



# Randomly Punctured Reed–Solomon Codes Achieve List-Decoding Capacity over Linear-Sized Fields\*

Omar Alrabiah<sup>†</sup>

UC Berkeley

Berkeley, USA

oalrabiah@berkeley.edu

Venkatesan Guruswami<sup>‡</sup>

UC Berkeley

Berkeley, USA

venkatg@berkeley.edu

Ray Li<sup>§</sup>

Santa Clara University

Santa Clara, USA

rayyli@cs.stanford.edu

## ABSTRACT

Reed–Solomon codes are a classic family of error-correcting codes consisting of evaluations of low-degree polynomials over a finite field on some sequence of distinct field elements. They are widely known for their optimal unique-decoding capabilities, but their list-decoding capabilities are not fully understood. Given the prevalence of Reed–Solomon codes, a fundamental question in coding theory is determining if Reed–Solomon codes can optimally achieve list-decoding capacity.

A recent breakthrough by Brakensiek, Gopi, and Makam, established that Reed–Solomon codes are combinatorially list-decodable all the way to capacity. However, their results hold for randomly-punctured Reed–Solomon codes over an exponentially large field size  $2^{O(n)}$ , where  $n$  is the block length of the code. A natural question is whether Reed–Solomon codes can still achieve capacity over smaller fields. Recently, Guo and Zhang showed that Reed–Solomon codes are list-decodable to capacity with field size  $O(n^2)$ . We show that Reed–Solomon codes are list-decodable to capacity with linear field size  $O(n)$ , which is optimal up to the constant factor. We also give evidence that the ratio between the alphabet size  $q$  and code length  $n$  cannot be bounded by an absolute constant.

Our techniques also show that random linear codes are list-decodable up to (the alphabet-independent) capacity with optimal list-size  $O(1/\varepsilon)$  and near-optimal alphabet size  $2^{O(1/\varepsilon^2)}$ , where  $\varepsilon$  is the gap to capacity. As far as we are aware, list-decoding up to capacity with optimal list-size  $O(1/\varepsilon)$  was not known to be achievable with any linear code over a constant alphabet size (even non-constructively), and it was also not known to be achievable for random linear codes over any alphabet size.

Our proofs are based on the ideas of Guo and Zhang, and we additionally exploit symmetries of reduced intersection matrices.

\*A full version of this work appears at [1].

<sup>†</sup>Research supported in part by a Saudi Arabian Cultural Mission (SACM) Scholarship, NSF CCF-2210823, and V. Guruswami's Simons Investigator Award.

<sup>‡</sup>Research supported by a Simons Investigator Award and NSF grants CCF-2210823 and CCF-2228287.

<sup>§</sup>Research supported by the NSF Mathematical Sciences Postdoctoral Research Fellowships Program under Grant DMS-2203067, and a UC Berkeley Initiative for Computational Transformation award.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0383-6/24/06

<https://doi.org/10.1145/3618260.3649634>

With our proof, which maintains a hypergraph perspective of the list-decoding problem, we include an alternate presentation of ideas from Brakensiek, Gopi, and Makam that more directly connects the list-decoding problem to the GM-MDS theorem via a hypergraph orientation theorem.

## CCS CONCEPTS

- **Theory of computation** → **Error-correcting codes; Generating random combinatorial structures.**

## KEYWORDS

List-decoding, Reed–Solomon codes, Random puncturing

### ACM Reference Format:

Omar Alrabiah, Venkatesan Guruswami, and Ray Li. 2024. Randomly Punctured Reed–Solomon Codes Achieve List-Decoding Capacity over Linear-Sized Fields. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24), June 24–28, 2024, Vancouver, BC, Canada*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3618260.3649634>

## 1 INTRODUCTION

An (*error-correcting*) code is simply a set of strings (*codewords*). In this paper, all codes are *linear*, meaning our code  $C \subset \mathbb{F}_q^n$  is a space of vectors over a finite field  $\mathbb{F}_q$ , for some prime power  $q$ . A *Reed–Solomon code* [52] is a linear code obtained by evaluating low-degree polynomials over  $\mathbb{F}_q$ . More formally, we define the code  $\text{RS}_{n,k}(\alpha_1, \dots, \alpha_n)$  to be the set

$$\{(f(\alpha_1), \dots, f(\alpha_n)) \in \mathbb{F}_q^n : f \in \mathbb{F}_q[X], \deg(f) < k\}. \quad (1)$$

The *rate*  $R$  of a code  $C$  is  $R \stackrel{\text{def}}{=} \log_q |C|/n$ , which, for a Reed–Solomon code, is  $k/n$ . Famously, Reed–Solomon codes are optimal for the *unique decoding problem* [52]: for any rate  $R$  Reed–Solomon code, for every received word  $y \in \mathbb{F}_q^n$ , there is at most one codeword within Hamming distance  $pn$  of  $y$ ,<sup>1</sup> and this *error parameter*  $p = \frac{1-R}{2}$  is optimal by the *Singleton bound* [58].

In this paper, we study Reed–Solomon codes in the context of *list-decoding*, a generalization of unique-decoding that was introduced by Elias and Wozencraft [15, 62]. Formally, a code  $C \subset \mathbb{F}_q^n$  is  $(p, L)$ -list-decodable if, for every received word  $y \in \mathbb{F}_q^n$ , there are at most  $L$  codewords of  $C$  within Hamming distance  $pn$  of  $y$ .

It is well known that the *list-decoding capacity*, namely the largest fraction of errors that can be list-decoded with small lists, is  $1 - R$  [33, Theorem 7.4.1]. Specifically, for  $p = 1 - R - \varepsilon$ , there are (infinite families of) rate  $R$  codes that are  $(p, L)$  list-decodable for a list-size  $L$  as small as  $O(1/\varepsilon)$ . On the other hand, for  $p = 1 - R + \varepsilon$ , if a rate

<sup>1</sup>The Hamming distance between two codewords is the number of coordinates on which they differ.

$R$  code is  $(p, L)$  list decodable, the list size  $L$  must be exponential in the code length  $n$ . Informally, a code that is list-decodable up to radius  $p = 1 - R - \varepsilon$  with list size  $O_\varepsilon(1)$ , or even list size  $n^{O_\varepsilon(1)}$  where  $n$  is the code length, is said to *achieve (list-decoding) capacity*.

The list-decodability of Reed–Solomon codes is important for several reasons. Reed–Solomon codes are the most fundamental algebraic error-correcting code. In fact, all known explicit constructions of codes achieving list-decoding capacity are based on algebraic constructions that generalize Reed–Solomon codes, for example, Folded Reed–Solomon codes [32, 45], Multiplicity codes [36, 44, 45], and algebraic-geometric codes [14, 37–39]. Thus, it is natural to wonder whether and when Reed–Solomon codes themselves achieve list-decoding capacity. Additionally, all Reed–Solomon codes are optimally *unique-decodable*, so (equivalently) they are optimally list-decodable  $L = 1$ , making them a natural candidate for codes achieving list-decoding capacity. Further, capacity-achieving Reed–Solomon codes would potentially offer advantages over existing explicit capacity-achieving codes, such as simplicity and potentially smaller alphabet sizes (which we achieve in this work). Lastly, list-decoding of Reed–Solomon codes has found several applications in complexity theory and pseudorandomness [10, 48, 59].

For these reasons, the list-decodability of Reed–Solomon codes is well-studied. As rate  $R$  Reed–Solomon codes are uniquely decodable up to the optimal radius  $\frac{1-R}{2}$  given by the Singleton Bound, the Johnson-bound [42] automatically implies that Reed–Solomon codes are  $(p, L)$ -list-decodable for error parameter  $p = 1 - \sqrt{R} - \varepsilon$  and list size  $L = O(1/\varepsilon)$ . Guruswami and Sudan [34] showed how to *efficiently* list-decode Reed–Solomon codes up to the Johnson radius  $1 - \sqrt{R}$ . For a long time, this remained the best list-decodability result (even non-constructively) for Reed–Solomon codes.

Since then, several results suggested Reed–Solomon codes could *not* be list-decoded up to capacity, and in fact, not much beyond the Johnson radius  $1 - \sqrt{R}$ . Guruswami and Rudra [31] showed that, for a generalization of list-decoding called *list-recovery*, Reed–Solomon codes are not list-recoverable beyond the (list-recovery) Johnson bound in some parameter settings. Cheng and Wan [12] showed that efficient list-decoding of Reed–Solomon codes beyond the Johnson radius in certain parameter settings implies fast algorithms for the discrete logarithm problem. Ben-Sasson, Kopparty, and Radhakrishnan [3] showed that full-length Reed–Solomon codes ( $q = n$ ) are not list-decodable much beyond the Johnson bound in some parameter settings.

Since then, an exciting line of work [9, 18, 23–25, 55, 57] has shown the existence of Reed–Solomon codes that could in fact be list-decoded beyond the Johnson radius. These works all consider *combinatorial* list-decodability of *randomly punctured* Reed–Solomon codes. By combinatorial list-decodability, we mean that the code is proved to be list-decodable without providing an algorithm to efficiently decode the list of nearby codewords. By randomly punctured Reed–Solomon code, we mean a code  $\text{RS}_{n,k}(\alpha_1, \dots, \alpha_n)$  where  $(\alpha_1, \dots, \alpha_n)$  are chosen uniformly over all  $n$ -tuples of pairwise distinct elements of  $\mathbb{F}_q$ . Several of these works [18, 23, 55] proved more general list-decoding results about randomly puncturing any code with good unique-decoding properties, not just Reed–Solomon codes.

In this line of work, a recent breakthrough of Brakensiek, Gopi, and Makam [9] showed, using notions of “higher-order MDS codes” [8, 54], that Reed–Solomon codes can be list-decoded up to capacity. In fact, they show, more strongly, that Reed–Solomon codes can be list-decoded with list size  $L$  with radius  $p = \frac{L}{L+1}(1 - R)$ , exactly meeting the *generalized Singleton bound* [57], resolving a conjecture of Shangguan and Tamo [57]. However, their results require randomly puncturing Reed–Solomon codes over an exponentially large field size  $2^{O(n)}$ , where  $n$  is the block length of the code.

A natural question is how small can the field size be in a capacity-achieving Reed–Solomon code. Brakensiek, Dhar, and Gopi [6, Corollary 1.7, Theorem 1.8] showed that the exponential-in- $n$  field size in [9] is indeed necessary to *exactly* achieve the generalized Singleton bound for  $L = 2$  – under the additional assumptions that the code is linear and MDS. These assumptions were removed in followup work [2], which also generalized the result to all  $L$  – but smaller field sizes remained possible if one allowed a small  $\varepsilon$  slack in the parameters. Recently, an exciting work of Guo and Zhang [25] showed that Reed–Solomon codes are list-decodable up to capacity, in fact up to (but not exactly at) the generalized Singleton bound, with alphabet size  $O(n^2)$ .

## 1.1 Our Results

*List-Decoding Reed–Solomon Codes.* Building on Guo and Zhang’s argument, we show that Reed–Solomon codes are list-decodable up to capacity and the generalized Singleton bound with linear alphabet size  $O(n)$ , which is evidently optimal up to the constant factor. Our main result is the following.

**THEOREM 1.1.** *Let  $\varepsilon \in (0, 1)$ ,  $L \geq 2$  and  $q$  be a prime power such that  $q \geq n + k \cdot 2^{10L/\varepsilon}$ . Then with probability at least  $1 - 2^{-Ln}$ , a randomly punctured Reed–Solomon code of block length  $n$  and rate  $k/n$  over  $\mathbb{F}_q$  is  $(\frac{L}{L+1}(1 - R - \varepsilon), L)$  average-radius list-decodable.*

As in previous works [9, 25], Theorem 1.1 gives *average-radius list-decodability*, a stronger guarantee than list-decodability: for any distinct  $L+1$  codewords  $c^{(1)}, \dots, c^{(L+1)}$  and any vector  $y \in \mathbb{F}_q^n$ , the average Hamming distance from  $c^{(1)}, \dots, c^{(L+1)}$  to  $y$  is at least  $\frac{L}{L+1}(1 - R - \varepsilon)$ . Taking  $L = O(1/\varepsilon)$  in Theorem 1.1, it follows that Reed–Solomon codes achieve list-decoding capacity even over linear-sized alphabets.

**COROLLARY 1.2.** *Let  $\varepsilon \in (0, 1)$  and  $q$  be a prime power such that  $q \geq n + k \cdot 2^{O(1/\varepsilon^2)}$ . Then with probability at least  $1 - 2^{-\Omega(n/\varepsilon)}$ , a randomly punctured Reed–Solomon code of block length  $n$  and rate  $k/n$  over  $\mathbb{F}_q$  is  $(1 - R - \varepsilon, O(\frac{1}{\varepsilon}))$  average-radius list-decodable.*

The alphabet size in [25] is  $2^{O(L^2/\varepsilon)} nk$ . Our main contribution is improving their alphabet size from quadratic to linear. As a secondary improvement, we also bring down the constant factor from  $2^{O(L^2/\varepsilon)}$  to  $2^{O(L/\varepsilon)}$ . We defer the proof overview of Theorem 1.1 to Section 3.1 after setting up the necessary notions in Section 2.

In our proof of Theorem 1.1, we maintain a hypergraph perspective of the list-decoding problem, which was introduced in [24]. Section 2.2 elaborates on the advantages of this perspective, which include (i) more compact notations, definitions, and lemma statements, (ii) our improved constant factor of  $2^{O(L/\varepsilon)}$ , (iii) an improved alphabet size in our random linear codes result below

(Theorem 1.3), and (iv) an alternate presentation of ideas from Brakensiek, Gopi, and Makam [9] that more directly connects the list-decoding problem to the GM-MDS theorem [13, 47, 63] via a hypergraph orientation theorem (see Appendix A of the full version of our work [1]).

*List-Decoding Random Linear Codes.* A random linear code of rate  $R$  and length  $n$  over  $\mathbb{F}_q$  is a random subspace of  $\mathbb{F}_q^n$  of dimension  $Rn$ . List-decoding random linear codes is well-studied [16, 26, 27, 29, 30, 46, 49, 51, 55, 56, 61, 64] and is an important question for several reasons. First, finding explicit codes approaching list-decoding capacity is a major challenge, and random linear codes provide a stepping stone towards explicit codes: a classic result says that uniformly random codes achieve list-decoding capacity [15, 62], and showing list-decodability of random linear codes can be viewed as a derandomization of the uniformly random construction. Mathematically, the list-decodability of random linear codes concerns a fundamental geometric question: to what extent do random subspaces over  $\mathbb{F}_q$  behave like uniformly random sets? In coding theory, list-decodable random linear codes are useful building blocks in other coding theory constructions [28, 40]. Lastly, the algorithmic question of decoding random linear codes is closely related to the Learning With Errors (LWE) problem in cryptography [53] and Learning Parity with Noise (LPN) problem in learning theory [4, 17].

The list-decodability of random linear codes is more difficult to analyze than uniformly random codes, because codewords do not enjoy the same independence as in random codes. Thus the naive argument that shows that random linear codes achieve list-decoding capacity [64] gives an exponentially worse list size of  $q^{1/\varepsilon}$  than for random codes ( $\varepsilon$  is the gap to the “ $q$ -ary capacity”,  $R = 1 - H_q(p)$ , where  $H_q(x) \stackrel{\text{def}}{=} x \log_q(q-1) - x \log_q(x) - (1-x) \log_q(1-x)$  is the  $q$ -ary entropy function). Several works have sought to circumvent this difficulty [16, 26, 27, 29, 46, 55, 56, 61] improving the list-size bound to  $O_q(1/\varepsilon)$ , matching the list-size of uniformly random codes.

However, these results are more relevant for smaller alphabet sizes  $q$ , and approaching the alphabet-independent capacity of  $p = 1 - R$  is less understood. In this setting, uniformly random codes are, with high probability, list-decodable to capacity with optimal alphabet size  $2^{O(1/\varepsilon)}$ <sup>2</sup> and optimal list size  $O(1/\varepsilon)$ .<sup>3</sup> However, it was not known whether random linear codes (or, in general, more structured codes) could achieve similar parameters. In particular, both of the following questions were open (as far as we are aware).

- Are rate  $R$  random linear codes  $(1 - R - \varepsilon, O(1/\varepsilon))$ -list-decodable with high probability? Previously, this was not known for *any* alphabet size  $q$ , even alphabet size growing with the length of the code. Previously, the best list size for

<sup>2</sup>This follows from the list-decoding capacity theorem [15, 62]. Over  $q$ -ary alphabets, the list-decoding capacity is given by  $p = H_q^{-1}(1 - R)$ , which is larger than  $1 - R - \varepsilon$  when  $q \geq 2^{\Omega(1/\varepsilon)}$ .

<sup>3</sup>For codes over smaller alphabets, the list size  $O(1/\varepsilon)$ , where  $\varepsilon$  is the gap to capacity, is believed to be optimal, but a proof is only known for large radius [35]. However, for approaching the alphabet independent capacity, the list size  $O(1/\varepsilon)$  is known to be optimal by the generalized Singleton bound [57].

random linear codes list-decodable to radius  $p = 1 - R - \varepsilon$  was at least  $2^{\Omega(1/\varepsilon)}$  [26, 56].<sup>4</sup>

- Do there exist *any* linear codes (even non-constructively) over constant-sized (independent of  $n$ ) alphabets that are  $(1 - R - \varepsilon, O(1/\varepsilon))$ -list-decodable?

Using the same framework as the proof of Theorem 1.3, we answer both questions affirmatively. We show that, with high probability, random linear codes approach the generalized Singleton bound, and thus capacity, with alphabet size close to the optimal.

**THEOREM 1.3.** *For all  $L \geq 1, \varepsilon \in (0, 1)$ , a random linear code over alphabet size  $q \geq 2^{10L/\varepsilon}$  and  $n$  sufficiently large is with high probability  $(\frac{L}{L+1}(1 - R - \varepsilon), L)$ -average-radius-list-decodable.*

Taking  $L = O(1/\varepsilon)$ , we get that random linear codes achieve capacity with optimal list size  $O(1/\varepsilon)$  and near-optimal alphabet size  $2^{O(1/\varepsilon^2)}$ .

**COROLLARY 1.4.** *For all  $\varepsilon > 0$ , a random linear code over alphabet size  $q \geq 2^{O(1/\varepsilon^2)}$  and  $n$  sufficiently large is with high probability  $(1 - R - \varepsilon, O(1/\varepsilon))$ -average-radius-list-decodable.*

The techniques developed in this work for the proof of Theorem 1.1 are important for obtaining the strong alphabet size guarantees of Theorem 1.3. One could also have adapted the proof of Guo and Zhang, but doing so in the same natural way would yield an alphabet size of  $O(n)$ . Further, our use of the hypergraph machinery, which gives a secondary improvement over [25] in constant factor in the alphabet size in Corollary 1.2, gives the primary improvement in the alphabet size in Corollary 1.4 from  $2^{O(1/\varepsilon^3)}$  to  $2^{O(1/\varepsilon^2)}$ .

As the proof of Theorem 1.3 is very similar to the proof of Theorem 1.1, we focus on Theorem 1.1 for brevity and clarity of presentation in Section 2 and Section 3 and refer the reader to Section 4 of the full version of our work [1] for the proof of Theorem 1.3.

*Alphabet Size Lower Bounds.* Above, we saw that random linear codes achieve list-decoding capacity with optimal list-size and near-optimal alphabet size. A natural question, asked by Guo and Zhang, is how large the alphabet size needs to be for Reed–Solomon codes to achieve capacity. We showed that  $q \geq n \cdot 2^{O(1/\varepsilon^2)}$  suffices. By the list-decoding capacity theorem [15, 62], having an exponential-type dependence on  $1/\varepsilon$  for subconstant  $\varepsilon < O(1/\log n)$  is necessary.

For approaching capacity with constant  $\varepsilon$ , Ben-Sasson, Kopparty, and Radhakrishnan [3] showed that, for any  $c \geq 1$ , there exist full-length Reed–Solomon codes that are not list-decodable much beyond the Johnson bound with list-sizes  $O(n^c)$ . Thus in order to achieve list-decoding capacity, one needs  $q > n$  in some cases. However, while full-length Reed–Solomon codes could not achieve capacity, perhaps it was possible that Reed–Solomon codes over field size, say  $q = 2n$  or even  $q = (1 + \gamma)n$ , could achieve capacity in all parameter settings. We observe that, as a corollary of [3], such a strong guarantee is not possible. We show that, for any  $c > 1$ , there exist a constant rate  $R = R(c) > 0$  and infinitely many field sizes  $q$  such that all Reed–Solomon codes of length  $n \geq q/c$  and rate  $R$  over  $\mathbb{F}_q$  are not list-decodable to capacity  $1 - R$  with list size  $n^c$ . Due to space constraints, we omit the proof and refer the reader to Appendix B of the full version of this work [1].

<sup>4</sup>[26] appears to give a list-size bound of  $O(q^{O_R(1)} / \varepsilon)$ , and [56] appears to give a list size bound that is at least  $q^{\log^2(1/\varepsilon)}$ , and we need  $q \geq 2^{\Omega(1/\varepsilon)}$ .

**PROPOSITION 1.5.** *Let  $\delta = 2^{-b}$  for some integer  $b \geq 3$ . There exist infinitely many  $q$  such that any Reed–Solomon code of length  $n \geq 4\delta^{0.99}q$  and rate  $\delta$  is not  $(1 - 2\delta, n^{\Omega(\log(1/\delta))})$ -list-decodable.*

## 1.2 Follow-up Work

The techniques in our paper have already been influential. In follow-up work, Brakensiek, Dhar, Gopi, and Zhang [7] used our argument to prove that Algebraic Geometry (AG) codes achieve list-decoding capacity over constant-sized alphabets. They prove this by combining our techniques with a generalized GM-MDS theorem, proved by Brakensiek, Dhar, Gopi [5].

## 2 PRELIMINARIES

### 2.1 Basic Notation

For positive integers  $t$ , let  $[t]$  denote the set  $\{1, 2, \dots, t\}$ . The *Hamming distance*  $d(x, y)$  between two vectors  $x, y \in \mathbb{F}_q^n$  is the number of indices  $i$  where  $x_i \neq y_i$ . For a finite field  $\mathbb{F}_q$ , we follow the standard notation that  $\mathbb{F}_q[X_1, \dots, X_n]$  denotes the ring of multivariate polynomials with variables  $X_1, \dots, X_n$  over  $\mathbb{F}_q$ , and  $\mathbb{F}_q(X_1, \dots, X_n)$  denotes the field of fractions of the polynomial ring  $\mathbb{F}_q[X_1, \dots, X_n]$ . By abuse of notation, we let  $X_{\leq i}$  or  $X_{[i]}$  to denote the sequence  $X_1, \dots, X_i$ , and we let, for example,  $X_{\leq i} = \alpha_{\leq i}$  to denote  $X_1 = \alpha_1, X_2 = \alpha_2, \dots, X_i = \alpha_i$ . Given a matrix  $M$  over the field of fractions  $\mathbb{F}_q(X_1, \dots, X_n)$  and field elements  $\alpha_1, \dots, \alpha_i \in \mathbb{F}_q$ , let  $M(X_{\leq i} = \alpha_{\leq i})$  denote the matrix over  $\mathbb{F}_q(X_{i+1}, X_{i+2}, \dots, X_n)$  obtained by setting  $X_{\leq i} = \alpha_{\leq i}$  in  $M$ .

### 2.2 Hypergraphs and Connectivity

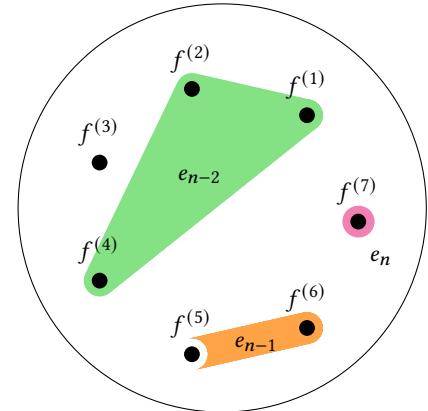
In this work, we maintain a hypergraph perspective of the list-decoding problem, which was introduced in [24]. We describe a bad list-decoding instance with a hypergraph where the  $L + 1$  bad codewords identify the vertices and the  $n$  evaluation points identify the hyperedges (Definition 2.1). While prior works described a bad list-decoding instance by  $L + 1$  sets indicating the agreements of the codewords with the received word, this hypergraph perspective gives us several advantages:

- (1) The constraints imposed by a bad list-decoding configuration yield a hypergraph that is *weakly-partition-connected*. This is a natural notion of hypergraph connectivity, which is well-studied in combinatorics [21, 22, 43] and optimization [11, 19, 20, 41], and which generalizes a well-known notion ( $k$ -partition-connectivity) for graphs [50, 60].<sup>5</sup> This connection allows us to have more compact notation, definitions, and lemma statements.
- (2) Because we work with weakly-partition-connected hypergraphs, we save a factor of  $L$  in Lemma 2.10 compared to the analogous lemma in [25]. This allows us to improve the constant factor in alphabet size for Reed–Solomon codes from  $2^{O(L^2/\epsilon)}$  in [25] to  $2^{O(L/\epsilon)}$  in Theorem 1.1.

<sup>5</sup>The notion of weakly-partition-connected sits between two other well-studied notions:  $k$ -partition-connected implies  $k$ -weakly-partition-connected implies  $k$ -edge-connected [43]. Each of these three notions generalizes an analogous notion on graphs. On graphs,  $k$ -partition-connected and  $k$ -weakly-partition-connected are equivalent.

- (3) For similar reasons, for random linear codes, the hypergraph perspective saves a factor of  $L$  in the alphabet size exponent, improving from  $2^{O(L^2/\epsilon)}$  to  $2^{O(L/\epsilon)}$  in Theorem 1.3.
- (4) With the hypergraph perspective, we can give a new presentation of the results in [9] and more directly connect the list-decoding problem to the GM-MDS theorem [13, 47, 63], as the heavy-lifting in the combinatorics is done using known results on hypergraph orientations. This is done in Appendix A of the full version of our work [1].

A hypergraph  $\mathcal{H} = (V, \mathcal{E})$  is given by a set of vertices  $V$  and a set  $\mathcal{E}$  of (*hyper*)edges, which are subsets of the vertices  $V$ . In this work, all hypergraphs have *labeled* edges, meaning we enumerate our edges  $e_i$  by distinct indices  $i$  from some set, typically  $[n]$ , in which case we may also think of  $\mathcal{E}$  as a tuple  $(e_1, \dots, e_n)$ . Throughout this paper, the vertex set  $V$  is typically  $[t]$  for some positive integer  $t$ . The *weight* of a hyperedge  $e$  is  $\text{wt}(e) \stackrel{\text{def}}{=} \max(0, |e| - 1)$ , and the *weight* of a set of hyperedges  $\mathcal{E}$  is simply  $\text{wt}(\mathcal{E}) \stackrel{\text{def}}{=} \sum_{e \in \mathcal{E}} \text{wt}(e)$ .



**Figure 1: Example edges from an agreement hypergraph  $\mathcal{H} = ([7], (e_1, \dots, e_n))$  (Definition 2.1) arising from a bad list-decoding configuration with polynomials  $f^{(1)}, \dots, f^{(7)} \in \mathbb{F}_q[X]$ , received word  $y \in \mathbb{F}_q^n$ , and evaluation points  $\alpha_1, \dots, \alpha_n$ . In the figure,  $e_{n-2} = \{1, 2, 4\}$  means  $f^{(1)}(\alpha_{n-2}) = f^{(2)}(\alpha_{n-2}) = f^{(4)}(\alpha_{n-2}) = y_{n-2}$ ,  $e_{n-1} = \{5, 6\}$  means  $f^{(5)}(\alpha_{n-1}) = f^{(6)}(\alpha_{n-1}) = y_{n-1}$ , and  $e_n = \{7\}$  means  $f^{(7)}(\alpha_n) = y_n$ .**

All hypergraphs that we consider in this work are *agreement hypergraphs* for a bad list-decoding configuration. See Figure 1 for an illustration.

**Definition 2.1 (Agreement Hypergraph).** Given vectors  $y, c^{(1)}, \dots, c^{(t)} \in \mathbb{F}_q^n$ , the *agreement hypergraph* has a vertex set  $[t]$  and a tuple of  $n$  hyperedges  $(e_1, \dots, e_n)$  where  $e_i \stackrel{\text{def}}{=} \{j \in [t] : c_i^j = y_i\}$ .

A key property of hypergraphs that we are concerned with is weak-partition-connectivity.

**Definition 2.2 (Weak Partition Connectivity).** A hypergraph  $\mathcal{H} = ([t], \mathcal{E})$  is  $k$ -weakly-partition-connected if, for every partition  $\mathcal{P}$  of the set of vertices  $[t]$ ,

$$\sum_{e \in \mathcal{E}} \max\{|\mathcal{P}(e)| - 1, 0\} \geq k(|\mathcal{P}| - 1) \quad (2)$$

where  $|\mathcal{P}|$  is the number of parts of the partition, and  $|\mathcal{P}(e)|$  is the number of parts of the partition that edge  $e$  intersects.

To give some intuition for weak partition connectivity, we state two of its combinatorial implications. First, if a graph is  $k$ -weakly-partition-connected, then it is  $k$ -edge-connected [43], which, by the Hypergraph Menger’s (Max-Flow-Min-Cut) theorem [43, Theorem 1.11], equivalently means that every pair of vertices has  $k$  edge-disjoint (hyper)paths between them.<sup>6</sup> Second, suppose we replace every hyperedge  $e$  with an arbitrary spanning tree of its vertices (which we effectively do in Definition 2.5). The resulting (non-hyper)graph is  $k$ -partition-connected,<sup>7</sup> which, by the Nash-Williams-Tutte Tree-Packing theorem [50, 60], equivalently means there are  $k$  edge-disjoint spanning trees (this connection was used in [24]).

The reason we consider weak-partition-connectivity is that a bad list-decoding configuration yields a  $k$ -weakly-partition-connected agreement hypergraph.

**LEMMA 2.3 (BAD LIST GIVES  $k$ -WEAKLY-PARTITION-CONNECTED HYPERGRAPH. SEE ALSO LEMMA 7.4 OF [24]).** *Suppose that vectors  $y, c^{(1)}, \dots, c^{(L+1)} \in \mathbb{F}_q^n$  are such that the average Hamming distance from  $y$  to  $c^{(1)}, \dots, c^{(L+1)}$  is at most  $\frac{L}{L+1}(n-k)$ . That is,  $\sum_{j=1}^{L+1} d(y, c^{(j)}) \leq L(n-k)$ . Then, for some subset  $J \subseteq [L+1]$  with  $|J| \geq 2$ , the agreement hypergraph of  $(y, c^{(j)} : j \in J)$  is  $k$ -weakly-partition-connected.*

Lemma 2.3 follows from the following result about weakly-partition-connected hypergraphs

**LEMMA 2.4.** *Let  $\mathcal{H} = (V, \mathcal{E})$  be a hypergraph with at least two vertices and with total edge weight  $\sum_{e \in \mathcal{E}} \text{wt}(e) \geq k \cdot (|V| - 1)$ , for some positive integer  $k$ . Then there exists a subset  $V' \subset V$  of at least two vertices such that the hypergraph  $\mathcal{H}' = (V', \{V' \cap e : e \in \mathcal{E}\})$  is  $k$ -weakly-partition-connected.*

**PROOF.** Let  $V'$  be an inclusion-minimal subset  $V' \subseteq [L+1]$  with  $|V'| \geq 2$  such that

$$\sum_{e \in \mathcal{E}} \text{wt}(e \cap V') \geq (|V'| - 1)k. \quad (3)$$

By assumption,  $V' = [L+1]$  satisfies (3), so  $V'$  exists (note that singleton subsets of  $[L+1]$  satisfy (3) with equality). Let  $\mathcal{H} = (V', \mathcal{E}')$  be the hypergraph with edge set  $\mathcal{E}' = \{V' \cap e : e \in \mathcal{E}\}$ . By minimality of  $V'$ , for all  $V'' \subsetneq V'$ , we have  $\sum_{e \in \mathcal{E}'} \text{wt}(e \cap V'') \leq (|V''| - 1)k$ . Now, consider a non-trivial partition  $\mathcal{P} = P_1 \sqcup \dots \sqcup P_p$  of  $V'$  where  $P_i \neq V'$  for all  $i \in [p]$  (as otherwise (2) trivially follows). We have

$$\begin{aligned} & \sum_{e \in \mathcal{E}'} \max\{|\mathcal{P}(e)| - 1, 0\} \\ &= \sum_{e \in \mathcal{E}', e \neq \emptyset} \left( -1 + \sum_{\ell=1}^p \mathbf{1}[|e \cap P_\ell| > 0] \right) \\ &= \sum_{e \in \mathcal{E}', e \neq \emptyset} \left( (|e| - 1) - \sum_{\ell=1}^p (|e \cap P_\ell| - \mathbf{1}[|e \cap P_\ell| > 0]) \right) \end{aligned}$$

<sup>6</sup>In general the converse is not true.

<sup>7</sup>In (non-hyper)graphs,  $k$ -partition-connectivity and  $k$ -weak-partition-connectivity are equivalent.

$$\begin{aligned} &= \sum_{e \in \mathcal{E}', e \neq \emptyset} \left( \max(|e| - 1, 0) - \sum_{\ell=1}^p \max(|e \cap P_\ell| - 1, 0) \right) \\ &= \sum_{e \in \mathcal{E}'} \text{wt}(e) - \sum_{\ell=1}^p \sum_{e \in \mathcal{E}'} \text{wt}(e \cap P_\ell) \\ &\geq (|V'| - 1)k - \sum_{\ell=1}^p (|P_\ell| - 1)k \\ &= (p - 1)k = (|\mathcal{P}| - 1)k. \end{aligned} \quad (4)$$

This holds for all partitions  $\mathcal{P}$  of  $V'$ , so  $\mathcal{H}'$  is  $k$ -weakly-partition-connected.  $\square$

**PROOF OF LEMMA 2.3.** Consider the agreement hypergraph  $([L+1], \mathcal{E})$  of  $y, (c^{(1)}, \dots, c^{(L+1)})$ . The total edge weight is

$$\begin{aligned} \sum_{e \in \mathcal{E}} \text{wt}(e) &\geq -n + \sum_{e \in \mathcal{E}} |e| \\ &= -n + \sum_{i=1}^n \sum_{j=1}^{L+1} \mathbf{1}[y_i = c_i^{(j)}] \\ &= -n + \sum_{j=1}^{L+1} (n - d(y, c^{(j)})) \geq Lk. \end{aligned} \quad (5)$$

By Lemma 2.4, there exists a subset  $J \subset [L+1]$  of at least two vertices such that  $\mathcal{H}' = (J, \{J \cap e : e \in \mathcal{E}\})$  — which is exactly the agreement hypergraph of  $(y, c^{(j)} : j \in J)$  — is  $k$ -weakly-partition-connected.  $\square$

**REMARK 1.** *The condition  $|J| \geq 2$  is needed later so that the reduced intersection matrix (defined below) is not a  $0 \times 0$  matrix, in which case the matrix does not help establish list-decodability.*

### 2.3 Reduced Intersection Matrices: Definition and Example

As in [25], we work with the reduced intersection matrix, though our proof should work essentially the same with a different matrix called the (non-reduced) *intersection matrix*, which was considered in [9, 24, 57].

**Definition 2.5 (Reduced intersection matrix).** The reduced intersection matrix  $\text{RIM}_{k,q,\mathcal{H}}$  associated with a prime power  $q$ , degree  $k$ , and a hypergraph  $\mathcal{H} = ([t], (e_1, \dots, e_n))$  is a  $\text{wt}(\mathcal{E}) \times (t-1)k$  matrix over the field of fractions  $\mathbb{F}_q(X_1, \dots, X_n)$ . For each hyperedge  $e_i$  with vertices  $j_1 < j_2 < \dots < j_{|e_i|}$ , we add  $\text{wt}(e_i) = |e_i| - 1$  rows to  $\text{RIM}_{\mathcal{H}}$ . For  $u = 2, \dots, |e_i|$ , we add a row  $r_{i,u} = (r^{(1)}, \dots, r^{(t-1)})$  of length  $(t-1)k$  defined as follows:

- If  $j = j_1$ , then  $r^{(j)} = [1, X_i, X_i^2, \dots, X_i^{k-1}]$
- If  $j = j_u$  and  $j_u \neq t$ , then  $r^{(j)} = [-1, X_i, X_i^2, \dots, X_i^{k-1}]$
- Otherwise,  $r^{(j)} = 0^k$ .

We omit  $k$  and  $q$  and write  $\text{RIM}_{\mathcal{H}}$  as they are typically understood.

**Example 2.6.** Recall the example edges of the agreement hypergraph  $\mathcal{H} = ([7], (e_1, \dots, e_n))$  in Figure 1. The edges  $e_{n-2}, e_{n-1}, e_n$  from  $\mathcal{H}$  contribute the following length  $(t-1)k$  rows to its reduced

intersection matrix:

$$\begin{bmatrix} V_{n-2} & -V_{n-2} & 0 & 0 & 0 & 0 \\ V_{n-2} & 0 & 0 & -V_{n-2} & 0 & 0 \\ 0 & 0 & 0 & 0 & V_{n-1} & -V_{n-1} \end{bmatrix} \quad (6)$$

Here  $V_i = [1, X_i, X_i^2, \dots, X_i^{k-1}]$  is a “Vandermonde row”, and  $0$  denotes the length- $k$  vector  $[0, 0, \dots, 0]$ . Note that each edge  $e$  contributes  $|e| - 1$  rows to the agreement matrix, and in particular  $e_n$  does not contribute any rows.

Reduced intersection matrices arise by encoding all agreements from a bad list-decoding configuration into linear constraints on the message symbols (the polynomial coefficients). These constraints are placed into one matrix that we call the reduced intersection matrix. The following lemma implies that, if every reduced intersection matrix arising from a possible bad list-decoding configuration has full column rank when  $X_1 = \alpha_1, \dots, X_n = \alpha_n$ , the corresponding Reed–Solomon code is list-decodable.

**LEMMA 2.7 (RIM OF AGREEMENT HYPERGRAPHS ARE NOT FULL COLUMN RANK).** *Let  $\mathcal{H}$  be an agreement hypergraph for  $(y, c^{(1)}, \dots, c^{(t)})$ , where  $c^{(j)} \in \mathbb{F}_q^n$  are codewords of  $RS_{n,k}(\alpha_1, \dots, \alpha_n)$ , not all equal to each other. Then the reduced intersection matrix  $\text{RIM}_{\mathcal{H}}(X_{[n]} = \alpha_{[n]})$  does not have full column rank.*

**PROOF.** By definition,

$$\text{RIM}_{\mathcal{H}}(X_{[n]} = \alpha_{[n]}) \cdot \begin{bmatrix} f^{(1)} - f^{(t)} \\ \vdots \\ f^{(t-1)} - f^{(t)} \end{bmatrix} = 0 \quad (7)$$

where  $f^{(1)}, \dots, f^{(t)} \in \mathbb{F}_q^k$  are the vectors of coefficients of the polynomials that generate codewords  $c^{(1)}, \dots, c^{(t)} \in \mathbb{F}_q^n$ . Since these vectors are not all equal to each other,  $\text{RIM}_{\mathcal{H}}(X_{[n]} = \alpha_{[n]})$  does not have full column rank.  $\square$

**REMARK 2 (SYMMETRIES OF REDUCED INTERSECTION MATRICES).** *From this definition, it should be clear that we can divide the variables  $X_1, \dots, X_n$  into at most  $2^L$  classes such that variables in the same class are exchangeable with respect to the reduced intersection matrix  $\text{RIM}_{\mathcal{H}}$ : if  $e_i$  and  $e_{i'}$  are the same hyperedge, then swapping  $X_i$  and  $X_{i'}$  yields the same reduced intersection matrix (up to row permutations). This observation, which was alluded to in [25], turns out to be crucial in our argument that allows us to improve the alphabet size in [25] from quadratic to linear.*

**REMARK 3.** *The pairwise distinctness requirement in the definition of average-radius-list-decodability (see Section 1.1) is nonetheless crucial in the proof of Theorem 1.1, despite the weaker requirement in Lemma 2.7. That is because we will eventually apply Lemma 2.7 on the subcollection of codewords given from Lemma 2.3, which can potentially be arbitrary. The guarantee that this subcollection of codewords is not all equal to each other would then follow from pairwise distinctness of the codewords in the original list.*

## 2.4 Reduced Intersection Matrices: Full Column Rank

The following theorem shows that reduced intersection matrices of  $k$ -weakly-partition-connected hypergraphs are nonsingular when

viewed as a matrix over  $\mathbb{F}_q(X_1, \dots, X_n)$ . This was essentially conjectured by Shangguan and Tamo [57] and essentially established by Brakensiek, Gopi, and Makam [9], who conjectured and showed, respectively, nonsingularity of the (non-reduced) intersection matrix under similar conditions. By the same union bound argument as in [57, Theorem 5.8], Theorem 2.8 already implies list-decodability of Reed–Solomon codes up to the generalized Singleton bound over exponentially large fields sizes, which is [9, Theorem 1.5]. For completeness, and to demonstrate how the hypergraph perspective more directly connects the list-decoding problem to the GM-MDS theorem, we include a proof of Theorem 2.8 in Appendix A of the full version of this work [1].

**THEOREM 2.8 (FULL COLUMN RANK. IMPLICIT FROM THEOREM A.2 OF [9]).** *Let  $n$  and  $k$  be positive integers and  $\mathbb{F}_q$  be a finite field. Let  $\mathcal{H}$  be a  $k$ -weakly-partition-connected hypergraph with  $n$  hyperedges and at least 2 vertices. Then  $\text{RIM}_{\mathcal{H}}$  has full column rank over the field  $\mathbb{F}_q(X_1, \dots, X_n)$ .*

**REMARK 4.** *We note that, [9] assumes throughout their paper that the alphabet size  $q$  is sufficiently large, but Theorem 2.8 follows from the weaker “ $q$  sufficiently large” version: For any fixed field size  $q$ , take  $Q$  to be a sufficiently large power of  $q$ . Then, by the “ $q$  sufficiently large” version of Theorem 2.8, matrix  $\text{RIM}_{Q, \mathcal{H}}$  has full column rank over the field  $\mathbb{F}_Q(X_1, \dots, X_n)$ . Hence, the determinant of some square full-rank submatrix of  $\text{RIM}_{Q, \mathcal{H}}$  is a nonzero polynomial in  $\mathbb{F}_Q[X_1, \dots, X_n]$ . The entries of  $\text{RIM}_{Q, \mathcal{H}}$  can all be viewed as polynomials over  $\mathbb{F}_q$ , so the corresponding full-rank submatrix of  $\text{RIM}_{Q, \mathcal{H}}$  has a determinant that is a nonzero polynomial in  $\mathbb{F}_q[X_1, \dots, X_n]$  — symbolically, the determinants are the same polynomials, as  $\mathbb{F}_q$  and  $\mathbb{F}_Q$  have the same characteristic. Hence, the matrix  $\text{RIM}_{Q, \mathcal{H}}$  has full column rank over the field  $\mathbb{F}_q(X_1, \dots, X_n)$ .*

## 2.5 Reduced Intersection Matrix: Row Deletions

As in [25], we consider row deletions from the reduced intersection matrix. The goal of this section is to establish Lemma 2.10, that the full-column-rank-ness of reduced intersection matrices are robust to row deletions.

**Definition 2.9 (Row deletion of reduced intersection matrix).** *Given a hypergraph  $\mathcal{H} = ([t], (e_1, \dots, e_n))$  and set  $B \subseteq [n]$ , define  $\text{RIM}_{\mathcal{H}}^B$  to be the submatrix of  $\text{RIM}_{\mathcal{H}}$  obtained by deleting all rows containing a variable  $X_i$  with  $i \in B$ .*

The next lemma appears in a weaker form in [25]. It roughly says that, given a reduced intersection matrix  $\text{RIM}_{\mathcal{H}}$  with some constant factor “slack” in the combinatorial constraints, we can omit a constant fraction of the rows without compromising the full-column-rank-ness of the matrix. Our version of this lemma saves roughly a factor of  $t \sim L$  compared to the analogous lemma [25, Lemma 3.11]. The reason is that the  $k$ -weakly-partition-connected condition is more robust to these row deletions (by a factor of roughly  $t$ ) than the condition in [25]. As such, our proof is also more direct.

**LEMMA 2.10 (ROBUSTNESS TO DELETIONS. SIMILAR TO LEMMA 3.11 OF [25]).** *Let  $\mathcal{H} = ([t], \mathcal{E})$  be a  $(k + \epsilon n)$ -weakly-partition-connected hypergraph with  $t \geq 2$ . For all sets  $B \subset [n]$  with  $|B| \leq \epsilon n$ , we have that  $\text{RIM}_{\mathcal{H}}^B$  is nonempty and has full column rank.*

**PROOF.** By definition of  $\text{RIM}_{\mathcal{H}}$ , the matrix with row deletions  $\text{RIM}_{\mathcal{H}}^B$  is the matrix  $\text{RIM}_{\mathcal{H}'}$ , where  $\mathcal{H}' = ([t], \mathcal{E}')$  is the hypergraph obtained from  $\mathcal{H}$  by deleting  $e_i$  for  $i \in B$ . By Theorem 2.8, it suffices to prove that  $\mathcal{H}'$  is  $k$ -weakly-partition connected. Indeed, consider any partition  $\mathcal{P}$  of  $[t]$ . We have

$$\begin{aligned} & \sum_{e \in \mathcal{E}'} \max\{|\mathcal{P}(e)| - 1, 0\} \\ &= \sum_{i \in [n]} \max\{|\mathcal{P}(e)| - 1, 0\} - \sum_{i \in B} \max\{|\mathcal{P}(e)| - 1, 0\} \\ &\geq (k + \varepsilon n) \cdot (|\mathcal{P}| - 1) - |B| \cdot (|\mathcal{P}| - 1) \\ &= k \cdot (|\mathcal{P}| - 1), \end{aligned} \quad (8)$$

as desired. The first inequality holds because  $\mathcal{H}$  is  $(k + \varepsilon n)$ -weakly-partition-connected, and, trivially, any edge  $e_i$  touches at most  $|\mathcal{P}|$  parts of  $\mathcal{P}$ .  $\square$

### 3 PROOF OF LIST-DECODABILITY WITH LINEAR-SIZED ALPHABETS

#### 3.1 Overview of the Proof

By Lemma 2.7 and Lemma 2.3, every bad list-decoding configuration admits a weakly-partition-connected agreement hypergraph whose reduced intersection matrix does not have full column rank. Thus, to prove Theorem 1.1, it suffices to show that, with high probability, every such reduced intersection matrix has full column rank. The main technical lemma for this section is the one stated below. Our main result, Theorem 1.1, follows by applying Lemma 2.3 and Lemma 2.7 with Lemma 3.1, and taking a union bound over all  $\sum_{t=2}^{L+1} 2^{tn}$  possible agreement hypergraphs.

**LEMMA 3.1.** *Let  $k$  be a positive integer and  $\varepsilon > 0$ . For any  $(k + \varepsilon n)$ -weakly-partition-connected hypergraph  $\mathcal{H} = ([t], (e_1, \dots, e_n))$  with  $t \geq 2$ , let  $\mathcal{B}_{\mathcal{H}}$  denote the event that the matrix  $\text{RIM}_{\mathcal{H}}(X_{[n]} = \alpha_{[n]})$  does not have full full column rank. For  $r = \lfloor \varepsilon n/2 \rfloor$ , we have*

$$\Pr_{\alpha_1, \dots, \alpha_n \sim \mathbb{F}_q \text{ distinct}} [\mathcal{B}_{\mathcal{H}}] \leq \binom{n}{r} 2^{tr} \cdot \left(\frac{(t-1)k}{q-n}\right)^r. \quad (9)$$

At a high level, the proof of Lemma 3.1 follows the same outline as [25]. For every tuple of evaluation points  $(\alpha_1, \dots, \alpha_n) \in \mathbb{F}_q^n$  for which  $\text{RIM}_{\mathcal{H}}$  does not have full column rank, we show that there is a *certificate*  $(i_1, \dots, i_r) \in [n]^r$  of distinct indices (Lemma 3.8), which intuitively ‘attests’ to the failure of the matrix  $\text{RIM}_{\mathcal{H}}$  to be full column rank. We then show that, for any certificate  $(i_1, \dots, i_r)$ , the probability that  $(\alpha_1, \dots, \alpha_n)$  has certificate  $(i_1, \dots, i_r)$  is exponentially small. (More precisely, it will at most be  $(\frac{(t-1)k}{q-n})^r$ . See Corollary 3.12.) We then show that there are not too many certificates (Corollary 3.10), and then union bound over the number of possible certificates to obtain the desired result (Lemma 3.1).

Our argument differs from [25] in how we choose our certificates. The argument of [25] allowed for up to  $n^r$  certificates. Our argument instead only needs  $\binom{n}{r} 2^{tr}$  many certificates, which is much smaller when  $r = \Omega(n)$  (the parameter regime of interest here) and overall allows us to save a factor of  $n$  in the alphabet size. Our savings comes from leveraging that there are at most  $2^t$  different ‘types’ of hyperedges (see Remark 2), and thus at most  $2^t$  different types of variables  $X_i$  in the reduced intersection matrix

$\text{RIM}_{\mathcal{H}}$ . This observation was alluded to in [25].<sup>8</sup> With this observation in mind, we assume, without loss of generality, that the edges of  $\mathcal{H}$  are ordered by their respective type (we can relabel the edges of  $\mathcal{H}$ , which effectively permutes the rows of  $\text{RIM}_{\mathcal{H}}$ ).

Our method of generating a certificate  $(i_1, \dots, i_r)$  for the evaluation sequence  $(\alpha_1, \dots, \alpha_n)$  (Algorithm 2) is similar to that of [25] at a high level—with each certificate  $i_1, \dots, i_r$ , we associate a sequence of  $(t-1)k \times (t-1)k$  submatrices  $M_1, \dots, M_r$  of  $\text{RIM}_{\mathcal{H}}$  (Algorithm 1) that are entirely specified by  $i_1, \dots, i_r$  as follows: since evaluating  $X_{[n]} = \alpha_{[n]}$  forces  $\text{RIM}_{\mathcal{H}}$  to not be full rank, then so will all of its  $(t-1)k \times (t-1)k$  submatrices. Thus if we sequentially ‘reveal’  $X_1 = \alpha_1, X_2 = \alpha_2, \dots$ , then at some point,  $M_j$  becomes singular exactly when we set  $X_{i_j} = \alpha_{i_j}$ —in fact,  $i_j$  is defined as such, so that we select  $M_1, i_1, M_2, i_2, \dots$ , in that order, but we emphasize that  $M_j$  can be computed from  $i_1, \dots, i_{j-1}$  without knowing  $\alpha_1, \dots, \alpha_n$ . Conditioned on  $M_j$  being non-singular with  $X_1 = \alpha_1, \dots, X_{i_j-1} = \alpha_{i_j-1}$ , the probability that  $M_j$  becomes singular when setting  $X_{i_j} = \alpha_{i_j}$  is at most  $\frac{(t-1)k}{q-n}$ :  $\alpha_{i_j}$  is uniformly random over at least  $q - n$  field elements, and the degree of  $X_{i_j}$  in the determinant of  $M_j$  is at most  $(t-1)k$  (and the determinant is nonzero by definition). Running conditional probabilities in the correct order, we conclude that the probability that a particular certificate  $i_1, \dots, i_r$  is generated is at most  $(\frac{(t-1)k}{q-n})^r$ , just as in [25].

Whereas [25] pick any matrix  $M_j$  that is obtained after removing the variables  $X_{i_1}, \dots, X_{i_{j-1}}$ , we do a more deliberate choice of matrices by leveraging the symmetries of  $\text{RIM}_{\mathcal{H}}$  (Remark 2). First, we ensure that we can keep a ‘bank’ of  $\Omega_t(r)$  unused variables of each of the  $O_t(1)$  types. Then, starting with a full column rank submatrix  $M$  of  $\text{RIM}_{\mathcal{H}}$  devoid of all variables in the ‘bank’, we start sequentially applying the evaluations  $X_1 = \alpha_1, X_2 = \alpha_2, \dots$ . Whenever  $M(X_{\leq i_1} = \alpha_{\leq i_1})$  turns singular, we find that the evaluation  $X_{i_1} = \alpha_{i_1}$  is what ‘caused’ it to become singular. We then go to the ‘bank’ to find a variable  $X_{i'_1}$  of the same type as  $X_{i_1}$  and ‘re-indeterminate’  $M$  by replacing all instances of  $X_{i_1}$  in  $M$  with  $X_{i'_1}$ . That way, we ensure that  $M$  is, in a sense, ‘reused.’ Furthermore, we ensure  $i'_1 > i_1$ , so that the matrix  $M(X_{\leq i_1} = \alpha_{\leq i_1})$  is now nonsingular, so we can keep going. Of course, if we end up reaching the end (i.e.  $M(X_{[n]} = \alpha_{[n]})$  is full column rank), then in fact,  $\text{RIM}_{\mathcal{H}}(X_{[n]} = \alpha_{[n]})$  is full column rank, and so the evaluations  $(\alpha_1, \dots, \alpha_n)$  were ‘good’ after all.

Otherwise, if the evaluations  $(\alpha_1, \dots, \alpha_n)$  were ‘bad’, then the submatrix  $M$  couldn’t have reached the end, and that can only happen if some specific type was completely exhausted from the bank. However, given the size of our initial bank, that must have meant that  $M$  was ‘re-indeterminate’ at least  $\Omega_t(r)$  times. When that happens, we collect the indices  $i_1, \dots, i_\ell$  that we gathered from this round, remove them from  $\text{RIM}_{\mathcal{H}}$ , and repeat the process again with a refreshed bank. Since we only need  $r$  indices, then we end up doing at most  $O_t(1)$  rounds. Because each round yields a strictly increasing sequence of indices of length at least  $\Omega_t(r)$ , then we up getting a certificate consisting of at most  $O_t(1)$  strictly increasing runs of total length  $r$ , of which there are at most  $\binom{n}{r} \cdot O_t(1)^r$ .

<sup>8</sup>Guo and Zhang [25] write ‘It is possible that achieving an alphabet size linear in  $n$  would require establishing and exploiting other properties of intersection matrices or reduced intersection matrices, such as an appropriate notion of exchangeability.’ We found this prediction to be insightful and true.

To be more concrete, when we generate the submatrix  $M = M_1$ , we ensure that any variable appearing in  $M_1$  has the same type as  $\Omega_t(r)$  variables that are *not* in  $M_1$  (but still in  $\text{RIM}_{\mathcal{H}}$ ). This creates a “bank” of variables of each type. Then, if  $X_{\leq i_1} = \alpha_{\leq i_1}$  was the point that made  $M_1$  singular, we can get  $M_2$  by replacing all copies of  $X_{i_1}$  with some  $X_{i'_1}$  that is of the same type and in the “bank.” Since variables  $i_1$  and  $i'_1$  are of the same type, they have analogous rows in the reduced intersection matrix  $\text{RIM}_{\mathcal{H}}$ , so this new matrix  $M_2$  is still a submatrix of  $\text{RIM}_{\mathcal{H}}$ . Therefore, we can pick up where we left off with  $M_1$  but with  $M_2$  instead. That is,  $M_2$  will in fact be full rank when we apply the evaluations  $X_{\leq i_1} = \alpha_{\leq i_1}$ . Thus the next index  $i_2$  on which  $M_2$  turns singular will be strictly greater than  $i_1$ . We then repeat the process in  $M_2$ , replacing  $X_{i_2}$  with some  $X_{i'_2}$  that is in the “bank” and of the same type, getting  $M_3$ , and so on. We can continue this process for  $\Omega_t(r)$  steps because of the size of the bank of each type, so we get an increasing run of length  $\Omega_t(r)$  in our certificate. After we run out of some type in our bank, we remove the used indices  $i_1, \dots, i_\ell$  from  $\text{RIM}_{\mathcal{H}}$  and repeat the process again with a refreshed bank. This continues for  $O_t(1)$  times only, as we only need  $r$  indices in the end.

We now finish the proof of Theorem 1.1, assuming Lemma 3.1. The rest of this section is devoted to proving Lemma 3.1.

**PROOF OF THEOREM 1.1, ASSUMING LEMMA 3.1.** By Lemma 2.3, if the code  $RS_{n,k}(\alpha_1, \dots, \alpha_n)$  is not  $\left(\frac{L}{L+1}(1-R-\varepsilon), L\right)$  average-radius list-decodable, then there exists a vector  $y$  and pairwise distinct codewords  $c^{(1)}, \dots, c^{(t)}$  with  $t \geq 2$  such that the agreement hypergraph  $\mathcal{H} = ([t], \mathcal{E})$  is  $(R+\varepsilon)n = (k+\varepsilon n)$ -weakly-partition-connected. By Lemma 2.7, the matrix  $\text{RIM}_{\mathcal{H}}(X_{[n]} = \alpha_{[n]})$  is not full column rank. That is, if we let  $\mathcal{B}_{\mathcal{H}}$  denote the event that the matrix  $\text{RIM}_{\mathcal{H}}(X_{[n]} = \alpha_{[n]})$  does not have full full column rank, then  $\mathcal{E}_{\mathcal{H}}$  occurs. Now, the number of possible agreement hypergraphs  $\mathcal{H}$  is at most  $\sum_{t=2}^{L+1} 2^{tn} \leq 2^{(L+2)n}$ . Thus by the union bound over possible agreement hypergraphs  $\mathcal{H}$  with Lemma 3.1, we have, for  $r = \lfloor \frac{\varepsilon n}{2} \rfloor$ ,

$$\begin{aligned} & \Pr_{\alpha_{[n]}} \left[ \text{RS}_{n,k}(\alpha_1, \dots, \alpha_n) \text{ not } \left(\frac{L}{L+1}(1-R-\varepsilon), L\right) \text{ list-decodable} \right] \\ & \leq \Pr_{\alpha_{[n]}} [\exists (k+\varepsilon n)\text{-w.p.c. hypergraph } \mathcal{H} \text{ s.t. } \mathcal{B}_{\mathcal{H}} \text{ occurs}] \\ & \leq 2^{(L+2)n} \max_{(k+\varepsilon n)\text{-w.p.c. } \mathcal{H}} \Pr_{\alpha_{[n]}} [\mathcal{B}_{\mathcal{H}}] \\ & \leq 2^{(L+2)n} \cdot \binom{n}{r} 2^{(L+1)r} \left(\frac{Lk}{q-n}\right)^r \\ & \leq \left(2^{(L+2)n/r} \cdot \frac{\varepsilon n}{r} \cdot 2^{L+1} \frac{Lk}{q-n}\right)^r \leq 2^{-Ln}, \end{aligned} \quad (10)$$

as desired. Here, we used that  $q = n + k \cdot 2^{10L/\varepsilon}$ .  $\square$

### 3.2 Setup for Proof of Lemma 3.1

We now devote the rest of this Section to proving Lemma 3.1.

**Types.** For a hypergraph  $\mathcal{H} = ([t], (e_1, \dots, e_n))$ , the *type* of an index  $i$  (or, by abuse of notation, the type of the variable  $X_i$ , or the edge  $e_i$ ) is simply the set  $e_i \subset [t]$ . There are  $2^t$  types, and by abuse of notation, we identify the types by the numbers  $1, 2, \dots, 2^t$  in an arbitrary fixed order with a bijection  $\tau : (\text{subsets of } [t]) \rightarrow [2^t]$ .

---

#### Algorithm 1: GetMatrixSequence

---

```

Input: indices  $i_1, \dots, i_{j-1} \in [n]$  for some  $j \geq 1$ .
Output:  $M_1, \dots, M_j$ , which are  $(t-1)k \times (t-1)k$  matrices
over  $\mathbb{F}_q(X_1, X_2, \dots, X_n)$ .
1  $B \leftarrow \emptyset, i_0 \leftarrow \perp, \ell_0 \leftarrow \perp$ 
2 for  $\ell = 1, \dots, j$  do
   //  $M_\ell$  depends only on  $i_1, \dots, i_{\ell-1}$ 
3   if  $\ell > 1$  then
      // Fetch new index from bank  $B$ 
5       $\tau \leftarrow$  the type of  $i_{\ell-1}$ 
6       $s \leftarrow$  number of indices among  $i_{\ell_0}, i_{\ell_0+1}, \dots, i_{\ell-1}$  that
      are type  $\tau$ 
7       $i'_{\ell-1} \leftarrow$  the  $s$ -th smallest element of  $B$  that has type  $\tau$ 
8      if  $i'_{\ell-1}$  is defined then
         $M_\ell \leftarrow$  the matrix obtained from  $M_{\ell-1}$  by
        replacing all copies of  $X_{i_{\ell-1}}$  with  $X_{i'_{\ell-1}}$ 
9   if  $M_\ell$  not yet defined then
   // Refresh bank  $B$ 
10   $B \leftarrow \emptyset$ 
11  for  $\tau = 1, \dots, 2^t$  do
   //  $B \leftarrow B \cup$ 
   // {top  $\lfloor r/2^t \rfloor$  type  $\tau$  indices in  $[n] \setminus \{i_1, \dots, i_{\ell-1}\}$ }
   // (if there are less than  $\lfloor r/2^t \rfloor$  indices of type  $\tau$ ,
   // then  $B$  contains all such indices)
13   $M_\ell \leftarrow$  lexicographically smallest nonsingular
    $(t-1)k \times (t-1)k$  submatrix of  $\text{RIM}_{\mathcal{H}}^{B \cup \{i_1, \dots, i_{\ell-1}\}}$ 
14   $\ell_0 \leftarrow \ell$  // new refresh index
15
16 return  $M_1, \dots, M_j$ 

```

---

We say a hypergraph is *type-ordered* if the hyperedges  $e_1, \dots, e_n$  are sorted according to their type:  $\tau(e_1) \leq \tau(e_2) \leq \dots \leq \tau(e_n)$ . Since permuting the labels of the edges of  $\mathcal{H}$  preserves the rank of  $\text{RIM}_{\mathcal{H}}$  (it merely permutes the rows of  $\text{RIM}_{\mathcal{H}}$ ), we can without loss of generality assume in Lemma 3.1 that  $\mathcal{H}$  is type-ordered.

**Global variables.** Throughout the rest of the section, we fix a positive integer  $k$ , parameter  $\varepsilon > 0$ , and  $\mathcal{H} = ([t], (e_1, \dots, e_n))$ , a type-ordered  $(k+\varepsilon n)$ -weakly-partition-connected hypergraph with  $t \geq 2$ . We also fix

$$r \stackrel{\text{def}}{=} \left\lfloor \frac{\varepsilon n}{2} \right\rfloor. \quad (11)$$

### 3.3 GetCertificate and GetMatrixSequence: Basic Properties

As mentioned at the beginning of this section, we design an algorithm, Algorithm 2, that attempts to generate a certificate  $(i_1, \dots, i_r) \in [n]^r$  for evaluation points  $\alpha_1, \dots, \alpha_n$ . It uses Algorithm 1, a helper function that generates the associated square submatrices  $M_1, \dots, M_r$  of  $\text{RIM}_{\mathcal{H}}$ . Below, we establish some basic properties of these algorithms.

**Algorithm 2: GetCertificate**


---

**Input:** Evaluation points  $(\alpha_1, \dots, \alpha_n) \in \mathbb{F}_q^n$ .  
**Output:** A “certificate”  $(i_1, \dots, i_r) \in [n]^r$ .

```

1 for  $j = 1, \dots, r$  do
2   //  $M_1, \dots, M_{j-1}$  stay the same,  $M_j$  is now
   // defined
3    $M_1, \dots, M_j = \text{GetMatrixSequence}(i_1, \dots, i_{j-1})$ 
4    $i_j \leftarrow$  smallest index  $i$  such that  $M_j(X_{\leq i} = \alpha_{\leq i})$  is
   // singular
5   if  $i_j$  not defined then
6     return  $\perp$ 
7   return  $(i_1, \dots, i_r)$ 

```

---

First, we establish that the matrices outputted by the function `GetMatrixSequence` are well-defined. After that, we show that `GetMatrixSequence` is an “online” algorithm.

**LEMMA 3.2 (OUTPUT IS WELL-DEFINED).** *For all sequence of indices  $i_1, \dots, i_{j-1}$ , if  $M_1, \dots, M_j$  are the matrices outputted by the function `GetMatrixSequence`( $i_1, \dots, i_{j-1}$ ), then  $M_1, \dots, M_j$  are well-defined.*

**PROOF.** If  $\ell$  is a refresh index, then we have  $|B \cup \{i_1, \dots, i_{\ell-1}\}| < |B| + r \leq 2r \leq \varepsilon n$ , so by Lemma 2.10,  $\text{RIM}_{\mathcal{H}}^{B \cup \{i_1, \dots, i_{\ell-1}\}}$  is nonempty and has full column rank. Thus  $M_\ell$  exists in Line 13. If  $\ell$  is not a refresh index,  $M_\ell$  is always well-defined by definition.  $\square$

**LEMMA 3.3 (ONLINE).** *Furthermore, `GetMatrixSequence` is a deterministic function of  $i_1, \dots, i_{j-1}$ , and it computes  $M_\ell$  “online”, meaning  $M_\ell$  depends only on  $i_1, \dots, i_{\ell-1}$  for all  $\ell = 1, \dots, j$  (and  $M_1$  is always the same matrix). In particular, `GetMatrixSequence`( $i_1, \dots, i_{j-1}$ ) is a prefix of `GetMatrixSequence`( $i_1, \dots, i_j$ ).*

**PROOF.** By definition and Lemma 3.2.  $\square$

**Definition 3.4 (Refresh index).** In `GetMatrixSequence`, in the outer loop over  $\ell$ , we say a *refresh index* is an index  $\ell$  obtained at Line 14 (i.e. when  $M_\ell$  is defined on Line 13). For example,  $\ell = 1$  is a refresh index.

Our first lemma shows that the new indices we are receiving from `GetMatrixSequence` are in fact new.

**LEMMA 3.5 (NEW VARIABLE).** *In `GetMatrixSequence`, in the outer loop iteration over  $\ell$  at Line 2, if we reach Line 8 of the function `GetMatrixSequence`, variable  $X_{i'_{\ell-1}}$  does not appear in  $M_{\ell_0}, M_{\ell_0+1}, \dots, M_{\ell-1}$ , where  $\ell_0$  is the largest refresh index less than  $\ell$ .*

**PROOF.** Let  $B$  be the set defined in Line 12 at iteration  $\ell_0$ . In iterations  $\ell' = \ell_0, \ell_0 + 1, \dots, \ell$ , the set  $B$  is the same, and  $i'_{\ell-1}$  is in this set  $B$  by definition. Thus, the variable  $X_{i'_{\ell-1}}$  does not appear in  $M_{\ell_0}$  by definition. For  $\ell' = \ell_0, \ell_0 + 1, \dots, \ell$ , the  $(\tau, s)$  pairs generated at Line 4 and Line 5 are pairwise distinct, so  $X_{i'_{\ell-1}}$  is not added to  $M_{\ell'}$  for  $\ell' = \ell_0 + 1, \dots, \ell - 1$  and thus is not in  $M_{\ell_0}, M_{\ell_0+1}, \dots, M_{\ell-1}$ .  $\square$

To show that the probability of a particular certificate  $(i_1, \dots, i_r)$  is small (Lemma 3.11, Corollary 3.12), we crucially need that  $i_1, \dots, i_r$  are pairwise distinct. The next lemma guarantees that.

**LEMMA 3.6 (DISTINCT INDICES).** *For any sequence of evaluation points  $(\alpha_1, \dots, \alpha_n) \in \mathbb{F}_q^n$ , the output of `GetCertificate`( $\alpha_1, \dots, \alpha_n$ ) is a sequence  $(i_1, \dots, i_r) \in [n]^r$  of pairwise distinct indices.*

**PROOF.** By definition of  $i_\ell$  at Line 3 of `GetCertificate`, variable  $X_{i_\ell}$  must be in  $M_\ell$ , so it suffices to show that  $M_\ell$  never contains any variable  $X_i$  for  $i \in \{i_1, \dots, i_{\ell-1}\}$ . We induct on  $\ell$ . If  $\ell$  is a refresh index, this is true by definition. If not, let  $\ell_0$  be the largest refresh index less than  $\ell$ . By induction,  $i_1, \dots, i_{\ell-2}$  are not in  $M_{\ell-1}$ , so we just need to show  $i'_{\ell-1}$  (the index replacing  $i_{\ell-1}$  in  $M_\ell$  at Line 8) is not any of  $i_1, \dots, i_{\ell-1}$ . None of  $i_1, \dots, i_{\ell_0-1}$  are in  $B$  by definition, so  $i'_{\ell-1}$  cannot be any of them.  $i'_{\ell-1}$  is not any of  $i_{\ell'}$  for  $\ell' = \ell_0, \dots, \ell-1$ , because  $X_{i_{\ell'}}$  is in  $M_{\ell'}$ , but  $X_{i'_{\ell-1}}$  is not, by Lemma 3.5.  $\square$

### 3.4 Bad Evaluation Points Admit Certificates

Here, we establish Lemma 3.8, that if some evaluation points make  $\text{RIM}_{\mathcal{H}}$  not full column rank, then `GetCertificate` outputs a certificate. First, in Lemma 3.7 we justify our matrix constructions, showing that the matrices in `GetMatrixSequence` are in fact submatrices of  $\text{RIM}_{\mathcal{H}}$ . Then in Lemma 3.8, we show that any tuple of bad evaluation points admits a certificate.

**LEMMA 3.7 (GetMatrixSequence GIVES SUBMATRICES OF RIM $_{\mathcal{H}}$ ).** *For all sequence of indices  $i_1, \dots, i_{j-1}$ , if  $M_1, \dots, M_j$  is the output of `GetMatrixSequence`( $i_1, \dots, i_{j-1}$ ), then  $M_1, \dots, M_j$  are  $(t-1)k \times (t-1)k$  submatrices of  $\text{RIM}_{\mathcal{H}}$ .*

**PROOF.** We proceed with induction on  $\ell = 1, \dots, j$ . First, if  $\ell$  is a refresh index, then  $M_\ell$  is a submatrix of  $\text{RIM}_{\mathcal{H}}$  by definition. In particular,  $M_1$  is a submatrix of  $\text{RIM}_{\mathcal{H}}$ , so the base case holds. Now suppose  $\ell$  is not a refresh index and  $M_{\ell-1}$  is a submatrix of  $\text{RIM}_{\mathcal{H}}$ . Matrix  $M_\ell$  is defined by replacing all copies of  $X_{i_{\ell-1}}$  with  $X_{i'_{\ell-1}}$ . To check that  $M_\ell$  is a submatrix of  $\text{RIM}_{\mathcal{H}}$ , it suffices to show that

- (i) for each row of  $\text{RIM}_{\mathcal{H}}$  containing  $X_{i_{\ell-1}}$ , replacing all copies of  $X_{i_{\ell-1}}$  with  $X_{i'_{\ell-1}}$  gives another row of  $\text{RIM}_{\mathcal{H}}$ , and
- (ii) the variable  $X_{i'_{\ell-1}}$  does not appear in  $M_{\ell-1}$ .

The first item follows from the fact that indices  $i_{\ell-1}$  and  $i'_{\ell-1}$  are of the same type, so (i) holds by definition of types and  $\text{RIM}_{\mathcal{H}}$  (see also Remark 2). The second item is Lemma 3.5. Thus,  $M_\ell$  is a submatrix of  $\text{RIM}_{\mathcal{H}}$ , completing the induction.  $\square$

**LEMMA 3.8 (BAD EVALUATIONS POINTS ADMIT CERTIFICATES).** *If  $(\alpha_1, \dots, \alpha_n) \in \mathbb{F}_q^n$  are evaluation points such that  $\text{RIM}_{\mathcal{H}}(X_{[n]} = \alpha_{[n]})$  does not have full column rank, `GetCertificate`( $\alpha_1, \dots, \alpha_n$ ) returns a certificate  $(i_1, \dots, i_r) \in [n]^r$  (rather than  $\perp$ ).*

**PROOF.** Suppose for contradiction that `GetCertificate` returns  $\perp$  at iteration  $j$  in the loop. Then there is no index  $i$  such that  $M_j(X_{\leq i} = \alpha_{\leq i})$  is singular. In particular,  $M_j(X_{[n]} = \alpha_{[n]})$  is non-singular and thus has full column rank. By Lemma 3.7,  $M_j$  is a submatrix of  $\text{RIM}_{\mathcal{H}}$ , so we conclude  $\text{RIM}_{\mathcal{H}}$  has full column rank.  $\square$

### 3.5 Bounding the Number of Possible Certificates

In this section, we upper bound the number of possible certificates. The key step is proving the following certificate structural result.

**LEMMA 3.9 (CERTIFICATE STRUCTURE).** *Given a sequence of evaluation points  $(\alpha_1, \dots, \alpha_n) \in \mathbb{F}_q^n$  such that  $\text{RIM}_{\mathcal{H}}(X_{[n]} = \alpha_{[n]})$  is not full column rank, the return value  $(i_1, \dots, i_r) = \text{GetCertificate}(\alpha_1, \dots, \alpha_n)$  satisfies  $i_{j-1} < i_j$  for all but at most  $2^t$  values  $j = 2, \dots, r$ .*

**PROOF.** Let  $(i_1, \dots, i_r)$  be the return of  $\text{GetCertificate}$ , and let  $M_1, \dots, M_r$  be the associated matrix sequence. By Lemma 3.3, we have  $M_1, \dots, M_j = \text{GetMatrixSequence}(i_1, \dots, i_{j-1})$  for  $j = 1, \dots, r$ . Recall an index  $\ell \in [r]$  is a *refresh index* if  $M_\ell$  is defined on Line 13 rather than Line 8. The lemma follows from two claims:

- (i) If  $\ell > 1$  is not a refresh index, then  $i_{\ell-1} < i_\ell$ .
- (ii) Any two refresh indices differ by at least  $r/2^t$ .

To see claim (i), let  $\ell_0$  be the largest refresh index less than  $\ell$ . By definition of a refresh index, the set  $B$  stays constant between when  $M_{\ell_0}$  is defined and when  $M_\ell$  is defined. From the definition of  $i_j$  at Line 3 in  $\text{GetCertificate}$ , we know that

- For  $i < i_{\ell-1}$  the matrix  $M_{\ell-1}(X_{\leq i} = \alpha_{\leq i})$  is nonsingular.
- The matrix  $M_\ell(X_{\leq i_\ell} = \alpha_{\leq i_\ell})$  is singular.

Suppose for contradiction that  $i_\ell < i_{\ell-1}$ . (Note that  $i_{\ell-1} \neq i_\ell$  by Lemma 3.6.) We contradict the first item by showing, using the second item, that  $M_{\ell-1}(X_{\leq i_\ell} = \alpha_{\leq i_\ell})$  is also singular. By the definition of  $\text{GetMatrixSequence}$ , since  $\ell$  is not a refresh index,  $M_\ell$  is defined in Line 8. By construction of  $B$  and  $i'_{\ell-1}$ , we know that  $i'_{\ell-1} > i_{\ell-1} > i_\ell$ . Thus, not only is  $M_\ell$  obtained from  $M_{\ell-1}$  by replacing all copies of  $X_{i_{\ell-1}}$  with  $X_{i'_{\ell-1}}$ , but  $M_\ell(X_{\leq i_\ell} = \alpha_{\leq i_\ell})$  is also obtained by replacing all copies of  $X_{i_{\ell-1}}$  with  $X_{i'_{\ell-1}}$  in  $M_{\ell-1}(X_{\leq i_\ell} = \alpha_{\leq i_\ell})$ . Moreover, the variable  $X_{i'_{\ell-1}}$  does not appear in  $M_{\ell-1}$  by Lemma 3.5. So we conclude that, as  $M_\ell(X_{\leq i_\ell} = \alpha_{\leq i_\ell})$  is singular, so is  $M_{\ell-1}(X_{\leq i_\ell} = \alpha_{\leq i_\ell})$ .

Now we show claim (ii). Suppose  $\ell_0$  and  $\ell_1$  are consecutive refresh indices. If a variable of type  $\tau$  appears in the matrix  $M_{\ell_0}$ , there must be exactly  $\lfloor r/2^t \rfloor$  indices of type  $\tau$  in  $B$  (if there were fewer, then  $B \cup \{i_1, \dots, i_{\ell-1}\}$  would contain all indices of type  $\tau$ , and the corresponding variables would not appear in  $\text{RIM}_{\mathcal{H}}^{B \cup \{i_1, \dots, i_{\ell-1}\}}$ ). Let  $\tau$  be the type of index  $i_{\ell_1-1}$ . Since  $\ell_1$  is a refresh index, the number of indices of type  $\tau$  among  $i_{\ell_0}, i_{\ell_0+1}, \dots, i_{\ell_1-1}$  must therefore be  $\lfloor r/2^t \rfloor + 1$ . In particular, this means  $\ell_1 - \ell_0 \geq \lfloor r/2^t \rfloor + 1 \geq r/2^t$ , as desired.  $\square$

**COROLLARY 3.10 (CERTIFICATE COUNT).** *The number of possible outputs to  $\text{GetCertificate}$  is at most  $\binom{n}{r} 2^{tr}$ .*

**PROOF.** The certificate consists of  $r$  distinct indices of  $[n]$  by Lemma 3.6. We can choose those in  $\binom{n}{r}$  ways. These indices are distributed between at most  $2^t$  increasing runs by Lemma 3.9. We can distribute these indices between the  $2^t$  increasing runs in at most  $(2^t)^r$  ways.  $\square$

### 3.6 Bounding the Probability of One Certificate

The goal of this section is to establish Corollary 3.12, which states that the probability of obtaining a particular certificate is at most  $(\frac{(t-1)k}{q-n})^r$ . The argument is implicit in [25], but we include a proof for completeness.

**LEMMA 3.11 (IMPLICIT IN [25]).** *Let  $i_1, \dots, i_r \in [n]$  be pairwise distinct indices, and  $M_1, \dots, M_r$  be  $(t-1)k \times (t-1)k$  submatrices*

*of  $\text{RIM}_{\mathcal{H}}$ . Over randomly chosen pairwise distinct evaluation points  $\alpha_1, \dots, \alpha_n \in \mathbb{F}_q$ , define the following events for  $j = 1, \dots, r$ :*

- $E_j$  is the event that  $M_j(X_{\leq i} = \alpha_{\leq i})$  is non-singular  $\forall i < i_j$ .
- $F_j$  is the event that  $M_j(X_{\leq i_j} = \alpha_{\leq i_j})$  is singular.

*The probability that all the events hold is at most  $(\frac{(t-1)k}{q-n})^r$ .*

**PROOF.** Note that the set of evaluation points  $\alpha_1, \dots, \alpha_n$  for which events  $E_j$  and  $F_j$  occur depends only on  $M_j$  and  $i_j$ . Furthermore, each of the events  $E_j$  and  $F_j$  depends only on  $M_i, i_j$ , and the evaluation points. Thus, by relabeling the index  $j$ , we may assume without loss of generality that  $i_1 < i_2 < \dots < i_r$ . We emphasize that we are *not* assuming that the output of  $\text{GetCertificate}$  satisfies  $i_1 < \dots < i_r$  (this is not true). We are instead just choosing how we 'reveal' our events  $E_j$  and  $F_j$ : starting with the smallest index in  $i_1, \dots, i_r$  and ending with the largest index in it.

We have

$$\begin{aligned} \Pr_{\alpha_{[n]}} \left[ \bigwedge_{j=1}^r (E_j \wedge F_j) \right] &= \prod_{j=1}^r \Pr_{\alpha_{[n]}} \left[ E_j \wedge F_j \mid \bigwedge_{\ell=1}^{j-1} (E_\ell \wedge F_\ell) \right] \\ &\leq \prod_{j=1}^r \Pr_{\alpha_{[n]}} \left[ F_j \mid \bigwedge_{\ell=1}^{j-1} (E_\ell \wedge F_\ell) \wedge E_j \right] \end{aligned} \quad (12)$$

Note that  $\bigwedge_{\ell=1}^{j-1} (E_\ell \wedge F_\ell) \wedge E_j$  depends only on  $\alpha_1, \dots, \alpha_{i_{j-1}}$ , and  $F_j$  depends only on  $\alpha_1, \dots, \alpha_{i_j}$ . For any  $\alpha_1, \dots, \alpha_{i_{j-1}}$  for which  $\bigwedge_{\ell=1}^{j-1} (E_\ell \wedge F_\ell) \wedge E_j$  holds, we have that  $M_j(X_{\leq i_{j-1}} = \alpha_{\leq i_{j-1}})$  is a  $(t-1)k \times (t-1)k$  matrix in  $\mathbb{F}_q(X_{i_j}, X_{i_j+1}, \dots, X_n)$  whose determinant is a nonzero polynomial of degree at most  $(t-1)k$  in each variable (the determinant contains at most  $t-1$  rows including  $X_{i_j}$ , each time with maximum degree  $k-1$ ). In particular, at most  $(t-1)k$  values of  $\alpha_{i_j}$  can make the determinant zero since, viewing the determinant as a polynomial in variables  $X_{i_j+1}, \dots, X_n$  with coefficients in  $\mathbb{F}_q[X_{i_j}]$ , any single nonzero coefficient becomes zero on at most  $(t-1)k$  values of  $\alpha_{i_j}$ . Conditioned on  $\alpha_1, \dots, \alpha_{i_{j-1}}$ , the field element  $\alpha_{i_j}$  is uniformly random over  $q - i_j + 1 \geq q - n$  elements. Thus, we have, for all  $\alpha_1, \dots, \alpha_{i_{j-1}}$  such that  $\bigwedge_{\ell=1}^{j-1} (E_\ell \wedge F_\ell) \wedge E_j$ ,

$$\Pr_{\alpha_{i_j}} \left[ F_j \mid \alpha_1, \dots, \alpha_{i_{j-1}} \right] \leq \frac{(t-1)k}{q-n}. \quad (13)$$

Since  $E_1 \wedge F_1 \wedge \dots \wedge E_{j-1} \wedge F_{j-1} \wedge E_j$  depends only on  $\alpha_{\leq i_{j-1}}$  and  $F_j$  depends only on  $\alpha_{\leq i_j}$ , we have

$$\Pr_{\alpha_{[n]}} \left[ F_j \mid \bigwedge_{\ell=1}^{j-1} (E_\ell \wedge F_\ell) \wedge E_j \right] \leq \frac{(t-1)k}{q-n}. \quad (14)$$

Combining with (12) gives the desired result.  $\square$

The key result for this section is a corollary of Lemma 3.11.

**COROLLARY 3.12 (PROBABILITY OF ONE CERTIFICATE).** *For any sequence  $i_1, \dots, i_r \in [n]$ , over randomly chosen pairwise distinct evaluation points  $\alpha_1, \dots, \alpha_n$ , we have*

$$\Pr \left[ \text{GetCertificate}(\alpha_1, \dots, \alpha_n) = (i_1, \dots, i_r) \right] \leq \left( \frac{(t-1)k}{q-n} \right)^r. \quad (15)$$

**PROOF.** By Lemma 3.6, we only need to consider pairwise distinct indices  $i_1, \dots, i_r$ , otherwise the probability is 0. Let  $M_1, \dots, M_r = \text{GetMatrixSequence}(i_1, \dots, i_r)$ . By Lemma 3.7, matrices  $M_1, \dots, M_r$

are all submatrices of  $\text{RIM}_{\mathcal{H}}$ . Thus, Lemma 3.11 applies. Let  $E_1, \dots, E_r, F_1, \dots, F_r$  be the events in Lemma 3.11. If  $\text{GetCertificate}(\alpha_1, \dots, \alpha_n) = (i_1, \dots, i_r)$ , then the definition of  $i_j$  in Line 3 of the function  $\text{GetCertificate}$  implies that events  $E_j$  and  $F_j$  both occur. By Lemma 3.11, the probability that all  $E_j$  and  $F_j$  hold is at most  $(\frac{(t-1)k}{q-n})^r$ , hence the result.  $\square$

### 3.7 Finishing the Proof of Lemma 3.1

PROOF OF LEMMA 3.1. Recall (Section 3.2) that we fixed  $\mathcal{H}$  to be a type-ordered  $(k + \epsilon n)$ -weakly-partition-connected hypergraph. By Lemma 3.8, if the matrix  $\text{RIM}_{\mathcal{H}}(X_{[n]} = \alpha_{[n]})$  does not have full column rank, then  $\text{GetCertificate}(\alpha_1, \dots, \alpha_n)$  is some certificate  $(i_1, \dots, i_r)$ . The probability that  $\text{GetCertificate}(\alpha_1, \dots, \alpha_n) = (i_1, \dots, i_r)$  is at most  $(\frac{(t-1)k}{q-n})^r$  by Corollary 3.12. By Corollary 3.10, there are at most  $\binom{n}{r} 2^r$  certificates. Taking a union bound over possible certificates gives the lemma.  $\square$

## ACKNOWLEDGEMENTS

We thank Mary Wootters and Francisco Pernice for helpful discussions about [9] and the hypergraph perspective of the list-decoding problem. We thank Karthik Chandrasekaran for helpful discussions about hypergraph connectivity notions and for the reference of Theorem A.3 in [19]. We thank Nikhil Shagrithaya and Jonathan Mosheiff for pointing out a mistake in the proof of Lemma 4.3 in an earlier version of the paper. We thank an anonymous reviewer for pointing out a mistake in the proof of Corollary A.4 in an earlier version of the paper.

## REFERENCES

- [1] Omar Alrabiah, Venkatesan Guruswami, and Ray Li. 2023. Randomly punctured Reed–Solomon codes achieve list-decoding capacity over linear-sized fields. *CoRR* abs/2304.09445 (2023). <https://doi.org/10.48550/ARXIV.2304.09445> arXiv:2304.09445
- [2] Omar Alrabiah, Venkatesan Guruswami, and Ray Li. 2024. AG codes have no list-decoding friends: Approaching the generalized Singleton bound requires exponential alphabets. In *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024, Alexandria, VA, USA, January 7-10, 2024*, David P. Woodruff (Ed.). SIAM, 1367–1378. <https://doi.org/10.1137/1.9781611977912.55>
- [3] Eli Ben-Sasson, Swastik Kopparty, and Jaikumar Radhakrishnan. 2010. Subspace polynomials and limits to list decoding of Reed–Solomon codes. *IEEE Trans. Inf. Theory* 56, 1 (2010), 113–120. <https://doi.org/10.1109/TIT.2009.2034780>
- [4] Avrim Blum, Adam Kalai, and Hal Wasserman. 2003. Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM* 50, 4 (2003), 506–519. <https://doi.org/10.1145/792538.792543>
- [5] Joshua Brakensiek, Manik Dhar, and Sivakanth Gopi. 2023. Generalized GM-MDS: Polynomial Codes are Higher Order MDS. *CoRR* abs/2310.12888 (2023). <https://doi.org/10.48550/ARXIV.2310.12888> arXiv:2310.12888
- [6] Joshua Brakensiek, Manik Dhar, and Sivakanth Gopi. 2023. Improved Field Size Bounds for Higher Order MDS Codes. In *IEEE International Symposium on Information Theory, ISIT 2023, Taipei, Taiwan, June 25–30, 2023*. IEEE, 1243–1248. <https://doi.org/10.1109/ISIT54713.2023.10206952>
- [7] Joshua Brakensiek, Manik Dhar, Sivakanth Gopi, and Zihan Zhang. 2023. AG codes achieve list decoding capacity over constant-sized fields. *CoRR* abs/2310.12898 (2023). <https://doi.org/10.48550/ARXIV.2310.12898> arXiv:2310.12898
- [8] Joshua Brakensiek, Sivakanth Gopi, and Visu Makam. 2022. Lower Bounds for Maximally Recoverable Tensor Codes and Higher Order MDS Codes. *IEEE Trans. Inf. Theory* 68, 11 (2022), 7125–7140. <https://doi.org/10.1109/TIT.2022.3187366>
- [9] Joshua Brakensiek, Sivakanth Gopi, and Visu Makam. 2023. Generic Reed–Solomon Codes Achieve List-Decoding Capacity. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20–23, 2023*, Barna Saha and Rocco A. Servedio (Eds.). ACM, 1488–1501. <https://doi.org/10.1145/3564246.3585128>
- [10] Jin-yi Cai, Aduri Pavan, and D. Sivakumar. 1999. On the Hardness of Permanent. In *STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science*, Trier, Germany, March 4–6, 1999, Proceedings (Lecture Notes in Computer Science, Vol. 1563), Christoph Meinel and Sophie Tison (Eds.). Springer, 90–99. [https://doi.org/10.1007/3-540-49116-3\\_8](https://doi.org/10.1007/3-540-49116-3_8)
- [11] Chandra Chekuri and Chao Xu. 2018. Minimum Cuts and Sparsification in Hypergraphs. *SIAM J. Comput.* 47, 6 (2018), 2118–2156. <https://doi.org/10.1137/18M1163865>
- [12] Qi Cheng and Daqing Wan. 2007. On the List and Bounded Distance Decodability of Reed–Solomon Codes. *SIAM J. Comput.* 37, 1 (April 2007), 195–209. <https://doi.org/10.1137/S0097539705447335>
- [13] Son Hoang Dau, Wentu Song, and Chau Yuen. 2014. On the existence of MDS codes over small fields with constrained generator matrices. In *2014 IEEE International Symposium on Information Theory, Honolulu, HI, USA, June 29 – July 4, 2014*. IEEE, 1787–1791. <https://doi.org/10.1109/ISIT.2014.6875141>
- [14] Zeev Dvir and Shachar Lovett. 2012. Subspace evasive sets. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 – 22, 2012*, Howard J. Karloff and Toniann Pitassi (Eds.). ACM, 351–358. <https://doi.org/10.1145/2213977.2214010>
- [15] Peter Elias. 1957. List decoding for noisy channels. *Wescon Convention Record, Part 2, Institute of Radio Engineers* (1957), 99–104.
- [16] Peter Elias. 1991. Error-correcting codes for list decoding. *IEEE Trans. Inf. Theory* 37, 1 (1991), 5–12. <https://doi.org/10.1109/18.61123>
- [17] Vitaly Feldman, Parikshit Gopalan, Subhash Khot, and Ashok Kumar Ponnuswami. 2006. New Results for Learning Noisy Parities and Halfspaces. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21–24 October 2006, Berkeley, California, USA, Proceedings*. IEEE Computer Society, 563–574. <https://doi.org/10.1109/FOCS.2006.51>
- [18] Asaf Ferber, Matthew Kwan, and Lisa Sauermann. 2022. List-Decodability With Large Radius for Reed–Solomon Codes. *IEEE Trans. Inf. Theory* 68, 6 (2022), 3823–3828. <https://doi.org/10.1109/TIT.2022.3148779>
- [19] András Frank. 2011. *Connections in combinatorial optimization*. Vol. 38. Oxford University Press Oxford.
- [20] András Frank and Tamás Király. 2008. A Survey on Covering Supermodular Functions. In *Research Trends in Combinatorial Optimization, Bonn Workshop on Combinatorial Optimization, November 3–7, 2008, Bonn, Germany*, William J. Cook, László Lovász, and Jens Vygen (Eds.). Springer, 87–126. [https://doi.org/10.1007/978-3-540-76796-1\\_6](https://doi.org/10.1007/978-3-540-76796-1_6)
- [21] András Frank, Tamás Király, and Zoltán Király. 2003. On the orientation of graphs and hypergraphs. *Discret. Appl. Math.* 131, 2 (2003), 385–400. [https://doi.org/10.1016/S0166-218X\(02\)00462-6](https://doi.org/10.1016/S0166-218X(02)00462-6)
- [22] András Frank, Tamás Király, and Matthias Kriesell. 2003. On decomposing a hypergraph into  $k$  connected sub-hypergraphs. *Discret. Appl. Math.* 131, 2 (2003), 373–383. [https://doi.org/10.1016/S0166-218X\(02\)00463-8](https://doi.org/10.1016/S0166-218X(02)00463-8)
- [23] Eitan Goldberg, Chong Shangguan, and Itzhak Tamo. 2022. Singleton-type bounds for list-decoding and list-recovery, and related results. In *IEEE International Symposium on Information Theory, ISIT 2022, Espoo, Finland, June 26 – July 1, 2022*. IEEE, 2565–2570. <https://doi.org/10.1109/ISIT50566.2022.9834849>
- [24] Zeyu Guo, Ray Li, Chong Shangguan, Itzhak Tamo, and Mary Wootters. 2021. Improved List-Decodability and List-Recoverability of Reed–Solomon Codes via Tree Packings: [Extended Abstract]. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7–10, 2022*. IEEE, 708–719. <https://doi.org/10.1109/FOCS52979.2021.00074>
- [25] Zeyu Guo and Zihan Zhang. 2023. Randomly Punctured Reed–Solomon Codes Achieve the List Decoding Capacity over Polynomial-Size Alphabets. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6–9, 2023*. IEEE, 164–176. <https://doi.org/10.1109/FOCS57990.2023.00019>
- [26] Venkatesan Guruswami, Johan Håstad, and Swastik Kopparty. 2011. On the List-Decodability of Random Linear Codes. *IEEE Trans. Inf. Theory* 57, 2 (2011), 718–725. <https://doi.org/10.1109/TIT.2010.2095170>
- [27] Venkatesan Guruswami, Johan Håstad, Madhu Sudan, and David Zuckerman. 2002. Combinatorial bounds for list decoding. *IEEE Trans. Inf. Theory* 48, 5 (2002), 1021–1034. <https://doi.org/10.1109/18.995539>
- [28] Venkatesan Guruswami and Piotr Indyk. 2001. Expander-Based Constructions of Efficiently Decodable Codes. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14–17 October 2001, Las Vegas, Nevada, USA*. IEEE Computer Society, 658–667. <https://doi.org/10.1109/SFCS.2001.959942>
- [29] Venkatesan Guruswami, Ray Li, Jonathan Mosheiff, Nicolas Resch, Shashwat Silas, and Mary Wootters. 2022. Bounds for List-Decoding and List-Recovery of Random Linear Codes. *IEEE Trans. Inf. Theory* 68, 2 (2022), 923–939. <https://doi.org/10.1109/TIT.2021.3127126>
- [30] Venkatesan Guruswami and Jonathan Mosheiff. 2022. Punctured Low-Bias Codes Behave Like Random Linear Codes. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 – November 3, 2022*. IEEE, 36–45. <https://doi.org/10.1109/FOCS54457.2022.00001>
- [31] Venkatesan Guruswami and Atri Rudra. 2006. Limits to List Decoding Reed–Solomon Codes. *IEEE Trans. Inform. Theory* 52, 8 (Aug. 2006), 3642–3649. <https://doi.org/10.1109/TIT.2006.878164>

[32] Venkatesan Guruswami and Atri Rudra. 2008. Explicit Codes Achieving List Decoding Capacity: Error-Correction With Optimal Redundancy. *IEEE Trans. Inf. Theory* 54, 1 (2008), 135–150. <https://doi.org/10.1109/TIT.2007.911222>

[33] Venkatesan Guruswami, Atri Rudra, and Madhu Sudan. 2022. Essential coding theory. *Draft available at <https://cse.buffalo.edu/faculty/atri/courses/coding-theory/book/>* (2022).

[34] Venkatesan Guruswami and Madhu Sudan. 1999. Improved decoding of Reed-Solomon and algebraic-geometry codes. *IEEE Trans. Inf. Theory* 45, 6 (1999), 1757–1767. <https://doi.org/10.1109/18.782097>

[35] Venkatesan Guruswami and Salil P. Vadhan. 2010. A Lower Bound on List Size for List Decoding. *IEEE Trans. Inf. Theory* 56, 11 (2010), 5681–5688. <https://doi.org/10.1109/TIT.2010.2070170>

[36] Venkatesan Guruswami and Carol Wang. 2013. Linear-Algebraic List Decoding for Variants of Reed-Solomon Codes. *IEEE Trans. Inf. Theory* 59, 6 (2013), 3257–3268. <https://doi.org/10.1109/TIT.2013.2246813>

[37] Venkatesan Guruswami and Chaoping Xing. 2012. Folded codes from function field towers and improved optimal rate list decoding. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, Howard J. Karloff and Toniann Pitassi (Eds.). ACM, 339–350. <https://doi.org/10.1145/2213977.2214009>

[38] Venkatesan Guruswami and Chaoping Xing. 2013. List decoding reed-solomon, algebraic-geometric, and gabidulin subcodes up to the singleton bound. In *Symposium on Theory of Computing Conference, STOC '13, Palo Alto, CA, USA, June 1-4, 2013*, Dan Boneh, Tim Roughgarden, and Joan Feigenbaum (Eds.). ACM, 843–852. <https://doi.org/10.1145/2488608.2488715>

[39] Brett Hemenway, Noga Ron-Zewi, and Mary Wootters. 2020. Local List Recovery of High-Rate Tensor Codes and Applications. *SIAM J. Comput.* 49, 4 (2020). <https://doi.org/10.1137/17M116149X>

[40] Brett Hemenway and Mary Wootters. 2018. Linear-time list recovery of high-rate expander codes. *Inf. Comput.* 261 (2018), 202–218. <https://doi.org/10.1016/J.ICO.2018.02.004>

[41] Kamal Jain, Mohammad Mahdian, and Mohammad R. Salavatipour. 2003. Packing Steiner trees. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 12-14, 2003, Baltimore, Maryland, USA*. ACM/SIAM, 266–274. <http://dl.acm.org/citation.cfm?id=644108.644154>

[42] Selmer M. Johnson. 1962. A new upper bound for error-correcting codes. *IRE Trans. Inf. Theory* 8, 3 (1962), 203–207. <https://doi.org/10.1109/TIT.1962.1057714>

[43] Tamás Király. 2003. *Edge-connectivity of undirected and directed hypergraphs*. Ph.D. Dissertation, Eötvös Loránd University.

[44] Swastik Kopparty. 2015. List-Decoding Multiplicity Codes. *Theory Comput.* 11 (2015), 149–182. <https://doi.org/10.4086/TOC.2015.V011A005>

[45] Swastik Kopparty, Noga Ron-Zewi, Shubhangi Saraf, and Mary Wootters. 2018. Improved Decoding of Folded Reed-Solomon and Multiplicity Codes. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, Mikkel Thorup (Ed.). IEEE Computer Society, 212–223. <https://doi.org/10.1109/FOCS.2018.00029>

[46] Ray Li and Mary Wootters. 2021. Improved List-Decodability of Random Linear Binary Codes. *IEEE Trans. Inf. Theory* 67, 3 (2021), 1522–1536. <https://doi.org/10.1109/TIT.2020.3041650>

[47] Shachar Lovett. 2018. MDS Matrices over Small Fields: A Proof of the GM-MDS Conjecture. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, Mikkel Thorup (Ed.). IEEE Computer Society, 194–199. <https://doi.org/10.1109/FOCS.2018.00027>

[48] Ben Lund and Aditya Potukuchi. 2020. On the List Recoverability of Randomly Punctured Codes. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2020, August 17-19, 2020, Virtual Conference (LIPIcs, Vol. 176)*, Jaroslaw Byrka and Raghu Meka (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 30:1–30:11. <https://doi.org/10.4230/LIPICS.APPROX/RANDOM.2020.30>

[49] Jonathan Mosheiff, Nicolas Resch, Noga Ron-Zewi, Shashwat Silas, and Mary Wootters. 2020. LDPC Codes Achieve List Decoding Capacity. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, Sandy Irani (Ed.). IEEE, 458–469. <https://doi.org/10.1109/FOCS46700.2020.00050>

[50] Crispin St. J. A. Nash-Williams. 1961. Edge-disjoint spanning trees of finite graphs. *Journal of the London Mathematical Society* 1, 1 (1961), 445–450.

[51] Aaron (Louie) Puterman and Edward Pyne. 2024. Pseudorandom Linear Codes Are List-Decodable to Capacity. In *15th Innovations in Theoretical Computer Science Conference, ITCS 2024, January 30 to February 2, 2024, Berkeley, CA, USA (LIPIcs, Vol. 287)*, Venkatesan Guruswami (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 90:1–90:21. <https://doi.org/10.4230/LIPICS.ITCS.2024.90>

[52] Irving S. Reed and Gustave Solomon. 1960. Polynomial Codes Over Certain Finite Fields. *J. Soc. Indust. Appl. Math.* 8, 2 (1960), 300–304. <https://doi.org/10.1137/0108018>

[53] Oded Regev. 2009. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM* 56, 6 (2009), 34:1–34:40. <https://doi.org/10.1145/1568318.1568324>

[54] Ron M. Roth. 2022. Higher-Order MDS Codes. *IEEE Trans. Inf. Theory* 68, 12 (2022), 7798–7816. <https://doi.org/10.1109/TIT.2022.3194521>

[55] Atri Rudra and Mary Wootters. 2014. Every list-decodable code for high noise has abundant near-optimal rate puncturings. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, David B. Shmoys (Ed.). ACM, 764–773. <https://doi.org/10.1145/2591796.2591797>

[56] Atri Rudra and Mary Wootters. 2018. Average-radius list-recoverability of random linear codes. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 644–662.

[57] Chong Shangguan and Itzhak Tamo. 2020. Combinatorial list-decoding of Reed-Solomon codes beyond the Johnson radius. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy (Eds.). ACM, 538–551. <https://doi.org/10.1145/3357713.3384295>

[58] Richard C. Singleton. 1964. Maximum distance  $q$ -nary codes. *IEEE Trans. Inf. Theory* 10, 2 (1964), 116–118. <https://doi.org/10.1109/TIT.1964.1053661>

[59] Madhu Sudan, Luca Trevisan, and Salil P. Vadhan. 2001. Pseudorandom Generators without the XOR Lemma. *J. Comput. Syst. Sci.* 62, 2 (2001), 236–266. <https://doi.org/10.1006/JCSS.2000.1730>

[60] William T. Tutte. 1961. On the problem of decomposing a graph into  $n$  connected factors. *Journal of the London Mathematical Society* 1, 1 (1961), 221–230.

[61] Mary Wootters. 2013. On the List Decodability of Random Linear Codes with Large Error Rates. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing (Palo Alto, California, USA) (STOC '13)*. ACM, New York, NY, USA, 853–860. <https://doi.org/10.1145/2488608.2488716>

[62] John M. Wozencraft. 1958. List Decoding. *Quarterly Progress Report, Research Laboratory of Electronics, MIT* 48 (1958), 90–95.

[63] Hikmet Yıldız and Babak Hassibi. 2019. Optimum Linear Codes With Support-Constrained Generator Matrices Over Small Fields. *IEEE Trans. Inf. Theory* 65, 12 (2019), 7868–7875. <https://doi.org/10.1109/TIT.2019.2932663>

[64] Victor Vasilievich Zyablov and Mark Semenovich Pinsker. 1981. List concatenated decoding. *Problemy Peredachi Informatsii* 17, 4 (1981), 29–33.

Received 31-OCT-2023; accepted 2024-02-11