EA Codes Approaching Singleton Bound (with Application to Field-Agnostic SNARKs)

Chongrong Li[1], Runtian Xu[1], Yun Li[2], Michael Dong[3], Alan Li[3], Yu Yu[1], and Yuncong Hu[1⋆]

[1] Shanghai Jiao Tong University
chongrongli@sjtu.edu.cn, 51255902065@stu.ecnu.edu.cn, yuyuathk@gmail.com, huyuncong@sjtu.edu.cn
[2] Ant Group
liyun24@antgroup.com
[3] Brevis Network
mdong@brevis.network,alan@brevis.network

Abstract. SNARKs based on error-correcting codes require codes that simultaneously support fast encoding and large relative distance. Reed–Solomon codes achieve the optimal rate–distance tradeoff given by the Singleton bound, but their fast encoding relies on FFT-friendly fields, limiting their applicability to field-agnostic constructions.

In this work, we revisit expand–accumulate (EA) codes, a simple family of linear codes with efficient encoding over arbitrary fields. Although extensively studied in prior work, existing distance guarantees are weak, and practical instantiations have relied on conjectured results. We prove that EA codes with exact-weight expansion matrices, over sufficiently large fields, achieve a rate–distance tradeoff that approaches the Singleton bound with high probability, resolving these conjectures.

Building on this result, we construct Flare, a field-agnostic polynomial commitment scheme based on EA codes. At its core is an efficient IOPP for EA codes that leverages code switching. For statements of size M, Flare achieves $O(M \log M)$ prover time and $O(\lambda \log^2 M)$ proof size, improving upon the $O(\sqrt{\lambda M})$ proof size of prior constructions based on EA codes (Block et al., CRYPTO 2024).

Comprehensive evaluations show that Flare achieves compact proofs while maintaining competitive prover efficiency. For proving circuits of size 2^{22}, Flare produces the smallest proofs of approximately 16 MiB. Compared with BaseFold (Zeilberger et al., CRYPTO 2024) and ERA (Baweja et al., ASIACRYPT 2026), it reduces proof size by 1.66× and 1.47×, respectively, while maintaining comparable prover efficiency. Compared with Brakedown (Golovnev et al., CRYPTO 2023), Flare produces 4.85× smaller proofs at the cost of only a 1.50× slower prover.

1 Introduction

SNARKs (Succinct Non-Interactive Arguments of Knowledge) are cryptographic primitives that allow a prover to convince a verifier of the validity of a statement

⋆ Corresponding author.

using a succinct proof that can be verified efficiently. These properties have led to their widespread use in applications including blockchain [21, 27, 29, 49], anonymous credentials [5, 23], and machine learning [35, 43, 47].

Modern SNARKs typically follow a design paradigm that combines a Polynomial Interactive Oracle Proof (PIOP) with a Polynomial Commitment Scheme (PCS). The PIOP reduces the validity of the statement to a few evaluation claims about polynomials provided by the prover, while the PCS provides such evaluation proofs: it allows the prover to commit to a polynomial f and later demonstrate that $f(r) = y$ for an evaluation point r and a claimed evaluation y. As the cost of the PCS often dominates the overall cost, an efficient PCS is essential for achieving high-performance SNARKs.

Recently, PCSs based on error-correcting codes [2, 3, 4, 11, 13, 17, 19, 50, 51] have been extensively studied due to their concrete efficiency and plausible post-quantum security. These schemes are constructed in a similar way: the prover sends the encoding of the witness under some linear code C, and the verifier checks that each encoding is close to a codeword of C by opening it at randomly chosen indices. The efficiency is thus determined by two properties of the underlying code. On the prover side, the cost is dominated by the encoding itself. On the verifier side, a higher relative distance increases the probability that each opening detects a deviation, thus requiring fewer openings and yielding a smaller proof and faster verification. Code-based PCS therefore call for linear codes that combine fast encoding with high relative distance.

In terms of distance, Reed–Solomon (RS) codes achieve the best possible rate–distance tradeoff by attaining the Singleton bound. Therefore, they have been widely used to construct efficient PCSs [2, 3, 4, 8]. However, the encoding efficiency of RS codes relies heavily on FFT-friendly fields, namely those containing sufficiently large subgroups of smooth order. This requirement substantially restricts the choice of the native computation field. For statements defined over an FFT-unfriendly field, an RS-based scheme must emulate the non-native arithmetic, which can incur significant overhead. For instance, [11] observes that an ECDSA verification circuits over the secp256k1 curve can be about $25\times$ smaller when expressed over its native field rather than over an FFT-friendly field.

This limitation has motivated the study of *field-agnostic* codes, namely codes that admit efficient encoding over arbitrary fields with high distance. Despite recent progress, no existing field-agnostic construction matches both the encoding efficiency and the distance guarantees offered by RS codes. For example, the Brakedown codes [17, 20] admits fast encoding over any field, but at rate $1/2$ over a 256-bit field it achieves a relative distance of only 0.1, far below the value 0.5 permitted by the Singleton bound. The BaseFold codes [51] achieve a relative distance of 0.728 at rate $1/8$, albeit at the cost of increased encoding time. Moreover, this distance still falls well short of the corresponding asymptotic Singleton bound of 0.875. Thus, a natural question arises:

Can we construct linear codes over arbitrary fields that simultaneously enjoy fast encoding and a rate–distance tradeoff approaching the Singleton bound?

1.1 Our Contribution

In this work, we answer the above question affirmatively by revisiting expand-accumulate (EA) codes, which have been widely used in various areas such as building pseudorandom correlation generators [12, 18] and SNARKs [11].

Intuitively, the EA code first multiplies the message by a sparse expansion matrix E, sampled according to a certain distribution, and then applies an accumulator A that computes prefix sums. Since E is sparse, both steps can be performed efficiently over any field. However, the distance guarantees in prior work are weak. For example, the provable distance in [11] over large fields is bounded by 10^{-6}, which would incur an impractical verifier overhead and proof size. Consequently, [11] had to rely on a conjectured distance in their evaluation.

For example, the analysis in [12] requires the sparse matrix E to be sampled entrywise from a Bernoulli distribution, resulting in a varying number of nonzero entries across rows. This irregularity makes the encoding algorithm less structured and incurs a higher overhead. Another work [11] requires concatenating two independently sampled EA codes, which at least doubles the encoding time. In practice, both works conjecture that the more efficient variant, in which all rows of E have the same weight, also achieves high distance.

Formally, let $\mathsf{BP}(n, N, t, \mathbb{F}_q)$ denote the *exact-weight* distribution over matrices $E \in \mathbb{F}_q^{n\times N}$, where the n rows of E are sampled independently and uniformly from the set of vectors in $\mathbb{F}_q^N$ of Hamming weight exactly t. Let A denotes the accumulator, i.e., the upper-triangular matrix with ones, and $\mathsf{d}(\cdot)$ denotes the relative distance of the code generated by a given matrix.

Conjecture 3.11 ([12]). There exist constants $R, \delta, C > 0$ such that the following holds. Let $N \in \mathbb{N}$, $n := \lfloor RN \rfloor$, and $t \in \mathbb{N}$ with $t \geq C \ln N$. Then

$$\Pr_{E\leftarrow\mathsf{BP}(n,N,t,\mathbb{F}_2)}[\mathsf{d}(EA) \geq \delta] \geq 1 - \frac{1}{\mathrm{poly}(N)}.$$

Conjecture 4.6 ([11]). Let $R := n/N$ be a constant. There exist constants $\gamma \geq 1$ and $\delta < 1/2$ such that, for $t := \gamma \log N$, we have

$$\Pr_{E\leftarrow\mathsf{BP}(n,N,t,\mathbb{F}_q)}[\mathsf{d}(EA) \geq \delta] \geq 1 - \frac{1}{\mathrm{poly}(N)}.$$

EA Codes Approaching Singleton Bound. We resolve these conjectures affirmatively by demonstrating that the EA code with exact-weight sampling achieves a large distance with high probability over a large finite field. More strongly, we demonstrate that the EA code achieves a rate–distance tradeoff arbitrarily close to the Singleton bound over sufficiently large fields.

Our main result is summarized as follows:

Theorem 1 (informal). *For every $0 < \delta < 1$ and $0 < R < 1 - \delta$, there exists $\gamma > 0$ such that the following holds. Let $n = \lfloor RN \rfloor$ and $t = \lceil \gamma \ln N \rceil$. If $\log q = \omega(\log N)$, then*

$$\Pr_{E\leftarrow\mathsf{BP}(n,N,t,\mathbb{F}_q)}[\mathsf{d}(EA) \leq \delta] \leq O(N^{(1-\gamma \ln(1/\delta))/2}).$$

Message length n	Analysis	γ	Distance δ
2^{10}	Our result	15	0.490
	[11]	18	0.100
2^{14}	Our result	11	0.490
	[11]	18	0.100
2^{18}	Our result	9	0.490
	[11]	18	0.100

Table 1: Comparison of concrete parameters for EA codes under different analyses. γ determines the row weight of the expansion matrix, lower is better.

Scheme	Rate	Distance δ	Singleton bound
EA (Our Result)		0.740	
QA [28]		0.727	
Brakedown [17]	1/4	0.100	0.750
Basefold [51]		0.569	
RAA [1]		0.500	

Table 2: Distance of different codes over a 256-bit prime field.

We compare our results with the conjectures from prior work [11] in Table 1. Our results substantially improve the conjectured distance guarantee while requiring a smaller value of γ. Specifically, prior work conjectured that, over a 256-bit prime field and with failure probability 2^{-100}, a rate-1/2 EA code with $\gamma \geq 18$ achieves relative distance 0.1. By contrast, under the same setting, our analysis shows that $\gamma = 9$ already suffices to guarantee a substantially larger relative distance of 0.49.

We also compare our results with several widely used codes over a 256-bit prime field, using a message length of 2^{20}. These include Basefold [51], Brakedown [17], QA [28], and RAA [1]. The results are summarized in Table 2. All codes except RAA are instantiated with failure probability at most 2^{-100}, whereas RAA achieves only a failure probability of 2^{-20}. We do not include ERA [6] in the table because it does not support rate 1/4. For reference, an ERA code with rate 1/4.4 achieves a relative distance of 0.231, which is still substantially below the corresponding Singleton bound of 0.773. Overall, among the schemes considered, our construction achieves the largest relative distance and is the closest to the Singleton bound.

The result established above applies to sufficiently large fields. However, in several applications of interest, such as PCGs and binary-field SNARKs, EA codes are instantiated over the binary field $\mathbb{F}_2$. As an independent contribution, in Appendix B, we analyze EA codes over small finite fields and show that they asymptotically approach the Gilbert–Varshamov (GV) bound. This matches the

best-known asymptotic distance guarantee for linear codes over small alphabets and resolves the conjecture of [12].

Efficient Field-Agnostic PCS from EA Codes. We construct a new field-agnostic PCS, Flare, based on our EA code. For statements of size M, Flare achieves $O(M \log M)$ prover time and $O(\lambda \log^2 M)$ proof size, improving upon the $O(\lambda\sqrt{M})$ proof size of prior work [11]. Our construction follows the code-switching paradigm [37,38], which has been used in a series of prior works [7,13, 19]. The core of our construction is an efficient IOP that simultaneously proves the proximity of a given word to an EA codeword and an evaluation claim on the multilinear extension of the corresponding message. To enable code switching, the prover supplies an encoding of the same message under an auxiliary code, while consistency with the EA codeword is enforced through a small number of random spot checks. By exploiting the expansion–accumulation structure of EA codes, we efficiently batch these consistency checks and reduce their verification to two sumcheck instances. We also extend the protocol to interleaved EA codes via random linear folding, enabling the code switching protocol to be applied with only a small additional overhead.

Implementation and Evaluation. We fully implement Flare and conduct comprehensive benchmarks against state-of-the-art field-agnostic PCSs. For polynomials of size 2^{22}, our standalone PCS benchmarks show that Flare produces 1.73× and 1.61× smaller proofs than BaseFold and ERA, respectively, while achieving 1.54× and 1.42× faster verification. Its opening is 2.83× slower than BaseFold but 1.18× faster than ERA. Compared with Brakedown, Flare reduces proof size by 4.67× and achieves 2× faster verification, at the cost of an 8.86× slower opening. In end-to-end SNARK benchmarks for R1CS instances with 2^{22} constraints, Flare achieves the smallest proofs and fastest verification among all evaluated schemes, with a proof size of 15.67 MiB and a verification time of 213.18 ms. Relative to BaseFold, Flare produces 1.66× smaller proofs and verifies 1.81× faster, with only a 1.27× slower prover. Relative to ERA, it achieves comparable prover efficiency while producing 1.47× smaller proofs and verifying 1.31× faster. Compared with Brakedown, Flare produces 4.85× smaller proofs and verifies 2.45× faster, with only a 1.50× slower prover.

1.2 Additional Related Work

We review several code families used in code-based PCSs and the properties of the resulting PCSs. In this section, we use M to denote the instance size.

Brakedown Codes. Brakedown [17] introduces a family of field-agnostic code with linear-time encoding based on the Spielman code [42]. The refined analysis of [20] shows that their relative distance is at most 0.2 and is typically below 0.1 for practical parameters. The original Brakedown PCS has $O(\sqrt{\lambda M})$ proof size and verifier time. BrakingBase [33] reduces this to $O(\lambda \log^2 M)$ by proposing an efficient code-switching protocol for Brakedown codes.

Random Foldable Codes. BaseFold [51] generalizes the FRI folding paradigm from Reed–Solomon codes to a broader class of foldable linear codes, and further introduces random foldable codes (RFCs) that support $O(M \log M)$ encoding

over any field. Its analysis demonstrates, RFCs achieves relative distance 0.728 at rate 1/8 over a 256-bit field. By interleaving the IOPP for RFCs with the sumcheck protocol, BaseFold constructs a field-agnostic PCS with $O(M \log M)$ commitment time and $O(\lambda \log^2 M)$ proof size and verifier time. However, the resulting PCS introduces large computational overhead due to the multi-round folding procedure and repeated Merkle commitments.

RAA Codes. Repeat-Accumulate-Accumulate (RAA) codes [16] are a family of linear-time encodable codes originally introduced over the binary field. Blaze [13] provides a refined analysis of binary RAA codes, which establishes a relative distance of 0.19 at rate 0.25. Blaze also develops a code-switching protocol for RAA codes, yielding a PCS over binary extension fields with $O(\lambda \log^2 M)$ proof size and verifier time. Several recent works [1, 22, 25] further extend and analyze RAA codes over prime fields, thereby making them applicable to field-agnostic PCS constructions. However, their analyses establish relative distance only up to 0.5. Moreover, the resulting PCSs have $O(\sqrt{\lambda M})$ proof size and verifier time, as they directly follow the Brakedown approach without code switching. We believe that the code-switching technique of Blaze can also be adapted to these constructions, potentially reducing both complexities to $O(\lambda \log^2 M)$.

Sketched Codes. Bolt [19] and Lightning [34] introduce *sketched codes*, which combine the parity-check matrix of a random LDPC code with an arbitrary linear code, thereby enabling extremely efficient encoding. A key limitation, however, is their relatively small relative distance. For example, over a 254-bit prime field, the relative distance is only up to 0.011 [34], largely because it is governed by the relative distance of the underlying random LDPC code. For PCS applications, Bolt proposes a code-switching protocol from sketched codes to RS codes, and is therefore not field-agnostic. However, the protocol may be made field-agnostic by replacing the RS codes with suitable field-agnostic codes, potentially yielding $O(\lambda \log^2 M)$ proof size and verifier time.

Quasi-Abelian Codes. Li et al. [28] construct random rank-one quasi-Abelian (QA) codes which can be encoded with $O(M \log M)$ time using the Walsh–Hadamard transformation. At rate 1/2 and message length 2^{20} over a 256-bit field, their distance guarantee is 0.4570. While this is favorable compared the other codes discussed here, it remains below the Singleton bound of 0.5. Moreover, following Brakedown, they construct a PCS directly from interleaved codes without code switching. As a result, their PCS has proof size and verifier time $O(\sqrt{\lambda M})$, which is larger than the $O(\lambda \log^2 M)$ achieved by our construction.

ERA Codes. Baweja et al. [6] propose Encode-Repeat-Accumulate (ERA) codes, a variant of RAA codes that first encodes the message using an arbitrary linear base code and then applies the standard RAA encoding procedure. When the base code admits linear-time encoding, ERA codes likewise achieve linear-time encoding. Their analysis shows a relative distance of 0.5039 at rate 1/8.8 over a 256-bit field, which remains well below the Singleton bound of roughly 0.88 targeted by our construction. For PCS, they develop a code-switching protocol for ERA codes that iteratively switches to codes with smaller message

lengths, yielding $O(\lambda \log M)$ query complexity. When instantiated with Merkle commitments, this results in an $O(\lambda \log^2 M)$ proof size and verifier time.

EA Codes. In concurrent work [24], Khabbazian studies the same class of exact-weight EA codes, under the name of fixed-row-weight EA codes. Some of their arguments are based on high-level ideas similar to ours, such as exploiting the geometric structure of the expansion matrix and the algebraic constraints imposed by the accumulator matrix. Their result addresses a somewhat different aspect of the problem by showing that exact-weight EA codes achieve a constant relative distance over arbitrary fields. However, their distance guarantee is more limited than ours, as their analysis only applies to relative distances below $1/2$. By contrast, we obtain an asymptotically optimal result by establishing that the rate–distance tradeoff of exact-weight EA codes approaches the Singleton bound over sufficiently large fields. Moreover, using a distinct approach, we show that these codes approach the GV bound over small fields, which also meets the best known results. Our analysis further yields explicit parameter choices and concrete distance guarantees for practical instantiations, which are not provided in their work. Finally, we develop an efficient code-switching protocol for EA codes, which leads to a PCS with $O(\lambda \log^2 M)$ proof size and verifier time.

2 Preliminaries

2.1 Notation

We use λ to denote the security parameter. We write $\mathbb{F}_q$ for the finite field of order q and $\mathbb{F}_q^\times := \mathbb{F}_q \setminus \{0\}$ for its multiplicative group of nonzero elements; when q is clear from the context, we simply write $\mathbb{F}$. Logarithms in base 2 and natural logarithms are denoted by log and ln, respectively. For a positive integer n, define $[n] := \{1, \ldots, n\}$. Unless otherwise specified, indices of vectors and matrices start from 1. For a vector x, we write $\mathsf{wt}(x)$ for its Hamming weight (i.e., the number of nonzero entries) and $\mathsf{supp}(x) := \{i : x_i \neq 0\}$ for its support. For vectors $x, y \in \mathbb{F}^n$, their relative distance is $\Delta(x, y) := \mathsf{wt}(x - y)/n$, and for a nonempty set $C \subseteq \mathbb{F}^n$ we define $\Delta(x, C) := \min_{y \in C} \Delta(x, y)$.

2.2 Multilinear Extension

For every $f : \{0,1\}^m \to \mathbb{F}$ there is a unique multilinear polynomial $\hat{f} : \mathbb{F}^m \to \mathbb{F}$ such that $\hat{f}(x) = f(x)$ for all $x \in \{0,1\}^m$. The polynomial $\hat{f}$ is called the *multilinear extension* (MLE) of f, and can be written as $\hat{f}(x) = \sum_{b \in \{0,1\}^m} f(b) \cdot \widetilde{eq}(x, b)$, where $\widetilde{eq}(x, b) := \prod_{i \in [m]} (x_i b_i + (1 - x_i)(1 - b_i))$. For a given $x \in \mathbb{F}^m$, the table $(\widetilde{eq}(x, b))_{b \in \{0,1\}^m}$ can be computed in $O(2^m)$ field operations [45, 46].

2.3 Interactive Proofs and IOPs

Interactive oracle proofs (IOPs) [9, 36] are a hybrid between interactive proofs and probabilistically checkable proofs. In this work we consider IOPs for promise

problems. Recall that a promise problem consists of two disjoint sets YES and NO. In an IOP for a promise problem $(\mathsf{YES}, \mathsf{NO})$, the goal is for the verifier to accept every $x \in \mathsf{YES}$ and reject every $x \in \mathsf{NO}$; there is no requirement for instances outside of $\mathsf{YES} \cup \mathsf{NO}$. We assume the verifier is given oracle access to the main input x, and we count the number of queries (in particular, this lets us capture the notion of IOP of proximity (IOPP) [36]). The prover, who is given x explicitly, is allowed to interact with the verifier over multiple rounds: in each round the prover sends an oracle proof string, and the verifier responds with a random string. For the sake of efficiency, it will also be convenient to allow the prover to send explicit (i.e., non-oracle) messages. At the end of the interaction, the verifier accepts or rejects. The protocol satisfies *completeness* if for every $x \in \mathsf{YES}$ the honest prover convinces the verifier with probability 1. The protocol has *soundness error* ϵ if for every $x \in \mathsf{NO}$ and every malicious prover P^*, the probability that the verifier accepts is at most ϵ. An IOP for a language L is an IOP for the promise problem $(\mathsf{YES}, \mathsf{NO})$ with $\mathsf{YES} = L$ and $\mathsf{NO} = \{0,1\}^* \setminus L$. Similarly, for $\delta > 0$, a δ-IOPP for a language L is an IOP for the promise problem $(\mathsf{YES}, \mathsf{NO})$ with $\mathsf{YES} = L$ and $\mathsf{NO} = \{x : \Delta(x, L) \geq \delta\}$.

Sumcheck protocol. We will use the celebrated sumcheck protocol [30], with a linear-time prover [44], as described next.

Theorem 2 (Sumcheck protocol). *Let $m, d \in \mathbb{N}$. There exists a protocol between a prover, who is given the description of an m-variate polynomial $f : \mathbb{F}^m \to \mathbb{F}$ of individual degree at most d, and a verifier, who gets as input $\alpha \in \mathbb{F}$ and either rejects or outputs a claim $(z, v) \in \mathbb{F}^m \times \mathbb{F}$. The protocol satisfies the following properties:*

- ***Complexity**: the protocol has m rounds; in each round the prover sends $d+1$ field elements and the verifier responds with a single field element. The verifier performs $O(d \cdot m)$ field operations. The prover runs in time $O((d+1)^m)$; if f is a product of d multilinear polynomials, the prover time decreases to $O(d \cdot 2^m)$.*
- ***Completeness**: if $\alpha = \sum_{x \in \{0,1\}^m} f(x)$, then, when interacting with the honest prover, the verifier always outputs (z, v) such that $f(z) = v$.*
- ***Soundness**: if $\alpha \neq \sum_{x \in \{0,1\}^m} f(x)$, then for any prover schemes, with probability at least $1 - \frac{dm}{|\mathbb{F}|}$, the verifier either rejects or outputs (z, v) such that $f(z) \neq v$.*

Batching polynomial evaluations. We will also use the protocol from [14], which reduces multiple evaluation claims to a single one.

Theorem 3. *Let $k \in \mathbb{N}$ and $t > 1$. There exists a protocol between a prover, who is given a k-variate multilinear polynomial f, and a verifier, who gets as input t evaluation claims $(z_1, v_1), \ldots, (z_t, v_t) \in \mathbb{F}^k \times \mathbb{F}$ and either rejects or outputs a claim $(r^*, v^*) \in \mathbb{F}^k \times \mathbb{F}$. The protocol satisfies the following properties:*

- ***Complexity**: the prover runs in time $O(t \cdot 2^k)$, the verifier runs in time $O(tk)$, and the communication is $O(k + \log t)$ field elements.*

– ***Completeness***: *if $f(z_i) = v_i$ for all $i \in [t]$, then, when interacting with the honest prover, the verifier always outputs (r^*, v^*) such that $f(r^*) = v^*$.*
– ***Soundness***: *if $f(z_i) \neq v_i$ for some $i \in [t]$, then, when interacting with any prover, with probability at least $1 - O\left(\frac{k+\log t}{|\mathbb{F}|}\right)$ the verifier either rejects or outputs (r^*, v^*) such that $f(r^*) \neq v^*$.*

2.4 Linear Codes

A linear error-correcting code with message length n and codeword length N over $\mathbb{F}$ is an n-dimensional linear subspace $C \subseteq \mathbb{F}^N$. Equivalently, there is a rank-n generator matrix $G \in \mathbb{F}^{n \times N}$ such that $C = \{xG : x \in \mathbb{F}^n\}$. The rate of the code is $R := n/N$ and its relative distance is $\mathsf{d}(C) := \min_{x,y \in C,\, x \neq y} \Delta(x, y)$. When C is generated by G, we also write $\mathsf{d}(G) := \mathsf{d}(C)$. The classical Singleton bound upper-bounds the achievable rate–distance tradeoffs.

Theorem 4 (Singleton bound). *Every linear code $C \subseteq \mathbb{F}^N$ of dimension $n \geq 1$ and relative distance δ satisfies $\delta \leq (N - n + 1)/N = 1 - R + 1/N$.*

Expand-accumulate codes. An expand-accumulate (EA) code [12] over $\mathbb{F}_q$ is generated by the product EA of two matrices: a sparse *expansion* matrix $E \in \mathbb{F}_q^{n \times N}$ and the *accumulator* $A \in \mathbb{F}_q^{N \times N}$, the upper-triangular matrix with $A_{j,i} = 1$ if and only if $j \leq i$. The encoding of a message $x \in \mathbb{F}_q^n$ is thus $\mathsf{EA}(x) := xEA$, obtained by expanding x into $y := xE$ and accumulating y into its prefix sums, $(yA)_j = \sum_{i \leq j} y_i$.

In this work, we consider the case in which the expansion matrix is sampled from the *exact-weight ensemble* $\mathsf{BP}(n, N, t, \mathbb{F}_q)$. Specifically, the n rows of E are independent, and each row is sampled uniformly from the set of vectors in $\mathbb{F}_q^N$ with Hamming weight exactly t. Equivalently, for each row, t distinct positions are chosen uniformly at random from $[N]$, and their entries are assigned independent uniformly random values from $\mathbb{F}_q^\times$.

Proximity gaps. We will make use of the proximity gap of linear codes with respect to random linear combinations [2, 10, 39]. Roughly speaking, given a set of vectors, either they are all simultaneously close to the code, or a random linear combination of these words is far from the code with high probability. In the first case, the vectors in fact share a common agreement set. Therefore, this property is now also referred to as *correlated agreement*.

Theorem 5 (Correlated agreement for linear codes [10, 13]). *Let $C \subseteq \mathbb{F}^N$ be a linear code, let $t \in \mathbb{N}$, and let $U \in \mathbb{F}^{N \times t}$. For all $\delta, \eta > 0$ with $\delta - \eta \leq \mathsf{d}(C)/3$, if $\Delta(U, C^t) \geq \eta$, then*

$$\Pr_{r \leftarrow \mathbb{F}^t}\left[\Delta(Ur, C) < \delta - \eta\right] \leq \frac{1}{\eta\,|\mathbb{F}|}.$$

3 EA Codes over Finite Fields

In this section we analyze the distance of EA codes over (large) finite fields. Recall first how the expansion matrix is sampled: under $\mathsf{BP}(n, N, t, \mathbb{F}_q)$, each

row of E independently chooses t distinct positions uniformly at random in $[N]$ and places a uniform value from $\mathbb{F}_q^\times$ at each of them. Now fix a message x of weight r. Its expansion xE is determined by the $M = rt$ active position–value pairs from the r rows selected by the support of x. Each active pair contributes its value, weighted by the corresponding coordinate of x, to its position in xE. Since xEA consists of the prefix sums of xE, each coordinate of it accumulates the contributions of all active pairs up to that coordinate. This gives rise to two distinct mechanisms by which the output xEA can have low weight: one geometric and the other algebraic.

Geometrically, the coordinates of xEA is kept constant over an *empty* interval that contains no active positions. If this value is zero, the entire interval contributes zero coordinates to the codeword. Thus, a low-weight codeword may arise from a few long empty intervals on which the starting value vanishes. To bound the probability of seeing long empty intervals, we need the joint distribution of the gaps between the active positions. Instead of studying the original distributions where the t positions are dependent with each other, we pass to an auxiliary model in which all M positions are sampled independently and uniformly. In this model, the probability of such large empty intervals can be bounded more easily, and the resulting bound transfers to the original ensemble at the cost of a mild conditioning factor.

Algebraically, a low-weight output can also arise from repeated cancellations that return the prefix sum to zero. Order the M active positions from left to right, and let the k-th *partial sum*, for $0 \le k \le M$, be the weighted sum of the first k values, with the 0-th partial sum equal to zero. Every coordinate of xEA equals one of these partial sums. If the k-th and k'-th partial sums vanish for some $k < k'$, then the weighted values at positions $k+1, \dots, k'$ must sum to zero. Hence, successive vanishing partial sums impose cancellation equations on pairwise disjoint blocks of values. Each such equation holds with probability at most $(q-1)^{-1}$, so the probability of many returns to zero decays rapidly with the number of required cancellations.

We prove our main theorem by combining the two mechanisms, treating low-weight and high-weight messages separately. For a low-weight message the number of active positions is small, so neither mechanism is by itself unlikely; instead, we exhibit a tradeoff between the two. If few partial sums vanish, a low-weight codeword requires unusually long empty intervals, whereas if many partial sums vanish, the resulting cancellation equations are collectively very unlikely over a large field. For a high-weight message the active positions are so numerous which makes long empty intervals already unlikely, and a low-weight codeword must instead create many disjoint cancellations. The main result of this section is summarized in the following theorem.

Theorem 6. *Let $0 < \delta < 1$ and $0 < R < 1-\delta$, and define $\varepsilon := 1-\delta-R$. Choose $\gamma > 0$ such that $\gamma \ln(1/\delta) > 1$, and set $n = \lfloor RN \rfloor$ and $t = \lceil \gamma \ln N \rceil$. Assume $\log q = \omega(\log n)$. Then for sufficiently large N, there exist constants $C_0, C_1 > 0$,*

such that

$$\Pr_{E\leftarrow\mathsf{BP}(n,N,t,\mathbb{F}_q)}[\mathsf{d}(EA)\le\delta]\le C_0N^{-\kappa}+\frac{C_0N(\ln N)^4}{q-1}+e^{-C_1N}+\exp\left(2N\ln 2-\frac{\varepsilon N}{4}\ln(q-1)\right),$$

where $\kappa=\frac{\gamma\ln(1/\delta)-1}{2}$. *Especially, this implies*

$$\Pr_{E\leftarrow\mathsf{BP}(n,N,t,\mathbb{F}_q)}[\mathsf{d}(EA)\le\delta]\le O(N^{-\kappa}).$$

Throughout the proof we write $D:=\lfloor\delta N\rfloor$, so that $\mathsf{wt}(xEA)\le\delta N$ if and only if $\mathsf{wt}(xEA)\le D$.

3.1 Analysis for Low-Weight Messages

In this part we bound the probability

$$\Pr_E\left[\exists\, x\in\mathbb{F}_q^n\setminus\{0^n\}:\ \mathsf{wt}(xEA)\le\delta N\ \wedge\ \mathsf{wt}(x)<cN/t\right], \tag{1}$$

where $c>0$ is a constant that will be chosen in later part.

Fix a message x of weight r and denote its support by $I\subseteq[n]$, $|I|=r$; put $M:=rt$. Since each row of E has exactly t nonzero entries, the rows indexed by I contribute M nonzero entries in total to the expansion xE, each specified by its *position* in $[N]$ and its *value* in $\mathbb{F}_q^\times$.

In the exact-weight ensemble, the t positions within a row are sampled without replacement. We first introduce an auxiliary *independent-position model*, in which each row samples t positions independently and uniformly on $[N]$, with each position carrying an independent uniform value in $\mathbb{F}_q^\times$. Let $\mathcal{E}_I$ be the event that, for every row indexed by I, the t sampled positions are distinct. Conditional on $\mathcal{E}_I$, the rows in I have exactly the exact-weight distribution. Hence, any bound proved in the independent-position model transfers to the original ensemble at the cost of the factor $\Pr(\mathcal{E}_I)^{-1}$.

Lemma 1. *If* $t<N$, *then* $\Pr(\mathcal{E}_I)^{-1}\le\exp\left(\frac{rt(t-1)}{2(N-t+1)}\right)$. *In particular,* $\Pr(\mathcal{E}_I)^{-1}=\exp(O(rt^2/N))$ *whenever* $t=O(\ln N)$.

Proof. For a single row, the probability that the t positions are pairwise distinct equals $N(N-1)\cdots(N-t+1)/N^t$. By the independence of the rows and the inequality $-\ln(1-u)\le u/(1-u)$ for $0\le u<1$,

$$-\ln\Pr(\mathcal{E}_I)=r\sum_{j=0}^{t-1}-\ln\left(1-\frac{j}{N}\right)\le r\sum_{j=0}^{t-1}\frac{j}{N-j}\le\frac{rt(t-1)}{2(N-t+1)}.$$

Let $X_1,\ldots,X_M\in[N]$ denote the positions in the independent-position model. For the spacing analysis, it is convenient to embed these discrete positions into a continuous interval, where collisions occur with probability zero. Let $U_1,\ldots,U_M$ be independent and uniform on $[0,1)$, and define $Y_a:=X_a-1+U_a$.

Then $Y_1, \ldots, Y_M$ are independent and uniform on $[0, N)$. Moreover, the embedding preserves the location information relative to the integer coordinates: for every $j \in [N]$ we have $X_a \le j$ if and only if $Y_a < j$. Let $0 < Y_{(1)} < \cdots < Y_{(M)} < N$ be their order statistics, and define the spacings

$$S_0 := Y_{(1)}, \qquad S_k := Y_{(k+1)} - Y_{(k)} \quad (1 \le k < M), \qquad S_M := N - Y_{(M)}.$$

Lemma 2. *The normalized spacings $\frac{1}{N}(S_0, \ldots, S_M)$ are uniform on the simplex $\Delta_M := \{(s_0, \ldots, s_M) : s_k \ge 0,\ \sum_{k=0}^{M} s_k = 1\}$.*

Proof. Write $V_a := Y_a/N$ for $1 \le a \le M$, so that $V_1, \ldots, V_M$ are independent and uniform on $[0, 1)$. Their order statistics $0 < V_{(1)} < \cdots < V_{(M)} < 1$ have joint density $M!$ on the ordered simplex $\{0 < v_1 < \cdots < v_M < 1\}$. This follows from the fact that the unordered variables have constant density one on $[0, 1]^M$, and each ordered configuration arises from $M!$ permutations.

Define the normalized spacings $Q_k := S_k/N$ for $k \in \{0, \ldots, M\}$. The map from the order statistics to the spacings is linear and invertible, with inverse $V_{(k)} = \sum_{j=0}^{k-1} Q_j$ for $1 \le k \le M$, and its Jacobian has absolute value one; it maps the ordered simplex onto Δ_M. Hence $(Q_0, \ldots, Q_M)$ has constant density $M!$ on Δ_M.

The spacings $S_0, \ldots, S_M$ represent, in the continuous model, the empty intervals between successive chosen positions. Since no new contribution is added within such an interval, the codeword is constant throughout it. Thus, a large total spacing on which this constant value is zero translates into a large fraction of zero coordinates in the codeword.

To describe the output algebraically, suppose that the k-th ordered position belongs to row $i(k)$ and carries value $\xi_k \in \mathbb{F}_q^\times$. Define the partial sums $Z_0 := 0$ and $Z_k := \sum_{a=1}^{k} x_{i(a)} \xi_a$ for $1 \le k \le M$. For $j \in [N]$, let $K(j) := \#\{a : X_a \le j\}$ be the number of positions encountered up to coordinate j. Since Z_k is the accumulated contribution of the first k ordered positions, the output at coordinate j equals the corresponding partial sum:

$$(xEA)_j = Z_{K(j)}. \tag{2}$$

For $0 \le k \le M$, put $G_k := \#\{j \in [N] : K(j) = k\}$, so that G_k is the number of coordinates on which the output equals Z_k. In particular, $\sum_{k=0}^{M} G_k = N$.

These definitions separate the two ingredients of a low-weight codeword. The quantities G_k measure how many zero coordinates can be supplied by empty space once the partial sum Z_k vanishes, whereas the indices k with $Z_k = 0$ record which partial sums vanish. We now show how these ingredients interact with each other.

Lemma 3. *If $\mathsf{wt}(xEA) \le D$, then there exist an integer $0 \le m \le M$ and a set $\mathcal{K} \subseteq [M]$ with $|\mathcal{K}| = m$ such that*

$$S_0 + \sum_{k \in \mathcal{K}} S_k \ge N - D - m, \quad \text{and} \quad Z_k = 0, \forall k \in \mathcal{K}. \tag{3}$$

Moreover, writing $\mathcal{K} = \{k_1 < \cdots < k_m\}$ *and* $k_0 := 0$*, the* m *pairwise disjoint nonempty index blocks* $B_\ell := \{k_{\ell-1}+1, \ldots, k_\ell\}$*,* $1 \le \ell \le m$*, satisfy* $\sum_{a \in B_\ell} x_{i(a)} \xi_a = 0$ *for every* $1 \le \ell \le m$*.*

Proof. By Eq. 2, $\sum_{k:Z_k=0} G_k = \#\{j : (xEA)_j = 0\} \ge N - D$. Let $\mathcal{K} := \{k \in [M] : Z_k = 0 \text{ and } G_k > 0\}$ and $m := |\mathcal{K}|$. Since $Z_0 = 0$, it follows that $G_0 + \sum_{k \in \mathcal{K}} G_k \ge N - D$. If $G_0 > 0$, the first occupied coordinate is $G_0 + 1$, and hence $S_0 \ge G_0$. If $k \ge 1$ and $G_k > 0$, the G_k integer coordinates with $K(j) = k$ lie between the k-th and the $(k+1)$-st ordered positions (or after the last one when $k = M$). Passing from the number of integer coordinates in this range to the continuous spacing loses at most one unit at the boundary, so $S_k \ge G_k - 1$. Therefore, $S_0 + \sum_{k \in \mathcal{K}} S_k \ge G_0 + \sum_{k \in \mathcal{K}} (G_k - 1) \ge N - D - m$.

For the algebraic part, the partial sums at both endpoints of each block vanish, so $Z_{k_\ell} - Z_{k_{\ell-1}} = 0$, and the left-hand side is exactly the sum $\sum_{a \in B_\ell} x_{i(a)} \xi_a$ of the weighted values at the $(k_{\ell-1}+1)$-st through k_ℓ-th ordered positions. The blocks are pairwise disjoint because they are consecutive intervals in the ordering, and each is nonempty because $k_\ell > k_{\ell-1}$.

The next lemma bounds the probability of simultaneous cancellations of all disjoint block equations.

Lemma 4. *Let* $(\xi_a)_{a \in \mathcal{A}}$ *be independent and uniform on* $\mathbb{F}_q^\times$*, let* $B_1, \ldots, B_g \subseteq \mathcal{A}$ *be pairwise disjoint nonempty sets, and let* $c_a \in \mathbb{F}_q^\times$ *be fixed. Then*

$$\Pr\left[\sum_{a \in B_\ell} c_a \xi_a = 0 \text{ for every } 1 \le \ell \le g\right] \le (q-1)^{-g}.$$

Proof. Select one index from each block and condition on all the remaining variables. Each block equation then determines its selected variable uniquely, if it can be satisfied at all. The selected variables are independent and uniform on a set of size $q - 1$, which gives the product bound.

We next control the geometric part of Lemma 3. For $0 \le m \le M$, set $\alpha_m := \min\{1, (D+m)/N\}$. The spacing condition in Eq. 3 requires $m+1$ selected spacings to occupy at least $N - D - m$ of the total length, or equivalently, the remaining $M - m$ spacings to occupy at most $D + m$. The following lemma bounds the probability of this event.

Lemma 5. *For every fixed* $\mathcal{K} \subseteq [M]$ *of size* m*,*

$$\Pr\left[S_0 + \sum_{k \in \mathcal{K}} S_k \ge N - D - m\right] \le \binom{M}{m} \alpha_m^{M-m},$$

where the right-hand side is interpreted as one when $m = M$*.*

Proof. If $m = M$, there is nothing to prove, so assume $m < M$ and put $r := M - m$. Let $\Sigma := \sum_{k \notin \{0\} \cup \mathcal{K}} S_k$, which is the sum of r spacings. By Lemma 2, the normalized spacing vector is uniform on Δ_M.

By symmetry, relabel the coordinates so that the r spacings appearing in Σ/N are the first r coordinates. Write these coordinates as $y_1, \dots, y_r$, and write m of the remaining coordinates as $z_1, \dots, z_m$, with the last coordinate determined by the condition that all $M+1$ coordinates sum to one. Since Δ_M has volume $1/M!$, we have

$$\begin{aligned}
\Pr\left[\frac{\Sigma}{N} \le \alpha_m\right] &= M! \int_{\substack{y_i \ge 0 \\ \sum_{i=1}^r y_i \le \alpha_m}} \int_{\substack{z_j \ge 0 \\ \sum_{j=1}^m z_j \le 1 - \sum_{i=1}^r y_i}} dz\, dy \\
&= \frac{M!}{m!} \int_{\substack{y_i \ge 0 \\ \sum_{i=1}^r y_i \le \alpha_m}} \left(1 - \sum_{i=1}^r y_i\right)^m dy \\
&\le \frac{M!}{m!} \operatorname{Vol}\left\{ y \in \mathbb{R}^r_{\ge 0} : \sum_{i=1}^r y_i \le \alpha_m \right\} \\
&= \frac{M!}{m!\, r!} \alpha_m^r = \binom{M}{m} \alpha_m^{M-m}.
\end{aligned}$$

The event in the statement implies $\Sigma/N \le \alpha_m$, and the claim follows.

Proposition 1. *Fix $I \subseteq [n]$ with $|I| = r$, put $M = rt$, and assume $t < N$. We have,*

$$\begin{aligned}
&\Pr_{E \leftarrow \mathsf{BP}(n,N,t,\mathbb{F}_q)} \left[\exists\, x \in \mathbb{F}_q^n :\ \mathsf{supp}(x) = I,\ \mathsf{wt}(xEA) \le D\right] \\
&\qquad \le \exp\left(\frac{rt(t-1)}{2(N-t+1)}\right) \sum_{m=0}^{M} \binom{M}{m}^2 \alpha_m^{M-m} \min\{1, (q-1)^{r-1-m}\}. \qquad (4)
\end{aligned}$$

Proof. We first work in the independent-position model. By Lemma 3, if a message x with $\mathsf{supp}(x) = I$ satisfies $\mathsf{wt}(xEA) \le D$, then for some $0 \le m \le M$ there exists a set $\mathcal{K} \subseteq [M]$ of size m such that the spacing condition in Eq. 3 holds and the corresponding m disjoint nontrivial block equations are satisfied.

Fix m and $\mathcal{K}$. The spacing event depends only on the sampled positions. Conditional on these positions, the m block equations are determined, and Lemma 4 shows that a fixed message satisfies them with probability at most $(q-1)^{-m}$. Since $\mathsf{wt}((\lambda x)EA) = \mathsf{wt}(xEA)$ for every $\lambda \in \mathbb{F}_q^\times$, it suffices to take a union bound over the $(q-1)^{r-1}$ projective classes of messages with support I. Thus, for the fixed $\mathcal{K}$, the algebraic contribution is at most $\min\{1, (q-1)^{r-1-m}\}$.

On the other hand, Lemma 5 bounds the probability of the spacing event by $\binom{M}{m}\alpha_m^{M-m}$. Therefore, for fixed m and $\mathcal{K}$, the probability that both the geometric and algebraic conditions hold for some message supported on I is at most

$$\binom{M}{m} \alpha_m^{M-m} \min\{1, (q-1)^{r-1-m}\}.$$

There are $\binom{M}{m}$ choices of $\mathcal{K}$. Taking a union bound over $\mathcal{K}$ and then summing over $0 \le m \le M$ gives, in the independent-position model,

$$\Pr\left[\exists\, x :\ \mathsf{supp}(x) = I,\ \mathsf{wt}(xEA) \le D\right] \le \sum_{m=0}^{M} \binom{M}{m}^2 \alpha_m^{M-m} \min\{1, (q-1)^{r-1-m}\}.$$

Finally, Lemma 1 transfers this bound to the exact-weight ensemble at the cost of the factor $\exp\left(\frac{rt(t-1)}{2(N-t+1)}\right)$, which yields Eq. 4.

It remains to sum Eq. 4 over all supports of low weight. For $1 \le r < cN/t$ and $0 \le m \le rt$, define

$$\tau_{r,m} := \binom{n}{r} \exp\left(\frac{rt(t-1)}{2(N-t+1)}\right) \binom{rt}{m}^2 \alpha_m^{rt-m} \min\{1, (q-1)^{r-1-m}\}.$$

By Proposition 1 and a union bound over the $\binom{n}{r}$ supports of each weight r, the probability Eq. 1 is at most $\sum_{r<cN/t} \sum_{m=0}^{rt} \tau_{r,m}$. We split the inner sum at $m = 2r$, corresponding to the two mechanism that can produce a low-weight code word. When $m < 2r$, few partial sums vanish, so most of the $N - D$ zero coordinates must come from long empty intervals. When $m \ge 2r$, there are sufficiently many disjoint cancellation equations that the probability of satisfying them simultaneously dominates the combinatorial counting factors.

Consider first the range $m < 2r$. Here we discard the projective factor, that is, we bound $\min\{1, (q-1)^{r-1-m}\} \le 1$, and estimate the remaining factors of $\tau_{r,m}$ in logarithmic form. First, $\ln \binom{n}{r} \le r \ln(eN/r) \le r \ln N + O(r)$. Second, $\frac{rt(t-1)}{2(N-t+1)} = O(rt^2/N) = o(r)$, because $t = O(\ln N)$. Third, for $1 \le m < 2r$ we have $\ln \binom{rt}{m} \le m \ln \frac{ert}{m} \le 2r \ln(et) = O(r \ln t)$, the middle expression being increasing in m on $[1, rt]$, and the case $m = 0$ is trivial; hence $\ln \binom{rt}{m}^2 = O(r \ln t)$. Fourth, since $D \le \delta N$ and $m < 2r < 2cN/t$, we have $\alpha_m \le (D+m)/N \le \delta + 2c/t$, which is smaller than one for all sufficiently large N; using $\ln(\delta + 2c/t) = \ln \delta + \ln(1 + \frac{2c}{\delta t})$, the bound $\ln(1+u) \le u$, and $m < 2r$, we obtain

$$\begin{aligned}
\ln \alpha_m^{rt-m} &\le (rt - m) \ln\left(\delta + \frac{2c}{t}\right) = (rt - m) \ln \delta + (rt - m) \ln\left(1 + \frac{2c}{\delta t}\right) \\
&\le rt \ln \delta + m \ln \frac{1}{\delta} + rt \cdot \frac{2c}{\delta t} \\
&\le -rt \ln \frac{1}{\delta} + 2r \ln \frac{1}{\delta} + \frac{2c}{\delta} r = -rt \ln \frac{1}{\delta} + O(r).
\end{aligned}$$

Collecting the four estimates and using $t \ge \gamma \ln N$,

$$\ln \tau_{r,m} \le r \ln N - rt \ln \frac{1}{\delta} + O(r \ln t) \le -\left(\gamma \ln \frac{1}{\delta} - 1\right) r \ln N + O(r \ln \ln N). \tag{5}$$

Fix any constant $0 < \kappa < \gamma \ln(1/\delta) - 1$; such a choice is possible because $\gamma \ln(1/\delta) > 1$. The implied constants in Eq. 5 are uniform over $1 \le r < cN/t$ and $0 \le m < 2r$, so for all sufficiently large N we have $\ln \tau_{r,m} \le -\kappa r \ln N$, that is, $\tau_{r,m} \le N^{-\kappa r}$, on the entire range. Since the inner range contains at most $2r$ values of m,

$$\sum_{r<cN/t} \sum_{m<2r} \tau_{r,m} \le \sum_{r \ge 1} 2r\, N^{-\kappa r} = \frac{2N^{-\kappa}}{(1 - N^{-\kappa})^2} \le C_0 N^{-\kappa} \tag{6}$$

for all sufficiently large N.

Now consider the range $m \ge 2r$, which is nonempty only when $t \ge 2$. We bound each factor of $\tau_{r,m}$ by an m-th power. Since $m \ge 2r$ implies $r - 1 -$

$m \le -m/2$, the projective factor satisfies $\min\{1, (q-1)^{r-1-m}\} \le (q-1)^{-m/2}$. Since $r \le m/2$, we have $\binom{n}{r} \le (eN)^r \le (eN)^{m/2}$. Since $rt/m \le t/2$, we have $\binom{rt}{m} \le (\frac{ert}{m})^m \le (\frac{et}{2})^m$, and hence $\binom{rt}{m}^2 \le (\frac{et}{2})^{2m}$. The conditioning factor obeys $\frac{rt(t-1)}{2(N-t+1)} \le \frac{mt^2}{4(N-t+1)} \le m$ for all sufficiently large N, because $t^2 = o(N)$, so it is at most e^m. Finally, $\alpha_m^{rt-m} \le 1$. Multiplying the five bounds,

$$\tau_{r,m} \le (eN)^{m/2} \left(\frac{et}{2}\right)^{2m} e^m \, (q-1)^{-m/2} = \theta^m, \qquad \theta := \frac{e^{7/2}\, t^2}{4} \sqrt{\frac{N}{q-1}}.$$

Because $t = O(\ln N)$ and $\log q = \omega(\log N)$, we have $\theta \to 0$ as $N \to \infty$; in particular, $\theta \le 1/2$ for all sufficiently large N. Summing the geometric series first over m and then over r,

$$\sum_{r \ge 1} \sum_{m \ge 2r} \theta^m = \sum_{r \ge 1} \frac{\theta^{2r}}{1-\theta} = \frac{1}{1-\theta} \cdot \frac{\theta^2}{1-\theta^2} \le \frac{\theta^2}{(1/2)(3/4)} = \frac{8}{3}\theta^2.$$

Since $\theta^2 = \frac{e^7}{16} \cdot \frac{t^4 N}{q-1}$ and $t = \lceil \gamma \ln N \rceil$ gives $t^4 = O((\ln N)^4)$, we conclude that

$$\sum_{r < cN/t} \sum_{m \ge 2r} \tau_{r,m} \le \frac{8}{3}\theta^2 \le \frac{C_0 N (\ln N)^4}{q-1}. \tag{7}$$

Combining Eq. 6 and Eq. 7, we conclude that for all sufficiently large N the probability Eq. 1 is at most $C_0 N^{-\kappa} + C_0 N (\ln N)^4/(q-1)$.

3.2 Analysis for High-Weight Messages

For a message of weight $r \ge cN/t$, there at least cN active positions. At this scale, the geometric mechanism can be ruled out first: with high probability, every such set of rows occupies a large number of distinct coordinates. Once this support-expansion event holds, a low-weight output can no longer be explained by empty space and must instead create many disjoint cancellation equations.

For $I \subseteq [n]$, let $U_I := \bigcup_{i \in I} \mathsf{supp}(E_i) \subseteq [N]$ denote the union of the supports of the rows indexed by I, and set $\eta := \varepsilon/8$. We first show that, for a sufficiently large constant c, all the unions U_I with $|I| \ge cN/t$ are simultaneously large with overwhelming probability.

Lemma 6. *There exists a constant $c = c(\delta, \varepsilon) > 0$ such that, with probability $1 - e^{-\Omega(N)}$, every $I \subseteq [n]$ with $r := |I| \ge cN/t$ satisfies*

$$|U_I| \ge D + r + 2\eta N. \tag{8}$$

Proof. Since $n \le RN$ and $R = 1 - \delta - \varepsilon$, for every $r \le n$ we have

$$\frac{D + r + 2\eta N}{N} \le \delta + R + \frac{\varepsilon}{4} = 1 - \frac{3\varepsilon}{4}.$$

Hence, setting $s_r := \lceil D + r + 2\eta N \rceil$, we have $s_r/N \le \alpha := 1 - \varepsilon/8 < 1$ for all sufficiently large N.

If Eq. 8 fails for some I of size r, then U_I is contained in some set $W \subseteq [N]$ of size s_r. For fixed I and W, the independence of the row supports gives

$$\Pr[U_I \subseteq W] = \left(\frac{\binom{s_r}{t}}{\binom{N}{t}}\right)^r \le \left(\frac{s_r}{N}\right)^{rt} \le \alpha^{rt} \le \alpha^{cN},$$

where the last step uses $rt \ge cN$. A union bound over the at most N values of r, the at most 2^n sets I, and the at most 2^N sets W yields

$$\Pr[\text{some } I \text{ violates Eq. 8}] \le \sum_{r=\lceil cN/t \rceil}^{n} \binom{n}{r}\binom{N}{s_r}\alpha^{rt} \le N\, 2^n 2^N \alpha^{cN} \le N\, 4^N \alpha^{cN},$$

which is $e^{-\Omega(N)}$ once c is chosen so large that $4\alpha^c < 1$.

Assume henceforth that the event in Lemma 6 holds. Fix $I \subseteq [n]$ with $|I| = r \ge cN/t$ and a message x with $\mathsf{supp}(x) = I$. Let $S := \mathsf{supp}(xEA)$, and $T := [N] \setminus S$. Suppose that $|S| \le D$, and write $T = \{j_1 < \cdots < j_\ell\}$. Partition the coordinates into the dishoint blocks

$$B_1 := [1, j_1], \qquad B_k := [j_{k-1}+1, j_k], 2 \le k \le \ell.$$

Since $(xEA)_{j_k} = 0$ for every $j_k \in T$, subtracting consecutive prefix-sum equations shows that the weighted sum of the active contributions within each B_k sum to zero. Call B_k *touched* if $B_k \cap U_I \ne \emptyset$. Each touched block therefore gives a nontrivial cancellation equation, while an untouched block contains no active position and contributes only the trivial identity $0 = 0$. Moreover, since the blocks are disjoint, the cancellation equations arising from touched blocks involve pairwise disjoint sets of random values.

Every coordinate in $U_I \cap T$ is the right endpoint of a distinct touched block. Hence, if g denote the number of touched blocks, then

$$g \ge |U_I \cap T| \ge |U_I| - |S| \ge r + 2\eta N. \tag{9}$$

Thus, the D term in the expansion bound Eq. 8 accounts for the coordinates on which the output may be nonzero, while the remaining $r + 2\eta N$ occupied coordinates force distinct nontrivial cancellation equations.

For a fixed projective message, the touched blocks are disjoint, so Lemma 4 bounds the probability that all g cancellation equations hold by $(q-1)^{-g}$. Taking a union bound over the $(q-1)^{r-1}$ projective messages supported on I and using Eq. 9, we obtain

$$\Pr\left[\exists\, x : \mathsf{supp}(x) = I \ \wedge\ z_j = 0, \forall j \in T \mid \{\mathsf{supp}(E_i)\}_{i=1}^n\right] \le (q-1)^{-2\eta N - 1}. \tag{10}$$

Finally, we sum Eq. 10 over all row sets I and all output supports S of size at most D, using the crude estimates $|\{I \subseteq [n]\}| \le 2^n$ and $|\{S \subseteq [N] : |S| \le D\}| \le 2^N$. Together with the failure probability of Lemma 6, this yields

$$\Pr\left[\exists\, x :\ \mathsf{wt}(x) \ge cN/t,\ \mathsf{wt}(xEA) \le D\right] \le e^{-\Omega(N)} + 4^N (q-1)^{-2\eta N - 1}. \tag{11}$$

We now combine the preceding bounds for low- and high-weight messages to complete the proof of the main theorem.

Proof (Proof of Theorem 6). The event $\mathsf{d}(EA) \le \delta$ means that some $x \in \mathbb{F}_q^n \setminus \{0^n\}$ satisfies $\mathsf{wt}(xEA) \le \delta N$, and hence $\mathsf{wt}(xEA) \le D$. Splitting according to

whether $\mathsf{wt}(x) < cN/t$ or $\mathsf{wt}(x) \geq cN/t$, and applying Eq. 6, Eq. 7, and Eq. 11, we obtain, for all sufficiently large N,

$$\Pr\left[\mathsf{d}(EA) \leq \delta\right] \leq C_0 N^{-\kappa} + \frac{C_0 N (\ln N)^4}{q-1} + e^{-\Omega(N)} + 4^N (q-1)^{-2\eta N - 1}.$$

The third term is at most $e^{-C_1 N}$ for a constant $C_1 > 0$ depending only on δ and ε, and the fourth term is at most $\exp(2N \ln 2 - 2\eta N \ln(q-1))$ with $2\eta = \varepsilon/4$, which is the last term in the theorem. Finally, when $\log q = \omega(\log n)$, we have $\frac{C_0 N (\ln N)^4}{q-1} = N^{-\omega(1)}$. Thus this term, as well as the two exponential terms, is negligible compared with $N^{-\kappa}$. Consequently, $\Pr\left[\mathsf{d}(EA) \leq \delta\right] = O(N^{-\kappa})$.

Remark 1. The result above requires the field size to be significantly larger than the message size—in fact, super-polynomially larger. This requirement is not restrictive in typical SNARK applications. Concretely, to achieve negligible soundness error, one usually works over a field of size $2^{\Theta(\lambda)}$, where the security parameter λ satisfies $\lambda = \omega(\log N)$ for instance size N. Consequently, the field size is $2^{\Theta(\lambda)} = 2^{\omega(\log N)}$, which precisely meets the requirement of the theorem.

4 Polynomial Commitment from EA Codes

We now construct multilinear polynomial commitments from EA codes.

Throughout this section, we identify $[n]$ with $\{0,1\}^{\log n}$ by mapping each $i \in [n]$ to the $\log n$-bit binary representation of $i-1$, and conversely. For a vector $w \in \mathbb{F}^N$ and a matrix $E \in \mathbb{F}^{n \times N}$, we henceforth identify w and E with their respective multilinear extensions $\hat{w}$ and $\hat{E}$, satisfying $\hat{w}(i) = w_i$ for every $i \in \{0,1\}^{\log N}$ and $\hat{E}(b,j) = E_{b,j}$ for every $(b,j) \in \{0,1\}^{\log n} \times \{0,1\}^{\log N}$.

Following [4,13,19], we consider the constrained relation that captures proximity to a codeword whose preimage satisfies a multilinear evaluation claim.

Definition 1. *For a linear code $\mathcal{C} : \mathbb{F}^n \to \mathbb{F}^N$, the* constrained relation *of $\mathcal{C}$ is*

$$\mathcal{R}_{\mathcal{C}} = \left\{((z,v), w, m) \in (\mathbb{F}^{\log n} \times \mathbb{F}) \times \mathbb{F}^N \times \mathbb{F}^n : w = \mathcal{C}(m) \wedge \hat{m}(z) = v\right\},$$

where (z,v) is the explicit input, w is the implicit input, and m is the witness.

For such relation and a proximity parameter $\delta > 0$, we consider the associated promise problem $(\mathsf{YES}, \mathsf{NO})$. An IOPP for $\mathcal{R}_{\mathcal{C}}$ is therefore an IOP for distinguishing YES instances from NO instances. For simplicity, we focus on the unique-decoding regime. The YES instances are the well-formed inputs,

$$\mathsf{YES} = \left\{(z, v, \mathcal{C}(m)) \in \mathbb{F}^{\log n} \times \mathbb{F} \times \mathbb{F}^N : m \in \mathbb{F}^n \wedge \hat{m}(z) = v\right\},$$

and the NO instances are the tuples $(z, v, y) \in \mathbb{F}^{\log n} \times \mathbb{F} \times \mathbb{F}^N$ such that either:

- $\Delta(y, \mathcal{C}) \geq \delta$; or
- there exists a unique codeword $\mathcal{C}(u)$ such that $\Delta(y, \mathcal{C}(u)) < \delta$, but $\hat{u}(z) \neq v$.

4.1 IOPP for EA Codes

In this part, we construct an IOPP for the constraint relation of EA codes and extend it to interleaved EA codes in next part.

Let $\mathsf{EA} : \mathbb{F}^n \to \mathbb{F}^N$ be an EA code with relative distance δ_{EA}. Recall that $\mathsf{EA}(x) = xEA$, where $E \in \mathbb{F}^{n \times N}$ is the sparse expansion matrix with $O(n \log n)$ non-zero entries and $A \in \mathbb{F}^{N \times N}$ is the accumulator matrix, i.e., $A(j, i) = 1$ if and only if $j \leq i$. Let $\mathcal{C}' : \mathbb{F}^n \to \mathbb{F}^{N'}$ be an auxiliary linear code with the same message length. We assume that $\mathcal{C}'$ has an IOPP for its constrained relation.

The protocol starts by letting the prover send an oracle $[\![\mathcal{C}'(x)]\!]$ claimed to contain the encoding of the witness x under $\mathcal{C}'$. The problem is that a malicious prover may send the $\mathcal{C}'$-encoding of an arbitrary message, so we must enforce consistency of the new oracle with the input word w, which is allegedly $\mathsf{EA}(x)$.

Consistency is enforced by letting the verifier sample ν uniformly random entries of the input oracle $[\![w]\!]$ and check that each of them agrees with the corresponding entry of the EA codeword $\mathsf{EA}(x)$. If w is η-far from $\mathsf{EA}(x)$, at least one sampled entry exposes a disagreement, except with probability $(1-\eta)^\nu$. Since the entries of $\mathsf{EA}(x)$ are not directly available to the verifier, each check is phrased as a sumcheck claim using the following identity:

$$\hat{w}(i_k) = \sum_{b \in \{0,1\}^{\log n}} \sum_{j \in \{0,1\}^{\log N}} \hat{x}(b) \cdot \hat{E}(b, j) \cdot \hat{A}(j, i_k), \quad \forall k \in [\nu]. \tag{12}$$

While one could invoke the batching protocol of Theorem 3 to reduce the ν claims to a single random claim $\hat{w}(r)$, we instead incorporate the batching procedure directly into the sumcheck, thereby avoiding the cost of invoking the batching protocol separately. Concretely, the verifier samples a random challenge θ, and both parties define $\sigma := \sum_{k=1}^{\nu} \theta^{k-1} \cdot \hat{w}(i_k)$ and $q := \sum_{k=1}^{\nu} \theta^{k-1} \cdot e_{i_k} \in \mathbb{F}^N$, where $e_i \in \mathbb{F}^N$ denotes the standard basis vector which has a single 1 on position i. The ν instances of Eq. 12 are then batched into the single claim

$$\sigma = \sum_{b \in \{0,1\}^{\log n} j, i \in \{0,1\}^{\log N}} \hat{x}(b) \cdot \hat{E}(b, j) \cdot \hat{A}(j, i) \cdot \hat{q}(i).$$

We start by performing the sumcheck on variable i. Let $u := xE \in \mathbb{F}^N$ and let $A_q := Aq \in \mathbb{F}^N$. The parties then run a sumcheck for

$$\sigma = \sum_{j \in \{0,1\}^{\log N}} \hat{u}(j) \cdot \widehat{A_q}(j),$$

which reduces the claim to the two evaluations $\hat{u}(r_1)$ and $\widehat{A_q}(r_1)$ at a random point $r_1 \in \mathbb{F}^{\log N}$. The verifier evaluates $\widehat{A_q}(r_1)$ on its own (looking forward, this only requires $O(\nu \log N)$ time). The claim on $\hat{u}(r_1)$ is then handled by a second sumcheck, over the expansion layer,

$$\hat{u}(r_1) = \sum_{b \in \{0,1\}^{\log n}} \hat{x}(b) \cdot \hat{E}(b, r_1),$$

which in turn reduces to evaluation claims on $\hat{x}(r_2)$ and $\hat{E}(r_2, r_1)$ for a random $r_2 \in \mathbb{F}^{\log n}$. For the latter, since E is a fixed public matrix, we assume that

there exists a protocol for verifying evaluation claims of the form $\hat{E}(\rho) = \tau$ with $\rho \in \mathbb{F}^{\log n + \log N}$, and the parties invoke it once, on the claimed value of $\hat{E}(r_2, r_1)$. Finally, the resulting claim on $\hat{x}(r_2)$ and the original claim $\hat{x}(z) = v$ are batched into a single evaluation claim using Theorem 3, which is exactly an instance of the constrained relation of $\mathcal{C}'$ and is proved by the IOPP for $\mathcal{C}'$ relative to the oracle $[\![\mathcal{C}'(x)]\!]$. The complete protocol is presented in Protocol 1.

PROTOCOL 1 *IOPP for EA codes.*

The prover claims $((z, v), w, x) \in \mathcal{R}_{\mathsf{EA}}$ to the verifier. The verifier holds the explicit input $(z, v) \in \mathbb{F}^{\log n} \times \mathbb{F}$ and has oracle access $[\![w]\!]$ to the implicit input $w \in \mathbb{F}^N$; the prover additionally holds the witness $x \in \mathbb{F}^n$.

1. The prover computes $\mathcal{C}'(x) \in \mathbb{F}^{N'}$ and sends the oracle $[\![\mathcal{C}'(x)]\!]$.
2. The verifier samples $\nu := \lceil -\lambda / \log_2(1-\eta) \rceil$ indices $i_1, \ldots, i_\nu$ independently and uniformly from $\{0,1\}^{\log N}$ and sends them to the prover; it queries $[\![w]\!]$ at these positions, obtaining the entries $\hat{w}(i_k)$ for $k \in [\nu]$.
3. The verifier samples $\theta \leftarrow \mathbb{F}$ and sends it to the prover. Both parties define $\sigma := \sum_{k=1}^{\nu} \theta^{k-1} \cdot \hat{w}(i_k)$ and $q := \sum_{k=1}^{\nu} \theta^{k-1} \cdot e_{i_k} \in \mathbb{F}^N$, where $e_i \in \mathbb{F}^N$ is the standard basis vector supported on position i.
4. Let $\hat{u}$ and $\widehat{A_q}$ be the MLEs of the vectors $u := xE$ and $A_q := Aq$, respectively. The parties run the sumcheck protocol for the claim
$$\sigma = \sum_{j \in \{0,1\}^{\log N}} \hat{u}(j) \cdot \widehat{A_q}(j). \tag{13}$$
Let $r_1 \in \mathbb{F}^{\log N}$ denote the random challenges sampled during the sumcheck. The prover sends τ_u, claimed to equal $\hat{u}(r_1)$; the verifier computes $\widehat{A_q}(r_1)$ itself and performs the final sumcheck check against $\tau_u \cdot \widehat{A_q}(r_1)$.
5. The parties run the sumcheck protocol for the claim
$$\tau_u = \sum_{b \in \{0,1\}^{\log n}} \hat{x}(b) \cdot \hat{E}(b, r_1). \tag{14}$$
Let $r_2 \in \mathbb{F}^{\log n}$ denote the random challenges sampled during the sumcheck. The prover sends τ_x and τ_E, claimed to equal $\hat{x}(r_2)$ and $\hat{E}(r_2, r_1)$, respectively. The parties run the evaluation protocol for $\hat{E}$ to verify the claim $\hat{E}(r_2, r_1) = \tau_E$. The verifier performs the final sumcheck check against $\tau_x \cdot \tau_E$.
6. The parties run the protocol of Theorem 3 to reduce the two evaluation claims $\hat{x}(z) = v$ and $\hat{x}(r_2) = \tau_x$ to a single claim $\hat{x}(r^*) = v^*$.
7. The parties run the IOP for the constrained relation of $\mathcal{C}'$ on the explicit input (r^*, v^*), the implicit input $[\![\mathcal{C}'(x)]\!]$ and the witness x.
8. The verifier accepts if all of the above checks accept; otherwise it rejects.

Theorem 7. *Let* $\mathsf{EA} : \mathbb{F}^n \to \mathbb{F}^N$ *be an EA code with relative distance* δ_{EA}*, and let* $\mathcal{C}' : \mathbb{F}^n \to \mathbb{F}^{N'}$ *be a linear code with the same message length. For* $\eta \in (0, \delta_{\mathsf{EA}}/2)$*, consider the promise problem* $(\mathsf{YES}, \mathsf{NO})$ *of* EA *with proximity parameter* η*. Assume* $\lambda = \omega(\log n)$ *and* $|\mathbb{F}| = 2^{\Theta(\lambda)}$*, and*

1. *an IOP for evaluation claims about the expansion matrix, i.e., an protocol that, on input $\rho \in \mathbb{F}^{\log n + \log N}$ and $\tau \in \mathbb{F}$, verifies that $\hat{E}(\rho) = \tau$, with soundness error ϵ_{sp}, proof query complexity q_{sp}, non-oracle communication C_{sp}, prover time $T_{P,\mathrm{sp}}$, and verifier time $T_{V,\mathrm{sp}}$, where the complexities are viewed as functions of the number of nonzero entries of E; and*
2. *an IOPP for the constrained relation of $\mathcal{C}'$, with soundness error $\epsilon_{\mathcal{C}'}$, proof query complexity $q_{\mathcal{C}'}$, non-oracle communication $C_{\mathcal{C}'}$, prover time $T_{P,\mathcal{C}'}$ (including the encoding of the message), and verifier time $T_{V,\mathcal{C}'}$, where the complexities are viewed as functions of the message length.*

Then Protocol 1 is an IOP for $(\mathsf{YES}, \mathsf{NO})$ *with the following parameters:*

- *Soundness error:* $O(2^{-\lambda}) + \epsilon_{\mathrm{sp}} + \epsilon_{\mathcal{C}'}$;
- *Input query complexity:* $O(\lambda)$;
- *Proof query complexity:* $q_{\mathcal{C}'}(n) + q_{\mathrm{sp}}(n \log n)$;
- *Non-oracle communication:* $O(\log n) + C_{\mathcal{C}'}(n) + C_{\mathrm{sp}}(n \log n)$ *field elements;*
- *Prover time:* $O(n \log n) + T_{P,\mathcal{C}'}(n) + T_{P,\mathrm{sp}}(n \log n)$;
- *Verifier time:* $O(\lambda \log n) + T_{V,\mathcal{C}'}(n) + T_{V,\mathrm{sp}}(n \log n)$.

Proof. Completeness. Suppose $w = \mathsf{EA}(x)$ and $\hat{x}(z) = v$. By Step 3, $\sigma == \sum_{j \in \{0,1\}^{\log N}} \hat{u}(j)\, \widehat{A_q}(j)$, where $w = uA$ and $u = xE$. Thus, Eq. 13 is a true claim. The honest prover's message satisfies $\tau_u = \hat{u}(r_1)$, so the final check of Step 4 passes. Since $u = xE$, we have $\hat{u}(r_1) = \sum_{b \in \{0,1\}^{\log n}} \hat{x}(b)\hat{E}(b, r_1)$, so Eq. 14 is a true claim, and $\tau_x = \hat{x}(r_2)$, $\tau_E = \hat{E}(r_2, r_1)$ pass the final check of Step 5, and the evaluation protocol for $\hat{E}$ accepts the true claim $\hat{E}(r_2, r_1) = \tau_E$ with probability 1. By the completeness of Theorem 3, the output claim $\hat{x}(r^*) = v^*$ is true, so $(\mathcal{C}'(x), r^*, v^*)$ is a YES instance of the constrained relation of $\mathcal{C}'$, and the IOP for $\mathcal{C}'$ accepts with probability 1. Hence the verifier always accepts.

Soundness. Fix a NO instance (w, z, v) and a malicious prover, and let $c_x \in \mathbb{F}^{N'}$ denote the oracle sent in Step 1. Suppose the verifier accepts; we bound the probability of this event by accumulating a small list of failure probabilities.

First, by the soundness of the IOP for $\mathcal{C}'$ (Step 7), except with probability $\epsilon_{\mathcal{C}'}$, the word c_x is close to a unique codeword $\mathcal{C}'(x^*)$ and the output claim of Step 6 is true, i.e., $\widehat{x^*}(r^*) = v^*$. Next, by the soundness of Theorem 3 applied to the polynomial $\widehat{x^*}$: if $\widehat{x^*}(z) \neq v$ or $\widehat{x^*}(r_2) \neq \tau_x$, then, except with probability $(2 \log n + 1)/|\mathbb{F}|$, the verifier rejects in Step 6 or the output claim is false—and we have just ruled out a false output claim. Hence, except with total probability $\epsilon_{\mathcal{C}'} + (2 \log n + 1)/|\mathbb{F}|$, we have $\widehat{x^*}(z) = v$ and $\widehat{x^*}(r_2) = \tau_x$.

By the soundness of the evaluation protocol for $\hat{E}$ (Step 5), except with probability ϵ_{sp} we may further assume $\tau_E = \hat{E}(r_2, r_1)$. Write $u^* := x^*E$. If $\tau_u \neq \widehat{u^*}(r_1) = \sum_{b \in \{0,1\}^{\log n}} \widehat{x^*}(b)\hat{E}(b, r_1)$, then, by the soundness of the sumcheck protocol, except with probability $2 \log n/|\mathbb{F}|$ the verifier rejects or the final check fails on the true values $\tau_x = \widehat{x^*}(r_2)$ and $\tau_E = \hat{E}(r_2, r_1)$; in either case the verifier rejects. Hence, except with an additional $2 \log n/|\mathbb{F}| + \epsilon_{\mathrm{sp}}$, we have $\tau_u = \widehat{u^*}(r_1)$.

Similarly, the polynomial $\widehat{A_q}$ is determined by θ and the sampled indices, and the verifier evaluates $\widehat{A_q}(r_1)$ itself. Write $w^* := \mathsf{EA}(x^*) = u^*A$. By the soundness

of the sumcheck protocol in Step 4, except with probability $2\log N/|\mathbb{F}|$ we have

$$\sigma = \sum_{j\in\{0,1\}^{\log N}} \widehat{u^*}(j)\,\widehat{A_q}(j) = \langle u^*, Aq\rangle = \langle w^*, q\rangle = \sum_{k=1}^{\nu} \theta^{k-1}\cdot \widehat{w^*}(i_k).$$

On the other hand, $\sigma = \sum_{k=1}^{\nu}\theta^{k-1}\hat{w}(i_k)$ by construction. Both sides are polynomials in θ of degree less than ν, and they are determined before θ is drawn (c_x is sent in Step 1, and the indices are sampled in Step 2). Hence, if the sampled positions do not all agree, the two sides coincide with probability at most $(\nu-1)/|\mathbb{F}|$ over the choice of θ. Except with this probability, we conclude $\hat{w}(i_k) = \widehat{w^*}(i_k)$ for all $k \in [\nu]$.

It remains to bound the probability of full agreement on a NO instance. There are two cases.

- w is η-far from every EA codeword. Then in particular $\Delta(w, w^*) \geq \eta$, and since the indices $i_1, \ldots, i_\nu$ are sampled independently and uniformly, all of them lie in the agreement set with probability at most $(1-\eta)^\nu \leq 2^{-\lambda}$.
- w is η-close to a codeword $\mathsf{EA}(x')$ (unique, since $\eta < \delta_{\mathsf{EA}}/2$) with $\widehat{x'}(z) \neq v$. If $x^* = x'$, then $\widehat{x^*}(z) \neq v$, contradicting the conclusion $\widehat{x^*}(z) = v$ established above (an event already accounted for). Otherwise $x^* \neq x'$, and by the triangle inequality $\Delta(w, w^*) \geq \delta_{\mathsf{EA}} - \eta > \eta$, so again all ν samples agree with probability at most $2^{-\lambda}$.

Collecting all failure probabilities yields the claimed soundness error $2^{-\lambda} + \epsilon_{\mathrm{sp}} + \epsilon_{\mathcal{C}'} + (\nu + 2\log N + 4\log n)/|\mathbb{F}|$. Under our assumption on the field size, the last term is negligible compared with the preceding terms. Hence, the overall soundness error is $O(2^{-\lambda} + \epsilon_{\mathrm{sp}} + \epsilon_{\mathcal{C}'})$.

Efficiency. Recall that $N = O(n)$ and that E has $O(n\log n)$ nonzero entries. The prover computes $u = xE$ in $O(n\log n)$ operations, forms $A_q = Aq$ in $O(N)$ operations by one suffix-sum scan (as $(A_q)_j = \sum_{i\geq j}\hat{q}(i)$), and runs the sumcheck of Step 4, on a product of two multilinear polynomials in $\log N$ variables, in $O(N)$ time. For Step 5, the prover builds the table $(\hat{E}(b, r_1))_{b\in\{0,1\}^{\log n}}$ in $O(n\log n)$ operations—an $O(N)$-time table of $(\widetilde{eq}(j, r_1))_{j\in\{0,1\}^{\log N}}$ followed by one pass over the nonzero entries of E—and runs the sumcheck in $O(n)$ time. The protocol of Theorem 3 costs the prover $O(n)$. Together with encoding x under $\mathcal{C}'$, running the IOP prover, and running the prover of the evaluation protocol for $\hat{E}$, the total prover time is $O(n\log n) + T_{P,\mathcal{C}'}(n) + T_{P,\mathrm{sp}}(n\log n)$.

We now bound the verifier's running time. We first estimate the time required to evaluate $\widehat{A_q}(r_1)$. Since q is supported on the sampled indices $i_1, \ldots, i_\nu$, for any $\alpha \in \mathbb{F}^{\log N}$, $\widehat{A_q}(\alpha) = \sum_{k=1}^{\nu}\theta^{k-1}\cdot\hat{A}(\alpha, i_k)$. By the structure of the matrix A, its MLE can be evaluated as

$$\hat{A}(\alpha,\beta) = \widetilde{eq}(\alpha,\beta) + \sum_{s=1}^{\log N}(1-\alpha_s)\cdot\beta_s\cdot\prod_{s'=1}^{s-1}\widetilde{eq}(\alpha_{s'},\beta_{s'}),$$

By maintaining the prefix products, each $\hat{A}(r_1, i_k)$ can be computed in $O(\log N)$ field operations, and hence evaluating $\widehat{A_q}(r_1)$ costs $O(\nu\log N)$. The remaining verification work consists of verifying the two sumchecks in Steps 4 and 5 and

the batch protocol in Step 6, each requires $O(\log n)$ operations, together with constant-many $\widetilde{eq}$ evaluations in dimension $\log n$, the verifier for the evaluation protocol of $\hat{E}$, and the IOP verifier. Therefore, the total verifier time is $O(\nu \log N + \log n) + T_{V,\mathcal{C}'}(n) + T_{V,\mathrm{sp}}(n \log n)$.

Finally, each sumcheck has individual degree two and thus communicates three field elements per round, i.e., $3(\log N + \log n)$ elements over Steps 4 and 5, and the protocol of Theorem 3 communicates $O(\log n)$ further elements. Together with the constant-many claimed values $\tau_u, \tau_x, \tau_E, v^*$, the non-oracle communication is $O(\log n)$ field elements, plus $C_{\mathcal{C}'}(n)$ and $C_{\mathrm{sp}}(n \log n)$ contributed by the sub-protocols. The verifier makes exactly ν queries to the input oracle (Step 2); the proof oracles are $[\![\mathcal{C}'(x)]\!]$, queried $q_{\mathcal{C}'}(n)$ times by the IOP, and those sent by the evaluation protocol for $\hat{E}$, queried $q_{\mathrm{sp}}(n \log n)$ times.

4.2 IOPP for Interleaved EA Codes

In this part, we consider the case of the *interleaved* code.

Let $t \in \mathbb{N}$ be the interleaving parameter. For a matrix $X \in \mathbb{F}^{n \times t}$, we denote by $\mathsf{EA}^t(X) \in \mathbb{F}^{N \times t}$ the matrix whose j-th column is the encoding of the j-th column of X; the relative distance between matrices counts the fraction of disagreeing *rows*. For a slack parameter $\eta > 0$, the interleaved promise problem $(\mathsf{YES}_t, \mathsf{NO}_t)$ is defined as follows:

- the YES_t instances are the tuples $(z, v, \mathsf{EA}^t(X))$ with $z \in \mathbb{F}^{\log n + \log t}$, $v \in \mathbb{F}$, and $X \in \mathbb{F}^{n \times t}$ such that $\hat{X}(z) = v$;
- the NO_t instances are the tuples (z, v, Y) such that either:
  - $\Delta(Y, \mathsf{EA}^t) \geq \delta_{\mathsf{EA}}/3 + \eta$; or
  - there exists a (unique) codeword $\mathsf{EA}^t(U)$ such that $\Delta(Y, \mathsf{EA}^t(U)) < \delta_{\mathsf{EA}}/3 + \eta$, but $\hat{U}(z) \neq v$.

The complete protocol is presented in Protocol 2.

Theorem 8. *Let $t \in \mathbb{N}$, let $\eta \in (0, \delta_{\mathsf{EA}}/6)$, and consider the promise problem $(\mathsf{YES}_t, \mathsf{NO}_t)$ with slack parameter η. Suppose Protocol 1 is invoked with proximity parameter $\delta_{\mathsf{EA}}/3$ and has soundness error ϵ_{EA} (Theorem 7). Then Protocol 2 is an IOP for $(\mathsf{YES}_t, \mathsf{NO}_t)$ with perfect completeness and soundness error $\epsilon_t \leq 1/\eta|\mathbb{F}| + 1/|\mathbb{F}| + \epsilon_{\mathsf{EA}}$. Assume $\lambda = \omega(\log n)$ and $|\mathbb{F}| = 2^{\Theta(\lambda)}$, we have $\epsilon_t \leq O(2^{-\lambda}) + \epsilon_{\mathsf{EA}}$. Moreover, beyond a single execution of the inner protocol on a length-n instance, the prover sends t additional field elements, and the verifier reads ν rows of Y and performs $O(\nu t)$ additional field operations.*

Proof. Completeness. Suppose $Y = \mathsf{EA}^t(X)$ and $\hat{X}(z) = v$. Since $\hat{X}(z_1, \cdot)$ is multilinear in its last $\log t$ variables, the MLE of g satisfies $\hat{g} = \hat{X}(z_1, \cdot)$, so $\hat{g}(z_2) = \hat{X}(z_1, z_2) = v$ and the check of Step 3 passes. Write $x := Xr$. By linearity of the MLE, for every $b \in \mathbb{F}^{\log n}$, $\hat{x}(b) = \sum_{j \in \{0,1\}^{\log t}} r_j \cdot \hat{X}(b, j)$. Moreover, by linearity of the code, $Yr = \mathsf{EA}^t(X)r = \mathsf{EA}(Xr) = \mathsf{EA}(x)$, so every emulated query is answered consistently with the exact codeword $\mathsf{EA}(x)$. Hence

PROTOCOL 2 *IOPP for interleaved EA codes.*

The prover claims $((z, v), W, X) \in \mathcal{R}_{\mathsf{EA}^t}$ to the verifier. The verifier holds the explicit input $(z, v) \in \mathbb{F}^{\log n + \log t} \times \mathbb{F}$ and has oracle access $[\![W]\!]$ to the implicit input $W \in \mathbb{F}^{N \times t}$, where a single query returns a full row of W; the prover additionally holds the witness $X \in \mathbb{F}^{n \times t}$.

1. Both parties parse z as $z = (z_1, z_2) \in \mathbb{F}^{\log n} \times \mathbb{F}^{\log t}$.
2. The prover sends, as a non-oracle message, the function $g : \{0,1\}^{\log t} \to \mathbb{F}$ defined by $g(j) := \hat{X}(z_1, j)$ for $j \in \{0,1\}^{\log t}$.
3. The verifier checks that $\hat{g}(z_2) = v$.
4. The verifier samples $r \leftarrow \mathbb{F}^t$ and sends it to the prover.
5. The parties run Protocol 1 with proximity parameter $\delta_{\mathsf{EA}}/3$ on the explicit input $(z_1, \langle r, g\rangle)$ (viewing g as a vector in $\mathbb{F}^t$) and the implicit input $Yr \in \mathbb{F}^N$, for which the prover uses the witness $x := Xr$; in the honest case, $Yr = \mathsf{EA}(Xr)$. Each query $\widehat{Yr}(i)$ of the inner verifier, for $i \in \{0,1\}^{\log N}$, is answered by querying the i-th row $Y_i \in \mathbb{F}^t$ of Y and returning $\langle Y_i, r\rangle$.
6. The verifier accepts if the check of Step 3 and the inner protocol accept; otherwise it rejects.

$(z_1, \langle r, g\rangle, \mathsf{EA}(x))$ is a YES instance of the (non-interleaved) promise problem, and the inner IOP accepts with probability 1 by Theorem 7.

Soundness. Fix a NO_t instance (z, v, Y) and a malicious prover, who sends some function g in Step 2. We distinguish the two cases of the definition of NO_t.

Case 1: Y is $(\delta_{\mathsf{EA}}/3 + \eta)$-far from EA^t. By the correlated agreement theorem for linear codes in Theorem 5, applied with distance parameter $\delta_{\mathsf{EA}}/3 + \eta$ and slack η (note that $\delta_{\mathsf{EA}}/3 \le \mathsf{d}(\mathsf{EA})/3$), a uniformly random combination satisfies $\Pr_{r \leftarrow \mathbb{F}^t}[\Delta(Yr, \mathsf{EA}) < \delta_{\mathsf{EA}}/3] \le \frac{1}{\eta|\mathbb{F}|}$. Hence, except with probability $1/(\eta|\mathbb{F}|)$, the word Yr is at least $\delta_{\mathsf{EA}}/3$-far from EA, so $(Yr, z_1, \langle r, g\rangle)$ is a NO instance of the inner promise problem (regardless of the claimed value), and the inner IOP accepts with probability at most ϵ_{EA}.

Case 2: Y is $(\delta_{\mathsf{EA}}/3 + \eta)$-close to $\mathsf{EA}^t(U)$ with $\hat{U}(z) \neq v$. Since $\delta_{\mathsf{EA}}/3 + \eta < \delta_{\mathsf{EA}}/2$, the codeword $\mathsf{EA}^t(U)$ is unique. Define the honest restriction $g' : \{0,1\}^{\log t} \to \mathbb{F}$ by $g'(j) := \hat{U}(z_1, j)$.

If $g = g'$, then $\hat{g}(z_2) = \hat{U}(z_1, z_2) = \hat{U}(z) \neq v$, so the verifier rejects in Step 3 with probability 1.

Otherwise $g \neq g'$. Since $\langle r, g - g'\rangle$ is a nonzero linear function of r, we have $\Pr_r[\langle r, g\rangle = \langle r, g'\rangle] \le 1/|\mathbb{F}|$. Condition on $\langle r, g\rangle \neq \langle r, g'\rangle$ and write $x' := Ur$, so that $\widehat{x'}(z_1) = \langle r, g'\rangle \neq \langle r, g\rangle$. Because every row on which Y agrees with $\mathsf{EA}^t(U)$ is preserved by the column combination, $\Delta(Yr, \mathsf{EA}(x')) \le \Delta(Y, \mathsf{EA}^t(U)) \le \delta_{\mathsf{EA}}/3 + \eta$. We claim that $(Yr, z_1, \langle r, g\rangle)$ is a NO instance of the inner promise problem. Indeed, if $\Delta(Yr, \mathsf{EA}) \ge \delta_{\mathsf{EA}}/3$, this holds by the first condition. Otherwise Yr is $\delta_{\mathsf{EA}}/3$-close to some codeword $\mathsf{EA}(x'')$; then $\Delta(\mathsf{EA}(x''), \mathsf{EA}(x')) \le \delta_{\mathsf{EA}}/3 + (\delta_{\mathsf{EA}}/3 + \eta) < \delta_{\mathsf{EA}}$, which forces $x'' = x'$, and the claimed evaluation satisfies $\langle r, g\rangle \neq \widehat{x'}(z_1)$, so the second condition holds. In either case, the inner IOP accepts with probability at most ϵ_{EA}.

Collecting the failure probabilities of the two cases yields $\epsilon_t \le 1/(\eta|\mathbb{F}|) + 1/|\mathbb{F}| + \epsilon_{\mathsf{EA}}$. Under the assumption on the field size, we have $\epsilon_t \le O(2^{-\lambda}) + \epsilon_{\mathsf{EA}}$

Efficiency. Beyond one execution of the inner protocol on a length-n instance, the prover sends the t values of g, and the verifier evaluates $\hat{g}(z_2)$ and $\langle r, g\rangle$ in $O(t)$ field operations. Each of the ν input queries of the inner protocol is answered by reading one row of Y (t field elements) and taking its inner product with r, in $O(t)$ operations per query.

4.3 Instantiation and Optimization

We instantiate the auxiliary code $\mathcal{C}'$ with the BaseFold's code, which has efficient protocol for the corresponding constrained relation. We record the parameters needed in Theorem 7.

Theorem 9 (BaseFold IOPP [51]). *Let $\mathcal{C}_{\mathsf{BF}} : \mathbb{F}^n \to \mathbb{F}^{N_{\mathsf{BF}}}$ be a random foldable code with $N_{\mathsf{BF}} = \Theta(n)$ and relative distance $\delta_{\mathsf{BF}} > 0$. Assume $\log n = O(\log \lambda)$ and $|\mathbb{F}| = 2^{\omega(\lambda)}$ For every proximity parameter $\delta \in (0, \delta_{\mathsf{BF}}/3)$, there is a δ-IOPP for the constrained relation $\mathcal{R}_{\mathcal{C}_{\mathsf{BF}}}$ with perfect completeness and the following parameters:*

- *Soundness error: $O(2^{-\lambda})$;*
- *Input query complexity: $O(\lambda)$;*
- *Proof query complexity: $O(\lambda \log n)$;*
- *Non-oracle communication: $O(\log n)$ field elements;*
- *Prover time: $O(n \log n)$;*
- *Verifier time: $O(\lambda \log n)$ field operations;*

To prove evaluation claims associated with the sparse expansion matrix E, we represent E as a sparse multilinear polynomial and leverage the Spark compiler of Spartan [40], which transforms a polynomial commitment scheme into one that efficiently supports sparse polynomials. We recall the relevant result below.

Theorem 10 (Spark [40]). *Let f be an m-variate multilinear polynomial whose evaluation table over $\{0,1\}^m$ contains at most s nonzero entries. If $s = \Omega(2^{m/2})$, then Spark yields a sparse polynomial commitment scheme for f whose asymptotic costs match those of the underlying polynomial commitment scheme instantiated for a polynomial of size s, with additional soundness error $s/|\mathbb{F}|$ and non-oracle proof size $O(m)$ field elements.*

IOPP Instantiation. Combining the above results, we instantiate the IOPP for the interleaved EA code in Protocol 2 with the following parameters:

- Soundness error: $O(2^{-\lambda})$;
- Input query complexity: $O(\lambda)$;
- Proof query complexity: $O(\lambda \cdot \log n)$;
- Non-oracle communication: $O(\log n + t)$ field elements;
- Prover time: $O(n \cdot (t \cdot \log n + \log^2 n))$;
- Verifier time: $O(\lambda \cdot (t + \log n))$.

For a polynomial of size M, we set $t = \Theta(\log M)$ and $n = M/t$. Therefore, we get the IOPP for the evaluation proof with prover time $O(M \log M)$, verifier time $O(\lambda \log M)$, input query complexity $O(\lambda)$, proof query complexity $O(\lambda \cdot \log M)$. Finally, we compile the IOPP into a PCS in the random oracle model using the standard transformation [9, 26, 32]. In particular, instantiating the oracle commitments with using Merkle trees [31], each oracle query incurs an $O(\log M)$-sized authentication path. Hence, the resulting PCS has proof size and verifier time $O(\lambda \log^2 M)$.

Amortized Opening. The most expensive component of the above protocol is the evaluation proof for the sparse expansion matrix E. When opening several distinct polynomials, the protocol requires evaluating E at multiple points. Rather than executing a separate evaluation proof for each claim, we batch all claims into a single one using the protocol in Theorem 3, and invoke the sparse-polynomial evaluation protocol only once on the resulting claim. This replaces multiple independent openings of E with a single opening, at the cost of only an inexpensive batching protocol, thereby amortizing the cost of opening E across multiple polynomial openings.

5 Evaluation

We implement Flare in Rust using the arkworks [15] ecosystem, with Blake3 for hashing and the secp256k1 field for field arithmetic. We benchmark Flare against five field-agnostic PCSs: Brakedown [17], BaseFold [51], ERA [6], RAA [1], and QA [28]. All experiments were conducted on a machine equipped with Intel Xeon Platinum 8358 with 128 GiB memory. We report average runtimes over 10 independent executions.

Parameters. For Brakedown, we adopt the improved analysis of [20], using inverse rate 2, relative distance 0.1, $\alpha = 0.3$, $c_n = 11$, and $d_n = 22$. We instantiate QA with inverse rate 2 and relative distance 0.4668, and RAA with inverse rate 4 and relative distance 0.3. For BaseFold, we use inverse rate 2 and relative distance 0.224. For ERA, the base code is instantiated as a Brakedown code with inverse rate 1.1 and relative distance 0.01, together with repetition factor $r = 4$. This yields an ERA code with inverse rate 4.4 and relative distance 0.226, which is then code-switched to a BaseFold code with inverse rate 8 and relative distance 0.728. For Flare, we use an EA code with inverse rate 2, relative distance 0.49, and $\gamma = 10$, and then code-switch it to a BaseFold code with inverse rate 4 and relative distance 0.585.

All codes except RAA are sampled so that the probability of failing to achieve the target distance is at most 2^{-100}. For the instance sizes considered here, the chosen RAA parameters guarantee a failure probability of at most 2^{-20}.

Optimizations. We apply vectorization, multithreading and cache-aware optimizations to all schemes where applicable. The parallel executions use 32 threads. For Flare, we implement the batched-opening optimization of Section 4.3, and instantiate the sparse PCS using Spark [40] with Shout [41, 48]. For all schemes,

Scheme	Commit Time (ms)	Open Time (ms)	Verification Time (ms)	Proof Size (MiB)
Brakedown	256.87	28.47	28.20	6.30
QA	190.93	16.41	16.70	2.69
RAA	95.82	23.47	18.93	3.44
BaseFold	120.17	89.28	21.76	2.34
ERA	86.54	297.53	20.06	2.18
Flare	191.00	252.30	14.15	1.35

Table 3: PCS performance for a polynomial of size 2^{22}.

we adaptively choose the interleaving parameter based on the number of queries and the polynomial size to minimize proof size.

5.1 Polynomial Commitment Scheme

We first benchmark the efficiency of the PCSs and report the cost of committing to a polynomial of size 2^{22} in Table 3. For Flare, we report the costs amortized over a batch opening of 8 distinct polynomials. For reference, without amortization, Flare requires 1285.97 ms for opening and 35.30 ms for verification, with a proof size of 3.46 MiB.

Despite the quasi-linear encoding complexity of the EA code, Flare retains a practical commitment cost owing to the simplicity of the EA encoding procedure. In particular, Flare takes 191.00 ms to commit, making it comparable to QA and outperforming Brakedown by 1.34×. RAA, BaseFold, and ERA achieve faster commitment, by 2×, 1.59×, and 2.21×, respectively.

Flare achieves the smallest proof size and the fastest verification among all evaluated schemes. Compared with BaseFold, Flare produces 1.73× smaller proofs and achieves 1.54× faster verification, at the cost of a 2.83× slower opening procedure. Compared with ERA, Flare simultaneously improves all three evaluation metrics: it achieves 1.18× faster opening, 1.42× faster verification, and 1.61× smaller proofs.

Compared with Brakedown, QA, and RAA, Flare trades opening time for substantially smaller proofs and more efficient verification. For example, relative to Brakedown, Flare produces 4.67× smaller proofs and achieves 2× faster verification, at the cost of an 8.86× slower opening procedure.

5.2 SNARKs

To demonstrate the gains of Flare in an end-to-end setting, we benchmark the PCSs within SNARKs, which naturally involve openings of multiple distinct polynomials. Table 4 reports the results for proving R1CS instances with 2^{22} constraints.

Flare achieves the smallest proof size and the fastest verification among all evaluated schemes. Compared with BaseFold and ERA, Flare maintains comparable prover performance while substantially improving proof size and verification

Scheme	Prover Time (s)	Verifier Time (ms)	Proof Size (MiB)
Brakedown	5.15	522.01	75.93
QA	5.28	307.12	33.36
RAA	4.78	276.32	42.62
BaseFold	6.07	384.82	26.07
ERA	7.93	278.45	22.97
Flare	7.71	213.18	15.67

Table 4: End-to-end SNARK performance for an arithmetic circuit of size 2^{22}.

efficiency. Specifically, relative to BaseFold, Flare incurs a modest 1.27× prover overhead, while producing 1.66× smaller proofs and achieving 1.81× faster verification. Relative to ERA, Flare is slightly faster in proving, while achieving 1.31× faster verification and 1.47× smaller proofs.

Compared with Brakedown, QA, and RAA, the prover-performance gap observed in the PCS benchmarks is substantially narrowed, while Flare's advantages in proof size and verification efficiency are retained. For instance, compared with Brakedown, Flare incurs only a 1.50× prover overhead, while producing 4.85× smaller proofs and achieving 2.45× faster verification.

References

1. Akhiani, P., Zhang, Y.: Distance of RAA codes over large finite fields (with applications in zkSNARKs and PCGs). Cryptology ePrint Archive, Report 2026/524 (2026), https://eprint.iacr.org/2026/524
2. Ames, S., Hazay, C., Ishai, Y., Venkitasubramaniam, M.: Ligero: Lightweight sublinear arguments without a trusted setup. In: Proceedings of the 24th ACM Conference on Computer and Communications Security. pp. 2087–2104. CCS '17, Association for Computing Machinery, New York, NY, USA (2017)
3. Arnon, G., Chiesa, A., Fenzi, G., Yogev, E.: STIR: reed-solomon proximity testing with fewer queries. In: Reyzin, L., Stebila, D. (eds.) Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part X. Lecture Notes in Computer Science, vol. 14929, pp. 380–413. Springer (2024). https://doi.org/10.1007/978-3-031-68403-6_12, https://doi.org/10.1007/978-3-031-68403-6_12
4. Arnon, G., Chiesa, A., Fenzi, G., Yogev, E.: WHIR: reed-solomon proximity testing with super-fast verification. In: Fehr, S., Fouque, P. (eds.) Advances in Cryptology - EUROCRYPT 2025 - 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4-8, 2025, Proceedings, Part IV. Lecture Notes in Computer Science, vol. 15604, pp. 214–243. Springer (2025). https://doi.org/10.1007/978-3-031-91134-7_8, https://doi.org/10.1007/978-3-031-91134-7_8
5. Baldimtsi, F., Chalkias, K.K., Ji, Y., Lindstrøm, J., Maram, D., Riva, B., Roy, A., Sedaghat, M., Wang, J.: zklogin: Privacy-preserving blockchain authentication with existing credentials. In: Luo, B., Liao, X., Xu, J., Kirda, E., Lie, D. (eds.) Proceedings of the 2024 on ACM SIGSAC Conference on Computer and

Communications Security, CCS 2024, Salt Lake City, UT, USA, October 14-18, 2024. pp. 3182–3196. ACM (2024). https://doi.org/10.1145/3658644.3690356, https://doi.org/10.1145/3658644.3690356
6. Baweja, A., Fenzi, G., Mishra, P., Mopuri, T.: Field-agnostic SNARKs with small proofs via encode-repeat-accumulate (ERA) codes. IACR Cryptol. ePrint Arch. **2026**, 864 (2026), https://eprint.iacr.org/2026/864
7. Baweja, A., Mishra, P., Mopuri, T., Shtepel, M.: FICS and FACS: fast iopps and accumulation via code-switching. IACR Cryptol. ePrint Arch. **2025**, 737 (2025), https://eprint.iacr.org/2025/737
8. Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Fast reed-solomon interactive oracle proofs of proximity. In: Proceedings of the 45th International Colloquium on Automata, Languages, and Programming. pp. 14:1–14:17. ICALP '18 (2018)
9. Ben-Sasson, E., Chiesa, A., Spooner, N.: Interactive oracle proofs. In: Proceedings of the 14th Theory of Cryptography Conference. pp. 31–60. TCC '16-B (2016)
10. Ben-Sasson, E., Kopparty, S., Saraf, S.: Worst-case to average case reductions for the distance to a code. In: Servedio, R.A. (ed.) 33rd Computational Complexity Conference, CCC 2018, San Diego, CA, USA, June 22-24, 2018. LIPIcs, vol. 102, pp. 24:1–24:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2018). https://doi.org/10.4230/LIPICS.CCC.2018.24, https://doi.org/10.4230/LIPIcs.CCC.2018.24
11. Block, A.R., Fang, Z., Katz, J., Thaler, J., Waldner, H., Zhang, Y.: Field-agnostic SNARKs from expand-accumulate codes. In: Proceedings of the 44th Annual International Cryptology Conference. pp. 276–307. CRYPTO '24, Springer (2024), https://eprint.iacr.org/2024/1871, extended version: Cryptology ePrint Archive, Report 2024/1871, revision dated March 1, 2026
12. Boyle, E., Couteau, G., Gilboa, N., Ishai, Y., Kohl, L., Resch, N., Scholl, P.: Correlated pseudorandomness from expand-accumulate codes. In: Proceedings of the 42nd Annual International Cryptology Conference. pp. 603–633. CRYPTO '22 (2022)
13. Brehm, M., Chen, B., Fisch, B., Resch, N., Rothblum, R.D., Zeilberger, H.: Blaze: Fast snarks from interleaved RAA codes. In: Fehr, S., Fouque, P. (eds.) Advances in Cryptology - EUROCRYPT 2025 - 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4-8, 2025, Proceedings, Part IV. Lecture Notes in Computer Science, vol. 15604, pp. 123–152. Springer (2025). https://doi.org/10.1007/978-3-031-91134-7_5, https://doi.org/10.1007/978-3-031-91134-7_5
14. Chen, B., Bünz, B., Boneh, D., Zhang, Z.: Hyperplonk: Plonk with linear-time prover and high-degree custom gates. In: Proceedings of the 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 499–530. EUROCRYPT '23, Springer Nature Switzerland, Cham (2023)
15. arkworks contributors: `arkworks` zksnark ecosystem (2022), https://arkworks.rs
16. Divsalar, D., Jin, H., McEliece, R.J.: Coding theorems for "turbo-like" codes. In: Proceedings of the 36th Annual Allerton Conference on Communication, Control, and Computing. pp. 201–210 (1998)
17. Golovnev, A., Lee, J., Setty, S.T.V., Thaler, J., Wahby, R.S.: Brakedown: Linear-time and field-agnostic snarks for R1CS. In: Proceedings of the 43rd Annual International Cryptology Conference. pp. 193–226. CRYPTO '23, Springer Nature Switzerland, Cham (2023)
18. Guo, X., Liu, H., Huang, Z., Cui, H., Zhang, W., Hong, C., Wang, X., Yang, K., Yu, Y.: Dory: Streaming PCG with small memory. In: IEEE Symposium on Security and Privacy, SP 2026, San Francisco, CA, USA, May 18-21, 2026.

pp. 1902–1920. IEEE (2026). https://doi.org/10.1109/SP63933.2026.00079, https://doi.org/10.1109/SP63933.2026.00079
19. Gurkan, K., Novakovic, A., Rothblum, R.D.: Bolt: Faster snarks from sketched codes. IACR Cryptol. ePrint Arch. **2026**, 310 (2026), https://eprint.iacr.org/2026/310
20. Haböck, U.: Brakedown's expander code. IACR Cryptol. ePrint Arch. **2023**, 769 (2023), https://eprint.iacr.org/2023/769
21. Hu, Y., Kumar, S., Popa, R.A.: Ghostor: Toward a secure {Data-Sharing} system from decentralized trust. In: NSDI'20. pp. 851–877. USENIX Association, Santa Clara, CA (2020)
22. Huang, K., Guo, Y., Qu, W., Zhang, J.: Bootstrapping-free blind PCS: achieving constant depth and linear prover. IACR Cryptol. ePrint Arch. **2026**, 487 (2026), https://eprint.iacr.org/2026/487
23. Kaviani, D., Setty, S.: Vega: Low-latency zero-knowledge proofs over existing credentials. In: IEEE Symposium on Security and Privacy, SP 2026, San Francisco, CA, USA, May 18-21, 2026. pp. 1951–1969. IEEE (2026). https://doi.org/10.1109/SP63933.2026.00182, https://doi.org/10.1109/SP63933.2026.00182
24. Khabbazian, M.: Linear distance for fixed-row-weight expand–accumulate codes over arbitrary fields. Cryptology ePrint Archive, Paper 2026/1753 (2026), https://eprint.iacr.org/2026/1753
25. Khabbazian, M.: Sharp minimum-distance lower tails for RAA codes. Cryptology ePrint Archive, Paper 2026/1075 (2026), https://eprint.iacr.org/2026/1075
26. Kilian, J.: A note on efficient zero-knowledge proofs and arguments. In: Proceedings of the 24th Annual ACM Symposium on Theory of Computing. pp. 723–732. STOC '92 (1992)
27. Kosba, A.E., Miller, A., Shi, E., Wen, Z., Papamanthou, C.: Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In: Proceedings of the 37th IEEE Symposium on Security and Privacy. pp. 839–858. S&P '16 (2016)
28. Li, Z., Liu, H., Xing, C., Yao, Y., Yuan, C.: More efficient SNARKs via quasi-abelian codes: Faster, smaller, and field-agnostic. IACR Cryptol. ePrint Arch. **2026**, 939 (2026), https://eprint.iacr.org/2026/939
29. Liu, T., Zhang, Z., Zhang, Y., Hu, W., Zhang, Y.: Ceno: Non-uniform, segment and parallel zero-knowledge virtual machine. J. Cryptol. **38**(2), 17 (2025). https://doi.org/10.1007/S00145-024-09533-2, https://doi.org/10.1007/s00145-024-09533-2
30. Lund, C., Fortnow, L., Karloff, H.J., Nisan, N.: Algebraic methods for interactive proof systems. Journal of the ACM **39**(4), 859–868 (1992)
31. Merkle, R.C.: A certified digital signature. In: Proceedings of the 9th Annual International Cryptology Conference. pp. 218–238. CRYPTO '89 (1989)
32. Micali, S.: Computationally sound proofs. SIAM journal on computing **30**(4), 1253–1298 (2000), preliminary version appeared in FOCS '94.
33. Nair, V., Sharma, A., Thankey, B.: BrakingBase: A linear-prover, polylogarithmic-verifier, field-agnostic polynomial commitment scheme. In: Advances in Cryptology – ASIACRYPT 2025 (2025), https://eprint.iacr.org/2024/1825, extended version: Cryptology ePrint Archive, Report 2024/1825
34. Qu, W., Guo, Y., Zhang, J.: Lightning, field-agnostic super-efficient polynomial commitment scheme. IACR Cryptol. ePrint Arch. **2026**, 258 (2026), https://eprint.iacr.org/2026/258
35. Qu, W., Sun, Y., Liu, X., Lu, T., Guo, Y., Chen, K., Zhang, J.: zkgpt: an efficient non-interactive zero-knowledge proof framework for llm inference. In: Proceedings of the 34th USENIX Conference on Security Symposium. SEC '25, USENIX Association, USA (2025)

36. Reingold, O., Rothblum, R., Rothblum, G.: Constant-round interactive proofs for delegating computation. In: Proceedings of the 48th ACM Symposium on the Theory of Computing. pp. 49–62. STOC '16 (2016)
37. Ron-Zewi, N., Rothblum, R.D.: Local proofs approaching the witness length [extended abstract]. In: Irani, S. (ed.) 61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020. pp. 846–857. IEEE (2020). https://doi.org/10.1109/FOCS46700.2020.00083, https://doi.org/10.1109/FOCS46700.2020.00083
38. Ron-Zewi, N., Rothblum, R.D.: Proving as fast as computing: succinct arguments with constant prover overhead. In: Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing. p. 1353–1363. STOC 2022, Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3519935.3519956, https://doi.org/10.1145/3519935.3519956
39. Rothblum, G.N., Vadhan, S.P., Wigderson, A.: Interactive proofs of proximity: delegating computation in sublinear time. In: Proceedings of the 45th ACM Symposium on the Theory of Computing. pp. 793–802. STOC '13 (2013)
40. Setty, S.: Spartan: Efficient and general-purpose zksnarks without trusted setup. In: Proceedings of the 40th Annual International Cryptology Conference. pp. 704–737. CRYPTO '20 (2020)
41. Setty, S.T.V., Thaler, J.: Twist and shout: Faster memory checking arguments via one-hot addressing and increments. IACR Cryptol. ePrint Arch. p. 105 (2025), https://eprint.iacr.org/2025/105
42. Spielman, D.A.: Linear-time encodable and decodable error-correcting codes. IEEE Transactions on Information Theory **42**(6), 1723–1731 (1996), preliminary version appeared in STOC '95.
43. Sun, H., Li, J., Zhang, H.: zkllm: Zero knowledge proofs for large language models. In: Luo, B., Liao, X., Xu, J., Kirda, E., Lie, D. (eds.) Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security, CCS 2024, Salt Lake City, UT, USA, October 14-18, 2024. pp. 4405–4419. ACM (2024). https://doi.org/10.1145/3658644.3670334, https://doi.org/10.1145/3658644.3670334
44. Thaler, J.: Time-optimal interactive proofs for circuit evaluation. In: Annual Cryptology Conference. pp. 71–89. Springer, Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
45. Thaler, J.: Proofs, arguments, and zero-knowledge. Found. Trends Priv. Secur. **4**(2-4), 117–660 (2022)
46. Vu, V., Setty, S., Blumberg, A.J., Walfish, M.: A hybrid architecture for interactive verifiable computation. In: Proceedings of the 34th IEEE Symposium on Security and Privacy. pp. 223–237. S&P '13 (2013)
47. Wang, L., Lou, H., Li, C., Yu, Y., Hu, Y.: zkagent: Verifiable agent execution via one-shot complete LLM inference proof. IACR Cryptol. ePrint Arch. **2026**, 199 (2026)
48. Wiese, G.: Twist and shout via logup*. Technical report, Powdr Labs (2025), https://www.powdr.org/papers/twist_shout_logup_star.pdf
49. Xie, T., Zhang, J., Cheng, Z., Zhang, F., Zhang, Y., Jia, Y., Boneh, D., Song, D.: zkBridge: Trustless cross-chain bridges made practical. CCS'22 (2022)
50. Xie, T., Zhang, Y., Song, D.: Orion: Zero knowledge proof with linear prover time. In: CRYPTO 2022. pp. 299–328. Springer Nature Switzerland, Cham (2022)
51. Zeilberger, H., Chen, B., Fisch, B.: Basefold: efficient field-agnostic polynomial commitment schemes from foldable codes. In: Annual International Cryptology Conference. pp. 138–169. Springer (2024)

AUXILIARY SUPPORTING MATERIAL

A AI Disclosure

The main results of this work—namely, that EA codes approach the Singleton bound and the PCS from EA codes—were developed primarily by the authors. Still, during the later stages of the project, we also had extensive discussions with GPT-5.6 Sol Pro to refine the results. Concretely, after we provided the model with the overall proof strategy, we used it to help derive tighter upper bounds and optimize certain constants for the distance analysis. The model was also used to assist with drafting parts of some proofs and with improving the exposition.

The extended result in Appendix B, showing that EA codes approach the GV bound over samll fields, was primarily discovered by GPT-5.6 Sol Pro. We provided the model with prior works by Boyle et al. [12] and Block et al. [11], and asked it to investigate whether EA codes could approach the GV bound. The model first produced an argument for the binary-field case, which is not included in the manuscript. We then asked it to investigate the corresponding statement over general finite fields, and it proposed a generalization of the binary-field argument. We also asked the model to optimize the analysis for low-weight messages, which determines the constant choice for code and implicates the practical efficiency of the codes. This leads to the refined bound in Lemma 14, which removes the polynomial factor in Lemma 13. Because the proof initially generated by the model was difficult to interpret, the authors independently verified its details and devoted substantial effort to reorganizing and rewriting the argument presented here.

We used GPT-5.6 Sol for coding and benchmarking.

The authors take full responsibility for the accuracy, integrity, and content of the manuscript, including all material developed with AI assistance.

B EA Codes over Small Finite Fields

Section 3 shows that EA codes approach the Singleton bound, but only over large fields: the failure probability in Theorem 6 is small only when $\log q = \omega(\log N)$. Several applications of EA codes instead operate over a fixed small field, most prominently $\mathbb{F}_2$. In this appendix we analyze the same exact-weight ensemble over an arbitrary finite field $\mathbb{F}_q$ and show that its rate–distance tradeoff approaches the Gilbert–Varshamov (GV) bound. Reaching a weaker bound than in the large-field regime is natural, as the GV bound remains the best rate–distance tradeoff known over small alphabets.

Over $\mathbb{F}_2$, Boyle et al. [12] proved that EA codes have constant relative distance with high probability when the entries of E are sampled independently from a Bernoulli distribution. Block et al. [11] later obtained a related result for juxtaposed EA codes, which combine an expansion matrix $E_1 \leftarrow \mathsf{BP}(n, N, t, \mathbb{F}_2)$ with another expansion matrix E_2 whose entries follows independent Bernoulli

distribution. Both works left open whether the single code generated by EA with $E \leftarrow \mathsf{BP}(n, N, t, \mathbb{F}_q)$ already has constant relative distance. Theorem 13 below resolves this question: the ensemble attains the q-ary GV region.

B.1 Additional Preliminaries

We first collect the probabilistic and information-theoretic tools used in this appendix, complementing Section 2.

Probability notation. We write $X \leftarrow \mathcal{D}$ when X is sampled from the distribution $\mathcal{D}$, and $X \overset{d}{=} Y$ when X and Y have the same distribution. Separate random draws are independent unless stated otherwise.

Bernoulli distribution. For $p \in [0,1]$, the Bernoulli distribution $\mathsf{Ber}(p)$ is supported on $\{0,1\}$ and assigns probability p to one. The binomial distribution $\mathsf{Bin}(m,p)$ is the distribution of a sum of m independent $\mathsf{Ber}(p)$ variables; if $W \leftarrow \mathsf{Bin}(m,p)$, then $\Pr[W = k] = \binom{m}{k} p^k (1-p)^{m-k}$ for $0 \le k \le m$, and $\mathbb{E}[s^W] = (1 - p + ps)^m$ for every $s \ge 0$.

Hypergeometric distribution. For integers $0 \le w, t \le N$, the hypergeometric distribution $\mathrm{Hyp}(N, w, t)$ is the distribution of the number of marked elements in a uniformly random t-subset of a population of N elements of which w are marked: if $L \leftarrow \mathrm{Hyp}(N, w, t)$, then $\Pr[L = \ell] = \binom{w}{\ell}\binom{N-w}{t-\ell}/\binom{N}{t}$, where a binomial coefficient is zero outside its natural range.

Entropy functions. With the convention $0 \ln 0 := 0$, we use the natural-log binary entropy $h_e(u) := -u \ln u - (1-u)\ln(1-u)$, the q-ary entropy $H_q(u) := u \log_q(q-1) - u \log_q u - (1-u)\log_q(1-u)$, and the binary relative entropy $D_{\mathrm{KL}}(a \| b) := a \ln \frac{a}{b} + (1-a) \ln \frac{1-a}{1-b}$, which is nonnegative.

Row distributions. For $0 \le w \le N$, let $\mathcal{S}_{w,q}^{(N)} := \{y \in \mathbb{F}_q^N : \mathsf{wt}(y) = w\}$ denote the q-ary Hamming sphere. The cardinality of $\mathcal{S}_{w,q}^{(N)}$ is $\binom{N}{w}(q-1)^w$. Let $\mu_{w,q}$ be the uniform distribution on $\mathcal{S}_{w,q}^{(N)}$. A vector $y \sim \mu_{w,q}$ can be sampled in two steps. First, choose a support $S \subseteq [N]$ of size w uniformly at random, then for each $i \in S$ choose y_i independently and uniformly from $\mathbb{F}_q^\times$, and set $y_i = 0$ for every $i \notin S$. The distribution $\mu_{w,q}$ is invariant under coordinate permutations and nonzero coordinate scalings.

For $\eta \in [0,1]$, let $\nu_{\eta,q}$ be the product distribution on $\mathbb{F}_q^N$ each of whose coordinates is zero with probability $1-\eta$ and equals each $a \in \mathbb{F}_q^\times$ with probability $\eta/(q-1)$. If $V \leftarrow \nu_{\eta,q}$, then $\mathsf{wt}(V) \leftarrow \mathsf{Bin}(N, \eta)$, and conditional on $\mathsf{wt}(V) = w$ the vector V follows $\mu_{w,q}$.

Convolution and point probabilities. For probability measures α, β on the additive group $\mathbb{F}_q^N$, define $(\alpha * \beta)(z) := \sum_{y \in \mathbb{F}_q^N} \alpha(z-y)\beta(y)$ and $\|\alpha\|_\infty := \max_{z \in \mathbb{F}_q^N} \alpha(z)$. Write α^{*r} for the r-fold convolution, then the sum of r independent samples from $\mu_{t,q}$ has the distribution $\mu_{t,q}^{*r}$. Bounding each factor $\alpha(z-y)$ by $\|\alpha\|_\infty$ gives $\|\alpha * \beta\|_\infty \le \|\alpha\|_\infty$, and summing point masses gives $\alpha(B) \le |B| \cdot \|\alpha\|_\infty$ for every $B \subseteq \mathbb{F}_q^N$.

Lemma 7 (Binomial entropy bounds). *For integers $0 \le k \le m$, we have $\frac{1}{m+1} e^{m h_e(k/m)} \le \binom{m}{k} \le e^{m h_e(k/m)}$.*

Proof. The endpoint cases $k \in \{0, m\}$ are immediate. For $0 < k < m$, put $p := k/m$, so that $p^k(1-p)^{m-k} = e^{-mh_e(p)}$ and hence $\binom{m}{k} e^{-mh_e(p)} = \Pr[W = k]$ for $W \leftarrow \mathsf{Bin}(m, p)$. Since probabilities are at most one, the upper bound follows. For the lower bound, the ratio $\Pr[W = j+1]/\Pr[W = j] = \frac{(m-j)p}{(j+1)(1-p)}$ is at least one precisely when $j + 1 \le (m + 1)p$, so the mass of $\mathsf{Bin}(m, p)$ increases up to $j = k$ and decreases afterwards. Thus k is a mode, and the modal mass is at least the average $\frac{1}{m+1}$ of the $m + 1$ masses.

Lemma 8 (q-ary Hamming-ball volume). *Let $0 \le D \le (q-1)N/q$ and $s := D/N$. Then $V_q(N, D) := \sum_{h=0}^{D} \binom{N}{h}(q-1)^h \le q^{NH_q(s)}$. In particular, if $D \le \delta N$ and $\delta \le (q-1)/q$, then $V_q(N, D) \le q^{NH_q(\delta)}$.*

Proof. The case $D = 0$ is immediate, so assume $D \ge 1$ and let $V \leftarrow \nu_{s,q}$. Since $s \le (q-1)/q$, we have $\frac{s}{q-1} \le 1 - s$, so every $y \in \mathbb{F}_q^N$ with $\mathsf{wt}(y) = h \le D$ satisfies

$$\Pr[V = y] = \Big(\frac{s}{q-1}\Big)^h (1-s)^{N-h} \ge \Big(\frac{s}{q-1}\Big)^D (1-s)^{N-D}.$$

Taking logarithms and substituting $D = Ns$, we obtain

$$\log_q V_q(N, D) \le N\left(s \log_q(q-1) - s \log_q s - (1-s)\log_q(1-s)\right) = NH_q(s)$$

Exponentiating both sides yields exactly the first claim. The last claim follows because $H_q'(u) = \log_q \frac{(q-1)(1-u)}{u} \ge 0$ for $0 < u \le (q-1)/q$, so H_q is nondecreasing there.

The standard random-linear-code argument together with Lemma 8 yields the asymptotic form of the GV bound; our main theorem shows that EA codes attain the same tradeoff.

Theorem 11 (Gilbert–Varshamov bound). *Let q be a prime power and $0 < \delta < (q-1)/q$. For every $R < 1 - H_q(\delta)$ and all sufficiently large N, there exists a linear code over $\mathbb{F}_q$ with block length N, dimension $\lfloor RN \rfloor$, and relative distance greater than δ.*

The next three lemmas control hypergeometric moments and quantify the loss incurred when the exact-weight distribution $\mu_{t,q}$ is replaced by the product distribution $\nu_{t/N,q}$.

Lemma 9 (Hypergeometric moment). *Let $L \leftarrow \mathrm{Hyp}(N, w, t)$ and $s \ge 1$. Then $\mathbb{E}[s^L] \le \big(1 + (s-1)\frac{w}{N}\big)^t \le \exp\big((s-1)\frac{tw}{N}\big)$.*

Proof. Form a population of N numbers, w of which equal s and $N - w$ of which equal one. The product of the entries of a uniformly random t-subset is distributed as s^L, so $\mathbb{E}[s^L]$ is the normalized elementary symmetric mean of order t of this population. By Maclaurin's inequality, this mean is at most the t-th power of the arithmetic mean, which means $\mathbb{E}[s^L] \le \big(1 + (s-1)\frac{w}{N}\big)^t$. The second inequality follows from $1 + u \le e^u$.

Lemma 10. *Let $1 \le t \le N/2$ and $X \leftarrow \mathsf{Bin}(N, t/N)$. Then $b_{N,t} := \Pr[X = t] \ge \frac{3}{20\sqrt{t}}$.*

Proof. Put $p := t/N$, so that $\mathbb{E}[X] = t$. Since $\lfloor (N+1)p \rfloor = t$, the value t is a mode of X by the ratio computation in the proof of Lemma 7. Moreover, $\mathrm{Var}(X) = t(1-p) \le t$, so Chebyshev's inequality gives $\Pr[|X - t| < 2\sqrt{t}] \ge \frac{3}{4}$. The interval $(t - 2\sqrt{t}, t + 2\sqrt{t})$ contains at most $4\sqrt{t} + 1 \le 5\sqrt{t}$ integers, each carrying probability at most the modal mass $b_{N,t}$. Hence $b_{N,t} \ge \frac{3}{4}\frac{1}{5\sqrt{t}} = \frac{3}{20\sqrt{t}}$

Lemma 11. *Let $1 \le t \le N/2$, let $U_1, \dots, U_r$ be independent samples from $\mu_{t,q}$, and let $V_1, \dots, V_r$ be independent samples from $\nu_{t/N,q}$. Then, for every event $\mathcal{E} \subseteq (\mathbb{F}_q^N)^r$, we have $\Pr[(U_1, \dots, U_r) \in \mathcal{E}] \le q_{N,t}^{-r} \Pr[(V_1, \dots, V_r) \in \mathcal{E}]$.*

Proof. A sample from $\nu_{t/N,q}$ conditioned on having Hamming weight exactly t is uniform on $\mathcal{S}_t^{(N)}$. Since the rows are independent,

$$\begin{aligned}\Pr[(U_1, \dots, U_r) \in \mathcal{E}] &= \Pr[(V_1, \dots, V_r) \in \mathcal{E} \mid \mathsf{wt}(V_i) = t \ \forall i] \\ &\le q_{N,t}^{-r} \Pr[(V_1, \dots, V_r) \in \mathcal{E}]\,.\end{aligned}$$

Finally, we recall the input–output weight enumerator (IOWE) of the accumulator: for $0 \le w, h \le N$, let $A_{w,h}^{N,q} := |\{y \in \mathbb{F}_q^N : \mathsf{wt}(y) = w,\ \mathsf{wt}(yA) = h\}|$.

Theorem 12 ([11, 16]). *For a finite field $\mathbb{F}_q$ and $w, h \in [N]$,*

$$A_{w,h}^{N,q} = \sum_{i=0}^{w-1} \binom{h-1}{k_i - 1} \binom{N-h}{\ell_i} \binom{h-k_i}{i} (q-1)^{k_i} (q-2)^i,$$

where $k_i := \lceil (w-i)/2 \rceil$ and $\ell_i := \lfloor (w-i)/2 \rfloor$. In particular, $A_{w,h}^{N,q} = 0$ unless $h \ge \lceil w/2 \rceil$; equivalently, $\mathsf{wt}(yA) \ge \lceil \mathsf{wt}(y)/2 \rceil$ for every $y \in \mathbb{F}_q^N$.

We refer to Block et al. [11] for the proof; the last claim also follows directly from the identity $y_j = (yA)_j - (yA)_{j-1}$, since each nonzero coordinate of yA can account for at most two nonzero coordinates of y.

B.2 Reduction and Proof Overview

Throughout this appendix, q is a fixed prime power, $E \leftarrow \mathsf{BP}(n, N, t, \mathbb{F}_q)$. To keep the notation light, a summation limit or weight threshold written as δN always means $\lfloor \delta N \rfloor$. All constants may depend on q and δ but never on N.

The event $\mathsf{d}(EA) \le \delta$ occurs exactly when $\mathsf{wt}(xEA) \le \delta N$ for some nonzero message $x \in \mathbb{F}_q^n$. There are $\binom{n}{r}(q-1)^r$ messages of weight r, so the union bound gives

$$\Pr_E[\mathsf{d}(EA) \le \delta] \le \sum_{r=1}^{n} \binom{n}{r} (q-1)^r \max_{\mathsf{wt}(x)=r} \Pr_E[\mathsf{wt}(xEA) \le \delta N].$$

Fix a nonzero message x of weight r, let $j_1, \dots, j_r$ be the elements of $\mathsf{supp}(x)$, $U_1, \dots, U_r$ be the corresponding rows of E, and set $W_r := \mathsf{wt}(xE)$. Multiplication by a nonzero scalar preserves $\mu_{t,q}$, so $x_{j_s} U_s \overset{d}{=} U_s$ and hence $xE = \sum_{s=1}^{r} x_{j_s} U_s \overset{d}{=} U_1 + \cdots + U_r$. By the conditional-uniformity observation of Section B.1, conditional on $W_r = w \ge 1$, the vector xE is therefore uniform on $\mathcal{S}_{w,q}^{(N)}$.

Accordingly, for $1 \le w \le N$, let $Y_{w,q}$ be uniform on $\mathcal{S}_{w,q}^{(N)}$ and define $Q_{N,q}(w,\delta) := \Pr[\mathsf{wt}(Y_{w,q}A) \le \delta N]$. Since $xE = 0$ forces $xEA = 0$, we set $Q_{N,q}(0,\delta) := 1$. Then every nonzero message x of weight r satisfies

$$\begin{aligned}\Pr_E[\mathsf{wt}(xEA) \le \delta N] &= \sum_{w=0}^{N} \Pr[W_r = w] \cdot \Pr[\mathsf{wt}(xEA) \le \delta N | W_r = w] \\ &= \sum_{w=0}^{N} \Pr[W_r = w] \cdot Q_{N,q}(w,\delta) \\ &= \mathbb{E}[Q_{N,q}(W_r,\delta)] \end{aligned} \tag{15}$$

The analysis thus splits into two independent tasks: upper bounds on the accumulator tail $Q_{N,q}(w,\delta)$, and distributional estimates for the expansion weight W_r.

For $0 < \delta < (q-1)/q$, define $a_{q,\delta} := \frac{q-2}{q-1}\delta + 2\sqrt{\frac{\delta(1-\delta)}{q-1}}$, $\lambda_{q,\delta} := \ln \frac{1}{a_{q,\delta}}$, and $\kappa_{q,\delta} := 1 - a_{q,\delta}$. Lemma 12 below shows that $0 < a_{q,\delta} < 1$, so $\lambda_{q,\delta} > 0$ and $0 < \kappa_{q,\delta} < 1$. The quantity $a_{q,\delta}$ is the decay rate of the accumulator tail: Section B.3 establishes $Q_{N,q}(w,\delta) \le C_{q,\delta} N^3 a_{q,\delta}^w$ for every w, together with a variant for $w \le N/2$ in which the polynomial factor is replaced by $e^{w^2/N}$.

Estimating $\mathbb{E}[a_{q,\delta}^{W_r}]$ requires different arguments in different ranges of r. Set $\varepsilon_N := \frac{(\ln t)^3}{t}$, $r_0 := \lfloor \frac{\varepsilon_N N}{t} \rfloor$, and $r_1 := \lceil \frac{N \ln t}{t} \rceil$. With the parameter choice of Theorem 13 below, $1 \le r_0 < r_1 < n$ for all sufficiently large N: indeed, $r_0 \ge N(\ln t)^3/t^2 - 1 \to \infty$, $r_0/r_1 \le \varepsilon_N / \ln t = (\ln t)^2/t \to 0$, and $r_1 \le N \ln t/t + 1 = o(N)$ while $n = \lfloor RN \rfloor$. We call the messages with $1 \le r \le r_0$ *low-weight*, those with $r_0 < r < r_1$ *medium-weight*, and those with $r_1 \le r \le n$ *high-weight*.

For low-weight messages, we estimate $\mathbb{E}[a_{q,\delta}^{W_r}]$ directly under the exact-weight distribution, exposing one row at a time and controlling its overlap with the support of the current sum; combined with the local accumulator bound, this yields a decay of order $a_{q,\delta}^{rt}$ per message. For medium-weight messages this estimate breaks down, since its collision correction outgrows its main term, so we compare the exact rows with independent product rows from $\nu_{t/N,q}$. The comparison costs a factor exponential in $O(r \ln t)$, which the definition of r_0 renders subdominant. For high-weight messages we estimate no moments at all: we show that the distribution of xE is nearly uniform in the L_∞ sense and multiply its maximum point mass by the volume of the set of bad accumulator inputs. The main result of this appendix is the following theorem.

Theorem 13. *Let q be a prime power, $0 < \delta < (q-1)/q$, and $0 < R < 1 - H_q(\delta)$. Choose $\gamma > 0$ such that $\gamma\lambda_{q,\delta} > 1$, and set $n := \lfloor RN \rfloor$ and $t := \lceil \gamma \ln N \rceil$. If $E \leftarrow \mathsf{BP}(n, N, t, \mathbb{F}_q)$, then*

$$\begin{aligned}\Pr_E[\mathsf{d}(EA) \le \delta] \le N^{\,1-\gamma\lambda_{q,\delta}+o(1)} &+ \exp\left(-\Omega\left(\frac{N(\ln\ln N)^3}{\ln N}\right)\right) \\ &+ q^{-N(1-R-H_q(\delta)-o(1))}. \end{aligned} \tag{16}$$

All three terms vanish as $N \to \infty$, so in this parameter regime the EA code has relative distance greater than δ with probability $1 - o(1)$ and attains the tradeoff of Theorem 11. The rest of this appendix proves Theorem 13.

B.3 Upper Bounds on $Q_{N,q}(w, \delta)$

Fix $w \geq 1$. Since A is invertible, $\mathsf{wt}(Y_{w,q}A) \geq 1$, and among the $\binom{N}{w}(q-1)^w$ vectors of $\mathcal{S}_{w,q}^{(N)}$ exactly $A_{w,h}^{N,q}$ satisfy $\mathsf{wt}(yA) = h$. Theorem 12 therefore gives

$$\begin{aligned} Q_{N,q}(w,\delta) &= \frac{1}{\binom{N}{w}(q-1)^w} \sum_{h=1}^{\delta N} A_{w,h}^{N,q} \\ &= \frac{1}{\binom{N}{w}(q-1)^w} \sum_{h=1}^{\delta N} \sum_{i=0}^{w-1} \binom{h-1}{k_i-1} \binom{N-h}{\ell_i} \binom{h-k_i}{i} (q-1)^{k_i} (q-2)^i, \end{aligned} \tag{17}$$

with $k_i = \lceil (w-i)/2 \rceil$ and $\ell_i = \lfloor (w-i)/2 \rfloor$. This subsection derives two upper bounds on Eq. 17. The global upper bound is valid for every w but carrying a polynomial factor in N. The local upper bound is only valid for $w \leq N/2$ in which the polynomial factor is replaced by the milder correction $e^{w^2/N}$.

Both bounds decay geometrically in w at the rate $a_{q,\delta} = f_q(\delta)$, where $f_q(x) := \frac{q-2}{q-1}x + 2\sqrt{\frac{x(1-x)}{q-1}}$ for $0 \leq x \leq (q-1)/q$. We record the properties of f_q that we use.

Lemma 12. *For $0 \leq x \leq (q-1)/q$, we have $f_q(x) = 1 - \big(\sqrt{1-x} - \sqrt{x/(q-1)}\big)^2$. Consequently, f_q is strictly increasing on $[0, (q-1)/q]$, with $f_q(0) = 0$ and $f_q\big((q-1)/q\big) = 1$; in particular, $0 < a_{q,\delta} < 1$ whenever $0 < \delta < (q-1)/q$.*

Proof. Expanding the square gives

$$1 - (1-x) - \frac{x}{q-1} + 2\sqrt{\frac{x(1-x)}{q-1}} = \frac{q-2}{q-1}\,x + 2\sqrt{\frac{x(1-x)}{q-1}} = f_q(x).$$

The function $g(x) := \sqrt{1-x} - \sqrt{x/(q-1)}$ is strictly decreasing on $[0, (q-1)/q]$, and $g\big((q-1)/q\big) = \sqrt{1/q} - \sqrt{1/q} = 0$, so $g \geq 0$ on the whole interval. Hence g^2 is strictly decreasing there, and $f_q = 1 - g^2$ is strictly increasing. The endpoint values follow by direct evaluation.

A direct Stirling approximation of the binomial coefficients in Eq. 17 leads to a rather loose estimate. Instead, the next lemma bounds the hypergeometric-type ratio appearing there by a binomial expression, which keeps the resulting bounds compact.

Lemma 13. *Suppose $0 < h < N$ and $0 \leq w \leq N$. Let $0 \leq m \leq h$, set $\ell := w - m$, and assume $0 \leq \ell \leq N - h$. Set $u := h/N$. Then*

$$\frac{\binom{h}{m}\binom{N-h}{\ell}}{\binom{N}{w}} \leq (N+1)(w+1)\binom{w}{m} u^m (1-u)^\ell.$$

Proof. The cases $w = 0$ and $w = N$ hold trivially. We therefore assume $1 \le w \le N - 1$ and set $\omega := w/N$.

Lemma 7 bounds the numerator from above and the denominator from below:

$$\frac{\binom{h}{m}\binom{N-h}{\ell}}{\binom{N}{w}} \le \frac{e^{h\,h_e(m/h)} \cdot e^{(N-h)\,h_e(\ell/(N-h))}}{\frac{1}{N+1}\,e^{N h_e(\omega)}} = (N+1)\,e^{-NI}, \tag{18}$$

where $I := h_e(\omega) - u\,h_e\big(\frac{m}{h}\big) - (1-u)\,h_e\big(\frac{\ell}{N-h}\big)$. The exponent I is a mutual information. Let the pair (B, S) of $\{0,1\}$-valued random variables have joint distribution

$$\Pr[B=1,S=1] = \frac{m}{N}, \qquad \Pr[B=1,S=0] = \frac{h-m}{N},$$
$$\Pr[B=0,S=1] = \frac{\ell}{N}, \qquad \Pr[B=0,S=0] = \frac{N-h-\ell}{N};$$

the four entries are nonnegative and sum to one by the assumptions on m and ℓ. The marginals are $\Pr[B=1] = \frac{m+(h-m)}{N} = u$ and $\Pr[S=1] = \frac{m+\ell}{N} = \omega$, and the conditionals given B are $\Pr[S=1 \mid B=1] = \frac{m}{h}$ and $\Pr[S=1 \mid B=0] = \frac{\ell}{N-h}$. Measuring entropies in nats, $H(S) = h_e(\omega)$ and $H(S \mid B) = u\,h_e(m/h) + (1-u)\,h_e(\ell/(N-h))$, so $I = H(S) - H(S \mid B) = I(B;S)$.

Expressing the mutual information instead through the conditionals of B given S, namely $\Pr[B=1 \mid S=1] = \frac{m}{w}$ and $\Pr[B=1 \mid S=0] = \frac{h-m}{N-w}$, we obtain

$$I(B;S) = \omega\,D_{\mathrm{KL}}\Big(\frac{m}{w}\,\Big\|\,u\Big) + (1-\omega)\,D_{\mathrm{KL}}\Big(\frac{h-m}{N-w}\,\Big\|\,u\Big) \ge \omega\,D_{\mathrm{KL}}\Big(\frac{m}{w}\,\Big\|\,u\Big)$$

by the nonnegativity of relative entropy. Since $N\omega = w$, this gives

$$e^{-NI} \le e^{-w D_{\mathrm{KL}}(m/w \| u)}. \tag{19}$$

On the other hand, expanding the definitions shows

$$w\,h_e\Big(\frac{m}{w}\Big) + m \ln u + \ell \ln(1-u) = -w\,D_{\mathrm{KL}}\Big(\frac{m}{w}\,\Big\|\,u\Big),$$

so the lower bound of Lemma 7 yields

$$\binom{w}{m} u^m (1-u)^\ell \ge \frac{1}{w+1}\,e^{w h_e(m/w)}\,u^m(1-u)^\ell = \frac{1}{w+1}\,e^{-w D_{\mathrm{KL}}(m/w\|u)}. \tag{20}$$

Combining Eq. 18, Eq. 19, and Eq. 20, the desired inequality yields,

$$\frac{\binom{h}{m}\binom{N-h}{\ell}}{\binom{N}{w}} \le (N+1)\,e^{-w D_{\mathrm{KL}}(m/w\|u)} \le (N+1)(w+1)\binom{w}{m}u^m(1-u)^\ell.$$

Proposition 2. *For $0 < \delta < (q-1)/q$, there is a constant $C_{q,\delta} > 0$ such that $Q_{N,q}(w,\delta) \le C_{q,\delta} N^3 a_{q,\delta}^w$ for all $N \ge 1$ and $0 \le w \le N$.*

Proof. Since $Q_{N,q}(0,\delta) = 1$, the case $w = 0$ holds for any $C_{q,\delta} \ge 1$, so assume $w \ge 1$. Consider the summand of Eq. 17 for fixed $1 \le h \le \delta N$ and $0 \le i \le w-1$, and abbreviate $k := k_i$, $\ell := \ell_i$, and $m := k + i$, so that $m + \ell = w$. Since $\binom{h-1}{k-1} \le \binom{h}{k}$ and $\binom{h}{k}\binom{h-k}{i} = \binom{h}{m}\binom{m}{i}$, we have $\binom{h-1}{k-1}\binom{h-k}{i} \le \binom{h}{m}\binom{m}{i}$. A nonzero

summand moreover requires $m \le h$ and $\ell \le N-h$, and $u := h/N$ satisfies $0 < u \le \delta$. Lemma 13 therefore gives

$$\begin{aligned}
&\frac{\binom{h-1}{k-1}\binom{N-h}{\ell}\binom{h-k}{i}(q-1)^k(q-2)^i}{\binom{N}{w}(q-1)^w} \\
&\quad\le (N+1)(w+1)\binom{w}{m}\binom{m}{i}u^m(1-u)^\ell\,\frac{(q-1)^k(q-2)^i}{(q-1)^w} \\
&\quad= (N+1)(w+1)\binom{w}{k,i,\ell}u^k c^i v^\ell, \qquad (21)
\end{aligned}$$

where $\binom{w}{k,i,\ell} := \frac{w!}{k!\,i!\,\ell!} = \binom{w}{m}\binom{m}{i}$, $c := \frac{(q-2)u}{q-1}$, and $v := \frac{1-u}{q-1}$; the last equality uses $m = k+i$ and $k+i+\ell = w$.

By the definitions of k_i and ℓ_i we have $k - \ell \in \{0,1\}$. If $k = \ell$, then $u^k v^\ell = (\sqrt{uv})^{k+\ell}$; if $k = \ell+1$, then $u^k v^\ell = (\sqrt{uv})^{k+\ell}\sqrt{u/v}$, and $\sqrt{u/v} = \sqrt{\frac{(q-1)u}{1-u}} \le \sqrt{\frac{(q-1)\delta}{1-\delta}}$ because $u \mapsto \frac{u}{1-u}$ is increasing and $u \le \delta$. In both cases, $u^k v^\ell \le C_0(\sqrt{uv})^{k+\ell}$ with $C_0 := \max\bigl\{1, \sqrt{(q-1)\delta/(1-\delta)}\bigr\}$.

The triples (k_i, i, ℓ_i) for $0 \le i \le w-1$ are pairwise distinct and satisfy $k_i + i + \ell_i = w$. Enlarging the sum over i to a sum over all triples of nonnegative integers summing to w and applying the multinomial theorem,

$$\begin{aligned}
\sum_{i=0}^{w-1}\binom{w}{k_i,i,\ell_i}u^{k_i}c^i v^{\ell_i} &\le C_0 \sum_{\substack{k,i,\ell\ge 0\\ k+i+\ell=w}}\binom{w}{k,i,\ell}(\sqrt{uv})^k c^i (\sqrt{uv})^\ell \\
&= C_0\bigl(c + 2\sqrt{uv}\bigr)^w. \qquad (22)
\end{aligned}$$

Moreover, $c + 2\sqrt{uv} = \frac{q-2}{q-1}u + 2\sqrt{\frac{u(1-u)}{q-1}} = f_q(u)$, so Lemma 12 together with $u \le \delta$ gives $c + 2\sqrt{uv} \le f_q(\delta) = a_{q,\delta}$. Combining this bound with Eq. 21 and Eq. 22, the contribution of each fixed h to Eq. 17 satisfies

$$\begin{aligned}
&\frac{1}{\binom{N}{w}(q-1)^w}\sum_{i=0}^{w-1}\binom{h-1}{k_i-1}\binom{N-h}{\ell_i}\binom{h-k_i}{i}(q-1)^{k_i}(q-2)^i \\
&\quad\le C_0(N+1)(w+1)\,a_{q,\delta}^w.
\end{aligned}$$

Summing over the at most N values of h and using $w+1 \le N+1$,

$$Q_{N,q}(w,\delta) \le C_0\,N(N+1)^2 a_{q,\delta}^w \le 4C_0 N^3 a_{q,\delta}^w,$$

which proves the proposition with $C_{q,\delta} := 4C_0$.

Proposition 2 applies to every w, but its polynomial prefactor is too expensive for the low-weight analysis, where the final bound is itself only polynomially small in N. One factor of N enters through Lemma 13; the next lemma removes it when w is small.

Lemma 14. *Suppose $0 \le h \le N$ and $0 \le w < N$. Let $0 \le m \le h$, set $\ell := w-m$, and assume $0 \le \ell \le N-h$. Set $u := h/N$ and $\Xi_N(w) := \exp\bigl(\frac{w(w-1)}{2(N-w+1)}\bigr)$.*

Then, with the convention $0^0 := 1$,

$$\frac{\binom{h}{m}\binom{N-h}{\ell}}{\binom{N}{w}} \le \Xi_N(w)\binom{w}{m}u^m(1-u)^\ell.$$

In particular, $\Xi_N(w) \le e^{w^2/N}$ *whenever* $w \le N/2$.

Proof. The cases $h = 0$ and $h = N$ hold trivially, so we assume $0 < h < N$. Writing $(x)_j := x(x-1)\cdots(x-j+1)$ for the falling factorial and inserting $\binom{h}{m} = \frac{(h)_m}{m!}$, $\binom{N-h}{\ell} = \frac{(N-h)_\ell}{\ell!}$, $\binom{N}{w} = \frac{(N)_w}{w!}$, $\binom{w}{m} = \frac{w!}{m!\,\ell!}$, $u^m = \frac{h^m}{N^m}$, and $(1-u)^\ell = \frac{(N-h)^\ell}{N^\ell}$,

$$\frac{\binom{h}{m}\binom{N-h}{\ell}}{\binom{N}{w}} \cdot \frac{1}{\binom{w}{m}u^m(1-u)^\ell} = \frac{(h)_m}{h^m}\cdot\frac{(N-h)_\ell}{(N-h)^\ell}\cdot\frac{N^w}{(N)_w} \le \frac{N^w}{(N)_w},$$

since the first two factors are at most one. Moreover, by $-\ln(1-x) \le \frac{x}{1-x}$ for $0 \le x < 1$,

$$\begin{aligned}\ln\frac{N^w}{(N)_w} = \sum_{j=0}^{w-1} -\ln\!\left(1-\frac{j}{N}\right) &\le \sum_{j=0}^{w-1}\frac{j}{N-j} \\ &\le \frac{1}{N-w+1}\sum_{j=0}^{w-1} j = \frac{w(w-1)}{2(N-w+1)},\end{aligned}$$

which proves the main inequality. If $w \le N/2$, then $2(N-w+1) \ge N$, so $\Xi_N(w) \le e^{w(w-1)/N} \le e^{w^2/N}$.

Proposition 3. *For* $0 < \delta < (q-1)/q$, *there is a constant* $C_{q,\delta} > 0$ *such that* $Q_{N,q}(w,\delta) \le C_{q,\delta}\, e^{w^2/N} a_{q,\delta}^w$ *for all* $N \ge 1$ *and* $0 \le w \le N/2$.

Proof. The case $w = 0$ is immediate, so assume $w \ge 1$, and keep the notation $k = k_i$, $\ell = \ell_i$, $m = k + i$, $u = h/N$, $c = \frac{(q-2)u}{q-1}$, and $v = \frac{1-u}{q-1}$ from the proof of Proposition 2. In place of the bound $\binom{h-1}{k-1} \le \binom{h}{k}$ used there, we keep the exact identity

$$\binom{h-1}{k-1}\binom{h-k}{i} = \frac{k}{h}\binom{h}{k}\binom{h-k}{i} = \frac{k}{h}\binom{h}{m}\binom{m}{i},$$

which retains the factor $\frac{k}{h} = \frac{k}{Nu}$. Repeating the derivation of Eq. 21 with Lemma 14 in place of Lemma 13 gives

$$\frac{\binom{h-1}{k-1}\binom{N-h}{\ell}\binom{h-k}{i}(q-1)^k(q-2)^i}{\binom{N}{w}(q-1)^w} \le \Xi_N(w)\,\frac{k}{Nu}\binom{w}{k,i,\ell}u^k c^i v^\ell. \tag{23}$$

Bounding $k \le w$ and summing Eq. 23 over i as in Eq. 22, the contribution of each fixed h to Eq. 17 is at most $C_0\,\Xi_N(w)\,\frac{w}{Nu} f_q(u)^w$ with the same constant C_0. Summing over h and writing $u = h/N$,

$$Q_{N,q}(w,\delta) \le C_0\,\Xi_N(w)\cdot\frac{w}{N}\sum_{h=1}^{\delta N}\frac{f_q(h/N)^w}{h/N}.$$

We claim that there is a constant $C_1 = C_1(q,\delta)$ such that, for all $N \ge 1$ and all $w \ge 1$,

$$\frac{w}{N}\sum_{h=1}^{\delta N}\frac{f_q(h/N)^w}{h/N} \le C_1\, a_{q,\delta}^w. \tag{24}$$

Granting the claim, Lemma 14 gives, for $1 \le w \le N/2$,

$$Q_{N,q}(w,\delta) \le C_0 C_1\, \Xi_N(w)\, a_{q,\delta}^w \le C_0 C_1\, e^{w^2/N} a_{q,\delta}^w,$$

so the proposition holds with $C_{q,\delta} := \max\{C_0C_1, 1\}$.

It remains to prove Eq. 24. Define $\psi_q(x) := \frac{x f_q'(x)}{f_q(x)}$ for $0 < x \le \delta$. As $x \to 0$ we have $f_q(x) = 2\sqrt{x/(q-1)} + O(x)$ and $f_q'(x) = \frac{1}{\sqrt{(q-1)x}} + O(1)$, so $\psi_q(x) \to \frac{1}{2}$, and we extend ψ_q continuously to $[0,\delta]$ by $\psi_q(0) := \frac{1}{2}$. In the notation of Lemma 12, $f_q' = -2gg' > 0$ on $(0,\delta]$, so ψ_q is continuous and positive on the compact interval $[0,\delta]$ and attains a minimum $b_{q,\delta} > 0$.

Suppose first that $w \ge 1/b_{q,\delta}$, and set $G_w(x) := \frac{f_q(x)^w}{x}$. Then

$$\frac{d}{dx}\ln G_w(x) = \frac{w f_q'(x)}{f_q(x)} - \frac{1}{x} = \frac{w\psi_q(x) - 1}{x} \ge 0 \qquad \text{on } (0,\delta],$$

so G_w is nondecreasing there. Hence $\frac{1}{N}G_w(h/N) \le \int_{h/N}^{(h+1)/N} G_w(x)\,dx$ for every $1 \le h < \lfloor \delta N \rfloor$, while the term $h = \lfloor \delta N \rfloor$ is at most $\frac{1}{N}G_w(\delta)$. Summing from 0 to δ yields:

$$\frac{1}{N}\sum_{h=1}^{\delta N} G_w(h/N) \le \int_0^\delta G_w(x)\,dx + \frac{G_w(\delta)}{N}.$$

For the integral, the substitution $d\big(f_q(x)^w\big) = w f_q(x)^{w-1} f_q'(x)\,dx$ gives

$$w\int_0^\delta \frac{f_q(x)^w}{x}\,dx = \int_0^\delta \frac{1}{\psi_q(x)}\, d\big(f_q(x)^w\big) \le \frac{f_q(\delta)^w}{b_{q,\delta}} = \frac{a_{q,\delta}^w}{b_{q,\delta}},$$

while $\frac{w}{N}G_w(\delta) = \frac{w}{N}\cdot\frac{a_{q,\delta}^w}{\delta} \le \frac{a_{q,\delta}^w}{\delta}$ because $w \le N$. Hence Eq. 24 holds with $C_1 = b_{q,\delta}^{-1} + \delta^{-1}$ whenever $w \ge 1/b_{q,\delta}$.

For the finitely many remaining integers $1 \le w < 1/b_{q,\delta}$, we use the bound $f_q(x) \le C_q\sqrt{x}$ on $[0,1]$ with $C_q := \frac{q-2}{q-1} + \frac{2}{\sqrt{q-1}}$, which follows from $x \le \sqrt{x}$ and $1 - x \le 1$. Since $\frac{w}{2} - 1 \ge -\frac{1}{2}$ and $\sum_{h=1}^{M} h^{-1/2} \le 2\sqrt{M}$,

$$\begin{aligned}\frac{w}{N}\sum_{h=1}^{\delta N}\frac{f_q(h/N)^w}{h/N} &\le \frac{wC_q^w}{N}\sum_{h=1}^{\delta N}\left(\frac{h}{N}\right)^{w/2-1}\\ &\le \frac{wC_q^w}{\sqrt{N}}\sum_{h=1}^{\delta N}\frac{1}{\sqrt{h}} \le 2wC_q^w\sqrt{\delta},\end{aligned}$$

uniformly in N. Enlarging C_1 so that additionally $C_1 \ge 2wC_q^w a_{q,\delta}^{-w}$ for these finitely many w proves Eq. 24 in general.

B.4 Analysis for Low-Weight Messages

Let $\mathsf{Bad}_{\mathrm{low}}$ denote the event that some $x \in \mathbb{F}_q^n$ with $1 \le \mathsf{wt}(x) \le r_0$ satisfies $\mathsf{wt}(xEA) \le \delta N$. A union bound over the message weights gives

$$\Pr_E[\mathsf{Bad}_{\mathrm{low}}] \le \sum_{r=1}^{r_0} \binom{n}{r} (q-1)^r \max_{\mathsf{wt}(x)=r} \Pr_E[\mathsf{wt}(xEA) \le \delta N].$$

For a message of weight $r \le r_0$ we have $W_r \le rt \le \varepsilon_N N$, which is below $N/2$ for all large N, so the reduction Eq. 15 can be combined with the local bound of Proposition 3. What remains is the moment $\mathbb{E}[a_{q,\delta}^{W_r}]$ under the exact-weight distribution, which we estimate by exposing the rows one at a time and controlling, the overlap of each new row with the support of the current sum.

Lemma 15. *Let $U_1, \dots, U_r$ be independent samples from $\mu_{t,q}$, put $W_0 := 0$ and $W_j := \mathsf{wt}(U_1 + \cdots + U_j)$ for $1 \le j \le r$, and let $a \in (0,1)$. Define $s_{q,a} := \frac{(q-2)a^{-1}+a^{-2}}{q-1}$. Then $s_{q,a} > 1$ and*

$$\mathbb{E}[a^{W_r}] \le a^{rt} \exp\left(\frac{(s_{q,a}-1)\, r(r-1)\, t^2}{2N} \right). \tag{25}$$

Proof. Since $0 < a < 1$, both a^{-1} and a^{-2} exceed one, so $s_{q,a} > \frac{(q-2)+1}{q-1} = 1$.

Fix $0 \le j < r$. Suppose that the first j rows $U_1, \dots, U_j$ have already been sampled. We write $w := W_j$ for the weight of the current sum. The row U_{j+1} places independent uniform values of $\mathbb{F}_q^\times$ on a uniform support T of size t, independently of the past. Let $L := |T \cap \mathrm{supp}(U_1 + \cdots + U_j)|$. Then $L \sim \mathrm{Hyp}(N, w, t)$. Adding U_{j+1} changes the current sum as follows:

1. Each of the $t - L$ coordinates in T outside the current support becomes nonzero.
2. At each of the L coordinates in the intersection, the current sum has some value $s \in \mathbb{F}_q^\times$. The new row contributes an independent uniform value $\xi \in \mathbb{F}_q^\times$. Hence, the coordinate becomes zero exactly when $\xi = -s$, which occurs with probability $1/(q-1)$. These events are independent across the L coordinates.
3. Every coordinate outside T remains unchanged.

Hence, conditionally on L, we have $W_{j+1} = w + (t - L) - Z$ with $Z \leftarrow \mathsf{Bin}\big(L, \frac{1}{q-1}\big)$, and the binomial moment formula gives

$$\mathbb{E}\big[a^{W_{j+1}} \mid U_1, \dots, U_j, L\big] = a^{w+t-L} \left(\frac{q-2}{q-1} + \frac{a^{-1}}{q-1} \right)^L = a^{w+t}\, s_{q,a}^L.$$

Averaging over L with Lemma 9 and using the deterministic bound $W_j \le jt$,

$$\begin{aligned} \mathbb{E}\big[a^{W_{j+1}} \mid U_1, \dots, U_j\big] &\le a^{W_j}\, a^t \exp\left((s_{q,a}-1) \frac{tW_j}{N} \right) \\ &\le a^{W_j}\, a^t \exp\left((s_{q,a}-1) \frac{jt^2}{N} \right). \end{aligned}$$

Taking expectations of both sides yields:

$$\begin{aligned}
&\mathbb{E}[a^{W_1}] \le a^t \mathbb{E}[a^{W_0}] \\
&\mathbb{E}[a^{W_2}] \le a^t \exp\big((s_{q,a}-1)t^2/N\big)\,\mathbb{E}[a^{W_1}] \\
&\cdots \\
&\mathbb{E}[a^{W_r}] \le a^t \exp\big((s_{q,a}-1)(r-1)t^2/N\big)\,\mathbb{E}[a^{W_{r-1}}]
\end{aligned}$$

Since $E[a^{W_0}] = 1$ and $\sum_{j=0}^{r-1} j = \frac{r(r-1)}{2}$, iterating over all $j = 0, \ldots, r-1$ immediately gets

$$\mathbb{E}[a^{W_r}] \le a^{rt} \exp\left((s_{q,a}-1)\frac{t^2}{N}\sum_{j=0}^{r-1} j\right),$$

which is Eq. 25.

Proposition 4. *Fix $0 < \delta < (q-1)/q$. There is a constant $C_{q,\delta} > 0$ such that every nonzero $x \in \mathbb{F}_q^n$ of weight r with $rt \le N/2$ satisfies*

$$\Pr_E[\mathsf{wt}(xEA) \le \delta N] \le C_{q,\delta} \exp\left(-\lambda_{q,\delta}\, rt + C_{q,\delta}\frac{(rt)^2}{N}\right). \tag{26}$$

Proof. As observed in Section B.2, $xE \stackrel{d}{=} U_1 + \cdots + U_r$ for independent $U_i \leftarrow \mu_{t,q}$, so $W_r \le rt \le N/2$. Writing C' for the constant of Proposition 3, Eq. 15 gives

$$\begin{aligned}
\Pr_E[\mathsf{wt}(xEA) \le \delta N] = \mathbb{E}[Q_{N,q}(W_r, \delta)] &\le C'\,\mathbb{E}\big[e^{W_r^2/N} a_{q,\delta}^{W_r}\big] \\
&\le C'\, e^{(rt)^2/N}\,\mathbb{E}\big[a_{q,\delta}^{W_r}\big].
\end{aligned}$$

Lemma 15 with $a = a_{q,\delta}$ and the bound $r(r-1)t^2 \le (rt)^2$ yield

$$\mathbb{E}\big[a_{q,\delta}^{W_r}\big] \le \exp\left(-\lambda_{q,\delta}\, rt + \frac{s_{q,a_{q,\delta}} - 1}{2} \cdot \frac{(rt)^2}{N}\right),$$

and taking $C_{q,\delta} := \max\big\{C', 1 + \frac{s_{q,a_{q,\delta}}-1}{2}\big\}$ proves Eq. 26.

For $r \le r_0$ and all large N we have $rt \le \varepsilon_N N \le N/2$, so Proposition 4 applies. Moreover, since $(rt)^2/N \le \varepsilon_N\, rt$, abbreviating $\lambda_N := \lambda_{q,\delta} - C_{q,\delta}\varepsilon_N$ yields

$$\Pr_E[\mathsf{wt}(xEA) \le \delta N] \le C_{q,\delta}\, e^{-(\lambda_{q,\delta} - C_{q,\delta}\varepsilon_N)rt} = C_{q,\delta}\, e^{-\lambda_N rt}.$$

Using $\binom{n}{r} \le n^r \le N^r$ and extending the sum to all $r \ge 1$,

$$\Pr_E[\mathsf{Bad}_{\mathrm{low}}] \le C_{q,\delta} \sum_{r=1}^{r_0} \binom{n}{r} (q-1)^r e^{-\lambda_N rt} \le C_{q,\delta} \sum_{r \ge 1} \big((q-1)Ne^{-\lambda_N t}\big)^r.$$

The logarithm of the base is

$$\ln(q-1) + \ln N - \lambda_{q,\delta} t + C_{q,\delta}\varepsilon_N t = \big(1 - \gamma\lambda_{q,\delta} + o(1)\big)\ln N,$$

since $t = \gamma \ln N + O(1)$ and $\varepsilon_N t = (\ln t)^3 = O\big((\ln\ln N)^3\big) = o(\ln N)$. As $\gamma\lambda_{q,\delta} > 1$, the base tends to zero; in particular, it is at most $\frac{1}{2}$ for all large N, so the geometric series is at most twice its first term. Therefore

$$\Pr_E[\mathsf{Bad}_{\mathrm{low}}] \le N^{1 - \gamma\lambda_{q,\delta} + o(1)},$$

which is the first term of Eq. 16.

B.5 Analysis for Medium-Weight Messages

Let $\mathsf{Bad}_{\mathrm{med}}$ denote the event that some $x \in \mathbb{F}_q^n$ with $r_0 < \mathsf{wt}(x) < r_1$ satisfies $\mathsf{wt}(xEA) \le \delta N$; as before,

$$\Pr_E[\mathsf{Bad}_{\mathrm{med}}] \le \sum_{r=r_0+1}^{r_1-1} \binom{n}{r} (q-1)^r \max_{\mathsf{wt}(x)=r} \Pr_E[\mathsf{wt}(xEA) \le \delta N].$$

The low-weight argument does not extend to this range. Combining Eq. 15 with the global bound of Proposition 2 still reduces the fixed-message probability to $C_{q,\delta} N^3 \,\mathbb{E}[a_{q,\delta}^{W_r}]$, and the polynomial factor is harmless here; the problem is the moment estimate itself. Near the upper endpoint $r = O(r_1)$, the main term of Eq. 25 has order $rt = O(N \ln t)$ in the exponent, while the collision correction has order $(rt)^2/N = O(N(\ln t)^2)$ and overwhelms it. We therefore replace the fixed-weight rows with independent product rows drawn from $\nu_{\eta,q}$, where $\eta := t/N$. Lemma 11 allows us to transfer events between these two row models with a multiplicative cost of $b_{N,t}^{-r} = e^{O(r \ln t)}$. Under the product-row model, the coordinates of the sum are independent. This allows us to compute the moment exactly. The resulting negative exponent has order rt. By the choice of r_0, we have $rt \ge \varepsilon_N N$, so this negative exponent dominates the transfer cost.

We first identify the distribution of a sum of product rows.

Lemma 16. *Let $0 \le \eta \le (q-1)/q$ and let $V_1, \dots, V_r$ be independent samples from $\nu_{\eta,q}$. Define $\beta := 1 - \frac{q}{q-1}\eta$, $\rho_r := \frac{q-1}{q}(1-\beta^r)$, and $c_q := \frac{q-1}{q}\big(1 - e^{-q/(q-1)}\big)$. Then $V_1 + \cdots + V_r \leftarrow \nu_{\rho_r,q}$ and $\rho_r \ge c_q \min\{\eta r, 1\}$; in particular, $\mathsf{wt}(V_1 + \cdots + V_r) \leftarrow \mathsf{Bin}(N, \rho_r)$.*

Proof. It suffices to analyze a single coordinate. Let $X_1, \dots, X_r$ be the entries of $V_1, \dots, V_r$ at a fixed coordinate. We put $S_r := X_1 + \cdots + X_r$ and $\pi := \frac{q}{q-1}\eta = 1 - \beta$, so that $\pi \in [0,1]$. A sample from the coordinate distribution of $\nu_{\eta,q}$ can be generated by drawing $B \leftarrow \mathsf{Ber}(\pi)$. If $B = 1$, output a uniform element of $\mathbb{F}_q$, otherwise output 0.

Indeed, this procedure outputs 0 with probability $1 - \pi + \frac{\pi}{q} = 1 - \eta$ and each $a \in \mathbb{F}_q^\times$ with probability $\frac{\pi}{q} = \frac{\eta}{q-1}$.

Generate $X_1, \dots, X_r$ in this way with independent coins $B_1, \dots, B_r$. Suppose $B_i = 1$ for some i. Suppose $B_i = 1$ for some i. Fix the values of all the coins and of the summands $(X_j)_{j \ne i}$. By independence, X_i remains uniform on $\mathbb{F}_q$. Adding a fixed element of $(\mathbb{F}_q, +)$ to a uniform element again yields a uniform element. Therefore, S_r is uniform on $\mathbb{F}_q$ whenever at least one coin equals one. This event has probability $1 - (1-\pi)^r = 1 - \beta^r$.

If all the coins equal zero, then $S_r = 0$. Consequently, $\Pr[S_r = 0] = \beta^r + \frac{1-\beta^r}{q} = \frac{1+(q-1)\beta^r}{q}$ and $\Pr[S_r = a] = \frac{1-\beta^r}{q}$ for each $a \in \mathbb{F}_q^\times$, which is exactly the coordinate distribution of $\nu_{\rho_r,q}$, as $\frac{\rho_r}{q-1} = \frac{1-\beta^r}{q}$. Independence across coordinates proves the first claim, and the weight statement follows from the definition of $\nu_{\rho_r,q}$.

For the lower bound on ρ_r, put $\alpha := \frac{q}{q-1}$. From $0 \le \beta \le 1$ and $1-x \le e^{-x}$ we get $\rho_r \ge \frac{q-1}{q}(1-e^{-\alpha\eta r})$. If $\eta r \ge 1$, the right-hand side is at least $\frac{q-1}{q}(1-e^{-\alpha}) = c_q$. If $\eta r \le 1$, the concavity of $s \mapsto 1-e^{-\alpha s}$ gives $1-e^{-\alpha\eta r} \ge (1-e^{-\alpha})\,\eta r$, therefore $\rho_r \ge c_q\,\eta r$.

Proposition 5. *Fix $0 < \delta < (q-1)/q$, and let $\eta = t/N$ and ρ_r be as in Lemma 16. There is a constant $C_{q,\delta} > 0$ such that, for all sufficiently large N, every nonzero $x \in \mathbb{F}_q^n$ of weight r satisfies*

$$\Pr_E\big[\mathsf{wt}(xEA) \le \delta N\big] \le C_{q,\delta}\, N^3 \left(\frac{20}{3}\sqrt{t}\right)^r \exp\big(-\kappa_{q,\delta}\,\rho_r N\big). \tag{27}$$

Proof. By scalar invariance, $xE \overset{d}{=} U_1 + \cdots + U_r$ for independent $U_i \leftarrow \mu_{t,q}$, and for all large N we have $t \le N/2$, so Lemma 11 gives

$$\Pr_E\big[\mathsf{wt}(xEA) \le \delta N\big] \le b_{N,t}^{-r}\, \Pr\big[\mathsf{wt}\big((V_1 + \cdots + V_r)A\big) \le \delta N\big]$$

for independent $V_i \leftarrow \nu_{\eta,q}$. By Lemma 16, the sum $V_1 + \cdots + V_r$ follows $\nu_{\rho_r,q}$, so its weight W is $\mathsf{Bin}(N,\rho_r)$-distributed and, conditional on $W = w \ge 1$, the sum is uniform on $\mathcal{S}_{w,q}^{(N)}$; exactly as in Eq. 15, the probability on the right-hand side equals $\mathbb{E}[Q_{N,q}(W,\delta)]$. Proposition 2, the binomial moment formula, and $1-x \le e^{-x}$ then give

$$\begin{aligned}
\mathbb{E}[Q_{N,q}(W,\delta)] &\le C_{q,\delta} N^3\, \mathbb{E}\big[a_{q,\delta}^W\big] \\
&= C_{q,\delta} N^3 \sum_{w=0}^{N} \binom{N}{w} \rho_r^w (1-\rho_r)^{N-w} a_{q,\delta}^w \\
&= C_{q,\delta} N^3 \big(1 - \rho_r + \rho_r a_{q,\delta}\big)^N \\
&= C_{q,\delta} N^3 \big(1 - \kappa_{q,\delta}\rho_r\big)^N \le C_{q,\delta} N^3 e^{-\kappa_{q,\delta}\rho_r N}.
\end{aligned}$$

Finally, Lemma 10 bounds the transfer cost by $b_{N,t}^{-r} \le \big(\frac{20}{3}\sqrt{t}\big)^r$, which proves Eq. 27.

Now let $r_0 < r < r_1$. Lemma 16 gives $\rho_r N \ge c_q \min\{rt, N\}$. One can find that both branches of the minimum are at least $\varepsilon_N N$ since $rt \ge (r_0+1)t > \varepsilon_N N$ and $\varepsilon_N \le 1$ for large N.

By Proposition 5, each fixed message of medium weight is therefore bad with probability at most $C_{q,\delta} N^3 \big(\frac{20}{3}\sqrt{t}\big)^r \cdot e^{-c_q \kappa_{q,\delta} \varepsilon_N N}$.

The counting factors are subdominant. Throughout the medium range, $r > \varepsilon_N N/t = N(\ln t)^3/t^2$, so $\frac{en}{r} \le \frac{eRt^2}{(\ln t)^3}$, and hence, using $\binom{n}{r} \le (en/r)^r$ and $r < r_1 \le \frac{N \ln t}{t} + 1$,

$$\begin{aligned}
\ln\left[\binom{n}{r}(q-1)^r \left(\frac{20}{3}\sqrt{t}\right)^r\right] &\le r\left(\ln\frac{en}{r} + \ln(q-1) + \ln\frac{20\sqrt{t}}{3}\right) \\
&= O_q(r \ln t) = O_q\left(\frac{N(\ln t)^2}{t}\right).
\end{aligned}$$

Since $\varepsilon_N N = N(\ln t)^3/t$, the negative exponent $c_q \kappa_{q,\delta} \varepsilon_N N$ exceeds this by a factor of order $\ln t$. The prefactor $C_{q,\delta} N^3$ and the number of summands (at

most N) add only a further $O(\ln N) = o(\varepsilon_N N)$ to the logarithm. Hence, for all sufficiently large N,

$$\Pr_E[\mathsf{Bad}_{\mathrm{med}}] \le \exp\!\left(-\frac{c_q \kappa_{q,\delta}}{2} \cdot \frac{N(\ln t)^3}{t}\right), \tag{28}$$

which accounts for the second term of Eq. 16.

B.6 Analysis for High-Weight Messages

Let $\mathsf{Bad}_{\mathrm{high}}$ denote the event that some $x \in \mathbb{F}_q^n$ with $r_1 \le \mathsf{wt}(x) \le n$ satisfies $\mathsf{wt}(xEA) \le \delta N$. In this range we estimate no moments. Instead, we show that once $r \ge r_1$ the distribution $\mu_{t,q}^{*r}$ of the expanded vector xE is nearly uniform in the L_∞ sense, and we bound the probability of a bad output by the cardinality of the bad set times the maximum point mass.

Lemma 17. *Suppose $t \to \infty$ and $t = o(N)$ as $N \to \infty$. Then $\|\mu_{t,q}^{*r}\|_\infty \le q^{-N+o(N)}$ for every $r \ge r_1$, where the $o(N)$ term is uniform in r.*

Proof. Since $\mu_{t,q}^{*r} = \mu_{t,q}^{*(r-r_1)} * \mu_{t,q}^{*r_1}$ and convolution contracts the L_∞ norm, it suffices to prove the claim for $r = r_1$.

Let $\eta := t/N$, let β be as in Lemma 16, and let $V_1, \dots, V_{r_1}$ be independent samples from $\nu_{\eta,q}$. For all large N we have $t \le N/2$. Then for every $z \in \mathbb{F}_q^N$, Lemma 11 gives

$$\mu_{t,q}^{*r_1}(z) \le b_{N,t}^{-r_1} \Pr[V_1 + \cdots + V_{r_1} = z]. \tag{29}$$

By Lemma 16, the coordinates of $V_1 + \cdots + V_{r_1}$ are independent, each with mass $1 - \rho_{r_1} = \frac{1+(q-1)\beta^{r_1}}{q}$ at zero and $\frac{\rho_{r_1}}{q-1} = \frac{1-\beta^{r_1}}{q}$ at each nonzero value. Since $\eta \to 0$, we have $\beta \ge 0$ for all large N, so the largest coordinate mass is the one at zero, and

$$\Pr[V_1 + \cdots + V_{r_1} = z] \le \left(\frac{1 + (q-1)\beta^{r_1}}{q}\right)^N.$$

Moreover, $\eta r_1 = \frac{t}{N}\left\lceil \frac{N \ln t}{t} \right\rceil \ge \ln t$, so $\beta^{r_1} \le e^{-\frac{q}{q-1}\eta r_1} \le t^{-q/(q-1)}$. Combining these bounds with Eq. 29 and Lemma 10,

$$\|\mu_{t,q}^{*r_1}\|_\infty \le \left(\frac{20}{3}\sqrt{t}\right)^{r_1} q^{-N}\left(1 + (q-1)t^{-q/(q-1)}\right)^N = q^{-N} e^{o(N)},$$

where the last step uses $r_1 \ln\!\left(\frac{20}{3}\sqrt{t}\right) = O\!\left(\frac{N(\ln t)^2}{t}\right) = o(N)$ and $N \ln\!\left(1 + (q-1)t^{-q/(q-1)}\right) \le (q-1)Nt^{-q/(q-1)} = o(N)$.

Let $\mathcal{B}_{\delta N,q} := \{y \in \mathbb{F}_q^N : \mathsf{wt}(yA) \le \delta N\}$. Since A is invertible, $y \mapsto yA$ is a bijection of $\mathbb{F}_q^N$, so Lemma 8 gives $|\mathcal{B}_{\delta N,q}| = V_q(N, \delta N) \le q^{NH_q(\delta)}$. For a fixed message x of weight $r \ge r_1$, scalar invariance gives $xE \leftarrow \mu_{t,q}^{*r}$, so the point-mass bound of Section B.1 and Lemma 17 yield

$$\begin{aligned}\Pr_E\!\left[\mathsf{wt}(xEA) \le \delta N\right] = \mu_{t,q}^{*r}\!\left(\mathcal{B}_{\delta N,q}\right) &\le \left|\mathcal{B}_{\delta N,q}\right| \cdot \left\|\mu_{t,q}^{*r}\right\|_\infty \\ &\le q^{-N(1-H_q(\delta)-o(1))}.\end{aligned}$$

There are fewer than $q^n \le q^{RN}$ nonzero messages in total, so

$$\Pr_E[\mathsf{Bad}_{\text{high}}] \le q^{-N(1-R-H_q(\delta)-o(1))},$$

and the exponent $1 - R - H_q(\delta) - o(1)$ is positive for all large N because $R < 1 - H_q(\delta)$. This is the third term of Eq. 16.

B.7 Completing the Proof

Proof (Proof of Theorem 13). Since $1 \le r_0 < r_1 < n$ for all sufficiently large N, every nonzero message weight falls into the low, medium, or high range, and the event $\mathsf{d}(EA) \le \delta$ is contained in $\mathsf{Bad}_{\text{low}} \cup \mathsf{Bad}_{\text{med}} \cup \mathsf{Bad}_{\text{high}}$. Combing the estimates for the three weight ranges gives

$$\Pr_{E \leftarrow \mathsf{BP}(n,N,t,\mathbb{F}_q)}\big[\mathsf{d}(EA) \le \delta\big] \le N^{1-\gamma\lambda_{q,\delta}+o(1)} + \exp\left(-\frac{c_q \kappa_{q,\delta}}{2} \cdot \frac{N(\ln t)^3}{t}\right) + q^{-N(1-R-H_q(\delta)-o(1))}.$$

Since $t = \gamma \ln N + O(1)$, the middle exponent is $\Omega\big(\frac{N(\ln \ln N)^3}{\ln N}\big)$, which matches the second term of Eq. 16. The assumptions $\gamma\lambda_{q,\delta} > 1$ and $R < 1 - H_q(\delta)$ make the first and third exponents negative, so the failure probability is $o(1)$.