Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 222 (2026)

Discrepancy for Random Linear Codes

Dean Doron* Tal Leonov* Jonathan Mosheiff* Henrique Navas†
Nicolas Resch‡ João Ribeiro†

Abstract

We show that random linear codes possess nearly optimal discrepancy-type properties in a broad range of settings. Our main results are two general discrepancy theorems: one controls all translates of a fixed test, and the other controls large families of Fourier-pseudorandom tests. As motivating applications of these two theorems, respectively, we show that:

- Random linear codes behave essentially like unstructured random codes for list-decoding from errors *above capacity*. More precisely, a random linear code $C \subseteq \mathbb{F}_q^n$ of rate $1 - \frac{1}{n}\log_q |B_\rho| + \varepsilon$, where $|B_\rho|$ is the volume of a radius-ρ Hamming ball in $\mathbb{F}_q^n$, satisfies
$$|C \cap B| = (1 \pm o(1))\frac{|C| \cdot |B|}{q^n}$$
simultaneously for all radius-ρ Hamming balls B in $\mathbb{F}_q^n$ with high probability.
This significantly strengthens the best prior results, which follow from work by Blinovsky (1987, 1990) and Levitin (1988) on the covering properties of random linear codes.
- Over prime fields, random linear codes behave essentially like unstructured random codes for zero-error list-recovery, and list-recovery from erasures, above capacity. More precisely, for a prime $q > 2$ and input list size $2 \leq \ell \leq q - 1$, a random linear code $C \subseteq \mathbb{F}_q^n$ of rate $1 - \log_q \ell + \varepsilon$ will satisfy
$$|C \cap S| = (1 \pm o(1))\frac{|C| \cdot \ell^n}{q^n}$$
simultaneously for all combinatorial rectangles $S = S_1 \times S_2 \times \cdots \times S_n$, where $|S_i| = \ell$ for all i, with high probability. An analogous result also holds when we can bound $|S_i|$ only for some of the i's.
In particular, we use this to show the abundance of locally leakage-resilient n-party linear *ramp* secret sharing schemes with any linear reconstruction threshold and sublinear threshold gap $O(n/\log n)$ over fields $\mathbb{F}_q$ of polynomial size $q = \Theta(n^\gamma)$ for a constant $\gamma \in (0, 1/5)$. Prior work on the existence of leakage-resilient linear secret sharing was stuck at reconstruction thresholds above $n/2$ for both threshold and ramp schemes.

The translate-family result, and hence the list-decoding application, applies over arbitrary finite fields, even when the field size grows with n. The list-recovery and leakage applications over prime fields hold under moderate-growth conditions on q, for example $q \leq n^{1/5-o(1)}$. Our results are obtained through a careful second-moment analysis of the evolution of intersection sizes as random generators are added to C one by one.

*Stein Faculty of Computer and Information Science, Ben-Gurion University. deand@bgu.ac.il, leonovt@post.bgu.ac.il, mosheiff@bgu.ac.il.

†Instituto de Telecomunicações and Departamento de Matemática, Instituto Superior Técnico, Universidade de Lisboa. h.campos.navas@tecnico.ulisboa.pt, jribeiro@tecnico.ulisboa.pt.

‡Informatics Institute, University of Amsterdam. n.a.resch@uva.nl.

ISSN 1433-8092

Contents

1 Introduction **3**

1.1 Our Contributions 5

1.2 Additional Context and Related Works 8

1.3 AI Disclosure 12

2 Main Results **12**

2.1 Notation and Conventions 12

2.2 Discrepancy Results for Random Linear Codes 14

2.3 Applications to List-Decoding and List-Recovery 16

3 Proof Overview **19**

4 Random Linear Codes are Smooth—Proof of Theorems 2.4 and 2.5 **25**

4.1 Smoothness via a Direct Second Moment Argument 25

4.2 Self Convolution and Variance of the Relative Deviation 28

4.3 Controlling the Deviation throughout the Construction Sequence 31

5 α-Concentration of Combinatorial Rectangles **42**

6 Improved Resilience of Linear Ramp Secret Sharing Schemes against Local Leakage Functions **44**

6.1 Leakage-Resilience of Linear Secret Sharing Schemes 45

6.2 Prior Work on Leakage-Resilience of Linear Secret Sharing Schemes and Our Result 46

6.3 Linear Ramp Secret Sharing from Linear Codes 47

6.4 Leakage Resilience of Massey Ramp Secret Sharing Schemes from Random Linear Codes 48

References **59**

A Deferred Proofs **63**

1 Introduction

A central theme in coding theory is to understand when random linear codes behave like unstructured uniformly random codes. This question is especially natural for list-decoding and list-recovery, two important relaxations of unique decoding.

A code $C \subseteq \mathbb{F}_q^n$ is called (ρ, L)-list-decodable (from errors) if, whenever a channel corrupts a ρ-fraction of symbols, there are at most L codewords that could have produced the received word. Combinatorially, this means that for every $z \in \mathbb{F}_q^n$,

$$|C \cap B_\rho(z)| \leq L,$$

where $B_\rho(z)$ denotes the Hamming ball of radius ρn centered at z.[1] Unique decoding is the special case $L = 1$.

List-recovery is a further relaxation of list-decoding: instead of receiving a single word in $\mathbb{F}_q^n$, the receiver obtains input lists $(S_1, \ldots, S_n)$, where each $S_i \subseteq \mathbb{F}_q$ has size at most ℓ, and must return all codewords that largely agree with these lists. We focus on the important special case of zero-error list-recovery. A code C is (ℓ, L)-list-recoverable if for every such tuple $(S_1, \ldots, S_n)$,

$$|C \cap (S_1 \times S_2 \times \cdots \times S_n)| \leq L.$$

Thus, list-decoding asks for small intersections with Hamming balls, while zero-error list-recovery asks for small intersections with combinatorial rectangles of bounded side lengths.

With these notions in mind, a basic question is: for a code of a given rate R, how small can the list size L be? Both list-decoding and list-recovery exhibit phase transitions: below a certain rate, constant list size is possible; above it, every code must have exponentially large list size. For list-decoding, fix a field $\mathbb{F}_q$ and an error radius ρ. For every $\varepsilon > 0$, the following are known:

- **Below capacity.** There exist codes of rate $1 - \frac{1}{n}\log_q |B_\rho| - \varepsilon$ that are $(\rho, L = O(1/\varepsilon))$-list-decodable from errors.
- **Above capacity.** Every code $C \subseteq \mathbb{F}_q^n$ of rate $1 - \frac{1}{n}\log_q |B_\rho| + \varepsilon$ has some center $z \in \mathbb{F}_q^n$ for which
$$|C \cap B_\rho(z)| \geq q^{\varepsilon n}.$$

We therefore call $1 - \frac{1}{n}\log_q |B_\rho|$ the list-decoding capacity.[2] The phase transition for zero-error list-recovery is analogous, with capacity $1 - \log_q \ell$.

Unstructured uniformly random codes achieve the random-code benchmark on both sides of capacity. A standard probabilistic-method argument shows that a uniformly random code of rate $1 - \frac{1}{n}\log_q |B_\rho| - \varepsilon$ is $(\rho, L = O(1/\varepsilon))$-list-decodable with high probability. Above capacity, a uniformly random code $C \subseteq \mathbb{F}_q^n$ of rate $1 - \frac{1}{n}\log_q |B_\rho| + \varepsilon$ satisfies, simultaneously for all $z \in \mathbb{F}_q^n$

[1] $B_\rho(z)$ consists of all $y \in \mathbb{F}_q^n$ at Hamming distance at most ρn from z, i.e., all y such that $|\{i : y_i \neq z_i\}| \leq \rho n$.

[2] Recall that $\frac{1}{n}\log_q |B_\rho| = h_q(\rho) - o(1)$, where

$$h_q(x) = x \log_q \frac{q-1}{x} + (1-x)\log_q \frac{1}{1-x}$$

is the q-ary entropy function. Hence, $1 - h_q(\rho)$ is also often termed the list-decoding capacity.

with high probability,

$$|C \cap B_\rho(z)| = (1 \pm o(1))\frac{|C| \cdot |B_\rho|}{q^n} = (1 \pm o(1))q^{\varepsilon n}. \tag{1}$$

Analogous statements hold for list-recovery.

A long line of work has asked whether random *linear* codes match the performance of unstructured random codes for list-decoding and list-recovery. Most prior work has focused on the *below-capacity* setting, where random linear codes are known to match uniformly random codes in some parameter regimes [GHSZ02, GHK11, CGV13, Woo13, RW14, LW20, AGL24, DMRR25], but not in others [Gur03, GLM+21, RY24, LMS25, CZ25, LS25]. Our understanding of the *above-capacity* setting is much poorer. This leads to the question that motivates our work:

> *How well do random linear codes approximate uniformly random codes in list-decoding and list-recovery at rates above capacity?*

Take list-decoding as a running example. Above capacity, list-decoding for random linear codes is naturally a *discrepancy* problem: the expected size for $|C \cap B_\rho(z)|$ is $|C| \cdot |B_\rho|/q^n \approx q^{\varepsilon n}$, and the goal is to show that a random linear code has approximately this many codewords in *every* Hamming ball. This is qualitatively different from the below-capacity setting. Below capacity, the benchmark $|C| \cdot |B_\rho|/q^n$ is exponentially small; moreover, for a linear code, the zero codeword always lies in $B_\rho(0)$, while for centers z with $0 \notin B_\rho(z)$ the fixed-center expectation is exponentially small. Thus the below-capacity problem is not a uniform multiplicative-discrepancy question in the same sense. As we discuss in detail later, this above-capacity viewpoint leads us to discrepancy-type questions that generalize list-decoding and list-recovery.

Toward proving Eq. (1) for random linear codes, one could first ask for the much weaker but still nontrivial guarantee that a random linear code $C \subseteq \mathbb{F}_q^n$ of rate $1 - \frac{1}{n}\log_q |B_\rho| + \varepsilon$ satisfies

$$|C \cap B_\rho(z)| \geq 1$$

for every $z \in \mathbb{F}_q^n$ with high probability. This is precisely the statement that C has covering radius at most ρn. Blinovsky [Bli87] and Levitin [Lev88] independently proved this over arbitrary fields almost forty years ago, building on earlier work of Goblick [Gob62], Cohen [Coh83], and Cohen and Frankl [CF85] which only established the *existence* of a linear code with the property above. Blinovsky's proof is nicely discussed in the book of Cohen, Honkala, Litsyn, and Lobstein on covering codes [CHLL97, Theorem 12.3.5], and it implies not only that C is ρ-covering with high probability, but also the stronger lower bound

$$|C \cap B_\rho(z)| \geq \frac{1 - o(1)}{n} \cdot \frac{|C| \cdot |B_\rho|}{q^n} \tag{2}$$

for all $z \in \mathbb{F}_q^n$. This is smaller than the optimal lower bound by a factor of roughly n. To the best of our knowledge, these were essentially the strongest previously known uniform guarantees on $|C \cap B_\rho(z)|$ for random linear codes at rates above capacity.

Beyond being a natural question about the geometry of subspaces of $\mathbb{F}_q^n$, discrepancy questions for random linear codes above capacity arise in several other settings. For example, zero-error list-recovery above capacity appears naturally in the study of leakage-resilient linear secret-sharing schemes. Roughly speaking, such schemes should ensure that leaking, say, one bit of information

about each share reveals almost nothing about the secret. There is a natural way, due to Massey, to turn a linear code into a secret-sharing scheme. For such a scheme, leakage resilience can be phrased as follows: for any functions $f_1, \ldots, f_n : \mathbb{F}_q \to \{0,1\}$, the distribution of $(f_1(c_1), \ldots, f_n(c_n))$, where $c = (c_1, \ldots, c_n)$ is sampled uniformly from C, should be close in total variation distance to the distribution of $(f_1(x_1), \ldots, f_n(x_n))$, where $x = (x_1, \ldots, x_n)$ is sampled uniformly from $\mathbb{F}_q^n$. For leakage bits $b_1, \ldots, b_n \in \{0,1\}$, put $S_i = f_i^{-1}(b_i)$. Then

$$\Pr[(f_1(c_1), \ldots, f_n(c_n)) = (b_1, \ldots, b_n)] = \frac{|(S_1 \times \cdots \times S_n) \cap C|}{|C|}.$$

The numerator is the intersection of the code with a combinatorial rectangle, exactly the kind of quantity controlled by zero-error list-recovery. In this application, assuming the sets S_i are not too small, the relevant intersection size is expected to be exponential; the point is not to make the list small, but to show that its exponential size is the correct one.

1.1 Our Contributions

We show that random linear codes match the performance of uniformly random codes in list-decoding (over any field) and in zero-error list-recovery and list-recovery from erasures (over prime fields), above capacity. In fact, we go further than that. We prove two main discrepancy theorems for random linear codes which show that random linear codes have strong above-capacity discrepancy properties for a broad class of combinatorial tests generalizing the above. Our new theorems also imply new results about leakage-resilient secret sharing.

We state simplified versions of our general theorems here, specialized to the main motivating examples. The body of the paper gives the more general formulations.

We first introduce a useful normalization. For a code $C \subseteq \mathbb{F}_q^n$ and a nonempty set $B \subseteq \mathbb{F}_q^n$, define the relative deviation of C with respect to B by

$$\delta_{C,B} := \frac{\Pr_{x \sim C}[x \in B]}{\Pr_{x \sim \mathbb{F}_q^n}[x \in B]} - 1 = \frac{q^n \cdot |C \cap B|}{|C| \cdot |B|} - 1. \tag{3}$$

Thus $|\delta_{C,B}| \le \tau$ is equivalent to

$$(1-\tau)\frac{|C| \cdot |B|}{q^n} \le |C \cap B| \le (1+\tau)\frac{|C| \cdot |B|}{q^n}.$$

List-Decoding and Beyond. Our first theorem applies to all translates of a single set. As we discuss below, this is a more general setting than list-decoding.

Theorem 1.1 (Informal; see Theorem 2.4)**.** *For any constant prime power q and every $n \in \mathbb{N}$ the following holds. Let $B \subseteq \mathbb{F}_q^n$ be nonempty, and for $z \in \mathbb{F}_q^n$ define $B_z := z + B$. Let*

$$\beta := 1 - \frac{1}{n}\log_q |B|,$$

and fix $\eta \in (0, 1-\beta)$. Let C be a random linear code of rate $R = \beta + \eta$. Then, with probability at least $1 - q^{-\Omega(\eta n)}$,

$$\forall z \in \mathbb{F}_q^n, \qquad |\delta_{C,B_z}| \le q^{-\Omega(\eta n)}.$$

For list-decoding, take $B = B_\rho(0)$. Then $\beta = 1 - \frac{1}{n}\log_q |B_\rho|$, and Theorem 1.1 says that a random linear code of rate $1 - \frac{1}{n}\log_q |B_\rho| + \eta$ has, with high probability, the correct intersection size

$$|C \cap B_\rho(z)| = (1 \pm q^{-\Omega(\eta n)})\frac{|C| \cdot |B_\rho(0)|}{q^n}$$

for every center $z \in \mathbb{F}_q^n$. In particular, this recovers and significantly strengthens Blinovsky's and Levitin's results [Bli87, Lev88], which only guarantee that $|C \cap B_\rho(z)| \geq \frac{1-o(1)}{n} \cdot \frac{|C| \cdot |B_\rho(0)|}{q^n}$ for all $z \in \mathbb{F}_q^n$.

We additionally remark that defining a "capacity" with respect to an arbitrary set B via the formula $1 - \frac{1}{n}\log_q |B|$ is not new to our work; it appeared in an earlier work of Loeliger [Loe94].

List-Recovery and Beyond. The translate theorem is well suited to list-decoding because all Hamming balls of a fixed radius are translates of one another. For list-recovery, the relevant rectangles are not all translates of a single set. Our second theorem handles large families of sets, provided the sets satisfy a Fourier pseudorandomness condition. We say that a set $B \subseteq \mathbb{F}_q^n$ is α-Fourier-concentrated if

$$\forall y \in \mathbb{F}_q^n, \qquad \left| \mathop{\mathbb{E}}_{Z \sim B} [\chi_y(Z)] \right| \leq \alpha^{\mathrm{wt}(y)},$$

where Z is uniform over B and χ_y is the Fourier character defined in Section 2.1. Informally, this condition says that B has little correlation with any nontrivial linear test, with the correlation decaying exponentially in the Hamming weight of the test.

Theorem 1.2 (Informal; see Theorem 2.5)**.** *For any constant prime power q and every $n \in \mathbb{N}$ the following holds. Let $\mathcal{B}$ be a family of nonempty α-Fourier-concentrated subsets of $\mathbb{F}_q^n$, and let*

$$\beta := 1 - \frac{1}{n} \min_{B \in \mathcal{B}} \log_q |B|.$$

Fix $\eta \in (0, 1-\beta)$, and let C be a random linear code of rate $\beta + \eta$. Then, with probability at least

$$1 - q^{-\Omega_\alpha(\eta n)} - |\mathcal{B}|\, q^{-\Omega_\alpha(n^2)},$$

we have

$$\forall B \in \mathcal{B}, \qquad |\delta_{C,B}| \leq q^{-\Omega(\eta n)}.$$

As an example, consider zero-error list-recovery with input list size ℓ. For the upper-bound guarantee it suffices to control rectangles with side length exactly ℓ, since smaller input lists can be padded. Thus consider rectangles $S_1 \times \cdots \times S_n$ with $S_i \in \binom{\mathbb{F}_q}{\ell}$. Over prime fields, standard Fourier estimates imply that such rectangles are α-Fourier-concentrated for some $\alpha = \alpha(q, \ell) < 1$, assuming $2 \leq \ell \leq q$. Therefore, in any regime where the family size $\binom{q}{\ell}^n$ is $q^{o(n^2)}$, Theorem 1.2 implies that a random linear code C of rate $1 - \log_q \ell + \eta$ satisfies, with high probability,

$$\forall S_1, \ldots, S_n \in \binom{\mathbb{F}_q}{\ell}, \qquad \left(1 - q^{-\Omega_{q,\ell}(\eta n)}\right) \frac{|C| \cdot \ell^n}{q^n} \leq |C \cap (S_1 \times \cdots \times S_n)| \leq \left(1 + q^{-\Omega_{q,\ell}(\eta n)}\right) \frac{|C| \cdot \ell^n}{q^n}.$$

In particular, the upper bound gives the expected above-capacity zero-error list-recovery guarantee for random linear codes in these regimes. See Theorem 2.12.

Extending zero-error list recovery, our result holds for list recovery from *erasures* as well. A code C is (ζ, ℓ, L)-list-recoverable if for every tuple $(S_1, \ldots, S_n)$ in which $(1-\zeta)n$ of the sets S_i satisfy $|S_i| = \ell$, it holds that

$$|C \cap (S_1 \times S_2 \times \cdots \times S_n)| \leq L.$$

In other words, list recovery from erasures allows for a ζ-fraction of the rectangle's sides to be "trivial". The family of such rectangles is still small enough, and α-Fourier concentrated, so Theorem 1.2 will imply that with high probability, a random linear code C of rate $(1-\zeta)(1-\log_q \ell)+\eta$ will intersect every such S at $(1 \pm o(1))|C|(\ell/q)^{(1-\zeta)n}$ points. See Theorem 2.14.

Application to Leakage-Resilient Linear Ramp Secret Sharing. Our Theorem 1.2 yields new results about the local leakage-resilience of linear secret sharing schemes. We present here a brief overview of this topic and leave a more careful discussion to Section 6.

Informally, a threshold secret sharing scheme encodes a secret into a tuple of n shares such that any appropriately large subset of shares determines the secret, while any appropriately small subset of shares reveals no information about the secret. For applications we often require that the secret sharing scheme be *linear*. Roughly speaking, this means that the secret and the shares are all elements of $\mathbb{F}_q$, and any linear combination of shares of two secrets yields a valid sharing of the linear combination of the two secrets. When q is much smaller than the number of parties n (i.e., when we aim for small shares), there must be a gap between the *reconstruction threshold* t_{rec} (all subsets of shares of size at least t_{rec} can reconstruct the secret) and the *privacy threshold* t_{priv} (all subsets of shares of size at most t_{priv} give no information about the secret) unless both t_{rec} and t_{priv} are unusually large [CCX13]. Such schemes with a gap between thresholds are called *ramp secret sharing schemes*, and they are the focus of our discussion.

Motivated by side-channel attacks, recently there has been significant interest in understanding the resilience of linear secret sharing schemes against attacks that leak a small amount of side information from *every* share. More precisely, in the most popular model of *bounded local leakage* [BDIR21], an adversary chooses local leakage functions $g_i : \mathbb{F}_q \to \{0,1\}$ for $i \in [n]$, and, letting $S = (S_1, \ldots, S_n)$ denote the n shares of some secret, learns the n-bit leakage

$$g(S) = (g_1(S_1), g_2(S_2), \ldots, g_n(S_n)).$$

Then, we want to establish that linear secret sharing schemes are information-theoretically resilient against this type of leakage, meaning that $g(S)$ and $g(S')$ should be very close in statistical distance for any sharings S and S' of two distinct secrets s and s', respectively.

Establishing leakage-resilience of secret sharing schemes is easier for larger reconstruction threshold. Therefore, a central research direction has been to establish the existence of leakage-resilient linear secret sharing schemes with small reconstruction threshold (and, in the context of ramp secret sharing, with a small non-trivial gap between the privacy and reconstruction thresholds). Prior results have only managed to show existence of such leakage-resilient schemes for reconstruction thresholds above $n/2$ [MPSW21, TX22, KK23, Kas24, Ngu24]. Subsequently to our work, Sun and Wootters [SW26a] and Nguyen [Ngu26] improved on the results above by showing that Shamir secret sharing is locally leakage-resilient for any reconstruction threshold $t \geq (1/2-\delta)n$ with $\delta > 0$ a small constant.

We use Theorem 1.2 to obtain an improved existential result about locally leakage-resilient linear ramp secret sharing schemes. To avoid complicating the exposition, we state here a simplified

version of our result and leave details for Section 6. Informally, it states that for arbitrary $R \in (0,1)$ there exist locally leakage-resilient n-party linear ramp secret sharing schemes with reconstruction threshold $t_{\text{rec}} \approx Rn$ and threshold gap $t_{\text{rec}} - t_{\text{priv}} \leq 3n/\log q$ over appropriately large fields $\mathbb{F}_q$ (more precisely, for q larger than some constant depending only on R and smaller than about $n^{1/5}$).

Theorem 1.3 (Informal version of Theorem 6.5. See also Theorem 6.14 and Corollary 6.15)**.** *Fix an arbitrary constant $R \in (0,1)$. Given $n \in \mathbb{N}$, let q be a prime satisfying*

$$2^{800/R} \leq q \leq n^{0.19}.$$

Then, for all sufficiently large n (depending only on R) there is a locally leakage-resilient n-party ramp secret sharing scheme over $\mathbb{F}_q$ with thresholds

$$Rn - \frac{1.01n}{\log q} \leq t_{\text{priv}} < t_{\text{rec}} \leq Rn + \frac{1.01n}{\log q}$$

and leakage-resilience error $\varepsilon \leq 4q^{-\frac{R^2 n^{3/5}}{3200}}$.

A Glimpse at our Techniques. While we will provide a much more detailed overview of our techniques below in Section 3, we briefly give a birds-eye view of our approach. To construct a random linear code, we imagine constructing a random linear code by iteratively sampling uniformly random basis vectors. That is, define $C_0 := \{0\}$, and inductively put $C_i := C_{i-1} + \text{span}\{u_i\}$, where u_i is sampled uniformly at random from $\mathbb{F}_q^n$: the final code is $C = C_k$, where $k = Rn$ is the target dimension.

The main task, then, is to control the growth of the intersections $|C_i \cap B|$ for increasing $i = 0, 1, \ldots, k$, for all sets B under consideration. A first, fairly obvious fact is that the growth rate is "correct" in expectation – after all, we are trying to argue that the $|C_k \cap B|$ is always roughly its expectation! The main challenge is to get a tail bound: we need to argue that it is *very unlikely* that any $|C_i \cap B|$ deviates far from its expectation.

The main tool at our disposal is the fact that the events "$x \in C$" (over distinct x) are "nearly" *pairwise* independent (there is a slight exception from colinear pairs, but this is negligible): thus, a second-moment method is natural. The challenge is that applying a second moment argument to *each* B under consideration and then applying a union bound is doomed to fail (there are more sets to union over than the reciprocal of the failure probability). Our main accomplishment is to identify a *single common* event that implies *all* the growth rates are adequately controlled; additionally, this event can be proved likely to occur via a basic second-moment argument.

1.2 Additional Context and Related Works

Having given a taste of our results and our techniques, we now discuss (additional) related works to help situate our contribution.

Covering Radius. In studying codes above capacity the parameter that has seen the most attention is its covering radius ρ. As discussed above around Eq. (2), independent work by Blinovsky [Bli87] and Levitin [Lev88] demonstrates that codes of rate $1 - \frac{1}{n}\log_q |B_\rho| + \eta$ are very likely

to satisfy $|C \cap B_\rho(z)| \geq \frac{1-o(1)}{n} \cdot \frac{|C| \cdot |B_\rho|}{q^n}$ for all $z \in \mathbb{F}_q^n$. Soon after, Blinovsky [Bli90a] improved the bound on the probability that a random linear code is ρ-covering.

The arguments behind these bounds (see [CHLL97, Theorem 12.3.5], which presents the proof from [Bli87]) indeed consider revealing the random linear codes basis vectors one at a time, analogously to us, and track the step-by-step evolution of the first moment of an appropriate potential.

However, a notable simplification in these works compared to our case is that establishing lower bounds on $|C \cap B_\rho(z)|$ is a *monotone* process, in the sense that adding additional basis vectors to C will not "undo" the process and make the lower bound worse. For example, the argument in [Bli87] and [CHLL97, Theorem 12.3.5] exploits this fact essentially by arguing that after a certain number of steps "almost all" points are covered, and then adding a handful of more basis vectors to handle the uncovered vectors (in particular, the union bound here can be much smaller). In contrast to this, since we establish a two-sided bound, we are not able to argue that certain vectors z are "satisfied," as every step could potentially cause an issue. In particular, while covering codes can be viewed as codes "above capacity," the defining property is not a real "discrepancy" problem as we focus on addressing here.

Covering and Local Properties. Covering also suggests a natural direction for extending the systematic study of random linear codes. A line of work originating in [MRRZ+24] develops a unified theory of threshold rates for local properties. Roughly, a property is local if it is governed by the presence or absence of prescribed configurations involving a bounded number of codewords, independent of the block length. For example, for fixed L, (ρ, L)-list-decodability is local: its failure is witnessed by $L+1$ distinct codewords contained in a single radius-ρn Hamming ball. This theory identifies general principles governing the threshold rates of broad classes of local properties simultaneously.

Covering lies outside this local framework, but fits naturally into an extension based on a universal–existential quantifier structure. Indeed, a code C has covering radius at most ρn precisely when

$$\forall z \in \mathbb{F}_q^n \quad \exists c \in C \quad c \in B_\rho(z).$$

Informally, we may call conditions of this form Π_2 *properties*: "for every test, there exists a codeword (or a bounded tuple of codewords) passing the test." Developing a unified theory of such properties is a natural next step beyond the theory of local properties. The present work grew in part out of an investigation of Π_2 properties of random linear codes.

Linear Hashing. Discrepancy questions for random linear codes above capacity do arise in the study of linear hashing [ADM+99, DD22, JKZ25, Bsh26]. Here one fixes a subset $S \subseteq \mathbb{F}_q^n$ and, identifying C with the kernel of a random linear map, asks that every translate $y+S$ have about the expected intersection size with C. The relevant regime is again just above the natural threshold: C is chosen so that $|C| \cdot |S| = q^{n+O(1)}$. The main difference here is that we look at codes with constant *relative* gap above capacity (i.e., $|C| \cdot |S| = q^{(n+\varepsilon)n}$): it is an interesting open problem to determine if our techniques can additionally work with this narrower gap. We remark that most of these works achieve bounds only over the *binary* field [ADM+99, JKZ25, Bsh26]; only the work of Dhar and Dvir [DD22] achieves bounds over arbitrary fields (as we do).

Potential-Based Arguments. Prior works have used potential-based arguments to study the list-decodability of random linear codes below capacity [GHSZ02, LW20], linear hashing [ADM+99, JKZ25, Bsh26], and the covering radius of random linear codes [Gob62, Coh83, CF85, DP86, Bli87, Lev88, Bli90a, Bli90b].[3] Briefly, all these arguments consider a potential of the form $\Phi(C) = \mathbb{E}_{X \in \mathbb{F}_2^n}\left[b^{|C \cap (X+S)|}\right]$, where S is the set under consideration ($B_\rho(0)$ for list-decoding, or an arbitrary set for linear hashing), and $b \geq 0$ is an appropriate base for the exponential (about $2^{n/L}$ for list-decoding, $b \leq 1$ for linear hashing, and $b = 0$ for covering problems with the convention that $0^0 = 1$, as nicely observed in Kumar's PhD thesis [Kum26, Section 3.1.2]). As in our work, the arguments in these papers consider constructing the codes by including a basis vector one at a time, and observing how this potential evolves. The crucial observation in the papers above on list-decoding and linear hashing is that for any fixed code C we have

$$\mathop{\mathbb{E}}_{U \sim \mathbb{F}_2^n}\left[\Phi(C + \langle U \rangle)\right] \leq \Phi(C)^2. \tag{4}$$

That is, on average over the newly included basis vector, the potential is squared. By iterating this, ignoring some technicalities, one obtains a bound of the form $\Phi_k \leq \Phi(\{0\})^{2^k}$, which is enough to show that *some* linear code exists with the desired property. It is also possible to refine this argument to obtain a statement that holds with high probability over the sampling of the random linear code. Such refinements were obtained for the covering problem in the late 80s [Bli87, Lev88, Bli90a, Bli90b], and much later for list-decoding and linear hashing [LW20, JKZ25, Bsh26].

Establishing Eq. (4) crucially uses the assumption that the field is binary. Over $\mathbb{F}_2$, every line consists of just two points; additionally, note that adding a random basis vector essentially amounts to adding a line at every codeword in a random direction. For random linear codes (over *any* field), the events "$x \in C$" (over distinct non colinear vectors x) are always *pairwise* independent; however, they fail to be even 3-wise independent (for example, conditioned on $x \in C$ and $y \in C$, it holds that $x + y \in C$ with probability 1). In particular, one obtains correlations in computing $\mathbb{E}_{U \sim \mathbb{F}_2^n}\left[\Phi(C + \langle U \rangle)\right]$ that appear hard to handle. Thus, this potential-based argument, while extremely elegant, is currently limited to $\mathbb{F}_2$ for list-decoding and linear hashing. In fact, a motivation for our work was to develop techniques allowing for an analysis over arbitrary fields. To the best of our knowledge, the only works that managed to apply the potential method beyond $\mathbb{F}_2$ focus on the covering radius of random linear codes. Cohen and Frankl [CF85] adapted the potential method of Goblick [Gob62] and Cohen [Coh83] to pinpoint the optimal covering radius of linear codes over arbitrary fields, and Blinovsky [Bli87] and Leviti [Lev88] later established a high probability analogue. As we discussed above, the covering property satisfies a notion of monotonicity that simplifies the analysis of the step-by-step argument compared to list-decoding and the two-sided bounds we consider here.

Finally, we remark that by combining the potential function-based argument of [GHSZ02, LW20] and the recent tail bound of Jaber, Kumar, and Zuckerman [JKZ25] one can prove a much weaker version of Theorem 1.1 *for binary codes only*, where the $2^{-\Omega(\eta n)}$ term is replaced by a constant larger than 2 (in particular, it only guarantees list size a constant factor larger than the expectation) and the probability bound is also weaker. Generalizing this type of argument to larger fields has eluded researchers for the last 25 years. For future work, it would be interesting to determine to what extent our techniques can be applied to random linear codes *below* capacity.

[3] The use of the potential method in coding theory seems to go back at least to 1962 [Gob62, Theorem 4.2], in the context of characterizing the optimal covering radius of linear codes. We refer to the historical note of Ribeiro [Rib26] for a more detailed discussion on the early development of this method.

Optimal Polynomial Intersection (OPI). This problem has recently attracted attention as a possible route to demonstrating quantum advantage [JSW+25, CT25, Cha25, KSG+25, Ros26, SW26b]. In one version of the problem,[4] one is given input lists $S_1, \ldots, S_n \subseteq \mathbb{F}_q$ and a description of a linear code C, and must output a codeword $c \in C$ maximizing $|\{i \in [n] : c_i \in S_i\}|$. Equivalently, one seeks a codeword in an output list for a list-recovery problem (with errors) at the smallest possible decoding radius. The parameter regimes of interest are above capacity: the corresponding list will contain exponentially many codewords. However, the question here is more algorithmic – *find* some codeword from this (exponentially large) set of possibilities – rather than combinatorial/discrepancy-theoretic, like the questions we study.

Random Linear Codes Below Capacity. Finally, while our main goal is to study random linear codes at rates *above* capacity, we remark that there has been a long line of work studying random linear codes *below* capacity (we already mentioned some such works above when discussing potential function-based arguments). Firstly, we consider the task of list-decoding: at rates $1 - \frac{1}{n}\log_q |B_\rho| - \varepsilon$, one aims to establish that random linear codes are with high probability $(\rho, O(1/\varepsilon))$-list-decodable, which matches the performance of uniformly random codes. It is worth mentioning a classical argument of Zyablov and Pinsker [ZP81] demonstrates they are $(\rho, q^{O(1/\varepsilon)})$-list-decodable; so the challenge is to show an exponentially smaller list-size is still sufficient. The aforementioned work [GHSZ02] derived the $O(1/\varepsilon)$ list-size, but only over the binary alphabet; later works managed to (largely) extend this to arbitrary, constant-sized fields [GHK11, CGV13, Woo13], although often with some parameter restrictions.

List-recovery has also seen more attention in recent years, both with errors[5] and without. For list-recovery in this setting, there has only been one positive result establishing that random linear codes can be list-recovered with list-size $O_{q,\ell}(1/\varepsilon)$ [DMRR25] (while to match plain random codes, one would hope for list-size $O(\ell/\varepsilon)$). However, there have been a number of *negative* results [GLM+21, RY24], in some cases showing that random linear codes provably perform *worse* than plain random codes. Notably, the restrictions for zero-error list-recovery arise over fields of small characteristic, but disappear over prime fields: similarly, to establish zero-error list-recovery bounds (and hence our application to leakage-resilient secret-sharing), we work over prime fields.

When q is large (at least $\exp(\Omega(1/\varepsilon))$), many recent works [ST23, BGM23, GZ23, AGL24, LMS25] have established that random linear codes (and, even more interestingly, random *Reed-Solomon* codes;[6] in this case $q \geq \Omega(n) \cdot \exp(\Omega(1/\varepsilon))$) of rate R are $(1-R-\varepsilon, O(1/\varepsilon))$-list-decodable. More generally, they can be shown to achieve the *generalized Singleton bound*, i.e., are $(\frac{L}{L+1}(1 - R - \varepsilon), L)$-list-decodable for any $L \geq 1$. For list-recovery over large alphabets ($q \geq \ell^{\Omega(1/\varepsilon)}$), recent works have shown matching lower [LS25] and upper [RW14, RW18, BCDZ26] bounds on the list-size required at decoding radius $1 - R - \varepsilon$: $L = \ell^{O(1/\varepsilon)}$. This is provably *worse* than plain random codes (where $L = O(\ell/\varepsilon)$ is sufficient); furthermore, we emphasize that this limitation is inherent over *every* finite field, not just small characteristic fields.

[4] We state here the version for general linear codes. The original formulation focused on Reed–Solomon codes, hence the word "polynomial," but the same question is natural for arbitrary linear codes.

[5] Briefly, a code C is (ρ, ℓ, L)-list-recoverable *from errors* if for any combinatorial rectangle $S_1 \times \cdots \times S_n$ with each $|S_i| \leq \ell$, we have $|\{c \in C : |\{i \in [n] : c_i \notin S_i\}| \leq \rho n\}| \leq L$. That is, we allow a ρ fraction of the coordinates c_i to not lie in the input list S_i.

[6] Reed-Solomon codes are defined by evaluating low-degree univariate polynomials at n distinct field elements; to generate a random Reed-Solomon code, one samples these evaluation points randomly.

1.3 AI Disclosure

We used ChatGPT 5.6 Sol to assist with editing, proofreading, and verification. All ideas and technical analysis are due to the authors, who take full responsibility for the correctness of all content.

2 Main Results

This section presents the formal versions of the paper's main discrepancy results. We begin by fixing notation and conventions, and then state two theorems on the discrepancy of random linear codes: the first controls all translations of a single test function, while the second applies to large families of Fourier-concentrated test functions. We then derive the corresponding applications to list-decoding and list-recovery.

The application to leakage-resilient secret-sharing schemes is treated separately in Section 6.

2.1 Notation and Conventions

Random Linear Codes

Fix a prime power q and integers $k, n \in \mathbb{N}$ with $k \le n$. In this work, a random linear code (RLC) of designed dimension k and length n over $\mathbb{F}_q$ is the image of a uniformly random matrix $G \in \mathbb{F}_q^{n\times k}$. Equivalently, its designed rate is k/n. We refer to this as the random generating matrix model.

Another common model samples C uniformly from the set of all k-dimensional subspaces of $\mathbb{F}_q^n$. The two models agree after conditioning on the event that G has full rank, since the image of a full-rank random matrix is a uniformly random k-dimensional subspace. Moreover, the latter event holds with probability at least $1 - \frac{q^{-(n-k+1)}}{1-1/q}$. Thus, whenever $n-k$ grows linearly with n, passing between the two models changes the stated probabilities only by a negligible term. This is the regime of the constant-rate applications below.

Fourier Transform

For functions $f : \mathbb{F}_q^n \to \mathbb{C}$, we use the expectation measure on the primal space and the counting measure on the Fourier side. The Fourier transform of f at $y \in \mathbb{F}_q^n$ is

$$\widehat{f}(y) = q^{-n} \sum_{x\in\mathbb{F}_q^n} f(x)\chi_y(x),$$

where χ_y is the additive character of $\mathbb{F}_q^n$ associated with y. Specifically, if $\mathbb{F}_q$ has characteristic p and $q = p^m$, then

$$\chi_y(x) = \exp\left(\frac{2\pi i}{p}\mathrm{tr}(\langle x, y\rangle)\right),$$

where $\langle x, y\rangle = \sum_{j=1}^n x_j y_j$ is the standard inner product over $\mathbb{F}_q$. Here $\mathrm{tr} : \mathbb{F}_q \to \mathbb{F}_p$ is the absolute field trace,

$$\mathrm{tr}(a) = a + a^p + a^{p^2} + \cdots + a^{p^{m-1}}.$$

Under these normalizations, for $r \geq 1$ we write

$$\|f\|_r = \left(q^{-n} \sum_{x \in \mathbb{F}_q^n} |f(x)|^r \right)^{1/r} \quad \text{and} \quad \left\| \widehat{f} \right\|_r = \left(\sum_{y \in \mathbb{F}_q^n} |\widehat{f}(y)|^r \right)^{1/r}.$$

We define convolution by

$$(f * g)(y) = q^{-n} \sum_{x \in \mathbb{F}_q^n} f(x) g(y - x).$$

With these choices, Parseval's identity takes the form $\|f\|_2 = \left\| \widehat{f} \right\|_2$, and convolution satisfies

$$\widehat{f * g} = \widehat{f} \cdot \widehat{g}.$$

A nonnegative function $f : \mathbb{F}_q^n \to \mathbb{R}_{\geq 0}$ is normalized if its expectation is 1, i.e.,

$$\|f\|_1 = \mathop{\mathbb{E}}_{x \sim \mathbb{F}_q^n} [f(x)] = 1.$$

We will use the following Fourier pseudorandomness condition.

Definition 2.1 (α-Fourier-concentration)**.** *A normalized function $f : \mathbb{F}_q^n \to \mathbb{R}_{\geq 0}$ is* α-Fourier-concentrated*, for $\alpha \in [0, 1]$, if for every $y \in \mathbb{F}_q^n$,*

$$\left| \widehat{f}(y) \right| \leq \alpha^{\mathrm{wt}(y)}.$$

Relative Deviation

In Eq. (3), we defined the relative deviation of a code C with respect to a set B. We now extend this notion from sets to nonnegative normalized functions. This formulation is more convenient when the tests are weighted functions rather than indicators.

For a nonempty set $B \subseteq \mathbb{F}_q^n$, let

$$f_B := \frac{q^n}{|B|} 1_B.$$

Then f_B is normalized, and the set-based relative deviation is recovered from the functional definition below.

Definition 2.2 (Relative deviation)**.** *Let $f : \mathbb{F}_q^n \to \mathbb{R}_{\geq 0}$ be normalized, and let $C \subseteq \mathbb{F}_q^n$ be nonempty. Write*

$$f(C) := \sum_{x \in C} f(x).$$

The relative deviation *of C with respect to f is*

$$\delta_{C,f} := \mathop{\mathbb{E}}_{x \sim C} [f(x)] - \mathop{\mathbb{E}}_{x \sim \mathbb{F}_q^n} [f(x)] = \frac{f(C)}{|C|} - 1.$$

In particular, for $f_B = (q^n/|B|)1_B$,

$$\delta_{C,f_B} = \frac{q^n|C \cap B|}{|C|\,|B|} - 1 = \delta_{C,B}.$$

Thus, $|\delta_{C,f_B}| \le \tau$ is equivalent to

$$(1-\tau)\frac{|C|\,|B|}{q^n} \le |C \cap B| \le (1+\tau)\frac{|C|\,|B|}{q^n}.$$

2.2 Discrepancy Results for Random Linear Codes

2.2.1 Small Relative Deviation for Translations of a Single Function

We first state a detailed version of Theorem 1.1. It shows that, for any fixed normalized nonnegative function f, a random linear code of rate slightly above the natural threshold for f has exponentially small relative deviation with respect to every translation of f.

For a normalized function f, define its capacity parameter by

$$\beta := \frac{1}{n}\log_q \|f\|_\infty .$$

This choice is motivated in detail in Section 3. For now, observe that if $B \subseteq \mathbb{F}_q^n$ is nonempty and $f = f_B = (q^n/|B|)1_B$, then

$$\beta = 1 - \frac{1}{n}\log_q |B|.$$

Moreover, for a code of rate R, the expected intersection of C with a uniformly random translate of B is

$$\frac{|C|\,|B|}{q^n} = q^{(R-\beta)n}.$$

Thus $R = \beta$ is the point at which this expected intersection size is 1, while rates $R = \beta + \eta$ correspond to the above-capacity regime where the expected intersection size is $q^{\eta n}$.

Definition 2.3 (Translation)**.** *For a function $f : \mathbb{F}_q^n \to \mathbb{C}$ and a vector $z \in \mathbb{F}_q^n$, define the* translation *$f_z : \mathbb{F}_q^n \to \mathbb{C}$ by*

$$f_z(x) = f(x+z).$$

Theorem 2.4 (Small deviation of a translation family)**.** *Fix a normalized function $f : \mathbb{F}_q^n \to \mathbb{R}_{\ge 0}$, and let $\beta = \frac{1}{n}\cdot\log_q \|f\|_\infty$. Let $\mathcal{F} = \{f_z : z \in \mathbb{F}_q^n\}$ be the family of all translations of f.*

Let $C \subseteq \mathbb{F}_q^n$ be a random linear code of rate $R = \beta + \eta$ (where $\eta > 0$ and $k = Rn$), and assume that n is large enough that $\eta n \ge 240 \log_q n$. Let $\varepsilon = \frac{\eta}{12}$. Then, with probability at least

$$1 - 4q^{1-\frac{\eta n}{3}} - \frac{q^{k-n+1}}{(q-1)^2} - q^{2n-\frac{\eta^2 n^2}{1440}}$$

over the choice of C, we have

$$\max_{z \in \mathbb{F}_q^n} |\delta_{C,f_z}| \le q^{-\varepsilon n} \quad .$$

2.2.2 Small Relative Deviation for a Family of Fourier-Concentrated Functions

We next state a detailed version of Theorem 1.2, dealing with general families of Fourier-concentrated functions.

Theorem 2.5 (Small deviation of Fourier-concentrated Functions)**.** *Fix a family $\mathcal{F}$ of normalized and α-Fourier-concentrated functions $f : \mathbb{F}_q^n \to \mathbb{R}_{\geq 0}$ (where $0 < \alpha < 1$), and set $\beta = \max_{f\in\mathcal{F}} \frac{1}{n} \cdot \log_q \|f\|_\infty$. Let $C \subseteq \mathbb{F}_q^n$ be a random linear code of rate $R = \beta + \eta \leq 1$ for some $\eta > 0$. Denote*

$$d = \max\left\{0, \left\lceil \log_2\left(\frac{\log_2\left((q-1)/(q^{\eta/100}-1)\right)}{\log_2(1/\alpha)}\right)\right\rceil\right\} .$$

Fix a parameter $t > 20\log_q(n)$ such that $2t/n < 2^{-d}\min\left(R - \frac{\eta}{50}, 2\eta\right)$ and let

$$\varepsilon = 2^{-d}\min\left(R - \frac{\eta}{50}, 2\eta\right) - \frac{2t}{n} .$$

Then, for sufficiently large n, with probability at least

$$1 - n(Rn+1)q^{-\frac{\eta n}{100}} - |\mathcal{F}| \cdot d \cdot q^{n-t^2/10}$$

over the choice of C, we have

$$|\delta_{C,f}| \leq q^{-\varepsilon n}$$

for every $f \in \mathcal{F}$.

The following corollary is a simplified and slightly weaker version of Theorem 2.5. It is intended to make the trade-off between the family size, the success probability, and the resulting deviation bound easier to parse.

Corollary 2.6 (Simplified version of Theorem 2.5)**.** *For every $n \in \mathbb{N}$, prime power q, and $0 < \alpha < 1$, the following holds. Let $\mathcal{F}$ be a family of normalized and α-Fourier-concentrated functions $f : \mathbb{F}_q^n \to \mathbb{R}_{\geq 0}$. Let*

$$\beta := \max_{f\in\mathcal{F}} \frac{1}{n}\log_q \|f\|_\infty .$$

Let $0 < \eta \leq 50\beta/51$, and let $C \subseteq \mathbb{F}_q^n$ be a random linear code of rate $R = \beta + \eta \leq 1$. Set

$$\lambda := \min\left\{1, \frac{\log(1/\alpha)}{2\log(100q/\eta)}\right\} .$$

Let $0 < \theta < \eta\lambda$. Then, assuming $n \geq n_0(q, \eta, \alpha, \theta)$, it holds with probability at least

$$1 - q^{-\frac{\eta n}{200}} - |\mathcal{F}|q^{n - \frac{\theta^2 n^2}{20}},$$

that

$$|\delta_{C,f}| \leq q^{-2(\eta\lambda-\theta)n}$$

for every $f \in \mathcal{F}$. In particular, it suffices to take

$$n_0 \geq \max\left\{\frac{c\log_q(1/\eta)}{\eta}, \frac{c}{\theta}\sqrt{\log_q\left(2 + \log\left(\frac{q}{\eta(1-\alpha)}\right)\right)}\right\} ,$$

where $c > 0$ is a sufficiently large universal constant.

Remark 2.7 (On the probability–deviation trade-off in Corollary 2.6)**.** *The parameter θ controls the trade-off between the success probability and the deviation bound. For a simple way to read the corollary, suppose that $\log_q |\mathcal{F}|$ grows at least linearly in n. If one only wants success probability $1 - o(1)$, one may take*

$$\theta := \frac{\sqrt{\log_q |\mathcal{F}|}}{n}\, g(n),$$

where $g(n) \to \infty$ sufficiently slowly that

$$g(n)\sqrt{\log_q |\mathcal{F}|} = o(\lambda n).$$

The resulting deviation bound is

$$|\delta_{C,f}| \le q^{-2(\eta\lambda-\theta)n} = q^{-2\eta\lambda n + 2g(n)\sqrt{\log_q |\mathcal{F}|}}.$$

In particular, if $|\mathcal{F}| \le q^{n^\kappa}$ for some fixed $1 \le \kappa < 2$, then the same calculation yields constants $\gamma = \gamma_{\kappa,\eta} > 0$ for which, in the corresponding parameter regimes, the bound takes the form

$$|\delta_{C,f}| \le q^{-n^\gamma},$$

for all sufficiently large n. For example, this applies in regimes such as $q \le 2^{n^\gamma}$ and $\alpha \le 1 - n^{-\gamma}$, after choosing γ sufficiently small as a function of κ and η.

We defer the proof of Corollary 2.6 to Section 4.3.3.

2.3 Applications to List-Decoding and List-Recovery

We now record two direct coding-theoretic consequences of Theorems 2.4 and 2.5. The definitions of list-decoding and zero-error list-recovery were given in the introduction; we repeat them here only to fix notation for the formal statements.

Definition 2.8 (List-decodable code)**.** *A code $C \subseteq \mathbb{F}_q^n$ is* (ρ, L)-list-decodable *if, for every $z \in \mathbb{F}_q^n$,*

$$|B_\rho(z) \cap C| \le L.$$

Recall that $1 - \frac{1}{n}\log_q |B_\rho|$ is the list-decoding capacity. Below this rate, uniformly random codes achieve constant list size with high probability; above this rate, the relevant benchmark is the expected intersection size with a Hamming ball. More precisely, a uniformly random code of rate $1 - \frac{1}{n}\log_q |B_\rho| + \gamma$ has about

$$\frac{|C| \cdot |B_\rho|}{q^n} \approx q^{\gamma n}$$

codewords in each Hamming ball with high probability. The next theorem shows that random linear codes satisfy the same above-capacity discrepancy guarantee.

Theorem 2.9 (Random linear codes achieve nearly optimal list size in list-decoding above capacity)**.** *Let $n \in \mathbb{N}$, let q be a prime power, let $0 < \rho < 1 - 1/q$, and let $0 < \gamma < \frac{1}{n}\log_q |B_\rho(0)|$, where $\gamma n \ge 240 \log_q n$. Set*

$$R := 1 - \frac{1}{n}\log_q |B_\rho(0)| + \gamma.$$

Let $C \subseteq \mathbb{F}_q^n$ be a random linear code of designed rate R. Then, with probability at least

$$1 - 4q^{1-\gamma n/3} - \frac{q^{-(1-R)n+1}}{(q-1)^2} - q^{2n-\gamma^2 n^2/1440} \quad ,$$

the following holds simultaneously for every $z \in \mathbb{F}_q^n$:

$$|B_\rho(z) \cap C| = \left(1 \pm q^{-\gamma n/12}\right) \frac{|C| \cdot |B_\rho(0)|}{q^n}. \tag{5}$$

In particular, C is (ρ, L)-list-decodable with

$$L = \left\lceil \left(1 + q^{-\gamma n/12}\right) q^{\gamma n} \right\rceil.$$

Remark 2.10. *Observe that the probability bound is $1-o(1)$ whenever $\gamma^2 n \to \infty$ and $(1-R)n \to \infty$*

Proof. We apply Theorem 2.4 to the normalized indicator of the Hamming ball centered at the origin,

$$f = \frac{q^n}{|B_\rho(0)|} \mathbf{1}_{B_\rho(0)}.$$

Its capacity parameter is

$$\beta_n := \frac{1}{n} \log_q \|f\|_\infty = 1 - \frac{1}{n} \log_q |B_\rho(0)|.$$

Thus Theorem 2.4 gives

$$\max_{z \in \mathbb{F}_q^n} |\delta_{C, f_z}| \le q^{-\gamma n/12}$$

with probability at least

$$1 - 4q^{1-\frac{\gamma n}{3}} - \frac{q^{Rn-n+1}}{(q-1)^2} - q^{2n-\frac{\gamma^2 n^2}{1440}} \quad .$$

The functions f_z are the normalized indicators of the balls $B_\rho(-z)$; as z ranges over $\mathbb{F}_q^n$, these are exactly all Hamming balls of radius ρn. Translating this relative-deviation bound gives Eq. (5).

Finally, $|C| \le q^{Rn}$ by construction, so

$$\frac{|C| \cdot |B_\rho(0)|}{q^n} \le q^{\gamma n}.$$

This yields the stated list-size bound. □

We now turn to zero-error list-recovery.

Definition 2.11 (Zero-error list-recoverable code)**.** *A code $C \subseteq \mathbb{F}_q^n$ is* (ℓ, L)-list-recoverable *if, for every collection of sets $S_1, S_2, \ldots, S_n \subseteq \mathbb{F}_q$ with $|S_i| \le \ell$ for all $i \in [n]$,*

$$|(S_1 \times S_2 \times \cdots \times S_n) \cap C| \le L.$$

We call ℓ the input list size *and L the* output list size.

For the upper-bound guarantee, it suffices to consider rectangles $S_1 \times \cdots \times S_n$ with $|S_i| = \ell$ for every i, since smaller input lists can be padded to size ℓ. The zero-error list-recovery capacity is $1 - \log_q \ell$: below this rate uniformly random codes achieve constant output list size, while above this rate the expected intersection with an ℓ-by-$\cdots$-by-ℓ rectangle is exponential. We focus on the nontrivial range $2 \le \ell \le q-1$, since $\ell = 1$ and $\ell = q$ give the trivial bounds $L = 1$ and $L = |C|$, respectively.

Theorem 2.12 (Random linear codes achieve nearly optimal output list size in zero-error list-recovery above capacity)**.** *There exist universal constants $c, c' > 0$ such that the following holds. Let $n \in \mathbb{N}$, let q be prime, and let $2 \le \ell \le q-1$. Let*

$$0 < \gamma \le \frac{1}{2} \min\{\log_q \ell,\, 1 - \log_q \ell\},$$

and set $R = 1 - \log_q \ell + \gamma$. Assume that

$$n \ge c' \cdot \frac{q^5 \log^2(q/\gamma)}{\gamma^2} \quad . \tag{6}$$

Let $C \subseteq \mathbb{F}_q^n$ be a random linear code of designed rate R. Then, with probability at least

$$1 - 2q^{-c\gamma n/(q^2 \log(q/\gamma))}, \tag{7}$$

the following events both hold:

1. *the code C is (ℓ, L)-list-recoverable with*

$$L = \left\lceil \left(1 + q^{-c\gamma n/(q^2 \log(q/\gamma))}\right) q^{\gamma n} \right\rceil \quad .$$

2. *For every combinatorial rectangle $S = S_1 \times \cdots \times S_n$ with $|S_i| = \ell$ for each i, we have*

$$|C \cap S| = \left(1 \pm q^{-c\gamma n/(q^2 \log(q/\gamma))}\right) \frac{|C| \ell^n}{q^n} \quad .$$

We give the proof of that theorem, which is a careful application of Corollary 2.6, in Appendix A.1.

Essentially the same theorem can also be proven for list recovery from *erasures*, wherein we allow some fraction of the coordinates to be fully corrupt, i.e., have $S_i = \mathbb{F}_q$.

Definition 2.13 (Erasure list-recoverable code)**.** *A code $C \subseteq \mathbb{F}_q^n$ is* (ζ, ℓ, L)-list-recoverable *if for every collection of sets $S_1, S_2, \ldots, S_n \subseteq \mathbb{F}_q$ with $|S_i| \le \ell$ for at least $(1-\zeta)n$ of the i's it holds that*

$$|(S_1 \times S_2 \times \cdots \times S_n) \cap C| \le L.$$

The erasure list-recovery capacity is $R^\star = (1-\zeta)(1-\log_q \ell)$, which coincides with the zero-error setting when $\zeta = 0$.

Theorem 2.14 (Random linear codes achieve nearly optimal output list size in erasure list-recovery above capacity)**.** *There exist universal constants $c, c' > 0$ such that the following holds. Let $n \in \mathbb{N}$, let q be prime, and let $2 \le \ell \le q-1$ and $\zeta \in (0,1)$. Let $\beta = (1-\zeta)(1-\log_q \ell)$,*

$$0 < \gamma \le \frac{1}{2}\min\{1-\beta, \beta\},$$

and set $R = \beta + \gamma$.

Let $C \subseteq \mathbb{F}_q^n$ be a random linear code of designed rate R. Then, assuming the lower bound on n given in Equation (6) above holds, then the following events both hold, with the probability bound given in Equation (7).

1. *The code C is (ζ, ℓ, L)-list-recoverable with*

$$L = \left\lceil \left(1 + q^{-c\gamma n/(q^2 \log(q/\gamma))}\right) q^{\gamma n} \right\rceil .$$

2. *For every combinatorial rectangle $S = S_1 \times \cdots \times S_n$ with $|S_i| = \ell$ for $(1-\zeta)n$ of the i's, and $|S_i| = q$ for the rest of them, we have*

$$|C \cap S| = \left(1 \pm q^{-c\gamma n/(q^2 \log(q/\gamma))}\right) |C| \left(\frac{\ell}{q}\right)^{(1-\zeta)n} .$$

We defer the (nearly identical) proof to Appendix A.1.

3 Proof Overview

We organize the proof around two motivating examples. We begin with binary list-decoding, where the relevant sets are Hamming balls of a fixed radius. This illustrates Theorem 2.4: all such balls are translates of the ball centered at the origin, so the family consists of the translations of a single normalized indicator function. We then turn to list-recovery over a constant-size alphabet, where the relevant sets are combinatorial rectangles

$$A_1 \times \cdots \times A_n, \qquad |A_i| = \ell, \qquad 2 \le \ell \le q-1.$$

This illustrates Theorem 2.5: these rectangle indicators are α-Fourier-concentrated for some $\alpha < 1$ bounded away from 1, and the number of such rectangles is only $q^{o(n^2)}$.

Hamming Balls and the Translation-Family Case

Let

$$B = B(0, \rho n) \subseteq \mathbb{F}_2^n$$

be the Hamming ball of radius ρn around the origin, and let

$$f_B = \frac{2^n}{|B|} \cdot 1_B$$

be its normalized indicator. The Hamming ball around a center $z \in \mathbb{F}_2^n$ is $B + z$, whose normalized indicator is the translate

$$(f_B)_z(x) = f_B(x + z).$$

Thus, in the binary list-decoding example, the family we need to control is precisely

$$\{(f_B)_z : z \in \mathbb{F}_2^n\},$$

the set of all translates of one function. This is the motivating special case of Theorem 2.4.

We expose the random code one generator at a time. That is, starting from $C_0 = \{0\}$, we consider the random sequence

$$C_0 \subseteq C_1 \subseteq \cdots \subseteq C_k, \qquad C_i = C_{i-1} + \operatorname{span}\{u_i\},$$

where the vectors $u_1, \ldots, u_k$ are chosen independently and uniformly from $\mathbb{F}_2^n$. The final code is $C = C_k$.

The reason for using this incremental construction is that the deviation δ has a simple expected evolution in one step. Suppose

$$C' = C + \operatorname{span}\{u\}.$$

Over $\mathbb{F}_2$, this means

$$C' = C \cup (C + u).$$

The old coset C contributes the old deviation, while the new coset $C + u$ behaves like a random translate of C. A direct calculation gives

$$\mathop{\mathbb{E}}_{u}\left[\delta_{C',f}\right] = \frac{1}{2}\delta_{C,f}.$$

Thus one new generator should ideally divide the deviation by 2.

This explains the rate threshold: If

$$\beta = \frac{1}{n}\log_2 \|f_B\|_\infty$$

then the initial deviation of a translate is at most on the order of $2^{\beta n}$. If $k = Rn = (\beta + \eta)n$ and each of the k steps divided the deviation by 2, then the final deviation would be roughly

$$2^{\beta n - k} = 2^{-\eta n},$$

which is exponentially small.

We formalize this ideal behavior using smoothness. Given a function g and a parameter Γ, we say that the construction sequence is (g, Γ)-smooth if

$$|\delta_{C_i,g}| \le 2^{k-i}\Gamma \qquad \text{for every } 0 \le i \le k.$$

This is a two-sided deviation bound. It says that the deviation at time i is no larger than what one would expect if the final deviation were Γ and each preceding step divided the deviation by 2. In the q-ary setting, the corresponding definition uses $q^{k-i}\Gamma$ in place of $2^{k-i}\Gamma$.

An iterative dimension-by-dimension viewpoint also appears in Blinovsky's high-probability covering-radius results [Bli87, Bli90a] (see [CHLL97, Theorems 12.3.5 and 12.3.10]), but, as discussed earlier, the discrepancy problem requires a more precise version of this process. For comparison, in a covering argument with g the normalized indicator of a ball, it is enough to prove the one-sided condition

$$\delta_{C_i,g} > -1$$

for some i, since this is exactly the statement that the ball contains at least one codeword of C_i. Once this happens, it remains true for every later code $C_j \supseteq C_i$. In the present problem we instead want $|\delta_{C_k,g}|$ to be small, meaning that the final code samples the ball with approximately the correct multiplicity. This property is not monotone: even if $|\delta_{C_i,g}|$ is small at some intermediate stage, later added codewords may oversample the same ball and make the deviation large again. Although the theorem only concerns the final code C_k, our proof method requires the stronger inductive invariant that the deviations are controlled throughout the construction.

A first attempt would be to prove smoothness directly for every ball $B+z$. Indeed, the second-moment smoothness lemma (see Lemma 4.5) shows that, for a fixed normalized function g and any $\Gamma > 2^{-k}g(0)$, the failure probability is at most

$$\frac{(k+1)2^{-k}}{(\Gamma - 2^{-k}g(0))^2}\,\|g\|_2^2 + 2^{k-n}.$$

In the overview, it is useful to read this as roughly $(k+1)2^{-k}\Gamma^{-2}\,\|g\|_2^2$, as long as Γ is not too close to $2^{-k}g(0)$. Since

$$\|(f_B)_z\|_2^2 \le \|(f_B)_z\|_\infty = \|f_B\|_\infty = 2^{\beta n},$$

this gives exponentially small failure probability for each fixed center z, provided Γ is chosen exponentially small but still larger than $2^{-k}(f_B)_z(0)$ with sufficient slack.

This direct argument is not enough. There are 2^n possible centers, while the second-moment estimate gives only a bound of order 2^{-cn} for some small constant $c > 0$. The latter is not large enough to support a union bound over all centers. Thus we need a way to upgrade the probability after paying for only one direct second-moment estimate. We apply the direct second-moment estimate only once, to a common auxiliary function, and then use that common event to control all translates with much better conditional probability.

In the binary case, the auxiliary function is the ordinary self-convolution

$$F_f = f * f.$$

For the ball function f_B, the value $F_{f_B}(w)$ is a normalized count of pairs

$$x, y \in B \qquad \text{with} \qquad x + y = w.$$

Thus controlling F_{f_B} on a code is the smoothed version of controlling

$$(B+B) \cap C.$$

The reason for introducing F_f is the one-step variance identity

$$\mathrm{Var}_u\left(\delta_{C',f}\right) = \frac{1}{4}\,\delta_{C,F_f}\ ,$$

readily proven as part of Lemma 4.12. Thus F_f measures the noise in the update of $\delta_{C,f}$. If the construction sequence is smooth for F_f, then the update for f is concentrated around its mean $\delta_{C_i,f}/2$ at every step. Consequently, unless the deviation of f is already small, the next step is very likely to reduce it by almost the expected factor 2.

We now bootstrap as follows. First, apply the second-moment bound only to the single function F_{f_B}. This is possible because F_{f_B} has essentially the same second-moment size as f_B:

$$\|F_{f_B}\|_2 = \|f_B * f_B\|_2 \le \|f_B\|_1 \|f_B\|_2 = \|f_B\|_2 ,$$

and hence

$$\|F_{f_B}\|_2^2 \le \|f_B\|_2^2 \le \|f_B\|_\infty = 2^{\beta n}.$$

Therefore Lemma 4.5 implies that, with probability at least $1 - 2^{-\Omega(n)}$, the construction sequence is

$$(F_{f_B}, \Gamma)\text{-smooth}$$

for a suitable exponentially small $\Gamma = 2^{-cn}$.

The crucial point is that this smoothness event is common to all balls. For every translate $(f_B)_z$, the binary self-convolution appearing in the variance identity is the same:

$$(f_B)_z * (f_B)_z = f_B * f_B.$$

Thus we do not need to prove smoothness separately for the self-convolution of each ball. We condition on the single event that F_{f_B} is smooth.

After this conditioning, the probability scale improves dramatically. The smoothness-transfer lemma, Lemma 4.14, says that if the sequence is (F_{f_B}, Γ)-smooth, then for any fixed translate $(f_B)_z$, the conditional probability that $(f_B)_z$ fails to be smooth is at most $2^{-\Omega(t^2)}$, where t is a slackness parameter. A union bound over all 2^n centers gives failure probability at most

$$2^{n-\Omega(t^2)}.$$

Taking $t = \Theta(n)$ makes this probability $2^{-\Omega(n^2)}$.

This proves the translation-family theorem in the motivating binary list-decoding case. The statement of Theorem 2.4 is more general: it applies to arbitrary normalized functions and to q-ary alphabets. For $q > 2$, the only change is that the ordinary self-convolution is replaced by the mirrored self-convolution

$$F_f = f * \check{f}, \qquad \check{f}(x) = f(-x) \ .$$

Combinatorial Rectangles and the General Family Case

We now turn to the setting that motivates Theorem 2.5. Think of zero-error list-recovery over a constant-size alphabet $\mathbb{F}_q$. A typical bad set is a combinatorial rectangle

$$S_1 \times S_2 \times \cdots \times S_n, \qquad S_i \subseteq \mathbb{F}_q, \quad |S_i| = \ell,$$

where

$$2 \le \ell \le q - 1.$$

Let f_S denote the normalized indicator of such a rectangle. The important facts about these functions are the following. First, they are α-Fourier-concentrated for some constant $\alpha < 1$ depending only on q and ℓ. Indeed, the Fourier transform factors coordinate by coordinate, and the one-dimensional nontrivial Fourier coefficients are bounded away from 1. Second, there are

$$\binom{q}{\ell}^n \le q^{o(n^2)}$$

such rectangles in the constant-alphabet regime under discussion. More generally, Theorem 2.5 is designed to handle any family of $q^{o(n^2)}$ normalized α-Fourier-concentrated functions.

The argument for Hamming balls breaks down because the functions in this family are not all translates of a single function. For different rectangles $f \in \mathcal{F}$, the mirrored self-convolutions F_f may be different. Thus we cannot apply the direct second-moment estimate to each F_f separately: that estimate gives only $q^{-\Omega(n)}$ failure probability for each fixed function, whereas the family size allowed in Theorem 2.5 is much larger than what such a direct union bound can handle.

The extra assumption we use is Fourier-concentration. The mirrored self-convolution improves Fourier-concentration: if f is α-Fourier-concentrated, then F_f is α^2-Fourier-concentrated. Iterating this operation, define

$$f^{(0)} = f, \qquad f^{(r+1)} = F_{f^{(r)}}.$$

Then $f^{(r)}$ is α^{2^r}-Fourier-concentrated. We choose d so that

$$1 + (q-1)\alpha^{2^d} \le q^{\eta/100}.$$

At this depth, the Fourier coefficients of $f^{(d)}$ decay quickly enough that smoothness follows from a mild code-dependent event.

The event we need is that the dual codes do not have too many vectors at any Hamming weight. More precisely, for

$$W_j = \{y \in \mathbb{F}_q^n : \mathrm{wt}(y) = j\},$$

we require

$$|C_i^\perp \cap W_j| \le \gamma q^{-i}|W_j| \qquad \text{for every } 0 \le i \le k \text{ and every } 1 \le j \le n,$$

where $\gamma = q^{\eta n/100}$. This event depends only on the construction sequence, not on the function f. It holds with high probability by a standard first-moment argument applied to the sets W_j.

Assume this dual weight event holds. If g is normalized and α'-Fourier-concentrated, then

$$|\widehat{g}(y)| \le (\alpha')^{\mathrm{wt}(y)}.$$

Using the Fourier identity

$$\delta_{C_i,g} = \sum_{y \in C_i^\perp \setminus \{0\}} \widehat{g}(y),$$

we obtain

$$|\delta_{C_i,g}| \le \sum_{y \in C_i^\perp \setminus \{0\}} |\widehat{g}(y)| \le \gamma q^{-i} \sum_{j=1}^{n} |W_j|(\alpha')^j \le \gamma q^{-i}\big(1 + (q-1)\alpha'\big)^n.$$

Thus, under the dual weight event, every sufficiently Fourier-concentrated function is smooth. Applying this to $g = f^{(d)}$ gives

$$C_0, \dots, C_k \text{ is } (f^{(d)}, \Gamma_d)\text{-smooth},$$

with

$$\Gamma_d = q^{-k}\gamma\big(1 + (q-1)\alpha^{2^d}\big)^n \le q^{-k+\eta n/50}.$$

We then build back from $f^{(d)}$ to the original function $f = f^{(0)}$. At each step we use the same smoothness-transfer lemma as in the translation-family case. If

$$f^{(r)} = F_{f^{(r-1)}}$$

is smooth with parameter Γ_r, then $f^{(r-1)}$ is smooth with parameter

$$\Gamma_{r-1} = q^t \max\left\{\|f^{(r-1)}\|_\infty q^{-k}, \sqrt{\Gamma_r}\right\}.$$

except with probability $q^{n-\Omega(t^2)}$.

The square root is the main quantitative loss. If $\Gamma_r \approx q^{-an}$, then this step gives

$$\Gamma_{r-1} \approx q^{-an/2}$$

up to the slack factor q^t. Thus each reverse step essentially halves the smoothness exponent. After d steps, this accounts for the factor 2^{-d} in the final exponent. Since the starting functions are α-Fourier-concentrated with $\alpha < 1$ bounded away from 1, the required number of iterations d does not grow with n.

The slack factor q^t is chosen to make the conditional failure probability small enough. Taking

$$t = \Theta(2^{-d}\eta n)$$

makes

$$q^{n-\Omega(t^2)} = q^{-\Omega(n^2)}.$$

This is strong enough to union bound over all $f \in \mathcal{F}$ and over all d levels of the convolution iteration, even when $|\mathcal{F}| = q^{o(n^2)}$.

Finally, the initial terms in the recurrence are small because $R = \beta + \eta$. For every relevant iterate,

$$\|f^{(r)}\|_\infty q^{-k} \le q^{-\eta n},$$

using the bound $\|f^{(r)}\|_\infty \le \|f\|_\infty \le q^{\beta n}$. Unfolding the recurrence for the Γ_r's gives a final smoothness parameter

$$\Gamma_0 \le q^{-\varepsilon n},$$

where

$$\varepsilon = \Theta(\eta 2^{-d})$$

with the constants made explicit in Theorem 2.5. Hence

$$|\delta_{C_k,f}| \le q^{-\varepsilon n}$$

for every function in the family. The dual weight event is common to the entire family, and the remaining bootstrapping failures have $q^{-\Omega(n^2)}$ probability, so the final union bound gives Theorem 2.5.

4 Random Linear Codes are Smooth—Proof of Theorems 2.4 and 2.5

In this section we prove Theorems 2.4 and 2.5. We begin with several definitions. As explained in Section 3, we view a given RLC as being constructed one dimension at a time, leading us to define a construction sequence.

Definition 4.1. *A* construction sequence *is a sequence of codes $C_0, \ldots, C_k \subseteq \mathbb{F}_q^n$ where $C_i = C_{i-1} + \mathrm{span}\{u_i\}$ for some vector $u_i \in \mathbb{F}_q^n$. When $u_1, \ldots, u_k$ are chosen uniformly at random, we say that this is a* random construction sequence*, and each C_i is an* RLC *of designed dimension i.*

We would like to talk about a code having small deviation with regard to a specific function not just for the full code but also throughout its construction sequence. This requirement is captured by the notion of smoothness, motivated in Section 3.

Definition 4.2. *Fix a non-zero function $f : \mathbb{F}_q^n \to \mathbb{R}_{\geq 0}$. A construction sequence $C_0, \ldots, C_k$ is* (f, Γ, j)-smooth *(for some real $\Gamma \geq 0$ and integer $0 \leq j \leq k$) if*

$$|\delta_{C_i, f}| \leq q^{k-i} \cdot \Gamma$$

for all $0 \leq i \leq j$. As a shorthand, we write (f, Γ)-smooth to mean (f, Γ, k)-smooth.

Definition 4.3. *For $0 \leq j \leq n$, let $W_j = \left\{x \in \mathbb{F}_q^n : wt(x) = j\right\}$, where $wt(x)$ denotes the Hamming weight.*

Let $C_0, \ldots, C_k$ be a construction sequence. We say that this sequence is γ-dual-typical *($\gamma \geq 1$) if*

$$\left|C_i^{\perp} \cap W_j\right| \leq \gamma \cdot q^{-i} \cdot |W_j|$$

for all $0 \leq i \leq k$ and $1 \leq j \leq n$.

In Section 4.1 we begin with a weaker probabilistic argument that shows that RLCs are smooth for any fixed given function, and also dual-typical, with probability $q^{-\Omega(n)}$. In Section 4.2 we introduce the averaged convolution function F_f and relate δ_{C_i, F_f} to the i-th step change in $\delta_{C_{i+1}, f}$. In Section 4.3 we use this relation to control the smoothness of f via that of F_f, and use this implication as the basis to prove Theorems 2.4 and 2.5.

4.1 Smoothness via a Direct Second Moment Argument

In this subsection we use a simple second-moment argument to show that RLCs are smooth and dual-typical, with probability $1 - q^{-\Omega(n)}$. The core of the argument is the following lemma.

Lemma 4.4. *Fix $n \in \mathbb{N}$ and $m \leq n$. Let $f : \mathbb{F}_q^n \to \mathbb{R}_{\geq 0}$ be normalized, and let C be the linear code generated by a uniformly random matrix $G \in \mathbb{F}_q^{n \times m}$. Then, for every $\Delta > q^{-m}|f(0) - 1|$,*

$$\Pr_G\left[|\delta_{C,f}| > \Delta\right] \leq \frac{q-1}{q^m\left(\Delta - q^{-m}|f(0)-1|\right)^2}\|f\|_2^2 + \frac{q^{m-n}}{q-1} \ .$$

Proof. We evaluate the relative deviation by first considering the multiset of all vectors generated by G. Define

$$X = \sum_{v \in \mathbb{F}_q^m} f(Gv) \ ,$$

and let $\delta_{G,f} = q^{-m}X - 1$ be the relative deviation over this multiset.

If G has full column rank, the multiset $\{Gv : v \in \mathbb{F}_q^m\}$ contains exactly q^m distinct vectors, meaning it coincides perfectly with the code subspace C. In this case, $\delta_{C,f} = \delta_{G,f}$. Therefore, the event that the set deviation exceeds Δ can only occur if either the multiset deviation exceeds Δ, or G fails to have full rank. By the union bound:

$$\Pr_G\left[|\delta_{C,f}| > \Delta\right] \le \Pr_G\left[|\delta_{G,f}| > \Delta\right] + \Pr_G\left[\operatorname{rank}(G) < m\right] \ .$$

We first bound the multiset deviation unconditionally over the choice of G. Since f is normalized, $\sum_{x\in\mathbb{F}_q^n} f(x) = q^n$. For any fixed $v \neq 0$, the vector Gv is uniformly distributed over $\mathbb{F}_q^n$, so $\mathbb{E}_G\left[f(Gv)\right] = 1$. The expectation of X is

$$\mathop{\mathbb{E}}_G[X] = f(0) + \sum_{v\in\mathbb{F}_q^m\setminus\{0\}} \mathop{\mathbb{E}}_G\left[f(Gv)\right] = f(0) + q^m - 1 \ .$$

Hence, the exact expected multiset deviation is

$$\mu = \mathop{\mathbb{E}}_G[\delta_{G,f}] = q^{-m}\mathop{\mathbb{E}}_G[X] - 1 = q^{-m}\left(f(0) - 1\right) \ .$$

We next bound the variance of X. Since $f(0)$ is a constant, we can expand the variance using the sum definition:

$$\operatorname{Var}_G(X) = \sum_{v,w\in\mathbb{F}_q^m\setminus\{0\}} \left(\mathop{\mathbb{E}}_G\left[f(Gv)f(Gw)\right] - \mathop{\mathbb{E}}_G\left[f(Gv)\right]\mathop{\mathbb{E}}_G\left[f(Gw)\right]\right) \ .$$

If v and w are linearly independent, the vectors Gv and Gw are independent and uniformly distributed in $\mathbb{F}_q^n$. Thus, $\mathbb{E}_G\left[f(Gv)f(Gw)\right] = \mathbb{E}_G\left[f(Gv)\right]\mathbb{E}_G\left[f(Gw)\right] = 1$, making the term exactly zero.

The only non-zero terms in the double sum occur when v and w are linearly dependent, meaning $w = \alpha v$ for some $\alpha \in \mathbb{F}_q^*$. In this case, $Gw = \alpha Gv$. Since f is non-negative, we can form an upper bound by dropping the subtracted product $\mathbb{E}_G\left[f(Gv)\right]\mathbb{E}_G\left[f(Gw)\right] = 1$:

$$\mathop{\mathbb{E}}_G\left[f(Gv)f(\alpha Gv)\right] - 1 \le \mathop{\mathbb{E}}_G\left[f(Gv)f(\alpha Gv)\right] = q^{-n}\sum_{x\in\mathbb{F}_q^n} f(x)f(\alpha x) \ .$$

By Cauchy–Schwarz, for every $\alpha \in \mathbb{F}_q^*$:

$$\sum_{x\in\mathbb{F}_q^n} f(x)f(\alpha x) \le \sqrt{\sum_{x\in\mathbb{F}_q^n} f(x)^2}\sqrt{\sum_{x\in\mathbb{F}_q^n} f(\alpha x)^2} = \sum_{x\in\mathbb{F}_q^n} f(x)^2 = q^n\left\|f\right\|_2^2 \ .$$

Substituting this back into the variance expansion, we sum only over the pairs $(v, \alpha v)$. There are $q^m - 1$ choices for $v \neq 0$ and $q - 1$ choices for $\alpha \in \mathbb{F}_q^*$, giving:

$$\operatorname{Var}_G(X) \le \sum_{v\in\mathbb{F}_q^m\setminus\{0\}}\sum_{\alpha\in\mathbb{F}_q^*} q^{-n}\left(q^n\left\|f\right\|_2^2\right) = (q^m - 1)(q-1)\left\|f\right\|_2^2 \le q^m(q-1)\left\|f\right\|_2^2 \ .$$

Scaling down to the relative deviation $\delta_{G,f}$, we get

$$\mathrm{Var}_G\left(\delta_{G,f}\right) = q^{-2m}\mathrm{Var}_G\left(X\right) \le \frac{q-1}{q^m}\left\|f\right\|_2^2 \ .$$

By Chebyshev's inequality, to bound the absolute deviation we measure the distance from the mean μ to the boundary of the interval $[-\Delta, \Delta]$. Provided $\Delta > |\mu|$, we have:

$$\Pr_G\left[|\delta_{G,f}| > \Delta\right] \le \Pr_G\left[|\delta_{G,f} - \mu| > \Delta - |\mu|\right] \le \frac{\mathrm{Var}_G\left(\delta_{G,f}\right)}{\left(\Delta - q^{-m}|f(0)-1|\right)^2} \ .$$

Substituting our variance bound yields:

$$\Pr_G\left[|\delta_{G,f}| > \Delta\right] \le \frac{q-1}{q^m\left(\Delta - q^{-m}|f(0)-1|\right)^2}\left\|f\right\|_2^2 \ .$$

Finally, we bound the probability that G fails to have full column rank. The matrix G is rank-deficient if and only if at least one of its m columns falls into the span of the previous columns. The i-th column falls into the span of the first $i-1$ columns with probability at most q^{i-1}/q^n. Applying the union bound over all m columns gives

$$\Pr_G\left[\mathrm{rank}(G) < m\right] \le \sum_{i=1}^{m}\frac{q^{i-1}}{q^n} = q^{-n}\frac{q^m-1}{q-1} \le \frac{q^{m-n}}{q-1} \ .$$

Adding this rank-deficiency bound to the multiset deviation bound completes the proof. □

Using Lemma 4.4 across every dimension in the construction sequence and union-bounding the failures, we can show that a random linear code is highly likely to be smooth with respect to f.

Lemma 4.5. *Let $f : \mathbb{F}_q^n \to \mathbb{R}_{\ge 0}$ be normalized, and let $C_0, \ldots, C_k$ be a random construction sequence. If*

$$\Gamma > q^{-k}|f(0)-1|,$$

then $C_0, \ldots, C_k$ is (f, Γ)-smooth with probability at least

$$1 - \frac{(k+1)(q-1)q^{-k}}{\left(\Gamma - q^{-k}|f(0)-1|\right)^2}\left\|f\right\|_2^2 - \frac{q^{k-n+1}}{(q-1)^2} \ .$$

Proof. Fix $0 \le i \le k$. Since C_i is an RLC of dimension i, applying Lemma 4.4 with

$$\Delta_i = q^{k-i}\Gamma$$

gives

$$\begin{aligned}
\Pr_{C_i}\left[|\delta_{C_i,f}| > q^{k-i}\Gamma\right] &\le \frac{q-1}{q^i\left(q^{k-i}\Gamma - q^{-i}|f(0)-1|\right)^2}\left\|f\right\|_2^2 + \frac{q^{i-n}}{q-1} \\
&= \frac{(q-1)q^i}{\left(q^k\Gamma - |f(0)-1|\right)^2}\left\|f\right\|_2^2 + \frac{q^{i-n}}{q-1} \\
&\le \frac{(q-1)q^k}{\left(q^k\Gamma - |f(0)-1|\right)^2}\left\|f\right\|_2^2 + \frac{q^{i-n}}{q-1} \\
&= \frac{(q-1)q^{-k}}{\left(\Gamma - q^{-k}|f(0)-1|\right)^2}\left\|f\right\|_2^2 + \frac{q^{i-n}}{q-1} \ .
\end{aligned}$$

A union bound over $i = 0, \ldots, k$ gives the bound on the sum of the first terms, and sums the rank-deficiency error as $\sum_{i=0}^{k} \frac{q^{i-n}}{q-1} \leq \frac{q^{k-n+1}}{(q-1)^2}$. □

We next show that a random construction sequence is dual-typical with high probability.

Lemma 4.6. *A random construction sequence $C_0, \ldots, C_k$ is γ-dual-typical with probability at least*

$$1 - \frac{n(k+1)}{\gamma}.$$

Proof. Fix $0 \leq i \leq k$ and $1 \leq j \leq n$, and write $X_{i,j} := |C_i^{\perp} \cap W_j|$. For every fixed nonzero $y \in W_j$, the event $y \in C_i^{\perp}$ is the event that $\langle y, u_s \rangle = 0$ for every $s \in [i]$. Since $y \neq 0$ and the generators u_s are independent and uniform in $\mathbb{F}_q^n$, this event has probability q^{-i}. Hence $\mathbb{E}[X]_{i,j} = q^{-i}|W_j|$.

By Markov's inequality,

$$\Pr\left[X_{i,j} > \gamma q^{-i}|W_j|\right] \leq \frac{1}{\gamma}.$$

A union bound over the $n(k+1)$ choices of (i, j) proves the claim. □

We finish this section by showing that a dual-typical code is smooth on all sufficiently Fourier-concentrated functions.

Lemma 4.7. *Suppose that the construction sequence $C_0, \ldots, C_k$ is γ-dual-typical and that $f : \mathbb{F}_q^n \to \mathbb{R}_{\geq 0}$ is normalized and α-Fourier-concentrated. Then $C_0, \ldots, C_k$ is $(f, q^{-k} \cdot \gamma\,(1 + (q-1)\alpha)^n)$-smooth.*

Proof. Observe that

$$\begin{aligned}
|\delta_{C_i,f}| &\leq \sum_{y \in C_i^{\perp} \setminus \{0\}} \left|\widehat{f}(y)\right| \\
&\leq \sum_{j=1}^{n} |C_i^{\perp} \cap W_j| \cdot \alpha^j \\
&\leq \sum_{j=1}^{n} \gamma \cdot q^{-i} \cdot |W_j| \cdot \alpha^j \\
&= \gamma \cdot q^{-i} \cdot \sum_{j=1}^{n} \binom{n}{j} \cdot (q-1)^j \cdot \alpha^j \\
&= \gamma \cdot q^{-i} \cdot ((1 + (q-1)\alpha)^n - 1) \ .
\end{aligned}$$

□

4.2 Self Convolution and Variance of the Relative Deviation

The direct second-moment bound from Section 4.1 yields a failure probability of $q^{-\Omega(n)}$, which is too weak to union-bound over exponentially large function families. To achieve the required $q^{-\Omega(n^2)}$ probability, we must sharply bound the variance of the relative deviation at each step of the code's construction. Remarkably, this one-step variance is entirely controlled by the code's deviation with respect to an auxiliary smoothed function, which we define here as the averaged convolution F_f.

Definition 4.8 (Self-Convolution Function F_f)**.** *Let $f : \mathbb{F}_q^n \to \mathbb{R}_{\geq 0}$ be a function. We define the* reflected function *$\check{f} : \mathbb{F}_q^n \to \mathbb{R}$ by $\check{f}(x) = f(-x)$. We then define the* self-convolution function *F_f as $F_f = f * \check{f}$. Equivalently, F_f can be defined by*

$$\begin{aligned} F_f(w) &= q^{-n} \sum_{x \in \mathbb{F}_q^n} f(x) \cdot \check{f}(w - x) \\ &= q^{-n} \sum_{x \in \mathbb{F}_q^n} f(x) \cdot f(x - w) \ . \end{aligned}$$

We next prove several simple properties of F_f.

Lemma 4.9 (Fourier transform of F_f)**.** *Let $f : \mathbb{F}_q^n \to \mathbb{R}_{\geq 0}$ be a normalized function. Then,*

$$\widehat{F_f}(y) = \left|\widehat{f}(y)\right|^2$$

for all $y \in \mathbb{F}_q^n$.

Proof. First, observe that

$$\widehat{\check{f}}(y) = q^{-n} \sum_{x \in \mathbb{F}_q^n} f(-x)\chi_y(x) = q^{-n} \sum_{x \in \mathbb{F}_q^n} f(x)\chi_y(-x) = q^{-n} \sum_{x \in \mathbb{F}_q^n} f(x)\overline{\chi_y(x)} = \overline{\widehat{f}(y)} \ .$$

Thus, by the Convolution Identity,

$$\widehat{F_f}(y) = \widehat{f}(y) \cdot \widehat{\check{f}}(y) = \widehat{f}(y) \cdot \overline{\widehat{f}(y)} = \left|\widehat{f}(y)\right|^2 \ .$$

□

Lemma 4.9 immediately yields the following corollary.

Lemma 4.10 (α-Fourier-concentration of F_f)**.** *Let $f : \mathbb{F}_q^n \to \mathbb{R}_{\geq 0}$ be a normalized function. If f is α-Fourier-concentrated, then its self-convolution function F_f is α^2-Fourier-concentrated.*

Lemma 4.11 (Translation Invariance of F_f)**.** *For any function $f : \mathbb{F}_q^n \to \mathbb{R}$ and any translation vector $z \in \mathbb{F}_q^n$, let f_z be its translation. Then $F_{f_z} = F_f$.*

Proof. We prove this directly using the expanded definition of the self-convolution. Evaluating F_{f_z} at any point $w \in \mathbb{F}_q^n$ yields:

$$F_{f_z}(w) = q^{-n} \sum_{x \in \mathbb{F}_q^n} f_z(x) \cdot f_z(x - w) = q^{-n} \sum_{x \in \mathbb{F}_q^n} f(x + z) \cdot f(x - w + z) \ .$$

Let us apply the change of variables $y = x + z$. As x iterates over all of $\mathbb{F}_q^n$, y also iterates over all of $\mathbb{F}_q^n$. Substituting y into the summation gives:

$$F_{f_z}(w) = q^{-n} \sum_{y \in \mathbb{F}_q^n} f(y) \cdot f(y - w) = F_f(w) \ .$$

Since $F_{f_z}(w) = F_f(w)$ for all w, we have $F_{f_z} = F_f$. □

The following lemma describes how δ_{C,F_f} controls the single-step change in $\delta_{C,f}$.

Lemma 4.12 (Deviation of F_f controls the change in deviation of f)**.** *Let $f : \mathbb{F}_q^n \to \mathbb{R}_{\geq 0}$ be normalized. Fix a linear code $C \subseteq \mathbb{F}_q^n$ and assume $\delta_{C,f} \neq 0$. Let $u \sim \mathbb{F}_q^n$ be a uniformly random vector, and let $C' = C + \mathrm{span}(u)$. Then, for any $\lambda > 0$:*

$$\Pr_u \left[\left| \delta_{C',f} \right| > |\delta_{C,f}| \left(\frac{1}{q} + \lambda \right) \right] \leq \frac{\left(\frac{q-1}{q} \right)^2 \delta_{C,F_f}}{\lambda^2 \cdot \delta_{C,f}^2} \ .$$

Proof. For any normalized function $f : \mathbb{F}_q^n \to \mathbb{R}$ (meaning $\widehat{f}(0) = \|f\|_1 = 1$), we can express its relative deviation over C using the Fourier transform. Utilizing the property that $\sum_{x \in C} \chi_y(x) = |C| \cdot \mathbf{1}_{[y \in C^\perp]}$, we have:

$$\delta_{C,f} = \frac{1}{|C|} \sum_{x \in C} f(x) - 1 = \sum_{y \in C^\perp} \widehat{f}(y) - 1 = \sum_{y \in C^\perp \setminus \{0\}} \widehat{f}(y) \ , \tag{8}$$

where the second equality is due to the Convolution Identity.

For the augmented code $C' = C + \mathrm{span}(u)$, its dual space is $(C')^\perp = \{y \in C^\perp \mid \langle y, u \rangle = 0\}$. Let $\xi_y(u)$ be the indicator random variable $\mathbf{1}_{[\langle y,u \rangle = 0]}$. Thus, for any vector $u \in \mathbb{F}_q^n$, the deviation over the augmented code becomes:

$$\delta_{C',f} = \sum_{y \in C^\perp \setminus \{0\}} \widehat{f}(y) \cdot \xi_y(u) \ .$$

Since u is chosen uniformly at random, for any non-zero $y \in \mathbb{F}_q^n$, the inner product $\langle y, u \rangle$ is uniformly distributed over $\mathbb{F}_q$. Therefore, $\mathbb{E}_u \left[\xi_y(u) \right] = 1/q$. By linearity of expectation,

$$\mathop{\mathbb{E}}_u \left[\delta_{C',f} \right] = \frac{1}{q} \sum_{y \in C^\perp \setminus \{0\}} \widehat{f}(y) = \frac{\delta_{C,f}}{q} \ .$$

Next, we bound the variance of $\delta_{C',f}$. We evaluate

$$\mathrm{Var}_u \left(\delta_{C',f} \right) = \sum_{y,y' \in C^\perp \setminus \{0\}} \widehat{f}(y) \cdot \overline{\widehat{f}(y')} \left(\mathop{\mathbb{E}}_u \left[\xi_y(u) \xi_{y'}(u) \right] - \mathop{\mathbb{E}}_u \left[\xi_y(u) \right] \mathop{\mathbb{E}}_u \left[\xi_{y'}(u) \right] \right) \ .$$

Observe that for linearly independent y and y',

$$\mathop{\mathbb{E}}_u \left[\xi_y(u) \xi_{y'}(u) \right] - \mathop{\mathbb{E}}_u \left[\xi_y(u) \right] \mathop{\mathbb{E}}_u \left[\xi_{y'}(u) \right] = q^{-2} - q^{-2} = 0 \ ,$$

so these terms vanish from the sum. For y and y' colinear, we have

$$\mathop{\mathbb{E}}_u \left[\xi_y(u) \xi_{y'}(u) \right] - \mathop{\mathbb{E}}_u \left[\xi_y(u) \right] \mathop{\mathbb{E}}_u \left[\xi_{y'}(u) \right] = q^{-1} - q^{-2} = \frac{q-1}{q^2} \ .$$

Thus,

$$\begin{aligned}
\mathrm{Var}_u\left(\delta_{C',f}\right) &= \frac{q-1}{q^2}\sum_{y\in C^\perp\setminus\{0\}}\sum_{c\in\mathbb{F}_q^*}\widehat{f}(y)\cdot\overline{\widehat{f}(cy)} \\
&\le \frac{q-1}{q^2}\sum_{y\in C^\perp\setminus\{0\}}\sum_{c\in\mathbb{F}_q^*}\left|\widehat{f}(y)\right|\cdot\left|\widehat{f}(cy)\right| && \text{(triangle inequality)} \\
&\le \frac{q-1}{2q^2}\sum_{y\in C^\perp\setminus\{0\}}\sum_{c\in\mathbb{F}_q^*}\left(\left|\widehat{f}(y)\right|^2+\left|\widehat{f}(cy)\right|^2\right) && \text{(AM-GM inequality)} \\
&= \frac{q-1}{2q^2}\left((q-1)\sum_{y\in C^\perp\setminus\{0\}}\left|\widehat{f}(y)\right|^2+(q-1)\sum_{y'\in C^\perp\setminus\{0\}}\left|\widehat{f}(y')\right|^2\right) && \text{(taking } y'=cy\text{)} \\
&= \frac{q-1}{2q^2}\cdot 2(q-1)\sum_{y\in C^\perp\setminus\{0\}}\left|\widehat{f}(y)\right|^2 \\
&= \left(\frac{q-1}{q}\right)^2\sum_{y\in C^\perp\setminus\{0\}}\left|\widehat{f}(y)\right|^2 \\
&= \left(\frac{q-1}{q}\right)^2\sum_{y\in C^\perp\setminus\{0\}}\widehat{F_f}(y) && \text{(by Lemma 4.9)} \\
&= \left(\frac{q-1}{q}\right)^2\delta_{C,F_f}\ . && \text{(by Eq. (8))}
\end{aligned}$$

Therefore, by Chebyshev's inequality,

$$\begin{aligned}
\Pr_u\left[|\delta_{C',f}|>|\delta_{C,f}|\left(\frac{1}{q}+\lambda\right)\right] &\le \Pr_u\left[\left|\delta_{C',f}-\mathop{\mathbb{E}}_u\left[\delta_{C',f}\right]\right|>\lambda|\delta_{C,f}|\right] \\
&\le \frac{\mathrm{Var}_u\left(\delta_{C',f}\right)}{\lambda^2\delta_{C,f}^2}\le\frac{\left(\frac{q-1}{q}\right)^2\delta_{C,F_f}}{\lambda^2\delta_{C,f}^2}.
\end{aligned}$$

□

4.3 Controlling the Deviation throughout the Construction Sequence

To extend our one-step variance bound across the entire k-step construction, we must carefully track the deviation's trajectory. The following lemma achieves this by categorizing each step as:

- Horrible—steps in which the deviation grows.
- Good—steps in which the deviation goes down significantly, or is already as small as expected to begin with.
- Bad—steps in which the deviation does not increase, but also does not drop significantly even though it is expected to.

The lemma demonstrates that a breakdown in f's smoothness requires an accumulation of too many "bad" steps or at least one "horrible" step. We bound the probability of the first sort of failure, namely, we essentially show that a failure due to t bad steps occurs with probability at most $q^{-\Omega(t^2)}$. Later, in Lemma 4.14, we take care of the other failure condition.

Lemma 4.13. *Let $f : \mathbb{F}_q^n \to \mathbb{R}_{\geq 0}$ be normalized, let $C_0, \ldots, C_k$ be a random construction sequence, let $\Gamma > 0$, and let $t \in \mathbb{N}$ satisfy $20 \log_q(n) \leq t \leq k$. Set*

$$\Lambda := \max\{\|f\|_\infty \cdot q^{-k}, \sqrt{\Gamma}\}.$$

- *Let $S_{F\text{-}smooth}(f)$ denote the event that $C_0, \ldots, C_k$ is (F_f, Γ)-smooth.*
- *Let $S_{monotone}(f)$ denote the event that $|\delta_{C_j,f}| \leq |\delta_{C_{j-1},f}|$ for all j such that $(C_0, \ldots, C_{j-1})$ is $\big(f, q^t \cdot \Lambda, j-1\big)$-smooth.*
- *Let $S_{smooth}(f)$ denote the event that $C_0, \ldots, C_k$ is $\big(f, q^t \cdot \Lambda\big)$-smooth.*

Then

$$\Pr_{C_0,\ldots,C_k}\left[S_{F\text{-}smooth}(f) \text{ and } S_{monotone}(f) \text{ and } \underline{not}\ S_{smooth}(f)\right] \leq q^{-\frac{1}{10}t^2}. \tag{9}$$

Proof of Lemma 4.13. The event $S_{\text{smooth}}(f)$ holds if and only if $|\delta_{C_i,f}| \leq q^{k-i+t} \cdot \Lambda$ for all $0 \leq i \leq k$.

Suppose that $S_{\text{monotone}}(f)$ holds but $S_{\text{smooth}}(f)$ fails. Let $0 \leq \ell \leq k$ be the minimum index for which $C_0, \ldots, C_k$ is not $(f, q^t\Lambda, \ell)$-smooth In particular,

$$q^{k-\ell+t}\Lambda < |\delta_{C_\ell,f}| \ . \tag{10}$$

For a given step $i \in [\ell]$ in the construction sequence, we make the following classification where $\lambda = \frac{t}{4qk}$:

1. The i-th step is horrible if $|\delta_{C_i,f}| > |\delta_{C_{i-1},f}|$;
2. The i-th step is good if either $|\delta_{C_i,f}| \leq |\delta_{C_{i-1},f}|(\frac{1}{q}+\lambda)$ or $|\delta_{C_i,f}| \leq |\delta_{C_{i-1},f}| \leq q^k\Lambda\left(\frac{1}{q}+\lambda\right)^{i-t/2}$;
3. The i-th step is bad if it is not horrible, $|\delta_{C_{i-1},f}| > q^k\Lambda(\frac{1}{q}+\lambda)^{i-t/2}$, and $|\delta_{C_i,f}| > |\delta_{C_{i-1},f}|(\frac{1}{q}+\lambda)$. In other words, the i-th step is bad if it is neither horrible nor good.

First, note that if the event $S_{\text{monotone}}(f)$ holds then no horrible steps may occur up to and including the ℓ-th step. Now, for $0 \leq i \leq \ell$ let g_i denote the number of good steps from among the steps $\{1, \ldots, i\}$. We claim that if no step is horrible then

$$|\delta_{C_i,f}| \leq q^k\Lambda\left(\frac{1}{q}+\lambda\right)^{g_i - t/2} \ . \tag{11}$$

We prove this by induction on i. This is clearly true for $i = 0$, since $\delta_{C_0,f} \leq q^k\Lambda$. Now, suppose Eq. (11) holds for i. We will show that it also holds for $i+1$. If the $(i+1)$-st step is bad then $g_{i+1} = g_i$, and so in this case

$$|\delta_{C_{i+1},f}| \leq |\delta_{C_i,f}| \leq q^k\Lambda\left(\frac{1}{q}+\lambda\right)^{g_i-t/2} = q^k\Lambda\left(\frac{1}{q}+\lambda\right)^{g_{i+1}-t/2} \ .$$

On the other hand, if the $(i+1)$-st step is good then $g_{i+1} = g_i + 1 \leq i+1$. We now consider two cases.

- If this step is good because

$$|\delta_{C_{i+1},f}| \le |\delta_{C_i,f}| \le q^k \Lambda \left(\frac{1}{q}+\lambda\right)^{i+1-t/2} ,$$

then we also get $|\delta_{C_{i+1},f}| \le q^k \Lambda (\frac{1}{q}+\lambda)^{g_{i+1}-t/2}$ since $g_{i+1} \le i+1$.

- If this step is good because $|\delta_{C_{i+1},f}| \le |\delta_{C_i,f}|(\frac{1}{q}+\lambda)$, then

$$|\delta_{C_{i+1},f}| \le q^k \Lambda \left(\frac{1}{q}+\lambda\right)^{g_i-t/2+1} = q^k \Lambda \left(\frac{1}{q}+\lambda\right)^{g_{i+1}-t/2} ,$$

since $g_{i+1} = g_i + 1$.

In both cases we get the desired upper bound on $|\delta_{C_{i+1},f}|$, proving Eq. (11).

We turn to show that the event combination $S_{\text{monotone}}(f)$ and $\neg S_{\text{smooth}}(f)$ implies an upper bound on the number of good steps g_ℓ. Let $c = 1 - \log_q\left(1+\frac{t}{4k}\right)$. Since $q \ge 2$, we strictly have $c > 0$. We can rewrite our step multiplier exactly as $\frac{1}{q}+\lambda = q^{-c}$. By Eq. (11), $|\delta_{C_\ell,f}| \le q^k \Lambda q^{-c(g_\ell - t/2)}$. Therefore:

$$q^{k-\ell+t}\Lambda < |\delta_{C_\ell,f}| \le q^{k-c(g_\ell-t/2)}\Lambda ,$$

where the left-hand inequality is due to Eq. (10). Comparing exponents yields:

$$c(g_\ell - t/2) < \ell - t \implies g_\ell < \frac{\ell - t}{c} + \frac{t}{2} .$$

Let b_ℓ be the number of bad steps up to step ℓ, so $g_\ell = \ell - b_\ell$. Combining these inequalities and isolating b_ℓ:

$$\ell - b_\ell < \frac{\ell - t}{c} + \frac{t}{2} \implies b_\ell > \ell\left(1-\frac{1}{c}\right) + t\left(\frac{1}{c}-\frac{1}{2}\right) .$$

Since $c < 1$, the term $(1-\frac{1}{c})$ is negative. To find the minimum required bad steps τ, we minimize the right-hand side by setting $\ell = k$. We explicitly set:

$$\tau = \max\left\{0, \left\lfloor k\left(1-\frac{1}{c}\right) + t\left(\frac{1}{c}-\frac{1}{2}\right)\right\rfloor\right\} .$$

We have shown that the conjunction of events $S_{\text{monotone}}(f) \wedge \neg S_{\text{smooth}}(f)$ implies that there exists some failure step $\ell \in [\tau, k]$ such that $b_\ell > \tau$, and crucially, the sequence is monotone up to step $\ell - 1$.

For $0 \le i \le k$, let T_i denote the event that $|\delta_{C_j,F_f}| \le q^{k-j} \cdot \Gamma$ for all $1 \le j \le i$. Note that the event $S_{F\text{-smooth}}(f)$ directly implies T_k. For an increasing sequence $1 \le i_1 < i_2 < \cdots < i_\tau \le k$, let $W_{i_1,\ldots,i_\tau}$ denote the event that the steps $i_1, \ldots, i_\tau$ are bad **and the sequence is monotone up to step** $i_\tau - 1$.

Therefore, the left-hand side of Eq. (9) is bounded from above by the union over all such failing sequences:

$$\Pr_{C_0,\ldots,C_k}\left[S_{F\text{-smooth}}(f) \text{ and } \bigcup_{i_1<\cdots<i_\tau} W_{i_1,\ldots,i_\tau}\right] .$$

We bound this by summing over all sequences of τ bad steps:

$$\Pr_{C_0,\ldots,C_k}\left[S_{F\text{-smooth}}(f) \text{ and } \bigcup_{i_1<\cdots<i_\tau} W_{i_1,\ldots,i_\tau}\right] \le \sum_{i_1,\ldots,i_\tau} \Pr[W_{i_1,\ldots,i_\tau} \cap T_k]$$

$$= \sum_{i_1,\ldots,i_\tau} \prod_{r=1}^{\tau} \Pr_{u_{i_r}}[W_{i_1,\ldots,i_r} \cap T_{i_r} \mid W_{i_1,\ldots,i_{r-1}} \cap T_{i_{r-1}}] \text{(conditioning fixes } C_{i_r-1}\text{, leaving only } u_{i_r} \text{ as random)}$$

$$\le \sum_{i_1,\ldots,i_\tau} \prod_{r=1}^{\tau} \Pr_{u_{i_r}}[W_{i_1,\ldots,i_r} \mid W_{i_1,\ldots,i_{r-1}} \cap T_{i_{r-1}}] \ . \tag{12}$$

To bound the probability $\Pr_{u_{i_r}}[W_{i_1,\ldots,i_r} \mid W_{i_1,\ldots,i_{r-1}} \cap T_{i_{r-1}}]$, fix a code C_{i_r-1} for which the conditioning holds. If $\delta_{C_{i_r-1},f} = 0$ then the i_r-th step cannot be bad, so the probability of $W_{i_1,\ldots,i_r}$ under the conditioning is 0. Suppose now that $\delta_{C_{i_r-1},f} \neq 0$. Then, Lemma 4.12 yields

$$\Pr_{u_{i_r}}[\text{step } i_r \text{ is bad} \mid C_{i_r-1}] \le \frac{\left(\frac{q-1}{q}\right)^2 |\delta_{C_{i_r-1},F_f}|}{\lambda^2 \cdot \delta^2_{C_{i_r-1},f}} \ .$$

The event T_{i_r-1} yields an upper bound of $q^{k-(i_r-1)}\Gamma = q^{k-i_r+1}\Gamma$ on the numerator. Since $\Lambda = \max\left\{\|f\|_\infty \cdot q^{-k}, \sqrt{\Gamma}\right\}$ and $\|f\|_\infty \ge |\delta_{C_0,f}|$, we have $\Gamma \le \Lambda^2$, so the numerator is at most $q^{k-i_r+1}\Lambda^2$. For the denominator, we need a lower bound on $|\delta_{C_{i_r-1},f}|$. By definition, the i_r-th step being bad implies that

$$|\delta_{C_{i_r-1},f}| > q^k \Lambda \left(\frac{1}{q} + \lambda\right)^{i_r - t/2} \ .$$

Note that we must have $i_r \ge t/2$. If $i_r < t/2$, the multiplier $(\frac{1}{q} + \lambda)^{i_r - t/2} > 1$, which would require $|\delta_{C_{i_r-1},f}| > q^k\Lambda$. This is impossible because the event $W_{i_1,\ldots,i_r}$ guarantees monotonicity up to $i_r - 1$, ensuring $|\delta_{C_{i_r-1},f}| \le |\delta_{C_0,f}| \le \|f\|_\infty \le q^k\Lambda$.

Since $i_r \ge t/2$ and $1/q + \lambda > 1/q$, we can safely lower-bound the value:

$$|\delta_{C_{i_r-1},f}| > q^k \Lambda \left(\frac{1}{q}\right)^{i_r - t/2} = q^{k - i_r + t/2}\Lambda \ .$$

Squaring this lower bound, the denominator is strictly greater than $\lambda^2 q^{2k-2i_r+t}\Lambda^2$. Therefore,

$$\begin{aligned}
\Pr_{u_{i_r}}[W_{i_1,\ldots,i_r} \mid W_{i_1,\ldots,i_{r-1}} \cap T_{i_{r-1}}] &\le \frac{\left(\frac{q-1}{q}\right)^2 \cdot q^{k-i_r+1}\Lambda^2}{\lambda^2 \cdot q^{2k-2i_r+t}\Lambda^2} \\
&\le \frac{1}{\lambda^2} q^{-k+i_r+1-t} \\
&\le \frac{q}{\lambda^2} q^{-t} \quad \text{(since } i_r \le k\text{)} \ .
\end{aligned}$$

Returning to Eq. (12), we substitute this uniform bound for all τ bad steps. Recalling our exact

substitution $\lambda = \frac{t}{4qk}$:

$$\begin{aligned}\Pr_{C_0,\ldots,C_k}[S_{F\text{-smooth}}(f) \text{ and } S_{\text{monotone}}(f) \text{ and } \underline{\text{not}}\ S_{\text{smooth}}(f)] &\le \sum_{i_1,\ldots,i_\tau} \Pr[W_{i_1,\ldots,i_\tau} \cap T_k] \\ &\le \sum_{i_1,\ldots,i_\tau} \prod_{r=1}^{\tau} \left(\frac{16q^3k^2}{t^2} q^{-t}\right) \\ &\le \binom{k}{\tau} \left(\frac{16q^3k^2}{t^2} q^{-t}\right)^{\tau} .\end{aligned}$$

We bound the binomial coefficient using the standard inequality $\binom{k}{\tau} \le k^\tau$. Substituting this into our probability bound isolates the exponent:

$$\begin{aligned}\Pr_{C_0,\ldots,C_k}[S_{F\text{-smooth}}(f) \text{ and } S_{\text{monotone}}(f) \text{ and } \underline{\text{not}}\ S_{\text{smooth}}(f)] &\le k^\tau \left(\frac{16q^3k^2}{t^2} q^{-t}\right)^{\tau} \\ &= \left(\frac{16q^3k^3}{t^2} q^{-t}\right)^{\tau} \\ &= q^{-\tau\left(t - \log_q\left(\frac{16q^3k^3}{t^2}\right)\right)} .\end{aligned}$$

Now we bound τ. Since $q \ge 2$, we have $\ln q \ge \ln 2 \approx 0.693$. Using the standard bound $\log_q(1+x) \le \frac{x}{\ln q}$, let $L = \log_q\left(1 + \frac{t}{4k}\right)$. We have $L \le \frac{t}{4k \ln q} \le \frac{t}{2.77k}$. For valid fractional parameters where $t \le k$, we have $L \le 0.36 < 1$, ensuring $c = 1 - L > 0$. Substitute $c = 1 - L$ into the required steps condition:

$$\begin{aligned}\tau \ge \left(1 - \frac{1}{1-L}\right) + t\left(\frac{1}{1-L} - \frac{1}{2}\right) &= k\left(\frac{-L}{1-L}\right) + t\left(\frac{1}{1-L} - \frac{1}{2}\right) \\ &= \frac{1}{1-L}\left(-kL + t - \frac{t(1-L)}{2}\right) \\ &\ge \frac{1}{1-L}\left(-\frac{t}{4\ln q} + \frac{t}{2} + \frac{tL}{2}\right) \\ &\ge t\left(\frac{1}{2} - \frac{1}{4\ln 2}\right) \ge 0.13t .\end{aligned}$$

Therefore, $\tau \ge \lfloor 0.13t \rfloor \ge 0.1t$ for sufficiently large t.

Substituting this simplified lower bound into our probability inequality yields the final dependence:

$$\begin{aligned}\Pr_{C_0,\ldots,C_k}[S_{F\text{-smooth}}(f) \text{ and } S_{\text{monotone}}(f) \text{ and } \underline{\text{not}}\ S_{\text{smooth}}(f)] &\le q^{-0.13t\left(t - \log_q\left(\frac{16q^3k^3}{t^2}\right)\right)} \\ &\le q^{-0.1t^2} = q^{-\Omega(t^2)} .\end{aligned}$$

which completes the proof. □

We now show that smoothness of F_f implies smoothness for f with very high probability. In Lemma 4.7 we showed this implication "modulo" the possibility of failure due to horrible steps.

Below we deal with this obstacle: we show that if f fails to be smooth then there exists some translation f_z that fails to be smooth but also avoids horrible steps. The claim then follows from Lemma 4.13 via a union bound on all translations of f.

Lemma 4.14. *In the setting of Lemma 4.13,*

$$\Pr_{C_0,\ldots,C_k}\left[S_{F\text{-}smooth}(f) \text{ and } \underline{not}\ S_{smooth}(f)\right] \le q^n \cdot q^{-0.1t^2} \ .$$

Proof. We first prove the following implication:

$$\text{If } \underline{\text{not}}\ S_{\text{smooth}}(f) \quad \text{then} \quad \exists z \in \mathbb{F}_q^n \ \ S_{\text{monotone}}(f_z) \text{ and } \underline{\text{not}}\ S_{\text{smooth}}(f_z) \ . \tag{13}$$

Assume $S_{\text{smooth}}(f)$ fails at a first time $0 \le j \le k$, meaning $\left|\delta_{C_j,f}\right| > q^{k-j+t}\Lambda$, where $\Lambda = \max\left\{\|f\|_\infty q^{-k}, \sqrt{\Gamma}\right\}$.

We construct a sequence of vectors $z_j, \ldots, z_0$ backward, starting with $z_j = 0$. For $0 \le i < j$, after choosing z_{i+1}, we choose $z_i = z_{i+1} + \gamma_i u_{i+1}$ with $\gamma_i \in \mathbb{F}_q$ maximizing

$$\left|\delta_{C_i, f_{z_{i+1}+\gamma u_{i+1}}}\right| \ .$$

Since $C_{i+1} = C_i + \text{span}\{u_{i+1}\}$, we can express the deviation over the larger code as an average over the q cosets:

$$\delta_{C_{i+1},f_{z_{i+1}}} = \frac{1}{q}\sum_{\gamma\in\mathbb{F}_q} \delta_{C_i, f_{z_{i+1}+\gamma u_{i+1}}} \ .$$

Because the absolute value of an average is bounded by its maximum component, our choice of γ_i implies:

$$\left|\delta_{C_i,f_{z_i}}\right| \ge \left|\delta_{C_{i+1},f_{z_{i+1}}}\right| \ .$$

Next, setting $z = z_0$, we replace the moving translations f_{z_i} with the single fixed translate f_z. For every $0 \le i \le j$, we have $z - z_i \in C_i$ by construction. Hence, translating by $z - z_i$ preserves the subspace C_i, and therefore:

$$\delta_{C_i,f_z} = \delta_{C_i,f_{z_i}} \ .$$

Together with the previous bound, this yields:

$$\left|\delta_{C_i,f_z}\right| \le \left|\delta_{C_{i-1},f_z}\right|$$

for all $1 \le i \le j$. Since $S_{\text{smooth}}(f_z)$ already fails at time j, the definition of $S_{\text{monotone}}(f_z)$ imposes no conditions after the first failure time, so the required monotonicity is fully satisfied.

Furthermore, this fixed translation inherits the original failure at time j. Since $z_j = 0$, it follows that $z - z_j = z \in C_j$. Thus:

$$\left|\delta_{C_j,f_z}\right| = \left|\delta_{C_j,f}\right| > q^{k-j+t}\Lambda \implies \neg S_{\text{smooth}}(f_z) \ ,$$

where we use the fact that $\|f_z\|_\infty = \|f\|_\infty$, ensuring the threshold Λ remains identical.

This establishes Eq. (13). The lemma then follows by applying a union bound:

$$\begin{aligned}
&\Pr_{C_0,\ldots,C_k}\left[S_{F\text{-smooth}}(f) \text{ and } \underline{\text{not}}\ S_{\text{smooth}}(f)\right] \\
&\le \Pr_{C_0,\ldots,C_k}\left[\exists z \in \mathbb{F}_q^n\ \ S_{F\text{-smooth}}(f_z) \text{ and } S_{\text{monotone}}(f_z) \text{ and } \underline{\text{not}}\ S_{\text{smooth}}(f_z)\right] \\
&\le \sum_{z\in\mathbb{F}_q^n} \Pr_{C_0,\ldots,C_k}\left[S_{F\text{-smooth}}(f_z) \text{ and } S_{\text{monotone}}(f_z) \text{ and } \underline{\text{not}}\ S_{\text{smooth}}(f_z)\right] \\
&\le q^n \cdot q^{-0.1t^2}\ .
\end{aligned}$$

The first inequality uses Eq. (13) and the fact that $F_f = F_{f_z}$ implies $S_{F\text{-smooth}}(f) = S_{F\text{-smooth}}(f_z)$. The final bound follows directly from Lemma 4.13. □

4.3.1 Proof of Theorem 2.4

We now prove Theorem 2.4, which controls the deviation for all translations of a single function f. The core insight is that every translation f_z shares the exact same averaged convolution F_f. By securing a global smoothness bound for this single shared F_f, we use Lemma 4.14 to bound the probability to $q^{-\Omega(n^2)}$, then using union bound over all q^n translations.

Proof of Theorem 2.4. Let $c = \frac{\eta}{3}$. Let $\Gamma = q^{-cn}$. Let $S_{F\text{-smooth}}(f)$ be the event that the construction sequence $C_0, \ldots, C_k$ is (F_f, Γ)-smooth. This requires $\left|\delta_{C_i,F_f}\right| \le q^{k-i}\Gamma$ for all $0 \le i \le k$. The lower bound $\delta_{C_i,F_f} \ge -q^{k-i}\Gamma$ holds deterministically since

$$\delta_{C_i,F_f} = \frac{F_f(C)}{|C|} - 1 = \widehat{F_f}(C^\perp) - 1 = \sum_{y\in C^\perp\setminus\{0\}} \widehat{F_f}(y) \ge 0\ .$$

Here, the second transition is Parseval's identity, the third is since F_f is normalized so $\widehat{F_f}(0) = 1$, and the inequality is because Lemma 4.9 implies $\widehat{F_f}(y) \ge 0$ for all y.

Applying Lemma 4.4 to C_i with $\Delta = q^{k-i}\Gamma$ yields:

$$\Pr\left[\delta_{C_i,F_f} > q^{k-i}\Gamma\right] \le \frac{q-1}{q^i\left(q^{k-i}\Gamma - |F_f(0)-1|/q^i\right)^2}\|F_f\|_2^2 + \frac{q^{i-n}}{q-1} \le \frac{4(q-1)q^{\beta n}}{q^{2k-i}\Gamma^2} + \frac{q^{i-n}}{q-1}\ .$$

where the second inequality comes from $q^{k-i}\Gamma = q^{\beta n - i + 2\eta n/3}$ tightly dominating the subtraction, $|F_f(0)-1|/q^i \le q^{\beta n - i}$, and substituting $\|F_f\|_2^2 \le q^{\beta n}$. Then we use union bound for $S_{F\text{-smooth}}(f)$:

$$\begin{aligned}
\Pr\left[\neg S_{F\text{-smooth}}(f)\right] &\le \sum_{i=0}^{k}\left(\frac{4(q-1)q^{\beta n}}{q^{2k-i}\Gamma^2} + \frac{q^{i-n}}{q-1}\right) \\
&\le \frac{4q^{\beta n+1}}{q^k\Gamma^2} + \frac{q^{k-n+1}}{(q-1)^2} \\
&= 4q\cdot q^{-\eta n/3} + \frac{q^{k-n+1}}{(q-1)^2}\ .
\end{aligned}$$

We turn to bound $\Pr\left[S_{\text{smooth}}(f_z)\right]$. Let $t = \theta n$ for $\theta = \eta/12$. By the hypothesis $\eta n \ge 240\log_q n$, we have $t = \eta n/12 \ge 20\log_q n$, and $t = \eta n/12 \le Rn = k$ since $\eta < R$. Hence the condition

$20 \log_q(n) \leq t \leq k$ required by Lemmas 4.13 and 4.14 is satisfied. For each $z \in \mathbb{F}_q^n$, define $S_{\text{smooth}}(f_z)$ as the event that the sequence is (f_z, Γ')-smooth, where

$$\Gamma' = q^t \max\left\{ \|f\|_\infty q^{-k}, \sqrt{\Gamma} \right\} .$$

Because $F_{f_z} = F_f$, the event $S_{F\text{-smooth}}(f_z)$ is equivalent to $S_{F\text{-smooth}}(f)$. By Lemma 4.14,

$$\Pr\left[S_{F\text{-smooth}}(f) \cap \neg S_{\text{smooth}}(f_z)\right] \leq q^{n-0.1t^2} .$$

We can bound the probability that any translation fails to be smooth by conditioning on the global smoothness of F_f:

$$\begin{aligned}
\Pr\left[\exists z \in \mathbb{F}_q^n \;\; \neg S_{\text{smooth}}(f_z)\right] &\leq \Pr\left[\left(\exists z \in \mathbb{F}_q^n \;\; \neg S_{\text{smooth}}(f_z)\right) \cap S_{F\text{-smooth}}(f)\right] + \Pr\left[\neg S_{F\text{-smooth}}(f)\right] \\
&\leq \sum_{z \in \mathbb{F}_q^n} \Pr\left[\neg S_{\text{smooth}}(f_z) \cap S_{F\text{-smooth}}(f)\right] + \Pr\left[\neg S_{F\text{-smooth}}(f)\right] \\
&\leq q^n \cdot q^{n-0.1t^2} + 4q \cdot q^{-\eta n/3} + \frac{q^{k-n+1}}{(q-1)^2} \\
&\leq q^{2n-\frac{\eta^2 n^2}{1440}} + 4q \cdot q^{-\eta n/3} + \frac{q^{k-n+1}}{(q-1)^2} .
\end{aligned}$$

To conclude, calculate Γ'. Since $\|f\|_\infty = q^{\beta n}$, yielding $\|f\|_\infty q^{-k} = q^{\beta n - Rn} = q^{-\eta n}$. We then get $\sqrt{\Gamma} = q^{-\eta n/6}$. Substituting these bounds:

$$\begin{aligned}
\Gamma' &\leq q^{\theta n} \max\left\{ q^{-\eta n}, q^{-\eta n/6} \right\} \\
&= q^{\frac{\eta n}{12}} \cdot q^{-\frac{\eta n}{6}} \\
&= q^{-\frac{\eta n}{12}} .
\end{aligned}$$

By setting $\varepsilon = \eta/12 > 0$, we achieve $\max_{z \in \mathbb{F}_q^n} |\delta_{C_k, f_z}| \leq q^{-\varepsilon n}$, which completes the proof.

□

4.3.2 Proof of Theorem 2.5

We are now ready to prove Theorem 2.5. The proof analyzes a sequence of functions $f^{(0)}, \ldots, f^{(d)}$ generated by repeated averaged convolution, starting with $f^{(0)} = f$. Because each convolution squares the concentration parameter, the top iterate $f^{(d)}$ is highly Fourier-concentrated, allowing its smoothness to be bounded by the code's dual-typicality. We then work backward, inductively applying Lemma 4.14 to transfer this smoothness step-by-step from $f^{(m)}$ down to the base function $f^{(0)}$.

Proof of Theorem 2.5. Let S_{typical} denote the event that $C_0, \ldots, C_k$ is γ-dual-typical where $\gamma = q^{\frac{\eta n}{100}}$ and $k = Rn$. By Lemma 4.6,

$$\Pr\left[S_{\text{typical}}\right] \geq 1 - \frac{n(k+1)}{\gamma} = 1 - n(Rn+1)q^{-\frac{\eta n}{100}} .$$

Choose $d = \max\left\{0, \left\lceil \log_2\left(\frac{\log\left((q-1)/(q^{\eta/100}-1)\right)}{\log(1/\alpha)}\right)\right\rceil\right\}$, so that $1 + (q-1)\alpha^{2^d} \leq q^{\frac{\eta}{100}}$.

Fix $f \in \mathcal{F}$. We claim that

$$\Pr\left[S_{\text{typical}} \text{ and } |\delta_{C,f}| > q^{-\varepsilon n}\right] \leq d \cdot q^{n-0.1t^2} \ . \tag{14}$$

Observe that Eq. (14) implies the theorem since

$$\begin{aligned}
\Pr\left[\exists f \in \mathcal{F} \mid |\delta_{C,f}| > q^{-\varepsilon n}\right] &\leq \Pr\left[\text{not } S_{\text{typical}}\right] + \Pr\left[S_{\text{typical}} \text{ and } \exists f \in \mathcal{F} \quad |\delta_{C,f}| > q^{-\varepsilon n}\right] \\
&\leq \Pr\left[\text{not } S_{\text{typical}}\right] + \sum_{f \in \mathcal{F}} \Pr\left[S_{\text{typical}} \text{ and } |\delta_{C,f}| > q^{-\varepsilon n}\right] \\
&\leq n(Rn+1)q^{-\frac{\eta n}{100}} + |\mathcal{F}| \cdot d \cdot q^{n-0.1t^2}
\end{aligned}$$

We turn to proving Eq. (14). Define the sequence $f^{(0)} := f$, and $f^{(m)} := F_{f^{(m-1)}}$ for $m \geq 1$. Applying Lemma 4.10 we get that $f^{(m)}$ is α^{2^m}-Fourier-concentrated.

Assuming S_{typical}, Lemma 4.7 implies $C_0, \ldots, C_k$ is $(f^{(d)}, \Gamma_d)$-smooth, with:

$$\Gamma_d = q^{-k} \cdot \gamma \cdot \left(1 + (q-1)\alpha^{2^d}\right)^n \leq q^{-k+\frac{\eta n}{100}+\frac{\eta n}{100}} = q^{-k+\frac{\eta n}{50}} \ .$$

Define the sequence $\Gamma_d \ldots \Gamma_0$ where for all $1 \leq m \leq d$:

$$\Gamma_{m-1} = q^t \max\left\{\|f^{(m-1)}\|_\infty q^{-k}, \sqrt{\Gamma_m}\right\}. \tag{15}$$

Define the event $T_{\text{smooth}}(f,i)$ as $C_0 \ldots C_k$ being $(f^{(i)}, \Gamma_i)$-smooth. Let $T_{\text{smooth}}(f)$ be the event that $T_{\text{smooth}}(f,i)$ holds for all $0 \leq i \leq d$. Applying Lemma 4.14 to $f^{(m-1)}$ with $\Gamma = \Gamma_m$ yields

$$\Pr\left[T_{\text{smooth}}(f,m) \cap \neg T_{\text{smooth}}(f,m-1)\right] \leq q^{n-0.1t^2} \ .$$

Therefore,

$$\Pr\left[S_{\text{typical}} \text{ and } |\delta_{C,f}| > \Gamma_0\right] \leq \Pr\left[S_{\text{typical}} \text{ and } \underline{\text{not}}\ T_{\text{smooth}}(f,0)\right] \ .$$

Suppose that the event on the right-hand side holds, and let $0 \leq m \leq d$ be the maximum index for which $T_{\text{smooth}}(f,i)$ *does not* hold. Recall that $m \neq d$ since S_{typical} implies $T_{\text{smooth}}(f,d)$. Thus,

$$\begin{aligned}
\Pr\left[S_{\text{typical}} \text{ and } \underline{\text{not}}\ T_{\text{smooth}}(f,0)\right] &\leq \Pr\left[\exists\ 0 \leq m \leq d-1 \quad T_{\text{smooth}}(f,m+1) \text{ and } \underline{\text{not}}\ T_{\text{smooth}}(f,m)\right] \\
&\leq \sum_{m=0}^{d-1} \Pr\left[T_{\text{smooth}}(f,m+1) \text{ and } \underline{\text{not}}\ T_{\text{smooth}}(f,m)\right] \\
&\leq \sum_{m=1}^{d} q^{n-0.1t^2} \\
&= d \cdot q^{n-0.1t^2} \ .
\end{aligned}$$

Now, we bound the error term Γ_0. By definition, C_k is $(f^{(0)}, \Gamma_0)$-smooth. Setting $i = k$, we obtain:

$$|\delta_{C_k,f}| \leq q^{k-k} \cdot \Gamma_0 = \Gamma_0 \ .$$

Unfolding the recurrent definition Eq. (15) and setting $A_j = \left\|f^{(j)}\right\|_\infty \cdot q^{-k}$:

$$\begin{aligned}
\Gamma_0 &= q^t \max\left\{A_0, \Gamma_1^{1/2}\right\} \\
&= q^t \max\left\{A_0, \left(q^t \max\left\{A_1, \Gamma_2^{1/2}\right\}\right)^{1/2}\right\} \\
&\vdots \\
&\le q^{2t} \cdot \max\left\{\Gamma_d^{2^{-d}}, \max_{0\le j\le d-1}\left(\left\|f^{(j)}\right\|_\infty q^{-k}\right)^{2^{-j}}\right\} . \qquad (16)
\end{aligned}$$

We bound the first term of the maximum on the right-hand side of Eq. (16). Recall that $\Gamma_d \le q^{-Rn+\frac{\eta n}{50}}$. Thus,

$$\begin{aligned}
q^{2t} \cdot \Gamma_d^{2^{-d}} &\le q^{2t} \cdot \left(q^{-Rn+\frac{\eta n}{50}}\right)^{2^{-d}} \\
&= q^{-n\left(R2^{-d} - \frac{\eta 2^{-d}}{50} - \frac{2t}{n}\right)} . \qquad (17)
\end{aligned}$$

We turn to bound the other terms in the right-hand side of Eq. (16). We bound the L_∞ norm of $f^{(j)}$. Recall that $f^{(j)} = F_{f^{(j-1)}} = f^{(j-1)} * \check{f}^{(j-1)}$, and so

$$\begin{aligned}
\left\|f^{(j)}\right\|_\infty &= \left\|f^{(j-1)} * \check{f}^{(j-1)}\right\|_\infty \\
&= \max_w q^{-n} \sum_{x\in\mathbb{F}_q^n} f^{(j-1)}(x)\check{f}^{(j-1)}(w-x) \\
&= \max_w q^{-n} \sum_{x\in\mathbb{F}_q^n} f^{(j-1)}(x) f^{(j-1)}(x-w) \\
&\le \max_w q^{-n} \left\|f^{(j-1)}\right\|_\infty \sum_{x\in\mathbb{F}_q^n} f^{(j-1)}(x) \\
&= q^{-n} \left\|f^{(j-1)}\right\|_\infty \left(q^n \left\|f^{(j-1)}\right\|_1\right) \\
&= \left\|f^{(j-1)}\right\|_\infty .
\end{aligned}$$

Then we get a bound on all $\left\|f^{(j)}\right\|_\infty$:

$$\left\|f^{(j)}\right\|_\infty \le \left\|f^{(j-1)}\right\|_\infty \le \cdots \le \left\|f^{(0)}\right\|_\infty = \|f\|_\infty .$$

Therefore, for all $j \ge 0$, we have $\left\|f^{(j)}\right\|_\infty \le \|f\|_\infty = q^{\beta n}$. Substituting this into our maximum yields:

$$\left\|f^{(j)}\right\|_\infty q^{-k} \le q^{\beta n} q^{-Rn} = q^{(\beta-R)n} = q^{-\eta n} .$$

Substituting this bound back into Eq. (16):

$$\begin{aligned}
q^{2t} \cdot \max_{0\le j\le d-1}\left(\left\|f^{(j)}\right\|_\infty q^{-k}\right)^{2^{-j}} &\le q^{2t} \cdot \max_{0\le j\le d-1}\left(q^{-\eta n}\right)^{2^{-j}} \\
&= q^{2t} \cdot q^{-\eta n 2^{-(d-1)}} \\
&= q^{-n\left(\eta 2^{-d+1} - \frac{2t}{n}\right)} . \qquad (18)
\end{aligned}$$

Set $\varepsilon_1 := R2^{-d} - \frac{\eta 2^{-d}}{50}$ and $\varepsilon_2 := \eta 2^{-d+1}$. Finally, letting $\varepsilon = \min(\varepsilon_1, \varepsilon_2) - \frac{2t}{n}$, and noting by our assumption on t that $\varepsilon > 0$, we have by Eqs. (16) to (18):

$$\begin{aligned}\Gamma_0 &\le \max\left\{q^{-n\left(\varepsilon_1 - \frac{2t}{n}\right)}, q^{-n\left(\varepsilon_2 - \frac{2t}{n}\right)}\right\} \\ &\le q^{-\varepsilon n} \ .\end{aligned}$$

Therefore, assuming S_{typical}, the event $T_{\text{smooth}}(f, 0)$ guarantees $|\delta_{C_k,f}| \le q^{-\varepsilon n}$. Because $T_{\text{smooth}}(f, 0)$ fails with probability at most $d \cdot q^{n-0.1t^2}$, we have established Eq. (14), which completes the proof. □

4.3.3 Proof of Corollary 2.6

Finally, we use Theorem 2.5 to prove the simplified version given in Corollary 2.6.

Proof of Corollary 2.6. Let d be the auxiliary integer from Theorem 2.5, and set $\rho_* := 2^{-d}$. We first show that $\rho_* \ge \lambda$.

Write

$$B := \log(1/\alpha) \quad \text{and} \quad L := \log\left(\frac{q-1}{q^{\eta/100} - 1}\right).$$

We claim that

$$\rho_* \ge \min\left\{1, \frac{B}{2L}\right\}.$$

Indeed, if $L \le B$, then $d = 0$ and $\rho_* = 1$. Otherwise, $d = \lceil \log_2(L/B) \rceil$, so $2^d \le 2L/B$, and hence $\rho_* = 2^{-d} \ge B/(2L)$.

Since

$$q^{\eta/100} - 1 = e^{\eta \log q/100} - 1 \ge \frac{\eta \log q}{100},$$

we have

$$\frac{q-1}{q^{\eta/100} - 1} \le \frac{100(q-1)}{\eta \log q} \le \frac{100q}{\eta}.$$

Thus $L \le \log(100q/\eta)$, and therefore

$$\rho_* \ge \min\left\{1, \frac{\log(1/\alpha)}{2\log(100q/\eta)}\right\} = \lambda.$$

Apply Theorem 2.5 with $t = \theta n$. Since $\eta \le 50\beta/51$, we have $\beta \ge 51\eta/50$, and hence

$$2^{-d}\left(R - \frac{\eta}{50}\right) = \rho_*\left(\beta + \frac{49}{50}\eta\right) \ge 2\eta\rho_*.$$

Therefore,

$$2^{-d}\min\left\{R - \frac{\eta}{50}, 2\eta\right\} = 2\eta\rho_*.$$

Since $\theta < \eta\lambda \le \eta\rho_*$, the condition on t in Theorem 2.5 is satisfied.

The theorem gives $\varepsilon = 2\eta\rho_* - 2\theta$, so for every $f \in \mathcal{F}$,

$$|\delta_{C,f}| \le q^{-2(\eta\rho_* - \theta)n} \le q^{-2(\eta\lambda - \theta)n}.$$

It remains to simplify the failure probability. By Theorem 2.5, this probability is at most

$$n(Rn+1)q^{-\eta n/100} + |\mathcal{F}|dq^{n-\theta^2 n^2/10} \quad .$$

Since $R \leq 1$ and $\eta \leq 50/51$, the first term is at most $q^{-\eta n/200}$ whenever $n \geq c\eta^{-1}\log_q(1/\eta)$, after increasing the universal constant c. For the second term, the definition of d, together with $q^{\eta/100} - 1 \geq \eta \log q/100$ and $\log(1/\alpha) \geq 1 - \alpha$, gives

$$d \leq c' \log(q/(\eta(1-\alpha)))$$

for a universal constant $c' > 0$. Thus the second lower bound on n in the statement implies $\log_q d \leq \theta^2 n^2/20$, again after increasing c, and hence $dq^{n-\theta^2 n^2/10} \leq q^{n-\theta^2 n^2/20}$. Therefore the failure probability is at most

$$q^{-\eta n/200} + |\mathcal{F}|q^{n-\theta^2 n^2/20},$$

as claimed. □

5 α-Concentration of Combinatorial Rectangles

In this section, we show α-Fourier-concentration for normalized indicator functions of combinatorial rectangles, which are the subsets relevant to zero-error list-recovery, or list-recovery from erasures. In full generality, we fix subsets $S_1, \ldots, S_n \subseteq \mathbb{F}_q$, and put $\ell_i := |S_i|$ for all $i \in [n]$. This argument will only apply to the special case that $q = p$, a prime integer. Some such restriction is necessary: for example, if $q = 2^t$, if each S_i is an $\mathbb{F}_2$-subspace of $\mathbb{F}_q$, then it is not hard to see that any such α-Fourier-concentration is not possible.[7]

The main tool which we will use is the following bound on sums over subsets of p-th roots of unity, proved by Benhamouda et al [BDIR21, Lemma 3.11]. Below, $\omega_p = e^{2\pi i/p}$ denotes a primitive p-th root of unity.

Lemma 5.1 (Bound on subset sums)**.** *Let $S \subseteq \mathbb{F}_p$ be a subset of size ℓ. Then*

$$\left|\frac{1}{\ell}\sum_{x\in S}\omega_p^x\right| \leq \frac{\sin(\pi\ell/p)}{\ell\sin(\pi/p)} \quad .$$

In fact, the proof establishes that the sum is maximized if the sum is over "consecutive" p-th roots of unity, i.e., it is of the form $\left|\sum_{j=0}^{\ell-1}\omega_p^{j+i}\right|$ for some $i \in \{0, 1, \ldots, p-1\}$. This is predictable, upon visualizing the p-th roots of unity as lying on the radius one disk in the complex plane: to maximize the magnitude of the sum, all the p-th roots of unity should "as much as possible" point in the same direction. Additionally we remark that the same result would hold for any m-th root of unity and any subset of $\mathbb{Z}_m$, for any positive integer m.

We can now easily prove the requisite α-Fourier-concentration.

[7] One could, for example, weaken the requirement to being that $\mathrm{char}(q) > \max_i \ell_i$; however, we do not pursue that here.

Lemma 5.2 (α-Fourier-concentration of rectangles)**.** *Let p be a prime, $n \in \mathbb{N}$, and let $S_1, \ldots, S_n \subseteq \mathbb{F}_p$ be subsets, of size $\ell_1, \ldots, \ell_n$, respectively. Let $f = \frac{p^n}{\ell_1 \cdots \ell_n} 1_{S_1 \times \cdots \times S_n}$ denote the normalized indicator for the combinatorial rectangle $S_1 \times \cdots \times S_n$. Then, f is α-Fourier-concentrated for*

$$\alpha := \max_{i \in [n]} \left\{ \frac{\sin(\pi \ell_i / p)}{\ell_i \sin(\pi/p)} \right\} .$$

If $2 \le \ell_i \le p$ for all $i \in [n]$, then $\alpha \le \cos(\pi/p) < 1$.

Proof. Let $\xi \in \mathbb{F}_p^n$ be given, and let $T = \mathrm{supp}(\xi) = \{i \in [n] : \xi_i \neq 0\}$. We need to prove

$$|\widehat{f}(\xi)| = \left| \mathop{\mathbb{E}}_{X \sim \mathbb{F}_p^n} \left[f(X)\, \omega_p^{-\langle X, \xi \rangle} \right] \right| \le \alpha^{|T|} .$$

Firstly, as f is normalized we have

$$\mathop{\mathbb{E}}_{X \sim \mathbb{F}_p^n} \left[f(X)\, \omega_p^{-\langle X, \xi \rangle} \right] = \mathop{\mathbb{E}}_{X \sim S_1 \times \cdots \times S_n} \left[\omega_p^{-\langle X, \xi \rangle} \right] . \tag{19}$$

As $S_1 \times \cdots \times S_n$ is a product set, the coordinates $X_1, \ldots, X_n$ of X are independent, so

$$(19) = \prod_{i=1}^{n} \mathop{\mathbb{E}}_{X_i \sim S_i} \left[\omega_p^{X_i \cdot \xi_i} \right] . \tag{20}$$

Now, for each $i \in [n]$ with $\xi_i = 0$ the expectation is clearly 1. Hence,

$$(20) = \prod_{i \in T} \mathop{\mathbb{E}}_{X_i \sim S_i} \left[\omega_p^{X_i \cdot \xi_i} \right] = \prod_{i \in T} \left(\frac{1}{\ell_i} \sum_{x_i \in S_i} \omega_p^{-x_i \cdot \xi_i} \right) = \prod_{i \in T} \left(\frac{1}{\ell_i} \sum_{x_i \in -\xi_i \cdot S_i} \omega_p^{x_i} \right) .$$

Taking absolute values and applying Lemma 5.1 allows us to upper bound $|\widehat{f}(\xi)|$ as

$$|\widehat{f}(\xi)| = \prod_{i \in T} \left| \frac{1}{\ell_i} \sum_{x_i \in -\xi_i \cdot S_i} \omega_p^{x_i} \right| \le \prod_{i \in T} \frac{\sin(\pi \ell_i / p)}{\ell_i \sin(\pi/p)} \le \alpha^{|T|} ,$$

where the last inequality recalls we defined α as $\max_{i \in [n]} \left\{ \frac{\sin(\pi \ell_i / p)}{\ell_i \sin(\pi/p)} \right\}$.

To see the last claim about the $\cos(\pi/p)$ upper bound on α when $2 \le \ell_i \le p$ for all i, it suffices to show that $\frac{\sin(\pi \ell / p)}{\ell \sin(\pi/p)}$ is decreasing in ℓ, since $\frac{\sin(2\pi/p)}{2 \sin(\pi/p)} = \cos(\pi/p)$. To see that $\frac{\sin(\pi \ell/p)}{\ell \sin(\pi/p)}$ is decreasing in ℓ, note that this is equivalent to

$$\frac{\sin(\pi(\ell+1)/p)}{\ell+1} \le \frac{\sin(\pi \ell/p)}{\ell},$$

and that $\sin(x)/x$ is decreasing for $x \in (0, \pi]$. □

Finally, in practice what we need to lower bound is $\log(1/\alpha)$. For $\alpha = \cos \frac{\pi}{p}$, we record below the estimate $\log \frac{1}{\cos \pi/p} \ge \Omega(1/p^2)$.

Lemma 5.3. *There is a constant $c > 0$, such that*

$$\log \frac{1}{\cos(\pi/p)} = \log \sec(\pi/p) \geq \frac{\pi^2}{8p^2 \ln 2}$$

for all $p > 2$.

Proof. For $|z| < \frac{\pi}{2}$, the Taylor series for $\sec z$ about $z = 0$ is given by

$$\sec z = \sum_{j=0}^{\infty} \frac{(-1)^j E_{2j}}{(2j)!} z^{2j}.$$

where the E_{2j} are the Euler numbers [AS48, p. 75, eq. 4.3.68]. Since all these coefficients are positive (the Euler numbers E_{2j} are positive iff j is even [AS48, pp. 804–805]), we in particular have $\sec z \geq 1 + \frac{z^2}{2}$. Hence, since $\frac{\pi}{p} < \frac{\pi}{2}$ by the assumption on p,

$$\log\left(\sec\left(\tfrac{\pi}{p}\right)\right) \geq \log\left(1 + \frac{\pi^2}{4p^2}\right) .$$

Then, since $\ln(1 + y) \geq \frac{y}{2}$ for all $y \in [0, 1]$, as $p > 2$ also guarantees $\frac{\pi^2}{4p^2} \leq 1$, we find

$$\log\left(1 + \frac{\pi^2}{4p^2}\right) \geq \frac{1}{\ln 2} \ln\left(1 + \frac{\pi^2}{4p^2}\right) \geq \frac{1}{\ln 2} \cdot \frac{1}{2} \cdot \frac{\pi^2}{4p^2} ,$$

so we get the claimed lower bound. □

6 Improved Resilience of Linear Ramp Secret Sharing Schemes against Local Leakage Functions

In this section we apply our results to the study of leakage-resilient secret sharing schemes. Roughly speaking, an (n, t)-threshold secret sharing scheme encodes a secret s into a set of n shares $S_1, \ldots, S_n$, each assigned to one of n parties, with the property that every subset of at least t parties can pool their shares and perfectly reconstruct s, while any subset of fewer than t parties learns no information about the secret. It is not hard to see that in any (n, t)-threshold secret sharing scheme the length (in bits) of each share must be at least the length of the secret.

To obtain small share size compared to the number of parties n, we must weaken the requirements above by allowing a gap between the *reconstruction threshold* t_{rec} (with the property that any subset of parties of size at least t_{rec} can reconstruct the secret) and the *privacy threshold* t_{priv} (with the property that any subset of parties of size at most t_{priv} learns no information about the secret) [CCX13]. This corresponds to the notion of *ramp* secret sharing. The notion of (n, t)-threshold secret sharing corresponds to the special case where $t_{\text{rec}} = t$ and $t_{\text{priv}} = t_{\text{rec}} - 1 = t - 1$. We present a formal definition below.

Definition 6.1 (Ramp secret sharing scheme)**.** *An $(n, t_{\text{rec}}, t_{\text{priv}})$*-ramp secret sharing scheme *with secret space $\mathcal{S}$ is a pair of algorithms* (Share, Rec) *with* Share *a randomized algorithm that receives as input a secret $s \in \mathcal{S}$ and outputs a share S_i for the i-th party, $i \in [n]$, and* Rec *a deterministic algorithm that receives a subset $T \subseteq [n]$ of parties and the shares $(S_i)_{i \in T}$ corresponding to parties in T and outputs a guess for the secret satisfying the following properties:*

- ***Privacy:*** *For any subset $T \subseteq [n]$ such that $|T| \leq t_{\text{priv}}$ and any two secrets $s, s' \in \mathcal{S}$, the random variables $\text{Share}_T(s)$ and $\text{Share}_T(s')$ (i.e., the shares of s and s' corresponding to parties in T, respectively) are identically distributed.*

- ***Reconstruction:*** *For any subset $T \subseteq [n]$ such that $|T| \geq t_{\text{rec}}$ and any secret $s \in \mathcal{S}$ we have*

$$\Pr[\text{Rec}(\text{Share}(s)_T, T) = s] = 1,$$

with the probability taken over the random coins of Share.

When $t_{\text{priv}} = t_{\text{rec}} - 1$, we say that $(\text{Share}, \text{Rec})$ *is an* (n, t_{rec})-threshold secret sharing scheme.

Linear secret sharing schemes are especially important objects in cryptography (for example, in secure multiparty computation). Informally, these are secret sharing schemes that operate over some finite field $\mathbb{F}_q$, and such that local linear combinations of shares of two secrets yield a correct sharing of the linear combination of the secrets. More precisely, we have the following definition.

Definition 6.2 (Linear ramp secret sharing scheme)**.** *We say that a ramp secret sharing scheme* $(\text{Share}, \text{Rec})$ *is* linear *if the following holds:*

- *The* Share *procedure maps a secret $s \in \mathbb{F}_q$ to shares $S_i \in \mathbb{F}_q^m$ for some integer m;*

- *For every $s \in \mathbb{F}_q$, the tuple* $\text{Share}(s)$ *is uniformly distributed over an affine subspace of $(\mathbb{F}_q^m)^n$;*

- *If $S = (S_1, \ldots, S_n) = \text{Share}(s)$ and $S' = (S'_1, \ldots, S'_n) = \text{Share}(s')$ for any two secrets $s, s' \in \mathbb{F}_q$, then $\alpha S + \beta S'$ is identically distributed to* $\text{Share}(\alpha s + \beta s')$*, for any $\alpha, \beta \in \mathbb{F}_q$.*

We remark that ramp secret sharing schemes are the best we can hope for in regimes where n is much larger than q, which is our main focus here. For example, we must always have $t_{\text{rec}} - t_{\text{priv}} \geq \frac{n - t_{\text{priv}} + 1}{q}$ [CCX13], regardless of the size of the secret and of linearity. In particular, if q is held constant and $t_{\text{priv}} \leq t_{\text{rec}} \leq cn$ for some constant $c < 1$, then necessarily $t_{\text{rec}} - t_{\text{priv}} = \Omega(n)$.

6.1 Leakage-Resilience of Linear Secret Sharing Schemes

Motivated by side-channel attacks in cryptography, Benhamouda, Degwekar, Ishai, and Rabin [BDIR21] studied a different mode of attack on secret sharing schemes. They considered an attacker that learns some side information about *every* share (as opposed to an attacker that learns full information about a small subset of shares). The goal is then to ensure that the secret remains hidden from such an adversary, in which case we call the secret sharing scheme *leakage-resilient*. Studying this requires specifying the leakage functions that the adversary can apply to the shares. We first provide a general definition of leakage-resilience against an abstract family of leakages, and then discuss the important special case on which we focus (bounded local leakage).

Definition 6.3 (Statistical distance)**.** *The* statistical distance *between two discrete random variables X and Y supported on a finite set $\mathcal{S}$, denoted by* $\text{SD}(X, Y)$*, is defined as*

$$\text{SD}(X, Y) = \max_{A \subseteq \mathcal{S}} |\Pr[X \in A] - \Pr[Y \in A]| = \frac{1}{2} \sum_{s \in \mathcal{S}} |\Pr[X = s] - \Pr[Y = s]|.$$

We may write $X \approx_\varepsilon Y$, and say that X and Y are ε-close*, when* $\text{SD}(X, Y) \leq \varepsilon$.

Definition 6.4 (Leakage-resilient ramp secret sharing scheme)**.** *An* $(n, t_{\text{rec}}, t_{\text{priv}}, \mathcal{G}, \varepsilon)$-leakage-resilient ramp secret sharing scheme *is an* $(n, t_{\text{rec}}, t_{\text{priv}})$*-ramp secret sharing scheme with the following additional property:*

- ***Leakage-resilience against*** $\mathcal{G}$***:*** *For any leakage function* $g \in \mathcal{G}$ *and any two secrets* $s, s' \in \mathcal{S}$ *with corresponding shares* $S = (S_1, \ldots, S_n) = \text{Share}(s)$ *and* $S' = (S'_1, \ldots, S'_n) = \text{Share}(s')$*, we have*

$$g(S) \approx_\varepsilon g(S').$$

When the remaining parameters are not relevant, we may also say that $(\text{Share}, \text{Rec})$ *is* $(\mathcal{G}, \varepsilon)$-leakage-resilient*.*

We will be interested in linear ramp secret sharing schemes with $m = 1$, i.e., whose secrets and shares are elements of $\mathbb{F}_q$. The most well-studied leakage model for secret sharing is *bounded local leakage*, which corresponds to independently applying a leakage function with bounded output length to each share. For simplicity, and because this already yields a highly non-trivial problem, we focus on the special case where leakage functions have 1-bit output. More precisely, this corresponds to the family $\mathcal{G}$ containing all functions $g = (g_1, \ldots, g_n)$ of the form

$$g(S) = (g_1(S_1), \ldots, g_n(S_n)),$$

for possibly distinct bounded-output functions $g_i : \mathbb{F}_q \to \{0, 1\}$. Secret sharing schemes that are $(\mathcal{G}, \varepsilon)$-leakage-resilient for this family $\mathcal{G}$ are commonly called *ε-locally leakage-resilient*, and we use this name too. If we wish to emphasize the other parameters, we will write "$(n, t_{\text{rec}}, t_{\text{priv}}, \varepsilon)$-locally leakage-resilient".

6.2 Prior Work on Leakage-Resilience of Linear Secret Sharing Schemes and Our Result

The study of the leakage-resilience of linear threshold secret sharing schemes has attracted significant attention ever since its introduction in [BDIR21]. Establishing the local leakage-resilience of a secret sharing scheme becomes easier when the reconstruction threshold grows. Given this, one basic question in this area is to pinpoint the smallest reconstruction threshold (as a function of the number of parties n) that allows for resilience against bounded local leakage.

General bounded local leakage. The state-of-the-art result for locally leakage-resilient linear threshold secret sharing schemes is due to Sun and Wootters [SW26a] and Nguyen [Ngu26], who showed that Shamir secret sharing is locally leakage-resilient for any threshold $t \geq (1/2 - \delta)n$, with $\delta > 0$ a small constant. The previous best result was due to Maji, Paskin-Cherniavsky, Suad, and Wang [MPSW21], who showed the existence of ε-locally leakage-resilient linear (n, t)-threshold secret sharing schemes with threshold $t = (\frac{1}{2} + \varepsilon)n$ and exponentially small ε, for any $\varepsilon > 0$, over an exponentially large field of prime order $q = 2^{\Theta(n)}$.[8] This is achieved by considering secret sharing schemes induced by random linear codes, as we discuss in more detail below. It has been conjectured that there exist locally leakage-resilient linear threshold secret sharing schemes with threshold $t = cn$ for any constant $c > 0$, and that this holds even for Shamir's secret sharing [BDIR21].

[8] It is easy to see that local leakage-resilience cannot be achieved over fields of characteristic 2 [BDIR21].

Motivated by the prospect of constructing leakage-resilient secret sharing schemes with small share size (for which a gap between the privacy and reconstruction thresholds is unavoidable, as discussed above), Tjuawinata and Xing [TX22] studied the leakage-resilience of linear ramp secret sharing schemes over constant-sized fields. They use algebraic-geometric codes to obtain explicit constructions of such schemes. Looking ahead, in the regime of our results where the field size q is bigger than some sufficiently large constant, their main results [TX22, Theorems 1 and 2 and Corollaries VI.2 and VI.4] only give non-trivial leakage-resilience error ε for reconstruction thresholds $t_{\mathrm{rec}} > n/2$.

Summarizing, the current state of affairs is that the existence of leakage-resilient linear threshold and ramp secret sharing schemes with reconstruction threshold $t_{\mathrm{rec}} < 0.49n$ (and a non-trivial gap between privacy and reconstruction thresholds in the case of ramp secret sharing) remains open. Our result, which we state below in simplified form for constant field size, makes progress on this question whenever the field size does not grow too fast, in which case we must necessarily focus on ramp secret sharing.

Theorem 6.5 (Simplified version of Theorem 6.14)**.** *Fix an arbitrary constant $R \in (0,1)$. Given $n \in \mathbb{N}$, let q be a prime satisfying*

$$2^{800/R} \leq q \leq n^{0.19}.$$

Then, for all sufficiently large n (depending only on R) there is an $(n, t_{\mathrm{rec}}, t_{\mathrm{priv}}, \varepsilon)$-locally leakage-resilient ramp secret sharing scheme over $\mathbb{F}_q$ with

$$Rn - \frac{1.01n}{\log q} \leq t_{\mathrm{priv}} < t_{\mathrm{rec}} \leq Rn + \frac{1.01n}{\log q}$$

and

$$\varepsilon \leq 4q^{-\frac{R^2 n^{3/5}}{3200}}.$$

Note that, in particular, the gap between reconstruction and privacy thresholds can be made $O(n/\log n)$ when $q \geq n^{\gamma}$ for some constant $\gamma \in (0, 0.19)$. The next sections are devoted to proving a more general version of Theorem 6.5. We begin by setting up the standard underlying framework for building linear ramp secret sharing schemes out of linear codes, which matches that used by [MPSW21].

6.3 Linear Ramp Secret Sharing from Linear Codes

Towards proving Theorem 6.5, we will study a classical construction of linear secret sharing schemes from error-correcting codes due to Massey [Mas95], which we proceed to describe. Let $C^+ \subseteq \mathbb{F}_q^{n+1}$ be a linear code with corresponding generator matrix $G^+ \in \mathbb{F}_q^{(n+1)\times(k+1)}$ in systematic form, which we assume to be of the form

$$G^+ = \begin{bmatrix} 1 & 0 \\ v & G \end{bmatrix},$$

where $v \in \mathbb{F}_q^n$ and $G \in \mathbb{F}_q^{n\times k}$. Then, the *Massey secret sharing scheme* associated with C^+ has secret space $\mathbb{F}_q$ and works as follows:

- To share a secret $s \in \mathbb{F}_q$, choose $x \in \mathbb{F}_q^{k+1}$ by setting $x_0 = s$ and sampling $x_1, \ldots, x_k \sim \mathbb{F}_q$. Then, compute $c^+ = (c_0^+, c_1^+, \ldots, c_n^+) = G^+ x$, and set the i-th share to be c_i^+ for $i \in [n]$. Note that $c_0^+ = s$.

- Given a subset of parties $T \subseteq [n]$ and the corresponding shares $(S_i)_{i \in T}$, find the unique codeword $c^+ \in C^+$ such that $c^+_i = S_i$ for all $i \in T$ and output $s = c^+_0$.

It is not hard to see that the minimum distance and dual distance of C^+ respectively control the reconstruction and privacy thresholds of the resulting ramp secret sharing scheme, as formalized in the next standard lemma. For completeness we provide a short proof in Appendix A.2.

Lemma 6.6. *Let $C^+ \subseteq \mathbb{F}_q^{n+1}$ be a linear code with distance d and dual distance $d^\perp$. Then, the Massey secret sharing scheme associated with C^+ is an $(n, t_{\mathrm{rec}} = n - d + 2, t_{\mathrm{priv}} = d^\perp - 2)$-ramp secret sharing scheme.*

Remark 6.7. *If C^+ is an MDS code of dimension k, then Lemma 6.6 yields a linear (n,k)-threshold secret sharing scheme. First, since $d = (n+1) - k + 1 = n - k + 2$, we get that $t_{\mathrm{rec}} = n - d + 2 = k$. Second, since the dual of C^+ is also MDS with dimension $k^\perp = n + 1 - k$, we get that $d^\perp = (n+1) - k^\perp + 1 = k + 1$, and so $t_{\mathrm{priv}} = d^\perp - 2 = k - 1 = t_{\mathrm{rec}} - 1$.*

If the field size q is sufficiently large compared to the block length n, then a random linear code will be MDS with high probability. This fact was leveraged by [MPSW21] to obtain locally leakage-resilient linear threshold secret sharing schemes from random linear codes over exponentially large fields.

To study the local leakage-resilience of a Massey secret sharing scheme based on a linear code C^+ the following alternative perspective is useful. Let $G \in \mathbb{F}_q^{n \times k}$ be the matrix obtained by removing the first column and first row of G^+. Then, G is the generator matrix of a code $C \subseteq \mathbb{F}_q^n$, which we call the *reduced code* of C^+. Denote the first column of G^+ by $v^+ = (1, v_1, \ldots, v_n)^\top$, and let $v = (v_1, \ldots, v_n)^\top$. Sharing a secret $s \in \mathbb{F}_q$ using the Massey secret sharing scheme associated with C^+ is equivalent to the following: sample $x \sim \mathbb{F}_q^k$, compute $c = (c_1, \ldots, c_n)^\top = s \cdot v + Gx$, and set the i-th share to be c_i. In other words, the shares of s are generated by sampling a vector uniformly at random from the affine subspace $s \cdot v + C$. This alternative perspective helps us see that the leakage-resilience of this secret sharing scheme is controlled by the discrepancy of the reduced code C, as formalized in the next lemma which already appears, at least implicitly, in prior work [BDIR21]. For completeness we provide a proof in Appendix A.2.

Lemma 6.8. *Let $(\mathrm{Share}, \mathrm{Rec})$ be the Massey ramp secret sharing scheme associated to a linear code $C^+ \subseteq \mathbb{F}_q^{n+1}$ with reduced code $C \subseteq \mathbb{F}_q^n$. Let $\mathcal{G}$ be a collection of leakage functions closed under input-translations (i.e., if $g \in \mathcal{G}$, then so is the function $x \mapsto g(x+z)$ for all $z \in \mathbb{F}_q^n$).*[9] *Suppose that*

$$\mathrm{SD}(g(C), g(U_n)) \leq \varepsilon$$

for all tuples of leakage functions $g \in \mathcal{G}$. Then, $(\mathrm{Share}, \mathrm{Rec})$ is $(2\varepsilon, \mathcal{G})$-leakage-resilient.

6.4 Leakage Resilience of Massey Ramp Secret Sharing Schemes from Random Linear Codes

We now combine the standard framework described in Section 6.3 with our new results to prove Theorem 6.5. In fact, we will prove the stronger statement that a random linear code yields a linear

[9] Note that this is satisfied by the family of local leakage functions.

ramp secret sharing scheme with the desired properties with high probability, and we will obtain a more general relationship between the various parameters.

We will make use of the concept of an *induced* rectangle, obtained as follows. Given a local leakage function $g = (g_1, \ldots, g_n) : \mathbb{F}_q^n \to \{0,1\}^n$, we say a rectangle $S_1 \times \cdots \times S_n$ is *induced by* g if it is obtained as the preimage of some leakage bits. In other words, for some $b_1, \ldots, b_n \in \{0,1\}$, we have $S_i = g^{-1}(b_i)$ for all $i \in [n]$. Note that any local leakage function partitions $\mathbb{F}_q^n$ into 2^n induced rectangles.

6.4.1 Leakage-Resilience of Low-Rate Random Linear Codes

We begin by proving our key lemma, which applies Theorem 2.5 to the case of "large" rectangles. Formally, for a parameter $\gamma > 0$ which is a bit smaller than the code's rate (and may be chosen arbitrarily small, at the cost of increasing q), we will first consider rectangles $S = S_1 \times \cdots \times S_n$ that are γ-large, defined as follows.

Definition 6.9 (γ-large and γ-small combinatorial rectangles)**.** *Let $S = S_1 \times \cdots \times S_n \subseteq \mathbb{F}_q^n$ be a combinatorial rectangle. Given $\gamma > 0$, we say that it is* γ-large *if $|S| \geq q^{(1-\gamma)n}$. If it is* not *γ-large, we call it* γ-small*.*

Before we present the formal argument for leakage-resilience, we give a brief overview. By applying Corollary 2.6, we can show that a random codeword lands in a large rectangle with approximately the correct probability (this is controlled by $|\delta_{f,C}|$, where f is the normalized indicator function for the rectangle). To derive our main result proving leakage-resilience against all local leakage functions, we will observe that the uniform distribution already puts so little probability mass on γ-small rectangles that their contribution to the statistical distance is negligible. We will then couple this result with Lemma 6.8 to show the existence of Massey ramp secret sharing schemes with reconstruction and privacy thresholds well below the 1/2-barrier appearing in prior works.

One nuisance is that a rectangle induced by the local leakage function $g = (g_1, \ldots, g_n)$ could have some sides of length 1. This causes an issue as such rectangles are not α-concentrated for $\alpha < 1$, preventing an immediate blackbox application of Theorem 2.5. To account for this, we first note that γ-large rectangles cannot have more than γn sides of length 1. Assuming that the dual distance of C is greater than γn (which is very likely for a random linear code of rate slightly more than γ), we then find that the code shortened to the coordinates with side-lengths > 1 is still distributed as a random linear code, and therefore we may apply Theorem 2.5 to this shortened code. Furthermore the "δ-values" in the shortened code exactly equal the δ-values from the original code, so if the δ-values in the shortened code are controlled, the same is true for the original code.

We now proceed to the formal proof. The following corollary of Theorem 2.5, which conveniently does not place any restriction on the relationship between β and η, will be useful. We defer the proof to Appendix A.2.

Corollary 6.10. *Consider a family $\mathcal{F}$ of normalized indicator functions of combinatorial rectangles whose sides all have length greater than 1. Set $\beta = \max_{f \in \mathcal{F}} \frac{1}{n} \cdot \log_q \|f\|_\infty$. Fix a prime field $\mathbb{F}_q$ of size $3 \leq q \leq \frac{n^\lambda}{\log n}$ for some constant $\lambda \in (0, 1/5)$. Let $C \subseteq \mathbb{F}_q^n$ be a random linear code of rate $R = \beta + \eta \leq 1$ for some constant $\eta > 0$. Then, for sufficiently large n, with probability at least*

$$1 - q^{-\frac{\eta n}{200}} - |\mathcal{F}| q^{-9n^{1+\lambda}}$$

over the choice of C, we have

$$|\delta_{C,f}| \leq q^{-\frac{\eta^2 \cdot n^{1-2\lambda}}{400}}$$

for every $f \in \mathcal{F}$.

We now show how to obtain a bound for *all* γ-large combinatorial rectangles, even those that might have sides of length 1 (and therefore technically fail to be α-concentrated with $\alpha < 1$).

Lemma 6.11. *Fix any $0 < \eta, \gamma < 1$ with $\gamma + \eta < 1$, and fix $\lambda \in (0, 1/5)$. Given $n \in \mathbb{N}$, let q be a prime satisfying*

$$\max\left(2^{\frac{1}{(1-\gamma)^{1+\lambda}}}, 2^{400/\eta}\right) \leq q \leq \frac{(n(1-\gamma))^{\lambda}}{\log(n(1-\gamma))}.$$

Then, the following holds for all sufficiently large n. Let $C \subseteq \mathbb{F}_q^n$ be a random linear code of rate $R = \gamma + \eta$. Then, with probability at least

$$1 - 2^n q^{-\frac{\eta n}{200}} - q^{-8(1-\gamma)^{1+\lambda} n^{1+\lambda}} - q^{-\frac{n\eta}{2}} \geq 1 - 2q^{-\frac{\eta n}{400}} \tag{21}$$

over the choice of C, for all γ-large combinatorial rectangles $S = S_1 \times \cdots \times S_n$,

$$|\delta_{C,S}| \leq q^{-\frac{\eta^2 n^{1-2\lambda}}{400}} \quad .$$

Proof. Let $S = S_1 \times \cdots \times S_n$ be a given γ-large combinatorial rectangle, i.e., it satisfies $|S| \geq q^{(1-\gamma)n}$. Let $T \subseteq [n]$ denote the subset of coordinates i for which $|S_i| > 1$. Observe that $|T| \geq (1-\gamma)n$. Then, letting $t = |\overline{T}|$, we get $t \leq \gamma n$.

Let us now consider the code C shortened to the coordinates in T. That is, define

$$C_T := \{c_T : c \in C, c_{\overline{T}} = 0\} \subseteq \mathbb{F}_q^{|T|} \ ,$$

where the notation $c_T = (c_i)_{i \in T}$ denotes the subvector obtained by just keeping the coordinates in T (and similarly for $c_{\overline{T}}$).

First, we recall the structure of its generator matrix. Let $G \in \mathbb{F}_q^{n \times k}$ be the generator matrix sampled for C (where $k = Rn = (\gamma + \eta)n$), where we recall that for our probability model that means G is uniformly random. Let G_T (respectively, $G_{\overline{T}}$) denote the matrix consisting of just the rows in T (respectively, $G_{\overline{T}}$). Note that G_T and $G_{\overline{T}}$ are uniform and independent.

To obtain a generator matrix C_T, one can do the following. Firstly, let $K_{\overline{T}} := \ker(G_{\overline{T}}) = \{x \in \mathbb{F}_q^k : G_{\overline{T}} x = 0\}$, let $a = \dim(K_{\overline{T}})$, and let $A \in \mathbb{F}_q^{a \times k}$ be a matrix whose columns form a basis for $K_{\overline{T}}$ (in other words, it maps $\mathbb{F}_q^a$ injectively into $K_{\overline{T}}$). Then, we find that $G_T A \in \mathbb{F}_q^{a \times n-t}$ is a generator matrix for C_T. Furthermore, conditioned on the even that $G_{\overline{T}}$ has rank $|\overline{T}| = t \leq \gamma n$, we have $a = k - t$ and the matrix $G_T A$ is distributed uniformly at random over $\mathbb{F}_q^{(k-t)\times(n-t)}$. That is, C_T is a random linear code according to our probability model.

Let now $\mathcal{F}_T$ consist of all normalized indicator functions $f_S = \frac{q^n}{|S|} 1_S$ of γ-large combinatorial rectangles $S = S_1 \times \cdots \times S_n$ with $\{i \in [n] : |S_i| > 1\} = T$. Observe $\sum_{T:|T| \geq (1-\gamma)n} |\mathcal{F}_T| \leq 2^{qn} = q^{nq/\log_2 q}$. Fix some $f \in \mathcal{F}_t$, and let S denote the combinatorial rectangle it corresponds to. Let a_i for $i \in \overline{T}$ be such that $S_i = \{a_i\}$.

Let us first consider a (fixed) linear code $D \subseteq \mathbb{F}_q^n$ with the property that $\overline{T}$ does not support a nonzero dual codeword, i.e., for all $x \in D^\perp \setminus \{0\}$ it holds that $\mathrm{supp}(x) := \{i \in [n] : x_i \neq 0\} \not\subseteq \overline{T}$. Then,

$$\Pr_{c\sim D}\left[\forall i \in \overline{T}, c_i = a_i\right] = \frac{|\{c \in D : \forall i \in T, c_i = a_i\}|}{|D|} = q^{-|\overline{T}|} = q^{-t} = \Pr_{x\sim\mathbb{F}_q^n}\left[\forall i \in \overline{T}, x_i = a_i\right] \ .$$

Now, fix any codeword $c^* \in D$ with $c^*_i = a_i$ for all $i \in \overline{T}$, and let $S^* = S^*_1 \times \cdots \times S^*_n$ be defined as $S^*_i := S_i - c^*_i = \{a_i - c^*_i : a_i \in S_i\}$. Then, since $S^*_i = \{0\}$ for all $i \in \overline{T}$, observe that

$$\Pr_{c\sim D}\left[\forall i \in T, c_i \in S_i \mid \forall i \in \overline{T}, c_i = a_i\right] = \Pr_{c\sim D_T}\left[c \in S^*_T\right] \ ,$$

where $S^*_T = \prod_{i\in T} S^*_i$. Thus, if $f^*_T = \frac{q^{n-t}}{|S^*_T|}1_{S^*_T}$ denotes the normalized indicator function for S^*_T (viewed as a subset of $\mathbb{F}_q^{|T|} = \mathbb{F}_q^{n-t}$), we find

$$\delta_{D,f} = \frac{\Pr_{c\sim D}\left[c \in S\right]}{\Pr_{x\sim\mathbb{F}_q^n}\left[x \in S\right]} - 1 = \frac{\Pr_{c\sim D}\left[c \in S\right]}{\Pr_{x\sim\mathbb{F}_q^n}\left[x \in S\right]} - 1 = \frac{\Pr_{c\sim D_T}\left[c \in S^*_T\right]}{\Pr_{x\in\mathbb{F}_q^{n-t}}\left[x \in S^*_T\right]} - 1 = \delta_{D_T,f^*_T} \ .$$

Consequently, if $\mathcal{E}_T$ denotes the event that our random linear code C does not support a dual codeword on $\overline{T}$, we have

$$\Pr_C\left[\delta_{C,f} = \delta_{C_T,f^*_T} \;\middle|\; \mathcal{E}_{\overline{T}}\right] = 1 \ .$$

Let now $\mathcal{F}^*_T$ consist of all functions of the form f^*_T obtained from an $f \in \mathcal{F}_T$ as above. We look to apply Corollary 2.6 to C_T with this $\mathcal{F}^*_T$. Let $\tau := t/n$. Conditioned on the event $\mathcal{E}_{\overline{T}}$, it follows that C_T is a random linear code of rate

$$R_T := \frac{k-t}{n-t} = \frac{R-\tau}{1-\tau} = \frac{\gamma+\eta-\tau}{1-\tau} = \frac{\gamma-\tau}{1-\tau} + \frac{\eta}{1-\tau} \ .$$

Additionally, given any $f^*_T \in \mathcal{F}^*_T$, supposing it is induced by a combinatorial rectangle $S^*_T = \prod_{i\in T} S^*_i$, since we have $|S^*_T| \geq q^{(1-\gamma)n} = q^{\frac{1-\gamma}{1-\tau}(n-t)}$, it follows that

$$\|f^*_T\|_\infty = \frac{|\mathbb{F}_q^{n-t}|}{|S^*_T|} \leq \frac{q^{n-t}}{q^{\frac{1-\gamma}{1-\tau}(n-t)}}$$

with equality for some $f^*_T \in \mathcal{F}^*_T$, and so

$$\beta_T = \max_{f^*_T\in\mathcal{F}^*_T} \frac{1}{n-t}\log_q \|f^*_T\|_\infty = \frac{1}{n-t}\log_q \frac{q^{n-t}}{q^{\left(\frac{1-\gamma}{1-\tau}\right)(n-t)}} = 1 - \frac{1-\gamma}{1-\tau} = \frac{\gamma-\tau}{1-\tau} \ .$$

Thus, we see that C_T has (design) rate $\beta_T + \frac{\eta}{1-\tau}$. Put $\eta_T := \frac{\eta}{1-\tau}$. Then, applying Corollary 6.10, if $C^*_T \subseteq \mathbb{F}_q^{n-t}$ is a random linear code of rate R_T, we get that

$$\Pr_{C^*_T}\left[\exists f^*_T \in \mathcal{F}^*_T \text{ s.t. } |\delta_{C^*_T,f^*_T}| \geq q^{-\frac{\eta_T^2(n-t)^{1-2\lambda}}{400}}\right] \leq q^{-\frac{\eta_T(n-t)}{200}} + |\mathcal{F}_T| q^{-9(n-t)^{1+\lambda}} \ . \tag{22}$$

We use this to obtain a bound on the probability that some $\delta_{C,f}$ is large for some $f \in \mathcal{F}$. First, using that

$$\eta_T^2 (n-t)^{1-2\lambda} = \frac{\eta^2}{(1-\tau)^2}(1-\tau)^{1-2\lambda} n^{1-2\lambda} \geq \eta^2 n^{1-2\lambda} \ ,$$

we obtain the uniform bound

$$q^{-\frac{\eta_T^2 (n-t)^{1-2\lambda}}{400}} \leq q^{-\frac{\eta^2 n^{1-2\lambda}}{400}} =: \Delta \ . \tag{23}$$

Now, recalling that conditioned on the event $\mathcal{E}_{\overline{T}}$ (which we recall is that the coordinates in $\overline{T}$ do not support a nonzero dual codeword) we have $\delta_{C,f} = \delta_{C_T, f_T}$, we get that

$$\begin{aligned}
\Pr_C \left[\exists f \in \mathcal{F} \text{ s.t. } |\delta_{C,f}| \geq \Delta\right] &\leq \sum_{\substack{T \subseteq [n] \\ |T| \geq (1-\gamma)n}} \Pr_C \left[\exists f \in \mathcal{F}_T \text{ s.t. } |\delta_{C,f}| \geq \Delta\right] \\
&\leq \sum_{\substack{T \subseteq [n] \\ |T| \geq (1-\gamma)n}} \Pr_C \left[\exists f \in \mathcal{F}_T \text{ s.t. } |\delta_{C,f}| \geq \Delta \mid \mathcal{E}_{\overline{T}}\right] + \sum_{\substack{T \subseteq [n] \\ |T| \geq (1-\gamma)n}} \Pr_C \left[\neg \mathcal{E}_{\overline{T}}\right] .
\end{aligned}$$

We proceed to bound both of the above sums. To bound the leftmost sum, observe that $q^{-\frac{\eta_T (n-t)}{200}} = q^{-\frac{\eta n}{200}}$, as $\eta_T = \frac{\eta}{1-\tau}$ and $n - t = (1-\tau)n$. Therefore, we have that

$$\begin{aligned}
\sum_{\substack{T \subseteq [n] \\ |T| \geq (1-\gamma)n}} \Pr_C \left[\exists f \in \mathcal{F}_T \text{ s.t. } |\delta_{C,f}| \geq \Delta \mid \mathcal{E}_{\overline{T}}\right] &= \sum_{\substack{T \subseteq [n] \\ |T| \geq (1-\gamma)n}} \Pr_{C_T^*} \left[\exists f_T^* \in \mathcal{F}_T^* \text{ s.t. } |\delta_{C_T^*, f_T^*}| \geq \Delta\right] \\
&\leq \sum_{\substack{T \subseteq [n] \\ |T| \geq (1-\gamma)n}} q^{-\frac{\eta n}{200}} + |\mathcal{F}_T| q^{-9(1-\gamma)^{1+\lambda} n^{1+\lambda}} \\
&\leq 2^n q^{-\frac{\eta n}{200}} + |\mathcal{F}| q^{-9(1-\gamma)^{1+\lambda} n^{1+\lambda}} \\
&\leq 2^n q^{-\frac{\eta n}{200}} + q^{-8(1-\gamma)^{1+\lambda} n^{1+\lambda}} \ .
\end{aligned} \tag{24}$$

The first inequality follows from Eqs. (22) and (23). The second inequality uses the fact that there are at most 2^n subsets $T \subseteq [n]$ of size $|T| \geq (1-\gamma)n$. The last line uses

$$|\mathcal{F}| \leq 2^{qn} \leq 2^{n^{1+\lambda}} \leq q^{(1-\gamma)^{1+\lambda} n^{1+\lambda}}, \tag{25}$$

and the final inequality in Eq. (25) uses the assumption $q \geq 2^{\frac{1}{(1-\gamma)^{1+\lambda}}}$.

As for the rightmost sum, we note that $x \in C^\perp$ if and only if $G^\top x = 0$, where $G \in \mathbb{F}_q^{n \times Rn}$ is

the uniformly random matrix generating C. Hence, we have

$$\begin{aligned}
\sum_{\substack{T\subseteq[n]\\|T|\geq(1-\gamma)n}} \Pr_C\left[\neg\mathcal{E}_T\right] &= \sum_{\substack{T\subseteq[n]\\|T|\geq(1-\gamma)n}} \Pr_C\left[\exists x\in C^\perp\setminus\{0\} \text{ s.t. } \mathrm{supp}(x)\subseteq\overline{T}\right] \\
&\leq \sum_{\substack{T\subseteq[n]\\|T|\geq(1-\gamma)n}} \sum_{\substack{x\in\mathbb{F}_q^n\setminus\{0\}\\ \mathrm{supp}(x)\subseteq\overline{T}}} \Pr_G\left[G^\top x=0\right] \\
&= \sum_{U\subseteq[n]:|U|\leq\gamma n} \left(q^{|U|}-1\right) q^{-(\gamma+\eta)n} \\
&\leq q^{-(\gamma+\eta)n} \sum_{u=1}^{\lfloor\gamma n\rfloor} \binom{n}{u} q^u \\
&\leq q^{-(\gamma+\eta)n}\cdot 2^n\cdot q^{\gamma n} \\
&= q^{-n(\eta-1/\log q)} \\
&\leq q^{-n\eta/2}\ . \qquad (26)
\end{aligned}$$

The last inequality follows from our assumption $q\geq 2^{400/\eta}$, which in particular guarantees that $\eta-1/\log q\geq\eta/2$. Thus, putting Eqs. (24) and (26) together, we conclude that

$$|\delta_{f,C}|\leq\Delta=q^{-\frac{\eta^2 n^{1-2\lambda}}{400}}$$

with probability at least

$$1-2^n q^{-\frac{\eta n}{200}}-q^{-8(1-\gamma)^{1+\lambda}n^{1+\lambda}}-q^{-\frac{n\eta}{2}}\ . \qquad (27)$$

Lastly, we observe that the terms $q^{-8(1-\gamma)^{1+\lambda}n^{1+\lambda}}$ and $q^{-\frac{n\eta}{2}}$ are both $o\left(2^n q^{-\frac{\eta n}{200}}\right)$. Therefore, for sufficiently large n the probability lower bound in Eq. (27) is at least

$$1-2^{n+1}q^{-\frac{\eta n}{200}}\ .$$

Additionally, since $q\geq 2^{400/\eta}$ by hypothesis, we get that $2^{n+1}q^{-\frac{\eta n}{200}}\leq 2q^{-\frac{\eta n}{400}}$, and so the probability lower bound in Eq. (27) is at least

$$1-2q^{-\frac{\eta n}{400}}.$$

This justifies the second lower bound on the failure probability given in Eq. (21) in the lemma statement. □

6.4.2 The Main Result for Massey Ramp Secret Sharing

The following theorem is an easy consequence of Lemmas 6.6, 6.8 and 6.11. It shows that the Massey ramp secret sharing scheme induced by a random linear code of any positive rate will be locally leakage-resilient with high probability.

We first state two well-known facts that will be useful in our analysis. The first one is easily implied by the proof of [GRS25, Section 4.2.2, Theorem 4.2].

Lemma 6.12. *Suppose that $C \subseteq \mathbb{F}_q^n$ is sampled from a distribution over linear codes of rate R such that for all $x \in \mathbb{F}_q^n \setminus \{0\}$ we have $\Pr[x \in C] \leq q^{k-n}$. Let $\theta > 0$ be arbitrary (possibly depending on n and q). Then, C has minimum distance at least $h_q^{-1}(1-R-\theta)n$ with probability at least $1 - q^{-\theta n}$.*

Second, we recall the following simple bound on the q-ary entropy function, for which a proof can be found in [GRS25, proof of Proposition 3.3.4]. We will use this bound to control the *inverse* of h_q, where we recall that the q-ary entropy function h_q has a well-defined continuous inverse on the interval $[0, 1-1/q]$.

Lemma 6.13. *For any positive integer q and real number $\rho \in (0, 1-1/q]$ we have*

$$h_q(\rho) \leq \rho + \frac{1}{\log q} \ .$$

Therefore, $h_q^{-1}(\rho + \frac{1}{\log q}) \geq \rho$ whenever $0 < \rho + \frac{1}{\log q} \leq 1 - \frac{1}{q}$.

Theorem 6.14. *Fix arbitrary $\gamma, \eta > 0$ satisfying $\gamma + \eta < 1$ and $\lambda \in (0, 1/5)$. Given $n \in \mathbb{N}$, let q be a prime satisfying*

$$\max\left(2^{\frac{1}{(1-\gamma)^{1+\lambda}}}, 2^{2/\gamma}, 2^{400/\eta}\right) \leq q \leq \frac{(n(1-\gamma))^\lambda}{\log(n(1-\gamma))} \ .$$

Furthermore, let $\theta \in \left(\frac{\log_q n}{n}, \min(\gamma, 1-(\gamma+\eta)) - \frac{1}{n+1}\right)$. Then, there exists $n_0 := n_0(\gamma, \eta, \lambda)$ such that, for all $n \geq n_0$, the following holds. Let $C \subseteq \mathbb{F}_q^n$ be a random linear code of dimension Rn with $R = \gamma + \eta$, let $v \in \mathbb{F}_q^n$ be sampled uniformly (and independently of G), and let $C^+ \subseteq \mathbb{F}_q^{n+1}$ be the code generated by

$$G^+ = \begin{bmatrix} 1 & 0 \\ v & G \end{bmatrix} \in \mathbb{F}_q^{(n+1)\times(k+1)} \ .$$

Then, the Massey ramp secret sharing scheme associated with C^+ is an $(n, t_{\mathrm{rec}}, t_{\mathrm{priv}}, \varepsilon)$-locally leakage-resilient ramp secret sharing scheme with

$$t_{\mathrm{rec}} \leq (1 - h_q^{-1}(1-R-\theta))n + 3 \leq \left(R + \theta + \frac{1}{\log q}\right) n + 2 \ ,$$
$$t_{\mathrm{priv}} \geq h_q^{-1}(R-\theta)n - 2 \geq \left(R - \theta - \frac{1}{\log q}\right) n - 2 \ ,$$

and

$$\varepsilon \leq 2^{n+1} q^{-\gamma n} + 2q^{-\frac{\eta^2 n^{1-2\lambda}}{800}} \leq 2q^{-\frac{\gamma n}{2}} + 2q^{-\frac{\eta^2 n^{1-2\lambda}}{800}}$$

with probability at least

$$1 - q^{-(1-R-\theta)n} - 2q^{-\theta(n+1)} - 2q^{-\frac{\eta n}{400}}$$

over the sampling of C^+.

Proof. Note that $C^+ \subseteq \mathbb{F}_q^{n+1}$ is a linear code of design rate R' satisfying $R \leq R' \leq R + \frac{1}{n+1}$. We first note that C^+ has rate R' with high probability. Indeed, this is the probability that G^+ has rank $k+1$, which is the same as the probability that G has rank k, which is

$$\prod_{i=0}^{k-1}\left(1 - \frac{1}{q^{n-i}}\right) \geq 1 - kq^{k-n} \geq 1 - q^{-(1-R)n + \log_q n} \geq 1 - q^{-(1-R-\theta)n} \tag{28}$$

since $\theta > \frac{\log_q n}{n}$ by hypothesis.

We next observe that if $x \in \mathbb{F}_q^{n+1} \setminus \{0\}$, then

$$\Pr_{v,G}\left[x \in C^+\right] = \Pr_{v,G}\left[\exists m \in \mathbb{F}_q^{k+1} \setminus \{0\} \text{ s.t. } G^+ m = x\right] \leq \sum_{m \in \mathbb{F}_q^{k+1}\setminus\{0\}} \Pr_{v,G}\left[G^+ m = x\right] \ . \tag{29}$$

Now, if one writes $m = (m_0, \tilde{m})$ with $m_0 \in \mathbb{F}_q$ and $\tilde{m} \in \mathbb{F}_q^k$ and $x = (x_0, \tilde{x})$ with $x_0 \in \mathbb{F}_q$ and $\tilde{x} \in \mathbb{F}_q^n$, we get that $G^+ m = x$ if and only if $m_0 v + G\tilde{m} = x$, which holds if and only if $m_0 = x_0$ and $G\tilde{m} = \tilde{x} - x_0 v$. If $\tilde{m} = 0$, then we must have $m_0 = x_0 \neq 0$ (as otherwise $m = 0$), and so $G\tilde{m} = \tilde{x} - x_0 v$ if and only if $v = \tilde{x}$, which occurs with probability q^{-n}. Otherwise, if $\tilde{m} \neq 0$, we get that $G\tilde{m}$ is uniformly random over $\mathbb{F}_q^n$, so again $\Pr[G\tilde{m} = \tilde{x} - x_0 v] = q^{-n}$. Consequently, the right-hand side of Eq. (29) is

$$\sum_{\substack{m \in \mathbb{F}_q^{k+1}\setminus\{0\} \\ m_0 = x_0}} q^{-n} \leq q^k \cdot q^{-n} = q^{(k+1)-(n+1)} \ .$$

We can then apply Lemma 6.12 to conclude that C^+ has distance $d \geq h_q^{-1}(1 - R' - \theta)(n+1) \geq h_q^{-1}(1 - R - \frac{1}{n+1} - \theta)(n+1)$ with probability at least

$$1 - q^{-\theta(n+1)}. \tag{30}$$

As for the dual of C^+, we have $x \in (C^+)^\perp$ if and only if $x^\top G^+ = 0$. Writing again $x = (x_0, \tilde{x})$ with $x_0 \in \mathbb{F}_q$ and $\tilde{x} \in \mathbb{F}_q^n$, this is equivalent to $\langle v, \tilde{x}\rangle = -x_0$ and $\tilde{x}^\top G = 0$. If $\tilde{x} = 0$, then $x_0 \neq 0$, so the probability that $\langle v, \tilde{x}\rangle = -x_0$ is 0 (we have $\Pr[\langle v, \tilde{x}\rangle = 0] = 1$). Else, $(\langle v, \tilde{x}\rangle, \tilde{x}^T G)$ is distributed uniformly over $\mathbb{F}_q^{k+1}$ (over the randomness of v and G), so it equals $(-x_0, 0)$ with probability $q^{-(k+1)}$. Hence,

$$\Pr_{v,G}\left[x \in (C^+)^\perp\right] \leq q^{-(k+1)} = \frac{q^{(n+1)-(k+1)}}{q^{n+1}} \ ,$$

so we may again apply Lemma 6.12 to conclude that $(C^+)^\perp$ has distance $d^\perp \geq h_q^{-1}(R' - \theta)(n+1) \geq h_q^{-1}(R - \theta)(n+1)$ with probability at least

$$1 - q^{-\theta(n+1)}. \tag{31}$$

Assuming these three favourable events hold, we conclude via Lemma 6.6 that the associated Massey ramp secret sharing scheme will have reconstruction threshold

$$t_{\text{rec}} = n - d + 2 \leq \left(1 - h_q^{-1}(1 - R - \tfrac{1}{n+1} - \theta)\right) n + 2 \leq \left(R + \theta + \tfrac{1}{\log q}\right) n + 3$$

and privacy threshold

$$t_{\text{priv}} = d^\perp - 2 \geq h_q^{-1}(R - \theta) n - 2 \geq \left(R - \theta - \tfrac{1}{\log q}\right) n - 2.$$

The rightmost bounds on t_{rec} and t_{priv} above follow from Lemma 6.13, noting that $\frac{1}{q} \leq R \leq 1 - \frac{1}{q}$ (with the upper bound holding by the lower bound on q from the theorem statement, which implies that $(1 - \frac{1}{q}) \log q > 2 \log(\frac{2}{1-2\gamma}) > R \log q$) and that $\theta < \min(R, 1 - R) - \frac{1}{n+1}$ by hypothesis.

Now, we consider the leakage-resilience property. Let $g = (g_1, \ldots, g_n)$ be a local leakage function, and let $\mathcal{S}_g$ denote the set of its induced rectangles (i.e., the set of all rectangles $S = S_1 \times \cdots \times S_n$ such that for each $i \in [n]$ we have $S_i = g_i^{-1}(b_i)$ for some $b_i \in \{0,1\}$). By Lemma 6.8, it suffices to bound

$$\begin{aligned}\mathrm{SD}(g(C), g(U_n)) &= \max_{\mathcal{A} \subseteq \mathcal{S}_g} \left(\Pr_{x \sim \mathbb{F}_q^n} [\exists S \in \mathcal{A} : x \in S] - \Pr_{x \sim C} [\exists S \in \mathcal{A} : x \in S] \right) \\ &= \max_{\mathcal{A} \subseteq \mathcal{S}_g} \left(\sum_{S \in \mathcal{A}} \left(\Pr_{x \sim \mathbb{F}_q^n} [x \in S] - \Pr_{x \sim C} [x \in S] \right) \right).\end{aligned}$$

By definition of statistical distance, we are then guaranteed the existence of a subset $\mathcal{A}_g \subseteq \mathcal{S}_g$ such that

$$\mathrm{SD}(g(C), g(U_n)) = \sum_{S \in \mathcal{A}_g} \left(\Pr_{x \sim \mathbb{F}_q^n} [x \in S] - \Pr_{x \sim C} [x \in S] \right) . \tag{32}$$

Partition $\mathcal{A}_g = \mathcal{A}_{g,s} \sqcup \mathcal{A}_{g,l}$, where $\mathcal{A}_{g,s}$ consists of all γ-small rectangles in $\mathcal{A}_g$, and $\mathcal{A}_{g,l}$ consists of the large rectangles. We can then rewrite Eq. (32) as

$$\mathrm{SD}(g(C), g(U_n)) = \sum_{S \in \mathcal{A}_{g,s}} \left(\Pr_{x \sim \mathbb{F}_q^n} [x \in S] - \Pr_{x \sim C} [x \in S] \right) + \sum_{S \in \mathcal{A}_{g,l}} \left(\Pr_{x \sim \mathbb{F}_q^n} [x \in S] - \Pr_{x \sim C} [x \in S] \right) .$$

The sum over $S \in \mathcal{A}_{g,s}$ can be bounded as

$$\sum_{S \in \mathcal{A}_{g,s}} \left(\Pr_{x \sim \mathbb{F}_q^n} [x \in S] - \Pr_{x \sim C} [x \in S] \right) \leq \sum_{S \in \mathcal{A}_{g,s}} \Pr_{x \sim \mathbb{F}_q^n} [x \in S] = \sum_{S \in \mathcal{A}_{g,s}} \frac{|S|}{q^n} \leq |\mathcal{A}_{g,s}| \, q^{-\gamma n} \leq 2^n \, q^{-\gamma n} , \tag{33}$$

where the second inequality follows from the fact that S being γ-small means that $|S| < q^{(1-\gamma)n}$.

As for the sum over $S \in \mathcal{A}_{g,l}$, let f_S denote the normalized indicator function for a combinatorial rectangle S. Recalling that $\delta_{C,f_S} = \frac{\Pr_{x \sim C}[x \in S]}{\Pr_{x \sim \mathbb{F}_q^n}[x \in S]} - 1$, we have

$$\Pr_{x \sim \mathbb{F}_q^n} [x \in S] - \Pr_{x \sim C} [x \in S] = \Pr_{x \sim \mathbb{F}_q^n} [x \in S] \, (-\delta_{C,f_S}) \leq \Pr_{x \sim \mathbb{F}_q^n} [x \in S] \cdot |\delta_{C,f_S}| .$$

Therefore,

$$\sum_{S \in \mathcal{A}_{g,l}} \left(\Pr_{x \sim \mathbb{F}_q^n} [x \in S] - \Pr_{x \sim C} [x \in S] \right) \leq \sum_{S \in \mathcal{A}_{g,l}} \Pr_{x \sim \mathbb{F}_q^n} [x \in S] \cdot |\delta_{C,f_S}| \leq \max_{S \in \mathcal{A}_{g,l}} |\delta_{C,f_S}| . \tag{34}$$

By Lemma 6.11 (using our hypotheses on γ, η, λ, and q), for sufficiently large n it holds that

$$\max_{S \in \mathcal{A}_{g,l}} |\delta_{C,f_S}| \leq q^{-\frac{\eta^2 n^{1-2\lambda}}{400}} \tag{35}$$

with probability at least

$$1 - 2q^{-\eta n/400} \tag{36}$$

over the sampling of C.

Combining Eqs. (33) to (35) and taking a union bound over the probability lower bounds for the four events in Eqs. (28), (30), (31) and (36), we obtain ε-leakage resilience for

$$\varepsilon \leq 2^{n+1} q^{-\gamma n} + 2q^{-\frac{\eta^2 n^{1-2\lambda}}{800}}$$

with probability at least

$$1 - q^{-(1-R-\theta)n} - 2q^{-\theta(n+1)} - 2q^{-\eta n/400},$$

as claimed. Lastly, using our hypothesis that $q \geq 2^{2/\gamma}$ we get that $2^n \leq q^{\frac{\gamma n}{2}}$, and so $2^{n+1} q^{-\gamma n} \leq 2q^{-\frac{\gamma n}{2}}$. That is, we can further bound

$$\varepsilon \leq 2q^{-\frac{\gamma n}{2}} + 2q^{-\frac{\eta^2 n^{1-2\lambda}}{800}},$$

as desired. □

In our study of "low-rate" leakage-resilient ramp secret sharing, we are mainly interested in how large we can take q and how small we can make the gap between the privacy and reconstruction thresholds. Since Theorem 6.14 has several parameters that make it hard to parse answers to these questions, we now present an example instantiation of this theorem that showcases these answers more transparently and which includes Theorem 6.5 as a special case. Roughly speaking, the instantiation below shows that we can take q as large as $n^{1/5-\nu}$ for an arbitrary constant $\nu > 0$ while simultaneously achieving reconstruction threshold $\approx Rn$ for an arbitrary constant $R > 0$ and threshold gap $\approx 2n/\log q$. Note that the threshold gap can be made $O(n/\log n)$ when q is polynomial in n.

Corollary 6.15. *Fix arbitrary constants $R \in (0,1)$ and $\nu \in (0, 1/5)$. Given $n \in \mathbb{N}$, let q be a prime satisfying*

$$2^{800/R} \leq q \leq n^{1/5-\nu}.$$

Then, for all sufficiently large n (depending only on R and ν) there is an $(n, t_{\text{rec}}, t_{\text{priv}}, \varepsilon)$-locally leakage-resilient ramp secret sharing scheme over $\mathbb{F}_q$ with

$$Rn - \frac{(1+\nu)n}{\log q} \leq t_{\text{priv}} < t_{\text{rec}} \leq Rn + \frac{(1+\nu)n}{\log q}$$

and

$$\varepsilon \leq 4q^{-\frac{R^2 n^{3/5}}{3200}}.$$

Proof. Fix $R \in (0,1)$ and $\nu \in (0, 1/5)$. We instantiate Theorem 6.14 with $\gamma = \eta = R/2$, $\lambda = \frac{1}{5} - \frac{\nu}{2}$, and $\theta = \frac{2\log_q n}{n}$. Note that for sufficiently large n we get that

$$\frac{\log_q n}{n} < \theta = \frac{2\log_q n}{n} < \min\left(\gamma, 1 - (\gamma+\eta) - \frac{1}{n+1}\right) = \min\left(\frac{R}{2}, 1 - R - \frac{1}{n+1}\right),$$

as required for the application of Theorem 6.14. Regarding the constraints on q for the application of this theorem, note that under our choices of γ, η, and λ above together with our hypotheses on q from the corollary statement we get that

$$\max\left(2^{\frac{1}{(1-\gamma)^{1+\lambda}}}, 2^{2/\gamma}, 2^{400/\eta}\right) = 2^{400/\eta} = 2^{R/800} \leq q \leq n^{1/5-\nu} < \frac{(n(1-\gamma))^{\lambda}}{\log(n(1-\gamma))}$$

for all sufficiently large n, since $\lambda = \frac{1}{5} - \frac{\nu}{2} > \frac{1}{5} - \nu$.

Then, Theorem 6.14 guarantees that the Massey ramp secret sharing scheme associated with the random linear code C^+ will be an $(n, t_{\mathrm{rec}}, t_{\mathrm{priv}}, \varepsilon)$-locally leakage-resilient ramp secret sharing scheme for

$$\left(R - \theta - \frac{1}{\log q}\right) n - 2 \leq t_{\mathrm{priv}} < t_{\mathrm{rec}} \leq \left(R + \theta + \frac{1}{\log q}\right) n + 2$$

and

$$\varepsilon \leq 2q^{-\frac{\gamma n}{2}} + 2q^{-\frac{\eta^2 n^{1-2\lambda}}{800}} \leq 2q^{-\frac{Rn}{4}} + 2q^{-\frac{R^2 n^{3/5}}{3200}} \leq 4q^{-\frac{R^2 n^{3/5}}{3200}},$$

recalling that $\gamma = \eta = R/2$ and $\lambda < 1/5$, with probability at least

$$1 - q^{-(1-R-\theta)n} - 2q^{-\theta(n+1)} - 2q^{-\frac{\eta n}{400}} \geq 1 - q^{\frac{(1-R)n}{2}} - \frac{1}{n^2} - 2q^{-\frac{Rn}{800}} > 0$$

over the sampling of C^+, where the inequalities hold for all sufficiently large n, recalling that $\theta = \frac{2\log_q n}{n}$.

It remains to simplify the bounds on t_{priv} and t_{rec} above for sufficiently large n. It suffices to note that

$$\left(\theta + \frac{1}{\log q}\right) n + 2 = 2\log_q n + \frac{n}{\log q} + 2 \leq \frac{n}{\log q}\left(1 + \frac{4\log n}{n}\right) \leq \frac{(1+\nu)n}{\log q},$$

where the first inequality uses the fact that $q \leq n$ and the last inequality holds for all sufficiently large n, since ν is a fixed constant. □

The simplified Theorem 6.5 follows directly from Corollary 6.15 by setting $\nu = 0.01$.

Acknowledgements

Dean Doron is supported in part by Israel Science Foundation grant #857/25 and by NSF-BSF grant #2022644. Tal Leonov is supported by Israel Science Foundation grant #3450/24. Jonathan Mosheiff is supported by Israel Science Foundation grant #3450/24 and an Alon Fellowship. Henrique Navas is funded by FCT - Fundação para a Ciência e a Tecnologia, I.P., under grant 2025.05868.BD and by national funds through FCT, I.P., and, when eligible, co-funded by EU funds under project/support UID/50008/2025 – Instituto de Telecomunicações, with DOI 10.54499/UID/50008/2025. Nicolas Resch is supported in part by an NWO (Dutch Research Council) Veni grant VI.Veni.222.347. João Ribeiro is funded by the European Union (LESYNCH, 101218842). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. João Ribeiro is also funded by national funds through FCT – Fundação para a Ciência e a Tecnologia, I.P., and, when eligible, co-funded by EU funds under project/support UID/50008/2025 – Instituto de Telecomunicações, with DOI 10.54499/UID/50008/2025.

References

[ADM+99] Noga Alon, Martin Dietzfelbinger, Peter Bro Miltersen, Erez Petrank, and Gábor Tardos. Linear hash functions. *Journal of the ACM (JACM)*, 46(5):667–683, 1999.

[AGL24] Omar Alrabiah, Venkatesan Guruswami, and Ray Li. 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*, pages 1458–1469, 2024.

[AS48] Milton Abramowitz and Irene A Stegun. *Handbook of mathematical functions with formulas, graphs, and mathematical tables*, volume 55. US Government printing office, 1948.

[BCDZ26] Joshua Brakensiek, Yeyuan Chen, Manik Dhar, and Zihan Zhang. Combinatorial bounds for list recovery via discrete Brascamp-Lieb inequalities. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, pages 365–376, 2026.

[BDIR21] Fabrice Benhamouda, Akshay Degwekar, Yuval Ishai, and Tal Rabin. On the local leakage resilience of linear secret sharing schemes. *Journal of Cryptology*, 34(2), February 2021. Preliminary version in CRYPTO 2018.

[BGM23] Joshua Brakensiek, Sivakanth Gopi, and Visu Makam. Generic Reed-Solomon codes achieve list-decoding capacity. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1488–1501, 2023.

[Bli87] Vladimir Markovich Blinovskii. Lower asymptotic bound on the number of linear code words in a sphere of given radius in $\mathbb{F}_q^n$. *Problemy Peredachi Informatsii*, 23(2):50–53, 1987. Translated in: Problems of Information Transmission, vol. 23, no. 2, pp. 130–132.

[Bli90a] Vladimir Markovich Blinovskii. Asymptotically exact uniform bounds for spectra of cosets of linear codes. *Problemy Peredachi Informatsii*, 26(1):99–103, 1990. Translated in: Problems of Information Transmission, vol. 26, no. 1, pp. 83–86, 1990.

[Bli90b] Vladimir Markovich Blinovskii. Covering the Hamming space with sets translated by vectors of a linear code. *Problemy Peredachi Informatsii*, 26(3):21–26, 1990. Translated in: Problems of Information Transmission, vol. 26, no. 3, pp. 196–201, 1990.

[Bsh26] Nader H Bshouty. A note on second-order expected maximum-load bounds for binary linear hashing. *arXiv preprint arXiv:2605.18335*, 2026.

[CCX13] Ignacio Cascudo, Ronald Cramer, and Chaoping Xing. Bounds on the threshold gap in secret sharing and its applications. *IEEE Transactions on Information Theory*, 59(9):5600–5612, 2013.

[CF85] Gérard Cohen and Peter Frankl. Good coverings of Hamming spaces with spheres. *Discrete Mathematics*, 56(2):125–131, 1985.

[CGV13] Mahdi Cheraghchi, Venkatesan Guruswami, and Ameya Velingker. Restricted isometry of Fourier matrices and list decodability of random linear codes. *SIAM Journal on Computing*, 42(5):1888–1914, 2013.

[Cha25] André Chailloux. OPI × soft decoders. *arXiv preprint arXiv:2511.22691*, 2025.

[CHLL97] Gérard Cohen, Iiro Honkala, Simon Litsyn, and Antoine Lobstein. *Covering Codes*, volume 54 of *North-Holland Mathematical Library*. Elsevier, 1997.

[Coh83] Gérard Cohen. A nonconstructive upper bound on covering radius. *IEEE Transactions on Information Theory*, 29(3):352–353, 1983.

[CT25] André Chailloux and Jean-Pierre Tillich. Quantum advantage from soft decoders. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 738–749, 2025.

[CZ25] Yeyuan Chen and Zihan Zhang. Explicit folded Reed-Solomon and multiplicity codes achieve relaxed generalized Singleton bounds. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25, page 1–12, New York, NY, USA, 2025. Association for Computing Machinery.

[DD22] Manik Dhar and Zeev Dvir. Linear hashing with ℓ_∞ guarantees and two-sided Kakeya bounds. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 419–428. IEEE, 2022.

[DMRR25] Dean Doron, Jonathan Mosheiff, Nicolas Resch, and João Ribeiro. List-recovery of random linear codes over small fields. *IEEE Transactions on Information Theory*, 2025.

[DP86] Philippe Delsarte and Philippe Piret. Do most binary linear codes achieve the Goblick bound on the covering radius? (corresp.). *IEEE Transactions on Information Theory*, 32(6):826–828, 1986.

[GHK11] Venkatesan Guruswami, Johan Håstad, and Swastik Kopparty. On the list-decodability of random linear codes. *IEEE Transactions on Information Theory*, 57(2):718–725, 2011.

[GHSZ02] Venkatesan Guruswami, Johan Hastad, Madhu Sudan, and David Zuckerman. Combinatorial bounds for list decoding. *IEEE Transactions on Information Theory*, 48(5):1021–1034, 2002.

[GLM+21] Venkatesan Guruswami, Ray Li, Jonathan Mosheiff, Nicolas Resch, Shashwat Silas, and Mary Wootters. Bounds for list-decoding and list-recovery of random linear codes. *IEEE Transactions on Information Theory*, 68(2):923–939, 2021.

[Gob62] Thomas J. Goblick. *Coding for a discrete information source with a distortion measure*. PhD thesis, MIT, October 1962. Available at https://dspace.mit.edu/entities/publication/fbba9539-8ae5-4d8a-bf89-6fb07116288a.

[GRS25] Venkatesan Guruswami, Atri Rudra, and Madhu Sudan. *Essential Coding Theory*. 2025. Draft available at https://cse.buffalo.edu/faculty/atri/courses/coding-theory/book/web-coding-book.pdf.

[Gur03] Venkatesan Guruswami. List decoding from erasures: bounds and code constructions. *IEEE Transactions on Information Theory*, 49(11):2826–2833, 2003.

[GZ23] Zeyu Guo and Zihan Zhang. Randomly punctured Reed-Solomon codes achieve the list decoding capacity over polynomial-size alphabets. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 164–176. IEEE, 2023.

[JKZ25] Michael Jaber, Vinayak M Kumar, and David Zuckerman. Linear hashing is optimal. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 245–255, 2025.

[JSW+25] Stephen P. Jordan, Noah Shutty, Mary Wootters, Adam Zalcman, Alexander Schmidhuber, Robbie King, Sergei V. Isakov, Tanuj Khattar, and Ryan Babbush. Optimization by decoded quantum interferometry. *Nature*, 646(8086):831–836, 2025.

[Kas24] Dustin Kasser. An improvement upon the bounds for the local leakage resilience of Shamir's secret sharing scheme. In *Theory of Cryptography Conference (TCC)*, page 395–422. Springer Nature Switzerland, December 2024.

[KK23] Ohad Klein and Ilan Komargodski. New bounds on the local leakage resilience of Shamir's secret sharing scheme. In *Advances in Cryptology – CRYPTO 2023*, page 139–170. Springer Nature Switzerland, 2023.

[KSG+25] Tanuj Khattar, Noah Shutty, Craig Gidney, Adam Zalcman, Noureldin Yosri, Dmitri Maslov, Ryan Babbush, and Stephen P. Jordan. Verifiable quantum advantage via optimized dqi circuits. *arXiv preprint arXiv:2510.10967*, 2025.

[Kum26] Vinayak M. Kumar. *Linearity, locality, and lower bounds: A unified perspective through randomness*. PhD thesis, UT Austin, 2026.

[Lev88] Lev B. Levitin. Covering radius of almost all linear codes satisfies the Goblick bound, 1988. Presented at a recent results session, 1988 IEEE International Symposium on Information Theory.

[LMS25] Matan Levi, Jonathan Mosheiff, and Nikhil Shagrithaya. Random Reed-Solomon codes and random linear codes are locally equivalent. In *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2097–2131. IEEE, 2025.

[Loe94] Hans-Andrea Loeliger. On the basic averaging arguments for linear codes. In *Communications and Cryptography: Two Sides of One Tapestry*, pages 251–261. Springer, 1994.

[LS25] Ray Li and Nikhil Shagrithaya. Near-optimal list-recovery of linear code families. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2025)*, pages 53–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2025.

[LW20] Ray Li and Mary Wootters. Improved list-decodability of random linear binary codes. *IEEE Transactions on Information Theory*, 67(3):1522–1536, 2020.

[Mas95] James L Massey. Some applications of coding theory in cryptography. *Codes and Ciphers: Cryptography and Coding IV*, pages 33–47, 1995.

[MPSW21] Hemanta K. Maji, Anat Paskin-Cherniavsky, Tom Suad, and Mingyuan Wang. Constructing locally leakage-resilient linear secret-sharing schemes. In *Advances in Cryptology – CRYPTO 2021*, page 779–808. Springer International Publishing, 2021.

[MRRZ+24] Jonathan Mosheiff, Nicolas Resch, Noga Ron-Zewi, Shashwat Silas, and Mary Wootters. Low-density parity-check codes achieve list-decoding capacity. 53(6):FOCS20–38–FOCS20–73, 2024.

[Ngu24] Hai H. Nguyen. Towards breaking the half-barrier of local leakage-resilient Shamir's secret sharing. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024*, pages 257–285, Cham, 2024. Springer Nature Switzerland.

[Ngu26] Hai H. Nguyen. Local leakage resilience of Shamir secret sharing: Below half threshold rate, October 2026. https://arxiv.org/abs/2610.17711.

[Rib26] João Ribeiro. How old is the potential method in coding theory?, 2026. https://jmlribeiro.github.io/how-old-potential-method.pdf.

[Ros26] Ansis Rosmanis. A nearly linear-time decoded quantum interferometry algorithm for the optimal polynomial intersection problem. *arXiv preprint arXiv:2601.15171*, 2026.

[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, 2014.

[RW18] Atri Rudra and Mary Wootters. Average-radius list-recoverability of random linear codes. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 644–662. SIAM, 2018.

[RY24] Nicolas Resch and Chen Yuan. Threshold rates of code ensembles: Linear is best. *IEEE Transactions on Information Theory*, 70(7):4823–4842, 2024.

[ST23] Chong Shangguan and Itzhak Tamo. Generalized Singleton bound and list-decoding Reed–Solomon codes beyond the Johnson radius. *SIAM Journal on Computing*, 52(3):684–717, 2023.

[SW26a] Yihang Sun and Mary Wootters. Improved local leakage resilience of Shamir secret sharing and worst-case optimal polynomial intersection, October 2026. https://arxiv.org/abs/2610.12357.

[SW26b] Yihang Sun and Mary Wootters. On worst-case optimal polynomial intersection, April 2026. https://arxiv.org/abs/2604.09533.

[TX22] Ivan Tjuawinata and Chaoping Xing. Leakage-resilient secret sharing with constant share size. *IEEE Transactions on Information Theory*, 68(12):8228–8250, 2022.

[Woo13] Mary Wootters. 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*, pages 853–860, 2013.

[ZP81] Victor Vasilievich Zyablov and Mark Semenovich Pinsker. List concatenated decoding. *Problemy Peredachi Informatsii*, 17(4):29–33, 1981.

A Deferred Proofs

A.1 Deferred Proofs from Section 2

Proof of Theorem 2.12. Observe that the second event in the statement implies the first one. Indeed, let $S = S_1 \times \cdots \times S_n$ be a sequence of sets $S_i \subseteq \mathbb{F}_q$, each of size at most ℓ. Then there are sets $S'_1, \ldots, S'_n$, such that $S_i \subseteq S'_i$ and $|S'_i| = \ell$ for all i. By the second event,

$$|S \cap C| \leq \left|S' \cap C\right| \leq \left(1 + q^{-c\gamma n/(q^2 \log(q/\gamma))}\right) \frac{|C|\ell^n}{q^n} \leq \left(1 + q^{-c\gamma n/(q^2 \log(q/\gamma))}\right) q^{\gamma n} \ ,$$

implying the first event. It thus suffices to prove that the second event holds with probability at least

$$1 - 2q^{-c\gamma n/(q^2 \log(q/\gamma))} \ .$$

Let $\mathcal{F}$ be the family of normalized indicators $f_S = (q^n/|S|)\mathbf{1}_S$ over rectangles defined by $S_1, \ldots, S_n$ each of size exactly ℓ. Then

$$|\mathcal{F}| = \binom{q}{\ell}^n \leq 2^{qn} \leq q^{qn},$$

and $\|f_S\|_\infty = (q/\ell)^n$, so the capacity parameter in Corollary 2.6 is $\beta = 1 - \log_q \ell$.

By Lemma 5.2, every $f_S \in \mathcal{F}$ is α-Fourier-concentrated with $\alpha = \cos(\pi/q)$. Therefore the parameter λ from Corollary 2.6 satisfies

$$\lambda = \min\left\{1, \frac{\log(1/\cos(\pi/q))}{2\log(100q/\gamma)}\right\} \geq \frac{c_1}{q^2 \log(q/\gamma)}$$

for a universal constant $c_1 > 0$, using Lemma 5.3 to lower bound $\log(1/\cos(\pi/q)) \geq \Omega(1/q^2)$.

Choose the theorem constant $c > 0$ so that $c \leq \min\{c_1/4, 1/200\}$, and set

$$\theta := \frac{c\gamma}{q^2 \log(q/\gamma)}.$$

Then $\theta < \gamma\lambda$, and

$$2(\gamma\lambda - \theta) \geq \frac{c\gamma}{q^2 \log(q/\gamma)}.$$

Also, since $\gamma \leq \frac{1}{2}(1 - \log_q \ell) = \beta/2$, the hypothesis $\gamma \leq 50\beta/51$ of Corollary 2.6 is satisfied.

Taking c' large enough, the assumed lower bound

$$n \geq c' \frac{q^5 \log^2(q/\gamma)}{\gamma^2}$$

implies the lower bound $n \geq n_0(q, \gamma, \alpha, \theta)$ required in Corollary 2.6. Hence Corollary 2.6, applied with $\eta = \gamma$, gives

$$|\delta_{C,f_S}| \leq q^{-c\gamma n/(q^2 \log(q/\gamma))}$$

for every $f_S \in \mathcal{F}$, except with failure probability at most

$$q^{-\gamma n/200} + |\mathcal{F}| q^{n-\theta^2 n^2/20}.$$

The first term is at most $q^{-c\gamma n/(q^2 \log(q/\gamma))}$ by the choice of c. For the second, using $|\mathcal{F}| \le q^{qn}$ and increasing c' if necessary, the lower bound on n gives $\theta^2 n/20 \ge q + 1 + \theta$. Therefore

$$|\mathcal{F}| q^{n-\theta^2 n^2/20} \le q^{qn+n-\theta^2 n^2/20} \le q^{-\theta n} = q^{-c\gamma n/(q^2 \log(q/\gamma))}. \tag{37}$$

Thus the desired deviation event holds with probability at least

$$1 - 2q^{-c\gamma n/(q^2 \log(q/\gamma))}.$$

On this event, every rectangle S with side length exactly ℓ satisfies

$$|C \cap S| = \left(1 \pm q^{-c\gamma n/(q^2 \log(q/\gamma))}\right) \frac{|C| \ell^n}{q^n} \quad .$$

□

Proof of Theorem 2.14. The proof is nearly identical to the proof of Theorem 2.12, so we only highlight the few differences. Letting S be any sequence of sets that corresponds to list recovery with ζ-fraction of erasures, we have that $|S| = q^{\zeta n} \ell^{(1-\zeta)n}$ (we can assume that exactly ζ-fraction of the lists satisfy $S_i = \mathbb{F}_q$ and the rest satisfy $|S_i| = \ell$). Keeping with the same notation as above, we have that $\|f_S\|_\infty = (q/\ell)^{(1-\zeta)n}$, so indeed $\beta = (1-\zeta)(1 - \log_q \ell)$. Our choice of γ still satisfies $R < 1$ and $\gamma \le \beta/2$, and note that the family $\mathcal{F}$ is still α-concentrated with the same α.

However, $|\mathcal{F}|$ has changed, so we need to verify that the bound on the failure probability still holds. A straightforward bound gives

$$|\mathcal{F}| = \binom{n}{\zeta n} \binom{q}{\ell}^{(1-\zeta)n} \le \sum_{i=0}^{n} \binom{n}{i} \binom{q}{\ell}^{n-i} = \left(\binom{q}{\ell} + 1\right)^n \le 2^{qn} \le q^{qn},$$

so Equation (37) still holds. The rest of the proof is unchanged. □

A.2 Deferred Proofs from Section 6

Proof of Lemma 6.6. To see the claim about reconstruction, note that since C^+ has minimum distance d, it corrects any $d-1$ erasures. Therefore, $(n+1) - (d-1) = n - d + 2$ shares are sufficient to recover $c^+ = Gx$, and so we recover the secret $s = c_0^+$. To see the claim about privacy, note that since C^+ has dual distance $d^\perp$, then the uniform distribution over C^+ is $(d^\perp - 1)$-wise independent. This means that if $c^+ \sim C^+$, then for any indices $1 \le i_1 < i_2 < \cdots < i_{d^\perp - 2} \le n$ we have that $(c_0^+, c_{i_1}^+, \ldots, c_{i_{d^\perp-2}}^+)$ is uniformly distributed over $\mathbb{F}_q^{d^\perp - 1}$. Therefore, the shares associated with parties $i_1, \ldots, i_{d^\perp - 2}$ are independent of the secret c_0^+. □

Proof of Lemma 6.8. Fix an arbitrary tuple of leakage functions $g = (g_1, \ldots, g_n)$. Fix also two secrets $s, s' \in \mathbb{F}_q$, with corresponding tuples of shares $S = (S_1, \ldots, S_n)$ and $S' = (S'_1, \ldots, S'_n)$, respectively. Recall that the shares of s (resp. s') correspond to a uniformly random sample from

$s \cdot v + C$ (resp. $s' \cdot v + C$). Let X and X' be independent and uniformly distributed over C. Define also $g_s(z) = g(s \cdot v + z)$ for any $s \in \mathbb{F}_q$. By the assumption that $\mathcal{G}$ is closed under input-translations, we have $g_s \in \mathcal{G}$. Then,

$$\begin{aligned}
\mathrm{SD}(g(S), g(S')) &\le \mathrm{SD}(g(s \cdot v + X), g(U_n)) + \mathrm{SD}(g(s' \cdot v + X'), g(U_n)) \\
&= \mathrm{SD}(g_s(X), g(U_n)) + \mathrm{SD}(g_{s'}(X'), g(U_n)) \\
&= \mathrm{SD}(g_s(X), g_s(U_n)) + \mathrm{SD}(g_{s'}(X'), g_{s'}(U_n)) \\
&= \mathrm{SD}(g_s(C), g_s(U_n)) + \mathrm{SD}(g_{s'}(C), g_{s'}(U_n)) \\
&\le 2\varepsilon.
\end{aligned}$$

The first inequality follows by the triangle inequality, the second equality uses the fact that $g(U_n)$ and $g_s(U_n)$ are identically distributed for any $s \in \mathbb{F}_q$, and the last inequality uses the hypothesis from the lemma statement applied to g_s and $g_{s'}$. □

Proof of Corollary 6.10. We apply Theorem 2.5 and refer to notation in the context of that theorem's statement.

First, by Lemma 5.2, $\mathcal{F}$ is a family of normalized α-concentrated functions for $\alpha \le \cos(\pi/q)$. Since $\log(1/\alpha) \ge \log(1/\cos(\pi/q)) \ge \frac{\pi^2}{8q^2 \ln 2}$ by Lemma 5.3 for all $q \ge 3$, we get that

$$\begin{aligned}
d &= \max\left\{0, \left\lceil \log_2\left(\frac{\log_2\left((q-1)/(q^{\eta/100}-1)\right)}{\log_2(1/\alpha)}\right)\right\rceil\right\} \\
&\le \max\left\{0, \left\lceil \log_2\left(\frac{8q^2 \ln 2}{\pi^2} \cdot \log_2\left(\frac{q-1}{q^{\eta/100}-1}\right)\right)\right\rceil\right\} \\
&\le \left\lceil \log_2\left(\frac{8q^2 \ln 2}{\pi^2} \cdot \log_2\left(\frac{q-1}{2^{\eta/100}-1}\right)\right)\right\rceil \\
&\le 1 + \log_2\left(\frac{8q^2}{\pi^2} \cdot \log_2\left(\frac{q-1}{2^{\eta/100}-1}\right)\right). \qquad (38)
\end{aligned}$$

Furthermore, observe that $\min(R - \frac{\eta}{50}, 2\eta) = \min(\beta + \frac{49\eta}{50}, 2\eta) \ge \frac{49\eta}{50}$. Choose $t = 10n^{\frac{1+\lambda}{2}}$. Then, using the observations above, for all sufficiently large n we have that

$$\begin{aligned}
\varepsilon &= 2^{-d} \min\left(R - \frac{\eta}{50}, 2\eta\right) - \frac{2t}{n} \\
&\ge 2^{-d} \cdot \frac{49\eta}{50} - \frac{2t}{n} \\
&\ge \frac{\pi^2 \eta}{32q^2 \log\left(\frac{q-1}{2^{\eta/100}-1}\right)} - 20n^{-\frac{1-\lambda}{2}} \\
&\ge \frac{\eta \cdot n^{-2\lambda}}{\log\left(\frac{1}{2^{\eta/100}-1}\right)} - 20n^{-\frac{1-\lambda}{2}} \\
&\ge (2^{\eta/100}-1)\eta \cdot n^{-2\lambda} - 20n^{-\frac{1-\lambda}{2}} \\
&\ge \frac{\eta^2 \cdot n^{-2\lambda}}{100 \log_2 e} - 20n^{-\frac{1-\lambda}{2}}.
\end{aligned}$$

The third inequality uses the fact that $q \leq \frac{n^\lambda}{\log n}$, and holds for large enough n. The last inequality uses the fact that $2^x - 1 \geq \frac{x}{\log_2 e}$ for all $x > 0$. Since $\lambda < 1/5$ we get that $n^{-\frac{1-\lambda}{2}} = o(n^{-2\lambda})$, and so for all sufficiently large n we have

$$\varepsilon \geq \frac{\eta^2 \cdot n^{-2\lambda}}{200 \log_2 e} \geq \frac{\eta^2 \cdot n^{-2\lambda}}{400}.$$

Concluding, from Theorem 2.5 we get that for sufficiently large n it holds that

$$|\delta_{C,f}| \leq q^{-\frac{\eta^2 \cdot n^{1-2\lambda}}{400}}$$

for every $f \in \mathcal{F}$ with probability at least

$$1 - n(Rn+1)q^{-\frac{\eta n}{100}} - |\mathcal{F}| \cdot d \cdot q^{n-10n^{1+\lambda}}$$

over the choice of C. It remains to argue the desired lower bound on the probability. Using Eq. (38), for large enough n we get that

$$d \leq 1 + \log_2\left(\frac{8q^2}{\pi^2} \cdot \log_2\left(\frac{q-1}{2^{\eta/100}-1}\right)\right) \leq 1 + \log_2\left(\frac{8q^2}{\pi^2} \cdot \left(\log_2 q + \frac{100}{\eta}\right)\right) \leq \log_2 n + \log_2\left(\frac{100}{\eta}\right),$$

from the fact that $\log_2\left(\frac{1}{2^x-1}\right) \leq \frac{1}{x}$ for all $x > 0$ and that $q \leq \frac{n^\lambda}{\log n}$ for $\lambda < 1/5$. Combining this observation with the fact that η is constant, we conclude that the probability is at least

$$1 - n(Rn+1)q^{-\frac{\eta n}{100}} - |\mathcal{F}| \cdot d \cdot q^{n-10n^{1+\lambda}} \geq 1 - q^{-\frac{\eta n}{200}} - |\mathcal{F}| \cdot q^{-9n^{1+\lambda}}$$

for all sufficiently large n. □

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il