Reed–Solomon Codes Beyond Johnson: Efficient Decoding and Smaller Cryptographic Proofs

Quang Dao* Scott Duke Kominers† Justin Thaler‡

November 2, 2026

Abstract

List decoding Reed–Solomon codes is a central problem in coding theory and, together with mutual correlated agreement (MCA), underpins the soundness of many succinct cryptographic proofs. In this work, we give precise quantitative bounds on Reed–Solomon list decoding and MCA up to capacity, provide deterministic decoding algorithms that remain efficient over cryptographically large fields, and specialize our bounds to reduce proof sizes in implemented proof systems. Our bounds apply to arbitrary prescribed evaluation domains; beyond Johnson, they require sufficiently large field characteristic.

Refining the hidden-derivative method of Brakensiek, Chen, Putterman, Zhang, and Zheng, we obtain an explicit agreement threshold $a_1(\rho)$ strictly below the Johnson bound $\sqrt{\rho}$, using only the first derivative, without requiring its evaluations in the received word. At agreement $a_1(\rho) + \eta_1$, where $\eta_1 > 0$, our sharper interpolation and candidate counts give list size $O_\rho(n/\eta_1^2)$ and MCA error $O_\rho(n^2/(q\eta_1^4))$ for codes of length n over $\mathbb{F}_q$. Higher derivatives give explicit quantitative bounds up to capacity, uniformly over all rates, sharpening prior MCA results in Zheng's unpublished manuscript and Jeronimo's work, concurrent with ours. At fixed rate, for agreement gap $\eta_0 > 0$ above Johnson, we improve the MCA error bound of Ben-Sasson et al. (BCHKS) from $O_\rho(n/(q\eta_0^5))$ to $O_\rho(n/(q\eta_0^3))$, in every characteristic.

Our decoders determine the values of the message and its derivatives at a point from the agreements with the received word, rather than enumerating them over the field as earlier hidden-derivative decoders do. With fast explicit field arithmetic, at fixed rate and a fixed positive gap above the respective threshold, deterministic decoding takes $\widetilde{O}(n \log q)$ bit operations above Johnson in every characteristic and $\widetilde{O}(n^3 \log q)$ above our first-order curve in sufficiently large characteristic. Decoding remains polynomial in n and $\log q$ at every fixed positive gap from capacity, under the same requirement of sufficiently large characteristic.

Our first-derivative bounds give the first proof-size reductions in existing proof-system implementations from provable Reed–Solomon proximity-gap bounds beyond the Johnson radius. Relative to the deployed parameters and at unchanged security targets, we save 79.4 KB (11.1%) for ProveKit passport proofs, 13.2 KB (4.63%) for ZisK compressed final proofs, and 59.8 KB (4.85%) for LambdaVM CPU subproofs. We have formally verified the list-decoding and MCA bounds and concrete parameters in ArkLib, a Lean library for verified cryptographic proofs.

*Carnegie Mellon University. qvd@andrew.cmu.edu.

†Harvard University and a16z crypto research. kominers@fas.harvard.edu.

‡a16z crypto research and Georgetown University. jthaler@a16z.com.

Contents

1 Introduction **4**
- 1.1 Our Results . . . 5
- 1.2 Related work . . . 7

2 Technical Overview **9**
- 2.1 Hidden-derivative interpolation . . . 10
- 2.2 From an equation to lists and MCA . . . 11
- 2.3 First-order list and MCA bounds . . . 12
- 2.4 Higher derivatives approach capacity . . . 13
- 2.5 Decoding without enumerating initial values . . . 14

3 Hidden-derivative interpolation and the first-order curve **15**
- 3.1 Reed–Solomon codes, lists, and line MCA . . . 16
- 3.2 Local constraints and the interpolation principle . . . 17
- 3.3 One equation for every challenge . . . 19
- 3.4 The first-order support . . . 21
- 3.5 The first-order agreement curve . . . 22

4 From an equation to lists and MCA **25**
- 4.1 Incidence and common-witness transfer . . . 25
- 4.2 Ordinary equations in every characteristic . . . 28
- 4.3 The Johnson theorem . . . 29

5 First-order list and MCA bounds **30**
- 5.1 The bounds above the first-order curve . . . 30
- 5.2 Regular and singular solutions . . . 31
- 5.3 The counting bound . . . 33
- 5.4 Proofs of the two bounds . . . 34
- 5.5 Uniform bounds at gap $6/25$ . . . 35

6 Larger families of received words **35**
- 6.1 Affine families . . . 36
- 6.2 Interleaved codes . . . 37
- 6.3 Multilinear folds . . . 38

7 First-order bounds and smaller proofs **39**
- 7.1 From certificates to error bounds . . . 39
- 7.2 ProveKit: passport and lookup proofs . . . 40
- 7.3 ZisK: a smaller compressed final proof . . . 41
- 7.4 LambdaVM: smaller CPU subproofs . . . 42

8 Higher derivatives up to capacity **44**
- 8.1 The uniform capacity theorem . . . 44
- 8.2 A weighted support and the derivative-order condition . . . 45
- 8.3 Counting the solutions of an order-d explainer . . . 46
- 8.4 List and line-MCA bounds at a fixed rate . . . 49
- 8.5 Proof of the uniform theorem . . . 51

9 Efficient decoding over large fields **52**
9.1 Arithmetic model and notation . . . 53
9.2 Constructing an explainer . . . 53
9.3 Agreement recovery in finite algebras . . . 55
9.4 Order zero: a finite family at a regular center . . . 57
9.5 Taylor families at positive order . . . 59
9.6 First-order recovery from a curve . . . 63
9.7 Higher-order recovery by isolated systems . . . 67
9.8 Uniform capacity decoding . . . 69

10 Limits and open questions **70**

A Rational-image and Taylor tools **72**

B The ordinary transfer and the Johnson constants **75**

C First-order finite certificates **79**

D Deferred proofs for higher derivatives **85**

E Arithmetic for the decoding algorithms **88**

F Application details **94**

G Notes on prior work **97**

H The main statements in Lean **98**

AI methodology statement **99**

1 Introduction

For decades, a central question in Reed–Solomon list decoding was how to close the gap between the Johnson radius and capacity. A Reed–Solomon code of length n and rate $\rho = k/n$ encodes a degree-$< k$ polynomial by its evaluations at n points [RS60]. Given possibly corrupted evaluations, list decoding [Eli57, Woz58] returns every such polynomial agreeing with the received word in sufficiently many positions, so a lower required agreement fraction permits more errors. The Guruswami–Sudan algorithm recovers every polynomial agreeing on more than roughly a $\sqrt{\rho}$ fraction of positions [Sud97, GS99]. This agreement threshold is known as the Johnson bound; its exact finite form is $\sqrt{(k-1)/n}$. Capacity asks for agreement only $\rho + \delta$, for any fixed $\delta > 0$, with list size polynomial in n [GRS25, Section 7.4]. The challenge was to achieve this on arbitrary prescribed evaluation sets, with an efficient algorithm to recover the list.

These bounds also affect the size of succinct cryptographic proofs. Many proof systems use FRI [BBHR18], STIR [ACFY24a], or WHIR [ACFY24b] to check Reed–Solomon encodings at a few random positions; sharper correlated-agreement bounds can justify fewer queries at the same soundness level [BCI+23, BGKS19, GMW25]. Because these protocols combine words using random coefficients, their analyses need more than a bound on each word's list: they need to relate candidates before and after a combination.

Mutual correlated agreement (MCA) provides this connection: except for a small set of challenges, agreements in the combined word can be traced back to shared agreements in its two inputs. For $f, g \in \mathbb{F}_q^n$ and uniformly random $z \in \mathbb{F}_q$, MCA requires that, except with small probability, every sufficiently close polynomial P for $f + zg$ has the form

$$P = F + zG,$$

where F, G satisfy the same degree bound and agree with f, g, respectively, wherever P agrees with $f + zg$. Thus P's full agreement set equals the witnesses' common agreement set; we refer to this as *full agreement-set line MCA*. An exceptional challenge can create extra agreements through cancellation of the two inputs' discrepancies from their witnesses.

Work on correlated agreement and its mutual form established these guarantees up to the Johnson radius [BCI+23, BCH+26, Hab25, BCGM25]. Beyond the radius, the agreement threshold, list size, and probability of an exceptional combination all enter the resulting soundness bounds.

In a recent breakthrough, Brakensiek, Chen, Putterman, Zhang, and Zheng (BCPZZ) achieved Reed–Solomon list decoding up to capacity over prime fields [BCP+26a]. Their hidden-derivative method builds on Naskręcki's submission to the Proximity Prize on better.codes [Nas26]: derivatives constrain candidate polynomials even though their evaluations are absent from the received word. Corresponding MCA guarantees now also reach capacity [Zhe26, Jer26]; Zheng's manuscript states its guarantee at low rates, and padding extends it to all rates (Section G.1).

These results leave two gaps for proof systems. First, smaller proofs need explicit list-size and MCA-error bounds at practical field sizes and code lengths. Second, some security proofs run the decoder itself: the knowledge-soundness analyses of FRI-based proof systems [BGK+23], STIR [ACFY24a, Lemma 7.3], and WHIR [ACFY24b, Theorem 7.5] extract a witness by list decoding the prover's committed words. Over a field of size $q = 2^{\Theta(\lambda)}$, where λ is the security parameter, this extraction is efficient only if decoding takes time polynomial in n and $\log q$. The decoders of BCPZZ and Jeronimo instead enumerate the missing derivative values over the field, at cost polynomial in q [BCP+26a, Corollary 5.1], [Jer26, Theorem 4.4]. This brings us to our goal:

Can we decode beyond Johnson efficiently over cryptographically large fields, and obtain agreement bounds strong enough to reduce proof sizes in existing systems?

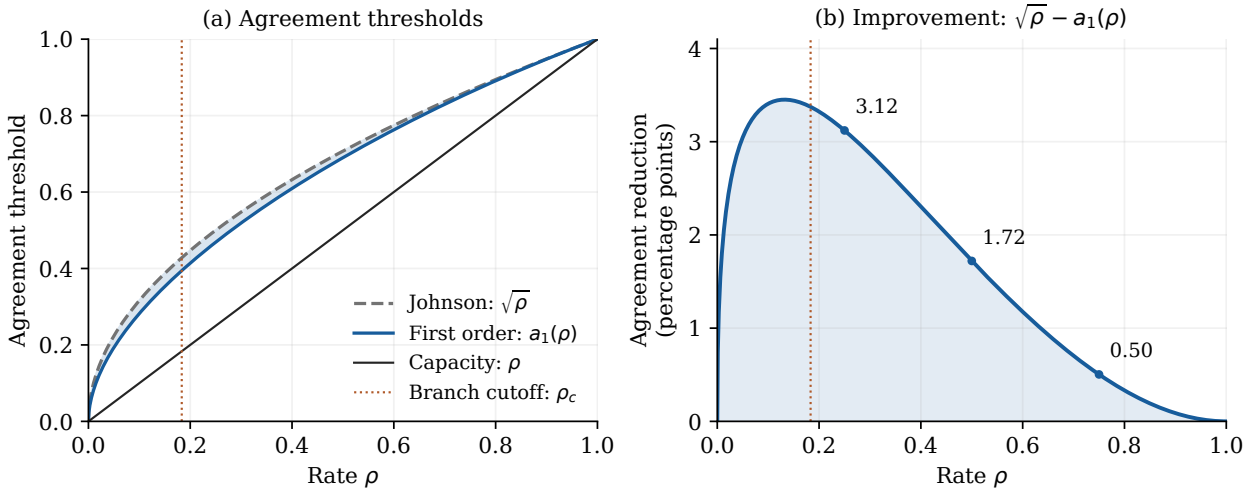

Figure 1: Left: the first-order threshold $a_1(\rho)$, Johnson's $\sqrt{\rho}$, and capacity ρ; the shaded region is the improvement. Right: $\sqrt{\rho} - a_1(\rho)$ in percentage points. The dotted line marks ρ_c, where a_1 changes branches.

1.1 Our Results

We make four main contributions in this paper. First, we show that one hidden derivative already gives explicit list-decoding and MCA bounds beyond Johnson, above a threshold curve $a_1(\rho)$; higher derivatives extend these bounds to capacity, and without derivatives we sharpen the MCA error at Johnson. Second, we give deterministic decoders that run in time polynomial in n and $\log q$ at each of these thresholds, in the characteristic our bounds require. Third, we use the first-derivative bounds to reduce query counts and proof sizes in ProveKit, ZisK, and LambdaVM at their existing security targets. Finally, we formally verify the mathematical bounds and concrete parameter certificates in Lean.

In the bounds below, subscripts on $O(\cdot)$ indicate the parameters on which its hidden constants may depend, and $\widetilde{O}(\cdot)$ suppresses logarithmic factors.

One derivative beyond Johnson. Write a for the required agreement fraction. Following Naskręcki and BCPZZ, we interpolate a polynomial equation $Q(X, P, P') = 0$ satisfied by every close message P; we restrict the monomials of Q and count the independent conditions they must satisfy. The resulting sufficient threshold is an explicit curve $a_1(\rho)$. For $\rho \geq \rho_c = 11 - 3\sqrt{13} \approx 0.183$,

$$a_1(\rho) = \frac{3\rho + 2\sqrt{\rho(5-\rho)(2-\rho)}}{8-\rho}, \tag{1}$$

and for smaller rates $a_1(\rho) = \sqrt{2\rho}\cos\left(\frac{1}{3}\arccos(\sqrt{\rho/8} - 1)\right)$ (Section 3.5). It lowers the threshold from 0.5 to approximately 0.4688 at rate $1/4$, and from 0.7071 to approximately 0.6899 at rate $1/2$ (Figure 1). At every rate, $\rho < a_1(\rho) < \sqrt{\rho}$, and $a_1(\rho) \sim \sqrt{\rho/2}$ as $\rho \to 0$, a factor $1/\sqrt{2}$ below Johnson. A lower agreement threshold reduces the query term of the soundness-error bound, so even a small improvement can save queries at a fixed security target, provided the algebraic error terms still meet their error targets (Section 7).

Write $\eta_1 = a - a_1(\rho) > 0$; below, B_∂ bounds the degree of Q in its derivative variable.

Agreement fraction a	List size	MCA exceptional count	Decoding time
Above Johnson $a = \sqrt{(k-1)/n} + \eta_0$ every characteristic	$O_\rho(\eta_0^{-1})$ classical [GS99]	$O_\rho(n\eta_0^{-3})$ *previously* $O_\rho(n\eta_0^{-5})$ [BCH+26]	$\widetilde{O}_{\rho,a}(n \log q)$ *previously* poly$(n, \log q)$ [CHK26]
One derivative $a = a_1(\rho) + \eta_1$ char $F = 0$ or $p > \max\{k-1, B_\partial\}$	$O_\rho(n\eta_1^{-2})$ *previously* earlier entries above if $a > \sqrt{\rho}$, else below with $\delta = a - \rho$	$O_\rho(n^2\eta_1^{-4})$ *previously* likewise	$\widetilde{O}_{\rho,a}(n^3 \log q)$ for $p > k-1$ *previously* likewise
Up to capacity $a \geq \rho + \delta$, $n \geq N(\delta)$ char $F = 0$ or $p > k-1$	$C_\delta n^{d(\delta)}$, where $d(\delta) = \lceil e^{1.5/\delta} \rceil$ if $\delta < 6/25$ and $d(\delta) = 1$ otherwise *previously* order $\exp(O_\rho(\delta^{-1} \log \delta^{-1}))$ [BCP+26a]	$C_\delta n^{d(\delta)+1}$ *previously* $n^{O_{\rho,\delta}(1)}$ at fixed rate, prime fields [Zhe26] (Section G.1); $n^{O_\delta(1)}$, prime $q \geq n$ or $p > \max\{k-1, B_\gamma\}$ [Jer26, Section 8.4]	$n^{O(d(\delta))}\, \widetilde{O}_\delta(\log q)$ for $n \geq N_\delta$ *previously* poly(n, q): prime $q \geq C_{\rho,\delta} n$ [BCP+26a]; prime $q \geq n$ [Jer26]

Table 1: Guarantees on any n distinct evaluation points, with earlier bounds where a direct comparison exists; $C_\delta, d(\delta), N(\delta)$ are as in Theorem 1.2, $N_\delta \geq N(\delta)$ is as in Theorem 9.14, B_γ is Jeronimo's jet-degree bound, with $\gamma = \delta$ and input-list bound $L = 1$, and a count $\mathsf{E}_{\mathrm{exc}}$ gives MCA error $\mathsf{E}_{\mathrm{exc}}/q$. The deterministic decoders (Corollaries 9.4 and 9.9 and Theorem 9.14) assume $q \geq n$ and the representation of Theorem 1.3.

Theorem 1.1 (First-order list decoding and MCA)**.** *Fix $0 < \rho < 1$ and $a = a_1(\rho) + \eta_1 < 1$, with $\eta_1 > 0$. Let $n \geq 2$, $2 \leq k \leq \rho n$, and $an \leq A \leq n$ be integers, where A is the required number of agreements. There is an explicit nonnegative integer $B_\partial = O_\rho((\eta_1 + 1/n)^{-1})$ such that the following holds. Over every field of characteristic zero or $p > \max\{k-1, B_\partial\}$, and every set of n distinct evaluation points, the agreement list has size $O_\rho(n/\eta_1^2)$ and full agreement-set line MCA has exceptional count $O_\rho(n^2/\eta_1^4)$. Over $\mathbb{F}_q$, divide the exceptional count by q for the soundness error.*

The characteristic condition comes from reconstruction: in characteristic p, both 0 and $(X-c)^p$ solve $P' = 0$ and vanish at c, so an initial value need not fix later coefficients; $p > k-1$ excludes this for degree-$< k$ messages, and $p > B_\partial$ gives separability in the derivative variable after removing repeated factors.

Higher derivatives approaching capacity. To approach capacity, we let the equation involve higher derivatives of P. The derivative order controls the exponents of n in our list-size, MCA, and decoding bounds, so we make it explicit and uniform over rates; first order suffices for $\delta \geq 6/25$ (Theorems 5.7 and 8.1).

Theorem 1.2 (Uniform capacity bounds, informal)**.** *Fix $0 < \delta < 1$, and put $d(\delta) = \lceil \exp(1.5/\delta) \rceil$ for $\delta < 6/25$ and $d(\delta) = 1$ otherwise. There are explicit constants $N(\delta), C_\delta$, depending only on δ and controlling the minimum length and counting bounds, such that for every $n \geq N(\delta)$, $1 \leq k \leq n$, and integer $k + \delta n \leq A \leq n$, the following holds. Over any field of characteristic zero or $p > k-1$, and any n distinct evaluation points, every received word has at most $C_\delta n^{d(\delta)}$ degree-$< k$ polynomials agreeing in at least A positions, and full agreement-set line MCA has exceptional count at most $C_\delta n^{d(\delta)+1}$. Over $\mathbb{F}_q$, the corresponding error is at most $C_\delta n^{d(\delta)+1}/q$.*

Quantitative MCA at Johnson. Without derivatives, at agreement $a = \sqrt{(k-1)/n} + \eta_0$ with $\eta_0 > 0$, our exceptional count is $O_\rho(n/\eta_0^3)$ in every characteristic, even those our first-order theorem excludes (Section 4.3).

System	Proof	Shipped	Johnson floor	Ours	Measured saving
ProveKit	Passport, BN254 (first round)	127	119	109	79.4 KB (11.1%)
ZisK	Compressed final proof	54	53	51	13.2 KB (4.63%)
LambdaVM	CPU subproof	219	217	208	59.8 KB (4.85%)

Table 2: Query counts at each system's security target and grinding; the Johnson floor is a query-term lower bound, attained only for ProveKit (Theorem 4.4). Measured savings are reductions in the listed proof objects relative to the shipped configurations; ProveKit's includes later rounds and internal proofs (Section 7).

Efficient decoding over large fields. Our decoders impose agreement with the received word before recovering coefficients (Section 2.5). Bit costs depend on how field elements are represented, and ours assume a polynomial basis: $q = p^e$, and $\mathbb{F}_q = \mathbb{F}_p[t]/(f(t))$ for an input irreducible f of degree e (Definition E.1 states what other representations must provide: fast field arithmetic, an indexing of the elements, and fast p-th roots).

Theorem 1.3 (Efficient decoding, informal)**.** *Fix $0 < \rho < 1$, $q \geq n$, distinct evaluation points in $\mathbb{F}_q$, and $1 \leq k \leq \rho n$. For $\sqrt{\rho} < a \leq 1$, the complete degree-$<$ k agreement list at threshold an can be found deterministically in $\widetilde{O}_{\rho,a}(n \log q)$ bit operations over every finite field given in such a polynomial basis. For $a_1(\rho) < a < 1$ and $p > k - 1$, the corresponding bound is $\widetilde{O}_{\rho,a}(n^3 \log q)$.*[1]

At derivative order d our decoders take $\widetilde{O}_{\rho,a,d}(n^{O(d+1)} \log q)$ (Corollary 9.12), and uniformly over rates at gap δ, for $p > k - 1$ and n at least some $N_\delta \geq N(\delta)$, $n^{O(\exp(1.5/\delta))}\widetilde{O}_\delta(\log q)$ (Theorem 9.14).

Families of received words. The proof systems we consider often combine many received words into one, and some reuse one random challenge across several combinations, so the agreement guarantees must hold for all of them at once. We extend line MCA to affine families and interleaved codes through the reductions of Bordage, Chiesa, Guan, and Manzur and of Jo [BCGM25, Jo26a], and separately bound combinations by powers of one challenge (Section 6). A depth-h multilinear fold, combining pairs by $(1 - r)u + rv$ with one random r per level, has error at most $h\mathsf{E}_{\mathrm{exc}}/q$ if each level has at most $\mathsf{E}_{\mathrm{exc}}$ exceptional challenges (Corollary 6.7).

Smaller proofs. We obtain the first proof-size reductions in existing proof-system implementations from provable Reed–Solomon proximity-gap bounds beyond Johnson. ProveKit, ZisK, and LambdaVM use large-characteristic challenge fields of about 254 and 192 bits, where first-order bounds apply (Table 2). The part beyond Johnson is 10 of ProveKit's 18 saved first-round queries and, against the Johnson floors, at least 2 of ZisK's 3 and 9 of LambdaVM's 11. ProveKit's lookup proofs shrink by a further 16.7 KB (10.3%).

Formal verification in Lean. We have formally verified our quantitative list-decoding and MCA bounds in ArkLib, a Lean library for verified cryptographic proofs [DA26]. The proofs cover both branches of the first-order curve, the uniform capacity bounds, and their affine, polynomial-curve, and interleaved extensions, together with the parameter inequalities and byte identities behind our ProveKit, ZisK, and LambdaVM reductions. Appendix H maps the main results to their Lean declarations.

1.2 Related work

The choice of code. In a proof system the protocol fixes the evaluation set, and the code is plain Reed–Solomon; this is the setting of BCPZZ and of this paper. Earlier results beyond Johnson relax one of these

[1] An earlier version [DKT26b] obtains $\widetilde{O}_{\rho,a}(n^2 \log q)$ with a refinement that we omit (Section 9.6.2).

two constraints. Some choose suitable evaluation sets [RW14, ST23, GLS+24, GST23]: Brakensiek, Gopi, and Makam prove capacity list bounds for random evaluation sets over large fields [BGM23], and later work reduces the field size to linear in n [GZ23, AGL24, AGG+25]. Others change the code: Parvaresh–Vardy, folded Reed–Solomon, and multiplicity codes reach capacity [PV05, GR08, KSY14, GW13, Kop15], and capacity MCA or correlated agreement is known for folded and random codes [GG25, GGSW26, YZ26].

Hidden derivatives up to capacity. The capacity results differ in how they reach every rate. BCPZZ first proved capacity at low rates and reached every fixed rate by padding the evaluation set, a reduction they credit to Alrabiah, Goyal, and Guruswami [BCP+26a, Section 5]. Mishra instead keeps the evaluation set and proves that a sufficient derivative order exists [Mis26], and Jeronimo makes one order serve every rate by enlarging the message degree bound and filtering the output [Jer26, Lemmas 2.2 and 4.1]. Jeronimo also treats polynomial combinations, list recovery, and a fixed multiplicative gap, where his order is polynomial in the inverse agreement [Jer26, Corollary B.1], as is ours for singleton lists, without his logarithmic factor (70). The orders differ as well: at a fixed rate and agreement $\rho + \delta$, BCPZZ's parameters give order $\exp(O_\rho(\delta^{-1}\log(1/\delta)))$ (Section G.2), whereas our order $\lceil \exp(1.5/\delta) \rceil$ is explicit and uniform over rates (Theorem 8.1).

These capacity results became public within three days.[2] The works disclose contemporaneous AI assistance, as does our methodology statement at the end of the paper.

Correlated agreement and MCA. BCIKS prove correlated agreement for all Reed–Solomon codes by decoding over a function field [BCI+23], Haböck and BCGM prove its mutual form [Hab25, BCGM25], and Ben-Sasson, Carmon, Haböck, Kopparty, and Saraf (BCHKS) bound the MCA error above Johnson by $O_\rho(n/(q\eta_0^5))$ in every characteristic [BCH+26, full version, Theorem 4.6]. Our $O_\rho(n/(q\eta_0^3))$ bound (Section 4.3) uses two degree-counting savings that BitWonka and jieyilong found for the Proximity Prize instance [Bit26a, jie26a].

Below Johnson, Jo and Chojecki obtain MCA bounds with different tradeoffs. Jo's theorem tolerates a fixed number of errors beyond the Johnson radius, with error $O_{r,h}(n^6/q)$ at rate $1/r$, so its relative improvement vanishes as n grows [Jo26b]; Chojecki obtains a fixed relative improvement, but with exceptional counts exponential in n [Cho26]. Our improvement is a fixed fraction, with error $O_\rho(n^2/(q\eta_1^4))$. For list decoding, Gabizon proves linear list-size bounds at a constant agreement fraction below Johnson in large characteristic [Gab26].

Unlike Jeronimo's symbolic interpolant, whose challenge degree is $O(n)$ [Jer26, Proposition 5.2 and Theorem 5.5], ours has challenge degree independent of n (66), and we count the agreement surplus and derivative degrees separately (Theorem 8.5 and Lemma C.1).

Extraction without list decoding. The extractors cited in the introduction list decode the prover's committed words. Arnon, Boneh, and Fenzi [ABF26, Lemmas 6.6 and 6.8] and Bünz, Chiesa, Fenzi, and Wang [BCFW25, Section 2.7] avoid list decoding for a toy protocol and for the accumulation scheme WARP: once a later message or the next round's witness fixes the combined codeword, MCA makes its agreement set common to the inputs, which are then recovered by erasure decoding on that set. The same strategy plausibly applies to FRI, STIR, and WHIR, which send the final polynomial in the clear, given MCA for each folding and batching step; to our knowledge, no such analysis has appeared. Its extraction error depends on the MCA error and the list size, which our bounds sharpen.

[2] On September 4, 2026, BCPZZ's low-rate paper and Zheng's low-rate MCA manuscript were public, Mishra's all-rate pull request was merged, and our ArkLib history had an initial all-rate list bound [BCP+26b, Zhe26, Mis26, Ark26b]. On September 5, BCPZZ posted their all-rate revision, Jeronimo submitted his all-rate paper, and our ArkLib history had line and affine MCA [BCP+26a, Jer26, Ark26a]. ECCC posted his paper on September 6.

Deterministic decoding. At order zero the obstacle to deterministic decoding is root finding. Fast Guruswami–Sudan interpolation and power-series lifting have near-linear arithmetic cost, but only a Las Vegas bit bound once root finding is charged [JNSV17], [NRS17, Theorem 1.2]. Received values avoid it: Augot and Pecquet use received pairs to initialize Hensel lifting [AP00], and Chatterjee, Harsha, and Kumar (CHK), crediting the Kalev–Ta-Shma multiplicity-code decoder [KTS22], decode deterministically in $\mathrm{poly}(n, \log q)$ time at every agreement above Johnson over every finite field [CHK26].

Root finding also enters the deterministic near-linear Johnson decoder of Goyal, Harsha, Kumar, and Shankar, through a routine whose base case factors univariate polynomials [GHKS24, Theorems 1.5 and 10.8], [Ale05, Appendix, Figure 5]. A recent manuscript produced by an OpenAI model claims deterministic factorization over prime fields $\mathbb{F}_p$ in time polynomial in the degree and $\log p$, with an exponent of about 10^{12}; it rests on a zero-free region for Hecke L-functions from a companion manuscript [Ope26, Theorems 1.1 and 11.1]. If the claim holds, such factoring steps become deterministic and polynomial-time over prime fields, but the claimed bound does not give near-linear bit complexity. Our zeroth-order decoder needs no root-finding oracle and improves CHK's running time at fixed rate and Johnson gap, while CHK allow every positive gap; in small characteristic we remove inseparability in the message variable following Lecerf [Lec08].

Beyond order zero, root finding alone no longer suffices, because the solutions of an explainer can form a family of positive dimension. For folded Reed–Solomon and multiplicity codes, whose symbols include the folded or derivative values, decoders select messages from affine candidate spaces [KRZSW23, GHKS24, AHS26, GHKS26]. Plain Reed–Solomon evaluations leave a nonlinear family instead, which the decoders of BCPZZ and Jeronimo enumerate by its initial derivative values; Section 2.5 explains how ours avoid this enumeration and which steps they inherit.

Proof-system parameters. Our applications use the coding-to-soundness connection of DEEP-FRI and later round-by-round analyses [BGKS19, GMW25]. Some implementations choose queries from conjectured capacity-level bounds or use different codes [RIS26, Win25, Plo26, Sta26a, Sta26b], where improving on Johnson alone does not shrink proofs. Section 10 compares our bounds with the Proximity Prize [ABF26]: on its fixed instance our certificate scores 66.15 bits, and a later prize submission 68.11 bits [Bit26d].

Organization. Sections 3 to 8 prove the bounds and apply them to proof systems, Section 9 computes the lists, and Section 10 discusses limits and open questions. The appendices give supporting tools, the parameter certificates, notes on prior work, and the main statements in Lean.

2 Technical Overview

The list bounds, the MCA bounds, and the decoders in this paper rest on the same three steps. First, we interpolate a polynomial equation that every close message satisfies together with its derivatives. Second, we describe the solutions of this equation by their initial Taylor coefficients, which places them in an algebraic family of small dimension and bounded degree. Third, we cut this family by the agreements with the received word, each a hyperplane in the space of message coefficients: once to bound the list, once to bound the exceptional challenges in MCA, and once to compute the list. The first two steps come from the hidden-derivative method of BCPZZ, which builds on Naskręcki, and from Kopparty's reconstruction for multiplicity codes; for the counting bounds, the third follows Zheng's symbolic solution families and incidence argument. Our bounds gain from sharper interpolation supports and candidate counts; Section 2.5 explains what the decoders add.

Fix a field F, distinct evaluation points $\alpha_1, \dots, \alpha_n \in F$, a received word $y \in F^n$, and an agreement threshold $k \le A \le n$. A *message* is a polynomial of degree at most $D = k - 1$, and it is *close* if it agrees with y in at least A positions. We use bounds $\rho \ge k/n$ on the rate and $a \le A/n$ on the agreement fraction, with

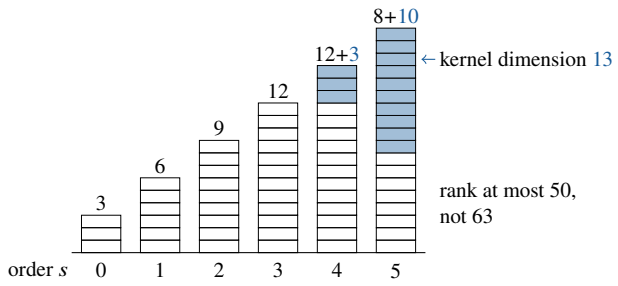

Figure 2: Local rank at one position, for $m = 6$ and $\deg_{Y_1} Q \le B_\partial = 2$. With $X = \alpha_i + T$ and $Y_0 = y_i + TU$, order $s < m$ has $(s+1)(B_\partial + 1)$ coefficients (cells), and Guruswami–Sudan with parameter Y_1 sets each to zero. The hidden-derivative constraint needs only $U = Y_1 + TV$, and then the multiples of $T^s(U - Y_1)^{m-s}$ vanish modulo T^m; shading counts their dimension. At most $63 - 13 = 50$ independent conditions remain.

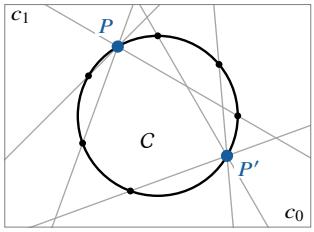

Figure 3: A message $P = c_0 + c_1 X$ for $k = 2$ is a point (c_0, c_1), and agreement at position i is the line $c_0 + \alpha_i c_1 = y_i$. No line contains the circle C, of degree $\Delta = 2$, so each meets it at most twice (dots). A close message on C (blue) lies on at least $A = 3$ of the $n = 6$ lines, so C contains at most $2 \cdot 6/3 = 4$ of them.

$0 < \rho < a < 1$. Interpolation and the results without derivatives hold in every characteristic. We take $k \ge 2$; constant codes are elementary (Lemma 3.3).

2.1 Hidden-derivative interpolation

Section 3 constructs one first-order equation that every close message satisfies, at every agreement above an explicit curve $a_1(\rho)$ below the Johnson bound, and keeps the challenge of a received line as a variable in it.

The equation. Guruswami–Sudan interpolation chooses a nonzero $Q(X, Y)$ with multiplicity m at every received pair [GS99]. If $\deg Q(X, P(X)) < mA$, then A agreements give mA roots and force $Q(X, P(X)) = 0$. The hidden-derivative construction of Naskręcki and of BCPZZ keeps this argument and adds a variable Y_1 for the derivative, of which no evaluation is received [Nas26, BCP+26a]. At an agreement $P(\alpha_i) = y_i$, backward Taylor expansion gives $P(\alpha_i + T) = y_i + TP'(\alpha_i + T) + T^2 V_P(T)$ with V_P a polynomial (Lemma 3.5). We therefore seek $Q(X, Y_0, Y_1)$ with

$$Q\big(\alpha_i + T,\ y_i + TY_1 + T^2 V,\ Y_1\big) \equiv 0 \pmod{T^m} \qquad (1 \le i \le n)$$

as an identity in the formal variables Y_1, V (Definition 3.6). Substituting $Y_1 = P'(\alpha_i + T)$ and $V = V_P(T)$ shows that $Q(X, P, P')$ vanishes to order m at every agreement (Lemma 3.7). If its degree is less than mA, then $Q(X, P, P') = 0$ for every close message; such a nonzero Q is a *hidden-derivative explainer* (Definition 3.8).

The source of the gain. Guruswami–Sudan interpolation imposes $\binom{m+1}{2}$ conditions at each position and reaches the Johnson threshold $a^2 > \rho$. Counted naively, the variable Y_1 multiplies both the number of coefficients and the number of conditions and does not lower the threshold. The gain comes from dependencies among the conditions at one position, which Naskręcki's construction already uses [Nas26]. Under the substitution above, the polynomial $L_i = Y_0 - y_i - (X - \alpha_i)Y_1$, linear in Y_0, Y_1, becomes $T^2 V$, so every product $(X - \alpha_i)^e L_i^j$ with $e + 2j \ge m$ meets all conditions at position i (Figure 2). Our new step bounds the local rank for supports that restrict the derivative exponent j_1 separately from the jet degree $j_0 + j_1$, the total degree in Y_0, Y_1. We take the monomials

$$X^x Y_0^{j_0} Y_1^{j_1}, \qquad j_1 \le B_\partial, \quad j_0 + j_1 \le B_{\text{jet}}, \quad x + D(j_0 + j_1) < mA$$

(Equation (19)), where the last inequality gives $\deg Q(X, P, P') < mA$. For these supports we bound the rank of the conditions at one position by

$$R_{\mathrm{loc}} = (B_\partial + 1)\left[\frac{m(m+1-B_\partial)}{2} + \frac{B_\partial(2B_\partial+1)}{6}\right] \qquad (B_\partial \le m/2)$$

(Lemma 3.14), while the naive count is $(B_\partial + 1)\binom{m+1}{2}$. With $\beta = B_\partial/m \le 1/2$, the dependencies remove an asymptotic fraction $\beta - 2\beta^2/3$ of the conditions. The coefficient gain is smaller than the naive factor $B_\partial + 1$, because a jet degree $t < B_\partial$ has only $t + 1$ monomials. An explainer exists whenever the number N_{coef} of coefficients exceeds nR_{loc} (Proposition 3.9), and optimizing the full piecewise rank bound against the coefficient count, with an asymptotically optimal ratio B_∂/m greater than $1/2$ on the low-rate branch, gives the curve $a_1(\rho)$: for every $a > a_1(\rho)$, explicit finite $(m, B_\partial, B_{\mathrm{jet}})$ satisfy $N_{\mathrm{coef}} > nR_{\mathrm{loc}}$ (Equation (24) and Lemma 3.16). For example, with $(n, k, A) = (1000, 250, 480)$ and $(m, B_\partial, B_{\mathrm{jet}}) = (24, 10, 46)$ there are 2,402,466 coefficients and at most $1000 \cdot 2365 = 2{,}365{,}000$ independent conditions, against 3,300,000 counted naively, although the agreement fraction 0.48 is below the Johnson threshold $\sqrt{249/1000} \approx 0.499$.

One equation for every challenge. For MCA we receive a line $f + zg$ and must relate the candidates at different challenges, which separate explainers for each word would not do. Following the function-field viewpoint of BCIKS and Zheng's symbolic interpolation [BCI+23, Zhe26], we keep the challenge as a variable Z and interpolate $Q(X, Z, Y_0, Y_1)$ over $F(Z)$ on the same support. The surplus of coefficients over conditions bounds its degree in Z:

$$\deg_Z Q \le \left\lfloor \frac{nR_{\mathrm{loc}}B_{\mathrm{jet}}}{N_{\mathrm{coef}} - nR_{\mathrm{loc}}} \right\rfloor, \qquad Q(X, z, Y_0, Y_1) \ne 0 \quad \text{for every } z \in \overline{F}$$

(Lemma 3.11 and Proposition 3.13). Every candidate close to $f + zg$ satisfies $Q(X, z, P, P') = 0$. We write B_Z for a bound on $\deg_Z Q$; at fixed rate and gap it is independent of n (Lemma 3.16).

2.2 From an equation to lists and MCA

An explainer is not yet a list bound, since its solutions can form a family of positive dimension, and MCA also asks for common witnesses outside a small set of challenges. Section 4 proves two tools that give both for any algebraic family of candidates, from its dimension and degree.

Counting by agreement. Zheng's incidence argument uses the received positions to count points of a solution family [Zhe26]. In coefficient space, agreement at position i is the hyperplane $P(\alpha_i) = y_i$. Any k of these hyperplanes meet in one point, so an irreducible curve of degree Δ lies in at most D of them and meets each of the others in at most Δ points (Figure 3). A close message lies on at least A hyperplanes, so counting incidences bounds the number of close messages on the curve by $\Delta(n - D)/(A - D) = \Delta\lambda_{\mathrm{agr}}$, and a family of dimension r incurs one such factor per dimension (Lemma 4.1).

Witness graphs. Fix $f, g \in F^n$, write $S_z(P)$ for the agreement set of P with $f + zg$, and $S(F_0, G_0)$ for the positions where F_0 agrees with f and G_0 with g. MCA asks for one set $\mathcal{E}$, fixed before the candidate, such that for $z \notin \mathcal{E}$ every message P with $|S_z(P)| \ge A$ has witnesses with

$$P = F_0 + zG_0, \qquad S_z(P) = S(F_0, G_0)$$

(Definition 3.1). Any k agreements of P determine a pair (F_0, G_0) with $P = F_0 + zG_0$: interpolate f and g on those positions. As z varies, the pairs $(z, F_0 + zG_0)$ trace a *witness graph* in the joint space of pairs (z, P). At a position i outside $S(F_0, G_0)$, the polynomial on this graph agrees with $f + zg$ exactly when $(F_0(\alpha_i) - f_i) + z(G_0(\alpha_i) - g_i) = 0$, and the left side is a nonzero polynomial of degree at most one in z. Each graph therefore has at most n challenges with an accidental agreement, but there can be $\binom{n}{k}$ graphs.

The common-witness transfer. To avoid summing over all graphs, Zheng's persistent-line argument organizes the witness graphs inside one algebraic family, and Jeronimo independently uses the same type of symbolic solution family [Zhe26, Jer26]. We keep this organization and sharpen the degrees and agreement thresholds that control the exceptional set. Suppose the close pairs (z, P) lie in a family of dimension at most two and joint degree J_1, and that its set of messages at a generic challenge, one treated as an indeterminate, has degree F_{reg}. We fix an integer L with $D + 1 \le L \le A$ and retain the witness graphs contained in this family that have at least L common agreements. A close pair violates MCA only if it is off every retained graph, or lies on one and has an agreement outside its common set. The incidence argument bounds the number of challenges of the two kinds by

$$\frac{n - L + 1}{A - L + 1}\,\lambda_{\mathrm{agr}}\,J_1 \qquad \text{and} \qquad (n - L)\,\frac{n - D}{L - D}\,F_{\mathrm{reg}}$$

(Lemma 4.2). The first bound has one incidence ratio for each of two successive cuts of the surface by agreement hyperplanes. In the second, $\frac{n-D}{L-D}F_{\mathrm{reg}}$ bounds the number of retained graphs: they give distinct points of the family at the generic challenge, each with at least L agreements, so the fixed-word count applies with threshold L. Each retained graph has at most $n - L$ positions outside its common set. Increasing L decreases the second bound and increases the first.

Ordinary equations and the Johnson theorem. For an *ordinary* equation $Q(X, z, P) = 0$, which involves no derivative of P and arises in the Johnson theorem and in the first-order count, we follow the function-field approach of BCIKS [BCI+23]: we track its solutions over the challenge, and a factorwise Frobenius pullback makes each inseparable factor separable, so the transfer holds in every characteristic (Lemma 4.3). Its first use is the Johnson theorem, which starts from the symbolic Guruswami–Sudan interpolant and factor aggregation of BCHKS [BCH+26]. At fixed rate and Johnson gap η_0, their exceptional bound has scale n/η_0^5 and ours has scale n/η_0^3 (Theorem 4.4). Following savings found by BitWonka and jieyilong [Bit26a, jie26a], we save one factor $1/\eta_0$ by reconstructing only the k coefficients of the message, although the interpolation bound allows a longer expansion, and another by bounding the degree in the challenge separately from the degree in the message.

2.3 First-order list and MCA bounds

Section 5 applies these tools to the first-order explainer. The main step is to parameterize the solutions of one differential equation by a family whose dimension and degrees are controlled.

Solutions as Taylor data. The explainer alone does not bound the list: for $Q = Y_1$, every constant polynomial is a solution. Kopparty's recurrence describes the solutions through their initial Taylor coefficients [Kop15]. If $P(x_0 + T) = c_0 + c_1 T + c_2 T^2 + \cdots$ at a center x_0, then $Q(x_0, c_0, c_1) = 0$, and the coefficient of T in $Q(X, P, P') = 0$ gives

$$2\,Q_{Y_1}(x_0, c_0, c_1)\,c_2 = -Q_X(x_0, c_0, c_1) - c_1\,Q_{Y_0}(x_0, c_0, c_1),$$

with subscripts for partial derivatives. The factor Q_{Y_1} is the *separant*. When it is nonzero at (x_0, c_0, c_1), this equation and its successors express $c_2, \ldots, c_D$ as rational functions of (c_0, c_1); step j divides by j times the separant, which requires characteristic zero or $p > D$. Following Zheng's symbolic solution families [Zhe26], we keep (c_0, c_1) as variables on the curve $Q(x_0, c_0, c_1) = 0$. The solutions with nonzero separant at x_0, the *regular* solutions, then lie in the image of one rational map from this curve to coefficient space.

Regular and singular solutions. The regular–singular separation of Naskręcki's construction covers the remaining solutions [Nas26]. Write the squarefree part of Q as a product $Q_{\mathrm{ord}}Q_{\mathrm{diff}}$ of polynomials, where Q_{ord} collects the factors independent of Y_1; for $p > B_\partial$, Q_{diff} is separable in Y_1. We take x_0 to be a new indeterminate, so that every solution of $Q_{\mathrm{diff}} = 0$ whose separant $\partial_{Y_1} Q_{\mathrm{diff}}(X, P, P')$ is a nonzero polynomial is regular at x_0. Every other solution satisfies the ordinary equation

$$Q_{\mathrm{ord}} \cdot \mathrm{Res}_{Y_1}\left(Q_{\mathrm{diff}}, \partial_{Y_1} Q_{\mathrm{diff}}\right) = 0,$$

to which the ordinary transfer applies in every characteristic. We count the degree of the regular family differently from prior work. Each factor of Q is counted once, whereas differentiating Q repeatedly counts a Taylor family at every stage (Lemma C.1), and the image-degree bound keeps the Y_1-degree B_∂ separate from B_{jet} (Lemma 5.3). The regular family then has fiber degree $F_{\mathrm{reg}} \le 4DB_{\mathrm{jet}}B_\partial$, and its joint degree J_1 is bounded by these degrees and B_Z (Lemma 5.5).

The two bounds. For one received word, the incidence count bounds the regular solutions, and the Y_0-degree of the ordinary equation, at most $2B_{\mathrm{jet}}B_\partial + B_{\mathrm{jet}}$, bounds the others. For MCA, we choose L in the transfer as $L = D + \lceil (A-D)/2 \rceil$, which makes both incidence ratios at most $2\lambda_{\mathrm{agr}}$. Since $F_{\mathrm{reg}} = O(n)$ and $J_1 = O(n^2)$ at fixed interpolation parameters, each count is $O(n^2)$; the ordinary equation is counted at a separate threshold. In characteristic zero or $p > \max\{k-1, B_\partial\}$, at fixed ρ and gap $\eta_1 = a - a_1(\rho) > 0$, where $B_{\mathrm{jet}}, B_\partial = O_\rho(1/\eta_1)$, Theorems 5.1 and 5.2 give

$$|\mathcal{L}(y, A)| \le 4D\lambda_{\mathrm{agr}}B_{\mathrm{jet}}B_\partial + 2B_{\mathrm{jet}}B_\partial + B_{\mathrm{jet}} = O_\rho(n/\eta_1^2), \qquad |\mathcal{E}| = O_\rho(n^2/\eta_1^4).$$

First order thus keeps the list size linear and the exceptional count quadratic in n, small enough for the proof systems of Section 7.

2.4 Higher derivatives approach capacity

To approach capacity, BCPZZ use higher hidden derivatives [BCP+26a]. We keep their local conditions but choose a support whose count is explicit in the rate, the agreement, and d (Section 8).

The weighted support. With Hasse derivatives $P^{[l]}$ (Definition 3.4), so that $P^{[1]} = P'$, an agreement $P(\alpha_i) = y_i$ gives

$$P(\alpha_i + T) \equiv y_i + \sum_{l=1}^{d} (-1)^{l+1} T^l P^{[l]}(\alpha_i + T) \pmod{T^{d+1}}$$

(Lemma 3.5). The local conditions substitute $Y_0 = y_i + \sum_{l=1}^{d} (-1)^{l+1} T^l Y_l + T^{d+1} V$ into $Q(X, Y_0, \ldots, Y_d)$ and reduce modulo T^m (Definition 3.6). The powers of $T^{d+1}V$ vanish modulo T^m from exponent $m/(d+1)$ on, which reduces the number of conditions. For an ambient dimension $K \ge \max\{k, d+1\}$ and a weight bound W, we take the monomials

$$X^x \prod_{l=0}^{d} Y_l^{j_l}, \qquad \sum_{l=1}^{d} l\, j_l \le W, \qquad x + (K-1)\sum_{l=0}^{d} j_l < mA$$

(Equation (56)). The first inequality weights each derivative variable by its order. The second gives degree less than mA after substituting any polynomial of degree less than K, in particular any message.

The order as a function of the gap. Take $K-1=\lfloor \rho n \rfloor$ and W proportional to m. We sum the remaining choices of x and j_0 over the weighted tuples $(j_1,\dots,j_d)$ and estimate the sum by a simplex integral. For large m, the ratio $N_{\mathrm{coef}}/(nR_{\mathrm{loc}})$ is then, up to a factor tending to one, at least an explicit function of ρ, a, and d that has the rate ρ as a prefactor and, for fixed $\rho < a$, grows like $d^{1-\rho/a}$ (Equation (61) and Lemma 8.3). So for large d the ratio exceeds one, and some finite m gives an order-d explainer. Because the rate stays in the bound, we can solve for a sufficient order, but the prefactor ρ makes that order grow as the rate tends to zero.

To obtain one order for every rate, we enlarge the interpolation message space, as Jeronimo does [Jer26], but only at low rate: given $A-k \geq \delta n$, we take $K-1=\lfloor 2\delta^2 n \rfloor$ when $k/n < \delta^2$ and $K-1=k$ otherwise. One order then suffices at every rate, $d(\delta)=\lceil e^{1.5/\delta} \rceil$ for $0<\delta<6/25$ and $n \geq N(\delta)$ (Theorem 8.1).

Lists and exceptional sets. Counting at order d follows Sections 2.2 and 2.3, with explicit degrees: Kopparty's recurrence reconstructs $\max\{k, d+1\}$ Taylor coefficients from the initial jet. A solution whose separant vanishes identically also solves the separant, a lower-degree equation, and passes to it; this separant chain has at most B_{jet} stages. At fixed parameters the regular solutions for one word lie on images of dimension at most d and degree $O(n^d)$ (Theorem 8.4), and the joint family over the challenge has degree $O(n^{d+1})$ (Theorem 8.5). For a support with $N_{\mathrm{coef}} > \gamma n R_{\mathrm{loc}}$, $\gamma > 1$, and gap $A-k \geq \delta n$, we get $|\mathcal{L}(y,A)| \leq C_L n^d$ and $|\mathcal{E}| \leq C_E n^{d+1}$, with explicit constants depending only on d, B_{jet}, γ and δ, for $p > \max\{k-1, d, B_{\mathrm{jet}}\}$ or characteristic zero (Proposition 8.7). With $d=d(\delta)$ only $p > k-1$ is needed: if $p \leq B_{\mathrm{jet}}$ then $k \leq B_{\mathrm{jet}}$, and for $n \geq N(\delta)$ the agreement is then above Johnson, where Theorem 4.4 applies (Theorem 8.1).

2.5 Decoding without enumerating initial values

The bounds above count the close messages; Section 9 computes them, deterministically and in time polynomial in n and $\log q$, over fields $\mathbb{F}_q$ with $q \geq n$ and characteristic $p > k-1$. The difficulty is extracting the messages from the equation. For Guruswami–Sudan the messages are the roots $Y=P(X)$ of $Q(X,Y)$, a finite set.

A differential equation can instead have a family of solutions of positive dimension (Section 2.3). Kopparty's decoder for multiplicity codes therefore enumerates the initial Taylor coefficients at a center x_0, such as $(P(x_0), P'(x_0)) \in \mathbb{F}_q^2$ at first order, and extends each choice by his recurrence [Kop15], and the capacity decoders of BCPZZ and Jeronimo keep this enumeration [BCP+26a, Jer26]. It costs a power of q: polynomial for a univariate multiplicity code, whose length is q, but not for a Reed–Solomon code with $n \ll q$. A small list does not help, since the enumeration runs over all solutions.

Our decoders avoid this enumeration by imposing the agreements $P(\alpha_i)=y_i$ before extracting any candidate. We describe the first-order case; higher orders follow the same route. In this sketch the explainer has jet degree at most a fixed B_{jet} and $\deg_X Q \leq Cn$, hidden constants may depend on B_{jet} and C, and $p > \max\{k-1, B_{\mathrm{jet}}\}$. If only $p > k-1$ holds, then $k \leq p \leq B_{\mathrm{jet}}$, and we decode such codes separately.

A curve of candidates. Let P be a close message. The separant chain of Q (Section 9.5) assigns P to an equation T that P solves with nonzero separant. If T has order zero, its solutions already form a finite set; suppose it has order one, so that $T_{Y_1}(X,P,P')$ is a nonzero polynomial of degree $O(n)$, nonzero at one of $O(n)$ prescribed centers x_0. There Kopparty's recurrence (Section 2.3) expresses every coefficient of P through $U=(P(x_0),P'(x_0))$. Keeping U as an unknown, as in the symbolic families of Zheng and Jeronimo [Zhe26, Jer26], gives one candidate $\widetilde{P}_U(X)=U_0+U_1(X-x_0)+\sum_{2\leq j<k}\tau_j(U)(X-x_0)^j$ for each point U of the plane curve $T(x_0,U_0,U_1)=0$ at which the separant is nonzero, where the τ_j are rational functions of degree $O(k)$. Every close message assigned to T and regular at x_0 is one of these candidates, but they form a curve, not a finite list.

Agreements cut the curve. Agreement at position i is a polynomial equation $g_i(U) = 0$ of degree $O(k)$, obtained from $\widetilde{P}_U(\alpha_i) = y_i$ by clearing denominators, and a close message satisfies at least A of them. As in Section 2.2, at most $k - 1$ of these vanish identically on the component of the curve through its point, so at least $t = A - k + 1$ of its agreement equations are *proper* on that component, and each proper equation meets it in $O(k)$ points. Every close message regular at this stage and center is therefore among the $O(nk)$ points where some proper g_i cuts the curve. Jeronimo uses this isolation to bound the list size [Jer26, Lemma 6.2]; we compute with it. At order $r < k$ the family has dimension r, and some r of the message's agreement equations, with the equation of the family, isolate it as a simple root of a square system (Lemma 9.10). Rojas's algorithm constructs a finite representation containing all isolated roots of such a system [Roj99].

Reading the messages off the received values. The cut points can have coordinates in extensions of large degree, and computing them would mean finding roots over $\mathbb{F}_q$. Deterministic root finding in time polynomial in both the degree and $\log q$ is a longstanding open problem (see Section 1.2 for a recent claim over prime fields). We never compute them. Instead we treat them together as the points of a finite algebra over $\mathbb{F}_q$ or a small extension of it.

In the simplest case the algebra is $\mathbb{F}_q[u]/G$, or its analogue over the extension, for a squarefree G; its points are the roots of G, and the candidates form one polynomial $\widetilde{P}(X)$ with coefficients in this algebra. Position i splits G into $G_0 = \gcd(G, \widetilde{P}(\alpha_i) - y_i)$, the roots whose candidate agrees with y_i, and G/G_0, the rest. Once a factor has agreed at k positions, every root of it gives the same polynomial, the interpolant of those k received values, which we compute over $\mathbb{F}_q$ and check against y. This extends the received-value idea of Augot–Pecquet and CHK [AP00, CHK26] from finite root sets to these algebraic families; algebras with repeated roots need one adjustment (Lemma 9.2). These steps already give a deterministic decoder, polynomial in n and $\log q$ at every fixed order and explainer degree.

Refinements. The core construction runs at $O(n)$ centers and uses one cut per close message, but a close message is cut by t proper equations. Keeping only the values of u above which at least t of them vanish shrinks the algebras from dimension $O(nk)$ to $O(nk/t)$, which is $O(k)$ above the first-order curve (Lemma 9.8). With fast coefficient formulas (Lemma 9.5), each center costs $\widetilde{O}(n^2)$ operations, giving $\widetilde{O}_{\rho,a}(n^3 \log q)$ bit operations for $p > k - 1$ (Corollary 9.9). At order zero, fast interpolation [JNSV17] and a relaxed power-series lift, based on van der Hoeven's relaxed multiplication [vdH02], give $\widetilde{O}_{\rho,a}(n \log q)$ above Johnson (Theorem 9.3), and at order r we try all r-subsets of positions (Theorem 9.11).

3 Hidden-derivative interpolation and the first-order curve

We construct the first-order equation on which the list and MCA bounds of Section 5 rest; Section 2.1 sketches the construction. For a received line $f + zg$ and an agreement fraction $a > a_1(\rho)$, Corollary 3.17 gives a nonzero polynomial $Q(X, Z, Y_0, Y_1)$ such that $Q(X, z, P, P^{[1]}) = 0$ for every challenge z and every message P agreeing with $f + zg$ in at least $A \geq an$ positions. Because Section 8 reuses the local constraints and the symbolic interpolation (Sections 3.2 and 3.3), we state them at general order before specializing to order one (Sections 3.4 and 3.5); the construction works in every characteristic.

Throughout, F denotes a field and $\mathbb{F}_q$ the field with q elements; $\mathbb{N} = \{0, 1, 2, \ldots\}$ and $[n] = \{1, \ldots, n\}$. We write $F[X]_{\leq D}$ for the polynomials of degree at most D, including 0, and take $\deg 0 = -\infty$. We write $x_+ = \max\{x, 0\}$, $\mathbf{1}_{\{\mathcal{P}\}}$ for the indicator of a condition, and Pr, $\mathbb{E}[\cdot]$, and vol for probability, expectation, and Euclidean volume. Logarithms are natural unless a base is shown, and a subscript in $O_\delta(\cdot)$ means that the implied constant may depend on that parameter.

3.1 Reed–Solomon codes, lists, and line MCA

Fix pairwise distinct evaluation points $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n) \in F^n$. The Reed–Solomon code of message dimension k is

$$\mathrm{RS}_{\boldsymbol{\alpha}}[n,k] := \big\{(P(\alpha_1), \dots, P(\alpha_n)) : P \in F[X],\ \deg P < k\big\}. \tag{2}$$

Its rate is k/n; a theorem stated at a rate parameter ρ assumes only $k \le \rho n$. We assume no multiplicative or additive structure of the evaluation set. For $u, v \in F^n$, put $\mathrm{Agr}(u,v) := \{i \in [n] : u_i = v_i\}$. For a received word $y \in F^n$, an agreement threshold A, and an ambient degree bound D, the agreement list is

$$\mathcal{L}_D(y,A) := \big\{P \in F[X]_{\le D} : |\mathrm{Agr}((P(\alpha_i))_{i=1}^n, y)| \ge A\big\}, \qquad \mathcal{L}(y,A) := \mathcal{L}_{k-1}(y,A). \tag{3}$$

The ambient bound D may exceed $k-1$; we keep it distinct from the message dimension k.

For MCA we receive a line $f + zg$ with $f, g \in F^n$, or more generally a *polynomial received curve* of degree $\ell \ge 1$, whose i-th coordinate is

$$w_i(z) = \sum_{t=0}^{\ell} z^t w_{t,i} \qquad (w_0, \dots, w_\ell \in F^n). \tag{4}$$

A line is the case $\ell = 1$ with $(w_0, w_1) = (f, g)$. MCA asks that, outside a few challenges, every candidate close to $w(z)$ come from witnesses that agree with $w_0, \dots, w_\ell$ on the same positions. For a polynomial P and witnesses $P_0, \dots, P_\ell$, write

$$S_z(P) = \{i : P(\alpha_i) = w_i(z)\}, \qquad S(P_0, \dots, P_\ell) = \{i : P_t(\alpha_i) = w_{t,i} \text{ for every } 0 \le t \le \ell\},$$

and write $S(F_0, G_0) = S(P_0, P_1)$ for a line.

Definition 3.1 (Full common-witness agreement)**.** A pair (z, P) has the *full common-witness property* for the received curve if there are degree-$< k$ polynomials $P_0, \dots, P_\ell$ such that

$$P = \sum_{t=0}^{\ell} z^t P_t, \qquad S_z(P) = S(P_0, \dots, P_\ell). \tag{5}$$

For a collection C of pairs (z, P), the property holds outside $\mathcal{E} \subseteq F$ if it holds for every pair in C with $z \notin \mathcal{E}$.

Definition 3.2 (Line MCA with full agreement sets)**.** Let $1 \le A \le n$ and $\mathsf{E}_{\mathrm{exc}} \ge 0$. The code $\mathrm{RS}_{\boldsymbol{\alpha}}[n,k]$ has *line MCA at agreement threshold A with exceptional bound* $\mathsf{E}_{\mathrm{exc}}$ if every received line $f + zg$ admits one set $\mathcal{E} \subseteq F$ with $|\mathcal{E}| \le \mathsf{E}_{\mathrm{exc}}$ such that the full common-witness property holds outside $\mathcal{E}$ for the collection of all pairs (z, P) with $\deg P < k$ and $|S_z(P)| \ge A$.

The set $\mathcal{E}$ depends on the line but is chosen before the candidate P. The equality of agreement sets is stronger than a bound on their sizes: it rules out agreements created only by cancellation at the chosen challenge. Over $\mathbb{F}_q$, $\mathsf{E}_{\mathrm{exc}}/q$ bounds the probability that a uniform challenge is exceptional. MCA for polynomial received curves of degree ℓ is defined in the same way.

Our main results assume $k \ge 2$; constant codes need neither interpolation nor geometry, since a collision argument applies to every polynomial received curve at once.

Lemma 3.3 (Constant-code polynomial-curve MCA)**.** *Let F be any field, let $1 \le A \le n$, and fix a received curve of degree $\ell \ge 1$. For every challenge, the number of constants $c \in F$ with $|S_z(c)| \ge A$ is at most $\lfloor n/A \rfloor$. Moreover, there is an exceptional set $\mathcal{E} \subseteq F$ such that*

$$|\mathcal{E}| \le \begin{cases} \ell\binom{n}{2}, & A = 1, \\ \left\lfloor \ell\binom{n}{2}/(A-1) \right\rfloor, & A \ge 2, \end{cases} \tag{6}$$

and every pair (z, c) with $z \notin \mathcal{E}$ and $|S_z(c)| \ge A$ has the full common-witness property of Definition 3.1 for $k = 1$.

Proof. Put $v_i = (w_{0,i}, \dots, w_{\ell,i})$. For a fixed challenge, the agreement sets of distinct constants are disjoint, which proves the list bound. Call a pair $\{i, j\}$ a collision at z if $v_i \neq v_j$ but

$$\sum_{t=0}^{\ell} z^t w_{t,i} = \sum_{t=0}^{\ell} z^t w_{t,j}.$$

Each such pair collides at at most ℓ challenges. If the desired conclusion fails for a constant c, its agreement set contains two distinct tuples and hence a collision. When $A \geq 2$, a partition of at least A agreeing positions into at least two tuple classes has at least $A - 1$ cross-class pairs, and these pairs are disjoint for distinct constants at the same challenge. Since there are at most $\ell\binom{n}{2}$ pairs of a collision and its challenge, double counting gives (6); for $A = 1$, count just one collision pair.

Outside the resulting exceptional set, all tuples on $S_z(c)$ are equal to some $(c_0, \dots, c_\ell)$. No position with a different tuple can evaluate to c, since that would be a collision. Thus $c = \sum_t z^t c_t$ and $S_z(c) = S(c_0, \dots, c_\ell)$. □

3.2 Local constraints and the interpolation principle

Guruswami–Sudan interpolation asks the interpolant to vanish to a high order m at every received pair, so that each agreement becomes a root of multiplicity m [GS99]. The hidden-derivative constructions of Naskręcki and BCPZZ [Nas26, BCP+26a] keep this mechanism and add variables for derivatives of the message, which we take in the standard Hasse convention [DKSS13].

Definition 3.4 (Hasse derivative)**.** For $P \in F[X]$, its j-th Hasse derivative $P^{[j]}$ is defined by the polynomial identity

$$P(X + T) = \sum_{j \geq 0} P^{[j]}(X) T^j. \tag{7}$$

Equivalently, $P^{[j]} = \sum_{i \geq j} \binom{i}{j} p_i X^{i-j}$ when $P = \sum_i p_i X^i$.

The definition is valid in every characteristic and gives $\deg P^{[j]} \leq D - j$ whenever $\deg P \leq D$, with a zero derivative when $j > D$.

An ordinary Reed–Solomon symbol is one value $P(\alpha_i)$; a multiplicity-code symbol also transmits derivatives at α_i [GW13, Kop15]. We receive only values, so the derivatives that enter an interpolation equation $Q(X, Y_0, \dots, Y_d)$ are hidden. The largest index d is the derivative order of Q, while the exponent of a variable Y_j is an algebraic degree: Y_1^2 uses the first derivative quadratically, not the second derivative.

Ordinary Taylor expansion uses derivatives at X to obtain the value at $X + T$. The local constraints use the inverse direction: derivatives at $X + T$ recover the value at X.

Lemma 3.5 (Backward Taylor identity)**.** *Every $P \in F[X]$ satisfies*

$$P(X) = \sum_{j \geq 0} (-T)^j P^{[j]}(X + T) \qquad \text{in } F[X, T]. \tag{8}$$

Consequently, for every $d \geq 0$,

$$P(X + T) \equiv P(X) + \sum_{j=1}^{d} (-1)^{j+1} T^j P^{[j]}(X + T) \pmod{T^{d+1}}. \tag{9}$$

Proof. Apply (7) with X replaced by $X + T$ and the increment replaced by $-T$. This gives (8). Isolating its $j = 0$ term and discarding the terms with $j > d$ gives (9). □

BCPZZ's backward-Taylor substitution turns agreement into multiplicity without derivative evaluations [BCP+26a, Section 3.1]. Fix an order $d \geq 0$ and a target multiplicity $m \geq 1$. For formal variables $Y = (Y_1, \ldots, Y_d)$, write $B_d(T, Y) = \sum_{j=1}^{d} (-1)^{j+1} T^j Y_j$; an empty sum is zero. By (9), at an agreement $P(\alpha) = y$ the difference between $P(\alpha + T) - y$ and $B_d(T, Y)$ at $Y_j = P^{[j]}(\alpha + T)$ is divisible by T^{d+1}. The constraints encode this divisibility with a new variable V.

Definition 3.6 (Local hidden-derivative constraints)**.** For $(\alpha, y) \in F^2$ and $Q \in F[X, Y_0, \ldots, Y_d]$, put

$$\Phi^{d,m}_{\alpha,y}(Q) := Q(\alpha + T, y + B_d(T, Y) + T^{d+1} V, Y_1, \ldots, Y_d) \pmod{T^m} \tag{10}$$

in $F[T, V, Y_1, \ldots, Y_d]/(T^m)$. The multiplicity-$m$, order-$d$ constraints at (α, y) are

$$\Phi^{d,m}_{\alpha,y}(Q) = 0. \tag{11}$$

These are homogeneous linear conditions on any fixed interpolation space: every coefficient in $V, Y_1, \ldots, Y_d$ must be divisible by T^m. For $d = 0$ they are the Guruswami–Sudan conditions: every coefficient of $U^a V^b$ with $a + b < m$ in $Q(\alpha + U, y + V)$ vanishes, which is bivariate multiplicity at least m at (α, y).

Lemma 3.7 (Local multiplicity [BCP+26a, Lemma 3.1])**.** *If $P(\alpha) = y$ and Q satisfies (11), then*

$$Q(\alpha + T, P(\alpha + T), P^{[1]}(\alpha + T), \ldots, P^{[d]}(\alpha + T)) \equiv 0 \pmod{T^m}. \tag{12}$$

Thus each agreement gives a root of multiplicity at least m after substituting the candidate and its derivatives into Q.

Proof. The backward Taylor identity makes

$$V_P(T) = \frac{P(\alpha + T) - y - \sum_{j=1}^{d} (-1)^{j+1} T^j P^{[j]}(\alpha + T)}{T^{d+1}} \tag{13}$$

a polynomial. Substitute $Y_j = P^{[j]}(\alpha + T)$ and $V = V_P(T)$ into (11). The second input to Q becomes $P(\alpha + T)$, proving the claim in every characteristic. □

For rank calculations we use the remainder variable $\widetilde{V} = T^d V$, so that $Y_0 = y + B_d(T, Y) + T\widetilde{V}$. Since $T^i \widetilde{V}^u Y^{\mathbf{j}} = T^{i+du} V^u Y^{\mathbf{j}}$, the coefficient of $T^i \widetilde{V}^u Y^{\mathbf{j}}$ in a row with $i + du < m$ is the coefficient of $T^{i+du} V^u Y^{\mathbf{j}}$ in (10). Rows with $i < u$ vanish, since each $\widetilde{V}$ comes with a factor of T, so both coordinate systems give the same matrix entries, kernel, rank, jet-degree grouping, and challenge degrees.

The interpolation principle. Once each agreement is a root of multiplicity m, enough agreements force a polynomial identity. Every construction in this paper produces a polynomial with the following property.

Definition 3.8 (Hidden-derivative explainer)**.** Fix a field F, distinct evaluation points $\alpha_1, \ldots, \alpha_n \in F$, a received word $y \in F^n$, integers $K \geq 1$ and $1 \leq A \leq n$, and an order bound $d \geq 0$. A nonzero polynomial $Q \in F[X, Y_0, \ldots, Y_d]$ is an *explainer* for y at agreement threshold A and degree bound K if every $P \in F[X]$ with $\deg P < K$ satisfies

$$\big|\{i : P(\alpha_i) = y_i\}\big| \geq A \quad \Longrightarrow \quad Q\big(X, P, P^{[1]}, \ldots, P^{[d]}\big) \equiv 0.$$

The implication is one-way: an explainer can have solutions far from y, so the counting arguments impose agreement again. The ambient bound K may exceed the message dimension k; an explainer at degree bound $K \geq k$ also explains the degree-$< k$ candidates.

An explainer exists as soon as an interpolation space has more coefficients than its local constraints have total rank. This is the Guruswami–Sudan dimension argument [GS99], in the hidden-derivative form of BCPZZ [BCP+26a]. We state it for an arbitrary space and local maps, so that each later construction only has to bound its dimension, local ranks, and specialization degree.

Proposition 3.9 (Global interpolation principle)**.** *Let $\mathcal{Q} \subseteq F[X, Y_0, \ldots, Y_d]$ be finite dimensional, and let Φ_i be linear maps on $\mathcal{Q}$, one for each received pair (α_i, y_i). Suppose that $Q \in \ker \Phi_i$ and $P(\alpha_i) = y_i$ imply*

$$(X - \alpha_i)^m \mid Q(X, P, P^{[1]}, \ldots, P^{[d]}).$$

Assume every $Q \in \mathcal{Q}$ specializes at every degree-$\leq D$ polynomial to degree less than mA. If

$$\dim \mathcal{Q} > \sum_{i=1}^{n} \operatorname{rank} \Phi_i, \tag{14}$$

then $\mathcal{Q}$ contains an explainer for y at degree bound $D + 1$ and agreement threshold A.

Proof. The stacked map has rank at most the sum of the local ranks, so its kernel contains a nonzero Q. At A distinct agreement positions, the specialization is divisible by $\prod_i (X - \alpha_i)^m$, of degree mA. Its degree is smaller than mA, so it vanishes identically. □

3.3 One equation for every challenge

For MCA, one equation must relate the candidates at different challenges; separate explainers for the words $f + zg$ give no such relation. Following the function-field approach of BCIKS [BCI+23] and Zheng's symbolic interpolation [Zhe26], we keep the challenge as a variable Z, solve the local constraints over $F[Z]$, and bound the degree of the solution in Z, which controls the exceptional counts of Section 4.

Definition 3.10 (Symbolic explainer for a received line)**.** Fix $f, g \in F^n$, an agreement threshold A, a degree bound $K \geq 1$, and an order bound $d \geq 0$. A nonzero polynomial $Q \in F[X, Z, Y_0, \ldots, Y_d]$ is a *symbolic explainer* for $f + Zg$ at these bounds if, for every $z \in F$ and every $P \in F[X]$ with $\deg P < K$,

$$|S_z(P)| \geq A \quad \Longrightarrow \quad Q\big(X, z, P, P^{[1]}, \ldots, P^{[d]}\big) \equiv 0.$$

For a polynomial received curve $w(Z) = \sum_{t=0}^{\ell} Z^t w_t$, a symbolic explainer for $w(Z)$ is defined in the same way, with $S_z(P)$ the agreement set with $w(z)$ from Section 3.1.

The definition requires Q to be nonzero only as a polynomial in all its variables. Our construction also makes $Q(X, z, Y)$ nonzero for every z over every field extension, so for each $z \in F$ the specialization $Q(X, z, Y)$ is itself an explainer for $f + zg$ in the sense of Definition 3.8. The local constraints form a linear system over $F[Z]$, and a kernel vector with low-degree entries gives Q of low degree in Z. The next lemma finds one by counting scalar coefficients; its separate costs for columns and rows give the sharper challenge-degree tests of Proposition C.2.

Lemma 3.11 (Graded degree bound for a polynomial kernel vector)**.** *Let F be a field, let I, J be finite sets, and let $\mathbf{A}(Z) \in F[Z]^{I \times J}$. Fix integer column costs c_j, row costs r_i, and a height H. Suppose*

$$\deg \mathbf{A}_{ij} \leq c_j - r_i, \qquad i \in I,\ j \in J,$$

where a negative upper bound means that the entry is zero. If

$$\sum_{j \in J} (H - c_j + 1)_+ > \sum_{i \in I} (H - r_i + 1)_+, \tag{15}$$

then there is a nonzero vector $v \in F[Z]^J$ satisfying $\mathbf{A}(Z)v = 0$, with

$$\deg v_j \leq H - c_j \quad \text{when } c_j \leq H, \qquad v_j = 0 \quad \text{when } c_j > H.$$

Its entries can be chosen to have greatest common divisor one in $F[Z]$. For that choice, $v(z) \neq 0$ for every z in every field extension of F.

Proof. For each column with $c_j \le H$, give v_j degree at most $H - c_j$; set the other entries to zero. The left side of (15) counts the scalar coefficients of these unknown polynomials. In row i, each product has degree at most $(c_j - r_i) + (H - c_j) = H - r_i$, giving at most $(H - r_i + 1)_+$ scalar equations. The strict surplus in (15) gives a nonzero kernel vector.

Divide its entries by their greatest common divisor. Since $F[Z]$ is an integral domain, the polynomial identities remain valid, and division does not increase any degree. Bézout's identity for the primitive entries shows that they have no common zero over any field extension. □

Corollary 3.12 (Equal costs)**.** *Let $\mathbf{A}(Z)$ have N_{coef} columns, rank $R < N_{\mathrm{coef}}$ over $F(Z)$, and entries of degree at most b. Then $\mathbf{A}(Z)v = 0$ for some $v \in F[Z]^{N_{\mathrm{coef}}}$ whose entries have greatest common divisor one and degree at most*

$$\left\lfloor \frac{Rb}{N_{\mathrm{coef}} - R} \right\rfloor; \tag{16}$$

in particular $v(z) \neq 0$ for every z in every extension of F.

Proof. Let β denote (16). Choose R rows spanning the row space over $F(Z)$, and apply Lemma 3.11 to them with $c_j = b$, $r_i = 0$, and $H = \beta + b$. The test (15) becomes $(N_{\mathrm{coef}} - R)(\beta + 1) > Rb$, which holds by the choice of β. The kernel vector annihilates the selected rows, hence every row. □

Proposition 3.13 (Constructing a symbolic explainer)**.** *Fix integers $d \ge 0$, $m \ge 1$, $K \ge 1$, and $1 \le A \le n$. Let $Q \subseteq F[X, Y_0, \dots, Y_d]$ be a finite monomial space invariant under translations in X and Y_0. Suppose its members have total degree at most B_{jet} in $Y_0, \dots, Y_d$, specialize at every degree-$< K$ polynomial to degree less than mA, and that $N_{\mathrm{coef}} = \dim Q$ exceeds nR_{loc}, where R_{loc} is an upper bound for the rank of the local constraint map $\Phi^{d,m}_{0,0}$ on Q. For any distinct $\alpha_i \in F$ and $f, g \in F^n$, there is $Q \in F[Z] \otimes_F Q$ such that*

$$\deg_Z Q \le \left\lfloor \frac{nR_{\mathrm{loc}}B_{\mathrm{jet}}}{N_{\mathrm{coef}} - nR_{\mathrm{loc}}} \right\rfloor, \qquad Q(X, z, Y) \neq 0 \quad \text{for every } z \in \overline{F},$$

and every degree-$< K$ polynomial P agreeing with $f + zg$ in at least A positions satisfies

$$Q(X, z, P, P^{[1]}, \dots, P^{[d]}) = 0. \tag{17}$$

The same holds for a received curve $w(Z) = \sum_{t=0}^{\ell} Z^t w_t$, with $f + zg$ replaced by $w(z)$ and

$$\deg_Z Q \le \left\lfloor \frac{\ell nR_{\mathrm{loc}}B_{\mathrm{jet}}}{N_{\mathrm{coef}} - nR_{\mathrm{loc}}} \right\rfloor. \tag{18}$$

These assertions hold after any extension of the base field.

Proof. Use the local constraints at $(\alpha_i, f_i + Zg_i)$. In their matrices, the challenge enters only through powers of $f_i + Zg_i$, whose exponents are bounded by the Y_0-degree in Q, hence by B_{jet}. Translation gives an automorphism of $Q \otimes_F F(Z)$, and the local map factors as translation followed by the map at $(0, 0)$. Its rank therefore equals the rank at $(0, 0)$, which is at most R_{loc}, also after scalar extension.

The stacked matrix has N_{coef} columns, rank $R \le nR_{\mathrm{loc}} < N_{\mathrm{coef}}$, and entries of degree at most B_{jet}. Corollary 3.12 with $b = B_{\mathrm{jet}}$ gives a primitive kernel vector of degree at most $\lfloor RB_{\mathrm{jet}}/(N_{\mathrm{coef}} - R) \rfloor$, which is at most the bound in the statement because $R/(N_{\mathrm{coef}} - R)$ increases with R. Its entries have no common root, so the resulting Q is nonzero at every specialization. Specializing $Z = z$ in the matrix identities shows that $Q(X, z, Y)$ satisfies every local constraint at $(\alpha_i, f_i + zg_i)$. The multiplicity and specialization-degree argument of Proposition 3.9 proves (17). For a received curve, the constraint entries at $(\alpha_i, w_i(Z))$ have Z-degree at most ℓB_{jet}; apply Corollary 3.12 with $b = \ell B_{\mathrm{jet}}$ to obtain (18). □

Bound	Degree of Q	Role in counting
B_∂	Y_1	regular families, resultants, and differentiations
B_{jet}	(Y_0, Y_1) jointly	algebraic degree of the solution families
B_Z	challenge Z	exceptional challenges

Table 3: The three degree bounds used in first-order interpolation and counting. Smaller bounds reduce the counts but leave fewer coefficients.

3.4 The first-order support

Proposition 3.13 reduces the construction to choosing a support whose specializations have degree below mA and whose dimension exceeds n times the local rank. At order one, the inherited local constraint at (α, y) is [Nas26, BCP+26a]

$$Q(\alpha + T, y + TY_1 + T^2V, Y_1) \equiv 0 \pmod{T^m}.$$

Naskręcki's construction already uses the linear dependencies among these constraints [Nas26]. Our new step bounds the local rank for supports in which the exponent of Y_1 is at most B_∂ and, separately, the jet degree, the total degree in (Y_0, Y_1), is at most B_{jet} (Table 3). The resulting bound $R_{\text{loc}}(m, B_\partial)$ in (20) does not depend on B_{jet}, so powers of Y_0 can add coefficients without raising it. A larger B_{jet} costs instead in the challenge degree of Proposition 3.13 and in the algebraic degree of the solution families counted in Section 5.

Fix $\delta > 0$ and integers $m \geq 1$, $B_\partial \geq 0$, and $B_{\text{jet}} \geq 1$. For $n \geq 2$, $1 \leq k \leq n$, and $k + \delta n \leq A \leq n$, put $K = \max\{k, 2\}$ and $D = K - 1 \geq 1$. Define the first-order support by

$$\mathcal{V}_{m,B_\partial,B_{\text{jet}}} = \operatorname{span}_F\{X^x Y_0^a Y_1^b : b \leq B_\partial,\ a + b \leq B_{\text{jet}},\ x + D(a + b) < mA\}. \tag{19}$$

The last inequality gives $\deg Q(X, P, P^{[1]}) < mA$ whenever $\deg P \leq D$. Define

$$\begin{aligned} R_{\text{loc}}(m, B_\partial) &= \sum_{s=0}^{m-1} \left[(s+1)(B_\partial + 1) - (2s - m + 1)_+(s + B_\partial - m + 1)_+\right] \\ &\overset{B_\partial \leq m/2}{=} (B_\partial + 1)\left[\frac{m(m + 1 - B_\partial)}{2} + \frac{B_\partial(2B_\partial + 1)}{6}\right], \end{aligned} \tag{20}$$

$$N_{\text{jet}}(t) = \min\{t, B_\partial\} + 1. \tag{21}$$

Here $N_{\text{jet}}(t)$ counts the monomials $Y_0^a Y_1^b$ with $a + b = t$ and $b \leq B_\partial$. At the received point $(0, 0)$, write $Y_0 = TU$ as in Figure 2. The summand of R_{loc} for T-degree s counts the $(s + 1)(B_\partial + 1)$ coefficients of $U^a Y_1^b$ with $a \leq s$ and $b \leq B_\partial$, minus the dimension of Naskręcki's kernel of multiples of $(U - Y_1)^{m-s}$. This kernel is nonzero only when $2s \geq m$ and $s + B_\partial \geq m$, so the dependencies reduce the count only at high T-degrees.

Lemma 3.14 (First-order dimension and rank)**.** *Over any field, the space in* (19) *is invariant under translations in X and Y_0, has total jet degree at most B_{jet}, and specializes to degree less than mA at every polynomial of degree less than K. Its order-one, multiplicity-m local constraint map has rank at most $R_{\text{loc}}(m, B_\partial)$, uniformly over the received point. Its coefficient count $N_{\text{coef}} = \dim \mathcal{V}$ satisfies*

$$N_{\text{coef}} = \sum_{t=0}^{B_{\text{jet}}} N_{\text{jet}}(t)(mA - Dt)_+ \geq n \sum_{t=0}^{B_{\text{jet}}} N_{\text{jet}}(t)\big(m\delta + (k/n)(m - t)\big)_+.$$

Proof. Translations lower x or a and preserve every support inequality. For $\deg P \leq D$, specialization of a monomial has degree at most $x + Da + (D - 1)b \leq x + D(a + b) < mA$. There are $N_{\text{jet}}(t)$ jet monomials of

total degree t. Each admits $(mA - Dt)_+$ nonnegative integer choices of x. Since $D \le k$ and $A \ge k + \delta n$, the dimension formula gives the stated lower bound term by term.

For the rank, we use the divisibility kernel that Naskręcki's construction [Nas26] finds in the local constraints, one Taylor degree at a time. Translation reduces the calculation to the received point $(0, 0)$. First substitute $X = T$ and $Y_0 = TU$, and reduce modulo T^m. At T-degree $s < m$, the image is contained in the rectangle spanned by $U^a Y_1^b$ with $0 \le a \le s$ and $0 \le b \le B_\partial$. The second substitution $U = Y_1 + \widetilde{V}$, followed by reduction modulo $\widetilde{V}^{m-s}$, imposes exactly the order-one local constraints at this T-degree. Its kernel on the whole rectangle consists of multiples of $(U - Y_1)^{m-s}$.

To count this kernel, note that multiplying a nonzero polynomial H by $(U - Y_1)^{m-s}$ increases both its separate U-degree and its separate Y_1-degree by $m - s$. Thus the kernel dimension on the rectangle is $(s - (m - s) + 1)_+ (B_\partial - (m - s) + 1)_+$. Subtracting it from $(s + 1)(B_\partial + 1)$ and summing over s proves the rank bound. The support can omit some monomials of the rectangle, so this count is an upper bound for its rank. The argument uses polynomial divisibility and inverts no binomial coefficients, hence holds in every characteristic. Evaluating the finite sum gives (20). □

3.5 The first-order agreement curve

We now turn the finite counts of Lemma 3.14 into an explicit threshold. Rescaled by m^3, the coefficient count and the local rank become integrals in the ratio B_∂/m, the densities Coef and Rank of (26). Optimizing this ratio gives the following curve (Figure 1), above which coefficients outnumber constraints. After locating the curve between capacity and Johnson, we choose explicit parameters above it, prove a finite margin with its growth in the gap (Lemma 3.16), and conclude with the symbolic explainer used in Section 5 (Corollary 3.17).

Definition 3.15 (First-order agreement curve)**.** Put

$$\rho_c := 11 - 3\sqrt{13} \approx 0.183346. \tag{22}$$

For $0 < \rho < \rho_c$, let u_ρ be the unique positive solution of

$$u^2(u + 3) = \sqrt{\rho/2}; \tag{23}$$

it exists because the left side increases strictly from zero to infinity for $u > 0$. Define

$$a_1(\rho) = \begin{cases} \sqrt{\rho/2}\,(1 + u_\rho), & 0 < \rho < \rho_c, \\ \dfrac{3\rho + 2\sqrt{\rho(5 - \rho)(2 - \rho)}}{8 - \rho}, & \rho_c \le \rho < 1. \end{cases} \tag{24}$$

Solving the cubic gives the equivalent $a_1(\rho) = \sqrt{2\rho}\cos\!\left(\frac{1}{3}\arccos(\sqrt{\rho/8} - 1)\right)$ on the lower branch. The two branches meet at $a_1(\rho_c) = 4 - \sqrt{13}$. At $a = a_1(\rho)$, the parameter λ_∂ chosen below, of which B_∂/m is the finite approximation, is at most $1/2$ on the upper branch and greater than $1/2$ on the lower branch, and both choices give $1/2$ at ρ_c. The two branches thus correspond to the two formulas for Rank in (26), which change at the ratio $1/2$ (proof of Lemma 3.16). On the upper branch, $a_1(\rho)$ is the positive root of the quadratic equation in a

$$\frac{a^2}{\rho} + \frac{3(1 - a)^2}{4(2 - \rho)} = 1, \tag{25}$$

and $a > a_1(\rho)$ exactly when the left side exceeds one: clearing denominators gives $(8-\rho)a^2 - 6\rho a - \rho(5-4\rho) = 0$, whose other root is negative.

For every $0 < \rho < 1$, we have $\rho < a_1(\rho) < \sqrt{\rho}$. On the upper branch, the left side of (25) is less than one at $a = \rho$ and greater than one at $a = \sqrt{\rho}$. On the lower branch, $u_\rho + 3 > 3$ gives $0 < u_\rho \le (\rho/2)^{1/4}/\sqrt{3}$, so

$$\sqrt{\rho/2} < a_1(\rho) \le \sqrt{\rho/2} + \frac{(\rho/2)^{3/4}}{\sqrt{3}}, \qquad \frac{a_1(\rho)}{\sqrt{\rho}} \longrightarrow \frac{1}{\sqrt{2}} \quad \text{as } \rho \to 0,$$

and $\rho < \sqrt{\rho/2}$ because $\rho < \rho_c < 1/2$. At low rates, the threshold is therefore smaller than Johnson's by a factor $1/\sqrt{2}$, while it remains a square-root threshold, well above the capacity threshold ρ. The improvement over the exact finite Johnson threshold $\sqrt{(k-1)/n}$ holds for sufficiently large n at each fixed rate.

Explicit parameters above the curve. The bounds of Section 5 need finite parameters and their growth in the gap, which we now choose. Fix $0 < \rho < a < 1$ with $a > a_1(\rho)$, write $\eta_1 = a - a_1(\rho)$, and let $n \ge 2/\rho$. The choices below depend only on (ρ, a, n), so they are fixed before the field, the evaluation set, and the received word. Put $\rho_n = \rho - 1/n$, so that the exact message-degree bound $D = k - 1$ satisfies $D \le \rho_n n$; this saving of one degree gives parameter and MCA bounds in terms of $\eta_1 + 1/n$ instead of η_1, which matters when the gap is smaller than $1/n$. To keep the parameters bounded as a approaches one, we use the clipped agreement fraction $a_{\mathrm{cert}} = \min\{a, (1 + a_1(\rho))/2\}$. The ratio $\lambda_\partial \approx B_\partial/m$ is the normalized bound on the exponent of Y_1, and we choose

$$\lambda_\partial = \begin{cases} \dfrac{a_1(\rho) - \sqrt{\rho/2}}{\rho}, & 0 < \rho < \rho_c, \\[2ex] \dfrac{3(1 - a_{\mathrm{cert}})}{2(2 - \rho)}, & \rho_c \le \rho < 1. \end{cases}$$

Two densities measure the coefficients and the constraints per received position on the scale m^3:

$$\mathsf{Coef}(r, a, \lambda_\partial) = \int_0^\infty \min\{u, \lambda_\partial\}(a - ru)_+ \, du = \frac{\lambda_\partial a^2}{2r} - \frac{a\lambda_\partial^2}{2} + \frac{r\lambda_\partial^3}{6} \quad (0 < \lambda_\partial < a/r), \tag{26}$$

$$\mathsf{Rank}(\lambda_\partial) = \int_0^1 \left[u\lambda_\partial - (2u-1)_+(u + \lambda_\partial - 1)_+\right] du = \begin{cases} \lambda_\partial/2 - \lambda_\partial^2/2 + \lambda_\partial^3/3, & \lambda_\partial \le 1/2, \\ \lambda_\partial/4 + 1/24, & \lambda_\partial \ge 1/2. \end{cases}$$

They are the limits of $\underline{\nu}_{\mathrm{coef}}/m^3$ and R_{loc}/m^3 below, with total jet degree and Taylor degree rescaled to $u = t/m$ and $u = s/m$. The two formulas for Rank differ by $(\lambda_\partial - 1/2)^3/3$, so Rank is their minimum and either formula is an upper bound for it. Put $\mathsf{Margin}_{\mathrm{int}}(r, a, \lambda_\partial) := \mathsf{Coef}(r, a, \lambda_\partial) - \mathsf{Rank}(\lambda_\partial)$ and

$$\mathsf{Margin}_{\mathrm{int}}^{(n)} = \mathsf{Margin}_{\mathrm{int}}(\rho_n, a_{\mathrm{cert}}, \lambda_\partial). \tag{27}$$

The remaining parameters are

$$C_{\mathrm{rank}} = 2\lambda_\partial + 3, \qquad m = \left\lceil \frac{4C_{\mathrm{rank}}}{\mathsf{Margin}_{\mathrm{int}}^{(n)}} \right\rceil, \qquad B_\partial = \lfloor \lambda_\partial m \rfloor, \qquad B_{\mathrm{jet}} = \lceil m a_{\mathrm{cert}}/\rho_n \rceil,$$

$$\underline{\nu}_{\mathrm{coef}} = \sum_{t=0}^{B_{\mathrm{jet}}} (\min\{t, B_\partial\} + 1)(m a_{\mathrm{cert}} - \rho_n t)_+,$$

$$B_Z = \max\left\{1, \left\lfloor \frac{R_{\mathrm{loc}}(m, B_\partial) B_{\mathrm{jet}}}{\underline{\nu}_{\mathrm{coef}} - R_{\mathrm{loc}}(m, B_\partial)} \right\rfloor \right\}. \tag{28}$$

Here $\underline{\nu}_{\mathrm{coef}}$ lower-bounds the number of coefficients per position. The next lemma proves positivity of both $\mathsf{Margin}_{\mathrm{int}}^{(n)}$ and $\underline{\nu}_{\mathrm{coef}} - R_{\mathrm{loc}}$.

Lemma 3.16 (Finite coefficient-constraint margin above the first-order curve)**.** *Let* $0 < \rho < a < 1$ *with* $a > a_1(\rho)$ *and* $n \geq 2/\rho$*. Then* $0 < \lambda_\partial < a_{\mathrm{cert}}/\rho$*, and* $\mathsf{Margin}_{\mathrm{int}}^{(n)} \geq \mathsf{Margin}_{\mathrm{int}}(\rho, a_{\mathrm{cert}}, \lambda_\partial)$*, where*

$$\mathsf{Margin}_{\mathrm{int}}(\rho, a_{\mathrm{cert}}, \lambda_\partial) = \begin{cases} \dfrac{\lambda_\partial}{2}\big(a_{\mathrm{cert}} - a_1(\rho)\big)\left[\dfrac{a_{\mathrm{cert}} + a_1(\rho)}{\rho} - \lambda_\partial\right], & 0 < \rho < \rho_c, \\ \dfrac{\lambda_\partial}{2}\left[\dfrac{a_{\mathrm{cert}}^2}{\rho} - 1 + \dfrac{3(1 - a_{\mathrm{cert}})^2}{4(2 - \rho)}\right], & \rho_c \leq \rho < 1, \end{cases} \tag{29}$$

is positive. The parameters in (28) *satisfy* $0 \leq B_\partial \leq B_{\mathrm{jet}}$ *and*

$$\underline{\nu}_{\mathrm{coef}} - R_{\mathrm{loc}}(m, B_\partial) \geq m^3 \mathsf{Margin}_{\mathrm{int}}^{(n)} - C_{\mathrm{rank}} m^2 \geq 3m^3 \mathsf{Margin}_{\mathrm{int}}^{(n)}/4 > 0.$$

For fixed ρ*, writing* $\bar{\eta}_1 = \eta_1 + 1/n$*, they have* $m, B_\partial, B_{\mathrm{jet}} = O_\rho(\bar{\eta}_1^{-1})$ *and* $B_Z = O_\rho(\bar{\eta}_1^{-2})$*.*

Proof. The upper branch. Let $\rho \geq \rho_c$. At $a = (1+\rho)/3$, the left side of (25) equals $1 + (\rho^2 - 22\rho + 4)/(36\rho)$. This is at most one for $\rho_c \leq \rho < 1$, since $\rho_c = 11 - 3\sqrt{13}$ is the smaller root of $\rho^2 - 22\rho + 4$. Thus $a_1(\rho) \geq (1+\rho)/3$, and $a_{\mathrm{cert}} > a_1(\rho) \geq (1+\rho)/3$, which gives $0 < \lambda_\partial \leq 1/2 < 1 < a_{\mathrm{cert}}/\rho$. With the first rank formula, the margin factors as

$$\mathsf{Margin}_{\mathrm{int}}(\rho, a, \lambda_\partial) = \frac{\lambda_\partial}{2}\left[\frac{a^2}{\rho} - 1 + (1-a)\lambda_\partial - \frac{2-\rho}{3}\lambda_\partial^2\right].$$

The bracket is a concave quadratic in λ_∂, maximized at $\lambda_\partial = 3(1-a)/(2(2-\rho))$; at $a = a_{\mathrm{cert}}$ this is our choice, and the bracket becomes that of (29). It is positive because $a_{\mathrm{cert}} > a_1(\rho)$. As $\lambda_\partial \to 0$, the bracket tends to the Johnson condition $a^2/\rho - 1$; a positive derivative-degree ratio adds the second term of (25), so the bracket can be positive for some $a < \sqrt{\rho}$.

The lower branch. Let $\rho < \rho_c$, and put $u = \sqrt{2\rho}\,\lambda_\partial$, so that $u = u_\rho$ and $\lambda_\partial < a_1(\rho)/\rho < a_{\mathrm{cert}}/\rho$. At $u = \sqrt{\rho/2}$, the left side of (23) is less than $\sqrt{\rho/2}$ exactly when $\sqrt{\rho/2}(\sqrt{\rho/2} + 3) < 1$, that is, when $\rho < \rho_c$. Hence $u_\rho > \sqrt{\rho/2}$ and $\lambda_\partial > 1/2$, so the second rank formula applies. The lower branch is the point where the margin and its derivative in λ_∂ both vanish:

$$\mathsf{Margin}_{\mathrm{int}} = 0, \qquad \frac{\partial \mathsf{Margin}_{\mathrm{int}}}{\partial \lambda_\partial} = \frac{(a - \rho\lambda_\partial)^2}{2\rho} - \frac{1}{4} = 0.$$

The second equation gives $a - \rho\lambda_\partial = \sqrt{\rho/2}$, and substituting into the first yields

$$a = \sqrt{\rho/2} + \rho\lambda_\partial, \qquad 4\rho\lambda_\partial^3 + 6\sqrt{2\rho}\,\lambda_\partial^2 = 1. \tag{30}$$

With $u = \sqrt{2\rho}\,\lambda_\partial$ and $a = \sqrt{\rho/2}(1+u)$, these equations reduce to (23). Thus our λ_∂ satisfies $\mathsf{Margin}_{\mathrm{int}}(\rho, a_1(\rho), \lambda_\partial) = 0$. Subtracting this from $\mathsf{Margin}_{\mathrm{int}}(\rho, a_{\mathrm{cert}}, \lambda_\partial)$ leaves only the terms of Coef involving a, and gives the first case of (29). Both factors are positive because $a_{\mathrm{cert}} > a_1(\rho)$ and $\lambda_\partial < a_1(\rho)/\rho$.

Finite length. Since Coef decreases in its first argument, $\rho_n < \rho$ gives the stated inequality for $\mathsf{Margin}_{\mathrm{int}}^{(n)}$. Integrating the decreasing positive-part function over unit boxes under the sum gives

$$\underline{\nu}_{\mathrm{coef}} \geq m^3 \mathsf{Coef}(\rho_n, a_{\mathrm{cert}}, \lambda_\partial).$$

The boxes with $0 \leq b \leq \lfloor \lambda_\partial m \rfloor$ cover the strip $0 \leq b \leq \lambda_\partial m$, and the total-degree cutoff includes every positive summand. For the rank, put $r_{\mathrm{dens}}(u) = u\lambda_\partial - (2u-1)_+(u + \lambda_\partial - 1)_+$. The summand at s is at most

$$m^2 r_{\mathrm{dens}}(s/m) + m(s/m + \lambda_\partial) + 1.$$

In general, r_{dens} has Lipschitz constant at most $1+\lambda_\partial$. Comparing the left Riemann sum with its integral and summing the remaining terms gives $R_{\mathrm{loc}}(m, B_\partial) \le m^3 \mathsf{Rank}(\lambda_\partial) + (2\lambda_\partial + 3)m^2$. The definitions of C_{rank} and m give $m\mathsf{Margin}_{\mathrm{int}}^{(n)} \ge 4C_{\mathrm{rank}}$ and prove the claim. Also $\lambda_\partial < a_{\mathrm{cert}}/\rho < a_{\mathrm{cert}}/\rho_n$ gives $B_\partial \le B_{\mathrm{jet}}$. If $B_\partial = 0$, the same counts give an ordinary interpolant, which is covered by the same bounds.

Growth. On either branch, (29) has a simple zero at $a = a_1(\rho)$. For $a_1(\rho) < a_{\mathrm{cert}} \le (1 + a_1(\rho))/2$, $\mathsf{Margin}_{\mathrm{int}} = \Theta_\rho(\eta_1)$, and λ_∂ stays between positive ρ-dependent constants. For $\rho_n \le r \le \rho$, we have $r \ge \rho/2$ and $\lambda_\partial < a_{\mathrm{cert}}/r$, so $-\partial_r \mathsf{Coef}(r, a_{\mathrm{cert}}, \lambda_\partial) = \lambda_\partial a_{\mathrm{cert}}^2/(2r^2) - \lambda_\partial^3/6 > \lambda_\partial a_{\mathrm{cert}}^2/(3r^2)$. This derivative is bounded above and below by positive constants depending only on ρ, and integrating over an interval of length $1/n$ proves $\mathsf{Margin}_{\mathrm{int}}^{(n)} = \Theta_\rho(\eta_1 + 1/n)$. Consequently $m, B_\partial, B_{\mathrm{jet}} = O_\rho(\bar\eta_1^{-1})$, $R_{\mathrm{loc}} = O_\rho(m^3)$, and $B_Z \le 1 + 4R_{\mathrm{loc}}(m, B_\partial)B_{\mathrm{jet}}/(3m^3\mathsf{Margin}_{\mathrm{int}}^{(n)}) = O_\rho(\bar\eta_1^{-2})$. □

The positive margin now gives the required symbolic explainer.

Corollary 3.17 (Explainer above the curve)**.** *Fix $0 < \rho < a < 1$ with $a > a_1(\rho)$, integers $n \ge 2$, $2 \le k \le \rho n$, and $an \le A \le n$, and choose $(m, B_\partial, B_{\mathrm{jet}}, B_Z)$ by (28). Over every field, for all distinct evaluation points and all $f, g \in F^n$, there is $Q \in F[Z] \otimes_F \mathcal{V}_{m,B_\partial,B_{\mathrm{jet}}}$ with $\deg_Z Q \le B_Z$ and $Q(X, z, Y) \ne 0$ for every z, satisfying (17) at degree bound k and threshold A.*

Proof. By Lemma 3.14, $N_{\mathrm{coef}} = \sum_t N_{\mathrm{jet}}(t)(mA - Dt)_+ \ge n\underline{\nu}_{\mathrm{coef}}$, since $A \ge a_{\mathrm{cert}}n$ and $D \le \rho_n n$; the local rank is at most R_{loc}. Lemma 3.16 gives $\underline{\nu}_{\mathrm{coef}} > R_{\mathrm{loc}}$. Proposition 3.13 then gives

$$\deg_Z Q \le \left\lfloor \frac{R_{\mathrm{loc}} B_{\mathrm{jet}}}{\underline{\nu}_{\mathrm{coef}} - R_{\mathrm{loc}}} \right\rfloor \le B_Z,$$

because $x/(N - x)$ increases in x and decreases in N. □

4 From an equation to lists and MCA

Because an explainer can have a positive-dimensional family of solutions (Section 2.2), we bound lists and find common witnesses for any algebraic family of bounded dimension and degree, then build such families for ordinary equations, which involve no derivatives. This gives Johnson MCA with $O_\rho(n/\eta_0^3)$ exceptional challenges in every characteristic (Theorem 4.4); Section 5 treats the first-order explainer.

4.1 Incidence and common-witness transfer

Lemma 4.1 bounds the points of a family that agree with a received word in many positions, and Lemma 4.2 gives common witnesses for the close candidates outside a small set of challenges. Both follow Zheng's incidence and persistent-line arguments [Zhe26], with sharper degrees and agreement thresholds.

The arguments use only two quantities of a family: its dimension, which counts free parameters, and its degree, which bounds how many components and isolated points remain after we impose agreement equations. All algebraic sets below are common zero sets of polynomials over an algebraically closed field E, in finitely many coordinates.

We recall the few notions these quantities require. An *irreducible component* is a maximal piece that is not a union of two proper algebraic subsets, and the *Zariski closure* of a set is the smallest algebraic set containing it. The degree of an irreducible s-dimensional set is the number of points, counted with multiplicity, in which its projective closure meets a general linear subspace of codimension s. The *total degree* of an algebraic set is the sum of the degrees of its distinct irreducible components, including components of smaller dimension. A rational map has coordinates that are ratios of polynomials, is used where its denominators are nonzero, and has as *image closure* the Zariski closure of its values; Section A bounds the degrees of such images.

To bound degrees after cutting, we use the refined Bézout inequality [Hei83, Ful98] in the following forms. A hypersurface defined by a nonzero polynomial of degree b has total degree at most b. A proper hyperplane section of an irreducible set has smaller dimension and total degree at most that of the set, and intersecting with a linear subspace does not increase total degree. More generally, successive cuts by polynomials of degree at most $b \geq 1$ multiply total degree by at most b at each cut, keeping a component when the cutting polynomial vanishes identically on it.

We now fix the spaces in which the candidates lie. In the coefficient space of degree-$< k$ polynomials, the condition $P(\alpha_i) = y_i$ is an *agreement hyperplane*. Any $c \leq k$ of them are independent by Vandermonde inversion, so together they lower the dimension by exactly c. For MCA we add the challenge and its powers as coordinates, which keeps agreement with a received curve linear. Given words $w_0, \ldots, w_\ell \in F^n$, a received curve $w(z) = \sum_{t \leq \ell} z^t w_t$, and $\ell \geq 1$, we work in the *lifted locus*

$$\mathbb{L}_\ell = \{(z, z^2, \ldots, z^\ell, P) : \deg P < k\}$$

over E, with coordinates $T_1 = z, T_2, \ldots, T_\ell$ and the coefficients of P. Agreement at position i is the affine hyperplane

$$H_i = \{P(\alpha_i) = w_{0,i} + z w_{1,i} + \textstyle\sum_{t=2}^{\ell} T_t w_{t,i}\},$$

which is linear in these coordinates; for a line, $\ell = 1$ and $\mathbb{L}_1$ is the joint space of pairs (z, P). Within $\mathbb{L}_\ell$, any k of the H_i meet in a *witness graph*, traced by witnesses P_t that interpolate w_t on those positions:

$$\Gamma_{P_0,\ldots,P_\ell} = \{(z, z^2, \ldots, z^\ell, \textstyle\sum_{t=0}^{\ell} z^t P_t) : z \in E\}, \qquad \deg P_t < k.$$

A graph contained in H_i is one whose witnesses agree with every $w_{t,i}$ at position i; we call these positions its common agreements. Each graph is isomorphic to the affine line through its z coordinate, so every positive-dimensional irreducible closed subset of it is the whole graph.

For a fixed received word, an irreducible curve of degree Δ has at most $\Delta(n - k + 1)/(T - k + 1)$ points with T agreements (Section 2.2), and the incidence bound has one such factor per dimension: for $r \geq 0$ and $k \leq T \leq n$, define

$$\Pi_r(T) := \prod_{i=1}^{r} \frac{n - k + i}{T - k + i}, \tag{31}$$

an empty product when $r = 0$. Each factor is at least one.

Lemma 4.1 (Dimension-sensitive agreement incidence)**.** *In the coefficient space of degree-$< k$ polynomials, an algebraic set of dimension at most r and total degree at most B has at most $B\Pi_r(T)$ points agreeing with a fixed word in at least T positions.*

In the joint space $\mathbb{L}_1$ of pairs (z, P), exclude the witness graphs with at least L common agreements, where $k \leq L \leq A \leq n$. For a set of dimension at most $r + 1$ and total degree at most B, the number of remaining points with at least A agreements is at most

$$B \frac{n - L + 1}{A - L + 1} \Pi_r(A).$$

The same bound holds on the lifted locus $\mathbb{L}_\ell$ for every $\ell \geq 1$, excluding the witness graphs with at least L common agreements.

Proof. We induct on the dimension of an irreducible component C; a zero-dimensional component is a point. If $\dim C = t \geq 1$ and c agreement hyperplanes contain C, Vandermonde inversion gives $c \leq k - t$: any $c \leq k$ evaluation equations leave dimension $k - c$. Each point with T agreements lies on at least $T - c$ of the

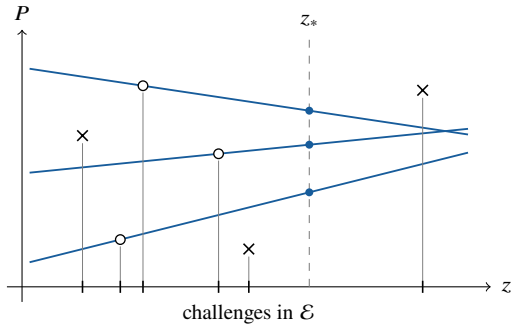

Figure 4: A family of pairs (z, P) for a line, projected with P drawn as one coordinate, so crossings need not be intersections. Blue lines are witness graphs in the family with at least L common agreements. Crosses are points off them with at least A agreements; open circles are accidental agreements, at most $n - L$ per graph. Their challenges in F and $\mathcal{E}_{\text{pre}}$ form $\mathcal{E}$. Over a transcendental z_*, distinct graphs give distinct points.

remaining hyperplanes. Their proper sections have smaller dimension and total degree at most $\deg C$, so induction first proves finiteness and then bounds their point incidences by

$$(T - c)N(C) \le (n - c)\deg(C)\Pi_{t-1}(T).$$

Since $(n - c)/(T - c) \le (n - k + t)/(T - k + t)$, this gives $N(C) \le \deg(C)\Pi_t(T)$. A section that lowers the dimension by more than one satisfies the same bound, because all incidence factors are at least one. Summing over components proves the fixed-word assertion.

In the joint space, imposing $c \le k$ agreement equations instead leaves dimension $k + 1 - c$. A component of dimension $t \ge 2$ lies in at most $k + 1 - t$ hyperplanes and contributes the factor $(n - k + t - 1)/(A - k + t - 1)$. At dimension one, a component contained in L agreement hyperplanes is one of the excluded graphs: any k of the equations determine its coefficient polynomials by Vandermonde inversion. Every other curve contributes a factor at most $(n - L + 1)/(A - L + 1)$. The same induction, with zero-dimensional components as its base, proves the joint bound and finiteness. On $\mathbb{L}_\ell$, solving c evaluation equations still leaves the challenge and $k - c$ free coefficients, so its dimension is $k + 1 - c$. This proves the final assertion by the identical induction. □

We now combine the two bounds into the common-witness transfer. In (32), the last two terms count the close pairs off the witness graphs with at least L common agreements and the accidental agreements on those graphs (Figure 4). The first term is a preliminary set $\mathcal{E}_{\text{pre}}$ of challenges excluded in advance, such as those at which the equation degenerates in Lemma 4.3, or the exceptional set of the ordinary equation in Lemma 5.5. The family may contain points that are not candidates; it need only contain every candidate whose challenge lies outside $\mathcal{E}_{\text{pre}}$.

Lemma 4.2 (Geometric transfer on a polynomial received curve)**.** *Let F be a field, let E be an algebraically closed extension, and let $1 \le k \le L \le A \le n$ and $\ell \ge 1$. Fix distinct $\alpha_1, \dots, \alpha_n \in F$ and words $w_0, \dots, w_\ell \in F^n$, and work in the lifted locus $\mathbb{L}_\ell$. Let $V_1, \dots, V_N \subseteq \mathbb{L}_\ell$ be closed algebraic sets, and for each ν let $r_\nu \ge 0$ be an integer and J_ν a number with*

$$\dim V_\nu \le r_\nu + 1, \qquad \deg V_\nu \le J_\nu,$$

where deg *is total degree. Let z_* be transcendental over E, and form the fibers after base change to an algebraic closure of $E(z_*)$. Suppose that the reduced algebraic set underlying $(V_\nu)_{z_*}$ has dimension at most r_ν and total degree at most F_ν.*

Let C be a collection of pairs (z, P) with $z \in F$ and $P \in F[X]$ of degree less than k. Suppose that, outside a finite preliminary set $\mathcal{E}_{\text{pre}} \subseteq F$, every $(z, P) \in C$ gives the point $(z, z^2, \dots, z^\ell, P) \in \bigcup_\nu V_\nu$. Then there is $\mathcal{E} \subseteq F$ of size at most

$$|\mathcal{E}_{\text{pre}}| + \frac{n - L + 1}{A - L + 1} \sum_\nu \Pi_{r_\nu}(A) J_\nu + \ell(n - L) \sum_\nu \Pi_{r_\nu}(L) F_\nu \tag{32}$$

such that the following holds. For every $(z, P) \in C$ *with* $z \notin \mathcal{E}$ *and*

$$|\{i : P(\alpha_i) = \textstyle\sum_{t=0}^{\ell} z^t w_{t,i}\}| \geq A,$$

there are $P_0, \ldots, P_\ell \in F[X]$ *of degree less than* k *such that*

$$P = \sum_{t=0}^{\ell} z^t P_t, \qquad \{i : P(\alpha_i) = \textstyle\sum_{t=0}^{\ell} z^t w_{t,i}\} = \{i : P_t(\alpha_i) = w_{t,i} \text{ for every } t\}.$$

Proof. Use the agreement hyperplanes H_i and witness graphs defined above, and consider the graphs contained in V_ν and in at least L of the H_i. Distinct graphs give distinct points over z_*: their coefficient polynomials in z cannot coincide at a transcendental element. The fixed-word incidence bound on the reduced fiber therefore gives at most $\Pi_{r_\nu}(L)F_\nu$ such graphs. The joint incidence bound gives at most

$$\frac{n - L + 1}{A - L + 1}\Pi_{r_\nu}(A)J_\nu$$

high-agreement points off these graphs. This count includes the components over finitely many challenges: they miss the transcendental fiber and contain no graph, since graphs project onto the challenge line.

Let $\mathcal{E} \subseteq F$ consist of $\mathcal{E}_{\mathrm{pre}}$, the base-field challenges of these high-agreement points, and the base-field challenges at which one of the counted graphs has an accidental agreement. Outside a graph's common set, an agreement equation is a nonzero polynomial in z of degree at most ℓ. There are at most $n - L$ such positions, so each graph adds at most $\ell(n - L)$ challenges to $\mathcal{E}$. Summation proves (32).

Outside $\mathcal{E}$, each candidate lies on one of the counted graphs and has no agreement outside its common set. Any k common positions recover $P_0, \ldots, P_\ell$ by interpolation over F, so the witnesses have coefficients in F; this proves the decomposition and the equality of agreement sets. □

4.2 Ordinary equations in every characteristic

Lemma 4.2 needs a family containing the close candidates. We build one for an *ordinary* equation $Q(X, z, P(X)) = 0$, which involves no derivative of P. Such equations arise in the Johnson theorem and, in Section 5, by eliminating Y_1 from a first-order equation and its Y_1-derivative. Following BCIKS and Zheng [BCI+23, Zhe26], we track solutions across the challenge; we also extend the argument to inseparable factors and keep the degrees in Y and Z separate.

In the following bound, B and H stand for the degrees of the equation in Y and Z; Lemma 4.3 uses it with $(B, H) = (B_{\mathrm{jet}}, B_Z)$, and Section 5 with other degrees. The threshold L is that of Lemma 4.2. For integers $D \geq 1$, $n \geq A > D$, $B \geq 1$, $H \geq 0$, a received-curve degree $\ell \geq 1$, and $D + 1 \leq L \leq A$, define

$$\Psi_D(B) = 1 + (2D - 1)(2B - 1) + 2(B - 2D - 1)_+,$$

$$\mathsf{E}^{(\ell)}_{\mathrm{ord}}(D, A, n; B, H; L) = (2B - 1)H + \frac{n - L + 1}{A - L + 1}\big(\ell B + H\Psi_D(B)\big) + \ell(n - L)B, \tag{33}$$

$$\mathsf{E}^{(\ell)}_{\mathrm{ord}}(D, A, n; B, H) := \mathsf{E}^{(\ell)}_{\mathrm{ord}}(D, A, n; B, H; D + 1). \tag{34}$$

Omitting L means $L = D + 1$.

Lemma 4.3 (Ordinary MCA transfer in every characteristic)**.** *Let* F *be a field, let* $1 \leq D \leq n - 2$*, and let* $D + 1 \leq L \leq A \leq n$*. Fix distinct* $\alpha_1, \ldots, \alpha_n \in F$ *and a received curve* $w(Z) = \sum_{t=0}^{\ell} Z^t w_t$*, where* $\ell \geq 1$ *and* $w_t \in F^n$*. Suppose the nonzero polynomial* $Q \in F[X, Z, Y]$ *satisfies*

$$\deg_Y Q \leq B_{\mathrm{jet}}, \qquad \deg_Z Q \leq B_Z, \qquad B_{\mathrm{jet}} \geq 1.$$

There is a set $\mathcal{E}_Q \subseteq F$ of size at most

$$\mathsf{E}^{(\ell)}_{\mathrm{ord}}(D, A, n; B_{\mathrm{jet}}, B_Z; L) \tag{35}$$

such that every $z \in F \setminus \mathcal{E}_Q$ and every $P \in F[X]_{\le D}$ satisfying

$$Q(X, z, P(X)) = 0, \qquad |S_z(P)| \ge A,$$

has the full common-witness property of Definition 3.1, with message dimension $D + 1$. No restriction on the characteristic or on ℓ is needed. The line case is $\ell = 1$.

The three terms of (33) count the challenges at which the content of Q in Y vanishes identically or a factor of Q acquires a repeated root in Y, the points off witness graphs, and the accidental agreements on those graphs. The threshold L balances the last two. The lemma concerns close solutions of Q; interpolation must separately show that every close message solves it.

The proof in Section B.1 first makes each factor separable in the message variable. In characteristic p, an irreducible factor that is a polynomial in Y^p has zero Y-derivative, so all its roots are repeated; a Frobenius pullback replaces X and Z by p^ν-th powers and takes a polynomial p^ν-th root, as in (95).

At a transcendental center, the order-zero Taylor recurrence then reconstructs each solution from its value there. The coefficient of each new unknown is the derivative of the pulled-back factor in the message variable, rather than j times the separant as at first order (Section 5.2), so no characteristic condition arises. The resulting family has joint degree at most $\ell B_{\mathrm{jet}} + B_Z \Psi_D(B_{\mathrm{jet}})$ and fiber degree at most B_{jet} over a transcendental challenge, so Lemma 4.2 applies.

4.3 The Johnson theorem

As a first use of Lemma 4.3, we take for Q the Guruswami–Sudan interpolant [GS99] at order zero, that is, without derivative variables, with the challenge kept symbolic as in BCHKS [BCH+26], who obtain $O_\rho(n/\eta_0^5)$ exceptional challenges at Johnson gap η_0. We keep their interpolant and change how the exceptional challenges of each factor are counted: reconstructing only the $D + 1$ message coefficients removes one factor $1/\eta_0$, and bounding the degrees in Y and Z separately in the image bound of Lemma 4.3 removes another. Both savings have antecedents for concrete parameters in the Proximity Prize submissions of BitWonka and jieyilong [Bit26a, jie26a].

Theorem 4.4 (Quantitative Johnson MCA in every characteristic)**.** *Let F be any field, let $\alpha_1, \dots, \alpha_n \in F$ be distinct, and suppose $1 \le D \le n - 2$. Put $\rho = (D + 1)/n$ and $\rho_- = D/n$. Choose $\eta_0 > 0$ with $a := \sqrt{\rho_-} + \eta_0 \le 1$, and define*

$$A = \lceil an \rceil, \qquad m = \max\left\{ \left\lceil \frac{\sqrt{\rho_-}}{2\eta_0} \right\rceil, 3 \right\}, \qquad t = m + \frac{1}{2}, \tag{36}$$

and

$$B_{\mathrm{jet}} = \left\lceil \frac{t}{\sqrt{\rho_-}} \right\rceil - 1, \qquad B_Z = \left\lceil \frac{t^2}{3\rho_-} \right\rceil - 1, \qquad \lambda_{\mathrm{agr}} = \frac{n - D}{A - D}. \tag{37}$$

Then every received word $y \in F^n$ satisfies the classical list bound

$$|\mathcal{L}_D(y, A)| \le \min\left\{ B_{\mathrm{jet}}, \left\lfloor \frac{n(A - D)}{A^2 - nD} \right\rfloor \right\}. \tag{38}$$

Moreover, for every $f, g \in F^n$ there is a set $\mathcal{E} \subseteq F$ such that

$$|\mathcal{E}| \le B_{\mathrm{cf}} := \mathsf{E}^{(1)}_{\mathrm{ord}}(D, A, n; B_{\mathrm{jet}}, B_Z) < \frac{8}{3} \frac{nt^3}{\rho_-}. \tag{39}$$

For every $z \notin \mathcal{E}$ and every $P \in F[X]_{\le D}$ whose full agreement set

$$S_z(P) = \{i : P(\alpha_i) = f_i + zg_i\}$$

has size at least A, there are $F_0, G_0 \in F[X]_{\le D}$ with

$$P = F_0 + zG_0, \qquad S_z(P) = \{i : F_0(\alpha_i) = f_i \text{ and } G_0(\alpha_i) = g_i\}.$$

All assertions are uniform over the prescribed evaluation set. If F is finite and the challenge is uniform in F, the exceptional probability is at most $B_{\mathrm{cf}}/|F|$. In particular, at fixed rate the exceptional count is $O_\rho(n/\eta_0^3)$, giving probability $O_\rho(n/(q\eta_0^3))$ over $\mathbb{F}_q$.

Proof. Apply Lemma 4.3 with $\ell = 1$ and $L = D + 1$ to the interpolant of Proposition B.1. Every close candidate satisfies this interpolant at its challenge, so the lemma gives the common witnesses; the factor count and pairwise Johnson argument following that proposition give the list bound. The rounding calculation in Section B.3 gives $B_{\mathrm{cf}} < 8nt^3/(3\rho_-)$. Since $\eta_0 \le 1$, we have $t \le 4/\eta_0$, and $\rho_- \ge \rho/2$ because $D \ge 1$; this gives the asymptotic form. □

The theorem excludes the degree-zero code $D = 0$, where ρ_- vanishes. At $D = n - 1$ every word is a codeword, and line MCA holds with $\mathcal{E} = \varnothing$ by interpolating f and g separately. Since Lemma 4.3 holds for received curves of every degree ℓ, the same argument extends to polynomial curves in every characteristic (Remark 6.1).

5 First-order list and MCA bounds

Above the first-order curve, Corollary 3.17 gives an explainer, and Section 4 bounds lists and MCA for candidates that lie in a family of known dimension and degree. This section places the explainer's solutions in such families, as sketched in Section 2.3, and proves the bounds of Section 5.1; Theorem 5.7 gives explicit constants, independent of the rate, at agreement $k + 6n/25$.

5.1 The bounds above the first-order curve

As in Table 3, the explainer has degree at most B_{jet} in (Y_0, Y_1) jointly, B_∂ in Y_1, and B_Z in the challenge.

Theorem 5.1 (First-order list decoding above the first-order curve)**.** *Fix $0 < \rho < a < 1$ with $a > a_1(\rho)$. Let $n \ge 2$, $2 \le k \le \rho n$, and $an \le A \le n$ be integers, and choose $(m, B_\partial, B_{\mathrm{jet}}, B_Z)$ by (28). Put $D = k - 1$ and $\lambda_{\mathrm{agr}} = (n - D)/(A - D)$. Let F be a field of characteristic zero or $p > \max\{k - 1, B_\partial\}$, and fix n distinct evaluation points in F. Every received word $y \in F^n$ satisfies*

$$|\mathcal{L}(y, A)| \le 4D\lambda_{\mathrm{agr}} B_{\mathrm{jet}} B_\partial + 2B_{\mathrm{jet}} B_\partial + B_{\mathrm{jet}} = O_\rho\left(\frac{n}{(\eta_1 + 1/n)^2}\right) = O_\rho(n/\eta_1^2). \tag{40}$$

Theorem 5.2 (First-order line MCA above the first-order curve)**.** *Fix $0 < \rho < a < 1$ with $a > a_1(\rho)$. Let $n \ge 2$, $2 \le k \le \rho n$, and $an \le A \le n$ be integers, choose $(m, B_\partial, B_{\mathrm{jet}}, B_Z)$ by (28), and put $D = k - 1$, $\lambda_{\mathrm{agr}} = (n - D)/(A - D)$, and*

$$b = B_{\mathrm{jet}}(2B_\partial + 1), \qquad h = B_Z(2B_\partial + 1).$$

Let F be a field of characteristic zero or $p > \max\{k-1, B_\partial\}$, and fix n distinct evaluation points in F. For every received line there is one exceptional set $\mathcal{E}$ with

$$\begin{aligned} |\mathcal{E}| &\le (2b-1)h + \lambda_{\mathrm{agr}}(h + b + 4Dbh) + (n-D-1)b \\ &\quad + (48D^2 B_Z + 16D)\lambda_{\mathrm{agr}}^2 B_{\mathrm{jet}} B_\partial + 8D(n-D-1)\lambda_{\mathrm{agr}} B_{\mathrm{jet}} B_\partial \\ &= O_\rho\left(\frac{n^2}{(\eta_1 + 1/n)^4}\right) = O_\rho(n^2/\eta_1^4) \end{aligned} \tag{41}$$

such that every $z \notin \mathcal{E}$ and every degree-$<k$ candidate P with $|S_z(P)| \ge A$ admit degree-$<k$ witnesses F_0, G_0 with $P = F_0 + zG_0$ and $S_z(P) = S(F_0, G_0)$. Over $\mathbb{F}_q$, divide the bound by q for the exceptional probability.

The two lines of (41) correspond to the two kinds of solutions separated below: the first counts solutions of an ordinary equation, and the second the regular first-order family. The two characteristic conditions also have separate sources. Taylor reconstruction through degree D divides by $1, \dots, D$, so it needs $p > k-1$; the factors of positive Y_1-degree must be separable, which holds when $p > B_\partial$. Neither condition is needed for the ordinary part, which Lemma 4.3 bounds in every characteristic, and when $B_\partial = 0$ the regular terms vanish.

5.2 Regular and singular solutions

Kopparty's recurrence describes a solution through its initial Taylor coefficients [Kop15]. At a Taylor center x_0, write $P(x_0 + T) = \sum_j c_j T^j$, so that the initial pair (c_0, c_1) gives the value and first derivative. For a differential equation $Q(X, P, P^{[1]}) = 0$, the *separant* is $\partial Q/\partial Y_1$. When its value at (x_0, c_0, c_1) is nonzero, the equation determines the later coefficients one at a time: the coefficient of the next unknown c_j is j times the separant, so $p > D$ suffices through degree D. Following Zheng [Zhe26], we keep the initial pair as variables on the curve $Q(x_0, c_0, c_1) = 0$. The solutions with nonzero separant at x_0, the *regular* solutions, then lie in the image of one rational map from this curve, and the incidence bound needs the degree of that image.

For the remaining solutions, we use Naskręcki's separation of regular and singular solutions, which eliminates Y_1 with a resultant [Nas26]. Write the squarefree part of Q as $Q_{\mathrm{ord}} Q_{\mathrm{diff}}$, where Q_{ord} collects the factors free of Y_1; every solution of $Q = 0$ solves one of the two. Take the center x_0 transcendental. If the separant $\partial Q_{\mathrm{diff}}/\partial Y_1$ is not identically zero along a solution of $Q_{\mathrm{diff}} = 0$, it is nonzero at x_0, so the solution is regular. Otherwise $P^{[1]}$ is a common root in Y_1 of Q_{diff} and its separant, so the solution satisfies the ordinary equation $Q_{\mathrm{ord}} \operatorname{Res}_{Y_1}(Q_{\mathrm{diff}}, \partial Q_{\mathrm{diff}}/\partial Y_1) = 0$.

Passing to the squarefree part is necessary: for $Q = Y_1^2$ in characteristic zero, every constant solution is singular and $\operatorname{Res}_{Y_1}(Q, \partial Q/\partial Y_1) = 0$, while the squarefree part Y_1 has separant 1.

Taylor and image-degree bounds. Lemma A.5 bounds the degree of a rational image by the degrees of the map. We keep the Y_1-degree separate from the total jet degree, since it also bounds the exponent of Y_1 in Taylor reconstruction. To see why this helps, consider polynomials in Y_0, Y_1 of total degree at most B and Y_1-degree at most $C \le B$. Their monomials fill a triangle with one corner removed, and the exact count is

$$N_\triangle(B, C) = \sum_{v=0}^{C} (B - v + 1) = (C+1)(B+1) - \frac{C(C+1)}{2}.$$

The missing corner also reduces the degree of a family parameterized by these polynomials:

Lemma 5.3 (Rational images with a separate derivative-degree bound)**.** *Let $G \in E[Z, Y_0, Y_1]$ be nonzero, of challenge degree at most H, total jet degree at most j, and Y_1-degree at most $r \le j$. Consider a rational map on a union of irreducible components of $V(G)$. Suppose that its numerators and common denominator*

have challenge degree at most D_Z, total jet degree at most b, and Y_1-degree at most c, where $D_Z \geq 1$ and $1 \leq c \leq b$. Assume the denominator is nonzero on each selected component. Put

$$D_{\text{fiber}}(j, r; b, c) = jc + r(b - c).$$

The image closure has degree at most

$$H(2bc - c^2) + 2D_Z D_{\text{fiber}}(j, r; b, c).$$

After fixing a generic challenge, the image of the remaining curve components has degree at most $D_{\text{fiber}}(j, r; b, c)$.

Proof. Use the monomial embedding and Hilbert-function comparison from Lemma A.5. For an irreducible factor G_0 with actual degrees $(d_Z, d_{\text{jet}}, d_\partial)$, the argument changes only in the dimension of the truncated source space. With bounds $(D_Z N, bN, cN)$ it is $(D_Z N + 1)N_\triangle(bN, cN)$, so multiplication by G_0 leaves quotient dimension at most

$$(D_Z N + 1)N_\triangle(bN, cN) - (D_Z N - d_Z + 1)N_\triangle(bN - d_{\text{jet}}, cN - d_\partial).$$

Its quadratic coefficient is

$$\frac{d_Z(2bc - c^2)}{2} + D_Z\big(d_{\text{jet}}c + d_\partial(b - c)\big).$$

The argument of Lemma A.5 turns twice this coefficient into the image-degree bound, including with base points or positive-dimensional fibers. When $c = b$, the full triangle gives the same coefficient. At a fixed generic challenge, $N_\triangle$ alone has linear coefficient $d_{\text{jet}}c + d_\partial(b - c)$, giving the curve bound. Finally sum over the selected factors: their three degrees add to at most (H, j, r), and both bounds are linear in these source degrees. □

We apply this lemma to the maps given by Taylor reconstruction, whose degrees are bounded as follows.

Lemma 5.4 (Degrees of a regular first-order Taylor map)**.** *Let $D \geq 1$, $\ell \geq 1$, and $\tau = \max\{0, 2D - 3\}$. Suppose a first-order equation $Q \in F[X, Z, Y_0, Y_1]$ has total jet degree at most j, Y_1-degree at most $1 \leq r \leq j$, and challenge degree at most H. Fix a Taylor center $x_0 \in F$, and put $Q_{x_0} = Q(x_0, Z, Y_0, Y_1)$ and $q = (\partial Q/\partial Y_1)(x_0, Z, Y_0, Y_1)$. Assume characteristic zero or $p > D$. On every component of $V(Q_{x_0})$ where q is nonzero, Taylor reconstruction through message degree D, with the coordinates $(Z, Z^2, \ldots, Z^\ell, P)$ retained, has numerator and common-denominator degrees at most*

$$b = 1 + \tau(j - 1), \qquad c = \min\{b, \tau(r - 1) + D\}, \qquad D_Z = \ell + \tau H \tag{42}$$

in total jet degree, Y_1-degree, and challenge degree, respectively.

Proof. Before specializing $X = x_0$, differentiation along a solution of $Q = 0$ is

$$\partial_X + Y_1\partial_{Y_0} - \frac{Q_X + Y_1 Q_{Y_0}}{Q_{Y_1}}\partial_{Y_1},$$

and its iterates give the Taylor coefficients. For order $s \geq 2$, induction gives the common denominator q^{2s-3} after specialization and numerator total jet degree at most $1 + (2s - 3)(j - 1)$. The numerator at order two has Y_1-degree at most $r + 1$, and each further differentiation adds at most $2r - 1$, so at order s it is at most

$$(2s - 3)(r - 1) + s = (2s - 3)r - s + 3.$$

If $D = 1$, no coefficient beyond the initial pair is reconstructed: the Taylor map is the identity on (Y_0, Y_1), so $(b, c, D_Z) = (1, 1, \ell)$. For $D \geq 2$, the largest denominator exponent through degree D is $2D - 3 = \tau$. Passing to the common denominator q^τ through degree D gives Y_1-degree at most $\tau(r - 1) + D$; the total-degree induction gives $1 + \tau(j - 1)$. Clearing the common denominator adds at most τH in the challenge, while the retained powers of Z add at most ℓ. Division by $s!$ for $s \leq D$ requires only $p > D$. □

5.3 The counting bound

We now count both kinds of solutions together, for any first-order equation with total jet, derivative, and challenge degrees at most (B, M, H); the rate theorems use $(B_{\mathrm{jet}}, B_\partial, B_Z)$.

Let $1 \le D \le n-2$, $D+1 \le A \le n$, $B \ge 1$, $1 \le M \le B$, and $H \ge 0$. Fix distinct evaluation points $\alpha_1, \dots, \alpha_n \in F$ and a received curve $w(Z) = \sum_{t=0}^{\ell} Z^t w_t$, with $\ell \ge 1$ and $w_t \in F^n$. Put $\lambda_{\mathrm{agr}} = (n-D)/(A-D)$. The count uses the following degree bounds:

$$\tau = \max\{0, 2D-3\}, \qquad u = 1 + \tau(B-1), \qquad v = \min\{u, \tau(M-1) + D\},$$
$$F_{\mathrm{reg}} = Bv + M(u-v), \qquad S = \max\{B, (2M-1)B - M^2\}, \qquad H_s = (2M-1)H. \tag{43}$$

Here u and v are the bounds b and c of Lemma 5.4 at $(j, r) = (B, M)$, so they bound the total jet degree and Y_1-degree of the Taylor map. With these map degrees, $F_{\mathrm{reg}} = D_{\mathrm{fiber}}(B, M; u, v)$ is the bound of Lemma 5.3 on the degree of the regular family at a generic challenge. The same lemma, with challenge degree H for the equation and $\ell + \tau H$ for the map, bounds the joint degree of the regular family of pairs (z, P) for the received curve by

$$J_\ell = H(2uv - v^2) + 2(\ell + \tau H)F_{\mathrm{reg}}.$$

Finally, S and H_s bound the degrees in Y_0 and Z of the ordinary equation satisfied by the remaining solutions.

Lemma 5.5 (Squarefree first-order counting)**.** *With these parameters, let $Q \in F[X, Z, Y_0, Y_1]$ be nonzero, with total jet degree at most B, Y_1-degree at most M, and challenge degree at most H. Assume characteristic zero or $p > \max\{D, M\}$. For each integer $D+1 \le L \le A$, outside a set of at most*

$$\min_{D+1 \le L_0 \le A} \mathsf{E}^{(\ell)}_{\mathrm{ord}}(D, A, n; S, H_s; L_0) + \frac{n-L+1}{A-L+1}\lambda_{\mathrm{agr}} J_\ell + \ell(n-L)\frac{n-D}{L-D}F_{\mathrm{reg}} \tag{44}$$

challenges, every degree-$\le D$ solution of $Q(X, z, P, P^{[1]}) = 0$ with at least A agreements has the full common-witness property of Definition 3.1, with message dimension $D+1$. For a nonzero fixed-word equation with the same jet-degree bounds, the number of close solutions is at most

$$\lambda_{\mathrm{agr}} F_{\mathrm{reg}} + S \le 4\lambda_{\mathrm{agr}} DBM + 2BM + B. \tag{45}$$

The three terms of (44) match those of (32). The exceptional challenges of the ordinary equation, at their own threshold L_0, form the preliminary set. For the regular family, the last two terms count high-agreement points off the witness graphs with at least L common agreements, and accidental agreements on them.

Proof of Lemma 5.5. Work over $\overline{F}$. Write the squarefree part of Q as $Q_{\mathrm{ord}} Q_{\mathrm{diff}}$, where Q_{ord} contains every factor independent of Y_1 and Q_{diff} is the product of the distinct factors of positive Y_1-degree. Put $j = \deg_{Y_0, Y_1} Q_{\mathrm{diff}}$, $r = \deg_{Y_1} Q_{\mathrm{diff}}$, and $b_{\mathrm{ord}} = \deg_{Y_0} Q_{\mathrm{ord}}$. Then $b_{\mathrm{ord}} + j \le B$, $r \le M$, and $\deg_Z Q_{\mathrm{ord}} + \deg_Z Q_{\mathrm{diff}} \le H$. If $r = 0$, every solution solves the ordinary equation Q_{ord}, whose degrees in Y_0 and Z are at most (S, H_s). Assume $r > 0$. Because $p > M$ in positive characteristic, Q_{diff} is separable over $\overline{F}(X, Z, Y_0)$: each irreducible factor has positive degree less than p, and the factors are distinct.

First consider the solutions P for which $\partial_{Y_1} Q_{\mathrm{diff}}(X, z, P(X), P^{[1]}(X))$ is nonzero as a polynomial in X. Extend the coefficient field to the algebraic closure of $\overline{F}(x_0)$, where x_0 is transcendental over $\overline{F}$, and use x_0 as the Taylor center. Since x_0 is transcendental, each such nonzero polynomial in X remains nonzero at $X = x_0$. Apply Lemma 5.4 over this field to obtain one rational map on the regular components of $Q_{\mathrm{diff}} = 0$. Its numerator and denominator bounds are u in total jet degree, v in Y_1, and $\ell + \tau H$ in the challenge, retaining the coordinates $(Z, Z^2, \dots, Z^\ell, P)$. Applying Lemma 5.3 to these components gives reduced generic-fiber degree at most F_{reg} and joint degree at most J_ℓ. The joint image has dimension at most two, and its generic challenge fiber has dimension at most one.

We now apply the transfer argument. After the linear intersection with degree-$\le D$ coefficient space, Lemma 4.2 applies to the regular family with $k = D+1$, $r_1 = 1$, $J_1 = J_\ell$, and $F_1 = F_{\mathrm{reg}}$. Since $\Pi_1(A) = \lambda_{\mathrm{agr}}$ and $\Pi_1(L) = (n-D)/(L-D)$, it gives the last two terms of (44). For a fixed received word, apply the Taylor construction separately to its nonzero equation. Its image degree is at most F_{reg}, so Lemma 4.1 bounds its regular candidates by $\lambda_{\mathrm{agr}} F_{\mathrm{reg}}$; this needs no degree bound on the special fibers of the symbolic family.

For the remaining solutions, form the nonzero polynomial

$$R = \mathrm{Res}_{Y_1}(Q_{\mathrm{diff}}, \partial_{Y_1} Q_{\mathrm{diff}}).$$

Writing $Q_{\mathrm{diff}} = \sum_{i=0}^{r} a_i Y_1^i$, the bound on total jet degree gives $\deg_{Y_0} a_i \le j - i$. As a polynomial in the a_i, every monomial of this resultant has degree $2r-1$, and the indices of its factors sum to r^2; these identities follow from $\mathrm{Res}(Q_{\mathrm{diff}}, Q'_{\mathrm{diff}}) = a_r \,\mathrm{Disc}(Q_{\mathrm{diff}})$. Consequently

$$\deg_{Y_0} R \le (2r-1)j - r^2, \qquad \deg_Z R \le (2r-1)\deg_Z Q_{\mathrm{diff}}.$$

If both specialized equations vanish, so does their resultant, including a specialization that lowers the Y_1 degree: the Sylvester matrix is specialized at its original degrees. Every remaining solution thus solves $Q_{\mathrm{ord}} R$, and the degree bounds give

$$\deg_{Y_0}(Q_{\mathrm{ord}} R) \le b_{\mathrm{ord}} + (2r-1)j - r^2 \le (2r-1)B - r^2 \le S,$$
$$\deg_Z(Q_{\mathrm{ord}} R) \le \deg_Z Q_{\mathrm{ord}} + (2r-1)\deg_Z Q_{\mathrm{diff}} \le H_s.$$

Here $(2r-1)B - r^2$ increases for $1 \le r \le M \le B$. The ordinary transfer, Lemma 4.3, applied to $Q_{\mathrm{ord}} R$ with its own threshold L_0, contributes the first term of (44). Its content term includes the challenges at which $Q_{\mathrm{ord}} R$ vanishes identically. For a fixed word, the same construction over $\overline{F}(X)$ gives at most S ordinary roots. This proves the first list inequality; the second uses $u \le 2DB$, $v \le 2DM$, and $S \le B + 2BM$.

The witnesses from both transfers are defined over F, also when F is infinite. □

Remark 5.6 (Equations without the derivative variable)**.** If $M = 0$, then $Q \in F[X, Z, Y_0]$ is an ordinary equation with $\deg_{Y_0} Q \le B$ and $\deg_Z Q \le H$. For each $D+1 \le L \le A$, Lemma 4.3 gives the common-witness conclusion outside $\mathsf{E}^{(\ell)}_{\mathrm{ord}}(D, A, n; B, H; L)$ challenges, in every characteristic. A nonzero fixed-word equation has at most B roots over $F(X)$, so at most B close solutions.

5.4 Proofs of the two bounds

Proof of Theorem 5.1. Take $g = 0$ and $z = 0$ in Corollary 3.17. The resulting nonzero fixed-word explainer has degrees at most $(B_{\mathrm{jet}}, B_\partial)$. Since $1 \le D \le n-2$ and $A > D$, the counting hypotheses hold.

If $B_\partial = 0$, Remark 5.6 bounds the list by B_{jet}, which is the displayed bound. Otherwise apply Lemma 5.5 with $(B, M) = (B_{\mathrm{jet}}, B_\partial)$; its list bound (45) is the displayed bound. For the asymptotic form, use $B_{\mathrm{jet}}, B_\partial = O_\rho((\eta_1 + 1/n)^{-1})$ from Lemma 3.16, $D \le n$, and $\lambda_{\mathrm{agr}} \le 1/(a-\rho) = O_\rho(1)$. □

Proof of Theorem 5.2. As in the previous proof, Corollary 3.17 gives a nonzero symbolic explainer $Q \in F[X, Z, Y_0, Y_1]$ for the received line, of degrees at most $(B_{\mathrm{jet}}, B_\partial, B_Z)$, that every close candidate satisfies at its challenge. By Lemma 3.16, $B_Z = O_\rho((\eta_1 + 1/n)^{-2})$.

We first bound the ordinary term. Since $2B_\partial + 1 \ge \max\{1, 2B_\partial - 1\}$, we have $b \ge S$ and $h \ge H_s$, and in particular $b \ge B_{\mathrm{jet}}$ and $h \ge B_Z$. The bound $\mathsf{E}^{(1)}_{\mathrm{ord}}$ is nondecreasing in both degree arguments, and $\Psi_D(b) \le 1 + 4Db$. Choosing $L_0 = D+1$ therefore bounds $\mathsf{E}^{(1)}_{\mathrm{ord}}(D, A, n; S, H_s; L_0)$, and also $\mathsf{E}^{(1)}_{\mathrm{ord}}(D, A, n; B_{\mathrm{jet}}, B_Z; L_0)$, by the first line of (41). The factor $2B_\partial + 1$ rather than $2B_\partial - 1$ keeps this valid at $B_\partial = 0$.

If $B_\partial = 0$, Remark 5.6 with $L = D + 1$ gives the exceptional set, and the second line of (41) vanishes. Otherwise apply Lemma 5.5 with $(B, M, H) = (B_{\text{jet}}, B_\partial, B_Z)$, $\ell = 1$, and $L = D + \lceil (A - D)/2 \rceil$. Both incidence factors $(n - L + 1)/(A - L + 1)$ and $(n - D)/(L - D)$ are at most $2\lambda_{\text{agr}}$. Since $u \le 2DB$ and $v \le 2DM$, we have $F_{\text{reg}} \le 4DBM$ and $J_1 \le 8DBM + 24D^2HBM$. Substituting these bounds in (44), with $n - L \le n - D - 1$, gives the second line.

To obtain the gap exponent, write $\bar{\eta}_1 = \eta_1 + 1/n$. The regular terms contribute $O_\rho(n^2\bar{\eta}_1^{-4})$. The ordinary terms have $b = O_\rho(\bar{\eta}_1^{-2})$ and $h = O_\rho(\bar{\eta}_1^{-3})$, so they contribute $O_\rho(n\bar{\eta}_1^{-5})$. Since $n\bar{\eta}_1 \ge 1$, we have $n\bar{\eta}_1^{-5} \le n^2\bar{\eta}_1^{-4}$, proving the fourth-power bound even when $n\eta_1 < 1$. □

5.5 Uniform bounds at gap $6/25$

The constants in Theorems 5.1 and 5.2 depend on the rate and on the gap η_1 above the first-order curve. The uniform capacity theorem of Section 8 needs, for its large-gap branch, constants that depend on neither. Here the gap is measured from capacity instead: the threshold satisfies $A \ge k + \delta n$ with $\delta \ge 6/25$. Two fixed interpolation supports, one for lists and one for MCA, give such constants (Section C.3).

Theorem 5.7 (Uniform first-order bounds at gap $6/25$ from capacity)**.** *Let $\delta \ge 6/25$. Let $n \ge 2$, $1 \le k \le n$, and let F be a field. If $k \ge 2$, assume that its characteristic is zero or greater than $k - 1$; no restriction is needed for $k = 1$. Fix distinct $\alpha_1, \dots, \alpha_n \in F$ and an integer threshold $k + \delta n \le A \le n$. Then the following hold.*

(i) *Every $y \in F^n$ has at most*

$$307n \tag{46}$$

degree-$< k$ polynomials agreeing with it in at least A positions.

(ii) *For every $f, g \in F^n$, there is a set $\mathcal{E} \subseteq F$ of size at most $1325775n^2$ such that every $z \notin \mathcal{E}$ and every degree-$< k$ polynomial P with $|S_z(P)| \ge A$ admit degree-$< k$ polynomials F_0, G_0 satisfying*

$$P = F_0 + zG_0, \qquad S_z(P) = S(F_0, G_0).$$

For $F = \mathbb{F}_q$, the exceptional probability is at most $1325775n^2/q$ for lines.

At gap $6/25$, the agreement fraction $\rho + 6/25$ lies below the asymptotic Johnson threshold $\sqrt{\rho}$ exactly when $0.16 < \rho < 0.36$; at rate $1/4$, the theorem needs $\lceil 0.49n \rceil$ agreements, against the limiting Johnson threshold $0.50n$. Against the finite Johnson threshold $\sqrt{(k-1)/n}$, the comparison holds on each closed subinterval of this window for sufficiently large n. For a particular code, Section C gives sharper parameters from integer inequalities at its own (n, k, A), and a second counting bound.

6 Larger families of received words

Having proved MCA for a received line in Sections 4 and 5, we now extend it to the larger families that the proof systems of Section 7 combine: affine combinations with independent challenges, interleaved words, and multilinear folds.

The existing reductions for affine families and interleaved words (Sections 6.1 and 6.2) are stated for *subset MCA*: if the combined word agrees with a codeword on a set of at least A positions, so does each coefficient word on that set. The folds of Section 6.3 need the full agreement-set form of Definition 3.2, because each level must pass the full common agreement set to the next, and interpolation recovers this form whenever $A \ge k$. At a good challenge, let S be a candidate's full agreement set, with $|S| \ge A \ge k$. Subset MCA gives degree-$< k$ witnesses for every coefficient word on S. Their prescribed combination agrees with

the candidate on S, so it equals the candidate by interpolation on k positions. Their common agreement set therefore contains S, and the identity gives the reverse inclusion. The same argument applies row by row to interleaved words.

Remark 6.1 (Polynomial curves)**.** A curve $w_0 + zw_1 + \cdots + z^\ell w_\ell$ is not a special case of the affine families below: an exceptional set that is small in $\mathbb{F}_q^\ell$ can contain the whole moment curve $\{(z, z^2, \ldots, z^\ell)\}$. We therefore bound such curves directly, keeping the powers of Z in the symbolic explainer, whose challenge degree grows by the factor ℓ in (18). The resulting bounds are Lemma 4.3 at order zero, in every characteristic, Lemma 5.5 and Proposition C.2 at first order, and Corollary 8.6 at general order. Corollary 6.6 transfers each of them to interleaved codes.

6.1 Affine families

The line-to-affine reduction loses only the factor $1/(1-1/q)$, not a factor for each independent coefficient. We give BCGM's argument for a code $C \subseteq W^n$ over a vector space W, with agreement on whole W-coordinates, so that it also applies to interleaved codes.

Lemma 6.2 (Line-to-affine MCA transfer [BCGM25, Lemma 7.1])**.** *Let W be a finite-dimensional $\mathbb{F}_q$-vector space, let $C \subseteq W^n$ be an $\mathbb{F}_q$-subspace, and let $1 \le A \le n$. Suppose that for every $f, g \in W^n$, at most an ε fraction of $z \in \mathbb{F}_q$ admit a set $S \subseteq [n]$ with $|S| \ge A$ such that*

$$(f + zg)|_S \in C|_S, \qquad f|_S \notin C|_S \text{ or } g|_S \notin C|_S.$$

For every $r \ge 1$ and $a, u_1, \ldots, u_r \in W^n$, the fraction of $t \in \mathbb{F}_q^r$ admitting such an S with

$$\left(a + \sum_{j=1}^{r} t_j u_j\right)|_S \in C|_S, \qquad a|_S \notin C|_S \text{ or some } u_j|_S \notin C|_S$$

is at most $\varepsilon/(1-1/q)$.

Proof. Let B be the bad parameters and choose a witnessing set S_t for each. In $W^{S_t}/C|_{S_t}$, some $[u_j]$ is nonzero: otherwise $[a] + \sum_j t_j[u_j] = 0$ would force $[a] = 0$, contradicting badness. Thus the map $\lambda \mapsto \sum_j \lambda_j[u_j]$ is nonzero, and its kernel has at most q^{r-1} elements. Averaging gives a direction λ for which the set B' of bad parameters with $(\sum_j \lambda_j u_j)|_{S_t} \notin C|_{S_t}$ has size at least $(1-1/q)|B|$. On every line $v + z\lambda$, each point of B' is bad for the received line with direction $\sum_j \lambda_j u_j$. The line hypothesis gives $\Pr_z[v + z\lambda \in B'] \le \varepsilon$. Averaging over uniform v yields $|B|/q^r \le |B'|/(q^r(1-1/q)) \le \varepsilon/(1-1/q)$. □

Corollary 6.3 (Mutual agreement for affine families)**.** *Suppose the degree-$<k$ Reed–Solomon code over $\mathbb{F}_q$ has full agreement-set line MCA at threshold $A \ge k$ with exceptional count $\mathsf{E}_{\mathrm{exc}}$. Let $s \ge 1$, and let $a, u_1, \ldots, u_s \in \mathbb{F}_q^n$. For $w(t) = a + \sum_{j=1}^s t_j u_j$, there is $\mathcal{E}_s \subseteq \mathbb{F}_q^s$ of size at most $\mathsf{E}_{\mathrm{exc}} q^s/(q-1)$ such that the following holds. For every $t \notin \mathcal{E}_s$ and every degree-$<k$ polynomial P with full agreement set*

$$S = \{i : P(\alpha_i) = w(t)_i\}, \qquad |S| \ge A,$$

there are degree-$<k$ polynomials $F_0, \ldots, F_s \in \mathbb{F}_q[X]$ with

$$P = F_0 + \sum_{j=1}^{s} t_j F_j, \qquad S = \{i : F_0(\alpha_i) = a_i \text{ and } F_j(\alpha_i) = u_{j,i} \text{ for every } j \ge 1\}.$$

In particular, unless common witnesses on at least A positions exist, the probability that $w(t)$ has A agreements with a codeword is at most $\mathsf{E}_{\mathrm{exc}}/(q-1)$ for uniform t, independently of s. For $s = 1$, the sharper bound $\mathsf{E}_{\mathrm{exc}}/q$ remains valid.

Proof. Take C to be the degree-$< k$ Reed–Solomon code on the given points, with $W = \mathbb{F}_q$. The line bound implies the hypothesis of Lemma 6.2 with $\varepsilon = \mathsf{E}_{\mathrm{exc}}/q$: any set S witnessing an agreement extends to the full agreement set of a polynomial, and the line bound gives witnesses on that full set, hence on S. The lemma therefore bounds the bad affine parameters by $\mathsf{E}_{\mathrm{exc}} q^s/(q-1)$. Outside them, the observation opening this section gives the polynomial identity and equality of the full sets. □

Theorems 4.4, 5.2, 5.7 and 8.1 give such a bound $\mathsf{E}_{\mathrm{exc}}$, each under its own hypotheses.

6.2 Interleaved codes

An interleaved word stacks s received words as rows, and we ask whether the scalar bounds still hold when agreement must occur in all rows at the same positions. For a code $C \subseteq F^n$, its width-s interleaving is

$$C^{\equiv s} := \{M \in F^{s\times n} : M_{j,*} \in C \text{ for every } j \in [s]\},$$

and agreement at a column means agreement in every row. For lists the answer needs no new counting: encoding each column as one element of an extension field preserves equality, so any scalar list bound valid over that extension transfers.

Corollary 6.4 (Interleaved lists)**.** *Let F be a field with distinct evaluation points $\alpha_1, \dots, \alpha_n$, and fix a dimension k, an integer threshold A, and a width $s \geq 1$. Suppose an extension E/F contains s elements linearly independent over F, and every scalar received word over E has at most B degree-$< k$ candidates at agreement threshold A. Then the interleaved list over F has at most B candidates at column agreement threshold A. In particular, a scalar bound valid over every extension of F transfers at every width. This applies to the first-order list bound: under the hypotheses of Lemma C.3, its scalar bound holds at every interleaving width, with agreement measured in columns, and for $B_\partial = 0$ it gives the bound B_{jet} in every characteristic.*

Proof. Choose linearly independent elements $\beta_1, \dots, \beta_s \in E$. The map $(P_1, \dots, P_s) \mapsto \sum_j \beta_j P_j$ is injective on degree-$< k$ row tuples. Encode each received column with the same elements. Linear independence makes equality in E equivalent to equality in every row, so this maps the interleaved list injectively into the scalar list. Apply its bound B. For arbitrary F, the extension $F(\xi)$ with a transcendental ξ contains the independent elements $1, \xi, \dots, \xi^{s-1}$. For finite $F = \mathbb{F}_q$, a basis of $E = \mathbb{F}_{q^s}$ gives the usual bijective encoding [ABF26, Lemma 2.21]; only injectivity is needed here. The first-order bound of Lemma C.3 holds over every extension of F, since extension preserves the characteristic. □

MCA for interleaved curves. For MCA, Jo's projection argument avoids a separate exceptional set for each row [Jo26a, Theorem 4.4 and Corollaries 4.5–4.6]: one scalar combination of the rows keeps every bad challenge bad, also for a polynomial received curve.

Lemma 6.5 (Interleaving preserves polynomial-family MCA)**.** *Let $C \subseteq \mathbb{F}_q^n$ be a linear code, $1 \leq A \leq n$, and $\ell \geq 1$. Suppose that, for every $w_0, \dots, w_\ell \in \mathbb{F}_q^n$, at most E_ℓ challenges z admit a set S of size at least A such that $(\sum_{t=0}^{\ell} z^t w_t)|_S \in C|_S$ but some $w_t|_S \notin C|_S$. Then the same bound holds for $C^{\equiv s}$ at every width $s \geq 1$, with agreement measured in columns.*

Proof. At each bad challenge z, choose a witnessing set S_z and let $u_{t,i}$ be the class of coefficient row i in $\mathbb{F}_q^{S_z}/C|_{S_z}$. For each row, $\sum_t z^t u_{t,i} = 0$, while some $u_{t,i}$ is nonzero. Thus the projections $\lambda \in \mathbb{F}_q^s$ satisfying $\sum_i \lambda_i u_{t,i} = 0$ for every t form a proper subspace K_z. At most q proper subspaces have union of size at most $1 + q(q^{s-1} - 1) < q^s$. One projection therefore avoids every K_z. It commutes with the polynomial combination and keeps every bad challenge bad for a single scalar received curve. Its bound E_ℓ proves the claim. □

Corollary 6.6 (MCA for interleaved Reed–Solomon codes)**.** *Suppose every scalar received polynomial curve of degree at most $\ell \geq 1$ has full agreement-set MCA at threshold $A \geq k$ with exceptional-count bound E_ℓ. Then the same bound holds for interleaved received curves of every width, with degree-$< k$ polynomial witnesses for all $\ell + 1$ coefficient arrays and equality of full common column-agreement sets. The exceptional probability is at most E_ℓ/q, independently of width. In particular, a scalar line bound $\mathsf{E}_{\mathrm{exc}}$ gives interleaved line error $\mathsf{E}_{\mathrm{exc}}/q$ and affine-family error $\mathsf{E}_{\mathrm{exc}}/(q-1)$, independently of both width and the number of affine parameters.*

Proof. Apply Lemma 6.5 to the subset MCA implied by the scalar bound. The observation opening Section 6, applied row by row, gives the polynomial identity and equality of full column-agreement sets. Lines use $\ell = 1$. For affine families, apply Lemma 6.2 with $W = \mathbb{F}_q^s$ to the interleaved code $C^{\equiv s} \subseteq W^n$, whose line hypothesis is the case $\ell = 1$ just proved, and recover full sets as before. □

6.3 Multilinear folds

A binary fold combines 2^h words in pairs, with a fresh challenge at each level, until one word remains. The pairs at one level share its challenge and form one interleaved line, so Corollary 6.6 bounds the exceptional challenges level by level. Because it preserves the full common agreement set, the levels compose.

This composition is an instance of a general principle of Bordage et al.: the tensor product of two generators with MCA has MCA, and the errors add [BCGM25, Lemma 4.4]. Writing $\ell_0(r) = 1 - r$ and $\ell_1(r) = r$, the fold coefficients $(\prod_{j=1}^h \ell_{\sigma_j}(r_j))_{\sigma\in\{0,1\}^h}$ form the h-fold tensor power of $r \mapsto (1-r, r)$, which is an invertible linear image of the line generator $r \mapsto (1, r)$. Bordage et al. obtain in this way MCA for the multilinear equality generator over every linear code, with error h times the line error [BCGM25, Corollary 8.6]. This improves the error of Baweja, Mishra, Mopuri, and Shtepel [BMMS25], who prove MCA for this generator almost up to the one-and-a-half Johnson bound.

Corollary 6.7 applies the same composition to the interleaved Reed–Solomon line bound of Corollary 6.6, which our theorems prove beyond the Johnson radius, and states the conclusion in the full agreement-set form.

Write $P|_\alpha = (P(\alpha_i))_{i=1}^n$ for a polynomial's evaluation word and Agr for the agreement set of Section 3.

Corollary 6.7 (MCA for multilinear folds)**.** *Let the degree-$< k$ Reed–Solomon code over $\mathbb{F}_q$ satisfy the full agreement-set line-MCA conclusion at threshold $A \geq k$ with exceptional-count bound $\mathsf{E}_{\mathrm{exc}}$, as in Corollary 6.6. Fix $h \geq 1$ and words $(u_\sigma)_{\sigma\in\{0,1\}^h}$. Sample independent uniform challenges $r_1, \ldots, r_h \in \mathbb{F}_q$, and define the folded words from the leaves upward by $W_\sigma = u_\sigma$ for $|\sigma| = h$ and*

$$W_\tau = (1 - r_{j+1})W_{\tau 0} + r_{j+1}W_{\tau 1} \qquad (|\tau| = j < h). \tag{47}$$

Outside at most $h\mathsf{E}_{\mathrm{exc}}q^{h-1}$ challenge tuples, every degree-$< k$ polynomial P agreeing with $W_\varnothing$ on at least A positions admits degree-$< k$ polynomials $(P_\sigma)_\sigma$ such that, writing $\ell_0(r) = 1 - r$ and $\ell_1(r) = r$,

$$P = \sum_{\sigma\in\{0,1\}^h} \left(\prod_{j=1}^h \ell_{\sigma_j}(r_j)\right) P_\sigma, \tag{48}$$

$$\mathrm{Agr}(P|_\alpha, W_\varnothing) = \bigcap_{\sigma\in\{0,1\}^h} \mathrm{Agr}(P_\sigma|_\alpha, u_\sigma). \tag{49}$$

The exceptional probability is therefore at most $h\mathsf{E}_{\mathrm{exc}}/q$. The same bound holds when the input words are themselves interleaved, independently of their width.

Proof. At depth j, the arrays $U_j = (W_{\tau 0})_{|\tau|=j}$ and $V_j = (W_{\tau 1} - W_{\tau 0})_{|\tau|=j}$ depend only on $r_{j+2}, \ldots, r_h$. Conditional on these challenges, r_{j+1} is uniform, so Corollary 6.6 gives one exceptional set of size at most

E_{exc} for $U_j + r_{j+1}V_j$. This set works simultaneously for every tuple of polynomial candidates. A union bound over the h levels gives

$$\Pr[\text{some level is exceptional}] \le h\mathsf{E}_{\text{exc}}/q. \tag{50}$$

Fix a challenge tuple outside these events and a candidate P for the folded word $W_\varnothing$. Let S be its full agreement set. Inductively maintain a tuple $(P_\tau)_{|\tau|=j}$ whose full common column-agreement set is S. Since $|S| \ge A$, interleaved MCA splits this entire tuple into its children and preserves that common set. Rewriting its witnesses for U_j, V_j as witnesses for the two children gives $P_\tau = (1 - r_{j+1})P_{\tau 0} + r_{j+1}P_{\tau 1}$. At the leaves these identities compose to (48), and the invariant gives (49). Individual rows may agree at additional positions; only their common set is preserved.

Each uniform challenge tuple has probability q^{-h}, giving the stated cardinality bound. For interleaved input words, include their original row index in the arrays at each level.

For $h = 0$ there is one word and no challenge, and the conclusion holds with $P_\varnothing = P$ and no exceptions, so the bound $h\mathsf{E}_{\text{exc}}/q$ holds for every $h \ge 0$. □

7 First-order bounds and smaller proofs

In the proof systems we consider, t queries contribute an error term of about $(A/n)^t$, where A is the agreement at which the list and MCA bounds are applied. Our finite first-order bounds apply at agreements A below the Johnson threshold, so fewer queries suffice, provided the algebraic terms, which grow as A decreases, still meet their error targets. Each removed query also removes its field values and authentication path, so the measured proofs shrink at unchanged security targets, and our query counts are below the Johnson floor (Table 4).

We use first order because it keeps the list and exceptional counts linear and quadratic in n; at second order they already grow to $O(n^2)$ and $O(n^3)$ (Theorems 8.4 and 8.5), and the resulting Λ/q and $\mathsf{E}_{\text{exc}}/q$ terms can cancel the query saving. The first-order bounds require the characteristic to exceed $k - 1$ and B_∂; the characteristics of BN254 and Goldilocks, near 2^{254} and 2^{64}, exceed every such value below [Wor26a, Pro26b, Lam26c].

7.1 From certificates to error bounds

For each queried code (n, k, A), our artifact searches for a support satisfying Proposition C.2 and evaluates the applicable counting bounds, in exact arithmetic [DKT26a]. A support passing these checks is a *certificate* for the code; it yields a list bound Λ and exceptional counts E_{exc} for lines and E_ℓ for curves of degree ℓ. The list and MCA bounds may use different supports, each under its own field hypotheses. We then recheck every affected error term at the actual field size.

Lowering the agreement changes error terms of three kinds. Arnon, Boneh, and Fenzi isolate the same three in their knowledge-soundness error $\epsilon_{\text{mca}} + \Lambda/q + a^t$ [ABF26, Construction 6.2 and Lemmas 6.6, 6.8]; we nevertheless recheck each system's own analysis term by term. The MCA term depends on how words are combined. Combining them along a line, by powers of one challenge, in an affine family, or by an h-level fold contributes $\mathsf{E}_{\text{exc}}/q$, E_ℓ/q, $\mathsf{E}_{\text{exc}}/(q-1)$, or $h\mathsf{E}_{\text{exc}}/q$, respectively (Section 6). Collision and identity terms depend on Λ, and the query term depends on grinding: a proof of work accepted with probability p_{grind} makes it $p_{\text{grind}}(A/n)^t$.

We choose the agreement and query count together. Writing ϵ_{alg} for the sum of the MCA term and the terms depending on Λ, we require $\epsilon_{\text{alg}} + p_{\text{grind}}(A/n)^t \le \epsilon$ when these terms share one error target ϵ; phases with separate targets are checked separately, without bounding their sum. Every comparison keeps each component's error target, its rule for combining errors, and its grinding convention. ZisK and LambdaVM keep their powers batching, including LambdaVM's optimized sampling [Lam26b], and the grinding factor

System and field	First code (n, k, A)	Queries	Proof bytes, before → after	Saved
ProveKit passport, BN254	$(2^{18}, 2^{16}, 123208)$	127/119/109	716 276.5 → 636 865.6	79.4 KB, 11.1%
ProveKit lookup, $\mathbb{F}_{p^3}$	$(2^{12}, 2^{10}, 1933)$	127/–/109	161 184.8 → 144 508.4	16.7 KB, 10.3%
ZisK final proof, $\mathbb{F}_{p^3}$	$(2^{19}, 2^{15}, 124136)$	54/53/51	286 013 → 272 783	13.2 KB, 4.63%
LambdaVM CPU, $\mathbb{F}_{p^3}$	$(2^{16}, 2^{15}, 45690)$	219/217/208	1 233 024 → 1 173 192	59.8 KB, 4.85%

Table 4: Query counts and measured proof sizes against the deployed parameters, with unchanged security targets and grinding. Here $\mathbb{F}_{p^3}$ is cubic Goldilocks, $p = 2^{64} - 2^{32} + 1$, and (n, k, A) gives the length, dimension, and our agreement threshold for the first queried code. Queries are listed as shipped/Johnson floor/ours, where the Johnson floor is the fewest queries whose query term alone meets the target at agreement $A_J = \lceil \sqrt{n(k-1)} \rceil$; ProveKit's counts are for the first round, and Table 7 gives the later ones. ProveKit's floor is attained using Theorem 4.4, while the ZisK and LambdaVM floors are lower bounds for any Johnson-based analysis (Sections 7.3 and 7.4). We did not compute a Johnson floor for the lookup proof. Sizes are measured against the shipped configuration: ProveKit means of complete compressed proof files over ten (passport) and thirty (lookup) proofs per configuration; one ZisK compressed final proof in native format, before any PLONK wrapping; and one serialized LambdaVM CPU subproof per configuration. One KB is 1000 bytes.

p_{grind} does not multiply their curve-MCA error term. LambdaVM's bound covers only the changed CPU component.

7.2 ProveKit: passport and lookup proofs

In ProveKit we lower the first-round query count of both a passport proof and a lookup proof from 127 to 109; the passport proof shrinks by 79.4 KB (11.1%) and the lookup proof by 16.7 KB (10.3%) (Table 4). ProveKit proves witness constraints with Spartan [Set20] and WHIR [ACFY24b], and World Foundation reports its integration into World ID, describing the passport benchmark as checking signatures, expiry, and minimum age without revealing the passport [Wor26d, Wor26c]. The two benchmarks use different commitments. We choose new agreements and query counts for the zkWHIR witness commitments of this public v1 benchmark, zkWHIR being the hiding variant of WHIR, and for the non-hiding WHIR commitments of a lookup benchmark [Wor26e, Wor26b], keeping each changed error term within its 2^{-128} target. Only the lookup proof adds an out-of-domain evaluation, to the first code of each witness commitment, 48 bytes in total.

The rechecked error terms. Consider a WHIR code with n rows and degree bound $k - 1$. Its rows contain eight field elements, representing a vector of length $8k$. At agreement $a = A/n$, our certificate gives a list bound Λ and a line exceptional count E_{exc}, and the interleaved list has the same bound as the scalar list (Corollary 6.4). WHIR combines the eight columns with the equality-polynomial weights of three random challenges, which are the weights of three successive binary folds, so Corollary 6.7 bounds their exceptional probability by $3\mathsf{E}_{\text{exc}}/q$. The folding and identity checks, and the test distinguishing candidate vectors by u out-of-domain evaluations, therefore satisfy

$$\epsilon_{\text{fold}} \le \frac{3\mathsf{E}_{\text{exc}} + 2\Lambda}{q}, \qquad \epsilon_{\text{OOD}} \le \binom{\Lambda}{2} \left(\frac{8k-1}{q} \right)^{u}. \tag{51}$$

The second bound holds because two distinct candidates differ by a nonzero polynomial of degree at most $8k - 1$; increasing u suppresses this term when the list is too large for one evaluation. The transition to the next code uses both the source query count t and the target list bound Λ': with u' out-of-domain evaluations

in the target, its constraint-combination term is at most

$$\frac{2(t+u')\Lambda'}{q}. \tag{52}$$

We recompute this term whenever either code changes, with the new target list rather than the Johnson list of the original parameters. The remaining vector and linear-form combination terms keep their original field-size bounds [Wor26e, Wor26f]. Each code keeps its grinding threshold τ: the verifier accepts a 64-bit hash at most τ, with probability $(\tau+1)/2^{64}$, so the query check is

$$\frac{\tau+1}{2^{64}}\left(\frac{A}{n}\right)^{t} \le 2^{-128}. \tag{53}$$

These probabilities use the uniform-field challenge model; the artifact bounds the implementation's field-sampling error, which is negligible compared with these targets [DKT26a].

The passport proof. Table 7 gives the first and later codes of both witness commitments; every passport code keeps one out-of-domain evaluation. An optimized Johnson analysis would recover part of our saving. At the same grinding probability and error target, even the finite Johnson threshold $\sqrt{(k-1)/n} = \sqrt{65535/262144}$ requires 119 queries in (53), and 119 is achievable: at agreement $1/2$, the finite Johnson bounds of Theorem 4.4 satisfy the first round's list, MCA, and out-of-domain checks. Our 109 queries therefore save ten first-round queries, or 8.4%, beyond this optimized Johnson count. The first outer round accounts for about 21.6 KB of the expected raw saving, of which about 12.1 KB lies beyond the optimized Johnson count. The later rounds and internal proofs account for the rest; we have not optimized Johnson query counts for them.

Each witness opening also contains an internal zkWHIR proof about its masking polynomials, whose parameters we also change. Its initial commitment batches two vectors and so requires a separate line-MCA bound $\mathsf{E}_{\mathrm{exc}}/q$ before the multilinear-fold bound $3\mathsf{E}_{\mathrm{exc}}/q$. All changed checks meet their existing error targets, including the grouped initial batching error $(\mathsf{E}_{\mathrm{exc}}+1)/q$; the masking randomness, privacy parameters, and out-of-domain counts are unchanged. Random query positions share different parts of their Merkle paths and compression adds variation, so we compare means over ten verified proofs per configuration on one input.

The lookup proof over cubic Goldilocks. The lookup benchmark shows where the separate derivative-degree bound B_∂ matters. It uses ProveKit's non-hiding backend [Wor26b, Wor26a] to prove by LogUp that 4000 values belong to a 2000-entry table; because lookup challenges make some witness values extension-field elements, those values take 24 bytes each. Beyond the first code, we also change the later codes and the separate blinding commitment (Table 7). In the second code, $(n,k,A) = (2048,128,453)$, the squarefree bound gives $\Lambda \le 8157620$, so that with one out-of-domain evaluation ϵ_{OOD} has 137.08 bits of security, whereas a bound through total jet degree alone with the same support gives only 125.16 bits, short of the 128-bit target. The artifact records the alternative list bounds [DKT26a].

7.3 ZisK: a smaller compressed final proof

ZisK's compressed final proof attests to a recursively aggregated zkVM execution [Zis26b].[3] We lower its query count from 54 to 51, shrinking its final proof by 13.2 KB (4.63%) (Table 4), and keep the released code, batching, and grinding parameters [Zis26c]. The change affects only this final proof, not the aggregated proofs before it.

[3] Its repository describes ongoing security and correctness audits toward production [Zis26a].

The existing batching and transcript. ZisK batches in two levels: it combines the words opened at a common point by powers of one challenge, and then combines these groups by powers of another. After absorbing the fixed opening-point weights into the words, one interleaved curve bound applies to all inner groups at once, and, conditional on the inner challenge, a scalar curve bound applies to the outer combination. The inner guarantee is uniform over candidate tuples, including those reconstructed using the outer challenge.

These curve bounds enter at the degrees of the two combinations: the nontrivial inner curves have degrees $1, 2, 103, 24$, and the outer curve has degree 4. At agreement $A/n = 124136/2^{19}$, the support $(m, B_\partial, B_{\text{jet}}) = (10, 6, 37)$ gives bounds that meet the targets for both batching levels and the three FRI folds. To treat the inner groups as one interleaved curve, we pad their coefficient arrays with zeros to degree 103 and apply Corollary 6.6. The witnesses for the padded coefficients are degree-$< k$ polynomials vanishing on the common agreement set; since $A \geq k$ they are zero, leaving witnesses for exactly the original words. A union bound over the inner and outer challenges gives batching probability $(\mathsf{E}_{103} + \mathsf{E}_4)/q < 2^{-130.17}$, and every fold has more than 140.12 bits; the grinding factor applies to neither.

The query term carries ZisK's grinding factor 2^{-22}. Under the protocol's computational grinding convention, with hash and transcript outputs modeled as independent uniform field elements, the query term is $2^{-22}(124136/2^{19})^{51} < 2^{-128}$. The implemented sampler allows repeated queries and slices query indices from the low 63 bits of 16 successive Goldilocks outputs [Pro26a], and its grinding predicate accepts 2^{42} of the p possible outputs. Bounding the sampling bias per output by $2^{64}/p$ gives the conservative query term

$$(124136/2^{19})^{51}\frac{2^{42}}{p}\left(\frac{2^{64}}{p}\right)^{16} < 2^{-128}, \qquad p = 2^{64} - 2^{32} + 1.$$

Each batching, folding, and query phase therefore meets its separate 2^{-128} target.

Below the finite Johnson threshold. ZisK's implemented Johnson analysis uses agreement $\sqrt{\rho} + \eta_0$, with $\rho = 1/16$ and a positive gap η_0, and selects 54 queries [Zis26c]. Even the sharper finite-length threshold $a_J = \sqrt{(k-1)/n}$ cannot reach 52: the query term alone is

$$2^{-22}a_J^{52} = 2^{-22}\left(\frac{2^{15}-1}{2^{19}}\right)^{26} > 2^{-128},$$

so this threshold requires at least 53 queries with the same grinding. Our certificate applies at a smaller agreement, where 51 queries suffice and the algebraic terms also meet their targets.

Measured size. We generate and verify a compressed final proof for a 679-step execution with the released prover. The transcript's queries form a prefix of the original sequence, so keeping the first 51 query openings gives a proof accepted by the matching 51-query verifier. The 13230-byte saving in Table 4 preserves all other contents, including public values. Both proofs, pinned configurations, and the procedure are in the artifact.

7.4 LambdaVM: smaller CPU subproofs

In LambdaVM we change one subproof of a modular proof, the CPU subproof, and leave the others as deployed. Its query count drops from 219 to 208 and its serialized size by 59.8 KB (4.85%), net of two evaluations that we add before LogUp to fix the CPU trace before the lookup challenges (Table 4) [Dao26b]. Unlike the ProveKit and ZisK changes, which adjust parameters, this adds a step to the protocol. LambdaVM is an experimental Rust RISC-V zkVM developed with LambdaClass and 3MI Labs [Lam26a], whose CPU, arithmetic, memory, and coprocessor subproofs are linked by LogUp [Lam26d]; the CPU constrains control flow and connects instruction decoding, arithmetic, and memory requests to those tables. Its default cubic-Goldilocks profile has half-rate codes, 20 grinding bits, 219 queries, and a 128-bit target [Lam26c, Lam26d]. We keep the field, rate, powers batching, grinding, and final codeword length 256.

Fixing the CPU trace before lookup. Below Johnson, one commitment can be close to several candidate traces, and the shared LogUp challenges must not help the prover choose among them. After the CPU main commitment and before LogUp, the verifier samples distinct points s_1, s_2 outside the CPU LDE domain and receives the values of all $w = 38$ main columns there: $2w$ extension-field elements, or 1856 serialized bytes. The later out-of-domain point z keeps the usual polynomial-identity check, but each main column's DEEP quotient becomes $(f - I_f)/((X - z)(X - s_1)(X - s_2))$, where I_f interpolates the three claimed values. Successful reconstruction then produces a polynomial of degree at most $T + 2$, where $T = 32768$ is the trace length, that agrees with the committed word and takes both early values. The list we need is therefore for dimension $T + 3$, and our certificate bounds it by $\Lambda = 113769033$, its rational list bound $367473974837/3230$ rounded up. Lemma F.1 shows that the early values then select at most one candidate, uniformly over the prover's claimed values, except with probability

$$\epsilon_{\text{early}} = \binom{\Lambda}{2}\left(\frac{T+2}{q-n-1}\right)^2 \approx 2^{-301.5}.$$

One point would not suffice: its collision bound $\binom{\Lambda}{2}(T + 2)/(q - n - 1) \approx 2^{-124.5}$ exceeds the target.

The local error bound. With the trace fixed before lookup, it remains to show that 208 queries meet the local 128-bit target. The CPU certificate uses agreement $a = A/n = 45690/65536$, support $(m, B_\partial, B_{\text{jet}}) = (32, 9, 44)$, and the squarefree curve bound (44), at dimension T for the MCA terms and, because of the two early evaluations, $T + 3$ for the list. Powers batching combines 51 DEEP terms, so the curve bound enters at degree 50. Let ϵ_{MCA} be the sum of this curve term and the fold terms, and let $q = p^3$. In the independent uniform transcript-output model with the existing computational 20-bit grinding convention [Lam26b], the $t = 208$ queries give the local bound

$$\epsilon_{\text{CPU}} \le \epsilon_{\text{MCA}} + 2^{-20}a^t + \frac{49\Lambda}{q} + \frac{(4T+6)\Lambda + 2n + 2}{q - n - T} + \epsilon_{\text{early}}. \tag{54}$$

Its five terms cover exceptional challenges, acceptance by all query tests, cancellation among the CPU constraints, failure of the out-of-domain identity test, and collision at the two early points; only the query term is multiplied by the grinding factor 2^{-20}. Section F.2 proves the lemma and derives each term. The bound has 128.063 bits of security (the negative base-two logarithm, rounded to three decimals), so the CPU subproof keeps its 128-bit local target. The reduction requires agreement below the finite Johnson threshold: every integer agreement $A \ge \sqrt{65536 \cdot 32767}$, the least being $A_J = 46341$, needs at least 217 queries with the same grinding, even from the query term alone.

Measured size and cost. We compare the two profiles on the same assembly benchmark with empty input. It executes 32768 instructions and produces one 32768-row CPU chunk; both complete proofs contain 26 table subproofs, and every non-CPU table keeps 219 queries. The release CLI of our fork generates, deserializes, and verifies both proofs under their profiles [Dao26b]. The sizes in Table 4 serialize the complete CPU subproof with `rkyv`, including its commitments, claims, and openings but excluding shared VM metadata. Each removed query saves 5608 serialized bytes and the early claim vectors add 1856, so the saving is $11 \cdot 5608 - 1856 = 59\,832$ bytes. The complete VM proof decreases by the same amount, from 36 868 888 to 36 809 056 bytes (about 0.162%). This is one measured A/B instance on the supported CPU trace shape.

We also measured latency, on an Apple M4 Max with 16 threads. Over 20 paired rounds, the change in proving time is −15.2 ms, with paired bootstrap 95% interval [−49.4, +18.9] ms; over 50 paired rounds, verification changes by +0.02 ms, with interval [−0.62, +0.52] ms. We detect no latency regression on this benchmark; the artifact records the procedure, component timings, and memory diagnostics [Dao26a].

8 Higher derivatives up to capacity

The first-order bounds hold only above the first-order curve, or uniformly at gaps $\delta \geq 6/25$ from capacity (Theorem 5.7). To reach every positive gap we hide d derivatives instead of one, as the earlier capacity results do (Section 1.2 compares how they reach every fixed rate). Theorem 8.1 makes one order $\lceil e^{1.5/\delta} \rceil$, multiplicity, and minimum length serve all rates, with list and line-MCA bounds polynomial in n. As at first order, we construct an explainer from BCPZZ's order-d local constraints (Definition 3.6) and count the Taylor images of its solutions, now in the symbolic-family framework of Zheng and Jeronimo [Zhe26, Jer26]; Section 2.4 sketches the argument.

8.1 The uniform capacity theorem

For each gap $0 < \delta < 6/25$, the uniform theorem fixes four parameters that depend on δ alone: the derivative order $d(\delta)$, the multiplicity $m(\delta)$, the bound $B_{\mathrm{jet}}(\delta)$ on the explainer's total degree in the jet variables $Y_0, \ldots, Y_d$, and the minimum length $N(\delta)$:

$$\begin{aligned} d(\delta) &= \lceil e^{1.5/\delta} \rceil, \qquad m(\delta) = \lceil 300 d(\delta)^2 \log(6d(\delta)) \rceil, \\ B_{\mathrm{jet}}(\delta) &= \left\lceil \frac{m(\delta)}{\delta^2} \right\rceil - 1, \qquad N(\delta) = \left\lceil \frac{4B_{\mathrm{jet}}(\delta)}{\delta^2} \right\rceil. \end{aligned} \tag{55}$$

Theorem 8.1 (Uniform finite choices up to capacity)**.** *For $0 < \delta < 6/25$, use the parameters in (55). Let F be a field, fix $n \geq N(\delta)$ distinct evaluation points in F, and let $1 \leq k < A \leq n$ satisfy $A - k \geq \delta n$.*

Interpolation. *For every received word $y \in F^n$ there is an order-$d(\delta)$ explainer of jet degree at most $B_{\mathrm{jet}}(\delta)$, with ambient dimension $k \leq K \leq n$, $d(\delta) < K$, constructed from a monomial support whose dimension and local rank satisfy $N_{\mathrm{coef}} > (151/150)nR_{\mathrm{loc}}$. This interpolation assertion is characteristic independent. For every received line, the same support gives a symbolic explainer of challenge degree less than $150B_{\mathrm{jet}}(\delta)$.*

Counting. *In characteristic zero or $p > k - 1$ when $k \geq 2$, and in every characteristic when $k = 1$, every received word has*

$$|\mathcal{L}(y, A)| \leq \max\left\{ B_{\mathrm{jet}}(\delta)^2 (2B_{\mathrm{jet}}(\delta)/\delta)^{d(\delta)}, \frac{4}{3\delta} \right\} n^{d(\delta)},$$

and every received line has the full common-witness property of Definition 3.1, for all degree-$< k$ candidates with at least A agreements, outside a set $\mathcal{E} \subseteq F$ with

$$|\mathcal{E}| \leq C_\delta n^{d(\delta)+1}, \qquad C_\delta = \max\{C_E, 343/3\},$$

where, with $d = d(\delta)$, $B_{\mathrm{jet}} = B_{\mathrm{jet}}(\delta)$, and $B_Z = 150B_{\mathrm{jet}}(\delta)$,

$$C_E = B_Z + 2^d B_{\mathrm{jet}}^{d+2} \delta^{-d} \left\lceil \frac{B_Z(d+1)(3d+5)}{\delta} + 3 \right\rceil.$$

Over $\mathbb{F}_q$, the exceptional probability is at most $C_\delta n^{d(\delta)+1}/q$.

For $6/25 \leq \delta < 1$, let $n \geq 2$, $1 \leq k \leq n$, and $k + \delta n \leq A \leq n$, and fix n distinct evaluation points in a field F. If $k \geq 2$, assume its characteristic is zero or greater than $k - 1$; no restriction is needed when $k = 1$. Then Theorem 5.7 gives list bound $307n$ and line exceptional bound $1325775n^2$.

The four parameters in (55) come from different steps of the proof. At every rate we interpolate with ambient degree bound $D \geq \delta^2 n$, padding it when $k < \delta^2 n$; this keeps every monomial of the support below jet degree $m(\delta)/\delta^2$ and gives $B_{\mathrm{jet}}(\delta)$. The order $d(\delta)$ gives $\delta \log d(\delta) \geq 3/2$, which Lemma D.1 turns into the bound $\Gamma \geq \frac{9}{40} e^{3/2} \approx 1.0084$ on the limiting coefficient-to-condition ratio (61), at every rate including the

padded ones. The multiplicity $m(\delta)$ is the closed choice of Lemma 8.3, at which rounding costs at most the factor $e^{-1677/10^6}$, so the ratio stays above $151/150$ and the challenge degree below $150B_{\mathrm{jet}}(\delta)$. The minimum length gives $\delta^2 n \geq m(\delta)$ and, when $k-1 < p \leq B_{\mathrm{jet}}(\delta)$, places the agreement fraction at least $\delta/2$ above $\sqrt{(k-1)/n}$, where Theorem 4.4 applies.

The constant C_E is (67) of Proposition 8.7 at the margin $\gamma = 151/150$, and the second entries of the two maxima come from the Johnson-regime argument used when $k-1 < p \leq B_{\mathrm{jet}}(\delta)$ (Section 8.5). Corollaries 6.3 and 6.6 carry this bound to interleaved received lines and, with denominator $q-1$, to affine families. For received curves of degree ℓ, the same support has challenge degree less than $150\ell B_{\mathrm{jet}}(\delta)$ by (18), and Corollary 8.6 gives an exceptional count linear in ℓ, in characteristic zero or $p > \max\{k-1, B_{\mathrm{jet}}(\delta)\}$ (recall $d(\delta) < B_{\mathrm{jet}}(\delta)$). Section 9.8 computes these lists.

8.2 A weighted support and the derivative-order condition

As at first order, an explainer exists once the number of coefficients exceeds the bound on the number of independent local conditions, $N_{\mathrm{coef}} > nR_{\mathrm{loc}}$, and the size of the surplus bounds the challenge degree (Proposition 3.13). With d derivative variables, we must also choose how much degree each may take. For integers $D \geq d$, $m, W \geq 1$, and an integer threshold $A_0 \leq A$, we use the translation-invariant support

$$\mathcal{V}_{m,W} = \operatorname{span}\left\{ X^x Y_0^{b_0} \prod_{j=1}^{d} Y_j^{b_j} : \sum_{j=1}^{d} j b_j \leq W, \quad x + D\left(b_0 + \sum_{j=1}^{d} b_j\right) < mA_0 \right\}, \tag{56}$$

with nonnegative integer exponents. The second inequality gives degree less than mA_0 after substituting any polynomial of degree at most D, the ambient degree bound. The first bounds the weighted derivative degree by W, which limits the derivative monomials reachable by the local map. Unlike the first-order support, $\mathcal{V}_{m,W}$ distributes the derivative degree among several orders, weighting Y_j by j.

Let $p_d(B) = \#\{b \in \mathbb{N}^d : \sum_j j b_j \leq B\}$, the number of partitions of integers at most B into parts of size at most d, and write $|b| = \sum_j b_j$. The dimension and a local-rank bound are

$$N_{\mathrm{coef}} = \sum_{\sum_j j b_j \leq W} \sum_{b_0 \geq 0} (mA_0 - D(b_0 + |b|))_+, \tag{57}$$

$$R_{\mathrm{loc}} \leq \sum_{s=0}^{m-1} \left\lceil \frac{m-s}{d+1} \right\rceil p_d(W+s). \tag{58}$$

The rank bound counts the rows of the local map that can be nonzero. In the coordinates $\widetilde{V} = T^d V$ of Section 3.2, the local map substitutes $Y_0 = y + \sum_{j=1}^{d} (-1)^{j+1} T^j Y_j + T\widetilde{V}$, so each factor Y_j it creates comes with T^j, and each factor $\widetilde{V}$ with T. A row $T^i \widetilde{V}^u Y^{b'}$ reachable from the support therefore satisfies

$$i \geq u, \qquad \sum_j j b'_j \leq W + i - u, \qquad i + du < m,$$

where the last condition says that $T^i \widetilde{V}^u = T^{i+du} V^u$ is nonzero modulo T^m. Because the weight of Y_j equals the power of T that accompanies it, the weighted degree of a row exceeds W by at most its excess T-degree $s = i - u$. Each power of $\widetilde{V}$ costs $d+1$ powers of T, since $i + du = s + (d+1)u$, so there are at most $\lceil (m-s)/(d+1) \rceil$ possible u and at most $p_d(W+s)$ possible b'; this divisor is where hiding d derivatives saves conditions.

At any order, the inequality $N_{\mathrm{coef}} > nR_{\mathrm{loc}}$ suffices for interpolation. To see which orders satisfy it at a given rate, fix $0 < \rho < a < 1$ with $D \leq \rho n$ and $A_0 \geq an$, put $w = W/m$, and let U be uniform on the weighted simplex S_w, where

$$S_B = \{\beta \in \mathbb{R}^d_{\geq 0} : \textstyle\sum_j j\beta_j \leq B\}$$

is the real relaxation of the exponent vectors b. Each point of S_W rounds down to an eligible b with no larger $|b|$, so comparing the sum over b and the sum over the Y_0-exponent with integrals gives

$$\frac{N_{\mathrm{coef}}}{nm^2\,\mathrm{vol}(S_W)} \geq \frac{1}{2\rho}\mathbb{E}\big[(a-\rho|U|)_+^2\big], \qquad \mathrm{vol}(S_W) = \frac{W^d}{(d!)^2}. \tag{59}$$

After normalization, the term $(a-\rho|U|)_+^2/(2\rho)$ counts the monomials in X and Y_0 that fit in the degree left once the derivative exponents are chosen; keeping it is stronger than replacing $|U|$ by its mean. To evaluate the expectation, set $v_i = \sum_{j=i}^{d} U_j$. Then $v_1 \geq \cdots \geq v_d \geq 0$, $\sum_i v_i \leq w$, and $|U| = v_1$. This change of variables has determinant one, so sorting a uniform point of the ordinary simplex shows that $|U|$ has the law of wM_d, with M_d as in the next lemma. Its proof is in Section D.1.

Lemma 8.2 (A simplex moment bound)**.** *Let $(T_1,\dots,T_d)$ be uniform on the simplex $T_i \geq 0$, $\sum_i T_i \leq 1$, and put $M_d = \max_i T_i$. For every integer $d \geq 500$, the variable $Z_d = dM_d - \log d$ satisfies*

$$\mathbb{E}\big[(\log 6 - Z_d)_+^2\big] > \frac{27}{10}. \tag{60}$$

We take $W = \lfloor mad/(\rho\log(6d))\rfloor$. Since $\rho w \leq ad/\log(6d)$ and $|U|$ has the law of wM_d, the degree left for X and Y_0 satisfies $a-\rho|U| \geq a(\log 6 - Z_d)/\log(6d)$ in distribution, so Lemma 8.2 gives $\mathbb{E}\big[(a-\rho|U|)_+^2\big] > \frac{27}{10}a^2/\log^2(6d)$. On the rank side, placing a unit cube at each lattice point bounds $p_d(W+s)$ by $\mathrm{vol}(S_W)e^{\lambda(s+d(d+1)/2)/m}$, where $\lambda = dm/W$. Dividing the two estimates and letting $m \to \infty$ bounds $N_{\mathrm{coef}}/(nR_{\mathrm{loc}})$ below by

$$\Gamma(\rho,a,d) = \frac{27}{20}\,\frac{\rho(d+1)}{(6d)^{\rho/a}}. \tag{61}$$

The factor $d+1$ comes from the count $\lceil (m-s)/(d+1)\rceil$ in (58), and $(6d)^{-\rho/a}$ is the limit of $e^{-\lambda}$, which records the rows with s close to m. Since Γ grows like $d^{1-\rho/a}$, some order suffices whenever $a > \rho$. The following finite version, proved in Section D.2, includes the rounding error at finite m; below order 500, the exact counts can be evaluated instead.

Lemma 8.3 (Finite partition ratio)**.** *Let $d \geq 500$, $0 < \rho < a \leq 1$, and $m \geq 1$ be an integer for which $W = \lfloor mad/(\rho\log(6d))\rfloor \geq 1$. Set*

$$w = W/m, \quad \lambda = d/w, \qquad \gamma_m = \frac{27}{20}\rho(d+1)\,\frac{e^{-\lambda(1+d(d+1)/(2m))}}{1+(d+1)\lambda/m}.$$

For the space (56), with $D \geq d$, $D \leq \rho n$, and $A_0 \geq an$, its dimension and local rank obey $N_{\mathrm{coef}} > \gamma_m nR_{\mathrm{loc}}$ whenever $R_{\mathrm{loc}} > 0$; its dimension is positive also when $R_{\mathrm{loc}} = 0$. Moreover $\gamma_m \to \Gamma(\rho,a,d)$ as $m \to \infty$. If $d \geq 519$ and $m = \lceil 300d^2\log(6d)\rceil$, then

$$\gamma_m > \Gamma(\rho,a,d)e^{-1677/10^6}. \tag{62}$$

Thus whenever $\Gamma > 1$, every sufficiently large m gives $N_{\mathrm{coef}} > nR_{\mathrm{loc}}$ and hence an explainer. With the closed choice of m, the stronger condition $\Gamma > e^{1677/10^6}(151/150)$ gives $N_{\mathrm{coef}} > (151/150)nR_{\mathrm{loc}}$ and hence challenge degree less than $150B_{\mathrm{jet}}$; this is the margin Theorem 8.1 uses.

8.3 Counting the solutions of an order-d explainer

With an explainer in hand, the list bound reduces to counting its close solutions. The regular solutions are counted as at first order: a regular solution is determined by its initial jet at a center, because Kopparty's recurrence [Kop15] recovers the remaining Taylor coefficients, so the regular solutions lie in the image

of a rational map; Lemma A.3 bounds its degree, and Lemma 4.1 bounds the points of its image with A agreements.

The other solutions are treated differently. At first order, Section 5 reduces them to an ordinary equation through the squarefree split $Q_{\mathrm{ord}}Q_{\mathrm{diff}}$ and a resultant in Y_1. At order d we use instead the separant chain (91), which repeatedly differentiates in the highest jet variable still present and assigns each solution to the first member that vanishes along it while the next member does not. Call such a member an *actual stage* (j, r), where j is its total jet degree and Y_r its highest jet variable. Since only the initial jet and the degree-$<$ k message matter, we reconstruct only $K_{\mathrm{rec}} = \max\{k, d+1\}$ coefficients, whatever the larger interpolation dimension. The chain needs the characteristic to exceed B_{jet}, and the recurrence through K_{rec} coefficients needs it to exceed $K_{\mathrm{rec}} - 1$; together these give the hypothesis of Theorem 8.4. The bound below also uses the incidence factor $\Pi_r(T)$ from (31).

Theorem 8.4 (Direct agreement bound for differential roots)**.** *Let $1 \le k \le n$ and $0 \le d < n$, and put $K_{\mathrm{rec}} = \max\{k, d+1\}$. Let $Q \in F[X, Y_0, \dots, Y_d]$ be nonzero of total jet degree at most B_{jet}, where $B_{\mathrm{jet}} \ge 1$, and assume the characteristic is zero or exceeds $\max\{k-1, d, B_{\mathrm{jet}}\}$. For distinct $\alpha_1, \dots, \alpha_n \in F$, $y \in F^n$, and $k \le A \le n$, write $\mathcal{L}$ for the set of degree-$<$ k solutions of $Q(X, P, P^{[1]}, \dots, P^{[d]}) = 0$ with at least A agreements. Then $\mathcal{L}$ is finite. If (j, r) runs over the actual stages of its separant chain and $\Delta_j = 1 + 2K_{\mathrm{rec}}(j-1)$, its size is at most*

$$\sum_{\text{actual stages } (j,r)} \Pi_r(A)\, j\Delta_j^r \le \Pi_d(A) \sum_{j=1}^{B_{\mathrm{jet}}} j\Delta_j^d \le B_{\mathrm{jet}}^2 \left(\frac{n[1 + 2K_{\mathrm{rec}}(B_{\mathrm{jet}} - 1)]}{A - k + 1} \right)^d .$$

If moreover $A - k \ge \delta n > 0$, then

$$|\mathcal{L}| \le B_{\mathrm{jet}}^2 (2B_{\mathrm{jet}}/\delta)^d n^d. \tag{63}$$

Proof. If Q has jet degree zero, it is a nonzero polynomial in X and has no solutions. Otherwise, the characteristic hypothesis makes every member of the separant chain nonzero, and the chain ends at a nonzero polynomial in X, so every solution belongs to some stage.

Choose a center c transcendental over $\overline{F}$ and work over $E = \overline{F(c)}$. Transcendence of c keeps the substituted separant nonzero at $X = c$, so this one center is regular for every candidate at its assigned stage.

Fix a stage of jet degree j and highest jet index $r \le d$, and put $\Delta_j = 1 + 2K_{\mathrm{rec}}(j-1)$. By Lemma A.3, applied with jet-degree bound j and $K = K_{\mathrm{rec}}$, its regular initial-jet locus has dimension at most r and total degree at most j, and its map to the K_{rec} Taylor coefficients has numerator and common-denominator degrees at most Δ_j. The denominator vanishes identically on none of the components we keep. Hence Lemma A.4 gives an image closure of dimension at most r and total degree at most $j\Delta_j^r$.

Intersect this image with the linear subspace of coefficient vectors representing polynomials of degree less than $k \le K_{\mathrm{rec}}$. Translation from ordinary to Taylor coefficients preserves this degree filtration, and a linear intersection does not increase dimension or total degree. In this message space, Lemma 4.1 bounds the points agreeing with y in at least A positions by $\Pi_r(A)\, j\Delta_j^r$. Every candidate assigned to this stage is one of these points.

The jet degrees of the actual stages are positive and strictly decrease along the chain, so they are distinct elements of $\{1, \dots, B_{\mathrm{jet}}\}$. Since $r \le d$, $\Delta_j \ge 1$, and every incidence factor is at least one, summing over the actual stages gives the first two bounds in the statement and proves finiteness. Finally,

$$\sum_{j=1}^{B_{\mathrm{jet}}} j\Delta_j^d \le \frac{B_{\mathrm{jet}}(B_{\mathrm{jet}} + 1)}{2} \Delta_{B_{\mathrm{jet}}}^d \le B_{\mathrm{jet}}^2 \Delta_{B_{\mathrm{jet}}}^d, \qquad \Pi_d(A) \le \left(\frac{n}{A - k + 1} \right)^d ,$$

which gives the coarse bound. If $A - k \ge \delta n$, then $1 + 2K_{\mathrm{rec}}(B_{\mathrm{jet}} - 1) \le 2B_{\mathrm{jet}} n$ and $A - k + 1 \ge \delta n$ give (63). □

For MCA we keep the challenge Z in the Taylor map, obtaining one family of images for a received curve; Lemma 4.2 turns bounds on its joint and generic-fiber degrees into a bound on the exceptional challenges.

Theorem 8.5 (Symbolic transfer for polynomial curves)**.** *Let F be a field, $1 \le k \le L \le A \le n$, $0 \le d < n$, and $K_{\mathrm{rec}} = \max\{k, d+1\}$, and let $B_{\mathrm{jet}} \ge 1$ and $B_Z \ge 0$ be integers. Fix $\ell \ge 1$, distinct evaluation points $\alpha_1, \dots, \alpha_n \in F$, and coefficient words $w_0, \dots, w_\ell \in F^n$, giving the received curve $w(z) = \sum_{t=0}^{\ell} z^t w_t$ and agreement sets $S_z(P)$. Suppose* $\operatorname{char} F = 0$ *or* $\operatorname{char} F > \max\{k-1, d, B_{\mathrm{jet}}\}$*. Let $Q \in F[X, Z, Y_0, \dots, Y_d]$ be nonzero, of total jet degree at most B_{jet} and challenge degree at most B_Z. Suppose every degree-$\prec k$ polynomial P with $|S_z(P)| \ge A$ satisfies the symbolic soundness condition (17). Put $\lambda_{\mathrm{off}} = (n-L+1)/(A-L+1)$, $B_{\mathrm{lift}} = 1 + 2K_{\mathrm{rec}}(B_{\mathrm{jet}} - 1)$, and*

$$D^{(\ell)}_{\mathrm{joint}} = B_{\mathrm{jet}} B^d_{\mathrm{lift}} [B_Z B_{\mathrm{lift}} + (d+1) B_{\mathrm{jet}} (\ell + 2K_{\mathrm{rec}} B_Z)].$$

There is an exceptional set $\mathcal{E} \subseteq F$ of size at most

$$B_Z + \lambda_{\mathrm{off}} \Pi_d(A) D^{(\ell)}_{\mathrm{joint}} + \ell(n-L)\Pi_d(L) B^2_{\mathrm{jet}} B^d_{\mathrm{lift}} \tag{64}$$

such that every pair (z, P) outside this set with $\deg P < k$ *and $|S_z(P)| \ge A$ has the full common-witness property of Definition 3.1. No characteristic restriction involves ℓ.*

Proof. Lemma D.2, applied with $K = K_{\mathrm{rec}}$, gives a set Z_0 of at most B_Z challenges and closed sets of dimension at most $d+1$ whose union contains every lifted root pair $(z, z^2, \dots, z^\ell, P)$ with $\deg P < K_{\mathrm{rec}}$ and $z \notin Z_0$. This union has joint degree at most $D^{(\ell)}_{\mathrm{joint}}$, and its reduced fiber over a transcendental challenge has dimension at most d and degree at most $B^2_{\mathrm{jet}} B^d_{\mathrm{lift}}$. Intersect it with the linear subspace $\deg P < k$, which increases neither degree. Symbolic soundness places every candidate in this intersection, so Lemma 4.2 applies with

$$r_1 = d, \qquad J_1 = D^{(\ell)}_{\mathrm{joint}}, \qquad F_1 = B^2_{\mathrm{jet}} B^d_{\mathrm{lift}}, \qquad \mathcal{E}_{\mathrm{pre}} = Z_0 \cap F.$$

It gives (64), witnesses over F, and equality of full agreement sets. Its accidental-agreement polynomials have degree at most ℓ in z, so no characteristic restriction involving ℓ is needed. □

Combining words with coefficients $1, z, \dots, z^\ell$ is called powers batching. Its challenge-degree bound grows by a factor ℓ, and (64) is linear in that bound, so the exceptional count grows only linearly in ℓ.

Corollary 8.6 (Powers batching from a certified support)**.** *Fix the parameters of Theorem 8.5, with characteristic zero or greater than* $\max\{k-1, d, B_{\mathrm{jet}}\}$*. Suppose a fixed support constructs a symbolic explainer satisfying that theorem's hypotheses for every scalar received curve of degree at most ℓ, and certifies $N_{\mathrm{coef}} \ge \gamma n R_{\mathrm{loc}}$ for some $\gamma > 1$. Then (18) gives*

$$B_Z \le \left\lfloor \frac{\ell B_{\mathrm{jet}}}{\gamma - 1} \right\rfloor.$$

Put $c = B_{\mathrm{jet}}/(\gamma - 1)$, and let C_{pow} be the coefficient of ℓ in (64) after substituting $B_Z = c\ell$. Then at most ℓC_{pow} challenges are exceptional: outside them, every degree-$\prec k$ candidate with at least A agreements has the full common-witness property of Definition 3.1. No further characteristic restriction depends on ℓ. Over $\mathbb{F}_q$, the error is at most $\ell C_{\mathrm{pow}}/q$, with the same guarantee at every interleaving width.

Proof. Each term of (64) is nondecreasing in B_Z and becomes linear in ℓ after this substitution. Apply Theorem 8.5 to every scalar received curve; Corollary 6.6 then gives the interleaved conclusion. □

8.4 List and line-MCA bounds at a fixed rate

This subsection combines the interpolation ratio of Section 8.2 with the counting theorems of Section 8.3. Besides d and δ, the resulting constants depend on the support only through its jet-degree bound B_{jet} and its margin $\gamma - 1$, which bounds the challenge degree by $B_{\mathrm{jet}}/(\gamma - 1)$ (Proposition 3.13). Proposition 8.7 states this for any certified support, and Corollary 8.8 applies it to the weighted support at a fixed rate, where $B_{\mathrm{jet}} = \lceil 2m/\rho \rceil$ still depends on the rate.

Proposition 8.7 (Bounds from a certified support)**.** *Let $1 \le d < n$, $m \ge 1$, and $1 \le k < A \le n$ be integers, and let $\delta > 0$ satisfy $A - k \ge \delta n$. Fix a field F and n distinct evaluation points in F. Let $Q \subseteq F[X, Y_0, \dots, Y_d]$ be a monomial space invariant under translations in X and Y_0, of total jet degree at most $B_{\mathrm{jet}} \ge 1$, whose specialization degree at degree-$\prec k$ polynomials is less than mA_0 for an integer $1 \le A_0 \le A$. Suppose that $N_{\mathrm{coef}} = \dim Q$ and an upper bound R_{loc} for the rank of the local constraint map $\Phi^{d,m}_{0,0}$ on Q satisfy $N_{\mathrm{coef}} > \gamma n R_{\mathrm{loc}}$ for some $\gamma > 1$, and that F has characteristic zero or $p > \max\{k - 1, d, B_{\mathrm{jet}}\}$. Then every received word has list size*

$$|\mathcal{L}(y, A)| \le C_L n^d, \qquad C_L = B_{\mathrm{jet}}^2 (2B_{\mathrm{jet}}/\delta)^d. \tag{65}$$

Put

$$B_Z = \max\{1, \lceil B_{\mathrm{jet}}/(\gamma - 1) \rceil\}. \tag{66}$$

Every received line has an exceptional set of size at most $C_E n^{d+1}$ outside which every degree-$\prec k$ candidate with at least A agreements has the full common-witness property of Definition 3.1, where

$$C_E = B_Z + 2^d B_{\mathrm{jet}}^{d+2} \delta^{-d} \left\lfloor \frac{B_Z (d+1)(3d+5)}{\delta} + 3 \right\rfloor. \tag{67}$$

Over $\mathbb{F}_q$, the exceptional probability is at most $C_E n^{d+1}/q$.

Proof. Apply Proposition 3.13 with dimension k and threshold A_0. For every received line it gives a symbolic explainer of jet degree at most B_{jet} and challenge degree at most $\lfloor n R_{\mathrm{loc}} B_{\mathrm{jet}}/(N_{\mathrm{coef}} - n R_{\mathrm{loc}}) \rfloor \le B_Z$, satisfying (17) for every degree-$\prec k$ candidate with at least $A \ge A_0$ agreements. Taking $g = 0$ and $z = 0$ gives an explainer for each received word. Since $K_{\mathrm{rec}} = \max\{k, d+1\} \le n$, (63) gives (65).

For line MCA, apply Theorem 8.5 with $\ell = 1$ and $L = k + \lfloor d(A-k)/(d+1) \rfloor$. Since $A - L + 1 > (A-k)/(d+1)$ and every factor of Π_d has the form $(n - k + i)/(T - k + i)$, its incidence factors obey

$$\lambda_{\mathrm{off}} \le \frac{d+1}{\delta}, \qquad \Pi_d(A) \le \delta^{-d}, \qquad \Pi_d(L) < 3\delta^{-d}.$$

The last estimate uses $(1 + 1/d)^d < 3$; integer rounding is absorbed by the $+i$ in each denominator of Π_d. Since $K_{\mathrm{rec}} \le n$ and $B_Z \ge 1$, the joint and fiber degrees satisfy

$$B_{\mathrm{lift}} \le 2B_{\mathrm{jet}} n, \quad D^{(1)}_{\mathrm{joint}} \le B_Z (3d+5) 2^d B_{\mathrm{jet}}^{d+2} n^{d+1}, \quad B_{\mathrm{jet}}^2 B_{\mathrm{lift}}^d \le 2^d B_{\mathrm{jet}}^{d+2} n^d.$$

Inserting these bounds into (64) and using $n - L \le n$ gives $C_E n^{d+1}$. □

At a fixed rate with $\Gamma(\rho, a, d) > 1$, the support (56) satisfies $N_{\mathrm{coef}} > \gamma n R_{\mathrm{loc}}$ for some $\gamma > 1$ once the multiplicity is large enough.

Corollary 8.8 (Lists and line MCA from the derivative-order condition)**.** *Fix $0 < \rho < a < 1$ and an integer $d \ge 500$ with $\Gamma(\rho, a, d) > 1$. There are computable integers $m, B_{\mathrm{jet}}, N_{\rho,a,d}$ and a certified ratio $\gamma > 1$ with the following property. Let n, k, A be integers with $n \ge N_{\rho,a,d}$, $1 \le k \le \rho n$, and $an \le A \le n$. Over any field of characteristic zero or $p > \max\{k - 1, d, B_{\mathrm{jet}}\}$, and for any n distinct evaluation points, every received word has list size at most $C_L n^d$, and every received line has an exceptional set of size at most $C_E n^{d+1}$ outside which all close candidates have the full common-witness property, where C_L, B_Z, and C_E are given by (65)–(67) with $\delta = a - \rho$. Over $\mathbb{F}_q$, the exceptional probability is at most $C_E n^{d+1}/q$.*

Proof. Increase m until $\gamma = \gamma_m$ in Lemma 8.3 is greater than one; this terminates since $\gamma_m \to \Gamma > 1$. Set $W = \lfloor mad/(\rho \log(6d)) \rfloor$ and $B_{\mathrm{jet}} = \lceil 2m/\rho \rceil$. For example, take

$$N_{\rho,a,d} = \left\lceil \max\left\{ \frac{d+1}{\rho}, B_{\mathrm{jet}} + 1, \frac{1}{1-\rho} \right\} \right\rceil.$$

Then, for $n \geq N_{\rho,a,d}$,

$$D = \lfloor \rho n \rfloor \geq \max\{d+1, \rho n/2\}, \quad D+1 \leq n, \quad n > B_{\mathrm{jet}}, \quad n \geq 2m.$$

Use $A_0 = \lceil an \rceil \leq A$. The space (56) has $N_{\mathrm{coef}} > \gamma n R_{\mathrm{loc}}$, total jet degree less than $mA_0/D \leq 2m/\rho \leq B_{\mathrm{jet}}$, and specialization degree less than mA_0 at every polynomial of degree at most D, in particular at every message, since $k \leq D$. Since $A - k \geq (a-\rho)n$, the claims follow from Proposition 8.7 with $\delta = a - \rho$. □

Remark 8.9 (Explicit derivative orders)**.** Rewriting Γ gives the exact identity

$$\Gamma(\rho, a, d) = \frac{9\rho}{40}(6d)^{(a-\rho)/a}\left(1 + \frac{1}{d}\right), \tag{68}$$

so

$$d \geq \max\left\{500, \frac{1}{6}\left(\frac{40}{9\rho}\right)^{a/(a-\rho)}\right\} \quad \Longrightarrow \quad \Gamma(\rho, a, d) \geq 1 + \frac{1}{d} > 1. \tag{69}$$

At additive gap $\delta = a - \rho$, the order is $\max\{500, (20/(27\rho))\exp(c(\rho)/\delta)\}$ with $c(\rho) = \rho \log(40/(9\rho))$. At a fixed multiplicative slack $\rho = (1-\vartheta)a$, the exponent $a/(a-\rho)$ equals $1/\vartheta$, and every integer

$$d \geq \max\left\{500, \frac{1}{6}\left(\frac{40}{9\rho}\right)^{1/\vartheta}\right\} \tag{70}$$

suffices; this grows only polynomially as $a \to 0$. Corollary 8.8 applies to both choices, with its length and characteristic hypotheses; Sections 1.2 and G.2 compare them with the orders of prior work.

Remark 8.10 (Bounded-input list recovery)**.** In list recovery, each position i has an input list $S_i \subseteq F$ of symbols, and the list consists of the degree-$< k$ polynomials P with $P(\alpha_i) \in S_i$ at no fewer than A positions. Following Jeronimo's point–symbol formulation [Jer26], the proof of Corollary 8.8 extends to $|S_i| \leq L_{\mathrm{in}}$, with total input size $M = \sum_i |S_i| \leq L_{\mathrm{in}} n$. Fix an integer $L_{\mathrm{in}} \geq 1$, reals $0 < \rho < a < 1$, and an integer $d \geq 500$. If $\Gamma(\rho, a, d) > L_{\mathrm{in}}$, there are explicit integers m, B_{jet}, $N_{\rho,a,d,L_{\mathrm{in}}}$ such that for $n \geq N_{\rho,a,d,L_{\mathrm{in}}}$, $1 \leq k \leq \rho n$, and $an \leq A \leq n$, the list has at most

$$B_{\mathrm{jet}}^2 \left(\frac{2L_{\mathrm{in}} B_{\mathrm{jet}}}{a-\rho}\right)^d n^d$$

members, in characteristic zero or $p > \max\{k-1, d, B_{\mathrm{jet}}\}$. By (68), the order

$$d \geq \max\left\{500, \frac{1}{6}\left(\frac{40L_{\mathrm{in}}}{9\rho}\right)^{a/(a-\rho)}\right\}$$

suffices. Three changes are needed. First, interpolation imposes one local system for each pair (α_i, y) with $y \in S_i$, so the stacked rank is at most $MR_{\mathrm{loc}} \leq L_{\mathrm{in}} n R_{\mathrm{loc}}$; the condition $\Gamma > L_{\mathrm{in}}$ means that a large enough m gives $N_{\mathrm{coef}} > L_{\mathrm{in}} n R_{\mathrm{loc}}$. Second, the incidence count of Lemma 4.1 uses the M hyperplanes $\{P(\alpha_i) = y\}$, $y \in S_i$, so its factor becomes $\prod_{j \leq r}(M-k+j)/(T-k+j)$. Third, the hyperplanes are labelled by their positions: two with the same position and different symbols are disjoint, so an irreducible component lies in at most one per position, Vandermonde inversion still shows that a component of dimension t lies in at most $k - t$ of them, and agreement still counts distinct positions.

8.5 Proof of the uniform theorem

Corollary 8.8 does not yet give rate-independent parameters: its jet-degree bound $\lceil 2m/\rho \rceil$ diverges as $\rho \to 0$, and so do its list and line constants, which grow with B_{jet}. Like Jeronimo [Jer26], we enlarge the explained message space, but only at low rates: when $k < \delta^2 n$ we interpolate at ambient degree bound $\lfloor 2\delta^2 n \rfloor$.

Proof of Theorem 8.1. Fix $0 < \delta < 6/25$ and abbreviate $d(\delta), m(\delta), N(\delta)$ by d, m, N. Then $d \geq 519$. Write $\rho_* = k/n$ and use one of the following ambient choices:

$$(\rho, D, a) = \begin{cases} (\rho_*, k, \rho_* + \delta), & \rho_* \geq \delta^2, \\ (2\delta^2, \lfloor 2\delta^2 n \rfloor, \delta), & \rho_* < \delta^2. \end{cases}$$

In both cases set $K = D + 1$, $A_0 = \lceil an \rceil$ and $W = \lfloor mad/(\rho \log(6d)) \rfloor$. The low-rate choice keeps $D \geq \delta^2 n$ even when k is small. The length hypothesis gives

$$k \leq D, \quad \delta^2 n \leq D \leq \rho n, \quad D \geq m > d, \quad K \leq n, \quad A_0 \leq A \leq n.$$

Here $N = \lceil 4B_{\mathrm{jet}}/\delta^2 \rceil \geq B_{\mathrm{jet}} + 1 = \lceil m/\delta^2 \rceil$, since $4/\delta^2 > 1$, so $\delta^2 n \geq m \geq 1$ and $\lfloor 2\delta^2 n \rfloor \geq \delta^2 n$; for $K \leq n$, use $k < A \leq n$ in the high-rate branch and $2\delta^2 < 1/8$ in the low-rate branch. This padding makes the jet-degree bound uniform over the rate:

$$D \geq \delta^2 n \quad \Longrightarrow \quad \deg_{\mathrm{jet}} Q < mA_0/D \leq m/\delta^2,$$

so the integer jet degree is at most $B_{\mathrm{jet}}(\delta) < n$.

In both branches $\log \Gamma(\rho, a, d) \geq 3/2 - \log(40/9)$ by Lemma D.1. Since $m = m(\delta)$ is the closed choice of Lemma 8.3, that lemma and the rational inequality in Lemma D.1 give

$$\log \frac{N_{\mathrm{coef}}}{nR_{\mathrm{loc}}} > \frac{3}{2} - \log(40/9) - \frac{1677}{10^6} > \log(151/150)$$

when $R_{\mathrm{loc}} > 0$; the case $R_{\mathrm{loc}} = 0$ is immediate. Proposition 3.13 now gives the interpolation claims in every characteristic, with challenge degree less than $150B_{\mathrm{jet}}(\delta)$.

For the counting claims with $k = 1$, Lemma 3.3 gives list bound n and line exceptional bound $(n-1)/(2\delta)$ in every characteristic, both smaller than the displayed bounds. Let $k \geq 2$. In characteristic zero or $p > B_{\mathrm{jet}}(\delta)$, apply Proposition 8.7 with $\gamma = 151/150$, so that (66) gives $B_Z = 150B_{\mathrm{jet}}(\delta)$. The support has specialization degree less than mA_0 at degree-$\prec k$ polynomials since $k \leq D$, and $d < B_{\mathrm{jet}}(\delta) < n$, so the proposition's characteristic hypothesis is $p > \max\{k - 1, B_{\mathrm{jet}}(\delta)\}$. It gives the list bound $B_{\mathrm{jet}}(\delta)^2 (2B_{\mathrm{jet}}(\delta)/\delta)^d n^d$ and the exceptional bound $C_E n^{d+1}$, which use the first entries of the two maxima.

It remains to treat $k - 1 < p \leq B_{\mathrm{jet}}(\delta)$. Then $k \leq B_{\mathrm{jet}}(\delta)$, and $n \geq N(\delta) \geq 4B_{\mathrm{jet}}(\delta)/\delta^2$ gives

$$\frac{k-1}{n} \leq \frac{\delta^2}{4}, \qquad \frac{A}{n} - \sqrt{\frac{k-1}{n}} \geq \frac{\delta}{2}.$$

Apply Theorem 4.4 with $D = k - 1$ and $a = A/n$. Its multiplicity is $m = 3$, and its bounds give

$$|\mathcal{L}(y, A)| \leq \frac{4}{3\delta}, \qquad |\mathcal{E}| < \frac{343}{3} \frac{n^2}{k-1} \leq \frac{343}{3} n^2.$$

Since $d \geq 1$, these are at most the stated bounds. The larger-gap branch is Theorem 5.7. □

9 Efficient decoding over large fields

Sections 3 to 8 bound the number of close messages; we now compute them. Fix distinct evaluation points $\alpha_1, \dots, \alpha_n \in \mathbb{F}_q$ with $q \geq n$, a received word $y \in \mathbb{F}_q^n$, and a threshold $k \leq A \leq n$. We want the list

$$\mathcal{L}(y, A) = \{P \in \mathbb{F}_q[X] : \deg P < k,\ \#\{i : P(\alpha_i) = y_i\} \geq A\},$$

computed deterministically in time polynomial in n and $\log q$. The field can be far larger than the code, so the running time may depend on q only through the cost of arithmetic in $\mathbb{F}_q$.

Results of this section. All our decoders are deterministic and return the complete list. In the first two bounds below, the rate ρ and the agreement fraction a are fixed above the respective threshold.

- Above Johnson, $\widetilde{O}_{\rho,a}(n \log q)$ bit operations in every characteristic (Theorem 9.3).
- Above the first-order curve, $\widetilde{O}_{\rho,a}(n^3 \log q)$ bit operations for $p > k - 1$ (Corollary 9.9).
- Given an explainer of order d with $\deg_X Q \leq Cn$ and total jet degree at most B_{jet}, $n^{O(d+1)}\widetilde{O}_{d,B_{\text{jet}},C}(\log q)$ bit operations for $p > \max\{k - 1, B_{\text{jet}}\}$, with no positive linear gap between A and k (Theorem 9.11). Combined with the explainers of Theorems 5.7 and 8.1, and with the order-zero decoder when k is bounded, this decodes at every fixed gap from capacity, uniformly over the rate, for $p > k - 1$ (Theorem 9.14).

Outline. Section 2.5 explains why the earlier hidden-derivative decoders cost a power of q however small the list [Jer26, Remark 6.4], and sketches ours at first order. Apart from the cases $k = 1$, $k \leq d$, and bounded length, which Section E.7 decodes directly, our decoders follow four steps (Figure 5), which the subsections below carry out.

1. *Initial values determine the message.* Every close message P solves a stage $T(X, Y_0, \dots, Y_r)$ of the separant chain, of order $r < k$, whose *separant* $T_{Y_r} = \partial T/\partial Y_r$ does not vanish identically along P (Section A.1). Provided the characteristic exceeds $k - 1$ and the jet degree, this separant is nonzero at one of $O(n)$ prescribed centers x_0, and there the first $r + 1$ Taylor coefficients $U_j = P^{[j]}(x_0)$ determine P by Kopparty's recurrence (Lemma A.2).

2. *Keep the initial values as unknowns.* Running the recurrence with U as variables subject to $T(x_0, U) = 0$ gives one symbolic candidate $\widetilde{P}_U(X)$ of degree less than k, whose first $r + 1$ Taylor coefficients at x_0 are U and whose later coefficients are rational functions of U (Section 9.5). At order $r \geq 1$ the candidates form a family of dimension r, a curve at first order.

3. *Agreement leaves finitely many candidates.* At positive order, every close message satisfies at least $t = A - k + 1$ agreement equations $\widetilde{P}_U(\alpha_i) = y_i$ that are not identically zero on its component of the family (Figure 3). At first order each of them cuts the curve in finitely many points (Section 9.6); at order $r \geq 2$, some r of them, with $T(x_0, U) = 0$, have the message's initial values as an isolated solution, and Rojas's algorithm represents these solutions (Section 9.7).

4. *Read the messages off the received values.* We store the finitely many points as a finite algebra and split it at each position i by a gcd into the part where the candidate agrees with y_i and the rest. A part that has agreed at k positions gives the interpolant of those k received values, which we keep if it agrees with y in at least A positions (Section 9.3). This extends the received-value method of Augot–Pecquet and CHK [AP00, CHK26] to our families and extracts no root.

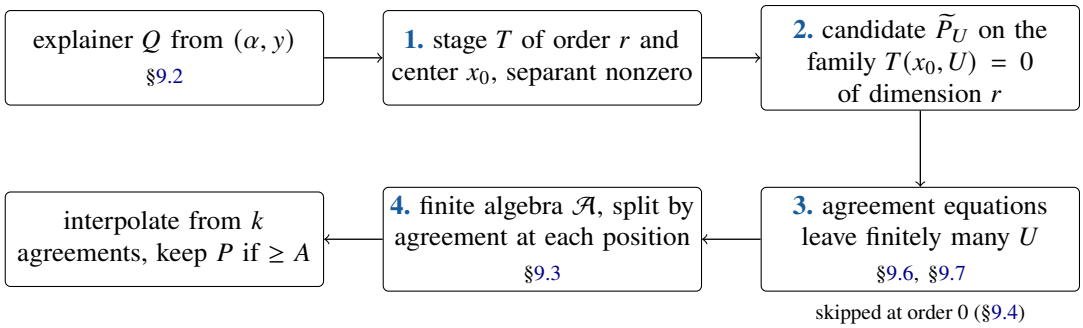

Figure 5: The decoders of this section; numbered boxes are the four steps of the outline.

Core constructions and refinements. These four steps already give a deterministic decoder with cost polynomial in n and $\log q$ for every fixed order and explainer degree. The subsections headed *Refinement* lower this cost, and Section E gives the field arithmetic behind the stated bit bounds.

Bounds used by the decoders. From the earlier sections the decoders use the explainers (Definition 3.8), with the supports in which those sections solve the local constraints of Definition 3.6 and the degree bounds of the resulting explainers (Theorems 5.1, 5.7 and 8.1 and Corollary 8.8), and the separant chain and Taylor lift of step 1, but never the list-size or MCA bounds.

9.1 Arithmetic model and notation

Decoders return each message as its k coefficients, and we count bit operations. The field is presented as in Definition E.1, for example by a polynomial basis $\mathbb{F}_p[t]/(f(t))$, so field operations cost $\widetilde{O}(\log q)$ bit operations. Univariate multiplication, division with remainder, multipoint evaluation, interpolation, and gcd of degree-N polynomials take $\widetilde{O}(N)$ field operations [vzGG13, Chapters 8–11]. The notation $\widetilde{O}$ hides polylogarithmic factors in n and $\log q$, and subscripts list the fixed parameters on which the constant depends. Searches for centers and coordinates test polynomially many prescribed field elements; exhaustive searches occur only in auxiliary fields of size polynomial in n (Section E.2).

The decoders use an explainer Q with $\deg_X Q \leq Cn$ and total degree at most B_{jet} in the jet variables $Y_0, \ldots, Y_d$, and we put $N_X = 1 + \deg_X Q$. Centers x_0 lie in a field $E \supseteq \mathbb{F}_q$ of degree over $\mathbb{F}_q$ bounded in terms of the parameters. For curves we use coordinates (u, v), write $h(u, v)$ for an equation monic in v and $b = \deg_v h$ for its degree in v, and write Δ for the dimension over E of a finite algebra. Throughout, $t = A - k + 1$ is the number of agreements beyond the $k - 1$ that a nonconstant family can satisfy identically.

We write $X = x_0 + Z$ for a Taylor center x_0; Lemma 5.4 writes the increment as T, but here T denotes a stage of the separant chain, and no MCA challenge occurs. The candidate is always denoted $\widetilde{P}$, and C is reserved for the constant in $\deg_X Q \leq Cn$. Table 5 records the notation used in more than one subsection.

9.2 Constructing an explainer

Every decoder in this section starts from an explainer, the first box of Figure 5. The earlier sections prove that a certified support contains one; here we compute it from the received pairs. In the core construction we solve a linear system: we assign one unknown coefficient to each monomial of the support, and each local constraint $\Phi^{d,m}_{\alpha_i, y_i}(Q) = 0$ is linear in these unknowns. The support inequalities put the number of unknowns above the rank of all constraints, so the homogeneous system has a nonzero solution. Gaussian elimination computes one in polynomial time, since the jet degree and order are fixed and the X-degree is $O(n)$.

Symbol	Meaning in Section 9
n, k, A; $t = A - k + 1$	Length, message dimension, agreement threshold; excess
d, r; B_{jet}, C, N_X	Explainer order, stage order; jet-degree bound, X-degree constant, $1 + \deg_X Q$
D	Ambient polynomial degree used by the interpolation support
E	Center field
x_0, Z, U; τ_j	Taylor center, increment, initial Taylor coefficients; rational Taylor coefficients
T, S; H	Stage and separant; specialized equation
u, v, h, b	Monic coordinates, family equation, and $b = \deg_v h$
s, B_0, N_j	Transformed separant, denominator, numerators
$\widetilde{P}$; g_i, e_i	Candidate; cleared and algebra-valued agreement residuals
G	Squarefree base modulus
$\mathcal{A}, \Delta$	Finite candidate algebra and total dimension over E

Table 5: Decoder notation used in more than one subsection. Other symbols are declared at their use.

Support and its jet monomials	Weight	Strict cutoff
Order zero: Y^j, $0 \le j \le \ell$	$j(k-1)$	mA
First order: $Y_0^a Y_1^b$, $b \le B_\partial$, $a + b \le B_{\mathrm{jet}}$	$D(a+b)$	mA
Tuned first order: the same jet monomials	$Da + (D-1)b$	mA
Higher order: $Y_0^{b_0} \prod_{j=1}^{d} Y_j^{b_j}$, $\sum_{j=1}^{d} j b_j \le W$	$D(b_0 + \sum_j b_j)$	mA_0

Table 6: Support data for fixed-word interpolation. The higher-order threshold satisfies $A_0 \le A$, and W bounds the weighted derivative degree $\sum_j j b_j$ in (56).

Refinement: fast module interpolation. For one fixed received word, multiplication by X acts on the truncated Taylor expansion at α_i as a Jordan block. This makes the linear system a polynomial-module interpolation problem. The minimal-basis algorithm of Jeannerod, Neiger, Schost, and Villard [JNSV17, Theorem 1.5] then finds a solution in time quasilinear in n.

That algorithm takes the support as a degree shift. Group terms with the same jet monomial: $Q = \sum_j f_j(X) M_j$. For every support we use, the powers of X that multiply a given M_j form an initial interval, which we describe by a weight w_j and a strict cutoff κ, requiring $\deg f_j + w_j < \kappa$. Table 6 lists these data. In it, D is the ambient polynomial degree of Section 3, which can exceed the actual message degree $k - 1$.

The first-order rows are the supports of (19) and (112); the higher-order row is (56). In the last row only monomials whose weight is below the cutoff need be listed; the certified parameters bound their total jet degree by B_{jet}. In the order-zero row, the strict cutoff mA is the inclusive weighted-degree bound $mA - 1$ used in Section 9.4.

Lemma 9.1 (Fast fixed-word interpolation)**.** *Fix $d, B \ge 0$, $m, \mu \ge 1$, and distinct jet monomials $M_1, \dots, M_\mu \in F[Y_0, \dots, Y_d]$ of total degree at most B. Let $w_1, \dots, w_\mu, \kappa$ be nonnegative integers and put*

$$\mathcal{V} = \left\{ \sum_{j=1}^{\mu} f_j(X) M_j : \deg f_j + w_j < \kappa \text{ whenever } f_j \neq 0 \right\}.$$

Assume $\kappa, w_j = O(n)$. Given $n \ge 1$ received pairs (α_i, y_i) over F, whose abscissae need not be distinct, one can deterministically find a nonzero $Q \in \mathcal{V}$ with $\Phi^{d,m}_{\alpha_i, y_i}(Q) = 0$ for every i, or certify that none exists, in $\widetilde{O}_{\mathrm{params}}(n)$ field operations. The assertion holds in every characteristic, and costs $\widetilde{O}_{\mathrm{params}}(n \log q)$ bits over a finite field of size q presented as in Definition E.1.

Proof of Lemma 9.1. We check the hypotheses of [JNSV17, Theorem 1.5] and count field operations. We write Z for the local Taylor increment of (10). The constraints at the i-th pair take values in

$$\mathcal{R}_i = (F[Z]/Z^m) \otimes_F \operatorname{span}_F\{V^u Y_1^{b_1} \cdots Y_d^{b_d} : u + \textstyle\sum_j b_j \le B\}. \tag{71}$$

The substitution (10) has total degree at most B in $V, Y_1, \ldots, Y_d$ and separates coefficients from monomials:

$$\Phi^{d,m}_{\alpha_i,y_i}(f_j(X)M_j) = f_j(\alpha_i + Z)\Phi^{d,m}_{\alpha_i,y_i}(M_j). \tag{72}$$

Thus the constraints are $F[X]$-linear, with X acting on each copy of $F[Z]/Z^m$ by the Jordan block at α_i. With $e_j = (\Phi^{d,m}_{\alpha_i,y_i}(M_j))_i$ in $\mathcal{R} = \bigoplus_i \mathcal{R}_i$, they form one module relation

$$\sum_{j=1}^{\mu} f_j \cdot e_j = 0. \tag{73}$$

Repeated abscissae keep separate summands and Jordan blocks, so their constraints remain distinct; the minimal-basis algorithm applies to multiplication matrices with repeated eigenvalues. The module has dimension $\sigma = nm\binom{B+d+1}{d+1}$, with its multiplication matrix given in Jordan form, and forming all e_j costs $O_{d,m,B}(n\mu)$ field operations. Apply the minimal interpolation basis algorithm of [JNSV17, Theorem 1.5] with shift $(w_1, \ldots, w_\mu)$, normalized by subtracting its minimum. Its cost is

$$\widetilde{O}\big(\mu^{\omega-1}(\sigma + \xi)\big), \qquad \xi = \sum_j (w_j - \min_h w_h) = O_{\text{params}}(n). \tag{74}$$

Since $\mu \le \binom{B+d+1}{d+1}$ is fixed, this is soft-linear in n. The returned basis has a row $(f_j)_j$ minimizing the shifted degree $\max_j(\deg f_j + w_j)$ among nonzero relations. Indeed, every relation is a polynomial combination of the basis rows, and the shifted leading coefficient matrix of the basis has full row rank. So leading terms cannot cancel: the shifted degree of a combination is the maximum, over rows with a nonzero multiplier, of the row's shifted degree plus the degree of its multiplier. Hence no nonzero relation has smaller shifted degree than the minimal row. If this minimum is below κ, then $Q = \sum_j f_j M_j$ is the required interpolant; otherwise none exists. No step divides by the characteristic. □

9.3 Agreement recovery in finite algebras

Every decoder below ends with step 4 of the outline, so we make it precise first, in Lemma 9.2. Each later subsection produces a finite commutative E-algebra $\mathcal{A}$ together with a candidate $\widetilde{P}(X) \in \mathcal{A}[X]_{<k}$, a polynomial whose coefficients lie in $\mathcal{A}$. The individual candidates are its specializations: each E-algebra homomorphism $\varphi : \mathcal{A} \to \overline{E}$, which we call a *geometric point*, gives a polynomial $\varphi(\widetilde{P}) \in \overline{E}[X]_{<k}$. If $\mathcal{A} = E[u]/G(u)$ with G squarefree, for example, the geometric points evaluate u at the roots of G. We want the base-field polynomials among the $\varphi(\widetilde{P})$ that agree with y in at least A positions. The routine splits $\mathcal{A}$ into factors and records, for each factor, the positions at which every geometric point of that factor agrees with y (Figure 6).

The algebras are *towers*

$$\mathcal{A} = (E[u]/G(u))[v]/h(u,v), \qquad \widetilde{P}(X) \in \mathcal{A}[X]_{<k}, \tag{75}$$

where G is monic and squarefree and h is monic of bounded positive v-degree b, not necessarily squarefree on the fibers. It has dimension $\deg G \cdot b$, and the univariate case $E[u]/G$ is $h = v$.

Repeated roots require care. For example, v vanishes at the only point of $E[v]/v^2$, although it is nonzero in the algebra. We want that point on the agreeing side of the split. A residual e is *nilpotent* on a factor if some power of it is zero there, and a *unit* if it has an inverse there. We split each algebra into a factor on which e is nilpotent, containing exactly the points where e vanishes, and a factor on which it is a unit, using a gcd as follows.

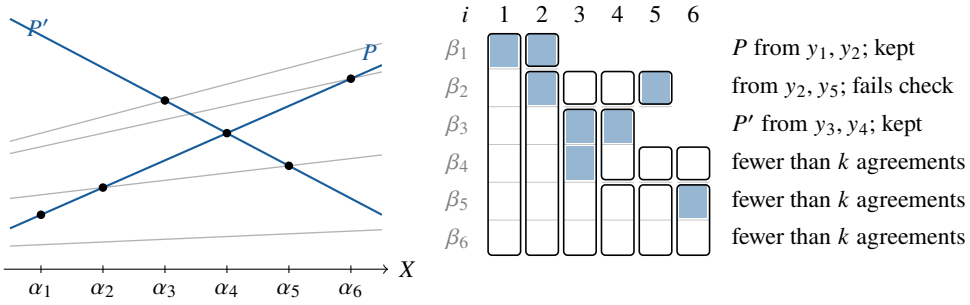

Figure 6: Agreement recovery (Algorithm 1) in $\mathcal{A} = E[u]/G$, with G squarefree of degree 6, $k = 2$, $A = 3$ and $n = 6$. Left: the candidate $\widetilde{P} \in \mathcal{A}[X]$ holds one line for each root β_j of G, drawn against the received word (schematic, since the roots may lie in extensions of $\mathbb{F}_q$). Right: row j marks where the line at β_j agrees with y, up to the column at which its factor is removed. Arithmetic in $\mathcal{A}$ acts on all rows at once, and no root is computed: at column i, each current factor g of G (outlined) splits into $\gcd(g, e_i)$, the rows that agree, and the rest, where $e_i = \widetilde{P}(\alpha_i) - y_i$. A factor with k agreements is removed, and the interpolant of its received values is kept if it has at least A agreements.

Algorithm 1 Exact recovery from finite algebraic candidate families

Require: finite algebras $\mathcal{A}$, each with a candidate $\widetilde{P}(X) \in \mathcal{A}[X]_{<k}$, and α, y, k, A

1: Evaluate every residual $e_i = \widetilde{P}(\alpha_i) - y_i$ by shared multipoint evaluation.
2: Give each algebra an empty agreement history; initialize an empty output set.
3: **for** $i = 1, \ldots, n$ and each current factor $\mathcal{R}$ with history J **do**
4: Transfer e_i to $\mathcal{R}$ and split into its nilpotent and unit factors (Lemma E.5).
5: Keep J on unit factors; append i on nilpotent factors, sharing history prefixes.
6: For each factor with k agreements, interpolate P from those received pairs and remove the factor.
7: Insert P into the output set if it has at least A agreements with the full word.
8: **end for**
9: Return the output set as duplicate-free coefficient vectors.

Computing the nilpotent and unit factors. Let e be a residual on a factor $\mathcal{R} = (E[u]/G)[v]/\bar{h}$ with $b = \deg_v \bar{h}$. Raising e to the fiber degree before taking the gcd captures each root where e vanishes with its full multiplicity in $\bar{h}$. To see this, work over an algebraic closure of a residue field of the squarefree base and write

$$\bar{h} = \prod_{\beta} (v - \beta)^{m_\beta}, \qquad m_\beta \le b. \tag{76}$$

If $e(\beta) = 0$, then $(v - \beta)^b$ divides e^b, so e^b contains the full primary factor $(v - \beta)^{m_\beta}$; otherwise e^b is coprime to that factor. Hence

$$h_{\mathrm{nil}} := \gcd(\bar{h}, e^b \bmod \bar{h}) = \prod_{e(\beta)=0} (v - \beta)^{m_\beta}. \tag{77}$$

The quotient by h_{nil} is the nilpotent factor and the quotient by $\bar{h}/h_{\mathrm{nil}}$ is the unit factor; we omit constant moduli. The Chinese remainder theorem shows that the two factors partition the points and that their dimensions add up to $\dim_E \mathcal{R}$.

The Euclidean algorithm in v must invert leading coefficients $a \in E[u]/G$, and the base need not be a field. If a is a zero divisor, split G into $G_0 = \gcd(G, \widetilde{a})$ and $G_1 = G/G_0$, where $\widetilde{a} \in E[u]$ represents a. Then a vanishes modulo G_0, where we delete the leading term, and is invertible modulo G_1, where we invert it

by extended gcd. This is the dynamic evaluation (D5) principle, whose complexity Dahan, Moreno Maza, Schost, and Xie analyze [DMSX05]; its depth is $O_b(1)$ since v-degrees are bounded, and it preserves the total base degree $\deg G_0 + \deg G_1 = \deg G$. Bounded powering, fast gcds, and extended gcds therefore cost $\widetilde{O}_b(\dim_E \mathcal{R})$ operations in E per split.

Later subsections also discard the points where the separant s of a family vanishes. This is the same split applied to s, keeping only the unit factor

$$h_{\mathrm{keep}} = h/\gcd(h, s^b \bmod h). \tag{78}$$

It retains the full primary factors where $s \neq 0$, including their nilpotents, in every characteristic.

Lemma 9.2 (Agreement recovery from a finite family)**.** *Let $1 \le k \le A \le n$ and let the received pairs lie in $\mathbb{F}_q$. Let a collection of towers (75) have fiber degrees at most a fixed b and total dimension Δ. Once its coefficient representations are available, exactly the base-field polynomials represented by its geometric points that have at least A agreements can be recovered deterministically in $\widetilde{O}_b((n+k)\Delta)$ operations in E, and $\widetilde{O}(k\Delta)$ base-field coefficient comparisons. If the collection covers every close message, this is the full list.*

Proof. By (77), the Chinese remainder theorem, and the dynamic-evaluation splits of the base, the split in line 4 partitions the geometric points of each current factor into those at which e_i vanishes, which lie on nilpotent factors, and the others, which lie on unit factors; it preserves total dimension. Since line 5 records i exactly on the nilpotent factors, every geometric point of a current factor agrees with y at every position in that factor's history. A geometric point whose candidate has at least A agreements therefore remains on a current factor until the history of its factor has k entries. Its candidate is then a degree-$< k$ polynomial agreeing with y at k distinct evaluation points, so Vandermonde invertibility identifies it with the interpolant $P \in \mathbb{F}_q[X]_{<k}$ of line 6, which is computed from received values alone. This conclusion holds at the geometric points only: $\widetilde{P} - P$ can still have nilpotent coefficients in that factor. The full check in line 7 discards interpolants with fewer than A agreements.

For the cost, the dimensions of the current factors sum to at most Δ, since every split is a direct-product decomposition and completed factors are removed. Each position therefore costs $\widetilde{O}_b(\Delta)$, and at most Δ interpolants need full-word checks. Lemma E.5 gives the costs of the shared evaluations and transfers and of the agreement histories. □

9.4 Order zero: a finite family at a regular center

At order zero the outline needs only steps 1, 2, and 4: the Guruswami–Sudan equation $Q(X, P) = 0$ already gives a finite family, since $P(x_0)$ is a root of the univariate polynomial $Q(x_0, Y)$, so no agreement cuts are needed. We keep all roots of $Q(x_0, Y)$ together in one finite algebra, lift them symbolically to power series, and let Lemma 9.2 produce the messages, so no root of $Q(x_0, Y)$ is ever computed. Section 1.2 compares the resulting decoder with earlier Johnson decoders.

Theorem 9.3 (Deterministic near-linear decoding above Johnson)**.** *Fix $0 < \rho < 1$ and $\sqrt{\rho} < a \le 1$. Over every finite field $\mathbb{F}_q$ presented as in Definition E.1, with $q \ge n$ and unrestricted characteristic, the complete degree-$< k$ agreement list at threshold an can be computed deterministically in $\widetilde{O}_{\rho,a}(n \log q)$ bit operations, for every prescribed set of n distinct evaluation points and $1 \le k < n$ with $(k-1)/n \le \rho$. The dependence on the positive Johnson gap is contained in the constant.*

The algorithm interpolates with multiplicity m and Y-degree at most ℓ, positive integers depending only on ρ, a such that

$$\sum_{j=0}^{\ell} (ma - j\rho)_+ > m(m+1)/2, \qquad N_* := 2\ell m + 1, \tag{79}$$

where $(x)_+ = \max\{x, 0\}$; the proof shows that they exist.

Algorithm 2 Zeroth-order decoding at a fixed Johnson gap

Require: data of Theorem 9.3, $A = \lceil an \rceil$, and (79)

1: If $k = 1$, use frequency counting; if $n < N_*$, use subset interpolation (Section E.7).
2: Interpolate Q with multiplicity m, Y-degree $\le \ell$, and $(1, k-1)$-weighted degree $\le mA - 1$ (Lemma 9.1).
3: Remove X-content and take its separable part T, which keeps every message factor (Lemma E.7); return $\varnothing$ if $\deg_Y T = 0$.
4: Set $\Omega_0 = \mathrm{lc}_Y(T)\,\mathrm{Res}_Y(T, T_Y)$ and choose $x_0 \in E$ with $\Omega_0(x_0) \neq 0$ (Section E.2).
5: Set $h(z) = T(x_0, z)/\mathrm{lc}_Y(T)(x_0)$ and $\mathcal{A} = E[z]/h$.
6: Lift z modulo Z^k to $\widetilde{P}_{x_0}(Z)$ solving $T(x_0 + Z, \widetilde{P}_{x_0}) = 0$; set $\widetilde{P}(X) = \widetilde{P}_{x_0}(X - x_0)$ (Lemma E.8).
7: Return agreement recovery on $(\mathcal{A}, \widetilde{P})$ (Algorithm 1).

Proof. **Interpolation (step 2).** For $\ell = \lfloor ma/\rho \rfloor$, the sum in (79) is $m^2a^2/(2\rho) + O_{\rho,a}(m)$; $a^2 > \rho$ therefore gives the inequality for large m. The number of monomials X^iY^j with $j \le \ell$ and $i + j(k-1) \le mA - 1$ is

$$\sum_{j=0}^{\ell} \max\{mA - j(k-1), 0\} \ge n \sum_{j=0}^{\ell} (ma - j\rho)_+ > nm(m+1)/2.$$

It exceeds the number of constraints, so the interpolant is nonzero. For a close message P, $Q(X, P)$ has at least mA zeros counted with multiplicity and degree at most $mA - 1$, hence is zero [GS99]. The exceptional cases in step 1 are exact since $A \ge k$, and cost $\widetilde{O}_{\rho,a}(\log q)$ when $n < N_*$.

A finite family (steps 3–6). We first remove the X-content of Q, which removes no message factor $Y - P(X)$. In small characteristic, making Q separable takes two different operations, as the following two examples show. For $Q = (Y - P(X))^p$, division by $\gcd(Q, Q_Y)$ would give 1 and lose the message. Conversely, $Y^p - X$ is reduced as a bivariate polynomial but inseparable in Y. Following the distinction in Lecerf [Lec08], we first take the bivariate radical Q_{rad}, then form the primitive part of $Q_{\mathrm{rad}}/\gcd(Q_{\mathrm{rad}}, (Q_{\mathrm{rad}})_Y)$.

The radical retains each irreducible factor once. Since it is reduced, the second gcd contains exactly the factors with zero Y-derivative. Every message factor $Y - P(X)$ has derivative 1, so it is retained. Gauss's lemma identifies the remaining factors with separable factors over $\mathbb{F}_q(X)$. This proves the claim of Lemma E.7 that T retains every factor $Y - P(X)$; its inverse-Frobenius implementation and arithmetic bound are given in Section E.4.

The resulting T has $\deg_Y T \le \ell$ and $\deg_X T \le mA - 1 < mn$; if its Y-degree is zero, no message factor remains and the list is empty. Otherwise separability makes $\Omega_0 = \mathrm{lc}_Y(T)\,\mathrm{Res}_Y(T, T_Y)$ nonzero, with $\deg \Omega_0 \le 2\ell mn$. If $q > \deg \Omega_0$, take $E = \mathbb{F}_q$. Otherwise take the quadratic extension $E = \mathbb{F}_{q^2}$ of Section E.2, which is large enough because $n \ge N_* = 2\ell m + 1$ gives $q^2 \ge n^2 > 2\ell mn \ge \deg \Omega_0$. In both cases E has at least $\deg \Omega_0 + 1$ elements, so testing that many distinct centers finds x_0 with $\Omega_0(x_0) \neq 0$.

At this center, the leading coefficient is nonzero and the resultant specializes to a nonzero value. Thus $h(z)$ has the same degree as T in Y and is coprime to $T_Y(x_0, z)$, so h is squarefree. A Bézout identity gives $a(z)h(z) + b(z)T_Y(x_0, z) = 1$; in $\mathcal{A} = E[z]/h$ this makes $b(z)$ an inverse of the separant. Over an algebraic closure, $\mathcal{A}$ is a product of copies of $\overline{E}$, one for each root of h, but the algorithm computes only in the quotient algebra.

Simultaneous lifting. In the order-zero recurrence every new coefficient is multiplied by the same unit $T_Y(x_0, z)$. It therefore determines a unique series $\widetilde{P}_{x_0} \in \mathcal{A}[Z]/(Z^k)$ with $\widetilde{P}_{x_0}(0) = z$ and $T(x_0 + Z, \widetilde{P}_{x_0}) \equiv 0 \pmod{Z^k}$, without dividing by a Taylor index or a factorial. For a close message P, the value $P(x_0)$ is a root of h. Applying the point $z \mapsto P(x_0)$ to this recurrence gives the unique lift $P(x_0 + Z) \bmod Z^k$. Since $\deg P < k$, shifting the degree-$< k$ representative back to X represents P exactly.

Other points can give truncations that solve no polynomial equation; agreement recovery checks the actual candidate and discards those below the threshold.

Refinement: near-linear time. JNSV interpolation (Lemma 9.1) costs $\widetilde{O}_{\rho,a}(n)$ field operations. All remaining polynomial degrees are $O_{\rho,a}(n)$ and all algebra dimensions at most ℓ. Bounded-size subresultants, batched center evaluation, division-free Taylor shifts, and agreement recovery each fit the same bound. The lift of step 6 is the relaxed Taylor recurrence of Lemma E.8, based on van der Hoeven's relaxed multiplication [vdH02]; it lifts all roots of h together in $\mathcal{A}$ with $\widetilde{O}_{\rho,a}(k)$ operations in E. Lemma E.7 and Section E.2 bound inverse Frobenius and field construction, giving $\widetilde{O}_{\rho,a}(n \log q)$ bits in total. □

Corollary 9.4 (Zeroth-order decoding at a fixed gap above Johnson)**.** *Fix $\eta_0 > 0$. Let $\mathbb{F}_q$, the evaluation points, and $1 \le k < n$ be as in Theorem 9.3, and write $\rho = k/n$. The complete degree-$<k$ agreement list at threshold an with $a = \sqrt{(k-1)/n} + \eta_0$ can be computed deterministically in $\widetilde{O}_{\rho,\eta_0}(n \log q)$ bit operations, in every characteristic.*

Proof. Since $\sqrt{\rho} - \sqrt{\rho - 1/n} \le 1/(n\sqrt{\rho})$, lengths $n \ge 2/(\eta_0\sqrt{\rho})$ have gap at least $\eta_0/2$ above $\sqrt{\rho}$. Apply Theorem 9.3 at that smaller threshold and retain only outputs with at least an agreements; shorter lengths use bounded subset interpolation. □

9.5 Taylor families at positive order

This subsection carries out steps 1 and 2 of the outline at positive order and writes down the agreement equations that the next two subsections use to make the family finite. Throughout we assume $p > \max\{k-1, B_{\text{jet}}\}$; if only $p > k-1$ holds, then $k \le p \le B_{\text{jet}}$, and Lemma E.9 computes the list for this bounded message dimension with the order-zero decoder, in every characteristic.

Choosing the equation and the center. Run the separant chain (91) from $T_0 = Q$: while T_j depends on a jet variable, let r_j be the largest index of such a variable, put $S_j = \partial T_j/\partial Y_{r_j}$, and set $T_{j+1} = S_j$. Each T_j is a *stage* of the chain. Assign a close message P to the first stage j with $T_j(X, P, \dots, P^{[r_j]}) = 0$ but $S_j(X, P, \dots, P^{[r_j]}) \ne 0$ as a polynomial in X. Such a stage exists, since each step lowers the total jet degree and the chain ends in a nonzero polynomial in X, which has no polynomial solutions. The nonzero specialization of S_j has degree at most a computable $\Theta_j = O_{\text{params}}(n)$ by (86), so it is nonzero at one of any $\Theta_j + 1$ distinct centers; Section E.2 constructs them.

The regular family at one center. Fix a stage $T(X, Y_0, \dots, Y_r)$ of total jet degree at most B_{jet} and a center x_0. A message is *regular* at x_0 if its separant is nonzero there; its initial jet $U = (U_0, \dots, U_r)$, the vector of its first $r+1$ Taylor coefficients at x_0, then satisfies

$$H(U) := T(x_0, U) = 0, \qquad S(U) := T_{Y_r}(x_0, U) \ne 0.$$

The nonzero separant S lets the equation determine the later Taylor coefficients. At first order, for example, the coefficient of Z in $T(x_0 + Z, U_0 + U_1 Z + \tau_2 Z^2 + \cdots, U_1 + 2\tau_2 Z + \cdots)$ is $T_X(x_0, U) + T_{Y_0}(x_0, U)U_1 + 2S(U)\tau_2$, and setting it to zero determines τ_2. Kopparty's recurrence continues this calculation to every later coefficient [Kop15]. Because we keep the initial values as variables, as in earlier symbolic solution families [Zhe26, Jer26], each coefficient is a rational function of U rather than a field element. With $r < k \le n$ and $p > \max\{k-1, B_{\text{jet}}\}$, truncating at the actual message bound k, even if interpolation used a larger ambient degree bound D, gives

$$\widetilde{P}_U(x_0 + Z) = \sum_{j<k} \tau_j(U) Z^j, \qquad \tau_j(U) = \frac{N_j(U)}{B_0(U)}, \qquad \tau_j(U) = U_j \ (j \le r). \tag{80}$$

To see the triangular structure in every order, write $Y(Z) = \sum_{j<k} \tau_j Z^j$. For $1 \le m < k - r$, the next coefficient occurs as

$$[Z^m]T(x_0 + Z, Y, Y^{[1]}, \dots, Y^{[r]}) = \binom{r+m}{r} S(U)\tau_{r+m} + R_m(U, \tau_{r+1}, \dots, \tau_{r+m-1}). \tag{81}$$

Indeed, τ_{r+m} first enters the r-th Hasse derivative at degree m; all other occurrences have higher Z-degree. The linear term in that variable has coefficient $S(U)$, giving the displayed multiplier. It is invertible when $S \neq 0$, because $r + m < k < p + 1$. Solving successively therefore constructs the rational coefficients in polynomial time.

Lemma A.3, applied with $K = k$, bounds the degrees of these rational coefficients: its proof solves (81) inductively and gives, for $1 \le m < k - r$,

$$\tau_{r+m} = \frac{\widehat{N}_m(U)}{S(U)^{2m-1}}, \qquad \deg \widehat{N}_m \le (2m-1)(B_{\text{jet}} - 1) + 1.$$

Thus $B_0 = S^{2k}$ is a common denominator for all τ_j with $j < k$, and each numerator $N_j = \tau_j B_0$ has degree at most $L = 1 + 2k(B_{\text{jet}} - 1)$; for $j \le r$ the numerator is $U_j S^{2k}$. Since the first $r + 1$ coefficients are the coordinates U, distinct initial jets give distinct candidates.

Every close message regular at x_0 equals its own truncation. Other points of the family need not give solutions of T, but a truncation with at least A agreements still belongs to the list. Hence we impose no equations on the coefficients of degrees $k, \dots, D$.

Equations for the family. To compute with this family we work with $H(U) = T(x_0, U)$, the equation satisfied by the initial jets. Three tests discard centers at which no jet is regular: $H = 0$, since then $S = \partial H/\partial U_r$ also vanishes; H a nonzero constant, since then no jet satisfies $H = 0$; and H_* constant, where H_* is the reduced product of the components of $H = 0$ on which S is not identically zero. We compute it without factoring: since $p > B_{\text{jet}}$,

$$H_{\text{rad}} = H/\gcd(H, H_{U_0}, \dots, H_{U_r}), \qquad H_* = H_{\text{rad}}/\gcd(H_{\text{rad}}, S).$$

The first quotient removes repeated factors and the second removes components on which the separant is identically zero.

If H_* is nonconstant, of degree b, we make it monic in one coordinate. Choose $w \in E^{r+1}$ at which the top homogeneous part of H_* is nonzero, by testing a product set of size depending only on r and B_{jet}. An invertible linear change of coordinates $U = M(u, v)$, with $u = (u_1, \dots, u_r)$ and w as the last column of M, followed by scaling, turns H_* into

$$h(u, v) = v^b + h_{b-1}(u)v^{b-1} + \cdots + h_0(u), \qquad \deg h_i \le b - i. \tag{82}$$

Monicity lets us treat the family as finitely many values of v above each u, even where these values collide. Regularity is still decided by the original separant, now $s(u, v) = S(M(u, v)) \bmod h$, which in general differs from h_v. We use N_j and B_0 for the preceding polynomials after substitution by M and reduction modulo h. Their v-degrees are less than b, and the bound $\deg h_i \le b - i$ shows that this reduction does not increase total degree.

Agreement equations. Step 3 cuts the family by agreement with the received word, so we write agreement as a polynomial condition on the family. With the common denominator $B_0 = s^{2k} \bmod h$, agreement at position i is the polynomial equation

$$g_i(u, v) := \sum_{j<k} N_j(u, v)(\alpha_i - x_0)^j - y_i B_0(u, v) = 0, \tag{83}$$

Algorithm 3 RegularTaylorFamily: the regular family at a center

Require: Stage T of order $r < k$, center $x_0 \in E$, $p > \max\{k-1, B_{\text{jet}}\}$, and the field-size condition of Lemma 9.5

Ensure: A chart $(h, s, B_0, N_0, \dots, N_{k-1})$ representing every solution regular at x_0, or an empty chart

1: Form $H = T(x_0, U)$ and $S = T_{Y_r}(x_0, U)$; return an empty chart if $H = 0$ or H is constant.
2: Compute H_* by the two gcd quotients above; return an empty chart if it is constant.
3: Choose monic coordinates $U = M(u, v)$ and reduce $S(M(u, v))$ modulo h to obtain s.
4: Run the triangular recurrence, directly or in the truncated parameter algebra of Lemma 9.5.
5: Return h, s and the reduced numerator polynomials with common denominator $B_0 = s^{2k} \bmod h$.

of total degree at most L. The formulas divide only by the separant, so they hold at every point with $s \neq 0$, including points where several values of v above the same u coincide. We keep the condition $s \neq 0$ throughout, because clearing the common denominator B_0 can create spurious agreements along $s = 0$.

The complete control flow. Algorithm 4 assembles the four steps of the outline into one decoder. For each stage of the separant chain and each of its candidate centers (step 1), RegularTaylorFamily (Algorithm 3) returns the *chart* of the stage at x_0 (step 2): the monic equation h, the transformed separant s, and the coefficient formulas N_j/B_0. A *chart point* is a point with $h = 0$ and $s \neq 0$, and a *close point* is a chart point whose candidate has at least A agreements. An empty chart means that no jet is regular at this stage and center.

How we carry out step 3 depends on the order of the stage. At order zero the chart is already finite, and we skip step 3. At order one, CurveCandidates (Section 9.6) cuts the curve by agreement equations, and at order $r \geq 2$, IsolatedSystemCandidates (Section 9.7) solves square systems. Both return finite algebras with candidates, and RecoverAgreement carries out step 4, returning polynomials over $\mathbb{F}_q$. An interpolation failure differs from an empty chart: the given support then has no nonzero interpolant, so it is not a certificate, and the decoder reports a certificate failure, never an empty list.

9.5.1 Refinement: fast construction of the coefficient formulas

The direct recurrence already runs in polynomial time. Lemma 9.5 constructs h, s, and the numerators in $\widetilde{O}_{r,B_{\text{jet}}}(k^{r+1} + \deg_X T)$ operations in E; multipoint evaluation of the scalar coefficient polynomials then forms all n agreement equations in $\widetilde{O}_{r,B_{\text{jet}}}(n^{r+1})$ operations. For the first bound, we compute in a finite ring where the separant is a unit and then recover the global numerator polynomials from their degree bounds.

Lemma 9.5 (Fast regular Taylor family)**.** *Fix r and B_{jet}. Let E be a finite field of characteristic p, let $0 \leq r < k$, and let*

$$T \in E[X, Y_0, \dots, Y_r], \qquad \deg_{Y_0,\dots,Y_r} T \leq B_{\text{jet}}, \qquad \deg_X T \leq D_X.$$

Fix $x_0 \in E$, and put

$$H(U) = T(x_0, U_0, \dots, U_r), \qquad S(U) = T_{Y_r}(x_0, U_0, \dots, U_r).$$

Suppose $p > \max\{k-1, B_{\text{jet}}\}$. Let H_* be the reduced union of the components of $H = 0$ on which S is not identically zero, and suppose H_* is nonconstant. There is a linear coordinate map

$$(U_0, \dots, U_r) = M(u_1, \dots, u_r, v)$$

and polynomials $B_0, N_0, \dots, N_{k-1} \in E[u, v]$ such that:

Algorithm 4 Complete decoding from a given explainer or a certified support

Require: Data and characteristic hypotheses of Theorem 9.11, with Q given or a certified support

Ensure: Exactly the close list as coefficient vectors; report CERTIFICATE-FAILURE if support interpolation fails

Decode $k = 1$, $k \le d$, and instances of bounded length by Section E.7.
If needed, interpolate Q by Lemma 9.1; report CERTIFICATE-FAILURE if no nonzero interpolant exists.
Initialize an empty collection of finite families; set $T = Q$.
while T depends on a jet variable **do**
    Set $r = \max\{j : T \text{ depends on } Y_j\}$ and $S = T_{Y_r}$.
    Choose $\Theta + 1$ distinct centers, where Θ bounds the degree of every nonzero specialized separant.
    for each chosen center x_0 **do**
        Obtain a chart by RegularTaylorFamily(T, x_0, k); skip an empty chart.
        Form the agreement polynomials g_i of (83) by multipoint evaluation.
        if $r = 0$ **then**
            Add $E[v]/h$ after removing $s = 0$, with the candidate given by the chart's coefficient formulas.
        else if $r = 1$ **then**
            Add the finite algebras from CurveCandidates with the threshold filter (Algorithm 5).
        else
            Add the finite algebras from IsolatedSystemCandidates (Algorithm 6).
        end if
    end for
    Set $T = S$ and continue to the next stage.
end while
Return RecoverAgreement on all finite families (Algorithm 1).

(i) *their v-degrees are less than $b := \deg H_*$ and their total degrees are at most*

$$L := 1 + 2k(B_{\mathrm{jet}} - 1);$$

(ii) *in the coordinate ring of H_*, $B_0 = S(M(u,v))^{2k}$. In its localization at S, the coefficient formulas give*

$$\widetilde{P}_{u,v}(x_0 + Z) = \sum_{j<k} \frac{N_j(u,v)}{B_0(u,v)} Z^j;$$

(iii) *this series recovers every regular degree-$< k$ polynomial solution of the differential equation defined by T whose initial jet lies on $H_* = 0$; here regular means that the same jet satisfies $S \neq 0$.*

If $|E|$ exceeds a constant depending only on r and B_{jet}, the coordinate map and all the displayed polynomials are constructed deterministically in

$$\widetilde{O}_{r,B_{\mathrm{jet}}}(k^{r+1} + D_X)$$

operations in E, with $O_{r,B_{\mathrm{jet}}}(k^{r+1})$ output field coefficients and no extension of E.

Proof. **Monic coordinates.** Since $p > B_{\mathrm{jet}}$, the bounded-degree gcd quotients above compute H_*, and the product-set search above gives M and the monic equation (82).

Making the separant invertible. We cannot divide by $s = S(M(u,v))$ in $E[u,v]/h$, so we expand around a point where it is a unit. The nonzero $\mathrm{Res}_v(h,s)$ has bounded degree, so a bounded product set contains $a \in E^r$ where it is nonzero. For $N_{\mathrm{par}} = L+1$, use

$$C_{\mathrm{par}} = E[\epsilon_1,\dots,\epsilon_r]/(\epsilon_1^{N_{\mathrm{par}}},\dots,\epsilon_r^{N_{\mathrm{par}}}), \qquad R = C_{\mathrm{par}}[v]/h(a+\epsilon,v).$$

This is free of rank b over C_{par}. A Bézout identity on the fiber at a inverts s there; Newton iteration extends the inverse through the nilpotent ideal (ϵ). This only requires s to be a unit, even when $h(a,v)$ has repeated roots.

Taylor recurrence and cost. Shift T to $x_0 + Z$ by division-free linear composition in $\widetilde{O}(D_X)$ operations, valid even for $D_X \geq p$. Apply Lemma E.8 over R with center x_0 and initial jet $U = M(a+\epsilon,v)$. The constant residual $T(x_0,U)$ is zero modulo h, the separant $T_{Y_r}(x_0,U)$ is s, and s and the binomial coefficients are units. By the uniqueness in that lemma, its output is the Taylor expansion of the family. Counting operations in E: since $\dim_E R = bN_{\mathrm{par}}^r = O_{r,B_{\mathrm{jet}}}(k^r)$, Kronecker substitution in $Z,\epsilon_1,\dots,\epsilon_r$ gives $M_R(\ell) = \widetilde{O}_{r,B_{\mathrm{jet}}}(\ell k^r)$, and with boundedly many jet monomials the recurrence costs $\widetilde{O}_{r,B_{\mathrm{jet}}}(k^{r+1})$.

Recovering the coefficient formulas. By Lemma A.3 with $K = k$, multiplying the Taylor coefficients by S^{2k} gives polynomials of degree at most L, and reduction modulo the monic h preserves this bound. Compute s^{2k} and its products with the Taylor coefficients in R. They are the images of B_0, N_j under the expansion $f(u) \mapsto f(a+\epsilon) \bmod (\epsilon_1^{N_{\mathrm{par}}},\dots,\epsilon_r^{N_{\mathrm{par}}})$, which is injective on polynomials of degree at most $L < N_{\mathrm{par}}$. Division-free shifts back to u recover B_0, N_j exactly, in $\widetilde{O}_{r,B_{\mathrm{jet}}}(k^{r+1})$ operations. The identities therefore hold in $(E[u,v]/h)[s^{-1}]$, hence at every point with $s \neq 0$, including ramified fibers. Each operation over the decoder's constant-degree extensions costs $\widetilde{O}(\log q)$ bits. □

9.6 First-order recovery from a curve

At first order the initial values of the regular family lie on a curve. This subsection carries out step 3 of the outline on it and proves Theorem 9.6.

Theorem 9.6 (First-order reconstruction with an explicit agreement gap)**.** *Fix $B_{\mathrm{jet}}, C \geq 1$. Let $\mathbb{F}_q$ be presented as in Definition E.1, with $q \geq n$, $1 \leq k < n$, $k \leq A \leq n$, and $t = A - k + 1$. If $k \geq 2$, assume $p > \max\{k-1, B_{\mathrm{jet}}\}$. Given distinct evaluation points $\alpha_i \in \mathbb{F}_q$ and a word $y \in \mathbb{F}_q^n$, suppose a given nonzero $Q \in \mathbb{F}_q[X,Y_0,Y_1]$ has $\deg_X Q \leq Cn$, $\deg_{Y_0,Y_1} Q \leq B_{\mathrm{jet}}$, and satisfies $Q(X,P,P') = 0$ for every degree-$< k$ polynomial with at least A agreements. Then the exact list, as duplicate-free coefficient vectors, is computable deterministically in*

$$\widetilde{O}_{B_{\mathrm{jet}},C}\left(\left(n^2k + \frac{n^3k}{t}\right)\log q\right)$$

bit operations. In particular, for $t \geq \varepsilon n$ the bound is $\widetilde{O}_{B_{\mathrm{jet}},C}((1+\varepsilon^{-1})n^2k\log q)$. The same bound holds when Q is not given but computed from a fixed certified interpolation support by Lemma 9.1. The range $k \leq A$ includes $A = k$, where $t = 1$, only because Q is given: the explainers constructed in the earlier sections have degree bounds that depend on the gap between A and k. Over a prime field no characteristic hypothesis is needed: $p = q \geq n > k-1$, and $p \leq B_{\mathrm{jet}}$ forces $n \leq B_{\mathrm{jet}}$.

9.6.1 Core construction: cutting the curve by agreement

Fix a center that passes the tests of Section 9.5, and write $h(u,v) = 0$ for the curve of initial values. The candidates vary along every component: the first two Taylor coefficients are $\tau_0 = U_0$ and $\tau_1 = U_1$, and they

Algorithm 5 CurveCandidates: finite algebras from a regular curve

Require: Order-one chart (h, s, B_0, N_j) at center x_0 from RegularTaylorFamily, received pairs, $k \le A$, and a choice of the union or threshold filter

Ensure: Finite towers with candidate $\widetilde{P}$ whose geometric points include every close point

1: Remove $\gcd_{E(u)[v]}(h, s)$ from h; return no families if $h = 1$.
2: Form $g_i = \sum_{j<k} N_j(\alpha_i - x_0)^j - y_i B_0$ for every i (83).
3: Reduce B_0, N_j, g_i modulo h and compute each lowest nonzero coefficient ν_i of $\mathrm{Res}_v(h, \zeta - g_i)$.
4: For the union filter set $G = \mathrm{rad}(\prod_i \nu_i)$; for the threshold filter set $G = \mathrm{rad}(\nu_{\ge t})$, $t = A - k + 1$.
5: Form $(E[u]/G)[v]/h$; return no families if $G = 1$.
6: Remove $s = 0$ by (78); invert B_0 on each retained factor.
7: Return these towers with $\widetilde{P}(x_0 + Z) = \sum_{j<k}(N_j/B_0)Z^j$.

recover (u, v) because M is invertible. Since k agreements determine a polynomial of degree less than k, at most $k - 1$ agreement equations hold identically on any one component.

Call an agreement *proper* on a component when its equation is not identically zero there; every close point has at least $t = A - k + 1$ proper agreements. A proper equation cuts its component in finitely many points, but its resultant with the whole curve need not detect them. For example,

$$h = v(v - 1), \qquad g = uv, \qquad \mathrm{Res}_v(h, g) = 0.$$

The equation holds identically on $v = 0$ and has a proper intersection at $u = 0$ on $v = 1$. Introducing a variable ζ for the value of g gives $\mathrm{Res}_v(h, \zeta - g) = \zeta(\zeta - u)$; its lowest nonzero ζ-coefficient detects this proper intersection.

Detecting proper agreements. For each position i we want a polynomial in u that vanishes at the u-coordinate of every proper agreement at i and ignores the components on which g_i vanishes identically. First remove $\gcd_{E(u)[v]}(h, s)$ from h, which discards any component on which the separant vanishes identically, and return the empty family if $h = 1$. The agreement equation at position i is the residual numerator $g_i = B_0(\widetilde{P}(\alpha_i) - y_i)$ of (83). Reduce B_0 and each g_i modulo the monic h before forming any resultant. This preserves their restrictions to the curve, makes their v-degrees less than b, and keeps the Sylvester matrices below of bounded size. Then set

$$\chi_i(u, \zeta) = \mathrm{Res}_v(h, \zeta - g_i), \qquad \nu_i(u) = \text{monic lowest nonzero } \zeta\text{-coefficient of } \chi_i.$$

We call ν_i the *detector* of position i. To see what it detects, let $h_i^{(0)}$ be the product of geometric components on which g_i vanishes identically, and $h_i^{(1)}$ the product of the others. Both products are polynomials in $E[u, v]$, monic in v. Then

$$\chi_i = \zeta^{\deg_v h_i^{(0)}} \mathrm{Res}_v(h_i^{(1)}, \zeta - g_i), \qquad \nu_i \doteq \mathrm{Res}_v(h_i^{(1)}, g_i),$$

where $\doteq$ means equality up to a nonzero scalar. The last resultant is nonzero. Because $h_i^{(1)}$ is monic in v, the resultant commutes with substituting $u = \beta$, even at ramified fibers or component intersections, so a proper agreement at (β, γ) forces $\nu_i(\beta) = 0$. If g_i vanishes on the whole curve, $\nu_i = 1$.

The union filter and its finite families. For the core construction use

$$G = \mathrm{rad}\left(\prod_{i=1}^{n} \nu_i\right).$$

Every close point has a proper agreement, so its u-coordinate is a root of G. The towers $(E[u]/G)[v]/h$ have dimension $O_{B_{\mathrm{jet}}}(nk)$: each detector has degree $O_{B_{\mathrm{jet}}}(k)$ and b is bounded. Remove their factors where $s = 0$ by the split (78), invert B_0, and substitute into the candidate formulas. These finite algebras have every close point among their geometric points, and RecoverAgreement returns the candidates with at least A agreements.

Repeating the construction at $O(n)$ centers for every separant-chain stage takes polynomial time, and every close message is represented by a geometric point of some returned algebra. Algorithm 5 records this construction. Its threshold option replaces this product by the polynomial $\nu_{\geq t}$ of Lemma 9.7 for $f_i = \nu_i$ (Section 9.6.2).

9.6.2 Refinement: the threshold filter

A close point has t proper agreements, not just one, so its u-coordinate is a root of at least t different detectors ν_i. Keeping only the values of u at which at least t detectors vanish, with every point above them, divides the dimension of the union filter's algebras by t, to $O_{B_{\mathrm{jet}}}(nk/t)$. The filter counts, for each value of u, the positions whose detectors vanish there, not the agreements of a single point: two points above the same u may agree at different positions. This only adds candidates, and agreement recovery checks each one.

Counting roots of a product would count the wrong quantity, because a multiple root of one detector would count several times. For $f_1 = u^5$, $f_2 = u - 1$, $f_3 = (u-1)^2$, and $t = 2$, only $u = 1$ is a root of two different inputs; the multiplicity five at $u = 0$ comes from a single position. We therefore sort the inputs' multiplicities at each root using gcd–lcm comparators, following Batcher's sorting network [Bat68].

Lemma 9.7 (Polynomial threshold)**.** *Let F be a field, let $f_1, \dots, f_n \in F[u]$ be nonzero monic polynomials, and put $\mathcal{D} = \sum_i \deg f_i$. For every $1 \leq t \leq n$, one can construct a nonzero monic polynomial $f_{\geq t}$ in $\widetilde{O}(\mathcal{D} + n)$ operations in F such that, for every $\beta \in \overline{F}$,*

$$f_{\geq t}(\beta) = 0 \quad \Longleftrightarrow \quad \#\{i : f_i(\beta) = 0\} \geq t.$$

Moreover, $t \deg f_{\geq t} \leq \mathcal{D}$.

Proof. Pad the inputs with 1 to a power of two and use a bitonic sorting network [Bat68], replacing each comparator by $(f, g) \mapsto (\gcd(f, g), fg/\gcd(f, g))$, with monic outputs. For every irreducible ϕ this sorts the valuations by minimum and maximum. The t-th largest output $f_{\geq t}$ therefore satisfies

$$v_\phi(f_{\geq t}) > 0 \iff \#\{i : v_\phi(f_i) > 0\} \geq t, \qquad t v_\phi(f_{\geq t}) \leq \sum_i v_\phi(f_i).$$

The first relation gives the root condition over $\overline{F}$; summing the second with weights $\deg \phi$ gives $t \deg f_{\geq t} \leq \mathcal{D}$. Each comparator preserves total degree, so each layer costs $\widetilde{O}(\mathcal{D} + n)$ by fast gcds, products, and exact divisions. The $O(\log^2 n)$ layers are absorbed by the soft notation. □

Lemma 9.8 (Agreement recovery on a bounded-degree curve)**.** *Let $h(u, v) \in E[u, v]$ be reduced, monic of v-degree b, and of total degree at most B_{jet}, and let $s \in E[u, v]$ have total degree bounded in terms of B_{jet}. Suppose a degree-$< k$ message family on the open set $h = 0$, $s \neq 0$ has coefficients $N_j(u, v)/B_0(u, v)$, where B_0 is invertible on this open set, $\deg_v N_j < b$, and the numerator and denominator degrees are $O_{B_{\mathrm{jet}}}(k)$. Assume $2 \leq k \leq A \leq n$ and that the coefficient map*

$$(u, v) \longmapsto \big(N_0/B_0, \dots, N_{k-1}/B_0\big)$$

is nonconstant on every geometric irreducible component of the open curve. Put $t = A - k + 1$. Its base-field polynomial images with at least A received agreements can be recovered exactly in

$$\widetilde{O}_{B_{\mathrm{jet}}}(k^2 + nk + n^2 k/t)$$

operations over E, given the coefficient polynomials. Here E is the base field or the quadratic extension constructed in Section E.2.

Proof. Run Algorithm 5 with the threshold filter, using the detectors ν_i constructed above. As explained before Lemma 9.7, every close point lies above a root of $\nu_{\geq t}$, and

$$\deg G \leq \deg \nu_{\geq t} \leq \frac{\sum_i \deg \nu_i}{t} = O_{B_{\mathrm{jet}}}(nk/t), \qquad G = \mathrm{rad}(\nu_{\geq t}).$$

The returned towers therefore contain every close point and have total dimension $\Delta = O_{B_{\mathrm{jet}}}(nk/t)$, and Algorithm 1 returns the exact list.

Cost. Multipoint evaluation forms the g_i in $\widetilde{O}_{B_{\mathrm{jet}}}(k^2 + nk)$ operations. The resultants have bounded-size Sylvester matrices with entries of degree $O_{B_{\mathrm{jet}}}(k)$; hence all ν_i have that degree and cost $\widetilde{O}_{B_{\mathrm{jet}}}(nk)$. The threshold and radical fit this bound by Lemmas E.6 and 9.7. Computing the coefficients and recovery cost $\widetilde{O}_{B_{\mathrm{jet}}}((n+k)\Delta)$, giving the assertion. □

Proof of Theorem 9.6. Constant codes use frequency counting, and the parameter-bounded field and length exceptions use Section E.7. Otherwise run Algorithm 4 with $d = 1$. Every close message is assigned to a stage of the separant chain and is regular at one of the $Cn + (B_{\mathrm{jet}} - 1)(k-1) + 1$ centers used there (Section 9.5); centers rejected by the tests of that subsection have no regular candidate. For each remaining pair of a stage and a center, Lemma 9.5 constructs the chart in $\widetilde{O}_{B_{\mathrm{jet}}}(N_X + k^2)$ operations, where $N_X = 1 + \deg_X Q$, and its numerators have degree $O_{B_{\mathrm{jet}}}(k)$. At a stage of order one, the first two Taylor coefficients are the coordinates U, so the coefficient map is nonconstant on every component, and Lemma 9.8 recovers the close messages regular at the center in $\widetilde{O}_{B_{\mathrm{jet}}}(k^2 + nk + n^2k/t)$ operations. A stage of order zero gives a finite family of dimension at most B_{jet}, from which agreement recovery costs $\widetilde{O}_{B_{\mathrm{jet}}}(n+k)$ (Lemma 9.2). The union of the recovered lists, deduplicated, is the exact list.

Each step of the chain lowers the total jet degree, so there are at most $B_{\mathrm{jet}} + 1$ stages, each with $O_{B_{\mathrm{jet}},C}(n)$ centers. With $N_X = O_C(n)$ and $k \leq n$, the total is $\widetilde{O}_{B_{\mathrm{jet}},C}(n^2k + n^3k/t)$ operations. Lemma 9.1 constructs a fixed-support explainer in soft-linear time, and Section E.2 constructs the centers and gives the bit conversion, including inverse Frobenius over extension fields. For the prime-field assertion, failure of $p > B_{\mathrm{jet}}$ implies $n \leq q = p \leq B_{\mathrm{jet}}$, so subset interpolation has parameter-bounded cost. □

An earlier version of this paper [DKT26b] also prepares coordinates at a single center for all messages, by normalizing the curve and covering its boundary by finite families. This avoids trying $O(n)$ centers and gives $\widetilde{O}(n^2 \log q)$ at a linear gap. We omit that refinement for a simpler argument.

Corollary 9.9 (Deterministic decoding above the first-order curve)**.** *Fix $0 < \rho < a < 1$ with $a > a_1(\rho)$. Let $n \geq 2$, $1 \leq k \leq \rho n$, and $an \leq A \leq n$ be integers. Over a field $\mathbb{F}_q$ presented as in Definition E.1, with $q \geq n$ and $p > k - 1$ when $k \geq 2$, the complete degree-$< k$ agreement list at threshold A is computable deterministically in*

$$\widetilde{O}_{\rho,a}(n^3 \log q)$$

bit operations. Constant codes need no characteristic restriction.

Proof. Choose the interpolation parameters of Theorem 5.1; their degree bounds depend only on (ρ, a). When $p > B_{\mathrm{jet}}$, construct the explainer by Lemma 9.1 and apply Theorem 9.6 with $t = A - k + 1 \geq (a - \rho)n$, so that $n^2k + n^3k/t = O_{\rho,a}(n^3)$. When $p \leq B_{\mathrm{jet}}$, the hypothesis gives $k \leq B_{\mathrm{jet}}$, and Lemma E.9 applies with $K_0 = B_{\mathrm{jet}}$ and $a_0 = a$. □

Algorithm 6 IsolatedSystemCandidates: finite families at higher order

Require: Regular chart (h, s, B_0, N_j) of order $2 \le r < k$ at center x_0, its agreement polynomials g_i of (83), and the fields of Section E.6

Ensure: Univariate finite algebras with candidates whose geometric points include every close point

1: Initialize an empty collection of finite families.
2: **for** each $J \subseteq \{1, \dots, n\}$ with $|J| = r$ **do**
3: Apply Rojas's algorithm to $h = 0$, $g_i = 0$ for $i \in J$.
4: For each representation take $G_{\mathrm{iso}} = \mathrm{rad}(F_{\mathrm{iso}})$ and remove its gcds with $s(\psi_1, \dots, \psi_{r+1})$ and $B_0(\psi_1, \dots, \psi_{r+1})$.
5: If the modulus is constant, continue; otherwise form $E'[\theta]/G_{\mathrm{iso}}$.
6: Substitute the ψ_j into N_j, B_0, invert B_0, and append the algebra with its candidate $\widetilde{P}$.
7: **end for**
8: Return the finite families, without computing roots of any G_{iso}.

9.7 Higher-order recovery by isolated systems

This subsection carries out step 3 of the outline at a stage of order $r \ge 2$. The regular family at a center then has dimension r, so one agreement equation, or a threshold on one coordinate, no longer leaves finitely many points. We use r agreement equations at once: together with the equation $h = 0$ of the family, they form a square system of $r + 1$ equations in the $r + 1$ initial coordinates. The next lemma shows that for each close point, some choice of r of its agreement equations makes it a simple isolated root of this system: the linear parts of its agreement equations at the point span all linear functions on the r directions along the family. Rojas's algorithm then represents these isolated roots, and we store them as finite algebras for agreement recovery.

A solution of a polynomial system is *isolated* if, over the algebraic closure, the point by itself is an irreducible component of the solution set. For a square system, a solution is *simple* if the Jacobian matrix is nonsingular there; a simple solution is isolated and has multiplicity one. In the next lemma, the cotangent space at a point of $h = 0$ is the space of linear functions on its tangent space there.

Lemma 9.10 (Agreement cuts isolate regular candidates)**.** *At a regular order-r chart point representing a degree-$< k$ polynomial, with $r < k$, any k agreeing evaluations have differentials spanning its r-dimensional cotangent space. Thus some r agreement equations, together with $h = 0$, isolate the point simply.*

Proof. Let (u_*, v_*) be the chart point, and let ξ be a tangent vector to $h = 0$ there. Differentiating the candidate coefficients defines

$$\dot{\widetilde{P}}_\xi(X) = \sum_{j<k} d\tau_j(\xi)(X - x_0)^j.$$

Because $\tau_j = U_j$ for $j \le r$ and M is invertible, a zero tangent polynomial forces every initial-coordinate differential to be zero and hence $\xi = 0$. Thus this map from tangent directions to degree-$< k$ polynomials is injective.

At an agreeing position, differentiating $g_i = B_0(\widetilde{P}(\alpha_i) - y_i)$ gives

$$dg_i(\xi) = B_0(u_*, v_*)\dot{\widetilde{P}}_\xi(\alpha_i).$$

The term involving dB_0 is zero because the residual vanishes at (u_*, v_*), and $B_0(u_*, v_*) \neq 0$. If k agreement differentials annihilate ξ, the tangent polynomial has k distinct roots, so it is zero and $\xi = 0$. Their common annihilator is therefore zero; equivalently, they span the cotangent space.

It remains to justify that this cotangent space has dimension r. At the corresponding initial jet, $S = H_{U_r} \neq 0$, so H has a nonzero linear term. Only one reduced factor can vanish there, and it occurs to multiplicity one;

otherwise all first partials of H would vanish. The other factors are units locally, so removing them and changing coordinates leaves $dh \neq 0$. This uses the original separant; the transformed separant s need not equal h_v. Choose r independent agreement differentials on this smooth hypersurface and adjoin dh. The resulting ambient Jacobian is nonsingular, proving simple isolation. □

Finite representations of isolated roots. We use the isolated-root algorithm of Rojas [Roj99, Main Theorem 1]. For $r+1$ equations in $r+1$ variables it returns a univariate polynomial $F_{\mathrm{iso}}(\theta)$ and coordinate polynomials $\psi_1(\theta), \dots, \psi_{r+1}(\theta)$ such that

$$\{(\psi_1(\theta), \dots, \psi_{r+1}(\theta)) : F_{\mathrm{iso}}(\theta) = 0\} \supseteq \{\text{isolated affine roots of the system}\}.$$

The representation need only contain the isolated roots, since the final agreement check discards extra candidates.

The system may also have components of positive dimension, for instance a component of $h = 0$ on which the chosen agreement equations hold identically. The unperturbed u-resultant of the system then vanishes identically, as the plain resultant does at first order, so this resultant alone gives no information. Rojas's algorithm instead subtracts from each equation a new variable times the corresponding equation of a fixed system with finitely many roots, and takes the lowest nonzero coefficient in the new variable of the resulting u-resultant [Roj99, Definition 2]. This is the same kind of device as the lowest nonzero ζ-coefficient in Section 9.6.1. The algorithm needs one generic parameter, and the number of bad values is bounded in terms of the degrees, so Section E.6 makes it deterministic by trying prescribed elements of a large enough extension field; it also gives the matrix-size, trial-count, and extension-field bounds.

Set $G_{\mathrm{iso}} = \mathrm{rad}(F_{\mathrm{iso}})$ and remove its gcds with $s(\psi_1, \dots, \psi_{r+1})$ and $B_0(\psi_1, \dots, \psi_{r+1})$. A represented point need not satisfy $h = 0$, so the identity $B_0 = s^{2k}$, which holds modulo h, cannot be used to invert B_0 there. Removing both gcds makes s and B_0 units in the resulting quotient algebra, and it discards no isolated root in the regular chart, where $h = 0$ and $s \neq 0$ give $B_0 \neq 0$. Substituting the coordinate polynomials into N_j/B_0 gives a finite algebra with candidate $\widetilde{P}$. Passing to the radical of F_{iso} loses no represented point, since a polynomial and its radical have the same roots.

Theorem 9.11 (Deterministic reconstruction in arbitrary fixed order)**.** *Fix integers $d \geq 0$, $B_{\mathrm{jet}} \geq 1$, and $C \geq 1$. Let $\mathbb{F}_q$ be presented as in Definition E.1, with $q \geq n$, $1 \leq k \leq n$, and $k \leq A \leq n$. If $k \geq 2$, assume $p > \max\{k-1, B_{\mathrm{jet}}\}$. Given distinct evaluation points $\alpha_1, \dots, \alpha_n \in \mathbb{F}_q$, a received word $y \in \mathbb{F}_q^n$, and a nonzero $Q \in \mathbb{F}_q[X, Y_0, \dots, Y_d]$ with*

$$\deg_X Q \leq Cn, \qquad \deg_{Y_0,\dots,Y_d} Q \leq B_{\mathrm{jet}},$$

where the second bound is the total jet degree, suppose every degree-$< k$ polynomial with at least A agreements solves its differential equation. Represent Q sparsely, with one coefficient and exponent tuple per nonzero monomial. Then the complete list can be found deterministically in

$$n^{O(d+1)}\, \widetilde{O}_{d,B_{\mathrm{jet}},C}(\log q)$$

bit operations. No positive linear gap between A and k is needed.

Proof. The decoder is Algorithm 4. Instances with $k = 1$, $k \leq d$, or bounded length are decoded by Section E.7; otherwise every order $r \leq d$ satisfies $r < k$.

Coverage. As in the proof of Theorem 9.6, each close message is regular at one of the centers used for its separant-chain stage, and Lemma 9.5 constructs the chart at each center that passes the tests of Section 9.5. At orders zero and one, the order-zero chart and CurveCandidates with the threshold filter (Lemma 9.8)

give finite families containing the message; these lower-order stages are needed even when the explainer has higher order. At order $r \geq 2$, Lemma 9.10 makes the message's chart point a simple isolated root of the square system

$$h = 0, \qquad g_i = 0 \quad (i \in J) \tag{84}$$

for some r-subset J of positions. This system has $r + 1$ variables and degree $O_{B_{\mathrm{jet}}}(n)$, so IsolatedSystemCandidates returns a family containing the point. Algorithm 1 then returns exactly the base-field polynomials meeting the threshold, removing extra candidates and duplicates.

Cost. There are $O(n)$ centers per stage and at most n^r systems per family. Construction, substitution, and recovery cost $n^{O(r+1)}$ per system. The sparse input has at most $(\lfloor Cn \rfloor + 1)\binom{B_{\mathrm{jet}}+d+1}{d+1}$ monomials, with $O_{d,B_{\mathrm{jet}},C}(\log(n+1))$-bit exponents, so reading it fits the bound. Section E.2 constructs the fields and gives the bit conversion. □

Corollary 9.12 (Fixed-order decoding at a rate-dependent gap)**.** *Fix $0 < \rho < a < 1$ and an integer $d \geq 500$ with $\Gamma(\rho, a, d) > 1$, and let $m, B_{\mathrm{jet}}, N_{\rho,a,d}$ be the finite parameters of Corollary 8.8. Let $n \geq N_{\rho,a,d}$, $1 \leq k \leq \rho n$, and $an \leq A \leq n$ be integers. Over a field $\mathbb{F}_q$ presented as in Definition E.1, with $q \geq n$ and $p > k - 1$ when $k \geq 2$, the complete degree-$< k$ agreement list at threshold A is computable deterministically in*

$$n^{O(d+1)}\, \widetilde{O}_{\rho,a,d}(\log q)$$

bit operations.

Proof. If $k \leq \max\{d, B_{\mathrm{jet}}\}$, apply Lemma E.9 with $K_0 = \max\{d, B_{\mathrm{jet}}\}$ and $a_0 = a$. Otherwise $p > k - 1 \geq \max\{d, B_{\mathrm{jet}}\}$, so the characteristic hypothesis of Corollary 8.8 holds. Its certified support has bounded jet degree and X-degree $O_{\rho,a,d}(n)$: construct the explainer by Lemma 9.1, then apply Theorem 9.11, whose hypotheses $p > \max\{k - 1, B_{\mathrm{jet}}\}$ and $k > d$ hold. If ρ, a have no finite input representation, fix rational bounds with the same positive margin in advance. □

9.8 Uniform capacity decoding

We now decode at every gap δ from capacity, uniformly over the rate, by combining the uniform explainers with the reconstruction routines above. The explainers come from Theorem 8.1 for gaps below $6/25$ and from the fixed first-order support of Theorem 5.7 for gaps at least $6/25$, so the constants depend on the gap but not on the rate. We treat the larger gaps first, with the first-order decoder.

Corollary 9.13 (Uniform first-order decoding at gap $6/25$)**.** *Over a field $\mathbb{F}_q$ presented as in Definition E.1, with $q \geq n$ and $p > k - 1$ when $k \geq 2$, the exact agreement list is computable deterministically in $\widetilde{O}(n^3 \log q)$ bit operations whenever $1 \leq k \leq n$ and $k + 6n/25 \leq A \leq n$, with an absolute hidden constant.*

Proof. The support $(m, B_\partial, B_{\mathrm{jet}}) = (12, 4, 22)$ in Theorem 5.7 gives $\deg_X Q < 12A \leq 12n$. For $p > 22$, apply Theorem 9.6 with $C = 12$ and $t \geq 6n/25$, so that $n^2 k + n^3 k/t = O(n^3)$. If $p \leq 22$, then $k \leq 22$, and Lemma E.9 applies with $K_0 = 22$ and $a_0 = 6/25$. Bounded lengths and constant codes use Section E.7. □

Theorem 9.14 (All-rate Reed–Solomon list decoding)**.** *For every fixed gap $\delta \in (0, 1)$, there are an integer N_δ and a deterministic decoder with the following guarantee. Let $n \geq N_\delta$, let $1 \leq k \leq n$, and let $\mathbb{F}_q$ be presented as in Definition E.1, with $q \geq n$. For $k \geq 2$, require $p > k - 1$. Constant codes require no characteristic assumption. Given any n distinct evaluation points $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n) \in \mathbb{F}_q^n$, any received word $y \in \mathbb{F}_q^n$, and any integer agreement threshold*

$$k + \delta n \leq A \leq n,$$

the decoder, given this presentation and $(n, k, A, \boldsymbol{\alpha}, y)$, returns exactly

$$\mathcal{L}(y, A) = \left\{P \in \mathbb{F}_q[X] : \deg P < k,\ \#\{i : P(\alpha_i) = y_i\} \geq A\right\}. \tag{85}$$

The output list size is $n^{O_\delta(1)}$, independently of q, and the running time is $n^{O(\exp(1.5/\delta))}\widetilde{O}_\delta(\log q)$ bit operations. The constants and exponents, as well as N_δ, depend only on δ, uniformly over the rate k/n and the evaluation set.

Proof. Fix δ. Enlarge N_δ to dominate the finite threshold in Theorem 8.1 and every parameter-dependent field, grid, and center condition used by the applicable reconstruction routine. If $k = 1$, frequency counting returns exactly the constants occurring at least A times, in every characteristic. We may therefore assume $2 \leq k < A \leq n$.

Suppose first that $0 < \delta < 6/25$. If $p \leq B_{\mathrm{jet}}(\delta)$, then $p > k - 1$ gives $k \leq B_{\mathrm{jet}}(\delta)$, and Lemma E.9 with $K_0 = B_{\mathrm{jet}}(\delta)$ and $a_0 = \delta$ recovers the exact list in $\widetilde{O}_\delta(n \log q)$ bit operations. For $p > B_{\mathrm{jet}}(\delta)$, the finite choices in Theorem 8.1 give an order-$d(\delta)$ explainer of jet degree at most $B_{\mathrm{jet}}(\delta)$ and X-degree $O_\delta(n)$, which Lemma 9.1 constructs in $\widetilde{O}_\delta(n \log q)$ bit operations. If $k \leq d(\delta)$, apply Lemma E.9 with $K_0 = d(\delta)$ and $a_0 = \delta$; otherwise Theorem 9.11 returns the complete list. Both branches fit the bound

$$n^{O(d(\delta)+1)}\widetilde{O}_\delta(\log q) = n^{O(\exp(1.5/\delta))}\widetilde{O}_\delta(\log q).$$

The list-size bound is the one proved in Theorem 8.1 and is independent of q.

For $6/25 \leq \delta < 1$, the requested threshold satisfies $A - k \geq \delta n \geq 6n/25$, so Corollary 9.13 applies at A itself and returns the complete list in $\widetilde{O}(n^3 \log q)$ bit operations, and Theorem 5.7(i) bounds the list by $307n$ under the same characteristic assumption. □

Taking $A = k + \lceil \delta n \rceil$ and writing $\rho = k/n$, the theorem corrects every pattern of at most

$$n - k - \lceil \delta n \rceil = \lfloor (1 - \rho - \delta) n \rfloor$$

errors whenever this radius is nonnegative. This is the usual capacity radius with gap δ. The threshold A is an integer input to the decoder; the fixed real parameter δ determines the derivative order, the interpolation parameters, and the minimum length N_δ. If $A > n$, the output is empty.

Both the list size and the decoding time are polynomial in the binary input length $n \log q$ for every fixed gap, without an upper bound on the field size. Below N_δ, the subset interpolation of Section E.7 still decodes exactly in polynomial time; the explicit list-size estimates keep their length hypotheses.

10 Limits and open questions

Our bounds hold at a fixed gap, under a characteristic hypothesis, and with polynomial dependence on the length; each qualification leaves an open question.

Fixing the gap before the length grows. Our capacity bounds fix the gap δ before letting the length grow, and the order of these limits matters. Crites and Stewart disprove the original capacity conjectures and propose entropy-corrected versions [CS25]. At error radius τ, the finite-alphabet condition is $H_q(\tau) \leq 1 - \rho$, where H_q is the q-ary entropy function. Since $H_q(\tau) \leq \tau + \log_q 2$, agreement $\rho + \delta$ meets this condition once $\log_q 2 < \delta$, which holds for every fixed δ and large n, since we assume $q \geq n$. This does not give a bound uniform in a shrinking gap: prime-field constructions with gap $\Theta_\rho(1/\log n)$ have exceptional counts $n^{c-o(1)}$ for arbitrarily large fixed c [KKH26, Kam26], and our exponent $d(\delta)$ and minimum length $N(\delta)$ in Theorem 1.2 also depend on the gap.

Characteristic. Our first-order and higher-order bounds assume a large characteristic, not merely a large field, and known constructions show why. Extending a small prime field enlarges the denominator q of an exceptional probability without changing the characteristic, and below Johnson the characteristic matters: characteristic-two and base-code constructions obstruct MCA [BCH+26, DG25, FS25], and Ben-Sasson, Kopparty, and Radhakrishnan use subspace polynomials over extension fields of fixed small characteristic to build received words with superpolynomially many Reed–Solomon codewords beyond Johnson [BKR10].

Over binary fields, Dao, Kominers, Thaler, and Zheng show that beyond-Johnson proximity gaps fail [DKTZ26]. At rate $1/4$, on additive domains filling a constant fraction of a binary field and over a suitable extension challenge field, combinations of two words approach the Johnson threshold for superpolynomially many challenges. This rules out any $\mathrm{poly}(n)/q$ exceptional probability below Johnson; in the same setting, one word has superpolynomially many codewords at agreement approaching Johnson. At rate $1/16$, every additive domain in a proper subfield of the challenge field has quadratically many exceptional challenges at the Johnson threshold itself.

Which characteristic is necessary for our bounds remains open, as does an evaluation of our characteristic-free Johnson bound for smaller proofs over binary fields.

Dependence on the length. At a fixed positive error radius, every linear code has MCA error at least $\min\{\lfloor \tau n \rfloor / q, 1\}$ in the range of [ABF26, Lemma 4.16], so our Johnson theorem has the optimal linear dependence on n. It remains open whether its dependence on the Johnson gap can be improved, and whether the quadratic dependence on n in our first-order MCA bound can be reduced.

Decoding time. Because Theorem 9.11 solves a square system for each r-subset of positions, the exponent of n in Theorem 9.14 grows with the derivative order $d(\delta)$; we do not know whether it must.

A fixed-length benchmark. Arnon, Boneh, and Fenzi's Proximity Prize [ABF26], hosted at `https://better.codes`, tests one instance: eight Reed–Solomon words of length 2^{18} and dimension 2^{17} over $\mathbb{F}_{p^6}$, on a multiplicative subgroup of KoalaBear, $p = 2130706433$ [Pro26c]. It scores b bits for a radius τ satisfying $(1-\tau)^{128} \le 2^{-b}$ and combined MCA and list error at most 2^{-128} [Pro26d].

Using only the finite bounds proved in this paper, a bounded parameter search finds a certificate at $A = 183210$ agreements, a score of 66.15 bits with $(\mathsf{E}_{\mathrm{exc}} + \Lambda)/q < 2^{-128}$, where $\mathsf{E}_{\mathrm{exc}}$ bounds the exceptional count and Λ the list size at interleaving width sixteen; our Johnson certificate attains 63.99 bits at $A = 185372$. The prize submission promoted on September 15, 2026 gives 68.11 bits at $A = 181284$ [Bit26d]. Its agreement fraction, about 0.691544, lies between our certificate's and $a_1(1/2) \approx 0.689898$.

Later submissions add arguments that our single-equation analysis does not use. They impose stronger interpolation constraints, combine several equations with their common factors removed, use the whole interpolation kernel, require Taylor coefficients to vanish, and count jointly over components [erd26a, Bit26b, rec26, Bit26c, jie26b, jac26a, sau26, jsi26, net26, i3426, jac26d, jac26b, jac26c, Bit26d], and an upper-track construction exploits the multiplicative structure of the domain [erd26b]. Better counting alone improves finite parameters but leaves unchanged the coefficient-to-constraint inequality that defines a_1, and whether these ingredients lower the limiting threshold at rates $1/2$ and below remains open. The artifact records our certificate and a concordance of submissions [DKT26a].

Acknowledgements

We thank Kai Zhe Zheng for close discussion and feedback, including a suggestion to name the central interpolation polynomial in BCPZZ the hidden-derivative explainer. Kominers thanks MIT Sloan's Technological Innovation, Entrepreneurship, and Strategic Management (TIES) Group for its hospitality during this work.

Disclosure

The second and third authors are Research Partners at a16z crypto, which is an investor in various blockchain-based platforms, as well as in the crypto ecosystem more broadly. General a16z disclosures are available at https://www.a16z.com/disclosures/.

A Rational-image and Taylor tools

This appendix proves Kopparty's unique lift of a solution from its first Taylor coefficients (Section A.1) and gives new bounds on the degrees of the resulting rational map and of its image, the family counted in Sections 4 and 8 (Section A.2). The decoders in Section 9 also use the lift, and Section 5.2 refines the bounds at order one.

A.1 Weighted degree and unique lifting

We describe a solution of a differential equation by its initial Taylor coefficients at one point and reconstruct the remaining coefficients, following Kopparty [Kop15, Theorem 4.3, Theorem 4.4, and Corollary 4.5]. We use the Hasse-derivative convention of Definition 3.4, and bound the degree of a specialization using the following weights.

Definition A.1 (Weighted degree)**.** For nonnegative integer weights $w_0, \dots, w_s$, the weighted degree of a monomial $Z_0^{e_0} \cdots Z_s^{e_s}$ is $\sum_{i=0}^{s} w_i e_i$. The weighted degree $\operatorname{wdeg}_{w_0,\dots,w_s} Q$ of a nonzero polynomial Q is the maximum weighted degree of one of its monomials with nonzero coefficient. We set $\operatorname{wdeg}_{w_0,\dots,w_s}(0) = -\infty$.

These weights record the degree lost under Hasse differentiation. If $d \le D$, $Q \in F[X, Y_0, \dots, Y_d]$, and $P \in F[X]_{\le D}$, then

$$\deg_X Q\big(X, P(X), P^{[1]}(X), \dots, P^{[d]}(X)\big) \le \operatorname{wdeg}_{1,D,D-1,\dots,D-d} Q, \tag{86}$$

because X contributes degree one and the substitution for Y_j contributes degree at most $D - j$.

In characteristic zero or $p > D$, each new Taylor coefficient enters the recurrence with a nonzero multiplier, so regular initial data determine at most one solution.

Lemma A.2 (Unique lift from a regular jet)**.** *Let F be a field of characteristic zero or $p > D$, let $0 \le r \le D$, and let*

$$T \in F[X, Y_0, \dots, Y_r].$$

Fix a center $x_0 \in F$ and an initial jet $U_0, \dots, U_r \in F$ such that

$$T(x_0, U_0, \dots, U_r) = 0 \quad \text{and} \quad \frac{\partial T}{\partial Y_r}(x_0, U_0, \dots, U_r) \neq 0. \tag{87}$$

There is at most one $P \in F[X]_{\le D}$ satisfying

$$T(X, P, P^{[1]}, \dots, P^{[r]}) \equiv 0 \tag{88}$$

and $P^{[j]}(x_0) = U_j$ for every $0 \le j \le r$. Moreover, this polynomial can be reconstructed, or its nonexistence detected, using a number of field operations polynomial in D, $\operatorname{wdeg}_{1,D,D-1,\dots,D-r} T$, and the sparse support size and binary exponent lengths of T. Here coefficients in the arbitrary field F are accessed through field operations and equality tests; no bit encoding of F is assumed.

Proof. Write $P(x_0 + Z) = \sum_{i=0}^{D} \tau_i Z^i$, with $\tau_i = U_i$ for $i \le r$. For $r < i \le D$, changing τ_i by γ changes the j-th Hasse derivative by

$$\gamma \binom{i}{j} Z^{i-j}. \tag{89}$$

The first affected coefficient in the differential equation is at Z^{i-r}, and its change is exactly

$$\gamma \binom{i}{r} \frac{\partial T}{\partial Y_r}(x_0, U_0, \dots, U_r). \tag{90}$$

Terms with two changed factors have order at least $2(i-r) > i-r$; a lower derivative or a coefficient τ_j with $j > i$ affects only higher orders. The displayed multiplier is nonzero by regularity and the characteristic hypothesis. Solving this linear equation successively therefore fixes every remaining coefficient. Truncated polynomial arithmetic and a final check of (88) reconstruct P or detect that it does not exist, within the asserted polynomial cost. □

The separant chain. The separant $\partial T/\partial Y_r$ can vanish identically along a solution, while the separant of a derivative of T does not. One deterministic chain assigns every solution to a stage where the lemma applies. Fix a nonzero $Q \in F[X, Y_0, \dots, Y_d]$ whose individual jet degrees are smaller than the positive characteristic, if any. Starting from $Q_0 = Q$, as long as Q_i depends on a jet variable, let

$$r_i := \max\{j : Q_i \text{ depends on } Y_j\}, \qquad Q_{i+1} := \frac{\partial Q_i}{\partial Y_{r_i}}. \tag{91}$$

Every member is nonzero by the degree hypothesis, and each derivative lowers total jet degree, so the chain ends at a nonzero $Q_t \in F[X]$. Every solution P therefore has a first stage i at which $Q_i(X, P, \dots, P^{[r_i]})$ vanishes identically in X while the specialized separant $Q_{i+1}(X, P, \dots, P^{[r_i]})$ does not. The degree hypothesis on Q only makes the chain terminate; it does not give unique lifting. Under the hypotheses of Lemma A.2, namely characteristic zero or $p > D$, $r_i \le D$, and $P \in F[X]_{\le D}$, choose a center x_0 at which this specialized separant is nonzero; such a center exists in F or in a finite extension E of F. The lemma, applied over E to $T = Q_i$ with $r = r_i$, then shows that the initial jet $(P^{[j]}(x_0))_{j \le r_i}$ determines P.

A.2 The regular Taylor chart

Lemma A.2 reconstructs one solution from given regular initial coefficients, or detects that none exists. Following Zheng and Jeronimo [Zhe26, Jer26], we instead keep the initial jet $U = (U_0, \dots, U_r)$ as variables. This gives one rational map for all regular solutions, the *Taylor chart*, which sends U to all reconstructed coefficients $\tau_j(U)$ and is defined where the separant at the center is nonzero.

Lemma A.3 (Regular Taylor chart)**.** *Let $0 \le r < K$ and let $T \in F[X, Y_0, \dots, Y_r]$ have total jet degree at most B_{jet}, where $B_{\mathrm{jet}} \ge 1$, and separant $T_{Y_r} = \partial T/\partial Y_r$. Assume the characteristic is zero or exceeds $K-1$. Choose a center $x_0 \in F$ such that $S(U) := T_{Y_r}(x_0, U)$ is nonzero, and put*

$$H(U) = T(x_0, U), \qquad \mathcal{U} = \{U \in F^{r+1} : H(U) = 0,\ S(U) \ne 0\}.$$

There are polynomials $B_0, N_0, \dots, N_{K-1}$ of degree at most $1 + 2K(B_{\mathrm{jet}} - 1)$, with B_0 nonzero wherever S is nonzero, such that the rational map

$$U \longmapsto (\tau_0(U), \dots, \tau_{K-1}(U)), \qquad \tau_j = N_j/B_0,$$

is injective on $\mathcal{U}$, has $\tau_j = U_j$ for $j \le r$, and recovers the Taylor coefficients at x_0 of every degree-$< K$ solution of $T = 0$ whose separant specialization is nonzero at x_0. The closure of $\mathcal{U}$ has dimension at most r and sum of component degrees at most B_{jet}.

Proof. Apply the triangular recurrence of Lemma A.2 with degree bound $K-1$, keeping the initial jet U as variables. For $1 \le m < K-r$, it gives

$$\binom{r+m}{r} S(U)\tau_{r+m} + R_m(U, \tau_{r+1}, \dots, \tau_{r+m-1}) = 0. \tag{92}$$

The leading multiplier is nonzero on $\mathcal{U}$, since $S \neq 0$ there and $\binom{r+m}{r} \neq 0$ by the characteristic assumption.

Induction gives a common form for each new coefficient:

$$\tau_{r+m} = \frac{\widehat{N}_m(U)}{S(U)^{2m-1}}, \qquad \deg \widehat{N}_m \le (2m-1)(B_{\text{jet}}-1)+1. \tag{93}$$

Indeed, a monomial $U^\beta \prod_{i<m} \tau_{r+i}^{e_i}$ in R_m satisfies $|\beta| + \sum_i e_i \le B_{\text{jet}}$ and $\sum_i i e_i \le m$. After substituting the earlier rational coefficients, its denominator exponent is $\sum_i (2i-1)e_i = 2\sum_i i e_i - \sum_i e_i \le 2m-2$. For the last inequality: if the monomial contains at least two factors τ_{r+i}, counted with multiplicity, subtracting their number gives the bound; if it contains exactly one, its index $i \le m-1$ gives $2i-1 \le 2m-3$; if it contains none, the exponent is zero. Clearing to S^{2m-2} bounds the numerator degree by

$$(2m-2)(B_{\text{jet}}-1) + |\beta| + \sum_i e_i \le (2m-1)(B_{\text{jet}}-1)+1.$$

Division by the leading binomial coefficient times S proves the induction step.

Use the same denominator $B_0 = S^{2K}$ for every coordinate. For $j \le r$, set $N_j = U_j S^{2K}$; for $j = r+m > r$, set $N_j = \widehat{N}_m S^{2K-(2m-1)}$, so that $N_j = \tau_j B_0$ for every $j < K$. Since $\deg S \le B_{\text{jet}} - 1$, all these numerators and the denominator have degree at most $1 + 2K(B_{\text{jet}}-1)$. This convention also covers $r = K-1$, where the recurrence has no step, and makes $B_0 \neq 0$ equivalent to $S \neq 0$. Since $\tau_j = U_j$ for $j \le r$, the map is injective on the whole open set $S \neq 0$. The Taylor coefficients of a regular solution satisfy the same recurrence, so induction proves recovery. Finally, H is nonzero because its U_r-derivative is S. Its hypersurface has degree at most B_{jet} and dimension at most r; passing to the closure of its principal open subset only discards components contained in $V(S)$. □

A.3 Degree of a rational image

Lemma A.4 (Degree of a rational image)**.** *Let E be algebraically closed, as in Section 4.1, and let $V \subseteq E^M$ have dimension at most s and total degree at most Δ. Let $\phi = (N_1/B_0, \dots, N_K/B_0)$, where all numerators and the common denominator have degree at most $b \ge 1$, and B_0 vanishes identically on no component of V. The closure of $\phi(V \setminus V(B_0))$ has dimension at most s and total degree at most Δb^s.*

Proof. For an irreducible component C of V, its image closure W is irreducible of dimension $s' \le \dim C$. The image is constructible and contains a dense open subset of W by Chevalley's theorem [The26]. Intersect W with s' generic affine hyperplanes; their intersection consists of $\deg W$ distinct points in that open subset. Pulling back each hyperplane gives a polynomial of degree at most b. On $C \setminus V(B_0)$, the resulting intersection is the disjoint union of the nonempty fibers above those points. Each fiber contains the nonempty open part of at least one irreducible component of the closed intersection, and one such component cannot meet two fibers on $B_0 \neq 0$, since its image in this finite set is irreducible. Refined Bézout (Section 4.1) therefore bounds the number of fibers by $\deg(C) b^{s'} \le \deg(C) b^s$. Summing over C proves the degree bound, including when the fibers are positive-dimensional. □

A.4 Rational images with separate degree bounds

For a rational map on a hypersurface $V(G)$, the next lemma keeps the challenge degree in Z separate from the jet degree in u, omitting the term $D_G^Z B_\phi^Z$ of Lemma A.4. In the Johnson theorem that term is of order DB_Z^2 already in the separable case, against $DB_Z B_{\mathrm{jet}}$, so omitting it saves one factor of the gap.

Lemma A.5 (Rational images with two degree bounds)**.** *Let E be algebraically closed and let $G \in E[Z, u_0, \dots, u_d]$ be nonzero with degree at most D_G^Z in Z and at most D_G^u in u. Let V_0 be a union of components of $V(G)$, and let a rational map ϕ have a common denominator that vanishes identically on no component of V_0. If its numerators and denominator have integer degree bounds $B_\phi^Z \geq 1$ in Z and $B_\phi^u \geq 1$ in u, its image closure has dimension at most $d+1$ and total degree at most*

$$D_G^Z (B_\phi^u)^{d+1} + (d+1) D_G^u B_\phi^Z (B_\phi^u)^d.$$

Proof. Take one irreducible factor G_0 of G, with degrees (D_0^Z, D_0^u), and put $s = d+1$. Let R_N consist of polynomials of degree at most $B_\phi^Z N$ in Z and total degree at most $B_\phi^u N$ in the s jet variables. Multiplication by G_0 injects the rectangle with degree bounds $(B_\phi^Z N - D_0^Z, B_\phi^u N - D_0^u)$ into $(G_0) \cap R_N$. For sufficiently large N, the image of R_N in $E[Z,u]/(G_0)$ therefore has dimension at most

$$(B_\phi^Z N + 1)\binom{B_\phi^u N + s}{s} - (B_\phi^Z N - D_0^Z + 1)\binom{B_\phi^u N - D_0^u + s}{s}.$$

The terms of degree $s+1$ cancel. The remaining leading term is

$$\frac{D_0^Z (B_\phi^u)^s + s D_0^u B_\phi^Z (B_\phi^u)^{s-1}}{s!} N^s.$$

Embed $V(G_0)$ using all monomials of bidegree at most (B_ϕ^Z, B_ϕ^u). Since $B_\phi^Z, B_\phi^u \geq 1$, these coordinates include Z and every u_i, so this is a closed embedding. Every polynomial of ordinary degree at most N in the embedding coordinates pulls back into R_N. Comparing the resulting Hilbert-function bound with the cumulative affine Hilbert polynomial bounds the degree of the embedded hypersurface by $D_0^Z (B_\phi^u)^s + s D_0^u B_\phi^Z (B_\phi^u)^{s-1}$.

The numerators and common denominator of ϕ are linear forms in these embedding coordinates. Applying Lemma A.4 with degree one bounds the image closure by the same quantity. Sum over the irreducible factors of G that define components of V_0, using $\sum D_0^Z \leq D_G^Z$ and $\sum D_0^u \leq D_G^u$. □

B The ordinary transfer and the Johnson constants

This appendix proves the parts of Theorem 4.4 deferred from Section 4: the ordinary transfer, Lemma 4.3, the order-zero symbolic interpolant of BCHKS, and the closed-form bound of (39) for the rounded integer parameters.

B.1 Proof of the ordinary transfer

Proof of Lemma 4.3. We construct the family of candidates over an algebraic closure $\overline{F}$; the final candidates and received words remain over F. Separate the content of Q in Y and write $Q = C(X,Z) \prod_i Q_i(X,Z,Y)^{e_i}$, where $e_i \geq 1$ and the Q_i are the distinct irreducible factors of positive Y-degree. Put $j_i = \deg_Y Q_i$, $h_i = \deg_Z Q_i$, and $h_C = \deg_Z C$. The squarefree product divides Q, so

$$\sum_i j_i \leq B_{\mathrm{jet}}, \qquad h_C + \sum_i h_i \leq B_Z, \qquad \sum_i h_i j_i \leq B_Z B_{\mathrm{jet}}. \tag{94}$$

Choose one nonzero coefficient of C as a polynomial in X. There are at most h_C challenges for which that coefficient, and hence possibly $C(X, z)$, vanishes identically. We include them in the preliminary set $\mathcal{E}_{\mathrm{pre}}$. Away from them, every solution of $Q(X, z, P) = 0$ solves $Q_i(X, z, P) = 0$ for at least one i.

For each factor Q_i, construct its separable Frobenius pullback as follows. In characteristic zero put $s_i = 1$ and $G_i = Q_i$. In characteristic $p > 0$, take the largest $s_i = p^{v_i}$ such that $Q_i(X, Z, Y) = G_i(X, Z, Y^{s_i})$, and put $b_i = \deg_U G_i = j_i/s_i$. Then $\partial_U G_i \neq 0$ by maximality of s_i. Applying inverse Frobenius (bijective on $\overline{F}$ and on every finite base field) to the coefficients of G_i gives an irreducible polynomial $\widetilde{G}_i(T, W, U)$, separable in U, of bidegree at most (h_i, b_i) in (W, U), and

$$\widetilde{G}_i(T, W, U)^{s_i} = Q_i(T^{s_i}, W^{s_i}, U). \tag{95}$$

Hence $Q_i(X, z, P(X)) = 0$ and $w^{s_i} = z$ imply

$$\widetilde{G}_i(T, w, P(T^{s_i})) = 0. \tag{96}$$

To parameterize the solutions, choose a Taylor center t_0 transcendental over $\overline{F}$ and work over an algebraic closure of $\overline{F}(t_0)$. By Gauss's lemma, $\widetilde{G}_i(t_0, W, U)$ is coprime to its U-derivative, so

$$R_i(W) = \mathrm{Res}_U(\widetilde{G}_i(t_0, W, U), \partial_U \widetilde{G}_i(t_0, W, U)) \neq 0, \tag{97}$$

$$\deg_W R_i \leq (2b_i - 1)h_i. \tag{98}$$

The roots yield at most this many original challenges $z = w^{s_i}$; add their base-field images to $\mathcal{E}_{\mathrm{pre}}$.

Away from these roots, the plane curve

$$C_i : \quad \widetilde{G}_i(t_0, W, u) = 0 \tag{99}$$

has nonzero separant $S = \partial_U \widetilde{G}_i(t_0, W, u)$. Put $K_i' = s_i D + 1$. The order-zero Taylor recurrence reconstructs the pulled-back root $P((t_0 + \xi)^{s_i}) = \sum_{r<K_i'} c_r \xi^r$ from $c_0 = u$ and gives, for $1 \leq r < K_i'$,

$$c_r = \frac{A_r(W, u)}{S(W, u)^{2r-1}}, \qquad \deg_u A_r \leq (2r - 1)(b_i - 1) + 1, \qquad \deg_W A_r \leq (2r - 1)h_i.$$

This is the denominator induction of Lemma A.3, with W as an extra variable. Its leading binomial coefficient is always $\binom{r}{0} = 1$, so no characteristic restriction occurs. With denominator $S^{2K_i'-3}$, the map to $(z, z^2, \ldots, z^\ell, P)$ has degree bounds

$$\beta_i = 1 + (2K_i' - 3)(b_i - 1), \qquad \alpha_i = s_i\ell + (2K_i' - 3)h_i \tag{100}$$

in u and W, respectively. The term $s_i\ell$ comes from the output coordinates $z, z^2, \ldots, z^\ell$, since $z = W^{s_i}$. The map outputs the original P: for $x_0 = t_0^{s_i}$,

$$P((t_0 + \xi)^{s_i}) = \sum_{r=0}^{D} P^{[r]}(x_0)\xi^{s_i r}, \tag{101}$$

so the coefficients with indices $s_i r$ determine P by an invertible linear change of coordinates. The reconstruction runs to degree $s_i D$, but the output P still has degree at most D. Hence every solution pair (z, P) of Q_i with z outside the resultant set lies in this rational image.

Applying Lemma A.5 with $d = 0$ gives an image of dimension at most one and degree

$$\begin{aligned} J_i &\leq h_i\beta_i + b_i\alpha_i \\ &= \ell j_i + h_i\big[1 + (2Ds_i - 1)(2b_i - 1)\big] \\ &\leq \ell j_i + h_i\Psi_D(j_i). \end{aligned} \tag{102}$$

The identity $s_i b_i = j_i$ turns the term $s_i \ell b_i$ of $b_i \alpha_i$ into ℓj_i.

To prove the final inequality, rearrange that bracket as

$$1 + (2Ds_i - 1)(2b_i - 1) = 1 + (2D - 1)(2j_i - 1) + 2(s_i - 1)(b_i - D). \tag{103}$$

If $s_i = 1$ or $b_i \le D$, its last term is nonpositive. Otherwise $s_i \ge 2$ and $b_i \ge D + 1$, and

$$(s_i - 1)(b_i - D) = j_i - \big(b_i + (s_i - 1)D\big) \le j_i - 2D - 1.$$

This proves the bound by $\Psi_D(j_i)$ in both cases. The function Ψ_D is increasing on the positive integers, so $\sum_i h_i \Psi_D(j_i) \le B_Z \Psi_D(B_{\text{jet}})$.

For a transcendental output challenge z, there is set-theoretically one w with $w^{s_i} = z$, and at most b_i possible values of u. The rational image is constructible and dense in each curve component of its closure. Its finite omitted subset and any vertical components miss this fiber, whose reduced degree is therefore at most b_i.

Let V be the union of the image closures, intersected with the degree-$\le D$ coefficient space. By (94),

$$\deg V \le D_{\text{joint}} := \sum_i J_i \le \ell B_{\text{jet}} + B_Z \Psi_D(B_{\text{jet}}), \tag{104}$$

and the reduced transcendental fiber has degree at most $\sum_i b_i \le B_{\text{jet}}$. Apply Lemma 4.2 with $(k, r_1, J_1, F_1) = (D+1, 0, D_{\text{joint}}, B_{\text{jet}})$. Since both incidence factors are empty products, it contributes $\frac{n-L+1}{A-L+1} D_{\text{joint}} + \ell(n-L) B_{\text{jet}}$ and gives witnesses over F with equality of agreement sets. Finally,

$$|\mathcal{E}_{\text{pre}}| \le h_C + \sum_i (2b_i - 1) h_i \le (2B_{\text{jet}} - 1) B_Z. \tag{105}$$

These bounds prove (35). □

B.2 The order-zero interpolant of BCHKS

The transfer needs the interpolant's degrees in Y and Z, and its soundness needs the degree of its specialization. Fix the notation of Theorem 4.4, and set real thresholds for the specialization degree, the Y-degree, and the joint degree in (Y, Z):

$$C_{\text{spec}} = t\sqrt{Dn} = tn\sqrt{\rho_-}, \qquad C_{\text{jet}} = t\sqrt{n/D} = \frac{t}{\sqrt{\rho_-}}, \qquad C_Z = \frac{t^2 n}{3D} = \frac{t^2}{3\rho_-}.$$

Thus $B_{\text{jet}} = \lceil C_{\text{jet}} \rceil - 1$ and $B_Z = \lceil C_Z \rceil - 1$.

Proposition B.1 (BCHKS order-zero symbolic interpolant)**.** *For every field F, every distinct $\alpha_1, \dots, \alpha_n \in F$, and every $f, g \in F^n$, there is a polynomial $Q \in F[X, Z, Y]$ such that*

$$Q(X, z, Y) \ne 0 \qquad \text{for every } z \text{ in every extension of } F, \tag{106}$$

$$\begin{aligned} \operatorname{wdeg}_{1,0,D} Q &< C_{\text{spec}}, & \deg_Y Q &\le B_{\text{jet}}, \\ \operatorname{wdeg}_{0,1,1} Q &< C_Z, & \deg_Z Q &\le B_Z. \end{aligned} \tag{107}$$

For every field extension E/F, every $z \in E$, and every $P \in E[X]_{\le D}$ with $|S_z(P)| \ge A$, one has

$$Q(X, z, P(X)) = 0. \tag{108}$$

Proof. Use the BCHKS monomial space

$$Q = \mathrm{span}_F\{X^i Y^j Z^e : i + Dj < C_{\mathrm{spec}},\ j + e < C_Z\}.$$

At each position r, require every coefficient of $U^a V^b$, for $a + b < m$, in $Q(\alpha_r + U, Z, f_r + Zg_r + V)$ to vanish as a polynomial in Z. Translation in Y preserves total Y, Z degree, so the scalar counts are

$$N_{\mathrm{var}} = \sum_{j=0}^{B_{\mathrm{jet}}} (\lceil C_{\mathrm{spec}} \rceil - Dj)(\lceil C_Z \rceil - j),$$
$$N_{\mathrm{eq}} = n \sum_{b=0}^{m-1} (m - b)(\lceil C_Z \rceil - b). \tag{109}$$

For $m \geq 3$, BCHKS's dimension estimate gives $N_{\mathrm{var}} > N_{\mathrm{eq}}$ for these cutoffs [BCH+26, full version, Lemma 3.1]. Primitive normalization as in Lemma 3.11 preserves the constraints and degree bounds and gives (106).

For soundness, the definition of m gives $m\eta_0 \geq \sqrt{\rho_-}/2$, and therefore

$$C_{\mathrm{spec}} = (m + \tfrac{1}{2})n\sqrt{\rho_-} \leq mn(\sqrt{\rho_-} + \eta_0) = man \leq mA.$$

If $P(\alpha_r) = f_r + zg_r$, the imposed bivariate multiplicity gives a root of multiplicity at least m in $Q(X, z, P(X))$ at α_r. Its degree is less than $C_{\mathrm{spec}} \leq mA$, so A agreements force it to vanish. This polynomial argument remains valid over every scalar extension, proving (108). □

For a fixed z, every member of $\mathcal{L}_D(f + zg, A)$ gives a distinct linear factor $Y - P(X)$ of $Q(X, z, Y)$ over $F(X)$. Consequently $|\mathcal{L}_D(f + zg, A)| \leq \deg_Y Q \leq B_{\mathrm{jet}}$. Applying the same construction with $g = 0$ proves the first term in (38) for every received word, by the Guruswami–Sudan factor argument [GS99]. For the classical pairwise bound, choose A agreements for each of L candidates. Their pairwise intersections have size at most D; Cauchy–Schwarz gives $L^2A^2/n \leq LA + L(L-1)D$. As $A^2 > nD$, this yields the other term $L \leq n(A - D)/(A^2 - nD)$ in (38).

B.3 Closed constants

To prove the closed form in (39), write $x = \sqrt{\rho_-}$ and $y = 1/n$. Since $1 \leq D \leq n - 2$,

$$y \leq \min\{x^2, (1 - x^2)/2\}. \tag{110}$$

The rounded parameters satisfy

$$B_{\mathrm{jet}} < \frac{t}{x}, \qquad B_Z < \frac{t^2}{3x^2}, \qquad \lambda_{\mathrm{agr}} \leq \frac{1 - x^2}{x - x^2} = \frac{1 + x}{x}.$$

Use $\Psi_D(B_{\mathrm{jet}}) \leq 1 + 4DB_{\mathrm{jet}}$ in (34). Keeping the -1 in $n - D - 1$ and combining its term with $\lambda_{\mathrm{agr}} B_{\mathrm{jet}}$ gives

$$\frac{B_{\mathrm{cf}}}{nt^3/x^2} < \frac{4(1 + x)}{3} + \left(\frac{2}{3x} + \frac{1 + x}{3tx} + \frac{1}{t^2}\right) y + \frac{x(1 - x^2)}{t^2}.$$

The right side decreases with $t \geq 7/2$. Substituting $t = 7/2$ and (110) reduces the claim to $U(x) < 8/3$ for

$$U(x) = \frac{4(1 + x)}{3} + \left(\frac{16}{21x} + \frac{26}{147}\right) \min\{x^2, (1 - x^2)/2\} + \frac{4x(1 - x^2)}{49}.$$

Indeed,

$$x \le 1/\sqrt{3}: \quad \frac{8}{3} - U(x) = \frac{196 - 320x - 26x^2 + 12x^3}{147} > 0,$$
$$x \ge 1/\sqrt{3}: \quad \frac{8}{3} - U(x) = \frac{(1-x)(-12x^3 - 25x^2 + 127x - 56)}{147x} > 0.$$

In the first case the numerator decreases to $562/3 - 316/\sqrt{3} > 0$; in the second, the cubic factor has derivative at least 41 and begins at $123/\sqrt{3} - 193/3 > 0$. This proves the strict bound in (39).

C First-order finite certificates

Section 7 and the constants of Theorem 5.7 need exact inequalities. We give a second counting bound and, for one code, integer inequalities that certify an explainer and bound its challenge degree. We then verify the two supports used in the proof of Theorem 5.7.

C.1 Counting by successive derivatives

Instead of factoring as in Lemma 5.5, we sort solutions by successive derivatives in Y_1: a solution along which $\partial_{Y_1}^i Q$ vanishes but $\partial_{Y_1}^{i+1} Q$ does not is a regular solution (Section 5.2) of $\partial_{Y_1}^i Q = 0$ at stage i. Stopping when Y_1 disappears leaves an ordinary equation and avoids a characteristic restriction on the total jet degree. In our parameter searches this improved the bound only for $k \le 3$, including the $k = 2$ rows of Table 7.

Let $2 \le k \le A \le n$, $D = k - 1$, $\ell \ge 1$, $B_{\text{jet}} \ge 1$, $0 \le B_\partial \le B_{\text{jet}}$, and $B_Z \ge 0$. Fix a received curve $w(Z) = \sum_{t=0}^{\ell} Z^t w_t$ with $w_t \in F^n$. We record degree bounds for each stage and for the final ordinary equation. Put $\lambda_{\text{agr}} = (n - D)/(A - D)$ and $\tau = \max\{0, 2D - 3\}$. For each possible value $0 \le d \le B_\partial$ of $\deg_{Y_1} Q$, and each $0 \le i < d$, set

$$j_i = B_{\text{jet}} - i, \quad r_i = d - i, \quad b_i = 1 + \tau(j_i - 1), \quad c_i = \min\{b_i, \tau(r_i - 1) + D\}, \quad D_{\text{fiber},i} = j_i c_i + r_i(b_i - c_i).$$

Here j_i and r_i bound the total jet degree and the Y_1-degree of $\partial_{Y_1}^i Q$. The quantities b_i and c_i bound the corresponding degrees of the reconstruction map, and $D_{\text{fiber},i}$ is the contribution of this stage to the reduced generic-fiber degree. Summing over the stages gives the joint and reduced generic-fiber degree bounds

$$D_{\text{joint}}(d) = \sum_{i<d} \left[B_Z(2b_i c_i - c_i^2) + 2(\ell + \tau B_Z) D_{\text{fiber},i} \right], \qquad D_{\text{fiber}}(d) = \sum_{i<d} D_{\text{fiber},i}.$$

The ordinary equation left after the last stage has degree $B_{\text{jet}} - d$ in Y_0. Its exceptional challenges are counted by $\mathsf{E}_{\text{ord}}^{(\ell)}$ from (33), with the convention $\mathsf{E}_{\text{ord}}^{(\ell)}(D, A, n; 0, B_Z; L_0) = B_Z$ for an equation independent of Y_0, and we minimize its threshold L_0 independently of the threshold L for the regular stages:

$$\mathsf{E}_d^{(\ell)}(L) = \min_{D+1 \le L_0 \le A} \mathsf{E}_{\text{ord}}^{(\ell)}(D, A, n; B_{\text{jet}} - d, B_Z; L_0) + \frac{n - L + 1}{A - L + 1} \lambda_{\text{agr}} D_{\text{joint}}(d) + \ell(n - L) \frac{n - D}{L - D} D_{\text{fiber}}(d). \tag{111}$$

Lemma C.1 (First-order counting with an ordinary tail)**.** *With these parameters, let $Q \in F[X, Z, Y_0, Y_1]$ be nonzero, of total jet degree at most B_{jet}, degree in Y_1 at most B_∂, and challenge degree at most B_Z. Put $d_Q = \deg_{Y_1} Q$. Assume characteristic zero or $p > \max\{D, d_Q\}$. For a received-word-independent support certificate, the sufficient characteristic assumption is $p > \max\{D, B_\partial\}$. For each z outside one set of size at most $\max_{0 \le d \le B_\partial} \min_{k \le L \le A} \mathsf{E}_d^{(\ell)}(L)$, every degree-$\prec$ k solution of $Q(X, z, P, P^{[1]}) = 0$ with at least A agreements with $w(z)$ has the full common-witness conclusion of Definition 3.1. For a nonzero fixed-word equation, its close polynomial solutions number at most $\lambda_{\text{agr}} D_{\text{fiber}}(B_\partial) + B_{\text{jet}} - B_\partial$.*

Proof of Lemma C.1. If $D = n - 1$, then the hypotheses force $k = A = n$. Interpolate the coefficient words separately as $P_0, \dots, P_\ell \in F[X]_{\le D}$. A candidate with A agreements agrees everywhere, so it equals $\sum_t z^t P_t$ and its full agreement set is the common agreement set of the P_t. Thus the empty exceptional set suffices. A fixed word likewise has at most one candidate, which is within the displayed list bound. We may therefore assume $D \le n - 2$ below.

Take $d = d_Q$. Its successive Y_1 derivatives end at the nonzero leading coefficient $U = (1/d!)\partial_{Y_1}^d Q \in F[X, Z, Y_0]$. For a solution of $Q = 0$, either there is an index $0 \le i < d$ such that $\partial_{Y_1}^i Q$ vanishes identically after substitution while $\partial_{Y_1}^{i+1} Q$ is nonzero as a polynomial in X, or the solution satisfies $U(X, z, P) = 0$. In the first case, the solution is a regular solution of $\partial_{Y_1}^i Q = 0$ at stage i; in the second, it solves the ordinary equation $U = 0$, of degree at most $B_{\mathrm{jet}} - d$ in Y_0.

At a regular stage, Lemma 5.4, with $(j, r, H) = (j_i, r_i, B_Z)$, gives the degree bounds $(b_i, c_i, \ell + \tau B_Z)$. Lemma 5.3 bounds the image degrees of stage i by the ith summands of $D_{\mathrm{joint}}(d)$ and $D_{\mathrm{fiber}}(d)$. No bound $p > B_{\mathrm{jet}}$ enters this calculation.

For $b = B_{\mathrm{jet}} - d \ge 1$, apply Lemma 4.3 to U with threshold L_0; its bound already counts the challenges where the content of U vanishes identically or a factor acquires a repeated root. For $b = 0$, exclude the at most B_Z challenges where U vanishes identically. This gives the first term of (111). For the order-one families, apply Lemma 4.2 with $k = D + 1$, $r_1 = 1$, and the summed joint and reduced generic-fiber degrees $D_{\mathrm{joint}}(d)$ and $D_{\mathrm{fiber}}(d)$. Its two incidence factors are exactly the displayed multipliers of those degrees. This proves the bound $\mathsf{E}_d^{(\ell)}(L)$ for $d = d_Q$; taking the maximum over $0 \le d \le B_\partial$ gives a bound that does not depend on the received word.

For a fixed word, the ordinary equation has at most $B_{\mathrm{jet}} - d$ polynomial roots over $F(X)$, and the order-one families contribute at most $\lambda_{\mathrm{agr}} D_{\mathrm{fiber}}(d)$. As d increases, the bound for each stage $i < d$ increases and the new stage contributes at least one; since $\lambda_{\mathrm{agr}} \ge 1$, $\lambda_{\mathrm{agr}} D_{\mathrm{fiber}}(d) + B_{\mathrm{jet}} - d$ is nondecreasing. This proves the list bound. □

The actual degree d_Q can be smaller than B_∂, and then the final ordinary equation has Y_0-degree at most $B_{\mathrm{jet}} - d_Q$, a larger bound than $B_{\mathrm{jet}} - B_\partial$. A bound uniform over received words therefore takes the maximum over d.

C.2 A finite certificate for one code

At a concrete (n, k, A), exact counts improve on Lemma 3.16. We sharpen the support and local-rank counts, then bound the challenge degree by also counting scalar coefficients in Z. The resulting integer inequalities do not depend on the evaluation points or the received word.

Fix $n \ge 2$, $1 \le k < A \le n$, $m \ge 1$, $B_\partial \ge 0$, and $B_{\mathrm{jet}} \ge 1$, and put $K = \max\{k, 2\}$ and $D = K - 1$. Replace B_∂ by $\min\{B_\partial, B_{\mathrm{jet}}\}$ throughout: no monomial of total jet degree at most B_{jet} can have a larger derivative exponent. Thus we may assume $0 \le B_\partial \le B_{\mathrm{jet}}$. One can use either the equal-weight support $\mathcal{V}_{m, B_\partial, B_{\mathrm{jet}}}$ of (19) or the larger derivative-weighted support

$$\mathcal{V}^{\mathrm{wt}}_{m, B_\partial, B_{\mathrm{jet}}} = \mathrm{span}\{X^x Y_0^a Y_1^b : b \le B_\partial,\ a + b \le B_{\mathrm{jet}},\ x + Da + (D-1)b < mA\}. \tag{112}$$

Since $\deg P^{[1]} \le D - 1$ when $\deg P \le D$, we give Y_1 weight $D - 1$. For each total jet degree $t = a + b$ and derivative exponent b, the last inequality in (112) leaves $(mA - Dt + b)_+$ choices of the X exponent. For the local rank, group the substituted monomials by their T-exponent s (Lemma 3.14) and total jet degree t. The

rank of a group is at most the smaller of its numbers of source and target coefficients. These two counts give

$$N_{\mathrm{coef}} = \sum_{t=0}^{B_{\mathrm{jet}}} \sum_{b=0}^{\min(t,B_\partial)} (mA - Dt + b)_+, \tag{113}$$

$$R_{\mathrm{loc}} = \sum_{s=0}^{m-1} \sum_{t=0}^{B_{\mathrm{jet}}} \mathbf{1}_{\{s+(D-1)t<mA\}} \min\{m-s, (\min(t,s) - \max(0, t-B_\partial) + 1)_+\}. \tag{114}$$

For the equal-weight support, replace the dimension summand by $(mA - Dt)_+$ and the rank summand by

$$\min\{m-s, (\min(t,s) - \max(0, t-B_\partial, s+Dt-mA+1) + 1)_+\}.$$

Unlike (20), these bounds retain the total-degree bound B_{jet} and the specialization cutoff, and do not need $B_\partial \le m/2$.

Counting coefficients in the challenge as well. In the local constraints of a degree-ℓ received curve, an entry from source jet degree t to output degree u has challenge degree at most $\ell(t-u)$. A coefficient of challenge degree at most $B_Z - \ell t$ therefore gives conditions of degree at most $B_Z - \ell u$, and Lemma 3.11 compares the two scalar counts; negative bounds omit a column or row.

Proposition C.2 (A finite first-order certificate)**.** *Use either support and its corresponding counts above, with $N_{\mathrm{coef}} > nR_{\mathrm{loc}}$. Let $N_{\mathrm{coef}}(t)$ count source monomials of total jet degree t, and let $R_{\mathrm{loc}}(t)$ be the sum over s of the rank summands at that degree. Thus $N_{\mathrm{coef}} = \sum_t N_{\mathrm{coef}}(t)$ and $R_{\mathrm{loc}} = \sum_t R_{\mathrm{loc}}(t)$. For a degree-$\ell$ received curve, any integer $B_Z \ge 0$ satisfying*

$$\sum_{t=0}^{B_{\mathrm{jet}}} N_{\mathrm{coef}}(t)(B_Z - \ell t + 1)_+ > n \sum_{t=0}^{B_{\mathrm{jet}}} R_{\mathrm{loc}}(t)(B_Z - \ell t + 1)_+ \tag{115}$$

certifies a sound first-order interpolant of challenge degree at most B_Z, total jet degree at most B_{jet}, and Y_1-degree at most $\min(B_\partial, B_{\mathrm{jet}})$, nonzero at every specialization. The construction holds over every field, uniformly over evaluation sets and received curves. Characteristic restrictions enter only when reconstructing and counting polynomial solutions.

There is also a column-only test. If $N_{Y_0}(a)$ counts monomials with Y_0 exponent a, it suffices that

$$\sum_{a=0}^{B_{\mathrm{jet}}} N_{Y_0}(a)(B_Z - \ell a + 1)_+ > nR_{\mathrm{loc}}(B_Z + 1). \tag{116}$$

One may use whichever test gives the smaller certified challenge-degree bound.

Proof of Proposition C.2. Both supports are translation invariant and specialize to degree less than mA. At the origin, a source monomial has local image $T^{x+a}(Y_1 + \widetilde{V})^a Y_1^b$ modulo terms of total $(T, \widetilde{V})$-degree at least m. Group by $s = x + a$ and total jet degree $t = a + b$, assigning $\widetilde{V}$ degree one. The possible a lie between $\max(0, t - B_\partial)$ and $\min(t, s)$. The target block has at most $m - s$ coordinates. For the derivative-weighted support the degree cutoff becomes $s + (D-1)t < mA$; for the equal-weight support it also requires $a \ge s + Dt - mA + 1$. This proves the rank bounds degree by degree.

In each homogeneous image at the origin, choose at most $R_{\mathrm{loc}}(t)$ output coordinates whose projection is injective on that image over F. Translation is an automorphism of the support, so every translated local image over $F(Z)$ is the scalar extension of this same fixed image. The selected coordinates therefore detect all local constraints, without challenge-dependent row operations. An entry from source degree t to output degree u is

zero when $u > t$ and otherwise has challenge degree at most $\ell(t-u)$. Apply Lemma 3.11 with column cost ℓt, row cost ℓu, and height B_Z. Its two coefficient counts are exactly the sides of (115).

For (116), entries in a column of Y_0-exponent a have challenge degree at most ℓa. Choose at most nR_{loc} original rows spanning the stacked matrix over $F(Z)$ and apply the same lemma with column cost ℓa and row cost zero. The resulting primitive kernel vector annihilates the whole matrix and is nonzero at every challenge. The multiplicity argument of Proposition 3.13 proves soundness in both cases. □

To find a challenge-degree bound, write $W_{\mathrm{coef}} = \sum_t tN_{\mathrm{coef}}(t)$ and $W_{\mathrm{rank}} = \sum_t tR_{\mathrm{loc}}(t)$. For $B_Z \geq \ell B_{\mathrm{jet}}$, the row test reduces to

$$(N_{\mathrm{coef}} - nR_{\mathrm{loc}})(B_Z + 1) > \ell(W_{\mathrm{coef}} - nW_{\mathrm{rank}}).$$

Hence the following value of B_Z passes the row test:

$$\max\left\{\ell B_{\mathrm{jet}}, \left\lfloor \frac{\ell(W_{\mathrm{coef}} - nW_{\mathrm{rank}})}{N_{\mathrm{coef}} - nR_{\mathrm{loc}}} \right\rfloor\right\}.$$

Smaller values can be checked on the linear pieces of (115), whose breakpoints are $\ell t - 1$; the test itself need not be monotone in B_Z. For $B_Z \geq \ell B_{\mathrm{jet}}$, the column test has the analogous formula with $W_{Y_0} = \sum_a aN_{Y_0}(a)$ in place of W_{coef} and no W_{rank} term.

Using a finite certificate. Given a support with $N_{\mathrm{coef}} > nR_{\mathrm{loc}}$, we choose the challenge degree and the counting thresholds that give the smallest certified bound. We take the smaller challenge degree certified by (115) or (116). For $k \geq 2$ and characteristic zero or $p > \max\{k-1, B_\partial\}$, we evaluate Lemma 5.5 at $(B, M, H) = (B_{\mathrm{jet}}, B_\partial, B_Z)$ and Lemma C.1, using the ordinary transfer directly if $B_\partial = 0$, and minimize the regular threshold L and the ordinary threshold L_0 independently. Because the explainer's actual Y_1-degree can be smaller than B_∂, we minimize the successive-derivative bound for each possible Y_1-degree $0 \leq d \leq B_\partial$ before taking the maximum. We then use the smaller of the bounds whose hypotheses hold. Neither bound restricts the characteristic in terms of the curve degree ℓ.

For fixed d, both objectives are convex in the threshold on $[k, A]$: their nonlinear terms are positive multiples of $1 + (n-A)/(A-L+1)$ and $(n-k+1)/(L-k+1) - 1$, and any remaining term is linear. Thus a binary search on the sign of the forward difference finds an integer minimizer in $O(\log n)$ exact comparisons. When $A = n$, choose $L = n$.

The line count $\mathsf{E}^{(1)}_{\mathrm{exc}}$ is the smaller of (44) at $\ell = 1$, minimized over L and L_0, and

$$\max_{0 \leq d \leq B_\partial} \min_{k \leq L \leq A} \mathsf{E}^{(1)}_d(L);$$

under the characteristic hypothesis, the line error is at most $2^{-\sigma}$ when $q \geq 2^\sigma \lceil \mathsf{E}^{(1)}_{\mathrm{exc}} \rceil$, and affine sampling uses denominator $q - 1$. The same minimization at degree ℓ gives E_ℓ.

The list bound from the same certificate. The support also bounds the candidates for one received word; Corollary 6.4 transfers this bound to arrays without a loss depending on the number of rows.

Lemma C.3 (A finite first-order list-size bound)**.** *Suppose the support hypotheses of Proposition C.2 hold, $k \geq 2$, and the characteristic is zero or greater than* $\max\{k-1, B_\partial\}$. *Let $D_{\mathrm{fiber}}(B_\partial)$ be the successive-stage bound in Lemma C.1, let F_{reg}, S be as in (43) with $(B, M) = (B_{\mathrm{jet}}, B_\partial)$, and put $\lambda_{\mathrm{agr}} = (n-k+1)/(A-k+1)$. For $B_\partial \geq 1$, every received word has list size at most*

$$\Lambda = \left\lfloor \min\left\{\lambda_{\mathrm{agr}} F_{\mathrm{reg}} + S,\ \lambda_{\mathrm{agr}} D_{\mathrm{fiber}}(B_\partial) + B_{\mathrm{jet}} - B_\partial\right\}\right\rfloor$$

at agreement threshold A. If $B_\partial = 0$, the list size is at most B_{jet} in every characteristic.

Proof. For $B_\partial \geq 1$, apply the fixed-word conclusions of Lemmas C.1 and 5.5 and take the smaller integer bound. For $B_\partial = 0$, the nonzero interpolant is bivariate with Y_0-degree at most B_{jet}. Every close polynomial satisfies $Q(X, P(X)) = 0$, and a nonzero polynomial over $F(X)$ has at most its degree many roots. Both arguments hold over every field extension. □

The bound B_{jet} applies only when the certified interpolant does not involve Y_1, so the certificates in Section 7, all with $B_\partial \geq 1$, use Λ.

C.3 A uniform first-order certificate

We prove Theorem 5.7 using the supports $(m, B_\partial, B_{\mathrm{jet}}) = (12, 4, 22)$ for lists and $(12, 4, 23)$ for MCA. The computational supplement includes further certificates [DKT26a].

Proof of Theorem 5.7. It suffices to prove the claim at $\delta = 6/25$: increasing the agreement threshold only removes candidates, so it cannot enlarge the list or the exceptional set. For $k = 1$, Lemma 3.3 gives list size at most $\lfloor n/A \rfloor$ and exceptional count at most $\lfloor \binom{n}{2}/(A-1) \rfloor$; here $A \geq 2$. This conclusion holds in every characteristic. For $k \geq 2$, first consider characteristics $p \leq 4$. The hypothesis $p > k - 1$ implies $k \leq p \leq 3$. Counting k-subsets of agreeing positions gives

$$|\mathcal{L}(y, A)|\binom{A}{k} \leq \binom{n}{k}, \qquad |\mathcal{L}(y, A)| \leq \left(\frac{n}{A-k+1}\right)^k < (25/6)^3 < 73 < 307n.$$

For MCA at $n \geq 278$, put $D = k - 1$ and $a = A/n$. Then $D/n \leq (6/25)^2/4$ and the Johnson gap is at least $3/25$. Theorem 4.4 uses $m = 3$ and gives $|\mathcal{E}| < (343/3)n^2/D \leq (343/3)n^2$. For $n < 278$, interpolate witnesses F_T, G_T on each k-subset T. At every position outside their common agreement set, the residual $(F_T(\alpha_i) - f_i) + z(G_T(\alpha_i) - g_i)$ is a nonzero affine polynomial in z. Exclude its roots for all T and i; there are at most

$$\binom{n}{k}(n-k) \leq n^4 < 278^2 n^2 < 1325775n^2$$

such challenges. Outside this set, choose T among any close candidate's agreements. Interpolation gives $P = F_T + zG_T$, and no position outside the common agreement set can agree. Thus the full agreement sets coincide.

It remains to prove both claims in characteristic zero or for $p > 4$. Here (19) has $K = k$ and $D = k - 1$. For the list bound, use the equal-weight support $(m, B_\partial, B_{\mathrm{jet}}) = (12, 4, 22)$. The rank bounds grouped by total jet degree are

$$(R_{\mathrm{loc}}(0), \ldots, R_{\mathrm{loc}}(22)) = (12, 22, 30, 36, 40, 35, 30, 25, 20, 16, 12, 9, 6, 4, 2, 1, 0, 0, 0, 0, 0, 0, 0).$$

They sum to 300, and $\sum_t t R_{\mathrm{loc}}(t) = 1570$. With $N_{\mathrm{jet}}(t) = \min(t, 4) + 1$ as in (21) and $\rho = k/n$, put

$$N_{\mathrm{pos}}(t) = N_{\mathrm{jet}}(t)(72/25 + \rho(12 - t))_+.$$

For each monomial $Y_0^a Y_1^b$ with $a + b = t$, the number of exponents x in the support, divided by n, is $(12A/n - \rho t + t/n)_+$. Using the agreement lower bound as in Lemma 3.14 with $m = 12$ and $\delta = 6/25$, it is at least $72/25 + \rho(12 - t) + t/n$, so the number of source coefficients of jet degree t is at least $nN_{\mathrm{pos}}(t)$. Since $\sum_t N_{\mathrm{jet}}(t) = 105$ and $\sum_t N_{\mathrm{jet}}(t)(12 - t) = 5$, dropping the positive part gives

$$N_{\mathrm{coef}} \geq n(1512/5 + 5\rho) > 300n.$$

This proves $N_{\mathrm{coef}} > nR_{\mathrm{loc}} = 300n$ at every rate.

Apply Lemma 5.5 with $D = k - 1$; its characteristic requirement is $p > \max\{k - 1, 4\}$. Put $\lambda_{\mathrm{agr}} = (n - D)/(A - D)$. Since $A - D \geq 1 + \delta n$,

$$\lambda_{\mathrm{agr}} \leq \frac{1}{\delta}, \qquad D\lambda_{\mathrm{agr}} \leq \frac{D(n-D)}{\delta n} \leq \frac{n}{4\delta}.$$

For this support, (43) gives the exact regular-fiber degree bound

$$F = \begin{cases} 22, & D = 1, \\ 294D - 410, & D \geq 2, \end{cases} \qquad S = 138.$$

The first line is the identity map on the initial Taylor pair. In both cases $F \leq 294D - 134$. Consequently

$$\Lambda \leq \lambda_{\mathrm{agr}}(294D - 134) + 138 \leq \frac{294n^2}{4(\delta n + 1)} + 4 < \frac{1225}{4}n < 307n.$$

Here $\lambda_{\mathrm{agr}} \geq 1$ gives the +4, while $D\lambda_{\mathrm{agr}} \leq D(n-D)/(\delta n + 1) \leq n^2/(4(\delta n + 1))$. The penultimate inequality is $16 < (30529/25)n$ after clearing the positive denominator, and holds for $n \geq 2$, over every extension field.

For MCA, enlarge the same equal-weight support to $(m, B_\partial, B_{\mathrm{jet}}) = (12, 4, 23)$ and set $H = 276$. The added monomials, of total jet degree 23, have local rank zero, so the rank remains 300. Use the same normalized lower bound $N_{\mathrm{pos}}(t)$, now through $t = 23$ and for $0 \leq \rho \leq 19/25$ (as $k \leq A - \delta n$). Both the coefficient margin and the height-276 row margin of (115) are continuous and piecewise affine in ρ. Their only possible changes of slope occur at

$$\left\{0, \frac{19}{25}\right\} \cup \left\{\frac{72}{25s} : 4 \leq s \leq 11\right\}.$$

Checking these ten points gives

$$\min_{\rho}\left(\sum_{t=0}^{23} N_{\mathrm{pos}}(t) - 300\right) = \frac{204}{55}, \qquad \min_{\rho}\sum_{t=0}^{23}(N_{\mathrm{pos}}(t) - R_{\mathrm{loc}}(t))(277 - t) = \frac{14}{5}.$$

Both minima occur at $\rho = 72/275$. Thus Proposition C.2 gives a symbolic explainer of challenge degree at most 276 for every rate. The exact normalized coefficient count has the additional nonnegative term t/n inside the positive part, so both margins remain valid. For the endpoint $k = 2$, the hypothesis $A \geq k + \delta n$ gives $A \geq 3$; hence $12A \geq 36$, and the specialization cutoff is inactive through jet degree 23 in the displayed rank blocks. This verifies the coefficient and row-height tests for $D = 1$.

To count the exceptional challenges, apply Lemma 5.5 to this explainer with $(B, M, H) = (23, 4, 276)$ and $\ell = 1$. Its hypotheses hold: $M \leq B$, $D \leq n - 2$ because $k \leq 19n/25$, and $p > \max\{D, 4\}$. Then (43) gives $S = 145$ and $H_s = 1932$. At $D = 1$ it gives $F_{\mathrm{reg}} = 23$ and $J_1 = 322$; for $D \geq 2$,

$$F_{\mathrm{reg}} = 309D - 431 \leq 309D, \qquad J_1 = 497628D^2 - 1421886D + 1013438 \leq 497628D^2.$$

The $D = 1$ values obey the same upper bounds $309D$ and $497628D^2$. We use these bounds uniformly in what follows. Choose $L = D + \lceil (A - D)/42 \rceil$. Then

$$\frac{n - L + 1}{A - L + 1} \leq \frac{42}{41}\lambda_{\mathrm{agr}}, \qquad \frac{n - D}{L - D} \leq 42\lambda_{\mathrm{agr}}.$$

Bound the ordinary term of (44) by its value at $L_0 = k$, using $\Psi_D(S) \leq 1 + 4DS$. The preceding inequalities, $\lambda_{\mathrm{agr}} \leq 1/\delta$, $D\lambda_{\mathrm{agr}} \leq n/(4\delta)$, and $n \geq 2$ give

$$\begin{aligned} \frac{|\mathcal{E}|}{n^2} &\leq \frac{(2S-1)H_s}{4} + \frac{H_s + S}{4\delta} + \frac{SH_s}{2\delta} + \frac{S}{2} + \frac{42}{41}\frac{497628}{16\delta^2} + 42\frac{309}{4\delta} \\ &= \frac{317855987}{246} < 1325775. \end{aligned} \tag{117}$$

□

D Deferred proofs for higher derivatives

This appendix proves the estimates that Section 8 states without proof: the bounds on the coefficient-to-constraint ratio of the weighted support behind Theorem 8.1, and the joint Taylor-image bound behind Theorem 8.5.

D.1 The simplex moment

Proof of Lemma 8.2. Write $H_d = \sum_{j=1}^{d} 1/j$ and $H_d^{(2)} = \sum_{j=1}^{d} 1/j^2$ for the first- and second-order harmonic sums. For $t \geq 0$, inclusion–exclusion and a translation of the simplex give

$$\Pr[M_d > t] = \sum_{j=1}^{d} (-1)^{j+1} \binom{d}{j} (1 - jt)_+^d.$$

Integrating this identity for the first two moments yields

$$\mathbb{E}[M_d] = \frac{H_d}{d+1}, \qquad \mathbb{E}\big[M_d^2\big] = \frac{H_d^2 + H_d^{(2)}}{(d+1)(d+2)}.$$

The required alternating sums follow by expanding $1 - (1-x)^d$ and integrating after division by x; a second integration gives the sum with denominator j^2. These formulas give the mean and variance of $Z_d = dM_d - \log d$. For $t \geq 0$, the union bound gives

$$\Pr[Z_d > t] \leq d \left(1 - \frac{\log d + t}{d}\right)_+^d \leq e^{-t}.$$

To use the variance, split the second moment at a threshold θ:

$$\mathbb{E}\big[(\theta - Z_d)_+^2\big] = (\theta - \mathbb{E}[Z_d])^2 + \operatorname{Var}(Z_d) - \mathbb{E}\big[(Z_d - \theta)_+^2\big].$$

At $\theta = \log 6$, the preceding tail bound gives

$$\mathbb{E}\big[(Z_d - \log 6)_+^2\big] = \int_0^\infty 2u \Pr[Z_d > \log 6 + u]\, du \leq \frac{1}{3}.$$

It remains to show

$$(\log 6 - \mathbb{E}[Z_d])^2 + \operatorname{Var}(Z_d) - \frac{1}{3} > \frac{27}{10}. \tag{118}$$

For $d \geq 500$, we use

$$H_d - \log d < \frac{29}{50}, \qquad H_d^{(2)} > \frac{41}{25}, \qquad \log 6 > \frac{179}{100}.$$

For the first, $H_d - \log d$ decreases because $\log(1 + 1/d) > 1/(d+1)$, and the finite bounds $H_{200} < 58781/10000$ and $\log 200 > 52981/10000$ suffice. For the second, comparing $\sum_{j=21}^{d} j^{-2}$ with $\int_{21}^{d+1} x^{-2}\, dx$ gives $H_d^{(2)} \geq \sum_{j=1}^{20} j^{-2} + 1/21 - 1/(d+1) > 41/25$ for $d \geq 500$. Put $c = 29/50$, $v = 41/25$, $b = 121/100$, and $t_d = \log d + c$. The exact mean and variance imply

$$\mathbb{E}[Z_d] \leq c - \frac{t_d}{d+1}, \qquad \operatorname{Var}(Z_d) \geq v(1 - 3/d) - t_d^2/d.$$

Since $\log 6 - c > b$, substituting these estimates gives

$$(\log 6 - \mathbb{E}[Z_d])^2 + \operatorname{Var}(Z_d) - \frac{1}{3} > b^2 + v - \frac{1}{3} - \frac{P(t_d)}{d}, \quad P(t) = t^2 - 2b\frac{500}{501}t + 3v.$$

The function $P(\log d + c)/d$ decreases for $d \geq 500$: its derivative in d has the sign of $P'(t_d) - P(t_d)$, which is negative for $t_d > 6$. Since $\log 500 < 6.22$ and P increases there,

$$b^2 + v - \frac{1}{3} - \frac{P(6.8)}{500} = \frac{67667389}{25050000} > \frac{27}{10}.$$

This proves (118) and the lemma. The numerical bounds on logarithms can be verified in exact rational arithmetic: for $0 < z < 1$, truncate $2\sum_{j\geq 0} z^{2j+1}/(2j+1)$ and bound its tail by $2z^{2s+1}/((2s+1)(1-z^2))$ after s terms. □

D.2 The finite dimension ratio

To prove the finite ratio, we compare the coefficient count (59) with the following bound on the local rank, obtained by placing a unit cube at each lattice point:

$$\frac{R_{\mathrm{loc}}}{m^2 \operatorname{vol}(S_W)} \leq \frac{1}{d+1} \int_0^1 \left(1 - t + \frac{d+1}{m}\right)\left(1 + \frac{t + d(d+1)/(2m)}{w}\right)^d dt. \tag{119}$$

Proof of Lemma 8.3. For fixed b, compare the decreasing sum over b_0 with its integral, using $D \leq \rho n$ and $A_0 \geq an$:

$$\sum_{b_0 \geq 0} (mA_0 - D(b_0 + |b|))_+ \geq \frac{n}{2\rho}(ma - \rho|b|)_+^2.$$

For each point in the weighted simplex S_W, flooring coordinates gives an eligible b, decreases $|b|$, and assigns volume at most one to each lattice point. Integrating proves (59). The ordered-coordinate change of variables in Section 8.2 identifies $|U| = wM_d$ in distribution. Since $w \leq ad/(\rho \log(6d))$, we have $a\lambda/\rho - \log d \geq \log 6$. Thus

$$\frac{N_{\mathrm{coef}}}{nm^2 \operatorname{vol}(S_W)} \geq \frac{\rho}{2\lambda^2}\mathbb{E}\left[(a\lambda/\rho - dM_d)_+^2\right] > \frac{27\rho}{20\lambda^2}.$$

For the rank, the disjoint unit cubes based at eligible integer b lie in $S_{B+d(d+1)/2}$, proving $p_d(B) \leq (B + d(d+1)/2)^d/(d!)^2$. Also $\lceil (m-s)/(d+1) \rceil \leq (m - s + d)/(d+1)$. For $t \in [s/m, (s+1)/m]$, the integrand in (119) is at least

$$\left(1 - \frac{s}{m} + \frac{d}{m}\right)\left(1 + \frac{s/m + d(d+1)/(2m)}{w}\right)^d.$$

Summing these intervals proves that inequality. Using $(1+x)^d \leq e^{dx}$, and then extending the nonnegative integral after the substitution $u = 1 - t$, bounds it by

$$\frac{e^{\lambda(1+d(d+1)/(2m))}}{d+1} \int_0^\infty \left(u + \frac{d+1}{m}\right) e^{-\lambda u}\, du = \frac{e^{\lambda(1+d(d+1)/(2m))}}{(d+1)\lambda^2}\left(1 + \frac{(d+1)\lambda}{m}\right).$$

Dividing the coefficient-count estimate by the rank estimate proves the ratio bound. For $R_{\mathrm{loc}} = 0$, the dimension is positive because the constant monomial lies in the space. The limit follows from $W/m \to ad/(\rho \log(6d))$.

For the explicit choice, write $L = \log(6d)$, $w_0 = ad/(\rho L)$, $\lambda_0 = (\rho/a)L$, and $\epsilon = 1/(mw_0) \leq 1/(300d^3)$. Then

$$\lambda \leq \frac{\lambda_0}{1-\epsilon}, \qquad \lambda - \lambda_0 \leq \frac{2L}{300d^3}, \qquad \frac{\lambda d(d+1)}{2m} \leq \frac{1 + 1/d}{600(1-\epsilon)}, \qquad \log\left(1 + \frac{(d+1)\lambda}{m}\right) \leq \frac{1 + 1/d}{300d(1-\epsilon)}.$$

For $d \geq 519$, $L < d$, and the sum of these three losses is at most

$$\frac{2}{300d^2} + \frac{1 + 1/d}{600(1 - 1/(300d^3))} + \frac{1 + 1/d}{300d(1 - 1/(300d^3))}.$$

This expression decreases with d and is less than $1677/10^6$ at $d = 519$. This proves (62). □

D.3 The rate-uniform margin

Lemma D.1 (A rate-uniform lower bound on Γ)**.** *Let* $0 < \delta < 6/25$ *and* $d = d(\delta) = \lceil e^{1.5/\delta} \rceil$*. If* $\delta^2 \le \rho \le 1-\delta$ *and* $a = \rho + \delta$*, or if* $(\rho, a) = (2\delta^2, \delta)$*, then*

$$\log \Gamma(\rho, a, d) \ge \frac{3}{2} - \log(40/9).$$

Moreover

$$\frac{3}{2} - \log(40/9) - \frac{1677}{10^6} > \log(151/150) + \frac{1}{50000}.$$

Proof. Put $C = 40/9$. In the first case,

$$(\rho + \delta) \log \Gamma \ge \delta \log d - G(\rho, \delta), \qquad G(\rho, \delta) = \rho \log(C/\rho) + \delta \log \frac{20}{27\rho}.$$

For $\delta^2 \le \rho \le 1 - \delta$, we claim $G(\rho, \delta) \le \log C$. On $[\delta, 1-\delta]$, the function is concave and its derivative at the right endpoint is at least $\log C - 4/3 > 0$, so its maximum is there. At that endpoint,

$$G(1-\delta, \delta) = \log C - \delta \log 6 - \log(1-\delta) \le \log C,$$

using $-\log(1-\delta) \le 4\delta/3 < \delta \log 6$. On $[\delta^2, \delta]$, the function is convex, so its maximum is at an endpoint; $\rho = \delta$ lies in the interval treated above, and the other endpoint satisfies

$$G(\delta^2, \delta) < \frac{\log C}{16} + \frac{\log 4}{8} + \frac{\log 4}{2} < 1 < \log C.$$

Here we used $\delta < 1/4$ and monotonicity of $-\delta \log \delta$ and $-\delta^2 \log \delta$ on this range. Since $\delta \log d \ge 3/2$ and $\rho + \delta \le 1$, the first claim follows in this case.

In the second case the same lower bound follows from the stronger estimate

$$\log \Gamma \ge \frac{3}{2\delta} - 3 + \log(27/10) + 2 \log \delta - 2\delta \log 6.$$

The right side has negative derivative for $0 < \delta \le 1/4$, and at $\delta = 1/4$ it exceeds $3/10 > 3/2 - \log(40/9)$.

For the rational inequality, combine the logarithms as $\log(604/135)$ and put $z = 469/739$. The rational Taylor bound gives

$$\log(604/135) < 2 \sum_{j=0}^{19} \frac{z^{2j+1}}{2j+1} + \frac{2z^{41}}{41(1-z^2)} < \frac{1498303}{10^6}.$$

The second inequality is an exact rational comparison. The remaining margin is exactly $3/2 - 1677/10^6 - 1498303/10^6 = 1/50000$. □

D.4 The joint Taylor image of a received curve

Lemma 4.2 needs the degrees of the family of Taylor images over all challenges and of its generic fiber. As in Zheng and Jeronimo [Zhe26, Jer26], the challenge stays a variable: separate challenge and initial-jet degrees bound the first, and a transcendental challenge bounds the second. A *root pair* is a pair (z, P) with $Q(X, z, P, P^{[1]}, \dots, P^{[d]}) = 0$.

Lemma D.2 (Joint envelope and generic fiber)**.** *Let* $1 \le K \le n$, $0 \le d < K$, $\ell \ge 1$, *and let* $Q \in F[X, Z, Y_0, \dots, Y_d] \setminus \{0\}$ *have total jet degree at most* $B_{\mathrm{jet}} \ge 1$ *and challenge degree at most* $B_Z \ge 0$. *Assume characteristic zero or greater than* $\max\{K-1, B_{\mathrm{jet}}\}$. *For each actual separant-chain stage of jet degree* j *and highest jet index* r, *put* $\Delta_j = 1 + 2K(j-1)$ *and*

$$J_{j,r}^{(\ell)} = B_Z \Delta_j^{r+1} + (r+1) j (\ell + 2KB_Z) \Delta_j^r, \qquad F_{j,r} = j\Delta_j^r. \tag{120}$$

There are an algebraically closed extension E/F, *a set* $Z_0 \subseteq \overline{F}$ *with* $|Z_0| \le B_Z$, *and stage images* $V_{j,r} \subseteq E^{\ell+K}$ *of dimension at most* $r+1$ *and total degree at most* $J_{j,r}^{(\ell)}$. *Their union contains every lifted root pair* $(z, z^2, \dots, z^\ell, P)$ *over* $\overline{F}$ *with* $\deg P < K$ *and* $z \notin Z_0$. *Over a challenge transcendental over* E, *each reduced geometric fiber has dimension at most* r *and degree at most* $F_{j,r}$. *In particular, writing* $B_{\mathrm{lift}} = 1 + 2K(B_{\mathrm{jet}} - 1)$, *the union has joint degree at most*

$$D_{\mathrm{joint}}^{(\ell)} = B_{\mathrm{jet}} B_{\mathrm{lift}}^d [B_Z B_{\mathrm{lift}} + (d+1) B_{\mathrm{jet}} (\ell + 2KB_Z)]$$

and reduced generic-fiber degree at most $B_{\mathrm{jet}}^2 B_{\mathrm{lift}}^d$.

Proof. Choose a center c transcendental over $\overline{F}$ and put $E = \overline{\overline{F}(c)}$. Successive highest-jet derivatives (91) of Q are nonzero and strictly decrease total jet degree, ending at $T_* \in F[X, Z] \setminus \{0\}$. Each has challenge degree at most B_Z, so $T_*(X, z)$ vanishes identically on a set Z_0 of size at most B_Z. If Q has jet degree zero, the empty union suffices.

Outside Z_0, every root pair is regular at some stage T, with separant T_{Y_r}. For its degrees (j, r), put $S(Z, U) = T_{Y_r}(c, Z, U)$; the locus

$$\mathcal{U}^\circ = \{(z, U) : T(c, z, U) = 0,\ S(z, U) \ne 0\}$$

has closure of bidegree at most (B_Z, j). Write $m_* = K - 1 - r$. Retaining Z in (92), for jet degree j, gives

$$\tau_{r+m} = \frac{\widehat{N}_m(Z, U)}{S(Z, U)^{2m-1}}, \quad \deg_U \widehat{N}_m \le (2m-1)(j-1) + 1, \quad \deg_Z \widehat{N}_m \le (2m-1) B_Z.$$

The last bound follows by clearing S^{2m-2} in the numerator induction, and a coefficient of T adds one further B_Z. With common denominator S^{2m_*-1}, the map to $(Z, Z^2, \dots, Z^\ell, \tau_0, \dots, \tau_{K-1})$ has numerator and denominator degrees at most $\ell + 2KB_Z$ in Z and Δ_j in U. If $m_* = 0$, denominator one satisfies the same bounds. Transcendence of c puts every regular root pair in this image. Lemma A.5 with $d = r$ gives the degree $J_{j,r}^{(\ell)}$ and dimension at most $r+1$.

Over a transcendental challenge, the source is an open part of a nonzero hypersurface of U-degree at most j, and the specialized map has degree at most Δ_j. Its image closure therefore has dimension at most r and degree at most $j\Delta_j^r$ by Lemma A.4. This is the reduced generic fiber of the joint image closure: the constructible image contains a dense open subset of each component of its closure [Mil09, Chapter 10]. Vertical components miss the generic fiber; on each remaining component, the omitted subset has smaller fiber dimension, so the image remains dense. Adding powers of Z does not change this fiber. There are at most B_{jet} stages, their jet degrees sum to at most B_{jet}^2, and $r \le d$. Summing gives the two coarse bounds. □

E Arithmetic for the decoding algorithms

This appendix gives the arithmetic behind the bit-cost bounds of Section 9.

E.1 Field presentations

Every decoding theorem assumes the following presentation of $\mathbb{F}_q$, which Section 9.1 summarizes. Write $q = p^e$ and $b_q = \lceil \log_2 q \rceil$.

Definition E.1 (Field presentation for bit costs)**.** A *presentation* of a finite field $F = \mathbb{F}_q$ for bit costs consists of p, e and presentation data of size $\widetilde{O}(b_q)$ with the following deterministic operations.

(i) Every element has a canonical $O(b_q)$-bit encoding. Equality, addition, subtraction, multiplication, inversion of a nonzero element, and embedding an element of the prime field each cost $\widetilde{O}(b_q)$ bit operations.

(ii) There is a bijection $\mathrm{elt}_F : \{0, \dots, q-1\} \to F$ whose value at any index is computed in $\widetilde{O}(b_q)$ bit operations.

(iii) After one reusable preprocessing step of cost $\widetilde{O}(pb_q)$, a batch of $N \geq 1$ elements can be mapped to their unique p-th roots in $\widetilde{O}(pNb_q)$ bit operations. The preprocessing is charged only if inverse Frobenius is invoked.

Primality, the cardinality and validity of the presentation, and distinctness of the evaluation points are input promises.

Polynomial bases, and quotient extensions by a given irreducible polynomial, meet these requirements for every e (Lemma E.3 and Corollary E.4); we do not charge for constructing the presentation from q. The factor p in clause (iii) adds no dependence on the field size: the bivariate radicals take p-th roots only when p is bounded, and a univariate root step of degree D_u occurs only when $p \leq D_u$ and takes roots of $O(D_u/p)$ coefficients (Lemma E.6).

Remark E.2 (Slower explicit presentations)**.** Correctness and the field-operation counts do not depend on Definition E.1. The same algorithms work with any effective finite-field representation after charging its actual costs for arithmetic, enumeration, and inverse Frobenius. For example, a basis with structure constants gives deterministic time polynomial in n and in the length of that presentation [Len91]. Such a description can be much larger than $\log q$, so it does not imply the quasilinear dependence on $\log q$.

Lemma E.3 (Closure under a given quotient extension)**.** *Let F_0 be presented as in Definition E.1, and let*

$$L = F_0[\theta]/(f(\theta)), \qquad f(\theta) = \sum_{i=0}^{m} c_i \theta^i \in F_0[\theta]$$

for a given monic irreducible polynomial f. Represent elements of L in the relative power basis $1, \theta, \dots, \theta^{m-1}$. Then L, in this basis, is presented as in Definition E.1, without conversion to an absolute polynomial basis.

Proof. The relative basis has $\widetilde{O}(m \log |F_0|) = \widetilde{O}(\log |L|)$ bits. Fast arithmetic and extended-gcd inversion cost $\widetilde{O}(m)$ operations in F_0 [Sch15]. Concatenated encodings are canonical, and mixed-radix enumeration applies elt_{F_0} to each digit within the same bit bound.

For inverse Frobenius, put $\beta = \theta^{1/p}$ and $f^{(-1)}(x) = \sum_i c_i^{1/p} x^i$. This is separable, and $f^{(-1)}(\beta) = 0$, so

$$R_f(x) = f^{(-1)}(x) \bmod (x^p - \theta) = \sum_{r=0}^{p-1} \left(\sum_j c_{pj+r}^{1/p} \theta^j \right) x^r, \qquad \gcd(R_f, x^p - \theta) = x - \beta.$$

One batch of $m+1$ inverse-Frobenius evaluations in F_0, this gcd, and the powers $1, \beta, \dots, \beta^{p-1}$ cost $\widetilde{O}(p \log |L|)$ bits. For $a = \sum_{i=0}^{m-1} a_i \theta^i$, the precomputed powers give

$$a^{1/p} = \sum_{r=0}^{p-1} \beta^r \left(\sum_j a_{pj+r}^{1/p} \theta^j \right),$$

as taking p-th powers verifies. For a batch of N elements, one batch of Nm coefficient roots in F_0, together with the coordinate scans and $O(pN)$ operations in L, costs $\widetilde{O}(pN \log |L|)$ bits. □

Corollary E.4 (Batched inverse Frobenius in a polynomial basis)**.** *Let $F = \mathbb{F}_p[\theta] = \mathbb{F}_p[t]/f(t)$ be given in a polynomial basis, where f is monic irreducible of degree e and $q = p^e$. This presentation satisfies Definition E.1. In particular, one can preprocess F in $\widetilde{O}(p \log q)$ bit operations and then compute the unique p-th roots of any batch of N elements deterministically in $\widetilde{O}(pN \log q)$ bit operations.*

Proof. In the integer representation of $\mathbb{F}_p$, modular arithmetic and Euclid's algorithm give clauses (i)–(ii), and inverse Frobenius is the identity. Apply Lemma E.3 with $F_0 = \mathbb{F}_p$, $m = e$. □

Applying the same lemma repeatedly covers a tower of a fixed number of quotient extensions by given irreducible polynomials, with constants depending on that number. The decoders use towers of height at most two over $\mathbb{F}_q$ (Sections E.2 and E.6).

E.2 Centers and auxiliary fields

Each decoder needs a center x_0 at which a nonzero polynomial of degree at most $\Theta = O_{\text{params}}(n)$, such as a specialized separant or the Ω_0 of Algorithm 2, does not vanish; $\Theta + 1$ distinct elements contain one. Since only $q \geq n$ is assumed, $\mathbb{F}_q$ can be too small. If $q > \Theta$, use the first $\Theta + 1$ values of $\text{elt}_{\mathbb{F}_q}$, not $0, 1, 2, \dots$, which lie in the prime field. Otherwise $q = O_{\text{params}}(n)$: construct $E = \mathbb{F}_q[\theta]/(f(\theta))$ with $f = \theta^2 - \gamma$ for a nonsquare γ, found with Euler's criterion, when q is odd, or $f = \theta^2 + \theta + \gamma$ with $\text{Tr}(\gamma) = 1$ when $q = 2^e$. The trace polynomial is nonzero of degree $2^{e-1} < q$, and the image of $x \mapsto x^2 + x$ is the trace-zero subspace of size $q/2$, proving existence and irreducibility.

At most q tests of $O(\log q)$ operations cost $\widetilde{O}_{\text{params}}(n)$ on this branch. Outside bounded lengths, $q^2 > \Theta$; mixed-radix enumeration and Lemma E.3 give centers and arithmetic in this field E, where all curve-filter and Taylor operations take place.

The Rojas calls of Section E.6 need a field $E' \supseteq E$ satisfying (123), with $\mathbb{F}_4 \subseteq E'$ when $p = 2$; take $E' = E$ when E qualifies. If E is too small, let E' be the extension of E of least degree satisfying it, doubling $[E' : E]$ if needed when $p = 2$ so that the absolute degree of E' is even. Then $|E'| = n^{O(r+1)}$, and $[E' : E]$ is bounded in terms of the parameters since $q \geq n$. Enumerating monic polynomials through elt_E and testing irreducibility constructs E' in $n^{O(r+1)}$ operations; this search never takes place in an unboundedly large field. If E is large enough but $p = 2$ and $[E : \mathbb{F}_2]$ is odd, adjoin a root of $x^2 + x + 1$, which is irreducible over E by the trace criterion above, since $\text{Tr}(1) = [E : \mathbb{F}_2] \bmod 2 = 1$. In every case $\log |E'| = O_{\text{params}}(\log q)$, and Lemma E.3 gives arithmetic and batched inverse Frobenius in E'; this is the source of the dependence on $\log q$ in Theorem 9.11.

E.3 Splitting finite algebras

The next lemma, cited in the proof of Lemma 9.2, counts the bookkeeping around the splits of Algorithm 1 and the localization (78).

Lemma E.5 (Arithmetic for agreement recovery)**.** *Fix $B \geq 1$. Let $E/\mathbb{F}_q$ be an explicitly represented extension of parameter-bounded degree. Let $G \in E[u]$ be monic and squarefree, and let*

$$h \in (E[u]/G)[v]$$

be monic of degree $1 \leq b \leq B$, with no reducedness assumption on its fibers. Put

$$\mathcal{A} = (E[u]/G)[v]/h, \qquad \Delta = \dim_E \mathcal{A} = b \deg G.$$

Given $\widetilde{P}(X) \in \mathcal{A}[X]_{<k}$ in the tower basis, distinct abscissae $\alpha_i \in \mathbb{F}_q$, received values $y_i \in \mathbb{F}_q$, and $1 \leq k \leq A \leq n$, the residuals, coordinate transfers, splits, agreement histories, and final checks of Algorithm 1 have deterministic total cost $\widetilde{O}_B((n+k)\Delta)$ operations in E, without irreducible factorization, root extraction, or a primitive element for $\mathcal{A}$. The same bound holds for a collection of towers of total dimension Δ. Each split sends the geometric points at which the residual vanishes to factors on which it is nilpotent and the other points to factors on which it is a unit, and preserves total dimension.

A given $s \in (E[u]/G)[v]$ can also be localized in $\widetilde{O}_B(\Delta)$ operations. The output towers represent exactly the geometric points with $G = h = 0$ and $s \neq 0$, retain their nilpotents, have total dimension at most Δ, and make s a unit. No comparison between b and the characteristic is needed.

Proof. The correctness of each split, the preservation of total dimension, and the cost $\widetilde{O}_B(\dim_E \mathcal{R})$ per split are proved in Section 9.3. Localizing at s is the split (78), so it has the same cost and properties; b enters only through the power s^b, never through a comparison with p. We count the remaining batching costs.

Residuals and their transfer to the current factors. We count operations in E. Shared multipoint evaluation of the Δ scalar coefficient polynomials of $\widetilde{P}$ in the tower basis forms all residuals in $\widetilde{O}((n+k)\Delta)$ operations. At position i, every current factor $\mathcal{R}$ needs e_i modulo its own base polynomial $G_{\mathcal{R}}$: remainder trees reduce the at most B scalar v-coefficients of e_i modulo all of them at once. Repeated moduli are allowed, and $\sum_{\mathcal{R}} \deg G_{\mathcal{R}} \leq \sum_{\mathcal{R}} \dim_E \mathcal{R} \leq \Delta$ over the current factors. Reduction modulo the fibers and splitting therefore cost $\widetilde{O}_B(\Delta)$ per position.

Histories and final checks. Store histories as linked lists: the two children of a split share the parent's list, and only the nilpotent child appends the index i. A history's k indices are read out only when its factor is finished. At most Δ factors are current at each position, so the total history cost is $\widetilde{O}_B(n\Delta)$. Finished factors are disjoint and have positive dimension, so at most Δ interpolants need full-word checks. Each check interpolates k pairs and evaluates at n points, so checks and deduplication cost $\widetilde{O}((n+k)\Delta)$, including collections of towers and the univariate case $b = 1$. □

E.4 Radicals and a separable message equation

The bivariate radical of Section 9.4 and the univariate radicals of CurveCandidates, IsolatedSystemCandidates, and Rojas's algorithm share one recursion. We count its inverse-Frobenius calls, the only part of their arithmetic that depends on the characteristic.

The radical recursion. Let $m_{\max}$ bound the multiplicities of the irreducible factors of f, and let ∂f denote all first partials of f. Radical removes multiplicities before SeparablePart tests separability in Y.

Lemma E.6 (Fast characteristic-safe radical)**.** *Let F be a perfect field of characteristic $p > 0$. If $f \in F[u]$ is monic of degree D_u, then Radical(f, D_u) computes* rad(f) *using $\widetilde{O}(D_u)$ operations in F, excluding inverse Frobenius, together with batched inverse Frobenius on $O(D_u)$ coefficients across the recursion.*

Algorithm 7 Radical computation and the separable message divisor

Require: A polynomial f over a perfect field and a multiplicity bound $m_{\max}$

Ensure: Its radical, without irreducible factorization

1: **procedure** RADICAL($f, m_{\max}$)
2: If f is constant, return 1.
3: Set $f_0 = f/\gcd(f, \partial f)$ and $f_1 = f/\gcd(f, f_0^{m_{\max}})$.
4: If f_1 is constant, return f_0.
5: Form $f_1^{1/p}$ by dividing exponents by p and applying inverse Frobenius to coefficients.
6: Return f_0 RADICAL($f_1^{1/p}, m_{\max}$).
7: **end procedure**
8: For primitive bivariate Q, SEPARABLEPART(Q, ℓ) computes $Q_{\mathrm{rad}} = \mathrm{RADICAL}(Q, \ell)$ and returns the primitive part of $Q_{\mathrm{rad}}/\gcd(Q_{\mathrm{rad}}, (Q_{\mathrm{rad}})_Y)$.

If $f \in F[X, Y]$ is primitive in Y with $\deg_Y f \le \ell$ and $\deg_X f \le D_X$, then RADICAL(f, ℓ) computes the product of the distinct irreducible factors of f using $\widetilde{O}_\ell(D_X + 1)$ operations and inverse Frobenius on $O_\ell(D_X + 1)$ coefficients. Over $F = \mathbb{F}_q$ presented as in Definition E.1, the total bit costs are $\widetilde{O}(D_u \log q)$ and $\widetilde{O}_\ell((D_X + 1) \log q)$.

Proof. For the analysis only, write $f = \lambda \prod_i \phi_i^{m_i}$ with distinct irreducibles ϕ_i; the algorithm does not compute this factorization. Every $m_i \le m_{\max}$: in the univariate case $m_{\max} = D_u$, and in the bivariate case primitivity gives $\deg_Y \phi_i \ge 1$, hence $m_i \le \ell$. Because F is perfect, an irreducible ϕ_i has a nonzero partial derivative, so $\gcd(f, \partial f)$ contains ϕ_i to exponent $m_i - 1$ if $p \nmid m_i$ and to exponent m_i if $p \mid m_i$. Thus f_0 contains exactly once the factors with $p \nmid m_i$, and the definition of f_1 removes all their copies. Every remaining multiplicity is divisible by p, so f_1 is a p-th power over the perfect field F, coprime to f_0, and

$$\mathrm{rad}(f) = f_0 \,\mathrm{rad}\big(f_1^{1/p}\big). \tag{121}$$

Induction proves correctness, and the algorithm never computes any multiplicity m_i. If f_1 is nonconstant, then $p \le m_{\max}$, and taking $f_1^{1/p}$ divides every degree by p.

In the univariate case, fast modular powering, gcds, and exact divisions cost $\widetilde{O}(D_u)$ operations in F per level, and $f_1^{1/p}$ has at most $D_u/p + 1$ coefficients. One batched inverse-Frobenius call, with preprocessing, costs $\widetilde{O}(p(D_u/p + 2) \log q) = \widetilde{O}(D_u \log q)$ bits. These costs sum geometrically over the recursion.

In the bivariate case, all gcds and exact quotients are taken over $F(X)[Y]$ and made primitive by clearing denominators and removing content. There are $O_\ell(1)$ recursion levels. All intermediate Y-degrees are bounded in terms of ℓ, including $\deg_Y f_0^\ell \le \ell^2$, and their X-degrees are $O_\ell(D_X + 1)$. Subresultant polynomial matrices have bounded size, so each level costs $\widetilde{O}_\ell(D_X + 1)$ field operations. At each root step $p \le \ell$, and inverse Frobenius is applied to $O_\ell(D_X + 1)$ coefficients, which clause (iii) of Definition E.1 charges $\widetilde{O}_\ell((D_X + 1) \log q)$ bits, including its preprocessing. □

Lemma E.7 (Separable preprocessing with bounded Y-degree)**.** *Let $Q \in F[X, Y]$ be nonzero and primitive in Y, with $\deg_Y Q \le \ell$ and $\deg_X Q \le D_X$, over a finite field $F = \mathbb{F}_q$ presented as in Definition E.1. One can compute a primitive divisor T of Q that is separable over $F(X)$ and retains every factor $Y - P(X)$, $P \in F[X]$, in $\widetilde{O}_\ell((D_X + 1) \log q)$ deterministic bit operations. In particular, $\deg_Y T \le \ell$ and $\deg_X T \le D_X$.*

Proof. Run SEPARABLEPART(Q, ℓ). By Lemma E.6, Q_{rad} is the product of the distinct irreducible factors of Q, computed within the stated bit bound. Section 9.4 shows that the quotient $Q_{\mathrm{rad}}/\gcd(Q_{\mathrm{rad}}, (Q_{\mathrm{rad}})_Y)$ is separable over $F(X)$ and retains every factor $Y - P(X)$. Its primitive part T divides Q, which proves the degree bounds, and the final gcd and content removal cost $\widetilde{O}_\ell(D_X + 1)$ field operations. □

E.5 Relaxed multiplication for the Taylor recurrence

The triangular recurrence (81) is derived in Section 9.5. We implement its successive steps using van der Hoeven's relaxed multiplication [vdH02]: previously computed coefficient blocks are multiplied together once, instead of recomputing each series product at every step.

Lemma E.8 (Fast triangular Taylor recurrence)**.** *Fix $r, B_{\mathrm{jet}} \geq 0$, let R be a commutative ring, and let $0 \leq r < k$. Suppose*

$$T \in R[X, Y_0, \ldots, Y_r]$$

has jet degree at most B_{jet}, the center $x_0 \in R$ and the initial jet $U = (U_0, \ldots, U_r) \in R^{r+1}$ make the constant residual $T(x_0, U)$ zero, and

$$S = T_{Y_r}(x_0, U) \quad \text{and} \quad \binom{r+m}{r} \quad (1 \leq m < k - r)$$

are units, with their inverses given. From the Z-coefficient streams of $T(x_0 + Z, Y_0, \ldots, Y_r)$ through degree $k - r - 1$, we can compute the unique $Y = \sum_{j<k} \tau_j Z^j$ with $\tau_j = U_j$ for $j \leq r$ satisfying

$$T(x_0 + Z, Y, Y^{[1]}, \ldots, Y^{[r]}) = 0 \pmod{Z^{k-r}}.$$

If $M_R(\ell)$ is the cost of multiplying length-ℓ streams over R, the cost is

$$O_{r, B_{\mathrm{jet}}}\left(\sum_{\substack{\ell = 1, 2, 4, \ldots \\ \ell \leq k}} \frac{k}{\ell} M_R(\ell) \right)$$

operations in the representation of R. When $r = 0$, the binomial condition holds in every characteristic.

Proof. The triangular identity (81) is polynomial, hence valid over R. Its unit leading multiplier successively determines $\tau_{r+1}, \ldots, \tau_{k-1}$ and proves uniqueness to this precision. We now bound the cost of the relaxed schedule.

For speed, evaluate $T(x_0 + Z, \cdot)$ by a fixed circuit and keep, at every multiplication gate, the Z-coefficients computed so far. For inputs a, b, only two terms involve the current coefficient:

$$[Z^m](ab) = a_0 b_m + a_m b_0 + \sum_{\substack{i, j \geq 1 \\ i + j = m}} a_i b_j. \tag{122}$$

The remaining sum is computed in blocks [vdH02]: for each dyadic ℓ, multiply the index block $[\ell, 2\ell)$ by itself and, in both orders, by the later blocks $[j\ell, (j+1)\ell)$, $j \geq 2$. These products partition the positive-index pairs, and a block ending at $(j+1)\ell - 1$ is first needed at index $(j+1)\ell$, when it is complete. Hence the gate coefficients at index m are affine in τ_{r+m}; we solve (81) for it and then complete the gate coefficients at index m. Each scale ℓ has $O(k/\ell)$ products of length $O(\ell)$, which gives the bound. □

E.6 Isolated roots of the higher-order systems

We check the hypotheses of Rojas's theorem [Roj99, Main Theorem 1] for the square systems (84) of Algorithm 6 and count its cost. These have $r+1$ equations in $r+1$ variables, of degree at most $\mathcal{D}_{\mathrm{sys}} = O_{B_{\mathrm{jet}}}(n)$, over a field $E' \supseteq E$. The final clause of the theorem adjoins the origin to each support to obtain the conclusion quoted in Section 9.7; our dense degree-$\mathcal{D}_{\mathrm{sys}}$ supports already contain it. The theorem applies in every

characteristic, even when positive-dimensional components are present, and counts arithmetic steps in a finite extension of the coefficient field [Roj99, Main Theorem 1, §5].

For dense supports the mixed volume is $\mathcal{M} = \mathcal{D}_{\mathrm{sys}}^{r+1}$, the start system $(x_1^{\mathcal{D}_{\mathrm{sys}}}, \dots, x_{r+1}^{\mathcal{D}_{\mathrm{sys}}})$ needs no fill search, and the matrices are Macaulay's, of size $N_{\mathrm{Mac}} = \binom{(r+1)\mathcal{D}_{\mathrm{sys}}+1}{r+1}$ in both dimensions [Roj99, §3.4]. Rojas's bound $O^*\big((r+1)^4\mathcal{M}^3 N_{\mathrm{Mac}}^2 N_{\mathrm{Mac}}^{2.376}\big)$ is then $(\mathcal{D}_{\mathrm{sys}}+1)^{O(r+1)}$ operations in E'. Field elements are needed only as at most $N_{\mathrm{Mac}}+1$ interpolation nodes and as the derandomizing parameter, with at most $(r+1)(2r+3)\binom{\mathcal{M}}{2}$ bad values [Roj99, §5.2]. We take both from the first elements of E' in the order of $\mathrm{elt}_{E'}$, trying successive parameters until the genericity test succeeds; this needs

$$|E'| > (r+2)^2 \mathcal{D}_{\mathrm{sys}}^{2(r+1)}, \tag{123}$$

with $\mathbb{F}_4 \subseteq E'$ when $p = 2$ (Step 3 uses a primitive cube root of unity).

Rojas charges a univariate squarefree part $O^*(\deg)$ field operations, which in characteristic p requires p-th roots. These occur only in Steps 2 and 3 and the genericity test: at most $2r+3$ polynomials of degree at most $\mathcal{M}$ per trial, over at most $(r+1)(2r+3)\binom{\mathcal{M}}{2}+1$ trials. Lemma E.6 computes each in $\widetilde{O}(\mathcal{M}\log|E'|)$ bit operations, for a total of $(\mathcal{D}_{\mathrm{sys}}+1)^{O(r+1)}\widetilde{O}(\log|E'|)$. The remaining steps of Algorithm 6, namely the radical of F_{iso} (again by Lemma E.6), the gcds with s and B_0, and the evaluation of N_j/B_0 at $(\psi_1, \dots, \psi_{r+1})$, a multivariate modular composition, cost $n^{O(r+1)}$ operations in E' per system.

E.7 Bounded lengths and small message degree

We decode directly the instances that the theorems of Section 9 exclude by a parameter-dependent bound. Each theorem fixes one threshold N_* for its length conditions (at order zero, that of (79)). When $n < N_*$, subset interpolation enumerates k-subsets of received positions, interpolates, checks agreement, and deduplicates; since $A \geq k$, the list is exact. There are at most 2^{N_*} subsets, and arithmetic costs $\widetilde{O}(\log q)$ bits even when q is arbitrarily large. For $k = 1$, frequency counting returns exactly the constants with at least A agreements in $\widetilde{O}(n)$ comparisons. The same subset algorithm solves the case $k \leq d$ of the order-d theorem at cost $\widetilde{O}_d(n^{k+1}\log q)$, within that theorem's bound; when $k > d$, every order $r \leq d$ used by the decoder satisfies $r < k$.

If a field-cardinality condition $q > q_0$ with $q_0 = O_{\mathrm{params}}(1)$ fails, as in Lemma 9.5, then $n \leq q \leq q_0$. Likewise, if the quadratic center field is too small, $n^2 \leq q^2 \leq \Theta = O_{\mathrm{params}}(n)$ bounds n. These branches enumerate received positions, not the field. At bounded message dimension, which Theorem 9.14 meets when $p \leq B_{\mathrm{jet}}$, long codes have low enough rate for the zeroth-order decoder.

Lemma E.9 (Bounded message dimension at fixed agreement)**.** *Fix $K_0 \geq 1$ and $0 < a_0 \leq 1$. For $k \leq K_0$ and $k \leq A \leq n$ with $A \geq a_0 n$, the exact agreement list can be recovered over every field presented as in Definition E.1, with $q \geq n$, in every characteristic, in $\widetilde{O}_{K_0,a_0}(n\log q)$ bit operations.*

Proof. If $n \geq 4K_0/a_0^2$, then $(k-1)/n \leq a_0^2/4 < a_0^2$. Apply Theorem 9.3 with rate parameter $a_0^2/4$ and agreement a_0, then keep only outputs with at least A agreements. The remaining lengths are bounded in terms of K_0, a_0 and use subset interpolation; constant codes use frequency counting. □

F Application details

This appendix gives two sets of details deferred from Section 7: the round-by-round ProveKit certificates, and the LambdaVM arguments behind the local error bound (54).

F.1 ProveKit's round-by-round certificates

Table 7 records every changed WHIR and zkWHIR code. All rows satisfy $A^2 < n(k-1)$, strictly below Johnson. The artifact checks each support against Proposition C.2 and the error terms (51)–(53) against their targets in exact arithmetic [DKT26a].

Code	n	k	A	$(m, B_\partial, B_{\text{jet}})$	t_0	t_1	u_0	u_1
Passport	262144	65536	123208	(384, 168, 688)	127	109	1	1
	131072	8192	29020	(48, 30, 170)	62	55	1	1
	65536	1024	6906	(40, 34, 310)	41	37	1	1
	32768	128	1782	(12, 12, 168)	31	29	1	1
	16384	16	463	(8, 4, 246)	24	23	1	1
	8192	2	85	(4, 4, 339)	20	18	1	1
Passport internal	4096	1024	1933	(96, 42, 174)	127	109	1	1
	2048	128	453	(68, 34, 241)	62	55	1	1
	1024	16	107	(16, 16, 114)	41	37	1	1
	512	2	19	(12, 12, 227)	31	26	1	1
Goldilocks lookup	4096	1024	1933	(96, 42, 174)	127	109	1	2
	2048	128	453	(38, 23, 133)	62	55	1	1
	1024	16	107	(16, 16, 114)	41	37	1	1
	512	2	19	(12, 12, 227)	31	26	1	1
Blinding	32	8	14	(64, 28, 116)	127	99	1	1

Table 7: Certificates for each changed code: (n, k, A), support $(m, B_\partial, B_{\text{jet}})$, original and revised query counts t_0, t_1, and original and revised out-of-domain counts u_0, u_1. The witness-commitment parameters apply to each of the two witness parts. Each passport witness has one internal zkWHIR proof; the Goldilocks blinding commitment keeps its original 62-query final code. The weakest of the changed passport error terms has 128.0015 bits of security; for the Goldilocks lookup, the first folding and identity term has 136.937 bits and the first query term 128.027 bits. These decimals summarize exact rational inequalities against 2^{-128}.

The passport circuit, as reconstructed from the v1 benchmark, has 294402 constraints and 522648 witness elements; lookup challenges separate the witness into parts of lengths 249345 and 273303, each committed with padded vector length 2^{19}. The lookup circuit has 8001 constraints and two witness parts of length 8001, each padded to 2^{13}. Besides the compressed means of Table 4, the raw passport sizes (uncompressed proof encoding) are 722305.0 and 641748.2 bytes, and the raw lookup sizes (transcript and authentication data) are 207314.7 and 183419.2 bytes, original and revised. The sample standard deviations (raw, compressed) are (1649.6, 2056.2) and (1616.6, 1818.9) bytes for the two passport configurations and (776.6, 967.4) and (730.7, 895.3) bytes for the two lookup configurations.

F.2 LambdaVM: the early evaluations and the local error terms

We prove that the early evaluations of Section 7.4 fix one CPU trace before the LogUp challenges, and then derive the other terms of the local bound (54). Fix the main commitment, and let C be the set of 38-column tuples of polynomials of degree at most $T+2$ with at least A simultaneous agreements with the committed word. Corollary 6.4 at dimension $T+3$ gives $|C| \le \Lambda$. This family is fixed before the early evaluations and LogUp challenges.

Lemma F.1 (Selecting the CPU trace before lookup)**.** *Sample two distinct points uniformly without replacement outside the LDE domain, assuming $q - n \ge 2$. Except with probability at most*

$$\epsilon_{\text{early}} = \binom{\Lambda}{2}\left(\frac{T+2}{q-n-1}\right)^2,$$

the following holds simultaneously for all claimed vectors, including those chosen after the early points s_1, s_2 are revealed:

$$\textit{for every } (v_1, v_2) \in (\mathbb{F}_q^{38})^2, \qquad \#\{f \in C : f(s_1) = v_1,\ f(s_2) = v_2\} \le 1.$$

If the later out-of-domain point is distinct from the early points and DEEP extraction succeeds with the degree-three main-column divisors, its reconstructed main-column tuple is this member.

Proof. Two members differ in a coordinate of degree at most $T + 2$. At each sample its conditional vanishing probability is at most $(T + 2)/(q - n - 1)$. Union-bounding over $\binom{\Lambda}{2}$ pairs gives the collision estimate for all claimed vectors simultaneously.

Later extraction yields a quotient polynomial h of degree less than T. Undoing its divisor reconstructs

$$f(X) = (X - z)(X - s_1)(X - s_2)h(X) + I_f(X).$$

Its degree is at most $T + 2$; it has the common agreements and both early values. The tuple therefore belongs to C and, outside the collision event, is uniquely determined. Without such a member, extraction cannot succeed. □

Reduction modulo $X^T - 1$ preserves the trace-domain values and bus contributions, giving one degree-$< T$ trace fixed before LogUp. The local error bound includes both the collision event and rejection when the out-of-domain point equals an early point.

The MCA and query terms. The first MCA term comes from powers batching, whose curve degree is one less than the number of DEEP terms. The CPU table has 38 main columns, 10 auxiliary columns, one next-row accumulator opening, and two composition parts, for $b = 51$ DEEP terms. Combining them with $1, \gamma, \ldots, \gamma^{b-1}$ therefore uses the curve theorem at degree 50; neither the number of DEEP terms nor the FRI degree bound changes. On a folded domain of size n_j we use agreement $A_j = \lceil An_j/n \rceil$, and the inequalities $2A_{j+1} \ge A_j$ let reconstruction pass through every binary fold, including the last uncommitted one. If E_{50} is the degree-50 curve count and $\mathsf{E}_{\mathrm{exc},j}$ are the line counts of the folds, then

$$\epsilon_{\mathrm{MCA}} = \frac{\mathsf{E}_{50}}{q} + \sum_j \frac{\mathsf{E}_{\mathrm{exc},j}}{q}.$$

The probability bounds rely on how each kind of randomness is sampled. Rejection sampling makes the field challenges uniform. The $t = 208$ query indices are sampled uniformly with replacement from the $n/2$ row pairs, and since at most $A/2$ pairs can have both rows agreeing, each query contributes at most A/n. The early points are sampled without replacement outside the LDE domain (Lemma F.1).

The constraint and out-of-domain terms. The out-of-domain term depends on the degree of the cleared CPU residual, which we bound first. The CPU table has maximum constraint degree 3, no transition end exemptions, and one LogUp boundary constraint. With two early points, every reconstructed trace coordinate has degree at most $T + 2$. A constraint product therefore has degree at most $3(T + 2)$, and clearing its denominator with $X^T - 1$ raises this to at most $4T + 6$. The two recovered composition parts have degree at most T; recombining them as $C_0(X^2) + XC_1(X^2)$ and multiplying by $X^T - 1$ gives the smaller bound $3T + 1$. Thus $D_{\mathrm{res}} = 4T + 6$ bounds the cleared CPU residual.

Fix the CPU trace and auxiliary commitments before sampling the uniform constraint-combination challenge β. An invalid trace candidate violates a constraint at a trace-domain point, so its rational constraint has a pole there. Combining the CPU constraints with powers of β still leaves a pole except at at most 49

challenge values. A union bound over the fixed trace-candidate list gives $49\Lambda/q$; the composition polynomial cannot cancel the pole.

All committed words, including the composition words, are fixed before the out-of-domain challenge. Apply the width-independent list bound to this full tuple at dimension $T + 3$ and common agreement A, giving at most Λ candidates. For each invalid tuple, the cleared out-of-domain residual is a nonzero polynomial of degree at most $4T + 6$. The out-of-domain sampler has at least $q - n - T$ choices after excluding the LDE and trace domains. At most $(4T + 6)\Lambda$ choices are roots of the residuals, at most $2n$ have their square in the LDE domain, and at most two coincide with the early points. Hence its failure probability is at most

$$\frac{(4T + 6)\Lambda + 2n + 2}{q - n - T},$$

which is the out-of-domain term in (54).

G Notes on prior work

The introduction cites two consequences of earlier capacity results that their authors leave implicit.

G.1 Padding the earlier MCA guarantee

Zheng's low-rate manuscript [Zhe26] proves that, for fixed $\theta \in (0, 1)$, sufficiently small $\varepsilon > 0$, and every sufficiently large N, a Reed–Solomon code of length N and dimension $\lfloor (1 - \theta)\varepsilon N \rfloor$ over a prime field with $q \geq N$ has at most $N^{O_{\theta,\varepsilon}(1)}$ exceptional challenges at agreement $\lceil \varepsilon N \rceil$. We derive its all-rate consequence by padding, as BCPZZ do for list decoding, crediting Alrabiah, Goyal, and Guruswami [BCP+26a, Section 5].

Fix $\rho, \delta > 0$ with $\rho + \delta \leq 1$ and let $k = \lfloor \rho n \rfloor$. Choose fixed $\theta \in (0, 1)$ with $\rho/(1 - \theta) < \rho + \delta$, then choose ε small enough for the low-rate theorem and for $c = (1 - \theta)\varepsilon < \rho$. Set $N = \lceil k/c \rceil$. For sufficiently large n,

$$N \geq n, \qquad \lfloor cN \rfloor = k, \qquad \lceil \varepsilon N \rceil \leq \lceil (\rho + \delta)n \rceil.$$

If $q \geq N$, append distinct evaluation points and extend both received words by zeros. An agreement set without common witnesses on the original coordinates still has none in the padded code: its polynomial degree bound is still k, and its agreement meets the padded threshold. Restricting any padded witness equality to the original coordinates would also witness the original candidate. Thus the original code has at most $N^{O_{\theta,\varepsilon}(1)} = n^{O_{\rho,\delta}(1)}$ exceptional challenges. When $n \leq q < N = O_{\rho,\delta}(n)$, a probability bound of the form $n^{O_{\rho,\delta}(1)}/q$ exceeds one, and so holds trivially, once its constant is large enough. Theorem 8.1 gives explicit parameters depending only on the gap.

G.2 The derivative order in BCPZZ's parameters

BCPZZ do not state their derivative order in terms of δ, so we derive it from their parameters. At fixed positive rate ρ and sufficiently small gap δ, their all-rate reduction sets $\theta = \delta/(2(\rho + \delta)) = \Theta_\rho(\delta)$. Choose its low-rate agreement parameter η to be half the upper bound

$$\left(\frac{\theta^3(1 - \theta)}{768}\right)^{(5+\theta)/(1-\theta)}$$

from their Theorem 4.1. For small enough δ, this also satisfies $\eta < \rho + \delta$, as their reduction requires. Then $\log(1/\eta) = O_\rho(\log(1/\delta))$, and their Algorithm 1 chooses $d = \lceil \eta^{-3/\theta} \rceil$ for this low-rate instance [BCP+26a, Algorithm 1, Theorem 4.1, and Section 5]. Hence

$$\log d = O_\rho\left(\delta^{-1} \log(1/\delta)\right).$$

```
theorem firstOrderBranch_finiteLength_finiteSlack_bounds
    (rho eta : ℝ) (n k A : ℕ)
    (hrho : 0 < rho) (hrhoOne : rho < 1) (heta : 0 < eta)
    (haOne : firstOrderBranchThreshold rho + eta < 1)
    (hk : 0 < k) (hkRate : (k : ℝ) ≤ rho * n)
    (hA : (firstOrderBranchThreshold rho + eta) * n ≤ A) (hAn : A ≤ n)
    {F : Type*} [Field F] (domain : Fin n ↪ F)
    (hchar : k = 1 ∨ ringChar F = 0 ∨
      max (k - 1) (firstOrderBranchFiniteLengthDerivativeCap rho eta n) < ringChar F) :
    (∀ received : Fin n → F,
      (closePolynomialSet domain received k A).Finite ∧
        ((closePolynomialSet domain received k A).ncard : ℝ) ≤
          7 * firstOrderBranchFiniteLengthMcaConstant rho ^ 3 * n /
            finiteLengthSlack eta n ^ 2) ∧
      ∀ f g : Fin n → F,
        ∃ exceptional : Finset F,
          (exceptional.card : ℝ) ≤
            140 * firstOrderBranchFiniteLengthMcaConstant rho ^ 6 * n ^ 2 /
              finiteLengthSlack eta n ^ 4 ∧
          ∀ z ∉ exceptional,
            ∀ P : F[X], P.degree < k →
            A ≤ (polynomialAgreementSet domain (fun i ↦ f i + z * g i) P).card →
            HasExactCorrelatedPair domain f g (RingHom.id F) k z P
```

Figure 7: The first-order list-decoding and MCA theorem in Lean.

This order and ours in Theorems 5.7 and 8.1 are sufficient, not shown to be necessary.

H The main statements in Lean

To find the Lean statement of a paper theorem, start with `ReedSolomon/PaperGuide.lean` at the validated mathematics revision [DA26]. It groups the results by their mathematical role and links to declarations whose comments explain the parameters and hypotheses. We read one joint list-decoding and MCA statement below; the artifact checks this excerpt against the source and publishes the other literal statements with a source manifest [DKT26a]. The artifact's implementation-source manifest records the repository, revision, and source paths behind the application parameters. Table 7 lists the ProveKit supports; the ZisK and LambdaVM checkers record each batching and fold certificate with its allocated error.

The RS-specific development spans approximately 75,000 source lines across 302 Lean modules, including concrete examples.

H.1 Reading the joint list and MCA statement

Figure 7 records the joint first-order theorem in namespace `ReedSolomon.FirstOrder`, the finite-length form of Theorems 5.1 and 5.2. Here `rho` bounds the rate, `eta` is η_1, and `firstOrderBranchThreshold` is $a_1(\rho)$ (Definition 3.15), including its low-rate branch. The hypotheses require $\eta_1 > 0$, $k \le \rho n$, $A \ge (a_1(\rho) + \eta_1)n$, and distinct evaluation points; `hchar` is the restriction $p > \max\{k-1, B_\partial\}$ or characteristic zero, waived for $k = 1$.

The first conclusion proves that `closePolynomialSet`, the set of *all* degree-$< k$ polynomials with at least A agreements, is finite and bounds its size. In the second, the exceptional set depends on f, g but precedes every challenge z and candidate P. The conclusion `HasExactCorrelatedPair` gives degree-$< k$

witnesses with $P = F_0 + zG_0$ whose common agreement set equals the candidate's full agreement set. Writing `finiteLengthSlack` as $\bar{\eta}_1 = \eta_1 + 1/n$, the list and the exceptional set $\mathcal{E}$ satisfy

$$|\mathcal{L}(y, A)| \le C_\rho \frac{n}{\bar{\eta}_1^2}, \qquad |\mathcal{E}| \le D_\rho \frac{n^2}{\bar{\eta}_1^4},$$

where C_ρ, D_ρ depend only on the rate bound, not on the field, code length, or received words.

Other quantitative results. The links below lead to the source statements; the source guide records their full declaration names.

Regime and hypotheses	Formalized guarantee
Zeroth order (Theorem 4.4) Every characteristic; positive Johnson gap	Complete finite list and Johnson bound; line MCA with full agreement sets
Fixed order (Corollary 8.8) $\Gamma(\rho, a, d) > 1$; the stated characteristic assumption	Lists from the derivative-order condition; line MCA from the same parameters
First-order curves Independent ordinary/regular tuning from Lemma 5.5 and Section C.1	Optimized finite curve certificate; squarefree recovery with independently minimized ordinary threshold
Gap from capacity (Theorem 8.1) Fixed $\delta > 0$; uniform over the rate	Uniform capacity lists (first order from gap $6/25$); uniform line, affine, and powers MCA
Interleaving and folds Width-independent transfer and shared-level union bound from Corollary 6.6 and Section 6.3	Arbitrary-width interleaved powers agreement; Reed–Solomon shared-level fold bound

The optimized first-order declarations prove that the minimum over the ordinary threshold $k \le L_0 \le A$ in (44) is attained, retains full agreement-set recovery, and is no larger than the choice $L_0 = k$. The bounded-input list-recovery extension of Remark 8.10 is proved in the paper only.

Finite-field MCA corollaries give errors $\mathsf{E}_{\mathrm{exc}}/q$ for lines and $\mathsf{E}_{\mathrm{exc}}/(q-1)$ for positive-dimensional affine families. The application guide covers the application certificates and byte identities.

AI methodology statement

GPT-5.6 Sol, GPT-6 Astra, and Claude Fable 5.1 discovered and refined the main results, including the list-decoding and MCA bounds and the concrete specializations, with our close prompting and guidance. Sol and Astra carried out the Lean formalization and wrote most of the manuscript under our direction. Claude Opus 5.5 assisted with the editorial revision of the manuscript. We take full responsibility for the results, proofs, and exposition. Section 1.2 discusses prior and contemporaneous work, including what those works disclose about their use of AI. The authors verified the correctness and originality of all content, including references.

References

[ABF26] Gal Arnon, Dan Boneh, and Giacomo Fenzi. Open problems in list decoding and correlated agreement. Cryptology ePrint Archive, Paper 2026/680, 2026. July 2026 revision.

[ACFY24a] Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, and Eylon Yogev. STIR: Reed–Solomon proximity testing with fewer queries. Cryptology ePrint Archive, Paper 2024/390, 2024.

[ACFY24b] Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, and Eylon Yogev. WHIR: Reed–Solomon proximity testing with super-fast verification. Cryptology ePrint Archive, Paper 2024/1586, 2024.

[AGG+25] Omar Alrabiah, Zeyu Guo, Venkatesan Guruswami, Ray Li, and Zihan Zhang. Random Reed–Solomon codes achieve list-decoding capacity with linear-sized alphabets. *Advances in Combinatorics*, 2025(8):1–39, 2025.

[AGL24] Omar Alrabiah, Venkatesan Guruswami, and Ray Li. Randomly punctured Reed–Solomon codes achieve list-decoding capacity over linear-sized fields. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1458–1469, 2024.

[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*, pages 880–898, 2026.

[Ale05] Michael Alekhnovich. Linear diophantine equations over polynomials and soft decoding of Reed–Solomon codes. *IEEE Transactions on Information Theory*, 51(7):2257–2265, 2005.

[AP00] Daniel Augot and Lancelot Pecquet. A hensel lifting to replace factorization in list-decoding of algebraic-geometric and Reed–Solomon codes. *IEEE Transactions on Information Theory*, 46(7):2605–2614, 2000.

[Ark26a] ArkLib contributors. All-rate line and affine mutual correlated agreement in Lean. Source, 2026. Public repository checkpoint, September 5, 2026.

[Ark26b] ArkLib contributors. Qualitative all-rate Reed–Solomon list bound in Lean. Source, 2026. Public repository checkpoint, September 4, 2026.

[Bat68] Kenneth E. Batcher. Sorting networks and their applications. In *Proceedings of the Spring Joint Computer Conference*, volume 32, pages 307–314, 1968.

[BBHR18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast Reed–Solomon interactive oracle proofs of proximity. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 14:1–14:17, 2018.

[BCFW25] Benedikt Bünz, Alessandro Chiesa, Giacomo Fenzi, and William Wang. Linear-time accumulation schemes. Cryptology ePrint Archive, Paper 2025/753, 2025. TCC 2025.

[BCGM25] Sarah Bordage, Alessandro Chiesa, Ziyi Guan, and Ignacio Manzur. All polynomial generators preserve distance with mutual correlated agreement. Cryptology ePrint Archive, Paper 2025/2051, 2025.

[BCH+26] Eli Ben-Sasson, Dan Carmon, Ulrich Haböck, Swastik Kopparty, and Shubhangi Saraf. On proximity gaps of Reed–Solomon codes. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC 2026)*, pages 1157–1167. ACM, 2026. https://doi.org/10.1145/3798129.3800827. Full version: https://eprint.iacr.org/2025/2055.

[BCI+23] Eli Ben-Sasson, Dan Carmon, Yuval Ishai, Swastik Kopparty, and Shubhangi Saraf. Proximity gaps for Reed–Solomon codes. *Journal of the ACM*, 70(5), 2023. Author-hosted preprint, July 2020.

[BCP+26a] Joshua Brakensiek, Yeyuan Chen, Aaron Putterman, Zihan Zhang, and Kai Zhe Zheng. Algorithmic list decoding of Reed–Solomon codes up to capacity. Technical Report TR26-164, revision 1, Electronic Colloquium on Computational Complexity, 2026. September 5, 2026. https://eccc.weizmann.ac.il/report/2026/164/.

[BCP+26b] Joshua Brakensiek, Yeyuan Chen, Aaron Putterman, Zihan Zhang, and Kai Zhe Zheng. Algorithmic list decoding of Reed–Solomon codes up to capacity in the low-rate regime, 2026. Original manuscript, September 4, 2026.

[BGK+23] Alexander R. Block, Albert Garreta, Jonathan Katz, Justin Thaler, Pratyush Ranjan Tiwari, and Michał Zając. Fiat–shamir security of FRI and related SNARKs. Cryptology ePrint Archive, Paper 2023/1071, 2023.

[BGKS19] Eli Ben-Sasson, Lior Goldberg, Swastik Kopparty, and Shubhangi Saraf. DEEP-FRI: Sampling outside the box improves soundness. Cryptology ePrint Archive, Paper 2019/336, 2019.

[BGM23] Joshua Brakensiek, Sivakanth Gopi, and Visu Makam. Generic Reed–Solomon codes achieve list-decoding capacity. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1488–1501, 2023.

[Bit26a] BitWonka. Proximity Prize: short-horizon johnson certificate, submission 64, 2026. Promoted August 21, 2026.

[Bit26b] BitWonka. Proximity Prize: submission 241, 2026.

[Bit26c] BitWonka. Proximity Prize: submission 255, 2026.

[Bit26d] BitWonka and Proximity Prize contributors. Proximity Prize: moving-coordinate degree bounds, submission 567, 2026. Promoted September 15, 2026; revision cdb451f13fdc6c84f5fe363e77ee13a89bd30974.

[BKR10] Eli Ben-Sasson, Swastik Kopparty, and Jaikumar Radhakrishnan. Subspace polynomials and limits to list decoding of Reed–Solomon codes. *IEEE Transactions on Information Theory*, 56(1):113–120, 2010.

[BMMS25] Anubhav Baweja, Pratyush Mishra, Tushar Mopuri, and Matan Shtepel. Query-optimal IOPPs for linear-time encodable codes. Cryptology ePrint Archive, Paper 2025/1588, 2025.

[CHK26] Soham Chatterjee, Prahladh Harsha, and Mrinal Kumar. Deterministic list decoding of Reed–Solomon codes. In *Proceedings of the 58th ACM Symposium on Theory of Computing (STOC)*, pages 2040–2051, 2026. Full version: ECCC TR25-170, revision 1, March 25, 2026.

[Cho26] Przemek Chojecki. Shortening bounds for Reed–Solomon MCA. Cryptology ePrint Archive, Paper 2026/1463, 2026.

[CS25] Elizabeth Crites and Alistair Stewart. On Reed–Solomon proximity gaps conjectures. Cryptology ePrint Archive, Paper 2025/2046, 2025.

[DA26] Quang Dao and ArkLib contributors. ArkLib: Reed–Solomon capacity and correlated agreement in Lean. https://github.com/quangvdao/ArkLib, 2026. Mathematical formalization and concrete application certificates. Revision 6aa8b6fe, November 2, 2026.

[Dao26a] Quang Dao. LambdaVM CPU Reed–Solomon anchor overhead: measurements and reproducibility, 2026. Apple M4 Max, 16 threads; September 14, 2026.

[Dao26b] Quang Dao. LambdaVM CPU RS anchors: complete-proof size experiment, 2026. Implementation 7310d1e8; measurement record c40323c8; September 14, 2026.

[DG25] Benjamin E. Diamond and Angus Gruen. On the distribution of the distances of random words. Cryptology ePrint Archive, Paper 2025/2010, 2025.

[DKSS13] Zeev Dvir, Swastik Kopparty, Shubhangi Saraf, and Madhu Sudan. Extensions to the method of multiplicities, with applications to Kakeya sets and mergers. *SIAM Journal on Computing*, 42(6):2305–2328, 2013.

[DKT26a] Quang Dao, Scott Duke Kominers, and Justin Thaler. Code and concrete certificates for Reed–Solomon list decoding and mutual correlated agreement. https://github.com/quangvdao/rs-beyond-johnson/tree/fd6d84edd50dcca7c0b3f388e0b4824d9daa082d, 2026. Exact-arithmetic parameter and application checkers, reference certificates, and links to pinned Lean sources.

[DKT26b] Quang Dao, Scott Duke Kominers, and Justin Thaler. Reed–Solomon codes beyond Johnson: Efficient decoding and smaller cryptographic proofs. Cryptology ePrint Archive, Paper 2026/2056, 2026.

[DKTZ26] Quang Dao, Scott Duke Kominers, Justin Thaler, and Kai Zhe Zheng. Counterexamples to beyond-Johnson proximity gaps over binary fields. Manuscript, October 2026.

[DMSX05] Xavier Dahan, Marc Moreno Maza, Éric Schost, and Yuzhen Xie. On the complexity of the D5 principle, 2005. Author manuscript, November 15, 2005.

[Eli57] Peter Elias. List decoding for noisy channels. Technical Report 335, Research Laboratory of Electronics, Massachusetts Institute of Technology, 1957.

[erd26a] erdkocak. Proximity Prize: submission 219, 2026.

[erd26b] erdkocak. Proximity Prize: submission 38. https://github.com/proximity-prize/proximity-prize/pull/38, 2026. Merged August 20, 2026.

[FS25] Giacomo Fenzi and Antonio Sanso. Small-field hash-based SNARGs are less sound than conjectured. Cryptology ePrint Archive, Paper 2025/2197, 2025. Revised April 2026.

[Ful98] William Fulton. *Intersection Theory*. Springer, second edition, 1998.

[Gab26] Ariel Gabizon. Linear list size bounds for Reed–Solomon beyond the Johnson radius. Cryptology ePrint Archive, Paper 2026/2048, 2026.

[GG25] Rohan Goyal and Venkatesan Guruswami. Optimal proximity gaps for subspace-design codes and (random) Reed–Solomon codes. Cryptology ePrint Archive, Paper 2025/2054, 2025. Revised March 2026.

[GGSW26] Rohan Goyal, Venkatesan Guruswami, Yihang Sun, and Mary Wootters. Locality of curve-decoding and improved proximity gaps. arXiv:2607.08516, 2026.

[GHKS24] Rohan Goyal, Prahladh Harsha, Mrinal Kumar, and Ashutosh Shankar. Fast list decoding of univariate multiplicity and folded Reed–Solomon codes. In *Proceedings of the 65th IEEE Symposium on Foundations of Computer Science*, pages 328–343, 2024.

[GHKS26] Rohan Goyal, Prahladh Harsha, Mrinal Kumar, and Ashutosh Shankar. Fast list recovery of univariate multiplicity codes. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2026)*, volume 392 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 67:1–67:18. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2026.

[GLS+24] Zeyu Guo, Ray Li, Chong Shangguan, Itzhak Tamo, and Mary Wootters. Improved list-decodability and list-recoverability of Reed–Solomon codes via tree packings. *SIAM Journal on Computing*, 53(2):389–430, 2024.

[GMW25] Albert Garreta, Nicolas Mohnblatt, and Benedikt Wagner. A simplified round-by-round soundness proof of FRI. Cryptology ePrint Archive, Paper 2025/1993, 2025.

[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.

[GRS25] Venkatesan Guruswami, Atri Rudra, and Madhu Sudan. Essential coding theory, 2025. Draft dated August 26, 2025.

[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.

[GST23] Eitan Goldberg, Chong Shangguan, and Itzhak Tamo. List-decoding and list-recovery of Reed–Solomon codes beyond the Johnson radius for every rate. *IEEE Transactions on Information Theory*, 69(4):2261–2268, 2023.

[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.

[GZ23] Zeyu Guo and Zihan Zhang. Randomly punctured Reed–Solomon codes achieve the list decoding capacity over polynomial-size alphabets. In *Proceedings of the 64th IEEE Symposium on Foundations of Computer Science*, pages 164–176, 2023.

[Hab25] Ulrich Haböck. A note on mutual correlated agreement for Reed–Solomon codes. Cryptology ePrint Archive, Paper 2025/2110, 2025.

[Hei83] Joos Heintz. Definability and fast quantifier elimination in algebraically closed fields. *Theoretical Computer Science*, 24(3):239–277, 1983.

[i3426] i34-9 and partylikeits1983 and 0xLucqs. Proximity Prize: Taylor-tail refinements, submission 551, 2026. Includes the preceding contributions linked in the submission.

[jac26a] jacklightChen. Proximity Prize: submission 343, 2026.

[jac26b] jacklightChen and Proximity Prize contributors. Proximity Prize: aggregate factor bounds, submission 563, 2026.

[jac26c] jacklightChen and Proximity Prize contributors. Proximity Prize: per-factor interpolation profiles, submission 565, 2026.

[jac26d] jacklightChen and Proximity Prize contributors. Proximity Prize: submission 561, 2026.

[Jer26] Fernando Granha Jeronimo. Algorithmic list decoding at capacity and optimal proximity gaps for Reed–Solomon codes. Technical Report TR26-169, Electronic Colloquium on Computational Complexity, 2026. Submitted September 5, 2026; published September 6, 2026.

[jie26a] jieyilong. Proximity Prize: universal taylor numerators and factor sums, submission 72, 2026. Promoted August 21, 2026.

[jie26b] jieyilong and Proximity Prize contributors. Proximity Prize: submission 303, 2026.

[JNSV17] Claude-Pierre Jeannerod, Vincent Neiger, Éric Schost, and Gilles Villard. Computing minimal interpolation bases. *Journal of Symbolic Computation*, 83:272–314, 2017.

[Jo26a] Sunghyeon Jo. Interleaving stability for mutual correlated agreement and curve decodability. Cryptology ePrint Archive, Paper 2026/891, 2026.

[Jo26b] Sunghyeon Jo. Reed–Solomon mutual correlated agreement beyond the Johnson radius. Cryptology ePrint Archive, Paper 2026/1432, 2026.

[jsi26] jsign. Proximity Prize: submission 506, 2026.

[Kam26] Antonio Kambiré. Proximity gaps conjecture fails near capacity over prime fields. arXiv:2604.09724, 2026.

[KKH26] Dmitry Krachun, Stepan Kazanin, and Ulrich Haböck. Failure of proximity gaps close to capacity. Cryptology ePrint Archive, Paper 2026/782, 2026.

[Kop15] Swastik Kopparty. List-decoding multiplicity codes. *Theory of Computing*, 11(5):149–182, 2015.

[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.

[KSY14] Swastik Kopparty, Shubhangi Saraf, and Sergey Yekhanin. High-rate codes with sublinear-time decoding. *Journal of the ACM*, 61(5):28:1–28:20, 2014.

[KTS22] Itay Kalev and Amnon Ta-Shma. Unbalanced expanders from multiplicity codes. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022)*, volume 245 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 12:1–12:14. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2022.

[Lam26a] LambdaVM contributors. Lambda VM: project overview and development status. Source, 2026. Accessed September 5, 2026.

[Lam26b] LambdaVM contributors. LambdaVM: Fiat–Shamir transcript and recursive verification. Source, 2026. Revision 8064a8efee4bd3edc9f064337d4e1d8bad54ae1a, inspected September 7, 2026.

[Lam26c] LambdaVM contributors. LambdaVM STARK proof options. Source, 2026. Revision 8064a8ef, inspected September 5, 2026.

[Lam26d] LambdaVM contributors. LambdaVM trace tables, proof construction, and command-line parameters. Source, 2026. Revision 8064a8ef, inspected September 7, 2026.

[Lec08] Grégoire Lecerf. Fast separable factorization and applications. *Applicable Algebra in Engineering, Communication and Computing*, 19:135–160, 2008.

[Len91] Hendrik W. Lenstra, Jr. Finding isomorphisms between finite fields. *Mathematics of Computation*, 56(193):329–347, 1991.

[Mil09] James S. Milne. Algebraic geometry, 2009. Version 5.20.

[Mis26] Pratyush Mishra. Generalize thm to any constant rate, 2026. Pull request 1 to kz99/rs-ld-mca, September 4, 2026; revision b1e346fc39780adb442ed2504a316b32702b97af.

[Nas26] Bartosz Naskręcki. Proximity Prize: submission 122. https://github.com/proximity-prize/proximity-prize/pull/122, 2026. Merged August 27, 2026.

[net26] nethoxa and jsign and jacklightChen. Proximity Prize: submission 517, 2026.

[NRS17] Vincent Neiger, Johan Rosenkilde, and Éric Schost. Fast computation of the roots of polynomials over the ring of power series. In *Proceedings of the 2017 ACM International Symposium on Symbolic and Algebraic Computation*, pages 349–356. ACM, 2017.

[Ope26] OpenAI. Deterministic polynomial factorization over prime fields. OpenAI Math Release preprint, October 4, 2026. Produced by an internal OpenAI model; not formally verified, 2026.

[Plo26] Plonky2 contributors. Plonky2 standard recursion configuration, 2026. Revision 5d9da5a6, inspected September 5, 2026.

[Pro26a] Proofman contributors. Goldilocks transcript query sampling and grinding verification, 2026. Revision 0f3fef8cd1897df469532996e72e0c84ef69d6fb, transcriptGL.cpp, lines 93–124, and stark_verify.hpp, lines 195–218; inspected September 14, 2026.

[Pro26b] Proofman contributors. STARK parameter generation, security calculator, and proof serialization. Source, 2026. Revision 0f3fef8c, inspected September 5, 2026.

[Pro26c] Proximity Prize contributors. Fixed interleaved Reed–Solomon benchmark and soundness certificates. Proximity Prize public Lean development, 2026. Profile at revision b34c0131; promoted lower certificate at revision 032154395, inspected September 5, 2026.

[Pro26d] Proximity Prize contributors. Proximity Prize: lower-bound benchmark specification, 2026.

[PV05] Farzad Parvaresh and Alexander Vardy. Correcting errors beyond the Guruswami–Sudan radius in polynomial time. In *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science*, pages 285–294, 2005.

[rec26] recmo. Proximity Prize: submission 245, 2026.

[RIS26] RISC Zero contributors. RISC Zero soundness models, version 3.0.6, 2026. Source snapshot accessed September 9, 2026.

[Roj99] J. Maurice Rojas. Solving degenerate sparse polynomial systems faster. *Journal of Symbolic Computation*, 28(1–2):155–186, 1999.

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

[RW14] Atri Rudra and Mary Wootters. Every list-decodable code for high noise has abundant near-optimal rate puncturings. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 764–773, 2014.

[sau26] saucegodbased. Proximity Prize: submission 492, 2026.

[Sch15] Éric Schost. Algorithms for finite field arithmetic. In *Proceedings of the 2015 ACM International Symposium on Symbolic and Algebraic Computation*, pages 7–15. ACM, 2015.

[Set20] Srinath Setty. Spartan: Efficient and general-purpose zkSNARKs without trusted setup. In *Advances in Cryptology – CRYPTO 2020*, volume 12172 of *Lecture Notes in Computer Science*, pages 704–737, 2020.

[ST23] Chong Shangguan and Itzhak Tamo. Generalized Singleton bound and list-decoding Reed–Solomon codes beyond the Johnson radius. *SIAM Journal on Computing*, 52(3):684–717, 2023.

[Sta26a] StarkWare. Stone prover and Cairo end-to-end parameters, 2026. Revision 1414a545, inspected September 5, 2026.

[Sta26b] StarkWare. Stwo Circle STARK prover, 2026. Revision 826591c6, inspected September 5, 2026.

[Sud97] Madhu Sudan. Decoding of Reed–Solomon codes beyond the error-correction bound. *Journal of Complexity*, 13(1):180–193, 1997.

[The26] The Stacks Project Authors. The stacks project, tag 054k: Chevalley's theorem, 2026.

[vdH02] Joris van der Hoeven. Relax, but don't be too lazy. *Journal of Symbolic Computation*, 34(6):479–542, 2002.

[vzGG13] Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. Cambridge University Press, third edition, 2013.

[Win25] Winterfell contributors. Winterfell proven and conjectured security estimates, version 0.13.1, 2025. Source snapshot accessed September 9, 2026.

[Wor26a] World Foundation. ProveKit backends and shared WHIR profile. Source, 2026. Revision 6481f961, inspected September 5, 2026.

[Wor26b] World Foundation. ProveKit Goldilocks lookup fixtures. Source, 2026. Source revision 6481f961, inspected September 7, 2026.

[Wor26c] World Foundation. ProveKit passport benchmark. Source, 2026. Source revision 420e99bc, inspected September 7, 2026.

[Wor26d] World Foundation. ProveKit: Privacy for the real world. Source, 2026. September 2, 2026; accessed September 5, 2026.

[Wor26e] World Foundation. ProveKit WHIR: version 0.1.1. Source, 2026.

[Wor26f] World Foundation. WHIR parameter generation and code-switch error bounds. Source, 2026. Revision 8804e80e, inspected September 5, 2026.

[Woz58] John M. Wozencraft. List decoding. *Quarterly Progress Report*, 48:90–95, 1958. Research Laboratory of Electronics, Massachusetts Institute of Technology.

[YZ26] Chen Yuan and Ruiqi Zhu. A syndrome-space approach to proximity gaps and correlated agreement for random linear codes and random Reed–Solomon codes. arXiv:2605.07595, 2026. Version 2, July 10, 2026.

[Zhe26] Zheng's rs-ld-mca project. Mutual correlated agreement for Reed–Solomon codes up to capacity in the low-rate regime. Source, 2026. Unpublished AI-generated manuscript distributed in Kai Zhe Zheng's repository, September 4, 2026.

[Zis26a] ZisK contributors. ZisK: project overview and release status. Source, 2026. Accessed September 5, 2026.

[Zis26b] ZisK contributors. ZisK: Proof output modes and byte encoding. Source, 2026. Revision b08d856c0f72b21d94fc49151deefaf8e7419b79, inspected September 7, 2026.

[Zis26c] ZisK contributors. ZisK v1.2.0-alpha proving parameters. Source, 2026. Public proving-key release, inspected September 7, 2026.