

Efficient Proximity-Gap SNARKs from Field-Separated Reed–Solomon Reserve Certificates

Przemek Chojceki
ulam.ai

June 27, 2026

Abstract

Reed–Solomon proximity SNARKs are most efficient when the proximity layer is run near capacity rather than at the Johnson radius, but a near-capacity soundness claim must budget exactly the coding-theoretic quantities consumed by the protocol reduction. This paper turns the corrected smooth-domain Reed–Solomon reserve theory of the companion manuscript into a protocol-facing certificate. The certificate separates six ledgers: generated-field list entropy; quotient-core obstructions at the actual (possibly dithered) dimension; a locator-fiber list bound, bridged explicitly to the interleaved list size $|\Lambda(C^{\equiv\mu}, \delta)|$ consumed by the reduction; the list-size-over-field budget $|\Lambda(C^{\equiv\mu}, \delta)|/q_{\text{line}} \leq 2^{-\lambda_{\text{list}}}$; mutual correlated agreement (MCA) or line-decoding over the field in which the line experiment is actually analyzed; and theorem-backed failure-ladder audits. Three field-accounting rules are central and easy to violate. First, an extension challenge field cannot be credited in the MCA or list denominators unless the protocol is analyzed over the corresponding extension code or an explicit extension-line transfer assumption is invoked. Second, the generated-field entropy floor binds regardless of the challenge field: every 2-smooth domain over a 31-bit FFT prime such as KoalaBear has generated field of 31 bits, so reserves below $\tau^*(\rho, 2^{31}) \approx 1/31$ at rate $1/2$ lie in the proven-large-list region no matter how large the extension used for challenges. Third, the self-contained universal field-size cap of the companion cap paper binds the line field itself: no reserve at or below 2^{-9} at the prize rates (2^{-10} at rate $1/16$) is certifiable at 2^{-128} over any field of size up to 2^{256} , and structurally none below $\approx H_2(\rho)/\log_2 q_{\text{line}}$ — a floor that enlarging the challenge field shifts but never clears. We give a corrected concrete comparison against the toy parameters of Arnon–Boneh–Fenzi: at the entropy-feasible reserve $\eta = 1/16$ over KoalaBear, the conjectural corrected layer reduces 128-bit argument size from 161 KiB to 97 KiB, and a field-size trade-off analysis shows that this small-field, larger-reserve configuration beats both a Goldilocks instantiation at $\eta = 1/32$ and a 128-bit-field instantiation at $\eta = 1/64$. Quotient hygiene—dithering $k = pn$ to $pn - 1$ —removes all exact dyadic quotient-core obstructions at rate loss $1/n$ and, by a maximal-remainder lemma, also defeats the natural remainder variant of the quotient-core construction. Without the final locator and line/MCA local-limit conjectures, the same machinery remains useful as an obstruction audit and as a guide to fallbacks: larger challenge fields with matching extension statements, dimension dithering, domain shattering, hybrid schedules, or theorem-backed folded and subspace-design codes.

This strengthened version adds a theorem-backed high-agreement ledger compiler: in the exact tangent range the line/CA/projective-line numerator is $r + 1$, interleaved lists contribute the single numerator unit 1, and finite power-curve branches contribute $d(r + 1)$ in the split moving-root model. These formulas give a concrete coding numerator certificate for the small-radius boundary layer while preserving the separation between protocol denominators, query/folding losses, and near-capacity local-limit assumptions.

Contents

1	Introduction	3
2	Protocol-facing code objects	4
2.1	Distances, agreement size, and reserve	4
2.2	Locator fibers and the list bridge	5
2.3	Interleaved lists	5
2.4	Extension fields	6
2.5	MCA, line-decoding, and curves	6
3	Corrected Reed–Solomon input	7
3.1	Generated-field list entropy	7
3.2	Quotient-core profile	8
3.3	Locator local-limit assumption	8
3.4	MCA floors and corrected MCA assumption	9
4	Reserve certificates	10
5	Quotient hygiene and dimension dithering	12
6	Fields and challenge accounting	13
7	Theorem-backed high-agreement ledger compiler	14
7.1	Same-denominator formulas	15
7.2	Different-denominator formulas	15
7.3	Worked boundary check: $\mathbb{F}_{17^{32}}$, $n = 512$, $k = 256$	15
7.4	Certificate row to print	16
8	Reserve-certified proximity layer	16
9	Efficiency estimates	18
9.1	Blowup from capacity-like reserves	18
9.2	Corrected toy comparison against the survey parameters	18
9.3	Hybrid schedule	19
9.4	Prover and verifier work	19
10	Modes and fallback strategy	19
10.1	Conjectural aggressive mode	19
10.2	Theorem-backed certificate mode	19
10.3	Obstruction-audit mode	20
10.4	Fallback hierarchy	20
10.5	Domain shattering	20
11	Open problems	20
12	Conclusion	21

1 Introduction

A Reed–Solomon proximity SNARK commits to evaluations of low-degree polynomials on a structured domain and then proves that a committed word is close to a code. Write

$$C = \text{RS}[\mathbb{F}_q, H, k], \quad |H| = n, \quad \rho = k/n,$$

where $H \leq \mathbb{F}_q^\times$ is usually a power-of-two multiplicative subgroup or coset. A verifier that reasons at distance

$$\delta = 1 - \rho - \eta$$

is operating η below Singleton capacity; we call η the *reserve*. Johnson-regime analyses [19, 16, 9] use roughly

$$\delta_J = 1 - \sqrt{\rho}, \quad \eta_J(\rho) = \sqrt{\rho} - \rho.$$

For moderate rates this reserve is a large constant. The natural efficiency question is whether a FRI/WHIR-like layer [13, 5] can safely use $\eta = o(1)$, thereby increasing rate, reducing low-degree-extension blowup, and shrinking oracle commitments.

The answer is not governed by a one-line entropy rule. Smooth domains have quotient structure. The companion manuscript [1] identifies two list-decoding lower-bound scales and matching MCA floors: a generated-field coefficient-pigeonhole obstruction, quotient-core obstructions at internal subgroup scales, a tangent MCA floor of order n/q , and quotient-periodic bad-slope floors, building on the no-slack disproof of [2]; concurrent near-capacity failures over prime-field multiplicative subgroups appear in [20, 21]. The survey of Arnon–Boneh–Fenzi [4] stresses, independently, that protocol soundness consumes both MCA and a list-size term for an *interleaved* code: in their toy protocol the round error contains

$$\varepsilon_{\text{mca}}(C, \delta) + \frac{|\Lambda(C^{\equiv 2}, \delta)|}{|\mathbb{F}|},$$

not merely a base-code list bound. This paper translates those code facts into an explicit SNARK parameter ledger. The governing principle is:

Do not run a high-rate proximity layer from a single inequality such as “ η exceeds the entropy gap.” Emit a certificate that proves or assumes the exact list and MCA objects consumed by the protocol reduction—including interleaving arity, extension-field lifting, the list-size-over-field term, and the generated-field entropy floor—and that states plainly which entries are theorem-backed and which are conjectural.

The practical benefits remain attractive. If a proximity protocol needs target distance δ_0 , Johnson operation imposes roughly

$$\rho \leq (1 - \delta_0)^2,$$

whereas a capacity-like layer would use

$$\rho \leq 1 - \delta_0 - \eta.$$

At $\delta_0 = 1/2$, Johnson gives $\rho \leq 1/4$ and blowup at least 4. A reserve $\eta = 1/64$ would allow $\rho = 31/64$ and blowup $64/31 \approx 2.06$, a $1.94\times$ oracle-length improvement. Three caveats attach immediately, and all are quantitative statements rather than hedges. First, $\eta = 1/64$ is *entropy-feasible only when the generated field has at least roughly 64 bits* (Design Rule 3.1); over a 31-bit FFT prime it lies deep inside the proven-large-list region. Second, for exact deployed-rate dimensions over

base prime fields of at most about 150 bits, the companion’s theorem-backed challenge cap places $\eta \leq 1/64$ outside the certified safe region for support-wise line MCA at 128-bit targets. Third, the self-contained universal cap of the companion cap paper [3] pins the threshold below $1 - \rho - 2^{-9}$ for *every* field below 2^{256} , and its reach at the deployed sextic parameters is already gap 2^{-7} — so reserves of $1/128$ or below are certifiably unsafe there regardless of the challenge field. A certificate must therefore say whether a parameter point is theorem-backed, conjectural with explicit named assumptions, or merely an obstruction audit.

Contributions. This paper makes six structural contributions to the proximity-layer design problem.

- (1) It defines locator fibers in a protocol-sufficient form and proves the bridge from fiber bounds to list bounds (Lemma 2.2), so the locator local-limit assumption can be consumed directly by a soundness proof.
- (2) It states the locator and MCA assumptions with $k = \rho n + O(1)$, so dimension dithering such as $k = \rho n - 1$ is not excluded by the very assumptions it is meant to exploit, and the quotient profile is computed at the actual pair (n, k) .
- (3) It separates q_{gen} from q_{line} and q_{chal} . The generated field controls locator entropy unconditionally; the line or challenge field controls random-line denominators only after the protocol is analyzed over that field, via a proved extension statement or the explicit Assumption 2.4.
- (4) It adds the interleaved-list budget λ_{list} , with the explicit inequality $\widehat{L}_\mu(\delta)/q_{\text{line}} \leq 2^{-\lambda_{\text{list}}}$, matching the list-size-over-field term that appears in concrete reductions [4].
- (5) It replaces a black-box “compiler theorem” by a ledger theorem (Theorem 4.3) that applies exactly to reductions whose code-level errors are the listed list/MCA/curve/query/fold terms, together with a worked instantiation on the Arnon–Boneh–Fenzi toy protocol (Example 4.5).
- (6) It distinguishes theorem-backed certificates from obstruction audits, and it corrects the efficiency comparison accordingly: the headline near-capacity numbers are computed at *entropy-feasible* reserves, with an explicit field-size trade-off table (section 9).

Scope. The paper treats the Reed–Solomon proximity layer and its coding-theoretic soundness ledger. Commitment schemes, hash assumptions, recursion, Fiat–Shamir, arithmetization, and Merkle-path accounting are recorded only insofar as they determine list arity, line/curve arity, challenge field, and query/folding errors.

2 Protocol-facing code objects

2.1 Distances, agreement size, and reserve

Let $C = \text{RS}[\mathbb{F}, H, k]$ with $|H| = n$. For a target radius $\delta < \delta_{\min}(C)$ define the required agreement size

$$a = a(\delta) := \lceil (1 - \delta)n \rceil$$

and the reserve in agreement coordinates

$$\sigma := a - k, \quad \eta := \sigma/n.$$

When $k = \rho n$ exactly and $\delta = 1 - \rho - \eta$, this matches the simpler rate-based notation. The agreement-size notation is safer because dimension dithering deliberately changes k by $O(1)$.

The exact finite-length entropy ledger is

$$R_{\text{ent}}(a; q_{\text{gen}}) := (a - k) \log_2 q_{\text{gen}} - \log_2 \binom{n}{a}. \quad (1)$$

Here q_{gen} is the size of the field generated by the evaluation domain, hence the field containing the locator coefficients of a -subsets of H . If H lies in a proper subfield, a larger ambient extension field does not improve (1) unless the code and the locator problem are genuinely lifted.

2.2 Locator fibers and the list bridge

Definition 2.1 (Feasible locator fiber). Let $U : H \rightarrow \mathbb{F}$ be a received word and let $a \geq k$. Define

$$\text{Fib}_U(a) := \{S \subseteq H : |S| = a \text{ and there exists } p \in \mathbb{F}[X]_{<k} \text{ with } p(x) = U(x) \text{ for all } x \in S\}.$$

Equivalently, writing $L_S(X) = \prod_{x \in S} (X - x)$ for the locator of S and $\hat{U}$ for the degree- $< n$ interpolant of U on H , we have $S \in \text{Fib}_U(a)$ if and only if $\deg(\hat{U} \bmod L_S) < k$. In the companion paper finer fibers are organized by locator coefficients or moments; this coarse feasible fiber is the object needed for the list bridge below.

Lemma 2.2 (Fiber-to-list bridge). *For every received word $U : H \rightarrow \mathbb{F}$ and every $a \geq k$,*

$$|\Lambda(C, 1 - a/n, U)| \leq |\text{Fib}_U(a)|.$$

Consequently, a uniform bound $|\text{Fib}_U(a)| \leq n^{B_L}$ over all U implies $|\Lambda(C, 1 - a/n)| \leq n^{B_L}$.

Proof. Every codeword $p \in \Lambda(C, 1 - a/n, U)$ agrees with U on at least a points; choose a canonical a -subset S_p of its agreement set, so $S_p \in \text{Fib}_U(a)$. If $p \neq p'$ had $S_p = S_{p'} = S$, then p and p' are degree- $< k$ polynomials agreeing on the $a \geq k$ points of S , hence $p = p'$. The map $p \mapsto S_p$ is therefore injective. $\square$

2.3 Interleaved lists

Protocol reductions rarely consume the base-code list size alone. If the reduction bundles μ committed words, or if an extension-code view turns one extension symbol into several base-field coordinates, the relevant object is an interleaved list.

Definition 2.3 (Protocol list arity). Let C_{line} be the code over the field from which the protocol's combination challenge is sampled. The protocol list arity is the smallest integer $\mu \geq 1$ such that the reduction's list term is bounded by

$$|\Lambda(C_{\text{line}}^{\equiv \mu}, \delta)| / q_{\text{line}}.$$

A certificate records an explicit number $\hat{L}_\mu(\delta)$ with

$$|\Lambda(C_{\text{line}}^{\equiv \mu}, \delta)| \leq \hat{L}_\mu(\delta).$$

A conservative derivation is the trivial product bound

$$\hat{L}_\mu(\delta) := \hat{L}_1(\delta)^\mu, \quad (2)$$

which is adequate for the constant arities ($\mu = 2$ or 3) of typical batching steps. Sharper Gopalan–Guruswami–Raghavendra interleaving bounds may be used instead, but the certificate should print the exact lemma and its hypotheses: the GGR factor is independent of μ but grows rapidly as $\delta \rightarrow \delta_{\min}(C)$, so near capacity the trivial bound can be the better choice for constant μ [14, 4].

2.4 Extension fields

Suppose a base code $C_B : B^k \rightarrow B^n$ is used with challenges in an extension F/B of degree e . The extension code $C_F : F^k \rightarrow F^n$ is not the same object as C_B with a larger denominator. The survey records the list-side identity

$$|\Lambda(C_F, \delta)| = |\Lambda(C_B^{\equiv e}, \delta)| \quad (3)$$

under the usual extension-code presentation [4]. Extension challenges thus move the list ledger into an interleaved base-code ledger; they do not increase the generated-field entropy reserve.

The MCA side needs an analogous transfer, and it is *not* automatic from a base-field MCA conjecture: the companion’s line/MCA local limit is stated for lines $f + zg$ with f, g over, and z ranging over, the generated field. After one folding round of a protocol with extension challenges, the committed words live over F and the line experiment runs over F . We therefore isolate the needed transfer as an explicit assumption rather than silently substituting denominators.

Assumption 2.4 (Extension-line MCA lift). *Let $C_B = \text{RS}[\mathbb{F}_{q_{\text{gen}}}, H, k]$ satisfy an MCA bound $\varepsilon_{\text{mca}}(C_B, \delta) \leq N_{\text{mca}}/q_{\text{gen}}$ for a numerator N_{mca} , and let C_F be its extension code over a challenge field F with $|F| = q_{\text{chal}}$. Then the same numerator transfers:*

$$\varepsilon_{\text{mca}}(C_F, \delta) \leq \frac{N_{\text{mca}}^{1+o(1)}}{q_{\text{chal}}}$$

for the line family $\{f + zg : z \in F\}$ with $f, g \in F^H$, and likewise for the line-decoding form.

Remark 2.5 (Why the lift is plausible, and what is genuinely open). Three observations delimit [Assumption 2.4](#). (i) The tangent floor over the extension is $\lfloor \delta n \rfloor / q_{\text{chal}}$, so the target shape is consistent: extension challenges suppress the floor exactly as the assumption predicts. (ii) The quotient-periodic bad slopes constructed in the companion are values of symmetric functions of subsets of H ; they lie in $\mathbb{F}_{q_{\text{gen}}}$, and their *count* does not grow under field extension, so the known lower-bound families are compatible with a field-independent numerator. (iii) Unconditionally, the one-bad-parameter-per-support bound gives $\varepsilon_{\text{mca}}(C_F, \delta) \leq 2^n/q_{\text{chal}}$ over *any* field, so the lift holds trivially once $q_{\text{chal}} \geq 2^{n+\lambda}$; the assumption is only needed in the practical window $\text{poly}(n) \leq q_{\text{chal}} \leq 2^{o(n)}$. What is genuinely open is whether *new* bad slopes appear from witnesses that exist only over the extension—lines whose anchor, direction, and explaining codewords all use extension coordinates. No such family is currently known, but none is excluded. [Open Problem 11.2](#) records this as an open problem.

If [Assumption 2.4](#) is neither invoked nor replaced by a proved extension-field theorem, the certificate must use the field of the line experiment actually covered by the cited statement. In particular, a numerator proved over $\mathbb{F}_{q_{\text{gen}}}$ cannot simply be divided by q_{chal} .

2.5 MCA, line-decoding, and curves

For a linear code $C \subseteq \mathbb{F}^n$, MCA at radius δ bounds the fraction of slopes $z \in \mathbb{F}$ for which the line $f + zg$ has a false support-wise explanation: a support S with $|S| \geq (1 - \delta)n$ on which $f + zg$ matches a codeword while (f, g) does not match $C^{\equiv 2}$. A closely related and often cleaner protocol interface is *line-decoding*: if C is $(\delta, a_{\text{LD}}, n + 1)$ line-decodable, then

$$\varepsilon_{\text{mca}}(C, \delta) \leq a_{\text{LD}}/|\mathbb{F}|,$$

and some interactive-oracle-proof analyses consume line-decoding directly, occasionally with tighter constants [15, 17, 18, 4]. The companion’s corrected MCA conjecture is itself phrased through

residue-line packing, which is essentially the line-decoding form; [Open Problem 11.3](#) asks to make that alignment exact.

Some batching steps use power curves

$$f_0 + \gamma f_1 + \cdots + \gamma^d f_d$$

rather than affine lines. A ledger theorem must therefore carry either a curve-MCA term or an explicit reduction from the affine-line and univariate-power cases to the full polynomial-generator family, as in Bordage–Chiesa–Guan–Manzur [7]. The certificate records this as a separate error term $\text{Adv}_{\text{curve}}$, which is zero when the protocol uses only affine lines.

Finally, implementations typically commit to an s -interleaved presentation $\text{IRS}[\mathbb{F}, H, k, s]$ for FFT locality. Interleaving degrades MCA by at most the interleaving factor,

$$\varepsilon_{\text{mca}}(C^{\equiv \nu}, \delta) \leq \nu \cdot \varepsilon_{\text{mca}}(C, \delta), \quad (4)$$

and below the unique-decoding radius the factor disappears entirely, $\varepsilon_{\text{mca}}(C^{\equiv \nu}, \delta) = \varepsilon_{\text{mca}}(C, \delta)$ for $\delta < \delta_{\min}(C)/2$ [4, 12]. The certificate carries the factor ν explicitly.

3 Corrected Reed–Solomon input

This section states the code-level input in a form that survives dimension dithering and extension-field accounting.

3.1 Generated-field list entropy

The coefficient-pigeonhole obstruction of the companion paper produces, below the entropy scale, an explicit received word with list size at least $q_{\text{gen}}^{-(a-k)} \binom{n}{a}$. Polynomial list size therefore requires

$$(a - k) \log_2 q_{\text{gen}} \gtrsim \log_2 \binom{n}{a}. \quad (5)$$

A finite certificate uses the exact ledger (1) with an additive margin:

$$R_{\text{ent}}(a; q_{\text{gen}}) \geq \Gamma_{\text{ent}}. \quad (6)$$

An implementation preferring a multiplicative margin

$$(a - k) \log_2 q_{\text{gen}} \geq (1 + \gamma_{\text{ent}}) \log_2 \binom{n}{a}$$

should convert it by setting $\Gamma_{\text{ent}} = \gamma_{\text{ent}} \log_2 \binom{n}{a}$. We use the additive convention throughout, with distinct symbols Γ_{ent} , Γ_Q , Γ_M for the entropy, quotient-profile, and MCA slacks.

It is useful to record the asymptotic feasibility threshold. Let $\tau^*(\rho, q)$ be the unique solution of $\tau \log_2 q = H_2(\rho + \tau)$; then (5) forces $\eta \geq (1 - o(1))\tau^*(\rho, q_{\text{gen}})$, and $\tau^*(\rho, q) = (H_2(\rho) + o(1))/\log_2 q$ for large q .

Design Rule 3.1 (Generated-field feasibility). Before any other tuning, check $\eta \geq (1 + \gamma_{\text{ent}})\tau^*(\rho, q_{\text{gen}})$ using the exact ledger (6). This floor depends only on the generated field and binds regardless of the challenge field. [Table 1](#) tabulates $\tau^*(1/2, q)$ for common FFT fields. Note that every 2-smooth domain over a KoalaBear-type tower lies in the 31-bit base field—all 2-power roots of unity of relevant order already live there—so $q_{\text{gen}} \approx 2^{31}$ for *any* extension degree.

Table 1: Entropy-feasibility floor $\tau^*(1/2, q_{\text{gen}})$ for common generated fields, and the smallest convenient dyadic reserve clearing it with multiplicative margin near 2.

Generated field	$\log_2 q_{\text{gen}}$	$\tau^*(1/2, q_{\text{gen}})$	feasible dyadic η (margin ≈ 2)
KoalaBear / BabyBear class	31	$0.0322 \approx 1/31$	$1/16$
Goldilocks	64	$0.0156 \approx 1/64$	$1/32$
128-bit 2-adic prime	128	$0.0078 \approx 1/128$	$1/64$
Stark-prime class	252	$0.0040 \approx 1/252$	$1/128$

Two consequences deserve emphasis. First, $\eta = 1/64$ over a 31-bit generated field has entropy ratio $(\eta \log_2 q_{\text{gen}}) / H_2(\rho + \eta) \approx 0.48$ at rate $1/2$: it is not merely uncertified but lies in the region where the pigeonhole theorem *proves* exponential lists. Second, $\eta = 1/64$ over Goldilocks sits exactly at the threshold (ratio 1.0007), with no usable margin; the first comfortable dyadic reserve there is $1/32$. These facts reshape the efficiency comparison of [section 9](#).

3.2 Quotient-core profile

For $M \mid n$, quotient cores become dangerous when k is aligned with the quotient scale and the slack is smaller than that scale: the companion constructs, for $M \mid k$, $k/M \leq n/M - 1$, and $1 \leq \sigma < M$, a received word with at least $\binom{n/M-1}{k/M}$ codewords at agreement $k + \sigma$ [[1](#), Thm. 3.5]. The exact-divisibility profile used by the certificate is

$$\mathcal{Q}_H(a, k) := \max_{\substack{M \mid \gcd(n, k), \ M > 1, \ a - k < M \\ k/M \leq n/M - 1}} \log_2 \binom{n/M - 1}{k/M}, \quad (7)$$

with value $-\infty$ if the maximum is empty. The condition $M > 1$ removes the vacuous $M = 1$ edge case (the underlying construction needs $1 \leq \sigma < M$, hence $M \geq 2$), which would otherwise appear when $a = k$ or $\eta < 1/n$.

A list certificate with target exponent B_L requires

$$\mathcal{Q}_H(a, k) \leq B_L \log_2 n + \Gamma_Q. \quad (8)$$

The profile is computed at the actual pair (n, k) *after* dimension dithering, not at an idealized exact value $k = \rho n$. In scale form, the quotient-core family forces $\sigma = \Omega(n / \log n)$ for unspecified polynomial list size and $\sigma \gtrsim (H_2(\rho) / B_L) n / \log_2 n$ for a fixed list exponent B_L .

3.3 Locator local-limit assumption

Assumption 3.2 (Field-aware locator local limit). *Fix $\rho \in (0, 1)$, $B_L > 0$, and $\varepsilon > 0$. There is a constant $C_{\rho, B_L, \varepsilon}$ such that the following holds. Let $H_n \leq \mathbb{F}_{q_n}^\times$ be generated-field smooth domains of order $n = 2^m$ (so $q_n = q_{\text{gen}}$), and let $k_n = \rho n + O(1)$. For every received word $U : H_n \rightarrow \mathbb{F}_{q_n}$,*

$$|\text{Fib}_U(k_n + \sigma_n)| \leq n^{B_L}$$

whenever

$$\sigma_n \geq C_{\rho, B_L, \varepsilon} \frac{n}{\log_2 n}, \quad \sigma_n \log_2 q_n \geq (1 + \varepsilon) \log_2 \binom{n}{k_n + \sigma_n},$$

and all active quotient-core profiles satisfy (8). Equivalently, after quotient-periodic classes are accounted for, locator fibers behave like random codimension- σ_n fibers on $(k_n + \sigma_n)$ -subsets.

The $O(1)$ freedom in k_n is intentional: it makes the assumption compatible with quotient hygiene, whose whole point is to replace exact deployed dimensions such as $k = \rho n$ by nearby dimensions such as $k = \rho n - 1$. Via [Lemma 2.2](#), the assumption yields $|\Lambda(C, 1 - \rho - \eta_n)| \leq n^{B_L}$ at the corrected reserve.

Theorem-backed lanes. The companion proves several special lanes and obstruction families. In characteristic zero, equal prefix moments on power-of-two roots of unity are exactly quotient-periodic, so characteristic-zero prefix fibers are polynomial once $\sigma \geq Cn/\log n$ [[1](#), Thm. 4.2, Cor. 4.3]. In split prime fields $p \equiv 1 \pmod{n}$, a Galois-amplified collision sieve [[1](#), Thm. 5.2, Cor. 5.3] transfers this structure whenever

$$p > \exp\left(C \frac{n \log n}{\sigma}\right),$$

which at the corrected reserve is the quasi-polynomial range $p > \exp(O((\log n)^2))$. These results are valuable for auditing and for monomial-prefix instances. They are not, by themselves, a full arbitrary-word polynomial-field list theorem.

3.4 MCA floors and corrected MCA assumption

The companion proves unavoidable MCA floors. The tangent floor [[1](#), Prop. 9.4] states that for $C = \text{RS}[\mathbb{F}, H, k]$, $\delta \in (0, 1 - \rho)$, and $\lfloor \delta n \rfloor \leq q$,

$$\varepsilon_{\text{mca}}(C, \delta) \geq \lfloor \delta n \rfloor / q, \tag{9}$$

where q is the size of the field over which the line experiment is actually run. A raw MCA target below n/q is therefore impossible without enlarging that field.

There is also a quotient-periodic floor. Let $\beta(\rho)$ denote the companion’s quotient value-set exponent; for deployed rates $\rho \in \{1/2, 1/4, 1/8, 1/16\}$, the ratios $H_2(\rho)/\beta(\rho)$ are approximately 1.262, 1.082, 1.024, and 1.009. Above the quotient norm threshold, the companion proves [[1](#), Prop. 19.1, Thm. 19.4]

$$\varepsilon_{\text{mca}}(C, 1 - \rho - \eta) \geq \frac{2^{(\beta(\rho)/H_2(\rho))\mathcal{Q}_H(a,k)(1-o(1))}}{q}. \tag{10}$$

A clean numerator $n^{1+o(1)}$ is therefore available only when quotient-periodic contributions are removed by parameter choice or explicitly paid for in the challenge budget.

Universal-cap floor. A third floor, proved self-contained in the companion cap paper [[3](#)], binds the line field itself. Whenever some divisor $N \mid n$ has $n/N \mid k$ and $\binom{N}{\rho N+2} \geq q_{\text{gen}}(q_{\text{line}}/k + 1)$, every radius $\delta \in [1 - \rho - 2/N, 1 - \rho)$ carries

$$\varepsilon_{\text{mca}}(C, \delta) \geq \frac{1}{2k} \left(1 - \frac{n}{q_{\text{line}}}\right), \tag{11}$$

where q_{line} is the field from which the line challenge is actually drawn. At the prize rates with $2^{10} \mid n$ ($2^{11} \mid n$ at rate $1/16$) this rules out every reserve $\eta \leq 2^{-9}$ (resp. 2^{-10}) over every field below 2^{256} , and structurally the mechanism reaches gap $\approx H_2(\rho)/\log_2 q_{\text{line}}$: the floor moves with the line field but is never removed by enlarging it, which quantitatively bounds the “increase q_{line} ” repair of [section 10](#). At the deployed sextic parameters the certified reach is gap 2^{-7} with error $> 2^{-22}$.

Assumption 3.3 (Line/MCA local limit over the actual line field). *Let $H_n \leq \mathbb{F}_{q_n}^\times$ be generated-field smooth domains of order n , let $k_n = \rho n + O(1)$, and let the line experiment be defined over $\mathbb{F}_{q_n}$ for the code used by the protocol. Assume $\text{poly}(n) \leq q_n \leq 2^{o(n)}$. If*

$$\sigma_n \geq C \frac{n}{\log_2 n}, \quad \sigma_n \log_2 q_n \geq (1 + \varepsilon) \log_2 \binom{n}{k_n + \sigma_n},$$

then, with $a_n = k_n + \sigma_n$ and $\eta_n = \sigma_n/n$,

$$\varepsilon_{\text{mca}}(\text{RS}[\mathbb{F}_{q_n}, H_n, k_n], 1 - \rho - \eta_n) \leq \frac{n^{1+o(1)} + 2^{(\beta(\rho)/H_2(\rho))\mathcal{Q}_{H_n}(a_n, k_n)(1+o(1))}}{q_n}. \quad (12)$$

The statement may equivalently be supplied in line-decoding form with numerator a_{LD} , using the implication $(\delta, a_{\text{LD}}, n+1)$ line-decodable $\Rightarrow \varepsilon_{\text{mca}} \leq a_{\text{LD}}/q_n$.

In a concrete certificate the $n^{1+o(1)}$ term is recorded as n^{A_M} for an explicit conservative exponent A_M , and the $(1 + o(1))$ in the quotient exponent is absorbed into the slack Γ_M .

Remark 3.4 (No automatic q_{chal} substitution). If [Assumption 3.3](#) is invoked over the generated field, then $q_n = q_{\text{gen}}$ in (12). To divide by $q_{\text{chal}} > q_{\text{gen}}$, one must either prove the bound over the extension code directly or invoke [Assumption 2.4](#). Conflating the two fields is the precise ambient-field accounting mistake that the generated-field pigeonhole already refutes on the list side.

4 Reserve certificates

A proximity implementation should output a small, checkable certificate for the code-level part of its security statement.

Definition 4.1 (Field-separated proximity reserve certificate). A certificate for a Reed–Solomon proximity layer consists of the tuple

$$\text{Cert} = (q_{\text{arith}}, q_{\text{gen}}, q_{\text{line}}, H, n, k, \rho, \delta, a, \sigma, \mu, \nu, e, B_L, A_M, \Gamma_{\text{ent}}, \Gamma_Q, \Gamma_M, C_{\text{loc}}, \lambda_{\text{list}}, \lambda_{\text{mca}}),$$

where μ is the protocol list arity of [Definition 2.3](#), ν is the implementation interleaving factor of the committed presentation (e.g. $\nu = s$ for $\text{IRS}[\cdot, \cdot, \cdot, s]$), and e is the extension degree $[\mathbb{F}_{q_{\text{line}}} : \mathbb{F}_{q_{\text{gen}}}]$ (with $e = 1$ when no lift is used); together with the following evidence.

- (a) **Generated-field ledger.** The exact value $R_{\text{ent}}(a; q_{\text{gen}})$ and a proof of (6), including the feasibility check of [Design Rule 3.1](#).
- (b) **Quotient profile.** The complete divisor list used to compute $\mathcal{Q}_H(a, k)$ in (7) at the actual (n, k) , and a proof of (8).
- (c) **Locator/list bound.** A theorem or assumption giving a base-list bound $\widehat{L}_1(\delta)$, usually $\widehat{L}_1(\delta) \leq n^{B_L}$ via [Assumption 3.2](#) and [Lemma 2.2](#).
- (d) **Interleaved-list bound.** A documented bound

$$|\Lambda(C_{\text{line}}^{\equiv \mu}, \delta)| \leq \widehat{L}_\mu(\delta), \quad (13)$$

obtained by the trivial product bound (2), by the extension identity (3), or by a sharper cited interleaving theorem with its hypotheses.

(e) **List-over-field budget.** The explicit inequality

$$\frac{\widehat{L}_\mu(\delta)}{q_{\text{line}}} \leq 2^{-\lambda_{\text{list}}}. \quad (14)$$

(f) **MCA or line-decoding budget.** A theorem or assumption for the line family actually used by the protocol, over the field actually used. For the corrected RS assumption this is recorded as

$$\frac{\nu \left(n^{A_M} + 2^{(\beta(\rho)/H_2(\rho))\mathcal{Q}_H(a,k)+\Gamma_M} \right)}{q_{\text{line}}} \leq 2^{-\lambda_{\text{mca}}}, \quad (15)$$

where the factor ν comes from (4) and may be replaced by a sharper theorem for $\varepsilon_{\text{mca}}(C^{\equiv\nu}, \delta)$, or removed below the unique-decoding radius [12].

(g) **Failure-ladder audit.** A scan of theorem-backed lower-bound rungs active at the chosen radius: tangent floors (9), restricted-sum and two-moment rungs, quotient-exact floors above their norm thresholds, and any deployed-rate challenge-cap theorem.

(h) **Extension status.** A flag saying whether C_{line} is the base code, an extension code, an interleaved base-code presentation, or a curve-generated family. If $q_{\text{line}} = q_{\text{chal}} > q_{\text{gen}}$, the certificate must point to the extension-list derivation (3) and to either a proved extension-line MCA theorem or [Assumption 2.4](#).

The certificate is cheap. On a dyadic domain it records the divisor list of $\gcd(n, k)$, two high-precision logarithms for the entropy ledger, and a failure scan that enumerates divisors of n against a small set of slack values; this is negligible next to proof generation, and it gives auditors a concrete account of why a high-rate layer avoids the known bad slices.

Design Rule 4.2 (Corrected reserve rule). Choose the smallest agreement reserve $\sigma = a - k$ for which all of the following hold:

$$R_{\text{ent}}(a; q_{\text{gen}}) \geq \Gamma_{\text{ent}}, \quad (16)$$

$$\sigma \geq C_{\text{loc}} \frac{n}{\log_2 n}, \quad (17)$$

$$\mathcal{Q}_H(a, k) \leq B_L \log_2 n + \Gamma_Q, \quad (18)$$

$$\delta < \delta_{\min}(C) \quad \text{and} \quad \eta > \eta_{\text{fail}}, \quad (19)$$

$$\widehat{L}_\mu(\delta)/q_{\text{line}} \leq 2^{-\lambda_{\text{list}}}, \quad (20)$$

$$\varepsilon_{\text{mca}}^{\text{used}}(C_{\text{line}}, \delta) \leq 2^{-\lambda_{\text{mca}}}. \quad (21)$$

Here η_{fail} is the largest gap already ruled out by a proved lower-bound ladder for the target soundness. Bad slopes propagate to larger radii, hence to smaller gaps, so a point with $\eta \leq \eta_{\text{fail}}$ is outside the certified safe region.

The phrase “smallest σ ” should be read with care: the entropy and local-limit conditions are monotone in the usual regime $q_{\text{gen}} \geq n$ at fixed k/n , while quotient-profile and failure-ladder conditions can jump at divisors. Implementations should enumerate candidate σ values around active divisor thresholds rather than binary-search a presumed-monotone predicate.

Theorem 4.3 (Ledger theorem for a reserve-certified layer). *Consider a FRI/WHIR/IOP layer whose code-level soundness reduction has been written in the form*

$$\text{Adv}_{\text{code}} \leq \text{Adv}_{\text{fold}} + \text{Adv}_{\text{curve}} + |\Lambda(C_{\text{line}}^{\equiv\mu}, \delta)| \cdot \text{Adv}_{\text{query}} + \frac{|\Lambda(C_{\text{line}}^{\equiv\mu}, \delta)|}{q_{\text{line}}} + \varepsilon_{\text{mca}}^{\text{used}}(C_{\text{line}}, \delta), \quad (22)$$

where any of the five terms may be absent for a given protocol. Assume the certificate of [Definition 4.1](#) satisfies [Design Rule 4.2](#) and that its cited assumptions or theorems apply to the actual code, interleaving arities μ and ν , curve family, and line field used in the reduction. Then

$$\text{Adv}_{\text{code}} \leq \text{Adv}_{\text{fold}} + \text{Adv}_{\text{curve}} + \widehat{L}_\mu(\delta) \cdot \text{Adv}_{\text{query}} + 2^{-\lambda_{\text{list}}} + 2^{-\lambda_{\text{mca}}}. \quad (23)$$

If the query/folding ledger is otherwise unchanged, the proximity layer may use rate $\rho = k/n$ at radius $\delta = 1 - a/n$ instead of reverting to the Johnson radius.

Proof. The locator evidence gives the base-list bound, [Lemma 2.2](#) converts locator fibers into code-word lists, and the interleaving evidence gives (13); substituting bounds term by term into (22) and applying (14) and (21) gives (23). The failure-ladder audit ensures the cited positive statements are not invoked at a parameter point already contradicted by a lower-bound theorem. $\square$

Remark 4.4 (Why this is a ledger theorem, not a universal compiler). Different IOPs consume different coding objects: base-code lists, interleaved lists, extension-code lists, affine-line MCA, power-curve MCA, or line-decoding. [Theorem 4.3](#) applies exactly to reductions rewritten in the form (22); it becomes a drop-in compiler only after that rewriting, which is itself [Open Problem 11.7](#).

Example 4.5 (Worked instantiation: the Arnon–Boneh–Fenzi toy protocol). The toy protocol of [4] tests that two committed words are close to constrained codewords by sending one combination challenge γ and spot-checking t positions. Its round-by-round knowledge errors are

$$\left(\varepsilon_{\text{mca}}(C, \delta) + \frac{|\Lambda(C^{\equiv 2}, \delta)|}{|\mathbb{F}|}, \quad (1 - \delta)^t \right).$$

This is (22) with $\mu = 2$, $\text{Adv}_{\text{fold}} = \text{Adv}_{\text{curve}} = 0$, no list multiplier on the query term, and $\text{Adv}_{\text{query}} = (1 - \delta)^t$ budgeted separately at λ_{query} bits. A corrected-reserve certificate therefore needs: $\widehat{L}_2(\delta) = \widehat{L}_1(\delta)^2 \leq n^{2B_L}$ by (2); the budget $n^{2B_L}/q_{\text{line}} \leq 2^{-\lambda_{\text{list}}}$; the MCA budget (15) with ν equal to the implementation interleaving of the commitment; and $t \geq \lambda_{\text{query}}/\log_2 \frac{1}{1-\delta}$. [Section 9](#) evaluates exactly this ledger at concrete parameters.

5 Quotient hygiene and dimension dithering

Quotient cores are exact-alignment phenomena. The cheapest practical repair is often to change the claimed degree bound by one coefficient.

Proposition 5.1 (Dyadic one-step hygiene). *Let $n = 2^m$ and let k_0 be even. If $k = k_0 - 1$, then $\gcd(n, k) = 1$. Hence the exact-divisibility profile (7) is empty for all dyadic quotients $M > 1$, and all exact quotient-core obstructions with $M \mid \gcd(n, k)$ disappear. The rate loss relative to k_0 is exactly $1/n$.*

Proof. The integer $k_0 - 1$ is odd, and the only odd divisor of 2^m is 1. $\square$

Exact divisibility is not the only worry. The quotient-core construction generalizes to dimensions with remainder $r = k \bmod M$: padding the witness polynomial by X^r produces, for $1 \leq \sigma < M$, a received word with $\binom{n/M-1}{\lfloor k/M \rfloor}$ nearby codewords at agreement $k + \sigma - r$ rather than $k + \sigma$.¹ The effective slack at gap η is therefore reduced by r , and the family stays dangerous whenever r is small. One-step dithering is in fact *maximally* robust against this variant:

¹With $k = M \lfloor k/M \rfloor + r$, take $Y = X^k L_T$ for a σ -subset T of one M -coset and, for each union of $\lfloor k/M \rfloor$ cosets U_A avoiding it, with locator $L_A = X^{k-r} + c_1 X^{k-r-M} + \dots$, set $P_A = L_T X^r (X^{k-r} - L_A)$. Then $\deg P_A \leq \sigma + r + (k - r - M) < k$ when $\sigma < M$, and P_A agrees with Y on $T \cup U_A$, a set of size $k + \sigma - r$.

Lemma 5.2 (Maximal-remainder hygiene at deployed dyadic rates). *Let $n = 2^m$, let $\rho = 2^{-b}$ with $1 \leq b < m$, and set $k_0 = \rho n = 2^{m-b}$ and $k = k_0 - 1$. For every dyadic quotient scale $M = 2^j$ with $1 \leq j \leq m - b$,*

$$k \equiv M - 1 \pmod{M}.$$

Thus the remainder variant of the quotient-core construction has maximal remainder at every quotient scale that divided the original deployed dimension k_0 , and in the active window $\sigma < M$ its agreement $k + \sigma - r = k + \sigma - M + 1 \leq k$ never exceeds trivial agreement.

Proof. For $j \leq m - b$, $M = 2^j$ divides k_0 , so $k = k_0 - 1 \equiv -1 \equiv M - 1 \pmod{M}$. With $r = M - 1$ and $\sigma < M$, the construction's agreement size is $k + \sigma - (M - 1) \leq k$, which produces no codewords beyond the agreement already guaranteed by interpolation. $\square$

Remark 5.3 (What the remainder lemma buys). [Proposition 5.1](#) removes the exact-divisibility family; [Lemma 5.2](#) removes its natural remainder generalization at the same stroke. This is attack removal for the theorem-backed quotient-core families, not a proof that all lists are small: the local-limit scale $\sigma \gtrsim n/\log n$ of [Assumption 3.2](#) may still be necessary. But it means the explicit superpolynomial quotient lists, the quotient-periodic MCA floor (10), and the deployed-rate challenge-cap mechanism (whose quotient machinery requires $n/N' \mid k$) are all disabled by a $1/n$ rate loss.

Design Rule 5.4 (Dimension dithering). On dyadic FFT domains, avoid dimensions with large 2-adic valuation. Start from the arithmetization degree bound k_0 and try $k = k_0 - 1$ when k_0 is even; more generally, choose the largest feasible $k \leq k_0$ with a small divisor profile. The certificate should print $\gcd(n, k)$ and all active divisors, not merely the nominal rate.

In practice a one-coefficient loss can be achieved by reserving one top-degree coefficient, padding the trace by one column, splitting a high-degree relation into two lower-degree relations, or rounding the evaluation domain upward and the degree bound downward. These repairs preserve the FFT domain; they only change the quotient arithmetic of the claimed degree bound. For mixed-smoothness domains the same rule applies prime by prime: minimize the large divisors of $\gcd(n, k)$, eliminating first the divisors M just above σ for which k/M has constant density in n/M .

6 Fields and challenge accounting

The corrected ledger distinguishes three fields.

Arithmetic field $\mathbb{F}_{q_{\text{arith}}}$.

The field in which the trace, constraints, and committed values are represented.

Generated field $\mathbb{F}_{q_{\text{gen}}}$.

The smallest field containing the evaluation domain, hence the locator coefficients used in the entropy and list problem.

Line or challenge field $\mathbb{F}_{q_{\text{line}}}$.

The field over which random linear combinations, line-decoding, MCA, power curves, DEEP points, or folding challenges are actually analyzed. Often $q_{\text{line}} = q_{\text{chal}}$, but only after the protocol proof is lifted to that field.

Design Rule 6.1 (No double-crediting of extension fields). Use q_{gen} in the entropy ledger (1) and the feasibility check of [Design Rule 3.1](#). Use q_{line} in (14) and (15) only for the code and line family actually covered by the reduction. If $q_{\text{line}} > q_{\text{gen}}$, the certificate must include the extension-code list derivation (3) and an MCA/line-decoding statement over the extension—a proved theorem or [Assumption 2.4](#).

Design Rule 6.2 (Challenge-size target). After fixing the interleaved-list bound and MCA numerator, choose q_{line} so that

$$\log_2 q_{\text{line}} \geq \max \left\{ \lambda_{\text{list}} + \log_2 \widehat{L}_\mu(\delta), \lambda_{\text{mca}} + \log_2 N_{\text{mca}} \right\},$$

$$N_{\text{mca}} := \nu \left(n^{A_M} + 2^{(\beta(\rho)/H_2(\rho))\mathcal{Q}_H(a,k)+\Gamma_M} \right).$$

When the quotient profile is empty and $N_{\text{mca}} \approx \nu n$, the MCA side asks for about $\lambda_{\text{mca}} + \log_2(\nu n)$ bits. If the interleaved list is bounded only by $n^{\mu B_L}$, the list side asks for $\lambda_{\text{list}} + \mu B_L \log_2 n$ bits and typically dominates.

The companion’s deployed-rate challenge cap [1, Prop. 19.7] is a warning, not a tuning annoyance. For exact deployed dimensions $\rho \in \{1/2, 1/4, 1/8, 1/16\}$ on smooth domains with $64 \mid n$, over base prime fields of size at most about 2^{150} (per-rate reaches between $2^{150.6}$ and $2^{161.8}$), the MCA threshold at target 2^{-128} is provably below $1 - \rho - 1/64$. Repairs, in increasing cost: dither k to disable the quotient attack family ([Lemma 5.2](#)); increase η ; sample line challenges from a larger field with a valid extension-line statement; split the proximity claim; or fall back to a Johnson or theorem-backed folded/subspace-design layer. Beyond the prime-field reach the universal cap of [3] keeps the pressure on: every field below 2^{256} has the threshold pinned below $1 - \rho - 2^{-9}$ at the prize rates (2^{-10} at rate 1/16), so the larger-challenge-field repair buys room only down to the $\approx H_2(\rho)/\log_2 q_{\text{line}}$ floor of (11), never past it.

7 Theorem-backed high-agreement ledger compiler

Status: theorem-backed coding ledger, not a deployed protocol theorem. The reserve certificate of the previous sections is designed for the near-capacity regime. There is, however, a complementary high-agreement regime where the coding numerators consumed by the protocol are exactly computable. We record it here because it is a useful boundary condition for any concrete SNARK/proximity reduction and a regression test for the field-accounting rules.

Let

$$C = \text{RS}[F, D, k], \quad |D| = n,$$

write $a = n - r$, and define

$$R_{\text{line}} = \left\lfloor \frac{n - k}{3} \right\rfloor, \quad R_{\text{list}} = \left\lfloor \frac{n - k}{2} \right\rfloor, \quad B(Q, \lambda) = \left\lfloor \frac{Q}{2^\lambda} \right\rfloor.$$

The companion slack paper proves the following coding ledger theorem.

Theorem 7.1 (high-agreement coding numerators). *For $0 \leq r \leq R_{\text{line}}$, the affine line, no-loss CA, and projective-line support-wise numerators are all exactly*

$$r + 1.$$

For every interleaving arity $\mu \geq 1$, if $0 \leq r \leq R_{\text{list}}$, then

$$|\Lambda(C^{\equiv\mu}, n - r)| = 1$$

for the worst-case interleaved list numerator. For the finite degree- d power-curve/evaluation-domain model in the split moving-root regime, if

$$r \leq R_d := \left\lfloor \frac{n - k}{d + 2} \right\rfloor,$$

then the curve numerator is $d(r + 1)$. Without the split moving-root check, the safe theorem-backed envelope is an upper bound by $d(r + 1)$ and a line lower bound by $r + 1$.

Remark 7.2 (projective denominators). If projective slopes are sampled uniformly from $\mathbb{P}^1(F)$, the numerator remains $r + 1$ but the denominator is $|F| + 1$, not $|F|$. If the protocol samples an affine chart, the denominator is the affine line field. The certificate must print which sampler is used.

7.1 Same-denominator formulas

Suppose a reduction consumes s line/CA/projective-line branches, finite power-curve branches of degrees $d_1, \dots, d_t$, and ℓ independent interleaved-list terms, all divided by the same challenge denominator Q , and suppose all terms lie in their high-agreement theorem-backed ranges. Then the coding numerator is

$$N_{\text{code}}(r) = \ell + s(r + 1) + \sum_{i=1}^t d_i(r + 1).$$

The exact same-denominator coding safety condition at target $2^{-\lambda}$ is

$$N_{\text{code}}(r) \leq B(Q, \lambda).$$

For a single line term this is $r + 1 \leq B(Q, \lambda)$; for one line term plus one interleaved-list term it is $r + 2 \leq B(Q, \lambda)$; and for one degree- d curve plus one list term it is $d(r + 1) + 1 \leq B(Q, \lambda)$.

7.2 Different-denominator formulas

If the line, list, and curve terms are paid by different fields, numerators must not be merged. The certificate must check the rational inequality

$$\sum_j \frac{N_j(r)}{Q_j} + \varepsilon_{\text{query}} + \varepsilon_{\text{fold}} + \varepsilon_{\text{crypto}} \leq 2^{-\lambda}.$$

This is the high-agreement analogue of the field-separation rules in [Section 6](#): q_{gen} , q_{line} , q_{chal} , projective line denominators, and curve parameter spaces remain distinct unless a theorem identifies them.

7.3 Worked boundary check: $\mathbb{F}_{17^{32}}$, $n = 512$, $k = 256$

For

$$C = \text{RS}[\mathbb{F}_{17^{32}}, H, 256], \quad |H| = 512,$$

we have

$$B(17^{32}, 128) = 6, \quad R_{\text{line}} = 85, \quad R_{\text{list}} = 128.$$

Thus, inside the high-agreement theorem-backed range,

line alone: largest safe radius $r = 5$,

whereas

line plus one interleaved-list term: largest safe radius $r = 4$.

Equivalently, an affine or projective line/CA/MCA term alone is below the 2^{-128} ledger at agreement $a = 507$, but a same-field line-plus-list ledger first becomes safe at $a = 508$. For a degree-2 finite power-curve term, the same-field curve-plus-list condition is

$$2(r + 1) + 1 \leq 6,$$

so the largest safe radius is $r = 1$. These are coding numerators only: query, folding, Fiat–Shamir, cryptographic binding, and sampler-to-model losses must be added separately.

7.4 Certificate row to print

A high-agreement certificate should print at least the following row:

code row	$n, k, \rho, D, q_{\text{gen}}, q_{\text{line}}, q_{\text{chal}}, \text{projective?}$
radius	$r = n - a$
line range	$r \leq R_{\text{line}}?$
list range	$r \leq R_{\text{list}}?$
line numerator	$r + 1$ if in range
list numerator	1 if in range
curve numerator	$d(r + 1)$ if split and in range; otherwise upper-envelope only
denominators	$Q_{\text{line}}, Q_{\text{proj}}, Q_{\text{curve}}, Q_{\text{list}}$
external terms	$\varepsilon_{\text{query}}, \varepsilon_{\text{fold}}, \varepsilon_{\text{crypto}}$
status	proved / conditional sampler model / unknown.

Remark 7.3 (how this strengthens the certificate paper). The high-agreement compiler does not replace the near-capacity reserve certificate. It gives an exact theorem-backed boundary layer. Therefore a protocol whose target radius lies in this layer can be certified by elementary coding numerators. A protocol whose radius lies outside this layer still needs the locator local-limit, all-line MCA, extension-line, quotient-hygiene, and interleaved-list machinery of the rest of the paper.

8 Reserve-certified proximity layer

Construction 8.1 (Reserve-certified RS proximity layer). Fix a target soundness allocation

$$(\lambda_{\text{list}}, \lambda_{\text{mca}}, \lambda_{\text{query}}, \lambda_{\text{fold}})$$

and a protocol reduction whose code-level terms have the form (22).

Step 1. Choose the target radius. Determine the largest protocol radius δ_0 tolerated by the folding/query analysis.

- Step 2. Check generated-field feasibility.** Determine q_{gen} from the domain family and test candidate reserves against [Design Rule 3.1](#). If the implementation field pins q_{gen} too low for the desired reserve, either accept the larger feasible reserve or change the domain family.
- Step 3. Choose a candidate rate.** Start from $\rho_0 = 1 - \delta_0$ and pick a feasible reserve η . Set $k \leq (1 - \delta_0 - \eta)n$.
- Step 4. Dither the dimension.** Apply [Design Rule 5.4](#); compute the actual $a = \lceil (1 - \delta_0)n \rceil$ and $\sigma = a - k$.
- Step 5. Compute generated-field entropy.** Check (16) and (17) exactly. If they fail, increase σ , enlarge the generated field, or lower k .
- Step 6. Enumerate quotient cores.** Compute (7) at the actual (n, k) . If (18) fails, dither again or increase σ past the active divisors.
- Step 7. Derive the interleaved-list bound.** Determine μ from the protocol proof. Derive $\widehat{L}_\mu$ by the trivial product bound, a GGR-type bound, or the extension identity (3).
- Step 8. Prove or assume line/MCA over the actual field.** Determine q_{line} and the factor ν . If $q_{\text{line}} = q_{\text{chal}} > q_{\text{gen}}$, record the extension-list and extension-MCA status. If batching uses power curves, record the curve-MCA status.
- Step 9. Scan proved failures.** Record tangent floors, quotient floors, challenge caps, and other lower-bound rungs. Set $\eta > \eta_{\text{fail}}$ as required by (19).
- Step 10. Finalize the soundness ledger.** Check (20) and (21); emit the certificate; add query/folding/hash/Fiat–Shamir terms outside the code certificate.

FRI interpretation. A FRI prover [13] benefits when low-degree-extension blowup $1/\rho$ falls. The corrected reserve changes the admissible code rate, not the ordinary FRI query analysis: keep the standard folding and sampling analysis and replace only the code-level list and MCA premises by the certificate. FRI folding combinations are affine lines, so $\text{Adv}_{\text{curve}} = 0$ for the folds themselves; batching steps may still need the curve entry.

WHIR interpretation. WHIR-style protocols [5] use random combinations and proximity reasoning in exactly the regime where MCA and interleaved lists matter. The certificate matches the toy analysis of [4]: it budgets both the MCA failure and the $|\Lambda(C^{\equiv\mu}, \delta)|/q_{\text{line}}$ list term. Quotient hygiene is particularly natural here because it changes only the claimed degree bound, never the multiplicative FFT domain.

DEEP and extension hybrids. DEEP and extension challenges [10] are valuable for the line/challenge denominators. They do not improve locator entropy unless the code and locator problem are themselves lifted, which for 2-smooth domains over small FFT primes is impossible without changing the domain family ([Design Rule 3.1](#)). A DEEP hybrid should therefore print both q_{gen} and q_{line} , the extension-list identity, and the extension-MCA status.

9 Efficiency estimates

9.1 Blowup from capacity-like reserves

For a target distance δ_0 , Johnson operation uses $\rho_J = (1 - \delta_0)^2$, while a capacity-like reserve η uses $\rho_{\text{cap}} = 1 - \delta_0 - \eta$. Table 2 keeps the illustrative $\eta = 1/64$ with an explicit status column; by Table 1, this reserve requires a generated field of at least roughly 128 bits (with margin), and the deployed-rate challenge cap must additionally be cleared by dithering or extension challenges.

Table 2: Illustrative blowup gains at reserve $\eta = 1/64$. Feasibility requires $\log_2 q_{\text{gen}} \gtrsim 128$ and a cleared challenge cap.

Target δ_0	Johnson rate	Corrected rate	Blowup improvement	Status
1/2	1/4	31/64	$4/(64/31) \approx 1.94\times$	conditional/audited
2/3	1/9	61/192	$9/(192/61) \approx 2.86\times$	conditional/audited
3/4	1/16	15/64	$16/(64/15) \approx 3.75\times$	conditional/audited

9.2 Corrected toy comparison against the survey parameters

The toy setting of [4] uses a sextic KoalaBear extension for challenges, message length $k = 2^{20}$, committed length $sn = 2^{21}$, and the argument-size proxy

$$t (256 \log_2 n + 2 \log_2 |B| \cdot s) \text{ bits}$$

for an s -interleaved presentation over base field B , minimized near $s = 8$, $n = 2^{18}$. Their theorem-backed Johnson-regime configuration at 128-bit security uses $t = 259$ and about 161.4 KiB; their theorem-backed capacity-like folded-RS configuration pays the subspace-design alphabet overhead and bottoms out at 281.2 KiB.

The generated field of every 2-smooth domain in this tower is KoalaBear itself, $\log_2 q_{\text{gen}} = 31$, regardless of the sextic challenge extension: by (3) the list problem reduces to the interleaved *base* code, whose locator coefficients live in $\mathbb{F}_p$. Design Rule 3.1 therefore forbids $\eta = 1/64$ here—its entropy ratio is 0.48, inside the proven-large-list region—and the smallest comfortable dyadic reserve is $\eta = 1/16$ (ratio 1.96). At $\rho = 1/2$ and $\delta = 1/2 - 1/16 = 7/16$, the per-query survival is $9/16$ and the query branch needs

$$t = \left\lceil \frac{128}{\log_2(16/9)} \right\rceil = 155,$$

giving $155 (256 \cdot 18 + 62 \cdot 8) \text{ bits} \approx 96.6 \text{ KiB}$. The challenge budget is comfortable: with the quotient profile emptied by dithering the per-row dimension to $k_{\text{row}} = 2^{17} - 1$ (total dimension $8(2^{17} - 1) = 2^{20} - 8$), $\nu = 8$, and $A_M \approx 1$, the MCA side asks for about $128 + \log_2(8n) \approx 149$ bits and the list side ($\mu = 2$, $B_L = 1$) for about $128 + 2 \log_2 n \approx 164$ bits, against 185.9 available—provided the extension-status entries (extension-list identity plus Assumption 2.4) are present, and noting that $\eta = 1/16$ at $n = 2^{18}$ presumes a local-limit constant $C_{\text{loc}} \leq 1.13$ in (17). The universal-cap rung is also cleared with margin: its structural reach at these parameters is $\text{gap} \approx H_2(1/2)/\log_2 q_{\text{line}} \approx 2^{-7.5}$, far below $\eta = 1/16$, and the certified failure interval $[1/2 - 2^{-7}, 1/2)$ does not approach $\delta = 7/16$ [3] — the cap bounds this configuration from outside rather than contradicting it.

Larger generated fields permit smaller reserves but pay more per opened element. Holding the survey’s structural proxy fixed and varying only the base-field bit width gives Table 3.

Table 3: Field-size/reserve trade-off on the survey’s toy proxy ($\rho = 1/2$, $n = 2^{18}$, $s = 8$, 128-bit target, query branch). Corrected rows are conditional on [Assumptions 2.4, 3.2 and 3.3](#); the Johnson and folded rows are theorem-backed.

Setting	$\log_2 q_{\text{gen}}$	feasible η	t	size
Survey IRS, Johnson radius	31	—	259	161.4 KiB
Survey folded RS, capacity-like	31	—	295	281.2 KiB
Corrected RS, KoalaBear	31	1/16	155	96.6 KiB
Corrected RS, Goldilocks-class	64	1/32	141	96.9 KiB
Corrected RS, 128-bit prime	128	1/64	134	108.9 KiB

Two lessons follow. First, the conjectural corrected route beats both theorem-backed routes by roughly $1.7\times$ against Johnson and $2.9\times$ against folded RS on this proxy. Second, and less obviously, the generated-field entropy floor creates a genuine trade-off in which the *small-field, larger-reserve* configuration wins: the query savings from shrinking η below 1/16 are smaller than the per-element opening cost of the wider base field needed to make such reserves entropy-feasible. The corrected theory thus does not simply say “use the smallest reserve”; it locates an interior optimum.

9.3 Hybrid schedule

A practical compromise is a hybrid schedule: use quotient-hygienic corrected RS in the early, large rounds where oracle length and LDE cost dominate, and switch to folded-RS or another subspace-design code in the final small rounds, where the $O(1/\eta^2)$ alphabet overhead of the theorem-backed capacity results is bounded [11, 15, 4]. This confines conjectural assumptions to the rounds where they buy the most and gives a theorem-backed landing zone near the end of the reduction.

9.4 Prover and verifier work

If an arithmetization degree d is fixed and blowup falls from b_J to b_{cap} , the FFT part of prover work scales roughly like

$$\frac{b_{\text{cap}} \log(b_{\text{cap}} d)}{b_J \log(b_J d)}.$$

Hashing and commitment work scale with oracle length. Extension-field challenges add verifier arithmetic and sometimes prover arithmetic, but when used only for combination coefficients or DEEP points their cost is far below returning the entire LDE to Johnson blowup.

10 Modes and fallback strategy

10.1 Conjectural aggressive mode

The fastest mode invokes [Assumptions 3.2 and 3.3](#) and, when $q_{\text{line}} > q_{\text{gen}}$, [Assumption 2.4](#). The security statement must name these assumptions and the exact line/curve family they cover. Appropriate for research systems and parameter exploration.

10.2 Theorem-backed certificate mode

This mode uses only proved positive results: Johnson-regime RS bounds [9, 8], subspace-design and folded-RS capacity theorems [11, 15], extension-code list identities, proved interleaving inequalities,

and the companion’s special lanes [1, Cor. 4.3, Cor. 5.3] (characteristic-zero prefix structure; Galois-amplified split-prime monomial-prefix bounds above $p > \exp(Cn \log n/\sigma)$). Slower, but with a conventional theorem-backed coding statement.

10.3 Obstruction-audit mode

This mode does not prove soundness. It scans for known bad regions: entropy deficits against [Design Rule 3.1](#), active quotient cores, tangent floors, quotient-exact floors above norm thresholds, deployed-rate and universal field-size challenge caps [3], and known CA/MCA lower-bound rungs [2, 1]. Passing the audit means the parameters avoid known attacks; it does not imply a proximity theorem, and a certificate produced in this mode must say so.

10.4 Fallback hierarchy

When a candidate high-rate certificate fails, repair in this order:

- (1) dither k to remove exact quotient cores and maximize remainders ([Proposition 5.1](#) and [Lemma 5.2](#));
- (2) increase q_{line} , but only together with the matching extension-code list derivation and extension-line statement, and only down to the universal-cap floor: no enlargement certifies gaps below $\approx H_2(\rho)/\log_2 q_{\text{line}}$ [3];
- (3) increase σ past active failure-ladder gaps;
- (4) split one large proximity claim into independently challenged lower-divisibility claims;
- (5) adopt the hybrid schedule, switching to folded/subspace-design codes in small rounds;
- (6) return to a Johnson or near-Johnson radius for the affected layer.

The hierarchy degrades smoothly: even when the conjectural constants disappoint, quotient hygiene and field-separated accounting remain valid and nearly free.

10.5 Domain shattering

Domain shattering splits a smooth subgroup into cosets or punctured pieces and runs separate proximity claims. It sacrifices some FFT regularity but destroys exact quotient-periodic structure. The far endpoint is not mysterious: randomly punctured Reed–Solomon codes achieve list-decoding capacity with high probability [6], and random-domain RS codes have MCA up to capacity with constant probability [15]. Shattering is therefore an interpolation between two regimes that are each understood—fully smooth FFT domains with their quotient pathology, and random evaluation sets with theorem-backed capacity behavior—and a clean theory of intermediate shattering is a promising route if fixed-prime local limits remain hard.

11 Open problems

Open Problem 11.1 (Polynomial-field locator local limit). *Prove [Assumption 3.2](#) for generated-field smooth domains with $q = \text{poly}(n)$ and arbitrary received words, with finite constants usable in (6) and (8).*

Open Problem 11.2 (Extension-line MCA). *Prove or refute [Assumption 2.4](#). The list side has the clean identity (3); the MCA side needs an equally explicit transfer theorem, or a counterexample exhibiting bad slopes whose witnesses exist only over the extension.*

Open Problem 11.3 (Line-decoding form of the corrected conjecture). *Restate the companion’s residue-line packing conjecture as a line-decoding statement with explicit $(\delta, a_{LD}, n + 1)$ parameters. This would align the assumption with protocol analyses that consume line-decoding directly and make the MCA conclusion a corollary.*

Open Problem 11.4 (Curve-MCA for batching). *For power-curve batching $\sum_i \gamma^i f_i$, prove a corrected-reserve theorem with the same quotient and field separation, or prove that affine-line MCA plus univariate-power MCA imply the needed polynomial-generator guarantee with concrete constants, sharpening [7] at the corrected reserve.*

Open Problem 11.5 (Sharp quotient-profile and local-limit constants). *Determine the finite-length constants in (8), (10), and C_{loc} in (17). The toy comparison of [section 9.2](#) presumes $C_{\text{loc}} \lesssim 1.1$; for proof systems, the difference between $C/\log n$ and $2C/\log n$ is a large concrete cost.*

Open Problem 11.6 (Sharp interleaved-list constants near capacity). *Determine the best interleaving bound for the concrete arities used in SNARKs near capacity. The trivial product bound is safe but overcharges λ_{list} linearly in μ ; GGR-style bounds are arity-independent but deteriorate as $\delta \rightarrow \delta_{\min}$.*

Open Problem 11.7 (Protocol-level reductions in ledger form). *Rewrite the FRI and WHIR soundness analyses so that their code-level premises are exactly the terms of (22). This would make [Theorem 4.3](#) a drop-in compiler rather than a design pattern, extending [Example 4.5](#) from the toy protocol to deployed systems.*

Open Problem 11.8 (Quotient-hygienic arithmetization). *Design AIR/R1CS/Plonkish layouts that naturally produce low-divisibility degree bounds without increasing constraint count, making $\mathcal{Q}_H(a, k)$ empty or logarithmic by construction.*

12 Conclusion

A high-rate Reed–Solomon proximity layer can be efficient only if its soundness ledger is exact about what the protocol consumes. The corrected certificate has six non-negotiable features: a generated-field entropy floor checked before all other tuning; quotient-core accounting at the actual, possibly dithered, dimension; a locator-fiber list bound bridged to the interleaved list actually used; an interleaved-list-over-field budget; MCA or line-decoding over the field in which the line experiment is analyzed; and a clear separation between theorem-backed proof and obstruction audit.

Under the final corrected locator and line/MCA local-limit assumptions, quotient-hygienic RS can plausibly replace Johnson-regime blowup by a near-capacity reserve in the large rounds of FRI/WHIR-style systems, with concrete argument-size gains near $1.7\times$ on the surveyed toy parameters—achieved, perhaps surprisingly, at a *small* generated field with a correspondingly larger reserve, because the entropy floor and per-element opening costs trade off at an interior optimum. The most robust practical improvements are low-cost and theorem-guided even before the conjectures are proved: check generated-field feasibility first, dither dimensions, never double-credit extension fields, budget λ_{list} explicitly, use extension challenges only with matching extension-code statements, and fall back to folded or subspace-design codes where a theorem-backed capacity guarantee justifies the alphabet overhead.

References

- [1] Przemysław Chojek. *Slack, Quotient Cores, and the Entropy Gap for Smooth-Domain Reed–Solomon Codes: List fibers, cyclotomic rigidity, Galois-amplified collisions, and the corrected slack mutual-correlated-agreement theory*. Companion manuscript, merged corrected edition, 2026.
- [2] Przemysław Chojek. *Capacity-Edge Obstructions to Reed–Solomon Mutual Correlated Agreement over Smooth Multiplicative Domains*. Companion manuscript, June 2026.
- [3] Przemysław Chojek. *A Universal Field-Size Cap for Mutual Correlated Agreement on Smooth Reed–Solomon Domains*. Companion manuscript, June 2026.
- [4] Gal Arnon, Dan Boneh, and Giacomo Fenzi. Open Problems in List Decoding and Correlated Agreement. Cryptology ePrint Archive, Report 2026/680, April 2026.
- [5] Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, and Eylon Yogev. WHIR: Reed–Solomon Proximity Testing with Super-Fast Verification. EUROCRYPT, 2025.
- [6] Omar Alrabiah, Venkatesan Guruswami, and Ray Li. Randomly Punctured Reed–Solomon Codes Achieve List-Decoding Capacity over Linear-Sized Fields. STOC, 2024.
- [7] Sarah Bordage, Alessandro Chiesa, Ziyi Guan, and Ignacio Manzur. All Polynomial Generators Preserve Distance with Mutual Correlated Agreement. Cryptology ePrint Archive, Report 2025/2051, 2025.
- [8] Eli Ben-Sasson, Dan Carmon, Ulrich Haböck, Swastik Kopparty, and Shubhangi Saraf. On Proximity Gaps for Reed–Solomon Codes. Cryptology ePrint Archive, Report 2025/2055, 2025.
- [9] Eli Ben-Sasson, Dan Carmon, Yuval Ishai, Swastik Kopparty, and Shubhangi Saraf. Proximity Gaps for Reed–Solomon Codes. FOCS, 2020.
- [10] Eli Ben-Sasson, Lior Goldberg, Swastik Kopparty, and Shubhangi Saraf. DEEP-FRI: Sampling Outside the Box Improves Soundness. ITCS, 2020.
- [11] Yeyuan Chen and Zihan Zhang. Explicit Folded Reed–Solomon and Multiplicity Codes Achieve Relaxed Generalized Singleton Bounds. STOC, 2025.
- [12] Benjamin Diamond and Angus Gruen. Proximity Gaps in Interleaved Codes. IACR Communications in Cryptology, 1(4), 2025. Cryptology ePrint Archive, Report 2024/1351.
- [13] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast Reed–Solomon Interactive Oracle Proofs of Proximity. ICALP, 2018.
- [14] Parikshit Gopalan, Venkatesan Guruswami, and Prasad Raghavendra. List Decoding Tensor Products and Interleaved Codes. SIAM Journal on Computing, 40(5):1432–1462, 2011.
- [15] Rohan Goyal and Venkatesan Guruswami. Optimal Proximity Gaps for Subspace-Design Codes and (Random) Reed–Solomon Codes. Cryptology ePrint Archive, Report 2025/2054, 2025.
- [16] Venkatesan Guruswami and Madhu Sudan. Improved Decoding of Reed–Solomon and Algebraic-Geometry Codes. IEEE Transactions on Information Theory, 45(6):1757–1767, 1999.

- [17] Ulrich Haböck. Basefold in the List Decoding Regime. Cryptology ePrint Archive, Report 2024/1571, 2024.
- [18] Ulrich Haböck. A Note on Mutual Correlated Agreement for Reed–Solomon Codes. Cryptology ePrint Archive, Report 2025/2110, 2025.
- [19] Selmer M. Johnson. A New Upper Bound for Error-Correcting Codes. IRE Transactions on Information Theory, 8(3):203–207, 1962.
- [20] D. Krachun, S. Kazanin, and U. Haböck. Failure of Proximity Gaps Close to Capacity. Cryptology ePrint Archive, Report 2026/782, 2026.
- [21] A. Kambiré. Proximity Gaps Conjecture Fails Near Capacity over Prime Fields. arXiv:2604.09724, 2026.