CONJECTURES AND BARRIERS FOR RS–MCA

PRZEMEK CHOJECKI

Abstract. As a companion to the proved bounds in *Shortening Bounds for Reed–Solomon MCA*, we formulate the complete finite Reed–Solomon mutual correlated agreement (MCA) problem and a pole-aware conjectural positive-density profile envelope. An exact first-match compiler and realized-image moment and incidence inequalities isolate the payments required for a safe certificate without confusing supports, pairs, rays, and affine slopes. We prove a ceiling-normalized moment obstruction, a projective incidence theorem and benchmark dimension diagnostic, and a moving-scale consequence of the previously proved implication from a Sidon payment through the Balog–Szemerédi–Gowers theorem and Boolean-cube growth. Exact unsafe edges, repository audits, and numerical margins lead to four direct adjacent conjectures—two MCA and two auxiliary list inequalities—whose safe sides still require primitive-fiber, residual-projection, algebraic-routing, and add-back payments. In a collision-nonbinding, subexponential-budget identity-candidate branch, we conjecture a non-oracular exhaustive atlas whose exact unsafe–safe bracket, through a proved crossing reduction, yields $\delta^*_{C_n,\mathrm{off},\sup} = 1-\rho_n-g^*(\rho_n,\log_2|\mathbb{B}_n|)+o(g^*_n)$; the general identity lower route instead uses an exact pole-adjusted target, and no matching bracket is claimed for the unrestricted smooth or circle problem.

1. Introduction

The preliminary Proximity Prize asks one to determine the complete security frontier of Reed–Solomon mutual correlated agreement (MCA) on structured evaluation domains [PP26, ABF26]. Its grand MCA challenge takes

$$C = \mathrm{RS}_{\mathbb{F}}(D,k) = \{(f(x))_{x\in D} : f\in\mathbb{F}[X],\ \deg f < k\}, \qquad n = |D|,$$

at a rate $k/n \in \{1/2, 1/4, 1/8, 1/16\}$. The preliminary statement phrases the task as finding the largest safe radius in the official interval $[0,1]$ at a target such as $\varepsilon^* = 2^{-128}$. That wording hides a finite staircase, and the soundness gap between an attack and a safe certificate is meaningful only after its integer and endpoint conventions have been fixed.

For a received pair $r = (r_0, r_1)$, the affine received line is $\{r_0 + \gamma r_1 : \gamma \in \mathbb{F}\}$. At agreement a, a slope γ is MCA-bad if there are a support $S \subseteq D$, with $|S| \geq a$, and $h \in \mathbb{F}[X]_{<k}$ for which $(r_0+\gamma r_1)|_S = h|_S$, while r_0 and r_1 are not both explained by degree-less-than-k polynomials on that same support. Thus MCA counts distinct affine slopes, not supports, locators, polynomials, pairs, or projective directions. For a nonempty challenge set $\Gamma \subseteq \mathbb{F}$, let $B^{\mathrm{MCA}}_{C,\Gamma}(a)$ be the maximum number of bad slopes in Γ, over all received pairs. At radius δ, put

$$a(\delta) = n - \lfloor \delta n \rfloor, \qquad \varepsilon_{\mathrm{MCA},\Gamma}(C,\delta) = \frac{B^{\mathrm{MCA}}_{C,\Gamma}(a(\delta))}{|\Gamma|}. \tag{1.1}$$

The official challenge is the full-field case $\Gamma = \mathbb{F}$. The challenge-restricted notation is our refinement for sampled variants; it does not change the support-wise witness predicate. We abbreviate $B^{\mathrm{MCA}}_C = B^{\mathrm{MCA}}_{C,\mathbb{F}}$.

Problem 1.1 (Complete finite Proximity Prize problem)**.** *Given the official data above, determine the complete safe set*

$$\mathsf{S}_{C,\Gamma}(\varepsilon^*) = \{\delta\in[0,1] : B^{\mathrm{MCA}}_{C,\Gamma}(a(\delta)) \leq B^*\}, \qquad B^* = \lfloor \varepsilon^*|\Gamma| \rfloor, \tag{1.2}$$

Date: July 19, 2026.

2020 *Mathematics Subject Classification.* 94B35, 11T71, 05B40, 11B30.

Key words and phrases. Reed–Solomon codes, mutual correlated agreement, profile certificates, additive energy, smooth domains, proximity gaps.

and not only a decimal approximation to one boundary point. When the safe set is nonempty, report its real supremum $\delta^*_{C,\Gamma,\mathrm{off,sup}}$*, its largest safe grid point* $\delta^*_{C,\Gamma,\mathrm{off,grid}} \in \{0, 1/n, \dots, 1\}$*, and whether the supremum is attained. When it is empty, report that fact separately; the number zero may be used as a placeholder but must not be read as saying that radius zero is safe. The concrete box emphasized in the accompanying preliminary problem discussion includes* $k \le 2^{40}$*,* $|\mathbb{F}| < 2^{256}$*, and target* 2^{-128} [ABF26]*; the definitions and endpoint bookkeeping here do not depend on those numerical caps.*

The complete answer has three branches. For a proper MDS code, proposition 2.1 proves that $B^{\mathrm{MCA}}_{C,\Gamma}(n) = 1$ and that the numerator is constant for $a \le k+1$. Hence $B^* = 0$ gives an empty safe set. If $B^{\mathrm{MCA}}_{C,\Gamma}(k+1) \le B^*$, the row is saturated: $\mathsf{S}_{C,\Gamma} = [0,1]$, both official thresholds equal 1, and the endpoint is safe. Otherwise $B^* \ge 1$, and if

$$a^*_C = \min\{a \in \{k+1, \dots, n\} : B^{\mathrm{MCA}}_{C,\Gamma}(a) \le B^*\}, \tag{1.3}$$

then

$$\mathsf{S}_{C,\Gamma} = \left[0, \frac{n - a^*_C + 1}{n}\right), \quad \delta^*_{C,\Gamma,\mathrm{off,grid}} = \frac{n - a^*_C}{n}, \quad \delta^*_{C,\Gamma,\mathrm{off,sup}} = \frac{n - a^*_C + 1}{n}, \tag{1.4}$$

and the supremal endpoint is unsafe. These empty, saturated, grid, and half-open cases are part of the finite Prize answer; a subcapacity threshold alone is incomplete.

The Prize statement uses a smooth evaluation domain. Our refined term *admissibly smooth* means that the row data include specified complete equal-fiber polynomial folding maps. Here a folding $\pi : D \to D_\pi$ is complete and c-to-one, for an integer $c \ge 1$, when $\pi(D) = D_\pi$ and every fiber $\pi^{-1}(y)$, $y \in D_\pi$, has exactly c elements. The principal examples are positive-density multiplicative domains, with power maps $x \mapsto x^c$, and Chebyshev twin-coset domains, with the corresponding maps T_c. We write $D \subseteq \mathbb{B} \subseteq \mathbb{F}$ when this locator and folding algebra is defined over a designated coefficient-and-folding field $\mathbb{B}$, while slopes and codewords live over $\mathbb{F}$. This field is part of the public row data; for the asymptotic prediction it is the field generated by D and the coefficients of the declared folding maps, not an arbitrarily enlarged intermediate field. A circle presentation enters only through an explicit coordinate bijection and nonzero diagonal scaling that give an agreement-preserving monomial equivalence with such an RS row [HLP24, Cho26a]. Geometric resemblance alone is not used. For a sequence, *positive density* means $\liminf_n |D_n|/|\mathbb{B}_n| > 0$; hence $\log_2|\mathbb{B}_n| = \log_2 n + O(1) = o(n)$.

The proved literature does not yet determine the grand frontier. BCIKS established the foundational Reed–Solomon proximity-gap framework [BCIKS20/23]. Haböck gave the quadratic predecessor of a Johnson-range MCA estimate [Hab25]; BCHKS state a linear-in-n Johnson-range estimate, which the companion theorem paper treats as an imported input, as well as complementary proximity-gap obstructions [BCHKS26]. Jo proves an all-test-size MDS circuit-incidence envelope and an exact agreement-set shortening theorem, yielding polynomial control at any fixed number of steps beyond Johnson at fixed reciprocal-integer rates [Jo26]. BCHKS use degree at most d; for $\mathrm{RS}_{<K}$, their strict Johnson hypothesis is $(N-E)^2 > N(K-1)$.

Recent attacks show why no identity-only capacity law can be universal. Crites and Stewart disprove several correlated-agreement and MCA up-to-capacity conjectures [CS25]. Krachun, Kazanin, and Haböck construct multiplicative-subgroup rows over prime fields with $2^{\Omega_\rho(1/\eta)}$ close points on one affine line at a gap $\eta = \Theta_\rho(1/\log n)$ from capacity [KKH26]. Gao, Yang, Xu, and Kan give a constructive conversion from list-decoding counterexamples to MCA lower bounds for a related code, including an RS specialization [GYXK26]. These are nonidentity lower mechanisms: they sharpen the attack side and the profile taxonomy, but do not furnish fixed-row safe certificates for the benchmarks below.

The preceding theorem paper, *Shortening Bounds for Reed–Solomon MCA*, version 9.2 [Cho26a], is the source of the unconditional RS–MCA results used here. It proves the official MDS plateau and exact endpoint conversion, the endpoint-exact linear-budget staircase in its below-half-distance range, finite MDS circuit and shortening certificates beyond Johnson, and a one-sided positive-radius safe-frontier certificate for positive exponential budgets. It also proves

the exact locator-prefix lower construction with collision-aware pole conversion, structured-domain and circle transfers under proved monomial equivalence, and the implication from an image-normalized Sidon payment to primitive Boolean-fiber flatness. The latter implication uses Balog–Szemerédi–Gowers and Boolean-cube difference growth; its Sidon premise and residual ray compiler remain open. It also determines complete safe sets for four explicit maximal-dimension $k = 2^{40}$ multiplicative rows in the tangent branch and four all-radius saturated rows [Cho26a, Corollary 1.5 and Proposition B.2]. Those exact instances do not determine the positive-density identity-profile regime studied here. This paper quotes those established frontiers instead of presenting them as new. The short endpoint and analytic implications are restated where needed to make the present compiler auditable.

The object of this companion is the missing classification. A safe proof must classify every bad witness, normalize boundary maps by their attained images, convert each count to the relevant projection, and add the payments without overlap. The exact first-match compiler in theorem 2.5 formalizes these requirements. The target-exact profile-closure meta-conjecture meta-conjecture 2.6 asks for one family-uniform, non-oracular atlas schema whose specializations determine every target comparison. It is a meta-conjecture about a proposed proof architecture, not a substitute for a statement about the true numerator.

The stable finite conjectures are therefore stated directly. Here $B_C^{\mathrm{list}}(a)$, defined formally in Section 2, denotes the maximum number of codewords of C agreeing with one received word on at least a coordinates. The benchmark codes C_{KB} and C_{M}, their fields, domains, and targets are defined in Section 5. At $n = 2^{21}$, $k = 2^{20}$, they are

$$\begin{aligned} B^{\mathrm{MCA}}_{C_{\mathrm{KB}}}(1116048) &\le 274980728111395087, \quad B^{\mathrm{MCA}}_{C_{\mathrm{M}}}(1116024) \le 16777215, \\ B^{\mathrm{list}}_{C_{\mathrm{KB}}}(1116047) &\le 274980728111395087, \quad B^{\mathrm{list}}_{C_{\mathrm{M}}}(1116023) \le 16777215. \end{aligned} \tag{1.5}$$

The top line contains the two MCA inequalities; the bottom line contains the two auxiliary agreeing-codeword inequalities. Each preceding agreement is rigorously unsafe, but the four displayed upper inequalities remain open (conjecture 5.1). Their exact margins, periodic audit, random-fiber heuristic, and falsifiers are recorded in Section 5.

The new theorems concern the interfaces and barriers. We prove realized-image moment and support-to-slope incidence inequalities, the necessary moment orders $94196, 94991, 641594, 680397$ for the stated ceiling-normalized route, a general projective incidence theorem, and a conditional correction-space diagnostic with dimensions near $9.1 \cdot 10^5$. We also isolate the moving-scale consequence of the Sidon–Balog–Szemerédi–Gowers–Boolean-cube implication already proved in [Cho26a, Section 10]. Closing any adjacent row still requires realized-image primitive-fiber control, residual slope or codeword payments, periodic/descent/extension/rank routing, and exact first-match add-back.

The pole-adjusted identity profile has a precise conjectural asymptotic shape. Write $H_2(x) = -x\log_2 x - (1-x)\log_2(1-x)$ for binary entropy, and let $\rho_n = k_n/n$, $\beta_n = \log_2|\mathbb{B}_n|$, and

$$g^*(r,b) = \sup(\{g \in [0, 1-r] : H_2(r+g) \ge bg\} \cup \{0\}). \tag{1.6}$$

For a declared target sequence $0 < \varepsilon_n \le 1$, put $q_n = |\mathbb{F}_n|$, $Q_n = q_n - n$, and $B_n^* = \lfloor \varepsilon_n q_n \rfloor$. On the nonzero pole-admissible branch $B_n^* \ge 1$ and $Q_n > k_n B_n^*$, let $T_{\mathrm{pole},n}$ be the exact threshold in (2.5), set $\tau_n^{\mathrm{pole}} = n^{-1}\log_2 T_{\mathrm{pole},n}$, and define

$$g_{\mathrm{pole},n} = \sup\left(\left\{g \in [0, 1-\rho_n^+] : H_2(\rho_n^+ + g) - \beta_n g \ge \tau_n^{\mathrm{pole}}\right\} \cup \{0\}\right), \qquad \rho_n^+ = \frac{k_n+1}{n}. \tag{1.7}$$

The exact ceiling test uses $T_{\mathrm{pole},n} - 1$, rather than $T_{\mathrm{pole},n}$; the normalized logarithmic levels differ by at most $1/n$. A matching exact certificate bracket of width $o(g_n^* n)$ at this crossing gives

$$\delta^*_{C_n,\mathrm{off},\sup} = 1 - \rho_n - g_{\mathrm{pole},n} + o(g_n^*), \qquad g_n^* = g^*(\rho_n, \beta_n). \tag{1.8}$$

If in addition the raw budget and the pole-conversion factor are subexponential,

$$\frac{1}{n}\log_2 B_n^* = o(1), \qquad \frac{1}{n}\log_2 \frac{Q_n - k_n}{Q_n - k_n B_n^*} = o(1), \tag{1.9}$$

then lemma 6.1 gives $g_{\mathrm{pole},n} = g_n^* + o(g_n^*)$. Thus the sharper prediction becomes

$$\delta^*_{C_n,\mathrm{off},\sup} = 1 - \rho_n - g_n^* + o(g_n^*). \tag{1.10}$$

At constant rate $\rho_n = \rho$ this implies the coarser formula

$$\boxed{\delta^*_{C_n,\mathrm{off},\sup} = 1 - \rho - g^*(\rho, \log_2|\mathbb{B}_n|) + o(1)}. \tag{1.11}$$

The stronger collision-nonbinding condition $k_n B_n^*/Q_n = o(1)$ implies the second condition in (1.9). Raw subexponential budget alone does not: there is no off-domain pole when $Q_n = 0$, and the certified pole lower bound never exceeds the budget when $0 < Q_n \le k_n B_n^*$, regardless of the locator-list size.

These conclusions do not require every multiplicative loss to be $2^{o(g_n^* n)}$. Near the right crossing, one more agreement changes the base-two logarithm of the identity mean by $-\beta_n + O(1)$; consequently a uniform aggregate factor $2^{o(n)}$ already moves the crossing by $o(n/\beta_n) = o(g_n^* n)$. Moving-scale control is a useful stronger statement, especially for finite ledgers, but is not necessary for (1.10).

Here the tangent branch is *nonbinding* when the previously proved linear-budget tangent edge has been tested and is not the active constraint in the crossing window. To formulate a lower-side hypothesis, fix uniformly from the public row-family data, before evaluating the unknown MCA numerator or the eventual safe/unsafe outcome, the non-oracular finite catalogue $\mathcal{L}^{\neq \mathrm{id}}$ containing every explicit nonidentity lower construction declared for that family in this manuscript or its pinned public row specification. The catalogue is part of the fixed family data; no subcatalogue may be selected when testing candidate status. For $\lambda \in \mathcal{L}^{\neq \mathrm{id}}$, write $P_{\lambda,n}(a) \le B_{C_n}^{\mathrm{MCA}}(a)$ for its certified slope count and put

$$a^{\mathrm{low}}_{\lambda,n} = \max(\{a \in \{k_n + 1, \dots, n\} : P_{\lambda,n}(a) > B_n^*\} \cup \{k_n\}). \tag{1.12}$$

Define $a^{\mathrm{low}}_{\mathrm{id},n}$ by the same formula using the exact pole-converted identity lower bound $P_{\mathrm{id},n}$ from (2.3). A pole-admissible sequence is *identity-candidate*, relative to this fixed catalogue, when the empty, saturation, and tangent branches do not bind, $a^{\mathrm{low}}_{\mathrm{id},n} \ge k_n + 1$, and

$$\max_{\lambda \in \mathcal{L}^{\neq \mathrm{id}}} a^{\mathrm{low}}_{\lambda,n} \le a^{\mathrm{low}}_{\mathrm{id},n} \tag{1.13}$$

for all sufficiently large n; the maximum is k_n if the catalogue is empty. This definition uses only realized lower constructions and does not assume an upper envelope.

Conjecture 1.2 (Identity dominance on candidate smooth branches)**.** *Let $C_n = \mathrm{RS}_{\mathbb{F}_n}(D_n, k_n)$ be an official full-field, positive-density admissibly smooth multiplicative or Chebyshev twin-coset sequence of fixed rate, with $B_n^* \ge 1$. Suppose the rows are nonsaturated, the tangent branch is nonbinding, $Q_n > k_n B_n^*$, and (1.9) holds, and the sequence is identity-candidate relative to the fixed declared catalogue above. Then there are integer sequences $a_{-,n} < a_{+,n}$ in $\{k_n + 1, \dots, n\}$ and a family-uniform, non-oracular profile atlas whose specialization at $a_{+,n}$ is witness-exhaustive and has exact first-match payments summing to at most B_n^*. Together with the pole-converted identity lower certificate at $a_{-,n}$, these give*

$$B_{C_n}^{\mathrm{MCA}}(a_{-,n}) > B_n^*, \qquad B_{C_n}^{\mathrm{MCA}}(a_{+,n}) \le B_n^*, \tag{1.14}$$

and

$$a_{\pm,n} = k_n + 1 + g_{\mathrm{pole},n} n + o(g_n^* n) = k_n + 1 + g_n^* n + o(g_n^* n). \tag{1.15}$$

Consequently the first safe agreement is $k_n + 1 + g_n^* n + o(g_n^* n)$, and the official supremum satisfies (1.10). The same assertion applies to a circle presentation only through a proved agreement-preserving monomial equivalence.

We call a sequence satisfying this conclusion *identity-dominant*. Thus identity dominance is the predicted conclusion, not part of the lower-side hypothesis. Section 6 proves the numerical conclusion from a matching exact certificate bracket; the conjectural content is witness-exhaustive closure with the asserted payments and bracket.

These are conditional identity-candidate predictions, not unconditional frontiers: quotient, partial-occupancy, planted, image-collapse, tangent, or ray-saturation profiles may move the crossing. Thus the general predicted answer is the right crossing of the complete image-normalized profile envelope, after separately testing the empty-budget, tangent, and saturation branches.

No theorem here improves the safe threshold in [Cho26a]. What is proved here is the compiler and barrier calculus, including the conditional moving-scale implication and the pole-aware crossing reduction; what remains open is the Sidon payment, the residual projection compiler, the other algebraic payments, the four adjacent inequalities, and uniform smooth/circle closure. The preliminary Prize welcomes partial progress, but any claimed resolution must survive scientific peer review [PP26].

2. Exact certificate compiler

Retain the fields, code, challenge set, and budget of Section 1. A word $u \in \mathbb{F}^D$ is *explained* on $S \subseteq D$ if $u|_S = f|_S$ for some $f \in \mathbb{F}[X]_{<k}$; a pair is simultaneously explained if both words are, possibly by different polynomials.

For an integer agreement a, an MCA witness for a received pair $r = (r_0, r_1)$ is a triple $w = (\gamma, S, h)$ such that

$$\gamma \in \Gamma, \qquad |S| \geq a, \qquad (r_0 + \gamma r_1)|_S = h|_S, \quad \deg h < k,$$

and (r_0, r_1) is not simultaneously explained on S. Let $\mathcal{W}_r(a)$ be the set of witnesses and let $Z_r(a) = \{\gamma : (\gamma, S, h) \in \mathcal{W}_r(a)\}$. Thus

$$B^{\mathrm{MCA}}_{C,\Gamma}(a) = \max_{r \in (\mathbb{F}^D)^2} |Z_r(a)|.$$

The following endpoint conversion is [Cho26a, Theorem 1.4, Remark 4.4 and Lemma B.1], restated because a subcapacity formula alone loses the all-radius branch.

Proposition 2.1 (Exact staircase conversion)**.** *Let C be a proper $[n, k]$ MDS code. Then $a \mapsto B^{\mathrm{MCA}}_{C,\Gamma}(a)$ is nonincreasing,*

$$B^{\mathrm{MCA}}_{C,\Gamma}(a) = B^{\mathrm{MCA}}_{C,\Gamma}(k+1) \quad (0 \leq a \leq k+1), \qquad B^{\mathrm{MCA}}_{C,\Gamma}(n) = 1. \tag{2.1}$$

If $B^* = 0$, no radius is safe. If $B^{\mathrm{MCA}}_{C,\Gamma}(k+1) \leq B^*$, the complete official safe set is $[0, 1]$. Otherwise $B^* \geq 1$, and the first safe agreement

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

is well defined. Its complete real safe set is

$$\left[0, \frac{n - a^* + 1}{n}\right), \tag{2.2}$$

whose right endpoint is unsafe; the largest safe grid radius is $(n - a^*)/n$.

Proof. This is the cited endpoint theorem and its exact integer-to-radius conversion. □

For a support S, write $Q_S(X) = \prod_{x \in S}(X - x)$. If $m = |S|$ and $K \leq m$, the first $w = m - K$ nonleading coefficients define the locator-prefix map

$$\Phi_w(S) = (c_1(S), \ldots, c_w(S)) \in \mathbb{B}^w, \qquad Q_S = X^m + c_1(S)X^{m-1} + \cdots + c_m(S).$$

Pigeonholing its fibers gives a received word with at least $\lceil \binom{n}{m} |\mathbb{B}|^{-w} \rceil$ agreeing codewords of dimension K. For $K = k + 1$, a collision-aware simple-pole construction converts the list to MCA slopes. In a full-field row, with $q = |\mathbb{F}| > n$, its exact consequence at agreement a is

$$L_{\mathrm{id}}(a) = \left\lceil \binom{n}{a} |\mathbb{B}|^{-(a-k-1)} \right\rceil, \qquad P_{\mathrm{id}}(a) = \left\lceil \frac{L_{\mathrm{id}}(a)(q-n)}{q - n + k(L_{\mathrm{id}}(a) - 1)} \right\rceil \leq B^{\mathrm{MCA}}_C(a). \tag{2.3}$$

These exact lower mechanisms and their nontriviality checks are established in [Cho26a, Proposition 8.1 and Theorem 8.2]. We use their certified unsafe agreements below, but never interpret the next, smaller lower value as a safe bound. In particular, the first agreement at which

$P_{\mathrm{id}}(a) \le B^*$ is only an identity-route candidate; identifying it with the first safe agreement requires a closed upper atlas.

The following elementary conversion is essential when the challenge field is not so large that pole collisions are negligible.

Lemma 2.2 (Exact pole conversion)**.** *Let $B \ge 1$, $Q = q - n > 0$, and*

$$M_{Q,k}(L) = \left\lceil \frac{LQ}{Q + k(L-1)} \right\rceil \qquad (L \in \mathbb{Z}_{\ge 1}).$$

Then

$$M_{Q,k}(L) > B \iff L(Q - kB) > B(Q - k). \tag{2.4}$$

If $Q \le kB$, then $M_{Q,k}(L) \le B$ for every $L \ge 1$. If $Q > kB$, the least integer L for which $M_{Q,k}(L) > B$ is

$$T_{\mathrm{pole}}(B; Q, k) = \left\lfloor \frac{B(Q-k)}{Q - kB} \right\rfloor + 1. \tag{2.5}$$

Consequently the certified pole lower bound in (2.3) exceeds B exactly when

$$L_{\mathrm{id}}(a) \ge T_{\mathrm{pole}}(B; Q, k) \iff \binom{n}{a} |\mathbb{B}|^{-(a-k-1)} > T_{\mathrm{pole}}(B; Q, k) - 1. \tag{2.6}$$

Proof. Since B is integral, $\lceil x \rceil > B$ if and only if $x > B$; clearing the positive denominator gives (2.4). If $Q \le kB$, then for $L \ge 1$, $L(Q - kB) \le Q - kB \le B(Q - k)$, so the strict inequality cannot hold. If $Q > kB$, then also $Q > k$, and division by $Q - kB$ gives (2.5). The last equivalence uses $L_{\mathrm{id}}(a) = \left\lceil \binom{n}{a} |\mathbb{B}|^{-(a-k-1)} \right\rceil$. □

Definition 2.3 (Admissibly smooth row)**.** *An RS row is* admissibly smooth *for a declared profile schema if its public data include a finite specified family of polynomial folding maps $\pi : D \to D_\pi$, each complete and c_π-to-one in the sense fixed in Section 1, with c_π, every scale, and every occupancy parameter used by the schema included in those data. No folding map may be chosen after a heavy fiber or bad received line is observed.*

Definition 2.4 (Declared profile rule)**.** *Fix public row data $P = (\mathbb{B} \subseteq \mathbb{F}, D, k, \Gamma, a, B^*)$. A declared profile rule is a finite ordered list $\mathcal{A}(P) = (\mathcal{P}_1, \dots, \mathcal{P}_s)$ of predicates on pairs (r, w), where $r \in (\mathbb{F}^D)^2$ and $w \in \mathcal{W}_r(a)$. The predicates and their order are specified from the public data before the values $|Z_r(a)|$, a maximizing received line, or the truth of the target inequality are known. An instantiation is* auditable *if every predicate is given by explicit algebraic equalities, inequalities, rank conditions, field-of-definition conditions, or occupancy conditions involving the received line, the support locator, and the declared folding maps on D.*

A concrete atlas must display each cell's parameters, predicate, order, projection to affine slopes, and proved integer payment. Defining a final cell as a residual may establish set-theoretic coverage, but it does not pay that residual. Singleton predicates chosen from the computed bad set, or predicates involving the unknown numerator or target outcome, are excluded. Thus the theorem below is only the exact compiler; exhaustion and payment are the substantive work.

For a rule and received line, let $\mathcal{C}_i(r) = \{w \in \mathcal{W}_r(a) : \mathcal{P}_i(r, w)\}$ and let $Z_i(r)$ be its projection to Γ. The rule is *witness-exhaustive* for r if $\mathcal{W}_r(a) = \bigcup_i \mathcal{C}_i(r)$. Its first-match slope cells are

$$Z_i^\circ(r) = Z_i(r) \setminus \bigcup_{j<i} Z_j(r).$$

A *payment* is a proved integer inequality $|Z_i^\circ(r)| \le U_i(r)$, uniform when it holds for every received line. The profile-closure meta-conjecture uses the stronger notion of an *admissible algebraic payment*. Such a payment is a specialization of a row-family-uniform formula

$$|Z_i^\circ(r)| \le F_i(P; \Xi_{i,1}(r), \dots, \Xi_{i,t_i}(r)),$$

where the integer formula F_i, the source-side statistics $\Xi_{i,j}$, and a proof of the inequality are declared before specialization to a maximizing received line. Permitted statistics are

explicitly named algebraic data: parameter-space sizes, degrees, ranks, field-orbit sizes, folding occupancies, attained images and certified fiber moments of predeclared maps, and proved incidence multiplicities. An attained-image or fiber statistic may enter only through a separately stated row-uniform estimate.

Neither F_i nor any $\Xi_{i,j}$ may invoke, recover, or enumerate $Z_i^\circ(r)$, $Z_r(a)$, their cardinalities, an analogous list-codeword projection, the true numerator, a maximizing received datum, or the truth of the target comparison. Optimization or exhaustive enumeration over received data, witnesses, slopes, or agreeing codewords is likewise excluded. Thus the one-cell choice $U_1(r) = |Z_r(a)|$, and any disguised equivalent, is not admissible. A set-theoretic residual cell is allowed, but it still requires an independently proved admissible payment.

Theorem 2.5 (Exact first-match compiler)**.** *Suppose a single declared rule $\mathcal{A}(P)$, fixed from the public row data, is witness-exhaustive for every received line and has valid first-match payments satisfying $\sum_i U_i(r) \le U$. Then $B_{C,\Gamma}^{\mathrm{MCA}}(a) \le U$. In particular, $U \le B^*$ is a safe certificate.*

Proof. The sets $Z_i^\circ(r)$ are disjoint and their union is $Z_r(a)$. Therefore $|Z_r(a)| = \sum_i |Z_i^\circ(r)| \le \sum_i U_i(r) \le U$. Maximize over r. □

For an auxiliary list row and received word $u \in \mathbb{F}^D$, define

$$\Lambda_C(u,a) = \{c \in C : |\{x \in D : u_x = c_x\}| \ge a\}, \qquad B_C^{\mathrm{list}}(a) = \max_{u \in \mathbb{F}^D} |\Lambda_C(u,a)|.$$

A declared list rule classifies triples (u, c, S) and projects its cells to agreeing codewords. If $Y_i^\circ(u)$ are the first-match codeword cells and $|Y_i^\circ(u)| \le V_i(u)$ with $\sum_i V_i(u) \le V$, the same disjoint-union proof gives $B_C^{\mathrm{list}}(a) \le V$. This projection change is essential: an MCA ray or slope payment cannot substitute for a required agreeing-codeword payment, or conversely. For an atlas at agreement a, set

$$\mathcal{E}_{\mathcal{A}}(P,a) = \max_r \sum_i U_i(P,r;a), \tag{2.7}$$

where the maximum ranges over received lines in the MCA case and received words in the list case.

Meta-conjecture 2.6 (Target-exact profile-closure program)**.** *For each declared unbounded-length admissibly smooth multiplicative or Chebyshev row family and each public target rule $P \mapsto B^*(P)$, there is one finite, auditable algebraic atlas schema fixed uniformly over the family. For every proper row P in the family and every $k+1 \le a \le n$, its specialization is witness-exhaustive and has admissible algebraic payments satisfying*

$$\mathcal{E}_{\mathcal{A}}(P,a) \le B^*(P) \quad \Longleftrightarrow \quad B_{C,\Gamma}^{\mathrm{MCA}}(a) \le B^*(P). \tag{2.8}$$

The schema has finite syntax independent of the row length. It may depend on the declared family and target rule, but not on a literal encoding of one finite row or agreement. The circle analogue is asserted only after a proved agreement-preserving monomial equivalence.

The atlas must be frozen before the numerator is evaluated and may not encode the eventual answer. In particular, one schema, rather than a separately hard-coded lookup table, must work at every agreement in every row of its declared family. Thus meta-conjecture 2.6 is a falsifiable proposal for proof architecture: an escaping witness family or an unpayable first-match cell refutes a proposed instantiation. Independently of this architecture, conjecture 5.1 consists of four explicit inequalities for the true MCA or list numerators.

For reference, in a full-field row with $q = |\mathbb{F}| > n$ and $B^* \ge 1$, the identity route singles out

$$a_{\mathrm{id}}^* = \min\{a \in \{k+1, \dots, n\} : P_{\mathrm{id}}(a) \le B^*\}. \tag{2.9}$$

When the empty, saturation, and tangent branches do not bind and every lower edge in the fixed declared nonidentity catalogue is no later than the corresponding identity edge, this is the identity candidate in the sense of (1.13). Its preceding agreement is unconditionally unsafe by (2.3); safety at a_{id}^* still requires a closed upper atlas and remains conjectural both for the direct rows in conjecture 5.1 and for the asymptotic families in conjecture 1.2.

The proposed atlas has four payment groups. Algebraic routing must cover the endpoint and tangent branches, declared foldings with full and partial occupancy, quotient, remainder, and planted factors, and all descent, extension-orbit, rank-loss, sparse, and auxiliary cells. Here auxiliary means in particular correlated-agreement and shortened or interleaved-list residuals; orbit stabilizers, extension degrees, and named rank multiplicities remain in the payment. Primitive cells require estimates on the attained image together with max-fiber and incidence degrees. Every residual chart requires a distinct-slope payment for MCA or a distinct codeword payment for a list row. Finally, first-match deduplication must add the integer payments with a reserve below the target. In particular, for a declared c-to-one folding $\pi : D \to D'$, the public occupancy datum is $\nu_\pi(S) = (|S \cap \pi^{-1}(y)|)_{y \in D'}$; "primitive" means only that the witness escaped the preceding declared cells, not that its boundary map is surjective or flat.

For multiplicative rows the declared maps are $x \mapsto x^c$; for Chebyshev twin-coset rows they are T_c. Circle presentations enter only after a coordinate bijection and nonzero diagonal scaling give a support-preserving monomial equivalence with an RS row [HLP24, Cho26a]. The calculus itself is characteristic-independent, but the norm-one-torus presentation used here is in odd characteristic; in characteristic two we claim only statements for explicit RS rows and declared algebraic maps.

3. Primitive-fiber barriers

The primitive problem begins only after the earlier algebraic cells have been declared. Its normalization must refer to the full reference slice before first-match deletion. Let Ω^0 be such a nonempty finite slice, let $\Omega^\circ \subseteq \Omega^0$ be a residual leaf, and let $\Phi : \Omega^0 \to A$ be a boundary map. Write

$$I = \Phi(\Omega^0), \qquad M = |\Omega^0|, \qquad L = |I|, \qquad \overline{N} = \frac{M}{L},$$

and, for $y \in I$, put $F_y = \Omega^\circ \cap \Phi^{-1}(y)$, $f_y = |F_y|$, and $W = \sum_y f_y = |\Omega^\circ|$. Thus $\overline{N}$ is the mean fiber size of the reference slice, not the mean over the smaller residual image. This is the normalization used by the finite incidence payment and by the Gowers/cube implication in [Cho26a].

For an integer $r \geq 2$, define the image-normalized residual moment

$$\Gamma_r = \frac{1}{L} \sum_{y \in I} \left(\frac{f_y}{\overline{N}} \right)^r . \tag{3.1}$$

The elementary maximum bound is

$$\max_y f_y \leq \overline{N} (L \Gamma_r)^{1/r}. \tag{3.2}$$

Indeed, the largest summand in (3.1) is at most the whole sum. This already shows why an ordinary-moment conjecture strong enough to make $\log(L\Gamma_r)/r = o(n)$ would itself prove max-fiber flatness; the Gowers/cube step is needed for a weaker, low-energy input instead.

The exact support-fiber-to-slope interface also needs an incidence theorem. Let Z° be the first-match slope projection of the leaf, and let $\mathcal{P} = \{(u, v) \in (\Omega^\circ)^2 : \Phi(u) = \Phi(v)\}$. Suppose an incidence relation $\mathcal{I} \subseteq Z^\circ \times \mathcal{P}$ has degree at least $H \geq 1$ at every slope and degree at most J at every ordered pair.

Proposition 3.1 (Exact primitive incidence payment)**.** *With the preceding notation,*

$$|Z^\circ| \leq \left\lfloor \frac{JW}{H} \, \overline{N} (L \Gamma_r)^{1/r} \right\rfloor . \tag{3.3}$$

An independent direct slope bound may be intersected with this one.

Proof. Double counting gives $H|Z^\circ| \leq J \sum_y f_y^2$. Since $\sum_y f_y^2 \leq W \max_y f_y$, (3.2) proves the result; the left side is an integer. □

The quantities H and J are not cosmetic. A bound for $\max_y f_y$, or even for all same-fiber pairs, does not bound slopes until one proves how many such pairs a slope supplies and how many slopes a pair may represent.

For the full fiber distribution $\Omega^\circ = \Omega^0$, set $N_y = |\Phi^{-1}(y)|$, $\mu_y = N_y/M$, and

$$R = L \max_y \mu_y, \qquad \Gamma_r^{\mathrm{full}} = L^{r-1} \sum_{y \in I} \mu_y^r.$$

Here R is the maximum-fiber overhead relative to M/L, and $\Gamma_r^{\mathrm{full}} = 1$ for a uniform distribution.

Proposition 3.2 (Exact moment sandwich)**.** *For every integer $r \geq 2$,*

$$\frac{\log_2 \Gamma_r^{\mathrm{full}}}{r-1} \leq \log_2 R \leq \frac{\log_2 L + \log_2 \Gamma_r^{\mathrm{full}}}{r}. \tag{3.4}$$

Consequently $R \leq 2^\Delta$ follows from $\log_2 L + \log_2 \Gamma_r^{\mathrm{full}} \leq r\Delta$.

Proof. The inequalities $\sum_y \mu_y^r \leq (\max_y \mu_y)^{r-1}$ and $(\max_y \mu_y)^r \leq \sum_y \mu_y^r$ give $\Gamma_r^{\mathrm{full}} \leq R^{r-1}$ and $R^r \leq L\Gamma_r^{\mathrm{full}}$, respectively. □

The image I cannot silently be replaced by the formal codomain A. Let $A_0 = |A|$, extend both f_y and N_y by zero on $A \setminus I$, and put $\overline{N}^{\mathrm{amb}} = M/A_0$. The residual ambient moment is

$$\Gamma_{r,\mathrm{res}}^{\mathrm{amb}} = \frac{1}{A_0} \sum_{y \in A} \left(\frac{f_y}{\overline{N}^{\mathrm{amb}}} \right)^r = \left(\frac{A_0}{L} \right)^{r-1} \Gamma_r. \tag{3.5}$$

For the full reference distribution, define

$$\Gamma_{r,\mathrm{full}}^{\mathrm{amb}} = \frac{1}{A_0} \sum_{y \in A} \left(\frac{N_y}{\overline{N}^{\mathrm{amb}}} \right)^r = A_0^{r-1} \sum_{y \in A} \left(\frac{N_y}{M} \right)^r \geq 1. \tag{3.6}$$

The last inequality follows from the power-mean inequality and $\sum_y N_y/M = 1$. Thus an ambient residual-moment estimate implies its image-normalized residual counterpart, but the converse pays the exact factor $(A_0/L)^{r-1}$. At exponent scale one may discard this factor only after proving $\log(A_0/L) = o(n)$ at the relevant order.

Two distinct failures are encoded here. First, the reference map itself may have a small image I, in which case M/L is much larger than the formal mean M/A_0. Second, first-match deletion may concentrate the residual leaf inside a few fibers even when the reference image is large. Resetting the denominator to $|\Phi(\Omega^\circ)|$ after this deletion would change the conjecture and sever the link to the preceding payments. The reference-slice normalization keeps both losses visible: image collapse appears through A_0/L, and residual concentration appears in the moment Γ_r.

Proposition 3.3 (Ambient moment-order barrier)**.** *Suppose a prefix map has formal codomain $\mathbb{B}^w$, and normalize its full distribution over all $Q = |\mathbb{B}|^w$ formal values, including empty fibers. Any argument using only the upper half of proposition 3.2 and $\Gamma_{r,\mathrm{full}}^{\mathrm{amb}} \geq 1$ to force ambient maximum-fiber overhead at most 2^Δ must have*

$$r \geq \left\lceil \frac{w \log_2 |\mathbb{B}|}{\Delta} \right\rceil. \tag{3.7}$$

Proof. That route requires $\log_2 Q + \log_2 \Gamma_{r,\mathrm{full}}^{\mathrm{amb}} \leq r\Delta$. Discarding the nonnegative second term and using $\log_2 Q = w \log_2 |\mathbb{B}|$ gives the claim. □

We now state the missing analytic input at a general resolution scale. Let $\mathbb{B}_N$ be a finite field, let $T_N \subseteq \mathbb{B}_N$ consist of N distinct points, let $1 \leq d_N \leq N$, and choose nonzero weights $\omega_N(t) \in \mathbb{B}_N^\times$. The weighted Vandermonde boundary map is

$$\Phi_N(x) = \sum_{t \in T_N} x_t \omega_N(t) (1, t, \ldots, t^{d_N - 1})^{\mathsf{T}}. \tag{3.8}$$

Fix $0 < \alpha \leq 1/2$, take $\Omega_N^0 = \{x \in \{0,1\}^{T_N} : \sum_t x_t = m_N\}$, where $m_N/N \in [\alpha, 1-\alpha]$, and let $\Omega_N^\circ \subseteq \Omega_N^0$ be a residual leaf selected by one fixed declared rule. Set

$$I_N = \Phi_N(\Omega_N^0), \quad L_N = |I_N|, \quad \overline{N}_N = \frac{|\Omega_N^0|}{L_N}, \quad F_{N,y} = \Omega_N^\circ \cap \Phi_N^{-1}(y).$$

For $F \subseteq \mathbb{Z}^{T_N}$, its additive energy is $E(F) = |\{(x_1, x_2, x_3, x_4) \in F^4 : x_1 + x_2 = x_3 + x_4\}|$. Compute this energy in the torsion-free group $\mathbb{Z}^{T_N}$, and put $\Delta_{N,y} = E(F_{N,y})/|F_{N,y}|^3$, with value zero for an empty fiber. Let $s_N \to \infty$, with $s_N \le N$. For $q \ge 2$ and $\sigma > 0$, define

$$\Gamma^{\mathrm{sid}}_{N,q,\sigma;s} = \frac{1}{L_N} \sum_{\Delta_{N,y} < e^{-\sigma s_N}} \left(\frac{|F_{N,y}|}{\overline{N}_N} \right)^q. \tag{3.9}$$

Problem 3.4 (Primitive Sidon payment)**.** *For primitive locator-prefix leaves in a fixed, predeclared multiplicative or Chebyshev atlas, prove that there are integer orders $q_N \ge 2$ such that, for every fixed $\sigma > 0$, uniformly over the declared leaves,*

$$\frac{\log L_N}{q_N} = o(s_N), \qquad \Gamma^{\mathrm{sid}}_{N,q_N,\sigma;s} \le e^{o(s_N q_N)}. \tag{3.10}$$

The primary asymptotic scale is $s_N = N$. As a stronger localization problem one may take $N = n$ and $s_n = g_n^* n$, under the necessary assumption $g_n^* n \to \infty$; this stronger scale is not required for the $o(g_n^*)$ frontier once all aggregate certificate losses are uniformly $2^{o(n)}$.

The adjective primitive is an input to this problem, not a conclusion of an equidistribution theorem. In particular, the rule defining the residual leaf must be fixed before a heavy fiber is known, and the moment is normalized by the image of the full reference slice.

We import the following consequence of polynomial Balog–Szemerédi–Gowers and Boolean-quasicube difference growth from [Cho26a, Proposition 10.4].

Proposition 3.5 (Imported BSG–Boolean-cube consequence)**.** *There is an absolute constant C_0 such that, for every finite set T, every nonempty $F \subseteq \{0,1\}^T$, and every $K \ge 1$,*

$$E(F) \ge \frac{|F|^3}{K} \quad \Longrightarrow \quad |F| \le K^{C_0}. \tag{3.11}$$

Here energy is computed after embedding the Boolean cube in $\mathbb{Z}^T$.

Corollary 3.6 (Moving-scale primitive flatness)**.** *If (3.10) holds at a scale $s_N \to \infty$, then*

$$\max_{y \in I_N} |F_{N,y}| \le e^{o(s_N)} \overline{N}_N. \tag{3.12}$$

For $s_N = N$, this is [Cho26a, Theorem 10.5]; the additional observation is that the same implication is valid at the displayed moving scale.

Proof. If the conclusion fails, then along a subsequence there are $\eta > 0$ and fibers of size at least $e^{\eta s_N} \overline{N}_N$. The ordinary moment from (3.1), with $r = q_N$, is therefore at least $L_N^{-1} e^{\eta s_N q_N} \ge e^{3\eta s_N q_N/4}$ for large N, by $\log L_N / q_N = o(s_N)$. Choose $0 < \sigma < \eta/(4C_0)$. The low-energy contribution is at most $e^{\eta s_N q_N/4}$, so the high-energy contribution is at least $e^{\eta s_N q_N/2}$. Some high-energy fiber then satisfies $|F_{N,y}|/\overline{N}_N \ge e^{\eta s_N/2}$. Since $I_N = \Phi_N(\Omega_N^0)$, one has $L_N \le |\Omega_N^0|$ and hence $\overline{N}_N \ge 1$; this gives $|F_{N,y}| \ge e^{\eta s_N/2}$. The same fiber has normalized energy at least $e^{-\sigma s_N}$, so proposition 3.5, with $K = e^{\sigma s_N}$, gives $|F_{N,y}| \le e^{C_0 \sigma s_N} < e^{\eta s_N/4}$, a contradiction. □

Corollary 3.7 (Analytic-to-slope interface)**.** *Assume problem 3.4 and suppose the primitive algebraic analysis proves an incidence relation with parameters $H_N, J_N \ge 1$. If $W_N = |\Omega_N^\circ|$ and Z_N° is the realized slope set, then*

$$|Z_N^\circ| \le \frac{J_N W_N}{H_N} e^{o(s_N)} \overline{N}_N. \tag{3.13}$$

Proof. Use the incidence double count from proposition 3.1 and corollary 3.6 in place of (3.2). □

This implication does not solve problem 3.4, prove reference-image coverage, supply the incidence degrees, or close the profile ledger. At the coarse scale it yields only an $e^{o(N)}$ loss, which does not fit a finite margin of a few bits, but would be asymptotically sufficient for the sharp identity crossing if every other loss were also uniform and $e^{o(n)}$. For the optional stronger specialization $s_n = g_n^* n$, the energy cutoff itself must be $e^{-\sigma g_n^* n}$; changing only the moment order while retaining the coarse cutoff is not a valid moving-scale argument. The phrase "Gowers/cube"

here means BSG followed by Boolean-quasicube difference growth, not a Gowers-uniformity norm or cube-counting theorem.

4. Ray-projection barriers

A same-fiber pair is often called a shift pair. Max-fiber control gives the support-level estimate $\sum_y \binom{N_y}{2} \le N_{\max} M/2$, and hence at most $\kappa M^2/(2L)$ when $N_{\max} \le \kappa M/L$. This does not supply the ray incidence required in proposition 3.1.

The next incidence bound identifies exactly what a one-parameter moving-root proof uses. It is stated for a general-dimensional family to show how quickly that argument deteriorates.

Theorem 4.1 (Projective moving-root incidence)**.** *Let $d \ge 1$, let $V \le \mathbb{F}_q[X]$ be a d-dimensional polynomial subspace, and let $G \subseteq D$ be its common zero set, of size g. Let $h \ge 1$, and suppose $\mathcal{A} \subseteq \mathbb{P}(V)$ and every $[f] \in \mathcal{A}$ has at least h zeros in $D \setminus G$. Then*

$$|\mathcal{A}|h \le (n-g)\frac{q^{d-1}-1}{q-1}. \tag{4.1}$$

For the affine analogue, let $\mathcal{F} = f_0 + V$ be an affine polynomial space with $\dim V = d$, let G be the common zero set in D of every member of $\mathcal{F}$, of size g, and let $\mathcal{A} \subseteq \mathcal{F}$. If every member of $\mathcal{A}$ has at least h zeros in $D \setminus G$, then $|\mathcal{A}|h \le (n-g)q^{d-1}$.

Proof. Count pairs $([f], x)$ with $[f] \in \mathcal{A}$, $x \in D \setminus G$, and $f(x) = 0$. Every projective member contributes at least h pairs. At a fixed $x \notin G$, evaluation is a nonzero linear functional on V, so its projective kernel has $(q^{d-1}-1)/(q-1)$ points. This proves (4.1). For the affine statement, evaluation at a point outside G is either a nonconstant affine linear functional, whose zero set has q^{d-1} points, or a nonzero constant, whose zero set is empty. The same double count applies. □

Corollary 4.2 (Moving-root pencil)**.** *Let $A, B \in \mathbb{F}_q[X]$ be linearly independent, let $G \subseteq D$ be their common zero set, of size g, and consider $L_{[s:t]} = sA + tB$. If every counted parameter has at least h zeros in $D \setminus G$, then there are at most $\lfloor (n-g)/h \rfloor$ counted parameters. If every counted member is squarefree, split over D, and has degree $j > g$, the bound is $\lfloor (n-g)/(j-g) \rfloor$.*

Proof. Use theorem 4.1 with $d = 2$. □

The pencil bound is strong because a point outside the base locus chooses at most one projective parameter. In dimension d, it instead chooses a projective hyperplane, and the factor in (4.1) is its exact one-coordinate worst-case multiplicity. Thus for large d the raw bound is essentially $(n-g)q^{d-2}/h$. A useful high-dimensional compiler must exploit several roots, restricted factorization or rank, or a lower-dimensional parametrization of the rays actually realized by MCA witnesses.

Proposition 4.3 (Dimension diagnostic for the declared correction spaces)**.** *At the four candidate adjacent agreements of Table 1, assume the declared prefix reduction parametrizes a residual chart by the full affine correction space $A \in \mathbb{F}[X]_{\le \omega - w - 1}$, where $\omega = n - a^*$ and $w = a^* - K$. Hence its affine dimension is $\omega - w = n - 2a^* + K$, equal respectively to*

$$913633, \qquad 913634, \qquad 913681, \qquad 913682. \tag{4.2}$$

Consequently the pencil case of theorem 4.1 cannot pay this declared full correction space without a further structural reduction.

Proof. Under the stated parametrization there is one free coefficient in each degree from zero through $\omega - w - 1$. Substitution of $n = 2097152$, the agreements and the values $K = k + 1$ for MCA and $K = k$ for the auxiliary list rows gives the four integers. A family of these dimensions is not a two-dimensional projective subspace or a one-dimensional affine pencil. □

The proposition verifies the correction-space arithmetic conditional on the declared parametrization. The release source package contains `verify_profile_barriers.py`, which independently recomputes (4.2). The lower-reserve artifacts in [Cho26b] record the separate unsafe-side inputs; they are not evidence for witness exhaustion. Proving that the declared

TABLE 1. Proved unsafe agreements, adjacent candidate agreements, and their candidate threshold intervals. Only the first and third rows are MCA rows.

row and target	proved unsafe a_0	candidate $a^* = a_0 + 1$	candidate threshold interval
KoalaBear MCA, 2^{-128}	1116047	1116048	$[0, 981105/2097152)$
KoalaBear list, 2^{-128}	1116046	1116047	$[0, 490553/1048576)$
Mersenne–31 MCA, 2^{-100}	1116023	1116024	$[0, 981129/2097152)$
Mersenne–31 list, 2^{-100}	1116022	1116023	$[0, 490565/1048576)$

charts exhaust the residual witnesses is part of problem 5.2, not a consequence of the dimension count.

Problem 4.4 (Balanced-core ray payment)**.** *Every residual balanced-core cell in the benchmark rows admits an explicit witness-exhaustive decomposition into structured charts such that each chart is either a projective pencil covered by corollary 4.2, a controlled finite union of such pencils, or a higher-dimensional chart with a proved distinct-slope payment. After first-match deduplication and exact incidence multiplicities, the total balanced-core payment fits the row allocation in problem 5.2.*

This is a slope statement, not merely a factorization statement. It is falsified by a single residual chart whose actual slope projection exceeds the allocated budget and which cannot be routed to an earlier paid cell.

5. FINITE BENCHMARK ROWS

We now record four adjacent rows that make the preceding obstructions quantitative. Set

$$n = 2^{21} = 2097152, \qquad k = 2^{20} = 1048576,$$

and define

$$p_{\mathrm{KB}} = 2^{31} - 2^{24} + 1 = 2130706433, \qquad p_{\mathrm{M}} = 2^{31} - 1 = 2147483647.$$

Let $D_{\mathrm{KB}} \subseteq \mathbb{F}_{p_{\mathrm{KB}}}^{\times}$ be the multiplicative subgroup of order n, and put $C_{\mathrm{KB}} = \mathrm{RS}_{\mathbb{F}_{p_{\mathrm{KB}}^6}}(D_{\mathrm{KB}}, k)$. Let D_{M} be the Chebyshev twin-coset domain of size n specified in [Cho26a, Cho26b], and put $C_{\mathrm{M}} = \mathrm{RS}_{\mathbb{F}_{p_{\mathrm{M}}^4}}(D_{\mathrm{M}}, k)$. Their integer budgets are

$$B^*_{\mathrm{KB}} = \left\lfloor p_{\mathrm{KB}}^6 / 2^{128} \right\rfloor = 274980728111395087, \qquad B^*_{\mathrm{M}} = \left\lfloor p_{\mathrm{M}}^4 / 2^{100} \right\rfloor = 16777215. \tag{5.1}$$

The Mersenne row is an auxiliary 2^{-100} comparison; at target 2^{-128} this challenge field has zero integer budget.

The list numerator and its separate first-match codeword projection were defined in Section 2. The two list rows below are auxiliary inputs, not Proximity Prize MCA thresholds.

The word "candidate" is essential. The exact prefix construction proves that each a_0 is unsafe, and hence that every smaller agreement is unsafe. At a^*, that particular lower route drops below the budget. This is not an upper bound for other witnesses.

For completeness, the exact lower-route numerators at a_0, in the order of Table 1, are

$$\begin{aligned} &13863474105832785 2652, \qquad 15770251823342597 5347, \\ &4281388998575706, \qquad 4870025984688527. \end{aligned} \tag{5.2}$$

Each is strictly larger than its row budget in (5.1). The MCA values use the dimension-$k+1$ locator-prefix list and the simple-pole conversion; the list values use dimension k directly. In the two MCA rows the pole-separation condition is satisfied, so no collision factor is lost at these displayed agreements [Cho26a, Proposition 8.6]; the auxiliary values are recorded in [Cho26b], and `verify_profile_barriers.py` in the release source package independently recomputes all four displayed integers.

The location of the one-step edge has a simple exact explanation. For the unrounded prefix mean, put $\mu_K(a) = \binom{n}{a} p^{-(a-K)}$, with $K = k+1$ in the MCA route and $K = k$ in the auxiliary

TABLE 2. Candidate-agreement diagnostics. The necessary order refers only to the ceiling-normalized ambient-moment route of proposition 3.3.

row	$P_{\mathrm{cand}}(a^*)$	margin bits	necessary order	extension headroom
KoalaBear MCA	57198030366	22.1969	94196	4807520
KoalaBear list	65065153468	22.0109	94991	4226236
Mersenne–31 MCA	1752700	3.2589	641594	9
Mersenne–31 list	1993678	3.0730	680397	8

list route. Then $\mu_K(a)/\mu_K(a+1) = p(a+1)/(n-a)$. In the four rate-one-half rows this is a drop of roughly 31 bits in one agreement step. The identity route therefore identifies the microscopic candidate scale, but the ratio supplies no safe-side bound and does not rule out a later nonidentity crossing.

Conjecture 5.1 (Benchmark adjacent safe inequalities)**.** *For the two MCA rows,*

$$B^{\mathrm{MCA}}_{C_{\mathrm{KB}}}(1116048) \le 274980728111395087, \qquad B^{\mathrm{MCA}}_{C_{\mathrm{M}}}(1116024) \le 16777215. \tag{5.3}$$

For the two auxiliary list rows,

$$B^{\mathrm{list}}_{C_{\mathrm{KB}}}(1116047) \le 274980728111395087, \qquad B^{\mathrm{list}}_{C_{\mathrm{M}}}(1116023) \le 16777215. \tag{5.4}$$

These are direct, falsifiable statements; they do not define their own proof method. If either MCA inequality is proved, the corresponding preceding unsafe certificate and monotonicity make the displayed a^* the first safe agreement. The exact integer-to-radius conversion then gives the associated half-open interval in Table 1. The same observation applies to the auxiliary list staircase.

To fix the notation used in the diagnostics, put

$$L_K(a) = \left\lceil \binom{n}{a} p^{-(a-K)} \right\rceil, \qquad P_{\mathrm{cand}}(a) = \begin{cases} \left\lceil \dfrac{L_{k+1}(a)(q-n)}{q-n+k(L_{k+1}(a)-1)} \right\rceil, & \text{MCA}, \\ L_k(a), & \text{auxiliary list}, \end{cases}$$

using the base prime p and challenge-field order q of the row. With $K = k+1$ for MCA and $K = k$ for list, let

$$\Delta_{\mathrm{ceil}} = \log_2 \frac{B^*}{P_{\mathrm{cand}}(a^*)}, \qquad w = a^* - K, \qquad H_{\mathrm{ext}} = \left\lfloor \frac{p^w B^*}{\binom{n}{a^*}} \right\rfloor.$$

The candidate values and remaining margins are as follows. They are reproduced by `verify_profile_barriers.py` in the release source package.

The periodic lower-family audit has deliberately limited coverage. In the repository labels, Gfloor and Gceil use $m = \lfloor a/c \rfloor$ and $m = \lceil a/c \rceil$, respectively; Rem adds the uncovered remainder, and Plant tests a planted quotient core. Gate G6 of the exact verifier scans all twenty-one dyadic scales $c = 2^j$, $0 \le j \le 20$, and these four lower families at the historical packet pairs: $(1116043, 1116044)$ for KoalaBear MCA, $(1116021, 1116022)$ for Mersenne–31 MCA, and the unchanged list pairs in Table 1. Gate G7 separately recomputes the moved MCA identity pairs in that table and the named Mersenne–31 MCA Gceil watch item at $c = 2048$, which covers agreement 1116160 and has exact mass $12769758 < 16777215 = B^*$, a deficit of 0.3938 bits. It does not perform a full twenty-one-scale G6 rescan at both moved MCA pairs and does not certify safety. The precise executable is `experimental/scripts/verify_frontier_adjacent_v13_rows.py` at pinned repository commit `e190193cebce`. The release source package contains a verbatim copy at that path and the four input packets under `experimental/data/certificates/frontier-adjacent/` [Cho26b]. These checks veto only the lower constructions they actually enumerate; they are neither an upper payment nor an exhaustion theorem.

There is also a useful but nonrigorous scale check. If M objects were allocated independently and uniformly among L realized fibers, with mean $\mu = M/L$, the Poisson extreme-value factor x would be predicted by

$$\mu(x \log x - x + 1) \approx \log L.$$

Under the additional full-image assumption, the four rows in table order give predicted maximum-to-mean factors approximately $1.0071, 1.0067, 2.54, 2.43$, while their target-to-mean factors are approximately $4.81\cdot10^6, 4.23\cdot10^6, 9.57, 8.42$. Thus the random-fiber model favors all four primitive inequalities, most strongly in the KoalaBear rows. It is not evidence against effective-image collapse or algebraic concentration and is never used in a safe-side proof.

The remaining safe-side obligations are not interchangeable. Algebraic routing must cover the dense tangent/common-line cell, the full folded and periodic union, and all field, rank, extension, sparse, correlated-agreement, shortened, and interleaved-list cells, with orbit stabilizers and every degree and rank-loss multiplicity retained. Primitive leaves require a finite image-normalized estimate and the exact incidence degrees H, J. Every residual chart then needs a slope payment for MCA or a codeword payment for a list row. Finally, the compiler must deduplicate and add all payments inside the exact integer reserve.

Here the ceiling-normalized margin is the quantity Δ_{ceil} defined above. A factor n^C costs $21C$ bits, so an unspecified polynomial loss cannot decide any Mersenne adjacent row and can consume the entire KoalaBear margin even for modest C. Applying proposition 3.3 at prefix depths $67471, 67471, 67447, 67447$ gives the four necessary orders in Table 2. This diagnoses only the ceiling-normalized ambient route; it is not a lower bound on every max-fiber theorem, and a smaller attained image creates an additional loss.

The extension column has another exact interpretation. The descended extension charts in the current ledger have prefix-normalized cost $\Delta_0 p^e$, where p is the base-field order, e is the chart dimension, and $\Delta_0 \geq 1$ is a degree factor. For each row its displayed headroom is H_{ext} as defined above. The largest displayed headroom is below 2^{23}, while both base primes exceed 2^{30}. Thus no positive-dimensional chart of this form fits: necessarily $e = 0$, and then Δ_0 must not exceed the corresponding headroom. This is a useful veto, not a payment for all extension witnesses.

Problem 5.2 (Finite row-sharp closure)**.** *For either MCA row of Table 1, display a witness-exhaustive declared atlas at the candidate a^* and prove payments whose exact sum is at most B^*. In particular, the proof must:*

(a) *pay or witness-preservingly route the dense tangent/common-line cell; the full quotient, remainder, planted, periodic, and partial-occupancy union; every field-descent and extension orbit with its stabilizer and degree; every rank-loss multiplicity; and all sparse, correlated-agreement, shortened, and interleaved-list residuals;*

(b) *prove realized-image primitive-fiber estimates and the incidence degrees in proposition 3.1;*

(c) *prove a residual compiler for distinct affine slopes or projective rays in MCA, including the balanced-core payment of problem 4.4;*

(d) *deduplicate by first match and add every integer payment, rounding, degree, and multiplicity with a reserve that keeps the total at most B^*.*

For either auxiliary list row, the corresponding closure problem uses first-match codeword projections throughout: the balanced core requires a direct codeword payment, and the ray-incidence item is inapplicable.

A counterexample to (5.3) is a received line with more than the displayed number of bad slopes; a counterexample to (5.4) is a received word with more than the displayed number of agreeing codewords. A counterexample to a proposed profile proof may be weaker: one unpaid first-match cell, a collapsed realized image, or an underestimated pair-to-slope multiplicity is enough to invalidate that proof without deciding the true inequality. Conversely, proving only problem 3.4 and problem 4.4 does not close a row; all four groups in problem 5.2 remain part of the certificate.

6. Target-aware asymptotic criterion

The asymptotic formula is best understood as a crossing calculation that a future certificate must realize on both sides. Let $C_n = \mathrm{RS}_{\mathbb{F}_n}(D_n, k_n)$, where $D_n \subseteq \mathbb{B}_n \subseteq \mathbb{F}_n$, and assume $k_n/n \to \rho \in (0,1)$. Throughout this section the challenge is the official full field, $\Gamma_n = \mathbb{F}_n$, and

$\mathbb{B}_n$ is the designated generated coefficient-and-folding field fixed in Section 1; a sampled challenge or an arbitrary enlargement of $\mathbb{B}_n$ is not covered by the formula below. On the nonzero-budget branch, put

$$B_n(a) = B_{C_n}^{\mathrm{MCA}}(a), \qquad \rho_n = \frac{k_n}{n}, \qquad \rho_n^+ = \frac{k_n+1}{n}, \qquad \beta_n = \log_2|\mathbb{B}_n|,$$
$$q_n = |\mathbb{F}_n|, \qquad Q_n = q_n - n, \qquad B_n^* = \lfloor \varepsilon_n q_n \rfloor \geq 1, \qquad \tau_n = \frac{1}{n}\log_2 B_n^*.$$

We abbreviate the corresponding official quantities from problem 1.1 by $\delta^*_{n,\mathrm{off,sup}}$ and $\delta^*_{n,\mathrm{off,grid}}$. Let $H_2(x) = -x\log_2 x - (1-x)\log_2(1-x)$. For $r \in (0,1)$ and $b \geq 0$, define

$$g^*(r,b) = \sup\left(\{g \in [0, 1-r] : H_2(r+g) \geq bg\} \cup \{0\}\right), \qquad g_n^* = g^*(\rho_n, \beta_n).$$

The exact locator-prefix mean for the dimension-k_n+1 lower route is

$$\overline{N}_{1,n}(a) = \binom{n}{a}|\mathbb{B}_n|^{-(a-k_n-1)}. \tag{6.1}$$

For an integer a, put $g = (a-k_n-1)/n$. Stirling's formula gives, uniformly away from rates zero and one,

$$\frac{1}{n}\log_2 \overline{N}_{1,n}(a) = H_2(\rho_n^+ + g) - \beta_n g + O\left(\frac{\log n}{n}\right). \tag{6.2}$$

The pole conversion now imposes an additional exact branch condition. If $Q_n = 0$, there is no off-domain pole. If $0 < Q_n \leq k_n B_n^*$, lemma 2.2 says that the certified simple-pole lower bound never exceeds B_n^*, regardless of the list size; this route then supplies no unsafe certificate. Suppose henceforth that $Q_n > k_n B_n^*$, and put

$$T_{\mathrm{pole},n} = T_{\mathrm{pole}}(B_n^*; Q_n, k_n), \qquad \tau_n^{\mathrm{pole}} = \frac{1}{n}\log_2 T_{\mathrm{pole},n}, \qquad A_n = \frac{Q_n - k_n}{Q_n - k_n B_n^*} \geq 1.$$

Since $B_n^* A_n < T_{\mathrm{pole},n} \leq B_n^* A_n + 1$,

$$0 < \tau_n^{\mathrm{pole}} - \tau_n - \frac{1}{n}\log_2 A_n \leq \frac{1}{n}\log_2\left(1 + \frac{1}{B_n^* A_n}\right) \leq \frac{1}{n}. \tag{6.3}$$

Thus $\tau_n = o(1)$ alone does not control the pole-adjusted level. The conditions $Q_n > k_n B_n^*$, $\tau_n = o(1)$, and $n^{-1}\log_2 A_n = o(1)$ imply $\tau_n^{\mathrm{pole}} = o(1)$. The stronger collision-nonbinding condition $k_n B_n^*/Q_n = o(1)$ gives $\tau_n^{\mathrm{pole}} = \tau_n + O(1/n)$; the error need not be $o(1/n)$ when B_n^* remains bounded.

The exact test for this certified pole lower bound to exceed B_n^* is $\overline{N}_{1,n}(a) > T_{\mathrm{pole},n} - 1$. Because $T_{\mathrm{pole},n} \geq 2$,

$$0 \leq \frac{1}{n}\log_2 \frac{T_{\mathrm{pole},n}}{T_{\mathrm{pole},n} - 1} \leq \frac{1}{n}. \tag{6.4}$$

It is therefore convenient to use the pole-adjusted analytic crossing

$$g_{\mathrm{pole},n} = \sup\left(\left\{g \in [0, 1-\rho_n^+] : H_2(\rho_n^+ + g) - \beta_n g \geq \tau_n^{\mathrm{pole}}\right\} \cup \{0\}\right). \tag{6.5}$$

For a real ξ, let $g_{\mathrm{pole},n}(\xi)$ denote the same right crossing with τ_n^{pole} replaced by $\tau_n^{\mathrm{pole}} + \xi$.

Lemma 6.1 (Stability of the pole-adjusted crossing)**.** *Assume $\rho_n \to \rho \in (0,1)$, $\beta_n \to \infty$, $Q_n > k_n B_n^*$, and $\tau_n^{\mathrm{pole}} = o(1)$. For every sequence $0 \leq \eta_n = o(1)$, uniformly for $|\xi| \leq \eta_n$,*

$$g_{\mathrm{pole},n}(\xi) = g_n^* + O_\rho\left(\frac{\tau_n^{\mathrm{pole}} + \eta_n + n^{-1}}{\beta_n}\right) = g_n^* + o(g_n^*). \tag{6.6}$$

Consequently a uniform multiplicative certificate loss $2^{\eta_n n}$, with $\eta_n = o(1)$, changes the crossing by $o(g_n^*)$.

Proof. Uniformly for r in a compact subinterval of $(0,1)$, the positive right zero of $H_2(r+g)-\beta g$ is $\Theta_\rho(\beta^{-1})$. On this scale its derivative in g is $-\beta+O_\rho(1)$. The mean-value theorem shows that changing the right side from zero to $\tau_n^{\mathrm{pole}}+\xi$ moves the crossing by $O_\rho((\tau_n^{\mathrm{pole}}+|\xi|)/\beta_n)$. Replacing ρ_n by $\rho_n^+=\rho_n+n^{-1}$ moves it by $O_\rho((n\beta_n)^{-1})$. Since $g_n^*=\Theta_\rho(\beta_n^{-1})$, the total error is $o(g_n^*)$. □

Theorem 6.2 (Certificate bracket and pole-adjusted crossing)**.** *Assume $Q_n > k_nB_n^*$. Suppose there are agreements $a_{-,n} < a_{+,n}$ in $\{k_n+1,\dots,n\}$ for which an exact lower certificate proves $B_n(a_{-,n}) > B_n^*$ and a closed exact upper certificate proves $B_n(a_{+,n}) \le B_n^*$. If*

$$a_{\pm,n} = k_n + 1 + g_{\mathrm{pole},n}n + o(n), \tag{6.7}$$

then the first safe agreement satisfies $a_n^* = k_n+1+g_{\mathrm{pole},n}n+o(n)$, and in a nonsaturated row the largest safe grid radius and official real supremum are both

$$1-\rho_n-g_{\mathrm{pole},n}+o(1). \tag{6.8}$$

If the two errors in (6.7) are $o(g_n^*n)$, $\tau_n^{\mathrm{pole}}=o(1)$, and $\beta_n=o(n)$, then

$$g_{\mathrm{pole},n} = g_n^* + o(g_n^*), \qquad \delta_{n,\mathrm{off,sup}}^* = 1-\rho_n-g_n^*+o(g_n^*). \tag{6.9}$$

Proof. Monotonicity places the first safe agreement in $(a_{-,n}, a_{+,n}]$. By proposition 2.1, in a nonsaturated row

$$\delta_{n,\mathrm{off,grid}}^* = 1-\frac{a_n^*}{n}, \qquad \delta_{n,\mathrm{off,sup}}^* = 1-\frac{a_n^*-1}{n},$$

which proves (6.8). For the sharper statement, $D_n \subseteq \mathbb{B}_n$ and $|D|_n = n$ imply $\beta_n \ge \log_2 n$. Apply lemma 6.1; since $\beta_n = o(n)$, one has $g_n^*n\to\infty$, so the endpoint discrepancy $1/n$ is $o(g_n^*)$. □

Corollary 6.3 (Explicit pole-adjusted crossing)**.** *Fix $0<\rho_0<1/2$, and suppose $\rho_n\in[\rho_0,1-\rho_0]$, $\beta_n\to\infty$, and $\tau_n^{\mathrm{pole}}=o(1)$. Uniformly in this range,*

$$g_n^* = \frac{H_2(\rho_n)}{\beta_n - H_2'(\rho_n)} + O_{\rho_0}(\beta_n^{-3}), \tag{6.10}$$

and

$$g_{\mathrm{pole},n} = \frac{H_2(\rho_n^+)-\tau_n^{\mathrm{pole}}}{\beta_n - H_2'(\rho_n^+)} + O_{\rho_0}(\beta_n^{-3}). \tag{6.11}$$

In the collision-nonbinding pole regime, the coefficient-and-folding field sets the leading displacement from capacity; the challenge field enters through the target budget and the pole-conversion condition.

Proof. Taylor expansion gives $H_2(\rho+g)=H_2(\rho)+H_2'(\rho)g+O_{\rho_0}(g^2)$. The root equation implies $g=O(\beta^{-1})$; substituting this estimate back and dividing by $\beta-H_2'(\rho)\asymp\beta$ gives (6.10). The same calculation at level τ_n^{pole}, with ρ_n^+, gives (6.11). A sharper expansion may retain the quadratic entropy term. □

Corollary 6.4 (Raw-budget form under subexponential pole conversion)**.** *Suppose $Q_n > k_nB_n^*$, $\tau_n=o(1)$, $n^{-1}\log_2 A_n=o(1)$, and $\beta_n=o(n)$. A sharp certificate bracket as in theorem 6.2 then gives*

$$\delta_{n,\mathrm{off,sup}}^* = 1-\rho_n-g^*(\rho_n,\log_2|\mathbb{B}_n|)+o(g_n^*). \tag{6.12}$$

In particular, $k_nB_n^*/Q_n = o(1)$ is a sufficient collision-nonbinding hypothesis for the pole-conversion part of this statement.

Proof. Equation (6.3) gives $\tau_n^{\mathrm{pole}}=o(1)$; now apply lemma 6.1 and theorem 6.2. □

The pole condition is a genuine restriction, not a technical restatement of subexponential budget. For example, if $\varepsilon_n=\varepsilon>0$, $k_n\sim\rho n$, Q_n/q_n is bounded away from zero, and $B_n^*\sim\varepsilon q_n$, then $k_nB_n^*/Q_n=\Theta(n)$; hence pole admissibility eventually fails. Such a sequence requires a target shrinking on the scale $1/k_n$ or a different unsafe mechanism before an identity-pole bracket can be proposed.

Because $D_n \subseteq \mathbb{B}_n$ has n distinct points, $\beta_n \geq \log_2 n$, and corollary 6.3 gives $g_n^* \asymp 1/\beta_n$. When $\beta_n = o(n)$, the radius correction tends to zero while $g_n^* n \asymp n/\beta_n$ diverges. A loss $2^{\eta_n n}$, with $\eta_n = o(1)$, shifts the crossing by $O(\eta_n/\beta_n) = o(g_n^*)$. Thus a uniform aggregate $e^{o(n)}$ loss is sufficient for the sharper identity displacement. This fact does not help with the finite margins in Section 5, where an unspecified polynomial factor may exceed the entire reserve. The stronger $e^{o(g_n^* n)}$ control supplied by a moving-scale argument gives finer localization but is not necessary for (6.12).

The theorem is a reduction, not the missing upper bound. A full profile envelope for a received line is the sum of actual first-match payments, normalized at their realized images and followed through their actual slope projections. The identity-candidate condition is entirely lower-side and supplies no upper bound. If a family-uniform closed atlas supplies the matching upper certificate at (6.5), the sequence is identity-dominant in the sense following conjecture 1.2. Quotient coefficient fields, partial occupancy, planted factors, image collapse, and ray saturation may move the crossing; then g_n^* is only the identity-profile candidate.

Problem 6.5 (Uniform asymptotic closure)**.** *For official full-field positive-density multiplicative and Chebyshev twin-coset sequences, and for circle presentations proved monomially equivalent to them, construct an explicit family-uniform, non-oracular, witness-exhaustive declared atlas and certificate bracket of width $o(n)$ around the right crossing of the full image-normalized envelope. In a pole-admissible identity-candidate branch with subexponential raw budget and pole-conversion factor, prove a uniform aggregate $2^{o(n)}$ loss and a bracket of width $o(g_n^* n)$, thereby establishing identity dominance. The proof must establish the Sidon payment of problem 3.4, the residual support-to-slope incidence, the balanced-core payment of problem 4.4, all remaining algebraic payments, and a reserve showing that their sum stays below the target.*

If this problem is solved in the stated pole-admissible identity-candidate window, corollary 6.4 gives the precise formula $\delta^*_{n,\mathrm{off},\sup} = 1-\rho_n-g^*(\rho_n, \log_2|\mathbb{B}_n|)+o(g_n^*)$. If $Q_n = 0$ or $0 < Q_n \leq k_n B_n^*$, this pole lower route gives no such conclusion. A zero integer budget gives an empty safe set, safety at $k_n + 1$ gives official threshold 1, and a binding tangent profile gives the linear-budget edge; these branches must be tested first.

7. Conclusion

The paper separates four stable numerical conjectures from the conjectural profile architecture that might prove them. The exact compiler identifies where primitive realized-image bounds, residual slope or codeword payments, algebraic routing, and first-match add-back must enter; the incidence and moment results show why none may be omitted, while the non-oracular payment condition prevents the architecture from encoding the desired answer. Known nonidentity attacks show why the identity-candidate hypothesis must test every explicitly declared nonidentity lower profile, while identity dominance remains the conclusion to be proved. The exact pole conversion makes pole admissibility necessary for its stated unsafe route; another lower mechanism may behave differently. No theorem here improves the safe MCA threshold of [Cho26a]; proving an MCA inequality in conjecture 5.1 would give an exact adjacent threshold, while problem 6.5 would yield the pole-adjusted identity formula and, under subexponential raw budget and pole-conversion factor, its simpler g^*-form. Any circle transfer still requires a proved monomial equivalence.

Acknowledgments. The author thanks Scott Hughes for contributions to the frontier-adjacent verification materials and the contributors to the `rs-mca` repository for experimental audits and discussions that informed the finite-ledger program. The release archive preserves the cited file paths and pinned commit so that the underlying contribution history remains auditable.

AI-use disclosure. OpenAI ChatGPT/Codex was used substantively to draft and restructure prose, propose and audit proof steps, check calculations, compare source versions, and edit LaTeX. The author supplied the research program, source manuscripts, repository materials, mathematical choices, and corrections, and retains sole responsibility for correctness and originality.

References

[ABF26] G. Arnon, D. Boneh, and G. Fenzi, Open problems in list decoding and correlated agreement, Cryptology ePrint Archive, Paper 2026/680, 2026, https://eprint.iacr.org/2026/680.

[BCHKS26] E. Ben-Sasson, D. Carmon, U. Haböck, S. Kopparty, and S. Saraf, On proximity gaps of Reed–Solomon codes, in *Proc. 58th ACM Symposium on Theory of Computing (STOC '26)*, 2026, 1157–1167, https://doi.org/10.1145/3798129.3800827; extended version ECCC TR25-169 and ePrint 2025/2055.

[BCIKS20/23] 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, https://doi.org/10.1145/3614423.

[CS25] E. Crites and A. Stewart, On Reed–Solomon proximity gaps conjectures, Cryptology ePrint Archive, Paper 2025/2046, 2025, https://eprint.iacr.org/2025/2046.

[Cho26a] P. Chojecki, Shortening bounds for Reed–Solomon MCA, companion manuscript, version 9.2, July 2026, public source at repository snapshot 9908454995f3.

[Cho26b] P. Chojecki, `rs-mca`: proof development and exact certificate ledgers, source repository, pinned snapshot e190193cebce, 2026, https://github.com/przchojecki/rs-mca, accessed July 19, 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, 2026, https://arxiv.org/abs/2607.10572.

[PP26] Ethereum Foundation, The Proximity Prize: $1M Reed–Solomon Challenge, preliminary version, 2026, https://proximityprize.org/, accessed July 19, 2026.

[Hab25] U. Haböck, A note on mutual correlated agreement for Reed–Solomon codes, Cryptology ePrint Archive, Paper 2025/2110, 2025, https://eprint.iacr.org/2025/2110.

[HLP24] U. Haböck, D. Levit, and S. Papini, Circle STARKs, Cryptology ePrint Archive, Paper 2024/278, 2024, https://eprint.iacr.org/2024/278.

[Jo26] S. Jo, Reed–Solomon mutual correlated agreement beyond the Johnson radius, Cryptology ePrint Archive, Paper 2026/1432, 2026, https://eprint.iacr.org/2026/1432.

[KKH26] D. Krachun, S. Kazanin, and U. Haböck, Failure of proximity gaps close to capacity, Cryptology ePrint Archive, Paper 2026/782, 2026, https://eprint.iacr.org/2026/782.

Ulam.ai

Email address: prz.chojecki@gmail.com

URL: https://github.com/przchojecki/rs-mca