SHORTENING BOUNDS FOR REED–SOLOMON MCA

PRZEMEK CHOJECKI

Abstract. We derive an explicit exponent Ψ_ρ that bounds the Reed–Solomon MCA bad-slope numerator at every fixed relative radius between Johnson and capacity. The resulting positive-relative-radius exponential-budget safe-frontier certificate strictly improves the smallest-test MDS exponent and gives a constant post-Johnson radius for every positive usable budget exponent. An all-test-size MDS circuit-incidence envelope gives exact large-field capacity plateaux, improved adjacent thresholds, and four explicit length-512, 2^{-128}-secure smooth multiplicative certificates beyond Johnson. An exact CA–MCA decomposition gives challenge-restricted, endpoint-exact linear-budget thresholds through asymptotically half the minimum distance, and monomial equivalence transfers the applicable bounds to circle presentations. A Gowers–cube argument proves primitive max-fiber flatness from an image-normalized Sidon payment at an accessible moment order. The unrestricted subexponential-budget smooth/circle frontier remains open because shortening has positive exponential cost and the analytic payment, residual ray compiler, profile add-back, and matching attacks are not yet available.

Contents

1. Introduction and main results 2
2. MCA witnesses and exact support reduction 9
3. Exact syndrome geometry 13
4. The mean-overlap regime 16
5. Prior Reed–Solomon completion below half distance 20
6. Agreement-set shortening beyond Johnson 22
7. Reproducible finite Prize instances 25
8. Exact locator prefixes and collision-aware pole lines 26
9. Further exact constructions 29
10. Primitive inverse from a Sidon payment at accessible moment order 44
11. Asymptotic and finite regimes 46
12. Conclusion 48
Appendix A. Conditional parameter-retained interpolation audit 49
Appendix B. Verification details 51
References 54

Date: July 17, 2026.
2020 *Mathematics Subject Classification.* 94B35, 11T71, 05B40.
Key words and phrases. Reed–Solomon codes, mutual correlated agreement, Proximity Prize, agreement-set shortening, MDS codes, soundness gaps, circle codes.
Affiliation: Ulam.ai. Public profile: https://github.com/przchojecki.

1. Introduction and main results

By a finite Reed–Solomon *row* we mean data $(\mathbb{F}, D, k)$, where $\mathbb{F}$ is a finite field, $D \subseteq \mathbb{F}$ is an *evaluation domain* of size n, and $1 \le k < n$. The associated code $C = \mathrm{RS}_{\mathbb{F}}(D,k) \subseteq \mathbb{F}^D$ consists of the evaluations on D of polynomials of degree less than k. Thus a support is simply a subset of the coordinate set D. A word $u \in \mathbb{F}^D$ is *explained on a support* $S \subseteq D$ if its restriction to S agrees with one such degree-less-than-k polynomial. When the evaluation points lie in a designated coefficient subfield, we write $D \subseteq \mathbb{B} \subseteq \mathbb{F}$; $\mathbb{B}$ controls locator coefficients and domain structure, while $\mathbb{F}$ controls codeword coefficients and slopes. We call a linear $[n,k]$ code *maximum-distance separable* (MDS) when its minimum distance is $n-k+1$, equivalently when puncturing to every k-set is an isomorphism onto $\mathbb{F}^k$. Such a k-set is an *information set*. The *support atlas* at agreement $a \ge k+1$ is the injection obtained by assigning to each bad slope a noncommon exact-a support. Given a punctured generator matrix and one additional word, the *augmented-word matroid* is the column matroid of the matrix obtained by appending that word as one further row. A rank-s matroid is *paving* when every set of at most $s-1$ elements is independent. The *tangent construction* is the explicit received-line construction that produces at least $\min\{|\Gamma|, n-a+1\}$ bad slopes at agreement a; it is proved in Proposition 4.2. A *profile* is a declared witness mechanism, and a *first-match atlas* is a disjoint assignment of witnesses to the first applicable profile. Such an atlas is *witness-exhaustive* only if every bad slope is assigned to a profile. A *ray* is a projective direction, whereas MCA counts distinct affine slopes in the declared challenge set; a *ray compiler* is a theorem that bounds this projection. Finally, the *soundness gap* in bits is the interval between a rigorously proved security lower bound and an available attack upper bound. These descriptive terms supply no bound unless the corresponding properties are proved. Throughout, $h(x) = -x\log x - (1-x)\log(1-x)$ denotes binary entropy in natural units, and $H_2(x) = h(x)/\log 2$ denotes binary entropy in bits, with the usual convention $0\log 0 = 0$.

A received pair $r = (r_0, r_1) \in (\mathbb{F}^D)^2$ determines the *received line* $\{r_0 + \gamma r_1 : \gamma \in \mathbb{F}\}$. The pair is simultaneously explained on S if each of the two received words has an explanation there, possibly by different polynomials. A finite slope γ is *support-wise MCA-bad at agreement* a if some support $S \subseteq D$, $|S| \ge a$, explains $r_0 + \gamma r_1$ but does not simultaneously explain the pair. The same support is used in both tests; this is the support-wise convention. A tuple (γ, S, h), where $\deg h < k$ and $h|_S = (r_0 + \gamma r_1)|_S$, is a *witness* to that bad slope when the pair is not simultaneously explained on S, or equivalently is *support-wise nontrivial* there. Finally, $B_C^{\mathrm{MCA}}(a)$ is the maximum, over all received pairs, of the number of distinct bad slopes. It counts slopes, not witnesses or supports.

For a general linear code $C \subseteq \mathbb{F}^D$, the phrase "u is explained on S" means $u|_S \in C|_S$; the MCA definition above then applies verbatim. For a nonempty slope set $\Gamma \subseteq \mathbb{F}$, its ordinary correlated-agreement numerator $B_{C,\Gamma}^{\mathrm{CA}}(a)$ is the same bad-slope maximum restricted to received pairs that are not simultaneously explained on any support of size at least a. These conventions make the MDS statement below well-defined before the detailed decomposition in Section 2.

This is the affine-line definition used by the Proximity Prize [PP26] and by Arnon–Boneh–Fenzi [ABF26, Definition 4.3]; Lemma 2.3 records its exact maximal-domain reformulation. For a real radius δ, a witnessing support has size at least $(1-\delta)n$. Since support size is integral, put

$$r(\delta) = \lfloor \delta n \rfloor, \qquad a(\delta) = n - r(\delta).$$

Thus the error is a right-continuous integer staircase. In particular, a statement about a "largest" real radius needs an endpoint convention; in the exact rows below the safe set is half open and has a supremum but no largest element.

Here is the asymptotic formulation used throughout. If $\varnothing \ne \Gamma \subseteq \mathbb{F}$ is the *challenge set*, meaning the set of slopes allowed in the test, let $B_{C,\Gamma}^{\mathrm{MCA}}(a)$ denote the same maximum with $\gamma \in \Gamma$, and for every $0 \le \delta \le 1$ put

$$e_{\mathrm{MCA},\Gamma}(C,\delta) = \frac{B_{C,\Gamma}^{\mathrm{MCA}}(n - \lfloor \delta n \rfloor)}{|\Gamma|}. \tag{1.1}$$

For the standard full-field challenge we abbreviate $e_{\mathrm{MCA}} = e_{\mathrm{MCA},\mathbb{F}}$. We keep two endpoint conventions separate. The customary *subcapacity* quantities are

$$\begin{aligned}\delta^*_{C,\mathrm{sub},\sup}(\varepsilon) &= \sup\{0 \le \delta < 1 - k/n : e_{\mathrm{MCA}}(C,\delta) \le \varepsilon\},\\ \delta^*_{C,\mathrm{sub},\mathrm{grid}}(\varepsilon) &= \max\left\{\frac{r}{n} : 0 \le r \le n-k-1, e_{\mathrm{MCA}}(C,r/n) \le \varepsilon\right\}.\end{aligned} \tag{SUB}$$

The official Prize interval is $[0,1]$, so we also put

$$\begin{aligned}\delta^*_{C,\mathrm{off},\sup}(\varepsilon) &= \sup\{0 \le \delta \le 1 : e_{\mathrm{MCA}}(C,\delta) \le \varepsilon\},\\ \delta^*_{C,\mathrm{off},\mathrm{grid}}(\varepsilon) &= \max\left\{\frac{r}{n} : 0 \le r \le n, e_{\mathrm{MCA}}(C,r/n) \le \varepsilon\right\}.\end{aligned} \tag{OFF}$$

For a restricted challenge set we insert Γ after C in this notation and replace e_{MCA} by $e_{\mathrm{MCA},\Gamma}$; omission of Γ always means the full-field challenge. For an empty safe set we assign the supremum value zero and leave the grid maximum undefined. Otherwise, if the first unsafe integer radius is r_0 before capacity, the grid value is $(r_0-1)/n$, the real supremum is r_0/n, and the latter endpoint is unsafe. The exceptional all-radius case, in which the official maximum is attained at one, is isolated below. By a *row sequence* we mean a family $C_n = \mathrm{RS}_{\mathbb{F}_n}(D_n,k_n)$ indexed so that $|D_n| = n$; the fields, domains, and dimensions may all vary with n. Every claimed upper bound is uniform over received pairs, whereas a lower construction need only exhibit the stated received pair. For such a sequence, challenge sets Γ_n, and targets $0 \le \varepsilon_n \le 1$, define

$$\begin{aligned}B^*_n &= \lfloor \varepsilon_n |\Gamma_n| \rfloor,\\ a^*_n &= \min\{a \in \{k_n+1,\dots,n\} : B^{\mathrm{MCA}}_{C_n,\Gamma_n}(a) \le B^*_n\},\\ \delta^*_{n,\mathrm{sub},\mathrm{grid}} &= 1 - \frac{a^*_n}{n}, \qquad \delta^*_{n,\mathrm{sub},\sup} = \frac{n - a^*_n + 1}{n},\end{aligned} \tag{1.2}$$

whenever the set defining a^*_n is nonempty. Hereafter we write $\delta^*_{n,\mathrm{sub},\mathrm{grid}} = 1 - a^*_n/n$ for the largest safe subcapacity grid radius; the corresponding continuous supremum is one grid step to the right whenever $a^*_n - 1$ is the first unsafe agreement. Every suppressed symbol $\delta^*_{n,*}$ is relative to the displayed pair (Γ_n,ε_n). Thus the asymptotic RS–MCA problem is to determine $\delta^*_{n,\mathrm{sub},\mathrm{grid}}$, or equivalently the first *safe agreement*: agreement a is safe when its bad-slope numerator is at most B^*_n, and unsafe otherwise. Since the numerator is nonincreasing in a, the resulting safe/unsafe boundary is the *MCA frontier*. This formulation retains the target and challenge size; replacing them by a zero-exponent slogan loses real information when the target is small.

Comparison with the literature and organization. For arbitrary prescribed Reed–Solomon domains, the proved positive theory has a linear bad-slope numerator on every compact subinterval below the Johnson radius $J_\rho = 1 - \sqrt{\rho}$. We use the list-correlated-agreement theorem of BCHKS [BCHKS25, Thm. 4.6], specialized to two received coordinates and reindexed from degree at most $k-1$ to dimension k. The source states this theorem and sketches the generalization following its statement; we use only its displayed $M = 1$ specialization and mark it as an imported input. Haböck's earlier note [Hab25] outlines the quadratic predecessor but is not load-bearing here. Jo [Jo26] proves an all-test-size circuit-incidence envelope for MDS codes and an exact agreement-set shortening theorem for Reed–Solomon codes. The former contains both the support-atlas and paving bounds. Jo also combines shortening with BCHKS to prove polynomial numerators at every fixed number of integer steps beyond Johnson for each fixed reciprocal-integer rate; we use his transfer below to derive the positive-relative-radius exponential-budget exponent Ψ_ρ. Below half the relative minimum distance, the BCHKS two-radius theorem, the exact CA–MCA decomposition, and the tangent construction already imply the $r+1$ staircase. Our contribution in that range is the finite, challenge-restricted, endpoint-exact formulation, not a new constant-rate phenomenon. We retain an independent augmented-word matroid interpretation of the smallest-test specialization of Jo's envelope and derive the positive-relative-radius exponential-budget safe-frontier certificate in the notation relevant to the Prize. A separate parameter-retained interpolation audit is confined to Appendix A; it is conditional on one explicit lift assumption and is not used by any unconditional main theorem.

TABLE 1. Proof status and the role of this paper.

Statement	Status	Attribution and role here
Positive-relative-radius exponential-budget safe-frontier certificate	derived here	Jo's shortening transfer and the BCHKS Johnson estimate give the exponent Ψ_ρ; this is one-sided and not a matching near-capacity threshold.
All-test-size MDS circuit envelope	imported and reproduced	Jo's rejecting-test proof is reproduced here; at $b = k + 1$ it has the augmented-word matroid interpretation recorded in Remark 2.14.
RS staircase below half the minimum distance	prior upper input; exact bookkeeping here	BCHKS plus the CA–MCA conversion supply the asymptotic phenomenon; this paper retains restricted challenges, integer floors, and open endpoints.
Linear MCA bound below Johnson	prior	BCHKS Theorem 4.6 supplies the Johnson-range input used in the exponential-budget certificate; Haböck outlines the earlier quadratic route.
Primitive inverse from a Sidon payment at accessible moment order	proved implication	Given $\log L_N/q_N = o(N)$, BSG and Boolean-cube difference growth eliminate the high-energy Boolean branch; proving the image-normalized Sidon payment and the residual ray compiler remain separate inputs.
Parameter-retained BCIKS–BCHKS calculation	conditional	The four candidate KoalaBear rows depend on the factor-lifting assumption isolated in Appendix A; no unconditional theorem uses it.
Unrestricted smooth/circle near-capacity frontier	open	The repository's profile, analytic-payment, ray, add-back, and reserve interfaces are not assumed in any theorem.

Capacity results for subspace-design codes and random-evaluation ensembles do not settle arbitrary smooth or circle rows [GG25, GGSW26]. Conversely, explicit near-capacity obstructions for structured domains [CS25, KKH26, Kam26] prevent a one-sided profile calculation from being promoted to a complete frontier. On the attack side, Gao, Yang, Xu, and Kan [GYXK26] constructively convert list-decoding counterexamples into MCA lower bounds for related linear codes, with an RS specialization that changes at most one evaluation point; this is complementary to our fixed-code finite-MDS upper bounds. In particular, the identity equation $H_2(\rho + g) - \beta g = \tau$ proposed in the source repository [Cho26] is an attack crossing unless a witness-exhaustive profile atlas, its Fourier or Sidon payments, the residual ray compiler, and a matching safe reserve have all been proved. One exact implication inside that program is established here. On a Boolean support leaf, an image-normalized Sidon payment at an accessible moment order q_N, meaning $\log L_N/q_N = o(N)$, forces any positive-rate max-fiber excess into a high-energy fiber; the Balog–Szemerédi–Gowers theorem and Boolean-cube difference growth then rule it out. This primitive inverse step is proved in Section 10, but its Sidon premise is not a consequence of smooth or circle folding alone and its conclusion is a support-fiber bound, not a distinct-slope bound. We therefore replace that conditional equality by two unconditional statements on their natural scales: an endpoint-exact linear-budget formula and a one-sided positive-relative-radius exponential-budget safe-frontier certificate.

Polynomial-generator preservation and syndrome-space results concern other sampler or code classes [BCGM25, YZ26], while WHIR supplies a protocol-level use of MCA [ACFY25]; none determines the arbitrary-domain finite MDS thresholds studied here.

Here *smooth* means possession of specified complete equal-fiber folding maps, not differentiability. A *circle presentation* means one of the circle evaluation codes defined in Section 9.3 and transferred only through a proved agreement-preserving monomial equivalence. All linear-algebraic, MDS, paving, syndrome, and ordinary Reed–Solomon conclusions are characteristic-independent unless stated otherwise. Circle uniformizations using $x^2 + y^2 = 1$, stereographic projection, or a norm-one torus require odd characteristic. In characteristic two we claim only

the explicitly stated Reed–Solomon results on additive domains in extension fields; "binary" describes $\mathbb{F}_2$-linear domain structure inside $\mathbb{F}_{2^m}$, not an n-point domain contained in the two-element field.

Scope of the unconditional claims. The exact linear-budget theorem, Jo's MDS circuit envelope, the one-sided shortening certificate, the all-radius completion, and every explicitly verified finite row are unconditional subject only to the displayed BCHKS and Jo inputs. The parameter-retained appendix is conditional on its displayed factor-lifting assumption. The primitive inverse theorem is unconditional as an implication from its stated Sidon payment at an accessible moment order, but it does not establish that premise on every smooth or circle leaf. No theorem in this paper asserts the unrestricted constant-rate, positive-relative-radius smooth/circle profile frontier.

AI-use disclosure. OpenAI ChatGPT/Codex, a GPT-5-based system, was used substantively to draft and restructure prose, propose and audit proof steps, check calculations, and edit LaTeX throughout this manuscript. The human author supplied the research program, source manuscripts, repository materials, mathematical choices, and corrections, and retains sole responsibility for correctness and originality. A more detailed use statement accompanies the source release as `AI_USE_v9.2.md`.

Principal original bound. Put $J_\rho = 1 - \sqrt{\rho}$. For $J_\rho \le \delta \le 1 - \rho$, set $\alpha = 1 - \delta$ and

$$\tau_\rho(\delta) = \frac{\rho - \alpha^2}{1 + \rho - 2\alpha}, \qquad \Psi_\rho(\delta) = H_2(\tau_\rho(\delta)) - \alpha H_2\left(\frac{\tau_\rho(\delta)}{\alpha}\right), \tag{SF1}$$

and extend Ψ_ρ by zero on $[0, J_\rho]$. For comparison, put

$$\Phi_\rho(\delta) = H_2(\rho) - (1 - \delta) H_2\left(\frac{\rho}{1 - \delta}\right), \qquad 0 \le \delta \le 1 - \rho. \tag{SF2}$$

Theorem 1.1 (Positive-relative-radius shortening exponent)**.** *Let $C_n = \mathrm{RS}_{\mathbb{F}_n}(D_n, k_n)$ with $k_n/n \to \rho \in (0,1)$, and let $\varnothing \ne \Gamma_n \subseteq \mathbb{F}_n$. Fix $J_\rho < \delta < 1 - \rho$ and put $a_n = n - \lfloor \delta n \rfloor$. Then*

$$\limsup_{n \to \infty} \frac{1}{n} \log_2\left(1 + B^{\mathrm{MCA}}_{C_n, \Gamma_n}(a_n)\right) \le \Psi_\rho(\delta). \tag{SF3}$$

Moreover, $0 < \tau_\rho(\delta) < \rho$ and

$$\Psi_\rho(\delta) < \Phi_\rho(\delta) \tag{SF4}$$

strictly between Johnson and capacity. The exponent extends continuously with $\Psi_\rho(J_\rho) = 0$ and $\Psi_\rho(1 - \rho) = H_2(\rho)$.

The proof is given in Section 6. The optimization leading to Ψ_ρ was carried out after the author learned of Jo's agreement-set shortening theorem [Jo26]: Jo proves the transfer and the fixed-integer-step consequence, while the positive-relative-radius entropy optimization and its Prize-budget formulation are derived here.

Theorem 1.2 (Imported circuit-incidence MDS envelope, Jo [Jo26, Thm. 4.2])**.** *Let $C \subseteq \mathbb{F}^n$ be a linear $[n,k]$ MDS code, let $\varnothing \ne \Gamma \subseteq \mathbb{F}$, and let $k + 1 \le a \le n$. Then*

$$B^{\mathrm{CA}}_{C,\Gamma}(a) \le B^{\mathrm{MCA}}_{C,\Gamma}(a) \le U_{n,k,\Gamma}(a) := \min\left\{ |\Gamma|, \min_{k+1 \le b \le a} \left\lfloor \frac{\binom{n}{b}}{\binom{a-1}{b-1}} \right\rfloor \right\}. \tag{1.3}$$

The full-field estimate is due to Jo; the restricted-challenge form follows by intersecting the bad-slope set with Γ. Its endpoint test sizes $b = k+1$ and $b = a$ recover, respectively, the paving and support-atlas bounds used below. The bound is invariant under coordinate permutations and nonzero diagonal scalings, and hence applies without loss to generalized Reed–Solomon codes and to the circle presentations proved monomially equivalent to them.

Theorem 1.3 (Imported BCHKS Johnson-range estimate, $M = 1$)**.** *Let N, d, E be integers with $0 < d < N$, $0 < E < N$, and $(N - E)^2 > Nd$. Put*

$$\rho_0 = \frac{d}{N}, \qquad \gamma_0 = \frac{E}{N}, \qquad \eta_0 = 1 - \sqrt{\rho_0} - \gamma_0,$$

$$m_B = \max\left\{\left\lceil \frac{\sqrt{\rho_0}}{\eta_0}\right\rceil, 3\right\}, \qquad x_B = m_B + \frac{1}{2},$$

and

$$\mathcal{J}(N,d,E) = \left(\frac{2x_B^5 + 3x_B\gamma_0\rho_0}{3\rho_0^{3/2}}\right)N + \frac{x_B}{\sqrt{\rho_0}}. \tag{B1}$$

For every Reed–Solomon code $C = \mathrm{RS}_{\mathbb{F}}(L, d+1)$ of length N and every nonempty $\Gamma \subseteq \mathbb{F}$,

$$B^{\mathrm{MCA}}_{C,\Gamma}(N-E) \le B^{\mathrm{MCA}}_{C}(N-E) \le \mathcal{J}(N,d,E). \tag{B2}$$

Proof. This is the affine-line $M = 1$ specialization of BCHKS [BCHKS25, Thm. 4.6], in the explicit normalization of Jo [Jo26, Thm. 2.3 and Eqs. (6)–(9)]. There the polynomial degree is at most d, hence the code dimension is $d+1$, and the strict Johnson condition is precisely $(N-E)^2 > Nd$. Restricting the slope set from $\mathbb{F}$ to Γ can only decrease the numerator. □

Theorem 1.4 (Official endpoint and all-radius saturation)**.** *Let $C \subseteq \mathbb{F}^n$ be a linear $[n,k]$ MDS code with $1 \le k < n$, let $\varnothing \ne \Gamma \subseteq \mathbb{F}$, let $0 \le \varepsilon \le 1$, and extend the threshold notation to $0 \le a \le n$. Then*

$$B^{\mathrm{MCA}}_{C,\Gamma}(a) = B^{\mathrm{MCA}}_{C,\Gamma}(k+1) \qquad (0 \le a \le k+1). \tag{OFF1}$$

Put $B^* = \lfloor \varepsilon|\Gamma| \rfloor$, and, when the defining set is nonempty, put

$$a^* = \min\{a \in \{k+1, \dots, n\} : B^{\mathrm{MCA}}_{C,\Gamma}(a) \le B^*\}.$$

If $B^{\mathrm{MCA}}_{C,\Gamma}(k+1) \le B^*$, the complete official safe set is $[0,1]$ and

$$\delta^*_{C,\Gamma,\mathrm{off},\mathrm{grid}}(\varepsilon) = \delta^*_{C,\Gamma,\mathrm{off},\mathrm{sup}}(\varepsilon) = 1. \tag{OFF2}$$

In particular, the sufficient condition

$$B^* \ge U_{n,k,\Gamma}(k+1) = \min\left\{|\Gamma|, \binom{n}{k+1}\right\} \tag{OFF3}$$

implies (OFF2). For the full-field target 2^{-128}, it is enough that $\lfloor 2^{-128}|\mathbb{F}| \rfloor \ge \binom{n}{k+1}$, equivalently $|\mathbb{F}| \ge 2^{128}\binom{n}{k+1}$.

If instead $B^{\mathrm{MCA}}_{C,\Gamma}(k+1) > B^*$ and a^* exists, then necessarily $a^* > k+1$, and the official and subcapacity transitions agree: the largest safe grid radius is $1 - a^*/n$, the complete real safe set is $[0, (n-a^*+1)/n)$, and its right endpoint is unsafe.

Proof. For every support S with $|S| \le k$, extend S to a k-element information set T. The puncturing map $C \to \mathbb{F}^T$ is an isomorphism, so its further projection $C \to \mathbb{F}^S$ is surjective. Hence both received words are explained on S, so no support-wise nontrivial witness has size at most k. This proves (OFF1). The first conclusion follows immediately. At $a = k+1$, Theorem 1.2 gives $B^{\mathrm{MCA}}_{C,\Gamma}(k+1) \le U_{n,k,\Gamma}(k+1)$, proving the sufficient condition. In the nonsaturated case, monotonicity and the identity $a(\delta) = n - \lfloor \delta n \rfloor$ give the stated grid point and half-open real interval. □

Corollary 1.5 (Four exact-plateau saturated rows)**.** *Let*

$$p = (2^{128} - 255)2^{128} + 1$$

and let $D \le \mathbb{F}_p^\times$ be the subgroup of order 128. For

$$C_\rho = \mathrm{RS}_{\mathbb{F}_p}(D, 128\rho), \qquad \rho \in \{1/2, 1/4, 1/8, 1/16\},$$

the complete official safe set at target 2^{-128} is $[0,1]$, so both official grid and supremum endpoints equal 1 for all four rates. Moreover, with $M_\rho = \binom{128}{128\rho+1}$,

$$B^{\mathrm{MCA}}_{C_\rho}(a) = M_\rho \qquad (0 \le a \le 128\rho + 1).$$

The field has fewer than 2^{256} elements and the evaluation domain has power-of-two order.

Proof. The Proth certificate in Proposition 6.5 proves that p is prime; since $128 \mid p-1$, the subgroup exists. The target budget is $2^{128}-255$, while

k	$\binom{128}{k+1}$
64	23582666872052266206656578733667004800
32	429907468073390739393985381161600
16	614965786737727286400
8	19062702032000

and every entry is smaller than that budget. Apply Theorem 1.4. The largest M_ρ is less than 2^{125}, while $p > 2^{255}$, so $p > \binom{M_\rho}{2}$ in every row. The exact plateau follows from Theorem 7.2. □

Remark 1.6 (The field-size quantifier). For fixed n and k, arbitrary scalar-field enlargement eventually forces (OFF3); the resulting endpoint 1 is a budget-saturation statement, not a near-capacity decoding theorem. If the phrase *sufficiently large* is intended to retain a prescribed field ceiling or a growing-length relation between q and n, the supplied row must instead be evaluated by the finite or asymptotic certificates below. In particular, selected saturated rows do not solve the uniform given-C problem.

The endpoint test sizes in (1.3) recover the support-atlas and paving bounds, while intermediate sizes can improve both at finite length. For Reed–Solomon rows and their monomially equivalent GRS or circle presentations—indeed, after Proposition 4.2, for every linear MDS code—put $L_{n,\Gamma}(a) = \min\{|\Gamma|, n-a+1\}$, the tangent construction and Theorem 1.2 give the finite security interval

$$\log_2 \frac{|\Gamma|}{U_{n,k,\Gamma}(a)} \le -\log_2 \frac{B^{\mathrm{MCA}}_{C,\Gamma}(a)}{|\Gamma|} \le \log_2 \frac{|\Gamma|}{L_{n,\Gamma}(a)}. \tag{1.4}$$

Thus replacing one rigorous upper numerator U_0 by a smaller U_1 earns exactly $\log_2(U_0/U_1)$ proved bits. The remaining interval has width at most $\log_2(U_1/L_{n,\Gamma}(a))$; this is the auditable pay-per-bit quantity, independent of any eventual prize-allocation rule. For rigorous upper numerators $U_0 \ge U_1 > 0$ and attack lower numerators $0 < L_0 \le L_1$, the exact uncertainty contraction is

$$\log_2 \frac{U_0}{L_0} - \log_2 \frac{U_1}{L_1} = \log_2 \frac{U_0}{U_1} + \log_2 \frac{L_1}{L_0}. \tag{SG}$$

This identity keeps proof-side and falsification-side gains separate and prevents the same closed bit from being counted twice.

Theorem 1.7 (Exact linear-budget threshold)**.** *Let $C_n = \mathrm{RS}_{\mathbb{F}_n}(D_n, k_n)$, let $k_n/n \to \rho \in (0,1)$, let $\varnothing \ne \Gamma_n \subseteq \mathbb{F}_n$, let $0 \le \varepsilon_n \le 1$, and put $B^*_n = \lfloor \varepsilon_n |\Gamma_n| \rfloor$. If $B^*_n = 0$, no agreement is safe; if $B^*_n = |\Gamma_n|$, then $a^*_n = k_n + 1$, $\delta^*_{n,\mathrm{sub,grid}} = 1 - (k_n+1)/n$, $\delta^*_{n,\mathrm{sub,sup}} = 1 - k_n/n$, and $\delta^*_{C_n,\Gamma_n,\mathrm{off,grid}} = \delta^*_{C_n,\Gamma_n,\mathrm{off,sup}} = 1$. Otherwise suppose that, for some $\eta > 0$, eventually*

$$1 \le B^*_n < |\Gamma_n|, \qquad B^*_n - 1 \le \left(\frac{1-\rho}{2} - \eta\right) n. \tag{1.5}$$

Then

$$a^*_n = n - B^*_n + 1, \qquad \delta^*_{n,\mathrm{sub,grid}} = \frac{B^*_n - 1}{n}, \tag{1.6}$$

and, under the continuous convention $r(\delta) = \lfloor \delta n \rfloor$, the complete official safe set is $[0, B^*_n/n)$; its endpoint is unsafe. Thus the official and subcapacity grid values agree and the official supremum is B^*_n/n. Under these otherwise hypotheses, if $B^*_n/n \to b < (1-\rho)/2$, then $\delta^*_{n,\mathrm{sub,grid}} \to b$ and $\delta^*_{C_n,\Gamma_n,\mathrm{off,sup}}(\varepsilon_n) \to b$.

Proof. The zero-budget case follows from the tangent numerator one at agreement n; the full-budget case follows from Theorem 1.4. In the remaining case, the uniform below-half staircase of Corollary 5.4, at radius $B^*_n - 1$, gives numerator B^*_n. At the next radius B^*_n, the tangent construction gives $\min\{|\Gamma_n|, B^*_n+1\} = B^*_n+1$. The range assumption also implies $B^*_n \le n-k_n-1$

eventually, so both radii belong to the admissible grid. Monotonicity proves the first-safe agreement and Lemma B.1 gives the half-open real interval. □

For the prize-facing fixed-bit scaling there is no entropy inversion to perform. If the challenge is the whole field, the target is $2^{-\lambda}$, $q_n/n \to \kappa$, and $b = \kappa 2^{-\lambda} < (1-\rho)/2$, then, whenever the endpoint cases are excluded,

$$a_n^* = n - \left\lfloor q_n 2^{-\lambda} \right\rfloor + 1, \qquad \delta^*_{n,\mathrm{sub,grid}} \longrightarrow b, \tag{1.7}$$

This is an exact asymptotic formula for $q_n = \Theta(n)$. If instead $q_n = 2^{o(n)}$ and the fixed-bit budget is $o(n)$, the exact grid threshold still has the form (1.6) and therefore tends to radius zero.

The shortening exponent of Theorem 1.1 converts directly into a target-budget certificate. For $\theta > 0$, define

$$D^{\mathrm{sub}}_{\rho,\mathrm{short}}(\theta) = \sup\{0 \le \delta < 1-\rho : \Psi_\rho(\delta) < \theta\}, \quad D^{\mathrm{off}}_{\rho,\mathrm{short}}(\theta) = \begin{cases} D^{\mathrm{sub}}_{\rho,\mathrm{short}}(\theta), & \theta \le H_2(\rho), \\ 1, & \theta > H_2(\rho). \end{cases} \tag{1.10}$$

In particular, every $\theta > 0$ moves the certified subcapacity radius strictly beyond Johnson.

Theorem 1.8 (Positive-relative-radius exponential-budget safe-frontier certificate)**.** *Let $C_n = \mathrm{RS}_{\mathbb{F}_n}(D_n, k_n)$, let $k_n/n \to \rho \in (0,1)$, let $\varnothing \neq \Gamma_n \subseteq \mathbb{F}_n$, let $0 \le \varepsilon_n \le 1$, and put $B_n^* = \lfloor \varepsilon_n |\Gamma_n| \rfloor$. Suppose*

$$\frac{1}{n}\log_2(1 + B_n^*) \longrightarrow \theta > 0. \tag{1.11}$$

Then a_n^* is defined for all sufficiently large n, and

$$\liminf_{n\to\infty} \delta^*_{n,\mathrm{sub,grid}} \ge D^{\mathrm{sub}}_{\rho,\mathrm{short}}(\theta), \qquad \liminf_{n\to\infty} \delta^*_{C_n,\Gamma_n,\mathrm{off,grid}}(\varepsilon_n) \ge D^{\mathrm{off}}_{\rho,\mathrm{short}}(\theta). \tag{1.12}$$

If $\theta \ge H_2(\rho)$, then $\delta^*_{n,\mathrm{sub,grid}} \to 1-\rho$. If the inequality is strict, then $a_n^* = k_n + 1$ eventually and the official safe set is $[0,1]$, so $\delta^*_{C_n,\Gamma_n,\mathrm{off,grid}}(\varepsilon_n) = \delta^*_{C_n,\Gamma_n,\mathrm{off,sup}}(\varepsilon_n) = 1$. At equality the subcapacity limiting radius is still capacity, but finite constants decide whether the official endpoint saturates at one. For $0 < \theta < H_2(\rho)$, (1.12) is a one-sided certificate, not an exact threshold.

Proof. The positive budget exponent gives $B_n^* \ge 1$ eventually, while Theorem 1.2 gives $B^{\mathrm{MCA}}_{C_n,\Gamma_n}(n) \le 1$; hence a_n^* exists. For every fixed $0 < \delta < J_\rho$, apply Theorem 1.3 with

$$N = n, \qquad d = k_n - 1, \qquad E = \lfloor \delta n \rfloor .$$

Here $d/N \to \rho$ and $1 - \sqrt{d/N} - E/N \to J_\rho - \delta > 0$, so the integer m_B in (B1) is bounded and $\mathcal{J}(N,d,E) = O_{\rho,\delta}(n)$. At $\delta = 0$, the same subexponential conclusion follows already from Theorem 1.2. Restricting from $\mathbb{F}_n$ to Γ_n can only decrease the numerator, and every such numerator is eventually below a budget of positive exponent. Beyond Johnson, Theorem 1.1 shows that every fixed radius with $\Psi_\rho(\delta) < \theta$ is safe. At Johnson itself, choose a nearby larger radius with the same strict inequality and use monotonicity. Taking suprema gives (1.12); the official bound follows because it contains the subcapacity safe grid, with the strict saturation branch supplied by Theorem 1.4. Every subcapacity grid radius is at most $1-(k_n+1)/n$, proving the matching subcapacity limit when $\theta \ge H_2(\rho)$. If $\theta > H_2(\rho)$, the envelope at $a = k_n + 1$ has exponent at most $H_2(\rho)$, so the first agreement is safe; then Theorem 1.4 gives the official endpoint one. □

Remark 1.9 (Imported Johnson-range input). The load-bearing Johnson-range input is BCHKS Theorem 4.6 in the standalone (N,d,E) normalization of Theorem 1.3. The source states the displayed theorem and sketches the list-correlated-agreement generalization. Thus the below-Johnson part of Theorem 1.8 and Corollary 1.10 depend on that published statement. Jo's shortening transfer and rejecting-test proof of the all-test-size MDS envelope are reproduced in Theorem 6.2 and Lemma 2.13.

Corollary 1.10 (Security and soundness exponents)**.** *Let $C_n = \mathrm{RS}_{\mathbb{F}_n}(D_n, k_n)$ with $k_n/n \to \rho \in (0,1)$, put $q_n = |\mathbb{F}_n|$, assume the full-field challenge, and suppose $\log_2 q_n/n \to \phi$. For a fixed radius δ, let $s_n(\delta) = -\log_2 e_{\mathrm{MCA}}(C_n, \delta)$. Then*

$$\frac{s_n(\delta)}{n} \longrightarrow \phi \quad (0 \le \delta < J_\rho), \qquad \liminf_{n\to\infty} \frac{s_n(\delta)}{n} \ge \big(\phi - \Psi_\rho(\delta)\big)_+ \quad (J_\rho \le \delta < 1-\rho). \tag{1.15}$$

The tangent attack gives $\limsup s_n(\delta)/n \le \phi$, so the remaining soundness-gap rate beyond Johnson is at most $\min\{\phi, \Psi_\rho(\delta)\}$. Relative to the smallest-test MDS exponent, shortening closes

$$\min\{\phi, \Phi_\rho(\delta)\} - \min\{\phi, \Psi_\rho(\delta)\}$$

bits per coordinate; this equals $\Phi_\rho(\delta) - \Psi_\rho(\delta)$ only when $\phi \ge \Phi_\rho(\delta)$. For a target $2^{-\lambda_n}$ with $\lambda_n \ge 0$, $\lambda_n/n \to \lambda$, and $\phi - \lambda > 0$, the certified limiting frontier is $D^{\mathrm{sub}}_{\rho,\mathrm{short}}(\phi - \lambda)$: every fixed radius strictly below it is eventually safe. If $\phi - \lambda > H_2(\rho)$, the official safe set is $[0,1]$ eventually. When $\phi > 0$, a fixed 128-bit target has $\lambda = 0$; for $q_n = \Theta(n)$, the exact linear-budget formula (1.7) is the relevant statement instead.

Proof. Below Johnson, the BCHKS linear upper numerator and the tangent polynomial lower numerator have zero base-two exponent, which gives the limit. For $\delta > J_\rho$, combine the field cap with Theorem 1.1. At $\delta = J_\rho$, monotonicity gives, for every fixed $\delta' > J_\rho$ and all large n,

$$B^{\mathrm{MCA}}_{C_n}(n - \lfloor J_\rho n \rfloor) \le B^{\mathrm{MCA}}_{C_n}(n - \lfloor \delta' n \rfloor).$$

Apply that theorem and let $\delta' \downarrow J_\rho$. The target statement follows because $\log_2(1 + \lfloor q_n 2^{-\lambda_n} \rfloor)/n \to \phi - \lambda$. □

The universal MDS envelope gives 128-bit multiplicative and circle rows beyond Johnson at length 512; see Propositions 6.5 and 9.23. The exact Proth rows of Section 7 are instances of the below-half formula and are included for endpoint and certificate verification, not as a new literature consequence. Fixed-length congruence classes give positive prime density, but that density tends to zero along growing dyadic prime-field lengths. Circle claims below always pass through an explicit agreement-preserving diagonal equivalence; no statement is made for an unidentified bivariate presentation.

2. MCA witnesses and exact support reduction

Throughout this section $\mathbb{F}$ is a finite field, $D \subseteq \mathbb{F}$ has size n, and $C = \mathrm{RS}_{\mathbb{F}}(D, k)$. Write $\mathbb{F}[X]_{<k} = \{f \in \mathbb{F}[X] : \deg f < k\}$, and let $\mathrm{ev}_S : \mathbb{F}[X]_{<k} \to \mathbb{F}^S$ be restriction of the evaluation map.

Definition 2.1 (Explanation and support-wise nontriviality)**.** *For $u \in \mathbb{F}^D$ and $S \subseteq D$, an* explanation of u on S *is a polynomial $f \in \mathbb{F}[X]_{<k}$ with $\mathrm{ev}_S(f) = u|_S$. We write $\mathrm{Expl}_C(u; S)$ when such an f exists. A pair $r = (r_0, r_1)$ is* simultaneously explained on S *when $\mathrm{Expl}_C(r_0; S)$ and $\mathrm{Expl}_C(r_1; S)$ both hold, with no requirement that the two explaining polynomials coincide. It is* support-wise nontrivial on S *when it is not simultaneously explained there. We denote this last condition by $\mathrm{NT}_C(r; S)$.*

Thus support-wise nontriviality is a condition on the same support that explains the affine combination. It does not mean that both coordinates are individually unexplained; failure of either explanation is enough.

Definition 2.2 (MCA-bad slope)**.** *Let $r = (r_0, r_1) \in (\mathbb{F}^D)^2$ and let a be an agreement threshold. A finite slope $\gamma \in \mathbb{F}$ is support-wise MCA-bad at agreement a if there are $S \subseteq D$, $|S| \ge a$, and $h \in \mathbb{F}[X]_{<k}$ such that $\mathrm{ev}_S(h) = (r_0 + \gamma r_1)|_S$ and $\mathrm{NT}_C(r; S)$ holds. Define*

$$B^{\mathrm{MCA}}_C(a) = \max_{r \in (\mathbb{F}^D)^2} |\{\gamma \in \mathbb{F} : \gamma \text{ is MCA-bad for } r \text{ at } a\}|.$$

Lemma 2.3 (Support-wise and maximal-domain formulations)**.** *For a received pair $r = (r_0, r_1)$ and an integer a, let*

$$\mathcal{E}_{C,a}(r) = \{S \subseteq D : |S| \geq a,\ r_0|_S \in C|_S,\ r_1|_S \in C|_S\},$$

and let $\mathcal{M}_{C,a}(r)$ be its inclusion-maximal members. Put $\mathrm{Agr}(u, c) = \{x \in D : u(x) = c(x)\}$. A slope γ is MCA-bad for r at agreement a if and only if there is $c \in C$ such that, for $A = \mathrm{Agr}(r_0 + \gamma r_1, c)$,

$$|A| \geq a \quad \text{and} \quad A \nsubseteq M \quad \text{for every } M \in \mathcal{M}_{C,a}(r). \tag{MDB}$$

If $a > n - d_{\min}(C)$, the last condition is equivalently $A \notin \mathcal{M}_{C,a}(r)$. Consequently, for $0 < \delta < 1$ and $a(\delta) = \lceil (1-\delta)n \rceil = n - \lfloor \delta n \rfloor$,

$$\varepsilon_{\mathrm{mca}}(C, \delta) = \frac{1}{|\mathbb{F}|} \max_r \#\{\gamma \in \mathbb{F} : \gamma \text{ satisfies (MDB)}\},$$

which is the quantity in [ABF26, Definition 4.3]*; our values at $\delta = 0, 1$ are its natural endpoint extensions.*

Proof. For every $S \subseteq D$ with $|S| \geq a$, simultaneous explanation on S is equivalent, by finiteness and restriction, to containment in some member of $\mathcal{M}_{C,a}(r)$. If S witnesses that γ is bad and c explains the line point on S, then $S \subseteq A = \mathrm{Agr}(r_0 + \gamma r_1, c)$. Were A contained in a maximal common domain, restriction would simultaneously explain the pair on S, a contradiction. Conversely, (MDB) says that the pair is not simultaneously explained on A, so $S = A$ is a support-wise witness.

Assume now $a > n - d_{\min}(C)$ and that the pair is simultaneously explained on A. The two explaining codewords and c are unique on A. Extending A to a maximal common domain therefore uses the same two codewords, whose affine combination is c; the extension is contained in A and hence equals it. Thus A itself is maximal. The probability identity follows because the official sampler chooses γ uniformly from $\mathbb{F}$. □

The official MCA maximum mixes two logically different sources: received pairs that are globally far from every codeword pair, and pairs that are close to one codeword pair but fail on a slope-dependent support. The following definitions and theorem separate them exactly.

Definition 2.4 (CA numerator and sparse mutual numerator)**.** *Assume $0 \leq a \leq n$ and put $t = n - a$. A received pair (r_0, r_1) is* column-far at agreement a *if it is not simultaneously explained on any support of size at least a. Let $B_C^{\mathrm{CA}}(a)$ be the maximum, over column-far pairs, of the number of slopes γ for which $r_0 + \gamma r_1$ is explained on some support of size at least a. If there is no column-far pair, this maximum is defined to be zero. Define the* sparse mutual numerator

$$\sigma_C(a) = \max\{\#\{\gamma : \gamma \text{ is MCA-bad for } (e_0, e_1)\} : |\mathrm{supp}(e_0) \cup \mathrm{supp}(e_1)| \leq t\}. \tag{SP1}$$

Thus $B_C^{\mathrm{CA}}(a)/|\mathbb{F}|$ is the ordinary affine-line correlated agreement error on the integer grid, while $\sigma_C(a)$ records the purely mutual obstruction after translation by a nearby codeword pair.

Theorem 2.5 (Exact sparsification of the mutual layer)**.** *For every linear code $C \subseteq \mathbb{F}^D$ and every agreement threshold $0 \leq a \leq n$,*

$$B_C^{\mathrm{MCA}}(a) = \max\{B_C^{\mathrm{CA}}(a), \sigma_C(a)\}. \tag{SP2}$$

Equivalently, after division by $|\mathbb{F}|$, the official MCA error is the maximum of the CA error and the sparse mutual error.

Proof. Both quantities on the right are obtained by maximizing the MCA-bad set over restricted classes of pairs, so the left side is at least their maximum. Conversely fix a received pair. If it is column-far, every support that explains a line point necessarily fails to explain the pair; its MCA-bad set is therefore its CA-bad set and has size at most $B_C^{\mathrm{CA}}(a)$.

Otherwise choose codewords $c_0, c_1 \in C$ and a common support of size at least a, and put $e_i = r_i - c_i$. The union of the supports of e_0, e_1 has size at most $n - a = t$. At each slope, translation sends an explaining codeword h to $h - c_0 - \gamma c_1$, and is a bijection between witness

supports for (r_0, r_1) and for (e_0, e_1). It also preserves simultaneous explanation on that support. Hence the two pairs have exactly the same MCA-bad slopes, and their number is at most $\sigma_C(a)$. □

For a nonempty challenge set $\Gamma \subseteq \mathbb{F}$, write $B^{\mathrm{CA}}_{C,\Gamma}(a)$ and $\sigma_{C,\Gamma}(a)$ for the same maxima with the bad slopes restricted to Γ. The proof just given is unchanged after intersecting every bad-slope set with Γ, and therefore gives the exact challenge-restricted identity

$$B^{\mathrm{MCA}}_{C,\Gamma}(a) = \max\{B^{\mathrm{CA}}_{C,\Gamma}(a), \sigma_{C,\Gamma}(a)\}. \tag{SP3}$$

Theorem 2.6 (Exact mutual layer below half the distance)**.** *Let $C \subseteq \mathbb{F}^D$ be a linear code of minimum distance $d_{\min}$, let $\varnothing \neq \Gamma \subseteq \mathbb{F}$, and let $0 \le r \le n$ satisfy $2r \le d_{\min} - 1$. Then*

$$\sigma_{C,\Gamma}(n-r) = \min\{|\Gamma|, r\}, \qquad B^{\mathrm{MCA}}_{C,\Gamma}(n-r) = \max\left\{B^{\mathrm{CA}}_{C,\Gamma}(n-r), \min\{|\Gamma|, r\}\right\}. \tag{HD1}$$

For a Reed–Solomon row the hypothesis is $2r \le n - k$.

Proof. Fix a sparse pair (e_0, e_1) whose support union E has size at most r, and let γ be bad on a support S of size at least $n - r$, explained there by a codeword $c \in C$. On $S \setminus E$ the sparse line point is zero, so $\mathrm{supp}(c) \subseteq E \cup (D \setminus S)$ has size at most $2r < d_{\min}$. Thus $c = 0$. If $S \cap E$ contained no nonzero pair coordinate, the zero codeword pair would simultaneously explain (e_0, e_1) on S, contrary to badness. Hence some $x \in S \cap E$ satisfies

$$(e_0(x), e_1(x)) \neq (0, 0), \qquad e_0(x) + \gamma e_1(x) = 0.$$

Necessarily $e_1(x) \neq 0$, and this coordinate determines the single slope $-e_0(x)/e_1(x)$. Therefore every sparse pair has at most $\min\{|\Gamma|, |E|\} \le \min\{|\Gamma|, r\}$ bad slopes.

For the reverse bound put $t = \min\{|\Gamma|, r\}$, choose distinct $x_1, \dots, x_t \in D$ and distinct $\gamma_1, \dots, \gamma_t \in \Gamma$, and set $e_1(x_i) = 1$, $e_0(x_i) = -\gamma_i$, with both words zero elsewhere. For $S_i = D \setminus \{x_j : j \neq i\}$, the line point of slope γ_i is zero on S_i, and $|S_i| = n - t + 1 \ge n - r$. The restriction of e_1 to S_i cannot be explained by a codeword: such a codeword would be nonzero at x_i and zero on $n - t$ coordinates, hence would have weight at most $t < d_{\min}$. Thus every γ_i is MCA-bad and equality holds. The second formula is (SP3); an RS code has $d_{\min} = n - k + 1$. □

Lemma 2.7 (Pair-interleaved Johnson bound)**.** *Let $C = \mathrm{RS}_{\mathbb{F}}(D, k)$, let $1 \le A \le n$, and suppose $A^2 > (k-1)n$. For every received pair $y \in (\mathbb{F}^D)^2$, the number of codeword pairs $(p_0, p_1) \in C^2$ simultaneously agreeing with y on at least A coordinates is at most*

$$J_2(A) = \left\lfloor \frac{n(A-k+1)}{A^2 - (k-1)n} \right\rfloor. \tag{HD2}$$

Proof. Let the list have size L, and for every member retain exactly A agreement coordinates. If d_x is the number of retained sets containing x, then $\sum_x d_x = LA$. Two distinct codeword pairs can agree coordinatewise on at most $k - 1$ points, since at least one nonzero difference polynomial has degree less than k. Therefore

$$\frac{L^2A^2/n - LA}{2} \le \sum_{x \in D} \binom{d_x}{2} \le \binom{L}{2}(k-1),$$

where the first inequality is Cauchy–Schwarz. Rearrangement gives $L(A^2 - (k-1)n) \le n(A - k + 1)$, and integrality gives (HD2). □

Theorem 2.8 (Self-contained half-Johnson CA and MCA bound)**.** *Let $C = \mathrm{RS}_{\mathbb{F}}(D, k)$, let $\varnothing \neq \Gamma \subseteq \mathbb{F}$, and let r be an integer with $0 \le r \le n$. Put $A = n - 2r$. If $A > 0$ and $A^2 > (k-1)n$, then*

$$B^{\mathrm{CA}}_{C,\Gamma}(n-r) \le \min\{|\Gamma|, 1 + (r+1)J_2(A)\}. \tag{HD3}$$

If in addition $2r \le n - k$, the same upper bound holds for $B^{\mathrm{MCA}}_{C,\Gamma}(n-r)$.

Proof. Fix a column-far received pair and let $Z \subseteq \Gamma$ be its CA-bad set. The assertion is immediate when $|Z| \le 1$, so fix an anchor $\gamma_0 \in Z$. For every $\gamma \in Z \setminus \{\gamma_0\}$, choose explaining codewords c_{γ_0}, c_γ and agreement sets of size at least $n - r$, and define

$$P_1 = \frac{c_{\gamma_0} - c_\gamma}{\gamma_0 - \gamma}, \qquad P_0 = c_{\gamma_0} - \gamma_0 P_1.$$

On the intersection of the two agreement sets, of size at least $n - 2r = A$, the codeword pair (P_0, P_1) agrees with the received pair. Hence all such pairs lie in a list of size at most $J_2(A)$ by Lemma 2.7.

Fix one pair in that list and let E be its column-error set, of size e. Column-farness gives $e \ge r+1$. If a slope γ rides this pair, meaning its explaining codeword is $P_0 + \gamma P_1$, then the nonzero affine coordinate errors vanish at at least $e - r$ coordinates of E. Each coordinate vanishes for at most one finite slope, so double counting gives at most $e/(e-r) \le r+1$ riders. Summing over the pair list and adding the anchor proves (HD3). Under $2r \le n-k$, Theorem 2.6 says that the other branch of MCA has only $\min\{|\Gamma|, r\}$ slopes. Since $n(A-k+1) - (A^2 - (k-1)n) = A(n-A) \ge 0$, one has $J_2(A) \ge 1$, so this tangent term is no larger than the displayed CA bound. □

Remark 2.9 (The exact remaining prize certificate). At a proposed safe agreement a_+, an unrestricted finite-row proof must establish both

$$B_C^{\mathrm{CA}}(a_+) \le B^*, \qquad \sigma_C(a_+) \le B^*, \qquad B^* = \left\lfloor 2^{-128} |\mathbb{F}| \right\rfloor.$$

An unsafe construction at $a_+ - 1$ does not prove either upper bound. This identity is the shortest exact description of what remains for the near-capacity rows not covered by Corollary 5.6; the older quadratic compiler Corollary 4.12 is a self-contained subrange.

The finite-slope restriction costs at most the point at infinity, which may also be treated by exchanging the two coordinates. The support-wise condition is essential: explanation and nontriviality are tested on the same set.

Definition 2.10 (Exact-agreement witness incidence)**.** *Assume $k+1 \le a \le n$. For a received line $r = (r_0, r_1)$, its exact-a witness incidence is*

$$\mathcal{W}_a(r) = \left\{ (\gamma, S, h) : \begin{array}{c} \gamma \in \mathbb{F}, \quad S \in \binom{D}{a}, \quad h \in \mathbb{F}[X]_{<k}, \\ \mathrm{ev}_S(h) = (r_0 + \gamma r_1)|_S, \quad \mathrm{NT}_C(r; S) \end{array} \right\}. \tag{2.1}$$

Its slope projection is $\pi_\gamma(\gamma, S, h) = \gamma$, and its exact bad-slope set is $Z_a(r) = \pi_\gamma(\mathcal{W}_a(r)) \subseteq \mathbb{F}$.

Lemma 2.11 (Reduction to exact agreement)**.** *For $a \ge k+1$, the set $Z_a(r)$ is precisely the set of slopes that are MCA-bad at threshold a in Definition 2.2.*

Proof. An exact-a witness is a threshold-a witness. Conversely, let (γ, S, h) witness badness with $|S| \ge a$. At least one of $r_0|_S, r_1|_S$ is outside $\mathrm{ev}_S(\mathbb{F}[X]_{<k})$; call it $u|_S$. Interpolate u on any fixed k-subset $K \subseteq S$. If the resulting polynomial agreed with u at every point of $S \setminus K$, it would explain u on S, a contradiction. Hence some $x \in S \setminus K$ makes u unexplained on $K \cup \{x\}$. Extend this set to an a-subset $S' \subseteq S$. Non-explanation persists on supersets, while the same h explains $r_0 + \gamma r_1$ on S'. □

Proposition 2.12 (Exact support-atlas upper bound)**.** *For every Reed–Solomon row, every nonempty challenge set $\Gamma \subseteq \mathbb{F}$, and every $k+1 \le a \le n$,*

$$B_{C,\Gamma}^{\mathrm{MCA}}(a) \le \min\left\{ |\Gamma|, \binom{n}{a} \right\}. \tag{2.2}$$

This is a literal finite-row bound: it has no polynomial or subexponential loss.

Proof. By Lemma 2.11, choose one noncommon exact-a support for every bad slope. One noncommon support carries at most one slope, since two distinct explained slopes would separately explain the two received words on that support. The chosen-support map is therefore injective. Intersect with the challenge set. □

For an MDS code on a coordinate set S, a b-coordinate *local test* is a subset $T \in \binom{S}{b}$; it accepts a word w when $w|_T \in C|_T$ and rejects it otherwise.

Lemma 2.13 (Rejecting-test abundance, Jo [Jo26, Lem. 4.1])**.** *Let $C \subseteq \mathbb{F}^S$ be an $[m,k]$ MDS code, let $w \notin C$, and let $k+1 \le b \le m$. Then at least*

$$\binom{m-1}{b-1} \tag{2.3}$$

sets $T \in \binom{S}{b}$ satisfy $w|_T \notin C|_T$. This number is best possible.

Proof. We reproduce Jo's induction. Use strong induction on m, simultaneously in k. For $k = 0$, choose x with $w(x) \neq 0$; every b-set containing x rejects. For $k \ge 1$, fix $x \in S$ and put $S' = S \setminus \{x\}$. If $w|_{S'} \in C|_{S'}$, choose $c \in C$ agreeing with w on S'. Since $w \notin C$, the words differ at x, and every b-set containing x rejects by MDS uniqueness on its other $b-1 \ge k$ coordinates.

Otherwise, if $b < m$, induction gives $\binom{m-2}{b-1}$ rejecting b-sets avoiding x; for $b = m$ this contribution is zero. Choose $c \in C$ with $c(x) = w(x)$, put $v = w - c$, and shorten at x:

$$C_x = \{d|_{S'} : d \in C,\ d(x) = 0\}.$$

This is $[m-1, k-1]$ MDS and $v|_{S'} \notin C_x$. A set $T \cup \{x\}$ accepts w precisely when T accepts v for C_x, so induction supplies $\binom{m-2}{b-2}$ further rejecting sets. Pascal's identity proves the bound. Sharpness is attained by $w = c + e_x$, where e_x is nonzero only at x: exactly the sets containing x reject. □

Proof of Theorem 1.2. Following Jo, the CA inequality is immediate. Fix a received pair, let $Z \subseteq \Gamma$ be its bad-slope set, and, for every $\gamma \in Z$, choose a witnessing support S_γ of size at least a. The direction word is outside $C|_{S_\gamma}$, whereas the affine combination at γ is inside it. For fixed $k+1 \le b \le a$, Lemma 2.13 gives at least $\binom{a-1}{b-1}$ subsets $T \in \binom{S_\gamma}{b}$ on which the direction rejects while the affine combination accepts. One fixed T cannot be owned by two distinct slopes, since subtraction would make the direction accept on T. Hence

$$|Z|\binom{a-1}{b-1} \le \binom{n}{b}. \tag{2.4}$$

Take the floor, minimize over b, and intersect with Γ. Coordinate permutations and invertible diagonal scalings preserve every punctured membership test. □

Remark 2.14 (Paving specialization). At $b = k+1$, the rejecting sets are the bases of the rank-$(k+1)$ augmented-word paving matroid, giving the independent matroid interpretation used in the earlier Chojecki manuscript [Cho26]. Jo's strengthening is to use every test size. The local lower bound is sharp for a one-coordinate perturbation, so further universal progress must exploit compatibility among different slopes.

3. Exact syndrome geometry

The *syndrome* of a word is the result of applying a parity-check matrix; it vanishes exactly on codewords. Thus syndrome space forgets the part of a received word already lying in the code and retains only its error class. This description separates possible error supports from distinct line parameters. Put $R = n - k$, the redundancy (and syndrome-space dimension), and choose the generalized Vandermonde parity-check matrix H with columns

$$h_x = \lambda_x (1, x, \dots, x^{R-1})^{\mathsf{T}}, \qquad \lambda_x = \left(\prod_{y \in D \setminus \{x\}} (x - y) \right)^{-1}. \tag{3.1}$$

Lagrange interpolation gives $\sum_{x \in D} \lambda_x x^j = 0$ for $0 \le j \le n-2$, so this matrix annihilates every degree-less-than-k evaluation. Every set of at most R columns is independent. For $E \subseteq D$, write $V_E = \mathrm{Span}\{h_x : x \in E\}$; it is exactly the set of syndromes of words supported on E. We write $s(u) = Hu^{\mathsf{T}}$ for the syndrome of u.

Proposition 3.1 (Syndrome-line normal form)**.** *Let $S = D \setminus E$. A word u is explained on S if and only if $s(u) \in V_E$. Consequently γ is MCA-bad on S for a pair (u_0, u_1) if and only if*

$$s(u_0) + \gamma s(u_1) \in V_E, \qquad \{s(u_0), s(u_1)\} \not\subseteq V_E. \tag{3.2}$$

For fixed E, there is at most one such finite slope.

Proof. If $u = c + e$, where $c \in C$ and e is supported on E, then $s(u) = s(e) \in V_E$. Conversely, if $s(u) = \sum_{x \in E} e_x h_x$, subtract the word supported on E with values e_x; the result has zero syndrome and belongs to C. Apply this equivalence to the line point and to both coordinates. Passing (3.2) to $\mathbb{F}^R/V_E$ shows that two distinct solutions would put both projected syndromes at zero, contradicting badness. This proof uses only that H is a full-rank parity-check matrix with kernel C; hence the normal form applies verbatim to every linear code when H is replaced by any such parity check. □

Since $\lambda_x \neq 0$, one has $[h_x] = [1 : x : \cdots : x^{R-1}]$; these are the points of the rational normal curve in projective syndrome space, while the weights λ_x merely choose vector representatives. A *secant span* is the linear span V_E of a selected set of those points. We call the meeting of a syndrome line with V_E *transverse* when the entire line is not contained in V_E, equivalently when its two generators are not both in V_E. The next result is a *compiler* in the sense that it translates the bad-slope problem, exactly and in both directions, into this line–secant incidence problem. A *ray* means a one-dimensional direction in syndrome or coefficient space; different raw witnesses may represent the same ray and hence the same slope.

Theorem 3.2 (Exact syndrome–secant compiler)**.** *Put $t = n - a$, and for a syndrome line $\ell_{y_0,y_1} = \{y_0 + \gamma y_1 : \gamma \in \mathbb{F}\}$ define*

$$\Theta_t(y_0, y_1) = |\{\gamma \in \mathbb{F} : \exists E \subseteq D,\ |E| \le t, y_0 + \gamma y_1 \in V_E,\ \{y_0, y_1\} \not\subseteq V_E\}|. \tag{3.3}$$

Then

$$B_C^{\mathrm{MCA}}(a) = \max_{y_0, y_1 \in \mathbb{F}^{n-k}} \Theta_t(y_0, y_1). \tag{3.4}$$

For each fixed E, the affine line has at most one transverse intersection parameter with V_E. For a Reed–Solomon code the spaces V_E are the spans of at most t points of the weighted rational normal curve in (3.1). Thus the higher-dimensional ray problem is exactly a transverse line–secant incidence problem, with repeated intersection parameters deduplicated.

Proof. Apply Proposition 3.1 with $S = D \setminus E$, and take the union over $|E| \le t$. The syndrome map is surjective, so maximizing over received pairs is the same as maximizing over (y_0, y_1). In the quotient $\mathbb{F}^{n-k}/V_E$, a transverse intersection solves $\overline{y}_0 + \gamma \overline{y}_1 = 0$ with the two projected vectors not both zero, and therefore has at most one finite solution. The last assertion is the Vandermonde description of the parity columns. □

A *parity-column circuit* is a minimally linearly dependent set of columns of H. The next lemma says that many distinct transverse rays cannot all be supported inside an independent union of columns.

Lemma 3.3 (Independent-union ray bound)**.** *Let every set of at most R parity columns be independent, put $t < R$, and fix a syndrome line. For each slope in a transverse set Z, choose an error set E_γ, $|E_\gamma| \le t$, witnessing (3.3). If $U = \bigcup_{\gamma \in Z} E_\gamma$ has size at most R, then $|Z| \le t+1$. Consequently, more than $t+1$ transverse slopes force the witness union to contain a parity-column circuit.*

Proof. If $|Z| \le 1$, there is nothing to prove; assume that Z contains two distinct slopes. Two distinct line points lie in V_U, hence $y_0, y_1 \in V_U$. Column independence gives unique coefficient lifts $e_0, e_1 \in \mathbb{F}^U$, and every chosen lift is $e_\gamma = e_0 + \gamma e_1$. Let U' be the set of s coordinates where the pair (e_0, e_1) is nonzero. If $s \le t$ and E_γ contained all of U', then $y_0, y_1 \in V_{U'} \subseteq V_{E_\gamma}$, contradicting the transversality of this particular witness. Hence some $x \in U' \setminus E_\gamma$ exists, and uniqueness of the lift gives $e_0(x) + \gamma e_1(x) = 0$. Each nonzero affine coordinate function has at most one root, so $|Z| \le s \le t$. If $s > t$, every chosen lift of weight at most t has at least $s - t$

zero coordinates. Double counting the pairs (γ, x) with $e_0(x) + \gamma e_1(x) = 0$ gives $|Z|(s-t) \le s$, whence $|Z| \le \lfloor s/(s-t) \rfloor \le t+1$. If $|U| \le R$, independence holds; the contrapositive proves the circuit assertion. □

Theorem 3.4 (Bounded residual-kernel ray compiler)**.** *Let H be a parity-check matrix over $\mathbb{F}$, let U be a set of columns, and put*

$$\kappa(U) = \dim_{\mathbb{F}} \ker(H_U : \mathbb{F}^U \longrightarrow \mathbb{F}^{n-k}).$$

Fix a syndrome line $y_0 + \gamma y_1$, an integer $t \ge 0$, and a set Z of finite transverse slopes. Suppose that for every $\gamma \in Z$ there is $E_\gamma \subseteq U$, $|E_\gamma| \le t$, such that

$$y_0 + \gamma y_1 \in V_{E_\gamma}, \qquad \{y_0, y_1\} \not\subseteq V_{E_\gamma}.$$

Then

$$|Z| \le (t+1)|\mathbb{F}|^{\kappa(U)}. \tag{$\mathrm{RC}_{\ker}$}$$

For a Reed–Solomon parity check of redundancy $R = n - k$, $\kappa(U) = \max\{0, |U| - R\}$. Hence a certified residual chart with $(|U| - R)_+ \log|\mathbb{F}| = o(n)$ has a direct subexponential ray payment with the displayed exact finite loss.

Proof. The assertion is immediate when $|Z| \le 1$. Otherwise two line points belong to V_U, so $y_0, y_1 \in V_U$. Choose lifts $b_0, b_1 \in \mathbb{F}^U$ with $H_U b_i = y_i$. For every γ, choose a lift c_γ supported in E_γ and put $z_\gamma = c_\gamma - b_0 - \gamma b_1 \in \ker H_U$.

Fix $z \in \ker H_U$, and consider the slopes for which $z_\gamma = z$. Put

$$U_z = \{x \in U : (b_0(x) + z(x), b_1(x)) \ne (0,0)\}, \qquad s = |U_z|.$$

Every active coordinate is a nonzero affine function of γ, and therefore vanishes for at most one slope. Hence the total number of active zero incidences in this kernel class is at most s. If $s > t$, the weight bound on c_γ forces at least $s - t$ active zeros, so

$$|Z_z|(s-t) \le s, \qquad |Z_z| \le \lfloor s/(s-t) \rfloor \le t+1.$$

If $s \le t$, transversality forces at least one active zero: without one, $U_z = \operatorname{supp} c_\gamma \subseteq E_\gamma$, while both $b_0 + z$ and b_1 are supported in U_z, putting y_0, y_1 in V_{E_γ}. Thus this kernel class contains at most $s \le t$ slopes. Summing over the $|\mathbb{F}|^{\kappa(U)}$ kernel elements proves the bound. The MDS column property gives the Reed–Solomon formula for $\kappa(U)$. □

Theorem 3.5 (Single MDS-circuit ray compiler)**.** *Let H be a Reed–Solomon parity check of redundancy R, let $U \subseteq D$ have size $R+1$, put $t < R$, and fix a syndrome line $y_0 + \gamma y_1$. If every slope in a transverse set Z has a witnessing error set $E_\gamma \subseteq U$ of size at most t, then*

$$|Z| \le \binom{R+1}{2}. \tag{$\mathrm{RC}_{\mathrm{circ}}$}$$

Thus one fixed MDS-circuit profile has a field-independent polynomial ray payment.

Proof. There is nothing to prove when $|Z| \le 1$. Otherwise two distinct line points lie in V_U; subtracting them and solving for the constant term gives $y_0, y_1 \in V_U$. Choose lifts $b_0, b_1 \in \mathbb{F}^U$. The circuit kernel is generated by a vector z whose coordinates are all nonzero, since a zero coordinate would give a proper dependent subset of the minimal circuit U. For each slope choose $c_\gamma \in \mathbb{F}$ so that $e_\gamma = b_0 + \gamma b_1 + c_\gamma z$ is supported in E_γ. For $x \in U$, set

$$f_x(\gamma) = -(b_0(x) + \gamma b_1(x))/z(x).$$

A zero coordinate is an equality $c_\gamma = f_x(\gamma)$, and at least $R + 1 - t \ge 2$ such equalities hold.

Every transverse slope makes two nonidentical functions $f_x, f_{x'}$ equal. Indeed, otherwise all vanishing coordinates belong to one class C of identical affine functions, say $f_x = \alpha + \beta\gamma$ on C, and no function outside C has that value at γ. Then $\operatorname{supp} e_\gamma = U \setminus C \subseteq E_\gamma$, while $b_0 + \alpha z$ and $b_1 + \beta z$ are also supported in $U \setminus C$. This puts both y_0, y_1 in V_{E_γ}, a contradiction. Two nonidentical affine functions agree at at most one finite slope. Charging each slope to one unordered pair proves the displayed bound. □

4. The mean-overlap regime

For a fixed received line, the exact bad-slope set is therefore the set of unique cancelling line parameters contributed by the transverse quotient equations. For each E, $|E| \le n-a$, project the two syndromes to $\mathbb{F}^R/V_E$. If the second projection is nonzero and the first is a scalar multiple of it, record the unique cancelling scalar; otherwise record nothing. Deduplicating these scalars gives the numerator. In particular, the number of supports, the number of pairs of supports, and the number of distinct slopes are three different quantities.

We call $3(n-a) \le d-1$ the *deep regime*: the allowed error radius $n-a$ is so small relative to the minimum codeword weight d that the union of three permitted error supports still cannot contain a nonzero codeword support.

Theorem 4.1 (Deep-regime upper bound)**.** *Let $C \subseteq \mathbb{F}^n$ be linear with minimum nonzero weight d, let $0 \le a \le n$, put $r = n-a$, and assume $3r \le d-1$. Then*

$$B_C^{\mathrm{MCA}}(a) \le r+1. \tag{3.5}$$

Proof. Fix (u_0, u_1), and let Z be its set of MCA-bad slopes at agreement a. If $|Z| \le r+1$, there is nothing to prove. Otherwise choose distinct $\gamma_1, \gamma_2 \in Z$, together with witnessing codewords p_1, p_2 and error sets E_1, E_2, each of size at most r, such that $u_0 + \gamma_i u_1 = p_i$ off E_i. Solving the two affine equations gives $c_0, c_1 \in C$ for which the error pair $e_i = u_i - c_i$ is supported on $T = E_1 \cup E_2$, $|T| \le 2r$.

For any $\gamma \in Z$, choose a witnessing explaining codeword p. The codeword $p - (c_0 + \gamma c_1)$ is supported on at most $3r \le d-1$ positions and is zero. Thus $\mathrm{wt}(e_0 + \gamma e_1) \le r$. Let T' be the support of the error pair. If $|T'| \ge r+1$, every $\gamma \in Z$ makes at least $|T'| - r$ nonzero affine coordinate functions $e_{0,x} + \gamma e_{1,x}$ vanish. Each coordinate has at most one root, so $|Z|(|T'| - r) \le |T'|$ and $|Z| \le r+1$.

It remains that $|T'| \le r$. If γ is MCA-bad on a set S, $|S| \ge n-r$, its explaining codeword equals $c_0 + \gamma c_1$, since their difference would otherwise have weight at most $2r \le d-1$. The set S must meet T', and at a point of $S \cap T'$ the nonzero pair satisfies $e_{0,x} + \gamma e_{1,x} = 0$. Such a point determines γ, giving at most $|T'| \le r$ bad slopes. □

Recall that for a nonempty challenge set $\Gamma \subseteq \mathbb{F}$, $B_{C,\Gamma}^{\mathrm{MCA}}(a)$ denotes the same maximum with slopes restricted to Γ.

The following "tangent floor" is an elementary family of neighboring coordinate-span intersections. The name does not assert a hidden differentiability or multiplicity condition.

Proposition 4.2 (Universal MDS tangent floor)**.** *For every linear $[n,k]$ MDS code, every $a \ge k+1$, and every nonempty $\Gamma \subseteq \mathbb{F}$,*

$$B_{C,\Gamma}^{\mathrm{MCA}}(a) \ge \min\{|\Gamma|, n-a+1\}. \tag{3.6}$$

Proof. Choose a full-rank parity-check matrix H for C. The MDS property is equivalent to independence of every set of at most $n-k$ columns; the syndrome-line normal form of Proposition 3.1 is linear algebra and applies to this H. Put $j = n-a$, $t = \min\{|\Gamma|, j+1\}$, and choose distinct $\gamma_1, \dots, \gamma_t \in \Gamma$ and independent parity columns $h_{x_1}, \dots, h_{x_t}$. Define $s_1 = \sum_i h_{x_i}$ and $s_0 = -\sum_i \gamma_i h_{x_i}$, and lift them to received words. For $E_i = \{x_\ell : \ell \ne i\}$, $s_0 + \gamma_i s_1 \in V_{E_i}$, whereas $s_1 \notin V_{E_i}$. Proposition 3.1 makes every γ_i bad on $D \setminus E_i$, which has size at least a. □

Corollary 4.3 (Exact deep numerator)**.** *If $C \subseteq \mathbb{F}^n$ is a linear $[n,k]$ MDS code, $1 \le k < n$, $k+1 \le a \le n$, $\varnothing \ne \Gamma \subseteq \mathbb{F}$, and $3(n-a) \le n-k$, then*

$$B_{C,\Gamma}^{\mathrm{MCA}}(a) = \min\{|\Gamma|, n-a+1\}.$$

Proof. The upper bound is Theorem 4.1, intersected with Γ. The hypotheses imply $a \ge k+1$, so Proposition 4.2 gives equality. □

Remark 4.4 (The zero-radius normalization). Under the formal affine sampler, Corollary 4.3 at $a = n$ gives $B_C^{\mathrm{MCA}}(n) = 1$, and hence $e_{\mathrm{MCA}}(C, 0) = 1/|\mathbb{F}|$ for every proper linear MDS code. This also follows directly: in $\mathbb{F}^n/C$, a noncontained affine line meets the zero coset in at most

one parameter, and the bound is attained by a line through the zero coset. A value $2/|\mathbb{F}|$ belongs to a different or loose convention and must not be imported into the affine definition of [ABF26].

The next lemma isolates the deterministic mechanism behind both extensions of the deep range.

Lemma 4.5 (Large-overlap collapse)**.** *Fix a received pair and an exact agreement $a = n - r \geq k+1$. Choose one noncommon exact witness support S_z and explaining polynomial p_z for every slope z in its bad set Z. If two distinct chosen supports satisfy*

$$|S_z \cap S_w| \geq k + r, \tag{MO1}$$

then $|Z| \leq r + 1$.

Proof. Put

$$c_1 = \frac{p_z - p_w}{z - w}, \qquad c_0 = p_z - z c_1.$$

On $S_z \cap S_w$, the two line equations imply that the received pair equals (c_0, c_1). Let C_0 be the full common support of this codeword pair and write $c = |C_0| \geq k + r$. For every $u \in Z$, $|S_u \cap C_0| \geq a + c - n \geq k$, so the explaining polynomial p_u equals $c_0 + uc_1$.

Subtract this codeword pair from the received pair. At every coordinate outside C_0 at least one of the two residual values is nonzero, and a coordinate can cancel at at most one finite slope. A noncommon witness for u must contain at least $\max\{1, a - c\}$ such cancellations: the support-size condition gives $a - c$, and one is still necessary when $a \leq c$. Double counting therefore gives

$$|Z| \leq \left\lfloor \frac{n - c}{\max\{1, a - c\}} \right\rfloor \leq r + 1. \tag{MO2}$$

Indeed, for $c \geq a$ the numerator is at most r; for $c = a - 1$ it is $r + 1$; and, with $x = a - c \geq 2$, the quotient is $(r + x)/x \leq r + 1$. □

Theorem 4.6 (Quadratic mean-overlap theorem)**.** *Let $C = \mathrm{RS}_{\mathbb{F}}(D, k)$, put $R = n - k$, let $0 \leq r \leq R - 1$, and let $\varnothing \neq \Gamma \subseteq \mathbb{F}$. If*

$$(n - r)^2 \geq n(k + r), \qquad \textit{equivalently} \qquad r^2 \geq n(3r - R), \tag{MO3}$$

then

$$B^{\mathrm{MCA}}_{C,\Gamma}(n - r) = \min\{|\Gamma|, r + 1\}. \tag{MO4}$$

The same upper bound holds for every linear $[n, k]$ MDS code.

Proof. The universal tangent construction gives the lower bound. For the upper bound, fix a received pair and its challenge-restricted bad set Z. If $|Z| \leq r + 1$ there is nothing to prove. Otherwise choose $r + 2$ slopes and distinct exact witness supports $S_1, \ldots, S_{r+2}$, each of size $a = n - r$. Suppose every pair had intersection at most $k + r - 1$, and put $d_x = \#\{i : x \in S_i\}$. Then

$$\sum_x \binom{d_x}{2} \leq \binom{r+2}{2}(k + r - 1), \tag{MO5}$$

whereas Cauchy–Schwarz and $a^2/n \geq k + r$ give

$$\sum_x \binom{d_x}{2} \geq \frac{(r+2)^2 a^2/n - (r+2)a}{2} \geq \frac{(r+2)^2(k+r) - (r+2)a}{2}. \tag{MO6}$$

The last lower bound exceeds the upper bound in (MO5) by

$$\frac{r+2}{2}(k + 3r - n + 1).$$

When $3r > R$ this is positive. When $3r \leq R$, the exact deep theorem already gives the desired upper bound. Thus in the remaining case some pair satisfies (MO1), and Lemma 4.5

contradicts $|Z| \geq r+2$. The exact-support reduction and large-overlap lemma use only the MDS information-set property, proving the stated extension. □

Corollary 4.7 (Exact quadratic staircase)**.** *In the setting of Theorem 4.6, put*

$$r_{\mathrm{quad}}(n,k) = \left\lfloor \frac{3n - \sqrt{n(5n+4k)}}{2} \right\rfloor. \tag{MO7}$$

For every $0 \leq r \leq r_{\mathrm{quad}}(n,k)$ and every nonempty challenge set Γ,

$$B^{\mathrm{MCA}}_{C,\Gamma}(n-r) = \min\{|\Gamma|, r+1\}.$$

If $k/n \to \rho$, then $r_{\mathrm{quad}}(n,k)/n \to \alpha_\rho = (3-\sqrt{5+4\rho})/2$.

Proof. The smaller root of $r^2 - 3nr + n(n-k) = 0$ is the real number inside the floor in (MO7). Apply Theorem 4.6; the limiting formula follows by division by n. □

Proposition 4.8 (Limit of pairwise-overlap information)**.** *Let $k/n \to \rho \in (0,1)$, $r/n \to \alpha \in (0,1)$, and suppose*

$$\alpha^2 < 3\alpha - (1-\rho). \tag{MO8}$$

For all sufficiently large n, there are $r+2$ distinct r-subsets T_i such that, whenever $j = 3r-(n-k) \leq r$, $|T_i \cap T_\ell| \leq j-1$ for every $i \neq \ell$. If $j > r$, that intersection condition is automatic. Hence pairwise support overlap alone cannot extend the exact theorem past the quadratic boundary; further progress must use the syndrome, determinant, or fiber incidence. This is a limitation of the method, not a construction of an MCA-bad line.

Proof. When $j \leq r$, choose $r+2$ independent uniform r-subsets. The intersection of a fixed pair is hypergeometric with mean $r^2/n = (\alpha^2 + o(1))n$, whereas $j = (3\alpha - (1-\rho) + o(1))n$. More explicitly,

$$\Pr(|T_1 \cap T_2| = s) = \frac{\binom{r}{s}\binom{n-r}{r-s}}{\binom{n}{r}}.$$

After writing $s = xn$, Stirling's formula gives the exponent

$$\alpha h(x/\alpha) + (1-\alpha)h((\alpha - x)/(1-\alpha)) - h(\alpha),$$

which is strictly concave and has its unique maximum zero at $x = \alpha^2$. The strict gap in (MO8), followed by summing at most $n+1$ terms, therefore gives probability $e^{-\Omega(n)}$ that one fixed intersection reaches j. A union bound over $O(n^2)$ pairs is smaller than one. A successful family is automatically distinct, because equal members have intersection $r \geq j$. When $j > r$, choose any $r+2$ distinct r-sets. □

The following complementary theorem can be stronger when the overlap defect $j = 3r - R$ is small. Its only new input is a packing inequality; in particular it uses no smoothness or Fourier estimate.

Theorem 4.9 (Exact overlap–packing criterion)**.** *Let $C = \mathrm{RS}_{\mathbb{F}}(D,k)$, put $R = n-k$, and fix an integer $0 \leq r \leq R-1$. Set $a = n-r$ and $j = 3r-R$. If either*

$$j \leq 0, \tag{OP0}$$

or

$$1 \leq j \leq r, \qquad \binom{n}{j} \leq (r+1)\binom{r}{j}, \tag{OP}$$

then

$$B^{\mathrm{MCA}}_{C}(n-r) = r+1. \tag{OP$'$}$$

The same upper bound holds for every linear $[n,k]$ MDS code; the matching lower bound holds whenever the field contains at least $r+1$ distinct slope parameters.

Proof. Condition (OP0) is Corollary 4.3, so assume (OP). Fix an arbitrary received pair (f, g), let Z be its bad-slope set, and use Lemma 2.11 to choose, for every $z \in Z$, an exact witness support S_z, $|S_z| = a$, and an explaining polynomial $p_z \in \mathbb{F}[X]_{<k}$. A single noncommon support cannot carry two distinct slopes, so the supports chosen for distinct elements of Z are distinct.

For the stated MDS extension, the same exact-support reduction is available: interpolate the unexplained component on any k-subset of the original witness support; because every k-subset is an information set, nonexplanation is detected at one further coordinate, and that $(k+1)$-set can be enlarged to size a.

First suppose that two of them have a large intersection:

$$|S_z \cap S_w| \geq n + k - a = k + r \qquad (z \neq w). \tag{OP1}$$

Then Lemma 4.5 gives $|Z| \leq r + 1$.

It remains to consider the complementary case, in which every pair of chosen supports satisfies $|S_z \cap S_w| \leq n + k - a - 1$. Put $T_z = D \setminus S_z$, so $|T_z| = r$. For distinct slopes,

$$|T_z \cap T_w| = n - 2a + |S_z \cap S_w| \leq 3r - R - 1 = j - 1. \tag{OP3}$$

No j-subset of D can therefore lie in two different T_z's. Counting pairs (z, J) with $J \in \binom{T_z}{j}$ yields

$$|Z|\binom{r}{j} \leq \binom{n}{j} \leq (r+1)\binom{r}{j},$$

and hence $|Z| \leq r + 1$. The universal tangent construction gives the reverse inequality. The proof used only linearity, the fact that agreement of two codewords on k coordinates forces equality, and the exact-support reduction; these are MDS properties. The coordinate tangent construction gives the stated MDS lower bound. □

Proposition 4.10 (General overlap–packing upper bound)**.** *Let $C \subseteq \mathbb{F}^n$ be a linear $[n, k]$ MDS code, let $\varnothing \neq \Gamma \subseteq \mathbb{F}$, let $0 \leq r \leq n - k - 1$, and suppose $1 \leq j = 3r - (n - k) \leq r$. Let $A(n, r, j)$ be the largest size of a family of r-subsets of an n-set whose distinct members intersect in at most $j - 1$ points. Then*

$$B^{\mathrm{MCA}}_{C,\Gamma}(n - r) \leq \min\{|\Gamma|, \max(r + 1, A(n, r, j))\} \tag{OP5}$$

and

$$A(n, r, j) \leq \left\lfloor \frac{\binom{n}{j}}{\binom{r}{j}} \right\rfloor. \tag{OP6}$$

Proof. For a fixed received pair, either two chosen exact witness supports have intersection at least $k + r$, in which case Lemma 4.5 gives at most $r + 1$ slopes, or every pair of complementary r-sets intersects in at most $j - 1$ points. The latter family has size at most $A(n, r, j)$. Finally, no j-subset can lie in two members of such a packing, so counting contained j-subsets proves (OP6). □

Corollary 4.11 (Quadratic staircase at the official rates)**.** *Let $D \subseteq \mathbb{F}$ have size $n = 2^e$, put $C = \mathrm{RS}_{\mathbb{F}}(D, k)$, let $\rho = k/n \in \{1/2, 1/4, 1/8, 1/16\}$, and put*

$$r_\rho(n) = \left\lfloor n \frac{3 - \sqrt{5 + 4\rho}}{2} \right\rfloor.$$

For every n-point Reed–Solomon domain over every field,

$$B^{\mathrm{MCA}}_{C}(n - r) = r + 1 \qquad (0 \leq r \leq r_\rho(n)). \tag{OP4}$$

Proof. This is Corollary 4.7, since the displayed integer is $r_{\mathrm{quad}}(n, k)$. □

Corollary 4.12 (Self-contained quadratic-range target compiler)**.** *In the setting of Corollary 4.11, put $q = |\mathbb{F}|$ and $B = \lfloor q/2^{128} \rfloor$. If*

$$1 \leq B \leq M_\rho(n) := \min\{r_\rho(n) + 1, n - k - 1\},$$

then

$$e_{\mathrm{MCA}}(C,\delta) \le 2^{-128} \quad \Longleftrightarrow \quad \lfloor \delta n \rfloor \le B-1. \tag{PW}$$

Thus every such row is solved throughout the field window $2^{128} \le q < 2^{128}(M_\rho(n)+1)$: its largest safe grid radius is $(B-1)/n$, while its real transition supremum is B/n and that endpoint is unsafe. If $q < 2^{128}$, even radius zero is unsafe.

Proof. For $r \le B-1$, Corollary 4.11 gives the exact numerator $r+1 \le B \le q/2^{128}$. At $r=B$, the tangent floor applies because $B \le n-k-1$, and gives at least $B+1 > q/2^{128}$ slopes. Monotonicity proves the equivalence and the endpoint statements. If $B=0$, the tangent floor at $r=0$ gives the exact numerator one. □

5. Prior Reed–Solomon completion below half distance

The quadratic theorem is self-contained and useful at finite parameters, but it is not the largest known RS range at constant rate. We now make the relevant implication of [BCHKS25] explicit at the integer numerator level. This section is deliberately separated from Section 4: the input below belongs to the literature, whereas Theorem 4.6 is the finite MDS argument of this paper.

Put $q = |\mathbb{F}|$. For a word u and a received pair (u_0,u_1), write $d(u,C)$ and $d_2((u_0,u_1),C^2)$ for their relative Hamming and relative column distances from C and C^2, respectively. For $0 \le \alpha < \beta < 1$, define the two-radius correlated-agreement error by

$$e_{\mathrm{CA}}(C,\alpha,\beta) = \frac{1}{q} \max_{\substack{u_0,u_1 \in \mathbb{F}^D \\ d_2((u_0,u_1),C^2) > \beta}} \#\{\gamma \in \mathbb{F} : d(u_0+\gamma u_1, C) \le \alpha\}. \tag{LC0}$$

The one-radius notation $e_{\mathrm{CA}}(C,\alpha)$ means that the far-pair radius is also α.

Theorem 5.1 (BCHKS two-radius bound)**.** *Let $C = \mathrm{RS}_{\mathbb{F}}(D,k)$, put $q=|\mathbb{F}|$, and $\Delta = (n-k+1)/n$. If $\Delta/3 \le \alpha \le \Delta/2 - 1/n$ and $\alpha < \beta < 1$, then*

$$e_{\mathrm{CA}}(C,\alpha,\beta) \le \frac{1}{q} \max\left\{ \frac{\Delta-\alpha}{\alpha(\Delta-2\alpha)}, \frac{\beta}{\beta-\alpha} \right\}. \tag{BCHKS}$$

Proof. BCHKS [BCHKS25, Thm. 1.3] uses degree at most $k-1$, hence dimension k and relative minimum distance Δ. Fix a received pair at column distance greater than β, and let Z be its set of α-close slopes. If

$$|Z| \ge \frac{\Delta-\alpha}{\alpha(\Delta-2\alpha)},$$

that theorem gives column distance at most $(1+1/(|Z|-1))\alpha$. The assumed distance greater than β therefore forces $|Z| < \beta/(\beta-\alpha)$. Otherwise $|Z|$ is already below the first displayed threshold. Taking the larger bound, then maximizing over the received pair and dividing by q, proves (BCHKS). This also verifies the dimension-k normalization; all denominators are positive on the stated range. □

Theorem 5.2 (Integral BCHKS completion)**.** *Let $C = \mathrm{RS}_{\mathbb{F}}(D,k)$, let $q=|\mathbb{F}|$, put $R = n-k$, and let $r \in \mathbb{Z}$ satisfy*

$$3r \ge R+1, \qquad 2r < R. \tag{LC1}$$

Then

$$B_C^{\mathrm{MCA}}(n-r) \le \left\lfloor \max\left\{ \frac{n(R+1-r)}{r(R+1-2r)},\ r+1 \right\} \right\rfloor. \tag{LC2}$$

Consequently, if

$$n(R-r) < r(r+2)(R-2r), \tag{LC3}$$

then, for every nonempty $\Gamma \subseteq \mathbb{F}$,

$$B_{C,\Gamma}^{\mathrm{MCA}}(n-r) = \min\{|\Gamma|, r+1\}. \tag{LC4}$$

Proof. Put $\delta_{\mathrm{fld}} = r/n$ and, for $0 < \theta < 1$, put $\delta_{\mathrm{int}} = (r+\theta)/n$. Hamming distances belong to $\{0, 1/n, \dots, 1\}$, so replacing the far-pair radius δ_{fld} by δ_{int} does not change the far-pair condition; this is the grid observation of [ABF26, Rem. 4.10]. The hypotheses in (LC1) give $\Delta/3 \le \delta_{\mathrm{fld}} \le \Delta/2 - 1/n$, and $\delta_{\mathrm{fld}} < \delta_{\mathrm{int}} < 1$. Hence Theorem 5.1 gives

$$q\, e_{\mathrm{CA}}(C, r/n) \le \max\left\{ \frac{n(R+1-r)}{r(R+1-2r)}, \frac{r+\theta}{\theta} \right\}. \tag{LC5}$$

On the integer grid, $q e_{\mathrm{CA}}(C, r/n) = B_C^{\mathrm{CA}}(n-r)$. The second inequality in (LC1) also places the radius below $d_{\min}/(2n)$, because $d_{\min} = R+1$. By Theorem 2.6,

$$B_C^{\mathrm{MCA}}(n-r) = \max\{B_C^{\mathrm{CA}}(n-r), \min\{q, r\}\}.$$

The right side of (LC5) is strictly larger than r, because $(r+\theta)/\theta > r$. It therefore bounds both branches of this maximum and hence bounds the integer $B_C^{\mathrm{MCA}}(n-r)$. For the full-field maximum one may alternatively invoke the known below-half equality [ABF26, Lem. 4.6]; the sparse decomposition is kept here because it is exact before maximizing over a restricted challenge set.

Letting $\theta \to 1^-$ proves that this integer is at most the real maximum in (LC2), and taking the floor proves (LC2). Under (LC3), positivity of the denominators and $r > 0$ give

$$\frac{n(R+1-r)}{r(R+1-2r)} \le \frac{n(R-r)}{r(R-2r)} < r+2.$$

Since $r+1$ is integral, the floor in (LC2) is at most $r+1$, so the full-field numerator is at most $r+1$. Restriction to Γ gives the upper bound $\min\{|\Gamma|, r+1\}$, and Proposition 4.2 gives the reverse bound. □

Remark 5.3 (Priority). The substance of the upper bound in Theorem 5.2 is prior work. In particular, [BCHKS25, Cor. 1.4] states a sharp lossless conclusion under its explicit distance and radius hypotheses, and [BCHKS25, Rem. 2.5] discusses the matching perturbation construction. Formula (LC5) is included to expose the integer limit and to make the challenge-subset and target compilers below completely transparent.

Corollary 5.4 (Positive density of exact radii)**.** *Let $C_n = \mathrm{RS}_{\mathbb{F}_n}(D_n, k_n)$ be an RS row sequence with $k_n/n \to \rho \in (0,1)$, let $\varnothing \ne \Gamma_n \subseteq \mathbb{F}_n$, and let $r_n \in \mathbb{Z}$ satisfy $r_n/n \to \tau$, and suppose $0 < \tau < (1-\rho)/2$. Then*

$$B_{C_n,\Gamma_n}^{\mathrm{MCA}}(n - r_n) = \min\{|\Gamma_n|, r_n + 1\} \tag{LC6}$$

for all sufficiently large n. More uniformly, for every $\eta > 0$, equality holds throughout

$$0 \le r \le \left(\frac{1-\rho}{2} - \eta \right) n$$

for every sufficiently large row whose rate tends to ρ. Consequently the exact staircase contains an initial interval of normalized length $(1-\rho)/2 - o(1)$, has lower density at least $(1-\rho)/2$ in the full Hamming-radius interval, and has asymptotic relative density one in the unique-decoding interval.

Proof. The quadratic theorem covers every sequence whose limiting radius lies below $\alpha_\rho = (3 - \sqrt{5+4\rho})/2$. In the remaining compact subinterval of $(0, (1-\rho)/2)$, the hypotheses (LC1) hold for large n, while

$$\frac{n(R-r)}{r(R-2r)} = O_{\rho,\eta}(1) \quad \text{and} \quad r+1 = \Theta_{\rho,\eta}(n).$$

Thus (LC3) holds and Theorem 5.2 applies. The uniform statement follows by the same compactness argument, using the deep theorem near radius zero. Counting integer radii gives the density assertions. □

Corollary 5.5 (Official-rate staircase)**.** *Let $D \subseteq \mathbb{F}$ have size $n = 2^e \geq 128$, put $C = \mathrm{RS}_{\mathbb{F}}(D, k)$, let $k = \rho n$, and $\rho \in \{1/2, 1/4, 1/8, 1/16\}$. Define*

$$r_\rho^\sharp = \begin{cases} n/4 - 3, & \rho = 1/2, \\ 3n/8 - 2, & \rho = 1/4, \\ 7n/16 - 2, & \rho = 1/8, \\ 15n/32 - 2, & \rho = 1/16. \end{cases} \tag{LC7}$$

Then for every nonempty $\Gamma \subseteq \mathbb{F}$,

$$B_{C,\Gamma}^{\mathrm{MCA}}(n - r) = \min\{|\Gamma|, r + 1\} \qquad (0 \leq r \leq r_\rho^\sharp). \tag{LC8}$$

Proof. Put $F(r) = r(r+1)(R-2r) - n(R-r)$, where $R = n - k$, and $L = \lceil (R+1)/3 \rceil$. The deep theorem covers $r < L$. On the positive real axis, $F'(r) = -6r^2 + 2(R-2)r + (R+n)$ has exactly one positive zero, so the minimum of F on $[L, r_\rho^\sharp]$ occurs at an endpoint. Since $R \geq n/2 \geq 64$,

$$F(L) \geq \frac{(R+1)(R+4)(R-8) - 9n(2R-1)}{27} > 0;$$

indeed the numerator is at least $R\{R(R-4) - 18n\} > 0$. Direct substitution at the other endpoint gives, in the order of the four rates,

$$\frac{n^2 - 84n + 288}{8}, \quad \frac{3n^2 - 104n + 128}{16}, \quad \frac{21n^2 - 464n + 512}{64}, \quad \frac{105n^2 - 1952n + 2048}{256},$$

all positive for $n \geq 128$. Therefore $F(r) \geq 0$ throughout the remaining integer interval. Since $r > 0$ and $R - 2r > 0$ there,

$$r(r+2)(R-2r) - n(R-r) = F(r) + r(R-2r) > 0,$$

which is (LC3). Apply Theorem 5.2. □

Corollary 5.6 (Almost-half-distance target compiler)**.** *In the setting of Corollary 5.5, let $0 \leq \varepsilon \leq 1$ and put $B = \lfloor \varepsilon |\Gamma| \rfloor$. If*

$$1 \leq B < |\Gamma|, \qquad B \leq r_\rho^\sharp + 1,$$

then the complete safe set is

$$\{\delta : e_{\mathrm{MCA},\Gamma}(C, \delta) \leq \varepsilon\} = [0, B/n), \tag{LC9}$$

and the endpoint is unsafe. For the full-field target 2^{-128}, this applies throughout $2^{128} \leq q < 2^{128}(r_\rho^\sharp + 2)$.

Proof. At every radius $r \leq B - 1$, Corollary 5.5 gives numerator $r + 1 \leq B$. At radius B, the tangent floor gives at least $B + 1$ bad slopes. Monotonicity and $r(\delta) = \lfloor \delta n \rfloor$ prove the assertion. □

6. Agreement-set shortening beyond Johnson

We record Jo's shortening mechanism because it preserves actual bad slopes and therefore bypasses support-to-ray conversion. The shortened evaluation set need not retain any smoothness.

Lemma 6.1 (Normalized shortening, Jo [Jo26, Lem. 3.1])**.** *Let $C = \mathrm{RS}_{\mathbb{F}}(D, k)$, let $S \subseteq D$ have size $t < k$, put $P_S(X) = \prod_{x \in S}(X - x)$, and let $I_S(w)$ be the degree-less-than- t polynomial interpolating $w|_S$, with $P_\varnothing = 1$ and $I_\varnothing = 0$. Define, for $x \in D \setminus S$,*

$$Q_S(w)(x) = \frac{w(x) - I_S(w)(x)}{P_S(x)}. \tag{SH1}$$

Then Q_S is linear and maps C into $C_S = \mathrm{RS}_{\mathbb{F}}(D \setminus S, k - t)$. If a slope is MCA-bad for C on a witness support $A \supseteq S$, then it is MCA-bad for the shortened line at agreement $|A| - t$ in C_S.

Proof. We reproduce Jo's normalization argument. For a code polynomial f, the polynomial $f - I_S(f)$ vanishes on S, so division by P_S leaves degree less than $k-t$. If the shortened received pair had explanations g_i on $A \setminus S$, then $I_S(u_i) + P_S g_i$ would explain each original word on all of A, contradicting the witness nontriviality. □

Theorem 6.2 (Agreement-set shortening transfer, Jo [Jo26, Thm. 3.2])**.** *Let $C = \mathrm{RS}_{\mathbb{F}}(D,k)$ have length n, let $k+1 \le a \le n$, and let $0 \le t < k$ with $t \le a$. For every nonempty challenge set $\Gamma \subseteq \mathbb{F}$,*

$$B^{\mathrm{MCA}}_{C,\Gamma}(a) \le \frac{\binom{n}{t}}{\binom{a}{t}} \max_{\substack{S \subseteq D \\ |S|=t}} B^{\mathrm{MCA}}_{C_S,\Gamma}(a-t). \tag{SH2}$$

Proof. This is Jo's double count, restricted to Γ. For a fixed received line, choose a witness support A_γ of size at least a for every bad $\gamma \in \Gamma$. Each A_γ contains at least $\binom{a}{t}$ shortening sets, and Lemma 6.1 preserves the same slope for every such set. Count pairs (γ, S) first by slope and then over the $\binom{n}{t}$ choices of S, and finally maximize over the received line. □

Corollary 6.3 (Fixed integer steps beyond Johnson, Jo [Jo26, Thm. 3.4])**.** *Fix integers $r \ge 2$ and $h \ge 1$. For all sufficiently large k, let $n = rk$, let $D \subseteq \mathbb{F}$ be arbitrary of size n, and put*

$$E_{r,h} = \left\lfloor n - \sqrt{n(k-1)} \right\rfloor + h, \qquad t_{r,h} = \left\lfloor \frac{2h\sqrt{r}}{(\sqrt{r}-1)^2} \right\rfloor + 1. \tag{SH3}$$

Then $E_{r,h}$ is the h-th integer error budget strictly beyond the exact Johnson boundary, and

$$B^{\mathrm{MCA}}_{C}(n - E_{r,h}) = O_{r,h}(k^6). \tag{SH4}$$

The assertion is uniform over the field and evaluation set.

The key identity behind the corollary is

$$(a-t)^2 - (n-t)(k-1-t) = a^2 - n(k-1) + t(n+k-1-2a). \tag{SH5}$$

The displayed $t_{r,h}$ makes the shortened side strictly Johnson-safe; the resulting normalized margin is $\Omega_{r,h}(k^{-1})$, and the BCHKS estimate is $O_{r,h}(k^6)$.

Proof of Theorem 1.1. Fix τ with $\tau_\rho(\delta) < \tau < \rho$ and put

$$t_n = \lfloor \tau n \rfloor, \quad E_n = \lfloor \delta n \rfloor, \quad N_n = n - t_n, \quad d_n = k_n - 1 - t_n, \quad A_n = a_n - t_n = N_n - E_n.$$

Because $\tau < \rho < \alpha$, while $k_n/n \to \rho$ and $a_n/n \to \alpha$, eventually

$$0 \le t_n \le \min\{a_n, k_n - 2\}. \tag{SH6}$$

Thus both the shortening transfer and Theorem 1.3 apply. Moreover,

$$(\alpha-\tau)^2 - (1-\tau)(\rho-\tau) = \alpha^2 - \rho + \tau(1+\rho-2\alpha) > 0, \tag{SH9}$$

so $A_n^2 > N_n d_n$ for all sufficiently large n. The limiting reduced rate $(\rho-\tau)/(1-\tau)$ is positive, and the strict surplus is bounded away from zero after normalization. Consequently $\mathcal{J}(N_n, d_n, E_n) = O_{\rho,\delta,\tau}(n)$, uniformly over the arbitrary shortened domains $D_n \setminus S$. Theorem 1.3 therefore makes the maximum shortened numerator in (SH2) polynomial in n. Stirling's formula therefore gives

$$\limsup_{n\to\infty} \frac{1}{n} \log_2\left(1 + B^{\mathrm{MCA}}_{C_n,\Gamma_n}(a_n)\right) \le H_2(\tau) - \alpha H_2(\tau/\alpha).$$

Let $\tau \downarrow \tau_\rho(\delta)$. This order of limits is essential: shortening at the critical density itself has no fixed Johnson margin.

Solving equality in (SH9) gives $\tau_\rho(\delta)$. Indeed $\rho < \alpha < \sqrt{\rho}$, so

$$1 + \rho - 2\alpha > 0, \qquad \rho - \tau_\rho(\delta) = \frac{(\alpha-\rho)^2}{1+\rho-2\alpha} > 0;$$

hence $0 < \tau_\rho(\delta) < \rho$. For fixed α,

$$\frac{d}{dx}\left[H_2(x) - \alpha H_2(x/\alpha)\right] = \log_2 \frac{1-x}{\alpha - x} > 0.$$

Comparison at $x = \tau_\rho(\delta)$ and $x = \rho$ proves (SF4); direct substitution at the two endpoints proves the final claims. □

Remark 6.4 (What shortening does not remove). Splitting a total shortening $t = t_1 + t_2$ into stages gives the same cost, because the two binomial ratios telescope to $\binom{n}{t}/\binom{a}{t}$. Thus constant relative improvement has the positive exponent $\Psi_\rho(\delta)$; the method does not by itself produce a subexponential numerator beyond Johnson.

Proposition 6.5 (Four finite rows beyond Johnson)**.** *Let*

$$p = (2^{128} - 255)2^{128} + 1, \qquad B = \left\lfloor p2^{-128} \right\rfloor = 2^{128} - 255,$$

and let $D \le \mathbb{F}_p^\times$ have order 512. The number p is prime. At target 2^{-128}, agreement-set shortening and the circuit-incidence envelope certify

(F512)

ρ	k	a	$(512-a)/512$	*certificate*	$\log_2 U <$
1/2	256	351	161/512	*shorten* 123	126.234492
1/4	128	233	279/512	*shorten* 67	127.290843
1/8	64	143	369/512	*shorten* 45	126.254265
1/16	32	52	115/128	*circuit* $b = 33$	127.310201.

In every row the rigorous numerator U is strictly smaller than B, so all four radii are safe and strictly beyond Johnson. The decimal logarithms are diagnostics only; the proof uses exact integer inequalities. These are upper certificates, not claims that the MCA numerator attains U, and failure one agreement lower would not prove unsafety.

Proof. Write $p = u2^{128} + 1$ with odd $u = 2^{128} - 255 < 2^{128}$. The exact congruence $3^{(p-1)/2} \equiv -1 \pmod p$ is a Proth certificate. Indeed every prime divisor ℓ of p satisfies $\ell \ge 2^{128} + 1$, whereas $p < 2^{256} + 1 < (2^{128} + 1)^2$, so a composite p could not have a prime divisor at most $\sqrt{p}$. Thus p is prime, and $512 \mid p - 1$ gives the subgroup.

For the first three rows put $E = 512 - a$, $N = 512 - t$, $d = k - 1 - t$, $A = a - t$, and $x = m + 1/2$, using

$$(t, m) = (123, 162), (67, 63), (45, 15), \qquad C = \frac{2}{3}x^5N^4 + xdN^2(E + 1).$$

The exact inequalities

$$A^2 > Nd, \qquad (m-1)^2A^2 < dm^2N, \qquad d(m+1)^2N \le m^2A^2$$

certify both the shortened Johnson condition and the ceiling parameter in Theorem 1.3. Combining that theorem with Theorem 6.2 gives

$$U = \frac{\binom{512}{t}}{\binom{a}{t}} \frac{C}{\sqrt{(dN)^3}}.$$

After squaring positive quantities, direct integer arithmetic gives

$$\left(\frac{\binom{512}{t}}{\binom{a}{t}} C \right)^2 < B^2(dN)^3$$

in all three cases, hence $U < B$. In the last row, Theorem 1.2 with $b = 33$ gives the exact integer upper bound $\left\lfloor \binom{512}{33}/\binom{51}{32} \right\rfloor = 2109546865085608674212113821347089722 91 < B$. Finally, the four exact inequalities $a^2 < 512(k-1)$, whose positive margins are respectively $7359, 10735, 11807, 13168$, prove that the four displayed radii exceed the exact MDS Johnson radii $1 - \sqrt{(k-1)/512}$, and hence also $1 - \sqrt{\rho}$. □

Jo's own challenge-scale post-Johnson rows and exact length-64 transitions are recorded in [Jo26, Thms. 5.1–5.2]; they are not duplicated here. The rows in Proposition 6.5 instead compare the shortening and circuit mechanisms on the common length-512 field already used elsewhere in this paper.

7. Reproducible finite Prize instances

Table 2. Explicit maximal-dimension prize-admissible rows. In each row $k = 2^{40}$, $B = \lfloor p/2^{128} \rfloor = r_\rho(n) + 1$, and (s, a_0) is part of the Proth certificate in Proposition 7.1.

rate	length	B	bits of p	(s, a_0)
1/2	2^{41}	389500552609	167	$(92, 3)$
1/4	2^{42}	1210584858040	169	$(93, 13)$
1/8	2^{43}	2879806199253	170	$(95, 5)$
1/16	2^{44}	6233898019554	171	$(97, 5)$

For completeness, the four field orders are

$$p_{1/2} = 132540169958804033333249306710494641010898987122689,$$
$$p_{1/4} = 411940680852499481698306614369841346700408394874881,$$
$$p_{1/8} = 979947269755402568812854322316630667196565607677953,$$
$$p_{1/16} = 212128557323758584829987561901119226267906543399731 3.$$

Their odd Proth coefficients u, in the same order, are

$$u_{1/2} = 26766274163673319604503, \qquad u_{1/4} = 41595378994516821279015,$$
$$u_{1/8} = 24737346889219389259839, \qquad u_{1/16} = 13387194060291799253121.$$

For $F_{n,k}(r) = r^2 - n(3r - (n-k))$, the exact boundary checks are

(PC0)

ρ	$F_{n,k}(B-1)$	$F_{n,k}(B)$
1/2	5154112775168	−663955886271
1/4	7590647904465	−3182321912768
1/8	13908181940112	−6720484728007
1/16	19335616403905	−20973145690236.

Proposition 7.1 (Exact arithmetic and primality certificates)**.** *For the four rows in Table 2, the displayed integer p satisfies*

$$p = u\,2^s + 1, \qquad u \text{ odd}, \qquad u < 2^s, \qquad a_0^{(p-1)/2} \equiv -1 \pmod p. \tag{PC1}$$

Consequently every p is prime. Moreover

$$n \mid p - 1, \qquad p < 2^{256}, \qquad B\,2^{128} \le p < (B+1)2^{128}, \tag{PC2}$$

where n and B are the corresponding entries of the table. They also satisfy $F_{n,k}(B-1) \ge 0 > F_{n,k}(B)$. Since $F'_{n,k}(r) = 2r - 3n < 0$ on $[0, n]$, these two signs locate its smaller root and give $B - 1 = r_\rho(n) = r_{\mathrm{quad}}(n, k)$. Thus $\mathbb{F}_p^*$ contains a subgroup of order n, and every row in Table 2 exists with the asserted slope budget.

Proof. All identities and inequalities in (PC0)–(PC2) are direct integer calculations; the values of (s, a_0), u, p, n, and B have all been printed above so that the calculation is independently reproducible. We recall why the certificate proves primality. If a prime ℓ divides p, then $a_0^{u2^{s-1}} \equiv -1 \pmod \ell$, whereas $a_0^{u2^s} \equiv 1 \pmod \ell$. The multiplicative order of a_0 modulo ℓ therefore has exact 2-part 2^s, so $2^s \mid \ell - 1$. Since $u < 2^s$, $p < 2^{2s} + 1 < (2^s + 1)^2$, so $2^s + 1 > \sqrt{p}$. A composite p would have a prime divisor at most $\sqrt{p}$, whereas every prime divisor is at least $2^s + 1$, a contradiction. This is Proth's theorem in the special form used here. Finally $s \ge \log_2 n$ gives $n \mid p - 1$, and cyclicity of $\mathbb{F}_p^*$ gives the required subgroup. □

Theorem 7.2 (Reed–Solomon endpoint, Jo [Jo26, Thm. 4.3])**.** *Let $C = \mathrm{RS}_{\mathbb{F}_q}(D, k)$, put $M = \binom{n}{k+1}$, and use the full-field challenge. For every $0 \le a \le k + 1$,*

$$\left\lceil \frac{Mq}{q + M - 1} \right\rceil \le B_C^{\mathrm{MCA}}(a) \le \min\{M, q\}. \tag{EP}$$

If $q > \binom{M}{2}$, *then*

$$B_C^{\mathrm{MCA}}(a) = M \qquad (0 \le a \le k+1). \tag{EP$'$}$$

Corollary 7.3 (Exact first adjacent row)**.** *Let* $C = \mathrm{RS}_{\mathbb{F}_q}(D,k)$, *assume* $k \le n-2$, *and put*

$$M = \binom{n}{k+1}, \qquad U = \left\lfloor \frac{M}{k+1} \right\rfloor.$$

Suppose $q > \binom{M}{2}$, *use the full-field challenge, let* $0 \le \varepsilon \le 1$, *put* $b = \lfloor \varepsilon q \rfloor$, *and, when the set is nonempty, define*

$$a^* = \min\{a \in \{k+1,\dots,n\} : B_C^{\mathrm{MCA}}(a) \le b\}.$$

In either of the following cases a^* *is defined, with*

$$\begin{aligned} b \ge M \quad &\Longrightarrow a^* = k+1, \\ U \le b < M \quad &\Longrightarrow a^* = k+2. \end{aligned} \tag{AD}$$

In the first branch the complete official safe set is $[0,1]$*; in the second branch the transition occurs before capacity.*

Proof. By Theorem 7.2, the numerator at $k+1$ is exactly M. At $a = k+2$, Theorem 1.2 with test size $k+1$ gives $B_C^{\mathrm{MCA}}(k+2) \le \lfloor M/(k+1) \rfloor = U$. Monotonicity proves both claims. □

Theorem 7.4 (Subfield confinement)**.** *Let* $\mathbb{B} \subseteq \mathbb{F}$, $D \subseteq \mathbb{B}$, *and let* $u_0, u_1 \in \mathbb{B}^D$. *If* $\gamma \in \mathbb{F} \setminus \mathbb{B}$ *and* $u_0 + \gamma u_1$ *is explained on* S, *then the pair is explained on the same* S. *Hence every MCA-bad slope of a* $\mathbb{B}$*-valued received line lies in* $\mathbb{B}$.

Proof. Choose a $\mathbb{B}$-basis $1, \beta_1, \dots$ of $\mathbb{F}$. Decompose the coefficients of the explaining polynomial and of γ in this basis. Because $D \subseteq \mathbb{B}$, comparison of basis components on S gives a codeword explaining u_1 from any nonzero nonbase component of γ, and then a codeword explaining u_0 from the base component. □

8. Exact locator prefixes and collision-aware pole lines

Assume $D \subseteq \mathbb{B} \subseteq \mathbb{F}$. For an m-set $S \subseteq D$, write

$$Q_S(X) = \prod_{x \in S} (X - x) = X^m + c_1(S)X^{m-1} + \cdots + c_m(S), \qquad \Phi_w(S) = (c_1(S), \dots, c_w(S)).$$

The monic polynomial Q_S, whose simple roots are precisely the points of S, is the *support locator*. Its *depth-w locator prefix* is the vector of the first w coefficients below the leading coefficient. For $z \in \mathbb{B}^w$, the inverse image $\Phi_w^{-1}(z)$ is the *prefix fiber over z*: the family of all m-element supports having that same prefix. Thus "fiber" in this paper always refers to the inverse image of a displayed map.

Proposition 8.1 (Exact prefix-list bijection)**.** *Let* $1 \le K \le m \le n$, *put* $w = m - K$, *and set* $U_z(X) = X^m + \sum_{i=1}^{w} z_i X^{m-i}$ *for* $z \in \mathbb{B}^w$. *The codewords of* $\mathrm{RS}_{\mathbb{F}}(D,K)$ *agreeing with* $U_z|_D$ *on at least* m *points are exactly*

$$(U_z - Q_S)|_D, \qquad S \in \Phi_w^{-1}(z). \tag{4.1}$$

No such codeword agrees on more than m *points. In particular, some received word has list size at least* $\lceil \binom{n}{m} |\mathbb{B}|^{-w} \rceil$.

Proof. For $S \in \Phi_w^{-1}(z)$, the leading term and next w coefficients cancel, so $\deg(U_z - Q_S) < m - w = K$. Conversely, if P, $\deg P < K$, agrees with U_z at m points, then $U_z - P$ is monic of degree m with those m roots and equals their locator. Coefficient comparison gives the prefix. A nonzero degree-m polynomial has no additional root. Pigeonhole the $\binom{n}{m}$ supports over $|\mathbb{B}|^w$ prefixes. □

The *simple-pole conversion* chooses an auxiliary point $\alpha \notin D$ and divides coordinatewise by $X - \alpha$. This turns one list of nearby polynomials into points on a received affine line for the dimension-k code. A *collision* is an equality $P_i(\alpha) = P_j(\alpha)$ for two distinct listed polynomials; the next formula pays such equalities rather than assuming all slopes are distinct.

Theorem 8.2 (Collision-aware simple-pole conversion)**.** *Let $C = \mathrm{RS}_{\mathbb{F}}(D,k)$, $C^+ = \mathrm{RS}_{\mathbb{F}}(D,k+1)$, $1 \le k < n$, and $q = |\mathbb{F}| > n$. If a word has L distinct codewords of C^+ agreeing on at least $m \ge k+1$ points, then one received line for C has at least*

$$M(L) = \left\lceil \frac{L(q-n)}{q-n+k(L-1)} \right\rceil \tag{4.2}$$

distinct MCA-bad slopes at agreement m.

Proof. Let the received word be $U \in \mathbb{F}^D$, and let the listed polynomials be $P_1,\dots,P_L$, each of degree at most k. For $\alpha \in \mathbb{F} \setminus D$, define the word coordinatewise by $f_\alpha(x) = U(x)/(x-\alpha)$ and set $g_\alpha(x) = -1/(x-\alpha)$. On an agreement set for P_i, the line point of slope $P_i(\alpha)$ equals $(P_i(x)-P_i(\alpha))/(x-\alpha)$, a degree-less-than-$k$ evaluation. The pair is not simultaneously explained on $k+1$ points: an explanation of g_α by G, $\deg G < k$, would make $(X-\alpha)G(X)+1$ vanish at $k+1$ points although its degree is at most k and its value at α is 1.

For $i \ne j$, the equality $P_i(\alpha) = P_j(\alpha)$ holds at at most k poles. Some pole has at most $k\binom{L}{2}/(q-n)$ colliding pairs. If its value multiplicities are $m_1,\dots,m_M$, then $\sum m_i^2 \le L + kL(L-1)/(q-n)$. Cauchy–Schwarz gives $L^2 \le M\sum m_i^2$, which yields (4.2). □

Corollary 8.3 (Exact good-pole reservoir)**.** *In Theorem 8.2, let $Q = q-n$, and for $\alpha \in \mathbb{F}\setminus D$ let C_α be the number of unordered pairs $\{i,j\}$ with $P_i(\alpha) = P_j(\alpha)$. For every integer $T \ge 0$, at least*

$$\max\left\{0, Q - \left\lfloor \frac{k\binom{L}{2}}{T+1} \right\rfloor\right\} \tag{PR1}$$

poles satisfy $C_\alpha \le T$, and every such pole produces at least

$$\left\lceil \frac{L^2}{L+2T} \right\rceil \tag{PR2}$$

distinct MCA-bad slopes. In particular at least $\max\{0, Q - k\binom{L}{2}\}$ poles are fully separating.

Proof. Every nonzero difference $P_i - P_j$ has at most k roots, so

$$\sum_{\alpha\in\mathbb{F}\setminus D} C_\alpha \le k\binom{L}{2}.$$

Every pole with $C_\alpha > T$ contributes at least $T+1$ to this sum, which gives (PR1). At a good pole let $m_1,\dots,m_M$ be the multiplicities of the distinct list values. Then $\sum_i m_i = L$ and $\sum_i m_i^2 = L + 2C_\alpha \le L + 2T$. Cauchy–Schwarz yields $M \ge L^2/(L+2T)$, proving (PR2); $T=0$ gives the last assertion. □

Remark 8.4 (One pole, one received line). The reservoir in Corollary 8.3 fixes the same pole when counting collisions, because different poles define different received lines. A cross-pole equation $P_i(\alpha) = P_j(\alpha')$ is one equation in two variables and generically defines a curve; a Bezout degree bound does not turn it into a finite collision set. Such cross-line collisions are neither controlled by the simple-pole argument nor needed by it.

Corollary 8.5 (Identity-prefix MCA floor)**.** *Let $C = \mathrm{RS}_{\mathbb{F}}(D,k)$, where $D \subseteq \mathbb{B} \subseteq \mathbb{F}$, $|D| = n$, and $q = |\mathbb{F}| > n$. For $k+1 \le m \le n$, put $w = m-k-1$ and $L_m = \lceil \binom{n}{m}|\mathbb{B}|^{-w} \rceil$. Then*

$$B_C^{\mathrm{MCA}}(m) \ge M(L_m). \tag{4.3}$$

Proposition 8.6 (Two deployed unsafe certificates)**.** *Put $n = 2^{21}$, $k = 2^{20}$, and define L_a and $M(L_a)$ by Corollary 8.5 and Eq. (4.2). The following full-field rows have exact integer attack evaluations:*

(DP1)

row and target	a_0	B^*	$M(L_{a_0})$	$M(L_{a_0+1})$
$p_{\mathrm{KB}}^6,\ 2^{-128}$	1116047	274980728111395087	138634741058327852652	57198030366
$(2^{31}-1)^4,\ 2^{-100}$	1116023	16777215	4281388998575706	1752700.

Here $p_{\mathrm{KB}} = 2^{31} - 2^{24} + 1$; the first domain is its multiplicative subgroup of order n, and the second is a Mersenne–31 Chebyshev x-domain of order n: the image of a disjoint twin coset in a norm-one torus under $u \mapsto (u + u^{-1})/2$, as formalized in Definition 9.13. Thus the first row is unsafe at radius 981105/2097152, by 8.9777 bits beyond the 128-bit budget, and the second is unsafe at radius 981129/2097152, by 27.9270 bits beyond its meaningful 100-bit budget. The last column says only that this particular attack falls below budget one agreement later; it is not a safety certificate.

Proof. For each row set $w = a - k - 1$ and evaluate

$$L_a = \left\lceil \binom{n}{a} p^{-w} \right\rceil, \qquad M(L_a) = \left\lceil \frac{L_a(q-n)}{q - n + k(L_a - 1)} \right\rceil$$

with $(p, q) = (p_{\mathrm{KB}}, p_{\mathrm{KB}}^6)$ and $(2^{31} - 1, (2^{31} - 1)^4)$, respectively. Exact integer arithmetic gives (DP1); the first two values exceed their exact budgets. Corollary 8.5 proves unsafety, and the radius and bit statements are the displayed integer differences and base-two ratios. □

For equal-size sets, the *Johnson distance* is $d_J(S, T) = |S \setminus T| = |T \setminus S|$. It counts the number of point replacements required to turn one support into the other.

Proposition 8.7 (Prefix rigidity and its packing limit)**.** *Two distinct m-sets in one depth-w prefix fiber have Johnson distance at least $w + 1$. Consequently, with $t = \lfloor w/2 \rfloor$, every fiber has size at most*

$$\frac{\binom{n}{m}}{\sum_{i=0}^{t} \binom{m}{i}\binom{n-m}{i}}. \tag{4.4}$$

At $w \asymp n/\log|\mathbb{B}|$, this packing loss is only $e^{o(n)}$ and does not supply the random factor $|\mathbb{B}|^{-w}$.

Prefix rigidity says that sharing many leading locator coefficients forces distinct supports to be far apart in this metric. Here the *polynomial-field regime* means $|\mathbb{B}| = n^{O(1)}$, equivalently $\log|\mathbb{B}| = O(\log n)$.

Proof. If S, T have Johnson distance $j \le w$, factor the locator of their intersection, of degree $m - j$, from $Q_S - Q_T$. The residual factor is nonzero, so the difference has degree at least $m - j \ge m - w$, whereas the common prefix makes its degree at most $m - w - 1$, a contradiction. Radius t Johnson balls about the fiber members are disjoint and each has the denominator in (4.4). For $w = n/\log|\mathbb{B}|$, the logarithm of that ball volume is $O(w \log(n/w)) = o(n)$ in the polynomial-field regime. □

Theorem 8.8 (Prefix fibers embed in received lines)**.** *Let $\mathcal{G} \subseteq \Phi_w^{-1}(z)$ be a fiber of N many m-sets, with $0 \le w \le m - 2$, and put $k = m - w - 1$. Every finite extension $\mathbb{F}/\mathbb{B}$ satisfying*

$$|\mathbb{F}| - n > k\binom{N}{2} \tag{4.5}$$

admits one received line for $\mathrm{RS}_{\mathbb{F}}(D, k)$ with at least N distinct MCA-bad slopes furnished by $\mathcal{G}$.

Proof. For distinct $S, T \in \mathcal{G}$, the polynomial $Q_S - Q_T$ is nonzero of degree at most k. Condition (4.5) leaves a pole $\alpha \in \mathbb{F} \setminus D$ avoiding the roots of all these differences. Thus the values $Q_S(\alpha)$ are pairwise distinct. By Proposition 8.1, the polynomials $U_z - Q_S$ form one list for the dimension-$k + 1$ code; the pole line makes their values, and hence the resulting MCA slopes, pairwise distinct. □

Theorem 8.9 (Exact list–line bijection at a separating pole)**.** *Let $D \subseteq \mathbb{B} \subseteq \mathbb{F}$, let $1 \le k < n$, let $m \ge k + 1$, and let $U \in \mathbb{B}^D$. Write*

$$\mathcal{L}_m(U) = \{P \in \mathbb{B}[X] : \deg P \le k,\ |\{x \in D : U(x) = P(x)\}| \ge m\}.$$

Suppose $\alpha \in \mathbb{F} \setminus D$ and the evaluation map $P \mapsto P(\alpha)$ is injective on $\mathcal{L}_m(U)$. Define

$$f_\alpha(x) = \frac{U(x)}{x - \alpha}, \qquad g_\alpha(x) = -\frac{1}{x - \alpha} \qquad (x \in D).$$

Then $P \mapsto P(\alpha)$ is a bijection from $\mathcal{L}_m(U)$ onto the set of finite MCA-bad slopes of the received line (f_α, g_α) for $\mathrm{RS}_{\mathbb{F}}(D,k)$ at agreement m. The corresponding explaining polynomial is

$$h_P(X) = \frac{P(X) - P(\alpha)}{X - \alpha},$$

and its agreement set with $f_\alpha + P(\alpha)g_\alpha$ is exactly the agreement set of P with U.

If $L = |\mathcal{L}_m(U)|$, such a pole exists in every finite extension satisfying

$$|\mathbb{F}| > n + k\binom{L}{2}. \tag{4.6}$$

Proof. For $P \in \mathcal{L}_m(U)$, the polynomial h_P has degree less than k, and on every point where $U = P$ one has

$$f_\alpha + P(\alpha)g_\alpha = \frac{P - P(\alpha)}{X - \alpha} = h_P.$$

The word g_α cannot be explained by a polynomial G of degree less than k on $k+1$ points, since then $(X-\alpha)G+1$ would have degree at most k, at least $k+1$ roots, and value one at α. Thus every displayed slope is MCA-bad.

Conversely, suppose $f_\alpha + \gamma g_\alpha$ is explained by $h \in \mathbb{F}[X]_{<k}$ on at least m points. Then $P(X) = (X-\alpha)h(X) + \gamma$ has degree at most k and agrees with U there. Interpolation on any $k+1$ of those points uses a Vandermonde system over $\mathbb{B}$ with right-hand side in $\mathbb{B}$, so $P \in \mathbb{B}[X]$. Hence $P \in \mathcal{L}_m(U)$ and $\gamma = P(\alpha)$. Injectivity proves the bijection, and the same calculation proves equality of the agreement sets.

For distinct $P, Q \in \mathcal{L}_m(U)$, the nonzero polynomial $P-Q$ has degree at most k, and therefore at most k roots. The union of D and the root sets of all $\binom{L}{2}$ differences has size at most $n + k\binom{L}{2}$; a larger extension contains an admissible pole. □

Corollary 8.10 (Exact prefix-fiber ray realization)**.** *Let $0 \le w \le m-2$, put $k = m-w-1$, and let $\mathcal{F}_z = \Phi_w^{-1}(z)$ be a complete depth-w fiber of m-subsets of D. If α separates the values $Q_S(\alpha)$, $S \in \mathcal{F}_z$, then the pole line associated with $U_z = X^m + \sum_{i=1}^{w} z_i X^{m-i}$ has exactly $|\mathcal{F}_z|$ MCA-bad slopes at agreement m. The maps*

$$S \longmapsto U_z(\alpha) - Q_S(\alpha) \longmapsto \left(U_z(\alpha) - Q_S(\alpha), S, \frac{U_z - Q_S - (U_z(\alpha) - Q_S(\alpha))}{X - \alpha}\right)$$

are bijections between the prefix fiber, the bad-slope set, and the exact-m witness incidence. In particular, the complete prefix fiber, its bad-slope set, and its exact-agreement witness incidence have the same cardinality.

Proof. The exact prefix-list bijection identifies $\mathcal{F}_z$ with the complete dimension-$(k+1)$ list for U_z. Apply Theorem 8.9. Since $U_z - (U_z - Q_S) = Q_S$ has precisely the roots in S, every agreement is exactly m and supplies exactly that support. □

9. Further exact constructions

The main threshold theorem needs only the syndrome correspondence and the mean-overlap argument. This section records complementary exact constructions that clarify the support-to-slope projection and the structured domains used in the examples. Every statement is a finite identity, inequality, or construction; no equidistribution or inverse theorem is assumed.

9.1. **Determinantal ray normalization.** The final projection from supports to slopes can be made explicit without postulating a support-pair multiplicity. The construction below applies to every Reed–Solomon row, in every characteristic. It also identifies the exact additional boundary maps that would govern extensions beyond the proved overlap range. No estimate for those maps is used in the main threshold theorem.

The support counts above are stated for agreement supports $S \in \binom{D}{a}$, whereas this subsection uses the complementary error supports $E = D \setminus S \in \binom{D}{n-a}$. Complementation is a bijection. If the domain is partitioned into fibers of a folding map, a fiber containing j agreement positions

contains its remaining positions as errors; exceptional positions are complemented as well. Thus every fiber-occupancy census transfers exactly to the error-support formulation.

Let $C = \mathrm{RS}_{\mathbb{F}}(D,k)$, put $R = n-k$, and use the weighted Vandermonde parity-check columns

$$h_x = \lambda_x(1, x, \dots, x^{R-1})^{\mathsf{T}}, \qquad \lambda_x \neq 0,$$

as in the syndrome–secant normal form. For $E \subseteq D$, write $V_E = \mathrm{Span}\{h_x : x \in E\}$. Fix an exact agreement $a \geq k+1$, and set

$$t = n-a, \qquad d = R-t = a-k. \tag{DR0}$$

In particular, $d-1 = a-k-1$ is precisely the identity locator-prefix depth. For $E \in \binom{D}{t}$, write its error locator in ascending order as

$$Q_E(X) = \prod_{x\in E}(X-x) = \sum_{\ell=0}^{t} q_\ell(E) X^\ell, \qquad q_t(E) = 1.$$

For $y = (\eta_0, \dots, \eta_{R-1}) \in \mathbb{F}^R$, define the *locator-recurrence boundary*

$$\mathcal{A}_E(y) = \left(\sum_{\ell=0}^{t} q_\ell(E)\eta_{j+\ell} \right)_{0\leq j<d} \in \mathbb{F}^d. \tag{DR1}$$

If $f(X) = \sum_{i=0}^{R-1} f_i X^i$, we use the coefficient pairing $\langle y, f\rangle = \sum_{i=0}^{R-1} y_i f_i$.

Lemma 9.1 (Locator recurrence and a secant span)**.** *For every $E \in \binom{D}{t}$ and $y \in \mathbb{F}^R$,*

$$y \in V_E \quad\Longleftrightarrow\quad \mathcal{A}_E(y) = 0. \tag{DR2}$$

Proof. Identify a polynomial of degree less than R with its coefficient vector in $\mathbb{F}^R$. Since $\langle h_x, f\rangle = \lambda_x f(x)$, every coefficient vector of a polynomial in $Q_E\mathbb{F}[X]_{<d}$ annihilates V_E. The MDS column property gives $\dim V_E = t$, while multiplication by the nonzero monic polynomial Q_E gives $\dim Q_E\mathbb{F}[X]_{<d} = d = R-t$. Consequently

$$V_E^{\perp} = \{\mathrm{coeff}(Q_E P) : P \in \mathbb{F}[X]_{<d}\}.$$

Testing against the basis $1, X, \dots, X^{d-1}$ of $\mathbb{F}[X]_{<d}$ gives exactly the d coordinates in (DR1). □

For a syndrome pair $y_0, y_1 \in \mathbb{F}^R$, and a support family $\Omega \subseteq \binom{D}{t}$, define its support-restricted finite bad-slope set by

$$Z_\Omega(y_0,y_1) = \left\{ \gamma \in \mathbb{F} : \begin{array}{l} \text{there is } E\in\Omega \text{ with } y_0+\gamma y_1 \in V_E, \\ \{y_0,y_1\} \not\subseteq V_E \end{array} \right\}. \tag{DR3}$$

For $E \in \Omega$, abbreviate

$$u_E = \mathcal{A}_E(y_0), \qquad v_E = \mathcal{A}_E(y_1).$$

Let $\mathcal{P}_j(y_1)$ be the universal pivot locus on which j is the first nonzero coordinate of v_E, and put $\Omega_j = \Omega \cap \mathcal{P}_j(y_1)$:

$$\begin{aligned} \mathcal{P}_j(y_1) &= \{E \in \tbinom{D}{t} : v_{E,i} = 0 \text{ for } 0 \leq i < j, \quad v_{E,j} \neq 0\}, \\ \Omega_j &= \Omega \cap \mathcal{P}_j(y_1), \qquad\qquad 0 \leq j < d. \end{aligned} \tag{DR4}$$

On this chart define the *determinantal ray boundary*

$$\Delta_j(E) = (v_{E,j}u_{E,\ell} - u_{E,j}v_{E,\ell})_{\substack{0\leq \ell<d \\ \ell\neq j}} \in \mathbb{F}^{d-1}. \tag{DR5}$$

When $d = 1$, the target in (DR5) is the trivial group and the displayed tuple is empty.

Theorem 9.2 (Exact determinantal ray compiler)**.** *With the notation above, an error support $E \in \Omega$ supplies a finite transverse slope if and only if*

$$v_E \neq 0, \qquad \mathrm{rank}\begin{pmatrix} u_E \\ v_E \end{pmatrix} \leq 1. \tag{DR6}$$

On the pivot chart Ω_j, this is equivalent to $\Delta_j(E) = 0$, and its unique slope is $\gamma_E = -u_{E,j}/v_{E,j}$. Consequently

$$Z_\Omega(y_0, y_1) = \bigcup_{j=0}^{d-1} \left\{ -\frac{u_{E,j}}{v_{E,j}} : E \in \Omega_j,\ \Delta_j(E) = 0 \right\}, \tag{DR7}$$

and

$$|Z_\Omega(y_0, y_1)| \le \sum_{j=0}^{d-1} |\Omega_j \cap \Delta_j^{-1}(0)|. \tag{DRC}$$

The construction uses $d \le n$ pivot charts and contains no unnamed finite or asymptotic multiplicity.

Proof. By Lemma 9.1, E supports the line parameter γ precisely when

$$u_E + \gamma v_E = 0. \tag{DR8}$$

If $v_E = 0$, a solution of (DR8) forces $u_E = 0$, so $y_0, y_1 \in V_E$ and the intersection is common rather than transverse. If $v_E \neq 0$, equation (DR8) has a solution if and only if u_E is a scalar multiple of v_E, which is (DR6); the scalar, and therefore the finite slope, is unique. The charts Ω_j partition the locus $v_E \neq 0$. On Ω_j, vanishing of all 2-by-2 minors using the pivot column j is equivalent to proportionality, and the proportionality factor is $-u_{E,j}/v_{E,j}$. This proves (DR7). Taking the cardinality of a union and then forgetting the slope projection gives (DRC). □

Proposition 9.3 (Exact slope-incidence linearization)**.** *With the notation of Theorem 9.2, put*

$$\mathcal{I}_\Omega = \{(\gamma, E) \in \mathbb{F} \times \Omega : u_E + \gamma v_E = 0,\ v_E \neq 0\}. \tag{DRC1}$$

Projection $\mathcal{I}_\Omega \to \Omega$ is injective, its slope projection is $Z_\Omega(y_0, y_1)$, and

$$|Z_\Omega(y_0, y_1)| \le |\mathcal{I}_\Omega| = \sum_{\gamma \in \mathbb{F}} |\{E \in \Omega : u_E + \gamma v_E = 0,\ v_E \neq 0\}|. \tag{DRC2}$$

Thus a fixed-slope zero-fiber estimate must be aggregated over all slopes. If every fixed-slope zero fiber were bounded by $|\Omega|/q^d$, summing over the q slopes would give $|\Omega|/q^{d-1}$, not $|\Omega|/q^d$.

Proof. For $v_E \neq 0$, the vector equation determines at most one γ, proving injectivity. The exact determinantal compiler says that its slope image is precisely the transverse bad-slope set. Counting the incidence first by E and then by γ gives (DRC2); the final scale is the corresponding uniform fiber calculation. □

Proposition 9.4 (The recurrence boundary is not pointwise additive)**.** *The exact proportionality condition above does not, by itself, turn the locator-recurrence boundary into a subset-sum map. Already for $E = \{x, y\}$, one coordinate of $\mathcal{A}_E(\eta)$ is*

$$F(x, y) = \eta_j xy - \eta_{j+1}(x + y) + \eta_{j+2}. \tag{DRC3}$$

If $\eta_j \neq 0$, this function cannot be written as $c+g(x)+g(y)$ whenever the four admissible two-sets $\{x, y\}, \{x', y\}, \{x, y'\}, \{x', y'\}$ lie in the chart and $(x - x')(y - y') \neq 0$. Consequently, a proposed Fourier reduction to additive maps $E \mapsto \sum_{t \in E} g_\gamma(t)$ requires a separate coordinate theorem; it is not a consequence of slope linearization.

Proof. For a two-set, $Q_E(X) = X^2 - (x + y)X + xy$, which gives (DRC3). Its mixed difference is

$$F(x, y) - F(x', y) - F(x, y') + F(x', y') = \eta_j (x - x')(y - y').$$

Every function $c+g(x)+g(y)$ has zero mixed difference, proving the claim whenever the displayed product is nonzero. □

The maps in (DR5) are explicit quadratic maps in the locator coefficients. They depend on the received syndrome pair, as they must in a worst-line theorem, but their number and degree are literal. The next corollary shows how they replace the abstract ray interface.

Theorem 9.5 (Fixed-core determinacy ray compiler)**.** *Let $C_0 \subseteq D$ have size $c \le k$, let $a \ge k+1$, and fix a received line $r = (r_0, r_1)$. Suppose that for every γ in a set Z of distinct finite MCA-bad slopes we have chosen an exact-a noncommon witness $(\gamma, S_\gamma, h_\gamma)$ with $C_0 \subseteq S_\gamma$. Then*

$$|Z| \le \binom{n-c}{k-c+1}. \tag{$\mathrm{RC}_{\mathrm{core}}$}$$

In particular, a fixed-core cell with $k - c = o(n)$ has a direct subexponential ray payment. This statement is uniform in the received line and valid in every characteristic.

Proof. Put $d = k - c$, let $Q_{C_0}(X) = \prod_{x \in C_0}(X - x)$, and choose $g_0, g_1 \in \mathbb{F}[X]_{<k}$ interpolating r_0, r_1, respectively, on C_0. Write $e_i = r_i - \mathrm{ev}_D(g_i)$. Since $h_\gamma - g_0 - \gamma g_1$ vanishes on C_0, there is a unique $P_\gamma \in \mathbb{F}[X]_{<d}$ such that

$$h_\gamma = g_0 + \gamma g_1 + Q_{C_0} P_\gamma.$$

For $x \in D \setminus C_0$, the witness equation is

$$e_0(x) + \gamma e_1(x) - Q_{C_0}(x) P_\gamma(x) = 0. \tag{$*$}$$

Homogenize the parameters as $[z : \gamma : p_0 : \cdots : p_{d-1}] \in \mathbb{P}^{d+1}$, and let the row indexed by $x \in D \setminus C_0$ be the linear functional

$$z e_0(x) + \gamma e_1(x) - Q_{C_0}(x) \sum_{j=0}^{d-1} p_j x^j.$$

The rows indexed by $S_\gamma \setminus C_0$ have rank exactly $d+1$. Indeed, the displayed witness parameter is a nonzero kernel vector, so the rank is at most $d+1$. If the rank were at most d, the kernel would contain an independent vector. Subtracting a multiple of the witness vector gives a nonzero kernel vector with $z = 0$. Its γ-coordinate cannot vanish: otherwise a nonzero polynomial of degree less than d would vanish at all $a - c \ge d+1$ points of $S_\gamma \setminus C_0$. Adding a nonzero multiple of this kernel vector to the witness vector therefore produces a second finite slope explained on the same support. By fixed-support slope uniqueness that support would be common, contrary to the witness hypothesis.

Choose, in a fixed order, the first $(d+1)$-subset $T_\gamma \subseteq S_\gamma \setminus C_0$ on which the rows have rank $d+1$. Their common projective kernel is the single witness parameter, so T_γ determines γ. Hence $\gamma \mapsto T_\gamma$ is injective into $\binom{D \setminus C_0}{d+1}$, proving the bound. If $d = o(n)$, then $\log \binom{n-c}{d+1} = o(n)$. □

Remark 9.6 (Cores are internal unless their census is small)*.* The theorem pays one fixed core. It does not permit an atlas to use every positive-density subset C_0 as a separate profile: that census can be exponential. A complete atlas either routes only an algebraically specified subexponential family of cores or keeps the actual core internal to a coarse intersection-size profile.

9.2. **Sharpness of support-only ray bounds.** A support S in a multiplicative group D is *multiplicatively aperiodic* when its stabilizer $\{g \in D : gS = S\}$ is trivial.

Theorem 9.7 (Aperiodic ray obstruction below the profile frontier)**.** *There is a sequence of smooth multiplicative rows*

$$D_n = \mathbb{B}_n^\times \subseteq \mathbb{F}_n^\times, \qquad \log|\mathbb{B}_n| = O(\log n) = o(n),$$

in characteristic two, with rates tending to $1/2$, and agreements $m_n = k_n + O(n/\log n)$, for which one received line has $2^{\Omega(n)}$ distinct finite MCA-bad slopes. Every threshold- m_n witness support of that line has trivial multiplicative stabilizer, and after enlarging $\mathbb{F}_n$ with only $\log|\mathbb{F}_n| = O(n)$, the bad slopes and exact-agreement supports are in bijection. Thus smoothness, quotient-aperiodicity, or one-parameter support-profile data, and exact agreement do not force retained-support occupancy greater than one. This is a sharpness result for occupancy-based compilers, not a counterexample to a direct distinct-slope theorem at the correctly normalized profile scale.

Proof. Let $q_r = 2^r$, put $\mathbb{B}_r = \mathbb{F}_{q_r}$, $D_r = \mathbb{B}_r^\times$, and $n_r = q_r - 1$. Set

$$m_r = 2^{r-1} - 1 = \frac{n_r - 1}{2}.$$

Fix $0 < \eta < \log 2$, take $w_r = \lfloor \eta n_r / \log q_r \rfloor$, and put $k_r = m_r - w_r - 1$. Then $k_r/n_r \to 1/2$ and $m_r - k_r = w_r + 1 = O(n_r/\log n_r)$.

Pigeonholing the m_r-set locators over their $q_r^{w_r}$ prefixes gives a fiber $\mathcal{F}_r$ satisfying

$$\log|\mathcal{F}_r| \geq \log \binom{n_r}{m_r} - w_r \log q_r \geq (\log 2 - \eta) n_r - o(n_r).$$

Moreover

$$\gcd(m_r, n_r) = \gcd(2^{r-1} - 1, 2^r - 1) = 2^{\gcd(r-1,r)} - 1 = 1.$$

If an m_r-support were invariant under a nontrivial subgroup of D_r of order $c > 1$, it would be a union of c-element orbits, forcing $c \mid m_r$ and $c \mid n_r$, a contradiction. Thus every member of $\mathcal{F}_r$ is aperiodic.

Choose a finite extension $\mathbb{F}_r/\mathbb{B}_r$ with

$$|\mathbb{F}_r| > n_r + k_r \binom{|\mathcal{F}_r|}{2}.$$

Since $|\mathcal{F}_r| \leq 2^{n_r}$, the least extension degree meeting this inequality has $\log|\mathbb{F}_r| = O(n_r)$. By Corollary 8.10, a separating pole gives exactly $|\mathcal{F}_r| = \exp(\Omega(n_r))$ distinct bad slopes, one for each aperiodic support, and every occupancy equals one. □

Remark 9.8 (Scope of the aperiodic obstruction). The construction lies below the entropy–profile crossing: its natural prefix scale is itself exponential. It therefore does not contradict a profile-envelope upper bound proved at the correct scale. It does rule out a field-independent polynomial ray bound, division of a one-parameter support count by the full support-slice size without a retained multiplicity theorem, and the assertion that quotient-aperiodicity alone forces ray collisions.

The opposite projection extreme occurs over the coefficient field itself. It shows that a large aperiodic prefix fiber need not supply a lower reserve after first-match saturation.

Theorem 9.9 (Aperiodic one-ray saturation)**.** *For an integer $\nu \geq 5$, put*

$$q = 2^\nu, \quad D = \mathbb{F}_q^\times, \quad n = q - 1, \quad m = 2^{\nu-1} - 1, \quad w = \left\lfloor n/\nu^2 \right\rfloor, \quad k = m - w - 1.$$

Refine the depth-w prefix by the locator constant coefficient:

$$\Xi_\nu(S) = (c_1(S), \ldots, c_w(S), c_m(S)) \qquad \left(S \in \binom{D}{m}\right).$$

There are $z \in \mathbb{F}_q^w$ and $c \in \mathbb{F}_q^\times$ for which

$$\text{(SAT1)} \qquad \mathcal{G}_{z,c} = \{S : \Xi_\nu(S) = (z, c)\}, \qquad |\mathcal{G}_{z,c}| \geq \left\lceil \binom{n}{m} q^{-(w+1)} \right\rceil = \exp((\log 2 - o(1))n).$$

Every support in $\mathcal{G}_{z,c}$ has trivial multiplicative stabilizer, but all witnesses arising from the exact prefix list project at the pole $\alpha = 0$ in the coefficient field $\mathbb{F}_q$ to the single MCA slope $-c$ for $\mathrm{RS}_{\mathbb{F}_q}(D, k)$.

More precisely, if $\mathcal{F}_z = \Phi_w^{-1}(z)$ and $\mathcal{C}_d = \{S \in \mathcal{F}_z : c_m(S) = d\}$, then the nonempty $\mathcal{C}_d$, $d \in \mathbb{F}_q^\times$, partition the supports of all exact-m witnesses of this pole line, and the witnesses supported on $\mathcal{C}_d$ have slope image exactly $\{-d\}$. Thus at most $q - 1 = \exp(o(n))$ explicit saturation cells pay the complete line even though one cell contains exponentially many aperiodic supports.

Proof. The map Ξ_ν has at most q^{w+1} values, so pigeonholing proves the first inequality in (SAT1). Moreover

$$\log \binom{n}{m} = n \log 2 - O(\log n), \qquad (w+1)\log q \le (n/\nu + \nu)\log 2 = o(n),$$

which proves its asymptotic form. Also

$$\gcd(m,n) = \gcd(2^{\nu-1}-1, 2^\nu - 1) = 2^{\gcd(\nu-1,\nu)} - 1 = 1.$$

An invariant support would be a union of orbits of a nontrivial subgroup whose order divides both m and n, so every m-support is multiplicatively aperiodic.

Let $U_z = X^m + \sum_{i=1}^{w} z_i X^{m-i}$ and $P_S = U_z - Q_S$. Prefix cancellation gives $\deg P_S \le k$, and $U_z - P_S = Q_S$ shows that P_S agrees with U_z exactly on S. At the only pole of the coefficient field outside D, namely zero,

$$P_S(0) = U_z(0) - Q_S(0) = -c_m(S).$$

The pole-line calculation therefore maps every $\mathcal{C}_d$ to $-d$. Conversely, an exact-m explanation on this line gives a polynomial $P = Xh + \gamma$ of degree at most k agreeing with U_z on m points. The exact prefix-list bijection makes it the unique $P_S = U_z - Q_S$ for some $S \in \mathcal{F}_z$, and then $\gamma = P(0) = -c_m(S)$. This proves coverage and the asserted one-slope images. □

Remark 9.10 (Two incompatible aperiodic inferences). Theorem 9.7 makes a separating extension pole inject an exponentially large aperiodic fiber into distinct slopes, while Theorem 9.9 makes a coefficient-field pole collapse such a refined fiber to one slope. Aperiodicity alone therefore implies neither ray expansion nor ray collision. A lower compiler needs a separating-pole estimate or a quantitative exclusion of earlier saturation cells; an upper compiler needs an actual determinant or ray-image bound.

9.3. **Smooth and circle evaluation domains.**

Characteristic two. The support reduction, tangent construction, MDS paving bound, and all-radius saturation theorem use only linear algebra and the MDS property, and therefore hold over every finite field. The domain transports require more care: $|\mathbb{F}_{2^s}^\times| = 2^s - 1$ is odd, so characteristic two has no nontrivial power-of-two multiplicative subgroup, while $x^2 + y^2 - 1 = (x + y - 1)^2$ is singular, so the circle equivalences used here are necessarily odd-characteristic statements. Characteristic-two analogues based on additive FFT domains or a nonsingular conic would require separate agreement-preserving maps. Thus the representation-free MDS results apply to characteristic-two Reed–Solomon rows, but no multiplicative-subgroup or circle transfer in that characteristic is asserted.

This subsection records the domain geometry behind the exact quotient constructions. The multiplicative and circle cases are treated simultaneously through uniform-fiber folding maps. The point is not merely that such maps exist, but that their fibers are complete inside the evaluation domain and their locators have a *lacunary* top part, meaning that prescribed gaps occur among the highest-degree coefficients. These two facts make quotient families visible in the witness moduli and keep their locator prefixes stable under planted cores.

Definition 9.11 (Complete-fiber folding map and smooth multiplicative coset)**.** *Let $D \subseteq \mathbb{B}$ be finite. A polynomial map $\varphi : D \to Q := \varphi(D)$ is a* c-fold complete-fiber folding map *if every $y \in Q$ has exactly c preimages in D. A support is* complete at this folding scale *if it is a union of whole fibers, equivalently if it is $\varphi^{-1}(E)$ for some $E \subseteq Q$. Passing from D to Q, and from a complete support to E, is the corresponding* quotient *or* folding.

A multiplicative coset *is a set $D = \theta H \subseteq \mathbb{B}^\times$, where $H \le \mathbb{B}^\times$ is a subgroup and $\theta \in \mathbb{B}^\times$. For every divisor $c \mid |H|$, the restriction*

$$\pi_c : D \longrightarrow D^{(c)}, \qquad \pi_c(x) = x^c, \qquad D^{(c)} = \{x^c : x \in D\},$$

is a c-fold complete-fiber folding map. We call D smooth at scale c *when this power folding is among the quotients retained by the row, and* smooth *for a family of scales when it is smooth at every scale in that family. Thus the adjective* smooth *here is an explicit uniform-fiber property,*

not a claim about nonsingularity of an algebraic variety. The locator of a complete support is $V_E(X^c)$; its nonleading terms occur only in degree gaps divisible by c. This controlled pattern of missing top coefficients is what lacunary *means in this section.*

Definition 9.12 (Circle domain and circle code)**.** *Let $\mathbb{B}_0$ have odd characteristic and let $\mathcal{C}/\mathbb{B}_0$ be the affine conic $\mathcal{C} : x^2 + y^2 = 1$. A* standard circle domain *is a finite set $E \subseteq \mathcal{C}(\mathbb{B}_0)$. If $\mathbb{F} \supseteq \mathbb{B}_0$, the degree-$w$* circle code *on E is*

$$\mathcal{C}_w(\mathbb{F}, E) = \{(f(P))_{P\in E} : f \in \operatorname{im}(\mathbb{F}[x, y]_{\le w} \to \mathbb{F}[x, y]/(x^2 + y^2 - 1))\}.$$

The degree is the total-degree filtration before passage to the quotient ring. A circle x-domain *is instead a set of first coordinates obtained from the torus construction below; an RS code on that x-domain and a bivariate circle code are different presentations until a diagonal equivalence such as Lemma 9.18 has been exhibited.*

Definition 9.13 (Norm-one torus, twin coset, and Chebyshev domain)**.** *Following the circle-domain geometry of* [HLP24]*, let $p \equiv 3 \pmod 4$, choose $i \in \mathbb{F}_{p^2}$ with $i^2 = -1$, and write*

$$\mathbb{U} = \ker\Big(N_{\mathbb{F}_{p^2}/\mathbb{F}_p} : \mathbb{F}_{p^2}^\times \to \mathbb{F}_p^\times\Big) = \{u \in \mathbb{F}_{p^2}^\times : u^{p+1} = 1\}.$$

The algebraic kernel of the norm morphism is the norm-one torus*; $\mathbb{U}$ is its cyclic group of $\mathbb{F}_p$-rational points. Put $\chi(u) = (u + u^{-1})/2 \in \mathbb{F}_p$. If $H \le \mathbb{U}$ and $g \in \mathbb{U}$ satisfies $g^2 \notin H$, then*

$$\mathcal{D}(g, H) = gH \sqcup g^{-1}H$$

is a disjoint union exchanged by $u \mapsto u^{-1}$; it is called a twin coset*. Its projected set*

$$D(g, H) = \chi(\mathcal{D}(g, H)) \subseteq \mathbb{F}_p$$

has cardinality $|H|$ and is the Chebyshev twin-coset x-domain*.*

For $c \ge 1$, let T_c be the Chebyshev polynomial normalized by

$$T_c\left(\frac{u + u^{-1}}{2}\right) = \frac{u^c + u^{-c}}{2}.$$

Thus $T_c \circ \chi = \chi \circ (u \mapsto u^c)$. Whenever $c \mid |H|$ and the two image cosets remain disjoint, the restriction of T_c to $D(g, H)$ is the Chebyshev folding map*; a Chebyshev twin-coset domain equipped with these complete-fiber maps is the circle analogue of a smooth multiplicative coset.*

Definition 9.14 (Map-smooth domain)**.** *Let $\mathbb{B} \subseteq \mathbb{F}$ be finite fields, let $\varphi \in \mathbb{B}[X]$ have degree $c \ge 1$, and let $D \subseteq \mathbb{B}$ have size n. The set D is (φ, c)-smooth if every nonempty fiber $\varphi^{-1}(y) \cap D$, $y \in \varphi(D)$, has exactly c elements. We write $Q = \varphi(D)$ and $N = n/c$.*

A multiplicative coset $D = \theta H \subseteq \mathbb{B}^\times$ is (X^a, a)-smooth for every divisor $a \mid |H|$. The circle case uses Chebyshev maps rather than power maps; see Lemma 9.17. The following lemma is the common quotient-support construction. It is included because it produces explicit large lists and because the same lacunarity controls their locator prefixes.

Lemma 9.15 (Graded complete-fiber prefix)**.** *Let D be (φ, c)-smooth, put $Q = \varphi(D)$ and $N = |Q| = n/c$, and let $C^+ = \mathrm{RS}_{\mathbb{F}}(D, K)$. For an integer*

$$d = \lceil K/c \rceil \le m \le N, \qquad w = m - d,$$

there is a word $u \in \mathbb{B}^D$ having at least

$$L = \left\lceil \binom{N}{m} |\mathbb{B}|^{-w} \right\rceil \tag{MS1}$$

distinct codewords of C^+ that agree with it on mc coordinates.

Proof. For each m-set $E \subseteq Q$, write

$$P_E(Y) = \prod_{b\in E} (Y - b) = Y^m + \sum_{j=1}^{w} z_j(E) Y^{m-j} + R_E(Y), \qquad \deg R_E \le d - 1.$$

Pigeonholing the $\binom{N}{m}$ sets by their vector $(z_1(E),\dots,z_w(E)) \in \mathbb{B}^w$ gives a fiber of size at least L; fix its value z. The word $u_z = \varphi^m + \sum_{j=1}^{w} z_j \varphi^{m-j}$ agrees on the mc points above E with the codeword $-R_E(\varphi)$. This is a codeword because $\deg R_E(\varphi) \le c(d-1) < K$. Distinct E give distinct codewords: each composed remainder has degree less than $K \le n$, so equality on the n-point domain first implies equality as polynomials. Composition with the nonconstant φ is injective, and the fixed leading prefix then gives $P_E = P_{E'}$. □

Proposition 9.16 (Collision-aware map-smooth attack)**.** *Let D be (φ, c)-smooth, let $q = |\mathbb{F}| > n$, and put $C = \mathrm{RS}_{\mathbb{F}}(D,k)$. For*

$$d = \lceil (k+1)/c \rceil \le m \le N, \quad w = m - d, \quad L = \left\lceil \binom{N}{m} |\mathbb{B}|^{-w} \right\rceil,$$

every agreement $k+1 \le A \le mc$ satisfies

$$B_C^{\mathrm{MCA}}(A) \ge \left\lceil \frac{L(q-n)}{q-n+k(L-1)} \right\rceil. \tag{MS2}$$

In particular, take $m = d+1$. If

$$\binom{N}{m} \ge |\mathbb{B}|\,(q/k+1), \tag{MS3}$$

then at every official radius $1 - mc/n \le \delta \le 1$,

$$e_{\mathrm{MCA}}(C,\delta) > \frac{1}{2k}\left(1 - \frac{n}{q}\right). \tag{MS4}$$

Proof. Apply Lemma 9.15 with $K = k+1$, then apply the collision-aware pole conversion Theorem 8.2. Its received pair is not simultaneously explained on any $k+1$ points, so the same mc-agreement list witnesses every displayed threshold. For the last assertion, (MS3) gives $L \ge q/k + 1$. With $x = q - n$,

$$\frac{Lx}{x + k(L-1)} > \frac{x}{2k},$$

because the cross-multiplied inequality is $kL + k > x$, which follows from $kL + k \ge q + 2k > q - n$. Divide by q. □

We now verify the complete-fiber property for the construction in Definition 9.13. Thus $\mathcal{D} = gH \sqcup g^{-1}H \subseteq \mathbb{U}$, $D = \chi(\mathcal{D})$, and $T_a(\chi(u)) = \chi(u^a)$.

Lemma 9.17 (Chebyshev smoothness)**.** *Let $\mathcal{D} = gH \cup g^{-1}H \subseteq \mathbb{U}$ be a twin coset, let $D = \chi(\mathcal{D})$, and let $a \mid |H|$. If $g^{2a} \notin H^{(a)} = \{h^a : h \in H\}$, then D is (T_a, a)-smooth. If $\mathrm{ord}(g) \nmid 2|H|$, this holds for all $a \mid |H|$.*

Proof. The map $u \mapsto u^a$ has kernel of size a on $\mathbb{U}$, and that kernel lies in H. Hence it maps gH and $g^{-1}H$ onto $g^a H^{(a)}$ and $g^{-a}H^{(a)}$, respectively, with fibers of size a. The hypothesis says these image cosets are disjoint, so the image is again a twin coset. Since χ is exactly two-to-one on both twin cosets and identifies only inverse points, the equation $T_a(x) = y$ on D has exactly a solutions for each $y \in T_a(D)$. Taking $a = |H|$ shows that the all-scale condition is equivalent to $g^{2|H|} \ne 1$, and the implication to every smaller divisor follows by raising to $|H|/a$. □

Lemma 9.18 (Diagonal uniformization of circle codes)**.** *Assume $\mathbb{F} \supseteq \mathbb{F}_{p^2}$. Let $E \subseteq \{(x,y) \in \mathbb{F}_p^2 : x^2 + y^2 = 1\}$, define $\iota(x,y) = x + iy$, put $E' = \iota(E) \subseteq \mathbb{U}$, and let $w \ge 0$ satisfy $2w+1 < |E|$. Let $\mathcal{C}_w(\mathbb{F}, E)$ be the circle code of Definition 9.12. Under the coordinate bijection ι and the diagonal multiplier $t_u = u^{-w}$, the code $\mathcal{C}_w(\mathbb{F}, E)$ is $t \circ \mathrm{RS}_{\mathbb{F}}(E', 2w+1)$. Consequently list sizes, ordinary correlated-agreement (CA) numerators, and mutual-correlated-agreement (MCA) numerators are preserved.*

Proof. The map $u = x + iy$ identifies $\mathbb{F}[x,y]/(x^2+y^2-1)$ with $\mathbb{F}[u,u^{-1}]$. The subspace of total degree at most w maps to $u^{-w}\mathbb{F}[u]_{\le 2w}$, and both spaces have dimension $2w+1$. Multiplication by the nonzero diagonal vector $(u^{-w})_{u \in E'}$ preserves coordinatewise equality and hence every agreement notion used here. Therefore it preserves list, CA, and MCA counts exactly. □

The preceding uniformization is convenient when i belongs to the challenge field. The following form works over every odd-characteristic field and will also be used for scalar extensions of odd degree.

Lemma 9.19 (Invariance under monomial equivalence)**.** *Let $C \subseteq \mathbb{F}^D$, let $\pi : D' \to D$ be a bijection, and let $(d_x)_{x\in D'} \in (\mathbb{F}^\times)^{D'}$. Define*

$$C' = \{(d_x c_{\pi(x)})_{x\in D'} : c \in C\}.$$

The same coordinate permutation and diagonal multiplication give a bijection between received pairs for C and C', preserving the bad-slope set on every support and for every challenge set. In particular $B^{\mathrm{MCA}}_{C,\Gamma}(a) = B^{\mathrm{MCA}}_{C',\Gamma}(a)$.

Proof. For every word u, codeword c, and support S, equality $u|_S = c|_S$ is equivalent after the displayed coordinate change to equality on $\pi^{-1}(S)$. The transformation is $\mathbb{F}$-linear, so it commutes with $r_0 + \gamma r_1$. It therefore preserves both the single-line-point explanation and simultaneous explanation of the pair on the corresponding support, slope by slope. □

Corollary 9.20 (Generalized Reed–Solomon rows)**.** *Every generalized Reed–Solomon code obtained from $\mathrm{RS}_{\mathbb{F}}(D,k)$ by a coordinate permutation and nonzero column multipliers has exactly the same CA and MCA numerators as that RS code. In particular all exact staircases and target compilers in this paper transfer verbatim.*

Proof. This is Lemma 9.19 applied to the defining monomial equivalence. □

Proposition 9.21 (Stereographic circle equivalence)**.** *Let $\mathbb{F}$ have odd characteristic, let $E \subseteq \{(x,y) \in \mathbb{F}^2 : x^2 + y^2 = 1\}$ avoid $(-1,0)$, and put*

$$t = \frac{y}{1+x}, \qquad T = t(E).$$

If $2w + 1 \le |E|$, then

$$\mathcal{C}_w(\mathbb{F}, E) = \mathrm{diag}\big((1+t^2)^{-w}\big)_{t\in T}\,\mathrm{RS}_{\mathbb{F}}(T, 2w+1) \tag{SC1}$$

after the coordinate bijection $E \to T$. Hence every list, CA, and support-wise MCA bad-slope set is preserved exactly.

Proof. The inverse parametrization is

$$x = \frac{1-t^2}{1+t^2}, \qquad y = \frac{2t}{1+t^2},$$

and $1 + t^2 = 2/(1+x) \neq 0$ on E. Modulo the circle equation, the degree-at-most-w filtered space is

$$\{f_0(x) + y f_1(x) : \deg f_0 \le w,\ \deg f_1 \le w-1\},$$

with dimension $2w+1$. After substitution and multiplication by $(1+t^2)^w$, every such function becomes a polynomial of degree at most $2w$. The substitution is injective in the localized coordinate ring and the dimensions agree, so all of $\mathbb{F}[t]_{\le 2w}$ occurs. Equation (SC1) follows, and Lemma 9.19 proves the last assertion. □

Corollary 9.22 (Exact circle staircase)**.** *Let E, w be as in Proposition 9.21, put $n = |E|$ and $k = 2w + 1 < n$. For every nonempty $\Gamma \subseteq \mathbb{F}$ and every $0 \le r \le n-k-1$,*

$$(n-r)^2 \ge n(k+r) \quad\Longrightarrow\quad B^{\mathrm{MCA}}_{\mathcal{C}_w(\mathbb{F},E),\Gamma}(n-r) = \min\{|\Gamma|, r+1\}. \tag{SC2}$$

The same equality holds under the hypotheses of Theorem 5.2. Consequently standard circle codes admitting such a projection chart have the positive-relative-radius staircase of Corollary 5.4.

Proof. Apply Lemma 9.19 and then Theorems 4.6 and 5.2 to the RS presentation in (SC1). □

Proposition 9.23 (A 128-bit circle certificate)**.** *Let $p_0 = 2^{127} - 1$, $q = p_0^2$, and work over $\mathbb{F}_q$. Then*

$$\left\lfloor q 2^{-128} \right\rfloor = 85070591730234615865843651857942052863. \tag{CP1}$$

There are two length-512 circle presentations certified beyond Johnson:
(CP2)

presentation	k	a	$(512-a)/512$	$U_{512,k,\mathbb{F}_q}(a)$
Chebyshev x-domain RS	32	53	459/512	81136417887908025931235146974888066266
degree-15 circle code	31	50	462/512	60827291905480389917403158602781524185.

Both upper numerators are at most (CP1)*, giving* 128.068311 *and* 128.483942 *proved bits, respectively.*

Proof. The Lucas–Lehmer criterion [Rie94, Sec. 4.2] proves that p_0 is prime, and its norm-one torus $\mathbb{U}$ over $\mathbb{F}_{p_0}$ has order $p_0+1=2^{127}$. Choose a generator z, put $H=\langle z^{2^{118}}\rangle$, and set $g=z$. Then $|H|=512$ and $g^2\notin H$. For the first row use the projection of the twin coset $gH\sqcup g^{-1}H$ and Lemma 9.17; for the second take the 512 circle points $E=\iota^{-1}(gH)$, where $\iota(x,y)=x+iy$, and use Lemma 9.18. The two displayed bounds are exactly $\left\lfloor \binom{512}{33}/\binom{52}{32}\right\rfloor$ and $\left\lfloor \binom{512}{32}/\binom{49}{31}\right\rfloor$. Their comparison with the exact budget proves safety. Both radii exceed their Johnson radii, and direct base-two logarithms give the printed bits. □

For a set $\mathcal{P}$ of primes, its *relative prime density* is

$$\operatorname{dens}_{\mathrm{pr}}(\mathcal{P})=\lim_{X\to\infty}\frac{\#\{p\le X : p\in\mathcal{P}\}}{\pi(X)},$$

when the limit exists. The prime number theorem in arithmetic progressions gives density $1/\varphi(M)$ to every reduced residue class modulo M [Dav00]. We write $\overline{\operatorname{dens}}_{\mathrm{pr}}$ for the same expression with lim sup in place of the limit.

Theorem 9.24 (Fixed-length positive-density prime families)**.** *Fix $n=2^m$, $m\ge 2$.*

(a) *The primes $p\equiv 1 \pmod n$, of relative prime density $1/\varphi(n)=2/n$, support the unique subgroup $H\le\mathbb{F}_p^\times$ of order n. Every exact RS conclusion of Theorems 4.6 and 5.2 holds on H and on each of its multiplicative cosets.*

(b) *The primes $p\equiv -1 \pmod{2n}$, of density $1/\varphi(2n)=1/n$, support a standard circle coset E of size n. For every odd $k=2w+1<n$, the circle code $\mathcal{C}_w(\mathbb{F}_{p^2},E)$ has both exact staircases.*

(c) *The primes $p\equiv -1 \pmod{4n}$, of density $1/\varphi(4n)=1/(2n)$, support a Chebyshev-smooth circle x-domain $D\subseteq\mathbb{F}_p$ of size n, and $\mathrm{RS}_{\mathbb{F}_p}(D,k)$ has both exact staircases for every $k<n$.*

Proof. Part (a) follows from cyclicity of $\mathbb{F}_p^\times$. For part (b), let $\mathbb{U}\le\mathbb{F}_{p^2}^\times$ be the norm-one group, of order $p+1$, choose $g\in\mathbb{U}$ of order $2n$, and put $H=\langle g^2\rangle$. The coset $E'=gH$ has size n, is stable under inversion, and avoids -1. Its inverse image under $u=x+iy$ consists of n points of the circle over $\mathbb{F}_p$. Lemma 9.18 identifies its circle code with a diagonal translate of $\mathrm{RS}_{\mathbb{F}_{p^2}}(E',2w+1)$.

For part (c), choose $g\in\mathbb{U}$ of order $4n$, put $H=\langle g^4\rangle$, and use the twin coset $gH\sqcup g^{-1}H$. Its χ-image has size n, and $\operatorname{ord}(g)=4n\nmid 2n$, so Lemma 9.17 proves smoothness at every dyadic scale. The density assertions are the prime number theorem in the three displayed reduced residue classes. □

Corollary 9.25 (Finite smooth and circle families beyond the prior window)**.** *Let $n=32,k=15,r=6$. For every prime $p\equiv 1 \pmod{32}$, let $H\le\mathbb{F}_p^\times$ have order 32 and put $C_p^+=\mathrm{RS}_{\mathbb{F}_p}(H,15)$. For every prime $p\equiv -1 \pmod{64}$, construct the circle coset E of Theorem 9.24(b) and put $C_p^-=\mathcal{C}_7(\mathbb{F}_{p^2},E)$. Then*

(SC3)
$$B^{\mathrm{MCA}}_{C_p^+}(26)=B^{\mathrm{MCA}}_{C_p^-}(26)=7.$$

At targets $6/p$ and $6/p^2$, respectively, both complete safe sets are $[0,3/16)$, with unsafe endpoint. The two supporting prime sets have relative densities 1/16 and 1/32.

Proof. Here $R=17$ and

$$(n-r)^2=26^2=676>672=32(15+6).$$

Thus Theorem 4.6 and Corollary 9.22 give the numerator seven for the smooth and circle rows. Radius five is in the deep range and has numerator six; the tangent floor at radius six and monotonicity give the stated safe sets. This example lies beyond the elementary one-third range because $3r = 18 > R$. It is not certified by the specialized BCHKS estimate (LC5), whose exact first term is

$$\frac{32(18-6)}{6(18-12)} = \frac{32}{3} > 7.$$

The older conservative R-based comparator is $176/15$. The respective densities are $1/\varphi(32) = 1/16$ and $1/\varphi(64) = 1/32$. □

Corollary 9.26 (Concrete Mersenne-31 circle rows)**.** *Let $p = 2^{31}-1$, the Mersenne prime used in [HLP24], and let $\mathbb{U}$ be its norm-one circle group. Choose a generator g, put $H_0 = \langle g^2\rangle$, and let $E \subseteq \{(x,y) \in \mathbb{F}_p^2 : x^2+y^2 = 1\}$ be the circle domain corresponding to gH_0, so $|E| = 2^{30}$. This domain avoids $(-1,0)$. Over $\mathbb{F}_{p^5}$, the circle code with $w = 2^{28}$ and $k = 2^{29}+1$ has, at target 2^{-128}, the exact safe set*

$$\left[0, \frac{1}{8} - 2^{-30}\right). \tag{SC4}$$

Moreover, if $H = \langle g^4\rangle$, the Chebyshev twin-coset $gH \sqcup g^{-1}H = gH_0$ has x-domain $D = \chi(gH_0)$ of size 2^{29}, and $\mathrm{RS}_{\mathbb{F}_{p^5}}(D, 2^{27})$ has exact safe set

$$\left[0, \frac{1}{4} - 2^{-29}\right). \tag{SC5}$$

The first code has rate slightly above $1/2$; the second has exact rate $1/4$ but is a circle-smooth x-domain rather than a literal multiplicative-domain Prize row.

Proof. The group $\mathbb{U}$ has order $p+1 = 2^{31}$. Since $-1 = g^{2^{30}} \in H_0$, neither 1 nor -1 belongs to the disjoint coset gH_0; in particular E avoids the pole $(-1,0)$ of the stereographic chart. Also $gH_0 = gH \sqcup g^{-1}H$. Inversion acts freely on this coset and $\chi(u) = (u+u^{-1})/2$ identifies precisely the pairs $\{u, u^{-1}\}$, proving $|D| = 2^{29}$. The relation $T_a \circ \chi = \chi \circ (u \mapsto u^a)$ and Lemma 9.17 give the asserted dyadic smoothness.

Binomial expansion gives $B = \lfloor p^5/2^{128} \rfloor = 2^{27}-1$. At radius $B-1$, the quadratic margin $(n-r)^2 - n(k+r)$ for the bivariate row is $162129591417176068 > 0$, so Corollary 9.22 and the tangent floor give (SC4). Here $p^5 \equiv 3 \pmod 4$, so $i \notin \mathbb{F}_{p^5}$; the applicable equivalence is the stereographic one in Proposition 9.21, not torus uniformization. For the x-domain row the same budget lies inside the quadratic staircase at rate $1/4$, since its margin at radius $B-1$ is $18014401193836548 > 0$. The identical target argument gives (SC5). □

Corollary 9.27 (Every odd-characteristic extension-field tower)**.** *Fix an odd prime p, put $n_e = 2^e$ and $q_e = p^{n_e}$, and assume $e \geq 7$. Then $4n_e \mid q_e - 1$. Over $\mathbb{F}_{q_e}$ there exist both a multiplicative subgroup RS row of length n_e and a Chebyshev-smooth circle x-domain row of that length. At every official rate and every $0 \leq r \leq r_\rho^\sharp$, both satisfy*

$$e_{\mathrm{MCA}}(C_e, r/n_e) = \frac{r+1}{q_e}. \tag{SC6}$$

Thus the construction exists in every odd characteristic and beats every fixed error target throughout every compact subinterval below half the minimum distance for large e.

Proof. The lifting-the-exponent identity gives

$$v_2(p^{2^e}-1) = v_2(p-1) + v_2(p+1) + e - 1 \geq e+2.$$

Cyclicity produces the multiplicative subgroup and an element of order $4n_e$. Since n_e is even, $q_e \equiv 1 \pmod 4$; choose $i \in \mathbb{F}_{q_e}$ with $i^2 = -1$. The map

$$u \longmapsto \left(\frac{u+u^{-1}}{2}, \frac{u-u^{-1}}{2i}\right)$$

identifies $\mathbb{F}_{q_e}^{\times}$ with the split circle over $\mathbb{F}_{q_e}$. If g has order $4n_e$ and $H = \langle g^4 \rangle$, the twin coset $gH \sqcup g^{-1}H$ projects to an n_e-point x-domain. The proof of Lemma 9.17, which uses only cyclicity and $T_a \circ \chi = \chi \circ (u \mapsto u^a)$, applies verbatim and proves Chebyshev smoothness at every dyadic scale. Apply Corollary 5.5. Since $q_e = p^{n_e}$, the normalized error tends to zero exponentially. □

Proposition 9.28 (No growing positive-density prime-field family)**.** *Let $\mathcal{P}$ be a set of primes and assign a dyadic integer $n_p \to \infty$ along $\mathcal{P}$. If $n_p \mid p-1$ for every $p \in \mathcal{P}$, or $n_p \mid p+1$ for every $p \in \mathcal{P}$, then $\mathcal{P}$ has upper relative prime density zero. If in addition p/n_p is bounded, there are only $O(\log X)$ possible primes of the indicated form below X.*

Proof. For every fixed E, all sufficiently large assigned lengths are divisible by 2^E. Hence the primes eventually lie in the single class $1 \pmod{2^E}$, or the single class $-1 \pmod{2^E}$, whose relative prime density is $1/\varphi(2^E) = 2^{1-E}$. Letting $E \to \infty$ proves the first assertion. Under the bounded-ratio hypothesis, $p \mp 1 = a2^e$ with a in a fixed finite set; there are only $O(\log X)$ choices of (a, e) below X. □

Remark 9.29 (Prize scope). The prime-density statements fix n; their densities tend to zero as n grows. The extension towers have exponential field size and soon leave the $q < 2^{256}$ Prize box. Finally the target $6/p^2$ in Corollary 9.25 varies with p. None of these facts may be rephrased as a positive-density growing prime-field solution at the fixed target 2^{-128}.

9.4. **Quotient–remainder profiles and partial occupancy.** The quotient and remainder variables can be separated exactly at the locator-prefix level. This removes a spurious factor coming from summing over marked remainder sets, but it also exposes a genuine scaled quotient-coefficient-field term which prevents an unrestricted comparison with the identity profile.

Definition 9.30 (Quotient/remainder profile and field drop)**.** *Let $\phi : D \to Q$ be a c-fold complete-fiber folding map in the sense of Definition 9.11. A* complete-fiber-plus-remainder support *is*

$$S = \phi^{-1}(E) \sqcup R,$$

where $E \subseteq Q$ and $R \subseteq D \setminus \phi^{-1}(E)$. The set R is the support remainder*: it records the selected points in fibers that are not declared complete. A* quotient/remainder profile *fixes ϕ, its degree c, the size $r = |R|$, the relevant coefficient field, and any planted remainder data, while E and the unfixed part of R range subject to the displayed disjointness. A fixed-R profile is the subprofile in which the remainder support itself is part of the label. The adjective* Euclidean *used below refers to separating the locator factor of degree $r < c$ from the composed quotient locator; it does not mean that an arbitrary support remainder has size less than c.*

We call $\mathbb{B}_\phi \subseteq \mathbb{B}$ a scaled quotient coefficient field *when, for some $\eta \in \mathbb{B}^{\times}$,*

$$Q = \phi(D) \subseteq \eta\mathbb{B}_\phi.$$

When this field is recorded in a profile, it is taken minimal among the subfields with this property. Then the j-th elementary quotient coefficient lies in $\eta^j\mathbb{B}_\phi$. A field drop *is the use of a proper such subfield. It is a* positive-rate field drop *along a sequence when*

$$1 - \frac{\log|\mathbb{B}_\phi|}{\log|\mathbb{B}|}$$

stays bounded away from zero and a positive-density number on the entropy scale of quotient-prefix slots is live. A field drop changes the number of available coefficient values and is therefore recorded explicitly.

Definition 9.31 (Balanced quotient core)**.** *A* split locator *is a monic squarefree polynomial all of whose roots lie in the evaluation domain D. After common locator factors, fixed quotient data, and planted remainder points have been removed, let A and B be two monic split locators of the same degree e. They form a* depth-w balanced core *if they have the same first w nonleading coefficients, equivalently*

$$\deg(A - B) \le e - w - 1.$$

It is a balanced quotient core *when the remaining locator factors are pullbacks through the same folding map ϕ. The word* balanced *records equality of degree and leading normalization; the* core *consists of the roots still allowed to move after common and planted factors are removed. If the resulting projective locator family is one-dimensional it is a split pencil $Q_0 + \lambda Q_1$; in higher dimension this definition by itself supplies no distinct-ray bound. The exact determinant-pivot compiler still applies, but controlling its boundary maps requires a separate quantitative upper bound on their fiber sizes.*

For a monic polynomial

$$A(X) = X^d + a_1 X^{d-1} + \cdots + a_d$$

and an integer $0 \le u \le d$, write $\mathrm{pref}_u(A) = (a_1, \dots, a_u)$. Equivalently, if $A^\vee(Z) := Z^d A(Z^{-1})$, then $\mathrm{pref}_u(A)$ is the coefficient vector of $A^\vee(Z)$ mod Z^{u+1}, with its constant coefficient omitted.

Theorem 9.32 (Exact quotient–remainder prefix normal form)**.** *Let $\mathbb{B} \subseteq \mathbb{F}$ be finite fields, let $D \subseteq \mathbb{B}$ have n distinct elements, and let $\phi \in \mathbb{B}[X]$ be monic of degree $c \ge 2$. Assume that the restriction*

$$\phi : D \longrightarrow Q := \phi(D)$$

has complete fibers of size c; in particular, $N := |Q| = n/c$. Fix integers m, r with $0 \le r < c$ and $cm + r \le n$. For

$$E \in \binom{Q}{m}, \qquad R \in \binom{D \setminus \phi^{-1}(E)}{r},$$

put

$$\begin{aligned} V_E(Y) &:= \prod_{y \in E} (Y - y) = Y^m + \sum_{j=1}^{m} v_j(E) Y^{m-j}, \\ P_R(X) &:= \prod_{x \in R} (X - x) = X^r + \sum_{i=1}^{r} p_i(R) X^{r-i}, \\ L_{E,R}(X) &:= P_R(X) V_E(\phi(X)). \end{aligned} \tag{QR1}$$

Then $L_{E,R}$ is the monic locator of the support $\phi^{-1}(E) \sqcup R$, of degree $A = cm + r$.

Let $0 \le w \le A$ and put $d = \lfloor w/c \rfloor$.

(i) *If $w \ge r$, every nonempty fiber of $(E, R) \mapsto \mathrm{pref}_w(L_{E,R})$ determines R uniquely and, after R is recovered, is exactly one quotient-prefix fiber. More precisely, for every prefix value z there are unique data R_z and $\eta_z \in \mathbb{B}^d$, whenever the fiber is nonempty, such that*

$$\begin{aligned} &\{(E, R) : \mathrm{pref}_w(L_{E,R}) = z\} \\ &\quad = \{(E, R_z) : E \in \tbinom{Q \setminus \phi(R_z)}{m},\ \mathrm{pref}_d(V_E) = \eta_z\}. \end{aligned} \tag{QR2}$$

Thus no factor counting possible remainder sets occurs in a fixed global prefix fiber.

(ii) *If $w < r$, then necessarily $w < c$, and the quotient set E is invisible to the first w coefficients. There is a fixed unitriangular bijection $T_{\phi,m,w} : \mathbb{B}^w \to \mathbb{B}^w$ such that*

$$\mathrm{pref}_w(L_{E,R}) = T_{\phi,m,w}(\mathrm{pref}_w(P_R)) \tag{QR3}$$

for every admissible (E, R). Consequently, with the convention $\binom{u}{v} = 0$ for $v > u$,

$$\begin{aligned} &\#\{(E, R) : \mathrm{pref}_w(L_{E,R}) = z\} \\ &\quad = \sum_{\substack{R \in \binom{D}{r} \\ \mathrm{pref}_w(P_R) = T_{\phi,m,w}^{-1}(z)}} \binom{N - |\phi(R)|}{m}. \end{aligned} \tag{QR4}$$

Thus this case is a remainder-prefix problem, multiplied by the number of quotient sets still available after R is fixed.

Proof. For a monic degree-e polynomial A, set $A^\vee(Z) = Z^e A(Z^{-1})$. Also set

$$\Phi(Z) := Z^c\phi(Z^{-1}) = 1 + \phi_1 Z + \cdots + \phi_c Z^c.$$

Taking reciprocal polynomials in (QR1) gives the exact formal-power-series identity

$$L^\vee_{E,R}(Z) = P^\vee_R(Z)\left(\Phi(Z)^m + \sum_{j=1}^{m} v_j(E) Z^{cj}\Phi(Z)^{m-j}\right). \tag{QR5}$$

All factors on the right have constant coefficient one except for the displayed scalar coefficients $v_j(E)$.

Suppose first that $w \geq r$. Since $r < c$, no term containing a v_j occurs below degree c. For $1 \leq i \leq r$, the coefficient of Z^i on the right of (QR5) is $p_i(R)$ plus an expression depending only on $p_1(R), \ldots, p_{i-1}(R)$, on m, and on Φ. The first r coefficients therefore recover $p_1(R), \ldots, p_r(R)$ successively. They recover the monic polynomial P_R, hence its root set R, uniquely.

We may now divide $L^\vee_{E,R}$ by the unit $P^\vee_R$ modulo Z^{w+1}. In the remaining factor of (QR5), the coefficient of Z^{cj} is $v_j(E)$ plus an expression depending only on $v_1(E), \ldots, v_{j-1}(E)$, on m, and on Φ. Hence the coefficients at degrees $c, 2c, \ldots, dc$ recover $v_1(E), \ldots, v_d(E)$ successively. Conversely, P_R and these d coefficients determine the right side of (QR5) modulo Z^{w+1}, because every term containing v_j with $j > d$ starts in degree $cj > w$. This proves (QR2); admissibility of E is exactly the condition $E \cap \phi(R) = \varnothing$.

If $w < r < c$, every term involving a v_j starts in degree at least $c > w$, so modulo Z^{w+1} the bracket in (QR5) is the fixed unit $\Phi(Z)^m$. Multiplication by this unit sends the first w coefficients of $P^\vee_R$ to the first w coefficients of $L^\vee_{E,R}$ by a unitriangular transformation. It is independent of E and invertible. For a fixed R, the admissible E are precisely the m-subsets of $Q \setminus \phi(R)$, of which there are $\binom{N-|\phi(R)|}{m}$. Summing over the indicated remainder-prefix fiber proves (QR4). □

Proposition 9.33 (Complete support factorization and its limitation)**.** *Under the complete-fiber hypotheses of Theorem 9.32, every support $S \subseteq D$ has the unique decomposition*

$$E(S) = \{y \in Q : \phi^{-1}(y) \subseteq S\}, \qquad R(S) = S \setminus \phi^{-1}(E(S)), \tag{QR5a}$$

and $Q_S(X) = P_{R(S)}(X)V_{E(S)}(\phi(X))$. The remainder meets every nonfull ϕ-fiber in at most $c-1$ points, but its total size may be $\Theta(n)$. The reciprocal identity (QR5) remains valid for this arbitrary remainder. When $|R(S)| \geq c$, however, the quotient coefficients and the remainder coefficients begin to interlace at degree c, so the triangular recovery theorem above does not give a comparison without an additional partial-occupancy atlas.

Proof. A fiber is put into $E(S)$ exactly when all of its points lie in S; all other selected points form $R(S)$. This proves existence, disjointness, the occupancy bound, and uniqueness. Multiplying the corresponding linear factors gives the locator identity. The reciprocal calculation used for (QR5) is purely multiplicative and does not require $\deg P_R < c$. If $\deg P_R \geq c$, both $p_c(R)$ and $v_1(E)$ occur in degree c, which proves the stated limitation of the triangular separation. □

Theorem 9.34 (Exact partial-occupancy disintegration)**.** *Let $D = D_0 \sqcup X$, where $|X| = b$, and let $\phi : D_0 \twoheadrightarrow Q$ have exactly $N = |Q|$ fibers, each of cardinality c. For $S \subseteq D$, put*

$$X_S = S \cap X, \qquad E(S) = \{y \in Q : \phi^{-1}(y) \subseteq S\}, \qquad R(S) = (S \cap D_0) \setminus \phi^{-1}(E(S)).$$

Write

$$t = |X_S|, \qquad m = |E(S)|, \qquad p = |\phi(R(S))|, \qquad r = |R(S)|.$$

Then this datum is unique, every fiber meeting $R(S)$ contains between one and $c-1$ points of $R(S)$, and $|S| = t + cm + r$. If $\Omega_{t,m,p,r}$ denotes the family of supports with these four parameters, then

$$|\Omega_{t,m,p,r}| = \binom{b}{t}\binom{N}{p}\binom{N-p}{m}[x^r]((1+x)^c - 1 - x^c)^p. \tag{PO1}$$

Consequently the nonempty families $\Omega_{t,m,p,r}$ partition $\binom{D}{a}$, and their exact add-back is

$$\sum_{\substack{t,m,p,r\\ t+cm+r=a}} |\Omega_{t,m,p,r}| = \binom{cN+b}{a}. \tag{PO2}$$

If ϕ is the restriction of a monic degree-c polynomial and every fiber $\phi^{-1}(y)$ is the complete simple root set in D_0 of $\phi(X)-y$, then their locators satisfy

$$Q_S = Q_{X_S} P_{R(S)} V_{E(S)}(\phi).$$

Proof. Full fibers, partial fibers, and empty fibers are mutually exclusive, so the displayed decomposition is canonical. Choose the t exceptional points, the p partial fibers, and the m full fibers in $\binom{b}{t}\binom{N}{p}\binom{N-p}{m}$ ways. On one partial fiber the weight enumerator of the allowed nonempty proper subsets is $(1+x)^c-1-x^c$; hence the coefficient in (PO1) counts the choices having total partial weight r. Summing (PO1) and extracting the coefficient of x^a gives

$$[x^a](1+x)^b\left(1+x^c+\left((1+x)^c-1-x^c\right)\right)^N = [x^a](1+x)^{b+cN} = \binom{b+cN}{a}.$$

Multiplying the linear factors belonging to the three disjoint parts of the support proves the locator identity. □

Proposition 9.35 (Partial-occupancy Fourier factorization)**.** *In the setting of Theorem 9.34, let G be a finite abelian group, let $g : D \to G$, and put $\Phi(S) = \sum_{x\in S} g(x)$. For a character $\chi \in \widehat{G}$ and a regular fiber $F_y = \phi^{-1}(y)$, define*

$$A_y(\chi) = \prod_{x\in F_y} \chi(g(x)), \qquad P_y(\chi; z) = \prod_{x\in F_y} (1 + z\chi(g(x))) - 1 - z^c A_y(\chi).$$

Then

$$\sum_{S\in\Omega_{t,m,p,r}} \chi(\Phi(S)) = [u^t v^m \zeta^p z^r] \prod_{x\in X} (1+u\chi(g(x))) \prod_{y\in Q} (1 + vA_y(\chi) + \zeta P_y(\chi; z)). \tag{PO3}$$

Consequently, for every $s \in G$,

$$|\{S \in \Omega_{t,m,p,r} : \Phi(S) = s\}| = \frac{1}{|G|} \sum_{\chi\in\widehat{G}} \overline{\chi(s)} \sum_{S\in\Omega_{t,m,p,r}} \chi(\Phi(S)). \tag{PO4}$$

At the trivial character, (PO3) *is exactly the support count* (PO1)*. Thus support add-back is unconditional, while a finite flatness theorem is exactly a bound for the nontrivial-character sum in* (PO4)*.*

Proof. For a fixed regular fiber the three terms in the second product encode, respectively, an empty fiber, a full fiber, and a nonempty proper subset. The variables v, ζ, z record the numbers m, p, r, while u records the exceptional weight t. Multiplicativity of χ turns the character of the support sum into the product of its selected-point characters, proving (PO3) by coefficient extraction. Formula (PO4) is character orthogonality on G. □

Remark 9.36 (Effective normalization of a partial-occupancy slice)*.* Fix a nonempty slice $\Omega_\lambda = \Omega_{t,m,p,r}$, choose $S_0 \in \Omega_\lambda$, and let

$$G_\lambda = \langle \Phi(S) - \Phi(S_0) : S \in \Omega_\lambda \rangle \le G.$$

The attained boundary values lie in the affine coset $\Phi(S_0) + G_\lambda$, and the effective inversion is

$$|\{S \in \Omega_\lambda : \Phi(S) = s\}| = \frac{1}{|G_\lambda|} \sum_{\chi\in\widehat{G_\lambda}} \overline{\chi(s-\Phi(S_0))} \sum_{S\in\Omega_\lambda} \chi(\Phi(S) - \Phi(S_0)). \tag{PO5}$$

Every character χ of G_λ extends to a character $\widetilde{\chi}$ of the ambient finite abelian group, and

$$\sum_{S\in\Omega_\lambda} \chi(\Phi(S) - \Phi(S_0)) = \overline{\widetilde{\chi}(\Phi(S_0))} \sum_{S\in\Omega_\lambda} \widetilde{\chi}(\Phi(S)).$$

The right side may therefore be evaluated from (PO3) and is independent of the chosen extension, because different extensions agree on all support differences. Equivalently, ambient characters group by their restriction to G_λ, and the annihilator multiplicity cancels exactly against the ambient denominator. This identifies the correct Fourier denominator; it does not assert that the realized image fills the affine group.

Remark 9.37 (What partial-occupancy add-back does and does not prove). The exponentially many local proper-subset choices are already contained inside the coarse slices $\Omega_{t,m,p,r}$; they incur no separate profile-count loss in (PO2). The theorem does not bound the nontrivial Fourier terms in (PO4), the size of a realized boundary image, or the projection of a witness cell to distinct slopes. Those are separate analytic and ray payments.

10. Primitive inverse from a Sidon payment at accessible moment order

We now isolate the exact part of the primitive-fiber argument supplied by additive combinatorics. It is intentionally stated at effective image scale. In this section a *primitive residual leaf* means a family remaining after specified algebraic first-match cells have been removed; the adjective is bookkeeping and does not itself assert Fourier flatness, image coverage, or a slope bound.

Definition 10.1 (Image-normalized fiber moments). *For every N, let T_N be a set of size N, let $\Omega_N^0 \subseteq \{0,1\}^{T_N}$ be a nonempty reference profile slice—the full family of support incidence vectors sharing the fixed profile parameters before any first-match deletion—let $\Omega_N^\circ \subseteq \Omega_N^0$ be a primitive residual leaf, and let $\Phi_N : \Omega_N^0 \to V_N$ be a map to a finite set. Write*

$$I_N = \Phi_N(\Omega_N^0), \qquad L_N = |I_N|, \qquad \overline{N}_N = \frac{|\Omega_N^0|}{L_N},$$

so $L_N \le |\Omega_N^0|$ and $\overline{N}_N \ge 1$. For $s \in I_N$, put $F_{N,s} = \Omega_N^\circ \cap \Phi_N^{-1}(s)$ and $f_{N,s} = |F_{N,s}|$. For every finite $F \subseteq \mathbb{Z}^{T_N}$, its additive energy is computed in the torsion-free group $\mathbb{Z}^{T_N}$ by

$$E(F) = \#\{(x_1,x_2,x_3,x_4) \in F^4 : x_1 + x_2 = x_3 + x_4\}.$$

Put $\Delta_{N,s} = E(F_{N,s})/f_{N,s}^3$ when $f_{N,s} > 0$, and $\Delta_{N,s} = 0$ when $f_{N,s} = 0$. For an integer $q \ge 2$ and $\sigma > 0$, define the ordinary and low-energy, or Sidon-heavy*, normalized moments by*

$$\Gamma_{N,q}^{\mathrm{ord}} = \frac{1}{L_N} \sum_{s \in I_N} \left(\frac{f_{N,s}}{\overline{N}_N}\right)^q, \tag{PI1}$$

$$\Gamma_{N,q,\sigma}^{\mathrm{sid}} = \frac{1}{L_N} \sum_{\substack{s \in I_N \\ \Delta_{N,s} < e^{-\sigma N}}} \left(\frac{f_{N,s}}{\overline{N}_N}\right)^q. \tag{PI2}$$

The terminology records that a large fiber with exponentially small $\Delta_{N,s}$ can carry large ordinary moment mass while remaining too additively sparse for a high-energy inverse theorem.

Lemma 10.2 (Exact ambient/image conversion). *Suppose $I_N \subseteq W_N$, put $A_N = |W_N|$, extend $f_{N,s}$ by zero on $W_N \setminus I_N$, and set $\overline{N}_N^{\mathrm{amb}} = |\Omega_N^0|/A_N$. If the ambient moments are defined by replacing $L_N, \overline{N}_N, I_N$ in (PI1)–(PI2) with $A_N, \overline{N}_N^{\mathrm{amb}}, W_N$, then*

$$(\Gamma_{N,q}^{\mathrm{ord}})^{\mathrm{amb}} = \left(\frac{A_N}{L_N}\right)^{q-1} \Gamma_{N,q}^{\mathrm{ord}}, \qquad (\Gamma_{N,q,\sigma}^{\mathrm{sid}})^{\mathrm{amb}} = \left(\frac{A_N}{L_N}\right)^{q-1} \Gamma_{N,q,\sigma}^{\mathrm{sid}}. \tag{PI3}$$

Thus an ambient-scale moment payment transfers to image scale, whereas the reverse transfer requires $A_N/L_N = e^{o(N)}$ at the relevant orders.

Proof. For every attained value,

$$\frac{f_{N,s}}{\overline{N}_N^{\mathrm{amb}}} = \frac{A_N}{L_N} \frac{f_{N,s}}{\overline{N}_N}.$$

Raise to the q-th power and combine the factor A_N^{-1} with L_N^{-1}. The same calculation applies after restricting the sum by the energy condition. □

Theorem 10.3 (BSG and Boolean-cube difference growth)**.** *There is an absolute constant $C \geq 1$ with the following property. If A is a nonempty finite subset of an abelian group, $K \geq 1$, and*

$$E(A) := \#\{(a_1, a_2, a_3, a_4) \in A^4 : a_1 + a_2 = a_3 + a_4\} \geq \frac{|A|^3}{K},$$

then some $A' \subseteq A$ satisfies

$$|A'| \geq K^{-C}|A|, \qquad |A' - A'| \leq K^C|A'|. \tag{PI4}$$

Moreover, every $U \subseteq \{0,1\}^d$, viewed in $\mathbb{Z}^d$, satisfies

$$|U - U| \geq |U|^{3/2}. \tag{PI5}$$

The first assertion is the energy form of the Balog–Szemerédi–Gowers theorem; polynomial dependence on K follows from Gowers's refinement [BS94, Gow98]. The second is the Boolean-quasicube specialization of the sumset inequality of Green–Matolcsi–Ruzsa–Shakan–Zhelezov [GMR+22]: their quasicube inequality, with its two auxiliary sets specialized to $-U$ and $\{0\}$, gives $|U - U| \geq |U|^{3/2}$. These are the only external additive-combinatorics inputs below.

Proposition 10.4 (High-energy Boolean fibers are small)**.** *There is an absolute constant C_0 such that, for every finite set T, every nonempty $F \subseteq \{0,1\}^T$, and every $K \geq 1$,*

$$E(F) \geq \frac{|F|^3}{K} \quad \Longrightarrow \quad |F| \leq K^{C_0}. \tag{PI6}$$

Proof. Apply Theorem 10.3 to obtain $F' \subseteq F$ satisfying (PI4). The cube bound gives $|F'|^{3/2} \leq |F' - F'| \leq K^C|F'|$, so $|F'| \leq K^{2C}$. Since $|F| \leq K^C|F'|$, we get $|F| \leq K^{3C}$; take $C_0 = 3C$. □

Theorem 10.5 (Primitive flatness from a Sidon payment at accessible moment order)**.** *In the setting of Definition 10.1, suppose there are integers $q_N \geq 2$ such that the moment order is* accessible*, by which we mean*

$$\frac{\log L_N}{q_N} = o(N) \tag{PI7}$$

and, for every fixed $\sigma > 0$,

$$\Gamma^{\mathrm{sid}}_{N,q_N,\sigma} \leq e^{o(Nq_N)}. \tag{PI8}$$

Then the primitive fibers are flat at effective image scale:

$$\max_{s \in I_N} f_{N,s} \leq e^{o(N)}\overline{N}_N. \tag{PI9}$$

Proof. Suppose (PI9) fails. Along a subsequence there are $\eta > 0$ and $s_N \in I_N$ such that $f_{N,s_N} \geq e^{\eta N}\overline{N}_N$. By (PI1) and (PI7),

$$\Gamma^{\mathrm{ord}}_{N,q_N} \geq L_N^{-1}e^{\eta N q_N} = \exp(\eta N q_N - \log L_N) \geq e^{3\eta N q_N/4}$$

for all sufficiently large N in this subsequence.

Let C_0 be the constant in Proposition 10.4 and fix $0 < \sigma < \eta/(4C_0)$. The payment (PI8) is at most $e^{\eta N q_N/4}$ for large N. After subtracting it from the ordinary moment, the contribution of image values with $\Delta_{N,s} \geq e^{-\sigma N}$ is at least $e^{\eta N q_N/2}$. Hence one such image value satisfies $f_{N,s}/\overline{N}_N \geq e^{\eta N/2}$. Since $\overline{N}_N \geq 1$, this gives $f_{N,s} \geq e^{\eta N/2}$. On the other hand, $\Delta_{N,s} \geq e^{-\sigma N}$ and Proposition 10.4, with $K = e^{\sigma N}$, give $f_{N,s} \leq e^{C_0\sigma N} < e^{\eta N/4}$, a contradiction. □

Remark 10.6 (What the Gowers/cube step proves). The theorem is an unconditional implication from a Sidon payment at an accessible moment order, namely (PI7)–(PI8); it is not an unconditional verification of that payment. In a smooth/circle application, (PI8) still requires an effective-image Fourier theorem—for example aggregate minor- and major-character bounds—or a direct low-energy moment estimate on every primitive residual leaf. Even after (PI9), a determinant or ray compiler is needed to turn support fibers into distinct MCA slopes, and a profile-envelope comparison is needed to reach a threshold.

The phrase "Gowers/cube step" refers here to the polynomial Balog–Szemerédi–Gowers theorem followed by Boolean-quasicube difference growth. No Gowers-uniformity norm or Gowers-cube counting theorem is used.

11. Asymptotic and finite regimes

Theorem 11.1 (Capacity-edge shallow-prefix exponent)**.** *Let $C_n = \mathrm{RS}_{\mathbb{F}_n}(D_n, k_n)$, where $D_n \subseteq \mathbb{B}_n \subseteq \mathbb{F}_n$, $|D_n| = n$, and $k_n/n \to \rho \in (0,1)$. Let $a_n \in \{k_n+1, \dots, n\}$, and put $w_n = a_n - k_n - 1 \geq 0$. Assume*

$$w_n \log_2 |\mathbb{B}_n| = o(n), \qquad |\mathbb{F}_n| > n, \tag{AS1}$$

and, for some $\phi \geq 0$,

$$\frac{1}{n}\log_2 |\mathbb{F}_n| \to \phi, \qquad \frac{1}{n}\log_2 (|\mathbb{F}_n| - n) \to \phi. \tag{AS2}$$

Then $a_n/n \to \rho$, and for the full-field challenge

$$\begin{aligned} \frac{1}{n}\log_2 B^{\mathrm{MCA}}_{C_n}(a_n) &= \min\{H_2(\rho), \phi\} + o(1), \\ \frac{1}{n}\log_2 e_{\mathrm{MCA}}\left(C_n, 1 - \frac{a_n}{n}\right) &= -(\phi - H_2(\rho))_+ + o(1). \end{aligned} \tag{AS3}$$

No smoothness, Fourier estimate, or auxiliary ray hypothesis is assumed.

Proof. Since $|\mathbb{B}_n| \geq n$, (AS1) gives $w_n = o(n)$, hence $a_n/n \to \rho$. Put $M_n = \binom{n}{a_n}$ and $L_n = \lceil M_n |\mathbb{B}_n|^{-w_n} \rceil$. Exact-support uniqueness gives $B^{\mathrm{MCA}}_{C_n}(a_n) \leq \min\{|\mathbb{F}_n|, M_n\}$, while the exact locator-prefix list and collision-aware pole compiler give

$$B^{\mathrm{MCA}}_{C_n}(a_n) \geq \left\lceil \frac{L_n(|\mathbb{F}_n| - n)}{|\mathbb{F}_n| - n + k_n(L_n - 1)} \right\rceil. \tag{AS4}$$

Stirling and (AS1) give $\log_2 L_n = H_2(\rho)n + o(n)$. For $x, y > 0$, $xy/(x + k_n y) \geq \frac{1}{2}\min\{y, x/k_n\}$, so (AS2) shows that the lower bound in (AS4) has exponent $\min\{H_2(\rho), \phi\}$, matching the upper bound. Subtract $\log_2 |\mathbb{F}_n|$ to obtain the error exponent. □

Corollary 11.2 (Sharp capacity-edge criterion)**.** *Under the hypotheses of Theorem 11.1, let $\varepsilon_n = 2^{-\lambda n + o(n)}$, where $\lambda \geq 0$. If*

$$\lambda < \phi - H_2(\rho), \tag{AS5}$$

then eventually

(AS6)

$$a_n^* = k_n + 1, \qquad \delta^*_{n,\mathrm{sub,grid}} = 1 - \frac{k_n + 1}{n} = 1 - \rho + o(1), \quad \delta^*_{n,\mathrm{sub,sup}} = 1 - \frac{k_n}{n},$$
$$\delta^*_{C_n,\mathrm{off,grid}} = \delta^*_{C_n,\mathrm{off,sup}} = 1.$$

If $\lambda > (\phi - H_2(\rho))_+$, every agreement sequence satisfying $(a_n - k_n - 1)\log|\mathbb{B}_n| = o(n)$ is unsafe. The equality case is constant-sensitive. Thus the strict capacity branch of Theorem 1.8 is exponent-sharp at the first agreement, but it does not supply a matching attack at a fixed radius away from capacity.

Proof. At $a = k_n + 1$, the support atlas gives $B^{\mathrm{MCA}}_{C_n}(a) \leq \binom{n}{k_n+1} = 2^{H_2(\rho)n + o(n)}$, whereas the target budget has exponent $\phi - \lambda$; (AS5) makes the first agreement safe. The official conclusion follows from Theorem 1.4. The converse is the strict reverse comparison in (AS3); when it applies at $a_n = k_n + 1$, the MDS plateau makes the entire official tail with $a(\delta) \leq k_n + 1$ unsafe. □

Corollary 11.3 (Exact binary additive-domain examples)**.** *Let $k = 2^{40}$, let D be any $\mathbb{F}_2$-linear subspace of the indicated field having the indicated size, and take the target 2^{-128}. In each row put $C = \mathrm{RS}_{\mathbb{F}}(D, k)$. The four rows*

k/n	$\mathbb{F}$	n	$\delta^*_{C,\mathrm{sub,sup}}(2^{-128})$
1/2	$\mathbb{F}_{2^{166}}$	2^{41}	1/8
1/4	$\mathbb{F}_{2^{168}}$	2^{42}	1/4
1/8	$\mathbb{F}_{2^{169}}$	2^{43}	1/4
1/16	$\mathbb{F}_{2^{170}}$	2^{44}	1/4

have exactly the displayed supremal MCA transitions, and every endpoint is unsafe. Since every transition occurs before capacity, the official supremum has the same value.

Proof. For $q = 2^m$, the slope budget is $B = \lfloor q2^{-128} \rfloor = 2^{m-128}$. In the four rows this is respectively $2^{38}, 2^{40}, 2^{41}, 2^{42}$, and $3(B-1) \le n-k$. Thus Corollary 4.3 gives numerator B at radius $B-1$, while Proposition 4.2 gives at least $B+1$ slopes at radius B. Translate the largest safe grid radius to the real closed-ball supremum B/n. □

Proposition 11.4 (Exact unsafe test)**.** *Let $C_n = \mathrm{RS}_{\mathbb{F}_n}(D_n, k_n)$, where $D_n \subseteq \mathbb{B}_n \subseteq \mathbb{F}_n$, $|D_n| = n$, $q_n = |\mathbb{F}_n| > n$, and $1 \le k_n < n$. Fix a nonempty challenge set $\Gamma_n \subseteq \mathbb{F}_n$, a target $\varepsilon_n \in [0,1]$, and an agreement $k_n + 1 \le a_n \le n$, and put*

$$B_n^* = \lfloor \varepsilon_n |\Gamma_n| \rfloor, \qquad L_n = \left\lceil \binom{n}{a_n} |\mathbb{B}_n|^{-(a_n - k_n - 1)} \right\rceil.$$

If

$$\left\lceil \frac{|\Gamma_n|}{q_n} \left\lceil \frac{L_n(q_n - n)}{q_n - n + k_n(L_n - 1)} \right\rceil \right\rceil > B_n^*, \tag{UT1}$$

then agreement a_n is unsafe at the target. The same statement holds with L_n replaced by any larger identity, quotient, Chebyshev, or remainder-profile list proved for the dimension-$k_n + 1$ code.

Proof. The identity list is Proposition 8.1, and Theorem 8.2 produces a bad-slope set $Z \subseteq \mathbb{F}_n$ of the inner-ceiling size. Replacing a line (r_0, r_1) by $(r_0 + \delta r_1, r_1)$ translates Z by $-\delta$. Averaging over $\delta \in \mathbb{F}_n$ gives $|Z||\Gamma_n|/q_n$ challenge slopes, so some translate has at least the outer-ceiling number. The profile variants use the same conversion after their list sizes have been established. □

Remark 11.5 (The target-normalized identity crossing). The source repository's entropy equation [Cho26] has a rigorous role on the attack side. Put

$$\rho_n = \frac{k_n + 1}{n}, \qquad \beta_n = \log_2 |\mathbb{B}_n|, \qquad b_n = \frac{1}{n} \log_2 \max\{1, B_n^*\}.$$

For an admissible integer agreement a_n, set $g_n = (a_n - k_n - 1)/n$ exactly and distinguish the raw locator mean $\widetilde{L}_n = \binom{n}{a_n} |\mathbb{B}_n|^{-(a_n - k_n - 1)}$ from $L_n = \left\lceil \widetilde{L}_n \right\rceil$. Stirling gives

$$\frac{1}{n} \log_2 \widetilde{L}_n = H_2(\rho_n + g_n) - \beta_n g_n + o(1), \qquad \frac{1}{n} \log_2 L_n = \big(H_2(\rho_n + g_n) - \beta_n g_n\big)_+ + o(1). \tag{IC1}$$

Whenever the pole conversion in Proposition 11.4 has only subexponential loss, the target-normalized raw attack crossing is the right endpoint

$$g_n^\dagger = \sup\left(\{g \in [0, 1 - \rho_n] : H_2(\rho_n + g) - \beta_n g \ge b_n\} \cup \{0\}\right). \tag{IC2}$$

For $b_n > 0$ this is the ordinary exponent comparison with the ceiled list. At $b_n = 0$ it deliberately uses the raw nonnegative root rather than the positive-part exponent, and constants in the finite test decide the boundary. The finite unsafe test remains the literal inequality (UT1); (IC2) locates it only up to asymptotic losses. If $b_n = o(1)$ and $\beta_n \to \infty$, the isolated root has the precise expansion

$$g_n^\dagger = \frac{H_2(\rho_n) - b_n}{\beta_n - H_2'(\rho_n)} + O_\rho(\beta_n^{-3}). \tag{IC3}$$

Using the zero-level equation is justified only when $b_n = o(1)$. None of these formulas identifies the true MCA threshold without a matching safe upper envelope; the smooth and circle quotient cells in the repository show why that additional comparison cannot be assumed.

Theorem 11.6 (Integrated finite frontier bracket)**.** *Let $C = \mathrm{RS}_{\mathbb{F}}(D,k)$, let $D \subseteq \mathbb{B} \subseteq \mathbb{F}$, put $q = |\mathbb{F}| > n$, fix $\varnothing \neq \Gamma \subseteq \mathbb{F}$, let $0 \le \varepsilon \le 1$, and let $B^* = \lfloor \varepsilon|\Gamma| \rfloor$. For $k+1 \le a \le n$, define*

$$L(a) = \left\lceil \binom{n}{a} |\mathbb{B}|^{-(a-k-1)} \right\rceil,$$

$$P(a) = \left\lceil \frac{|\Gamma|}{q} \left\lceil \frac{L(a)(q-n)}{q-n+k(L(a)-1)} \right\rceil \right\rceil, \tag{SB1}$$

$$U(a) = \min\left\{ |\Gamma|, \min_{k+1 \le b \le a} \left\lfloor \frac{\binom{n}{b}}{\binom{a-1}{b-1}} \right\rfloor \right\}.$$

Whenever the defining set is nonempty, write

$$a^* = \min\{a \in \{k+1,\dots,n\} : B^{\mathrm{MCA}}_{C,\Gamma}(a) \le B^*\}.$$

If $k+1 \le a_- < a_+ \le n$ satisfy

$$P(a_-) > B^*, \qquad U(a_+) \le B^*, \tag{SB2}$$

then the first safe nontrivial agreement exists and

$$a_- < a^* \le a_+. \tag{SB3}$$

If $a_+ = a_- + 1$, then $a^* = a_+$ exactly. In particular,

$$\min\left\{ |\Gamma|, \binom{n}{k+1} \right\} \le B^* \implies a^* = k+1, \qquad \delta^*_{C,\Gamma,\mathrm{off,grid}} = \delta^*_{C,\Gamma,\mathrm{off,sup}} = 1. \tag{SB4}$$

These are unconditional finite comparisons; (SB2) *is the literal target reserve and contains no asymptotic placeholder. Any stronger upper bound whose hypotheses have been checked may be intersected with $U(a)$; Theorem A.3 qualifies only after Assumption A.1 is discharged. Any stronger realized-list construction may replace $L(a)$ in $P(a)$.*

Proof. The exact prefix list and pole-line construction give $B^{\mathrm{MCA}}_{C,\Gamma}(a) \ge P(a)$ by Proposition 11.4. The circuit-incidence MDS envelope gives $B^{\mathrm{MCA}}_{C,\Gamma}(a) \le U(a)$ by Theorem 1.2. Thus a_- is unsafe and a_+ is safe; monotonicity proves (SB3) and the adjacent conclusion. For (SB4), the first permitted agreement is already safe, and Theorem 1.4 gives the complete official safe set $[0,1]$. □

12. Conclusion

The exact threshold depends first on the integer slope budget, not on a radius slogan. Uniformly below half the relative minimum distance, $B^* = \lfloor \varepsilon|\Gamma| \rfloor$ gives first safe agreement $n - B^* + 1$, largest safe grid radius $(B^*-1)/n$, and unsafe continuous endpoint B^*/n. This packages the proved BCHKS two-radius input, the exact sparse MCA layer, and the tangent construction with all floors exposed. If the first safe nontrivial agreement is instead $k+1$, the MDS plateau changes the official answer: every radius in $[0,1]$ is safe and the official endpoint is one. In the distinct exponential-budget regime, Jo's shortening transfer and the BCHKS Johnson theorem give the explicit subcapacity certificate $D^{\mathrm{sub}}_{\rho,\mathrm{short}}(\theta)$, strictly improving the smallest-test MDS exponent between Johnson and capacity. For $\theta > H_2(\rho)$ the subcapacity grid tends to $1-\rho$ while the official endpoint equals one; at equality the official outcome remains constant-sensitive. Below that branch the result is one-sided without a matching constant-radius attack.

Jo's all-test-size circuit bound subsumes the support and paving envelope. The augmented-word interpretation recorded here describes its smallest-test specialization in matroid language and helps expose a finite security interval rather than a one-sided score. Together with shortening it gives multiplicative and circle certificates beyond Johnson at length 512. The deployed prefix calculations move the unsafe side at KoalaBear and Mersenne–31, but their adjacent failure does not certify safety.

The primitive inverse theorem discharges the high-energy branch once an image-normalized Sidon payment is available at an accessible moment order. At energy threshold $e^{-\sigma N}$, its cube estimate gives $|F| \le e^{C_0 \sigma N}$; because the payment is assumed for every fixed $\sigma > 0$, an alleged

excess rate $\eta > 0$ is contradicted by choosing $\sigma < \eta/(4C_0)$. The payment and residual ray compiler remain substantive row inputs. Staged shortening telescopes to the same binomial cost as one-stage shortening, so it does not remove those interfaces in the subexponential-budget, positive-relative-radius regime.

All smooth and circle transfers are made through complete-fiber maps or an explicit monomial equivalence. Fixed-length residue classes have positive prime density, while no growing positive-density dyadic prime-field family is claimed. The identity entropy crossing remains a precise attack locator after target normalization; equality with the true near-capacity frontier still requires the missing first-match profile exhaustion, image-normalized Sidon/Fourier payment, residual ray compiler, and profile add-back with a safe target reserve. These open requirements are why the remaining difference is reported as a soundness gap, with proved-bit gains computed from (1.4).

Appendix A. Conditional parameter-retained interpolation audit

The circuit-incidence bound is strongest in short, low-rate rows over very large fields. Near Johnson at deployed length, a parameter-retained interpolation audit is potentially much stronger, but its exact lift is not a source-verbatim theorem. This appendix therefore isolates the single required assumption before retaining the three degrees in the BCHKS–BCIKS argument. Three finite details matter: ceilings in the monomial count, inseparability in small characteristic, and the top-degree inequality used by the Hensel lift. We state all three explicitly.

Write $q = |\mathbb{F}|$ and $C = \mathrm{RS}_{\mathbb{F}}(D, K)$, so K is the dimension and code polynomials have degree less than K. To avoid any degree–dimension ambiguity, the interpolation below uses the conservative weight $i + Kj$, rather than the slightly sharper $i + (K-1)j$. Let $A = n - r \geq K + 2$ be the agreement. For a positive integer m and positive real parameters D_X, D_Y, D_Z, put $U = \lceil D_X \rceil$, $V = \lceil D_Y \rceil$, and $W = \lceil D_Z \rceil$.

Assumption A.1 (Parameter-retained factor lift)**.** *Assume*

$$V \geq m, \quad W \geq V, \quad U > K(V-1),$$

$$D_X < mA, \qquad \mathrm{char}(\mathbb{F}) > V - 1, \tag{RF1}$$

$$(A - K - 1)(2U - 1) > (n - K - 1)(2K + 1),$$

$$q > 2UD_Y. \tag{RF2}$$

Let $u_0, u_1 \in \mathbb{F}^D$, and let $0 \neq Q \in \mathbb{F}[X, Y, Z]$ have $(1, K, 0)$-weighted degree less than D_X, Y-degree less than D_Y, and $(0, 1, 1)$-weighted degree less than D_Z. For distinct γ in a set S, suppose there are polynomials P_γ, $\deg P_\gamma < K$, satisfying $Q(X, P_\gamma(X), \gamma) = 0$, and chosen A-sets $A_\gamma \subseteq D$ on which $P_\gamma = u_0 + \gamma u_1$. If

$$|S| > 2UD_Y^2 D_Z + (r+1)D_Y, \tag{RF3}$$

then for some $\gamma_0 \in S$ there are $v_0, v_1 \in \mathbb{F}[X]$, $\deg v_i < K$, such that $u_i = v_i$ on the chosen support A_{γ_0}, for $i = 0, 1$.

Remark A.2 (Source and reduction of the retained-lift hypothesis). The intended source is the constrained-support form of the function-field Hensel factor-lifting argument in [BCIKS20, Secs. 5.2.5–5.2.7 and App. A], with the degree summation of [BCHKS25, Sec. 3.2]. We spell out the parameter map because it is where finite-rounding errors arise. Factor the content-free part of Q into irreducibles separable in Y, specialize at a common regular point x_0, and factor each specialization into Y-separable factors. The characteristic condition in (RF1) removes the purely inseparable branch. The excluded starting points are the zeros of a nonzero discriminant-content polynomial of degree less than $2UD_Y$, so the last condition in (RF2) supplies x_0. For a factor pair with degrees a_i, b_{ij}, c_i, the cited lifting argument either charges at most $2Ua_ib_{ij}c_i$ slopes or produces a factor $Y - v_0(X) - Zv_1(X)$; after the lift, at most $a_ib_{ij}c_i$ slopes are deleted. The integer truncation length is controlled by $2U - 1$, not by the real number $2D_X$; condition (RF2) is precisely the top-degree comparison that forces $\deg v_i < K$. We have made this comparison one unit more conservative than the source degree convention.

Summing over the factor pairs uses $\sum_i a_i < D_Y$, $\sum_j b_{ij} = a_i$, and the content-free weighted-degree budget $\sum_i a_i c_i < D_Y D_Z$. If no useful factor exists, these charges together with the content roots and the $r+1$ constrained-collinearity allowance total at most the right side of (RF3). Otherwise the lifted linear factor gives $P_\gamma = v_0 + \gamma v_1$ on a set T of more than $r+1$ slopes. Put $B = \{x : (u_0(x), u_1(x)) \neq (v_0(x), v_1(x))\}$. For each $x \in B$, the two affine combinations agree at at most one slope, while for each $\gamma \in T$ they may disagree only outside A_γ, at no more than r coordinates. Hence $|B|(|T|-1) \leq r|T|$, and the integrality together with $|T| > r+1$ gives $|B| \leq r$. If none of the chosen supports simultaneously explained the pair, every A_γ would meet B; at such a point the two affine combinations agree, so one point of B can serve at most one slope. Thus $|T| \leq |B| \leq r$, a contradiction. This is the constrained-agreement step behind [BCHKS25, Thm. 4.3]. What remains for unconditional use is a line-by-line verification that the cited Hensel claims permit this simultaneous arbitrary-parameter map, including the displayed truncation and degree ledger. Because that combined statement is not printed verbatim in either source, we isolate it as an assumption instead of silently promoting the synthesis to a theorem.

Theorem A.3 (Conditional retained-degree MCA bound)**.** *Assume Assumption A.1. In addition to* (RF1)–(RF2)*, suppose*

$$\sum_{j=0}^{V-1} (U - Kj)(W - j) > n \sum_{s=0}^{m-1} (W - s)(m - s). \tag{RF4}$$

Then, for every challenge set $\Gamma \subseteq \mathbb{F}$*,*

$$B^{\mathrm{MCA}}_{C,\Gamma}(A) \leq \min\left\{|\Gamma|, \left\lceil 2UD_Y^2 D_Z + (r+1)D_Y \right\rceil\right\}. \tag{RF5}$$

Proof. For a received pair (u_0, u_1), seek $Q = \sum q_{ijh} X^i Y^j Z^h$ with $i + Kj < D_X$, $j < D_Y$, and $j + h < D_Z$, vanishing to multiplicity m at every $(x, u_0(x) + Zu_1(x))$, $x \in D$. The number of coefficients is exactly the left side of (RF4). The Hasse derivative of Y-order s contributes $m - s$ derivative conditions, each a polynomial identity in Z with $W - s$ coefficients; hence the right side is an upper bound on the number of homogeneous linear constraints. The strict inequality produces a nonzero interpolant.

For each MCA-bad slope choose an exact A-support and an explaining polynomial P_γ. Multiplicity gives at least mA zeros of $Q(X, P_\gamma(X), \gamma)$, while its degree is less than $D_X < mA$; it is therefore zero. If the number of bad slopes exceeded (RF5), Assumption A.1 would give a chosen support on which both received coordinates are explained, contradicting the defining MCA nontriviality. Intersecting with Γ proves the result. □

Remark A.4 (Proof tier). The interpolation count and the MCA contradiction above are proved in full conditional on Assumption A.1. The unresolved import is precisely the parameter-retained combination of [BCIKS20, Claims 5.7, 5.9–5.10, Lem. A.1 and Claim A.2]. BCHKS records the same useful-factor step and degree summation in [BCHKS25, Sec. 3.2]. The arbitrary-parameter statement above is not printed verbatim as one numbered theorem in either source. The BCIKS claim and appendix numbering here refers to the extended ePrint version. The numerical rows below use the conservative dimension weight and are conditional on this single exposed proof obligation, not on the repository's experimental profile ledgers.

Corollary A.5 (Four conditional near-Johnson KoalaBear rows)**.** *Let* $p_{\mathrm{KB}} = 2^{31} - 2^{24} + 1$, $\mathbb{F} = \mathbb{F}_{p_{\mathrm{KB}}^6}$, $q = |\mathbb{F}| = p_{\mathrm{KB}}^6$, $n = 2^{21}$, *and let* $D \leq \mathbb{F}^*_{p_{\mathrm{KB}}}$ *have order* n*. Take the full-field challenge* $\Gamma = \mathbb{F}$*, and for each of the four dimensions* $K = \rho n$ *put* $C = \mathrm{RS}_{\mathbb{F}}(D, K)$*. Use* $D_X = U - 1 + 2^{-64}$, $D_Y = V - 1 + 2^{-64}$, *and* $D_Z = W - 1 + 2^{-64}$*, with the following convention-safe parameters.*

(RF6)

ρ	r	m	U	V	W
1/2	611982	119	176735230	169	27525
1/4	1045433	104	109378776	209	29028
1/8	1352390	90	67028580	256	31500
1/16	1569744	78	41137824	314	34101

Conditional on Assumption A.1, all hypotheses of Theorem A.3 hold. If R_{ret} denotes the resulting integer upper numerator, the tangent lower numerator $r+1$ leaves a security interval of width at most $g_{\tan} = \log_2(R_{\mathrm{ret}}/(r+1))$. The four conditional upper numerators, security bits, and residual gaps are

(RF7)

ρ	R_{ret}	$\log_2(q/R_{\mathrm{ret}})$	$g_{\tan}$
1/2	274589064742726105	128.002056	38.706920
1/4	274721012201264929	128.001363	37.935074
1/8	274578888391530706	128.002110	37.562917
1/16	274861787390229386	128.000624	37.349385

In particular every numerator is at most $\lfloor q2^{-128} \rfloor = 274980728111395087$, subject to the retained-lift assumption.

Proof. Here $p_{\mathrm{KB}} = 127 \cdot 2^{24} + 1$ and $3^{(p_{\mathrm{KB}}-1)/2} \equiv -1 \pmod{p_{\mathrm{KB}}}$, so Proth's theorem proves primality. Also $p_{\mathrm{KB}} > V-1$ and $2^{21} \mid p_{\mathrm{KB}}-1$. For the rows in table order, the strict interpolation margins in (RF4) are $4889934, 13182624, 11133440, 4204064$, and the left-minus-right margins in (RF2) are

$$152123705899212, \quad 113730027157979, \quad 63736189920080, \quad 32093320774290.$$

The other conditions follow from $U = mA$,

$$mA - D_X = 1 - 2^{-64} > 0, \qquad \begin{array}{c|cccc} \rho & 1/2 & 1/4 & 1/8 & 1/16 \\ \hline U - K(V-1) & 574462 & 326872 & 181860 & 112288. \end{array}$$

Substitution in (RF5), followed by exact ceiling, gives the four integers in (RF7). The bit columns are base-two logarithms, and the last column is $\log_2(R_{\mathrm{ret}}/(r+1))$. □

Remark A.6 (What the conditional bits do and do not close). Conditional on the retained lift, the computation crosses 128 bits at four near-Johnson radii. The tangent lower attacks still leave intervals of roughly 37.3–38.7 bits there. These are therefore major safe-side reductions of the soundness gap, not exact threshold determinations and not beyond-Johnson certificates. The circuit-incidence theorem and the conditional retained bound occupy complementary regimes.

Appendix B. Verification details

This appendix records elementary coding-theoretic and integer calculations that can be checked independently of the geometric proofs. The short derivations fix the normalization conventions. Exact integer checks accompany this source release in two scripts: `verify_paving_mca_v9_2.py` checks the unconditional paving, all-radius saturation, circle, primality, field-budget, and Johnson comparisons, while `verify_retained_bchks_v9_2.py` checks only the arithmetic consequences of the retained-lift assumption in Appendix A. It does not discharge that assumption. Commands and SHA-256 digests are recorded in `REPRODUCIBILITY_v9.2.md`. General proof-development materials are pinned at repository commit `02728b2`; the two scripts are ancillary files of this source release and are not claimed to occur in that earlier snapshot. The deployed prefix values are direct evaluations of (4.2).

B.1. **Reed–Solomon and MDS conventions.** Let $D = \{x_1, \dots, x_n\} \subseteq \mathbb{F}$ consist of distinct points and let $1 \le k < n$. Evaluation on D is injective on $\mathbb{F}[X]_{<k}$: a nonzero polynomial of degree below k cannot vanish at the $n \ge k$ distinct points of D. It follows that $\mathrm{RS}_{\mathbb{F}}(D,k)$ has dimension k. If two codewords are distinct, their difference is the evaluation of a nonzero polynomial of degree at most $k-1$, hence has at most $k-1$ zero coordinates and at least $n-k+1$ nonzero coordinates. Conversely, the evaluation of $\prod_{i=1}^{k-1}(X - x_i)$ has exactly $k-1$ zero coordinates. Thus the minimum distance is

$$d_{\min} = n - k + 1.$$

This proves directly that the code is MDS.

For $x \in D$, put

$$\lambda_x = \left(\prod_{y \in D \setminus \{x\}} (x - y) \right)^{-1}.$$

The Lagrange interpolation formula for a polynomial f of degree at most $n-1$ shows that its leading coefficient is $\sum_{x \in D} \lambda_x f(x)$. Applying this to $f(X) = X^j$, $0 \le j \le n-2$, gives

$$\sum_{x \in D} \lambda_x x^j = 0.$$

Let $R = n - k$ and let H be the R-by-n matrix whose column at x is

$$h_x = \lambda_x (1, x, \dots, x^{R-1})^{\mathsf{T}}.$$

If $P \in \mathbb{F}[X]_{<k}$, then the j-th entry of $H(P(x))_{x \in D}^{\mathsf{T}}$ is $\sum_x \lambda_x x^j P(x) = 0$, because $j + \deg P \le n-2$. Thus H annihilates the code. Every R-by-R submatrix of H is a nonzero diagonal scaling of a Vandermonde matrix, so it is invertible. Hence H has rank R, its kernel has dimension k, and its kernel is exactly $\mathrm{RS}_{\mathbb{F}}(D, k)$.

The same Vandermonde calculation proves that every set of at most R columns is linearly independent. In particular, for $E \subseteq D$, $|E| \le R$, the space $V_E = \mathrm{Span}\{h_x : x \in E\}$ has dimension $|E|$. A word u agrees with a codeword outside E if and only if its syndrome lies in V_E. Indeed, writing $u = c + e$ with e supported on E gives one direction. Conversely, if $Hu^{\mathsf{T}} = \sum_{x \in E} e_x h_x$, subtracting the word whose coordinates on E are e_x leaves zero syndrome and therefore a codeword. This is the precise error-support form of the syndrome normalization used in Proposition 3.1.

For completeness, the fixed-support slope uniqueness used repeatedly in the paper is also an MDS-free linear observation. Suppose a support S explains $u_0 + \gamma u_1$ and $u_0 + \gamma' u_1$ for distinct γ, γ'. Subtracting the two explaining codewords and dividing by $\gamma - \gamma'$ explains u_1 on S; subtracting γ times this explanation from the first explains u_0 there. Thus S is common. Consequently a noncommon support contributes at most one finite MCA-bad slope.

Lemma B.1 (Monotonicity and exact endpoint conversion)**.** *For fixed C and nonempty Γ, the integer numerator $B^{\mathrm{MCA}}_{C,\Gamma}(a)$ is nonincreasing in a. Let $0 \le \varepsilon \le 1$, let $B^* = \lfloor \varepsilon |\Gamma| \rfloor$, and suppose $a_0 > k+1$ is the first safe nontrivial agreement, while $a_0 - 1$ is unsafe. Then the official and subcapacity transitions agree: the largest safe grid radius is $1 - a_0/n$, whereas the safe real radii form*

$$\left[0, \frac{n - a_0 + 1}{n}\right),$$

so their supremum is one grid step larger and is not attained. If C is MDS and $a_0 = k+1$, the subcapacity safe set is $[0, 1 - k/n)$, its largest grid point is $1 - (k+1)/n$, and the official safe set is $[0, 1]$.

Proof. Every witness at agreement at least $a+1$ is also a witness at agreement at least a, proving monotonicity. For a real radius δ, the closed Hamming ball permits precisely $r = \lfloor \delta n \rfloor$ errors and therefore requires agreement $n - r$. The condition $n - \lfloor \delta n \rfloor \ge a_0$ is equivalent to $\lfloor \delta n \rfloor \le n - a_0$, or $\delta < (n - a_0 + 1)/n$. The largest grid point in that interval is $(n - a_0)/n = 1 - a_0/n$. In the final case, Eq. (OFF1) extends safety from agreement $k+1$ to every smaller threshold, which is exactly the official tail from capacity through radius one. □

This lemma explains the half-open intervals in the main theorem. In particular, at target 2^{-128}, writing $B = \lfloor q/2^{128} \rfloor$, exact numerator B at error radius $B - 1$ and numerator at least $B + 1$ at radius B give the safe set $[0, B/n)$, rather than a closed interval.

B.2. **The quadratic boundary.** Put $R = n - k$ and

$$F_{n,k}(r) = r^2 - n(3r - R).$$

The mean-overlap hypothesis is $F_{n,k}(r) \ge 0$. Its two roots are

$$\frac{3n \pm \sqrt{9n^2 - 4nR}}{2} = \frac{3n \pm \sqrt{n(5n + 4k)}}{2}.$$

Since $F_{n,k}(0) = nR > 0$ and $F'_{n,k}(r) = 2r - 3n < 0$ for $0 \le r \le n$, the admissible initial interval ends at

$$r_{\mathrm{quad}}(n,k) = \left\lfloor \frac{3n - \sqrt{n(5n+4k)}}{2} \right\rfloor.$$

If $k/n \to \rho$, division by n gives

$$\frac{r_{\mathrm{quad}}(n,k)}{n} \longrightarrow \alpha_\rho = \frac{3 - \sqrt{5+4\rho}}{2}.$$

The derivative observation also shows that checking $F_{n,k}(B-1) \ge 0 > F_{n,k}(B)$ proves the exact identity $r_{\mathrm{quad}}(n,k) = B - 1$; no floating-point approximation is involved.

At the four challenge rates the limiting constants are, for orientation,

ρ	α_ρ	decimal value
1/2	$(3-\sqrt{7})/2$	$0.1771243444\ldots$
1/4	$(3-\sqrt{6})/2$	$0.2752551286\ldots$
1/8	$(6-\sqrt{22})/4$	$0.3273960600\ldots$
1/16	$(6-\sqrt{21})/4$	$0.3543560762\ldots$

The decimals are not used in any proof; the exact signs in (PC0) locate the finite boundaries.

B.3. **Proth certificates and budget divisions.** We spell out the certificate format. Suppose $p = u2^s + 1$, where u is odd and $u < 2^s$, and suppose an integer a_0 satisfies

$$a_0^{(p-1)/2} \equiv -1 \pmod{p}.$$

If a prime ℓ divides p, then the order of a_0 modulo ℓ has exact 2-part 2^s. Hence $2^s \mid \ell - 1$, so $\ell \ge 2^s + 1$. The inequality $u < 2^s$ gives $p = u2^s + 1 < 2^{2s} + 1$, and therefore $2^s > \sqrt{p} - 1$. If p were composite it would have a prime divisor at most $\sqrt{p}$, contradicting the lower bound $\ell \ge 2^s + 1 > \sqrt{p}$. Thus p is prime. This is Proth's theorem in exactly the form needed here; see [Rie94, Sec. 4.1].

For a positive integer N, write $\mathrm{bits}(N) = \lceil \log_2 N \rceil$. All certificate inputs are collected below. The equality in the last column is a residue equality modulo the corresponding p.

ρ	s	u	a_0	bits(p)	$a_0^{(p-1)/2} \bmod p$
1/2	92	26766274163673319604503	3	167	$p-1$
1/4	93	41595378994516821279015	13	169	$p-1$
1/8	95	24737346889219389259839	5	170	$p-1$
1/16	97	13387194060291799253121	5	171	$p-1$

Together with the decimal values of p printed before Proposition 7.1, this table is a complete Proth certificate: the reader checks the single exact identity $p - 1 = u2^s$, the parity and size of u, and one modular exponentiation. Modular exponentiation is an integer operation and does not rely on a probabilistic primality test.

For a completely mechanical check, first compute $t = a_0^u \bmod p$ by binary square-and-multiply, then square t exactly $s - 1$ times modulo p. The certificate accepts precisely when the result is $p - 1$. The algorithm uses $O(\log u + s)$ modular multiplications. The remaining row checks use only division with remainder. Writing

$$p = B2^{128} + r_p,$$

the four exact remainders are

ρ	r_p
1/2	138154108384248438678742263 3985
1/4	292153849271349744876193316 8641
1/8	249568711919932663419663443 5585
1/16	204408659286801990991343391 86689

and each lies strictly between 0 and $2^{128} = 340282366920938463463374607431768211456$. Therefore the displayed B is exactly $\lfloor p/2^{128} \rfloor$.

Finally, $n = 2^e$ with $e = 41, 42, 43, 44$, respectively, while $s = 92, 93, 95, 97$. Thus $n \mid 2^s \mid p-1$, and the cyclic group $\mathbb{F}_p^\times$ contains a unique subgroup of order n. Each field has the printed bit length, all smaller than 256, and $k = 2^{40}$ gives the required four rates. The exact boundary signs are

ρ	$F_{n,k}(B-1)$	$F_{n,k}(B)$
1/2	5154112775168	−663955886271
1/4	7590647904465	−3182321912768
1/8	13908181940112	−6720484728007
1/16	19335616403905	−20973145690236

and hence $B - 1 = r_{\mathrm{quad}}(n, k)$ in every row.

Proposition B.2 (Independent certificate verification)**.** *The data in Table 2 define four prime fields, four power-of-two multiplicative subgroups of the stated lengths, and slope budgets* $B = \lfloor p2^{-128} \rfloor$*. In every row the exact safe set is* $[0, B/n)$*, and the endpoint is unsafe.*

Proof. The Proth calculation proves primality, the divisibility calculation produces the subgroup, and the remainder calculation gives the target budget. The two signs locate the quadratic boundary at $B - 1$. Apply Corollary 4.12 and then Lemma B.1. Every step is an equality or inequality between the printed integers. □

References

[ABF26] G. Arnon, D. Boneh, and G. Fenzi, Open problems in list decoding and correlated agreement, IACR Cryptology ePrint Archive, Report 2026/680, 2026, revised July 6, 2026.

[ACFY25] G. Arnon, A. Chiesa, G. Fenzi, and E. Yogev, WHIR: Reed–Solomon proximity testing with super-fast verification, in *EUROCRYPT 2025*, 214–243; IACR ePrint 2024/1586.

[BS94] A. Balog and E. Szemerédi, A statistical theorem of set addition, *Combinatorica* 14 (1994), no. 3, 263–268.

[BCHKS25] E. Ben-Sasson, D. Carmon, U. Haböck, S. Kopparty, and S. Saraf, On proximity gaps for Reed–Solomon codes, IACR Cryptology ePrint Archive, Report 2025/2055, 2025.

[BCIKS20] E. Ben-Sasson, D. Carmon, Y. Ishai, S. Kopparty, and S. Saraf, Proximity gaps for Reed–Solomon codes, *J. ACM* 70 (2023), no. 5, 1–57, doi:10.1145/3614423; conference version in *61st IEEE Symposium on Foundations of Computer Science (FOCS 2020)* and extended version, IACR ePrint 2020/654.

[BCGM25] S. Bordage, A. Chiesa, Z. Guan, and I. Manzur, All polynomial generators preserve distance with mutual correlated agreement, IACR Cryptology ePrint Archive, Report 2025/2051, 2025, revised May 2026; conference version in CCC 2026.

[Cho26] P. Chojecki, `rs-mca`: proof development and exact certificate ledgers, source repository, pinned snapshot `02728b208ea7`, 2026, https://github.com/przchojecki/rs-mca, accessed July 17, 2026.

[CS25] E. Crites and A. Stewart, On Reed–Solomon proximity gaps conjectures, IACR Cryptology ePrint Archive, Report 2025/2046, 2025.

[Dav00] H. Davenport, *Multiplicative Number Theory*, third ed., Graduate Texts in Mathematics 74, Springer, New York, 2000.

[PP26] Ethereum Foundation, The Proximity Prize: $1M Reed–Solomon Challenge, preliminary version, 2026, https://proximityprize.org/, accessed July 17, 2026.

[GYXK26] Y. Gao, H. Yang, Y. Xu, and H. Kan, List-decoding counterexamples yield lower bounds on mutual correlated agreement error, arXiv:2607.10572, version 1, July 2026.

[Gow98] W. T. Gowers, A new proof of Szemerédi's theorem for arithmetic progressions of length four, *Geom. Funct. Anal.* 8 (1998), no. 3, 529–551.

[GG25] R. Goyal and V. Guruswami, Optimal proximity gaps for subspace-design codes and (random) Reed–Solomon codes, IACR Cryptology ePrint Archive, Report 2025/2054, 2025, revised March 2026.

[GGSW26] R. Goyal, V. Guruswami, Y. Sun, and M. Wootters, Locality of curve-decoding and improved proximity gaps, arXiv:2607.08516, 2026.

[GMR+22] B. Green, D. Matolcsi, I. Z. Ruzsa, G. Shakan, and D. Zhelezov, A weighted Prékopa–Leindler inequality and sumsets with quasicubes, in *Analysis at Large*, Springer, 2022, 125–129; arXiv:2003.04077.

[Hab25] U. Haböck, A note on mutual correlated agreement for Reed–Solomon codes, IACR Cryptology ePrint Archive, Report 2025/2110, 2025.

[HLP24] U. Haböck, D. Levit, and S. Papini, Circle STARKs, IACR Cryptology ePrint Archive, Report 2024/278, 2024.

[Jo26] S. Jo, Reed–Solomon mutual correlated agreement beyond the Johnson radius, IACR Cryptology ePrint Archive, Report 2026/1432, 2026, https://eprint.iacr.org/2026/1432.

[Kam26] A. Kambiré, Proximity gaps conjecture fails near capacity over prime fields, arXiv:2604.09724, 2026.
[KKH26] D. Krachun, S. Kazanin, and U. Haböck, Failure of proximity gaps close to capacity, IACR Cryptology ePrint Archive, Report 2026/782, 2026.
[Rie94] H. Riesel, *Prime Numbers and Computer Methods for Factorization*, second ed., Progress in Mathematics 126, Birkhäuser, Boston, 1994.
[YZ26] C. Yuan and R. Zhu, A syndrome–space approach to proximity gaps and correlated agreement for random linear codes and random Reed–Solomon codes, arXiv:2605.07595, version 2, revised July 10, 2026.

Ulam.ai
Email address: prz.chojecki@gmail.com
URL: https://github.com/przchojecki