Chosen Commitments and Certified Decoding Regions in the Adaptive Threshold Schnorr Attack

Hüseyin Eren Demirtaş and Oğuz Yayla

Middle East Technical University

Abstract. Crites and Stewart reduce an adaptive forgery attack on threshold Schnorr signatures with N parties and threshold τ to one call to a sparse linear search oracle, which they show to be bounded distance decoding of a diagonally scaled Reed–Solomon code and whose sparse answer, a witness, names the parties to corrupt. Their analysis samples the coefficient vector behind the adversary's commitment uniformly at random and leaves open whether a chosen vector does better. We bound the success of an adversary that follows their template with a chosen vector and recovers witnesses only inside regions built from the radii that seven decoding theorems, three for Reed–Solomon and four for folded Reed–Solomon codes, certify under their own setup hypotheses. When the commitment is fixed through a known coefficient vector before a fresh hash challenge, the attack target is, conditional on the history that fixed the vector, a uniform point of a line. For distinct nonzero Shamir points, our main theorem shows that every such line contains at most M targets with a witness in a finite collection of coordinate and block weight regions fixed before the targets, each permitting fewer than τ nonzero coordinates, for an explicit integer M depending only on N, τ and the collection. We call the adversary the wrapper when it makes at most Q fresh challenges and no signing query, chooses each coefficient vector from its history before the challenge and keeps it afterwards, and lets its solvers make no further signature hash queries. It accepts only a witness of the same target and representation that lies in the regions of the collection within its corruption cap, corrupts parties only once it holds such a witness, and then forges from that witness and stops. Hence the wrapper succeeds with probability at most QM/p, with p the group order, even if its solvers recover every witness in the regions at no cost. On the six pairs (N, τ) of the original attack, fixed collections built from partitions into consecutive blocks contain all seven certified regions and have M below 2^{24}, so at the field size proxy $p = 2^{252}$ the bound stays below 2^{-100} after 2^{128} challenges. For uniform targets we prove much smaller bounds on the witness probability in these regions. For the folded Guruswami–Wang and Ashvinkumar–Habib–Srivastava block regions on partitions into consecutive blocks fixed before the target, this probability is negligible for polynomially many parties, exponentially large fields and τ/N bounded away from zero. These results also classify the studied decoder theorems. The hardness of the oracle, the success of other strategies and the adaptive security of threshold Schnorr schemes are open.

Keywords: threshold signatures · Schnorr signatures · adaptive security · Reed–Solomon codes · list decoding · folded Reed–Solomon codes

1 Introduction

Consider a threshold Schnorr signature scheme with N parties and signing threshold τ, whose secret key is shared with Shamir's scheme over a prime field $\mathbb{F}_p$. Crites and Stewart

E-mail: eren.demirtas@metu.edu.tr (Hüseyin Eren Demirtaş), oguz@metu.edu.tr (Oğuz Yayla)

This work is licensed under a "CC BY 4.0" license.
Date of this document: 2026-10-24.

show that an adaptive adversary against such schemes, a class that includes FROST and its variants, can forge a signature with one call to a sparse linear search oracle [CS25, Theorems 1 and 2 and Section 2.1]. The public Shamir points determine the $\tau \times N$ Vandermonde matrix H. The adversary commits to a vector $\alpha \in \mathbb{F}_p^\tau$, receives a hash challenge $c \in \mathbb{F}_p$, and asks the oracle for a vector $a \in \mathbb{F}_p^N$ with

$$Ha = \alpha + cK_0, \qquad \mathrm{wt}(a) \le f,$$

where $K_0 = (1, 0, \ldots, 0)^\top$ and f is the number of parties it may corrupt. The support of a names the parties to corrupt, and their secret shares complete the forgery. Crites and Stewart also prove that this oracle is exactly bounded distance decoding of a diagonally scaled Reed–Solomon code of length N, dimension $k = N - \tau$, and rate $\beta = k/N$ [CS25, Theorem 3, Lemmas 4 and 5, and Corollary 2]. Under the hypotheses of theorem 2.3, a probabilistic polynomial time solver with non-negligible witness recovery probability in the original uniform target experiment yields an adaptive forgery, so the decoding question has direct cryptographic meaning.

Crites and Stewart draw α uniformly, which makes the target $v = \alpha + cK_0$ uniform in $\mathbb{F}_p^\tau$. The attack is then useful only at weights where a uniform target has a witness with non-negligible probability. A first moment bound, obtained by counting the expected number of size f supports whose columns span the target, shows that witness probability at least $2^{-\kappa}$ forces the gap $g = \tau - f$ to satisfy $g \le (N + \kappa)/\log_2 p$ (lemma 7.5). Numerical evaluations replace p by the field size proxy 2^{252}. At the lengths considered by Crites and Stewart and $p \approx 2^{252}$, the useful weights therefore lie within a few units of the capacity radius $\tau = N - k$. Crites and Stewart observe that their attack region corresponds roughly to list decoding capacity, whereas efficient decoding of ordinary Reed–Solomon codes is known only up to the Johnson bound, and its feasibility beyond that bound is open [CS25, Section 7]. For the largest pair in their table, $(N, \tau) = (1024, 683)$ with $p \approx 2^{252}$, one call to an oracle solving the sparse search problem gives success with probability close to one at $f \in \{680, 681, 682\}$ [CS25, Table 1]. At this pair the unique decoding radius of Berlekamp–Welch is 341, and the Guruswami–Sudan list decoder certifies 433 errors. The largest budget certified by the studied folded decoders covers at most 512 coordinates, even for an optimally packed support.

The source leaves two points of this picture open. The first is the solver, since Crites and Stewart give no efficient algorithm for the oracle at the useful weights and the known hardness results for Reed–Solomon decoding stop short of this setting (section 1.1.3). The second is the choice of α, whose uniform sampling simplifies the analysis, and Crites and Stewart remark that an adversary who does not choose α at random might succeed with higher probability or obtain instances of a special form [CS25, Remark 6]. An adversary can moreover repeat the experiment, because a failed attempt uses one fresh signature hash query and no corruption, in addition to its solver computation. We therefore ask the following question.

> With what probability does an adversary succeed that makes Q fresh challenges, chooses every α freely before its challenge, recovers witnesses only where current Reed–Solomon decoder theorems certify recovery, and corrupts parties only once it holds such a witness?

Our answer rests on a counting theorem. Once the adversary has fixed its coefficient vector, its target is a uniform point of a line with p points. At most an explicit integer M_{line} of these points admit a witness in decoder regions that are declared before the targets and permit fewer corruptions than the signing threshold. We call such an adversary the wrapper when it follows the rules listed below, and the wrapper succeeds within Q challenges with probability at most $Q \cdot M_{\mathrm{line}}/p$, whatever vectors it chooses before its

challenges. On the six parameter pairs of the original attack, M_{line} stays below 2^{24} for seven decoder theorems. In the original uniform experiment the same regions are hit far less often, and we prove much smaller bounds there, including one at every polynomial length for the Guruswami–Wang and Ashvinkumar–Habib–Srivastava block regions of aligned partitions fixed before the target.

We study seven decoder routes. The coordinate metric routes are Berlekamp–Welch [Bla08, Sections 3.11 and 3.12], heterogeneous interleaved Reed–Solomon decoding [Yay11], and Guruswami–Sudan [GS99]. The folded routes are the decoders of Guruswami and Rudra [GR08], of Guruswami and Wang [GW13], of Kopparty, Ron-Zewi, Saraf, and Wootters (KRZSW) [KRZSW23], and of Ashvinkumar, Habib, and Srivastava (AHS) [AHS26]. Each theorem certifies recovery inside a region of error patterns, under its own setup, output and success hypotheses. The interleaved guarantee is distributional and holds with high probability over random errors, and the Guruswami–Rudra route enters our menu through an optimistic upper bound on its certified budgets. For the coordinate routes the region bounds the number of erroneous coordinates, and for the folded routes it bounds the number of erroneous blocks of a fixed partition of the coordinates. We grant the adversary perfect and free recovery inside all of these regions, that is, every witness there at no cost, and we bound how often a target has a witness there at all.

Our main result is a statement about lines. Once the history has fixed α, the fresh challenge is uniform conditional on that history, so the target is a uniform point of the line $\alpha + \mathbb{F}_p K_0$ (lemma 3.2). For pairwise distinct nonzero Shamir points, the direction K_0 lies in no span of fewer than τ columns of H (lemma 3.3). Witnesses for two different points of one line must therefore occupy at least τ coordinates together, and an elementary bound on families of sets with small pairwise intersections (lemma 3.4) limits how many points of a line can have a witness of small weight. Theorem 3.5 makes this precise. Let $\mathcal{M}$ be a finite menu, fixed independently of the targets, that consists of a coordinate region of weight below τ and of block regions that each cover fewer than τ coordinates. Let $G_{\mathcal{M}}$ be the set of targets with a witness in the menu. Then

$$\left|G_{\mathcal{M}} \cap (\alpha + \mathbb{F}_p K_0)\right| \le M_{\mathrm{line}}(\mathcal{M}) \qquad \text{for every } \alpha \in \mathbb{F}_p^{\tau},$$

where $M_{\mathrm{line}}(\mathcal{M})$ is an explicit integer that depends on N, τ, and the menu, but not on p. Corollary 3.7 turns the count into a probability. The wrapper makes at most Q fresh challenges and no signing query, commits through a known coefficient vector α chosen from its complete history, fixes that vector before the challenge and keeps it afterwards, and lets its solvers run without further signature hash queries. It corrupts parties only once it holds an accepted witness, and it forges only from witnesses in the menu that respect its corruption cap, where the partitions of the menu are fixed before the target. It succeeds with probability at most

$$\min\left\{1,\ \frac{Q \cdot M_{\mathrm{line}}(\mathcal{M})}{p}\right\}.$$

The proof conditions on the history of each challenge and takes a union bound, so the attempts may depend on one another and the solver may run for any time.

For one fresh challenge, the upper bound $M_{\mathrm{line}}(\mathcal{M})/p$ is $M_{\mathrm{line}}(\mathcal{M})$ times the exact $1/p$ success probability of plain guessing. Whether this factor is necessary and whether any strategy attains it are open questions. Guessing is itself a strategy of the wrapper and succeeds with probability exactly $1/p$ per challenge. The adversary fixes a vector a in the menu with $\mathrm{wt}(a) \le f$, guesses a value c', and commits to $\alpha = Ha - c'K_0$, so that a is a witness exactly when the challenge equals c'. The same observation explains why the bound is much larger than in the uniform experiment. There each span of codimension

d_{span} contains the target with probability $p^{-d_{\mathrm{span}}}$, so every coordinate by which a region falls short of the threshold costs a factor of p. A line through a known witness always contains one covered point, so no bound that holds for every line can fall below $1/p$ per challenge. The theorem shows that the number of covered points nevertheless stays below a quantity independent of the field.

We evaluate the theorem for the seven routes on the six parameter pairs of the original attack (section 6.4). The menu contains the Guruswami–Sudan coordinate region, which includes the Berlekamp–Welch, interleaved, and KRZSW regions on these pairs. For every divisor m of N it also contains the aligned partition into blocks of m consecutive coordinates with the largest Guruswami–Wang budget, whose region contains every AHS region. At each divisor where the optimistic upper bound on the Guruswami–Rudra budgets described above is nonnegative, the same partition appears once more with that bound as its budget. The adversary may search this whole menu after it sees the challenge, because the count covers the union. The count uses only the Shamir points, whereas reading a block region as the region of a folded decoder also requires the ordered evaluation domain of that decoder. The counts are

$$M_{\mathrm{line}} = 1083, \quad 815, \quad 331, \quad 3281, \quad 85955, \quad 10287538$$

for $(N, \tau) = (64, 43)$, $(128, 86)$, $(196, 131)$, $(512, 342)$, $(768, 513)$, and $(1024, 683)$, and all of them are below 2^{24}. For a prime p the bound is $Q \cdot M_{\mathrm{line}}/p$. Evaluated at the field size proxy 2^{252}, it is below 2^{-228} per challenge and below 2^{-100} after 2^{128} challenges on all six pairs.

For the uniform experiment of Crites and Stewart the same regions are hit far less often, and we prove the corresponding bounds as supporting results. A repetition bound shows that Q uniform trials succeed with probability at most Q times the one trial bound (corollary 7.4). For the seven routes on the six pairs the one trial bounds are below 2^{-2770} at the proxy. The first moment bound also gives a necessary condition on the gap for a target success probability. For success probability one half within 2^{128} uniform trials, no positive gap satisfies it on the two smallest pairs. On the other four the largest gap that satisfies it is, as a fraction of N, more than seven times smaller than the least real deficit of the common Guruswami–Wang and AHS radius. For that common radius we prove a statement at all lengths, namely that the block partition of the decoder, fixed independently of the target, admits a witness inside the certified budget with probability at most $\binom{L}{B_{\mathfrak{D}}} p^{-d_{\mathfrak{D}}}$ for $\mathfrak{D} \in \{\mathrm{GW}, \mathrm{AHS}\}$ (theorem 7.10). Here L is the number of blocks, $B_{\mathfrak{D}}$ the route's certified budget, and $d_{\mathfrak{D}} = \tau - mB_{\mathfrak{D}}$ the number of coordinates by which the packed budget falls short of τ. For a family with N polynomial in the security parameter λ, rate bounded away from one, and $\log_2 p$ at least linear in λ, this probability is negligible over the aligned partitions of all divisors of N, at every polynomial length (corollary 7.11). The certified radii of the three coordinate routes stay below $N - \sqrt{N(k-1)}$. In growing fixed rate families with $N \le p$ and $\kappa = o(N \log_2 p)$, every weight with witness probability at least $2^{-\kappa}$ eventually lies above all three radii, with a margin of at least $(\sqrt{\beta} - \beta)N - o(N)$ (proposition 7.7). In an unbounded fixed rate family with $\log_2 p \to \infty$ and $\kappa = \mathcal{O}(\log_2 p)$, witness probability at least $2^{-\kappa}$ and exact weight coverage by the common Guruswami–Wang and AHS radius, even with optimal packing, can hold together only if $N \ge 4\beta(1-\beta)(\log_2 p)^2(1 - o(1))$ (theorem 7.8 and corollary 7.9). Finally, on all 24 tested rows of the attack paper, no theorem admissible parameter choice of any route covers a tested weight (section 8).

These statements bound one strategy and classify the studied theorems. An adversary that departs from the rules of definition 3.1, for instance by accepting a witness outside the menu, lies outside corollary 3.7. The hardness of the oracle, the success of such adversaries and the adaptive security of threshold Schnorr schemes are open questions.

The inherited attack, its reduction to Reed–Solomon decoding, the cited decoder theorems, the qualitative capacity observation of Crites and Stewart, and the elementary counting and union bound arguments on which the main theorem rests are known material, and we use them as such.

1.1 Related work

1.1.1 Adaptive security of threshold Schnorr signatures

The adaptive security landscape is complementary to the decoder question. The analysis of FROST isolates the related low dimensional vector representation problem and gives security results under parameter dependent hardness assumptions. The key unique framework gives broader necessity and impossibility results [CKK+25, CCKM26]. Subsequent constructions such as FaFROST+ and ms-FROST change the protocol interface or model assumptions to obtain adaptive security results [BGL+25, BCL+26]. The original FROST protocol and its specification in RFC 9591 provide protocol context only [KG20, CKGW24].

1.1.2 Decoder literature and theorem types

The coordinate metric routes studied here draw on classical unique decoding, probabilistic interleaved decoding, and Guruswami–Sudan (GS) list decoding [BKY03, SSB09, Yay11, GS99, Gur06]. Guruswami and Rudra showed that folded Reed–Solomon (FRS) codes give an explicit route to list decoding capacity [GR08]. Guruswami–Wang give a linear algebraic decoder that returns a low dimensional affine candidate space for arbitrary permitted short lengths [GW13]. Later work improves folded RS list recovery, combinatorial list size, and complete algorithmic runtime under different parameter conditions [KRZSW23, Tam24, Sri25, CZ25, AHS26]. In particular, AHS give both a deterministic procedure for pruning an affine candidate space in a general linear code and a fully polynomial randomized FRS decoder [AHS26, Theorems 3.5 and 4.4].

These theorem types answer different parts of the oracle and attack criteria of section 4.3. A radius formula, a random error failure bound, a combinatorial list size bound, and a complete decoder runtime therefore each settle a different question, and we read every certified radius together with the setup, metric, list enumeration, runtime, and attack distribution hypotheses of its theorem.

1.1.3 Hardness and random error results for Reed–Solomon decoding

Two lines of work bound what a decoder for the attack regime could look like, and we record how far each of them reaches. Guruswami and Vardy prove that maximum likelihood decoding of Reed–Solomon codes is NP-hard [GV05, Section 1]. Their Section 3 gives an explicit sequence of codes that are fixed in advance and may be preprocessed without bound, but that admit no polynomial size family of maximum likelihood decoding circuits unless NP is contained in P/poly [GV05, Section 3]. Their hardness therefore persists for some codes whose evaluation points are fixed. The original uniform attack experiment of Crites and Stewart fixes the public Shamir points in advance and then draws the target uniformly over $\mathbb{F}_p^\tau$, by lemma 7.2. The Guruswami–Vardy result is about worst case received words, so it leaves the average case question of that experiment open. Crites and Stewart make this distinction for their uniform attack distribution [CS25, Section 7].

In their extended abstract, Cheng and Wan study the radius at which the Reed–Solomon first moment crosses one. In their field size notation q, let $\hat{g}(n, k, q)$ denote the least integer t with $\binom{n}{t}q^{k-t} < 1$. With $n = N$, $q = p$ and $k = N - \tau$, the substitution $t = N - f$ gives $\binom{n}{t}q^{k-t} = \binom{N}{f}p^{f-\tau}$, which is the support span first moment of section 4.5. Their radius $n - \hat{g}$ is therefore the integer threshold defined by this strict first moment inequality.

The positive exact list exponents reported in section 8 are precisely the tested rows lying strictly beyond that radius. They prove that a list decoder at radius $n - \hat{g}$ running in time $q^{\mathcal{O}(1)}$ yields a discrete logarithm algorithm over $\mathbb{F}_{q^{\hat{g}-k}}$ in time $q^{\mathcal{O}(1)}$. They also prove that bounded distance decoding of $[q,\, 3h+4]_q$ at radius $q - 4h - 4$ with $q > (h+2)^4$ yields a discrete logarithm algorithm over $\mathbb{F}_{q^h}$ [CW07, Theorems 1 and 2].

Both statements stop short of the present setting, for two reasons. First, both measure running time as $q^{\mathcal{O}(1)}$, that is, polynomial in the field cardinality. Under the bit model of section 4.4, where $\log_2 p = \Theta(\lambda)$ and $N = \mathrm{poly}(\lambda)$, a $q^{\mathcal{O}(1)}$ bound is exponential in the security parameter, and the conclusion is already satisfied by generic square root discrete logarithm methods. Second, their bounded distance reduction, which is the direction matching $P_{H,f}$, treats codes of length $n = q$ and dimension $3h + 4$ with $h < q^{1/4} - 2$, so its rate tends to zero and its length equals the field size. The attack rows instead have $N \ll p$ at rate close to $1/3$. The one witness question for a constant rate code with $N \ll p$ therefore lies outside both reductions, and this is the precise sense in which section 9 records it as open.

A third line, due to Rudra and Uurtamo, treats random rather than worst case errors. They show that for a code of relative distance Δ and a $(\Delta - \gamma)$ fraction of random errors, the transmitted codeword is with high probability the only one in the ball, provided the alphabet satisfies $q \geq 2^{\Omega(1/\gamma)}$. They also show that this alphabet condition cannot in general be weakened to $q = 2^{o(1/\gamma)}$ [RU10, Theorem 1 and Section 3.2]. Their accompanying Reed–Solomon enumeration procedure applies only to error weights at most $n - 4k$ [RU10, Corollary 2 and Section 3.1]. The inequality $4k > N$ holds on all six publication parameter pairs, so that procedure is vacuous there. More fundamentally, the random error model presupposes a distinguished transmitted codeword, whereas the original uniform attack experiment samples a uniform target and treats the nearby codeword as an output of the search. The objection raised in section 1.1.4 for the semi-adversarial channel model therefore applies verbatim to the random error model.

1.1.4 Distributional and structural transfer

Brakensiek, Chen, Dhar, and Zhang study a channel $\Psi : C \to \Delta(\Sigma^n)$ that starts from a distinguished transmitted codeword c. For each c, their semi-adversarial model fixes non-adversarial coordinates I, agreement coordinates $J \subseteq I$, and adversarial values outside I, while the other coordinates $I \setminus J$ are uniform. They give high probability, near linear time unique decoders for IRS and suitable folded RS codes under explicit tradeoffs between total and adversarial errors [BCDZ26, Definition 1.7 and Theorems 1.13 and 1.14].

The original uniform attack experiment follows a different order, since it first samples a uniform syndrome $v^\star$, chooses a representative u of its syndrome coset and then asks for any pair (c, a) with $c = u - a \in \ker(H)$ and $\mathrm{wt}(a) \leq f$. The nearby codeword and its support are therefore outputs of the search, and many support events can overlap. The supportwise construction in lemma 4.7 proves the particular random codeword and random error statement needed for the IRS theorem used here. Whether the union of attack witnesses, the virtual interleaving construction, or the folded block geometry obeys the semi-adversarial channel model above is a further question. Applying those new decoders would require a transfer theorem that preserves the code structure, the channel inequalities, the conditioning, and the complete probability composition.

2 Setting and attack interface

This section fixes the algebraic setting shared by the geometric theorem and its cryptographic wrapper. We first record the Shamir data and the sparse relation, then the

inherited threshold Schnorr interface and the finite menus of witness regions whose covered targets the main theorem counts, and we close with the solver to forgery interface.

2.1 Shamir points and the sparse relation

Fix a prime p, integers $1 \le \tau \le N < p$, and pairwise distinct nonzero public Shamir points $z_1, z_2, \ldots, z_N \in \mathbb{F}_p$. Let

$$K_i = (1, z_i, \ldots, z_i^{\tau-1})^\top, \qquad K_0 = (1, 0, \ldots, 0)^\top,$$
$$H = [K_1\, K_2 \cdots K_N], \qquad k = N - \tau.$$

Any τ columns of H form a Vandermonde matrix on pairwise distinct points and are therefore linearly independent, that is, H is full spark. The additional requirement that every z_i is nonzero is the usual convention of Shamir's scheme [Sha79], because the shared secret is the value of the sharing polynomial at 0, and it separates K_0 from every span of fewer than τ columns in lemma 3.3. For $a \in \mathbb{F}_p^N$, the support $\mathrm{supp}(a)$ is the set of indices of its nonzero coordinates and its weight is $\mathrm{wt}(a) = |\mathrm{supp}(a)|$. For a target $v \in \mathbb{F}_p^\tau$ and an integer $0 \le f \le \tau$, the sparse relation is

$$P_{H,f}(v): \quad \text{find } a \in \mathbb{F}_p^N \text{ such that } Ha = v,\ \mathrm{wt}(a) \le f. \tag{2.1}$$

The source Problem P asks for an f element set whose columns span v [CS25, Definition 1]. The two formulations are equivalent, because coefficients on such a set form a witness and, conversely, the support of any witness can be enlarged to size f by assigning zero coefficients to the added positions.

2.2 Threshold Schnorr interface and coefficient commitments

The cryptographic objects are those of the Crites–Stewart attack [CS25, Definitions 2 and 10, the Claim after Theorem 1, and the proofs of Theorems 1 and 2] on threshold versions of the Schnorr signature scheme [Sch91]. Assume a prime order group $\mathbb{G} = \langle g \rangle$ of order p, a signature hash $\mathcal{H}_{\mathsf{sig}}$ with values in $\mathbb{F}_p$ that is modeled as a random oracle [BR93], public verification shares, adaptive corruption access to secret shares, and Schnorr verification

$$g^s = R \cdot PK^{\mathcal{H}_{\mathsf{sig}}(R,\, PK,\, M)}$$

for a signature (R, s) on a message M. The hash input, the sign convention and the verification equation are those of the Crites–Stewart interface. These are standing scheme and model assumptions.

Let GoodKG be the event that the sampled key material is Shamir shared, that is, there is a polynomial

$$q(Z) = \sum_{j=0}^{\tau-1} x_j \cdot Z^j$$

with coefficients in $\mathbb{F}_p$ such that $sk_i = q(z_i)$ and $PK_i = g^{q(z_i)}$ for $1 \le i \le N$, and $PK = g^{x_0}$. For a scheme whose key generation or distributed key generation has this form, $\Pr[\mathsf{GoodKG}] = 1$. Write $x = (x_0, x_1, \ldots, x_{\tau-1})^\top$, so that $sk_i = K_i^\top x$ for every party and $x_0 = K_0^\top x$. The public coefficient commitments are $X_j = g^{x_j}$ for $0 \le j < \tau$, and in particular $X_0 = PK$. The recovery of these commitments in the next lemma and the forgery construction built on it belong to Crites and Stewart, and we use them as the inherited interface for the wrapper of section 3.1.

Lemma 2.1 (Public coefficient recovery)**.** *On* GoodKG*, all* $X_j = g^{x_j}$*,* $0 \le j < \tau$*, are computable in polynomial time from any* τ *public verification shares, without a discrete logarithm.*

Proof. For distinct indices $T = (i_1, i_2, \dots, i_\tau)$, let

$$V_T = \big(z_{i_\ell}^{\,j}\big)_{\substack{1 \le \ell \le \tau \\ 0 \le j < \tau}}.$$

The Vandermonde matrix V_T is invertible because the points are pairwise distinct. Since $V_T x = (sk_{i_1}, sk_{i_2}, \dots, sk_{i_\tau})^\top$, every coefficient is the known linear combination $x_j = \sum_{\ell=1}^{\tau} (V_T^{-1})_{j,\ell} \cdot sk_{i_\ell}$ of the shares, and exponentiation gives

$$X_j = \prod_{\ell=1}^{\tau} PK_{i_\ell}^{(V_T^{-1})_{j,\ell}}.$$

This uses polynomially many field operations and group exponentiations in τ and $\log p$. □

2.3 Finite menus of witness regions

The geometric theorem counts targets that admit a witness in a declared finite union of coordinate and block regions. We make these regions precise before any decoder enters, because the counting argument uses only the Shamir data above.

A partition $\mathcal{P}$ of the coordinate set $[N] = \{1, 2, \dots, N\}$ into $L = N/m$ blocks of equal size m, where $m \mid N$, defines the block weight

$$\mathrm{wt}_{\mathcal{P}}(a) \coloneqq \#\{\, \mathcal{B} \in \mathcal{P} : \ \mathcal{B} \cap \mathrm{supp}(a) \neq \varnothing \,\}, \qquad a \in \mathbb{F}_p^N.$$

By definition, a block counts once however many of its coordinates lie in the support. Consequently $\mathrm{wt}(a) \le m \cdot \mathrm{wt}_{\mathcal{P}}(a)$, so the block weight bound $\mathrm{wt}_{\mathcal{P}}(a) \le B$ implies the coordinate weight bound $\mathrm{wt}(a) \le mB$.

Definition 2.2 (Finite menu and its witness region)**.** A finite menu $\mathcal{M}$ consists of an integer $0 \le F < \tau$ and a finite indexed collection $(\mathcal{P}_\iota, B_\iota)_{\iota \in I}$, where each $\mathcal{P}_\iota$ partitions $[N]$ into $L_\iota = N/m_\iota$ blocks of equal size m_ι, and where each integer budget satisfies $B_\iota \ge 0$ and $m_\iota B_\iota < \tau$. The menu is fixed independently of every target. Its witness region and its target image are

$$\mathcal{W}_{\mathcal{M}} \coloneqq \{a \in \mathbb{F}_p^N : \ \mathrm{wt}(a) \le F\} \ \cup \bigcup_{\iota \in I} \{a \in \mathbb{F}_p^N : \ \mathrm{wt}_{\mathcal{P}_\iota}(a) \le B_\iota\},$$
$$G_{\mathcal{M}} \coloneqq \{Ha : \ a \in \mathcal{W}_{\mathcal{M}}\}.$$

Thus $G_{\mathcal{M}}$ consists of the targets v admitting $v = Ha$ for some $a \in \mathcal{W}_{\mathcal{M}}$, and the region is a finite union of one coordinate ball and finitely many block balls. The region is defined by the budgets of the menu alone. The corruption cap f of equation (2.1) enters through the wrapper of section 3.1, which accepts a witness only if $\mathrm{wt}(a) \le f \le \tau - 1$. Dropping that filter enlarges an event, which is legitimate for an upper bound.

The menu is a purely combinatorial object, whose definition uses only the coordinate set $[N]$, the coordinate budget F, the partitions and their block budgets. When a block region is later read as the certified region of a folded decoder theorem, that source's conditional setup requirements, including its ordered power domain, are hypotheses of that reading.

2.4 Solver to forgery interface

The decoding questions of this paper have cryptographic meaning through the forgery construction of Crites and Stewart, which uses the public coefficient commitments above. We state that construction once, in the form that its solver call needs, as the interface between the sparse relation of equation (2.1) and the adaptive security of the scheme.

Throughout this subsection we keep the standing scheme and model assumptions of section 2.2, that is, the prime order group, the signature hash modeled as a random oracle, the public verification shares, the adaptive corruption access and Schnorr verification. We fix a corruption bound $0 \le f \le \tau - 1$. The security notion is adaptive TS-UF-0 security, in which an adversary receives the public key and the public verification shares, may make signing queries and may corrupt at most f parties adaptively. It wins when it outputs a valid signature on a message that was never submitted to a signing query [CS25, Definition 8 and Figure 2], [CKK+25, Definition 3, Figure 3, and Definition 4]. The event GoodKG of section 2.2, which says that the sampled key material has the Shamir form, is the one sampled event of this setting, and its probability is one whenever key generation guarantees that form. It describes secret key material hidden from the adversary of appendix A.1, so this adversary takes no decision on GoodKG, and the event enters only through the joint validity probability of theorem 2.3.

A solver is a probabilistic algorithm $\mathcal{S}(H,\, v,\, f)$ that returns a vector $a \in \mathbb{F}_p^N$ or the failure symbol $\perp$. Its output is valid for the target v when the predicate

$$\mathsf{Valid}_{H,f}(v,\, a) \coloneqq \left[\, a \in \mathbb{F}_p^N,\ Ha = v,\ \mathrm{wt}(a) \le f \,\right]$$

holds, so the failure symbol is never valid. Everything that turns a target into one output belongs to the solver, that is, decoding, the production or pruning of a complete list, filtering, the selection of one candidate and its validation, together with every rerun and every random choice. The solver makes no signature hash query, and since validity is the only requirement on its output and the finite menus of section 2.3 restrict only the wrapper of section 3.1, it may use any algorithm.

In the attack experiment of Crites and Stewart, which we call the fresh one query experiment, the adversary computes the elements $X_0, X_1, \ldots, X_{\tau-1}$ of lemma 2.1 from the public verification shares. It then samples $\alpha \in \mathbb{F}_p^\tau$ uniformly and a message $M^\star$ that has not appeared in any signature hash query, makes one fresh hash query and puts

$$R^\star = \prod_{j=0}^{\tau-1} X_j^{\alpha_j}, \qquad c^\star = \mathcal{H}_{\mathsf{sig}}(R^\star,\, PK,\, M^\star), \qquad v^\star = \alpha + c^\star K_0.$$

The first steps of this experiment are those of a trial of the uniform alpha wrapper of definition 7.1, which draws α and makes the fresh query in the same way. The proof of lemma 7.2 uses only these steps, namely the uniform draw of α after the preceding history and the freshness of the hash input, so it shows that $v^\star$ is uniform in $\mathbb{F}_p^\tau$ conditional on the history before α is drawn. The wrapper of section 3.1 changes this sampling rule, since there the adversary chooses α, and lemma 3.2 shows that its target is then uniform on a line.

Theorem 2.3 (Crites–Stewart solver to forgery interface)**.** *Consider a family indexed by λ, with N and the bit length $b_p = \lceil \log_2 p \rceil$ of p polynomially bounded in λ and with $0 \le f \le \tau - 1$. Generate $v^\star$ by the fresh one query experiment above, and define*

$$\delta_{\mathcal{S}}(\lambda) \coloneqq \Pr\left[\mathsf{GoodKG} \ \wedge\ \mathsf{Valid}_{H,f}\left(v^\star, \mathcal{S}(H, v^\star, f)\right)\right], \tag{2.2}$$

where the probability is over all randomness of the experiment, that is, the key material, α, the hash answer and the randomness of $\mathcal{S}$. If the complete implementation of $\mathcal{S}$ has probabilistic polynomial bit complexity, then there is a probabilistic polynomial time (PPT) adversary with adaptive TS-UF-0 advantage at least $\delta_{\mathcal{S}}(\lambda)$. Hence a non-negligible value contradicts adaptive TS-UF-0 security of the scheme family.

The construction and its proof belong to Crites and Stewart [CS25, Definition 10 and proofs of Theorems 1 and 2], and appendix A.1 gives the explicit adversary and

appendix A.2 the complete proof of theorem 2.3 in the notation of this paper. The probability $\delta_{\mathcal{S}}$ is joint, that is, it counts an execution when the key material has the Shamir form and the solver returns a valid witness for this target, so it is at most $\Pr[\mathsf{GoodKG}]$ and at most the probability that a witness exists. Polynomial time and non-negligibility are properties of the family indexed by λ, and a single fixed parameter row is one member of such a family, so it cannot establish either property by itself.

For a decoder route $\mathfrak{D}$, let $\mathcal{Q}$ be the relation realized by the fixed route, which is $P_{H,f}$ or a block restricted relation of equation (4.5), and let $E_{\mathcal{Q}}$ be the event that $v^{\star}$ has a witness for $\mathcal{Q}$, as in definition 4.3. Let $D_{\mathfrak{D}}$ be the event that the implementation returns and validates a witness, again in the experiment above. A validated witness is valid, so $\delta_{\mathcal{S}} \geq \Pr[\mathsf{GoodKG} \cap E_{\mathcal{Q}} \cap D_{\mathfrak{D}}]$, and when $\mathsf{GoodKG} \cap E_{\mathcal{Q}}$ has positive probability the chain rule gives

$$\begin{aligned}\Pr[\mathsf{GoodKG} \cap E_{\mathcal{Q}} \cap D_{\mathfrak{D}}] = \Pr[\mathsf{GoodKG}] \cdot \Pr[E_{\mathcal{Q}} \mid \mathsf{GoodKG}] \\ \cdot \Pr[D_{\mathfrak{D}} \mid \mathsf{GoodKG} \cap E_{\mathcal{Q}}].\end{aligned} \tag{2.3}$$

The factors are probabilities of the attack experiment itself, and the identity is the chain rule, which holds whatever the dependence between the events. The joint probability on the left bounds $\delta_{\mathcal{S}}$ from below in every case, even when a conditioning event has probability zero.

For the heterogeneous IRS route, whose guarantee holds on average over the attack distribution, the composition uses the additive bound of theorem 4.8. The support events overlap, so the product of the existence lower bound of equation (4.14) and the success rate of the IRS source would need an independence statement beyond the IRS theorem. For the folded routes, the setup assumption of section 5 fixes the ordered evaluation domain before the experiment starts, which enters as a hypothesis on the parameters and contributes no probability factor, because no distribution over setups is defined. A folded route returns a witness of the block restricted relation of equation (4.5), and every such witness also satisfies $P_{H,f}$, so it is a valid output and suffices for the forgery, although the route realizes only the restricted relation. Finally, the bit complexity in the theorem is that of the complete implementation, including the output, list production or pruning and validation. For a route that first computes an affine basis, as Guruswami–Wang do, it therefore includes the completion of section 5.4 as well as the field operations of the basis.

The implication runs only from a solver to a forgery, that is, a solver of probabilistic polynomial bit complexity with non-negligible validity probability yields a PPT forger. The routes studied in this paper certify no such solver on the tested regimes, so the hardness of the oracle $P_{H,f}$, the success of the attack with other solvers and the adaptive security of the scheme are open questions. Theorem 2.3 bounds the advantage of one inherited adversary from below, whereas corollary 3.7 bounds the success probability of the wrapper of definition 3.1 from above in its declared regions.

3 Chosen targets and the line point bound

This section defines the attack wrapper and proves the wrapper corollary, an upper bound on its success. The wrapper follows the Crites–Stewart template with two changes. The adversary may choose α from its history, provided that the choice is fixed before the fresh hash answer, and it accepts only witnesses that lie in the regions of a fixed menu. We first prove that the target of each trial is a uniform point on a line through the chosen vector.

We then bound, uniformly over all lines in direction K_0, the number of targets on the line that lie in the image $G_{\mathcal{M}}$ of a fixed menu. The bound takes three steps, namely two lemmas and the main theorem that combines them. A separation lemma shows that K_0

lies in no span of fewer than τ Shamir columns, so that witnesses for two distinct points of a line must jointly occupy at least τ coordinates. An elementary bound on families of sets with small pairwise intersections then counts the possible supports. The two lemmas combine into the count $M_{\mathrm{line}}(\mathcal{M})$ of the main theorem, and the wrapper corollary follows by conditioning on the history of each trial.

3.1 Trials and chosen targets

Definition 3.1 (Chosen prechallenge alpha wrapper)**.** Fix a valid Shamir setup on GoodKG with pairwise distinct nonzero points, the public coefficient commitments $X_0, X_1, \ldots, X_{\tau-1}$ of lemma 2.1, a finite menu $\mathcal{M}$ as in definition 2.2, a corruption cap $0 \le f \le \tau - 1$, and a trial budget $Q \ge 1$. The wrapper makes no signing query, and it makes no corruption query before it accepts its first valid witness. It performs at most Q trials. In trial i it proceeds as follows.

1. From the preceding history and its private randomness, it chooses and fixes an explicit coefficient vector $\alpha_i \in \mathbb{F}_p^\tau$ and a message M_i that has not appeared in any previous signature hash query, and it forms $R_i = \prod_{j=0}^{\tau-1} X_j^{\alpha_{i,j}}$.
2. It makes exactly one fresh signature hash query $c_i = \mathcal{H}_{\mathsf{sig}}(R_i, PK, M_i)$ and puts $v_i = \alpha_i + c_i K_0$. The representation α_i of R_i is not changed after this answer.
3. It runs any solvers it likes on v_i. The solvers make no signature hash query. They may try several members of $\mathcal{M}$ after seeing v_i, and recovery inside the region may be perfect and free.
4. It accepts an output $a \in \mathbb{F}_p^N$ only if $Ha = v_i$ for this same target and representation, $a \in \mathcal{W}_{\mathcal{M}}$, and $\mathrm{wt}(a) \le f$.
5. On acceptance it corrupts exactly the parties in $\mathrm{supp}(a)$, puts $s_i = \sum_{j=1}^{N} a_j \cdot sk_j$, outputs the forgery $(M_i, (R_i, s_i))$, and stops. Otherwise it stops or starts trial $i + 1$.

The wrapper succeeds if some trial accepts.

The forgery identity behind the last step is verified in corollary 3.7, and this subsection determines the distribution of the target v_i. Freshness refers to the complete hash input (R_i, PK, M_i), and a previously unqueried message is the wrapper's sufficient rule for it. On every execution that reaches the hash query of trial i, we write $\mathcal{F}_i$ for the history just before the answer c_i, the finite string that records the public input of the wrapper, every earlier query and answer and the private randomness used so far. Thus $\mathcal{F}_i$ is a random variable with countably many values, and each value h with $\Pr[\mathcal{F}_i = h] > 0$ determines the values of α_i, M_i and R_i on the event $\mathcal{F}_i = h$.

Lemma 3.2 (Uniform point on the chosen line)**.** *In trial i of definition 3.1, let h be a value of $\mathcal{F}_i$ with $\Pr[\mathcal{F}_i = h] > 0$, and let α be the value of α_i that it determines. Conditional on $\mathcal{F}_i = h$, the target v_i is uniform on the line $\ell_\alpha := \alpha + \mathbb{F}_p K_0$, which has exactly p points. That is, for every $c \in \mathbb{F}_p$,*

$$\Pr[\, v_i = \alpha + cK_0 \mid \mathcal{F}_i = h \,] = \frac{1}{p}.$$

Proof. Let M and R be the values of M_i and R_i that h determines. The input (R, PK, M) has not been queried before, because M is new. The event $\mathcal{F}_i = h$ depends only on the key material, the private randomness and the oracle answers recorded in h, and none of these is the answer of $\mathcal{H}_{\mathsf{sig}}$ at (R, PK, M). In the random oracle model that answer is uniform in $\mathbb{F}_p$ and independent of all of them, so $\Pr[\, c_i = c \mid \mathcal{F}_i = h \,] = 1/p$ for every $c \in \mathbb{F}_p$. Since $K_0 \neq 0$, the map $c \mapsto \alpha + cK_0$ is a bijection from $\mathbb{F}_p$ onto ℓ_α. Hence the line has exactly p points, and $v_i = \alpha + c_i K_0$ is uniform on it conditional on $\mathcal{F}_i = h$. □

The wrapper corollary uses two features of this statement. First, it is conditional on the entire past, so the vector α_i may depend on everything the adversary has seen, including earlier challenges and failed trials. Second, it is a statement about one trial at a time, and the corollary builds on these conditional laws and a union bound alone, so the success events of different trials may depend on one another.

The uniform alpha model, in which α_i is sampled uniformly and v_i becomes uniform in all of $\mathbb{F}_p^\tau$, is a different sampling rule with its own repetition bound. The wrapper above samples only the p points of a line chosen by the adversary. The bound of corollary 3.7 must therefore come from the geometry of such lines, which the following subsections supply.

3.2 Separation of the direction from small spans

Lemma 3.3 (Shamir line separation)**.** *Let p be prime, $1 \le \tau \le N < p$, and let $z_1, z_2, \dots, z_N \in \mathbb{F}_p$ be pairwise distinct and nonzero, with H and K_0 as in section 2. For $U \subseteq [N]$, let H_U denote the submatrix of H formed by the columns indexed by U, and let a_U denote the restriction of $a \in \mathbb{F}_p^N$ to the coordinates in U. Then the following hold.*

1. *For every $U \subseteq [N]$ with $|U| < \tau$, $K_0 \notin \operatorname{span}(H_U)$.*
2. *For every $\alpha \in \mathbb{F}_p^\tau$ and every such U, the line $\alpha + \mathbb{F}_p K_0$ meets $\operatorname{span}(H_U)$ in at most one point.*
3. *If $a, a' \in \mathbb{F}_p^N$ satisfy $Ha = \alpha + cK_0$ and $Ha' = \alpha + c'K_0$ with $c \ne c'$, then $|\operatorname{supp}(a) \cup \operatorname{supp}(a')| \ge \tau$.*

Proof. For the first claim, define $P_U(X) = \prod_{i \in U}(X - z_i)$, a polynomial of degree $|U| \le \tau - 1$, and let $\pi \in \mathbb{F}_p^\tau$ be its coefficient vector padded with zeros to length τ. For every column, $\pi^\top K_i = P_U(z_i)$, which vanishes for $i \in U$. Hence π annihilates $\operatorname{span}(H_U)$. On the other hand, $\pi^\top K_0 = P_U(0) = \prod_{i \in U}(-z_i) \ne 0$, because every z_i is nonzero. A vector annihilated by π cannot equal K_0, so $K_0 \notin \operatorname{span}(H_U)$.

For the second claim, suppose that $\alpha + cK_0$ and $\alpha + c'K_0$ both lie in $\operatorname{span}(H_U)$ with $c \ne c'$. Their difference $(c - c')K_0$ lies in the span, and division by $c - c' \ne 0$ places K_0 in $\operatorname{span}(H_U)$, which contradicts the first claim.

For the third claim, let $W = \operatorname{supp}(a) \cup \operatorname{supp}(a')$. The difference $a - a'$ is supported in W and satisfies $H(a - a') = (c - c')K_0$, so $K_0 = (c - c')^{-1} \cdot H(a - a') \in \operatorname{span}(H_W)$. By the first claim, $|W| \ge \tau$. □

For $\tau \ge 2$, the nonzero hypothesis is needed in addition to the full spark property of H. If $z_i = 0$, then $K_i = K_0$, so the span of the single column K_i contains the entire line $\mathbb{F}_p K_0$ although $1 < \tau$.

3.3 Families of sets with bounded pairwise intersections

The following elementary count is the only combinatorial tool of the proof. For integers $0 \le \nu < w \le n$ and $0 \le j \le \nu$, define

$$D_j \coloneqq w^2 - n\nu + j \cdot (n + \nu - 2w), \qquad J_j \coloneqq \left\lfloor \frac{(n-j)(w-\nu)}{D_j} \right\rfloor, \qquad T_j \coloneqq \left\lfloor \frac{\binom{n}{j} J_j}{\binom{w}{j}} \right\rfloor, \tag{3.1}$$

where J_j and T_j require $D_j > 0$. We also define

$$T_{\text{pack}} \coloneqq \left\lfloor \frac{\binom{n}{\nu+1}}{\binom{w}{\nu+1}} \right\rfloor,$$

and

$$T_{\mathrm{set}}(n,\, w,\, \nu) \coloneqq \min\left(\left\{\binom{n}{w},\, T_{\mathrm{pack}}\right\} \cup \{T_j :\, 0 \le j \le \nu,\; D_j > 0\}\right).$$

The first two members make this set nonempty, even when no D_j is positive.

Lemma 3.4 (Bounded intersection families)**.** *Let $0 \le \nu < w \le n$, and let $\mathcal{C}$ be a family of distinct w element subsets of an n element set such that $|C \cap C'| \le \nu$ for all distinct $C,\, C' \in \mathcal{C}$. Then $|\mathcal{C}| \le T_{\mathrm{set}}(n,\, w,\, \nu)$.*

Proof. We verify that each quantity in the minimum bounds $|\mathcal{C}|$. The bound $\binom{n}{w}$ counts all w element subsets. For the packing bound, fix a $(\nu+1)$ element subset J. Two distinct members containing J would intersect in at least $\nu+1$ elements, so J lies in at most one member. Every member contains $\binom{w}{\nu+1}$ such subsets, and counting the pairs $(J,\, C)$ with $J \subseteq C \in \mathcal{C}$ gives $|\mathcal{C}| \cdot \binom{w}{\nu+1} \le \binom{n}{\nu+1}$. Since $|\mathcal{C}|$ is an integer, $|\mathcal{C}| \le T_{\mathrm{pack}}$.

Now fix $0 \le j \le \nu$ with $D_j > 0$ and a j element subset J, and let A_J be the number of members containing J. If $A_J = 0$ there is nothing to prove, so let $A_J > 0$ and delete J from each of these members. The residual sets have size $w - j$, they lie in a universe of $n - j \ge 1$ elements because $j \le \nu < w \le n$, and they pairwise intersect in at most $\nu - j$ elements, because the original intersections contained J. For each element t of the residual universe, let x_t be the number of residual sets containing t. Counting memberships gives $\sum_t x_t = A_J(w-j)$, and counting ordered pairs of distinct residual sets through a common element gives

$$\sum_t x_t(x_t - 1) \le A_J(A_J - 1)(\nu - j).$$

The Cauchy–Schwarz inequality gives $\sum_t x_t^2 \ge \left(\sum_t x_t\right)^2/(n-j) = A_J^2(w-j)^2/(n-j)$. Since $\sum_t x_t(x_t-1) = \sum_t x_t^2 - \sum_t x_t$, the two displays combine into

$$\frac{A_J^2(w-j)^2}{n-j} - A_J(w-j) \le A_J(A_J-1)(\nu-j).$$

Dividing by $A_J > 0$ and multiplying by $n - j > 0$ yields $A_J \cdot \big((w-j)^2 - (n-j)(\nu-j)\big) \le (n-j)(w-\nu)$, and expanding the coefficient on the left shows that it equals D_j. Therefore $A_J D_j \le (n-j)(w-\nu)$, and since $D_j > 0$ and A_J is an integer, $A_J \le J_j$. Finally, every member contains $\binom{w}{j}$ subsets of size j, so counting the pairs $(J,\, C)$ with $J \subseteq C \in \mathcal{C}$ gives $|\mathcal{C}| \cdot \binom{w}{j} = \sum_J A_J \le \binom{n}{j} J_j$. Taking the integer part yields $|\mathcal{C}| \le T_j$. □

The lemma is an elementary incidence count. Its term $T_0 = J_0$ is Johnson's classical bound for constant weight codes [Joh62, Theorem 3, Equation (7)], the terms with $j \ge 1$ apply that bound to the families shortened by j common elements, and T_{pack} is an incidence packing bound. Guruswami's survey presents the related Johnson bound for list decoding [Gur06, Section 3.1], and the count of the main theorem below is a list size bound of this kind. A witness a for a point of the line $\alpha + \mathbb{F}_p K_0$ gives a codeword of $\{x \in \mathbb{F}_p^N :\; Hx \in \mathbb{F}_p K_0\}$, a code of minimum distance τ, at distance $\mathrm{wt}(a)$ from a fixed a_0 with $Ha_0 = \alpha$, and different points give different codewords. We record the lemma with its complete proof because the main theorem depends on the exact values of the floors and not only on their order of magnitude.

3.4 Count definitions and the main theorem

We specialize the set bound to the two kinds of region in a menu. For $0 \le F < \tau$, for $m \mid N$ with $L = N/m$, and for an integer $B \ge 0$ with $mB < \tau$, define

$$\begin{aligned} M_{\mathrm{coord}}(N,\, \tau,\, F) &\coloneqq T_{\mathrm{set}}(N,\, N-F,\, k), \\ M_{\mathrm{block}}(N,\, \tau,\, m,\, B) &\coloneqq \min\left\{\binom{L}{B},\; T_{\mathrm{set}}\big(L,\, L-B,\, \lfloor k/m \rfloor\big),\; M_{\mathrm{coord}}(N,\, \tau,\, mB)\right\}, \end{aligned} \tag{3.2}$$

and, for a finite menu $\mathcal{M}$ as in definition 2.2, define

$$M_{\mathrm{line}}(\mathcal{M}) \coloneqq M_{\mathrm{coord}}(N,\, \tau,\, F) + \sum_{\iota \in I} M_{\mathrm{block}}(N,\, \tau,\, m_\iota,\, B_\iota). \tag{3.3}$$

The arguments lie in the domain of lemma 3.4. For M_{coord} one has $0 \le k < N - F \le N$ exactly because $F < \tau$. For the block term, $m \mid N$ and $N = \tau + k$ give

$$L - \lfloor k/m \rfloor = \lceil \tau/m \rceil,$$

and $mB < \tau$ gives $B \le \lceil \tau/m \rceil - 1$, so $0 \le \lfloor k/m \rfloor < L - B \le L$ and $0 \le mB < \tau$. These counts are integers determined by N, τ and the menu alone, so the finite application evaluates $M_{\mathrm{line}}(\mathcal{M})$ for the canonical menu. In the theorem below they bound the number of targets of a line that admit a witness in a region of the menu or, for $M_{\mathrm{line}}(\mathcal{M})$, in the whole menu.

Theorem 3.5 (Fixed menu line bound)**.** *Let p be prime, let $1 \le \tau \le N < p$, and let the Shamir points $z_1, z_2, \ldots, z_N \in \mathbb{F}_p$ be pairwise distinct and nonzero. Let $\mathcal{M}$ be a finite menu in the sense of definition 2.2, fixed independently of the targets. Then for every $\alpha \in \mathbb{F}_p^\tau$,*

$$\left| G_{\mathcal{M}} \cap (\alpha + \mathbb{F}_p K_0) \right| \le \min\{\, p,\ M_{\mathrm{line}}(\mathcal{M}) \,\}.$$

Proof. Fix α and write $\ell = \alpha + \mathbb{F}_p K_0$. Since the map $c \mapsto \alpha + cK_0$ is a bijection from $\mathbb{F}_p$ onto ℓ, it suffices to bound the set

$$C_{\mathcal{M}} \coloneqq \{\, c \in \mathbb{F}_p :\ \alpha + cK_0 \in G_{\mathcal{M}} \,\},$$

which is the union of the coordinate set $C_F \coloneqq \{c :\ \exists\, a,\ \mathrm{wt}(a) \le F,\ Ha = \alpha + cK_0\}$ and of the block sets $C_\iota \coloneqq \{c :\ \exists\, a,\ \mathrm{wt}_{\mathcal{P}_\iota}(a) \le B_\iota,\ Ha = \alpha + cK_0\}$ for $\iota \in I$. We bound each of these sets and then take the union.

Coordinate transfer. For every $c \in C_F$ choose one witness a_c with $\mathrm{wt}(a_c) \le F$ and $Ha_c = \alpha + cK_0$, and enlarge its support to a set $S_c \subseteq [N]$ with exactly F elements by adding arbitrary indices if needed, which is possible because $F < \tau \le N$. For distinct $c, c' \in C_F$, the third claim of lemma 3.3 gives $|\operatorname{supp}(a_c) \cup \operatorname{supp}(a_{c'})| \ge \tau$, and the inclusions $\operatorname{supp}(a_c) \subseteq S_c$ and $\operatorname{supp}(a_{c'}) \subseteq S_{c'}$ give $|S_c \cup S_{c'}| \ge \tau$. In particular $S_c \ne S_{c'}$, because a coincidence would give a union of only $F < \tau$ elements. The complements $[N] \setminus S_c$ are therefore distinct sets of the common size $N - F$ with

$$\left| ([N] \setminus S_c) \cap ([N] \setminus S_{c'}) \right| = N - |S_c \cup S_{c'}| \le N - \tau = k.$$

Lemma 3.4 with $(n, w, \nu) = (N,\, N - F,\, k)$ gives

$$|C_F| \le T_{\mathrm{set}}(N,\, N - F,\, k) = M_{\mathrm{coord}}(N,\, \tau,\, F).$$

Block transfer. Fix ι and write $\mathcal{P} = \mathcal{P}_\iota$, $m = m_\iota$, $L = L_\iota$ and $B = B_\iota$. We give three bounds on $|C_\iota|$. First, every a with $\mathrm{wt}_{\mathcal{P}}(a) \le B$ has its support inside the union U of some B blocks of $\mathcal{P}$, and $|U| = mB < \tau$. There are $\binom{L}{B}$ such unions, and by the second claim of lemma 3.3 each $\mathrm{span}(H_U)$ meets ℓ in at most one point. Hence $|C_\iota| \le \binom{L}{B}$. Second, for every $c \in C_\iota$ choose one witness a_c and enlarge the set of blocks meeting its support to a set $\mathcal{B}_c$ of exactly B blocks by adding arbitrary blocks if needed, which is possible because $B < L$. For distinct $c, c' \in C_\iota$ the coordinate supports have a union of at least τ elements by the third claim of lemma 3.3, and a union of b blocks contains at most mb coordinates, so $\mathcal{B}_c \cup \mathcal{B}_{c'}$ contains at least $\lceil \tau/m \rceil$ blocks. Since $B \le \lceil \tau/m \rceil - 1$, the enlarged block sets are distinct, and their complements within $\mathcal{P}$ are distinct sets of $L - B$ blocks with pairwise intersection at most $L - \lceil \tau/m \rceil = \lfloor k/m \rfloor$. Lemma 3.4 with $(n, w, \nu) = (L,\, L - B,\, \lfloor k/m \rfloor)$ gives $|C_\iota| \le T_{\mathrm{set}}(L,\, L - B,\, \lfloor k/m \rfloor)$. Third, $\mathrm{wt}_{\mathcal{P}}(a) \le B$ implies $\mathrm{wt}(a) \le mB < \tau$, so $C_\iota \subseteq C_{mB}$, and the coordinate transfer with F replaced by mB gives $|C_\iota| \le M_{\mathrm{coord}}(N, \tau, mB)$. All three quantities bound the same set, so $|C_\iota| \le M_{\mathrm{block}}(N, \tau, m, B)$.

Menu union. Since $C_{\mathcal{M}} = C_F \cup \bigcup_{\iota \in I} C_\iota$, the bounds add up to $|C_{\mathcal{M}}| \le M_{\text{line}}(\mathcal{M})$, and overlaps between the sets only cause overcounting. Finally, $C_{\mathcal{M}} \subseteq \mathbb{F}_p$ has at most p elements, which gives the minimum. The argument is combinatorial throughout. □

Remark 3.6 (Half threshold boundaries). Two special cases of the coordinate transfer clarify the small counts at the half threshold. If $2F < \tau$, then two witnesses of weight at most F have a support union of at most $2F < \tau$ elements, so by lemma 3.3 at most one point of any line lies in the image of the coordinate region. If $2F = \tau$ and $F > 0$, then distinct covered points force supports of size exactly F that are pairwise disjoint, so at most $\lfloor N/F \rfloor$ points of a line are covered. In the equality case the bound $\lfloor N/F \rfloor$ cannot be replaced by one. For disjoint $U, V \subseteq [N]$ of size F, the τ columns of $[H_U \ -H_V]$ are linearly independent, so the equation $H_U a - H_V b = K_0$ has a solution, and the line through $H_V b$ contains the second covered point $H_U a = H_V b + K_0$. At $F = 0$ or $B = 0$ the region is $\{0\}$, so at most one point is covered. The count equals one exactly for the line $\mathbb{F}_p K_0$ through the origin, and every other line contains no covered point. These boundary arguments explain individual terms without replacing the definition of M_{line} in equation (3.3), which is the quantity evaluated in the applications. They need not improve it. When $2F = \tau$ and $F > 0$, for example, the term $j = 0$ of that definition already equals $\lfloor N/F \rfloor$, because then $D_0 = F^2$ and $N(w - \nu) = NF$.

3.5 The wrapper corollary

Corollary 3.7 (Chosen prechallenge alpha)**.** *Consider the wrapper of definition 3.1 with a finite menu $\mathcal{M}$, a cap $0 \le f \le \tau - 1$ and at most Q trials, on a fixed valid Shamir setup with pairwise distinct nonzero points. Then*

$$\Pr[\,\textit{the wrapper succeeds within } Q \textit{ trials} \mid \mathsf{GoodKG}\,] \le \min\left\{1,\ \frac{Q \cdot M_{\text{line}}(\mathcal{M})}{p}\right\},$$

and every accepted witness yields a valid forgery on an unsigned message within the corruption cap.

Proof. All probabilities in this proof are conditional on GoodKG. This event depends only on the key material, which is independent of the random oracle, so lemma 3.2 holds conditional on GoodKG and on each value h as well. Let E_i be the event that trial i is reached and accepts a witness. On an execution that stops before trial i, no fresh query is made in that trial, and the acceptance indicator of E_i is zero. On an execution that reaches trial i, acceptance requires some $a \in \mathcal{W}_{\mathcal{M}}$ with $Ha = v_i$, so E_i is contained in the hit event $\{v_i \in G_{\mathcal{M}}\}$, and the additional checks $\mathrm{wt}(a) \le f$ and the representation condition only shrink E_i. Fix a value h of $\mathcal{F}_i$ with $\Pr[\mathcal{F}_i = h] > 0$, so that trial i is reached and its fresh query is actually made, and let α be the value of α_i that h determines. The menu is fixed independently of the targets, so $G_{\mathcal{M}} \cap \ell_\alpha$ is a fixed set, and theorem 3.5 bounds its size by $M_{\text{line}}(\mathcal{M})$. By lemma 3.2, v_i is uniform on the p points of ℓ_α conditional on $\mathcal{F}_i = h$. Therefore

$$\Pr[\,E_i \mid \mathcal{F}_i = h\,] \le \Pr[\,v_i \in G_{\mathcal{M}} \cap \ell_\alpha \mid \mathcal{F}_i = h\,] \le \frac{M_{\text{line}}(\mathcal{M})}{p}.$$

Up to an event of probability zero, E_i is contained in the union of the disjoint events $\mathcal{F}_i = h$ over these countably many values h, so $\Pr[E_i] = \sum_h \Pr[\mathcal{F}_i = h] \cdot \Pr[\,E_i \mid \mathcal{F}_i = h\,] \le M_{\text{line}}(\mathcal{M})/p$ for every $i \le Q$. The success event is contained in $\bigcup_{i \le Q} E_i$, and the union bound gives the claimed inequality after clipping at one. The argument conditions on the actual history of each reached trial, so the choice of α_i and the decision to continue may depend on everything seen before, and the bound holds whatever the dependence between the events E_i.

For the forgery, let trial i accept a. The wrapper corrupts exactly the parties in $\operatorname{supp}(a)$, and these are its only corruptions. Since $|\operatorname{supp}(a)| = \operatorname{wt}(a) \le f$, the set of all parties ever corrupted respects the cap. On GoodKG, $sk_j = K_j^\top x$ for every party, and the Shamir identity gives

$$s_i = \sum_{j=1}^{N} a_j \cdot sk_j = (Ha)^\top x = (\alpha_i + c_i K_0)^\top x = \alpha_i^\top x + c_i \cdot x_0.$$

Exponentiation gives $g^{s_i} = \prod_{j=0}^{\tau-1} X_j^{\alpha_{i,j}} \cdot PK^{c_i} = R_i \cdot PK^{c_i}$, which is the Schnorr verification equation for (R_i, s_i) on M_i with challenge $c_i = \mathcal{H}_{\mathsf{sig}}(R_i, PK, M_i)$. The message M_i was never signed, because the wrapper makes no signing query. This is the forgery construction of Crites and Stewart [CS25, proofs of Theorems 1 and 2]. □

The corollary is an upper bound for the wrapper of definition 3.1 only. It holds even when recovery inside the region is perfect and free, so additional running time or solver failure cannot raise it. It also holds for every trial budget Q, including an exponential one, because its proof bounds each trial conditionally on the history of that trial and combines the trials by a union bound. Here Q counts fresh hash queries, and the corollary makes no claim about running time. Whether the factor $M_{\text{line}}(\mathcal{M})$ is necessary and whether any attack attains the bound are left open. The inequality bounds how often a target of the wrapper has a witness in a declared region, and the behavior of a decoder outside that region therefore plays no part in it. The solver to forgery reduction, which constructs an attack from an efficient solver, runs in the opposite direction.

4 Decoder framework and coordinate routes

The line point bound of section 3 counts targets and needs no decoder. This section develops the decoder framework on which the comparison of the seven decoder routes in section 5.5 rests. We first restate, in the notation of section 2, the inherited equivalence between the sparse relation and bounded distance decoding of a scaled Reed–Solomon code, together with the parameter dictionary. We then fix the block restricted relation of an aligned partition and state what a decoder theorem must certify before it counts as a realization of either relation. We also record the bit cost of turning a decoder into an oracle and distinguish support spans from exact decoder lists. The certified regions of the seven decoders and the canonical menu build on these contracts. Section 4.6 closes this section by applying them to the three coordinate routes.

4.1 Reed–Solomon normalization and the parameter dictionary

We keep the Shamir data of section 2.1, with its dimension k, and the sparse relation $P_{H,f}$ of equation (2.1). In coding theoretic terms the length, the error budget, the required number of agreements and the rate are

$$n = N, \qquad e = f, \qquad A = N - f, \qquad \beta = \frac{k}{N}.$$

The Reed–Solomon code $\mathrm{RS}(\mathbb{F}_p, z, k)$ consists of the vectors $(q(z_1), q(z_2), \ldots, q(z_N))$ for the polynomials $q \in \mathbb{F}_p[X]$ of degree below k. These are the evaluation codes of Reed and Solomon [RS60], here evaluated at an arbitrary set of pairwise distinct points as in [CS25, Section 7.1], and the connection between Shamir's secret sharing and these codes is due to McEliece and Sarwate [MS81]. Crites and Stewart prove the following diagonal equivalence and the associated bounded distance reduction [CS25, Theorem 3, Lemmas 4 and 5, and Corollary 2]. We include a concise proof to fix the orientation of the diagonal scaling used throughout the decoder analysis.

Proposition 4.1 (GRS normalization)**.** *Define*

$$w_i = \frac{1}{\prod_{j \neq i}(z_i - z_j)}, \qquad D_w = \mathrm{diag}(w_1, w_2, \ldots, w_N).$$

Then

$$C_H \coloneqq \ker(H) = D_w \ \mathrm{RS}(\mathbb{F}_p, z, N - \tau). \tag{4.1}$$

Moreover, multiplication by D_w preserves the Hamming weight and the Hamming distance d_H.

Proof. We first prove the inclusion. Fix $q \in \mathbb{F}_p[X]$ with $\deg q < k$ and let $\hat{c} \in \mathbb{F}_p^N$ have the coordinates $\hat{c}_i = w_i \cdot q(z_i)$. The coefficient of X^{N-1} in the Lagrange interpolation formula for any polynomial P of degree at most $N-1$ is

$$\sum_{i=1}^{N} w_i \cdot P(z_i).$$

For $0 \leq j < \tau$, the polynomial $P_j(X) = X^j \cdot q(X)$ has degree at most $(\tau-1)+(k-1) = N-2$. Its X^{N-1} coefficient is therefore zero, and

$$(H\hat{c})_j = \sum_{i=1}^{N} w_i \cdot z_i^j \cdot q(z_i) = 0.$$

Hence $D_w \ \mathrm{RS}(\mathbb{F}_p, z, k) \subseteq \ker(H)$. We next compare dimensions. The code $\mathrm{RS}(\mathbb{F}_p, z, k)$ has dimension k, and D_w is invertible because every w_i is nonzero. Any τ columns of H form a square Vandermonde matrix on pairwise distinct points, so $\mathrm{rank}(H) = \tau$. Hence by the rank nullity theorem the inclusion is actually an equality. This proves equation (4.1). Finally, every w_i is nonzero, so diagonal scaling preserves support and Hamming distance. □

Proposition 4.2 (Oracle decoding equivalence)**.** *Fix $v \in \mathbb{F}_p^\tau$ and any $u \in \mathbb{F}_p^N$ satisfying $Hu = v$. Then $P_{H,f}(v)$ is solvable if and only if some $c_H \in C_H$ satisfies*

$$d_H(u, c_H) \leq f.$$

Equivalently, the normalized received word

$$y = D_w^{-1} u \tag{4.2}$$

lies within coordinate distance f of $\mathrm{RS}(\mathbb{F}_p, z, N - \tau)$.

Proof. If $Ha = v$, then $c_H = u - a \in \ker(H)$ and $d_H(u, c_H) = \mathrm{wt}(a)$. Conversely, a nearby $c_H \in \ker(H)$ gives the witness $a = u - c_H$. Substituting $c_H = D_w c$ from proposition 4.1 and using metric preservation gives

$$d_H(u, c_H) = d_H(D_w^{-1}u, c),$$

which proves the normalized statement. □

For a witness a, the normalized decoding error has weight $\mathrm{wt}(a) \leq f$, so its codeword agrees with the received word in at least $N - f$ coordinates. The attack relevant near capacity rows are most transparent in this agreement notation. If $f = \tau - g$, then the required number of agreements is

$$A = N - f = N - \tau + g = k + g. \tag{4.3}$$

Thus the decoder seeks a degree $< k$ polynomial agreeing with the received word on only $k+g$ unknown positions. If those positions were known, k would determine the polynomial and the other g would verify it. The unknown agreement set makes this a genuine near capacity search problem and explains why f can be close to τ while lying far beyond the certified coordinate decoder radii.

4.2 The block restricted relation of a fixed partition

Folded decoders control the number of erroneous blocks of a fixed partition rather than the number of erroneous coordinates, so their guarantees apply to a restricted relation. Fix $m \mid N$, let $L = N/m$, and partition the ordered coordinates into

$$\mathcal{B}_j = \Big\{(j-1)m+1,\ (j-1)m+2,\ \ldots,\ jm\Big\}, \qquad j = 1,\ 2,\ \ldots,\ L.$$

For $a \in \mathbb{F}_p^N$, define

$$\mathrm{wt}_m(a) \coloneqq \#\{\, j : \mathrm{supp}(a) \cap \mathcal{B}_j \neq \varnothing \,\}.$$

Then

$$\left\lceil \frac{\mathrm{wt}(a)}{m} \right\rceil \leq \mathrm{wt}_m(a) \leq \mathrm{wt}(a). \tag{4.4}$$

For this aligned partition $\mathcal{P} = \{\mathcal{B}_1, \mathcal{B}_2, \ldots, \mathcal{B}_L\}$, the block weight wt_m is the block weight $\mathrm{wt}_{\mathcal{P}}$ of section 2.3, and the partition is one of the partitions admitted by definition 2.2. The blocks are consecutive ranges of coordinate indices, so the partition depends on the order in which the Shamir points $z_1, z_2, \ldots, z_N$ are listed, and that order is part of the definition.

For a block budget B, imposed on top of the coordinate cap of the attack, the block restricted relation is

$$P_{m,B}(v): \quad \text{find } a \in \mathbb{F}_p^N \text{ such that } Ha = v, \quad \mathrm{wt}(a) \leq f, \quad \mathrm{wt}_m(a) \leq B. \tag{4.5}$$

Every $P_{m,B}$ witness solves $P_{H,f}$, but the converse need not hold. The relation keeps the cap $\mathrm{wt}(a) \leq f$ of the attack, whereas the witness region $\mathcal{W}_{\mathcal{M}}$ of definition 2.2 is defined by the budgets alone, as section 2.3 records.

4.3 Decoder realizations and attack certificates

A decoder theorem speaks about solvable search instances, whereas a decoder assisted attack must also account for how often the attack experiment produces such an instance. In the next definition, $D_{\mathfrak{D}}$ is the event that decoder route $\mathfrak{D}$ returns a valid witness.

Definition 4.3 (Theorem supported oracle realization)**.** Let $\mathcal{Q}_\lambda(v)$ be either $P_{H,f}(v)$ (defined in (2.1)) or a block restricted relation $P_{m,B}(v)$ (defined in (4.5)), and define $E_{\mathcal{Q}} \coloneqq \{\mathcal{Q}_\lambda(v) \neq \varnothing\}$.

A decoder route is a *pointwise theorem supported realization* of $\mathcal{Q}$, with conditional correctness $q_{\mathfrak{D}}(\lambda)$, if its reduction, setup, and radius hypotheses are proved and, for every solvable admissible input, its complete procedure returns a valid witness with probability at least $q_{\mathfrak{D}}(\lambda)$. Decoder execution, required list enumeration, filtering, and verification must all be polynomial in the cryptographic input length. The realization is PPT when $q_{\mathfrak{D}}$ is non-negligible.

If the guarantee instead holds only after v and auxiliary randomness are sampled from a stated distribution,

$$\Pr[D_{\mathfrak{D}} \mid E_{\mathcal{Q}}] \geq q_{\mathfrak{D}}(\lambda),$$

we call it a *distributional realization* and record that distribution explicitly.

To obtain a decoder assisted attack through theorem 2.3, we need a pointwise or explicitly distributional realization of a relation $\mathcal{Q}$ whose outputs solve $P_{H,f}$. On the stated attack distribution and under the stated setup conditioning, a solvability lower bound and its justified composition with decoder success must give a non-negligible probability of a valid witness. The complete solver to forgery procedure must also have polynomial bit

complexity. The frequency of solvable inputs enters through these requirements alone. In the pointwise case,

$$\Pr[E_{\mathcal{Q}} \cap D_{\mathfrak{D}}] \geq q_{\mathfrak{D}} \cdot \Pr[E_{\mathcal{Q}}],$$

whatever the dependence between the two events. A distributional theorem may instead require a direct joint argument, as in the IRS route below.

4.4 End to end bit complexity

For a cryptographic family indexed by λ, we require $b_p(\lambda) = \Theta(\lambda)$, $N(\lambda) = \mathrm{poly}(\lambda)$, and polynomially represented auxiliary parameters. A polynomial number of $\mathbb{F}_p$ operations then has polynomial bit cost. In contrast, $p^s > 2^{s(b_p-1)}$, so p^s is not polynomial in b_p for any fixed $s > 0$. A cost $N^{\mathcal{O}(1/\epsilon)}$ is polynomial only for a uniformly bounded exponent, whereas $\mathrm{poly}(N, b_p, 1/\epsilon)$ is polynomial for inverse polynomial ϵ.

A complete list of size $\mathcal{L}$ also requires $\Omega(\mathcal{L} \cdot k \cdot b_p)$ output bits, which may obstruct explicit complete list decoding. Enumerating the complete list and finding one valid witness are different tasks, and this output requirement applies to the first task only.

Proposition 4.4 (Polynomial decoder to oracle overhead)**.** *Fix τ columns T of H. The reduction data can be preprocessed in $\mathcal{O}(\tau^3 + N^2)$ field operations. For each target, representative construction and normalization cost $\mathcal{O}(\tau^2 + N)$. A route requiring the randomized coset representative of lemma 4.7 additionally samples a uniform kernel element, at a cost of $\mathcal{O}(Nk)$ field operations. Conversion, filtering, and verification of a complete list of size $\mathcal{L}$, whose entries are given by the k coefficients of their polynomials, cost $\mathcal{O}(\mathcal{L} \cdot N \cdot (k + \tau)) = \mathcal{O}(\mathcal{L} \cdot N^2)$ further field operations.*

Proof. Fix $T \subseteq [N]$ of cardinality τ. Every τ columns of H are independent, so H_T is invertible. Precompute H_T^{-1} by Gaussian elimination and all Lagrange weights

$$w_i = \left(\prod_{j \neq i} (z_i - z_j) \right)^{-1}.$$

This costs $\mathcal{O}(\tau^3 + N^2)$ field operations. For a target v, set

$$u_T = H_T^{-1} v, \qquad u_{[N]\setminus T} = 0.$$

Then $Hu = v$, and $y = D_w^{-1} u$ is the normalized RS received word. Representative construction and normalization cost $\mathcal{O}(\tau^2 + N)$. A route that needs the randomized coset representative adds a uniform element of C_H to u. By equation (4.1) such an element is D_w applied to a uniform codeword, so it suffices to sample the k coefficients of a polynomial of degree below k, evaluate it at the N points, and apply the diagonal map, which costs $\mathcal{O}(Nk)$ field operations.

For every decoded polynomial of degree below k, evaluate it at the N points to obtain $c \in \mathrm{RS}(\mathbb{F}_p, z, k)$, form $a = u - D_w c$, check the coordinate and any block weight conditions, and verify $Ha = v$. Naive evaluation costs $\mathcal{O}(Nk)$ and naive matrix vector verification costs $\mathcal{O}(N\tau)$ per candidate. A list of size $\mathcal{L}$ therefore adds $\mathcal{O}(\mathcal{L}N(k + \tau)) = \mathcal{O}(\mathcal{L}N^2)$ field operations. □

4.5 Support spans versus exact decoder lists

For uniform $v \in \mathbb{F}_p^\tau$, define

$$X_f(v) \coloneqq \#\{F \subseteq [N] : |F| = f,\ v \in \operatorname{span}(H_F)\},$$

$$Z_j(v) \coloneqq \#\{a \in \mathbb{F}_p^N : Ha = v,\ \operatorname{wt}(a) = j\},$$

$$Z_{\le f}(v) \coloneqq \sum_{j=0}^{f} Z_j(v).$$

Here X_f counts support spans, whereas Z_j counts coefficient vectors and hence potential complete decoder outputs. There are $\binom{N}{j}(p-1)^j$ vectors of exact weight j, and each of them satisfies $Ha = v$ with probability $p^{-\tau}$ for uniform v. Hence, for $0 \le j \le N$,

$$\mathbb{E}[Z_j] = \binom{N}{j}\frac{(p-1)^j}{p^\tau}, \qquad \mathbb{E}[Z_{\le f}] = \frac{1}{p^\tau}\sum_{j=0}^{f}\binom{N}{j}(p-1)^j. \tag{4.6}$$

If $f \le \tau$, every f column submatrix of H is independent, so a size f support span containing v has one coefficient representation, whose exact support has some size $j \le f$ and lies in $\binom{N-j}{f-j}$ supports of size f. With the case $j = f$ of equation (4.6) this gives

$$X_f(v) = \sum_{j=0}^{f}\binom{N-j}{f-j}Z_j(v), \qquad \mathbb{E}[Z_f] = \left(1 - \frac{1}{p}\right)^f \mu_f, \tag{4.7}$$

where $\mu_f = \binom{N}{f}p^{f-\tau} = \mathbb{E}[X_f]$ is the first moment of the support span count. These are elementary counting identities. Crites and Stewart compute the first moment μ_f [CS25, proof of Theorem 2]. The exact decoder list counts coefficient vectors, and its expected size is $\mathbb{E}[Z_{\le f}]$ of equation (4.6).

For any u with $Hu = v$, the map $a \mapsto u - a$ bijects the witnesses counted by $Z_{\le f}(v)$ with codewords of C_H within distance f of u. Since diagonal normalization preserves distance, $Z_{\le f}$ is the exact complete coordinate decoder list size, and its explicit output has a cost. Assume $k \ge 1$, and let a randomized coordinate decoder output the complete radius f list with probability at least π_{comp} on every input. For a fixed target v a complete output contains all $Z_{\le f}(v)$ distinct codewords, so averaging first over the coins of the decoder and then over uniform v gives at least $\pi_{\mathrm{comp}}\,\mathbb{E}[Z_{\le f}]$ expected output records. A codeword of C_H is D_w applied to a codeword of $\mathrm{RS}(\mathbb{F}_p, z, k)$ and is represented by the k coefficients of its polynomial, so in this coefficient representation the expected output length is

$$\Omega\left(\pi_{\mathrm{comp}}\, k\, b_p\, \mathbb{E}[Z_{\le f}]\right) \quad \text{bits.} \tag{4.8}$$

A superpolynomial average complete list therefore rules out an expected polynomial explicit complete list implementation with non-negligible pointwise completeness probability. This bound is about enumerating the list, and section 4.4 keeps that task apart from the search for one witness.

4.6 Coordinate decoder routes

This subsection applies the contracts of the preceding subsections to Berlekamp–Welch (BW), to heterogeneous interleaved Reed–Solomon (IRS) decoding, and to Guruswami–Sudan (GS) list decoding. For each route we record the certified radius, the output form and the cost, and then the kind of guarantee that it gives for $P_{H,f}$.

For the Shamir data of section 2.1, assume $1 \le f \le \tau - 1$, $2 \le k < N$ and $N \le p$. The routes in this subsection operate in ordinary coordinate Hamming distance after proposition 4.1, so their decoded error weight is exactly the witness weight.

4.6.1 Berlekamp–Welch

The normalized code has parameters $[N, k, \tau + 1]$. Berlekamp–Welch therefore uniquely decodes every error of weight

$$f \le F_{\mathrm{BW}}, \qquad F_{\mathrm{BW}} \coloneqq \left\lfloor \frac{\tau}{2} \right\rfloor .$$

We use Berlekamp–Welch decoding in this region [Bla08, Sections 3.11 and 3.12]. Inside this region the returned codeword yields a witness through proposition 4.2, so the route realizes $P_{H,f}$ pointwise and deterministically. Solving its error locator system by Gaussian elimination costs $\mathcal{O}(N^3)$ field operations, and its one word output and proposition 4.4 preserve bit polynomial time.

4.6.2 Heterogeneous IRS

Yayla lifts the normalized $\mathrm{RS}(\mathbb{F}_p, z, k)$ instance to a heterogeneous interleaving with component dimensions

$$k_i = i(k-1) + 1, \qquad i = 1, 2, \ldots, r$$

[Yay11, Theorem 3.3.2]. Here the integer $r \ge 1$ is a parameter of the decoder alone.

Definition 4.5 (Exact IRS parameter regions)**.** Define

$$\mathcal{R}_1(N, k) \coloneqq \{r \ge 1 : (k-1)r(r+1) \le 2(N-1)\},$$

$$\mathcal{R}_2(N, k) \coloneqq \{r \ge 1 : (k-1)r(r+1) > 2(N-1),\ r(k-1) \le N-1\},$$

where r is integral, and let

$$g_1(N, k; r) \coloneqq \frac{r(N-1)}{r+1} - \frac{r(k-1)}{2}, \tag{4.9}$$

$$g_2(N, k; r) \coloneqq N - r(k-1) - 1. \tag{4.10}$$

Proposition 4.6 (Exact IRS specialization)**.**
For $r \in \mathcal{R}_1$, Yayla's theorem covers $f < g_1(N, k; r)$ with decoder success probability at least

$$1 - \eta_1(N, k, p; r), \qquad \eta_1 \coloneqq g_1(N, k; r)/p.$$

For $r \in \mathcal{R}_2$, it covers $f \le g_2(N, k; r)$ with probability at least

$$1 - \eta_2(N, k, p; r), \qquad \eta_2 \coloneqq g_2(N, k; r)/p.$$

These probabilities require random messages and random error values. The stated complexity is $\mathcal{O}(N \cdot (Nr)^3)$ field operations.

Proof. For $r \in \mathcal{R}_1(N, k)$, Yayla's first regime has lifted code rate

$$R_0(r) = \frac{(k-1)(r+1) + 2}{2N}$$

and requires

$$\frac{f}{N} < \frac{r}{r+1}\big(1 - R_0(r)\big).$$

Multiplication by N and simplification give

$$f < \frac{r(N-1)}{r+1} - \frac{r(k-1)}{2} = g_1(N,k;r).$$

The accompanying failure term is g_1/p, and the source range is equivalent to $(k-1)r(r+1) \leq 2(N-1)$. In the second regime,

$$\frac{f}{N} \leq 1 - \frac{r(k-1)+1}{N}$$

is $f \leq g_2(N,k;r) = N - r(k-1) - 1$, with failure term g_2/p and dimension condition $r(k-1) \leq N-1$. These are the exact integer forms of [Yay11, Theorem 3.3.2]. □

Let $\lfloor x \rfloor_< \coloneqq \lceil x \rceil - 1$ be the largest integer strictly below x. The finite projected ceiling is

$$F_1^{\mathrm{IRS}}(N,\,k) \coloneqq \max_{r \in \mathcal{R}_1} \lfloor g_1(N,\,k;\,r) \rfloor_<, \tag{4.11}$$

$$F_2^{\mathrm{IRS}}(N,\,k) \coloneqq \max_{r \in \mathcal{R}_2} g_2(N,\,k;\,r), \tag{4.12}$$

$$F_{\mathrm{IRS}}(N,\,\tau) \coloneqq \max\{F_1^{\mathrm{IRS}}(N,\,N-\tau), F_2^{\mathrm{IRS}}(N,\,N-\tau)\}, \tag{4.13}$$

with the second branch omitted when $\mathcal{R}_2$ is empty. The two rounding conventions preserve the strict first regime and weak second regime inequalities.

The maximizing value of r is small for a structural reason. As a real function,

$$g_1'(r) = \frac{N-1}{(r+1)^2} - \frac{k-1}{2}, \qquad g_1''(r) = -\frac{2(N-1)}{(r+1)^3} < 0,$$

so its unique continuous maximizer is

$$r_1^\star = \sqrt{\frac{2(N-1)}{k-1}} - 1.$$

The discrete optimum is adjacent to $r_1^\star$, up to ties from strict integer rounding, while g_2 is decreasing. For each of the six publication parameter pairs, $1.45 < r_1^\star < 1.51$, $\mathcal{R}_1 = \{1,\,2\}$, and $\mathcal{R}_2 = \{3\}$. Increasing r also adds components of larger dimension, up to $k_r = r(k-1)+1$, and this dimension penalty, which is the term $-r(k-1)/2$ in g_1 and the term $-r(k-1)$ in g_2, dominates beyond the displayed optimum. Accepting a larger failure probability cannot independently extend the radius because both quantities are fixed by the same source theorem.

Yayla's IRS decoder runs Algorithm 3 for successive agreement parameters t, and Algorithm 3 returns one reconstructed tuple or FAILURE, so the decoder outputs at most one tuple [Yay11, Section 3.2 and Algorithm 3]. Since the theorem hypotheses imply $r(k-1) \leq N-1$, the $\mathcal{O}(N^4 \cdot r^3)$ cost of this decoder [Yay11, Theorem 3.3.2] is at most $\mathcal{O}(N^7)$ field operations. Its application therefore turns on its radius and on the transfer of its source distribution to the attack target.

4.6.3 Distribution transfer and additive composition

Let $\mathcal{T}_f \coloneqq \{F \subseteq [N] : |F| = f\}$, and define

$$E_F \coloneqq \{v \in \mathrm{span}(H_F)\}, \qquad E_f \coloneqq \bigcup_{F \in \mathcal{T}_f} E_F.$$

In the original uniform attack experiment, the target is uniform in $\mathbb{F}_p^\tau$ by lemma 7.2, and the distribution transfer and the additive bound below are statements about this experiment. Full spark then gives $\Pr[E_F] = p^{f-\tau}$ for every $F \in \mathcal{T}_f$, and these probabilities add up to the first moment μ_f of section 4.5. Under $N\tau \le 2^\lambda$ and $p \ge 2^{2\lambda}$, the second moment calculation of Crites and Stewart [CS25, Theorem 2 and its proof] gives

$$\Pr[E_f] \ge L_{\mathrm{att}} \coloneqq \frac{\binom{N}{f}}{p^{\tau-f}(1+\vartheta_\lambda) + \binom{N}{f}}, \qquad \vartheta_\lambda = \mathcal{O}(2^{-\lambda}). \tag{4.14}$$

Lemma 4.7 (Supportwise IRS distribution transfer)**.**

Fix $F \in \mathcal{T}_f$. Choose $u_0(v)$ with $Hu_0(v) = v$, sample $\tilde{c}$ uniformly from C_H, and set $u = u_0(v) + \tilde{c}$. Conditioned on E_F, let a_F be the unique vector supported in F satisfying $Ha_F = v$. Then $D_w^{-1}(u - a_F)$ is uniform in $\mathrm{RS}(\mathbb{F}_p, z, k)$, the normalized values $(D_w^{-1}a_F)_F$ are uniform in $\mathbb{F}_p^f$, and the two variables are independent.

Proof. Conditioned on E_F, the bijection $H_F : \mathbb{F}_p^f \to \operatorname{im} H_F$ makes a_F uniform. For every fixed a_F, the vector $u_0(v) - a_F + \tilde{c}$ is uniform in C_H, independently of a_F. The invertible map D_w^{-1} preserves both properties. □

Zero coefficients may be removed from the declared support, because the actual nonzero error values are still uniform and the error weight only decreases. For an admissible r, let $\eta_{\mathrm{IRS}}(r)$ be the corresponding η_1 or η_2. We write $[x]_+ = \max\{x, 0\}$ for the positive part.

Theorem 4.8 (Additive bound for the IRS assisted attack)**.** *Under the hypotheses of equation (4.14), suppose an admissible r covers f, and use the randomized coset representative of lemma 4.7. Then*

$$\Pr[\ \textit{IRS returns a valid sparse witness}\ ] \ge [\, L_{\mathrm{att}} - \eta_{\mathrm{IRS}}(r) \cdot \mu_f \,]_+. \tag{4.15}$$

Proof. For each support F, lemma 4.7 provides conditional on E_F the random codeword and random error values that proposition 4.6 requires, so the event that E_F occurs and the decoder fails on the associated codeword has probability at most $\eta_{\mathrm{IRS}}(r) \cdot p^{f-\tau}$. A union bound over $\binom{N}{f}$ supports gives total loss at most $\eta_{\mathrm{IRS}}(r) \cdot \mu_f$. Subtracting this from equation (4.14) and truncating at zero proves the claim. □

The support events overlap, so the loss is additive, whereas the expression $L_{\mathrm{att}}(1 - \eta_{\mathrm{IRS}})$ would require an independence statement that the sources leave unproved. A valid sparse witness exists only on E_f, which is the event $E_{\mathcal{Q}}$ of definition 4.3 for $\mathcal{Q} = P_{H,f}$, and $\Pr[E_f] \ge L_{\mathrm{att}} > 0$, so dividing equation (4.15) by $\Pr[E_f] \le 1$ gives $\Pr[D_{\mathfrak{D}} \mid E_f] \ge [\, L_{\mathrm{att}} - \eta_{\mathrm{IRS}}(r) \cdot \mu_f \,]_+$ for the IRS route $\mathfrak{D}$. The guarantee is therefore distributional in the sense of definition 4.3, for the uniform target of lemma 7.2 and the uniform coset sample of lemma 4.7, and a zero lower bound is inconclusive.

4.6.4 Guruswami–Sudan

The original Guruswami–Sudan theorem for degree at most $k-1$ gives the finite agreement condition

$$N - f > \sqrt{N(k-1)}$$

[GS99, Theorem 8 and Proposition 9]. Hence

$$F_{\mathrm{GS}}^{\mathrm{fin}} \coloneqq N - \left\lfloor \sqrt{N(N-\tau-1)} \right\rfloor - 1. \tag{4.16}$$

Inside this radius the Guruswami–Sudan decoder enumerates the complete list, whose source bound is

$$\mathcal{L}_{\mathrm{GS}} = \mathcal{O}\left(\sqrt{(k-1)N^3}\right) \le \mathcal{O}(N^2).$$

The source gives an $\mathcal{O}(N^{15})$ implementation under its stated field size qualification and $\mathcal{O}(N^3 \cdot \delta^{-6})$ when $(N-f)^2 = (1+\delta)\,N\,(k-1)$ [GS99, Theorem 12]. For cryptographic prime fields, its randomized root finding option is polynomial in $\log p$ [GS99, Section 1]. Each returned codeword is converted and verified through proposition 4.4, giving a pointwise complete realization of $P_{H,f}$ inside the radius. The rate form specialization $(N-f)^2 > N(N-\tau)$ is a conservative simplification, and every finite overlap statement uses equation (4.16).

5 Folded decoding and output contracts

Folded Reed–Solomon (FRS) decoding groups several consecutive evaluations of the same polynomial into one symbol and exploits the resulting algebraic relations to approach the list decoding capacity with a controlled list [GR08]. Relative to ordinary RS list decoding, this construction changes both the code setup and the decoding metric.

This section develops the conditional decoder map for the block restricted relation, the exact count of block compatible supports, the witness existence analysis for a uniform target, and how the coordinate weight limits the gap below capacity. Section 5.4 then states the contracts of the four folded decoders, and section 5.5 compares the contracts of all seven routes.

5.1 Setup, block metric, and restricted oracle

We use the aligned partition, $L = N/m$, and block weight wt_m fixed in section 4, with the support bounds of equation (4.4). Coding sources use q for the field size, which we specialize to the attack scalar field order p, and every $q(X)$ below is an explicitly named polynomial. For route $\mathfrak{D}$, let $\mathsf{SetupComp}_{\mathfrak{D}}(N, \tau, p, z)$ denote diagonal normalization and the source's ordered evaluation domain hypotheses. Original GR uses a primitive γ and all powers up to its native length [GR08, Definition 2.1 and Theorem 4.2]. GW permits a consecutive prefix of length $N \le p-1$ [GW13, Definition 1 and Theorem 7]. KRZSW permits distinct block bases among powers γ^{mi} [KRZSW23, Section 2 and Corollary 3.2]. AHS requires $p > N$, an element of order at least N, and a consecutive power prefix [AHS26, Definition 2.5 and Theorem 4.4].

Each route therefore imposes structure on the Shamir identifiers and their ordering. For the capability analysis in this paper, we make the standing conditional setup assumption that p is prime, $p > N$, and the ordered Shamir evaluation domain has the consecutive power structure required by the decoder route under consideration. Since $\mathbb{F}_p^{\times}$ is cyclic, this assumption supplies a primitive element and enough distinct powers for the tested lengths. We accordingly classify the four setup predicates as *conditional*, that is, as hypotheses that a concrete deployment would have to meet.

This domain needs a genuine prime field, and the numerical proxy 2^{252} of section 6.4 stands only for its size. Original GR additionally uses the internal prefix transfer result of proposition 5.7, whose list and runtime keep their dependence on p.

A relative block radius $1-\beta-\epsilon$ lies $\epsilon > 0$ below the capacity value $1-\beta = \tau/N$, and we call this gap ϵ the slack of the radius. In the inclusive rate form the radius corresponds to the block budget

$$B_{\mathrm{FRS}}(N, \tau; m, \epsilon) := \left\lfloor \frac{\tau - \epsilon N}{m} \right\rfloor. \tag{5.1}$$

Proposition 5.1 (Coordinate to block projection)**.** *Assume compatible setup. A fixed vector a lies inside some positive slack rate form FRS radius exactly when*

$$m \cdot \mathrm{wt}_m(a) < \tau.$$

When only $\mathrm{wt}(a) \le f$ *is known, the condition* $f < \tau/m$ *covers every such support. An exact* f *support can instead be packed into* $\lceil f/m \rceil$ *blocks, and this best geometry can have positive slack only if* $m\lceil f/m \rceil < \tau$.

Proof. A vector lies inside the rate form radius at slack ϵ exactly when $\mathrm{wt}_m(a)/L \le 1-\beta-\epsilon$. Since $L = N/m$ and $1-\beta = \tau/N$, this means $m \cdot \mathrm{wt}_m(a) \le \tau - \epsilon \cdot N$, which holds for some $\epsilon > 0$ exactly when $m \cdot \mathrm{wt}_m(a) < \tau$. The two other statements follow from equation (4.4), whose lower bound is attained by a support packed into $\lceil f/m \rceil$ blocks. □

The corresponding block restricted relation is equation (4.5), and the next theorem gives its conditional decoder realization.

Theorem 5.2 (Conditional folded realization)**.** *Let* $B = B_{\mathrm{FRS}}(N, \tau; m, \epsilon)$*, assume compatible setup, and suppose a complete FRS decoder enumerates every folded codeword within block distance* B*. If* $P_{m,B}(v)$ *is solvable, compute any* u *with* $Hu = v$*, decode the folding of* $y = D_w^{-1}u$*, and retain every returned* $c \in \mathrm{RS}(\mathbb{F}_p, z, k)$ *satisfying*

$$d_H(y, c) \le f, \qquad \mathrm{wt}_m(y - c) \le B.$$

Then $a = u - D_w c$*, for any retained* c*, is a valid* $P_{m,B}(v)$ *witness, and at least one such candidate is retained.*

Proof. For the conditional decoder map, take a restricted witness $a^\star$ and any u with $Hu = v$. Then $c_H^\star = u - a^\star$ lies in C_H, so

$$c^\star = D_w^{-1} c_H^\star \in \mathrm{RS}(\mathbb{F}_p, z, k), \qquad D_w^{-1}u - c^\star = D_w^{-1}a^\star.$$

Diagonal scaling preserves coordinate and block support. Complete folded decoding therefore returns $c^\star$. Conversely, each retained c gives $a = u - D_w c$ with $Ha = v$ and the required weights. □

5.2 Support counts, lists, and witness probability

Let $\mathcal{T}_{m,B,f}$ be the exact f supports touching at most B blocks

$$\mathcal{T}_{m,B,f} \coloneqq \Big\{ F \subseteq [N] : |F| = f,\ \#\{j : F \cap \mathcal{B}_j \neq \varnothing\} \le B \Big\},$$

and write $M_{m,B,f} = |\mathcal{T}_{m,B,f}|$.

Proposition 5.3 (Exact block compatible support count)**.**

$$M_{m,B,f} = \sum_{b=\lceil f/m \rceil}^{\min\{B,\,L\}} \binom{L}{b} \sum_{j=0}^{b} (-1)^j \binom{b}{j} \binom{m(b-j)}{f}. \tag{5.2}$$

In particular, $M_{m,B,f} > 0$ *exactly when* $f \le mB$.

Proof. For a support touching exactly b folded blocks, choose the blocks in $\binom{L}{b}$ ways. The number of exact f subsets of the union of the b selected blocks that meet each of these blocks is, by inclusion–exclusion over the blocks that a subset misses,

$$\sum_{j=0}^{b} (-1)^j \binom{b}{j} \binom{m(b-j)}{f}.$$

Summing over $\lceil f/m \rceil \le b \le \min\{B, L\}$ proves the exact formula for $M_{m,B,f}$. Such a support exists exactly when $\lceil f/m \rceil \le B$, equivalently $f \le mB$. □

The density

$$\pi_{m,B,f} \coloneqq \frac{M_{m,B,f}}{\binom{N}{f}} \tag{5.3}$$

is a ratio of two counts and no independently justified factor by which an unrestricted witness probability may be multiplied, so the probability bounds below use the restricted count $M_{m,B,f}$ and its moments. The expectations and probabilities in the rest of this subsection use a target v uniform in $\mathbb{F}_p^\tau$, as in the original uniform attack experiment of lemma 7.2.

For

$$Z_{m,B,j}(v) \coloneqq \#\{a : Ha = v,\ \mathrm{wt}(a) = j,\ \mathrm{wt}_m(a) \le B\},$$

exact counting gives

$$\mathbb{E}[Z_{m,B,j}] = M_{m,B,j}\frac{(p-1)^j}{p^\tau}, \qquad \mathbb{E}[Z_{m,B,\le f}] = \frac{1}{p^\tau}\sum_{j=0}^{f} M_{m,B,j}(p-1)^j. \tag{5.4}$$

These are the average numbers of outputs satisfying both the block and coordinate weight bounds. A complete folded decoder at block radius B first returns the unfiltered block radius list, which may include errors of coordinate weight as large as mB. Its exact average block list size is

$$\mathbb{E}[Z_{m,B}^{\mathrm{block}}] = \frac{1}{p^\tau}\sum_{b=0}^{B}\binom{L}{b}(p^m-1)^b. \tag{5.5}$$

Thus the support density and the restricted moments in equation (5.4) describe only the supports and outputs that the weight filters admit, whereas a complete decoder returns the whole list counted in equation (5.5). The two output lower bounds of section 4.5, on records and on bits in equation (4.8), therefore hold for a complete folded decoder with this unfiltered list in place of $Z_{\le f}$, that is, before the coordinate weight filter.

Theorem 5.4 (Restricted witness existence lower bound)**.** *Assume $0 \le B \le L$, $1 \le f \le \tau - 1$, $f \le mB$, $N\tau < p$, and uniform $v \in \mathbb{F}_p^\tau$. Then*

$$\Pr[P_{m,B}(v) \text{ is solvable}] \ge \frac{M_{m,B,f}}{p^{\tau-f}/(1 - N\tau/p) + M_{m,B,f}}. \tag{5.6}$$

Proof. For F in $\mathcal{T}_{m,B,f}$, write $V_F = \mathrm{span}(H_F)$ and let

$$X(v) = \sum_{F \in \mathcal{T}_{m,B,f}} \mathbf{1}_{\{v \in V_F\}}.$$

Full spark gives

$$\mu \coloneqq \mathbb{E}[X] = M_{m,B,f}p^{f-\tau}.$$

Write $d = \tau - f$. For exact f supports F, G, let $h = |F \setminus G| = |G \setminus F|$. Since $|F \cup G| = f + h$,

$$\dim(V_F \cap V_G) = 2f - \min\{\tau, f + h\}.$$

Thus the two membership indicators have zero covariance for $h \ge d$. For $h < d$, their covariance is at most $p^{-(d+h)}$. For fixed F, there are at most $\binom{f}{h}\binom{N-f}{h}$ unrestricted size f sets G at this distance. Dropping the block restriction on G gives

$$\mathrm{Var}(X) \le \mu\sum_{h \ge 0}\binom{f}{h}\binom{N-f}{h}p^{-h}.$$

The ratio of consecutive summands is at most

$$\frac{f(N-f)}{p} \le \frac{N\tau}{p} < 1,$$

and hence

$$\mathrm{Var}(X) \le \frac{\mu}{1 - N\tau/p}.$$

The second moment inequality now gives

$$\Pr[X > 0] \ge \frac{\mu^2}{\mu^2 + \mu/(1 - N\tau/p)} = \frac{M_{m,B,f}}{p^{\tau-f}/(1 - N\tau/p) + M_{m,B,f}}.$$

Finally, any solution of coordinate weight below f and block weight at most B can be enlarged with zero coefficients to an exact f set inside at most B blocks because $f \le mB$. Conversely, $v \in V_F$ supplies coefficients supported on F. Therefore $X > 0$ is exactly the restricted oracle solvability event. □

The result is a statement about witness existence, and throughout this paper we denote its lower bound by

$$L_{\text{FRS-att}}(N,\, \tau,\, p,\, f;\, m,\, B) \coloneqq \frac{M_{m,B,f}}{p^{\tau-f}/(1 - N\tau/p) + M_{m,B,f}}.$$

Theorem 5.5 (Folded attack distribution success)**.** *Assume compatible setup, $N\tau < p$, a fresh uniform attack target, $B = B_{\mathrm{FRS}}(N,\, \tau;\, m,\, \epsilon)$, and the hypotheses of theorem 5.4. If a complete decoder returns the radius B list, conditional on solvability, with probability at least $q^{\mathrm{thm}}_{\mathrm{FRS}}$, and the complete procedure has bit complexity polynomial along the parameter family under consideration, then it returns a valid sparse witness with probability at least*

$$q^{\mathrm{thm}}_{\mathrm{FRS}} \cdot L_{\text{FRS-att}}(N,\, \tau,\, p,\, f;\, m,\, B). \tag{5.7}$$

Proof. Let E be the event that $P_{m,B}(v)$ is solvable and D the event that the complete procedure returns and validates a witness. On E, whenever the decoder returns the complete radius B list, theorem 5.2 shows that at least one candidate is retained and that each retained candidate gives a valid $P_{m,B}(v)$ witness, so $\Pr[D \mid E] \ge q^{\mathrm{thm}}_{\mathrm{FRS}}$. By theorem 5.4, $\Pr[E] \ge L_{\text{FRS-att}}(N,\, \tau,\, p,\, f;\, m,\, B)$, which is positive because $f \le mB$ gives $M_{m,B,f} > 0$ in proposition 5.3. The chain rule $\Pr[D] \ge \Pr[E \cap D] = \Pr[E] \cdot \Pr[D \mid E]$ then gives the stated bound. □

The proof uses the chain rule, which holds whatever the dependence between E and D. Under the standing Crites–Stewart scheme interface, a returned restricted witness also satisfies the solver antecedent in theorem 2.3. A sampled decoder setup, in place of the fixed one assumed here, would bring its compatibility event into the probability accounting.

5.3 Coordinate weight and the gap below capacity

Near the threshold, an exact weight witness leaves only a small slack, as the next proposition shows.

Proposition 5.6 (Slack forced by an exact weight witness)**.** *If a witness a lies inside the rate form radius with slack $\epsilon > 0$, then*

$$\epsilon \le \frac{\tau - m\,\mathrm{wt}_m(a)}{N} \le \frac{\tau - \mathrm{wt}(a)}{N}.$$

Consequently, $\mathrm{wt}(a) = \tau - g$ *forces* $\epsilon \le g/N$.

Proof. Coverage at slack ϵ means $m \cdot \mathrm{wt}_m(a) \le \tau - \epsilon \cdot N$, which is the first inequality. The second follows from $m \cdot \mathrm{wt}_m(a) \ge \mathrm{wt}(a)$, itself a consequence of equation (4.4). Substituting $\mathrm{wt}(a) = \tau - g$ gives $\epsilon \le g/N$. □

5.4 Folded decoder contracts

We specialize the four folded routes under the compatible setup assumed at the start of this section, and for each route we record its radius, output and complete recovery cost. The witness probability bound of theorem 5.10 below carries the uniform target in its own hypotheses.

5.4.1 Original Guruswami–Rudra

Let $d = k - 1$, with interpolation arity s_{int} and multiplicity r_{GR}. Define

$$n_0 \coloneqq L(m - s_{\mathrm{int}} + 1),$$

$$D_{\mathrm{GR}} \coloneqq \left\lfloor \left(d^{s_{\mathrm{int}}} n_0 \prod_{j=0}^{s_{\mathrm{int}}} (r_{\mathrm{GR}} + j) \right)^{1/(s_{\mathrm{int}}+1)} \right\rfloor + 1. \tag{5.8}$$

The weighted degree D_{GR} is the choice in Equation (5) of Guruswami and Rudra. With this degree choice, Lemma 4.1 and the derivation preceding their Theorem 4.2 give the following sufficient agreement condition for an integer t [GR08, Lemma 4.1, Equation (5) and the derivation preceding Theorem 4.2][1]

$$t \geq \left(L \left(\frac{d}{m - s_{\mathrm{int}} + 1} \right)^{s_{\mathrm{int}}} \prod_{j=1}^{s_{\mathrm{int}}} \left(1 + \frac{j}{r_{\mathrm{GR}}} \right) \right)^{1/(s_{\mathrm{int}}+1)} + 2. \tag{5.9}$$

The resulting block budget is $B_{\mathrm{GR}} = L - t$. The equivalent integer predicate used by the artifact of appendix C is justified in the proof of proposition 5.7. For integer $t \geq 2$, it reads

$$(t-2)^{s_{\mathrm{int}}+1} \cdot (m - s_{\mathrm{int}} + 1)^{s_{\mathrm{int}}} \cdot r_{\mathrm{GR}}^{s_{\mathrm{int}}} \geq L \cdot d^{s_{\mathrm{int}}} \prod_{j=1}^{s_{\mathrm{int}}} (r_{\mathrm{GR}} + j).$$

For each (m, s_{int}), the artifact selects the smallest finite r_{GR} attaining the best integer agreement, and later maximizes mB_{GR}. The finite agreement condition also gives

$$t(m - s_{\mathrm{int}} + 1) \cdot r_{\mathrm{GR}} > D_{\mathrm{GR}},$$

which is the strict inequality used in the agreement implication of GR Lemma 4.1.

Proposition 5.7 (Original GR prefix transfer)**.** *Let p be prime, let $\gamma \in \mathbb{F}_p^\times$ be primitive, and assume*

$$1 \leq N \leq p - 1, \quad m \mid N, \quad 1 \leq d < N, \quad 1 \leq s_{\mathrm{int}} \leq m, \quad r_{\mathrm{GR}} \in \mathbb{Z}_{\geq 1}.$$

Let $\bar{y} \in (\mathbb{F}_p^m)^L$ be a received folded word and let $t \in \{0, 1, \ldots, L\}$ be an integer. On the consecutive prefix $1, \gamma, \ldots, \gamma^{N-1}$, if

$$\lfloor D_{\mathrm{GR}}/d \rfloor < p \tag{5.10}$$

and t satisfies equation (5.9), the Original GR interpolation and root procedure outputs every degree at most d message polynomial whose folded encoding agrees with $\bar{y}$ in at least t blocks. Interpolation is polynomial in N and $r_{\mathrm{GR}}^{s_{\mathrm{int}}}$, the list has size at most $p^{s_{\mathrm{int}}}$, and the cited root procedure costs $p^{\mathcal{O}(s_{\mathrm{int}})}$.

[1] In the cited arXiv version, Equation (6) places the block length inside the s_{int}th power. We use the preceding derivation with $n_0 = L(m - s_{\mathrm{int}} + 1)$, which places the block length outside that power.

Proof. The source owned components are GR Lemma 4.1's interpolation count and agreement implication, GR Lemma 3.5's extension field identity, and GR Theorem 3.6's root procedure. The internal component is the verification that these arguments hold on a consecutive prefix under the explicit root safety predicate of equation (5.10). Correctness transfers to the prefix, and the output and root costs keep the dependence on the field size that they have in the source.

With $d = k - 1$ and n_0 as in equation (5.8), write the radical in D_{GR} as

$$\chi_{\mathrm{GR}} = \left(d^{s_{\mathrm{int}}} n_0 \prod_{j=0}^{s_{\mathrm{int}}} (r_{\mathrm{GR}} + j) \right)^{1/(s_{\mathrm{int}}+1)}, \qquad D_{\mathrm{GR}} = \lfloor \chi_{\mathrm{GR}} \rfloor + 1.$$

For integer $t \geq 2$, the finite agreement condition is equivalent to

$$(t-2)^{s_{\mathrm{int}}+1} (m - s_{\mathrm{int}} + 1)^{s_{\mathrm{int}}} r_{\mathrm{GR}}^{s_{\mathrm{int}}} \geq L d^{s_{\mathrm{int}}} \prod_{j=1}^{s_{\mathrm{int}}} (r_{\mathrm{GR}} + j),$$

which is the exact integer predicate evaluated by the artifact. The root safety condition of equation (5.10) holds by hypothesis. The retained powers $1, \gamma, \ldots, \gamma^{N-1}$ are distinct. Moreover,

$$D_{\mathrm{GR}}^{s_{\mathrm{int}}+1} > d^{s_{\mathrm{int}}} n_0 \prod_{j=0}^{s_{\mathrm{int}}} (r_{\mathrm{GR}} + j),$$

which is the sufficient monomial count inequality in [GR08, Lemma 4.1]. It depends on the number n_0 of retained interpolation tuples, not on the native length identity, so the interpolation step of that lemma produces a nonzero $Q(X, Y_1, \ldots, Y_{s_{\mathrm{int}}})$ of weighted degree at most D_{GR}.

Each agreeing folded block supplies $m - s_{\mathrm{int}} + 1$ tuples with multiplicity r_{GR}. Set

$$c = (m - s_{\mathrm{int}} + 1) r_{\mathrm{GR}}.$$

The radical in the finite agreement bound, before its final $+2$, equals χ_{GR}/c. Thus that bound gives

$$ct \geq \chi_{\mathrm{GR}} + 2c > \lfloor \chi_{\mathrm{GR}} \rfloor + 1 = D_{\mathrm{GR}}.$$

For every close message polynomial F, the substituted polynomial has more zeros counted with multiplicity than its degree, and hence

$$Q\big(X, F(X), F(\gamma X), \ldots, F(\gamma^{s_{\mathrm{int}}-1} X)\big) = 0.$$

For the root stage, let $E(X) = X^{p-1} - \gamma$, and let Q_0 be the polynomial obtained by dividing all coefficient polynomials of Q by their largest common power of E. At least one coefficient stays nonzero modulo E. Primitivity makes E irreducible by GR Lemma 3.5, so reduction gives a nonzero polynomial over $\widetilde{\mathbb{F}} = \mathbb{F}_p[X]/(E)$. Since $d < N \leq p - 1$,

$$F(\gamma^j X) = F(X)^{p^j} \pmod{E(X)}.$$

The root safety inequality bounds the total degree in Y by a value below p, so exponent vectors have unique base p encodings after substituting

$$(Y_1, \ldots, Y_{s_{\mathrm{int}}}) = (Y, Y^p, \ldots, Y^{p^{s_{\mathrm{int}}-1}}).$$

The resulting nonzero univariate polynomial has degree below $p^{s_{\mathrm{int}}}$. The source root procedure therefore returns at most $p^{s_{\mathrm{int}}}$ candidates in $p^{\mathcal{O}(s_{\mathrm{int}})}$ time. Degree and distance filtering preserves every close codeword. □

Proposition 5.7 therefore gives the complete prefix recovery cost

$$T_{\mathrm{GR}}^{\mathrm{prefix}} = \mathrm{poly}(N, r_{\mathrm{GR}}^{s_{\mathrm{int}}}, b_p) + p^{\mathcal{O}(s_{\mathrm{int}})}.$$

On the source native family $N = m\lfloor (p-1)/m \rfloor$, one has $p \leq N + m$. For bounded s_{int} and polynomially bounded $r_{\mathrm{GR}}^{s_{\mathrm{int}}}$, this is polynomial in the source code length. For $b_p = \Theta(\lambda)$, $N = \mathrm{poly}(\lambda)$ and $N \ll p$, this cost lies outside cryptographic bit PPT, because prefix correctness leaves the list and root costs dependent on p. The capacity corollary's $(L/\epsilon^2)^{\mathcal{O}(\epsilon^{-1}\log(1/\beta))}$ bound is likewise not uniformly polynomial after substituting $\epsilon = \mathcal{O}(1/N)$ [GR08, Theorem 4.4].

5.4.2 Guruswami–Wang

For $1 \leq s_{\mathrm{GW}} \leq m$, GW return, in

$$\mathcal{O}(N^2 + s_{\mathrm{GW}} k^2) \tag{5.11}$$

field operations, an affine space of dimension $h \leq s_{\mathrm{GW}} - 1$ containing every codeword within

$$\rho_{\mathrm{GW}} \coloneqq \frac{s_{\mathrm{GW}}}{s_{\mathrm{GW}} + 1} \cdot \left(1 - \frac{k}{L(m - s_{\mathrm{GW}} + 1)}\right), \qquad B_{\mathrm{GW}} \coloneqq \lfloor L \cdot \rho_{\mathrm{GW}} \rfloor \tag{5.12}$$

when the radius is nonnegative [GW13, Lemmas 3 and 4, Observation 5, Lemma 6, and Theorem 7]. The source output is an affine basis, from which the close codewords still have to be enumerated.

Proposition 5.8 (Elementary coordinate subset pruning)**.** *Let $z_1, z_2, \ldots, z_N \in \mathbb{F}_p$ be distinct, and let $\mathcal{A} \subseteq \mathbb{F}_p[X]_{<k}$ be an affine polynomial space of dimension h. Given a received folded word $\bar{y} \in (\mathbb{F}_p^m)^L$ and an integer $t \in \{0, 1, \ldots, L\}$, if $mt \geq k$, all members of $\mathcal{A}$ whose folded encodings agree with $\bar{y}$ in at least t blocks can be enumerated in*

$$\binom{N}{h} \cdot \mathrm{poly}(N, k, m, \log p)$$

bit operations, and their number is at most $\binom{N}{h}$.

Proof. GW Lemmas 3 and 4, Observation 5, Lemma 6, and Theorem 7 own the affine space containment and basis computation, and the complete close codeword list is left to the coordinate subset enumeration below, the internal, elementary completion step.

Let $z_1, \ldots, z_N \in \mathbb{F}_p$ be distinct. Let $\mathcal{A} = F_0 + U \subseteq \mathbb{F}_p[X]_{<k}$ be an affine space with direction basis $U_1, \ldots, U_h$. Fix a received folded word $\bar{y}$ and an integer $t \in \{0, \ldots, L\}$, and write

$$F_x = F_0 + \sum_{j=1}^{h} x_j U_j.$$

Suppose every close member agrees with $\bar{y}$ in at least t folded blocks and $mt \geq k$. Its mt ordinary evaluation equations in $x_1, \ldots, x_h$ have rank h. Otherwise, a nonzero polynomial in U would vanish at $mt \geq k$ distinct points despite having degree below k. Some h agreement equations therefore determine x uniquely.

Enumerate the subsets of cardinality h of the N ordinary coordinates, solve each full rank system, and verify the folded distance. This returns every close member in

$$\binom{N}{h} \mathrm{poly}(N, k, m, \log p)$$

bit operations. Assigning each close member its first independent subset gives the same upper bound on list size. For $h = 0$, one verifies the unique affine point.

For the GW radius, write $R = k/N$ and $a_{\mathrm{GW}} = 1 - \rho_{\mathrm{GW}}$. Expanding the definition gives

$$a_{\mathrm{GW}} = \frac{1}{s_{\mathrm{GW}} + 1} + \frac{s_{\mathrm{GW}}}{s_{\mathrm{GW}} + 1} \frac{mR}{m - s_{\mathrm{GW}} + 1}.$$

Let $h_0 = m - s_{\mathrm{GW}} + 1$. The product $(s_{\mathrm{GW}} + 1)h_0(a_{\mathrm{GW}} - R)$ equals

$$h_0 + R(s_{\mathrm{GW}}^2 - 1 - m).$$

If the coefficient of R is negative, nonnegative GW radius gives $R \leq h_0/m$, so the product is at least $h_0(s_{\mathrm{GW}}^2 - 1)/m > 0$ for $s_{\mathrm{GW}} \geq 2$. For $s_{\mathrm{GW}} = 1$, the difference is $(1 - R)/2 > 0$. If the coefficient is nonnegative, the product is at least h_0. Hence $a_{\mathrm{GW}} > R$, and every covered word has $mt > k$.

The exact relative block distance of the folded code is

$$\Delta_{\mathrm{FRS}} = 1 - \frac{\lfloor (k-1)/m \rfloor}{L}.$$

A nonzero polynomial of degree below k vanishes at no more than $k-1$ ordinary evaluation points and hence on no more than $\lfloor (k-1)/m \rfloor$ complete blocks. Conversely, the product of the linear factors in that many complete blocks has degree at most $k - 1$ and attains the bound. Therefore

$$\Delta_{\mathrm{FRS}} > 1 - R > \rho_{\mathrm{GW}}.$$

For

$$\epsilon_{\mathrm{pr}} = \frac{\Delta_{\mathrm{FRS}} - \rho_{\mathrm{GW}}}{2},$$

AHS Theorem 3.5 returns every member of the encoded GW affine space at distance strictly below

$$\Delta_{\mathrm{FRS}} - \epsilon_{\mathrm{pr}} = \rho_{\mathrm{GW}} + \frac{\Delta_{\mathrm{FRS}} - \rho_{\mathrm{GW}}}{2} > \rho_{\mathrm{GW}}.$$

Exact distance filtering thus preserves the complete non-strict GW radius list, including boundary candidates [AHS26, Definitions 2.3 and 2.4 and Theorem 3.5]. □

This exhaustive pruning is an elementary internal step of cost $N^{\mathcal{O}(h)}$. The nonnegative GW radius implies $mt > k$, so proposition 5.8 applies to the GW affine space, and $h \leq s_{\mathrm{GW}} - 1$ gives the cost $N^{\mathcal{O}(s_{\mathrm{GW}})} \operatorname{poly}(\log p)$, uniformly bit polynomial for bounded s_{GW}.

The preceding proof gives the exact folded distance Δ_{FRS} and shows that it exceeds ρ_{GW}. With $\epsilon_{\mathrm{pr}} = (\Delta_{\mathrm{FRS}} - \rho_{\mathrm{GW}})/2$, AHS Theorem 3.5 gives a second, source supported completion at each fixed positive gap [AHS26, Theorem 3.5], of cost

$$\operatorname{poly}(N,\, s_{\mathrm{GW}},\, k,\, \log p) + N \cdot \operatorname{polylog}(p) \cdot \left(\frac{2}{\epsilon_{\mathrm{pr}}}\right)^{\mathcal{O}(2^h)}. \tag{5.13}$$

For a shrinking gap family, the uniform bound is the elementary pruning above.

5.4.3 KRZSW

For fixed $\epsilon_{\mathrm{K}} > 0$, KRZSW give a randomized complete decoder of the folded code with degree bound $k-1$ up to relative block distance $1-(k-1)/N-\epsilon_{\mathrm{K}}$ [KRZSW23, Sections 2.2 and 2.5 and Corollary 3.2], provided

$$16/\epsilon_{\mathrm{K}}^2 \leq m. \tag{5.14}$$

The list size is $(1/\epsilon_{\mathrm{K}})^{\mathcal{O}(\epsilon_{\mathrm{K}}^{-1} \cdot \log(1/\epsilon_{\mathrm{K}}))}$, and the complete list is returned with probability at least 0.99 in time polynomial in b_p, m, k, L and the corresponding non-fully polynomial inverse slack factor [KRZSW23, Corollary 3.2 and proof of Theorem 3.1, Section 3.3].

Using source degree $k-1$, define $\delta_b := \tau + 1 - mb$. An integer budget b lies within this radius exactly when $\epsilon_{\mathrm{K}} \le \delta_b/N$, and equation (5.14) gives $4/\sqrt{m} \le \epsilon_{\mathrm{K}}$, so such a slack exists exactly when $\delta_b > 0$ and $m\delta_b^2 \ge 16N^2$, where the squared inequality is weak because the source distance is inclusive. For fixed $m \mid N$ for which the set below is nonempty, the exact largest certified budget is

$$B_{\mathrm{K}} := \max\{b \in \{0, 1, \ldots, L\} : \delta_b > 0,\ m\delta_b^2 \ge 16N^2\}. \tag{5.15}$$

If this set is empty, that divisor contributes no KRZSW region. For $f = \tau - g$, coverage requires $\epsilon_{\mathrm{K}} \le (g+1)/N$. Combining this with equation (5.14) and $m \le N$ yields

$$N \le \frac{(g+1)^2}{16}. \tag{5.16}$$

It fails for every tested $N \ge 64$ and $g \le 4$.

5.4.4 AHS and modern list bounds

Srivastava's $\mathcal{O}(1/\epsilon^2)$ and Chen–Zhang's optimal $\mathcal{O}(1/\epsilon)$ scale results are combinatorial list bounds [Sri25, CZ25], and the randomized complete decoder comes from AHS [AHS26, Theorem 4.4]. For $1 \le s \le m$, define

$$\rho_{s,m}(\beta) := \frac{s}{s+1}\left(1 - \frac{m\beta}{m-s+1}\right), \qquad B_{\mathrm{AHS}} := \lceil L\rho_{s,m}(\beta)\rceil - 1. \tag{5.17}$$

The ceiling minus one structure preserves the source's strict radius. To avoid a collision with the rate $\beta = k/N$, let $\delta_{\mathrm{dec}} \in (0, 1)$ denote the user chosen decoder failure tolerance, which AHS Theorem 4.4 calls β. For a fixed received word, this probability is over the internal randomness of the pruning algorithm. The proof of AHS Theorem 4.4 uses

$$(s(s-1)+1)\cdot \ln(s/\delta_{\mathrm{dec}})$$

independent pruning trials, up to integer rounding. The complete list has size at most s, is returned with probability at least $1 - \delta_{\mathrm{dec}}$, and costs

$$T_{\mathrm{basis}} + \mathcal{O}\left(ms^5 \log(s/\delta_{\mathrm{dec}})L\right) \tag{5.18}$$

field operations, where $T_{\mathrm{basis}} = \widetilde{\mathcal{O}}(mL)\,\mathrm{poly}(s)$ [AHS26, Theorems 2.7, 2.8, and 4.4].

Here T_{basis} is the cost of the first stage, which computes, from the received folded word, a basis for the affine space guaranteed to contain all codewords in the stated radius. The second term of equation (5.18) is the cost of the subsequent randomized pruning. In our notation $mL = N$, so AHS Theorem 2.7 gives $T_{\mathrm{basis}} = \widetilde{\mathcal{O}}(N)\,\mathrm{poly}(s)$ field operations.

The exact capacity deficit is

$$\Delta_{s,m}(\beta) := (1-\beta) - \rho_{s,m}(\beta) = \frac{1-\beta}{s+1} + \frac{\beta s(s-1)}{(s+1)(m-s+1)} \ge 0. \tag{5.19}$$

Proposition 5.9 (AHS exact weight necessary gap)**.** *If an exact weight $f = \tau - g$ support, even optimally packed, lies strictly inside the AHS radius, then*

$$g > N\Delta_{s,m}(\beta) \ge 2\sqrt{\beta(\tau + 2\beta)} - 3\beta. \tag{5.20}$$

Proof. We use the deficit $\Delta_{s,m}(\beta)$ of equation (5.19). Let $f = \tau - g$. An exact weight support touches at least $b_{\min} = \lceil f/m \rceil$ blocks, so

$$\frac{b_{\min}}{L} \ge \frac{f}{N} = (1-\beta) - \frac{g}{N}.$$

Strict AHS coverage yields

$$g > N\Delta_{s,m}(\beta).$$

Since $m - s + 1 \le N$,

$$N\Delta_{s,m}(\beta) \ge \frac{\tau}{s+1} + \beta\frac{s(s-1)}{s+1}.$$

Writing $x = s + 1$, the right side is

$$\beta x - 3\beta + \frac{\tau + 2\beta}{x}.$$

The arithmetic geometric mean inequality gives

$$N\Delta_{s,m}(\beta) \ge 2\sqrt{\beta(\tau + 2\beta)} - 3\beta.$$

□

On the six publication parameter pairs, the rightmost expression in equation (5.20) is at least 6.585, so no tested $g \le 4$ has exact weight overlap, even before setup compatibility is imposed. In the capacity specialization $s \ge 1/\epsilon$, $m \ge s^2$ [AHS26, proof of Corollary 4.5], the bound $m \le N$ gives $\epsilon \ge 1/\sqrt{N}$, and since exact weight coverage requires $\epsilon N < g$, the two inequalities force

$$N < g^2. \tag{5.21}$$

A budget excluding an exact weight can still admit a smaller witness. Equation (7.2), the uniform block bound of proposition 7.3, applies to the AHS budget, and theorem 7.10 gives its quantitative consequences, including $mB_{\mathrm{AHS}} < \tau$.

Theorem 5.10 (Restricted modern FRS support scarcity)**.** *Let $1 \le f = \tau - g \le \tau - 1$, $m \mid N$, $L = N/m$, $0 \le B \le L$, ϵ, $C > 0$, $\sigma_0 \in \{0, 1\}$, and $f \le mB$. Suppose*

$$mB \le \tau + \sigma_0 - \epsilon N, \qquad m \ge C/\epsilon^2.$$

Let $h_{\mathrm{pk}} \coloneqq mB - f$. For uniform $v \in \mathbb{F}_p^\tau$,

$$M_{m,B,f} \le 2^L \binom{f + h_{\mathrm{pk}}}{h_{\mathrm{pk}}}, \tag{5.22}$$

$$\Pr[P_{m,B}(v) \text{ is solvable}] \le 2^{(g+\sigma_0)^2/(C\cdot N)}(N+1)^{g+\sigma_0}p^{-g}. \tag{5.23}$$

Proof. The hypotheses $f \le mB$ and $mB \le \tau + \sigma_0 - \epsilon N$ give $0 \le h_{\mathrm{pk}} \le g + \sigma_0 - \epsilon N$. Choose the touched blocks in at most 2^L ways. Their union has at most $mB = f + h_{\mathrm{pk}}$ coordinates, and hence

$$M_{m,B,f} \le 2^L \binom{mB}{f} = 2^L \binom{f + h_{\mathrm{pk}}}{h_{\mathrm{pk}}}.$$

Since $h_{\mathrm{pk}} \ge 0$, one has $\epsilon N \le g + \sigma_0$, so

$$L = N/m \le N\epsilon^2/C \le \frac{(g+\sigma_0)^2}{CN}.$$

Also $h_{\mathrm{pk}} \le g + \sigma_0$ and $f + h_{\mathrm{pk}} \le N$, giving $\binom{f+h_{\mathrm{pk}}}{h_{\mathrm{pk}}} \le (N+1)^{g+\sigma_0}$. Exact f enlargement reduces restricted solvability to a union of exact f support spans. Each contains a uniform target with probability $p^{f-\tau} = p^{-g}$. Therefore

$$\Pr[P_{m,B}(v) \text{ is solvable}] \le 2^{(g+\sigma_0)^2/(CN)}(N+1)^{g+\sigma_0}p^{-g}.$$

□

KRZSW has $(C, \sigma_0) = (16, 1)$, and the AHS capacity corollary has $(1, 0)$. For $p = 2^{\Omega(\lambda)}$, polynomial N, and fixed C, σ_0, the final bound is $2^{-\Omega(g\lambda)}$. These results keep setup, block support transfer, finite decoder admissibility, complete output, and attack distribution witness existence apart, and section 5.5 collects the consequences.

5.5 Comparison of decoder contracts

Tables 1 and 2 answer two questions. The first table records the certified metric, source output, and complete output and time contract. The second table states the resulting oracle relation and then projects the exact theorem conditions onto the tested attack weights.

In both tables "bit PPT" includes the conversion and verification overhead of proposition 4.4. The finite exact weight projections of the second table are statements about the original uniform attack experiment and the 24 tested rows described in section 8.

Table 1: Decoder theorem contracts.

Route	Metric or region	Source output	Output and time contract
BW	Coordinate weight $f \le \lfloor \tau/2 \rfloor$	One codeword	Deterministic $\mathcal{O}(N^3)$ field operations. Family level bit PPT.
IRS	Coordinate weight after the heterogeneous lift and inside the source regimes	One tuple or failure	Distributional guarantee under the source distribution, with $\mathcal{O}(N^4 r^3) \le \mathcal{O}(N^7)$ field operations.
GS	Coordinate weight $N - f > \sqrt{N(k-1)}$	Complete list of size $\mathcal{O}(N^2)$	Polynomial interpolation and root finding. The source gives an $\mathcal{O}(N^{15})$ bound with its field size qualification.
Original GR	Block budget $B_{\mathrm{GR}} = L - t$ from equation (5.9)	Complete list of size at most $p^{s_{\mathrm{int}}}$	Interpolation is polynomial in N and $r_{\mathrm{GR}}^{s_{\mathrm{int}}}$. The root stage is $p^{\mathcal{O}(s_{\mathrm{int}})}$, so the prefix transfer gives no short prefix bit PPT certificate.
GW	Folded radius ρ_{GW}	Affine message space of dimension $h \le s_{\mathrm{GW}} - 1$	Basis in $\mathcal{O}(N^2 + s_{\mathrm{GW}} k^2)$ field operations. The coordinate subset pruning of proposition 5.8 costs $N^{\mathcal{O}(s_{\mathrm{GW}})} \operatorname{poly}(\log p)$ and outputs at most $\binom{N}{h}$ words. AHS gives its own fixed gap bound in equation (5.13).
KRZSW	Route specific exact block budget from equation (5.15)	Complete list bounded by $\epsilon_{\mathrm{K}}^{-\mathcal{O}(\epsilon_{\mathrm{K}}^{-1} \cdot \log(1/\epsilon_{\mathrm{K}}))}$	Bit polynomial for a fixed positive gap ϵ_{K}. The cited theorem is not fully polynomial in $1/\epsilon_{\mathrm{K}}$.
AHS	Exact folded radius $\rho_{s,m}(\beta)$	Complete list of size at most s	$T_{\mathrm{basis}} + \mathcal{O}(ms^5 \log(s/\delta_{\mathrm{dec}})L)$ field operations.

Table 2: Oracle relations and finite exact weight projections.

Route	Theorem supported oracle relation	Finite exact weight projection
BW	Pointwise deterministic bit PPT realization of $P_{H,f}$ inside the unique radius.	No tested exact weight lies inside the certified radius.
IRS	Distributional guarantee of theorem 4.8, without a pointwise oracle theorem.	No tested exact weight lies inside the certified radius, so the additive bound is not reached.
GS	Pointwise complete bit PPT realization of $P_{H,f}$ inside the finite radius.	No tested exact weight lies inside the radius. List size is not the finite obstruction.
Original GR	Complete $P_{m,B}$ correctness under the internal prefix transfer, without a short prefix bit PPT certificate.	No selected block budget covers a tested exact weight.
GW	Conditional pointwise bit PPT realization of $P_{m,B}$ for bounded s_{GW} through the coordinate subset pruning of proposition 5.8.	No selected budget covers a tested exact weight. Witness frequency is bounded in section 7.
KRZSW	Conditional randomized bit PPT realization of $P_{m,B}$ for a fixed positive gap.	No selected budget covers a tested exact weight.
AHS	Conditional randomized pointwise bit PPT realization of $P_{m,B}$.	No selected budget covers a tested exact weight. Witnesses of lower weight are bounded in section 7.5.

The Guruswami–Sudan runtime entry needs one qualification, because its $\mathcal{O}(N^{15})$ bound is stated for fields of cardinality at most 2^N [GS99, Theorem 12], a hypothesis that holds for the three longest tested pairs and fails for the three shortest at the field size proxy 2^{252}. The bit polynomial classification of that row therefore rests on the source's randomized root finding option, which is polynomial in $\log p$ [GS99, Section 1], rather than on the $\mathcal{O}(N^{15})$ figure.

Every oracle status is conditional on the cited setup, metric, and output hypotheses, and every attack status also includes the attack induced distribution and probability composition requirements of section 4.3. Both tables list algorithmic decoder routes only. The combinatorial list size bounds of Srivastava and Chen–Zhang in section 5.4 are therefore absent from them, whereas AHS is present, because its theorem includes a complete list algorithm under its stated hypotheses.

Each row uses theorem components compatible with its code family. In the GW row, AHS Theorem 3.5 supplies source supported completion at fixed gap, and the internal result in proposition 5.8 supplies the uniform bounded dimension family level certificate. The setup and support transfer requirements of theorem 5.2 apply to every folded row, also when a polynomial list bound is available.

6 Canonical menu and finite application

The line bound of theorem 3.5 holds for every finite menu, and its value depends on the regions that the menu declares. This section fixes the menu of the finite application and evaluates it. We first fix one aligned partition for each divisor of N and record the nested Guruswami–Wang and AHS block regions. We then enclose the block regions certified

by the Original GR agreement condition in an optimistic envelope. Next we define the menu and state its count as an equality, and section 6.4 evaluates that count on the six parameter pairs of the original attack.

Throughout this section we keep the setting of section 2 and assume $k \geq 2$, so that $d = k - 1 \geq 1$. The counts below use only the Shamir data, and the ordered power domain of the standing conditional setup assumption of section 5 enters only when a block region is read as the certified region of a folded decoder theorem, where it is a hypothesis of that reading.

6.1 Fixed partitions and the nested Guruswami–Wang and AHS regions

For each positive divisor m of N, including $m = 1$, let $\mathcal{P}_m$ denote the aligned partition $\{\mathcal{B}_1, \mathcal{B}_2, \ldots, \mathcal{B}_L\}$ of section 4 with block size m and $L = N/m$ blocks, so that $\mathrm{wt}_{\mathcal{P}_m} = \mathrm{wt}_m$. Its blocks follow the order $z_1, z_2, \ldots, z_N$ of the points, which is fixed once. There is exactly one such partition for each divisor, and it is fixed independently of every target. We call these fixed aligned partitions canonical. A solver may select a region of the menu after it has seen the target, as the wrapper of definition 3.1 allows, but the menu itself is fixed beforehand.

For the coordinate ball, let $F \coloneqq F_{\mathrm{GS}}^{\mathrm{fin}}$, the finite Guruswami–Sudan ceiling of equation (4.16). For the block balls, let $\rho_{s,m}(\beta)$ be the radius of equation (5.17). For a divisor m and an integer $1 \leq s \leq m$ we consider the budget $\lfloor L \cdot \rho_{s,m}(\beta) \rfloor$ and retain it only when it is nonnegative. Let $\mathcal{D}_{\mathrm{GW}}$ be the set of divisors m of N that admit at least one retained budget, and keep the largest retained budget at each such divisor, that is,

$$\begin{aligned} B_{\mathrm{GW}}(m,\, s) &\coloneqq \lfloor L \cdot \rho_{s,m}(\beta) \rfloor, && 1 \leq s \leq m, \\ B_{\mathrm{GW},m}^{\max} &\coloneqq \max\{\, B_{\mathrm{GW}}(m,\, s) :\ 1 \leq s \leq m,\ B_{\mathrm{GW}}(m,\, s) \geq 0 \,\}, && m \in \mathcal{D}_{\mathrm{GW}}. \end{aligned} \tag{6.1}$$

Since $m\beta = k/L$, the radius $\rho_{s,m}(\beta)$ is the Guruswami–Wang radius ρ_{GW} of equation (5.12) at $s_{\mathrm{GW}} = s$, and $B_{\mathrm{GW}}(m,\, s)$ is the budget B_{GW} defined there. A divisor without a retained budget is absent from $\mathcal{D}_{\mathrm{GW}}$ and contributes no region, whereas a budget that is equal to zero is retained.

At the same pair $(m,\, s)$, the AHS budget of equation (5.17) satisfies

$$B_{\mathrm{AHS}} = \lceil L\rho_{s,m}(\beta) \rceil - 1 \leq \lfloor L\rho_{s,m}(\beta) \rfloor = B_{\mathrm{GW}}(m,\, s).$$

The two sides agree unless $L\rho_{s,m}(\beta)$ is an integer. In that case the AHS budget is one block smaller, which is the strict AHS boundary. Within one fixed partition the block balls are nested, because $\mathrm{wt}_{\mathcal{P}_m}(a) \leq B$ implies $\mathrm{wt}_{\mathcal{P}_m}(a) \leq B'$ whenever $B \leq B'$. The ball of budget $B_{\mathrm{GW},m}^{\max}$ therefore contains every retained Guruswami–Wang ball and every nonnegative AHS ball of $\mathcal{P}_m$. These are containments of regions, and the two algorithms keep their own outputs, running times and success probabilities.

6.2 The Original Guruswami–Rudra optimistic envelope

The Original GR agreement condition of equation (5.9) depends on the multiplicity r_{GR}. We bound every certified budget by a value independent of the multiplicity, which gives the optimistic envelope of this subsection. For a divisor m and an integer $1 \leq s_{\mathrm{int}} \leq m$, let $u_{m,s_{\mathrm{int}}}$ be the smallest positive integer u with

$$u^{s_{\mathrm{int}}+1} \cdot (m - s_{\mathrm{int}} + 1)^{s_{\mathrm{int}}} > L \cdot (k-1)^{s_{\mathrm{int}}},$$

which exists because the left side is unbounded in u. We omit the negative values of $L-2-u_{m,s_{\mathrm{int}}}$. Let $\mathcal{D}_{\mathrm{GR}}$ be the set of divisors m of N that admit at least one nonnegative value, and keep the largest nonnegative value at each such divisor, that is,

$$\begin{aligned} \overline{B}_{m,s_{\mathrm{int}}} &:= L-2-u_{m,s_{\mathrm{int}}}, && 1 \le s_{\mathrm{int}} \le m, \\ \overline{B}_{m} &:= \max\{\, \overline{B}_{m,s_{\mathrm{int}}} :\ 1 \le s_{\mathrm{int}} \le m,\ \overline{B}_{m,s_{\mathrm{int}}} \ge 0 \,\}, && m \in \mathcal{D}_{\mathrm{GR}}. \end{aligned} \tag{6.2}$$

As before, a divisor outside $\mathcal{D}_{\mathrm{GR}}$ is absent.

Lemma 6.1 (Original GR optimistic envelope)**.** *Let $m \mid N$ with $L = N/m$, let $1 \le s_{\mathrm{int}} \le m$, and let $r_{\mathrm{GR}} \ge 1$ be a finite integer multiplicity. If an integer agreement t satisfies equation (5.9), then $u_{m,s_{\mathrm{int}}} \le t-2$, and therefore*

$$B_{\mathrm{GR}} = L - t \le \overline{B}_{m,s_{\mathrm{int}}}.$$

Proof. By equation (5.9), the number $t-2$ is at least a nonnegative real number. Raising both sides to the power $s_{\mathrm{int}}+1$ and multiplying by $(m-s_{\mathrm{int}}+1)^{s_{\mathrm{int}}} > 0$ gives

$$(t-2)^{s_{\mathrm{int}}+1} \cdot (m-s_{\mathrm{int}}+1)^{s_{\mathrm{int}}} \ge L \cdot d^{s_{\mathrm{int}}} \prod_{j=1}^{s_{\mathrm{int}}} \left(1 + \frac{j}{r_{\mathrm{GR}}}\right) > L \cdot d^{s_{\mathrm{int}}}.$$

The second inequality is strict because $L \cdot d^{s_{\mathrm{int}}} > 0$ and every factor $1 + j/r_{\mathrm{GR}}$ exceeds one for a finite positive multiplicity. The right side is positive, so $t-2$ is a positive integer, and it satisfies the inequality that defines $u_{m,s_{\mathrm{int}}}$. Minimality gives $u_{m,s_{\mathrm{int}}} \le t-2$, hence $L - t \le L - 2 - u_{m,s_{\mathrm{int}}} = \overline{B}_{m,s_{\mathrm{int}}}$. □

Consequently, every block budget $B_{\mathrm{GR}} = L - t \ge 0$ that equation (5.9) certifies at a divisor m, at any arity $1 \le s_{\mathrm{int}} \le m$ and with any finite multiplicity, is at most $\overline{B}_m$. In particular such a divisor lies in $\mathcal{D}_{\mathrm{GR}}$, and the block ball of budget $\overline{B}_m$ of $\mathcal{P}_m$ contains every such certified Original GR ball. The envelope ignores the root safety predicate of equation (5.10). Ignoring it can only enlarge the region, which is legitimate for an upper bound. For the agreement condition alone the envelope is exact. The inequality that defines $u_{m,s_{\mathrm{int}}}$ is strict, and $\prod_{j=1}^{s_{\mathrm{int}}}(1 + j/r_{\mathrm{GR}})$ decreases to one as r_{GR} grows, so whenever $\overline{B}_{m,s_{\mathrm{int}}} \ge 0$ the agreement $t = u_{m,s_{\mathrm{int}}} + 2 \le L$ satisfies equation (5.9) for every sufficiently large finite multiplicity. Because of the ignored root safety predicate, whose left side does not decrease as r_{GR} grows, an envelope budget need not be certified by proposition 5.7 with any multiplicity at a given prime p. The output and the running time of Original GR keep the complete list and root costs of proposition 5.7.

6.3 The canonical menu and its count

Definition 6.2 (Canonical envelope menu)**.** Assume that $0 \le F < \tau$, that $m \cdot B^{\max}_{\mathrm{GW},m} < \tau$ for every $m \in \mathcal{D}_{\mathrm{GW}}$, and that $m \cdot \overline{B}_m < \tau$ for every $m \in \mathcal{D}_{\mathrm{GR}}$. The canonical envelope menu $\mathcal{M}_{\mathrm{can}}$ is the finite menu in the sense of definition 2.2 with the coordinate budget $F = F^{\mathrm{fin}}_{\mathrm{GS}}$ and with the index set

$$I_{\mathrm{can}} := \{\, (\mathrm{GW},\, m) :\ m \in \mathcal{D}_{\mathrm{GW}} \,\} \ \cup\ \{\, (\mathrm{GR},\, m) :\ m \in \mathcal{D}_{\mathrm{GR}} \,\},$$

whose members are

$$\big(\mathcal{P}_{(\mathrm{GW},m)},\, B_{(\mathrm{GW},m)}\big) := \big(\mathcal{P}_m,\, B^{\max}_{\mathrm{GW},m}\big), \qquad \big(\mathcal{P}_{(\mathrm{GR},m)},\, B_{(\mathrm{GR},m)}\big) := \big(\mathcal{P}_m,\, \overline{B}_m\big).$$

The three conditions are the hypotheses $0 \le F < \tau$ and $m_\iota B_\iota < \tau$ of definition 2.2, and the block budgets are nonnegative by construction. The three conditions are assumptions on the parameters, and when one of them fails the menu is undefined as it stands, since its members are fixed by the budgets above. The route tags keep the two families apart, so a divisor that lies in both $\mathcal{D}_{\mathrm{GW}}$ and $\mathcal{D}_{\mathrm{GR}}$ contributes two members with the same partition $\mathcal{P}_m$.

For this menu the count of equation (3.3) reads

$$M_{\mathrm{line}}(\mathcal{M}_{\mathrm{can}}) = M_{\mathrm{coord}}(N, \tau, F) + \sum_{m \in \mathcal{D}_{\mathrm{GW}}} M_{\mathrm{block}}\big(N, \tau, m, B^{\max}_{\mathrm{GW},m}\big) + \sum_{m \in \mathcal{D}_{\mathrm{GR}}} M_{\mathrm{block}}\big(N, \tau, m, \overline{B}_m\big). \tag{6.3}$$

This is the specialization of equation (3.3) to $\mathcal{M}_{\mathrm{can}}$, an equality whose every term is one of the quantities of equation (3.2), with its three way minimum unchanged.

The sums over divisors keep all of their terms, since their largest terms alone need not bound the union. Block balls of one partition are nested, which is why one budget for each route and divisor suffices. Balls of different partitions need not be nested, because a support that meets few blocks of $\mathcal{P}_m$ can meet many blocks of $\mathcal{P}_{m'}$. A solver may choose any member after seeing the target, so the targets covered on a line form the union of the sets covered by the members. Theorem 3.5 bounds that union by the sum of the member counts, and overlaps only cause overcounting.

The count keeps every tagged member as a term of its own. At a divisor that lies in both sets, the Guruswami–Wang ball and the Original GR ball belong to the same partition and are nested, yet both members are kept and both terms are added, and the divisor $m = 1$ contributes its terms like every other divisor. Merging nested or coinciding balls, or dropping $m = 1$, could only lower the count or leave it unchanged, so keeping every term is the conservative choice.

Under the conditions of definition 6.2, $\mathcal{M}_{\mathrm{can}}$ is a finite menu fixed independently of the targets. Theorem 3.5 therefore applies to it, so every line in direction K_0 contains at most $\min\{\, p,\ M_{\mathrm{line}}(\mathcal{M}_{\mathrm{can}}) \,\}$ targets of $G_{\mathcal{M}_{\mathrm{can}}}$. Corollary 3.7 then bounds the success probability of the wrapper of definition 3.1 with this menu, a cap $0 \le f \le \tau - 1$ and at most Q trials by $\min\{\, 1,\ Q \cdot M_{\mathrm{line}}(\mathcal{M}_{\mathrm{can}})/p \,\}$, with the clipping at one included. The cap is the acceptance condition of the wrapper rather than a restriction of the region of definition 2.2, which imposes no cap. On the six pairs of the original attack, the coordinate ball of radius F also covers the Berlekamp–Welch, heterogeneous IRS and KRZSW regions, as section 6.4 verifies.

6.4 Finite application to the six parameter pairs

We now evaluate the canonical envelope menu of this section on the six parameter pairs of the original attack, $(N, \tau) = (64, 43)$, $(128, 86)$, $(196, 131)$, $(512, 342)$, $(768, 513)$ and $(1024, 683)$, with dimensions $k = 21$, 42, 65, 170, 255 and 341, which turns corollary 3.7 into explicit numbers. The accompanying artifact [DY26] records the exact budgets and the integer certificates for this evaluation. The probability statement below assumes a prime $p > N$ and pairwise distinct nonzero Shamir points. Reading a block region as the region of a folded decoder also assumes that the points are ordered as the compatible power domain of the standing conditional folded setup of section 5.

The menu is defined on every pair. Each pair has $k \ge 2$, and equation (4.16) gives $F = F^{\mathrm{fin}}_{\mathrm{GS}} = 28$, 55, 83, 217, 326 and 433, below τ. The recorded certificates list every divisor $m \in \mathcal{D}_{\mathrm{GW}}$ with its budget $B^{\max}_{\mathrm{GW},m}$ of equation (6.1) and every divisor $m \in \mathcal{D}_{\mathrm{GR}}$ with its value $\overline{B}_m$ of equation (6.2). The largest products $m \cdot B^{\max}_{\mathrm{GW},m}$ are 32, 64, 98, 256, 384 and 512, and the largest products $m \cdot \overline{B}_m$ are 26, 53, 81, 224, 360 and 496, all below $\tau = 43$, 86, 131, 342, 513 and 683, respectively. The three conditions of definition 6.2 therefore hold, and $\mathcal{M}_{\mathrm{can}}$ is a finite menu fixed independently of the targets.

Its region $\mathcal{W}_{\mathcal{M}_{\mathrm{can}}}$ contains the certified region of each of the seven studied routes at every divisor of N and every admissible decoder parameter. Coordinate balls are nested, so a coordinate ball lies inside the ball of radius F whenever its radius is at most F. The Berlekamp–Welch radius $F_{\mathrm{BW}} = \lfloor \tau/2 \rfloor$ takes the values 21, 43, 65, 171, 256 and 341. On all six pairs $\mathcal{R}_1 = \{1, 2\}$ and $\mathcal{R}_2 = \{3\}$, as recorded in section 4.6.2, and equation (4.13) gives $F_{\mathrm{IRS}} = 21$, 43, 65, 171, 257 and 341, and both radii are at most F on every pair. A vector whose block weight in $\mathcal{P}_m$ is at most B has at most mB nonzero coordinates, so the KRZSW ball of a budget B_{K} lies inside the ball of radius F whenever $mB_{\mathrm{K}} \le F$. Over the divisors at which the set of equation (5.15) is nonempty, the largest product mB_{K} is 0 on the three smallest pairs and 128, 256 and 256 on the three largest, again at most F, and within one partition the KRZSW balls are nested in the largest budget. Every KRZSW ball therefore lies in the coordinate ball, which is the Guruswami–Sudan ball itself. At a common pair (m, s) the AHS budget is at most the Guruswami–Wang budget, one block smaller exactly when $L\rho_{s,m}(\beta)$ is an integer, so a divisor with a nonnegative AHS budget lies in $\mathcal{D}_{\mathrm{GW}}$. Section 6.1 shows that the ball of budget $B^{\max}_{\mathrm{GW},m}$ contains every retained Guruswami–Wang ball and every nonnegative AHS ball of $\mathcal{P}_m$. By lemma 6.1 and the paragraph after it, the Original GR member of a divisor contains every ball that equation (5.9) certifies there with a finite multiplicity. The union of the seven regions therefore lies in $\mathcal{W}_{\mathcal{M}_{\mathrm{can}}}$, so the wrapper may choose among all seven routes after it sees the challenge. Berlekamp–Welch, the IRS regions, KRZSW and AHS lie inside existing members, so the menu needs no member of its own for them.

Table 3 lists, for each pair, the coordinate ceiling, the three terms of equation (6.3) and their total $M_{\mathrm{line}}(\mathcal{M}_{\mathrm{can}})$, together with two strict binary exponents. The totals are exact evaluations of equation (6.3), which adds the coordinate term and the two sums over divisors, whose terms are the three way minima of equation (3.2) with their single final floor. Every tagged member and every retained divisor is kept, including the divisor $m = 1$, for the reason given in section 6. A divisor outside $\mathcal{D}_{\mathrm{GW}}$ or $\mathcal{D}_{\mathrm{GR}}$ contributes no term, whereas a retained budget zero contributes the count one. The recorded certificates give each block summand with the three bounds of equation (3.2) and their minimum, and the coordinate term with the index j that attains its minimum together with D_j and J_j of equation (3.1). At (64, 43), for instance, the coordinate term is 1051, the seven Guruswami–Wang summands add up to 14 and the five Original GR summands to 18, so $M_{\mathrm{line}}(\mathcal{M}_{\mathrm{can}}) = 1051 + 14 + 18 = 1083$. Each certified term is one of the quantities whose minimum defines its summand, so the certificates alone confirm that the totals of the table are valid bounds. The sum uses the defined terms, and sharper counts for individual terms, such as those of remark 3.6, stay outside it. The largest total 10287538 is below $2^{24} = 16777216$, so all six totals are below 2^{24}.

Since $\mathcal{M}_{\mathrm{can}}$ is a finite menu fixed independently of the targets, corollary 3.7 applies to the wrapper of definition 3.1 with this menu, a cap $0 \le f \le \tau - 1$ and at most Q trials. Conditional on GoodKG, it succeeds within Q trials with probability at most $\min\{1, Q \cdot M_{\mathrm{line}}(\mathcal{M}_{\mathrm{can}})/p\}$, in which only the factor $1/p$ depends on the field, and every accepted witness yields a valid forgery on an unsigned message within the corruption cap.

Table 3: Line point bounds and strict success exponents for the chosen alpha wrapper at the 2^{252} field size proxy.

(N, τ)	$F_{\mathrm{GS}}^{\mathrm{fin}}$	M_{coord}	GW sum	GR sum	M_{line}	ℓ_{bd}	
						$Q = 1$	$Q = 2^{128}$
$(64, 43)$	28	1,051	14	18	1,083	241	113
$(128, 86)$	55	764	23	28	815	242	114
$(196, 131)$	83	253	37	41	331	243	115
$(512, 342)$	217	1,137	958	1,186	3,281	240	112
$(768, 513)$	326	10,410	10,271	65,274	85,955	235	107
$(1024, 683)$	433	1,581	259,183	10,026,774	10,287,538	228	100

The last two columns evaluate this bound for one challenge and for $Q = 2^{128}$ challenges at the field size proxy $p = 2^{252}$, which stands for a prime $p \approx 2^{252}$. If M_{line} has bit length b, the two entries are $\ell_{\mathrm{bd}} = 252 - b$ and $\ell_{\mathrm{bd}} = 124 - b$, and since $M_{\mathrm{line}} < 2^b$ the bound is strictly below $2^{-\ell_{\mathrm{bd}}}$. All six totals are below 2^{24}, so at the proxy the bound is below 2^{-228} per challenge and below 2^{-100} after 2^{128} challenges on every pair, and the pair (1024, 683), whose total has 24 bits, has exactly these exponents.

On the three smallest pairs the coordinate term supplies most of the count, whereas on the three largest pairs the Original GR sum supplies the largest term. At (1024, 683) its summands at $m = 8$ and $m = 16$ alone are 7103626 and 2922274. These summands come from the deliberately optimistic envelope of section 6. The remarks after corollary 3.7 on perfect and free recovery and on the trial budget Q apply to these entries as well, and an estimate of the attack cost needs in addition an achieved success probability and the work per trial.

The line point counts of the Berlekamp–Welch and IRS regions illustrate remark 3.6. For odd τ, that is, on (64, 43), (196, 131), (768, 513) and (1024, 683), one has $2F_{\mathrm{BW}} = \tau - 1 < \tau$, so at most one point of a line is covered. For even $\tau = 86$ and 342, one has $2F_{\mathrm{BW}} = \tau$, so at most $\lfloor N/F_{\mathrm{BW}} \rfloor = 2$ points are covered, and the construction in that remark shows that two covered points occur. The IRS ceiling equals F_{BW} on every pair except (768, 513), where $F_{\mathrm{IRS}} = 257$ exceeds $\tau/2$. There the term $j = 0$ of equation (3.1) with $(n, w, \nu) = (768, 511, 255)$ has $D_0 = 65281$ and gives the bound $\lfloor 196608/65281 \rfloor = 3$. Hence the Berlekamp–Welch regions cover at most 1, 2, 1, 2, 1 and 1 points of any line, and the IRS regions at most 1, 2, 1, 2, 3 and 1, well below the coordinate terms of the table.

Remark 6.3 (Shortened terms at the coordinate ceiling). The coordinate term applies equation (3.1) with $(n, w, \nu) = (N, N - F, k)$, so that $D_0 = (N - F)^2 - Nk$. At the ceiling $F = F_{\mathrm{GS}}^{\mathrm{fin}}$ this value is -48, -47, 29, -15, -476 and 97 on the six pairs. It is negative on (64, 43), (128, 86), (512, 342) and (768, 513), so on these four pairs the unshortened term $j = 0$ is unavailable, and terms with positive j attain the reported coordinate count, at $j = 5$, 4, 2 and 5. On (196, 131) and (1024, 683) the term $j = 0$ is defined and equals $\lfloor 196 \cdot 48/29 \rfloor = 324$ and $\lfloor 1024 \cdot 250/97 \rfloor = 2639$, and the term $j = 1$ attains the smaller counts 253 and 1581. The certificates (j, D_j, J_j) of the six coordinate terms are (5, 17, 52), (4, 49, 78), (1, 64, 146), (2, 169, 377), (5, 219, 651) and (1, 280, 913). At (64, 43), for instance, $J_5 = \lfloor 59 \cdot 15/17 \rfloor = 52$ and $T_5 = \lfloor \binom{64}{5} \cdot 52/\binom{36}{5} \rfloor = 1051$. The count is the minimum of $\binom{N}{N-F}$, T_{pack} and every available T_j, as defined in equation (3.1), and the remark only records which term attains it at these six ceilings.

7 Uniform targets and supporting bounds

The line bound of section 3 is about a coefficient vector chosen by the adversary, whereas this section treats the original sampling rule, in which the vector is uniform, so that the target is uniform in $\mathbb{F}_p^\tau$ and elementary first moment bounds apply. For this sampling rule and a single block restricted relation, the folded analysis has already proved the witness existence lower bound of theorem 5.4 and the support scarcity bound of theorem 5.10. We first define the uniform variant of the wrapper and bound the probability that a uniform target has a witness in a coordinate or block region. We then bound repeated trials and derive the condition that a required success probability imposes on the gap $g = \tau - f$. The next two subsections compare the viable weights with the certified coordinate radii and, through the gap that an optimally packed support leaves below a folded radius, with the deficit of the common Guruswami–Wang and AHS radius, which yields a condition on the length. Section 7.5 then proves a witness scarcity statement at all lengths for the block regions of that common radius, and the last subsection evaluates these bounds on the six parameter pairs of the original attack.

7.1 The uniform sampling rule and witness region bounds

Definition 7.1 (Uniform alpha wrapper)**.** The uniform alpha wrapper is the wrapper of definition 3.1 with one change in the first step of every reached trial i. The wrapper first selects a message M_i that has not appeared in any previous signature hash query. It then samples α_i uniformly from $\mathbb{F}_p^\tau$, with fresh randomness that is independent of the preceding history, and forms $R_i = \prod_{j=0}^{\tau-1} X_j^{\alpha_{i,j}}$ from the public coefficient commitments X_j of lemma 2.1. The single fresh query $c_i = \mathcal{H}_{\mathsf{sig}}(R_i, PK, M_i)$, the target $v_i = \alpha_i + c_i K_0$, the acceptance rule and every other requirement are those of definition 3.1.

This is the original uniform experiment with repeated fresh queries. A trial is one fresh target, and a decoder run again on an unchanged target with new coins belongs to the same trial. The commitments R_i of different trials need not be distinct, but the hash inputs (R_i, PK, M_i) must be fresh, which the new message guarantees.

We distinguish two histories. The history $\mathcal{F}_i$ of section 3.1 is taken just before the answer c_i and already determines α_i. For a reached trial i of the uniform alpha wrapper, let $\mathcal{G}_i$ denote the history just before α_i is drawn and after M_i is selected, the finite string that records the public input, every earlier query and answer and the private randomness used so far. Like $\mathcal{F}_i$, it is a random variable with countably many values, and α_i is drawn after it. Uniformity of α_i is therefore a statement conditional on each value of $\mathcal{G}_i$ of positive probability, whereas each value of $\mathcal{F}_i$ of positive probability fixes α_i.

Lemma 7.2 (Uniform attack target)**.** *In a reached trial i of the uniform alpha wrapper of definition 7.1, let h be a value of $\mathcal{G}_i$ with $\Pr[\mathcal{G}_i = h] > 0$. Conditional on $\mathcal{G}_i = h$, the target v_i is uniform in $\mathbb{F}_p^\tau$. That is, for every $\theta \in \mathbb{F}_p^\tau$,*

$$\Pr[\, v_i = \theta \mid \mathcal{G}_i = h \,] = p^{-\tau}.$$

Proof. Fix a value $\xi \in \mathbb{F}_p^\tau$. The vector α_i is drawn with fresh randomness after $\mathcal{G}_i$, so the event $\{\mathcal{G}_i = h,\ \alpha_i = \xi\}$ has probability $p^{-\tau} \cdot \Pr[\mathcal{G}_i = h] > 0$. On this event the hash input (R_i, PK, M_i) takes one fixed value, which has not been queried before, because M_i is new. The event $\{\mathcal{G}_i = h,\ \alpha_i = \xi\}$ depends only on the key material, the private randomness and the oracle answers recorded in h, together with the draw of α_i, and none of these is the answer of $\mathcal{H}_{\mathsf{sig}}$ at that input. In the random oracle model the answer c_i is therefore uniform in $\mathbb{F}_p$ conditional on this event. Index the coordinates of a vector of $\mathbb{F}_p^\tau$ by $0, 1, \ldots, \tau - 1$. Since $K_0 = (1, 0, \ldots, 0)^\top$, the target $v_i = \alpha_i + c_i K_0$ agrees with α_i

outside the coordinate 0. The event $v_i = \theta$ therefore requires $\alpha_{i,j} = \theta_j$ for $1 \le j \le \tau - 1$, and then the single answer $c_i = \theta_0 - \alpha_{i,0}$. Summing over the vectors $\xi \in \mathbb{F}_p^\tau$ with $\xi_j = \theta_j$ for $1 \le j \le \tau - 1$ gives

$$\begin{aligned}\Pr[\, v_i = \theta \mid \mathcal{G}_i = h \,] &= \sum_{\xi} \Pr[\, \alpha_i = \xi \mid \mathcal{G}_i = h \,] \cdot \Pr[\, c_i = \theta_0 - \xi_0 \mid \mathcal{G}_i = h,\ \alpha_i = \xi \,] \\ &= p \cdot p^{-\tau} \cdot p^{-1} = p^{-\tau},\end{aligned}$$

because exactly p vectors ξ meet the condition, each has probability $p^{-\tau}$, and each requires one answer of conditional probability p^{-1}. Crites and Stewart state this uniformity for a single target [CS25, proofs of Theorems 1 and 2], and the argument above derives it for repeated trials conditional on the history. □

The region bounds below use only the fact that H is full spark, that is, that any τ of its columns are linearly independent, as recorded in section 2.1. The coordinate bound is the first moment argument used by Crites and Stewart [CS25, proofs of Theorems 1 and 2], and the block bound is its elementary specialization to a fixed partition, proved here.

Proposition 7.3 (Uniform target bounds for witness regions)**.** *Let H be full spark and fixed before the target, and let v be uniform in $\mathbb{F}_p^\tau$. For an integer $0 \le F < \tau$, the event that some $a \in \mathbb{F}_p^N$ satisfies $Ha = v$ and $\mathrm{wt}(a) \le F$ has probability at most*

$$U_{\mathrm{coord}}(N, \tau, F) \coloneqq \min\left\{1,\ \binom{N}{F} \cdot p^{F-\tau}\right\}. \tag{7.1}$$

For a fixed partition $\mathcal{P}$ of $[N]$ into $L = N/m$ blocks of equal size m and an integer $B \ge 0$ with $mB < \tau$, the event that some $a \in \mathbb{F}_p^N$ satisfies $Ha = v$ and $\mathrm{wt}_{\mathcal{P}}(a) \le B$ has probability at most

$$U_{\mathrm{block}}(N, \tau, m, B) \coloneqq \min\left\{1,\ \binom{L}{B} \cdot p^{mB-\tau}\right\}. \tag{7.2}$$

Proof. For the coordinate bound, let a satisfy $Ha = v$ and $\mathrm{wt}(a) \le F$. Its support lies in a set $S \subseteq [N]$ of exactly F elements, which exists because $F < \tau \le N$, and then $v \in \mathrm{span}(H_S)$. Full spark and $F < \tau$ give $\dim \mathrm{span}(H_S) = F$, so a uniform v lies in this span with probability $p^{F-\tau}$. There are $\binom{N}{F}$ sets S, and the union bound gives $\binom{N}{F} \cdot p^{F-\tau}$. A probability is at most one, which gives the minimum. For $F = 0$ the only set is empty, its span is $\{0\}$, and the bound is the probability $p^{-\tau}$ of the zero target.

For the block bound, let a satisfy $Ha = v$ and $\mathrm{wt}_{\mathcal{P}}(a) \le B$. Its support lies in the union S of some B blocks of $\mathcal{P}$, and $|S| = mB < \tau$. Full spark gives $\dim \mathrm{span}(H_S) = mB$, so v lies in this span with probability $p^{mB-\tau}$. There are $\binom{L}{B}$ such unions, and the union bound and the clipping at one give the claim. For $B = 0$ the union is empty, and the argument is the same. □

These bounds are statements about the existence of a witness before any decoding, so imposing the cap $\mathrm{wt}(a) \le f$, or requiring that a decoder recovers the witness, can only shrink the event. If the event of a coordinate route with ceiling $F_{\mathfrak{D}}$ retains the cap, its effective radius is $\min\{f, F_{\mathfrak{D}}\}$.

For a finite menu $\mathcal{M}$ as in definition 2.2, the event $v \in G_{\mathcal{M}}$ is the union of the events of its members. Within one fixed partition the block balls are nested, so the largest retained budget covers the smaller ones, and across the declared members the union bound gives the sum of the member bounds. The uniform target bound U and the line count $M_{\mathrm{line}}(\mathcal{M})$ of equation (3.3) belong to different sampling rules, and in this section U is obtained from the member bounds of equations (7.1) and (7.2), so that it bounds the whole union $G_{\mathcal{M}}$ as corollary 7.4 requires.

7.2 Repetition and the first moment necessary condition

Corollary 7.4 (Repetition with uniform alpha)**.** *Consider the uniform alpha wrapper of definition 7.1 with a finite menu $\mathcal{M}$, a cap $0 \le f \le \tau - 1$ and at most $Q \ge 1$ trials, on a fixed valid Shamir setup with pairwise distinct nonzero points. Let $U \in [0, 1]$ be a deterministic number such that a target v uniform in $\mathbb{F}_p^\tau$ satisfies $\Pr[\, v \in G_{\mathcal{M}} \,] \le U$. Then*

$$\Pr[\, \textit{the wrapper succeeds within } Q \textit{ trials} \mid \mathsf{GoodKG} \,] \le \min\{\, 1,\ Q \cdot U \,\}, \tag{7.3}$$

and every accepted witness yields a valid forgery on an unsigned message within the corruption cap.

Proof. All probabilities in this proof are conditional on GoodKG. This event depends only on the key material, which is independent of the random oracle, so lemma 7.2 holds conditional on GoodKG and on each value h as well. Let E_i be the event that trial i is reached and accepts a witness. On an execution that stops before trial i, the indicator of E_i is zero. On an execution that reaches trial i, acceptance requires some $a \in \mathcal{W}_{\mathcal{M}}$ with $Ha = v_i$, so E_i is contained in the event $v_i \in G_{\mathcal{M}}$, and the cap and the representation condition only shrink it. Fix a value h of $\mathcal{G}_i$ with $\Pr[\mathcal{G}_i = h] > 0$, taken in a reached trial i before α_i is drawn. The menu is fixed independently of the targets, so $G_{\mathcal{M}}$ is a fixed set, and by lemma 7.2 the target v_i is uniform in $\mathbb{F}_p^\tau$ conditional on $\mathcal{G}_i = h$. Therefore

$$\Pr[\, E_i \mid \mathcal{G}_i = h \,] \le \Pr[\, v_i \in G_{\mathcal{M}} \mid \mathcal{G}_i = h \,] \le U.$$

Up to an event of probability zero, E_i is contained in the union of the disjoint events $\mathcal{G}_i = h$ over these countably many values h, so $\Pr[E_i] = \sum_h \Pr[\mathcal{G}_i = h] \cdot \Pr[\, E_i \mid \mathcal{G}_i = h \,] \le U$ for every $i \le Q$. The success event is contained in $\bigcup_{i \le Q} E_i$, and the union bound gives the claimed inequality after clipping at one. The bound holds whatever the dependence between the events E_i, between supports, between the regions of the menu and in the randomness of the solvers. The forgery and the corruption accounting are those of the proof of corollary 3.7, which hold for every rule that chooses α_i. □

The number U is an upper bound on the success probability of one trial, and the remarks after corollary 3.7 on perfect and free recovery and on the trial budget Q apply here as well. An estimate of the attack cost requires in addition an achieved success probability and the work per trial, two quantities beyond U and its reciprocal.

Applying the same repetition argument to the capped solvability event gives a necessary condition on the gap $g = \tau - f$, for which the following lemma supplies the single target first moment bound. The first moment is the expectation $\mathbb{E}[X_f] = \mu_f$ of the support span count X_f of section 4.5. The expected number of coefficient vector witnesses of exact weight f is the slightly smaller $\mathbb{E}[Z_f] = (1 - 1/p)^f \mu_f$ of equation (4.7).

Lemma 7.5 (First moment viability)**.** *Let H be full spark, let $1 \le f = \tau - g \le \tau$, and let v be uniform in $\mathbb{F}_p^\tau$. If $\Pr[P_{H,f}(v) \ne \varnothing] \ge 2^{-\kappa}$, then*

$$g \log_2 p \le \log_2 \binom{N}{\tau - g} + \kappa, \qquad \textit{and hence} \qquad g \le \frac{N + \kappa}{\log_2 p}. \tag{7.4}$$

Proof. A witness of weight at most f has a support that extends to a set S of size exactly f, and then $v \in \operatorname{span}(H_S)$. Conversely such a membership supplies a witness supported on S. Solvability is therefore the event $X_f(v) > 0$ for the support span count of section 4.5. Full spark and $f \le \tau$ give $\dim \operatorname{span}(H_S) = f$, so $\Pr[v \in \operatorname{span}(H_S)] = p^{f-\tau} = p^{-g}$ and $\mathbb{E}[X_f] = \binom{N}{f} p^{-g}$. Markov's inequality in the direction required for a necessary condition gives

$$\Pr[P_{H,f}(v) \ne \varnothing] = \Pr[X_f > 0] \le \mathbb{E}[X_f] = \binom{N}{f} 2^{-g \log_2 p}.$$

Taking base two logarithms yields the first inequality, and $\binom{N}{f} \le 2^N$ yields the second. The argument uses the first moment alone, since a second moment lower bound points in the wrong direction for a necessary condition. □

The lemma is a statement about the relation $P_{H,f}$ and admits $f = \tau$, that is, $g = 0$. The wrappers of this paper keep the cap $f \le \tau - 1$, so this endpoint of the lemma lies outside them.

Corollary 7.6 (Necessary condition for a success target)**.** *Let $f = \tau - g$ with $1 \le g \le \tau - 1$, let $Q \ge 1$ be an integer and let $\eta_\star \in (0, 1]$. If the uniform alpha wrapper of definition 7.1 with a finite menu $\mathcal{M}$, the cap f and at most Q trials succeeds with probability at least $\eta_\star$ conditional on* GoodKG*, then*

$$\begin{aligned} \eta_\star &\le Q \cdot \binom{N}{\tau - g} \cdot p^{-g}, \\ g \log_2 p &\le \log_2 \binom{N}{\tau - g} + \log_2 Q + \log_2(1/\eta_\star). \end{aligned} \tag{7.5}$$

The two inequalities are equivalent.

Proof. Every accepted witness a satisfies $Ha = v_i$ and $\mathrm{wt}(a) \le f$, so acceptance in a reached trial i implies $P_{H,f}(v_i) \ne \varnothing$ for this cap f. For a target v uniform in $\mathbb{F}_p^\tau$, the proof of lemma 7.5 gives $\Pr[P_{H,f}(v) \ne \varnothing] \le \mathbb{E}[X_f] = \binom{N}{f} \cdot p^{-g}$. The conditional history argument in the proof of corollary 7.4, applied to the event $P_{H,f}(v_i) \ne \varnothing$ in place of the event $v_i \in G_{\mathcal{M}}$, bounds the success probability by $Q \cdot \binom{N}{f} \cdot p^{-g}$. Hence $\eta_\star \le Q \cdot \binom{N}{\tau-g} \cdot p^{-g}$, and taking base two logarithms gives the second inequality. The argument uses the first moment alone. □

The proof bounds the capped solvability event and is independent of the menu. A smaller menu can only lower the success probability, while for a menu with regions of weight larger than f the hit event $v \in G_{\mathcal{M}}$ can contain targets outside the event bounded here. The event counted by the first moment contains every witness of weight at most f, also when a decoding comparison names the exact weight f. The condition uses the exact logarithm $\log_2 p$, and for $\eta_\star = 1/2$ its second inequality is the first inequality of equation (7.4) with $\kappa = \log_2 Q + 1$.

We call a gap $g \in \{1, 2, \ldots, \tau - 1\}$ passing if it satisfies equation (7.5). If some gap passes, let $g_{\max}$ denote the largest passing gap, and if no positive gap passes, $g_{\max}$ is undefined, since the value 0 would correspond to the excluded cap $f = \tau$. For consecutive gaps $1 \le g \le \tau - 2$ and $N < p$,

$$\frac{\binom{N}{\tau-g-1} \cdot p^{-(g+1)}}{\binom{N}{\tau-g} \cdot p^{-g}} = \frac{\tau - g}{p \cdot (N - \tau + g + 1)} < 1,$$

so the bound $Q\binom{N}{\tau-g}p^{-g}$ decreases strictly with g. The passing gaps therefore form an initial segment of $\{1, 2, \ldots, \tau - 1\}$, and the largest passing gap determines all of them.

7.3 Coordinate family separation

The first moment condition of lemma 7.5 also locates every viable weight relative to the certified radii of the three coordinate routes of section 4.6, namely the Berlekamp–Welch ceiling F_{BW} of section 4.6.1, the projected IRS ceiling F_{IRS} of equation (4.13) and the finite Guruswami–Sudan ceiling $F_{\mathrm{GS}}^{\mathrm{fin}}$ of equation (4.16). The parameters are those of that section, with $2 \le k < N$ and $N \le p$, and this is the standing range of the proposition.

The proposition bounds the three ceilings by one expression in N and k and compares that expression with the cap on the gap g. In a growing fixed rate family it obtains a distance between the certified radii and the viable weights that is a positive fraction of N. Both conclusions follow from the envelope of equation (7.6) and from lemma 7.5 alone.

Proposition 7.7 (Finite and asymptotic separation of the coordinate routes)**.** *Let*

$$F_{\mathrm{coord}} \coloneqq \max\{F_{\mathrm{BW}}, F_{\mathrm{IRS}}, F_{\mathrm{GS}}^{\mathrm{fin}}\}.$$

For every parameter pair in the standing range with $2 \le k < N$*,*

$$F_{\mathrm{coord}} < N - \sqrt{N(k-1)}. \tag{7.6}$$

Suppose H *is full spark,* v *is uniform, and* $1 \le f = \tau - g \le \tau$ *satisfies* $\Pr[P_{H,f}(v) \neq \varnothing] \ge 2^{-\kappa}$*. If*

$$\sqrt{N(k-1)} - k \ge \frac{N+\kappa}{\log_2 p}, \tag{7.7}$$

then $f > F_{\mathrm{coord}}$*.*

Moreover, consider a sequence of such instances with $N \le p$*,* $N \to \infty$*,*

$$\frac{k}{N} \longrightarrow \beta \in (0,1), \qquad \frac{\kappa}{N \log_2 p} \longrightarrow 0.$$

Then equation (7.7) holds for all sufficiently large members of the sequence, and

$$\frac{f}{N} \longrightarrow 1-\beta, \qquad \limsup_{N\to\infty} \frac{F_{\mathrm{coord}}}{N} \le 1 - \sqrt{\beta} < 1 - \beta. \tag{7.8}$$

Proof. The finite GS formula gives

$$F_{\mathrm{GS}}^{\mathrm{fin}} < N - \sqrt{N(k-1)}.$$

For BW,

$$F_{\mathrm{BW}} \le \frac{N-k}{2} < N - \sqrt{N(k-1)},$$

where strictness follows from $(N+k)^2 - 4N(k-1) = (N-k)^2 + 4N > 0$. For the IRS ceiling of equation (4.13), with the two parameter regions of definition 4.5, write

$$\xi_N \coloneqq \frac{k-1}{N-1}.$$

First let r belong to the first IRS parameter region. Dividing its real radius by $N-1$ gives

$$\frac{g_1(N,k;r)}{N-1} = \frac{r}{r+1} - \frac{r\xi_N}{2}.$$

For $x = \sqrt{\xi_N}$,

$$\begin{aligned} 1 - x - \left(\frac{r}{r+1} - \frac{rx^2}{2}\right) &= \frac{1}{r+1} + \frac{rx^2}{2} - x \\ &\ge \frac{1}{r+1} - \frac{1}{2r} = \frac{r-1}{2r(r+1)} \ge 0. \end{aligned}$$

The middle inequality evaluates the quadratic at its minimum $x = 1/r$. The strict integer ceiling in the first region therefore satisfies

$$F_1^{\mathrm{IRS}}(N,k) \le (N-1)(1 - \sqrt{\xi_N}). \tag{7.9}$$

Now let r belong to the second region. Its defining strict inequality is

$$\xi_N r(r+1) > 2.$$

If $r\sqrt{\xi_N} \le 1$, then

$$\xi_N r(r+1) \le \frac{r+1}{r} \le 2,$$

which is a contradiction. Hence $r\xi_N > \sqrt{\xi_N}$, and

$$\frac{g_2(N,k;r)}{N-1} = 1 - r\xi_N < 1 - \sqrt{\xi_N}. \tag{7.10}$$

Taking the maximum over both regions in equations (7.9) and (7.10) yields

$$F_{\mathrm{IRS}} \le (N-1) - \sqrt{(N-1)(k-1)} < N - \sqrt{N(k-1)}.$$

Indeed, the difference between the final two right hand sides is

$$1 - \frac{\sqrt{k-1}}{\sqrt{N} + \sqrt{N-1}} > 0.$$

When $\mathcal{R}_2(N, k)$ is empty, $F_{\mathrm{IRS}} = F_1^{\mathrm{IRS}}(N, k)$ and equation (7.9) alone gives the same bound. This establishes equation (7.6).

By lemma 7.5,

$$g \le \frac{N+\kappa}{\log_2 p}.$$

Under equation (7.7),

$$f = N - k - g \ge N - k - \frac{N+\kappa}{\log_2 p} \ge N - \sqrt{N(k-1)} > F_{\mathrm{coord}}.$$

For the asymptotic statement, $N \le p$ and $N \to \infty$ imply $\log_2 p \to \infty$. Hence

$$\frac{\sqrt{N(k-1)} - k}{N} \longrightarrow \sqrt{\beta} - \beta > 0, \qquad \frac{N+\kappa}{N \log_2 p} \longrightarrow 0.$$

Thus equation (7.7) holds eventually. The viability bound also gives $g/N \to 0$, so $f/N = 1 - k/N - g/N \to 1 - \beta$. Finally, equation (7.6) gives $\limsup F_{\mathrm{coord}}/N \le 1 - \sqrt{\beta}$, and $\sqrt{\beta} - \beta > 0$ supplies the positive relative gap. □

The proposition classifies the three cited coordinate theorem regions, of which BW and finite GS realize $P_{H,f}$ pointwise inside their radii, while the heterogeneous IRS region keeps the distributional guarantee of theorem 4.8. The separation is a statement about the radii of these three theorems, and other solvers or theorem regions are outside it.

7.4 Folded packing and bounded range comparison

Throughout this subsection H is full spark, the target v is uniform in $\mathbb{F}_p^\tau$, the standing compatible folded setup of section 5 is assumed, $0 < \beta < 1$, and the exact weight $f = \tau - g$ lies in the range $1 \le f \le \tau$ of lemma 7.5. For a divisor m of N with $L = N/m$ and an integer $1 \le s \le m$, the pair (m, s) is admissible for Guruswami–Wang when the radius $\rho_{s,m}(\beta)$ of equation (5.17), which is the radius of equation (5.12) at $s_{\mathrm{GW}} = s$, is nonnegative. It is admissible for AHS when the strict budget $\lceil L\rho_{s,m}(\beta) \rceil - 1$ of equation (5.17) is nonnegative. For the common Guruswami–Wang and AHS radius, the first moment condition caps the gap g from above, and the radius deficit of equation (5.19) bounds it from below.

An exact weight f support meets at least $\lceil f/m \rceil$ blocks by equation (4.4), and some such support attains this number, as in proposition 5.1. Let $\epsilon_{\text{pack}} \coloneqq (\tau - m\lceil f/m \rceil)/N$. Since $1 - \beta = \tau/N$, the common radius has $L\rho_{s,m}(\beta) = (\tau - \Delta_{s,m}(\beta) \cdot N)/m$, so its deficit plays the role of the slack in equation (5.1). An integer b and a real number x satisfy $b \le \lfloor x \rfloor$ exactly when $b \le x$, and $b \le \lceil x \rceil - 1$ exactly when $b < x$. With $b = \lceil f/m \rceil$ and $x = L\rho_{s,m}(\beta)$, the inclusive Guruswami–Wang budget and the strict AHS budget therefore give

$$\begin{aligned} \text{GW:} \quad & \lceil f/m \rceil \le \lfloor L \cdot \rho_{s,m}(\beta) \rfloor \quad \iff \quad \Delta_{s,m}(\beta) \le \epsilon_{\text{pack}}, \\ \text{AHS:} \quad & \lceil f/m \rceil \le \lceil L \cdot \rho_{s,m}(\beta) \rceil - 1 \quad \iff \quad \Delta_{s,m}(\beta) < \epsilon_{\text{pack}}. \end{aligned} \tag{7.11}$$

The two budgets agree unless $L\rho_{s,m}(\beta)$ is an integer, and then the AHS budget is one block smaller. For such parameters the two routes have different regions, and for all parameters each route keeps its own output and running time contract. Since $m\lceil f/m \rceil \ge f$, one has $\epsilon_{\text{pack}} \le g/N$, and coverage of an optimally packed support of exact weight f requires $g \ge N\Delta_{s,m}(\beta)$ for Guruswami–Wang and $g > N\Delta_{s,m}(\beta)$ for AHS. When the success condition of corollary 7.6 is required as well, $g \le g_{\max}$, so the deficit of a covering radius is at most the optimistic ceiling $g_{\max}/N$, strictly for AHS, and this ceiling need not be approached by any admissible AHS parameter choice. Packing decides whether some support of exact weight f fits the region. Such a support need not carry a witness for the sampled target, so this condition leaves open the coverage of other supports and of witnesses of smaller weight.

The proof of proposition 5.9 bounds $N\Delta_{s,m}(\beta)$ from below using only equation (5.19) and $m - s + 1 \le N$. With these coverage requirements, that lower bound gives the radius requirement

$$g \ge N\Delta_{s,m}(\beta) \ge 2\sqrt{\beta(\tau + 2\beta)} - 3\beta, \tag{7.12}$$

with the first inequality strict for the AHS radius, as in proposition 5.9, and non-strict for the Guruswami–Wang radius. The right side grows like $\Theta(\sqrt{N})$ at fixed rate, whereas the cap of lemma 7.5 is $(N + \kappa)/\log_2 p$.

Theorem 7.8 (Analytic bounded range non-coverage)**.** *Assume the standing compatible folded setup, and let $m \mid N$ and $1 \le s \le m$ be admissible. Suppose an exact weight row $f = \tau - g$ satisfies $\Pr[P_{H,f}(v) \ne \varnothing] \ge 2^{-\kappa}$ for uniform v, and that its optimally packed support is covered by the radius $\rho_{s,m}(\beta)$. Then*

$$\frac{N + \kappa}{\log_2 p} \ge 2\sqrt{\beta(\tau + 2\beta)} - 3\beta, \tag{7.13}$$

with strict inequality in the AHS case. Consequently, if the reverse inequality holds, then no row with that solvability level is covered, for any admissible m and s.

Proof. Lemma 7.5 bounds g above by $(N + \kappa)/\log_2 p$, and equation (7.12) bounds the same g below by $2\sqrt{\beta(\tau + 2\beta)} - 3\beta$. Combining the two bounds gives equation (7.13), and the strictness follows from that of the AHS radius. □

The two routes differ only at the boundary of equation (7.13). Coverage by the Guruswami–Wang radius requires the left side to be at least the right side, and coverage by the AHS radius requires it to be strictly larger. A left side below the right side therefore excludes Guruswami–Wang coverage of every row at that solvability level, and a left side at most the right side excludes AHS coverage.

The theorem is a compatibility condition between two necessary requirements on the same gap. Solvability at level $2^{-\kappa}$ caps g from above and coverage of an optimally packed exact weight support bounds g from below, so both hold together only when the cap reaches the radius requirement. The statement applies to the common radius for every

admissible (m, s), and its reach is the range of lengths at which the two requirements conflict. Original GR and KRZSW keep the contracts of section 5.4.

Corollary 7.9 (Closed form length threshold)**.** *Fix $0 < \beta < 1$, assume $\kappa = \mathcal{O}(\log_2 p)$, and consider a family in which $N \to \infty$ and $\log_2 p \to \infty$. Then coverage as in theorem 7.8 of an exact weight row $f = \tau - g$ with $\Pr[P_{H,f}(v) \neq \varnothing] \geq 2^{-\kappa}$ for uniform v requires*

$$N \geq 4\beta(1-\beta)(\log_2 p)^2(1 - o(1)). \tag{7.14}$$

Replacing $\binom{N}{f} \leq 2^N$ by the exact first moment sharpens the leading condition to

$$N \geq \frac{4\beta(1-\beta)}{H_2(\beta)^2} \cdot (\log_2 p)^2(1 - o(1)),$$

where H_2 is the binary entropy.

Proof. We solve equation (7.13) for N. Its left side comes from the estimate $\binom{N}{f} \leq 2^N$. Write

$$\ell_p \coloneqq \log_2 p, \qquad A \coloneqq \beta(1-\beta), \qquad c \coloneqq \kappa + 3\beta\ell_p.$$

After substituting $\tau = N(1-\beta)$, simultaneous first moment viability and optimally packed exact weight coverage require

$$\frac{N + \kappa}{\ell_p} \geq 2\sqrt{AN + 2\beta^2} - 3\beta. \tag{7.15}$$

For AHS the inequality is strict. Adding 3β gives

$$\frac{N + c}{\ell_p} \geq 2\sqrt{AN + 2\beta^2},$$

again with strict inequality for AHS. Both sides are positive. Squaring therefore preserves the relation and yields

$$P_{\mathrm{len}}(N) \coloneqq (N + c)^2 - 4\ell_p^2(AN + 2\beta^2), \tag{7.16}$$

with $P_{\mathrm{len}}(N) \geq 0$ for GW and $P_{\mathrm{len}}(N) > 0$ for AHS. Define

$$D \coloneqq A^2\ell_p^2 - Ac + 2\beta^2.$$

If $D > 0$, the two roots of P_{len} are

$$N_\pm = 2A\ell_p^2 - c \pm 2\ell_p\sqrt{D}. \tag{7.17}$$

Since P_{len} is monic, the resulting necessary condition is

$$\begin{aligned} \text{GW}: &\quad N \leq N_- \text{ or } N \geq N_+, \\ \text{AHS}: &\quad N < N_- \text{ or } N > N_+. \end{aligned}$$

If $D = 0$, the GW inequality holds for every N, while the strict AHS inequality excludes the repeated root. If $D < 0$, then $P_{\mathrm{len}}(N) > 0$ for every real N, so this relaxed inequality alone excludes no length.

A finite statement keeps both branches, because the small branch also satisfies the necessary condition. For the cryptographic asymptotic, fix $0 < \beta < 1$, assume $\kappa = \mathcal{O}(\ell_p)$, and let both N and ℓ_p tend to infinity. Then

$$D = A^2\ell_p^2\big(1 + \mathcal{O}(1/\ell_p)\big),$$

and expansion of equation (7.17) gives

$$N_- = \mathcal{O}(1), \qquad N_+ = 4A\ell_p^2 + \mathcal{O}(\ell_p).$$

An unbounded family eventually leaves the small branch. Consequently,

$$N \geq 4\beta(1-\beta)(\log_2 p)^2(1-o(1))$$

is necessary for GW coverage, with the corresponding strict large root condition for AHS. This proves equation (7.14).

The coefficient can be sharpened by retaining the exact first moment. Write $C \coloneqq N/\ell_p^2$ for each member of the family. The entropy bound $\binom{N}{j} \leq 2^{NH_2(j/N)}$ for $0 \leq j \leq N$, the tangent line bound for the concave function H_2 at $1-\beta$ and the symmetry of binary entropy give

$$\log_2 \binom{N}{\tau - g} \leq NH_2(\beta) + g \log_2 \frac{1-\beta}{\beta}.$$

The radius requirement equation (7.12) and $\beta\tau = \beta(1-\beta)N$ give

$$g \geq 2\sqrt{\beta(1-\beta)N} - 3\beta.$$

Combining the first estimate with

$$g \log_2 p \leq \log_2 \binom{N}{\tau - g} + \kappa$$

gives $g\big(\ell_p - \log_2 \frac{1-\beta}{\beta}\big) \leq NH_2(\beta) + \kappa$, and the factor of g is positive for large ℓ_p. Inserting the lower bound on g, dividing by ℓ_p^2 and using $\kappa = \mathcal{O}(\ell_p)$ shows that $x \coloneqq \sqrt{C}$ satisfies

$$H_2(\beta)x^2 - 2\sqrt{\beta(1-\beta)}\big(1 + \mathcal{O}(1/\ell_p)\big)x + \mathcal{O}(1/\ell_p) \geq 0.$$

For large ℓ_p this quadratic has two real roots, of which the smaller is $\mathcal{O}(1/\ell_p)$ and the larger is $\big(2\sqrt{\beta(1-\beta)}/H_2(\beta)\big)\big(1 + \mathcal{O}(1/\ell_p)\big)$. A value of x at most the smaller root gives $N = x^2\ell_p^2 = \mathcal{O}(1)$, so an unbounded family eventually has x at least the larger root. Hence the sharper necessary leading condition is

$$C \geq \frac{4\beta(1-\beta)}{H_2(\beta)^2}\,(1-o(1)).$$

This is a necessary asymptotic condition on the length, and a decoding or attack threshold would need a sufficient condition as well. □

Since $b_p = \lceil \log_2 p \rceil$, the threshold has scale $\Theta(b_p^2)$. The inequalities themselves are stated with the exact logarithm $\log_2 p$, and the rounded bit length enters only this scale. Above the large root the two necessary conditions are compatible, and the comparison of this subsection is silent there, while corollary 7.11 of the next subsection holds at those lengths for the canonical Guruswami–Wang and AHS block regions. Like the proposition, the theorem and its corollary classify theorem regions.

7.5 Global uniform target GW/AHS witness scarcity

The region bounds of proposition 7.3 hold for any declared menu, and this subsection specializes them to the block regions certified by the common radius of Guruswami–Wang and AHS, for which we prove a statement at all lengths that needs no success condition. We first define the two integer budgets of that radius and bound the probability that a partition fixed independently of the target admits a witness inside either budget. We then take the union over the canonical partitions and show that the resulting probability is negligible under cryptographic scaling.

Throughout this subsection we assume $0 < \beta < 1$, that is, $\tau \le N - 1$. The Shamir matrix H is full spark, and this is the only property of it that the theorem below uses. For a divisor m of N and an integer $1 \le s \le m$, we use the common radius $\rho_{s,m}(\beta)$ of equation (5.17), which section 6 identifies with the Guruswami–Wang radius at $s_{\mathrm{GW}} = s$, and its deficit $\Delta_{s,m}(\beta)$ of equation (5.19). The block partition is the aligned partition $\mathcal{P}_m$ fixed there, with block weight wt_m.

The deficit limits the exact weight supports that fit the common radius, by proposition 5.9 and equation (7.11), and it also controls lower weight witnesses. The argument is about witness existence before either decoder is invoked, so it applies to both routes although their outputs differ. GW returns an affine candidate space, completed by the routes of section 5.4, whereas AHS returns a randomized complete list under its own contract.

For $\mathfrak{D} \in \{\mathrm{GW}, \mathrm{AHS}\}$, define

$$B_{\mathfrak{D}} \coloneqq \begin{cases} \lfloor L\rho_{s,m}(\beta) \rfloor, & \mathfrak{D} = \mathrm{GW}, \\ \lceil L\rho_{s,m}(\beta) \rceil - 1, & \mathfrak{D} = \mathrm{AHS}, \end{cases} \qquad d_{\mathfrak{D}} \coloneqq \tau - mB_{\mathfrak{D}}. \tag{7.18}$$

A pair (m, s) whose budget is negative certifies no region for that route and is absent from everything below, whereas a budget equal to zero is legitimate and is retained. The number $d_{\mathfrak{D}}$ counts the coordinates by which the packed budget $mB_{\mathfrak{D}}$ falls short of τ.

Theorem 7.10 (Canonical fixed partition GW/AHS witness scarcity)**.** *Let $H \in \mathbb{F}_p^{\tau \times N}$ be full spark, let $v \leftarrow \mathbb{F}_p^{\tau}$ be uniform, and fix $m \mid N$, $L = N/m$, $1 \le s \le m$, and $\mathfrak{D} \in \{\mathrm{GW}, \mathrm{AHS}\}$. Assume $B_{\mathfrak{D}} \ge 0$, and fix the corresponding ordered block partition independently of v. Then*

$$d_{\mathfrak{D}} \ge N \cdot \Delta_{s,m}(\beta) > 0 \tag{7.19}$$

for GW, with strict first inequality for AHS, and

$$\begin{aligned} \Pr[\exists a : Ha = v,\ \mathrm{wt}_m(a) \le B_{\mathfrak{D}}] &\le \binom{L}{B_{\mathfrak{D}}} p^{-d_{\mathfrak{D}}} \\ &\le 2^{-L((1-\beta)\log_2 p/2 - 1)}. \end{aligned} \tag{7.20}$$

Moreover, let

$$c_\beta \coloneqq \min\left\{ \frac{\sqrt{\beta(1-\beta)}}{2}, \frac{1-\beta}{4} \right\}.$$

If

$$N \ge \frac{9\beta}{1-\beta}, \qquad \log_2 p \ge \max\left\{ \frac{2}{\sqrt{\beta(1-\beta)}}, \frac{4}{1-\beta} \right\},$$

then the probability in equation (7.20) is at most

$$2^{-c_\beta (\log_2 p)\sqrt{N}}. \tag{7.21}$$

Proof. The GW radius in equation (5.12) equals $\rho_{s,m}(\beta)$, because $k/(L(m - s + 1)) = m\beta/(m - s + 1)$. The two integer conventions in equation (7.18) therefore give $mB_{\mathrm{GW}} \le N\rho_{s,m}(\beta)$ and $mB_{\mathrm{AHS}} < N\rho_{s,m}(\beta)$. Subtracting from $\tau = N(1-\beta)$ proves the inequalities of equation (7.19), whose right side is positive because the identity of equation (5.19) and $0 < \beta < 1$ give $\Delta_{s,m}(\beta) \ge (1-\beta)/(s+1) > 0$.

The partition is fixed independently of the uniform target, H is full spark, $B_{\mathfrak{D}} \geq 0$, and equation (7.19) gives $mB_{\mathfrak{D}} < \tau$, so the block bound of equation (7.2) applies at $B = B_{\mathfrak{D}}$. Its value is at most the unclipped expression $\binom{L}{B_{\mathfrak{D}}} p^{mB_{\mathfrak{D}}-\tau} = \binom{L}{B_{\mathfrak{D}}} p^{-d_{\mathfrak{D}}}$, which proves the first line of equation (7.20). Finally,

$$\Delta_{s,m}(\beta) \geq \frac{1-\beta}{s+1} \geq \frac{1-\beta}{m+1}$$

implies $d_{\mathfrak{D}} \geq (1-\beta)Lm/(m+1) \geq (1-\beta)L/2$. Combining this inequality with $\binom{L}{B_{\mathfrak{D}}} \leq 2^L$ gives the second line of equation (7.20).

We next retain both sources of deficit, under the two size hypotheses of the statement. If $L \leq \sqrt{N}$, then equation (7.19) and the deficit bound $N\Delta_{s,m}(\beta) \geq 2\sqrt{\beta(\tau+2\beta)} - 3\beta$ in the proof of proposition 5.9 give

$$d_{\mathfrak{D}} \geq 2\sqrt{\beta(\tau+2\beta)} - 3\beta \geq \sqrt{\beta(1-\beta)N}.$$

That deficit bound is derived there from the identity of equation (5.19) and $m - s + 1 \leq N$ alone, so it holds independently of the exact weight premise of that proposition. The second inequality holds because $\tau = N(1-\beta)$ and $N \geq 9\beta/(1-\beta)$. Thus

$$\begin{aligned}\log_2 \Pr[\exists a :\ Ha = v,\ \mathrm{wt}_m(a) \leq B_{\mathfrak{D}}] &\leq L - d_{\mathfrak{D}} \log_2 p \\ &\leq -\frac{\sqrt{\beta(1-\beta)}}{2}(\log_2 p)\sqrt{N}.\end{aligned}$$

The first step uses the first line of equation (7.20) with $\binom{L}{B_{\mathfrak{D}}} \leq 2^L$, and the second uses the preceding display, $L \leq \sqrt{N}$ and $\log_2 p \geq 2/\sqrt{\beta(1-\beta)}$. If $L > \sqrt{N}$, the estimate $d_{\mathfrak{D}} \geq (1-\beta)L/2$ instead gives

$$\begin{aligned}\log_2 \Pr[\exists a :\ Ha = v,\ \mathrm{wt}_m(a) \leq B_{\mathfrak{D}}] &\leq -\left(\frac{1-\beta}{2}\log_2 p - 1\right) L \\ &\leq -\frac{1-\beta}{4}(\log_2 p)\sqrt{N}.\end{aligned}$$

Here the second step uses $\log_2 p \geq 4/(1-\beta)$ and $L > \sqrt{N}$. Taking the weaker constant proves equation (7.21). □

The event of equation (7.20) is defined by the block budget alone, without the coordinate cap $\mathrm{wt}(a) \leq f$ of the attack. It therefore contains the event that the block restricted relation $P_{m,B}$ of equation (4.5) with $B = B_{\mathfrak{D}}$ is solvable, for every cap f, and the bounds of the theorem hold for that event as well. The event counts every witness inside the budget, including witnesses whose weight is below the exact weight of an attack row.

The theorem treats one pair (m, s), whereas the corollary below ranges over all positive divisors m of N, including $m = 1$, with their canonical partitions $\mathcal{P}_m$ from section 6. We call a pair (m, s) with $m \mid N$ and $1 \leq s \leq m$, taken with its canonical partition $\mathcal{P}_m$, a canonical pair, and we call it theorem admissible for the route $\mathfrak{D}$ when $B_{\mathfrak{D}} \geq 0$. The canonical envelope menu of definition 6.2 instead keeps one Guruswami–Wang budget for each divisor and adds a coordinate ball and Original GR envelopes.

Corollary 7.11 (Canonical GW and AHS witness scarcity)**.** *Consider a family indexed by λ in which $N = \mathrm{poly}(\lambda)$, $1 - \beta \geq c_0 > 0$, and $\log_2 p \geq c_1 \cdot \lambda$, for fixed constants c_0, $c_1 > 0$. Let v be uniform in $\mathbb{F}_p^\tau$, and fix the canonical partitions of all divisors $m \mid N$ independently of v. The probability that some theorem admissible canonical pair admits a witness inside its certified GW or AHS block budget is negligible in λ.*

Proof. For every real x, $\lceil x \rceil - 1 \le \lfloor x \rfloor$, so the AHS event for a fixed (m, s) is contained in the corresponding GW event. There are at most $\sum_{m \mid N} m \le N^2$ canonical pairs, including $m = 1$. Under the stated scaling, $(1 - \beta) \log_2 p/2 - 1 \ge c_0 c_1 \lambda/2 - 1$, and the right side is positive for all sufficiently large λ. For those λ, a union bound and equation (7.20), using $L \ge 1$, therefore give

$$\Pr[\text{some theorem admissible GW or AHS event occurs}] \le N^2 \cdot 2^{-((1-\beta)\log_2 p/2 - 1)},$$

which is negligible under the stated hypotheses, because N is polynomial in λ and the constants c_0 and c_1 of the preceding lower bound are fixed and positive. The union bound holds whatever the dependence between the events. □

The corollary covers every polynomial length under its scaling. The pair (m, s) may be chosen after the target is known, because the union bound covers every canonical pair and the canonical partitions are fixed independently of the target. Coordinate permutations chosen with knowledge of the target lie outside the corollary.

7.6 Numerical illustration for uniform targets

We close this section by evaluating its necessary conditions and witness region bounds on the six parameter pairs of section 6.4, with the field size p replaced by the numerical proxy 2^{252} under the convention of that subsection. Every comparison below is decided with exact integers and fractions, and the accompanying artifact [DY26] records the exact values.

Table 4 applies corollary 7.6 with the trial budgets $Q = 2^{32}$, 2^{64} and 2^{128} and the success target $\eta_\star = 1/2$, that is, with $\kappa = \log_2 Q + 1$ in equation (7.4). At the proxy the first inequality of equation (7.5) becomes the integer comparison $2^{252g} \le 2Q \cdot \binom{N}{\tau - g}$. Every gap at which the uniform alpha wrapper reaches the success target must pass it, and the table lists the largest passing gap $g_{\max}$ for each budget together with the rate ceiling $g_{\max}/N$. These are therefore necessary bounds, and a passing gap need not attain the target. The passing gaps form an initial segment of $\{1, 2, \ldots, \tau - 1\}$, as shown after corollary 7.6, so a passing gap followed by a failing one is the largest. A dash marks a budget at which no positive gap passes, so that $g_{\max}$ is undefined there, which happens on the two smallest pairs for all three budgets and on (196, 131) for $Q = 2^{32}$ and 2^{64}. Where a gap passes, the largest passing gap is at most four and the rate ceiling is at most $1/196$.

The last three columns compare these ceilings with two deficits of the common Guruswami–Wang and AHS radius, which have different optimizers. The least real deficit $\Delta_{\min} = \min_{m,s} \Delta_{s,m}(\beta)$ of equation (5.19) runs over all divisors m of N and all $1 \le s \le m$. The table rounds it to six decimals and lists the configuration (m, s) that attains it, and the artifact keeps the exact fraction. On every pair its minimizer has $m = N$, so that $L = 1$ and $L\rho_{s,m}(\beta) < 1$, and both integer budgets of equation (7.18) are zero there. The configuration of least real deficit therefore certifies only the budget zero, whereas the packed deficit $(\tau - \max mB)/N$ uses the configuration of largest certified packed capacity. On every pair this capacity is $N/2$ for both routes, a value that reflects rounding to whole blocks, as explained in remark 8.1, and it is attained already at $s = 4$, which gives the packed deficits 11/64, 11/64, 33/196, 43/256, 43/256 and 171/1024.

Table 4: Necessary gap bounds for success probability at least $\eta_\star = 1/2$ in Q uniform trials, and GW/AHS radius deficits at the 2^{252} field size proxy.

(N, τ)	$g_{\max}$, $Q =$			$g_{\max}/N$, $Q =$			$\Delta_{\min}$	(m, s)	$(\tau - \max mB)/N$
	2^{32}	2^{64}	2^{128}	2^{32}	2^{64}	2^{128}			
$(64, 43)$	–	–	–	–	–	–	0.109375	$(64, 9)$	11/64
$(128, 86)$	–	–	–	–	–	–	0.079411	$(128, 14)$	11/64
$(196, 131)$	–	–	1	–	–	1/196	0.064972	$(196, 17)$	33/196
$(512, 342)$	1	2	2	1/512	1/256	1/256	0.040834	$(512, 29)$	43/256
$(768, 513)$	2	3	3	1/384	1/256	1/256	0.033478	$(768, 37)$	43/256
$(1024, 683)$	3	3	4	3/1024	3/1024	1/256	0.029078	$(1024, 42)$	171/1024

Every rate ceiling of the table lies below the least real deficit of its pair, by a factor of more than seven, and far below the packed deficit. By equation (7.11), the Guruswami–Wang radius covers an optimally packed support of exact weight $f = \tau - g$ only when $\Delta_{s,m}(\beta) \le \epsilon_{\text{pack}} \le g/N$, and the AHS radius only when the first inequality is strict. Hence no pair (m, s) of either route covers an optimally packed support of exact weight $\tau - g$ for any gap g that passes at these budgets. These comparisons are necessary conditions evaluated on the six pairs, so a sufficient attack parameter and a conclusion about other lengths lie outside them.

Table 5 bounds, per trial, the probability that a uniform target has a witness in the certified regions of each route. Each entry is a strict exponent ℓ_{bd}, so the bound lies below $2^{-\ell_{\text{bd}}}$, and it is computed from the exact sum before the conversion. The Berlekamp–Welch, IRS and Guruswami–Sudan columns are the coordinate bound of equation (7.1) at the ceilings F_{BW}, F_{IRS} and $F_{\text{GS}}^{\text{fin}}$. The Guruswami–Wang and AHS columns are menu unions, that is, sums of the block bound of equation (7.2) over all divisors at the largest nonnegative budget of the route, which suffices because the block balls of one partition are nested. The Original GR and KRZSW columns are coordinate envelopes of the whole route. Every witness that such a route certifies has at most as many nonzero coordinates as the largest packed capacity mB of the route's envelope, so the bound of equation (7.1) at that capacity covers all parameters of the route without a count of its configurations. For Original GR this bound uses $F_{\text{GR}}^{\text{env}}$, the largest packed capacity $m \cdot \overline{B}_m$ of the optimistic envelope of lemma 6.1, which is at least every certified Original GR capacity and takes the values 26, 53, 81, 224, 360 and 496. For KRZSW the bound uses $F_{\text{K}}^{\text{env}}$, the largest packed capacity mB_{K} over the certified KRZSW budgets, with the values 0, 0, 0, 128, 256 and 256.

The last column, headed "All", is the aggregate, which bounds the event that a target has a witness in any of the seven regions. The three coordinate balls are nested in the largest of them, which on these pairs is the Guruswami–Sudan ball, and the AHS menu lies inside the Guruswami–Wang menu. The aggregate therefore adds the Guruswami–Wang menu bound, the Guruswami–Sudan coordinate bound and the Original GR and KRZSW envelopes. It is a conservative bound for choosing among all seven routes after the target is known.

Table 5: Strict exponents of witness existence upper bounds per uniform trial at the 2^{252} field size proxy.

(N, τ)	BW	IRS	GS	GR	GW	KRZSW	AHS	All
$(64, 43)$	5488	5488	3720	4224	2770	10835	2770	2770
$(128, 86)$	10721	10721	7689	8194	5540	21671	5540	5540
$(196, 131)$	16456	16456	11907	12412	8312	33011	8312	8312
$(512, 342)$	42626	42626	31001	29234	21658	53517	21658	21658
$(768, 513)$	64063	63810	46373	37795	32478	64063	32478	32478
$(1024, 683)$	85249	85249	61998	46106	43062	106778	43062	43062

On every pair the exponent of the aggregate equals that of the Guruswami–Wang menu, since the three added terms leave the bit length of the exact sum unchanged. By corollary 7.4 with this bound as U, a uniform alpha wrapper whose menu region lies in the union of these seven regions succeeds within $Q = 2^{128}$ trials with probability below 2^{-2642}, 2^{-5412}, 2^{-8184}, 2^{-21530}, 2^{-32350} and 2^{-42934} on the six pairs at the proxy.

The aggregate and the count of section 6.4 both bound a whole menu that covers the seven regions, each under its own sampling rule. The aggregate bounds the probability that a target uniform in $\mathbb{F}_p^\tau$ has a witness in the regions and uses coordinate envelopes for Original GR and KRZSW. The line count bounds the covered points of one chosen line and uses the block envelopes of section 6 for Original GR. Both are union bounds over all members of their menus, because a solver may choose any member after seeing the target, and the bound of one selected configuration or the largest bound of a single route need not bound that union.

8 Further finite decoder comparisons

This section compares the 24 attack rows of Crites and Stewart and broader diagnostic grids with the certified decoder regions, using the accompanying artifact [DY26], and appendix B gives the eight detailed tables. Every reported numerical value is an exact arithmetic evaluation at the field size proxy $p = 2^{252}$, which stands for a prime $p \approx 2^{252}$ as in section 6.4. The rows belong to the original uniform attack experiment, that is, to the uniform sampling rule of definition 7.1. The two evaluations of the wrapper bounds stand beside the results they evaluate, in section 6.4 for the chosen prechallenge alpha wrapper of definition 3.1 and in section 7.6 for uniform targets.

8.1 Attack paper rows

For the numerical comparison we suppress the negligible correction in the Crites–Stewart lower bound of equation (4.14) and use

$$\widehat{L}_{\mathrm{att}}(N,\, \tau,\, p,\, f) \coloneqq \frac{\binom{N}{f}}{p^{\tau-f} + \binom{N}{f}}. \tag{8.1}$$

This is the proxy used to reproduce Table 1 of [CS25], and since it omits the explicit $1 + \mathcal{O}(2^{-\lambda})$ term of the theorem, it serves only for numerical evaluation and does not replace the theorem's lower bound. For a uniform target, define

$$x_{\mathrm{att}} \coloneqq -\log_2 \widehat{L}_{\mathrm{att}}, \qquad \ell_f \coloneqq \log_2 \mathbb{E}[Z_f] = \log_2 \left(\binom{N}{f} \frac{(p-1)^f}{p^\tau} \right). \tag{8.2}$$

The first quantity is a lower bound proxy exponent, and the second is the base two exponent of the exact average list formula of equation (4.6) for weight f coefficient vectors, evaluated at the field size proxy. This average is the expected size of the exact weight list and need not equal the list size on any fixed target, and by equation (4.7) its exponent differs from the support span first moment exponent by the exact term $f \cdot \log_2(1 - 1/p)$.

For each parameter pair, table 6 gives one row for x_{att} and one for ℓ_f, with $f = \tau - g$ and $g \in \{1, 2, 3, 4\}$. A positive ℓ_f means that the average exact weight list exceeds one, and a negative value means that it is below one. An x_{att} entry displayed as 0.00 is a positive exponent rounded to zero, and the proxy there is close to one and below it. The computed exact list exponent is positive in precisely six displayed cells,

$$(512,\, 342,\, g = 1), \qquad (768,\, 513,\, g \in \{1, 2\}), \qquad (1024,\, 683,\, g \in \{1, 2, 3\}).$$

This finite output scale observation illustrates the output bound of equation (4.8), which bounds the cost of enumerating the complete list, and it leaves open whether an algorithm that seeks one witness must enumerate that list or run for exponential time.

In table 7, the Berlekamp–Welch, optimized IRS and finite Guruswami–Sudan columns are the largest integral coordinate error weights certified by their respective theorems, namely F_{BW} of section 4.6.1, F_{IRS} of equation (4.13) and $F_{\mathrm{GS}}^{\mathrm{fin}}$ of equation (4.16). The symbol $r^\star$, by contrast, records every admissible value of the IRS interleaving parameter r that attains the displayed IRS ceiling, and r is a parameter of the decoder and not of the threshold signature scheme. The final column records $\tau - 4$, the smallest of the four tested corruption levels and not a decoder radius.

The table uses the less restrictive original finite condition $(N-f)^2 > N(N-\tau-1)$ behind equation (4.16), whereas the rate form Guruswami–Sudan envelope $(N - f)^2 > N(N - \tau)$ would give a ceiling at most as large and hence a weaker non-overlap test. The finite condition raises the ceiling by one in four of the six rows and leaves it unchanged in the other two, so the non-overlap test below uses the larger finite radius. For the six attack paper pairs (N, τ), the smallest tested corruption level is $f = \tau - 4$, and exact evaluation gives

$$\tau - 4 > F_{\mathrm{GS}}^{\mathrm{fin}} \geq \max\{F_{\mathrm{BW}}, F_{\mathrm{IRS}}\}$$

in every row. Consequently, none of the 24 tested corruption levels lies inside any of these three certified coordinate decoder radii, and since F_{IRS} is the maximum over the admissible interleaving parameters, no admissible parameter choice of the three coordinate routes covers a tested row.

These large values of f should be read with the agreement interpretation of section 4.1. When $f = \tau - g$, the normalized decoding instance asks for agreement on $k + g$ unknown positions, so the rows are genuine decoding requests near the interpolation threshold.

8.2 Diagnostic sweep

The broader diagnostic sweep tests whether the coordinate decoder separation persists away from the six attack paper pairs, using

$$N \in \{64,\, 128,\, 256,\, 512,\, 1024\}, \qquad \beta_0 \in \left\{\tfrac{1}{4},\, \tfrac{1}{3},\, \tfrac{1}{2},\, \tfrac{2}{3}\right\}, \qquad k = \lfloor \beta_0 \cdot N \rfloor, \qquad \tau = N - k,$$

and the three target levels $\xi \in \{2^{-128},\, 2^{-64},\, 1/2\}$. For each parameter target combination, define

$$f_{\min}(N, \tau, p; \xi) \coloneqq \min\left\{0 \leq f < \tau : \widehat{L}_{\mathrm{att}}(N, \tau, p, f) \geq \xi\right\}, \tag{8.3}$$

when the set is nonempty, so that $f_{\min}$ is the first corruption level at which the numerical lower bound proxy reaches ξ. For comparison, define

$$F_{\mathrm{coord}}^{\max} \coloneqq \max\left\{F_{\mathrm{BW}}, F_{\mathrm{IRS}}, F_{\mathrm{GS}}^{\mathrm{fin}}\right\}.$$

The artifact evaluates all 60 parameter target combinations using exact integer and rational arithmetic, and table 8 records $F^{\max}_{\mathrm{coord}}$ and the three values of $f_{\min}$ for each of the 20 base rows, with an absent entry where the set in equation (8.3) is empty. On this grid, $F^{\max}_{\mathrm{coord}} = F^{\mathrm{fin}}_{\mathrm{GS}}$ for every base row. The set is empty in 27 combinations, and in the other 33 combinations

$$f_{\min} > F^{\max}_{\mathrm{coord}},$$

with smallest difference 25. The numbers of base rows reaching 2^{-128}, 2^{-64} and $1/2$ are respectively 13/20, 12/20 and 8/20, with coordinate decoder overlap 0/20 at each target.

Every corruption level reaching a target is at least its $f_{\min}$, so none of the tested combinations simultaneously reaches its selected probability target and lies in a theorem supported Berlekamp–Welch, optimized IRS or finite Guruswami–Sudan radius. The equality between the Guruswami–Sudan ceiling and $F^{\max}_{\mathrm{coord}}$ is specific to this grid. The additive IRS bound of theorem 4.8 leaves this finite conclusion unchanged, because every target reaching row already has $f_{\min} > F^{\max}_{\mathrm{coord}} \geq F_{\mathrm{IRS}}$, so the IRS radius condition fails before the additive probability bound becomes relevant.

An empty target set says only that the lower bound proxy stays below ξ for every $f < \tau$, which gives no upper bound on witness existence. The observation applies to the cited theorem regions on this finite grid, and proposition 7.7 treats growing families by its own argument.

8.3 Folded support geometry versus theorem reach

Folded diagnostics report restricted witness existence bounds and support densities for block compatible choices. With the support density $\pi_{m,B,f}$ of equation (5.3), the restricted lower bound $L_{\text{FRS-att}}$ of theorem 5.4 and the exact restricted average list of equation (5.4), define

$$x_{\mathrm{FRS}} \coloneqq -\log_2 L_{\text{FRS-att}}, \qquad \ell_{m,B,f} \coloneqq \log_2 \mathbb{E}[Z_{m,B,f}] = \log_2 \left(M_{m,B,f} \cdot \frac{(p-1)^f}{p^\tau} \right).$$

The density is the fraction of exact f coordinate supports that touch at most B folded blocks, and it is a ratio of two counts rather than a factor by which the unrestricted witness probability may be multiplied. The exponent x_{FRS} describes the restricted second moment witness existence lower bound, not decoder success. The exponent $\ell_{m,B,f}$ is the base two logarithm of the exact average number of block compatible, weight f coefficient vectors solving a uniform target at a given field cardinality, and its displayed numerical value evaluates that formula at the proxy.

The generated folded support dataset contains 456 rows covering both diagnostic budget types, all four gaps and every tested divisor $m > 1$ of each attack paper length, and table 9 shows its projection at $m = 2$ and $g \in \{1, 4\}$ for the one coordinate margin budget

$$B = \left\lfloor \frac{\tau - 1}{m} \right\rfloor, \tag{8.4}$$

which ensures

$$mB \leq \tau - 1, \qquad \epsilon_{\max} \coloneqq \frac{\tau - mB}{N} > 0,$$

at the relaxed capacity envelope level. This budget describes geometry only, and it leaves open whether a cited finite decoder theorem accepts the resulting m and $\epsilon_{\max}$. An absent entry of the table means that no support of exact weight f fits the budget, which happens at $g = 1$ for (128, 86) and (512, 342) because there $\lceil f/2 \rceil$ exceeds B.

For a concrete reading, take $(N, \tau, g) = (1024, 683, 1)$. Then

$$f = 682, \quad m = 2, \quad L = 512, \quad B = 341, \quad mB = f.$$

Every compatible support must therefore fill all coordinates in its 341 touched blocks, and hence

$$M_{2,341,682} = \binom{512}{341}.$$

Although its support density is only approximately 3.28×10^{-142}, one has

$$\ell_{2,341,682} = \log_2 \left(\binom{512}{341} \cdot \frac{(p-1)^{682}}{p^{683}} \right) \approx 213.76.$$

Thus the restricted second moment lower bound is extremely close to one and x_{FRS} rounds to 0.00, while the average restricted exact weight list is approximately $2^{213.76}$. These values are compatible because a tiny fraction of the enormous support universe can still produce a very large absolute candidate family. The two conclusions describe existence and average output scale for the uniform target model under the hypotheses of theorem 5.4, and the conditional success probability of a folded decoder is a different quantity.

8.4 Finite instantiations of the folded decoder theorems

The preceding calculation treats support geometry only. The artifact also performs an exhaustive finite search for Original GR, Guruswami–Wang, KRZSW and AHS over the admissible parameter domain stated below, and every radius comparison in it uses exact integer or rational arithmetic. Under the standing conditional folded setup of section 5, let $\mathcal{A}_{\mathfrak{D}}(N, \tau)$ be the set of candidates θ, each consisting of a divisor $m \mid N$ and source range auxiliary parameters for route $\mathfrak{D}$, that satisfy that route's finite radius predicates of section 5.4. The Original GR set also enforces the internal root safety condition of equation (5.10). When this set is nonempty, the selected candidate is

$$\theta^{\star}_{\mathfrak{D}} \in \operatorname*{arg\,max}_{\theta \in \mathcal{A}_{\mathfrak{D}}(N,\tau)} \{m(\theta) \cdot B_{\mathfrak{D}}(\theta)\}, \tag{8.5}$$

with a deterministic tie break recorded in the generated data. Thus the selection maximizes the certified packed coordinate budget, an objective that differs from minimum runtime, maximum witness probability and a complete attack certificate, and the generated record retains the full selected candidate.

The comparison weight is

$$f_0 \coloneqq \tau - 4,$$

the smallest of the four tested weights. For every benchmark pair, all four route specific admissible sets above are nonempty, although the KRZSW budget is zero on the three shortest pairs and therefore certifies no block error there. Table 10 lists the selected candidate of every route on every pair, with its block budget, its packed budget mB, the deficit $f_0 - mB$ and the certified list bound where the source supplies one. Let $\mathfrak{D}$ range over Original GR, Guruswami–Wang, KRZSW and AHS, and define the per pair envelope

$$E(N, \tau) \coloneqq \max_{\mathfrak{D}} \{m(\theta^{\star}_{\mathfrak{D}}) \cdot B_{\mathfrak{D}}(\theta^{\star}_{\mathfrak{D}})\}. \tag{8.6}$$

Thus $E(N, \tau)$ is the largest selected theorem certified packed coordinate budget among the four routes. Since each selected candidate maximizes the packed budget over the whole admissible set of its route, a perfectly packed support of weight f_0 lies outside the budget of every admissible candidate of every route whenever $f_0 - E(N, \tau) > 0$. Table 11 records $E(N, \tau)$, the routes attaining it and the deficit $f_0 - E(N, \tau)$ on the six pairs.

The displayed envelopes are all equal to $N/2$, which reflects the granularity of the block budget at the tested lengths. For an inclusive radius, $mB = m\lfloor L\rho \rfloor$ with $L = N/m$, so the certified packed fraction $mB/N = \lfloor L\rho \rfloor / L$ is the best approximation to the route radius from below by a rational with denominator L, and the strict AHS budget gives the best such approximation strictly below the radius. A large radius requires a large auxiliary parameter and hence, through $s \leq m$, a large folding, whereas a small rounding loss requires a large number L of blocks. The maximum over divisors of N balances these two effects, and its value depends on the divisor structure of the tested length. The equal envelopes and the deficits below are therefore finite observations on the tested lengths rather than an asymptotic separation.

Remark 8.1 (Granularity of the certified packed budget). Consider a folded route with relative block radius ρ and folding $m \mid N$, and let $L = N/m$. An inclusive radius, as for Guruswami–Wang, gives $B = \lfloor L\rho \rfloor$ and

$$N\rho - m < mB \leq N\rho.$$

A strict radius, as for AHS, gives $B = \lceil L\rho \rceil - 1$ and

$$N\rho - m \leq mB < N\rho.$$

The rounding loss is therefore at most m, and a route can convert a large radius into a large packed budget only when m is small relative to $N\rho$.

Every displayed deficit is positive, so an optimally packed exact weight f_0 support exceeds the corresponding theorem certified packed envelope by the displayed number of coordinates, and hence it exceeds every admissible route budget. This test is optimistic with respect to support geometry, because it uses the best packing of a support into blocks, and it tests block budgets only.

The last column of table 11 records $\mathcal{L}_{\min}$, the smallest list bound certified by any route attaining the envelope. Both Guruswami–Wang and AHS attain the displayed envelope, and the value $\mathcal{L}_{\min} = 4$ comes from the AHS route, so it is a list bound for AHS only. To display the finite scale through the visible factors of the asymptotic theorems, write s_{A} for the selected AHS parameter, write s_{GW} and $h_{\max}$ for the selected Guruswami–Wang parameter and the bound on its affine dimension $h \leq s_{\mathrm{GW}} - 1$ of equation (5.12), and define

$$W_{\mathrm{basis}} \coloneqq N^2 + s_{\mathrm{GW}}k^2, \qquad C_{\mathrm{sub}} \coloneqq \max_{0 \leq h \leq h_{\max}} \binom{N}{h}, \qquad W_{\mathrm{prune}} \coloneqq m s_{\mathrm{A}}^5 L.$$

The Guruswami–Wang basis stage costs $\mathcal{O}(W_{\mathrm{basis}})$ field operations by equation (5.11), and its coordinate subset pruning of proposition 5.8 examines at most C_{sub} coordinate subsets, with polynomial solve and verify work per subset. The AHS runtime of equation (5.18) is

$$T_{\mathrm{basis}} + \mathcal{O}\left(W_{\mathrm{prune}} \log(s_{\mathrm{A}}/\delta_{\mathrm{dec}})\right)$$

field operations, where $T_{\mathrm{basis}} = \widetilde{\mathcal{O}}(N)\,\mathrm{poly}(s_{\mathrm{A}})$ is the affine basis stage. Thus the three quantities displayed in table 12 are exact evaluations of visible theorem factors, not exact running times.

All selected Guruswami–Wang and AHS rows have $s = 4$ and $h_{\max} = 3$, so $C_{\mathrm{sub}} = \binom{N}{3}$ and $W_{\mathrm{prune}} = 4^5 N = 1024N$. For example, at $(N, \tau) = (128, 86)$ the Guruswami–Wang basis scale is 23,440, followed by at most 341,376 subset systems, while the visible AHS pruning term is $131{,}072 \log(4/\delta_{\mathrm{dec}})$ in addition to T_{basis}. The generated route record uses $\delta_{\mathrm{dec}} = 1/100$, and with the natural logarithm of the AHS proof the displayed factor is approximately 785,313, before the hidden constant in the $\mathcal{O}$ term. Both routes stay

within their polynomial time contracts at the selected parameters, but they certify only $mB = 64 < 82 = f_0$. The shortfall at weight 82 is therefore one of certified reach, since the completeness, runtime and list size guarantees of the theorems stop at this budget. The other per route bounds, including the $p^{s_{\mathrm{int}}}$ output of Original GR, appear in table 10.

For Guruswami–Wang and AHS, the deficits above persist without any restriction to parameters that keep a route inside its polynomial time regime. By the deficit bound in the proof of proposition 5.9, which uses only equation (5.19) and $m - s + 1 \le N$, every admissible $m \mid N$ and $1 \le s \le m$ satisfies

$$N\rho_{s,m}(\beta) = \tau - N\Delta_{s,m}(\beta) \le \tau - \left(2\sqrt{\beta(\tau + 2\beta)} - 3\beta\right),$$

which bounds the real radius before block integrality is imposed and irrespective of running time. On the six pairs the right side is 36.41, 76.32, 118.78, 321.66, 487.88 and 653.82, in each case below the corresponding f_0. Within this common radius family, accepting a larger list or a longer running time still leaves the tested weights out of reach, because every admissible parameter choice gives a real radius $N\rho_{s,m}(\beta)$ below f_0. The two routes share this radius family and keep their different output and runtime contracts.

Since $E(N, \tau) < f_0 = \tau - 4$ for every displayed pair, no admissible theorem budget of the four folded routes covers any of the four tested exact weights $\tau - g$, $g \in \{1, 2, 3, 4\}$. Together with the coordinate comparison of section 8.1, no theorem admissible parameter choice of any of the seven routes covers one of the 24 tested rows. The decoding procedures may still be executed on the received words, but the cited theorems no longer guarantee that their stated outputs contain every required close codeword at those weights. This comparison leaves their out of radius behavior and the existence of lower weight restricted witnesses undecided.

For Guruswami–Wang, KRZSW and AHS, the finite obstruction on these rows is radius coverage rather than an unbounded selected list. The Original GR prefix transfer of proposition 5.7 passes its finite root safety check on every route level row, and along the short prefix cryptographic family of section 5 its complete root stage retains the $p^{\mathcal{O}(s_{\mathrm{int}})}$ factor and therefore supplies no family level bit PPT certificate. The output and time distinctions appear in table 1, and the oracle and finite statuses appear in table 2.

These computations also illustrate theorem 7.8 on the tested grid, since the exhaustive search over every divisor and every route admissible auxiliary parameter yields positive coverage deficits on all six benchmark pairs, and the deficits persist when the polynomial time restriction is dropped. The theorem itself composes lemma 7.5 with the radius requirement of equation (7.12) and is proved without the benchmark rows.

Every folded row above is conditional on the compatible ordered power domain assumed for the exhaustive search. Such a row with polynomial decoding certifies an attack only when the block restricted solvability event and its composition with decoder success also satisfy the attack requirements of section 4.3. The positive deficits in table 11 are finite decoder applicability observations, and the support bounds of section 8.3 bear on attack distribution viability.

8.5 Fixed partition witness scarcity on the benchmark pairs

The artifact evaluates the fixed partition statement of theorem 7.10 with a second selection objective. Let $\mathcal{A}_{\mathfrak{D}}(N, \tau)$ be the admissible parameter set of section 8.4 for $\mathfrak{D} \in \{\mathrm{GW}, \mathrm{AHS}\}$, with the budgets $B_{\mathfrak{D}}$ of equation (7.18), and define

$$U_1 \coloneqq \max_{\substack{\mathfrak{D}\in\{\mathrm{GW},\mathrm{AHS}\} \\ (m,s)\in\mathcal{A}_{\mathfrak{D}}(N,\tau)}} \binom{N/m}{B_{\mathfrak{D}}} \cdot p^{mB_{\mathfrak{D}}-\tau}. \tag{8.7}$$

This maximization differs from maximizing mB. For a fixed m, the events obtained by varying s are nested in B, so with the largest admissible Guruswami–Wang budget $B^{\max}_{\mathrm{GW},m}$ at that m, as in equation (6.1), we define

$$U_{\mathrm{menu}} \coloneqq \min\left\{1,\ \sum_{\substack{m\mid N\\ B^{\max}_{\mathrm{GW},m}\geq 0}} \binom{N/m}{B^{\max}_{\mathrm{GW},m}} \cdot p^{mB^{\max}_{\mathrm{GW},m}-\tau}\right\}. \tag{8.8}$$

Table 13 reports the candidate attaining U_1, its syndrome codimension $d^\star = \tau - m^\star B^\star$ and both negative base two logarithms, and Guruswami–Wang and AHS select the same worst candidate on these six pairs. This candidate can differ from the maximum mB selection of equation (8.5) within a capacity tie. Candidates of equal packed capacity mB share the factor $p^{mB-\tau}$, so the probability aware objective ranks them by the binomial coefficient $\binom{N/m}{B}$, while the capacity objective resolves the tie by its deterministic rule. At (768, 513) the probability aware candidate is $(m, s) = (24, 5)$ and the maximum mB selection is (32, 4), from a capacity tie among several divisors. Both attain the packed capacity $N/2$, so each budget is half the number N/m of blocks, and this central binomial coefficient is larger for the finer partition $m = 24$. The distinction changes no conclusion here, but the bound that holds uniformly over all single candidates is U_1, since the bound of the candidate of largest packed capacity can be smaller than U_1.

The menu bound U_{menu} is the quantity behind the Guruswami–Wang column of table 5, which reports it as a strict integer exponent, whereas table 13 gives its negative base two logarithm rounded to two decimals. Both U_1 and U_{menu} are upper bounds on restricted witness existence for a uniform target, not conditional decoder success probabilities, and the finite calculations illustrate theorem 7.10, which is proved independently of them.

9 Conclusion and scope

The main result of this paper is geometric. For pairwise distinct nonzero Shamir points and any finite menu of coordinate and block regions fixed independently of the targets, theorem 3.5 bounds the number of targets on every line in direction K_0 that admit a witness in the menu by the explicit integer $M_{\mathrm{line}}(\mathcal{M})$ of equation (3.3). This integer depends on the code parameters and the menu but not on the field. The wrapper of definition 3.1 follows the Crites–Stewart template under its own rules. In each trial it chooses an explicit coefficient vector from its history and a message that was never queried, makes one fresh signature hash query and keeps that representation of the commitment afterwards. Its solvers make no further signature hash query and may search the whole menu, and the wrapper accepts only a witness of the same target that lies in the menu and within the corruption cap. It makes no signing query, corrupts nobody before it accepts a witness, and on acceptance corrupts exactly the support of that witness, outputs the forgery and stops. By lemma 3.2 its target is a uniform point of a line chosen by the adversary. Corollary 3.7 bounds its success within Q trials by $\min\{1, Q \cdot M_{\mathrm{line}}(\mathcal{M})/p\}$, with recovery inside the regions perfect and free, through a proof that conditions on the history of every trial and holds with arbitrary dependence between the trials.

The finite application evaluates this bound on the canonical envelope menu of definition 6.2. That menu contains the finite Guruswami–Sudan coordinate region, into which the Berlekamp–Welch, interleaved and KRZSW regions fall on the six benchmark pairs, and, at each divisor of N that admits one, the largest Guruswami–Wang budget, whose region contains every AHS region, and an optimistic Original GR envelope. By the evaluation of section 6.4, the count M_{line} is below 2^{24} on all six pairs, so for a prime p the

wrapper succeeds within Q trials with probability at most $Q \cdot M_{\text{line}}/p < Q \cdot 2^{24}/p$. At the field size proxy $p = 2^{252}$ this bound is below 2^{-100} after 2^{128} challenges. Reading a block region as the certified region of a folded decoder keeps the conditional ordered evaluation domain of section 5 as a hypothesis, whereas the count itself uses only the Shamir points.

The original uniform sampling rule gives much smaller bounds, because there every coordinate by which a region falls short of the threshold costs a factor of p. Corollary 7.4 bounds Q uniform trials by Q times the one trial bound, and for the seven routes on the six pairs the one trial bounds of table 5 are below 2^{-2770} at the proxy. Corollary 7.6 turns a required success probability into a necessary condition on the gap $g = \tau - f$, and section 7 compares the resulting slack with the exact packing of the common Guruswami–Wang and AHS radius. Beyond the tested pairs, proposition 7.7 shows that in growing fixed rate families with $N \leq p$ and $\kappa = o(N \log_2 p)$, every weight with witness probability at least $2^{-\kappa}$ eventually lies above the three certified coordinate radii. Theorem 7.8 and corollary 7.9 show that in an unbounded fixed rate family with $\log_2 p \to \infty$ and $\kappa = \mathcal{O}(\log_2 p)$, witness probability at least $2^{-\kappa}$ and exact weight coverage by the common Guruswami–Wang and AHS radius can hold together only at lengths at least quadratic in $\log_2 p$. Theorem 7.10 with corollary 7.11 bounds the probability that a canonical partition fixed independently of the target admits a witness inside the certified Guruswami–Wang or AHS budget. That probability is negligible at every polynomial length under cryptographic field scaling and rate bounded away from one. On the explicit grid of section 8, no theorem admissible parameter choice of any of the seven routes covers one of the 24 tested rows, and the six positive exact list exponents there illustrate the complete output question without deciding it.

On the decoding side the supported conclusion is therefore one of non-certification. Under their cited hypotheses, the Berlekamp–Welch, interleaved, Guruswami–Sudan and folded theorem families studied here do not yield a theorem supported polynomial time realization of the unrestricted relation $P_{H,f}$ on the tested regimes. Each of the seven routes misses, on every tested row, at least one of the four requirements of a complete decoder assisted attack certificate in the sense of section 4.3, namely decoder applicability, witness existence on the attack distribution, a justified probability composition and polynomial total running time. AHS attains the displayed folded envelope with a complete list bound of four, so its shortfall is a matter of certified reach rather than output size, whereas Guruswami–Wang attains the same envelope with the complete pruning bound $\binom{N}{3}$ of proposition 5.8. The conclusion applies to the theorem set studied here, and a future decoder, or a future analysis of an existing one, could change it.

These results bound the success of the wrapper of definition 3.1 and classify the studied theorems, and the implication of theorem 2.3 runs from a solver to a forgery. The hardness of the oracle $P_{H,f}$, the success of other attack strategies and the adaptive security of threshold Schnorr schemes are therefore open questions.

9.1 Limitations

The principal limitation is the open question of the average case complexity of finding one sparse witness for a uniform attack induced target in a constant rate code with $N \ll p$. This is a one witness search, in which an algorithm need not enumerate all nearby codewords, and it lies outside the cited decoder theorems and the hardness results of section 1.1.3. On the attack side, even a future polynomial time oracle algorithm would still need a non-negligible solvability event and a justified solver to forgery composition on the attack distribution, as equation (2.3) records.

The analytic radius comparison of theorem 7.8 is bounded in range, since it composes two necessary conditions that exclude each other below a length quadratic in $\log_2 p$,

equivalently of scale b_p^2, and is silent above that length. Beyond that length the canonical fixed partition Guruswami–Wang and AHS branch is still covered by corollary 7.11, which bounds its restricted witness probability at every polynomial length under its scaling hypotheses, whereas the unrestricted oracle, target dependent coordinate permutations, Original GR and KRZSW lie outside that corollary.

9.2 Open problems

The folded branch leaves a concrete transfer problem, namely whether the unrestricted sparse oracle $P_{H,f}$ can be reduced, with useful probability and polynomial overhead, to a block compatible oracle $P_{m,B}$ for parameters covered by modern folded Reed–Solomon decoder theorems. The semi-adversarial literature of section 1.1.4 suggests a second transfer problem, in which one must either construct, from the uniform target attack experiment, a distinguished codeword and received word channel satisfying the relevant semi-adversarial theorem, or analyze the decoder directly on the attack distribution. Such a theorem must keep the conditioning on witness existence, handle overlapping supports, preserve any virtual interleaving or folded code setup, and account for the complete runtime and output validation probability. This transfer is beyond the cited semi-adversarial results as they stand.

Finally, extending the complete multi-route comparison into an end to end concrete family would require an explicit prime field family realizing the standing setup assumption uniformly, together with uniform control of success probability, list enumeration and runtime. The separation of proposition 7.7, the bounded range comparison of theorem 7.8 and the scarcity bound of corollary 7.11 leave three parts of such a family open, namely a compatible folded setup, an unrestricted oracle transfer for folded decoding and a route uniform end to end runtime theorem.

Acknowledgments

Use of AI tools. The research questions, mathematical results, proofs and decisions in this paper are our own. We used AI tools, namely Anthropic Claude (Opus 5.5 and Fable 5.1) and OpenAI GPT models (GPT-6 Astra, GPT-6.1 Sol and GPT-5.6 Sol), to study the literature and the decoding theorems, to draft and revise parts of the manuscript text from our results and instructions, to check proofs, citations and numerical claims, and to write the code of the accompanying artifact, which we reviewed and validated with its test suite. We reviewed all text and code and take full responsibility for the content. Records of the instructions given to these tools are available on request.

A Solver to forgery derivation

This appendix gives the explicit adversary behind theorem 2.3 and the complete proof of that theorem. The construction is the attack of Crites and Stewart [CS25, Definition 10 and proofs of Theorems 1 and 2], written in the notation of section 2. It uses the standing assumptions and the event GoodKG of section 2.2, the public coefficient recovery of lemma 2.1 and the uniformity argument of lemma 7.2.

A.1 Explicit adversary

Let $\mathcal{S}$ be a solver as in section 2.4. The adversary $\mathcal{A}_{\mathcal{S}}$ receives the public key PK and the public verification shares and runs as follows. It first applies the formula of lemma 2.1 to τ verification shares and obtains elements $X_0, X_1, \ldots, X_{\tau-1}$, which are the coefficient

commitments on GoodKG. It then samples $\alpha \in \mathbb{F}_p^\tau$ uniformly and a message $M^\star$ that has not appeared in any signature hash query, forms $R^\star = \prod_{j=0}^{\tau-1} X_j^{\alpha_j}$, makes its only signature hash query $c^\star = \mathcal{H}_{\mathsf{sig}}(R^\star, PK, M^\star)$ and puts $v^\star = \alpha + c^\star K_0$. Afterwards it runs $\mathcal{S}$ once on $(H, v^\star, f)$ and receives a single output $\widehat{a}$. It checks $\widehat{a} \in \mathbb{F}_p^N$, $H\widehat{a} = v^\star$ and $\mathrm{wt}(\widehat{a}) \le f$, and it aborts when one of these checks fails. Only when all three hold does it corrupt exactly the parties in $F = \mathrm{supp}(\widehat{a})$, compute $s^\star = \sum_{i \in F} \widehat{a}_i \cdot sk_i$ from the revealed shares and output the forgery $(M^\star, (R^\star, s^\star))$.

The adversary makes no signing query and exactly one signature hash query, and it makes its at most f corruption queries only after the solver has returned a valid output. The solver makes no signature hash query, and every rerun, list computation or random choice that it performs happens inside the single call and counts toward its cost. The target $v^\star$ is produced by the fresh one query experiment of section 2.4, so the uniformity argument of lemma 7.2 applies to it.

A.2 Proof of the interface theorem

Proof of theorem 2.3. We show that every execution in the event of equation (2.2) is a winning execution of $\mathcal{A}_{\mathcal{S}}$ in the adaptive TS-UF-0 experiment, and that $\mathcal{A}_{\mathcal{S}}$ is PPT. The adversary carries out the one query experiment and the solver call of equation (2.2) with the same randomness, so the event of equation (2.2) is an event of its executions.

Consider an execution in which GoodKG holds and $\widehat{a} = \mathcal{S}(H, v^\star, f)$ is valid. On GoodKG the computed elements are the commitments $X_j = g^{x_j}$, where $x = (x_0, x_1, \ldots, x_{\tau-1})^\top$ is the coefficient vector of the Shamir polynomial, and $sk_i = K_i^\top x$ for every party. Since $\widehat{a}$ vanishes outside F, the value $s^\star$ computed by the adversary is the sum over all parties, and validity gives $H\widehat{a} = v^\star$, so

$$\begin{aligned} s^\star &= \sum_{i=1}^{N} \widehat{a}_i \cdot sk_i = \sum_{i=1}^{N} \widehat{a}_i \cdot K_i^\top x = (H\widehat{a})^\top x = (v^\star)^\top x \\ &= \alpha^\top x + c^\star \cdot K_0^\top x = \alpha^\top x + c^\star x_0. \end{aligned}$$

Since $X_0 = PK$, it follows that

$$g^{s^\star} = g^{\alpha^\top x} \cdot g^{c^\star x_0} = \Big(\prod_{j=0}^{\tau-1} X_j^{\alpha_j} \Big) \cdot PK^{c^\star} = R^\star \cdot PK^{c^\star}.$$

With $c^\star = \mathcal{H}_{\mathsf{sig}}(R^\star, PK, M^\star)$ this is the Schnorr verification equation of section 2.2 for the signature $(R^\star, s^\star)$ on $M^\star$. The adversary makes no signing query, so $M^\star$ was never signed and the forgery is fresh. It corrupts the $|F| = \mathrm{wt}(\widehat{a}) \le f$ parties of the support and no other party, so it respects the corruption bound of the experiment. The execution therefore wins, and the advantage of $\mathcal{A}_{\mathcal{S}}$ is at least $\delta_{\mathcal{S}}(\lambda)$. Executions outside this event may also win, and the bound leaves them uncounted.

Outside the solver, the adversary inverts one $\tau \times \tau$ Vandermonde matrix of public points with $\mathcal{O}(\tau^3)$ field operations, performs $\mathcal{O}(\tau^2)$ group exponentiations for the commitments and τ more for $R^\star$, and makes one signature hash query. After the solver has returned, it uses $\mathcal{O}(N\tau)$ field operations to check $H\widehat{a} = v^\star$ and $\mathrm{wt}(\widehat{a}) \le f$, makes at most f corruption queries and uses $\mathcal{O}(f)$ field operations for $s^\star$. A field operation costs polynomially many bit operations in b_p, and an exponentiation takes $\mathcal{O}(b_p)$ group operations, which have polynomial cost in the scheme family because its Schnorr verification performs the same operations. Since N and b_p are polynomially bounded in λ and $\tau \le N$, these steps take time polynomial in λ, and together with the complete implementation of $\mathcal{S}$, which has probabilistic polynomial bit complexity by hypothesis, the adversary is PPT. If $\delta_{\mathcal{S}}$

is non-negligible, this PPT adversary has non-negligible advantage, which contradicts adaptive TS-UF-0 security of the scheme family. □

B Supporting finite tables

This appendix collects the eight generated tables of the finite evaluation on the rows of the attack paper and on the diagnostic grids. They are outputs of the accompanying artifact [DY26], described in appendix C, and section 8 interprets them. The first three tables cover the coordinate routes, where table 6 lists the proxy and list formula exponents of the attack rows, table 7 the certified coordinate ceilings on those rows, and table 8 the coordinate diagnostic sweep. The other five cover the folded routes, where table 9 lists support and exact list diagnostics at selected budgets and table 10 the selected folded theorem budgets at $f_0 = \tau - 4$. Table 11 lists their maximum over the routes, table 12 the finite work indicators of the selected Guruswami–Wang and AHS candidates, and table 13 the fixed partition restricted witness exponents that accompany section 7.5.

Table 6: Attack-paper proxy and list-formula exponents.

(N, τ)	quantity	$g = 1$	$g = 2$	$g = 3$	$g = 4$
$(64, 43)$	x_{att}	195.84	446.97	698.20	949.52
	ℓ_f	-195.84	-446.97	-698.20	-949.52
$(128, 86)$	x_{att}	137.87	388.92	640.02	891.17
	ℓ_f	-137.87	-388.92	-640.02	-891.17
$(196, 131)$	x_{att}	75.41	326.45	577.53	828.64
	ℓ_f	-75.41	-326.45	-577.53	-828.64
$(512, 342)$	x_{att}	0.00	37.25	288.28	539.32
	ℓ_f	213.76	-37.25	-288.28	-539.32
$(768, 513)$	x_{att}	0.00	0.00	53.80	304.82
	ℓ_f	448.22	197.21	-53.80	-304.82
$(1024, 683)$	x_{att}	0.00	0.00	0.00	69.29
	ℓ_f	683.76	432.75	181.74	-69.29

Table 7: Certified coordinate-error ceilings on the attack-paper rows.

(N, τ)	k	F_{BW}	F_{IRS}	$r^\star$	$F_{\mathrm{GS}}^{\mathrm{fin}}$	$\tau - 4$
$(64, 43)$	21	21	21	1, 2	28	39
$(128, 86)$	42	43	43	2	55	82
$(196, 131)$	65	65	65	1, 2	83	127
$(512, 342)$	170	171	171	2	217	338
$(768, 513)$	255	256	257	2	326	509
$(1024, 683)$	341	341	341	1, 2	433	679

Table 8: Coordinate-decoder diagnostic sweep.

N	β_0	k	τ	$F_{\mathrm{coord}}^{\max}$	$f_{\min}(2^{-128})$	$f_{\min}(2^{-64})$	$f_{\min}(1/2)$
64	1/4	16	48	33	–	–	–
64	1/3	21	43	28	–	–	–
64	1/2	32	32	19	–	–	–
64	2/3	42	22	12	–	–	–
128	1/4	32	96	65	–	–	–
128	1/3	42	86	55	–	–	–
128	1/2	64	64	38	63	–	–
128	2/3	85	43	24	–	–	–
256	1/4	64	192	129	191	191	–
256	1/3	85	171	109	170	170	–
256	1/2	128	128	75	127	127	–
256	2/3	170	86	47	85	85	–
512	1/4	128	384	257	382	383	383
512	1/3	170	342	217	340	340	341
512	1/2	256	256	150	254	254	254
512	2/3	341	171	94	169	169	170
1024	1/4	256	768	513	765	765	765
1024	1/3	341	683	433	679	680	680
1024	1/2	512	512	300	508	508	508
1024	2/3	682	342	188	338	339	339

Table 9: Selected folded-support and exact-list diagnostics.

(N,τ)	g	m	B/L	mB	$\pi_{m,B,f}$	x_{FRS}	$\ell_{m,B,f}$
$(64,43)$	1	2	21/32	42	1.61×10^{-9}	225.06	-225.06
$(64,43)$	4	2	21/32	42	3.45×10^{-6}	967.67	-967.67
$(128,86)$	1	–	–	–	0	–	–
$(128,86)$	4	2	42/64	84	1.88×10^{-15}	940.09	-940.09
$(196,131)$	1	2	65/98	130	9.12×10^{-28}	165.23	-165.23
$(196,131)$	4	2	65/98	130	4.67×10^{-23}	902.82	-902.82
$(512,342)$	1	–	–	–	0	–	–
$(512,342)$	4	2	170/256	340	2.29×10^{-67}	760.69	-760.69
$(768,513)$	1	2	256/384	512	9.99×10^{-107}	0.00	96.09
$(768,513)$	4	2	256/384	512	2.84×10^{-100}	635.51	-635.51
$(1024,683)$	1	2	341/512	682	3.28×10^{-142}	0.00	213.76
$(1024,683)$	4	2	341/512	682	2.21×10^{-135}	516.60	-516.60

Table 10: Selected folded-theorem budgets at $f_0 = \tau - 4$.

(N, τ)	route	parameters	B/L	mB	f_0	$f_0 - mB$	$\log_2 \mathcal{L}$
$(64, 43)$	GR	$(m, s, r) = (1, 1, 80)$	$26/64$	26	39	13	252.0
$(64, 43)$	GW	$(m, s, h_{\max}) = (32, 4, 3)$	$1/2$	32	39	7	15.3
$(64, 43)$	KRZSW	$m = 64, \epsilon = 11/16$	$0/1$	0	39	39	†
$(64, 43)$	AHS	$(m, s) = (32, 4)$	$1/2$	32	39	7	2.0
$(128, 86)$	GR	$(m, s, r) = (1, 1, 65)$	$53/128$	53	82	29	252.0
$(128, 86)$	GW	$(m, s, h_{\max}) = (32, 4, 3)$	$2/4$	64	82	18	18.4
$(128, 86)$	KRZSW	$m = 64, \epsilon = 87/128$	$0/2$	0	82	82	†
$(128, 86)$	AHS	$(m, s) = (32, 4)$	$2/4$	64	82	18	2.0
$(196, 131)$	GR	$(m, s, r) = (1, 1, 56)$	$81/196$	81	127	46	252.0
$(196, 131)$	GW	$(m, s, h_{\max}) = (49, 4, 3)$	$2/4$	98	127	29	20.2
$(196, 131)$	KRZSW	$m = 49, \epsilon = 33/49$	$0/4$	0	127	127	†
$(196, 131)$	AHS	$(m, s) = (49, 4)$	$2/4$	98	127	29	2.0
$(512, 342)$	GR	$(m, s, r) = (8, 2, 57)$	$28/64$	224	338	114	504.0
$(512, 342)$	GW	$(m, s, h_{\max}) = (32, 4, 3)$	$8/16$	256	338	82	24.4
$(512, 342)$	KRZSW	$m = 128, \epsilon = 215/512$	$1/4$	128	338	210	†
$(512, 342)$	AHS	$(m, s) = (32, 4)$	$8/16$	256	338	82	2.0
$(768, 513)$	GR	$(m, s, r) = (24, 3, 217)$	$15/32$	360	509	149	756.0
$(768, 513)$	GW	$(m, s, h_{\max}) = (32, 4, 3)$	$12/24$	384	509	125	26.2
$(768, 513)$	KRZSW	$m = 256, \epsilon = 43/128$	$1/3$	256	509	253	†
$(768, 513)$	AHS	$(m, s) = (32, 4)$	$12/24$	384	509	125	2.0
$(1024, 683)$	GR	$(m, s, r) = (16, 3, 810)$	$31/64$	496	679	183	756.0
$(1024, 683)$	GW	$(m, s, h_{\max}) = (32, 4, 3)$	$16/32$	512	679	167	27.4
$(1024, 683)$	KRZSW	$m = 128, \epsilon = 107/256$	$2/8$	256	679	423	†
$(1024, 683)$	AHS	$(m, s) = (32, 4)$	$16/32$	512	679	167	2.0

Table 11: Maximum selected folded-theorem budgets at $f_0 = \tau - 4$.

(N, τ)	f_0	$\max_{\mathfrak{D}} mB_{\mathfrak{D}}$	maximizing route(s)	$f_0 - \max_{\mathfrak{D}} mB_{\mathfrak{D}}$	$\mathcal{L}_{\min}$
$(64, 43)$	39	32	GW, AHS	7	4
$(128, 86)$	82	64	GW, AHS	18	4
$(196, 131)$	127	98	GW, AHS	29	4
$(512, 342)$	338	256	GW, AHS	82	4
$(768, 513)$	509	384	GW, AHS	125	4
$(1024, 683)$	679	512	GW, AHS	167	4

Table 12: Finite work indicators for the selected GW and AHS candidates.

(N, τ)	GW $(m, s, h_{\max})$	W_{basis}	C_{sub}	AHS (m, s)	W_{prune}
$(64, 43)$	$(32, 4, 3)$	5,860	41,664	$(32, 4)$	65,536
$(128, 86)$	$(32, 4, 3)$	23,440	341,376	$(32, 4)$	131,072
$(196, 131)$	$(49, 4, 3)$	55,316	1,235,780	$(49, 4)$	200,704
$(512, 342)$	$(32, 4, 3)$	377,744	22,238,720	$(32, 4)$	524,288
$(768, 513)$	$(32, 4, 3)$	849,924	75,202,816	$(32, 4)$	786,432
$(1024, 683)$	$(32, 4, 3)$	1,513,700	178,433,024	$(32, 4)$	1,048,576

Table 13: Fixed-partition GW/AHS restricted-witness upper-bound exponents.

(N, τ)	$(m^\star, s^\star)$	$B^\star$	$d^\star$	$-\log_2 U_1$	$-\log_2 U_{\text{menu}}$
$(64, 43)$	$(32, 4)$	1	11	2771.00	2771.00
$(128, 86)$	$(32, 4)$	2	22	5541.42	5541.00
$(196, 131)$	$(49, 4)$	2	33	8313.42	8313.00
$(512, 342)$	$(32, 4)$	8	86	21658.35	21658.34
$(768, 513)$	$(24, 5)$	16	129	32478.84	32478.83
$(1024, 683)$	$(32, 4)$	16	171	43062.84	43062.84

C Artifact and reproducibility

The numerical results of this paper come from the accompanying artifact [DY26], which evaluates the finite formulas and theorem parameter searches of the paper and generates all eleven numerical tables. Every decision is made with exact integers and rational numbers, and logarithms and decimals are presentation values. The strict exponent ℓ_{bd} of a positive bound $u/2^v$ is v minus the bit length of u, so the bound lies strictly below $2^{-\ell_{\mathrm{bd}}}$. The artifact runs no decoder and measures no running time. It uses the field size proxy $p = 2^{252}$ of section 6.4 and, like the paper, takes the ordered compatible folded domain of section 5 as a standing conditional assumption. The README of the artifact maps each of the eleven numerical tables to its location in this paper and to the output that produces it, and gives the commands that reproduce and validate these outputs.

References

[AHS26] Vikrant Ashvinkumar, Mursalin Habib, and Shashank Srivastava. Algorithmic improvements to list decoding of folded Reed–Solomon codes. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2026)*, pages 880–898. SIAM, 2026. URL: https://arxiv.org/abs/2508.12548v2, doi:10.1137/1.9781611978971.35.

[BCDZ26] Joshua Brakensiek, Yeyuan Chen, Manik Dhar, and Zihan Zhang. Unique decoding of Reed–Solomon and related codes for semi-adversarial errors. In *53rd International Colloquium on Automata, Languages, and Programming (ICALP 2026)*, volume 374 of *LIPIcs*, pages 43:1–43:19. Schloss Dagstuhl, 2026. URL: https://arxiv.org/abs/2504.10399v3, doi:10.4230/LIPIcs.ICALP.2026.43.

[BCL+26] Renas Bacho, Yanbo Chen, Julian Loss, Stefano Tessaro, and Chenzhi Zhu. Adaptively secure partially non-interactive threshold Schnorr signatures in the AGM. In *Advances in Cryptology – EUROCRYPT 2026, Part I*, volume 16541 of *LNCS*, pages 33–62. Springer, 2026. URL: https://eprint.iacr.org/2025/1953, doi:10.1007/978-3-032-25291-3_2.

[BGL+25] Ruben Baecker, Paul Gerhart, Davide Li Calsi, Luigi Russo, Dominique Schröder, and Arkady Yerukhimovich. Fully adaptive FROST with identifiable aborts from AOMDL. Cryptology ePrint Archive, Paper 2025/1950, 2025. To appear at ASIACRYPT 2026. URL: https://eprint.iacr.org/2025/1950.

[BKY03] Daniel Bleichenbacher, Aggelos Kiayias, and Moti Yung. Decoding of interleaved Reed–Solomon codes over noisy data. In *Automata, Languages and Programming – ICALP 2003*, volume 2719 of *LNCS*, pages 97–108. Springer, 2003. doi:10.1007/3-540-45061-0_9.

[Bla08] Richard E. Blahut. *Algebraic Codes on Lines, Planes, and Curves: An Engineering Approach*. Cambridge University Press, 2008. doi:10.1017/CBO9780511543401.

[BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *Proceedings of the 1st ACM Conference on Computer and Communications Security*, pages 62–73. ACM, 1993. doi:10.1145/168588.168596.

[CCKM26] Michele Ciampi, Elizabeth Crites, Chelsea Komlo, and Mary Maller. On the adaptive security of key-unique threshold signatures. In *Advances in Cryptology – CRYPTO 2026, Part II*, volume 16801 of *LNCS*, pages 68–100. Springer, 2026. URL: https://eprint.iacr.org/2025/943, doi:10.1007/978-3-032-35374-0_3.

[CKGW24] Deirdre Connolly, Chelsea Komlo, Ian Goldberg, and Christopher A. Wood. The flexible round-optimized Schnorr threshold (FROST) protocol for two-round Schnorr signatures. RFC 9591, June 2024. URL: https://www.rfc-editor.org/info/rfc9591, doi:10.17487/RFC9591.

[CKK+25] Elizabeth Crites, Jonathan Katz, Chelsea Komlo, Stefano Tessaro, and Chenzhi Zhu. On the adaptive security of FROST. In *Advances in Cryptology – CRYPTO 2025, Part VI*, volume 16005 of *LNCS*, pages 480–511. Springer, 2025. URL: https://eprint.iacr.org/2025/1061, doi:10.1007/978-3-032-01887-8_16.

[CS25] Elizabeth Crites and Alistair Stewart. A plausible attack on the adaptive security of threshold Schnorr signatures. In *Advances in Cryptology – CRYPTO 2025, Part VI*, volume 16005 of *LNCS*, pages 457–479. Springer, 2025. All locators in this work refer to the full version, Cryptology ePrint Archive, Paper 2025/1001. URL: https://eprint.iacr.org/2025/1001, doi:10.1007/978-3-032-01887-8_15.

[CW07] Qi Cheng and Daqing Wan. On the list and bounded distance decodability of Reed–Solomon codes. *SIAM Journal on Computing*, 37(1):195–209, 2007. The statements used here are those of the extended abstract, arXiv:math/0405082v1 (2004). URL: https://arxiv.org/abs/math/0405082v1, doi:10.1137/S0097539705447335.

[CZ25] Yeyuan Chen and Zihan Zhang. Explicit folded Reed–Solomon and multiplicity codes achieve relaxed generalized Singleton bounds. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC 2025)*, pages 1–12. ACM, 2025. URL: https://arxiv.org/abs/2408.15925, doi:10.1145/3717823.3718114.

[DY26] Hüseyin Eren Demirtaş and Oğuz Yayla. Exact numerical artifact for threshold-Schnorr sparse-oracle decoding. Software, 2026. Version 1.0.0. Available from the authors.

[GR08] Venkatesan Guruswami and Atri Rudra. Explicit codes achieving list decoding capacity: Error-correction with optimal redundancy. *IEEE Transactions on Information Theory*, 54(1):135–150, 2008. URL: https://arxiv.org/abs/cs/0511072v2, doi:10.1109/TIT.2007.911222.

[GS99] Venkatesan Guruswami and Madhu Sudan. Improved decoding of Reed–Solomon and algebraic-geometry codes. *IEEE Transactions on Information Theory*, 45(6):1757–1767, 1999. Statement numbers cited in this work were taken from the authors' preprint dated May 11, 1999, whose Lemma 7, Theorem 8, Proposition 9 and Theorem 12 carry the same labels and statements in this journal version. doi:10.1109/18.782097.

[Gur06] Venkatesan Guruswami. Algorithmic results in list decoding. *Foundations and Trends in Theoretical Computer Science*, 2(2):107–195, 2006. doi:10.1561/0400000007.

[GV05] Venkatesan Guruswami and Alexander Vardy. Maximum-likelihood decoding of Reed–Solomon codes is NP-hard. *IEEE Transactions on Information Theory*, 51(7):2249–2256, 2005. Section numbers cited in this work refer to the arXiv version, cs/0405005. URL: https://arxiv.org/abs/cs/0405005, doi:10.1109/TIT.2005.850102.

[GW13] Venkatesan Guruswami and Carol Wang. Linear-algebraic list decoding for variants of Reed–Solomon codes. *IEEE Transactions on Information Theory*, 59(6):3257–3268, 2013. URL: https://eccc.weizmann.ac.il/report/2012/073/, doi:10.1109/TIT.2013.2246813.

[Joh62] S. M. Johnson. A new upper bound for error-correcting codes. *IRE Transactions on Information Theory*, 8(3):203–207, 1962. doi:10.1109/TIT.1962.1057714.

[KG20] Chelsea Komlo and Ian Goldberg. FROST: Flexible round-optimized Schnorr threshold signatures. In *Selected Areas in Cryptography – SAC 2020*, volume

12804 of *LNCS*, pages 34–65. Springer, 2020. URL: https://eprint.iacr.org/2020/852, doi:10.1007/978-3-030-81652-0_2.

[KRZSW23] Swastik Kopparty, Noga Ron-Zewi, Shubhangi Saraf, and Mary Wootters. Improved list decoding of folded Reed–Solomon and multiplicity codes. *SIAM Journal on Computing*, 52(3):794–840, 2023. Theorem and corollary numbers cited in this work refer to the arXiv preprint version 1, arXiv:1805.01498v1 (3 May 2018), titled "Improved decoding of Folded Reed–Solomon and Multiplicity Codes". URL: https://arxiv.org/abs/1805.01498v1, doi:10.1137/20M1370215.

[MS81] R. J. McEliece and D. V. Sarwate. On sharing secrets and Reed–Solomon codes. *Communications of the ACM*, 24(9):583–584, 1981. doi:10.1145/358746.358762.

[RS60] I. S. Reed and G. Solomon. Polynomial codes over certain finite fields. *Journal of the Society for Industrial and Applied Mathematics*, 8(2):300–304, 1960. doi:10.1137/0108018.

[RU10] Atri Rudra and Steve Uurtamo. Two theorems in list decoding. arXiv:1001.1781, 2010. URL: https://arxiv.org/abs/1001.1781.

[Sch91] C. P. Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991. doi:10.1007/BF00196725.

[Sha79] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979. doi:10.1145/359168.359176.

[Sri25] Shashank Srivastava. Improved list size for folded Reed–Solomon codes. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2025)*, pages 2040–2050. SIAM, 2025. URL: https://arxiv.org/abs/2410.09031, doi:10.1137/1.9781611978322.64.

[SSB09] Georg Schmidt, Vladimir R. Sidorenko, and Martin Bossert. Collaborative decoding of interleaved Reed–Solomon codes and concatenated code designs. *IEEE Transactions on Information Theory*, 55(7):2991–3012, 2009. URL: https://arxiv.org/abs/cs/0610074, doi:10.1109/TIT.2009.2021308.

[Tam24] Itzhak Tamo. Tighter list-size bounds for list-decoding and recovery of folded Reed–Solomon and multiplicity codes. *IEEE Transactions on Information Theory*, 70(12):8659–8668, 2024. URL: https://arxiv.org/abs/2312.17097, doi:10.1109/TIT.2024.3402171.

[Yay11] Oğuz Yayla. *On Decoding Interleaved Reed–Solomon Codes*. PhD thesis, Middle East Technical University, September 2011. URL: https://hdl.handle.net/11511/21239.