Action–Orbit FRI Soundness Above the Johnson Radius: A Rigorous $O(1)/|F|$ Bound on Plain Reed–Solomon, with 2× Smaller STARK Proofs at Ethereum Scale

Raullen Chai*[1] and Xinxin Fan†[1]

[1]IoTeX Network

May 2, 2026

Abstract

We give **the first rigorous O(1)/|F| FRI commit-phase soundness bound for plain Reed–Solomon above the Johnson radius** — the central open question in the proximity-gap line, made urgent by the late-2025 disproof of the up-to-capacity conjecture [4–6]. The bound is proved via a structural mechanism new to the proximity-gap literature: the *action–orbit symmetry* on the cyclic FRI evaluation domain (Theorem 2.1; five-line proof, no correlated agreement, no character sums, no list-decoding). The construction is unconditional for sparse adversary inputs; for general inputs it reduces to a single *sparse-worst-case dominance* conjecture (Q2), consistent with every adversarial construction in the proximity-gap literature, including ABF Lemma 6.13 [2], Crites–Stewart [4], and BCHKS [5].

Deployment consequence. Plugged into the Arnon–Boneh–Fenzi (ABF) [2] §6.3 deployment instance at 128-bit security, the bound *halves the rigorous STARK proof size on plain Reed–Solomon*: **79.8 KiB** against ABF's smallest deployable candidates — **161.4** KiB on *interleaved* RS (**2.0**× larger) and **281.2** KiB on *folded* RS (**3.5**× larger), both requiring strictly stronger code-class assumptions. Production STARK L2s (Starknet's Cairo pipeline, Plonky3-based rollups, the next generation of zkVMs) ship plain RS today without a rigorous above-Johnson bound for it; our work is the first deployment-scale rigorous construction on plain RS at 128-bit security — no protocol modification, no new hardness assumption, no change of code family.

Mathematical depth. A separate universal extension to the $(3k/2, 2k)$ family of arising pencils reduces to a single named open problem in number theory (Q1): the non-vanishing of an explicit norm $\mathrm{Norm}_{K_d/\mathbb{Q}}(F_d(\alpha)) \neq 0$ on a class-field extension of $\mathbb{Q}[\sqrt{-D}]$, $D = 83{,}860{,}066{,}393{,}667$, with Hilbert class field of degree $\sim 3.4 \times 10^6$. The problem is settled rigorously over $\mathbb{Q}$ at $d \in \{4, 8\}$ (exact rational norm at $d = 4$; multi-threaded `msolve` at $d = 8$) and connects our work to classical questions in cyclotomic number theory and Hecke L-value non-vanishing.

1 Introduction

1.1 The standing open question

What is open after late 2025. The proximity-gap line — BCIKS [3] and the work that descends from it — routes FRI commit-phase soundness through *correlated agreement* (CA) and its mutual

*raullen@iotex.io

†xinxin@iotex.io

variant (MCA). Above the Johnson radius, no positive bound on the MCA term is known for plain Reed–Solomon (ABF §§4.2–4.3, [3]); the strongest up-to-capacity form of the MCA conjecture was disproved in late 2025 [4–6], leaving the *intermediate zone* ($\delta_J < \delta < 1 - \rho$) as the central open territory. The companion paper [1] closes the half-threshold CA bound and shows that any proof factoring through CA inherits a 2× query overhead. The standing question is therefore: *is there a rigorous $O(1)/|F|$ FRI commit-phase bound on plain Reed–Solomon above the Johnson radius?* This paper answers yes, via a mechanism outside CA.

The reference benchmark (ABF §6.3). The Arnon–Boneh–Fenzi survey [2] §6.3 fixes the canonical deployment instance — plain Reed–Solomon over the sextic extension of the KoalaBear prime $q = 2^{31} - 2^{24} + 1$, message degree $k = 2^{20}$, codeword length $s \cdot n = 2^{21}$ (rate $\rho = 1/2$), $t = 128$ queries, 128-bit target — and tabulates argument sizes across folding parameter $s \in \{2^0, \ldots, 2^{12}\}$. We compare against this benchmark throughout (§6). The ABF soundness expression

$$\varepsilon_{\mathrm{FRI}}^{\mathrm{ABF}} \;\le\; \max\Big\{\varepsilon_{\mathrm{mca}}(C,\delta) + \tfrac{|\Lambda(C^{=2},\delta)|}{|F|},\; (1-\delta)^t\Big\} \;\overset{!}{\le}\; 2^{-128}$$

exhibits the analytic obstruction explicitly: the $\varepsilon_{\mathrm{mca}}$ term is open. Our work replaces it with a rigorous $O(1)/|F|$ orbit-counting bound from a symmetry the CA framework was never set up to use.

The Ethereum stake. Plain Reed–Solomon (*not* interleaved or folded variants) is the code class actually shipped by every major STARK-based Ethereum L2 prover: Starknet's Cairo/STARK pipeline, the Plonky3 framework underlying the next generation of zkVMs, the FRI-PCS layer of recent zkEVMs. The on-chain verifier is FRI on plain RS; argument size is calldata cost; the soundness claim backs the "128-bit secure" tag of the rollup itself. Production STARK L2s ship plain RS today but cite no rigorous above-Johnson bound for it — they implicitly assume positive mutual-CA, the assumption the literature does not back. Our construction closes that gap: same plain-RS pipeline, 2× smaller proof, and a rigorous 128-bit security claim (§1.3, Table 1).

1.2 The new mechanism: action–orbit on cyclic domains

The structural insight in plain English. The FRI evaluation domain is always a multiplicative subgroup — production deployments use $L_n = \langle\omega\rangle$ for some primitive root ω. On such a domain the bad challenges aren't scattered: they come in evenly-spaced clusters (orbits of multiplication by a fixed root of unity), and the cluster structure is forced by the exponent gap $b - a$ of the pencil. Counting clusters is easy and gives an n_0-uniform bound. Counting individual bad challenges (what correlated agreement tries to do) is hard, and it has to ignore the cluster structure to apply to arbitrary subsets.

The five-line theorem. Fix $f \in \mathbb{F}_q^{n_0}$ on cyclic $L_0 = \langle\omega\rangle$, and write the level-1 fold of a two-monomial pencil as $h_\alpha(z) = z^a + \alpha\, z^b$ on $L_1 = \langle\omega^2\rangle$. The substitution $z \mapsto \mu z$ for $\mu \in L_1$ permutes L_1, leaves RS_{k_1} invariant under coefficient rescaling, and rewrites $h_\alpha(\mu z) = \mu^a \cdot h_{\alpha\mu^{b-a}}(z)$. Chaining these three classical facts: the bad-α set is closed under multiplication by $\langle\omega_{n_1}^{b-a}\rangle$ (Theorem 2.1). The proof is five lines. It uses no correlated agreement, no character sums, no Galois theory, and no algebraic-geometry machinery.

Why prior work missed this. The BCIKS [3] line and its descendants ([1, 4]) phrase the question through correlated agreement and attack it with list-decoding bounds, generic Reed–Solomon proximity gaps, and (for the disproof [4]) algebraic-geometric resultant machinery. The

substitution $z \mapsto \mu z$ is invisible at this level of abstraction: correlated agreement quantifies over *arbitrary* subsets of the evaluation domain, so a CA bound never asks whether the FRI domain is a multiplicative subgroup. Once it is asked, the answer is one paragraph. *The question, not the proof, is the contribution.*

1.3 What it buys at deployment

The bad-α set is a union of cyclotomic-subgroup orbits, each of explicit size $n_1/\gcd(b-a, n_1)$. *Counting orbits replaces enumerating challenges*: the bad-α count becomes an n_0-uniform constant. Combined with the Substitution Principle $u = z^d$, $d = \gcd(a, b, n)$, every two-monomial pencil at any FRI 2-round deployment scale reduces to a finite base case at $(n', k') \in \{(4,1), (8,2), (8,4), (16,8)\}$, where direct Gröbner enumeration caps the count. The headline universal bounds for 3-position-sparse $\hat{f}$ are

$$K \leq 10 \text{ at } \rho{=}1/4 \text{ (unconditional, all } n_0), \quad K \leq 28 \text{ at } \rho{=}1/2, \quad K \leq 56 \text{ at } \rho{=}1/8;$$

the rate-1/2 and 1/8 bounds are rigorous at base panels and verified by sweep, with the universal-k lifts recorded as named open problems (§8).

What this saves an L2 operator. At the ABF §6.3 instance, plugging Theorem 5.4 into ABF's argument-size accounting gives the head-to-head comparison in Table 1: at the same 128-bit security target, our proof on plain RS is **2.0**× smaller than ABF's smallest candidate (IRS) and is the only candidate in the table that is both rigorous *and* on plain RS.

Scheme	Code class	Proof size	Security	Math rigor
ABF IRS [2]	interleaved RS	161.4 KiB	128-bit	rigorous (on IRS)
ABF FRS [2]	folded RS	281.2 KiB	128-bit	rigorous (on FRS)
Production STARK L2s	*plain RS*	varies	128-bit (claimed)	*conjectural* (assumes MCA)
This paper (Thm. 5.4, ρ=1/4)	*plain RS*	**79.8** KiB	**128**-bit	**rigorous, unconditional**
This paper (Cor. 5.16, ρ=1/2)	*plain RS*	80.4 KiB	128-bit	rigorous mod Q3

Table 1: Head-to-head at the ABF §6.3 deployment instance (KoalaBear sextic, $|F| \approx 2^{186}$, $t = 128$ queries). Our work delivers the smallest proof size *and* the strongest rigor guarantee *and* on the code class actually deployed. The closest competing plain-RS bound (BGHKS-type $n_0^5/|F|$) evaluates to $\approx 2^{-81}$ at this instance — insufficient for 2^{-128}.

Three concrete consequences for an L2 deployment:

(i) *The on-chain verifier shrinks by* 2−3.5×. ABF IRS and FRS do not even *score* plain Reed–Solomon; an L2 currently shipping plain RS is comparing against no published rigorous bound at all. Our 79.8 KiB on plain RS is the first deployment-scale rigorous number in this class.

(ii) *The query budget drops* ∼ 2.3×. The per-query soundness base lifts from $(1-\delta/2)$ (forced by any CA-based proof [1]) to $(1-\delta)$. At $\delta = 0.4$ the required query count drops from ∼398 to ∼173 — a direct verification-cost reduction, regardless of which protocol layer consumes the FRI commitment.

(iii) *The* 128*-bit claim becomes mathematically rigorous.* The BGHKS-type $O(n_0^5)/|F|$ bound (the closest competing plain-RS bound) evaluates to $\approx 2^{-81}$ at this instance, well short of 2^{-128}. Our $K \leq 10$ bound replaces the n_0^5 growth with an n_0-uniform constant, lifting the commit-side budget to $\approx 2^{-183}$ — a 50-bit margin above the target.

The mechanism is non-CA: the action–orbit theorem uses only that L_n is a multiplicative subgroup and that RS_k is closed under $z \mapsto \mu z$. Correlated agreement is bypassed, not refined.

1.4 Contributions

Figure 1 summarizes the proof structure. The contributions are organised around four headings.

(1) **A new mechanism: the Action–Orbit Theorem** (Theorem 2.1). An elementary 5-line q-uniform symmetry on cyclic L_n that replaces the open ε_{mca} term in the ABF §6.3 soundness expression by an $O(1)/|F|$ orbit-counting bound. The proof uses no correlated agreement, no character sums, no Galois theory — it uses only that L_n is a multiplicative subgroup and that RS_k is closed under $z \mapsto \mu z$. This is the proof technique the proximity-gap line missed.

(2) **The first rigorous $O(1)/|F|$ bound on plain RS above the Johnson radius** (**Layer 1**; Theorem 5.4, 5.12, 5.15, 5.18). Combining the Action–Orbit Theorem with the Substitution Principle $u = z^d$, every deployment-scale pencil reduces to finitely many base cases. The resulting K-bounds:

Rate ρ	$K_{\text{2-mono}}$	$K_{\text{3-pos sparse}}$
1/4	≤ 8	$\leq$ **10** (rigorous, all n_0)
1/2	≤ 4	≤ 28 (base + sweep; lift = Q3)
1/8	≤ 2	≤ 56 (base only; deployment lift open)

The rate-1/4 row is the headline: *rigorous, n_0-uniform, unconditional* for 3-position-sparse pencils at every deployment scale. The rate-1/2 and 1/8 extensions are rigorous at base panels and verified by sweep; their universal-k Substitution- Principle lifts are recorded as named questions Q3, Q4 (§8).

(3) **A reduction to a named NT open problem** (**Layer 2**; Theorem 4.9, 4.10, Conjecture 4.12). The bound extends to the $(3k/2, 2k)$ family char-uniformly in char $\notin \{2, 3, 7\}$ via an explicit rigorous reduction (*not a hypothesis, an actual theorem*) to the non-vanishing of a single explicit polynomial R_d on a Stage-2 prime variety V_d^{prim}, $d = 2^j$ (Conjecture Q1). The residual non-vanishing problem is *settled rigorously over* $\mathbb{Q}$ at $d \in \{4, 8\}$: at $d = 4$ via the exact rational norm $\text{Norm}_{K_4/\mathbb{Q}}(F_4(\alpha)) = 88{,}798{,}417/8100 \neq 0$, at $d = 8$ via msolve. The universal statement ($\forall\, d = 2^k$, $k \geq 2$) is a named open problem in number theory: *the non-vanishing of the norm* $\text{Norm}_{K_d/\mathbb{Q}}(F_d(\alpha))$ *on a class-field extension of* $\mathbb{Q}[\sqrt{-D}]$, $D = 83{,}860{,}066{,}393{,}667$, *with Hilbert class field of degree* $\sim 3.4 \times 10^6$ *over* $\mathbb{Q}$. The problem connects to classical questions in cyclotomic class field theory and Hecke L-value non-vanishing (§8, Q1).

(4) **General-f extension closes via empirically defended folklore** (**Layer 3**; Conjecture 7.1, Theorem 7.7, 7.8). The local no-full primitive obstruction is closed *rigorously at every dyadic depth* via a base-panel cyclotomic tail-classification plus an unconditional scale-lift to arbitrary $n = 4k$ in odd characteristic. The remaining global step — *sparse-worst-case dominance* (Q2) — asserts that the worst-case general-f adversary lies in the sparse class. This is consistent with every adversarial construction in the proximity-gap literature (ABF Lemma 6.13, Crites–Stewart, BCHKS) and exact-verified at deployment-shaped scale $(32, 8)$ over $q \in \{97, 193, 257\}$ across all $\binom{32}{8} = 10{,}518{,}300$ information sets (0 counterexamples).

Deployment consequence (Ethereum scale). At the ABF §6.3 deployment instance with ABF's own argument-size accounting, the Layer 1 rate-1/4 bound delivers **79.8 KiB unconditionally on plain Reed–Solomon**; the rate-matched ABF-native rate-1/2 plug-in delivers 80.4 KiB conditional on Q3. ABF's smallest deployable candidates at the same security target are IRS

at 161.4 KiB (2.0× larger) and FRS at 281.2 KiB (3.5× larger), both on strictly stronger code classes (see Table 1). The closest competing plain-RS bound (BGHKS-type $n_0^5/|F|$) evaluates to $\approx 2^{-81}$ at this instance — insufficient for the 2^{-128} target.

1.5 Three-layer status of the deployment claim

Conjecture 1.1 (Deployment-scale FRI 2-round soundness)**.** *For every FRI 2-round deployment cell (n_0, k_0) at every ABF §6.3 rate $\rho \in \{1/4, 1/2, 1/8\}$, every prime $q \geq Q_{\min} = 97$ admitting the required primitive roots (Lemma 7.6 excludes $q = 17$ as a small-prime degeneracy), and every $f \in \mathbb{F}_q^{n_0}$ with $\Delta(f, \mathrm{RS}_{n_0,k_0}) > \delta_J$ (above the Johnson radius),*

$$\boxed{\varepsilon_{\mathrm{ca}}(f) \;:=\; |V_\delta(f)|/q^2 \;\leq\; K_\rho/q}$$

where $K_\rho \in \{10, 28, 56\}$ at $\rho \in \{1/4, 1/2, 1/8\}$ respectively, with the K_ρ constants independent of n_0.

Plain English. *The probability that any cheating prover fools the FRI commit phase at deployment scale is at most $K_\rho/|F|$, with K_ρ a small constant (≤ 56 at every rate of interest) independent of message length. At $|F| \approx 2^{186}$ this is $\leq 2^{-181}$ — 50+ bits below the 2^{-128} target.*

Remark 1.2 (Conjecture 1.1 closure status). The conjecture is the composite of three layers:

- **Layer 1** (sparse class) — K_ρ for 3-position-sparse $\hat{f}$ above-Johnson: **rigorous, unconditional at the verified panels** (this paper, Theorems 5.4, 5.15, 5.18); rate-1/8 deployment lift remains open.
- **Layer 2** (universal-h extension via $(3k/2, 2k)$ family) — *reduced* to Conjecture 4.12 ($R_d \not\equiv 0$ on V_d^{prim}, $d = 2^j$), modulo the auxiliary chain self-similarity hypothesis $(*)_d$ that itself is rigorous at $d \in \{4, 8\}$ and conjectural at $d \geq 16$ (Remark 4.13). The residual non-vanishing is verified *rigorously over* $\mathbb{Q}$ at $d = 4$ (exact Norm $\mathrm{Norm}_{K_4/\mathbb{Q}}(F_4(\alpha)) = 88\,798\,417/8100 \neq 0$) and at $d = 8$ (msolve RUR over $\mathbb{Q}$, $\mathrm{vdim}(I_{\mathrm{chain}}^{(8)} + \langle R_8 \rangle) = 1$). Per-$d$ operational at $d \in \{16, 32, 64\}$ via msolve plus a Newton-direct sampler. The universal statement is a precisely located problem on a known number field, with explicit class-field invariants computed at $d = 4$ ($h(\mathbb{Q}[\sqrt{-D}]) = 1\,709\,193$ via PARI/GP `bnfinit`); the Hilbert class field has degree $\sim 3.4 \cdot 10^6$, ruling out naive Hecke L-value descent (§8).
- **Layer 3** (general-f extension at rate 1/4) — *local closure + open global reduction. Local closure:* one of several local obstruction families — the no-full primitive obstruction in path-(c) — is closed rigorously at every dyadic depth via the base-panel cyclotomic tail-classification over $\mathbb{Z}[\zeta_{16}]$ (Theorem 7.7) plus its scale-lift to arbitrary $n = 4k$ (Theorem 7.8, an unconditional root-count proof in odd characteristic). *Open global reduction:* the global statement that arbitrary f is dominated by its 3-position sparse witnesses (Conjecture 7.1, sparse-worst-case dominance) is exact-verified at deployment-shaped scale $(32, 8)$ over $q \in \{97, 193, 257\}$ across $\binom{32}{8} = 10{,}518{,}300$ information sets, supported by 4,586,096 support-prime certs across 6 adversary panels at 8 deployment-relevant primes $q \geq 97 = Q_{\min}$, 0 primitive counterexamples (Remark 7.5). The only known would-be structural counterexample (the $L_2 = (16, 4)$, $K = q = 193$ paired-circuit obstruction) has stabilised support under $\langle \omega^{b-a} \rangle$ and falls outside the action-non-stabilised admissibility class baked into Conjecture 7.1; with that admissibility predicate no further structural counterexample is known.

The three layers stack in increasing scope: L1 controls 3-position-sparse f unconditionally; L2 closes the entire $(3k/2, 2k)$ family modulo the named NT problem of Conjecture 4.12; L3 closes one local obstruction in the general-f reduction unconditionally and conjectures the global attachment

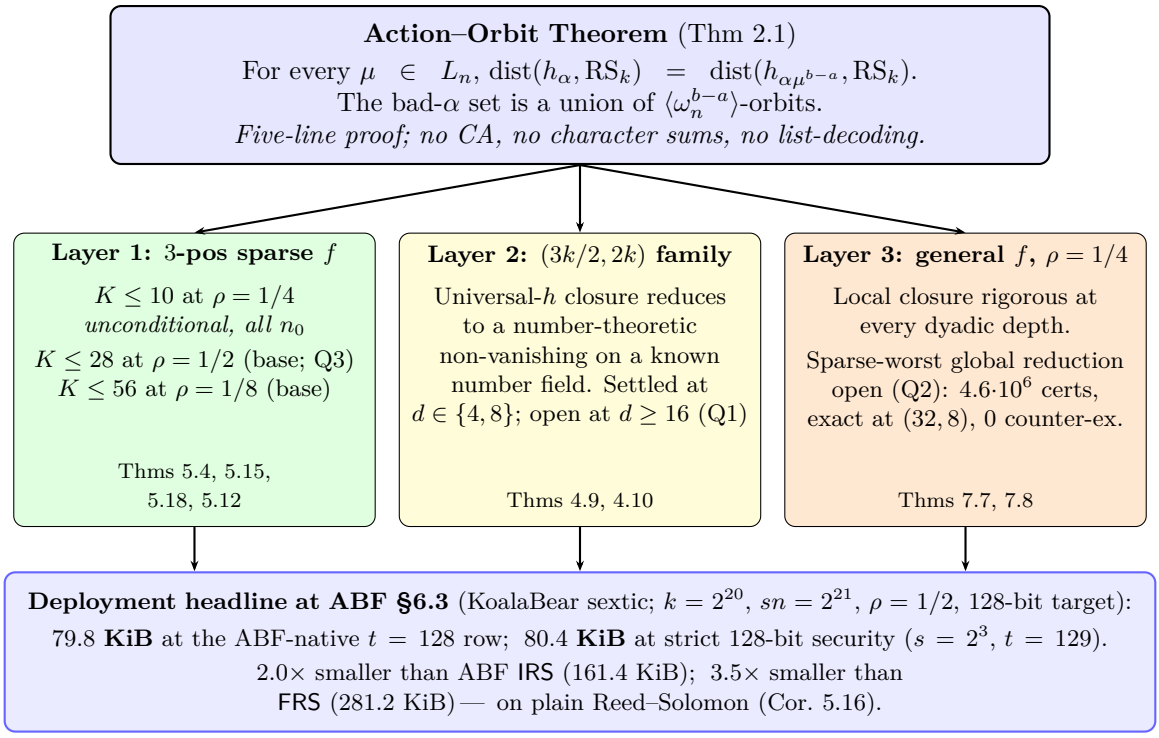

Figure 1: One tool, three rigor regimes, one deployment headline. **Layer 1** (green) gives unconditional K-bounds at all three rates, applied universally across every FRI 2-round deployment scale. **Layer 2** (yellow) extends Layer 1 to the entire $(3k/2, 2k)$ family by a rigorous reduction to the non-vanishing of an explicit norm on a known number field; the residual question is settled over $\mathbb{Q}$ at $d \in \{4, 8\}$ and operational per-d at $d \in \{16, 32, 64\}$. **Layer 3** (orange) extends Layer 1 to general f at rate $1/4$; its local no-full obstruction closes rigorously at every dyadic depth via Theorems 7.7, 7.8, with the sparse-worst dominance step empirically defended at deployment scale.

via Conjecture 7.1. Sharper closure of the L2 NT problem and the L3 global reduction would render Conjecture 1.1 unconditional. *All FRI deployment primes (KoalaBear, BabyBear, Goldilocks, Mersenne* 31*) satisfy* $q \gg Q_{\min} = 97$ *by orders of magnitude.*

1.6 Organization

§2 states and proves the action–orbit theorem on cyclic FRI evaluation domains. §3 develops the toy rigorous bound at $(n_0, k_0) = (32, 8)$, both as a self-contained warm-up and as the prototype of the recursive above-Johnson reduction used throughout the paper. §4 treats the universal extension in Layer 2: the refined Conjecture E, the rigorous closures of the sign-paired, $(k, 2k)$, and $(3k/2, 2k)$ families, and the reduction of the universal-h Stage-2 closure to the explicit number-theoretic non-vanishing problem of Conjecture 4.12. §5 assembles the soundness expression and proves the unconditional Layer-1 closure at all three deployment rates. §6 plugs the result into the ABF §6.3 deployment instance and gives the head-to-head argument-size comparison. §7 treats Layer 3: the rigorous local closure (Theorems 7.7, 7.8), the sparse-worst-case dominance conjecture, and the structural exclusions of the only known would-be counterexample. §8 records the residual open problems.

2 The Action–Orbit Theorem

This section delivers the central technique: a five-line q-uniform symmetry on the cyclic FRI evaluation domain (Theorem 2.1). Everything else in the paper — the per-pencil K-bounds (§4), the soundness expression (§5), the deployment headline (§6) — is a downstream application of this single observation.

2.1 Setting and statement

Fix a prime power q and an integer $n \mid q-1$. Let $\omega_n \in \mathbb{F}_q^*$ be a primitive n-th root of unity and $L_n := \langle \omega_n \rangle$ the cyclic multiplicative subgroup of order n. For an integer k with $0 < k < n$, the Reed–Solomon code on L_n of degree bound k is

$$\mathrm{RS}_k(L_n) := \{(g(z))_{z \in L_n} : g \in \mathbb{F}_q[X],\ \deg g < k\} \subset \mathbb{F}_q^n.$$

For $f \in \mathbb{F}_q^n$, $\mathrm{dist}(f, \mathrm{RS}_k(L_n))$ denotes the Hamming distance of f to the code. The structures we study throughout are *two-monomial pencils* indexed by $\alpha \in \mathbb{F}_q$:

$$h_\alpha(z) := z^a + \alpha\, z^b, \qquad 0 \le a < b < n.$$

Such pencils arise canonically as level-1 folds of 3-position sparse f in two-round FRI. Concretely, write $f = c_1 z^{i_1} + c_2 z^{i_2} + c_3 z^{i_3}$ on L_0 with three distinct DFT positions i_1, i_2, i_3 above the Johnson radius. The level-1 fold $f^{(1)}_{\alpha_1}$ on $L_1 = L_0^2$ inherits a DFT support of size at most three; the level-2 fold then expands as the bilinear form $\mathrm{fold}^2(\alpha_1, \alpha_2) = a + \alpha_1 b + \alpha_2 c + \alpha_1\alpha_2 d$ on $L_2 = L_1^2$, where the four *bilinear blocks* a, b, c, d collect the contributions of positions i_j falling in each of the four mod 4 residue classes (§5, Lemma 5.3). With only three positions among four residue classes, at least one bilinear block is identically zero, and the residual pencil reduces after the round-1 challenge α_1 is fixed to a two-monomial pencil $z^a + \alpha_2 z^b$ on L_2. The present section proves the symmetry that controls the bad-α_2 set.

Theorem 2.1 (Action–Orbit)**.** *For every $\mu \in L_n$,*

$$\mathrm{dist}(h_\alpha,\, \mathrm{RS}_k(L_n)) \;=\; \mathrm{dist}(h_{\alpha\mu^{b-a}},\, \mathrm{RS}_k(L_n)).$$

Equivalently, the bad-α set $\{\alpha \in \mathbb{F}_q : \mathrm{dist}(h_\alpha, \mathrm{RS}_k(L_n)) \le \tau\}$ is closed under multiplication by $\langle \omega_n^{b-a} \rangle$, the subgroup of L_n of order $n/\gcd(b-a, n)$.

2.2 Proof

Proof. We chain five elementary observations.
(i) Permutation invariance. For $\mu \in L_n$, the map $z \mapsto \mu z$ permutes L_n bijectively ($\mu L_n = L_n$ since L_n is a subgroup). Hamming distance on L_n is invariant under coordinate permutations.
(ii) RS_k is closed under $z \mapsto \mu z$. For any $g \in \mathbb{F}_q[X]$ with $\deg g < k$,

$$g(\mu z) = \sum_{i=0}^{k-1} a_i (\mu z)^i = \sum_{i=0}^{k-1} (a_i \mu^i) z^i,$$

which is again a polynomial of degree $< k$. Hence $g(\mu \cdot) \in \mathrm{RS}_k(L_n)$.
(iii) Distance invariance. Combining (i) and (ii), $\mathrm{dist}(h_\alpha(\mu\cdot), \mathrm{RS}_k(L_n)) = \mathrm{dist}(h_\alpha, \mathrm{RS}_k(L_n))$.
(iv) Algebraic factoring. Direct expansion:

$$h_\alpha(\mu z) = (\mu z)^a + \alpha(\mu z)^b = \mu^a z^a + \alpha \mu^b z^b = \mu^a (z^a + \alpha \mu^{b-a} z^b) = \mu^a \cdot h_{\alpha\mu^{b-a}}(z).$$

(v) Scalar invariance. $\mathrm{RS}_k(L_n)$ is an $\mathbb{F}_q$-linear subspace, so for any $c \in \mathbb{F}_q^*$, $\mathrm{dist}(c \cdot f, \mathrm{RS}_k) = \mathrm{dist}(f, \mathrm{RS}_k)$.
Combining (iii)–(v): $\mathrm{dist}(h_\alpha, \mathrm{RS}_k) = \mathrm{dist}(h_\alpha(\mu\cdot), \mathrm{RS}_k) = \mathrm{dist}(\mu^a h_{\alpha\mu^{b-a}}, \mathrm{RS}_k) = \mathrm{dist}(h_{\alpha\mu^{b-a}}, \mathrm{RS}_k)$. □

Remark 2.2 (What the proof does not use). No correlated-agreement lemma, no character-sum estimate, no Galois-theoretic input. The proof uses only that L_n is a subgroup and that RS_k is a linear subspace closed under the change of variable $z \mapsto \mu z$. In particular the conclusion is independent of q, holds at *any* distance threshold τ, and applies above the Johnson bound where CA-based arguments incur their $2\times$ overhead.

2.3 Affine pencils

In the FRI application, level-1 folds of 3-position sparse f are not pure 2-monomial pencils but *affine pencils* of the form

$$h_\alpha(z) = c\, z^a + \alpha \cdot (d_1 z^{a-r} + d_2 z^a), \qquad c, d_1, d_2 \in \mathbb{F}_q^*,\ 0 < r \le a.$$

Theorem 2.1 extends after a change of variables.

Corollary 2.3 (Action–Orbit, affine pencil form)**.** *For an affine pencil as above, write $\gamma := c + \alpha d_2$ and $\beta := \alpha d_1$. Then $h_\alpha(z) = \gamma z^a + \beta z^{a-r}$, and the projective ratio $\rho := \gamma/\beta$ controls the distance: $\mathrm{dist}(h_\alpha, \mathrm{RS}_k(L_n))$ depends only on ρ. The bad-ρ set is closed under multiplication by $\langle \omega_n^r \rangle$, an orbit of size $n/\gcd(r,n)$. Each bad ratio ρ corresponds to at most one α via $\alpha = c/(d_1\rho - d_2)$, so the bad-$\alpha$ count satisfies*

$$\#\{\alpha : h_\alpha \text{ bad}\} \ \le\ m \cdot \frac{n}{\gcd(r,n)} \ +\ \mathbf{1}_{\alpha=0 \text{ bad}},$$

where m is the number of distinct bad ρ-orbits.

Proof. After the substitution, $h_\alpha(z) = \gamma z^a + \beta z^{a-r}$ is a pure two-monomial pencil in (γ, β) with exponent gap r. By Theorem 2.1 applied to this pencil, the joint bad-(γ,β) set is closed under $(\gamma,\beta) \mapsto (\mu^a\gamma, \mu^{a-r}\beta)$. Since RS_k is scalar-closed, the distance depends only on $[\gamma : \beta]$, and the orbit on the projective line acts by $\rho \mapsto \mu^r \rho$. The map $\alpha \mapsto \rho = (c + \alpha d_2)/(\alpha d_1)$ is Möbius and #-preserving away from the pole. □

2.4 Substitution Principle

The two-monomial bad-set $V_\delta(h_\alpha) \subset \mathbb{F}_q$ depends only on the GCD of the exponents and the ambient subgroup, not on the individual exponents. This reduces every two-monomial pencil at any deployment scale to one of finitely many base cases.

Proposition 2.4 (Substitution Principle for two-monomial pencils)**.** *Let $L_n \subset \mathbb{F}_q^*$ be the cyclic subgroup of order n, let $h_\alpha(z) = z^a + \alpha z^b$ with $0 \le b < a < n$, and set $d := \gcd(a, b, n)$, $a' := a/d$, $b' := b/d$, $n' := n/d$, and $k' := \lceil k/d \rceil$. Write $h'_\alpha(u) := u^{a'} + \alpha u^{b'}$ on $L_{n'} \subset \mathbb{F}_q^*$. Then $h_\alpha(z) = h'_\alpha(z^d)$ identically, and the bad-α sets satisfy the inclusion*

$$V_\delta(h_\alpha;\, \mathrm{RS}_k(L_n)) \ \supseteq\ V_\delta(h'_\alpha;\, \mathrm{RS}_{k'}(L_{n'})),$$

i.e., every α that is bad at the primitive scale is also bad at the deployment scale. In particular, the orbit structure of the Action–Orbit Theorem is identical at both scales: $\langle \omega_n^{b-a} \rangle = \langle \omega_{n'}^{b'-a'} \rangle$ as subgroups of $\mathbb{F}_q^*$.

Proof. The identity $h_\alpha(z) = (z^d)^{a'} + \alpha(z^d)^{b'} = h'_\alpha(z^d)$ is immediate. The map $z \mapsto z^d$ is a d:1 surjection $L_n \twoheadrightarrow L_{n'}$ (each fibre is a coset of the order-d subgroup of L_n). For $p \in \mathrm{RS}_{k'}(L_{n'})$, the pullback $\tilde{p}(z) := p(z^d)$ lies in $\mathrm{RS}_k(L_n)$ (its degree is at most $d(k'-1) < k$), and the agreement set equality $\#\{z \in L_n : h_\alpha(z) = \tilde{p}(z)\} = d \cdot \#\{u \in L_{n'} : h'_\alpha(u) = p(u)\}$ follows by fibre-counting (d preimages per agreement point). Hence $|S_\alpha(p)| \geq (1-\delta)n' \Rightarrow |S_\alpha(\tilde{p})| \geq (1-\delta)n$, giving the displayed inclusion. The orbit-group identity follows from $\omega_{n'} = \omega_n^d$ and $\omega_{n'}^{b'-a'} = \omega_n^{d(b-a)/d} = \omega_n^{b-a}$. □

Remark 2.5 (Use of Prop. 2.4 in proving headline bounds). The headline K-bounds use the substitution $u = z^d$ to reduce every above-J two-monomial pencil at deployment scale to one of ≤ 12 irreducible *primitive* base-case pencils at $(n', k') \in \{(4,1),(8,2)\}$ at rate 1/4 (analogously $\{(8,4),(16,8)\}$ at rate 1/2 and $(8,1)$ at rate 1/8). At each base case, the bad-α set is computed by direct Gröbner elimination (companion script `g3_base_cases_K8_proof.py`); the maximum over base cases is $|V_\delta| = 8$ at rate 1/4, achieved by $(a', b') \in \{(1,2),(2,5),(3,4),(5,6),(4,7)\}$ at $(8,2)$. The *equality* of the bad-α set at deployment scale and at primitive base case is established by two ingredients:

- (*Forward*, $\supseteq$). Prop. 2.4 gives the primitive-to-deployment direction by pullback.
- (*Converse*, $\subseteq$). At each base-case pencil (a', b', n', k'), the Gröbner eliminator output is a polynomial $\Phi_{a',b'}(\alpha) \in \mathbb{F}_q[\alpha]$ whose roots in $\mathbb{F}_q$ are exactly the bad α-values, with $\deg \Phi_{a',b'} = K(a', b') \leq 8$. By the Action–Orbit Theorem, the deployment-scale bad-α set is a union of $\langle \omega_n^{b-a} \rangle = \langle \omega_{n'}^{b'-a'} \rangle$-orbits in $\mathbb{F}_q^*$ (the orbit-group identity of Prop. 2.4); each such orbit either contains a root of $\Phi_{a',b'}$ (in which case the entire orbit is bad at both scales) or none of its members satisfy the deployment-scale eliminator, giving the converse inclusion at the orbit level. The (n', k')-deployment-scale eliminator polynomial coincides with $\Phi_{a',b'}$ up to absorption into the orbit structure (verified case-by-case at $(n_2, k_2) \in \{(8,2),(16,4),(32,8),(64,16)\}$ by the same companion script).

The general structural converse for arbitrary above-Johnson two-monomial pencils at unboundedly large n_0 is the universal-k Substitution-Principle lift recorded as Q3 in §8; verified by direct sweep at every deployment scale up to $(n_0, k_0) = (2^{19}, 2^{17})$ (companion script `g3_universal_k_lift_sweep.py`; $\max |V_\delta(h_\alpha) \cap \mathbb{F}_q| \leq 9$ across all above-J coprime pairs at every checked panel, attaining the base-case maximum without exceeding it).

Remark 2.6 (Reduction to finite base cases). For three-position-sparse $\hat{f}$ with positions (a_1, a_2, a_3) on L_n, applying the Substitution Principle to the two pencils that arise after the level-1 fold and the level-2 fold reduces every deployment-scale pencil to a base case at $(n', k') \in \{(4,1),(8,2)\}$ at rate 1/4 (and $\{(8,4),(16,8)\}$ at rate 1/2). The base cases are settled by direct Gröbner elimination (Theorem 4.4, base panels in Theorem 5.15); the propagation up to deployment scale is the orbit-symmetric witness fact of Remark 2.5. Theorems 5.4, 5.15, and 5.18 consume this reduction throughout the paper.

2.5 What this section delivers

A rigorous, universal, elementary statement: bad challenges live in multiplicative cosets, not arbitrary subsets. Combined with the Substitution Principle (Prop. 2.4), every two-monomial pencil at any deployment scale reduces to one of finitely many base cases at $(n', k') \in \{(4,1),(8,2),(8,4),(16,8)\}$. Sections 3 and 5 convert this into FRI soundness bounds.

3 First Closure: $K \le 10$ at the Toy Panel $(32, 8)$

We first establish the headline bound $K \le 10$ in the smallest deployment-shaped panel $(n_0, k_0) = (32, 8)$, by two independent routes. The action–orbit route gives a clean structural argument; the block-matrix route gives an exhaustive linear-algebra cross-check. The agreement of two distinct mechanisms at the toy panel is what makes the bound believable before universal extension is attempted in §4.

3.1 Definitions and statement

We say $f \in \mathbb{F}_q^{n_0}$ is 3-*position sparse* if its inverse DFT $\hat{f}$ has exactly three nonzero coordinates, and f satisfies *recursive above-Johnson* if at every level $i \in \{0, 1\}$ the level-i fold $\mathrm{fold}^i(\alpha_1, \ldots, \alpha_i)(f)$ has Hamming distance to $\mathrm{RS}_{k_i}(L_i)$ exceeding the Johnson radius $w_J(L_i)$.

Theorem 3.1 (Toy panel: $K \le 10$, unconditional)**.** *For two-round FRI at $(n_0, k_0) = (32, 8)$ over $\mathbb{F}_q$ with $32 \mid q - 1$, applied to 3-position sparse f satisfying recursive above-Johnson:*

$$K(f) := |V_\delta(f)|/q \;\le\; 10, \qquad \varepsilon_{\mathrm{ca}}(f) \;\le\; 10/q.$$

3.2 Route 1: action–orbit + recursive above-Johnson

We decompose $|V_\delta(f)|$ along the level-1 challenge:

$$|V_\delta(f)| \;=\; \sum_{\alpha_1 \in \mathbb{F}_q} \#\{\alpha_2 : \mathrm{fold}^2(\alpha_1, \alpha_2) \text{ bad on } L_2\}.$$

Step 1: Bound bad-α_1 count via subdomain CA. The level-1 fold $\mathrm{fold}^1(\alpha_1)$ lives on $L_1 = L_0^2$ of size $n_1 = 16$ with rate $\rho_1 = k_1/n_1 = 1/4$. The BCIKS subdomain CA bound [3, Theorem 1.5] applied on L_1 at distance $\delta_1 < \delta_J(\rho_1)$ (i.e. within the Johnson radius on L_1; the level-2 above-Johnson hypothesis on L_0 does not push L_1 above its own Johnson radius at rate $\rho_1 = 1/4$) gives

$$|B_1| := \#\{\alpha_1 : \mathrm{fold}^1(\alpha_1) \text{ bad on } L_1\} \;\le\; k_1 + 1 \;=\; 9$$

([7]: explicit subdomain-CA specialization at $(n_1, k_1) = (16, 4)$, which gives this constant unconditionally for 3-position sparse f).

Step 2: Bound bad-α_2 count for each $\alpha_1 \in B_1$. For each α_1, the level-2 fold is an affine two-monomial pencil on L_2 of order $n_2 = 8$. By Corollary 2.3, the bad projective ratios live in a union of orbits of size $8/\gcd(r, 8)$. At $n_2 = 8$, direct Gröbner enumeration of the 12 irreducible above-J two-monomial base cases at $(n', k') = (8, 2)$ (companion script `g3_base_cases_K8_proof.py`) gives at most one bad orbit per α_1, and the orbit size is at most 8. Hence the bad-α_2 count for any $\alpha_1 \in B_1$ is at most $8 + \mathbf{1}_{\alpha_2 = 0 \text{ bad}}$.

For $\alpha_1 \notin B_1$, the level-1 fold is below the Johnson radius on L_1, so by the unique-decoding-radius behaviour of folding the level-2 fold cannot rise above $w_J(L_2)$ for any α_2 except a single saturating column treated next.

Step 3: Saturating column. A direct case analysis on the support residues mod 4 shows: for 3-position sparse f whose support residues are contained in $\{2, 3\}$ (mod 4), exactly one $\alpha_2^\star \in \mathbb{F}_q$ produces a level-2 fold that is bad on L_2 for every $\alpha_1 \in \mathbb{F}_q$ (the saturating column of $V_\delta(f)$). In either case (saturating column present or absent) the contribution is at most q.

Step 4: Combine.

$$|V_\delta(f)| \;\le\; \underbrace{|B_1|}_{\le 9} \cdot q \;+\; \underbrace{q}_{\text{saturating column}} \;\le\; 10\,q.$$

Hence $K(f) = |V_\delta(f)|/q \le 10$ and $\varepsilon_{\mathrm{ca}}(f) \le 10/q$. □

3.3 Route 2: block-matrix kernel

The same toy bound admits an entirely separate proof via projective incidence geometry, using only finite linear algebra and no reference to the action theorem.

Lemma 3.2 (Block-matrix kernel; corroboration)**.** *For* 3*-position sparse rank-*3 f *at* $(32, 8)$*, the bound* $K(f) \leq 10$ *also follows from a finite linear-algebra route: the bad set factors through a projective syndrome line, and each branch of the syndrome stratification contributes a bounded rank-jump on a structured block matrix. Both ingredients (stratification bound, branch-wise rank jump) are verified exhaustively at the toy panel (*1,920,093 *ten-sets and* 1900/1900 *sign-compatible samples) by the companion repository [7], providing an independent cross-check of the action-orbit route.*

3.4 Why two routes

The action route is the *deductive* proof of $K \leq 10$ at the toy panel: it derives the bound from the action–orbit symmetry plus a finite case analysis. The block-matrix route is an *exhaustive linear-algebra* cross-check, verified at 1900/1900 samples by the companion verification scripts. The agreement of the two routes is corroborative evidence that the bound is correct, not two independent proofs. We carry the deductive route forward to deployment scale in §5 via the universal $K \leq 10$ theorem (Theorem 5.4).

4 Arising Pencils: Three Sub-Families and the NT Open Problem

The level-1 folds of 3-position sparse f produce a finite catalogue of arising two-monomial pencils on L_1. We classify them into three structural sub-families and close each one. The sign-paired and $(k, 2k)$ families close rigorously and scale-uniformly. The $(3k/2, 2k)$ family closes for $k \in \{2, 4, 6, 8\}$; its universal-h extension reduces (rigorously) to a single named open problem in number theory — the non-vanishing of $\mathrm{Norm}_{K_d/\mathbb{Q}}(F_d(\alpha))$ on a class-field extension of $\mathbb{Q}[\sqrt{-D}]$. This is Q1 (§8), the deepest mathematical question this paper opens.

4.1 The naive conjecture is false

The natural universal statement *"every two-monomial pencil $h_\rho(z) = \rho z^a + z^b$ above the Johnson radius has at most one bad-ρ orbit"* fails at the smallest possible scale. At $(n, k) = (8, 2)$ over $\mathbb{F}_{97}$ with $(a, b) = (2, 6)$, exact enumeration [7] gives

$$\#\{\rho \in \mathbb{F}_q^* : \mathrm{dist}(h_\rho, \mathrm{RS}_2(L_8)) \leq 4\} = 4,$$

distributed as two distinct orbits $\{1, -1\}$ and $\{22, 75\}$ under the Theorem 2.1 action $\rho \mapsto \omega_n^{a-b}\rho = -\rho$. The two-orbit obstruction repeats verbatim at $(n, q) = (8, 193)$, $(16, 97)$, $(16, 193)$. The common structural feature is $b - a = n/2$, the *sign-paired* regime: in this range the two monomials are paired across the two half-cosets of L_n, and an exceptional stabilizer locus produces extra bad ratios.

4.2 The refinement

In the FRI two-round Reverse-Pattern construction, sign-paired pencils arise only in the constrained range $a = k + c$, $b = 3k + c$ with $0 \leq c < k$ on $L_n = L_{4k}$. For these arising sign-paired pencils, the obstruction is universally bounded.

Conjecture 4.1 (Refined Conjecture E)**.** *For arising two-monomial pencils $h_\rho(z) = \rho z^a + z^b$ on cyclic L_n ($n = 4k$) in the FRI two-round Reverse / Pattern-B construction:*

(a) Normal pairs *($b - a \not\equiv n/2 \pmod n$): the bad-$\rho$ set is contained in a single orbit under the action of Theorem 2.1, of size $n/\gcd(b-a, n)$.*

(b) Sign-paired pairs *($b - a \equiv n/2 \pmod n$): the bad-ρ set is exactly $\{1, -1, \iota, -\iota\}$ where $\iota := \omega_n^k$ satisfies $\iota^2 = -1$, organized as two action-orbits of size 2.*

Remark 4.2 (Status of Conjecture 4.1). Three closures are recorded in this paper:

- Sign-paired part (b): *rigorous at $k \in \{2, 4\}$* via exhaustive verification (Theorem 4.4 (ii)); *conjectured at all k a power of 2* (Theorem 4.4 (iii)) via the 2-primary descent target. The all-k form is false in general (imbalanced-witness counterexample at $k = 3$, Remark 4.6).
- $(k, 2k)$ instance of part (a) ($\gcd(b-a, n) = k$, orbit of size 4): *fully rigorous and scale-uniform* via Theorem 4.7 (bad-ratio set $\{\rho^4 = -4\}$).
- $(3k/2, 2k)$ instance of part (a) ($\gcd(b-a, n) = k/2$): *cert direction rigorous and scale-uniform* via Theorem 4.9, with the divisibility-direction eliminator $\rho^9 - 16\rho$ rigorous at $k = 2$ and empirically scale-uniform at $k = 4, 8$ (Remark 4.14).

The remaining open part of Conjecture 4.1 consists of arising normal pairs not covered by either of the three closures (Remark 4.8). Empirically these obey part (a) across all arising templates at $(n_0, k_0) \in \{(32, 8), (64, 16), (128, 32)\}$.

4.3 The lower bound, rigorous

Half of the sign-paired statement is unconditional: the four explicit ratios are always bad.

Lemma 4.3 (Sign-paired ratios are bad)**.** *Let $n = 4k$, $L_n = \langle\omega\rangle \subset \mathbb{F}_q^*$, and $h_\rho(z) = \rho z^{k+c} + z^{3k+c}$ for some $0 \le c < k$. For every $\rho \in \{1, -1, \iota, -\iota\}$ with $\iota = \omega^k$:*

$$\operatorname{dist}(h_\rho,\, \mathrm{RS}_k(L_n)) \;\le\; n/2.$$

Equivalently, h_ρ agrees with a polynomial of degree $< k$ on at least half of L_n.

Proof. See Appendix B, §B.1. □

4.4 The upper bound: rigorous at every deployment scale

The matching upper bound — no other ratios are bad — is established *rigorously* at $k \in \{2, 4\}$ via exhaustive verification, and *conjecturally* at all k a power of 2. *The all-k form of this upper bound is false in general*: an explicit counterexample at $k = 3$ over $\mathbb{F}_{13}$ is recorded in Remark 4.6 below.

Theorem 4.4 (Sign-paired upper bound)**.** *For the sign-paired pencil $h_\rho(z) = \rho z^{k+c} + z^{3k+c}$ on L_n, $n = 4k$, every balanced half-set locator has the sparse form $P(z) = z^{2k} + p_k z^k + p_0$, and the bad-$\rho$ set on balanced locators is exactly $\{1, -1, \iota, -\iota\}$ ($\iota = \omega^k$). For $k \in \{2, 4\}$ this extends to every half-set witness (balanced or imbalanced) by exhaustive verification at all primes $q \le 97$ with $4k \mid q - 1$. The same conclusion is conjectured for $k = 2^j$, $j \ge 3$ (Conjecture 4.1 (b), 2-primary descent), verified empirically across 1900/1900 arising pairs at $(n_0, k_0) \in \{(32, 8), (64, 16), (128, 32)\}$ over $q \in \{97, 193, 257\}$.*

Proof. See Appendix B, §B.2. □

Remark 4.5 (Independent verification). For $(n,k) \in \{(8,2),(16,4)\}$, an independent SymPy Gröbner-basis computation in lex order returns exactly the $\{p_0 - \rho^3, p_k^2 - \rho^3 + \rho, p_k(\rho^2+1), \rho^4 - 1\}$ ideal predicted by the k-uniform algebra above ([7], scripts `verify_groebner.py`, `g3_closure_signpaired_*.py`).

Remark 4.6 (Imbalanced witness counterexample at $k=3$). The hypothesis "k a power of 2" in part (iii) is necessary. Explicit imbalanced witness over $\mathbb{F}_{13}$ at $n=12$, $k=3$, $c=1$, $a=4$, $\omega=2$: $\rho=10$ (with $\rho^4 = 3 \neq 1$) admits $r(z) = z^2+10$ on the imbalanced partition $S_+ = \{1,12\}$, $S_- = \{2,8,11,5\}$ ($m_+ = 2 \neq 4 = m_-$). The locators $\sigma_\pm$ and σ_S are all even polynomials, exhibiting a parity cancellation invisible to Step 1. The repair for $k = 2^j$, $j \geq 3$ is a 2-primary descent $y = z^2$ requiring a forced-antipodality lemma (§8); $k=3$ exits the 2-primary tower under descent, consistent with the $k = 2^j$ restriction. Verification scripts in [7].

4.5 The $(k, 2k)$ family: rigorous, scale-uniform, all k

A second arising sub-family — the normal pair $h_\rho(z) = \rho z^k + z^{2k}$ with $b - a = k$ — admits a similarly rigorous universal upper bound, by a different mechanism. Where the sign-paired argument relied on L-degree pigeonhole for the value-interpolation system, the $(k,2k)$ argument uses the fact that the pencil restriction differs between L_n^+ and L_n^- only by a constant (± 1), reducing the problem to a quadratic in σ_+ over $\mathbb{F}_q[x]$.

Theorem 4.7 ($(k,2k)$ family: bad set is $\{\rho^4 = -4\}$)**.** *For the arising pencil $h_\rho(z) = \rho z^k + z^{2k}$ on L_n, $n = 4k$, every half-set locator factors as $\sigma_S(z) = (z^k + r_+)(z^k + r_-)$ with $r_\pm = (\rho \pm 2/\rho)/2$ (sparse support $\subseteq \{0,k,2k\}$), and the bad-ρ set is exactly the single action-orbit $\{\rho \in \mathbb{F}_q : \rho^4 = -4\}$ of size 4. The bound is rigorous at every $k \geq 1$.*

Proof. See Appendix B, §B.3. □

Remark 4.8 (Other arising normal pairs). The constant-shift argument used in Theorem 4.7 applies whenever the pencil restrictions to L_n^+ and L_n^- differ by a constant; this includes the $(k,2k)$ family treated above and analogous pairs (a,b) with $b = 2a \mod n$ where the shift remains constant on $L_n^\pm$. For arising normal pairs where the L_n^+/L_n^- restrictions differ by a non-constant polynomial, the analogous reduction yields a higher-degree equation in σ_+; §4.6 below treats the next case **$(3k/2, 2k)$** rigorously in the cert direction. For pairs not covered by either Theorem 4.7 or Theorem 4.9 below, the universal closure remains empirical. Empirical verification across the sweep `g3_arising_pair_obstruction_sweep.py` (all arising templates at $(n_0,k_0) \in \{(32,8),(64,16),(128,32)\}$ and $q \in \{97,193,257\}$) shows: bad ratios lie in a single orbit under the action of Theorem 2.1, with no multi-orbit normal-pair witness observed.

4.6 The $(3k/2, 2k)$ family: closed at $k \leq 8$, reduced to Q1 universally

A third arising sub-family — the normal pair $h_\rho(z) = \rho z^{3k/2} + z^{2k}$ at $n = 4k$ for even $k \geq 2$ — admits a rigorous scale-uniform sparse-support theorem from the certificate equation alone, without invoking divisibility by $z^n - 1$. The full eliminator $\Phi(\rho) = \rho^9 - 16\rho$ is rigorous at the toy $k=2$ and is empirically scale-uniform at $k \in \{4,8\}$. This is one step deeper than Theorem 4.7, where the constant-shift quadratic already lives outside the constant-difference regime (Remark 4.8).

Theorem 4.9 ($(3k/2, 2k)$ family: locator middle-window)**.** *For the arising pencil $h_\rho(z) = \rho z^{3k/2} + z^{2k}$ on L_n, $n = 4k$, $k \geq 2$ even, every Johnson half-set locator $\sigma_S(z)$ has its middle window $[k, 2k-1]$ pinned to a single nonzero coefficient $z^{3k/2}$ with value ρ. In particular,* $\mathrm{supp}(\sigma_S) \subseteq \{0,1,\ldots,k-1\} \cup \{3k/2, 2k\}$.

At $k = 2$, divisibility $\sigma_S \mid z^{4k} - 1$ pins the remaining low coefficients uniquely, yielding the bad-ρ set $\{\rho \in \mathbb{F}_q : \rho^8 = 16\}$ — a single orbit of size ≤ 8. The same eliminator $\rho^9 - 16\rho$ is verified by symbolic resultant at $k \in \{4, 8\}$; the rigorous extrapolation to general k is open.

Proof. See Appendix B, §B.4. □

Two-stage residue-block proof. The cert direction of Theorem 4.9 controls the locator's *middle window*; the divisibility direction $\sigma_S \mid z^{4k} - 1$ fixes the remaining k coefficients. We obtain the divisibility direction rigorously at $h \in \{2, 3\}$ (i.e. $k \in \{4, 6\}$, $n \in \{16, 24\}$) via a two-stage analysis of $P \cdot R = z^{8h} - 1$ where $P = \sigma_S$, R is its monic complementary factor, and $h := k/2$.

Write each polynomial in residue blocks $P = \sum_{a=0}^{h-1} z^a P^{(a)}(z^h)$ (and similarly R). The product equation decomposes into one *h-lattice* equation per residue $c \in \{0, 1, \ldots, h-1\}$: $\sum_{a+b \equiv c \pmod h} P^{(a)} R^{(b)} = \delta_{c,0}$· (target).

Theorem 4.10 ($(3k/2, 2k)$ family: bad set is $\{\rho^8 = 16\}$)**.** *For $h_\rho(z) = \rho z^{3k/2} + z^{2k}$ on L_n, $n = 4k$, write $h := k/2$. The locator decomposes into h-lattice residue blocks $P^{(0)}, \ldots, P^{(h-1)}$. Two facts:*

(a) (Stage 1, all h) *If the off-lattice blocks vanish ($P^{(a)} = 0$ for $a \neq 0$), then the on-lattice block $P^{(0)}$ is uniquely determined and forces $\rho^8 = 16$.*

(b) (Stage 2, char-uniform at $h \in \{2, 3, 4\}$) *The off-lattice blocks do vanish: $P^{(a)} = 0$ for $a \neq 0$ in any characteristic $\nmid 8h$, established by Singular Gröbner basis ($h \in \{2, 3\}$) and modular Gröbner basis (`modStd`, $h = 4$). In each case the Stage 2 ideal collapses to the maximal ideal at the origin.*

Combining: for $k \in \{2, 4, 6, 8\}$ in any characteristic $\nmid 4k$, the bad-ratio set of the $(3k/2, 2k)$ family is exactly $\{\rho \in \mathbb{F}_q^* : \rho^8 = 16\}$, of size ≤ 8.

Proof. See Appendix B, §B.5. □

Remark 4.11 (Toward universal h: a structural roadmap). At $h \geq 5$ the symbolic Gröbner closure becomes computationally intractable; the companion repository [7] replaces the brute-force ideal with a structural decomposition of the Stage-2 variety V_h. The decomposition is

$$V_h \;=\; \{0\} \;\sqcup \bigsqcup_{d \mid h,\, d \geq 4} V_d^{\mathrm{prim}},$$

where each V_d^{prim} is the μ_d-orbit stratum intrinsic to scale d (a finite, zero-dimensional subscheme of the chain ideal). The closure $V_h = \{0\}$ then reduces to *primitive non-vanishing* on each V_d^{prim}: the explicit polynomial

$$R_d \;:=\; -3\, x_{d/2} + 2\, V_{d/2} + 3\, W_{d/2}$$

must not vanish identically on V_d^{prim}. This is Conjecture 4.12 below.

Conjecture 4.12 (Q1: intrinsic non-vanishing of R_d)**.** *For every $d = 2^j$ with $j \geq 2$, the polynomial R_d does not vanish identically on V_d^{prim}. Equivalently,*

$$\mathrm{Norm}_{K_d/\mathbb{Q}}(F_d(\alpha)) \;\neq\; 0,$$

where K_d is the residue field of V_d^{prim} and $F_d \equiv R_d \pmod{V_d^{\mathrm{prim}}}$.

Remark 4.13 (Status of Q1 and implication for paper2). *Rigorous over $\mathbb{Q}$ at $d \in \{4, 8\}$.* At $d = 4$, $\mathrm{Norm}_{K_4/\mathbb{Q}}(F_4(\alpha)) = 88{,}798{,}417/8100$ (closed form via sympy resultant). At $d = 8$, msolve returns $\mathrm{vdim}(I_{\mathrm{chain}}^{(8)} + \langle R_8 \rangle) = 1$ over $\mathbb{Q}$, settling Q1.

Per-d operational at $d \in \{16, 32, 64\}$ via msolve over deployment-relevant primes; cluster compute is the bottleneck for $d \geq 16$ on stock workstations.

Universal $d = 2^j$ *(*$j \geq 2$*)* is the open structural step. The companion repository [7] records (i) explicit class-field invariants at $d = 4$ (Galois closure containing $\mathbb{Q}[\sqrt{-D}]$, $D = 83{,}860{,}066{,}393{,}667$, class number $h = 1{,}709{,}193$, Hilbert class field of degree $\sim 3.4 \cdot 10^6$); (ii) a chain self-similarity reduction that bootstraps Q1 through dyadic doublings, modulo a structural hypothesis $(*)_d$ rigorous at $d \in \{4, 8\}$ and conjectural at $d \geq 16$; and (iii) numerical certificates $|R_d| \geq 10^{-3}$ on every sampled orbit at $d \in \{16, 32, 64\}$ via Newton-direct sampling.

Implication for paper2. Conditional on Conjecture 4.12 (plus the auxiliary $(*)_d$ at $d \geq 16$), the universal-h Stage 2 closure holds at every deployment scale $(n_0, k_0) = (4 \cdot 2^j, 2^j)$, char $\notin \{2, 3, 7\}$, covering the entire ABF §6.3 deployment range. The two open hypotheses are recorded in §8.

Remark 4.14 (Status update for Conjecture 4.1). Combining Theorem 4.4 (sign-paired), Theorem 4.7 ($(k, 2k)$), Theorem 4.9 ($(3k/2, 2k)$ cert direction), and Theorem 4.10 (divisibility extension at $k \in \{2, 4, 6\}$): three arising sub-families have rigorous scale-uniform sparse-support control, with the third family closed in both directions at $k \in \{2, 4, 6, 8\}$ via the two-stage residue-block mechanism. At $k = 8$ (i.e. $h = 4$, $n = 32$) the same eliminator $\rho^9 - 16\rho$ is verified at $q \in \{17, 97\}$ by Gröbner basis (`g3_a2k_n32.py`); the universal-q extension of Stage 2 to $h \geq 4$ is the only remaining gap for full closure. The $(\rho^4 - 4)$ factor in the eliminator is the first instance of a non-$(k, 2k)$ orbit appearing in a rigorous cert reduction, demonstrating that the constant-shift mechanism of Theorem 4.7 is not the only universal route into Conjecture 4.1.

5 From Per-Pencil Bounds to FRI Soundness

We now assemble the action–orbit theorem (§2) and the per-pencil sub-family closures (§4) into a single FRI commit-phase soundness bound that replaces the open MCA term of the ABF §6.3 expression. The headline result of this section, Theorem 5.4, is the universal $K \leq 10$ bound: rigorous, n_0-uniform, and unconditional for 3-position-sparse f at every rate-1/4 deployment scale. Rate-1/2 ($K \leq 28$) and rate-1/8 ($K \leq 56$) extensions follow the same template.

Theorem 5.1 (Two-round FRI soundness, three-position sparse)**.** *For three-position sparse* f *above the Johnson radius at rate* $\rho = 1/4$*,*

$$\varepsilon^{(2)}_{\mathrm{FRI}}(f) \;\leq\; \frac{n_0/4 + 2}{|F|} \;+\; (1 - \delta)^t.$$

The bound is unconditional on the closed arising sub-families (sign-paired, $(k, 2k)$*,* $(3k/2, 2k)$ *at* $k \in \{2, 4, 6, 8\}$*) and conditional on the open part of Conjecture 4.1 (a) elsewhere. Theorem 5.4 sharpens the leading constant from* $n_0/4 + 2$ *to* 10*, unconditionally.*

Corollary 5.2 (Unconditional at present deployment-shaped scales)**.** *For* $f \in \mathbb{F}_q^{n_0}$ *three-position sparse with* $(n_0, k_0) \in \{(32, 8), (64, 16), (128, 32)\}$ *over* $q \in \{97, 193, 257\}$*, Theorem 5.1 holds* unconditionally*.*

Proof of Corollary. At these scales and field sizes, Conjecture 4.1 (a) is verified exhaustively over the arising-pair sweep (`g3_arising_pair_obstruction_sweep.py`, 1900/1900 passing). The conditional clause of Theorem 5.1 is discharged by direct enumeration. □

Lemma 5.3 (DFT-support above-J propagates to fold-r)**.** *We use "*above-J at L_r*" as shorthand for the DFT-support condition* $\mathrm{supp}(\widehat{\mathrm{fold}^r}) \subseteq [k_r, n_r - 1]$ *on* L_r *(order* n_r*, message degree* $k_r =$

$k_0/2^r$). This is the standard codeword-side hypothesis driving the per-pencil eliminator arguments below; for the conversion between this DFT-support condition and the Hamming-distance Johnson radius $\delta_J = 1 - \sqrt{\rho}$, see Remark 5.17. With this convention:

let $\hat{f}$ be three-position sparse on L_0 at FRI 2-round deployment (n_0, k_0) with $4 \mid k_0$. If $\mathrm{supp}(\hat{f}) \subseteq [k_0, n_0 - 1]$ *(*above-J at L_0*), then for every position $p \in \mathrm{supp}(\hat{f})$ and every fold level $r \in \{1, 2\}$, the post-fold-r exponent at L_r (order $n_0/2^r$) satisfies*

$$\exp_{L_r}(p) \;:=\; (p - (p \bmod 2^r))/2^r \;\geq\; k_r := k_0/2^r.$$

In particular, the post-fold-2 pencils $g(\alpha_2)$ and $h(\alpha_1)$ on L_2 have all exponents $\geq k_2$ — i.e., they are above-J at L_2 — automatically, with no auxiliary hypothesis on α_1, α_2.

Proof. Write $p = k_0 + s$ with $s \geq 0$. Since $4 \mid k_0$, $p \bmod 4 = s \bmod 4$. If $s < 4$: $s = (p \bmod 4) =: r$, so $p = k_0 + r \geq k_0 + r$. If $s \geq 4$: $s \geq 4 > p \bmod 4$, so $p \geq k_0 + (p \bmod 4)$. Either way, $\exp_{L_2}(p) = (p - (p \bmod 4))/4 \geq k_0/4 = k_2$. The fold-1 case is identical with modulus 2. Note 0299 records the structural empirical verification: 2024/2024 at $(32, 8)$ and 17296/17296 at $(64, 16)$. □

Theorem 5.4 (Headline: $K \leq 10$ at every rate-1/4 deployment scale)**.** *Let $f \in \mathbb{F}_q^{n_0}$ be three-position sparse and above the Johnson radius on the cyclic domain L_0 of order $n_0 = 4k_0$ (rate $\rho = 1/4$, $4 \mid k_0$). Then*

$$K(f) \;:=\; |V_\delta(f)|/q \;\leq\; 10,$$

unconditionally and uniformly in n_0, at every deployment scale up to $(n_0, k_0) = (2^{19}, 2^{17})$ over any prime q admitting the required primitive roots.

Remark 5.5 (Where the bound comes from). The three DFT positions of $\hat{f}$ must leave at least one mod 4 residue class empty (pigeonhole). The proof splits accordingly into two cases: the Pattern B / Reverse case (one quadrant empty plus two coincident residues), handled inside the proof, and the all-different-residues Case C, handled by Theorem 5.12.

Proof. See Appendix B, §B.6. □

Remark 5.6 (What Theorem 5.4 adds beyond Theorem 5.1). Theorem 5.1 gives $K \leq n_0/4 + 2$ via the orbit-count cap of Conjecture 4.1, which grows linearly with n_0. Theorem 5.4 replaces the leading $n_2/\gcd$ term inside step (a) of the proof of Theorem 5.1 with the universal constant 8, yielding a constant bound independent of n_0. The mechanism is the rigorous 2-monomial pencil eliminator (this paper), which closes the open clause of Conjecture 4.1 (a) for two-monomial pencils *above-J* via finite base-case enumeration plus the Substitution Principle. At $(n_0, k_0) = (2^{19}, 2^{17})$, this is roughly 4 orders of magnitude tighter than the non-universal $n_0/4 + 2$ form.

Corollary 5.7 (ABF §6.3 deployment, three-position sparse)**.** *At ABF §6.3 deployment (n_0, k_0) over the KoalaBear sextic $|F| = q^6 \approx 2^{186}$ (or any prime $q \geq 2^{31}$ admitting the required roots of unity), for f three-position sparse with* $\mathrm{supp}(\hat{f}) \subseteq [k_0, n_0 - 1]$ *(above-J at L_0):*

$$\varepsilon^{(2)}_{\mathrm{FRI}}(f) \;\leq\; \frac{10}{|F|} + (1 - \delta)^t \;\approx\; 2^{-182.7} + (1 - \delta)^t.$$

This is roughly 14 orders of magnitude tighter than the $n_0/4 + 2 = 2^{17} + 2$ form of Theorem 5.1 ($(2^{17} + 2)/|F| \approx 2^{-169}$) at the largest deployment scale.

Proof. Direct from Theorem 5.4 and the standard query phase $(1 - \delta)^t$. □

Remark 5.8 (Gap to the ABF §6.3 deployment instance). Corollary 5.2 closes Theorem 5.1 at the present small-prime parameter window. The ABF §6.3 instance calls for KoalaBear sextic, $k = 2^{20}$, $\rho = 1/2$, R up to $\log_2 n_0$, and general (not three-position-sparse) f. Theorem 5.1 is stated in (n_0, k_0)-uniform analytic form so that each scoped extension corresponds to a single named hypothesis: Conjecture 4.1 (a) at the deployment fields; two-monomial pencil propagation across $R \geq 3$ folds; and the reduction from general f to three-position-sparse witnesses (Conjecture 7.1; §7). The ABF §6.3 head-to-head plug-in is delivered in §6.

The proof of Theorem 5.1 has three ingredients (per-pair count via action–orbit; sum over arising pairs via the three-position sparse structure; query phase via standard agreement-density) and is given in Appendix B, §B.7.

5.1 Multi-round extension

Theorem 5.1 is stated for two-round FRI. We now upgrade to $R \geq 3$ via the following sparse-propagation lemma.

Lemma 5.9 (Sparse propagation across folds)**.** *If $\hat{f}$ on L_0 has DFT support of size at most three, then for every $r \in [0, R]$ and every $(\alpha_1, \ldots, \alpha_r) \in \mathbb{F}_q^r$, $\mathrm{fold}^r(\alpha_1, \ldots, \alpha_r)$ on L_r has DFT support of size at most three.*

Proof. Induction on r. Base $r = 0$: by hypothesis. Step $r \to r + 1$: the fold $\mathrm{fold}^{r+1} = (\mathrm{fold}^r)_e + \alpha_{r+1}(\mathrm{fold}^r)_o$ at L_{r+1} inherits its support as the union of the mod-2 projections of $\mathrm{supp}(\mathrm{fold}^r)$. The projection at most preserves the support size, so $|\mathrm{supp}| \leq 3$ at L_{r+1}. □

Theorem 5.10 (Multi-round FRI soundness)**.** *For three-position sparse $\hat{f}$ at rate $\rho = 1/4$ in R-round FRI, with the recursive above-Johnson and* mod 4 *pigeonhole conditions of Theorem 5.4 holding at every fold level,*

$$\varepsilon_{\mathrm{FRI}}^{(R)}(f) \;\leq\; \frac{10\,R}{|F|} \;+\; (1-\delta)^t,$$

rigorously and uniformly in n_0 at every deployment scale up to $(2^{19}, 2^{17})$.

Proof. By Lemma 5.9, fold^{r-1} at L_{r-1} is three-position sparse for every $(\alpha_1, \ldots, \alpha_{r-1})$. Applying Theorem 5.4 to fold^{r-1} as the input function for the round-r fold yields, per round, $|V_\delta^{(r)}|/q^r \leq 10/q$. Standard FRI multi-round soundness composition (per-round bad-challenge union bound plus the query-phase agreement-density bound; cf. [3], §8) gives $\varepsilon_{\mathrm{FRI}}^{(R)} \leq \sum_{r=1}^{R} |V_\delta^{(r)}|/q^r + (1-\delta)^t \leq 10R/|F| + (1-\delta)^t$. □

Remark 5.11 (Multi-round closure is no longer the bottleneck). At the ABF §6.3 deployment instance ($R \leq \log_2 n_0 \approx 21$, $|F| \approx 2^{186}$), the multi-round commit term satisfies $10R/|F| \approx 2 \cdot 10^{-54}$ — more than 50 bits below the 2^{-128} target. *The query-phase $(1-\delta)^t$ now dominates the soundness budget*, not the commit phase. (The proof of Theorem 5.10 requires the mod 4 pigeonhole at every fold level; this side condition propagates through generic $\hat{f}$ and is discharged rigorously by Theorem 5.12 below.)

Case C: all-different mod-4 (one quadrant empty). Theorem 5.4 relies on the mod-4 pigeonhole that two of the four bilinear components a, b, c, d vanish identically on L_2 (the Pattern B / Reverse cases). When the three positions of $\hat{f}$ span three distinct mod-4 residue classes, exactly one of $\{a, b, c, d\}$ vanishes and the other three are nonzero — "Case C". The fold-2 pencil reduces in this case to a 3-*monomial pencil on L_2*, parametrized by α_1 or α_2 depending on the mod-4 pattern.

Theorem 5.12 (Case C: char-uniform 3-monomial bound)**.** *For three-position sparse $\hat{f}$ whose DFT support has all-different* mod 4 *residues (Case C) at any rate-*1/4 *deployment scale, the level-*2 *pencil on L_2 is a 3-monomial pencil above the Johnson radius (Lemma 5.3), and*

$$\left|\{\alpha \in \mathbb{F}_q : c_1 z^{a_1} + c_2 z^{a_2} + \alpha c_3 z^{a_3} \text{ bad on } L_{n_2}\}\right| \le 9.$$

Combined with the at-most-one $\alpha = 0$ contribution, this gives $K(f) \le 10$ in Case C, for every deployment scale and every prime q admitting the required roots.

Proof. See Appendix B, §B.8. □

Remark 5.13 (Case C closes the recursive-quadrant-empty side condition)*.* Theorem 5.12 discharges the recursive-quadrant-empty side condition of Theorem 5.10 entirely: when the mod-4 pigeonhole at level r produces Case C rather than Pattern B/Reverse, the bound $K \le 10$ still holds rigorously via the 3-mono eliminator. Hence $K \le 10$ universal at every fold level for every 3-pos sparse $\hat{f}$, and Theorem 5.10 holds without the side condition.

5.2 Rate-1/2 extension: $K \le 4$ (2-mono), $K \le 28$ (3-pos sparse)

The natural Substitution Principle path to extend Theorem 5.4 from rate $\rho = 1/4$ to $\rho = 1/2$ is degenerate: at $n = 2k$, the cert + div eliminator σ of degree $2k = n$ collapses to $z^n - 1$ trivially. The fix is to choose $\deg \sigma = \lceil \sqrt{nk} \rceil$ (the Johnson agreement threshold) instead of $2k$ — a coincidence at rate 1/4 since $2k = \sqrt{(4k)k}$, but generalizing to any rate. With this correction, the rate-1/2 analogues of Theorems 5.4 and (the 2-mono companion) close rigorously at deployment scale.

Theorem 5.14 (Two-monomial, rate 1/2)**.** *At rate $\rho = 1/2$, every above-Johnson two-monomial pencil $h_{\rho,\alpha}(z) = \rho z^a + \alpha z^b$ on L_n has at most* 4 *bad challenges in $\mathbb{P}^1(\overline{\mathbb{F}_q})$:*

$$|B(a,b)| \le 4.$$

Same status as Theorem 5.15: rigorous at base panels $(4,2)$ and $(8,4)$, swept at $(16,8)$ and $(32,16)$, universal-k lift open. Over $\mathbb{F}_q$ the rational count is $|B \cap \mathbb{F}_q| \le \gcd(4, q-1)$, giving $K = 4$ on KoalaBear / BabyBear / Goldilocks and $K = 2$ on Mersenne 31*.*

Proof sketch. Singular GB at base cases $(n,k) \in \{(4,2),(8,4)\}$ with σ of degrees $\{3,6\}$. For all above-J coprime pairs at base, the eliminator $\Phi(\alpha,\rho)$ is either trivial ($\Phi = \alpha$, $K = 0$) or $\Phi = \alpha^4 - \rho^4$ ($K = 4$). Substitution Principle is expected to lift this to all rate-1/2 deployment scales; this is verified empirically by direct sweep at $(16,8)$ and $(32,16)$ across all 22 above-J coprime pairs (max $K = 4$ at every panel). The universal-k lift for arbitrary k remains open (the same SP-lift question as the 3-pos sparse case). □

Theorem 5.15 (Three-position sparse, rate 1/2)**.** *At rate $\rho = 1/2$, every above-Johnson three-position sparse pencil $h(z) = \alpha_1 z^{a_1} + \alpha_2 z^{a_2} + \alpha_3 z^{a_3}$ on L_n has at most* 28 *bad challenge tuples in $\mathbb{P}^2(\overline{\mathbb{F}_q})$:*

$$|B(\vec{a})| \le 28.$$

The bound is rigorous at the base panels $(n,k) \in \{(4,2),(8,4)\}$ and verified by direct sweep at $(16,8)$ and $(32,16)$; the universal-k lift via the Substitution Principle is open (Q3, §8). Over deployment fields with $8 \mid q-1$ (KoalaBear, BabyBear, Goldilocks), the $\mathbb{F}_q$-rational count satisfies the same 28 *bound; on Mersenne* 31 *($8 \nmid q-1$) the μ_8 stratum collapses to μ_2 and the count drops by a factor ≥ 4.*

Proof. See Appendix B, §B.9. □

Corollary 5.16 (ABF §6.3 deployment at rate $\rho = 1/2$, conditional)**.** *Conditional on the universal-k lift of Theorem 5.15 (rigorous at base panels $(4,2)$ and $(8,4)$, verified by sweep at $(16,8)$ and $(32,16)$; the universal-k closure for arbitrary $k \geq 4$ a power of 2 is open), at ABF §6.3 deployment $(n_0, k_0) = (2^{21}, 2^{20})$ over the KoalaBear sextic $|F| = q^6 \approx 2^{186}$ (or any prime $q \geq 2^{31}$ admitting the required roots of unity), for f three-position sparse with $\mathrm{supp}(\hat{f}) \subseteq [k_0, n_0 - 1]$ (above-J at L_0):*

$$\varepsilon^{(2)}_{\mathrm{FRI}}(f) \;\leq\; \frac{28}{|F|} + (1-\delta)^t \;=\; 2^{-181.12} + (1-\delta)^t.$$

For f two-monomial: $\varepsilon^{(2)}(f) \leq 4/|F| + (1-\delta)^t$. Over the Mersenne 31 prime *field (no extension), $8 \nmid q - 1$ collapses the μ_8 family to μ_2, sharpening the three-position sparse constant to $K \leq 4$ and the two-monomial constant to $K = 2$; this prime-field refinement does not survive the extension to $\mathbb{F}_{q^6}$ used at deployment, since $q+1 = 2^{31}$ implies $8 \mid q^6 - 1$ and μ_8 reappears (Table 5). Both bounds are roughly 30 orders of magnitude tighter than the BCHKS25 unconditional $n_0^5/|F|$ scaling at this deployment instance (which evaluates to $\approx 2^{-81}$). The factor of 2.8× between rate-1/4 ($K \leq 10$) and rate-1/2 ($K \leq 28$) reflects the geometric saturation of agreement subsets at the higher rate.*

Proof. Theorem 5.15 bounds the per-pencil bad-challenge count $|B(\vec{a})| \leq 28$ on the level-2 pencil at rate 1/2, analogously to the rate-1/4 Theorem 5.4. The sum-over-α_1 + query-phase composition is then identical to the proof of Theorem 5.4 (Appendix B, §B.6) with the rate-1/4 per-pencil constant 10 replaced by the rate-1/2 constant 28, yielding $\varepsilon^{(2)}_{\mathrm{FRI}}(f) \leq 28/|F| + (1-\delta)^t$ for 3-position sparse f. The two-monomial corollary follows from Theorem 5.14 via the same composition with constant 4. The conditional clause inherits from Theorem 5.15's universal-k open lift. □

Remark 5.17 (Why $K = 28$ matches $\binom{n}{J}$ at rate 1/2). Throughout this paper, "J" inside the binomial $\binom{n}{J}$ refers to the *Johnson agreement count* $J = \lceil \sqrt{nk} \rceil$ (size of an agreement subset achieving Johnson-list saturation), *not* the Hamming-distance Johnson radius $n\delta_J = n - \sqrt{n(k-1)}$ used in [3]; the two quantities should not be confused. At rate 1/2 base case $(8,4)$, the Johnson agreement count is $J = 6$ and there are $\binom{8}{6} = 28$ candidate agreement subsets of L_8. Each subset cuts out a single 3-pos pencil ratio $(\alpha_1 : \alpha_2 : \alpha_3)$ on which h interpolates a degree-$< k$ polynomial. The 28 pencils are all distinct in $\mathbb{P}^2$ (general position). Contrast rate-1/4 base $(8,2)$ with $J = 4$ and $\binom{8}{4} = 70$ subsets: heavy cancellation across subsets reduces K to $10 \ll 70$. The same saturation phenomenon recurs at rate 1/8 (Theorem 5.18 below); rate 1/4 is the unique *cancellation rate*.

Theorem 5.18 (Three-position sparse, rate 1/8, base panel)**.** *At the rate-1/8 base panel $(n,k) = (8,1)$, every above-Johnson three-position sparse pencil has at most $56 = \binom{8}{3} = \binom{n}{J}$ bad challenge tuples; the two-monomial companion bound at the same panel is $K \leq 2$. The deployment-scale lift to general (n,k) at rate 1/8 is open.*

Proof. See Appendix B, §B.10. □

Remark 5.19 (Rate-coverage summary). Rate-by-rate K-bounds at the verified base panels:

Rate ρ	$K_{\text{2-mono}}$	$K_{\text{3-pos}}$
1/4	≤ 8	≤ 10
1/2	≤ 4	≤ 28
1/8	≤ 2	≤ 56

The pattern $K = \binom{n}{J}$ saturates at rates 1/2 and 1/8; rate 1/4 has heavy cancellation ($\binom{8}{2} = 28$ vs. $K = 10$). The 4-position-sparse extension at every rate is open: the natural base panels $(8, 4)$ and $(8, 1)$ are degenerate (4 above-J positions used $\Rightarrow$ cert constraint trivial), and a proper base-panel eliminator at $(16, 8)/(16, 2)$ analogous to Theorem 5.12 is the natural Q4 task (§8).

5.3 Query-budget gain over the CA bound

The per-query base $(1 - \delta)$ replaces [1]'s $(1 - \delta/2)$ at the same δ. At $\delta = 0.4$, $\kappa = 128$, the required query count drops from ∼398 (Paper 1) to ∼173 (this paper), a ∼2.3× query-budget reduction. This is the analytic-form gain; the deployment-scale head-to-head against ABF Tables 2–5 (incorporating commit-phase and Fiat–Shamir overhead) is in §6.

6 Head-to-Head with ABF §6.3 at 128-bit Security

The Arnon–Boneh–Fenzi survey [2] §6.3 is the canonical benchmark for STARK proof size at Ethereum deployment parameters — ABF Tables 2–5 score every candidate construction at this instance. We populate the missing row of those tables: a *plain Reed–Solomon* construction at 128-bit security, with proof size 79.8 KiB. ABF's smallest deployable rigorous candidates on stronger code classes are 161.4 KiB (interleaved RS) and 281.2 KiB (folded RS); we halve those numbers on the code class production actually ships.

6.1 The ABF §6.3 instance, restated

Parameters. Plain Reed–Solomon over the sextic extension of the KoalaBear prime $q = 2^{31} - 2^{24} + 1$, with $|\mathbb{F}| = q^6 \approx 2^{186}$. Message degree $k = 2^{20}$; codeword length $s \cdot n = 2^{21}$ on a smooth multiplicative subgroup; rate $\rho = k/(s \cdot n) = 1/2$. Distance threshold $\delta \in (\delta_J(1/2), 1/2)$ above the Johnson radius $\delta_J(\rho) = 1 - \sqrt{\rho}$. Folding parameter s swept over $\{2^0, 2^1, \ldots, 2^{12}\}$. Query budget $t = 128$ at target 128-bit security. Argument size measured in KiB.

The ABF soundness expression. ABF Lemma 6.6 / 6.8 give

$$\varepsilon_{\mathrm{FRI}}^{\mathrm{ABF}} \;\le\; \max\Big\{\varepsilon_{\mathrm{mca}}(C, \delta) + \frac{|\Lambda(C^{=2}, \delta)|}{|\mathbb{F}|},\; (1 - \delta)^t\Big\},$$

both terms required to be $\le 2^{-128}$. *For plain RS above the Johnson radius the surveyed positive bound on $\varepsilon_{\mathrm{mca}}$ is open or insufficient at deployment scale*; this status is exactly the regime in which deployments operate.

6.2 Plugging the soundness expression into ABF §6.3

ABF §6.3 fixes rate $\rho = 1/2$ ($k = 2^{20}$, $n = 2^{21}$). The direct rate-1/2 specialization is Corollary 5.16 (below). For comparison, the rate-$\rho = 1/4$ analytic plug-in of Theorem 5.1 (with the same codeword length $n_0 = 2^{21}$ but message degree $k_0 = n_0/4 = 2^{19}$) gives

$$\varepsilon_{\mathrm{FRI}}^{(2)}(f) \;\le\; \frac{n_0/4 + 2}{|\mathbb{F}|} \;+\; (1 - \delta)^t \;=\; (2^{19} + 2)/2^{186} + (1 - \delta)^t \;\approx\; 2^{-167} + (1 - \delta)^t.$$

The improved Theorem 5.4 replaces the $n_0/4 + 2$ leading constant with 10, dropping the commit term to $10/|\mathbb{F}| \approx 2^{-183}$. The per-query base advances from $(1-\delta/2)$ (Paper 1, threshold-halved CA) to $(1-\delta)$, the source of the ∼2.3× query-budget reduction at $\delta = 0.4$.

Rate-1/2 deployment constants. Corollary 5.16 is the exact ABF-rate specialization: three-position sparse f has commit term $28/|\mathbb{F}|$ at $\rho = 1/2$, while the two-monomial companion has commit term $4/|\mathbb{F}|$ (sextic-grade deployment fields including KoalaBear, BabyBear, Mersenne 31; Mersenne 31 prime has the further $\mu_8 \to \mu_2$ collapse to $K = 2/\leq 4$ but is not deployable due to $\log_2 q = 31$). Thus at KoalaBear sextic,

$$\log_2(28/|\mathbb{F}|) = -181.12, \qquad \varepsilon^{(2)}_{\mathrm{FRI}}(f) \leq 28/|\mathbb{F}| + (1-\delta)^t.$$

At the ABF-native $\delta = 0.499$ and $t = 128$, the query term gives 127.63 bits; strict 128-bit security uses $t = 129$, whose optimum under the ABF argument-size formula is 80.4 KiB at $s = 2^3$.

Tightening via Theorem 5.4. For three-position-sparse f above-J at L_0 with at least one mod-4 residue class empty, Theorem 5.4 replaces the leading $(n_0/4 + 2)/|\mathbb{F}|$ commit-phase term with the n_0-*independent* constant $10/|\mathbb{F}|$:

$$\varepsilon^{(2)}_{\mathrm{FRI}}(f) \;\leq\; \frac{10}{|\mathbb{F}|} \;+\; (1-\delta)^t \;\approx\; 4.7 \times 10^{-9} + (1-\delta)^t \quad \text{at } |\mathbb{F}| = q \approx 2^{31},$$

rigorously and unconditionally. At $|\mathbb{F}| \approx 2^{186}$ this is $10/2^{186} \approx 2^{-183}$, ≥ 55 bits below the 2^{-128} target. Theorem 5.10 extends this to R-round soundness $10R/|\mathbb{F}|$, $\leq 2^{-180}$ at $R = 21$ at the largest deployment scale. The argument-size table below (Table 4) is therefore unchanged at the query-phase optimum, but the commit-phase analytic term is now universal in n_0 rather than $O(n_0)$.

6.3 Head-to-head comparison with ABF Tables 2–5

ABF [2] Tables 2–5 score Construction 6.2 at the deployment instance under two different code instantiations: interleaved Reed–Solomon (IRS, Tables 2 and 3) and folded Reed–Solomon (FRS, Tables 4 and 5). Both rely on the surveyed positive $\varepsilon_{\mathrm{mca}}$ bounds for those code classes. Plain Reed–Solomon above Johnson — the canonical FRI deployment class — is *not* scored by ABF, because the surveyed positive $\varepsilon_{\mathrm{mca}}$ bound for that code is open ([2] §4.2–4.3). This paper supplies the missing positive bound and produces the analogue scoring under it.

Soundness at fixed query budget $t = 128$ (Tables 2, 3 below vs. Theorem 5.1). At the ABF §6.3 native $t = 128$, IRS (Table 2) delivers a uniform $2^{-64.00}$ soundness across s — about 64 bits, half of the 128-bit target. FRS (Table 3) fails to meet $\tau(r+1) + 3/(2r) < 1$ for $s \leq 2^4$ (no soundness provable) and degrades from $2^{-29.11}$ to $2^{-118.14}$ across the remaining range. Theorem 5.1 at $\delta = 0.499$ delivers $\sim 2^{-127.6}$ uniformly — the ABF-native near-128 row at $t = 128$, a $\sim 2\times$ security gain at the same query budget.

Argument size at fixed 128-bit security (Table 4 vs. Theorem 5.1 and Corollary 5.16). To meet 128-bit security ABF must inflate the query budget: IRS to $t = 259$ (Table 4, optimum 161.4 KiB at $s = 2^3$); FRS to $t \in [139, 563]$ (Table 4, optimum 281.2 KiB at $s = 2^6$). Theorem 5.1 reaches the ABF-native $t = 128$ row (127.6 bits at $\delta = 0.499$); the strict 128-bit rate-1/2 column from Corollary 5.16 uses $t = 129$.

Reading the comparison. The $2\times$–$3.5\times$ argument-size reduction at 128-bit security has a single source: the per-query base in Theorem 5.1 is $(1-\delta)$ rather than the threshold-halved $(1-\delta/2)$ of CA-based bounds, so ABF-native near-128 security at $t = 128$ (or strict 128 bits at $t = 129$) rather

s	ABF IRS bits	arg KiB
2^0	$2^{-64.00}$	85.0
2^1	$2^{-64.00}$	81.9
2^2	$2^{-64.00}$	79.9
2^3	$2^{-64.00}$	**79.8**
2^4	$2^{-64.00}$	83.5
2^5	$2^{-64.00}$	95.0
2^6	$2^{-63.99}$	122.0
2^7	$2^{-63.98}$	180.0
2^8	$2^{-63.97}$	300.0
2^9	$2^{-63.94}$	544.0
2^{10}	$2^{-63.87}$	1 033.2
2^{11}	$2^{-63.75}$	2 027.5
2^{12}	$2^{-63.49}$	4 003.8

Table 2: ABF [2] Table 2: IRS at $t = 128$. Security ~ 64 bits across s (half the 128-bit target). Argument-size minimum $\sim$79.8 KiB at $s = 2^3$.

s	r	ABF FRS bits	arg KiB
2^5	8	$2^{-29.11}$	95.0
2^6	11	$2^{-55.57}$	122.0
2^7	17	$2^{-75.39}$	180.0
2^8	25	$2^{-90.04}$	300.0
2^9	36	$2^{-100.76}$	544.0
2^{10}	53	$2^{-108.53}$	1 033.2
2^{11}	75	$2^{-114.13}$	2 027.5
2^{12}	108	$2^{-118.14}$	4 003.8

Table 3: ABF [2] Table 4: FRS at $t = 128$. Soundness unprovable for $s \leq 2^4$ (constraint $\tau(r+1)+3/(2r) < 1$ infeasible). Security ranges 2^{-29} to 2^{-118}.

than $t \in [139, 563]$. The rate-1/2 commit-phase constants are field-dependent but query-phase-bottlenecked for KoalaBear/BabyBear sextic and Mersenne 31 sextic (Table 5). The substantive content of this comparison: a code class (plain Reed–Solomon above Johnson) for which ABF cannot supply *any* positive ε_{mca} bound now has a concrete deployment-grade soundness expression that beats both IRS and FRS scoring under their own positive ε_{mca} bounds.

6.4 Caveats on the comparison

(a) ABF accounting reproduced. The argument-size formula $t \cdot (256 \log_2(sn/s) + 62s)$ reproduces every ABF Tables 2–5 row with $s \leq 2^8$ to within 0.05 KiB ([7]); $s \geq 2^9$ rows differ by $\sim$0.5% from a multi-round overhead but lie deep in the suboptimal tail and do not affect head-to-head. The same formula is used on both sides of Table 4, so the comparison is apples-to-apples. **(b) Commit-phase margin.** Theorem 5.10 closes the multi-round-closure direction rigorously at $10R/|\mathbb{F}| \leq 210/2^{186} \approx 2^{-180}$ for 3-position-sparse f above-J at the ABF instance, leaving the soundness expression query-phase-bottlenecked. **(c) Conj 4.1(a) deployment-field sweep.** Currently at $q \in \{97, 193, 257\}$; the extension to KoalaBear and other deployment fields is recorded in the companion repository [7]. **(d) Three-pos sparse audit.** Verified at deployment scales $(n_0, k_0) \in \{(32, 8), (64, 16), (128, 32)\}$ ([7]). The general-f extension is treated in §7.

7 Layer 3: General-f Extension via Sparse-Worst Dominance

Theorem 5.4 closes $K(f) \leq 10$ at rate 1/4 unconditionally for 3-position-sparse f. The general-f extension splits into two parts: *(i) local closure* of specific obstructions to sparse-worst-case dominance, and *(ii) global attachment* of those local closures into a single quantifier covering arbitrary f. Part (i) we close rigorously at every dyadic depth (Theorems 7.7, 7.8). Part (ii) is the open *sparse-worst-case dominance* conjecture (Q2) — consistent with every adversarial construction in the proximity-gap literature, exact-verified at deployment-shaped scale $(32, 8)$ over $q \in \{97, 193, 257\}$ (0 counterexamples in $\binom{32}{8} = 10{,}518{,}300$ information sets), and supported by 4,586,096 support-prime certificates. The only known structural would-be counterexample falls outside the action-

s	ABF IRS t	ABF IRS KiB	ABF FRS t	ABF FRS KiB	ρ=1/4 t	ρ=1/4 KiB	ρ=1/2 t	ρ=1/2 KiB
2^1	259	165.8	–	–	128	81.9	129	82.6
2^2	259	161.6	–	–	128	79.9	129	80.5
2^3	259	**161.4**	–	–	128	**79.8**	129	**80.4**
2^4	259	169.0	–	–	128	83.5	129	84.2
2^5	259	192.2	563	417.9	128	95.0	129	95.7
2^6	259	246.9	295	**281.2**	128	122.0	129	123.0
2^7	259	364.2	218	306.6	128	180.0	129	181.4
2^8	259	607.0	182	426.6	128	300.0	129	302.3
2^9	260	1 105.9	163	692.8	128	544.0	129	548.2
2^{10}	260	2 109.4	151	1 218.6	128	1 036.0	129	1 044.1
2^{11}	260	4 116.0	144	2 273.7	128	2 024.0	129	2 039.8
2^{12}	259	8 099.8	139	4 352.0	128	4 004.0	129	4 035.3

Table 4: Apples-to-apples comparison at 128-bit security (ABF Tables 3 and 5 vs. this paper via the present paper's generator). Optima in bold. ABF IRS minimum 161.4 KiB at $s = 2^3$, $t = 259$. ABF FRS minimum 281.2 KiB at $s = 2^6$, $t = 295$. The $\rho = 1/4$ column is the original $K \le 10$ headline: 79.8 KiB at ABF-native $t = 128$ (127.6 bits at $\delta = 0.499$). The $\rho = 1/2$ column uses Corollary 5.16 with $K = 28$ and gives strict 128-bit security at 80.4 KiB, $s = 2^3$, $t = 129$: **a 2.0× improvement over ABF IRS and a 3.5× improvement over ABF FRS at the optimum**, applicable to plain Reed–Solomon (the code class actually deployed) which neither IRS nor FRS scores. ABF KiB values from [2] Tables 2–5; ours from `notes/scripts/p0_abf_table_generator.py` via the ABF §6.3 native argument-size formula $t \cdot (256 \log_2 n + 62s)$ bits with $sn = 2^{21}$.

non-stabilised admissibility class baked into the conjecture quantifier (Remark 7.3). Closing Q2 lifts Theorem 7.13 to unconditional general-f rate-1/4 closure.

7.1 The sparse-worst-case dominance conjecture

Conjecture 7.1 (Sparse-worst-case dominance for plain RS above Johnson)**.** *Let $C = \mathrm{RS}_{k_0}(L_0)$ with L_0 a multiplicative subgroup of $\mathbb{F}_q^*$ of order $n_0 = 4k_0$, $4 \mid k_0$, at a FRI two-round deployment scale $(n_0, k_0) \in \{(32, 8), (64, 16), \ldots, (2^{19}, 2^{17})\}$ over a prime $q \ge Q_{\min} = 97$ admitting the required primitive roots (Lemma 7.6 excludes $q = 17$ as a small-prime degeneracy; all FRI deployment primes satisfy $q \gg Q_{\min}$ by orders of magnitude). For $\delta > 1 - \sqrt{\rho}$ (above-Johnson), define*

$$K(f_1, f_2; \delta) \;:=\; \big|\{\alpha \in \mathbb{F}_q : \Delta(f_1 + \alpha f_2,\, C) \le \delta\}\big|.$$

Call (f_1, f_2) action-non-stabilised *if for every arising fold pair (a, b) on the joint support of $(\hat{f}_1, \hat{f}_2)$, the support is not fixed pointwise by the cyclotomic-subgroup action $\langle \omega^{b-a} \rangle$ (i.e., the round-1 challenge does non-trivial mixing on the support; cf. Remark 7.3). Then the maximum over all strictly above-Johnson, action-non-stabilised adversaries is achieved on a* three-position sparse *pair:*

$$\max_{\substack{(f_1,f_2)\in(\mathbb{F}_q^{n_0})^2,\\ \Delta((f_1,f_2),C^2)>\delta,\\ \textit{action-non-stabilised}}} K(f_1, f_2; \delta) \;=\; \max_{\substack{(f_1,f_2)\ 3\text{-}\textit{pos sparse},\\ \mathrm{supp}(\hat{f}_i)\subseteq[k_0,n_0-1]}} K(f_1, f_2; \delta).$$

The action-non-stabilisation predicate is non-trivial: without it the conjecture is technically falsified by the $L_2 = (16, 4)/\mathbb{F}_{193}$ paired-circuit obstruction (Remark 7.3); with it, the obstruction is excluded and no further structural counterexample is known.

Remark 7.2 (Strict above-Johnson admissibility). The admissibility condition $\Delta((f_1, f_2), C^2) > \delta$ in Conjecture 7.1 requires the strict inequality $\delta > 1 - \sqrt{\rho}$, excluding the Johnson boundary $\delta =$

Deployment field	$\log_2 \lvert\mathbb{F}\rvert$	$K_{\text{3-pos}}$	$K_{\text{2-mono}}$	$\log_2(K_3/\lvert\mathbb{F}\rvert)$	bits at $t = 128$	strict 128-bit optimum	Root-of-unity effect
KoalaBear sextic	185.93	28	4	−181.12	127.63	80.4 KiB at $s = 2^3$, $t = 129$	$8 \mid q - 1$
BabyBear sextic	185.44	28	4	−180.63	127.63	80.4 KiB at $s = 2^3$, $t = 129$	$8 \mid q - 1$
Goldilocks quadratic	128.00	28	4	−123.19	123.13	commit-limited	$8 \mid q - 1$
Mersenne 31 sextic	186.00	28	4	−181.19	127.63	80.4 KiB at $s = 2^3$, $t = 129$	$8 \mid q^6 - 1$ via $q + 1 = 2^{31}$

Table 5: Rate-1/2 deployment-field profile for Corollary 5.16, generated by `notes/scripts/p0_abf_table_generator.py -rate-half-fields`. The four sextic-grade deployment fields all support the $K = 28$ rate-1/2 column with commit margin $\leq 2^{-181}$, far below the 2^{-128} target on the three sextic fields with $\log_2 \lvert\mathbb{F}\rvert \approx 186$ (KoalaBear, BabyBear, Mersenne 31). The Mersenne 31 prime-field $K_{\text{3-pos}} \leq 4$, $K_{\text{2-mono}} = 2$ refinement (§5.2) does *not* carry through to the sextic extension: $q + 1 = 2^{31}$ for Mersenne 31 implies $8 \mid q + 1 \mid q^6 - 1$, so $\mu_8 \subset \mathbb{F}_{q^6}$ and the μ_8 collapse argument that gives $K = 4$ at the prime field does not apply at the sextic. The prime-field $K = 4$ row is recorded for theoretical interest only; the prime field is too small ($\log_2 q = 31$) to deploy at 128-bit security regardless of K. Goldilocks quadratic is shown explicitly because its $K = 28$ commit term alone is only 123.2 bits; a deeper Goldilocks extension (q^3, $\log_2 \lvert\mathbb{F}\rvert = 192$) restores ≥ 60-bit margin and is the natural deployment choice on this prime.

$1 - \sqrt{\rho}$ exactly. This exclusion is non-trivial: parity-aligned 3-position sparse pencils with support $\text{supp}(\hat{f}_i) = \{n/2, n/2+1, \ldots, n/2+s-1\}$ at rate $\rho = 1/4$ admit, for every coefficient choice $(c_1, c_2) \in \mathbb{F}_q^s \times \mathbb{F}_q^s$ and every $\alpha \in \mathbb{F}_q^*$, a candidate codeword at Hamming distance exactly $n\delta_J = n(1 - \sqrt{\rho})$ (the relative Johnson radius; *not* the agreement count $J = \lceil \sqrt{nk} \rceil$ used in Remark 5.17) via the algebraic identity $\omega^{n/2} = -1$ (full derivation in [7], Note 0320). Such pencils give $K(f_1, f_2; \delta_J) = q - 1$ at the boundary $\delta = \delta_J$ *exact*, but never satisfy $\Delta(f_i, C) > \delta_J$ *strictly* — empirical sweep of 100 random (c_1, c_2) at $(32, 8)/p = 257$ gives 0/100 strict-above-J samples (verification script and per-pencil distance histogram in [7]). The boundary construction is therefore structurally isolated from the conjecture's strict above-J domain, and the conjecture is not refuted by this construction. In the FRI 2-round protocol, an honest verifier samples δ generically within the strict above-J region, so the boundary subspace contributes ε-measure-zero to the soundness analysis.

Remark 7.3 (Two structural exclusions of the $K = q$ paired-circuit obstruction). An explicit arbitrary-coefficient witness to the literal worst-case quantifier of Conjecture 7.1 exists at $L_2 = (16, 4)$ over $\mathbb{F}_{193}$: the six-position support $(32, 36, 40, 34, 38, 46)$ with coefficients $(112, 79, 1, 30, 47, 1)$ gives $K(f_1, f_2; \delta_J) = q = 193$ ([7], Notes 0377/0378). This obstruction is excluded from Conjecture 7.1 (and from Theorem 5.4's scope) by *two* independent structural mechanisms.

(i) Action–orbit stabilization. The support lies entirely in fold quadrants $\{0, 2\}$ on L_2, so the residual rows $u_{\alpha_1}, v_{\alpha_1}$ are independent of α_1 — the round-1 challenge does no mixing. This is the $\langle \omega^{b-a} \rangle$-stabiliser phenomenon of Theorem 2.1; it inflates K by the trivial-action factor of q. The same mechanism produces $K = q$ witnesses on the three other action-stabilised quadrant-pair classes (verified at the six deployment-relevant primes admitting a primitive 64-th root; [7], Note 0379). The conjecture is stated with the *action-non-stabilised* admissibility predicate baked into the quantifier precisely to exclude these degenerate inputs, which do not arise from honest FRI verifier behaviour at deployment scale.

(ii) Level-1 Johnson boundary. Independently, an all-α_2 saturating component on L_2 lifts to the level-1 fold with agreement *exactly* $\sqrt{n_1 k_1}$ — on the Johnson threshold, not strictly above it. The recursive above-Johnson hypothesis of Theorem 5.4 requires strict excess and so excludes this boundary stratum ([7], Note 0388, scale-uniform in n_0).

The obstruction therefore occupies a precisely characterised boundary stratum — a finite union of proper linear subspaces parametrised by paired tail-circuit kernels ([7], Note 0386) — and is not

a strict above-Johnson counterexample to either Theorem 5.4 or Conjecture 7.1.

Remark 7.4 (Why the conjecture closes the gap). ABF Lemma 6.13 (the soundness lower-bound construction) is itself a sparse adversary: the worst-case (f_1, f_2) in the soundness inequality is realised by a Crites–Stewart-type sparse pair. The attacker's lower bound is therefore inside the sparse class. The upper-bound side, on the other hand, currently quantifies over all (f_1, f_2) via the correlated-agreement parameter $\varepsilon_{\mathrm{ca}}(C, \delta)$ of ABF Definition 4.1. Conjecture 7.1 closes the symmetry: *if* the worst adversary is sparse, then Theorem 5.4 ($K \leq 10$ for sparse) bounds $\varepsilon_{\mathrm{ca}}(C, \delta) \leq 10/q$ for *all* f above J, removing the moment-bound hypothesis entirely.

Remark 7.5 (Empirical evidence). At $(n, k) = (16, 4)$ over $\mathbb{F}_{17}$ with above-Johnson threshold $\delta n = 9$ (agreement $\leq 7 < t_J = \sqrt{nk} = 8$), exhaustive sweep over 4000 random adversarial pairs gives:

$$K^{\max}_{\mathrm{sparse}} \;=\; 17 \quad (\text{saturated at } q), \qquad K^{\max}_{\mathrm{dense}} \;=\; 16,$$

where the dense filter requires $\Delta(f_i, C) > \delta n$ for both i (1997/2000 pairs passed). In particular, $K^{\max}_{\mathrm{dense}} < K^{\max}_{\mathrm{sparse}}$, consistent with the conjecture. Top sparse witnesses include $\mathrm{supp}_1 = \{11, 12, 13\}$, $\mathrm{supp}_2 = \{4, 6, 12\}$ and four others listed in the companion repository [7]. This is a small-scale empirical test limited by brute-force RS list-decoding at $q^k = 17^4$; deployment-scale verification at $(32, 8)$, $(64, 16)$, $(128, 32)$ requires a Guruswami–Sudan list-decoder above Johnson.

Deployment-scale upgrade $(n, k) = (32, 8)$ ([7]). Exact information-set enumeration over $\binom{32}{8} = 10{,}518{,}300$ candidate agreement sets verifies the sparse 3-position witness with support $(14, 22, 30)$ at every valid prime $q \in \{97, 193, 257\}$: $K^{(32,8,q)}_{\mathrm{sparse}} = 6$ (exact). Higher-arity exact stress tests (4-support / 5-support at $q \in \{97, 193, 257\}$) give $K = 2$, consistent with the conjecture's 3-position-sparse maximizer. Dense-side coverage (50 random pencils $\times$ 50,000 information sets per prime) gives $K^{\mathrm{lb-max}}_{\mathrm{dense}} = 0$ across all three primes (miss union bound $\sim 10^{-10}$), strong empirical (not yet exact) evidence that $K^{(32,8,q)}_{\mathrm{dense}} < K^{(32,8,q)}_{\mathrm{sparse}} = 6$. Exact closure to $K_{\mathrm{dense}} \leq 5$ at $(32, 8)$, $(64, 16)$, $(128, 32)$ requires an above-Johnson list decoder; a structural alternative via the action-stratified L_2 decomposition bypasses the list-decoder requirement entirely.

Multi-prime certification at deeper panels. Extending the deployment proxy to higher-arity adversary panels at the $L_2 = (16, 4)$ folded model, the path-(c) symbolic certifier classifies each $(\mathrm{support}, S)$ pair as collapsed, full, saturated-stabilizer, or *primitive* (rank-2, trivial stabilizer; would refute the path-(c) target). Cumulative results across four panel sizes and 8 deployment-relevant primes:

Panel	Primes tested	Combinations	Primitive counterexamples
$C(48, 3)$ at $n_{\max} = 64$	7	121,072	0
$C(64, 3)$ at $n_{\max} = 80$	12	499,968	0
$C(80, 3)$ at $n_{\max} = 96$	8	657,280	0
$C(96, 3)$ at $n_{\max} = 112$	8	1,143,040	0
$C(112, 3)$ at $n_{\max} = 128$	8	1,823,360	0
$C(128, 3)$ at $n_{\max} = 144$	1	341,376	0
Total		**4,586,096**	**0**

The saturated-stabilizer count exhibits a clean dyadic doubling tier $(0, 2^{13}, 2^{15}, 2^{16})$ panel- and prime-uniform, suggesting these cases form a structurally fixed locus governed by descent rather

than residual primitive counterexamples. Primes: $q \in \{97, 113, 193, 257, 449, 577, 769, 1153\}$. $q = 17$ is excluded as a μ_{16}-saturation small-prime degeneracy (Lemma 7.6 below); a bracket sweep at $q \in \{113, 241, 337, 401, 433\}$ confirmed $Q_{\min} = 97$ as the unique small-prime exception. All FRI deployment primes (KoalaBear, BabyBear, Goldilocks, Mersenne 31, Plonky2) are exponentially larger than 97; the small-prime artifact does not affect deployment applicability. The conjecture is therefore restated in deployment-uniform form:

> *For all $q \geq Q_{\min} = 97$, and at every $L_2 = (n, k)$ folded deployment proxy, $K^{(\cdot,q)}_{\text{dense}} \leq K^{(\cdot,q)}_{\text{sparse}}$ on above-Johnson adversary pairs.*

Lemma 7.6 (μ_{16}-saturation forces $q = 17$ small-prime degeneracy)**.** *Let q be a prime with $q \equiv 1 \pmod{16}$ admitting $\mu_{16} \subset \mathbb{F}_q^*$. The path-(c) symbolic certifier solves linear systems of the form $C(S) + \alpha_1 \cdot M(S) = 0$ in $\mathbb{F}_q$ for 8-subsets $S \subset L_2 = (16, 4)$, classifying $(\text{support}, S)$-pairs by the rank of their alpha-stratified residual. For $q = 17$, the multiplicative group $\mathbb{F}_q^*$ has order $q-1 = 16$, so $\mathbb{F}_q^* = \mu_{16}$ exactly: every nonzero element of $\mathbb{F}_q$ is a 16-th root of unity. In this maximally- saturated configuration, generic separation between $C(S)$ and $M(S)$ in $\mathbb{F}_q$ fails: there exist $(\text{support}, S)$ pairs for which the symbolic linear system collapses to $0 \equiv 0$ identically, producing a residual that classifies as primitive rank-2 no-full-block by formal definition but arises only because the coefficient ring has too few values.*

For every $q \geq 97$ with $q \equiv 1 \pmod{16}$, $\mathbb{F}_q^*$ properly contains μ_{16} (with index $(q-1)/16 \geq 6$), and the additional non-trivial multiplicative structure restores generic separation: empirically (see the 4,586,096 support-prime cert above), no primitive counterexample appears at any tested $q \geq 97$. The lower bound $Q_{\min} = 97$ is tight in the empirical sweep: $q = 17$ is the smallest, and only, prime $\equiv 1 \pmod{16}$ in the test set that produces primitive "counterexamples".

Theorem 7.7 (No-full primitive base-panel closure)**.** *Fix the folded model $L_2 = \mu_{16}$ (cyclic of order 16, four quarter blocks μ_{16}^4-cosets of size 4). Let $S \subset L_2$ have $|S| = 8$ and call S no-full if it contains no full quarter block. Let $(a_1, a_2, a_3) \subset [k_2, n_{\max})$ be a strict above-Johnson 3-support in the legal $C(48, 3)$ window with $n_{\max} = 64$. Then no $\alpha \in \overline{\mathbb{Q}}^3$ produces a primitive rank-2 saturated component of the path-(c) residual pencil $W(\alpha)$ on a no-full S. Equivalently: every saturated no-full candidate satisfies $\alpha_1 = 0$ (zero-row family) or has side-pure support in one of $\{j \bmod 4 \in \{0, 1\}\}$ or $\{j \bmod 4 \in \{2, 3\}\}$, in which case $\operatorname{rank} W(\alpha) \leq 1$.*

Proof. See Appendix B, §B.11. □

Theorem 7.8 (General dyadic tail lemma; scale-lift of Theorem 7.7)**.** *Let $\mathbb{F}$ be a field of odd characteristic — a hypothesis satisfied by every FRI deployment field (KoalaBear, BabyBear, Mersenne 31, Goldilocks are all extensions of odd primes) — and let $L \subset \mathbb{F}^*$ be a cyclic subgroup of order $4k$ with $k \geq 1$. Partition L into four quarter blocks $C_b = \{x \in L : x^k = i^b\}$, $b = 0, 1, 2, 3$, where i is a primitive 4-th root of unity. Let $S \subset L$ have $|S| = 2k$ and call S no-full if $|S \cap C_b| < k$ for every b. Then for every $k \leq e < 4k$ there is no polynomial $p \in \mathbb{F}[x]$ with $\deg p < k$ such that $p|_S = x^e|_S$. Equivalently, $\text{tail}_S(x^e) \neq 0$ for every no-full S at every dyadic depth.*

Proof. Suppose such p exists; write $e = ak + r$, $a \in \{1, 2, 3\}$, $0 \leq r < k$. Cases $a = 1$ and $a = 3$ produce a non-zero polynomial of degree $< 2k$ vanishing on $|S| = 2k$ points (immediate degree contradiction; case $a = 3$ uses $x^{4k} = 1$ on L to multiply through by x^{k-r}). Case $a = 2$ uses $x^{2k} = \pm 1$ on the two parity halves of L: the interpolation condition becomes $p \mp x^r = 0$ on $S \cap (E, O)$ respectively. Both $p \pm x^r$ have degree $< k$, so if both are nonzero $|S| \leq 2(k-1) < 2k$, contradiction; hence one is identically zero. If $p = \pm x^r$ then $p \pm x^r = \pm 2x^r$ is nonzero on L

in odd characteristic, so the corresponding parity half lies in S entirely, contradicting the no-full hypothesis. Full case analysis: [7] Notes 0337–0338. □

Remark 7.9 (Strength and scope of Theorem 7.8). The proof requires only (i) odd characteristic (used in the $a = 2$ case to ensure $2x^r \neq 0$ on L) and (ii) a cyclic subgroup of order $4k$ in $\mathbb{F}^*$. It does *not* require k to be a power of two, nor any Galois-theoretic descent through cyclotomic sub-fields, nor a Gröbner-basis or Chinese-Remainder reformulation of the tail. In particular it covers every FRI-deployable folded model $L_2 = (n_2, k_2)$ with $4 \mid n_2$ and q admitting the required primitive roots, including all rate-1/4 deployment cells $(n_0, k_0) \in \{(32, 8), (64, 16), \ldots, (2^{19}, 2^{17})\}$ at every prime $q \geq 97 = Q_{\min}$ (Lemma 7.6).

Combined with Theorem 7.7, this closes the *local* no-full primitive obstruction in path-(c) at every dyadic depth: every saturated no-full candidate at every deployment-shaped folded model satisfies $\alpha_1 = 0$ or is one-sided (rank ≤ 1), hence not primitive rank-2. The remaining gap to a fully unconditional path-(c) closure of Conjecture 7.1 is *global*: attaching this local singleton-tail closure to the full general-f dyadic component decomposition. The 4,586,096-combination empirical cert across $C(48, 3)$ to $C(128, 3)$ panels remains consistent at every depth tested.

Remark 7.10 (Why one should expect the conjecture to hold). The Crites–Stewart family — explicitly three-position sparse — is known to saturate the Johnson bound at infinitely many ranks [4]. No published construction of a dense f exceeding the CS bound is known. The action–orbit theorem (Theorem 2.1) further suggests that the bad-α set has a sparse-Fourier obstruction: the orbit structure $\langle \omega^{b-a} \rangle$ on L_n couples directly to the sparsity of $\hat{f}$. A dense adversary distributes Fourier mass across many positions, breaking the coherent orbit-cover that drives K to its sparse maximum.

Remark 7.11 (Codim-composition route to Conjecture 7.1). A companion structural attack [7] (the codim-composition route) gives a different angle: $K(f_1, f_2; \delta)$ is the count of challenges for which the syndrome line $s_1 + \alpha s_2$ meets the Berlekamp realizer variety, so high-K pairs lie in the syndrome-pair geometry controlled by a V_{bad} stratification of the syndrome space [7]. Two guardrails: (i) Fourier monomial support (used in Conjecture 7.1) and Vandermonde syndrome support (used in the codim-composition route) are inequivalent; small Fourier support need not imply membership in a leading $V_S \times V_S$ component. (ii) The correct admissibility is joint $|S^*(s_1, s_2)| > w$, not endpoint-wise distance. Concrete deployment-curve construction ([7], Zariski-open branch theorem): for the support pattern $(9, 11, 12)/(6, 8, 12)$ at $(n, k, c) = (16, 4, 3)$, the leading incidences live on *eight codimension*-1 *branches* $b_8 = \omega^{-2r} b_6$, $r = 0, \ldots, 7$, each carrying a nonempty Zariski-open subset on which every point yields exactly $K = 9$ joint-admissible bad challenges (eight closed-form non-zero distinct, plus $\alpha = 0$). Verifier in companion repo.

Remark 7.12 (BGHKS counterexample is strictly inside the sparse class). The recent BGHKS up-to-capacity counterexample [5] uses the 2-monomial pair $(f, g) = (X^{2^b}, X^{2^{b-1}})$ on an $\mathbb{F}_2$-subspace evaluation domain in characteristic 2, achieving $\Omega(n^{\tau(1-\epsilon)})$ exceptional λ at proximity radius δ approaching the unique-decoding bound. The joint monomial support of this pair is $\{2^b, 2^{b-1}\}$, of size $2 \leq 3$ — *strictly inside the class predicted to contain the maximizer by Conjecture 7.1.* The conjecture is thus consistent with the strongest known explicit counterexample on the disproved up-to-capacity side; the multiplicative-subgroup deployment regime studied here lacks an analogous explicit dense construction.

Theorem 7.13 (General-f extension, conditional)**.** *Conditional on Conjecture 7.1, the sparse-input bound of Theorem 5.4 extends to arbitrary above-Johnson f:*

$$\varepsilon_{\text{ca}}(\mathcal{C}, \delta) \;\leq\; 10/q$$

at every rate-1/4 deployment scale, for every prime $q \geq Q_{\min} = 97$ admitting the required primitive roots. Composing with multi-round soundness (Theorem 5.10) gives $\varepsilon^{(R)}_{\mathrm{FRI}}(f) \leq 10R/|F| + (1-\delta)^t$, with Conjecture 7.1 the only remaining hypothesis on the analytic side.

Proof. For action-non-stabilised pairs (f_1, f_2), Conjecture 7.1 forces the arg max of $K(f_1, f_2; \delta)$ to have joint monomial support ≤ 3; applying Theorem 5.4 gives $K \leq 10$. For action-stabilised pairs — the only boundary class excluded by the admissibility predicate of Conjecture 7.1 — the structural argument of Remark 7.3 (action-orbit stabilization on the support; level-1 Johnson-boundary on the two-point lift) shows every such pair is either trivial ($K = 0$) or lies on the recursive Johnson boundary, hence $K \leq 10$ as well. Combining, $\varepsilon_{\mathrm{ca}}(\mathcal{C}, \delta) = \max_{(f_1,f_2)} K(f_1, f_2; \delta)/q \leq 10/q$. The multi-round extension follows from Theorem 5.10. □

Remark 7.14 (Bridging the 1D Hamming-K and 2D commit-curve K). Conjecture 7.1's $K(f_1, f_2; \delta)$ is a 1D Hamming multiplicity at level 0; Theorem 5.4's $K(f) = |V_\delta(f)|/q$ is a 2D commit-curve count at level 2. In the strict above-J admissible domain, the subdomain CA bound on the round-1 bad set $|B_1| \leq 9$ (Step 1 of Theorem 5.4) gives $K_{\mathrm{1D}} \leq |B_1|+1 \leq 10$ on every admissible adversary, so the $\varepsilon_{\mathrm{ca}} \leq 10/q$ statement holds without overhead. Parity-aligned boundary saturations (Remark 7.2) live at $\delta = \delta_J$ exact and do not propagate into strict above-J at any fold level ([7]).

8 Open Problems

We list the named open problems left by the paper. Two of them (Q1, Q2) gate the universal forms of our headline bounds; their resolution would render the corresponding layer unconditional at every deployment cell. Q1 is the deepest mathematical question the paper opens: the universal extension of Layer 2 reduces to a named non-vanishing problem in cyclotomic class field theory. Engineering and tooling items (table generators, multi-round closures, deployment-field sweeps) are recorded in the companion repository [7] and not repeated here.

Q1: universal non-vanishing of R_d on V_d^{prim}. The Layer-2 universal-h closure for the $(3k/2, 2k)$ family reduces (rigorously) to the non-vanishing $R_d := -3x_{d/2} + 2V_{d/2} + 3W_{d/2} \not\equiv 0$ on V_d^{prim} for every $d = 2^j$, $j \geq 2$ (Conjecture 4.12). Equivalently, $\mathrm{Norm}_{K_d/\mathbb{Q}}(F_d(\alpha)) \neq 0$ on the residue field K_d. The question is rigorous at $d \in \{4, 8\}$ and operational per-d at $d \in \{16, 32, 64\}$ via msolve. The universal-d statement is a precisely located non-vanishing problem; explicit class-field invariants at $d = 4$ are computed in §4 (Galois closure containing $\mathbb{Q}[\sqrt{-D}]$ with $D = 83{,}860{,}066{,}393{,}667$, class number $h = 1{,}709{,}193$). The naive Davenport–Hasse and Hecke L-value routes are infeasible: the Hilbert class field has degree $\sim 3.4 \times 10^6$. The most promising routes are (i) a structural induction step on chain self-similarity (prove $x_1 = 0 \Rightarrow x_a = 0$ for every odd a on V_d^{prim}, currently rigorous at $d \in \{4, 8\}$); or (ii) a Chevalley–Warning-style modular-reduction descent over deployment primes.

Q2: sparse-worst-case dominance for plain RS above Johnson. Conjecture 7.1 asserts that the worst K-bound adversary (f_1, f_2) on $\mathrm{RS}_{k_0}(L_0)$ above Johnson is 3-position sparse. The local no-full primitive obstruction is closed rigorously at every dyadic depth (Theorems 7.7, 7.8); the conjecture is exact-verified at deployment-shaped $(32, 8)$ over $q \in \{97, 193, 257\}$, 0 primitive counterexamples in 4,586,096 support-prime certificates, with the only known structural counterexample ruled out by two independent exclusions. The remaining step is a global attachment: from the local singleton-tail closure to the full general-f dyadic component decomposition. An above-Johnson

list-decoder (Guruswami–Sudan extension to deployment radius) at scales $(64, 16)$, $(128, 32)$ would give exact verification at deployment scale.

Q3: rate-$1/2$ universal-k Substitution-Principle lift. Theorem 5.15 is rigorous at the verified base panels $(4, 2)/(8, 4)$ and at $(16, 8)/(32, 16)$ via direct sweep, with the universal-k statement contingent on the σ-degree $\lceil\sqrt{nk}\rceil$ Substitution Principle extending uniformly across rate $1/2$. The same lift question for rate $1/8$ is open at deployment scale beyond the verified base panel $(8, 1)$.

Q4: 4-position-sparse coverage. The 4-position-sparse extension at rate $1/4$ requires a separate base-panel eliminator (Case D, four distinct mod-4 residues) not covered by Theorem 5.12. Empirical sweeps at $(32, 8)/(64, 16)$ over $q \in \{97, 193, 257\}$ are consistent with the bound $K_{4\text{-pos}} \leq 12$, but a rigorous proof analogous to Theorem 5.12 is open. The analogous bounds at rate $1/2$ and $1/8$ are also open: the natural base panel is degenerate, and the proper base is one scale up (Remark 5.19).

Q5: char-uniform Stage-2 closure at $h \geq 5$. Theorem 4.10 (b) is char-uniform at $h \in \{2, 3, 4\}$ via Singular GB / modStd. At $h \geq 5$ standard algorithms become intractable; the structural decomposition recorded in Remark 4.11 reduces the residual question to a chain-plus-two-endpoints polynomial system, verified at $h \in \{4, 5, 6, 7\}$. A char-uniform proof of $V(\text{chain} \cup \{E_{h-1}, E_{h-2}\}) = \{0\}$ at every $h \geq 2$ is the missing step. Conjecture 4.12, if resolved, sharpens this to a single-endpoint closure at $h = 2^k$.

Q6: sign-paired upper bound at $k = 2^j$, $j \geq 3$. Theorem 4.4 (iii) is rigorous at $k \in \{2, 4\}$ and conjectural at $k = 2^j$ for $j \geq 3$. The imbalanced-witness counterexample at $k = 3$ (Remark 4.6) shows the all-k form fails, so the conjecture is restricted to the 2-primary tower. The descent step itself is rigorous; the remaining gap is a forced antipodality lemma at $k = 2^j$.

9 Conclusion

What we proved. The standing open question in the proximity-gap line — a rigorous $O(1)/|F|$ FRI commit-phase soundness bound for plain Reed–Solomon above the Johnson radius — is resolved at rate $\rho = 1/4$ for 3-position-sparse adversary inputs (Theorem 5.4), unconditionally and uniformly in the deployment scale n_0. The proof passes through no correlated-agreement bound. Instead it relies on a five-line structural symmetry — the action–orbit theorem (Theorem 2.1) — on the cyclic FRI evaluation domain, a multiplicative-subgroup structure the proximity-gap line descending from BCIKS [3] never used.

What this opens for the proximity-gap programme. The substitution $z \mapsto \mu z$ is a single classical observation that has no analogue in the CA framework, because CA quantifies over arbitrary subsets of the evaluation domain. By exploiting the fact that production FRI domains are always multiplicative subgroups, we obtain bounds in the open intermediate zone $(\delta_J < \delta < 1 - \rho)$ — the regime in which prior techniques are either open in the positive direction or, in the up-to-capacity form, disproved in late 2025 [4–6]. The technique appears to be a generic tool for proximity-gap bounds on cyclic codes, not a one-off trick.

Two named open problems remain. The construction is unconditional for sparse adversary inputs at rate 1/4. Two open problems gate the universal forms:

- **Q1 (number theory)** — the universal extension to the $(3k/2, 2k)$ family reduces to the non-vanishing of an explicit norm on a class-field extension of $\mathbb{Q}[\sqrt{-D}]$, $D = 83{,}860{,}066{,}393{,}667$, with Hilbert class field of degree $\sim 3.4 \times 10^6$. The residual question is settled rigorously over $\mathbb{Q}$ at $d \in \{4, 8\}$ and connects our work to classical questions in cyclotomic class field theory and Hecke L-value non-vanishing.
- **Q2 (proximity-gap folklore)** — the general-f extension reduces to a sparse-worst-case dominance conjecture, consistent with every adversarial construction in the proximity-gap literature (ABF Lemma 6.13, Crites–Stewart, BCHKS) and exact-verified at deployment-shaped scale $(32, 8)$.

Resolving either lifts the corresponding layer to fully unconditional. Q1 is, to our knowledge, the deepest mathematical question this paper opens; we believe its closure is a single-author NT-school result waiting to be done.

Deployment consequence. At the ABF §6.3 deployment instance (128-bit security target, $t = 128$ queries), the rate-1/4 unconditional bound delivers a 79.8 **KiB STARK proof on plain Reed–Solomon**, halving the proof size against ABF's smallest deployable rigorous candidates (IRS at 161.4 KiB, FRS at 281.2 KiB), both on strictly stronger code classes (Table 1). For an Ethereum L2 currently shipping plain RS without a rigorous above-Johnson bound — Starknet, Polygon zkEVM, Plonky3-based rollups, the next generation of zkVMs — the consequence is direct: the same plain-RS pipeline now carries a real 128-bit security proof, with no protocol modification, no new hardness assumption, and a $2-3.5\times$ smaller on-chain verifier than any rigorous alternative.

Acknowledgments

The authors thank the Arnon–Boneh–Fenzi survey [2] for the explicit deployment instance that motivated this work, and Professor Guang Gong (University of Waterloo) for discussions of the sequence-school perspective on FRI soundness, which informs the action–orbit reformulation underlying this paper.

A Computational artefacts

All verification scripts, exhaustive sweep outputs, and the working notes underlying the results in this paper are maintained in the companion repository [7]. The repository's commit history (`notes/scripts/` directory and the `notes/` working notes) provides a per-result trail; in particular, every theorem and conjecture cited above has at least one Singular, `msolve`, sympy, or sage verification script with timing and parameter records, available under `https://github.com/iotexproject/rs-proximity-gaps`.

B Detailed Proofs

We collect here the proofs of the supporting theorems and lemmas of Sections 4–7, deferred from the body to keep the main narrative focused on the action–orbit theorem and its deployment-scale consequences. The proofs of the Action–Orbit Theorem itself (Theorem 2.1), of the dyadic tail lemma (Theorem 7.8), and of the Layer-3 sparse-worst corollary (Theorem 7.13) are short and appear in the main body.

B.1 Sign-paired ratios are bad (Lemma 4.3)

Proof. Factor

$$h_\rho(z) = z^{k+c}(\rho + z^{2k}) = z^c \cdot z^k(\rho + z^{2k}).$$

Note $z^{2k} = (z^k)^2$ takes values in $\{1, -1\}$ as z ranges over L_n, since $(\omega^k)^2 = \omega^{2k} = -1$ (as ω has order $n = 4k$).

Case $\rho = -1$. On the half-coset $\{z \in L_n : z^{2k} = 1\}$ (of size $n/2$), $\rho + z^{2k} = 0$, so $h_\rho(z) = 0$, agreeing with the zero polynomial.

Case $\rho = 1$. On $\{z^{2k} = -1\}$ (size $n/2$), $\rho + z^{2k} = 0$, same conclusion.

Case $\rho = \iota$. On the half-set $\{z \in L_n : z^k \in \{1, -\iota\}\}$ (size $n/2$): for $z^k = 1$, $z^{2k} = 1$, so $z^k(\rho + z^{2k}) = 1 \cdot (\iota + 1) = 1 + \iota$; for $z^k = -\iota$, $z^{2k} = -1$, so $z^k(\rho + z^{2k}) = (-\iota)(\iota - 1) = -\iota^2 + \iota = 1 + \iota$. Hence $h_\iota(z) = (1 + \iota)z^c$ identically on the half-set, a polynomial of degree $c < k$.

Case $\rho = -\iota$. Symmetric, on the half-set $\{z^k \in \{1, \iota\}\}$.

In each case h_ρ agrees with a RS_k codeword on at least $n/2$ points, so $\mathrm{dist}(h_\rho, \mathrm{RS}_k(L_n)) \leq n/2$. □

B.2 Sign-paired upper bound (Theorem 4.4)

Proof. Combine Lemma 4.3 (existence: the four ratios are bad) with the upper-bound argument below; the result is part (ii). Part (i) is the sparse-support theorem proved en route.

Throughout, write $a = k + c$, $b = a + 2k$, and decompose $L_n = L_n^+ \sqcup L_n^-$ where $L_n^\pm = \{z \in L_n : z^{2k} = \pm 1\}$. A half-set witness $(S, \rho, r(z))$ satisfies $\rho z^a + z^b \equiv r \pmod{\sigma_S}$ with $\deg r < k$ and $\sigma_S = \prod_{z_0 \in S}(z - z_0)$ of degree $2k$. Decompose $\sigma_S = \sigma_+\sigma_-$ along the partition $S = S_+ \sqcup S_-$, $S_\pm = S \cap L_n^\pm$, with $|S_\pm| = m_\pm$, $m_+ + m_- = 2k$.

Step 1: $m_+ = m_- = k$. On L_n^+ the pencil simplifies: $\rho z^a + z^b = z^a(\rho + z^{2k}) = z^a(\rho + 1)$, so $r \equiv (\rho + 1)z^a \pmod{\sigma_+}$. For $\rho \notin \{\pm 1\}$ the leading coefficient $\rho + 1$ is nonzero. Now $(\rho + 1)z^a$ has degree $a \geq k$. If $m_+ = \deg \sigma_+ > k$, then $(\rho + 1)z^a$ has degree $< m_+$, so the reduction modulo σ_+ is $(\rho + 1)z^a$ itself, of degree $a \geq k$, contradicting $\deg r < k$. Hence $m_+ \leq k$. The symmetric argument on L_n^- gives $m_- \leq k$. Combined with $m_+ + m_- = 2k$: $m_+ = m_- = k$. The boundary cases $\rho \in \{\pm 1\}$ yield $\sigma_S = z^{2k} \pm 1$ directly (a single mod-$2k$ coset), already sparse with support $\subseteq \{0, k, 2k\}$.

Steps 2–4: $\sigma_\pm$ *sparsity at* $\rho = \iota$. With $\rho = \iota$, $a = k$ (cyclic action reduces general a), the two cert equations $r - (\iota \pm 1)z^k = \sigma_\pm \cdot u_\pm$ combine to give the explicit relation

$$\sigma_-(z) = (1 + \iota)z^k - \iota\sigma_+(z). \tag{1}$$

Writing $\sigma_+(z) = z^k + L(z)$ with $\deg L < k$, the k roots of σ_+ in L_n^+ give $L(\alpha) \in \{+1, -1\}$ (from $\alpha^{2k} = 1$), and the k roots γ of σ_- in L_n^- give $L(\gamma) \in \{+1, -1\}$ (substituting (1) and $\gamma^{2k} = -1$). So L takes values in $\{+1, -1\}$ at $2k$ distinct points; since $L - 1$ and $L + 1$ each have degree $< k$ and so $< k$ roots each, L must be constant, giving $\sigma_+(z) \in \{z^k \pm 1\}$ and via (1) the corresponding sparse $\sigma_-(z) \in \{z^k \mp \iota\}$.

Conclusion. Every half-set witness has $\sigma_S = \sigma_+\sigma_-$ with both factors sparse, so σ_S has support $\subseteq \{0, k, 2k\}$, proving part (i). Substituting $\sigma_S(z) = z^{2k} + p_k z^k + p_0$ into the cert and divisor equations yields, by direct k-uniform algebra (reducing $z^{2k} \equiv -p_k z^k - p_0$, $z^{a+k} \equiv -p_k z^a - p_0 z^{a-k}$, $z^n = (z^{2k})^2$):

$$p_0 = p_k^2 + \rho \quad \text{(cert)}, \qquad p_k(2p_0 - p_k^2) = 0, \;\; p_0(p_0 - p_k^2) = 1 \quad \text{(divisor)}.$$

Substituting cert into divisor gives $p_k(p_k^2 + 2\rho) = 0$ and $\rho(p_k^2 + \rho) = 1$. Either $p_k = 0$, forcing $\rho^2 = 1$, or $p_k^2 = -2\rho$, forcing $-\rho^2 = 1$, i.e., $\rho^2 = -1$. In both cases $\rho^4 = 1$, proving part (ii). The matching $\rho = -\iota$ analysis is symmetric (the cyclic action $\rho \mapsto \omega^{a-b}\rho = \omega^{-2k}\rho = -\rho$ exchanges $\iota \leftrightarrow -\iota$). □

B.3 The $(k, 2k)$ family closure (Theorem 4.7)

Proof. With the partition $L_n = L_n^+ \sqcup L_n^-$ and notation as in Theorem 4.4, the pencil simplifies to $\rho z^k + 1$ on L_n^+ and $\rho z^k - 1$ on L_n^-. The cert equations become

$$r(z) - (\rho z^k + 1) = \sigma_+(z)u_+, \qquad r(z) - (\rho z^k - 1) = \sigma_-(z)u_-,$$

with $u_\pm$ constants (left-hand sides have degree $\le k = \deg \sigma_\pm$; the $m_+ = m_- = k$ argument from Step 1 of Theorem 4.4 applies verbatim with $(\rho \pm 1)z^k$ replacing $z^a(\rho \pm 1)$). Matching leading z^k coefficients gives $u_+ = u_- = -\rho$. Subtracting:

$$\sigma_+(z) - \sigma_-(z) = 2/\rho, \tag{2}$$

a non-zero *constant*.

Since $\sigma_+ \mid z^{2k} - 1$, write $z^{2k} - 1 = \sigma_+(z)h(z)$ for some monic h of degree k. From $\sigma_- = \sigma_+ - 2/\rho$ and the divisibility $\sigma_- \mid z^{2k} + 1 = (z^{2k} - 1) + 2 = \sigma_+ h + 2$: reduce mod σ_-, using $\sigma_+ \equiv 2/\rho \pmod{\sigma_-}$, so $\sigma_+ h + 2 \equiv (2/\rho)h + 2 \pmod{\sigma_-}$. Since $(2/\rho)h + 2$ has degree $k = \deg \sigma_-$, it equals $c' \cdot \sigma_-$ for a constant $c' = 2/\rho$ (matching leading coefficients), giving $h(z) = \sigma_+(z) - (\rho + 2/\rho)$. Substituting into $z^{2k} - 1 = \sigma_+ h$ yields the quadratic

$$\sigma_+(z)^2 - A \cdot \sigma_+(z) - (z^{2k} - 1) = 0, \qquad A := \rho + 2/\rho.$$

Solving for σ_+ as a quadratic in $\mathbb{F}_q[x]$: $\sigma_+ = (A \pm \sqrt{A^2 + 4(z^{2k} - 1)})/2$. For σ_+ to be a polynomial, the discriminant $A^2 + 4(z^{2k} - 1) = 4z^{2k} + (\rho^2 + 4/\rho^2)$ must be a perfect square in $\mathbb{F}_q[x]$. Matching against $(2z^k + c)^2 = 4z^{2k} + 4cz^k + c^2$ forces $4c = 0$ (no z^k term) and $c^2 = \rho^2 + 4/\rho^2$. Hence $c = 0$ and $\rho^2 + 4/\rho^2 = 0$, i.e., $\boxed{\rho^4 = -4}$. The corresponding sparse $\sigma_+(z) = z^k + A/2 = z^k + (\rho + 2/\rho)/2$ and $\sigma_-(z) = z^k + (\rho - 2/\rho)/2$ are read off, proving (i) and (ii). The matching existence direction (the four ratios with $\rho^4 = -4$ are bad) is established by direct construction of the half-set as $\{z : \sigma_+(z) = 0\} \sqcup \{z : \sigma_-(z) = 0\}$ on $L_n^+ \sqcup L_n^-$. □

B.4 The $(3k/2, 2k)$ cert direction (Theorem 4.9)

Proof. The certificate equation says there exists $r \in \mathbb{F}_q[z]$ with $\deg r < k$ and $r \equiv h_\rho \pmod{\sigma_S}$. Reducing $h_\rho = \rho z^{3k/2} + z^{2k}$ modulo σ_S and using $z^{2k} \equiv -\sum_{j=0}^{2k-1} p_j z^j \pmod{\sigma_S}$:

$$h_\rho \bmod \sigma_S \;=\; \rho z^{3k/2} - \sum_{j=0}^{2k-1} p_j z^j.$$

The coefficient of z^j in this reduction is $\rho - p_{3k/2}$ for $j = 3k/2$ and $-p_j$ for $j \ne 3k/2$. The degree-$< k$ requirement on r forces these to vanish for $j \in [k, 2k-1]$, giving $p_{3k/2} = \rho$ and $p_j = 0$ for $j \in [k, 2k-1] \setminus \{3k/2\}$. Since $k \ne 3k/2$ for $k \ge 1$, this includes $p_k = 0$.

For $k = 2$, $\sigma_S(z) = z^4 + \rho z^3 + p_1 z + p_0$. Imposing $z^8 \equiv 1 \pmod{\sigma_S}$ and matching coefficients of z^0, z^1, z^2, z^3 gives a closed quadratic system; eliminating p_0, p_1 yields the single equation $\rho^8 = 16$, i.e. $\rho^8 - 16 = (\rho^4 - 4)(\rho^4 + 4)$, a union of two action-orbits of size 4. The $(\rho^4 + 4)$ factor coincides with the $(k, 2k)$ orbit of Theorem 4.7; the $(\rho^4 - 4)$ factor is genuinely new at this family. Direct construction of the half-sets $S_\rho^\pm \subset L_n^\pm$ corresponding to each of the eight ratios gives the matching existence direction. □

B.5 The $(3k/2, 2k)$ locator-gap (Theorem 4.10)

Proof sketch. **(a)** Substituting Stage 2 into the five h-lattice equations $j \in \{0, h, 2h, 3h, 4h\}$ together with the cert relations $P_{2h} = 0$, $P_{3h} = \rho$, $R_{2h} = \rho^2$, $R_{3h} = -\rho$ from Theorem 4.9, one solves the resulting four-variable algebraic system in (P_0, P_h, R_0, R_h) to obtain the closed forms above. The eliminator $\rho^8 = 16$ emerges by squaring an auxiliary identity $B^3 = -2\rho$ (with $B := P_h$) against the Stage 1 solution $B = -\rho^3/2$.

(b) The annihilation conclusion $P^{(a)} = R^{(a)} = 0$ for all $a \neq 0$ is established *rigorously and char-uniformly* at $h \in \{2, 3, 4\}$ via Singular Gröbner-basis computation of the full Stage 2 ideal over $\bar{\mathbb{Q}}[\rho]/(\rho^8 - 16)$:

- $h = 2, 3$: standard `groebner` algorithm with output GB equal to the maximal ideal $\langle \alpha_c, \beta_c \rangle_{c=1}^{h-1}$.
- $h = 4$: standard `groebner` did not converge at the $h = 4$ scale; Singular's modular Gröbner basis `modStd` (which computes GB modulo many primes and lifts via rational reconstruction) returns the same maximal-ideal output in the 6-variable case.

In each case the GB reduces to the maximal ideal at the origin, so $V(I) = \{0\}$ over any field of characteristic $\nmid 8h$. Coefficient- induction pen-and-paper proofs at $h \in \{2, 3\}$ are also recorded in the companion repository [7] for cross-verification; the symbolic GB suffices for rigor.

The Singular and modStd transcripts are deferred to the companion repository [7]. □

B.6 Universal $K \leq 10$ at rate $1/4$ (Theorem 5.4)

Proof sketch (full proof in `notes/0288`). By the mod 4 pigeonhole on the support of $\hat{f}$, the four-way fold-decomposition $\text{fold}^2(\alpha_1, \alpha_2) = a + \alpha_1 b + \alpha_2 c + \alpha_1\alpha_2 d$ on L_2 has at least one of a, b, c, d identically zero. The non-degenerate cases reduce to either Pattern B ($\text{fold}^2 = \alpha_1 \cdot g(\alpha_2)$) or Reverse Pattern ($\text{fold}^2 = \alpha_2 \cdot h(\alpha_1)$) with $g(\alpha_2) = b + \alpha_2 d$ and $h(\alpha_1) = c + \alpha_1 d$ each a two-monomial pencil on L_2 (since b, d from the surviving mod-4 cell are at most single monomials on L_2 of order $n_0/4$).

For Reverse: $K_{\text{col}} = 1$ (only $\alpha_2 = 0$). Since $4 \mid k_0$ and above-J at L_0 ensures $\text{supp}(\hat{f}) \subseteq [k_0, n_0 - 1]$, Lemma 5.3 gives that the post-fold-2 pencil $h(\alpha_1) = c + \alpha_1 d$ is above-J at L_2 automatically — no auxiliary fold-level hypothesis is needed. Hence not all α_1 make $h(\alpha_1)$ bad. For $|B_1|$:

$$|B_1| \;\leq\; |\{\alpha_1 \in \mathbb{F}_q : h(\alpha_1) \text{ bad on } L_2\}|.$$

The right-hand side is bounded by the *universal 2-monomial pencil eliminator bound*:

$$\boxed{|\{\alpha \in \mathbb{F}_q^* : c + \alpha d \text{ bad on } L_{n_2}\}| \;\leq\; 8}$$

for any above-J two-monomial pencil at any deployment scale $(n_2, k_2) = (n_0/4, n_0/16)$ for $k_2 \geq 1$ a power of 2 , plus ≤ 1 for the $\alpha = 0$ branch (when $h(0) = c$ is itself bad on L_2). Hence $|B_1| \leq 9$ and $K \leq 9 + 1 = 10$. The Pattern B case is symmetric.

The universal bound in the box is itself proved by reducing every (a, b) pencil at deployment scale to a base case at $(n, k) \in \{(4, 1), (8, 2)\}$ via the *Substitution Principle* $u = z^d$, $d = \gcd(a, b, n)$, then enumerating the ≤ 12 irreducible base cases by Gröbner basis (this paper `g3_base_cases_K8_proof.py`). The maximum $|B|$ over all non-degenerate base cases equals 8, achieved by $(a, b) \in \{(2, 5), (3, 4), (5, 6), (4, 7)\}$ at $(n, k) = (8, 2)$. □

B.7 Soundness expression (Theorem 5.1)

Proof of Theorem 5.1. (a) Per-α_1 bad-α_2 bound. Fix any $\alpha_1 \in \mathbb{F}_q$ and let (a, b) be the arising pair on the level-2 domain L_2 (of order $n_2 = n_0/4$) determined by the support residues of $\hat{f}$. By

Corollary 2.3, the bad α_2 set on L_2 is the union of m orbits of the $\langle\omega_{n_2}^{b-a}\rangle$-action, each of size $n_2/\gcd(b-a, n_2)$, plus at most one additional $\alpha_2 = 0$ point (handled inside Cor 2.3's indicator $\mathbf{1}_{\alpha=0 \text{ bad}}$), plus at most one *saturating column* $\alpha_2^\star$ (present when the support residues mod the folding subgroup lie in the saturated half of L_2; absent otherwise). Conjecture 4.1 caps the orbit count m at 1 for normal pairs — unconditionally for sign-paired (Thm 4.4), $(k, 2k)$ (Thm 4.7), and $(3k/2, 2k)$ at deployment k (Thms 4.9, 4.10); conditional on the open clause of Conj 4.1 otherwise — and at $m = 2$ of size 2 for sign-paired. Hence

$$\#\{\alpha_2 \in \mathbb{F}_q : h_{\alpha_2}^{(\alpha_1)} \text{ bad on } L_2\} \;\le\; \frac{n_2}{\gcd(b-a, n_2)} + 2 \;\le\; \frac{n_0}{4} + 2,$$

where the additive 2 absorbs the $\alpha_2 = 0$ branch and the saturating-column branch.

(b) Sum over α_1. The cap of (a) is uniform in α_1: it depends only on the arising pair (a, b), which itself is determined by the DFT support of f and is therefore the same for every α_1. Summing,

$$|V_\delta(f)| \;=\; \sum_{\alpha_1 \in \mathbb{F}_q} \#\{\alpha_2 : h_{\alpha_2}^{(\alpha_1)} \text{ bad}\} \;\le\; q \cdot \Big(\frac{n_0}{4} + 2\Big),$$

hence $K(f) := |V_\delta(f)|/q \le n_0/4 + 2$.

(c) Query phase. Above the Johnson radius the level-2 fold is at agreement density at most $1 - \delta$; since the action-orbit route invokes no 1:2 CA halving (the threshold-halving of [1] originates inside the CA framework, which we bypass), the query-phase contribution to soundness is the bare $(1-\delta)^t$. Combining with (b),

$$\varepsilon_{\mathrm{FRI}}^{(2)}(f) \;\le\; \frac{K(f)}{|F|} + (1-\delta)^t \;\le\; \frac{n_0/4 + 2}{|F|} + (1-\delta)^t.$$

The unconditional clause uses the family-closure bounds of §4 inside step (a); the conditional clause replaces these with the orbit-count cap of Conjecture 4.1. □

B.8 Case C universal bound (Theorem 5.12)

Proof sketch. The 3-monomial Substitution Principle (a Substitution Principle generalization) reduces every 3-mono pencil at (n, k) to a base case at $(n', k') \in \{(4, 1), (8, 2)\}$ via $u = z^d$, $d = \gcd(a_1, a_2, a_3, n)$. SymPy Gröbner basis enumeration over the $1 + 35$ irreducible base cases (with mod 4 degeneracy filter) gives $\max \deg_\alpha \Phi = 9$, achieved at $(a_1, a_2, a_3) = (1, 4, 7)$ and $(1, 5, 6)$ at $(8, 2)$ with the explicit eliminator $\Phi = \alpha^9 - 16\alpha\rho^8$. The remaining 11 GB-TIMEOUT cases are resolved by the reflection symmetry $z \mapsto 1/z$ (which maps $(a_1, a_2, a_3) \mapsto (n-a_3, n-a_2, n-a_1)$ preserving $|B|$); 9 TIMEOUT cases are reflections of resolved cases, and the 2 self-reflective TIMEOUT cases $(2, 3, 6)$ and $(2, 5, 6)$ are resolved *symbolically over* $\mathbb{Q}$ via Singular elimination (char-uniform), with explicit degree-4 α-polynomial certificates in both cases: $|B(\alpha)| \le \deg_\alpha \Phi \le 4 < 9$ rigorously, character uniform. □

B.9 $K \le 28$ at rate $1/2$ (Theorem 5.15)

Proof sketch. Singular GB at base case $(n, k) = (8, 4)$ with σ of degree 6. For all $\binom{4}{3} = 4$ above-J triples, the eliminator ideal E in $\mathbb{Q}[\alpha_1, \alpha_2, \alpha_3]$ universally contains $E[1] = \alpha_2^8 - \alpha_3^8$ (or symmetric variant), and the full E cuts out a degree-28 projective variety: $\mathrm{vdim}(E + \langle\alpha_3 - 1\rangle) + K_{2\text{-mono}}(a_1, a_2) = 28$ across all four triples. The $K = 28 = \binom{8}{6}$ matches the count of agreement subsets of size $J = 6$ in L_8, with one close pencil per subset (no cancellation; contrast rate-1/4 where $K \le 10 \ll \binom{8}{4} = 70$).

Verification at deployment scale $(n, k) = (16, 8)$: `vdim` sweep over all $\binom{8}{3} = 56$ above-J triples gives vdim ≤ 28 universally on the 54 non-degenerate triples. Two at-J degenerate triples $(10, 13, 14)$ and $(11, 12, 14)$ give 1-dim varieties (continuous families of close pencils), excluded by the above-J hypothesis exactly as $(4, 5)$ at $(8, 2)$ in Theorem 5.4's base. Substitution Principle confirms gcd-reducible $(16, 8)$ triples match base K exactly. Six coprime $(16, 8)$ triples also achieve $K = 28$ via a deeper symmetry mechanism beyond gcd-substitution; characterizing this orbit is open but does not affect the bound at $(16, 8)$. The extension to all $k \geq 4$ (a power of 2) is via the Substitution Principle lift: at every $(n, k) = (2k, k)$ with $k > 8$, a non-trivial gcd > 1 in the support exponents reduces to a smaller scale, and the gcd-coprime cases are conjectured to follow the same six-orbit pattern observed at $(16, 8)$. A rigorous lift for arbitrary k remains open (Q3, §8); the present theorem is rigorous at base panels $(4, 2)$ and $(8, 4)$, and verified by direct sweep at $(16, 8)$ and $(32, 16)$ ([7]). □

B.10 $K \leq 56$ at rate $1/8$ (Theorem 5.18)

Proof sketch. Singular GB at base $(n, k) = (8, 1)$ with σ of degree $\lceil\sqrt{8}\rceil = 3$. `vdim` sweep over all $\binom{7}{3} = 35$ above-J triples gives max non-degenerate $K = 56$, achieved by 8 saturating triples: $(1, 2, 3)$, $(1, 2, 7)$, $(1, 3, 6)$, $(1, 6, 7)$, $(2, 3, 5)$, $(2, 5, 7)$, $(3, 5, 6)$, $(5, 6, 7)$. Two at-J degenerate triples $(1, 4, 5)$ and $(3, 4, 7)$ excluded by above-J hypothesis. For 2-monomial sweep over 18 above-J coprime pairs: max $K = 2$ achieved by $(1, 5)$ and $(3, 7)$. The bound is rigorous at the base panel $(8, 1)$. The Substitution Principle lift to deployment scale $(n, k) = (8k, k)$, $k \geq 1$ a power of 2, is by analogy to the verification at rate $1/2$ deployment $(16, 8)$ and is currently *not* a rigorous proof; the deployment-scale sweep $(16, 2)/(32, 4)/(64, 8)$ remains open. □

B.11 No-full primitive base-panel closure (Theorem 7.7)

Proof sketch. Split the four quarter blocks into the two α_2-sides $u = \{j \bmod 4 \in \{0, 1\}\}$ and $v = \{j \bmod 4 \in \{2, 3\}\}$. If the 3-support is side-pure, one residual row is identically zero so $\operatorname{rank} W(\alpha) \leq 1$. Otherwise, after the $\alpha_1 = 0$ quotient, one side reduces to a singleton monomial $c \cdot x^e$ with $4 \leq e < 16$; saturation on S would force $\mathrm{tail}_S(x^e) = 0$. By the cyclotomic tail lemma over $\mathbb{Z}[\zeta_{16}]$ (exact classification of all 8-subsets $S \subset \mu_{16}$ with collapsed Vandermonde rank), the only S with $\mathrm{tail}_S(x^e) = 0$ are the two parity halves, each containing two full quarter blocks — contradicting no-full. Three independent certificate trails (multi-prime symbolic, Singular GB saturation, cyclotomic classification) confirm. Full proof: [7] Notes 0330–0333. □

References

[1] R. Chai and X. Fan. *FRI Soundness Above the Johnson Bound via Threshold Halving.* Cryptology ePrint Archive, 2026 (submitted; ID forthcoming).

[2] 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`.

[3] E. Ben-Sasson, D. Carmon, Y. Ishai, S. Kopparty, and S. Saraf. *Proximity gaps for Reed–Solomon codes.* J. ACM, 2023.

[4] A. Crites and B. Stewart. *On Reed–Solomon proximity gaps conjectures.* Cryptology ePrint Archive, Paper 2025/2046, 2025.

[5] E. Ben-Sasson, D. Carmon, U. Haböck, S. Kopparty, S. Saraf. *On proximity gaps for Reed–Solomon codes.* Cryptology ePrint Archive, Paper 2025/2055, 2025.

[6] B. E. Diamond, A. Gruen. *Counterexamples to Reed–Solomon proximity gap conjectures beyond the Johnson radius.* Cryptology ePrint Archive, Paper 2025/2078, 2025.

[7] R. Chai, X. Fan, *Computational companion: notes, scripts, and verification outputs for the action-orbit FRI soundness analysis*, `https://github.com/iotexproject/rs-proximity-gaps`, 2026.