

A Linear Proximity Gap for Batched Reed–Solomon Codes over Galois Rings

Luis Angel Gonzalez

Rexbit Exchange Research
`angel@rexbit.exchange`

Abstract. A proximity gap controls how often a random combination of fixed words can be close to a Reed–Solomon code without the words sharing a common agreement set with codewords. We prove such a gap over $R = \text{GR}(2^k, m)$ at every $\delta < (1 - \rho)/2$. For a degree- s curve on an exceptional evaluation domain of length ℓ , more than $s\ell$ good residue parameters force common agreement outside at most $\lfloor \delta\ell \rfloor$ positions. Independent affine batching has bad-challenge probability at most $\ell/2^m$, independent of the batch width. The proof reduces each instance to the residue field, lifts decoded coefficients, and divides the residual through the injection between successive quotient rings. Puncturing lowers the field threshold by the same amount spent on excluded challenges, so neither an additional $s\delta\ell$ term nor a factor k is needed. We also prove a weighted version and extend the probability bounds to challenges with uniform residue, including full-ring sampling. The only external proximity input is the finite-field theorem; the argument does not use a Galois-ring Polishchuk–Spielman lemma. Applying these algebraic bounds to a commitment or folding protocol requires its separate consistency, extraction and transcript reductions.

Keywords: Reed–Solomon codes · proximity gaps · correlated agreement · Galois rings · unique decoding · p -adic descent

Table of Contents

A Linear Proximity Gap for Batched Reed–Solomon Codes over Galois Rings	1
<i>Luis Angel Gonzalez</i>	

1 Introduction

Succinct non-interactive arguments increasingly target arithmetic over the integers modulo a power of two, $\mathbb{Z}_{2^k}$, the native arithmetic of commodity processors and of much secure-computation tooling. A productive route to proving statements over $\mathbb{Z}_{2^k}$ lifts them to a *Galois ring* $R = \text{GR}(2^k, m)$, a finite local ring with residue field $\mathbb{F}_{2^m}$ that unifies $\mathbb{Z}_{2^k}$ ($m = 1$) and the binary field $\mathbb{F}_{2^m}$ ($k = 1$). Recent polynomial commitments and transparent arguments operate directly over Galois rings, mirroring the field-based pipeline built on Reed–Solomon proximity testing.

The workhorse of that pipeline is the *fold*: a length-halving reduction iterated $\log n$ times. Its binding soundness rests on a *proximity gap* with *correlated agreement*. In the batched form one takes words $g_0, \dots, g_s \in R^E$ and forms the degree- s curve

$$w(\alpha) = \sum_{j=0}^s \alpha^j g_j, \quad (1)$$

testing $w(\alpha)$ for proximity to a Reed–Solomon code at a random challenge α . The proximity gap asserts that if $w(\alpha)$ is close to the code for many challenges, then the components $g_0, \dots, g_s$ agree, *on one common coordinate set*, with codewords. The governing quantity is the number of “bad” challenges the test tolerates: the smaller it is relative to the challenge space, the better the binding error.

The gap between field and ring. Over a finite field $\mathbb{F}_q$, at the unique-decoding radius $\delta < (1 - \rho)/2$, the proximity gap for the curve (1) is *linear*: the bad set has size at most sn [2, Thm. 1.5]. Over a Galois ring, the published analysis of [7] states a quadratic threshold $s\ell^2$. For the illustrative parameters

$$(s, \ell, m) = (64, 2^{21}, 155) \quad (2)$$

the quadratic expression is 2^{-107} , while the curve bound proved here is 2^{-128} and the affine bound is 2^{-134} . These are upper bounds on a specified bad-challenge event. They are not matching attacks, necessary error floors, measurements, or complete system security levels.

This work. We prove the linear threshold by a *2-adic descent* (Section 4; Figure 1). Section 3 separately examines the numerical slack in the published bivariate condition; that examination does not discharge its ring divisibility hypotheses. Our proof avoids that step. Reducing modulo 2 lands in the residue field $\mathbb{F}_{2^m}$, where correlated agreement at unique decoding is linear; the field solution lifts to R by exact division by 2, legitimate because

$$\text{ann}_R(2^t) = 2^{k-t}R, \quad (3)$$

and an induction on the residue length closes the argument. Three accounting lemmas keep each punctured instance inside the unique-decoding radius and make the k layers share a single agreement-set budget, so no factor of k appears.

1.1 Contributions

We work over $R = \text{GR}(2^k, m)$, $k \geq 1$, with maximal ideal $\mathfrak{m} = (2)$, residue field $\mathbb{F}_{2^m}$, Teichmüller challenge set $\mathcal{T}$, and an exceptional evaluation domain $E \subseteq R$, $|E| = \ell$.

- (C1) **A linear proximity gap at the unique-decoding radius (Theorem 6).** For the curve (1) and $\delta < (1 - \rho)/2$, more than $L = s\ell$ challenges $\alpha \in \mathcal{T}$ with $\Delta(w(\alpha), \text{RS}[R, E, \rho]) \leq \delta\ell$ force correlated agreement of $g_0, \dots, g_s$ on a common set of size at least $(1 - \delta)\ell$, with proximity loss 0. The threshold is *linear* in ℓ and *independent of k* , sharpening the quadratic $s\ell^2$ of [7]. A variant batching the words as an affine space (Theorem 11) removes the s -factor from this algebraic upper bound.
- (C2) **An independent ring proof (Sections 3 and 4).** The descent cancels only units or uses the injection between successive quotient rings. It therefore avoids the primitive-reduction and zero-divisor cancellation obligations of a bivariate ring proof.
- (C3) **The 2-adic descent and its accounting (Section 4).** We reduce the ring statement to the field statement layer by layer and supply the three lemmas that make the reduction tight: *budget reduction* (Lemma 7; the punctured instance stays at the full radius $(1 - \rho)/2$), *cancellation challenges* (Lemma 9; at most s challenges per carved coordinate, so the k layers share one budget), and the *coefficient-lift carry* (Lemma 8; the lifted residual is again a degree- $< \rho\ell$ codeword that composes across layers, using only that reduction is a ring homomorphism).
- (C4) **Weighted and distributional forms.** Theorem 18 carries rational coordinate weights through the same descent. Corollary 14 extends the unweighted bounds to uniform-residue challenges. Section 5 states the separate protocol obligations needed before these bounds imply binding soundness.

Table 1 positions the result. We are explicit about what is and is not new (Section 1.3): a linear gap for *general linear codes* over Galois rings is known for a *single* random combination with the challenge drawn from the entire ring [8];

Table 1. Correlated-agreement proximity gaps for the batched combination of $s + 1$ words. “Challenge” is the set the random coefficient ranges over; ℓ is the block length; k is the ring length ($\text{GR}(2^k, m)$). Thresholds count parameters, including vectors in the affine row, where $q = 2^m$. The unique-decoding (UD) radius is $(1 - \rho)/2$; the Johnson radius is $1 - \sqrt{\rho}$.

Setting	Object	Challenge	Radius	Threshold
Field [2]	curve $\sum \alpha^j g_j$	$\mathbb{F}_q$	UD	$O(s\ell)$
Galois ring [8]	line $g_0 + \alpha g_1$	full ring	Johnson	$O(1)$ in ℓ
Galois ring [7]	curve $\sum \alpha^j g_j$	$\mathcal{T}$	UD	$s\ell^2$
This work	curve $\sum \alpha^j g_j$	$\mathcal{T}$	UD	$s\ell$
This work (Theorem 11)	affine space	$\mathcal{T}^s$	UD	ℓq^{s-1}

that bound is vacuous once the challenge is restricted to the Teichmüller set. Our statement is for the batched degree- s curve with the challenge in $\mathcal{T}$, at the unique-decoding radius, and it is what improves the batched bound of [7].

1.2 Technical overview

The finite-field curve theorem of [2] supplies the count $s\ell$ at unique decoding. Its field proof uses Berlekamp–Welch interpolation and a divisibility lemma. We use the resulting field theorem as an explicit external input, rather than extending those divisibility operations to a ring with zero-divisors.

Our descent avoids the bivariate object entirely (Figure 1). Reducing modulo 2 maps the instance into $\mathbb{F}_{2^m}$, where the field gap applies and yields field codewords and a common agreement set whose complement W_0 has size at most $\delta\ell$. We lift each field codeword by a fixed choice of its $< \rho\ell$ coefficients; on the agreement set the residual is divisible by 2, and dividing by 2 produces a word over $R/2^{k-1}R$ to which the same argument applies. The subtlety is bookkeeping. A naive accounting re-spends the full budget $\delta\ell$ at each of the k layers and degrades the radius to $(1 - \rho)/3$. We instead track the *new* disagreement set W_l carved at layer l ; the per-layer error budget reduces to $e_l - |W_l|$, which keeps every punctured instance at the full radius $(1 - \rho)/2$, and the discarded challenges are bounded by $s|W_l|$. Both totals telescope:

$$\sum_{l=0}^{k-1} |W_l| \leq \delta\ell \quad \text{and} \quad \sum_{l=0}^{k-1} s|W_l| \leq s\delta\ell, \quad (4)$$

neither carrying a factor of k . At layer l , the field threshold is $s(\ell - b_l)$, so $|S_0| > s\ell$ and $|S_l| \geq |S_0| - sb_l$ keep the field input applicable without adding $s\delta\ell$ to the initial threshold.

1.3 Related work

Reed–Solomon proximity testing originates with FRI [1]. Proximity gaps and correlated agreement up to the Johnson radius over fields are due to Ben-Sasson,

Carmon, Ishai, Kopparty and Saraf [2], whose unique-decoding bound is linear in the block length; we use their parameterized-curve theorem as the base case of the descent. The limits of large-radius gaps over characteristic-two fields were charted later [3]. A linear proximity gap for Reed–Solomon codes over finite fields within the 1.5-Johnson bound is given by Gao, Kan and Li [6]. Over Galois rings, Jia, Li, Xing, Yao and Yuan [7] proved the $(1-\rho)/2$ gap for the curve with the quadratic threshold $s\ell^2$, by generalizing Berlekamp–Welch decoding and the Polishchuk–Spielman lemma [10]; this is the bound we sharpen.

Wu et al. [8] state a proximity theorem for general linear codes over Galois rings, under their code hypotheses. Their statement is for a *single* random combination $u^* + \alpha u$ (the line, $s = 1$), with the challenge α drawn from the *entire* ring, at a (double-)Johnson radius. Writing their threshold in the present notation,

$$T_{\text{Wu}} = \frac{2|R|}{|\mathbb{F}_{2^m}| \gamma^3} = \frac{2 \cdot 2^{m(k-1)}}{\gamma^3} \gg 2^m = |\mathcal{T}| \quad (k \geq 2), \quad (5)$$

so over the Teichmüller set the antecedent $|S| > T_{\text{Wu}}$ is unsatisfiable and the statement is vacuous (at the parameters (2), $T_{\text{Wu}} \approx 2^{m \cdot 63}$ against $|\mathcal{T}| = 2^{155}$). This is a comparison of that counting bound after restricting its challenge set, not an impossibility of full-ring sampling or a criterion for a fold to be well defined. Our uniform-residue extension is proved separately in Corollary 14. BaseFold [9] gives field-agnostic commitments from foldable codes; list decoding of (folded) Reed–Solomon codes over Galois rings was studied by Yuan and Zhu [11]. The chain-ring structure we use is classical [12].

Roadmap. Section 2 fixes notation, the proximity-gap definition, and the field base case. Section 3 localizes the quadratic factor in the prior ring proof. Section 4 gives the descent and proves the main theorem. Section 5 records the binding-error consequence and the residual quantitative question. Section 6 concludes; deferred proofs are in Section A.

2 Preliminaries

2.1 Galois rings

For integers $k, m \geq 1$ fix a monic $h \in \mathbb{Z}_{2^k}[X]$ of degree m with $\bar{h} \in \mathbb{F}_2[X]$ irreducible. The *Galois ring* is $R = \text{GR}(2^k, m) = \mathbb{Z}_{2^k}[X]/(h(X))$, a finite local ring with maximal ideal $\mathfrak{m} = (2)$, residue field $R/\mathfrak{m} \cong \mathbb{F}_{2^m}$, and $|R| = 2^{km}$ [12]. The ideal is nilpotent, $\mathfrak{m}^k = 0$, so 2 is a nilpotent zero-divisor when $k \geq 2$. The annihilator structure (3) has the consequence we use repeatedly:

$$\text{multiplication by 2 is injective} \quad R/2^{k-t-1}R \hookrightarrow R/2^{k-t}R \quad (0 \leq t \leq k-2). \quad (6)$$

An element is a unit iff it is nonzero modulo $\mathfrak{m}$; we write $\overline{(\cdot)}: R \twoheadrightarrow \mathbb{F}_{2^m}$ for the reduction homomorphism and $R^\times$ for the units.

2.2 Teichmüller and exceptional sets

There is a unique *Teichmüller set* $\mathcal{T} \subset R$ of 2^m elements, mapped bijectively onto $\mathbb{F}_{2^m}$ by reduction, closed under multiplication, with $\mathcal{T}^\times = \mathcal{T} \setminus \{0\}$ cyclic of odd order $2^m - 1$ [12]. It is the set of roots of $X^{2^m} - X$ in R ; since that polynomial has unit derivative, every residue lifts uniquely (Hensel), so

$$\overline{(\cdot)}|_{\mathcal{T}}: \mathcal{T} \xrightarrow{\sim} \mathbb{F}_{2^m} \text{ is a bijection.} \quad (7)$$

A subset $E \subseteq R$ is *exceptional* if $e - e' \in R^\times$ for all distinct $e, e' \in E$; equivalently E injects into $\mathbb{F}_{2^m}$ under reduction. The Teichmüller set is the maximal exceptional set. The proximity-gap *challenge* ranges over $\mathcal{T}$; an evaluation domain need only be exceptional.

2.3 Reed–Solomon codes and the Hamming metric

Fix an exceptional $E \subseteq R$, $|E| = \ell$, and a rate $\rho \in (0, 1)$ with $\rho\ell \in \mathbb{Z}$. The *Reed–Solomon code* is

$$\text{RS}[R, E, \rho] = \{(f(a))_{a \in E} : f \in R[X], \deg f < \rho\ell\} \subseteq R^\ell. \quad (8)$$

For $u, v \in R^\ell$ the Hamming distance $\Delta(u, v)$ counts the coordinates where they differ (we use absolute distance, with $\Delta(u, C) = \min_{c \in C} \Delta(u, c)$). By Lemma 1, $\text{RS}[R, E, \rho]$ has minimum distance $(1 - \rho)\ell + 1$. We use the strictly interior relative radius $\delta < (1 - \rho)/2$ and the integer budget $e = \lfloor \delta\ell \rfloor$.

Lemma 1 (Root bound on exceptional sets). *Let $E \subseteq R$ be exceptional and $0 \neq f \in R[X]$ with $\deg f < |E|$. Then f has at most $\deg f$ roots in E .*

Proof. If $a \in E$ is a root, division by the monic $X - a$ is exact over R , giving

$$f(X) = (X - a)q(X), \quad \deg q = \deg f - 1.$$

For a second root $b \in E$ we get $0 = f(b) = (b - a)q(b)$ with $b - a \in R^\times$, hence $q(b) = 0$; induct on $\deg f$. $\square$

Corollary 2 (Unique decoding, list size one). *For $\delta < (1 - \rho)/2$, every $w \in R^\ell$ has at most one codeword of $\text{RS}[R, E, \rho]$ within distance $\delta\ell$.*

Proof. Two such codewords differ by a degree- $< \rho\ell$ polynomial that agrees on more than $\ell - 2\delta\ell > \rho\ell$ points of E ; by Lemma 1 the difference is zero. $\square$

2.4 Proximity gaps and the field base case

Definition 3 (Correlated agreement for the batched curve). *Words $g_0, \dots, g_s \in R^E$ are in correlated δ -agreement with $C = \text{RS}[R, E, \rho]$ if there are a single set $A \subseteq E$ with $|A| \geq (1 - \delta)\ell$ and codewords $C_0, \dots, C_s \in C$ such that $g_j|_A = C_j|_A$ for all j . For the curve (1) and a challenge set $\Theta \subseteq R$ we write*

$$S_\Theta = \{\alpha \in \Theta : \Delta(w(\alpha), C) \leq \delta\ell\}. \quad (9)$$

The base case of the descent is the field instance. We use the parameterized-curve theorem of [2], which at unique decoding is linear with no δ factor.

Theorem 4 (Field correlated agreement for curves at unique decoding). *Let $\mathbb{F} = \mathbb{F}_{2^m}$, $\bar{E} \subseteq \mathbb{F}$ with $|\bar{E}| = \ell$, $1 \leq \rho\ell < \ell$, $\rho\ell \in \mathbb{Z}$, $s \geq 1$, and $0 \leq \delta < (1 - \rho)/2$. For $\bar{g}_0, \dots, \bar{g}_s \in \mathbb{F}^{\bar{E}}$, put $\bar{w}(z) = \sum_{j=0}^s z^j \bar{g}_j$ and*

$$\bar{S} = \{z \in \mathbb{F} : \Delta(\bar{w}(z), \text{RS}[\mathbb{F}, \bar{E}, \rho]) \leq \delta\ell\}.$$

If $|\bar{S}| > s\ell$, there are codewords $\bar{C}_0, \dots, \bar{C}_s$ and one set $\bar{A} \subseteq \bar{E}$ with $|\bar{A}| \geq \ell - \lfloor \delta\ell \rfloor$ such that $\bar{g}_j|_{\bar{A}} = \bar{C}_j|_{\bar{A}}$ for every j .

Proof. For positive radius, apply [2, Thm. 6.1] in the July 3, 2021 ePrint version. At unique decoding its threshold is $s\ell$. That version uses degree at most $d - 1$ and rate d/ℓ , where $d = \rho\ell$.

At radius zero, take the quotient vector space by the code. If its polynomial is nonzero, a linear functional detects a nonzero coefficient. The resulting scalar polynomial has degree at most s and hence at most s roots. More than $s\ell$ zero evaluations therefore force every coefficient class to vanish. $\square$

Remark 5 (The curve threshold is $s\ell$: derivation of the base constant). Theorem 4 is the counting form of the parameterized-curve statement of [2, Thms. 1.5 and 6.1], which reads: for the degree- s curve $\text{curve}(\mathbf{g}) = \{\sum_{j=0}^s z^j g_j : z \in \mathbb{F}\}$ and $\delta \in (0, 1 - \sqrt{\rho})$,

$$\Pr_{z \in \mathbb{F}} [\Delta(\text{curve}(z), V)/n \leq \delta] > s \cdot \epsilon \implies \text{correlated agreement}, \quad (10)$$

where $\epsilon = \epsilon(q, n, \rho, \delta)$ is the proximity-gap error parameter of [2, Thm. 1.2]. At the unique-decoding radius $\delta < (1 - \rho)/2$ this parameter is $\epsilon = \epsilon_U = n/q$ [2, Eq. (1.1)], with *no* factor δ (the δn figure there is only a tightness example, not the proven upper bound). Converting the probability in (10) to a count over the q challenges, the threshold is therefore

$$|\bar{S}| > s \cdot \epsilon_U \cdot q = s \cdot \frac{n}{q} \cdot q = s\ell \quad (n = \ell), \quad (11)$$

exactly the form of Theorem 4: *linear in ℓ , constant 1, independent of δ* . This is the external input for the curve descent. The affine form uses the corresponding field theorem stated in Theorem 10. Neither input assumes ring divisibility. Quantitative refinements in [3] have their own radius and agreement-loss hypotheses; no uniform constant bound up to the decoding boundary is inferred here.

3 The published bivariate condition and an independent route

The published statement of [7] gives a curve threshold $s\ell^2$ at $\delta \leq (1 - \rho)/2$. The calculation here concerns the numerical requirements of that proof strategy, separately from whether its ring divisibility steps apply.

The matrix count. The source's matrix analysis uses McCoy's rank criterion and annihilators to produce a nontrivial kernel vector. Its stated count is

$$|S_{\mathcal{T}}| > s(e+1), \quad (12)$$

which is linear in ℓ for $e = \lfloor \delta \ell \rfloor$. A nontrivial kernel vector does not by itself certify that a specified component polynomial is primitive or regular.

The displayed bivariate condition. The source's sufficient numerical condition is

$$\frac{e + \rho\ell - 1}{\ell - e} + \frac{s(e+1)}{|S_{\mathcal{T}}| - se} < 1. \quad (13)$$

At the boundary $e = (1 - \rho)\ell/2$,

$$\frac{e + \rho\ell - 1}{\ell - e} = 1 - \Theta(1/\ell), \quad \rho + \delta \leq (1 + \rho)/2 < 1 \quad \text{in the field comparison.} \quad (14)$$

The first term leaves slack of order $1/\ell$, so satisfying this particular inequality at the boundary requires a challenge count of order $s\ell^2$. This is a restriction of the displayed sufficient condition, not a lower bound on the actual error.

Why the descent is independent. Reducing a nonzero polynomial modulo 2 can give the zero polynomial, and cancelling a nonzero ring element need not be valid: in $\mathbb{Z}/4\mathbb{Z}$, $2 \cdot 1 = 2 \cdot 3$ although $1 \neq 3$. A bivariate ring proof using these steps must establish the relevant primitive or regular hypotheses. Slice-divisibility witnesses can also be nonunique; a conclusion about given witness families must be distinguished from an existential choice of compatible quotients. These observations do not refute the final proximity theorem. The descent below does not depend on a repair of those steps: it proves the ring bound from the stated finite-field inputs.

4 The 2-adic descent

Throughout, $0 \leq \delta < (1 - \rho)/2$ and $e = \lfloor \delta \ell \rfloor$, so $2e < (1 - \rho)\ell$. We build, for $l = 0, 1, \dots, k-1$, nested sets $A_l \subseteq E$ and challenge sets $S_{l+1} \subseteq S_l$, together with words over $R/2^{k-l}R$, maintaining the *partition invariant*

$$b_l := \ell - |A_{l-1}|, \quad e_l = e - b_l, \quad b_l \leq e, \quad (15)$$

with $A_{-1} = E$, $b_0 = 0$, $e_0 = e$. Figure 1 summarizes one layer.

Theorem 6 (Linear gap for batched curves). *Let $R = \text{GR}(2^k, m)$, $k \geq 1$, and let $E \subseteq R$ be exceptional with $|E| = \ell < q = 2^m$. Set $C = \text{RS}[R, E, \rho]$ with $1 \leq \rho\ell < \ell$, $\rho\ell \in \mathbb{Z}$, and fix $0 \leq \delta < (1 - \rho)/2$, $s \geq 1$. For fixed words $g_0, \dots, g_s \in R^E$ and the curve (1), the threshold*

$$L = s\ell \quad (16)$$




Fig. 1. One nonterminal layer ($k - l > 1$) of the 2-adic descent. Reduction mod 2 produces a field instance to which the linear curve gap (Theorem 4) applies; the field codewords are lifted and the residual is divided by 2, yielding the next layer over $R/2^{k-l-1}R$ on the punctured domain A_l . The error budget e_l and the discarded coordinates b_l satisfy the invariant (15).

is independent of k : $|S_{\mathcal{T}}| > L$ implies correlated agreement outside at most $\lfloor \delta \ell \rfloor$ positions. If common agreement does not hold, a fresh uniform $\alpha \in \mathcal{T}$ therefore satisfies $\Pr[\Delta(w(\alpha), C) \leq \delta \ell] \leq \min\{1, s\ell/q\}$.

The proof occupies the rest of the section. At layer l the words $g_j^{(l)}$ live on A_{l-1} over $R_{k-l} := R/2^{k-l}R$ (with $g_j^{(0)} = g_j$). For $\alpha \in S_l$, applying Corollary 2 over R_{k-l} gives a unique codeword $c_\alpha^{(l)}$ with disagreement set

$$D_\alpha^{(l)} := \{x \in A_{l-1} : w^{(l)}(\alpha)(x) \neq c_\alpha^{(l)}(x)\}, \quad |D_\alpha^{(l)}| \leq e_l. \quad (17)$$

Project and agree. Reduce modulo 2. By Lemma 25 (Section A) the field bad-set on A_{l-1} has size at least $|S_l|$, so if $|S_l| > s(\ell - b_l)$ then Theorem 4 applies at radius e_l and yields field codewords $\bar{C}_j^{(l)}$ whose *common discrepancy support*

$$W_l := \{x \in A_{l-1} : \bar{g}_j^{(l)}(x) \neq \bar{C}_j^{(l)}(x) \text{ for some } j\} \quad (18)$$

satisfies $|W_l| \leq e_l$ —Theorem 4 returns agreement on a set of size $\geq (\ell - b_l) - e_l$, and W_l is contained in its complement—while by construction $\bar{g}_j^{(l)} = \bar{C}_j^{(l)}$ for all j on $A_l := A_{l-1} \setminus W_l$, and every $x \in W_l$ is a genuine discrepancy of some j (the property Lemma 9 uses). Here Theorem 4 is re-invoked at each layer on the *punctured* domain $\bar{A}_{l-1} \subseteq \mathbb{F}_{2^m}$ of length $\ell - b_l$, i.e. at the larger relative rate $\rho' = \rho\ell/(\ell - b_l) \geq \rho$. Two facts make this re-invocation harmless. First, the threshold: the per-layer field threshold is $s(\ell - b_l)$. Excluding at most sb_l challenges from an initial set larger than $s\ell$ leaves more than this threshold. The base count is rate-independent; puncturing changes the rate but preserves the dimension $\rho\ell$. Thus each carved coordinate lowers the field threshold by the same amount spent on its possible cancellations. The remaining requirement is the radius:

the punctured instance still lies within the unique-decoding radius of $C|_{A_{l-1}}$, whose minimum distance is $(1 - \rho)\ell - b_l + 1$, by Lemma 7 (ii). The bound $|W_l| \leq e_l$ is the absolute support count: on the length- $(\ell - b_l)$ domain A_{l-1} at radius e_l , Theorem 4 returns a common agreement set of size $\geq (\ell - b_l) - e_l$, so its complement W_l has at most $(\ell - b_l) - ((\ell - b_l) - e_l) = e_l$ points.

Lemma 7 (Budget reduction and full-radius preservation). (i) For $\alpha \in S_l$ such that no $x \in W_l$ is a cancellation point—that is,

$$\sum_{j=0}^s \alpha^j (\bar{g}_j^{(l)} - \bar{C}_j^{(l)})(x) \neq 0 \quad \text{for all } x \in W_l \quad (19)$$

—the carry $w^{(l+1)}(\alpha)$ defined below has at most $e_l - |W_l|$ errors on A_l against $c_\alpha^{(l+1)}$. (ii) The instance on A_{l-1} lies within the unique-decoding radius of the punctured code $C|_{A_{l-1}}$ for every l .

Proof. (ii) The code $C|_{A_{l-1}}$ has length $\ell - b_l$ and dimension $\rho\ell$ —the dimension is preserved because $b_l \leq e \leq (1 - \rho)\ell/2$ gives $\ell - b_l \geq \ell(1 + \rho)/2 > \rho\ell$, so any $\rho\ell$ of the remaining (exceptional) points still determine a degree- $< \rho\ell$ codeword uniquely (Lemma 1)—and minimum distance $(1 - \rho)\ell - b_l + 1$, hence unique-decoding radius $\lfloor ((1 - \rho)\ell - b_l)/2 \rfloor$. Using $e_l = e - b_l$,

$$e_l < \frac{(1 - \rho)\ell - b_l}{2} \iff 2e < (1 - \rho)\ell + b_l, \quad (20)$$

which holds since $2e < (1 - \rho)\ell$ and $b_l \geq 0$, and in fact gains slack as b_l grows. The naive choice $e_l = e$ would instead require $2e + b_l \leq (1 - \rho)\ell$, forcing $\delta \leq (1 - \rho)/3$ as $b_l \uparrow e$.

(i) On $A_l \setminus D_\alpha^{(l)}$ both lines split off the same lifted codeword:

$$w^{(l)}(\alpha) = \sum_j \alpha^j C_j^{(l)} + 2w^{(l+1)}(\alpha), \quad (21)$$

$$c_\alpha^{(l)} = \sum_j \alpha^j C_j^{(l)} + 2c_\alpha^{(l+1)}. \quad (22)$$

(The split (22) presumes the numerator $c_\alpha^{(l)} - \sum_j \alpha^j C_j^{(l)}$ is divisible by 2; this divisibility is Lemma 8, whose proof equates two field codewords by list-1 decoding and does *not* invoke this part (i), so the dependence is a forward reference, not a circularity.) Subtracting, $2(w^{(l+1)} - c_\alpha^{(l+1)})(x) = 0$, and since multiplication by 2 is the injective map $R_{k-l-1} \hookrightarrow R_{k-l}$ of (6) (not a self-map of R_{k-l-1}), $(w^{(l+1)} - c_\alpha^{(l+1)})(x) = 0$ there. Thus the carry's errors on A_l lie in $D_\alpha^{(l)} \cap A_l$. For the chosen α , every $x \in W_l$ disagrees already modulo 2: reducing (22) modulo 2 and using that $\overline{(\cdot)}$ is a ring homomorphism ($\overline{\alpha^j} = \overline{\alpha}^j$) gives the field-codeword identity $\bar{c}_\alpha^{(l)} = \sum_j \bar{\alpha}^j \bar{C}_j^{(l)}$; then by (19), $\bar{w}^{(l)}(\alpha)(x) \neq \bar{c}_\alpha^{(l)}(x)$, so $W_l \subseteq D_\alpha^{(l)}$ and

$$|D_\alpha^{(l)} \cap A_l| = |D_\alpha^{(l)}| - |W_l| \leq e_l - |W_l|. \quad \square$$

Lift and carry. Lift each $\bar{C}_j^{(l)}$ to a codeword $C_j^{(l)} \in C \bmod 2^{k-l}$ by choosing *any* lift of its $< \rho\ell$ coefficients to R_{k-l} —a fixed choice made before any α . On A_l , the agreement (18) gives $g_j^{(l)} - C_j^{(l)} \in 2R_{k-l}$. When $k - l > 1$, define the *carry*

$$g_j^{(l+1)} := \frac{g_j^{(l)} - C_j^{(l)}}{2} \quad \text{on } A_l, \quad (23)$$

a well-defined word over R_{k-l-1} .

Lemma 8 (Coefficient-lift carry). *At a layer with $k - l > 1$, for the single α -independent lift $\{C_j^{(l)}\}$ and every $\alpha \in S_l$ simultaneously,*

$$c_\alpha^{(l+1)} := \frac{c_\alpha^{(l)} - \sum_j \alpha^j C_j^{(l)}}{2} \quad (24)$$

is a codeword of $\text{RS}[R_{k-l-1}, E, \rho]$ (degree $< \rho\ell$). The argument uses reduction as a ring homomorphism, field unique decoding and the injection between consecutive quotient rings; neither Teichmüller multiplicativity nor Hasse derivatives are required.

Proof. Let $P_\alpha, \hat{P}_j$ be the degree- $< \rho\ell$ polynomials of $c_\alpha^{(l)}$ and $C_j^{(l)}$ over R_{k-l} , and set

$$D_\alpha := P_\alpha - \sum_{j=0}^s \alpha^j \hat{P}_j, \quad \deg D_\alpha < \rho\ell. \quad (25)$$

Reducing (25) modulo 2 and using that $(\overline{\cdot})$ is a ring homomorphism ($\overline{\alpha^j} = \overline{\alpha}^j$) gives $\overline{D_\alpha} = \overline{c_\alpha^{(l)}} - \sum_j \overline{\alpha}^j \overline{C_j^{(l)}}$. Both $\overline{c_\alpha^{(l)}}$ and $\sum_j \overline{\alpha}^j \overline{C_j^{(l)}}$ are degree- $< \rho\ell$ field codewords within e_l ($\leq \delta\ell$) of $\overline{w^{(l)}}(\alpha)$ —the former as the reduction of the unique $c_\alpha^{(l)}$, the latter because the discrepancy set on A_{l-1} is W_l , with $|W_l| \leq e_l$ —so by Corollary 2 over the field, applied to the punctured code $\overline{C}|_{A_{l-1}}$ of minimum distance $(1 - \rho)\ell - b_l + 1$ (Lemma 7 (ii)), where the list size is one at radius e_l , they coincide as polynomials. Hence $\overline{D_\alpha} \equiv 0$, i.e. every coefficient of D_α lies in $\mathfrak{m} = (2)$, so

$$D_\alpha = 2Q_\alpha, \quad \deg Q_\alpha < \rho\ell \text{ over } R_{k-l-1}. \quad (26)$$

Evaluating Q_α on E gives $c_\alpha^{(l+1)}$. Its degree is less than $\rho\ell$, so this is a codeword over R_{k-l-1} with the required dimension.

It remains to record how the per-layer codewords compose across the k layers; the cross-layer assembly is carried out in the proof of Theorem 6, and we isolate here the single per-layer identity it uses. On A_l the agreement (18) gives $\overline{g_j^{(l)}} = \overline{C_j^{(l)}}$, so $g_j^{(l)} - C_j^{(l)} \in 2R_{k-l}$, and rearranging the carry definition (23) yields

$$g_j^{(l)} = C_j^{(l)} + 2g_j^{(l+1)} \quad \text{on } A_l, \quad (27)$$

the digit relation that unrolls in Theorem 6. The lift $C_j^{(l)}$ is fixed before any α , so its α -independence is automatic. The choice of lift of the $< \rho\ell$ coefficients of $\overline{C_j^{(l)}}$ to R_{k-l} is *not* unique: its ambiguity is the maximal ideal $2R_{k-l}$, and a different choice changes the carry $g_j^{(l+1)}$ and hence all downstream digits $g_j^{(l')}$ ($l' > l$). This is harmless because the assembly in Theorem 6 reconstitutes a single degree- $< \rho\ell$ R -codeword C_j with $C_j|_A = g_j|_A$ for *any* such choices, via the self-correcting digit identity (33); no uniqueness of the per-layer lift is claimed or needed. $\square$

Lemma 9 (Cancellation challenges). *At most $s|W_l|$ challenges $\alpha \in \mathcal{T}$ are cancellation points of some $x \in W_l$.*

Proof. For fixed $x \in W_l$ consider the univariate polynomial over $\mathbb{F}_{2^m}$

$$P_x(z) := \sum_{j=0}^s z^j (\bar{g}_j^{(l)} - \bar{C}_j^{(l)})(x), \quad \deg P_x \leq s. \quad (28)$$

It is nonzero because $x \in W_l$ is an exact discrepancy. If $s \geq q$ the claimed root count is trivial; otherwise Lemma 1 bounds its roots by s . By the bijection (7) these lift to at most s challenges in $\mathcal{T}$. Summing over $x \in W_l$ gives the bound. $\square$

Proof (of Theorem 6). Set $S_0 = S_{\mathcal{T}}$ and suppose $|S_0| > s\ell$. At every layer maintain the words and decoders described above and

$$|S_l| \geq |S_0| - sb_l > s(\ell - b_l). \quad (29)$$

The initial instance has $b_0 = 0$ and $e_0 = e$. The radius inequality (20) and (29) permit the field step and give the exact support W_l with $|W_l| \leq e_l$.

If $k - l > 1$, remove the cancellation challenges and form the carries. By Lemma 9, at most $s|W_l|$ challenges are removed. By Lemmas 7 and 8, every retained challenge has a degree- $< \rho\ell$ decoder within $e_l - |W_l|$ of the carry. Set $b_{l+1} = b_l + |W_l|$ and $e_{l+1} = e - b_{l+1}$. The error budget remains nonnegative, and

$$|S_{l+1}| \geq |S_0| - sb_{l+1} > s(\ell - b_{l+1}).$$

Thus the field threshold decreases by exactly the cost of excluded challenges. At the last layer $k - l = 1$, stop after finding its field agreement set; no division in the zero ring is used.

The sets W_l are disjoint, and at the last layer define $b_k = b_{k-1} + |W_{k-1}|$. Then

$$\sum_{l=0}^{k-1} |W_l| = b_k \leq e = \lfloor \delta\ell \rfloor. \quad (30)$$

The possible exclusions have total at most

$$\sum_{l=0}^{k-1} s|W_l| = sb_k \leq se. \quad (31)$$

They do not require an additional se in the initial threshold: (29) already accounts for them against the shrinking domain.

Let $A = A_{k-1}$. On A , for $0 \leq l < k - 1$,

$$g_j^{(l)} = C_j^{(l)} + 2g_j^{(l+1)}, \quad (32)$$

while $g_j^{(k-1)} = C_j^{(k-1)}$ on A in the final field layer. Unrolling these identities gives

$$g_j = \sum_{l=0}^{k-1} 2^l C_j^{(l)} \quad \text{on } A. \quad (33)$$

Each $C_j^{(l)}$ is the evaluation of a degree- $< \rho\ell$ polynomial over R_{k-l} . Lift its coefficients to R and set

$$C_j = \sum_{l=0}^{k-1} 2^l C_j^{(l)}. \quad (34)$$

A coefficient's re-lift ambiguity lies in $2^{k-l}R$ and is killed by 2^l , so this defines a degree- $< \rho\ell$ polynomial over R . Equation (33) proves its agreement with g_j on A , and $|A| = \ell - b_k \geq \ell - e$ proves the conclusion.

Intermediate lifts and carries can vary, but the final decoded codeword for each fixed g_j is unique: any two outputs are within e of g_j and therefore agree on at least $\ell - 2e > \rho\ell$ exceptional points. Lemma 1 equates them. The probability bound is the contrapositive count divided by $|\mathcal{T}| = q$, with the trivial cap of 1. $\square$

4.1 An affine-space variant: removing the s -factor

The threshold of Theorem 6 is $O(s\ell)$ because the degree- s curve (1) is driven by a *single* challenge α , whose $q = 2^m$ values must absorb the s -fold batching. Combining the same words as an *affine space*, with s independent challenges $z = (z_1, \dots, z_s) \in \mathcal{T}^s$,

$$w(z) = g_0 + \sum_{j=1}^s z_j g_j, \quad (35)$$

removes the s -factor from the algebraic bad-challenge bound. The required field input is explicit.

Theorem 10 (Field affine correlated agreement at unique decoding).

Under the field, code and radius hypotheses of Theorem 4, let $\bar{w}(z) = \bar{g}_0 + \sum_{j=1}^s z_j \bar{g}_j$, $z \in \mathbb{F}_q^s$. If more than ℓq^{s-1} vectors satisfy $\Delta(\bar{w}(z), C) \leq \delta\ell$, then the $\bar{g}_j$ have common agreement with codewords outside at most $\lfloor \delta\ell \rfloor$ positions.

Proof. For positive radius use [2, Thm. 1.6, Sec. 6.3] in the July 3, 2021 ePrint version, with $\epsilon_U = \ell/q$. The affine map from $\mathbb{F}_q^s$ onto $\bar{g}_0 + \text{span}\{\bar{g}_1, \dots, \bar{g}_s\}$ has equally sized fibers. Thus uniform coefficient vectors give uniform elements of that affine space, even when the directions are dependent. At radius zero, the preimage of the code is an affine subspace of $\mathbb{F}_q^s$. A proper such subspace has at most q^{s-1} points; otherwise the offset and all directions belong to the code. $\square$

Theorem 11 (Linear affine-batched gap; the s -factor removed). *Let R, E, C, δ be as in Theorem 6, and let $s \geq 1$. For fixed $g_0, \dots, g_s \in R^E$ and the affine combination (35) with $z \in \mathcal{T}^s$, if*

$$|\{z \in \mathcal{T}^s : \Delta(w(z), C) \leq \delta\ell\}| > L^{\text{aff}} = \ell q^{s-1} \quad (q = 2^m), \quad (36)$$

then $g_0, \dots, g_s$ are in correlated δ -agreement with C on a common set of size $\geq (1 - \delta)\ell$, with proximity loss 0. If common agreement fails, the probability of a good fresh uniform vector is at most $L^{\text{aff}}/q^s = \ell/2^m$, independent of s . This is a bound on the specified batching event.

The proof is the descent of Theorem 6 with the cancellation count replaced; the other two lemmas transfer because they never see the coefficient structure.

Lemmas 7 and 8 are coefficient-agnostic. Lemma 7 uses only the additive split: on $A_l \setminus D_z^{(l)}$,

$$w^{(l)}(z) = \left(C_0^{(l)} + \sum_j z_j C_j^{(l)}\right) + 2w^{(l+1)}(z), \quad c_z^{(l)} = \left(C_0^{(l)} + \sum_j z_j C_j^{(l)}\right) + 2c_z^{(l+1)},$$

of the same form as (21)–(22); subtracting and using (6) puts the carry's errors in $D_z^{(l)} \cap A_l$, and the radius bound (20) involves only e_l, b_l, ρ, ℓ . Lemma 8 uses only that reduction is a ring homomorphism and field list-1 decoding (Corollary 2): with $D_z = P_z - \hat{P}_0 - \sum_j z_j \hat{P}_j$ one has $\overline{D_z} = \overline{c_z^{(l)}} - (\overline{C_0^{(l)}} + \sum_j \overline{z_j} \overline{C_j^{(l)}})$, both summands degree- $< \rho\ell$ field codewords within e_l of $\overline{w^{(l)}}(z)$, hence equal by Corollary 2, so $\overline{D_z} = 0$, $D_z = 2Q_z$ with $\deg Q_z < \rho\ell$, and the digits assemble into $C_j = \sum_l 2^l C_j^{(l)}$. Neither argument depends on whether the coefficients are the powers α^j of Theorem 6 or the independent z_j here; the support bound $|W_l| \leq e_l$ and the telescoping (30) are likewise unchanged.

Lemma 12 (Affine cancellation is a hyperplane). *At most q^{s-1} vectors $z \in \mathcal{T}^s$ are cancellation points of a fixed $x \in W_l$.*

Proof. The cancellation condition at x reads, over $\mathbb{F} = \mathbb{F}_{2^m}$,

$$\sum_{j=1}^s \overline{z_j} (\overline{g_j^{(l)}} - \overline{C_j^{(l)}})(x) = -(\overline{g_0^{(l)}} - \overline{C_0^{(l)}})(x). \quad (37)$$

As $x \in W_l$ is in the common support, $(\overline{g_j^{(l)}} - \overline{C_j^{(l)}})(x) \neq 0$ for some $j \in \{0, \dots, s\}$. If some coefficient with $j \geq 1$ is nonzero, (37) is a nondegenerate affine equation in $\overline{z} \in \mathbb{F}^s$ with exactly q^{s-1} solutions; if only the $j = 0$ term is nonzero it reads $(\text{nonzero}) = 0$, with none. The coordinatewise Teichmüller bijection (7) gives $\mathcal{T}^s \xrightarrow{\sim} \mathbb{F}^s$, so at most q^{s-1} vectors of $\mathcal{T}^s$ satisfy it. $\square$

Proof (of Theorem 11). Run the descent of Theorem 6, using Theorem 10 with threshold $(\ell - b_l)q^{s-1}$ and Lemma 12. At most $b_l q^{s-1}$ vectors have been excluded before layer l . An initial good set larger than ℓq^{s-1} therefore leaves

$$|S_l| \geq |S_0| - b_l q^{s-1} > (\ell - b_l) q^{s-1}.$$

The field theorem remains applicable. The radius, exact support, coefficient carry and final assembly are unchanged; at the last field layer stop without division. The final common set has at least $\ell - e$ positions. Dividing the contrapositive count by q^s gives the probability bound. $\square$

Remark 13 (The affine comparison). At (2), the affine upper bound is $\ell/2^m = 2^{-134}$; the degree-64 curve bound is 2^{-128} . Affine batching uses s coefficient draws rather than one. Its transcript expansion and arithmetic costs depend on the application. In particular, a folding protocol must admit the specified affine combination and its evaluation reduction before this bound can replace that protocol's analysis.

Corollary 14 (Challenges with uniform residue). *For fixed words under the hypotheses of Theorem 6 that do not have common agreement at its stated radius, the scalar probability bound $\min\{1, s\ell/q\}$ holds for any challenge in R whose residue is uniform in $\mathbb{F}_q$. For the affine combination, a vector whose joint residue is uniform in $\mathbb{F}_q^s$ has bound ℓ/q . The distribution of higher digits may be arbitrary. In particular, uniform full-ring draws and uniform Teichmüller draws both satisfy the hypotheses. The words must be fixed before the fresh challenge is drawn.*

Proof. The counting proofs also work for any selected good challenges with distinct residues. At every layer the field parameters remain distinct; the discrepancy polynomial has at most s scalar roots, or the affine discrepancy equation has at most q^{s-1} vector roots. Coefficient lifting, division and assembly use only that reduction is a ring homomorphism. They do not require a multiplicative choice of residue representatives, or a product-shaped section in the vector case.

Collect the residues having at least one good lift and select one good lift for each. If there were more than $s\ell$ scalar residues or ℓq^{s-1} vector residues, the preceding counting proofs would give common agreement. Contrapositively, only that many residues can have any good lift when common agreement fails. Uniform residual distribution gives the claimed probabilities. $\square$

4.2 Weighted and mutual correlated agreement

Some folding reductions measure agreement against coordinate weights. The weight and component words must be fixed before the fresh challenge. The theorem in this section supplies the weighted algebraic input; a protocol must separately prove that its adaptive transcript and path measure satisfy those hypotheses. The factor M below is a cost of this rational-weight argument, not a claimed lower bound on every weighted analysis.

Definition 15 (Weighted agreement). *A weight of denominator M , where $M \geq 1$ is an integer, is a map $\mu : E \rightarrow [0, 1] \cap M^{-1}\mathbb{Z}$. Throughout, $n = \ell = |E|$. For $u, v \in R^E$ and $B \subseteq E$, put*

$$\begin{aligned} \mu(B) &= \frac{1}{n} \sum_{x \in B} \mu(x), \\ \text{agr}_\mu(u, v) &= \mu(\{x : u(x) = v(x)\}), \\ \text{agr}_\mu(u, C) &= \max_{v \in C} \text{agr}_\mu(u, v). \end{aligned} \tag{38}$$

The normalized masses lie on the grid $(Mn)^{-1}\mathbb{Z}$ and $\mu(E) \leq 1$. On a punctured domain we continue to use the ambient normalization $1/n$ unless another normalization is explicitly stated. The constant weight $\mu \equiv 1$ has $M = 1$ and gives $\text{agr}_\mu(u, v) = 1 - \Delta(u, v)/n$.

A direct weighted field input. The next statement follows from the unweighted finite-field inputs Theorems 4 and 10. Its proof supplies the constants and does not assume equality of complete agreement sets for different accepting challenges.

Theorem 16 (Field weighted correlated agreement at unique decoding). *Let $D \subseteq \mathbb{F}_q$ have length N , and let C be the code of degree- $< d$ polynomials, where $1 \leq d < N$. Let $\mu : D \rightarrow [0, 1] \cap M^{-1}\mathbb{Z}$, normalized here by $1/N$. Fix $\beta \in (MN)^{-1}\mathbb{Z}$, $\beta \leq 1$, with $1 - \beta < (1 - d/N)/2$, and fixed words $f_0, \dots, f_s$, $s \geq 1$. For $w(z) = \sum_{j=0}^s z^j f_j$, let $S = \{z \in \mathbb{F}_q : \text{agr}_\mu(w(z), C) \geq \beta\}$. If $|S| > sMN$, there are codewords $v_0, \dots, v_s$ and a common set $B \subseteq D$ of mass at least β with $f_j|_B = v_j|_B$ for every j . For $w(z) = f_0 + \sum_{j=1}^s z_j f_j$, $z \in \mathbb{F}_q^s$, the same conclusion holds if $|S| > MNq^{s-1}$.*

Proof. Put $c = s$ for the curve and $c = q^{s-1}$ for the affine case. The inequality $\mu(x) \leq 1$ implies that weighted agreement at least β entails ordinary agreement on at least βN positions. Since $|S| > cMN \geq cN$, the relevant unweighted field theorem applies at radius $1 - \beta$ and supplies common codewords v_j on at least βN positions. For each $z \in S$, both the accepting decoder and the combination of these v_j lie within $(1 - \beta)N$ of $w(z)$. Unique decoding equates them.

Let B be the exact common agreement locus of all f_j, v_j . For each $x \notin B$, the discrepancy is a nonzero polynomial of degree at most s in the scalar parameter, or a nonzero affine form in the vector parameter. At most c parameters cancel it. Sum the weighted agreement of the accepting combinations:

$$|S|\beta \leq |S|\mu(B) + c(\mu(D) - \mu(B)) \leq |S|\mu(B) + c.$$

If $\mu(B) < \beta$, granularity gives $\beta - \mu(B) \geq 1/(MN)$, and therefore $|S| \leq cMN$. This contradicts the hypothesis. Hence $\mu(B) \geq \beta$. $\square$

For scalar parameters the strict threshold is equivalently $|S| \geq sMN + 1$. The earlier sufficient thresholds $s(MN + 1)$ and $(MN + 1)q^{s-1}$ also imply the revised hypotheses. An agreement threshold not on the grid can first be rounded up to the grid without changing the accepting set; the theorem then gives at least the originally requested mass.

Lemma 17 (Weighted unique decoding). *Let μ have denominator M , let $A \subseteq E$, and let $w \in R^A$, with masses normalized by the ambient n . (i) A codeword agreeing with w in mass at least β agrees on at least $n\beta$ positions. (ii) If $2\beta - \mu(A) > \rho$, at most one degree- $< \rho n$ codeword has such agreement on A .*

Proof. For (i), each agreeing coordinate contributes at most $1/n$. For (ii), two such codewords agree with each other on mass at least $2\beta - \mu(A) > \rho$, and therefore on more than ρn positions. The exceptional-set root bound equates their polynomials. $\square$

Theorem 18 (Weighted correlated agreement at unique decoding). *Let R, E, C, n, δ be as in Theorem 6, with $n = \ell$. Fix a weight μ of integer*

denominator $M \geq 1$, a threshold $\alpha \in (Mn)^{-1}\mathbb{Z}$ with $1 - \delta \leq \alpha \leq 1$, and words $u_0, \dots, u_s \in R^E$, $s \geq 1$. For the curve $w(z) = \sum_{j=0}^s z^j u_j$, define

$$S = \{z \in \mathcal{T} : \text{agr}_\mu(w(z), C) \geq \alpha\}. \quad (39)$$

If $|S| > sMn$, then there are codewords $v_0, \dots, v_s \in C$ with

$$\mu(\{x \in E : u_j(x) = v_j(x) \text{ for every } j\}) \geq \alpha. \quad (40)$$

For the affine combination $w(z) = u_0 + \sum_{j=1}^s z_j u_j$, $z \in \mathcal{T}^s$, the same conclusion follows from $|S| > Mnq^{s-1}$. In particular, the previous sufficient conditions $|S| \geq s(Mn + 1)$ and $|S| \geq (Mn + 1)q^{s-1}$ remain valid.

Proof. Use $c = s$ or $c = q^{s-1}$ according to the parameterization. Run the descent on the fixed words and weight. At layer l put $A_{-1} = E$, $b_l = n - |A_{l-1}|$, $N_l = n - b_l$ and $R_{k-l} = R/2^{k-l}R$. Maintain that each retained parameter $z \in S_l$ has a degree- $< \rho n$ codeword agreeing with $w^{(l)}(z)$ on mass at least α , using ambient normalization, and

$$|S_l| \geq |S_0| - cb_l > cMn - cb_l \geq cM(n - b_l).$$

Project and agree. Reduction can only increase agreement, and distinct selected Teichmüller parameters remain distinct after reduction. On the punctured domain of length N_l , set $\rho_l = \rho n / N_l$ and $\alpha_l = \alpha n / N_l$. The value α_l lies on the grid $(MN_l)^{-1}\mathbb{Z}$. Initially $N_0 = n$; at subsequent layers the construction below gives $N_l \geq \alpha n$, so $\alpha_l \leq 1$. Moreover

$$1 - \alpha_l < (1 - \rho_l)/2 \iff N_l + \rho n < 2\alpha n,$$

which holds because $N_l \leq n$ and $\alpha \geq 1 - \delta > (1 + \rho)/2$. In particular $N_l > \rho n$, so puncturing preserves the code dimension. The invariant gives the count required by Theorem 16. Its conclusion, translated back to ambient normalization, yields field codewords $\bar{C}_j^{(l)}$ and the exact support

$$\begin{aligned} W_l &= \{x \in A_{l-1} : \bar{u}_j^{(l)}(x) \neq \bar{C}_j^{(l)}(x) \text{ for some } j\}, & A_l &= A_{l-1} \setminus W_l, \\ \mu(A_l) &\geq \alpha, & b_{l+1} &= n - |A_l| \leq (1 - \alpha)n \leq \delta n. \end{aligned} \quad (41)$$

A set returned by the field theorem is contained in this exact agreement locus. Replacing it by the exact locus preserves its mass and ensures that every carved coordinate is a genuine discrepancy.

Coefficient carry and agreement mass. Fix lifts of the field codeword coefficients for the whole layer. On A_{l-1} , $2\alpha - \mu(A_{l-1}) \geq 2\alpha - 1 > \rho$. The reduction of an accepting decoder and the combination of the lifted field codewords both have agreement mass at least α ; Lemma 17 equates their field polynomials. Their difference is consequently divisible by 2 coefficientwise. If $k - l > 1$, divide by the injection $R_{k-l-1} \rightarrow 2R_{k-l}$ and form the carries, as in Lemma 8. If $k - l = 1$, stop after the field agreement step without division.

Cancellation and induction. For each $x \in W_l$, a nonzero residue polynomial or affine form has at most c roots. Exclude those parameters for every carved coordinate, losing at most $c|W_l|$ in total. For a retained parameter, all of W_l belonged to its error set $D_z^{(l)}$. Division preserves equality elsewhere, so the carry's error mass is at most

$$\mu(D_z^{(l)} \cap A_l) = \mu(D_z^{(l)}) - \mu(W_l) \leq \mu(A_{l-1}) - \alpha - \mu(W_l) = \mu(A_l) - \alpha.$$

Thus the next carry again agrees in mass at least α . The count invariant is preserved because

$$\begin{aligned} |S_l| &\geq |S_0| - cb_l > cMn - cb_l \\ &= cM(n - b_l) + c(M - 1)b_l \geq cMN_l. \end{aligned} \tag{42}$$

Here exclusions count ordinary coordinates, not weighted mass. All layers share one pool; no factor k is incurred.

Assembly. On the final common set $A = A_{k-1}$, the digit identities and last field-layer equality give $u_j = \sum_{l < k} 2^l C_j^{(l)}$. After coefficient lifts back to R , the ambiguity at layer l is killed by 2^l . The resulting v_j have degree $< \rho n$. Since $\mu(A) \geq \alpha$ and all u_j agree with their v_j there, (40) follows. $\square$

Corollary 19 (Uniform-residue weighted probabilities). *If the common-agreement conclusion of Theorem 18 fails, a fresh scalar challenge with uniform residue has good-event probability at most $\min\{1, sMn/q\}$. A vector whose joint residue is uniform has bound $\min\{1, Mn/q\}$. The component words and weight must be fixed before this draw.*

Proof. The proof of Theorem 18 works for selected good challenges with distinct residues, exactly as in Corollary 14. Choose one good lift per potentially good residue. The number of such residues is at most sMn in the scalar case and Mnq^{s-1} in the vector case; otherwise the strict threshold would force common agreement. Divide by q or q^s . $\square$

Remark 20 (The denominator and the cost of composing layers). The threshold is linear in n when M is bounded. Taking $M = n$ on a full-length domain gives the quadratic bound sn^2 from this analysis, without establishing that such a loss is necessary. For the constant weight $\mu \equiv 1$, $M = 1$, the strict threshold $|S| > sn$ exactly recovers Theorem 6.

If a separate folding reduction supplies weights of denominator $M_i = 2^i$ on lengths $n_i = n_0/2^i$, then $M_i n_i = n_0$. The local scalar bound is sn_0/q , or n_0/q for affine batching. A union bound over r such layers is therefore rsn_0/q or rn_0/q , respectively, capped at 1 if necessary. These are conditional sums of local bounds, not an unweighted geometric series or a proof of a compiler's round-by-round game. The protocol must supply its weight construction, prefix chronology and reduction to these events.

Corollary 21 (Mutual form at unique decoding). *Under the hypotheses of Theorem 18, a single set of mass at least α serves every component word*

simultaneously. The reconstructed codeword for each u_j is unique. For every accepting combination in the original set S , its unique decoder is the corresponding combination of these reconstructed codewords.

Proof. The descent already supplies the common set and individual agreement masses. Since $2\alpha - \mu(E) > \rho$, Lemma 17 makes each reconstructed codeword unique. For any accepting parameter, its decoder and the combination of those reconstructed codewords both agree with that combination word in mass at least α ; the same uniqueness argument equates them as polynomials. This argument covers every original accepting parameter, including ones excluded from the descent's retained sets. It does not assert that the complete agreement sets of different accepting combinations are identical. $\square$

5 Application bounds and their limits

For fixed component words without common agreement, Theorem 6 bounds the probability of a close random combination. A folding protocol must separately specify how the parent word decomposes into those components. It must prove that a far parent gives components without common agreement for the actual child code and radius. The definition of a fold alone does not prove this.

For example, suppose such a reduction has been proved at each fixed transcript prefix, with child length ℓ_{i+1} and degree- s scalar combination, and the challenge is fresh with uniform residue. Only under those premises does the local bound

$$\Pr[\Delta(\text{Fold}_\alpha f, C_{i+1}) \leq \delta \ell_{i+1} \mid \text{fixed prefix with } \Delta(f, C_i) > \delta \ell_i] \leq \frac{s \ell_{i+1}}{2^m} \quad (43)$$

follow. An affine adapter uses $\ell_{i+1}/2^m$. The radius must satisfy the theorem's strict inequality for the child code. A claimed fold identity, weighted path invariant, terminal test, authentication reduction or extractor requires its own proof; none follows from the proximity count alone.

If all the required unweighted reductions hold and lengths halve, their numerical sum is geometric. If a protocol instead uses surviving-path weights with $M_i = 2^i$ on $\ell_i = \ell/2^i$, then Theorem 18 gives a count proportional to $M_i \ell_i = \ell$ in each layer. Summing r such bounds costs a factor r , as detailed in Remark 20. A maximum of local transition bounds and a sum of interactive failure probabilities are different quantities. A compiler's round-by-round hypothesis must be stated for its actual game and cannot be inferred merely from either number.

Remark 22 (An illustrative numerical comparison). At (2) the curve calculation is

$$\frac{s\ell}{2^m} = \frac{2^{27}}{2^{155}} = 2^{-128}, \quad \frac{\ell}{2^m} = 2^{-134} \quad \text{for affine batching.} \quad (44)$$

Both hold for every $\delta < 1/4$ when $\rho = 1/2$. The older quadratic expression at the same geometry is 2^{-107} . These are comparisons of stated upper expressions for a specified batching event. They do not establish a matching attack or a

global Pareto ordering of protocols. An overall security target also budgets query errors, other algebraic checks, hash binding, and the applicable non-interactive compiler.

Corollary 23 (Coefficient representation of a combined RS word).

Suppose a consistency phase sends one combined codeword of degree $< d = \rho\ell$ and opens t columns of b committed rows. Sending its d coefficients, which the verifier evaluates on E , represents the same valid codewords as sending all ℓ evaluations. The raw per-phase symbol count is then $d + tb$ rather than $\ell + tb$, with the same authentication and evaluation-link checks. For a binary Merkle tree without path sharing there are additionally at most $t\lceil\log_2 \ell\rceil$ authentication digests. Multiple phases must be counted separately, and coefficients and committed data may have different element widths.

Proof. Evaluation on an exceptional set is injective for degree- $< d$ polynomials by Lemma 1. Thus coefficients and valid codewords are in bijection and give the same values at every queried coordinate. The stated counts enumerate the transmitted entries and paths. $\square$

Remark 24 (Choosing a challenge ring). The present theorems specify algebraic denominators and assumptions, not a deployment recommendation. Enlarging the residue field, changing curve batching to affine batching, adjusting matrix geometry, or changing the commitment protocol are distinct design choices. A smaller algebraic upper bound alone changes no messages and does not prove faster verification or a concrete post-quantum guarantee. Nor does the failure of one multiplicative-domain construction exclude other transparent commitments.

6 Conclusion

A residue descent proves the degree- s curve threshold $s\ell$ and the affine bad-challenge probability $\ell/2^m$ over Galois rings. The decisive accounting uses the actual length of each punctured domain, so cancelled challenges do not add another $s\delta\ell$ term. Coefficientwise division takes place between consecutive quotient rings, and the decoded codewords are unique at the stated radius. Rational weights and arbitrary higher-digit sampling with uniform residue are treated explicitly. The remaining protocol questions are the concrete folding and evaluation reductions, extraction, authentication, zero knowledge and non-interactive compilation. Extending common agreement beyond unique decoding requires different arguments and is not established by this descent.

A Deferred proofs

Lemma 25 (Residue projection of the challenge count). *Let $\overline{(\cdot)}: R \twoheadrightarrow \mathbb{F}_{2^m}$ denote reduction modulo $\mathfrak{m} = (2)$. Put $C = \text{RS}[R, E, \rho]$ and write $\overline{C} = \text{RS}[\mathbb{F}_{2^m}, \overline{E}, \rho]$ for its residue code. For any $u, v \in R^E$,*

$$\Delta(\overline{u}, \overline{v}) \leq \Delta(u, v), \quad (45)$$

$\overline{C}$ is the reduction of C . For a fixed curve (1) and any subset S of its good Teichmüller challenges, the field good set contains $\overline{S}$, with $|\overline{S}| = |S|$.

Proof. If $u_x = v_x$ in R then $\overline{u}_x = \overline{v}_x$, which gives (45): reduction cannot increase the number of disagreements. Reduction commutes with evaluation and sends a degree- $< \rho\ell$ polynomial over R to one over $\mathbb{F}_{2^m}$, so $\overline{C} = \text{RS}[\mathbb{F}_{2^m}, \overline{E}, \rho]$. Hence if $w(\alpha)$ is $\delta\ell$ -close to C then $\overline{w}(\overline{\alpha}) = \overline{w}(\alpha)$ is $\delta\ell$ -close to $\overline{C}$, so every $\alpha \in S$ reduces to an element of the field bad-set. The Teichmüller section (7) is a bijection, so distinct challenges in $\mathcal{T}$ have distinct reductions; therefore $|\overline{S}| = |S|$. $\square$

The bound of Lemma 7 (ii) instantiates Lemma 1 over each quotient R_{k-l} , whose Reed–Solomon code shares the exceptional domain E and hence the minimum distance $(1 - \rho)\ell + 1$.

References

1. E. Ben-Sasson, I. Bentov, Y. Horesh, M. Riabzev. Fast Reed–Solomon interactive oracle proofs of proximity. In *ICALP 2018*.
2. E. Ben-Sasson, D. Carmon, Y. Ishai, S. Kopparty, S. Saraf. Proximity gaps for Reed–Solomon codes. *J. ACM* 70(5), 2023. IACR ePrint 2020/654.
3. E. Ben-Sasson, D. Carmon, U. Haböck, S. Kopparty, S. Saraf. On proximity gaps for Reed–Solomon codes. IACR ePrint 2025/2055.
4. G. Arnon, A. Chiesa, G. Fenzi, E. Yogev. WHIR: Reed–Solomon proximity testing with super-fast verification. IACR ePrint 2024/1586.
5. U. Haböck. A note on mutual correlated agreement for Reed–Solomon codes. IACR ePrint 2025/2110.
6. Y. Gao, H. Kan, Y. Li. Linear proximity gaps for linear codes within the 1.5 Johnson bound. IACR ePrint 2024/1810, version of March 30, 2025.
7. Y. Jia, S. Li, C. Xing, Y. Yao, C. Yuan. Polynomial commitments for Galois rings and applications to SNARKs over $\mathbb{Z}_{2^k}$. In *CRYPTO 2025*.
8. Z. Wu, X. Zhang, Y. Deng, Y. Wei, Z. Zhang, L. Yang. Polylogarithmic polynomial commitment scheme over Galois rings. 2025.
9. H. Zeilberger, B. Chen, B. Fisch. BaseFold: efficient field-agnostic polynomial commitment schemes from foldable codes. In *CRYPTO 2024*.
10. A. Polishchuk, D. A. Spielman. Nearly-linear size holographic proofs. In *STOC 1994*.
11. C. Yuan, R. Zhu. List decoding of Reed–Solomon and folded Reed–Solomon codes over Galois rings. 2025.
12. Z.-X. Wan. *Lectures on Finite Fields and Galois Rings*. World Scientific, 2003.