A Universal Field-Size Cap for Mutual Correlated Agreement on Smooth Reed–Solomon Domains

Przemek Chojecki
ulam.ai

June 12, 2026

Abstract

We prove a field-size-universal upper bound on the proximity parameter achievable in the grand mutual-correlated-agreement (MCA) challenge of Arnon–Boneh–Fenzi [ABF26]. Let $\mathbb{F}$ be any finite field with $|\mathbb{F}| < 2^{256}$, let D be a smooth multiplicative evaluation domain of order n contained in a subfield $B \subseteq \mathbb{F}$ (with $B = \mathbb{F}$ allowed), and let $C = \mathrm{RS}[\mathbb{F}, D, k]$ have $k = \rho n \le 2^{40}$ and $\rho \in \{1/2, 1/4, 1/8, 1/16\}$. If $2^{10} \mid n$ for $\rho \in \{1/2, 1/4, 1/8\}$, and $2^{11} \mid n$ for $\rho = 1/16$, then

$$\varepsilon_{\mathrm{mca}}(C, \delta) \;>\; \frac{1}{2k}\Big(1 - \frac{n}{|\mathbb{F}|}\Big) \;\ge\; 2^{-86} \;\gg\; 2^{-128}$$

for every $\delta \in [1 - \rho - 2^{-9}, 1 - \rho)$ at rates $1/2, 1/4, 1/8$, and for every $\delta \in [1 - \rho - 2^{-10}, 1 - \rho)$ at rate $1/16$. The same lower bound holds for $\varepsilon_{\mathrm{ca}}(C, \delta)$ on the CS-admissible subinterval ending at $1 - \rho - 1/n$. If $|\mathbb{F}| \ge 2n$, the displayed lower bound improves to 2^{-42}. Consequently $\delta^*_C(2^{-128}) \le 1 - \rho - 2^{-9}$ for $\rho \in \{1/2, 1/4, 1/8\}$ and $\delta^*_C(2^{-128}) \le 1 - \rho - 2^{-10}$ for $\rho = 1/16$, throughout the $|\mathbb{F}| < 2^{256}$ challenge envelope.

The mechanism composes a pigeonhole lower bound on quotient-locator list fibers—a slack-two and subfield-pigeonhole sharpening of [Cho26a, Main Thm. (d)]—with the correlated-agreement-to-list-decoding theorem of Crites–Stewart [CS25], used directly in the integer-radius form implied by [CS25, Thm. 2]. The composition itself was first carried out, conditionally and with weaker constants, in the universal-cap theorem of [Cho26b] (error 2^{-42} at gap 2^{-11}, assuming $|\mathbb{F}| \ge 2n$); the present note sharpens the gap, removes the field-size condition, and extends the cap to extension fields. No value-set or equidistribution input is used, so the per-fiber collision problem isolated in [Cho26b] is routed around rather than solved. We instantiate the bound at the KoalaBear-sextic parameters of [ABF26, §6.3], including the interleaved rows of their Tables 2–3, give an independent slacked variant through [BCHKS25, Thm. 1.9], and prove a two-radius refinement of the subfield-confinement theorem of [Cho26b], showing that over extension fields every explicit base-rational locator line is inert: the CS25-certified MCA failure must be witnessed by genuinely $\mathbb{F}$-valued lines.

1 Introduction

Proximity testing for Reed–Solomon codes underlies the soundness of modern SNARKs, and the sharpest available soundness accounting is phrased through *mutual correlated agreement* (MCA): a random linear combination $f_1 + \gamma f_2$ being δ-close to the code should force f_1, f_2 to be simultaneously explained by codewords on a *common* agreement set, except with small probability $\varepsilon_{\mathrm{mca}}(C, \delta)$ over γ. The survey of Arnon, Boneh and Fenzi [ABF26] crystallizes the

state of the art into a *grand MCA challenge*: for $C = \mathrm{RS}[\mathbb{F}, L, k]$ over a smooth domain L, with rate $\rho \in \{1/2, 1/4, 1/8, 1/16\}$, degree $k \le 2^{40}$, and field size $|\mathbb{F}| < 2^{256}$, determine the largest proximity parameter δ^* at which $\varepsilon_{\mathrm{mca}}(C, \delta) \le 2^{-128}$ can hold.

Positive results reach the Johnson bound $1 - \sqrt{\rho}$. On the negative side, [Cho26b] showed $\delta^*_C(2^{-128}) < 1 - \rho - 1/64$ for all prime fields of size up to roughly 2^{150} (per-rate reaches between $2^{150.6}$ and $2^{161.8}$), by exhibiting quotient-locator witnesses whose bad-slope sets are value sets of restricted symmetric sums over small subgroups; it further proved a *conditional* universal cap over every field of size at most 2^{256} — error 2^{-42} at gap 2^{-11}, assuming $|\mathbb{F}| \ge 2n$ — by composing quotient-core lists with the same Crites–Stewart conversion used here (its universal-cap theorem). Above the value-set reach, what remained open was the *error-one* question, identified there with the *per-fiber collision problem*: controlling, at a *fixed* large prime, the collisions of characteristic-zero locator values under reduction. Extension fields, the setting actually deployed ([ABF26, §6.3] works over a sextic extension of KoalaBear), posed an additional obstruction: by the subfield-confinement theorem of [Cho26b], refined in Section 6, every explicit witness of [Cho26a, Cho26b] is B-rational and contributes only $p^{-5} < 2^{-154}$ over the deployed extension.

This note sharpens the universal cap — gap 2^{-9} (resp. 2^{-10}) in place of 2^{-11}, error 2^{-86} with no condition relating n and $|\mathbb{F}|$, and extension fields handled directly — for every field below 2^{256} at once, by changing what is counted.

Theorem 1.1 (informal; see Theorem 4.1 and Corollary 5.1)**.** *Let $\mathbb{F}$ be any finite field with $|\mathbb{F}| < 2^{256}$, let $D \subseteq \mathbb{F}^\times$ be a smooth multiplicative evaluation domain of order n, and let $C = \mathrm{RS}[\mathbb{F}, D, k]$ have $k = \rho n \le 2^{40}$ with $\rho \in \{1/2, 1/4, 1/8, 1/16\}$. Assume $2^{10} \mid n$ for $\rho \in \{1/2, 1/4, 1/8\}$, and $2^{11} \mid n$ for $\rho = 1/16$. Then*

$$\varepsilon_{\mathrm{mca}}(C, \delta) > 2^{-86} \gg 2^{-128}$$

for every $\delta \in [1 - \rho - 2^{-9}, 1 - \rho)$ at rates $1/2, 1/4, 1/8$, and for every $\delta \in [1 - \rho - 2^{-10}, 1 - \rho)$ at rate $1/16$. The same lower bound holds for $\varepsilon_{\mathrm{ca}}(C, \delta)$ on the CS-admissible subinterval ending at $1 - \rho - 1/n$. If $|\mathbb{F}| \ge 2n$, the lower bound is $> 2^{-42}$. Hence $\delta^*_C(2^{-128}) \le 1 - \rho - 2^{-9}$ for $\rho \in \{1/2, 1/4, 1/8\}$ and $\delta^*_C(2^{-128}) \le 1 - \rho - 2^{-10}$ for $\rho = 1/16$.

The proof composes two ingredients. The first (Lemma 3.1) is a pigeonhole lower bound on list sizes at a single received word: the quotient locators of [Cho26a] attach to each ℓ-subset A of the order-N quotient coset $Q = D^a$ a codeword c_A at distance $\le 1 - \rho - t/N$ from the word $u_{z_A} = x^{k+ta} + z_A x^{k+(t-1)a}$, where $z_A = -e_1(A)$; pigeonholing the $\binom{N}{\ell}$ subsets over their $\le |B|$ slope values produces one word u_z carrying a list of size $\binom{N}{\ell}/|B|$. Two refinements of [Cho26a, Main Thm. (d)] matter here: we run the construction at *slack two* ($t = 2$), so that the listed codewords have degree $\le k$ and live in $\mathrm{RS}[\mathbb{F}, D, k+1]$ without any divisibility demand on $k+1$ (which is odd in all challenge instantiations, while a is a 2-power); and we pigeonhole over the *subfield* B containing D rather than over $\mathbb{F}$, which is what makes the extension-field case go through. The second ingredient is the correlated-agreement-to-list-decoding theorem of Crites–Stewart [CS25, Thm. 2]: in the integer-error range $f < n - k - 1$, if $\varepsilon_{\mathrm{ca}}(C, \delta)$ is below $\eta(\frac{1}{k} - \frac{n}{k|\mathbb{F}|})$, then *every* list of $\mathrm{RS}[\mathbb{F}, D, k+1]$ at radius δ has size at most $\lceil |\mathbb{F}|\, \varepsilon_{\mathrm{ca}}/(1-\eta) \rceil$. Taking $\eta = \frac{1}{2}$, the heavy fiber of Lemma 3.1 violates this whenever $\binom{N}{\ell} \ge |B|(|\mathbb{F}|/k+1)$ — an inequality with hundreds of bits of slack throughout the challenge envelope — and the contrapositive yields $\varepsilon_{\mathrm{ca}}(C, \delta) > \frac{1}{2k}(1 - n/|\mathbb{F}|)$ throughout the CS-admissible range. The MCA cap then persists up to capacity by support-wise MCA monotonicity.

No value-set counting, exponential-sum estimate, or equidistribution hypothesis enters: pigeonhole fibers exist over every field unconditionally, and [CS25] converts list mass into

correlated-agreement error directly. The cost of this route, relative to the error-one results of [Cho26a, Cho26b] below 2^{150}, is an error of order $1/k$ rather than $1-o(1)$, and a gap 2^{-9} at the first three prize rates (uniformly 2^{-10} across all four), rather than the fixed gap $1/64$; Section 7 quantifies the trade-off and what remains open.

Beyond the universal cap (Corollary 5.1, with a two-tier summary combining it with [Cho26b]), we prove: a deployed-parameter corollary at the KoalaBear-sextic instantiation of [ABF26, §6.3], with $\varepsilon_{\mathrm{mca}} > 2^{-22}$ at gap 2^{-7} and the same $\varepsilon_{\mathrm{ca}}$ lower bound on the CS-admissible subinterval (Corollary 5.4), extended to every interleaved row of [ABF26, Tables 2–3] via an interleaving-transfer lemma (Lemma 3.5 and Corollary 5.5); an independent slacked variant through [BCHKS25, Thm. 1.9] for robustness (Proposition 5.6); and a subfield-confinement lemma (Lemma 6.1) showing that for lines with B-valued words over an extension $\mathbb{F}/B$, every CA- or MCA-bad slope lies in B. The latter has a structural consequence (Corollary 6.2): at deployed parameters the lines certifying Corollary 5.4 are necessarily *not* B-rational, so the failure proved here is fiber-borne and currently non-constructive; exhibiting explicit certifying lines is posed as Open Problem 7.1.

Roadmap. Section 2 fixes definitions, aligned with [ABF26, Defs. 4.1, 4.3], states the direct CS25 conversion used for the main cap, and records the separate BCHKS/ABF slacked fallback. Section 3 proves the fiber and interleaving lemmas, Section 4 the main theorem, Section 5 the challenge-level and deployed corollaries, Section 6 subfield confinement. Section 7 discusses scope, limits, calibration against the conjectural picture of [Cho26b], and open problems. Throughout, companion results of [Cho26a, Cho26b] are cited by their printed theorem names rather than numbers, since the numbering of those manuscripts is not yet frozen.

2 Preliminaries

Codes and distances. Throughout, $B \subseteq \mathbb{F}$ are finite fields, $q := |\mathbb{F}|$, and $D = \alpha H \subseteq B^\times$ is a multiplicative coset of a cyclic subgroup H of order n; we write D for the evaluation domain (denoted L in [ABF26]). The subgroup case is $\alpha = 1$, and taking $B = \mathbb{F}$ is always allowed. For $k \le n$,

$$\mathrm{RS}[\mathbb{F}, D, k] \;:=\; \big\{(p(x))_{x\in D} \;:\; p \in \mathbb{F}[X],\ \deg p < k\big\} \subseteq \mathbb{F}^n, \qquad \rho := k/n.$$

For $u, v \in \mathbb{F}^n$, $\Delta(u,v)$ is the relative Hamming distance, $\Delta(u,C) := \min_{c\in C}\Delta(u,c)$, and the relative minimum distance of $C = \mathrm{RS}[\mathbb{F}, D, k]$ is $\delta_{\min}(C) = 1 - \rho + 1/n$. For $S \subseteq D$, $\Delta_S(u,v)$ and $\Delta_S(u,C)$ denote the corresponding restricted relative distances on S. Lists are denoted

$$\Lambda(C, \delta, u) := \{c \in C : \Delta(u,c) \le \delta\}, \qquad \Lambda(C,\delta) := \max_{u\in\mathbb{F}^n} |\Lambda(C,\delta,u)|.$$

For $s \ge 1$ the s-interleaved code is $C^{\equiv s} := \{(c^{(1)}, \dots, c^{(s)}) : c^{(i)} \in C\} \subseteq (\mathbb{F}^n)^s$ with the column distance $\Delta_s(F,G) := \frac{1}{n}\big|\{j : \exists i,\ f_j^{(i)} \ne g_j^{(i)}\}\big|$, and $\Delta_s(F, C^{\equiv s}) := \min_{\mathbf{c}\in C^{\equiv s}} \Delta_s(F, \mathbf{c})$.

Definition 2.1 (correlated agreement error; [ABF26, Def. 4.1] up to notation)**.** For $0 \le \delta_{\mathrm{fld}} \le \delta_{\mathrm{int}} < 1$,

$$\varepsilon_{\mathrm{ca}}(C, \delta_{\mathrm{fld}}, \delta_{\mathrm{int}}) \;:=\; \max_{f_1, f_2 \in \mathbb{F}^n} \Pr_{\gamma\leftarrow\mathbb{F}}\Big[\, \Delta(f_1 + \gamma f_2,\, C) \le \delta_{\mathrm{fld}} \;\wedge\; \Delta_2\big((f_1,f_2),\, C^{\equiv 2}\big) > \delta_{\mathrm{int}} \,\Big],$$

and the no-proximity-loss version is $\varepsilon_{\mathrm{ca}}(C,\delta) := \varepsilon_{\mathrm{ca}}(C,\delta,\delta)$. A slope γ realizing the event for a fixed pair (f_1, f_2) is called *CA-bad* for that pair.

Definition 2.2 (mutual correlated agreement error, support-wise; [ABF26, Def. 4.3] up to notation)**.** For $\delta \in (0,1)$,

$$\varepsilon_{\mathrm{mca}}(C,\delta) := \max_{f_1,f_2 \in \mathbb{F}^n} \Pr_{\gamma \leftarrow \mathbb{F}} \left[\begin{array}{l} \exists S \subseteq D,\ |S| \geq (1-\delta)n,\ \Delta_S(f_1 + \gamma f_2, C) = 0, \\ \text{and } \Delta_S\big((f_1,f_2), C^{\equiv 2}\big) > 0 \end{array} \right].$$

Equivalently, γ is MCA-bad for (f_1,f_2) if some large agreement set S explains the point $f_1+\gamma f_2$, but the same set S does not simultaneously explain f_1 and f_2 by two codewords of C.

Definition 2.3 (challenge threshold)**.** Following the grand MCA challenge of [ABF26, §1], for $\varepsilon^* \in (0,1)$ set $\delta^*_C(\varepsilon^*) := \sup\{\delta \in (0, 1-\rho) : \varepsilon_{\mathrm{mca}}(C,\delta) \leq \varepsilon^*\}$, with $\sup \emptyset := 0$.

Fact 2.4 (cf. [ABF26, Fact 4.5])**.** *For every $\delta \in (0,1)$: $\varepsilon_{\mathrm{ca}}(C,\delta) \leq \varepsilon_{\mathrm{mca}}(C,\delta)$.*

Proof. Fix (f_1,f_2) and let γ be CA-bad: $\Delta(f_1+\gamma f_2, C) \leq \delta$ and $\Delta_2((f_1,f_2), C^{\equiv 2}) > \delta$. Choose $S \subseteq D$ of size at least $(1-\delta)n$ on which $f_1+\gamma f_2$ agrees with a codeword of C. Since (f_1,f_2) is not δ-close to $C^{\equiv 2}$, there is no set of size at least $(1-\delta)n$ on which f_1 and f_2 are simultaneously explained by codewords of C; in particular this fails on the chosen S. Thus γ is MCA-bad for the same pair. Taking maxima over pairs gives the claim. □

Lemma 2.5 (support-wise MCA monotonicity)**.** *For every linear code $C \subseteq \mathbb{F}^D$, the function*

$$\delta \longmapsto \varepsilon_{\mathrm{mca}}(C,\delta)$$

is nondecreasing.

Proof. Fix a pair (f_1,f_2) and suppose that γ is MCA-bad at radius δ. Then there is a set $S \subseteq D$ with

$$|S| \geq (1-\delta)n$$

such that $f_1+\gamma f_2$ is explained by C on S, but (f_1,f_2) is not simultaneously explained by two codewords of C on the same S. If $\delta' \geq \delta$, then

$$|S| \geq (1-\delta)n \geq (1-\delta')n,$$

so the same set S witnesses that γ is MCA-bad at radius δ'. Taking the maximum over pairs gives the claim. □

Imported theorems. The main composition below uses Crites–Stewart directly, with its integer-radius admissibility condition; see Remark 2.8. We also keep the slacked BCHKS/ABF fallback as a separate imported route.

Theorem 2.6 (Crites–Stewart conversion, relative-radius form)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, let $C^+ := \mathrm{RS}[\mathbb{F}, D, k+1]$, let $q := |\mathbb{F}|$ and $n := |D|$. Let $\delta \in (0,1)$ and set $f_\delta := \lfloor \delta n \rfloor$. Assume*

$$f_\delta < n - k - 1.$$

For any $\eta \in [0,1)$, if

$$\varepsilon_{\mathrm{ca}}(C,\delta) \leq \eta\Big(\frac{1}{k} - \frac{n}{kq}\Big),$$

then

$$\Lambda(C^+,\delta) \leq \left\lceil \frac{q\,\varepsilon_{\mathrm{ca}}(C,\delta)}{1-\eta} \right\rceil.$$

Proof. Let $\epsilon := \varepsilon_{\mathrm{ca}}(C, \delta)$. Since Hamming distances are integral, the event $\Delta(\cdot, C) \le \delta$ is the same as having Hamming distance at most $f_\delta = \lfloor \delta n \rfloor$. Likewise the "far" condition in the no-proximity-loss correlated-agreement definition is the corresponding integer condition at radius f_δ.

Theorem 2 of Crites–Stewart [CS25] says that if $\mathrm{RS}(\mathbb{F}, D, k)$ satisfies correlated agreement over lines with $f < n - k - 1$ errors and error parameter $\epsilon < (q-n)/(kq)$, then $\mathrm{RS}(\mathbb{F}, D, k+1)$ is $(f/n, L)$-list decodable with

$$L = \left\lceil \frac{\epsilon q(q-n)}{q - n - k\epsilon q} \right\rceil.$$

Our hypothesis gives

$$\epsilon \le \eta \frac{q-n}{kq} < \frac{q-n}{kq},$$

because $\eta < 1$. Therefore CS25 applies with $f = f_\delta$. Moreover,

$$q - n - k\epsilon q \ge (1-\eta)(q-n),$$

and hence

$$\frac{\epsilon q(q-n)}{q - n - k\epsilon q} \le \frac{q\epsilon}{1-\eta}.$$

Thus

$$\Lambda(C^+, \delta) \le \left\lceil \frac{q\, \varepsilon_{\mathrm{ca}}(C, \delta)}{1-\eta} \right\rceil,$$

as claimed. □

Theorem 2.7 ([BCHKS25, Thm. 1.9], as stated in [ABF26, Thm. 5.2])**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, let $\rho = k/n$, and let $\delta \in (0, 1-\rho)$ satisfy $\delta + 2/n \le 1 - \rho - 1/n$. If*

$$\varepsilon_{\mathrm{ca}}\Big(C,\ \delta + \tfrac{2}{n},\ 1 - \rho - \tfrac{1}{n}\Big) \ <\ \frac{1}{2n}, \qquad \textit{then} \qquad \Lambda(C, \delta) \ <\ |\mathbb{F}|.$$

Remark 2.8 (on imports and radius conventions). Theorem 2.6 is now used directly in the form implied by Crites–Stewart [CS25, Thm. 2], rather than through the ABF26 restatement. The important admissibility condition is the integer-radius condition

$$\lfloor \delta n \rfloor < n - k - 1.$$

Consequently, the no-loss CA lower bound proved below should be read only in the CS-admissible range. The Proximity-Prize MCA cap is unaffected: we prove the lower bound at the endpoint

$$\delta_N := 1 - \rho - \frac{2}{N}$$

and then use support-wise MCA monotonicity to extend the MCA failure to all larger sub-capacity radii.

Theorem 2.7 remains a separate slacked fallback imported through the BCHKS/ABF route. It is not needed for the main field-size-universal MCA cap.

3 Locator fibers and interleaving transfer

Fix $N \mid n$, set $a := n/N$, and let $Q := D^a := \{x^a : x \in D\} \subseteq B^\times$, a multiplicative coset of order N. The a-th power map $D \to Q$ is surjective and its fibers $S_b := \{x \in D : x^a = b\}$ have exactly a elements each. Since $a \mid n \mid |B| - 1$, the characteristic of $\mathbb{F}$ does not divide a, so $X^a - b$ is separable for $b \neq 0$ and its full root set inside D is S_b whenever $b \in Q$. For a set A, $e_j(A)$ denotes the j-th elementary symmetric function of its elements.

Lemma 3.1 (locator fibers are lists)**.** *Let $B \subseteq \mathbb{F}$, let $D \subseteq B^\times$ be a multiplicative coset of order n, and let k satisfy $a \mid k$ and $\rho = k/n$. Write $\ell_1 := \rho N + 1$ and $\ell_2 := \rho N + 2$.*

(i) *If $\ell_1 \le N$, there exists $z \in B$ such that, with $u_z := \big(x^{k+a} + z\,x^k\big)_{x \in D} \in \mathbb{F}^n$,*

$$\big|\Lambda\big(\mathrm{RS}[\mathbb{F}, D, k],\ 1 - \rho - \tfrac{1}{N},\ u_z\big)\big| \ \ge\ \binom{N}{\ell_1}\Big/|B|.$$

(ii) *If $\ell_2 \le N$, there exists $z \in B$ such that, with $u_z := \big(x^{k+2a} + z\,x^{k+a}\big)_{x \in D} \in \mathbb{F}^n$,*

$$\big|\Lambda\big(\mathrm{RS}[\mathbb{F}, D, k+1],\ 1 - \rho - \tfrac{2}{N},\ u_z\big)\big| \ \ge\ \binom{N}{\ell_2}\Big/|B|.$$

In both cases the words u_z are B-valued and all listed codewords lie in $\mathrm{RS}[B, D, k]$ (*resp.* $\mathrm{RS}[B, D, k+1]$).

Proof. We prove (ii); part (i) is the identical computation one degree rung lower and is [Cho26a, Main Thm. (d)] sharpened by pigeonholing over B in place of $\mathbb{F}$.

Let $A \subseteq Q$ with $|A| = \ell_2$, and put

$$L_A(X) \ :=\ \prod_{b \in A}\big(X^a - b\big) \ =\ \sum_{j=0}^{\ell_2}(-1)^j e_j(A)\,X^{a(\ell_2 - j)} \ =\ X^{k+2a} - e_1(A)\,X^{k+a} + R_A(X),$$

where $a\ell_2 = k + 2a$, $a(\ell_2 - 1) = k + a$, and R_A collects the terms with $j \ge 2$, so $\deg R_A \le a(\ell_2 - 2) = k$. The polynomial L_A is squarefree with root set $S_A := \bigcup_{b \in A} S_b \subseteq D$ of size $a\ell_2 = k + 2a$.

Set $z_A := -e_1(A) \in B$ and $c_A := \big(-R_A(x)\big)_{x \in D}$. Since $\deg R_A \le k$, we have $c_A \in \mathrm{RS}[B, D, k+1] \subseteq \mathrm{RS}[\mathbb{F}, D, k+1]$. For every $x \in S_A$,

$$0 = L_A(x) = x^{k+2a} + z_A\,x^{k+a} + R_A(x), \qquad \text{i.e.} \qquad u_{z_A}(x) = c_A(x),$$

so u_{z_A} and c_A agree on at least $k + 2a$ points of D and $\Delta(u_{z_A}, c_A) \le 1 - \rho - 2/N$.

The map $A \mapsto c_A$ is injective on the ℓ_2-subsets of Q sharing a common slope value: if $z_A = z_{A'}$ and $c_A = c_{A'}$, then R_A and $R_{A'}$ are polynomials of degree $\le k < n$ agreeing on all n points of D, hence equal as polynomials; then $L_A = L_{A'}$, so $S_A = S_{A'}$, and $A = \{x^a : x \in S_A\} = A'$.

Finally, $z_A = -e_1(A)$ takes at most $|B|$ values as A ranges over the $\binom{N}{\ell_2}$ subsets; some $z \in B$ is attained by at least $\binom{N}{\ell_2}/|B|$ of them, and the corresponding codewords c_A are pairwise distinct elements of $\Lambda(\mathrm{RS}[\mathbb{F}, D, k+1],\ 1 - \rho - 2/N,\ u_z)$. □

Remark 3.2 (why slack two). Part (i) applied to $C^+ = \mathrm{RS}[\mathbb{F}, D, k+1]$ directly would require $a \mid k+1$; in every challenge instantiation a is a power of two and $k+1$ is odd, so this fails. At slack two the listed codewords have degree $\le k$, i.e. lie in C^+, while the divisibility demand stays on k. This is exactly the interface needed by Theorem 2.6, whose list conclusion concerns C^+ while its hypothesis concerns C.

Remark 3.3 (the subfield pigeonhole elsewhere). The same one-line strengthening — pigeonholing locator data over the field of definition B rather than over $\mathbb{F}$ — also upgrades the coefficient-pigeonhole bound of [Cho26b] to the generated field: the prefix map Φ_σ there has image in B^σ whenever $D \subseteq B$, since the elementary symmetric prefixes of subsets of D are B-valued. This is the form consumed by the "necessary" half of the list-profile conjecture of [Cho26b], which states the entropy floor at the generated field $\tau^*(\rho, q_D)$; Lemma 3.1 supplies the missing subfield step.

Remark 3.4 (lists, not lines). The words u_z of Lemma 3.1 lie on the B-rational line $f + zg$ with $f = (x^{k+ta})_{x\in D}$, $g = (x^{k+(t-1)a})_{x\in D}$, but the lemma is purely a statement about lists at a received word and no line structure is used in Section 4. This matters: by Corollary 6.2 below, over a proper extension $\mathbb{F} \supsetneq B$ the lines that end up certifying the CA failure cannot be these (or any) B-rational lines.

Lemma 3.5 (interleaving transfer). *For every linear code $C \subseteq \mathbb{F}^n$, every $s \ge 1$ and every $\delta \in (0,1)$:*

$$\varepsilon_{\mathrm{ca}}\big(C^{\equiv s}, \delta\big) \;\ge\; \varepsilon_{\mathrm{ca}}(C, \delta), \qquad \varepsilon_{\mathrm{mca}}\big(C^{\equiv s}, \delta\big) \;\ge\; \varepsilon_{\mathrm{mca}}(C, \delta).$$

Proof. For $h \in \mathbb{F}^n$ write $h^\Delta := (h, \dots, h) \in (\mathbb{F}^n)^s$. For any codeword tuple $\mathbf{c} = (c^{(1)}, \dots, c^{(s)}) \in C^{\equiv s}$, the columns where h^Δ and $\mathbf{c}$ agree are $\bigcap_i \{j : c^{(i)}_j = h_j\} \subseteq \{j : c^{(1)}_j = h_j\}$, so $\Delta_s(h^\Delta, C^{\equiv s}) \ge \Delta(h, C)$; taking $c^{(1)} = \cdots = c^{(s)}$ gives equality: $\Delta_s(h^\Delta, C^{\equiv s}) = \Delta(h, C)$.

Fix (f_1, f_2) attaining $\varepsilon_{\mathrm{ca}}(C, \delta)$ (resp. $\varepsilon_{\mathrm{mca}}$) and consider the pair (f_1^Δ, f_2^Δ) for $C^{\equiv s}$, with the slope γ acting row-wise, so $f_1^\Delta + \gamma f_2^\Delta = (f_1 + \gamma f_2)^\Delta$. By the displayed equality, $\Delta_s\big((f_1 + \gamma f_2)^\Delta, C^{\equiv s}\big) \le \delta$ iff $\Delta(f_1 + \gamma f_2, C) \le \delta$. For the correlated condition, the pair (f_1^Δ, f_2^Δ) viewed as a $2s$-row interleaved word has rows f_1 (s times) and f_2 (s times); by the same column argument its distance to $(C^{\equiv s})^{\equiv 2} = C^{\equiv 2s}$ equals $\Delta_2((f_1, f_2), C^{\equiv 2})$. For the MCA condition, a common set S explains f_1^Δ and f_2^Δ iff it explains f_1 and f_2. Hence γ is CA-bad (resp. MCA-bad) for (f_1^Δ, f_2^Δ) iff it is so for (f_1, f_2), and the maxima over pairs can only grow. □

4 The main theorem

Theorem 4.1 (universal cap). *Let $B \subseteq \mathbb{F}$ be finite fields, let $q = |\mathbb{F}|$, let $D \subseteq B^\times$ be a multiplicative coset of order n, and let k have $\rho = k/n \in (0,1)$; set $C := \mathrm{RS}[\mathbb{F}, D, k]$. Let $N \mid n$ satisfy $a := n/N \mid k$ and $(1-\rho)N \ge 3$. If*

$$\binom{N}{\rho N + 2} \ge |B|\Big(\frac{q}{k} + 1\Big), \tag{1}$$

then, writing

$$\delta_N := 1 - \rho - \frac{2}{N},$$

we have

$$\varepsilon_{\mathrm{ca}}(C, \delta) > \frac{1}{2k}\Big(1 - \frac{n}{q}\Big)$$

for every CS-admissible radius

$$\delta_N \le \delta < 1 - \rho - \frac{1}{n}.$$

In particular,

$$\varepsilon_{\mathrm{mca}}(C, \delta) > \frac{1}{2k}\Big(1 - \frac{n}{q}\Big)$$

for every

$$\delta_N \le \delta < 1 - \rho.$$

Consequently,

$$\delta^*_C(\varepsilon^*) \le 1 - \rho - \frac{2}{N}$$

for every

$$\varepsilon^* \le \frac{1}{2k}\Big(1 - \frac{n}{q}\Big).$$

Proof. Note first that $\rho N = k/a \in \mathbb{Z}$, $\ell_2 = \rho N + 2 \le N - 1$ by $(1-\rho)N \ge 3$, and $\delta_N = 1 - \rho - 2/N > 0$ for the same reason. Also $k + 1 \le n$.

Fix

$$\delta \in [\delta_N,\, 1 - \rho - \tfrac{1}{n}).$$

Then

$$\lfloor \delta n \rfloor < n - k - 1,$$

so Theorem 2.6 is applicable.

Suppose toward a contradiction that

$$\varepsilon_{\mathrm{ca}}(C, \delta) \le \frac{1}{2k}\Big(1 - \frac{n}{q}\Big).$$

Apply Theorem 2.6 with $\eta = \frac{1}{2}$. Its hypothesis holds, so

$$\Lambda(C^+, \delta) \le \Big\lceil 2q \cdot \varepsilon_{\mathrm{ca}}(C, \delta) \Big\rceil \le \Big\lceil \frac{q-n}{k} \Big\rceil < \frac{q-n}{k} + 1 \le \frac{q}{k} + 1.$$

On the other hand, by Lemma 3.1(ii) and monotonicity of lists in the radius, since $\delta \ge \delta_N$,

$$\Lambda(C^+, \delta) \ge \big|\Lambda(C^+, \delta_N, u_z)\big| \ge \binom{N}{\ell_2} \Big/ |B| \ge \frac{q}{k} + 1$$

by (1). This is a contradiction. Hence

$$\varepsilon_{\mathrm{ca}}(C, \delta) > \frac{1}{2k}\Big(1 - \frac{n}{q}\Big)$$

for every $\delta \in [\delta_N, 1 - \rho - 1/n)$.

In particular, the above holds at $\delta = \delta_N$. By Fact 2.4,

$$\varepsilon_{\mathrm{mca}}(C, \delta_N) \ge \varepsilon_{\mathrm{ca}}(C, \delta_N) > \frac{1}{2k}\Big(1 - \frac{n}{q}\Big).$$

By support-wise MCA monotonicity, Lemma 2.5, the same lower bound for $\varepsilon_{\mathrm{mca}}(C, \delta)$ holds for every

$$\delta \in [\delta_N, 1 - \rho).$$

Thus no sub-capacity radius

$$\delta \geq 1 - \rho - \frac{2}{N}$$

is ε^*-admissible whenever

$$\varepsilon^* \leq \frac{1}{2k}\Big(1 - \frac{n}{q}\Big),$$

which proves the claimed bound on $\delta_C^*(\varepsilon^*)$. □

Remark 4.2 (constants). Choosing $\eta = 1 - \theta$ in Theorem 2.6 improves the conclusion to $\varepsilon_{\mathrm{ca}} > \frac{1-\theta}{k}(1 - \frac{n}{q})$ at the price of strengthening (1) to $\binom{N}{\rho N+2} \geq |B|\big(\frac{q}{\theta k} + 1\big)$; given the hundreds of bits of slack in (1) throughout the challenge envelope (Table 1), the error constant is effectively $(1 - o(1))/k$. We fix $\eta = \frac{1}{2}$ for cleanliness.

Remark 4.3 (scope). No smoothness is assumed beyond the existence of a single divisor $N \mid n$ with $a = n/N \mid k$ and (1). The theorem applies verbatim to 2-adic, mixed-radix, or any other FFT-friendly multiplicative domain, and — via the subfield pigeonhole — to any extension field $\mathbb{F} \supseteq B$ with $D \subseteq B$.

The CS25 conversion is an integer-radius theorem with the condition $\lfloor \delta n \rfloor < n - k - 1$. Therefore the no-loss CA lower bound is stated only for

$$\delta < 1 - \rho - \frac{1}{n}.$$

The MCA threshold cap is stronger in the range needed for the Proximity Prize: once the support-wise MCA lower bound is proved at

$$\delta_N = 1 - \rho - \frac{2}{N},$$

monotonicity of support-wise MCA extends it to every larger $\delta < 1 - \rho$.

5 Consequences for the grand MCA challenge

Corollary 5.1 (universal field-size cap for the challenge envelope)**.** *Let $\rho \in \{\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}\}$ and set*

$$N_\rho := 1024 \quad (\rho \in \{\tfrac{1}{2}, \tfrac{1}{4}, \tfrac{1}{8}\}), \qquad N_{1/16} := 2048.$$

Let $\mathbb{F}$ be any finite field with $q < 2^{256}$, let $B \subseteq \mathbb{F}$ be any subfield, let $D \subseteq B^\times$ be a multiplicative coset of order n with $N_\rho \mid n$, and let $k = \rho n \leq 2^{40}$. Then $C = \mathrm{RS}[\mathbb{F}, D, k]$ satisfies

$$\varepsilon_{\mathrm{mca}}(C, \delta) > \frac{1}{2k}\Big(1 - \frac{n}{q}\Big) \geq 2^{-86} \gg 2^{-128}$$

for every

$$\delta \in \Big[1 - \rho - \frac{2}{N_\rho},\, 1 - \rho\Big).$$

Moreover, the same lower bound holds for $\varepsilon_{\mathrm{ca}}(C, \delta)$ throughout the CS-admissible subinterval

$$\delta \in \Big[1 - \rho - \frac{2}{N_\rho},\, 1 - \rho - \frac{1}{n}\Big).$$

If $q \geq 2n$, the lower bound improves to $\geq 2^{-42}$. In particular,

$$\delta_C^*(2^{-128}) \;\leq\; 1 - \rho - 2^{-9} \quad (\rho \in \{\tfrac{1}{2}, \tfrac{1}{4}, \tfrac{1}{8}\}), \qquad \delta_C^*(2^{-128}) \;\leq\; 1 - \rho - 2^{-10} \quad (\rho = \tfrac{1}{16}).$$

Proof. We verify the hypotheses of Theorem 4.1 with $N = N_\rho$. Divisibility: $a = n/N_\rho$ divides k because $k/a = \rho N_\rho \in \mathbb{Z}$ (equal to $512, 256, 128, 128$ respectively); $(1-\rho)N_\rho \geq 128 \geq 3$. For (1), the entropy bound $\binom{N}{\ell} \geq 2^{N\mathrm{H}_2(\ell/N)}/(N+1)$ together with the conservative evaluations of H_2 recorded in Table 1 gives $\log_2 \binom{N_\rho}{\rho N_\rho + 2} \geq 552$ in all four cases, while

$$|B|\Big(\frac{q}{k}+1\Big) \;\leq\; q\Big(\frac{q}{k}+1\Big) \;\leq\; q^2+q \;<\; 2^{512}+2^{256} \;<\; 2^{513} \;\leq\; 2^{552},$$

so (1) holds with at least 39 bits to spare (and over 500 bits at rate $\frac{1}{2}$). Since $D \subseteq \mathbb{F}^\times$, we have $n \leq q-1$; for fixed n, the function $1-n/q$ is minimized at $q = n+1$, hence $1-n/q \geq 1/(n+1)$. Also $n = k/\rho \leq 16k \leq 2^{44}$. Therefore

$$\frac{1}{2k}\Big(1-\frac{n}{q}\Big) \geq \frac{1}{2k(n+1)} > 2^{-86}.$$

If $q \geq 2n$ then $1 - n/q \geq 1/2$, giving the sharper $1/(4k) \geq 2^{-42}$. The gap is $2/N_\rho = 2^{-9}$ for $\rho \in \{1/2, 1/4, 1/8\}$ and $2/N_\rho = 2^{-10}$ for $\rho = 1/16$. Theorem 4.1 gives the MCA lower bound throughout the stated sub-capacity interval, and gives the CA lower bound on the CS-admissible subinterval ending at $1 - \rho - 1/n$. □

ρ	N_ρ	$\ell_2 = \rho N_\rho + 2$	$\mathrm{H}_2(\ell_2/N_\rho) \geq$	$\log_2 \binom{N_\rho}{\ell_2} \geq$	needed	gap $2/N_\rho$
1/2	1024	514	0.99998	1013	513	2^{-9}
1/4	1024	258	0.8143	823	513	2^{-9}
1/8	1024	130	0.5490	552	513	2^{-9}
1/16	2048	130	0.34105	687	513	2^{-10}

Table 1: Numerics for Corollary 5.1: $\log_2 \binom{N}{\ell} \geq N\mathrm{H}_2(\ell/N) - \log_2(N+1)$, rounded down. The "needed" column is the universal upper bound $\log_2\big(|B|(q/k+1)\big) < 513$ valid throughout the envelope $q < 2^{256}$, $k \geq 1$.

Remark 5.2 (hypothesis $N_\rho \mid n$). For the 2-adic domains of the challenge this holds whenever $n \geq 2^{10}$ (resp. 2^{11}), i.e. for all but very small instances. For mixed-radix smooth n, replace N_ρ by any divisor $N \mid n$ of comparable size with $\rho N \in \mathbb{Z}$; the entropy count in Table 1 is insensitive to the exact value of N at fixed $\log_2 N$.

Remark 5.3 (two-tier answer to the grand challenge). Combining Corollary 5.1 with the grand-challenge cap of [Cho26b, Prop. "Grand-challenge cap"] (gap $\frac{1}{64}$, prime fields up to $\approx 2^{150}$, via value-set coverage) yields the current state of the negative side for smooth multiplicative domains:

$$\delta_C^*(2^{-128}) \;\leq\; \begin{cases} 1-\rho-\frac{1}{64}, & \mathbb{F} \text{ prime},\ q \lesssim 2^{150}, \\ 1-\rho-2^{-9}, & \rho \in \{\frac{1}{2}, \frac{1}{4}, \frac{1}{8}\},\ \text{every } \mathbb{F},\ q < 2^{256}, \\ 1-\rho-2^{-10}, & \rho = \frac{1}{16},\ \text{every } \mathbb{F},\ q < 2^{256}. \end{cases}$$

Together with the positive results collected in [ABF26, Table 1], this pins the grand MCA threshold between the Johnson region and the explicit capacity-edge caps above, throughout the stated challenge envelope. The second and third tiers supersede the conditional cap of [Cho26b, Thm. "Universal cap"] ($\delta^* < 1-\rho-2^{-11}$ under $|\mathbb{F}| \geq 2n$), which was the previous state in this band.

Corollary 5.4 (deployed parameters: KoalaBear sextic)**.** *Let $B = \mathbb{F}_p$ with $p = 2^{31} - 2^{24} + 1$, let $\mathbb{F} = \mathbb{F}_{p^6}$ (so $q = p^6 \approx 2^{185.93}$), let $D \subseteq B^\times$ be the subgroup of order $n = 2^{21}$, and let $k = 2^{20}$, $\rho = \frac{1}{2}$, as in [ABF26, §6.3]. Then*

$$\varepsilon_{\mathrm{mca}}\big(\mathrm{RS}[\mathbb{F}_{p^6}, D, 2^{20}],\, \delta\big) > \big(1 - 2^{-164}\big)\, 2^{-21} > 2^{-22}$$

for every

$$\delta \in \Big[\frac{1}{2} - 2^{-7},\, \frac{1}{2}\Big).$$

Moreover, the same lower bound holds for $\varepsilon_{\mathrm{ca}}$ throughout the CS-admissible subinterval

$$\delta \in \Big[\frac{1}{2} - 2^{-7},\, \frac{1}{2} - 2^{-21}\Big).$$

Proof. Apply Theorem 4.1 with $N = 256$, $a = 2^{13} \mid k$, $(1-\rho)N = 128 \geq 3$, gap $2/N = 2^{-7}$. Hypothesis (1): $\log_2 \binom{256}{130} \geq 256 \cdot \mathrm{H}_2(130/256) - \log_2 257 \geq 256 \cdot 0.99982 - 8.01 > 247$, while $|B|(q/k+1) < 2^{31}\,(2^{166}+1) < 2^{198}$. The error bound is $\frac{1}{2k}(1 - n/q)$ with $n/q = 2^{21}/p^6 < 2^{-164}$. Theorem 4.1 gives the stated MCA bound on $[\frac{1}{2} - 2^{-7}, \frac{1}{2})$ and the stated CA bound on $[\frac{1}{2} - 2^{-7}, \frac{1}{2} - 2^{-21})$. □

Corollary 5.5 (all interleaved rows of [ABF26, Tables 2–3])**.** *In the setting of Corollary 5.4, for each $s = 2^j$ with $0 \leq j \leq 12$ let $C_s := \mathrm{RS}[\mathbb{F}_{p^6}, D_s, k_s]$ with $|D_s| = n_s = 2^{21}/s$ and $k_s = 2^{20}/s$ (rate $\frac{1}{2}$, $s\, n_s = 2^{21}$ as deployed). Then for every $\delta \in \big[\frac{1}{2} - 2^{-7}, \frac{1}{2}\big)$,*

$$\varepsilon_{\mathrm{mca}}\big(C_s^{\equiv s}, \delta\big) > \frac{s}{2^{21}}\big(1 - 2^{-164}\big) \geq 2^{-21}\big(1 - 2^{-164}\big).$$

Moreover, the same lower bound holds for $\varepsilon_{\mathrm{ca}}(C_s^{\equiv s}, \delta)$ on the CS-admissible subinterval

$$\delta \in \Big[\frac{1}{2} - 2^{-7},\, \frac{1}{2} - \frac{1}{n_s}\Big).$$

Every row of the deployed parameter family therefore fails the 2^{-128} target at gap 2^{-7} below capacity.

Proof. For each $s \leq 2^{12}$: $256 \mid n_s$ (as $n_s \geq 2^9$), $a_s = n_s/256 = 2^{13}/s$ divides $k_s = 2^{20}/s$, and $(1-\rho)\cdot 256 = 128 \geq 3$. Hypothesis (1) for C_s: $|B|(q/k_s + 1) < 2^{31}(2^{166}\, s + 1) < 2^{198}\, s \leq 2^{210} \leq 2^{247} \leq \binom{256}{130}$. Theorem 4.1 gives the MCA lower bound for C_s throughout $[\frac{1}{2} - 2^{-7}, \frac{1}{2})$, and the CA lower bound throughout $[\frac{1}{2} - 2^{-7}, \frac{1}{2} - 1/n_s)$. Lemma 3.5 transfers the corresponding lower bounds to the s-interleaved code. □

Proposition 5.6 (independent slacked variant via Theorem 2.7)**.** *In the setting of Theorem 4.1, assume instead $(1-\rho)N \geq 2$, $n \geq 3N$, and*

$$\binom{N}{\rho N + 1} \;\geq\; |B| \cdot q.$$

Then

$$\varepsilon_{\mathrm{ca}}\Big(C,\; 1 - \rho - \frac{1}{N} + \frac{2}{n},\; 1 - \rho - \frac{1}{n}\Big) \;\geq\; \frac{1}{2n}.$$

In particular, with $N = N_\rho$ as in Corollary 5.1 and $n \geq 3N_\rho$, every challenge-envelope instantiation satisfies this slacked CA lower bound, and $\frac{1}{2n} \geq 2^{-45} \gg 2^{-128}$.

Proof. Set $\delta_0 := 1 - \rho - \frac{1}{N}$. By $(1-\rho)N \geq 2$, $\delta_0 \in (0, 1-\rho)$. By $n \geq 3N$, we also have $\delta_0 + 2/n \leq 1 - \rho - 1/n$, so the parameters in Theorem 2.7 are admissible. By Lemma 3.1(i), some u_z carries at least $\binom{N}{\rho N+1}/|B| \geq q$ codewords of C within radius δ_0, so $\Lambda(C, \delta_0) \geq q$. The contrapositive of Theorem 2.7 gives

$$\varepsilon_{\mathrm{ca}}\Big(C, \delta_0 + \frac{2}{n}, 1 - \rho - \frac{1}{n}\Big) \geq \frac{1}{2n}.$$

For the numerical claim, the worst case of $\log_2 \binom{N_\rho}{\rho N_\rho + 1}$ over the four rates is at $\rho = 1/8$, where the entropy bound used in Table 1 certifies $\log_2 \binom{1024}{129} \geq 1024\, \mathrm{H}_2(129/1024) - \log_2 1025 \geq 549$ (the exact value is ≈ 554.7); hence $\binom{N_\rho}{\rho N_\rho + 1} \geq 2^{549} > 2^{512} > |B|q$ throughout $q < 2^{256}$. Finally $n = k/\rho \leq 16k \leq 2^{44}$, so $1/(2n) \geq 2^{-45}$. □

Remark 5.7. Proposition 5.6 is logically independent of Theorem 2.6 and yields a weaker error ($\frac{1}{2n}$ versus $\frac{1}{2k}$) in the *maximally slacked* form of correlated agreement — the internal radius sits at capacity minus $\frac{1}{n}$ — which is the weakest failure mode one can certify. That even this fails by 83 orders of magnitude (in bits) above the 2^{-128} target adds robustness to Corollary 5.1 against any repair of the imported theorems.

6 Subfield confinement and the shape of certifying lines

The deployed instantiations run the protocol over an extension $\mathbb{F}/B$ while keeping the domain in the base: $D \subseteq B$. The subfield-confinement theorem of [Cho26b, Thm. "Subfield confinement"] pins down where bad slopes can live for lines whose words are base-rational: all support-wise MCA-bad slopes lie in B. The following lemma refines this by tracking two radii, so that it also covers the proximity-loss form of correlated agreement (Definition 2.1); part (b) is the statement of [Cho26b], reproved here for completeness in the present notation.

Lemma 6.1 (subfield confinement; refines [Cho26b])**.** *Let $B \subseteq \mathbb{F}$ be finite fields, $D \subseteq B^\times$, $C = \mathrm{RS}[\mathbb{F}, D, k]$, and let $f, g \in B^n \subseteq \mathbb{F}^n$ be B-valued words. Then:*

(a) *for any $0 \leq \delta_{\mathrm{fld}} \leq \delta_{\mathrm{int}} < 1$, every CA-bad slope for (f, g) at radii $(\delta_{\mathrm{fld}}, \delta_{\mathrm{int}})$ lies in B;*

(b) *for any $\delta \in (0, 1)$, every MCA-bad slope for (f, g) at radius δ lies in B.*

Consequently the bad-slope set of any B-valued pair has density at most $|B|/q$ in $\mathbb{F}$. The same holds for pairs $(\lambda f, \lambda g)$ with $\lambda \in \mathbb{F}^\times$ and f, g B-valued, by $\mathbb{F}$-linearity of C.

Proof. Fix a B-basis $1 = \beta_0, \beta_1, \ldots, \beta_{m-1}$ of $\mathbb{F}$. Every word $w \in \mathbb{F}^n$ decomposes uniquely as $w = \sum_i \beta_i w_i$ with $w_i \in B^n$, and every codeword $c \in \mathrm{RS}[\mathbb{F}, D, k]$ decomposes as $c = \sum_i \beta_i c_i$ with $c_i \in \mathrm{RS}[B, D, k]$: write the coefficient vector of the underlying polynomial in the basis and evaluate at $D \subseteq B$.

Let $z \in \mathbb{F} \setminus B$, say $z = \sum_i z_i \beta_i$ with $z_{i^*} \neq 0$ for some $i^* \geq 1$. We first record the support-wise consequence. Suppose that $c \in C$ agrees with $f + zg$ on some set $S \subseteq D$. Since $f_j, g_j \in B$, the basis components of $(f + zg)_j$ are $f_j + z_0 g_j$ (component 0) and $z_i g_j$ (components $i \geq 1$). Agreement on S holds componentwise, so on S:

$$f + z_0 g = c_0, \qquad z_{i^*}\, g = c_{i^*}.$$

Hence g agrees on S with $z_{i^*}^{-1}c_{i^*} \in C$, and then $f = (f+z_0g)-z_0g$ agrees on S with $c_0-z_0z_{i^*}^{-1}c_{i^*} \in C$. Thus every support on which the line point $f + zg$ is explained by C also explains the pair (f, g) on that same support.

For CA, if $\Delta(f + zg, C) \le \delta_{\mathrm{fld}}$, choose such an S with $|S| \ge (1 - \delta_{\mathrm{fld}})n$; the preceding paragraph gives $\Delta_2((f, g), C^{\equiv 2}) \le \delta_{\mathrm{fld}} \le \delta_{\mathrm{int}}$, so z is not CA-bad. For MCA, the preceding paragraph rules out every possible bad support S in the support-wise definition. Hence z is not MCA-bad either. □

Corollary 6.2 (certifying lines over extensions are genuinely $\mathbb{F}$-valued)**.** *At the parameters of Corollary 5.4, $|B|/q = p^{-5} < 2^{-154}$. Hence any pair (f, g) whose CA- or MCA-bad-slope density at some radius*

$$\delta \in \left[\frac{1}{2} - 2^{-7}, \frac{1}{2}\right)$$

exceeds 2^{-154} contains a word that is not B-valued, even after removing a common scalar factor. In particular, any pair attaining the MCA lower bound $> 2^{-22}$ of Corollary 5.4 is genuinely $\mathbb{F}$-valued. On the CS-admissible subinterval

$$\delta \in \left[\frac{1}{2} - 2^{-7}, \frac{1}{2} - 2^{-21}\right),$$

the same statement applies to pairs attaining the CA lower bound. All explicit failure witnesses currently known on these domains (the locator lines of [Cho26a, Cho26b] and the fiber words of Lemma 3.1) are B-rational and are therefore inert here: the failure certified by Corollary 5.4 is fiber-borne and, at present, non-constructive.

Proof. Immediate from Lemma 6.1 and Corollary 5.4: a B-valued pair, up to common scalar, has CA- and MCA-bad-slope density at most $|B|/q = p^{-5} < 2^{-154} < 2^{-22}$. Thus any pair witnessing the deployed lower bound cannot be of this form. □

Remark 6.3. Lemma 6.1 and Corollary 6.2 sharpen the division of labor between the two grand challenges of [ABF26] over deployed extensions: the *list* side does not deflate ($\mathrm{RS}[B, D, k] \subseteq \mathrm{RS}[\mathbb{F}, D, k]$, so all base-field list lower bounds, including Lemma 3.1, persist verbatim), while every sub-2^{-128} *MCA* attack by explicit base-rational witnesses is killed by confinement — and yet MCA still fails at 2^{-22}, through lines whose existence Theorem 2.6 extracts non-constructively from the surviving list mass. The deflation and decoupling corollaries of [Cho26b] record the same division at the no-loss radius, using the same corrected rounding $p^{-5} = 2^{-154.94\ldots} < 2^{-154}$ as in Corollary 6.2. Exhibiting such lines explicitly is Open Problem 7.1.

7 Discussion

Values versus fibers. Above the per-rate value-set reach of [Cho26b] (roughly 2^{150}–2^{162}), the value route — showing that the slope multiset $e_1(\ell^{\wedge}Q)$ covers a 2^{-128} fraction of $\mathbb{F}$ at a fixed large prime — runs into the per-fiber collision problem ([Cho26b, Rem. "The remaining band"]): characteristic-zero locator values can collide arbitrarily badly under reduction at a worst-case prime. [Cho26b] therefore closed the band at the δ^* grade only conditionally and with weaker constants — error 2^{-42} at gap 2^{-11}, assuming $|\mathbb{F}| \ge 2n$, via quotient-core lists and the same conversion Theorem 2.6 — leaving the *error-one* grade open. The present route also never counts values, and improves each parameter of that cap: pigeonhole guarantees one heavy fiber regardless of how the values $e_1(A)$ distribute, Theorem 2.6 converts heavy fibers

into correlated-agreement error directly, and the subfield pigeonhole carries the construction to extension fields with no condition relating n and $|\mathbb{F}|$. The cost relative to the error-one results of [Cho26a, Cho26b] below 2^{150} is quantitative: error $\approx \frac{1}{2k}$ instead of $1 - o(1)$, and gap 2^{-9} (uniformly 2^{-10} across all four rates) instead of $\frac{1}{64}$. For the prize question — is $\varepsilon_{\mathrm{mca}} \le 2^{-128}$ achievable above the stated gaps? — this cost is immaterial.

Limits of the method. Hypothesis (1) needs $N\,\mathrm{H}_2(\rho) \gtrsim 2\log_2 q - \log_2 k$, so the smallest certifiable gap is

$$\frac{2}{N} \;\approx\; \frac{2\,\mathrm{H}_2(\rho)}{2\log_2 q - \log_2 k} \;\approx\; \frac{\mathrm{H}_2(\rho)}{\log_2 q},$$

and the certified error is capped at $\approx \frac{1}{k}$ by the shape of Theorem 2.6 (any $\eta < 1$). Since q enters only as the size of the field from which the line challenge γ is drawn, this is also a protocol-facing floor: enlarging the challenge field cannot push a CS25-certifiable failure below gap $\approx \mathrm{H}_2(\rho)/\log_2 q_{\mathrm{line}}$, which quantitatively bounds the "larger challenge field" fallback in certificate frameworks. The construction also needs $N \le n$, i.e. $\log_2 q \lesssim \mathrm{H}_2(\rho)\, n/2$; this covers the entire cryptographic regime $\mathrm{poly}(n) \le q \le 2^{o(n)}$, but for $q \ge 2^{\mathrm{H}_2(\rho)n/2}$ the method is vacuous and the value-route caps of [Cho26a, Cho26b] are the only ones known. Within the band, the *error-one* question at gap $\frac{1}{64}$ (Open Problem 7.2) remains exactly as open as before: the present result bounds δ^*, not the failure profile near capacity.

Comparison with the cyclotomic sieve. The sieve of [Cho26a] produces, for infinitely many primes, failures of error $(\log p)^{-6}$ at gap $\Theta(1/\log p)$, but its prime set is given by Siegel–Walfisz and is ineffective. Corollary 5.1 trades error strength for universality: every field below 2^{256}, fully effective, and elementary given Theorem 2.6. The two results triangulate the failure landscape from different directions.

Pressure on the packing conjecture. By the exact normal form of [Cho26b], $\varepsilon_{\mathrm{mca}}(C, \delta) = \max_t \Lambda_t^{\mathrm{NC}}(\delta)/q$, where Λ^{NC} is a noncontained residue-line packing number. Corollary 5.4 therefore forces $\max_t \Lambda^{\mathrm{NC}} \ge (q-n)/2k \approx 2^{164}$ at gap 2^{-7} on the deployed sextic — superpolynomially many noncontained residue lines at $n = 2^{21}$. This puts no direct pressure on the final MCA conjecture of [Cho26b, Conj. "Final floor- and quotient-corrected MCA asymptotic"]: at gap 2^{-7} both of its corrected-reserve hypotheses fail ($\sigma = 2^{14}$ lies below $Cn/\log_2 n \approx 2^{16.6}$, and $\sigma \log_2 q_{\mathrm{gen}} = 2^{14} \cdot 31 \approx 2^{19}$ is far below $\log_2 \binom{n}{k+\sigma} \approx 2^{21}$), so the conjectured ceiling of the form $\left(n^{1+o(1)} + 2^{(\beta/\mathrm{H}_2)\,Q_{\mathrm{H}_2}(\eta)}\right)/q$ is simply not asserted there — the packing mass extracted through Theorem 2.6 lives in the region the framework already marks as below-floor. The genuine structural question this note raises for the framework of [Cho26b] is whether CS25-extracted fiber mass persists *above* the corrected reserve: whether, at reserves where the conjecture does apply, the fiber route still produces packings beyond the $n^{1+o(1)}$ aperiodic term plus the quotient-periodic term, which would force an additional fiber-borne term into the conjectured ceiling. This is open.

Verification caveat. Theorem 2.6 is derived directly from [CS25, Thm. 2], including its integer-radius condition $f < n - k - 1$, while Theorem 2.7 remains a separate slacked fallback imported through the BCHKS/ABF route. Both underlying sources are recent, and readers relying on the present results should consult them directly. The main composition is deliberately shallow — only the displayed CS25 consequence is used — and Proposition 5.6 provides a second,

independent route; still, readers should treat Corollary 5.1 as conditional on the correctness of the direct CS25 conversion recorded in Theorem 2.6.

Open problems.

Open Problem 7.1. Exhibit explicit pairs (f, g) over $\mathbb{F}_{p^6}$ (KoalaBear sextic, parameters of Corollary 5.4) with MCA-bad-slope density $> 2^{-22}$ at gap 2^{-7}. On the CS-admissible subinterval, ask for the corresponding CA-bad-slope density as well. By Corollary 6.2 such pairs exist and are not B-rational; the natural candidates are residue-line normal forms of [Cho26b] with denominator $E \in \mathbb{F}[X]$ not defined over B.

Open Problem 7.2. Error one in the band: does $\varepsilon_{\mathrm{mca}}(C, 1 - \rho - \frac{1}{64}) = 1 - o(1)$ hold for all primes $2^{150} \le p \le 2^{256}$? This is the per-fiber collision problem proper; Corollary 5.1 caps δ^* there but says nothing above error $\approx \frac{1}{2k}$.

Open Problem 7.3. Improve the threshold of Theorem 2.6: any strengthening of the hypothesis range $\varepsilon_{\mathrm{ca}} \le \eta/k$ toward $\varepsilon_{\mathrm{ca}} \le \eta \cdot \mathrm{polylog}(q)/k$, or of the list conclusion, immediately tightens the certified error of Theorem 4.1 toward $1/\mathrm{poly}\log$; the gap, by contrast, is governed by the binomial in (1) and would need a denser fiber construction to shrink.

Open Problem 7.4. Beyond multiplicative smoothness: Lemma 3.1 uses only that D is a union of complete fibers of a low-degree map with lacunary power structure. Carry the composition to isogeny-smooth domains — in particular the circle-group domains of Circle STARKs over $p = 2^{31} - 1$, where Chebyshev locators $\prod_b (T_{2^j}(x) - x_b)$ play the role of $\prod_b (X^a - b)$ — to obtain the same universal cap for that deployed family.

Acknowledgements. This note composes results of Crites–Stewart [CS25] and of [BCHKS25] with the locator framework of [Cho26a, Cho26b], following the synthesis of Arnon–Boneh–Fenzi [ABF26]; it sharpens the conditional universal-cap theorem and refines the subfield-confinement theorem of [Cho26b].

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. https://eprint.iacr.org/2026/680.

[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. https://eprint.iacr.org/2025/2055.

[Cho26a] P. Chojecki. Capacity-edge obstructions to Reed–Solomon mutual correlated agreement over smooth multiplicative domains. Manuscript, 2026.

[Cho26b] P. Chojecki. Slack, quotient cores, and the entropy gap for smooth-domain Reed–Solomon codes: list fibers, cyclotomic rigidity, Galois-amplified collisions, and the corrected slack mutual-correlated-agreement theory. Manuscript, merged corrected edition, 2026.

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