{q-1}{q}$.

Before we prove Lemma 5.2, we show how it implies Theorem 2.14.

Proof of Theorem 2.14. Our goal is to show that if C is a random s -LDPC code as in the statement of Theorem 2.14, then with high probability there are no codewords in C of relative weight less

than δ . In the following, we assume without loss of generality that δn is an integer. Now

$$\Pr[C \text{ has relative distance less than } \delta] \leq \sum_{i=1}^{\delta n} P_{\frac{i}{n}} \left| \left\{ u \in \mathbb{F}_q^n \mid \text{wt}(u) = \frac{i}{n} \right\} \right| \quad (14)$$

$$\begin{aligned} &\leq \sum_{i=1}^{\delta n} P_{\frac{i}{n}} q^{nh_q(\frac{i}{n})} \\ &\leq \sum_{i=1}^{\delta n} q^{(\varphi(\frac{i}{n})(1-R)+h_q(\frac{i}{n}))n} \end{aligned} \quad (15)$$

$$= \sum_{i=1}^{\delta n} q^{nh_q(\frac{i}{n}) \left(\frac{(1-R)\varphi(\frac{i}{n})}{h_q(\frac{i}{n})} + 1 \right)} \quad (16)$$

$$\leq \sum_{i=1}^{\delta n} q^{nh_q(\frac{i}{n}) \left(\frac{(1-R)\varphi(\delta)}{h_q(\delta)} + 1 \right)}. \quad (17)$$

Above, (14) follows from the union bound, (15) from Item 1 of Lemma 5.2, and (17) from Item 3 of Lemma 5.2. By Item 2 of Lemma 5.2,

$$\frac{(1-R)\varphi(\delta)}{h_q(\delta)} + 1 = \frac{(1-R) \cdot \left(\log_q \left(1 + (q-1) \left(1 - \frac{q}{q-1} \delta \right)^s \right) - 1 \right)}{h_q(\delta)} + 1.$$

Recall our hypothesis that the rate of the code satisfies $R \leq 1 - h_q(\delta) - \varepsilon$, and so $1 - R \geq h_q(\delta) + \varepsilon$. Noting that $\log_q \left(1 + (q-1) \left(1 - \frac{q}{q-1} \delta \right)^s \right) - 1 \leq 0$, we may thus bound the right hand side from above by

$$\begin{aligned} &\frac{(h_q(\delta) + \varepsilon) \cdot \left(\log_q \left(1 + (q-1) \left(1 - \frac{q}{q-1} \delta \right)^s \right) - 1 \right)}{h_q(\delta)} + 1 \\ &= \left(1 + \frac{\varepsilon}{h_q(\delta)} \right) \cdot \left(\log_q \left(1 + (q-1) \left(1 - \frac{q}{q-1} \delta \right)^s \right) - 1 \right) + 1 \\ &= \left(1 + \frac{\varepsilon}{h_q(\delta)} \right) \cdot \log_q \left(1 + (q-1) \left(1 - \frac{q}{q-1} \delta \right)^s \right) - \frac{\varepsilon}{h_q(\delta)} \\ &\leq \left(1 + \frac{\varepsilon}{h_q(\delta)} \right) \frac{(q-1)}{\ln(q)} \left(1 - \frac{q\delta}{q-1} \right)^s - \frac{\varepsilon}{h_q(\delta)}. \end{aligned}$$

Thus, as long as s is sufficiently large in terms of δ, ε and q , we conclude that

$$\frac{(1-R)\varphi(\delta)}{h_q(\delta)} + 1 \leq -\frac{\varepsilon}{2h_q(\delta)} \leq -\frac{\varepsilon}{2}.$$

Hence, the right-hand side of (17) is upper bounded by

$$\sum_{i=1}^{\delta n} q^{-\frac{nh_q(\frac{i}{n})\varepsilon}{2}}.$$

This sum is dominated by its first term, so it is at most $O(n^{-\Omega(1)})$. □

Remark 5.3 (The choice of s). *An inspection of the proof above shows that it suffices to take $s \gtrsim \ln(q/\varepsilon)/\delta$. Thus, this part of the proof requires that $s_0 \gtrsim \ln(q/\varepsilon)/\delta$.*

Remark 5.4 (Polynomially small failure probability). *In the proof, we see that the failure probability, while $o(1)$, is only polynomially small in n . In fact, this is tight: it is not hard to see that an s -random LDPC code C (for $s = O(1)$) contains a codeword of weight 2 with probability $n^{-O(1)}$.*

5.2 The function φ and proof of Lemma 5.2, Items 1 and 2

Let $\lambda \in \left(0, \frac{q-1}{q}\right]$ such that λn is an integer, and let $u \in \mathbb{F}_q^n$ have weight λn . Let H_1, \dots, H_t be the layers of the the parity-check matrix H of C , as in Figure 2. Note that the matrices H_1, \dots, H_t are identically and independently distributed. In particular, the events $\Pr(H_i u = 0)$ are independent. Hence,

$$P_\lambda = \Pr[u \in C] = \Pr[Hu = 0] = \Pr[H_1 u = 0]^t. \quad (18)$$

Since the distribution of H_1 is invariant to permutation of coordinates, this last probability does not depend on the vector u as long as it is of relative weight λ . Hence,

$$\Pr[H_1 u = 0] = \Pr[H_1 \bar{u} = 0] = \Pr[F\bar{u} = 0],$$

where \bar{u} is uniformly sampled from the set of all vectors of weight λ in \mathbb{F}_q^n . Therefore,

$$P_\lambda = \Pr[F\bar{u} = 0]^t,$$

where F is as in Figure 2.

We turn to bound this expression. Let $\beta \in \left(0, \frac{q-1}{q}\right]$. Denote by $\mu_q(\beta)$ the distribution on \mathbb{F}_q which is 0 with probability $1 - \beta$ and uniform on \mathbb{F}_q^* with probability β . When β is clear from context, we shorthand $\mu_q = \mu_q(\beta)$. Let $v \in \mathbb{F}_q^n$ be a random vector whose entries are i.i.d. random variables sampled according to μ_q , which we denote by $v \sim \mu_q^n$. Observe that the distribution of v , conditioned on $\text{wt}(v) = \lambda$, is identical to the distribution of \bar{u} . Hence, by Bayes' rule,

$$\Pr[F\bar{u} = 0] = \Pr[Fv = 0 \mid \text{wt}(v) = \lambda] = \Pr[\text{wt}(v) = \lambda \mid Fv = 0] \cdot \frac{\Pr[Fv = 0]}{\Pr[\text{wt}(v) = \lambda]} \leq \frac{\Pr[Fv = 0]}{\Pr[\text{wt}(v) = \lambda]} \quad (19)$$

where the probabilities are over the choice of $v \sim \mu_q(\beta)^n$.

We proceed to bound the right-hand side of (19). For the denominator, note that

$$\Pr[\text{wt}(v) = \lambda] = \binom{n}{\lambda n} \beta^{\lambda n} (1 - \beta)^{(1-\lambda)n} \geq q^{-D_{\text{KL}_q}(\lambda \parallel \beta)n} \quad (20)$$

where above $D_{\text{KL}_q}(x \parallel y)$ denotes the KL Divergence,

$$D_{\text{KL}_q}(x \parallel y) = -x \log_q \frac{y}{x} - (1-x) \log_q \frac{1-y}{1-x} \text{ for } x \in [0, 1] \text{ and } y \in (0, 1).$$

We next focus on the numerator. The following notation will be useful:

Definition 5.5. *For $k \in \mathbb{N}$, let*

$$\mathbb{V}_q^k = \left\{ w \in \mathbb{F}_q^k : \sum_{i=1}^k w_i = 0 \right\}.$$

Let $f_1, \dots, f_{\frac{n}{s}}$ denote the rows of the matrix F . Note that the vectors $f_1, \dots, f_{\frac{n}{s}}$ have disjoint supports, so the products $f_i v$ are independently and identically distributed. Hence, $\Pr[Fv = 0] = \Pr[f_1 v = 0]^{\frac{n}{s}}$. Observe that the distribution of v is symmetric to multiplication of each entry by a nonzero element of \mathbb{F}_q . Consequently,

$$\Pr_{v \sim \mu_q^n} [Fv = 0] = \Pr_{v \sim \mu_q^n} [f_1 v = 0]^{\frac{n}{s}} = \Pr_{v \sim \mu_q^n} \left[\sum_{i=1}^s v_i = 0 \right]^{\frac{n}{s}} = \left(\Pr_{w \sim \mu_q^s} [w \in \mathbb{V}_q^s] \right)^{n/s}. \quad (21)$$

The following lemma gives a closed form for this last expression.

Lemma 5.6.

$$\Pr_{w \sim \mu_q^s} [w \in \mathbb{V}_q^s] = \frac{1 + (q-1) \left(1 - \frac{q\beta}{q-1}\right)^s}{q}.$$

Proof. We proceed by induction. The base case ($s = 0$) is immediate. Now suppose that the statement holds for $s-1$ and let $\pi : \mathbb{F}_q^s \rightarrow \mathbb{F}_q^{s-1}$ denote the projection onto the first $s-1$ coordinates. Then

$$\begin{aligned} \Pr_{w \sim \mu_q^s} [w \in \mathbb{V}_q^s] &= \Pr_{w \sim \mu_q^s} [\pi(w) \in \mathbb{V}_q^{s-1}] \cdot \Pr_{w \sim \mu_q^s} [w_s = 0] + \Pr_{w \sim \mu_q^s} [\pi(w) \notin \mathbb{V}_q^{s-1}] \cdot \Pr_{w \sim \mu_q^s} \left[w_s = - \sum_{i=1}^{s-1} w_i \mid \pi(w) \notin \mathbb{V}_q^{s-1} \right] \\ &= \frac{1 + (q-1) \left(1 - \frac{q\beta}{q-1}\right)^{s-1}}{q} \cdot (1 - \beta) + \left(1 - \frac{1 + (q-1) \left(1 - \frac{q\beta}{q-1}\right)^{s-1}}{q} \right) \cdot \frac{\beta}{q-1} \\ &= \frac{1}{q} + \left(1 - \frac{q\beta}{q-1} \right)^s \left(\frac{q-1}{q} \right), \end{aligned}$$

which establishes the inductive hypothesis for s . □

Motivated by the computations above, we can define the following useful shorthands:

Definition 5.7. For $\lambda, \beta \in (0, \frac{q-1}{q}]$, define

$$Z(\beta) = \Pr_{w \sim \mu_q^s} [w \in \mathbb{V}_q^s] = \frac{1 + (q-1) \left(1 - \frac{q\beta}{q-1}\right)^s}{q}, \quad (22)$$

$$\psi(\lambda, \beta) = sD_{\text{KL}q}(\lambda \parallel \beta) + \log_q Z(\beta)$$

From Equations (18), (19), (20) and (21), we conclude that

$$\begin{aligned} \log_q P_\lambda &= t \log_q \Pr[F\bar{u} = 0] \leq tn \left(D_{\text{KL}q}(\lambda \parallel \beta) + \frac{\log_q \left(1 + (q-1) \left(1 - \frac{q\beta}{q-1}\right)^s \right) - 1}{s} \right) \\ &= (1-R)n \left(sD_{\text{KL}q}(\lambda \parallel \beta) + \log_q \left(1 + (q-1) \left(1 - \frac{q\beta}{q-1}\right)^s \right) - 1 \right) \\ &= (1-R)n\psi(\lambda, \beta) \end{aligned} \quad (23)$$

for every $\beta \in (0, \frac{q-1}{q}]$. Above, we have used the choice $t = (1-R)s$.

This motivates the following definition:

Definition 5.8. Let Z and ψ be as in Definition 5.7. Define:

$$\varphi(\lambda) = \inf_{\beta \in (0, \frac{q-1}{q}]} \psi(\lambda, \beta).$$

Definition 5.8, along with (23), implies that $\log_q P_\lambda \leq \varphi(\lambda)$, which establishes Item 2 of Lemma 5.2. Next we establish Item 1 of Lemma 5.2. This follows from Definition 5.8, since

$$\varphi(\lambda) \leq \psi(\lambda, \lambda) = \log_q \left(1 + (q-1) \left(1 - \frac{q\lambda}{q-1} \right)^s \right) - 1,$$

using the fact that $D_{\text{KL}q}(\lambda \parallel \lambda) = 0$.

This almost completes the proof of Lemma 5.2, except for Item 3, which we establish in the next section using calculus.

5.3 Proof of Item 3 of Lemma 5.2

In this section we prove Item 3, which will establish Lemma 5.2 and hence Theorem 2.14.

Remark 5.9 (Difference between [Gal62] and this proof). *This is the part of the proof where the technical similarity between our proof and Gallager's breaks down. The part of [Gal62] which corresponds to our Item 3 consists of an intricate analytic argument which does not seem (to us) to generalize to larger alphabets. Thus, our proof has to rely on a different, more general, argument, which we give below.*

Before proving Item 3 of Lemma 5.2, we need to better understand the relation between a given $\lambda \in (0, \frac{q-1}{q}]$, and the β which minimizes the expression $\psi(\lambda, \beta)$.

Lemma 5.10. *Let $\lambda \in (0, \frac{q-1}{q}]$. Then, $\psi(\lambda, \beta)$ is minimized by a unique $\beta \in (0, \frac{q-1}{q}]$. This β is the only solution for*

$$\mathbb{E}_{w \sim \mu_q(\beta)} [\text{wt}(w) \mid w \in \mathbb{V}_q^s] = \lambda.$$

Proof. We compute the derivative.

$$\begin{aligned} \frac{d \log_e Z(\beta)}{d\beta} &= \frac{1}{\Pr_{w \sim \mu_q^s}[w \in \mathbb{V}_q^s]} \cdot \frac{d \left(\Pr_{w \sim \mu_q^s}[w \in \mathbb{V}_q^s] \right)}{d\beta} \\ &= \frac{1}{\Pr_{w \sim \mu_q^s}[w \in \mathbb{V}_q^s]} \cdot \sum_{w \in \mathbb{V}_q^s} \frac{d \left(\frac{\beta}{q-1} \right)^{s \cdot \text{wt}(w)} (1-\beta)^{s \cdot (1-\text{wt}(w))}}{d\beta} \\ &= \frac{\sum_{w \in \mathbb{V}_q^s} \left(\left(\frac{\beta}{q-1} \right)^{s \cdot \text{wt}(w)} (1-\beta)^{s \cdot (1-\text{wt}(w))} \cdot s \cdot \left(\frac{\text{wt}(w)}{\beta} - \frac{1-\text{wt}(w)}{1-\beta} \right) \right)}{\Pr_{w \sim \mu_q^s}[w \in \mathbb{V}_q^s]} \\ &= s \cdot \left(\frac{\mathbb{E}_{w \sim \mu_q^s} [\text{wt}(w) \mid w \in \mathbb{V}_q^s]}{\beta} - \frac{1 - \mathbb{E}_{w \sim \mu_q^s} [\text{wt}(w) \mid w \in \mathbb{V}_q^s]}{1-\beta} \right). \end{aligned} \quad (24)$$

Also, it is not hard to see that

$$\frac{\partial D_{\text{KL}q}(\lambda \parallel \beta)}{\partial \beta} = \log_q e \cdot \left(\frac{1-\lambda}{1-\beta} - \frac{\lambda}{\beta} \right).$$

Consequently,

$$\begin{aligned}
\frac{\partial \psi(\lambda, \beta)}{\partial \beta} &= s \frac{\partial D_{\text{KL}q}(\lambda \parallel \beta)}{\partial \beta} + \frac{d \log_q Z(\beta)}{d\beta} \\
&= \log_q e \cdot \left(\frac{s(1-\lambda)}{1-\beta} - \frac{s\lambda}{\beta} + \frac{d \log_e Z(\beta)}{d\beta} \right) \\
&= s \cdot \log_q e \cdot \left(\mathbb{E}_{w \sim \mu_q^s} [\text{wt}(w) \mid w \in \mathbb{V}_q^s] - \lambda \right) \left(\frac{1}{1-\beta} + \frac{1}{\beta} \right).
\end{aligned}$$

We conclude that $\frac{\partial \psi(\lambda, \beta)}{\partial \beta}$ has the same sign as $\mathbb{E}_{w \sim \mu_q^s} [\text{wt}(w) \mid w \in \mathbb{V}_q^s] - \lambda s$. The lemma now follows from the following claim:

Claim 5.11. *As β increases in the range $(0, \frac{q-1}{q}]$ the function $\mathbb{E}_{w \sim \mu_q^s} [\text{wt}(w) \mid w \in \mathbb{V}_q^s]$ strictly increases from 0 to $\frac{q-1}{q}$.*

Proof. Due to (22) and (24),

$$\begin{aligned}
\mathbb{E}_{w \sim \mu_q^s} [\text{wt}(w) \mid w \in \mathbb{V}_q^s] &= \left(\frac{d \log_e Z(\beta)}{s \cdot d\beta} + \frac{1}{1-\beta} \right) \beta(1-\beta) \\
&= \left(\frac{\frac{dZ(\beta)}{d\beta}}{s \cdot Z(\beta)} + \frac{1}{1-\beta} \right) \beta(1-\beta) \\
&= \left(\frac{-q \left(1 - \frac{q\beta}{q-1}\right)^{s-1}}{1 + (q-1) \left(1 - \frac{q\beta}{q-1}\right)^s} + \frac{1}{1-\beta} \right) \beta(1-\beta) \\
&= \beta \frac{1 - \left(1 - \frac{q\beta}{q-1}\right)^{s-1} \cdot (1+q\beta)}{1 + (q-1) \left(1 - \frac{q\beta}{q-1}\right)^s}, \tag{25}
\end{aligned}$$

and the claim readily follows. □

The proof of the lemma is thus concluded. □

Lemma 5.10 and Claim 5.11 justify the following definition:

Definition 5.12. *For $\lambda \in (0, \frac{q-1}{q}]$, denote the $\beta \in (0, \frac{q-1}{q}]$ which minimizes $\psi(\lambda, \beta)$ by $\beta(\lambda)$. The inverse of this function is denoted $\lambda(\beta)$.*

By Lemma 5.10 and Equation (25),

$$\lambda(\beta) = \beta \frac{1 - \left(1 - \frac{q\beta}{q-1}\right)^{s-1}}{1 + (q-1) \left(1 - \frac{q\beta}{q-1}\right)^s}. \tag{26}$$

Remark 5.13. *Unfortunately, there are good reasons to suspect that the function $\beta(\lambda)$ has no closed-form expression (see, e.g., the discussion about backward mapping in [WJ08, Sec. 3.4.2]), so we prefer to work with its inverse.*

It is convenient to extend the definition of these functions to the closed interval $[0, \frac{q-1}{q}]$ by taking limits, namely, $\lambda(0) = \beta(0) = 0$, and

$$\begin{aligned}\varphi(0) &= \lim_{\lambda \rightarrow 0} \varphi(\lambda) = \lim_{\lambda \rightarrow 0} \psi(\lambda, \beta(\lambda)) \lim_{\beta \rightarrow 0} \psi(\lambda(\beta), \beta) = \lim_{\beta \rightarrow 0} D_{\text{KL}_q}(\lambda(\beta) \parallel \beta) + \log_q Z(\beta) \\ &= \lim_{\beta \rightarrow 0} D_{\text{KL}_q}(\lambda(\beta) \parallel \beta) = \lim_{\beta \rightarrow 0} -\lambda(\beta) \log_q \beta = 0.\end{aligned}$$

We are now able to prove Item 3 of Lemma 5.2.

Proof of Lemma 5.2, Item 3. Let $\alpha(\lambda) = \frac{\varphi(\lambda)}{h_q(\lambda)}$. The claim follows immediately from the four following claims:

Claim 5.14. $\alpha(\frac{q-1}{q}) = -1$.

Claim 5.15. $\alpha(\lambda) < -1$ for some $\lambda \in (0, \frac{q-1}{q})$.

Claim 5.16. There exists $\varepsilon > 0$ such that $\alpha(\lambda) > -\frac{s}{2}$ for all $\lambda \in (0, \varepsilon)$.

Claim 5.17. For each $y \in (-\frac{s}{2}, -1]$, the equation $\alpha(\lambda) = y$ has at most one solution $\lambda \in (0, \frac{q-1}{q}]$.

Indeed, Claims 5.14 and 5.17 show that $\alpha(\lambda) \neq 1$ for $\lambda < \frac{q-1}{q}$. Since α is continuous, it is either upper bounded or lower bounded by -1 in the whole range $(0, \frac{q-1}{q}]$. Claim 5.15 implies the former. By Claim 5.17, if $-\frac{s}{2} < \alpha(\lambda_0) < -1$ for some $\lambda_0 \in (0, \frac{q-1}{q})$, then α must be strictly increasing in the range $[\lambda_0, \frac{q-1}{q}]$. The lemma now follows from Claim 5.16. We proceed to prove these claims.

Proof of Claim 5.14. Note that $\alpha(\frac{q-1}{q}) = \varphi(\frac{q-1}{q})$. Due to Item 2,

$$\varphi\left(\frac{q-1}{q}\right) \leq -1.$$

In the reverse direction,

$$\begin{aligned}\varphi(\lambda) &= \min_{\beta} \psi(\lambda, \beta) = \min_{\beta} (s \cdot D_{\text{KL}_q}(\lambda \parallel \beta) + \log_q Z(\beta)) \\ &\geq \min_{\beta} (s \cdot D_{\text{KL}_q}(\lambda \parallel \beta)) - 1 \geq -1\end{aligned}$$

for all λ . The first inequality above holds since $Z(\beta) \geq \frac{1}{q}$, due to (22). \square

Proof of Claim 5.15. By Item 1,

$$\alpha(\lambda) \leq \frac{\log_q \left(1 + (q-1) \left(1 - \frac{q}{q-1} \lambda\right)^s\right) - 1}{h_q(\lambda)}. \quad (27)$$

Let $\lambda = \frac{q-1}{q} - \varepsilon$. As ε tends from above to 0, the numerator of (27)'s right-hand side is $-1 + \Theta(\varepsilon^s)$, while the denominator is $1 - \Theta(\varepsilon^2)$. Thus, for ε small enough, (27) yields $\alpha(\lambda) < -1$. \square

Proof of Claim 5.16. Let

$$\bar{Z}(\beta) = \Pr_{w \sim \mathcal{B}_q^s} \left(w \in \mathbb{V}_q^s \wedge \text{wt}(w) \leq \frac{2}{s} \right) = (1 - \beta)^s + \binom{s}{2} (1 - \beta)^{s-2} \beta^2$$

and

$$\bar{\psi}(\beta, \lambda) = s D_{\text{KL}_q}(\lambda \parallel \beta) + \log_q \bar{Z}(\beta).$$

Clearly, $\bar{\psi}(\beta, \lambda)$ is a lower bound on $\psi(\beta, \lambda)$, so

$$\varphi(\lambda) \geq \min_{\beta \in (0, \frac{q-1}{q}]} \bar{\psi}(\lambda, \beta).$$

Note that

$$\frac{\partial \bar{\psi}(\lambda, \beta)}{\partial \beta} = \frac{s}{\beta(1 - \beta)} \left(\frac{2(s-1)}{\left(\frac{1-\beta}{\beta}\right)^2 + \binom{s}{2}} - \lambda \right),$$

Hence, for $\lambda < \frac{2}{s}$, the minimum of $\bar{\psi}(\lambda, \beta)$ is attained at $\beta_0 = \frac{y}{1+y}$, where

$$y = \left(\frac{\lambda}{2(s-1) - \binom{s}{2}\lambda} \right)^{\frac{1}{2}}.$$

Therefore,

$$\alpha(\lambda) = \frac{\varphi(\lambda)}{h_q(\lambda)} \geq \frac{\bar{\psi}(\lambda, \beta_0)}{h_q(\lambda)} = \frac{s}{2} \left(-1 + \frac{\lambda (\log_q(2(s-1) - \binom{s}{2}\lambda) - \log_q(1 - \lambda s)) + (1 - \lambda) \log_q(1 - \lambda)}{h_q(\lambda)} \right).$$

For λ small enough, the right-hand side is clearly larger than $-\frac{s}{2}$. \square

Proof of Claim 5.17. Denote $\beta^* = \beta(\lambda)$. Let $y \in (-\frac{s}{2}, -1]$, and define the function $\varphi_y(\lambda) = \varphi(\lambda) - y h_q(\lambda)$. We seek to show that $\varphi_y(\lambda)$ has at most one root in the range $(0, \frac{q-1}{q}]$. This is a consequence of the following three statements, proven below:

1. $\frac{d\varphi_y(\lambda)}{d\lambda}$ has at most one extremal point in the open interval $(0, \frac{q-1}{q})$.
2. $\frac{d\varphi_y(\lambda)}{d\lambda}(\frac{q-1}{q}) = 0$.
3. $\varphi_y(0) = 0$.

Indeed, the first statement implies that $\frac{d\varphi_y(\lambda)}{d\lambda}$ has at most two roots in the interval $(0, \frac{q-1}{q}]$. The second statement says that one of these roots is at $\frac{q-1}{q}$, so $\frac{d\varphi_y(\lambda)}{d\lambda}$ has at most one root in $(0, \frac{q-1}{q})$. Consequently $\varphi_y(\lambda)$ has at most one extremal point and two roots in $[0, \frac{q-1}{q}]$. Due to the third statement, one of these roots is 0, so there can only be one root in $(0, \frac{q-1}{q}]$. We turn to prove these statements.

Statement 3 is trivial. For Statement 2, note that in the derivative

$$\frac{d\varphi(\lambda)}{d\lambda} = \frac{\partial \psi(\lambda, \beta)}{\partial \beta} \Big|_{\beta=\beta^*} \cdot \frac{d\beta^*}{d\lambda} + \frac{\partial \psi(\lambda, \beta)}{\partial \lambda} \Big|_{\beta=\beta^*},$$

the first term vanishes since ψ has a minimum at (λ, β^*) . Hence,

$$\frac{d\varphi(\lambda)}{d\lambda} = \frac{\partial\psi(\lambda, \beta)}{\partial\lambda} \Big|_{\beta=\beta^*} = s \frac{\partial D_{\text{KL}q}(\lambda \parallel \beta)}{\partial\lambda} \Big|_{\beta=\beta^*} = s \log_q \frac{\lambda(1-\beta^*)}{(1-\lambda)\beta^*}.$$

In particular, $\beta(\frac{q-1}{q}) = \frac{q-1}{q}$, so

$$\frac{d\varphi_y(\lambda)}{d\lambda} \Big|_{\lambda=\frac{q-1}{q}} = \frac{d\varphi(\lambda)}{d\lambda} \Big|_{\lambda=\frac{q-1}{q}} - y \frac{dh_q(\lambda)}{d\lambda} \Big|_{\lambda=\frac{q-1}{q}} = 0,$$

since, in the last transition, the two terms vanish.

We turn to Statement 1. Define the new variable $x = 1 - \frac{q\beta^*}{q-1}$. Note the following useful relations, the second of which follows from Equation (26):

$$\beta^* = \frac{q-1}{q}(1-x) \tag{28}$$

and

$$\frac{\lambda}{1-\lambda} = \frac{\beta^*}{1-\beta^*} \cdot \frac{1-x^{s-1}}{1+(q-1)x^{s-1}}. \tag{29}$$

By (28) and (29),

$$\begin{aligned} \frac{d\varphi_y(\lambda)}{d\lambda} &= s \frac{\partial D_{\text{KL}q}(\lambda \parallel \beta)}{\partial\lambda} \Big|_{\beta=\beta^*} - y \frac{dh_q(\lambda)}{d\lambda} \\ &= s \log_q \frac{\lambda(1-\beta^*)}{(1-\lambda)\beta^*} + y \log_q \frac{\lambda}{1-\lambda} \\ &= s \log_q \frac{1-\beta^*}{\beta^*} + (s+y) \log_q \frac{\lambda}{1-\lambda} \\ &= -y \log_q \frac{1+(q-1)x}{(q-1)(1-x)} + (s+y) \log_q \frac{1-x^{s-1}}{1+(q-1)x^{s-1}}. \end{aligned}$$

Now,

$$\frac{d^2\varphi_y(\lambda)}{dx d\lambda} \cdot \ln q = \frac{-yq}{(1+(q-1)x)(1-x)} - \frac{(s+y)(s-1)qx^{s-2}}{(1-x^{s-1})(1+(q-1)x^{s-1})}.$$

This second derivative vanishes when

$$\frac{-(s+y)}{y} = \frac{(1-x^{s-1})(1+(q-1)x^{s-1})}{(s-1)(1+(q-1)x)(1-x)x^{s-2}}.$$

Equivalently,

$$\frac{-(s+y)}{y} = \frac{1}{s-1} \sum_{i=0}^{s-2} \frac{x^{-i} + (q-1)x^{i+1}}{1+(q-1)x}. \tag{30}$$

By examining each term of this sum separately, it is straightforward to verify that the right-hand side of (30) is a convex function of x , which tends to ∞ (resp. 1) as $x \rightarrow 0$ (resp. $x \rightarrow 1$). Since $y > -\frac{s}{2}$, the left-hand side of (30) is larger than 1, so there is a unique $x \in (0, 1)$ which solves (30). Statement 1 follows. \square

This establishes Item 3 of Lemma 5.2. \square

Having completed the proof of Lemma 5.2, we have finished the proof of Theorem 2.14.

Acknowledgements

The first author would like to thank Yael Hachohen and Nati Linial for useful conversations. The second author would like to thank Venkat Guruswami for helpful feedback on a draft of this work.

References

- [AEL95] Noga Alon, Jeff Edmonds, and Michael Luby. Linear time erasure codes with nearly optimal recovery. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pages 512–519. IEEE, 1995.
- [Bol01] Béla Bollobás. *Random Graphs, Second Edition*, volume 73 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 2001.
- [CGV13] Mahdi Cheraghchi, Venkatesan Guruswami, and Ameya Velingker. Restricted isometry of fourier matrices and list decodability of random linear codes. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 432–442, 2013.
- [CS⁺04] Imre Csiszár, Paul C Shields, et al. Information theory and statistics: A tutorial. *Foundations and Trends® in Communications and Information Theory*, 1(4):417–528, 2004.
- [DHK⁺19] Irit Dinur, Prahladh Harsha, Tali Kaufman, Inbal Livni Navon, and Amnon Ta Shma. List decoding with double samplers. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2134–2153. SIAM, 2019.
- [DL12] Zeev Dvir and Shachar Lovett. Subspace evasive sets. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 351–358. ACM, 2012.
- [Eli57] Peter Elias. List decoding for noisy channels. *Wescon Convention Record, Part 2*, pages 94–104, 1957.
- [FKNP19] Keith Frankston, Jeff Kahn, Bhargav Narayanan, and Jinyoung Park. Thresholds versus fractional expectation-thresholds. *arXiv preprint arXiv:1910.13433*, 2019.
- [Gal62] Robert G. Gallager. Low-density parity-check codes. *IRE Trans. Information Theory*, 8(1):21–28, 1962.
- [GHK11] Venkatesan Guruswami, Johan Håstad, and Swastik Kopparty. On the list-decodability of random linear codes. *IEEE Trans. Information Theory*, 57(2):718–725, 2011.
- [GHSZ02] Venkatesan Guruswami, Johan Håstad, Madhu Sudan, and David Zuckerman. Combinatorial bounds for list decoding. *IEEE Trans. Information Theory*, 48(5):1021–1034, 2002.
- [GK16] Venkatesan Guruswami and Swastik Kopparty. Explicit subspace designs. *Combinatorica*, 36(2):161–185, 2016.

- [GLM⁺20] Venkatesan Guruswami, Ray Li, Jonathan Mosheiff, Nicolas Resch, Shashwat Silas, and Mary Wootters. Bounds for list-decoding and list-recovery of random linear codes. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2020)*, 2020.
- [GR08] Venkatesan Guruswami and Atri Rudra. Explicit codes achieving list decoding capacity: Error-correction with optimal redundancy. *IEEE Transactions on Information Theory*, 54(1):135–150, 2008.
- [Gur03] Venkatesan Guruswami. List decoding from erasures: Bounds and code constructions. *IEEE Transactions on Information Theory*, 49(11):2826–2833, 2003.
- [Gur06] Venkatesan Guruswami. Iterative decoding of low-density parity check codes (a survey). *arXiv preprint cs/0610022*, 2006.
- [GW13] Venkatesan Guruswami and Carol Wang. Linear-algebraic list decoding for variants of reed–solomon codes. *IEEE Transactions on Information Theory*, 59(6):3257–3268, 2013.
- [GX12] Venkatesan Guruswami and Chaoping Xing. Folded codes from function field towers and improved optimal rate list decoding. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 339–350. ACM, 2012.
- [GX13] Venkatesan Guruswami and Chaoping Xing. List decoding Reed-Solomon, Algebraic-Geometric, and Gabidulin subcodes up to the Singleton bound. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 843–852. ACM, 2013.
- [HRW17] Brett Hemenway, Noga Ron-Zewi, and Mary Wootters. Local list recovery of high-rate tensor codes & applications. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 204–215. IEEE, 2017.
- [HW18] Brett Hemenway and Mary Wootters. Linear-time list recovery of high-rate expander codes. *Information and Computation*, 261:202–218, 2018.
- [Kop15] Swastik Kopparty. List-decoding multiplicity codes. *Theory of Computing*, 11(5):149–182, 2015.
- [KRRZ⁺19] Swastik Kopparty, Nicolas Resch, Noga Ron-Zewi, Shubhangi Saraf, and Shashwat Silas. On list recovery of high-rate tensor codes. *Electronic Colloquium on Computational Complexity (ECCC)*, 2019.
- [KRSW18] Swastik Kopparty, Noga Ron-Zewi, Shubhangi Saraf, and Mary Wootters. Improved decoding of folded reed-solomon and multiplicity codes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 212–223. IEEE, 2018.
- [KRU13] Shrinivas Kudekar, Tom Richardson, and Rüdiger L Urbanke. Spatially coupled ensembles universally achieve capacity under belief propagation. *IEEE Transactions on Information Theory*, 59(12):7761–7813, 2013.

- [LM20] Nati Linial and Jonathan Mosheiff. On the weight distribution of random binary linear codes. *Random Structures & Algorithms*, 56(1):5–36, 2020.
- [LN94] Rudolf Lidl and Harald Niederreiter. *Introduction to finite fields and their applications*. Cambridge university press, 1994.
- [LW18] Ray Li and Mary Wootters. Improved list-decodability of random linear binary codes. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- [O’D14] Ryan O’Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.
- [RW14] Atri Rudra and Mary Wootters. Every list-decodable code for high noise has abundant near-optimal rate puncturings. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 764–773. ACM, 2014.
- [RW18] Atri Rudra and Mary Wootters. Average-radius list-recovery of random linear codes. In *Proceedings of the 2018 ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2018.
- [SS94] Michael Sipser and Daniel A Spielman. Expander codes. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 566–576. IEEE, 1994.
- [Tan81] R Tanner. A recursive approach to low complexity codes. *IEEE Transactions on information theory*, 27(5):533–547, 1981.
- [WJ08] Martin J. Wainwright and Michael I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, 2008.
- [Woo13] Mary Wootters. On the list decodability of random linear codes with large error rates. In *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 853–860, 2013.
- [Woz58] Jack Wozencraft. List decoding. *Quarter Progress Report*, 48:90–95, 1958.
- [Zém01] Gillés Zémor. On expander codes. *IEEE Transactions on Information Theory*, 47(2):835–837, 2001.
- [ZP81] Victor Vasilievich Zyablov and Mark Semenovich Pinsker. List concatenated decoding. *Problemy Peredachi Informatsii*, 17(4):29–33, 1981.