

Compare Before Clearing: Exact Integral-Comparison Frontiers for Lattice Extraction

Xiang Wang¹ and Shihui Fu²

¹ Guizhou University

ee.wangx24@gzu.edu.cn

² Shandong University

shihuifu@sdu.edu.cn

Abstract. Lattice extraction often produces openings normalized by challenge differences, whereas an inconsistency must ultimately yield a short integral SIS relation. Clearing each extracted branch before comparison removes every denominator obstruction carried by that branch, including factors irrelevant to the mismatch that is eventually tested.

We formalize direct integral comparison for generic polynomial block systems. If block a has width r_a and the two extraction centers differ on J , the minimum worst-case coefficient degree is $\max\{\max_a r_a, \sum_{a \in J} r_a\}$. Within a branch-separated polynomial integralize-then-compare architecture, it is $2 \sum_a r_a$. The coordinate case gives $\max\{1, h\}$ and $2L$, where $h = |J|$.

Exact conditional resampling obtains the required partially synchronized successful executions without a reciprocal-success loss. The coordinate schedule uses at most $2L+1$ additional retry invocations in unconditional expectation.

Two cases illustrate the bounds. For Cyclo-style coordinate folding, one unsynchronized coordinate has the same certified radius as same-root synchronization. For two independently extracted Eschin-style Vandermonde stars, direct comparison has degree $\binom{k+1}{2}$ in the anchor-universal polynomial-linear model. The degree is k^2 within the stated branch-separated integralize-then-compare architecture.

Keywords: Lattice extraction · Knowledge extraction · Module-SIS · Vandermonde systems

1 Introduction

Folding compresses several proof obligations into a new running instance that can itself be folded again. This idea underlies modern constructions of incrementally verifiable computation and recursive proofs [1,2]. Soundness then asks an extractor to recover consistent source information or expose a binding violation.

Over a field, fork extraction can divide by a nonzero challenge difference. A lattice reduction ultimately needs something more rigid: a short integral SIS

relation. Clearing the normalization too early can therefore enlarge the certificate even when the underlying identities are equivalent over the fraction field.

The effect is already visible with two coordinates. In branch $b \in \{0, 1\}$, suppose

$$U_1^{(b)} = X_1^{(b)} / \Delta_1^{(b)}, \quad U_2^{(b)} = X_2^{(b)} / \Delta_2^{(b)}.$$

Clearing a whole branch first introduces

$$P_b = \Delta_1^{(b)} \Delta_2^{(b)}.$$

If the branches disagree only in U_1 , branchwise comparison still carries $P_0 P_1$. Direct comparison uses

$$X_1^{(0)} \Delta_1^{(1)} - X_1^{(1)} \Delta_1^{(0)},$$

which never involves Δ_2 . The branchwise route has cleared denominator structure that the actual mismatch never uses. How much clearing is actually necessary?

Let J be the coordinates on which two extraction centers differ and let $h = |J|$. In the coordinate-linear polynomial comparer model, the answer is

$$d_{\text{coord}}^*(h) = \max\{1, h\}. \quad (1)$$

The same-root case is $h = 0$. One residual center disagreement still costs degree one. Beyond that, comparison pays for the h coordinates that remain unsynchronized.

For blocks of widths r_a , the corresponding result is

$$d_{\text{direct}}^*(J) = \max \left\{ \max_a r_a, \sum_{a \in J} r_a \right\} \quad (2)$$

whereas, within the branch-separated polynomial integralize-then-compare architecture,

$$d_{\text{branch}}^* = 2 \sum_a r_a. \quad (3)$$

These equalities refer to the comparer models defined in Section 2.3.

The two terms in the direct bound come from different mismatch cases. A nonconstant normalized-coordinate mismatch already has a local cross-product certificate. If those coordinates agree but the two centers still differ, the baseline mismatch must be bridged across the changed coordinates. In the block theorem, generic matrices replace scalar challenge differences and maximal minors play the same algebraic role.

The algebra only requires the two successful extraction objects to agree on a chosen part of their challenge state. Section 4 obtains such pairs by resampling from the corresponding conditional challenge law and accounting for retries before conditioning on success. The coordinate schedule has additive raw loss and at most $2L + 1$ expected additional retry invocations.

We study two rather different examples. Cyclo gives a coordinate application in which the amount of retained synchronization can be varied directly [3]. Vandermonde extraction gives a dense structured system in which the natural grading is instead the number of challenge-difference factors [4].

Contributions. Our contributions are the model-specific comparison degrees in Eqs. (1)–(3), a conditional-resampling argument that supplies the partially synchronized executions they require, and the Cyclo and Vandermonde instantiations.

Organization. Sections 2–3 give the comparison algebra and degree bounds. Section 4 gives the extraction argument. Sections 5 and 6 treat the two examples; full proofs and arithmetic certificates appear in the appendices.

2 Integral Comparison

An extracted star represents source openings through short numerators and invertible challenge differences. Two stars need not have the same numerical center. The comparison problem is therefore to turn a disagreement between their normalized openings into an integral kernel relation while clearing only the denominator structure needed by that disagreement.

2.1 Coordinate setup

Let R be a commutative unital ring, let M, N be right R -modules, and let $A : M \rightarrow N$ be right R -linear. We use a norm on M and multiplication-operator bounds on R satisfying

$$\|xr\| \leq \|x\|\|r\|_{\text{op}}, \quad \|rr'\|_{\text{op}} \leq \|r\|_{\text{op}}\|r'\|_{\text{op}}. \quad (4)$$

Fix public images $t_0, t_1, \dots, t_L \in N$. For each branch $b \in \{0, 1\}$, a baseline at center $s^{(b)} = (s_1^{(b)}, \dots, s_L^{(b)}) \in R^L$ satisfies

$$A(V^{(b)}) = t_0 + \sum_{i=1}^L t_i s_i^{(b)}, \quad \|V^{(b)}\| \leq B_V. \quad (5)$$

An i th coordinate fork in branch b replaces $s_i^{(b)}$ by $\tilde{s}_i^{(b)}$. If $V^{(i,b)}$ is its extracted output, write

$$X_i^{(b)} = V^{(i,b)} - V^{(b)}, \quad \Delta_i^{(b)} = \tilde{s}_i^{(b)} - s_i^{(b)}. \quad (6)$$

The baseline and fork equations give

$$A(X_i^{(b)}) = t_i \Delta_i^{(b)}. \quad (7)$$

We assume that every $\Delta_i^{(b)}$ used below is a unit and that

$$\|X_i^{(b)}\| \leq B_X, \quad \|\Delta_i^{(b)}\|_{\text{op}} \leq B_\Delta. \quad (8)$$

The normalized coordinates are

$$U_i^{(b)} = X_i^{(b)} (\Delta_i^{(b)})^{-1}, \quad U_0^{(b)} = V^{(b)} - \sum_{i=1}^L U_i^{(b)} s_i^{(b)}. \quad (9)$$

Thus $A(U_i^{(b)}) = t_i$ for every $i \in \{0, \dots, L\}$. Inverses identify the normalized tuple; none of the bounds below uses an operator norm for an inverse.

The centers differ on the residual active set

$$J = \{i \in [L] : s_i^{(0)} \neq s_i^{(1)}\}, \quad h = |J|, \quad \varepsilon_i = s_i^{(0)} - s_i^{(1)}. \quad (10)$$

For the homogeneous estimate assume $\|\varepsilon_i\|_{\text{op}} \leq B_\varepsilon$ for $i \in J$. Exact same-center comparison is the case $J = \emptyset$, or $h = 0$.

2.2 Local and baseline mismatches

A disagreement in one nonconstant coordinate is already local. It can be compared without using either center or any denominator from another coordinate.

Lemma 1 (Local coordinate comparison). *For $i \in [L]$, define*

$$K_i = X_i^{(0)} \Delta_i^{(1)} - X_i^{(1)} \Delta_i^{(0)}. \quad (11)$$

Then $A(K_i) = 0$ and

$$K_i = (U_i^{(0)} - U_i^{(1)}) \Delta_i^{(0)} \Delta_i^{(1)}. \quad (12)$$

Consequently $K_i \neq 0$ exactly when $U_i^{(0)} \neq U_i^{(1)}$, and

$$\|K_i\| \leq 2B_X B_\Delta. \quad (13)$$

Proof. Right linearity and commutativity give

$$A(K_i) = t_i \Delta_i^{(0)} \Delta_i^{(1)} - t_i \Delta_i^{(1)} \Delta_i^{(0)} = 0.$$

Substituting $X_i^{(b)} = U_i^{(b)} \Delta_i^{(b)}$ proves Eq. (12). A product of units acts bijectively on M , which proves the nonzeroness equivalence. Finally, Eq. (4) and the triangle inequality give Eq. (13).

Suppose now that the two stars agree in every normalized coordinate U_i , $i \geq 1$. Their baselines still have different images when J is nonempty. A bridge clears one selected fork denominator for each coordinate where the centers differ.

An orientation is a map $\sigma : J \rightarrow \{0, 1\}$. Put

$$P_\sigma = \prod_{j \in J} \Delta_j^{(\sigma_j)} \quad (14)$$

and define

$$K_{0,\sigma} = (V^{(0)} - V^{(1)}) P_\sigma - \sum_{i \in J} X_i^{(\sigma_i)} \varepsilon_i \prod_{j \in J \setminus \{i\}} \Delta_j^{(\sigma_j)}. \quad (15)$$

For $J = \emptyset$, the product is 1 and the sum is empty.

Lemma 2 (Oriented baseline bridge). Assume $U_i^{(0)} = U_i^{(1)}$ for every $i \in [L]$. For every orientation σ ,

$$A(K_{0,\sigma}) = 0, \quad K_{0,\sigma} = (U_0^{(0)} - U_0^{(1)})P_\sigma. \quad (16)$$

Hence $K_{0,\sigma}$ is nonzero exactly when $U_0^{(0)} \neq U_0^{(1)}$.

More generally, if

$$\|V^{(b)}\| \leq B_V^{(b)}, \quad \|X_i^{(b)}\| \leq B_{X,i}^{(b)}, \quad \|\Delta_i^{(b)}\|_{\text{op}} \leq B_{\Delta,i}^{(b)}, \quad \|\varepsilon_i\|_{\text{op}} \leq B_{\varepsilon,i},$$

then

$$\begin{aligned} \|K_{0,\sigma}\| &\leq (B_V^{(0)} + B_V^{(1)}) \prod_{j \in J} B_{\Delta,j}^{(\sigma_j)} \\ &\quad + \sum_{i \in J} B_{X,i}^{(\sigma_i)} B_{\varepsilon,i} \prod_{j \in J \setminus \{i\}} B_{\Delta,j}^{(\sigma_j)}. \end{aligned} \quad (17)$$

Under Eqs. (5)–(8) and $\|\varepsilon_i\|_{\text{op}} \leq B_\varepsilon$, this becomes

$$\|K_{0,\sigma}\| \leq 2B_V B_\Delta^h + hB_X B_\varepsilon B_\Delta^{h-1}, \quad (18)$$

where the second summand is absent when $h = 0$.

Proof. The full calculation is given in Appendix A. The image cancellation uses $A(V^{(0)} - V^{(1)}) = \sum_{i \in J} t_i \varepsilon_i$ and $A(X_i^{(\sigma_i)}) = t_i \Delta_i^{(\sigma_i)}$. Substitution of $X_i^{(\sigma_i)} = U_i \Delta_i^{(\sigma_i)}$ then leaves precisely the first identity in Eq. (16). The product P_σ is a unit, so the factorization also proves nonzeroness. Applying Eq. (4) term by term gives the two norm bounds.

The local and baseline cases give one deterministic safe radius for stars at different centers:

$$B_{\text{fib}}(h) = \max\{2B_X B_\Delta, 2B_V B_\Delta^h + hB_X B_\varepsilon B_\Delta^{h-1}\}. \quad (19)$$

At $h = 0$, the second entry is $2B_V$; thus same-root comparison is the zero-defect endpoint of the same bound.

The coordinate bridge clears a mismatch class in the cokernel of the fork-difference matrix. Appendix B.2 gives the general construction; the polynomial model below makes its degree quantitative through maximal minors.

2.3 Polynomial comparer models

Exact degree statements require a fixed algebraic model. We use a generic polynomial coefficient ring in which all denominator entries are independent indeterminates and center variables have degree zero. A polynomial module-linear comparer is an S -linear combination of the formal baseline and fork-numerator

generators whose image under the universal anchor map is zero. Its denominator-variable degree is the largest total denominator degree among its nonzero coefficients.

The *generic block-linear polynomial comparer model* permits coefficients from both branches in the same kernel expression. Its universal certificate has the form

$$K = p(V^{(0)} - V^{(1)}) + \sum_a (X_a^{(0)} c_a^{(0)} + X_a^{(1)} c_a^{(1)}), \quad (20)$$

where p is a scalar polynomial and the $c_a^{(b)}$ are polynomial column vectors. The universal anchor identity is equivalent, block by block, to

$$p\varepsilon_a + D_a^{(0)} c_a^{(0)} + D_a^{(1)} c_a^{(1)} = 0. \quad (21)$$

The model must provide local comparers for normalized block mismatches and a baseline bridge for every residual active block set J . Its degree is the largest denominator-variable degree among the nonzero coefficients $p, c_a^{(0)}, c_a^{(1)}$ required by either case.

The *branch-separated polynomial integralize-then-compare architecture* first maps each branch, using only that branch's variables, to an integral representative whose image is a scalar multiple of t_0 . Only these two integral representatives are then eliminated through a common multiplier. This branch-separation condition is part of every exact lower bound for that architecture.

Sections 3.1–3.3 use degree in independent denominator variables. The Vandermonde geometry later uses challenge-difference factor degree. Both grade integral comparison, but they are distinct polynomial gradings.

3 Comparison Degree in the Generic Model

Throughout this section, exactness refers to the generic polynomial module-linear model and denominator-variable grading of Section 2.3. Algebraically independent denominator variables make the constructive upper bounds match the lower bounds below.

3.1 Coordinate case

For coordinate stars, let $d_{\text{coord}}^*(h)$ be the minimum worst-case denominator-variable degree of a complete direct comparer system when the two centers differ in exactly h coordinates. Completeness requires both a local comparer for any normalized-coordinate mismatch and a bridge for a mismatch of the constant normalized coordinate after all other normalized coordinates agree.

Theorem 1 (Coordinate comparison degree). *For $L \geq 1$ and $0 \leq h \leq L$, the generic coordinate-linear polynomial comparer model has*

$$d_{\text{coord}}^*(h) = \max\{1, h\}. \quad (22)$$

Proof. The local and baseline cases give separate bounds. Lemma 1 has coefficient degree one. For a residual set J of size h , any oriented bridge in Eq. (15) has coefficient degree at most h . These comparers give the upper bound $\max\{1, h\}$.

For the local lower bound, restrict to one block and suppress the baselines. The two generic denominator variables form a 1×2 matrix, and every nonzero polynomial syzygy has degree at least one. For the synchronization-defect lower bound, restrict to a pure constant-coordinate mismatch while all nonconstant normalized coordinates agree. The baseline coefficient must lie in

$$\bigcap_{i \in J} (\delta_i^{(0)}, \delta_i^{(1)}).$$

The variable pairs are disjoint, so this intersection is their product and has initial denominator degree h . The detailed reductions are Lemmas 4 and 5 in Appendix A. Every complete comparer system therefore has degree at least both 1 and h .

The formula includes both endpoints. Same-root comparison has $h = 0$ but still pays the degree-one local coupling floor; a one-coordinate center defect has the same direct degree.

3.2 Block case

Coordinate denominators are scalar. A block of width r_a replaces the two scalars by two generic square matrices. The maximal minors needed to couple their normalized openings have degree r_a .

Fix block widths $r_1, \dots, r_m \geq 1$ and a field k . Let $\mathcal{S}$ be the polynomial ring over k generated by independent center variables and by the entries of

$$D_a^{(0)}, D_a^{(1)} \in \mathcal{S}^{r_a \times r_a}, \quad a \in [m].$$

Denominator entries have degree one and center variables have degree zero. Let t_0 together with the entries of the rows $\mathbf{t}_a = (t_{a,1}, \dots, t_{a,r_a})$ form a basis of a free target module. The universal anchor map is specified on formal baselines v_b and formal rows of fork numerators $\mathbf{x}_a^{(b)}$ by

$$\begin{aligned} \mathcal{A}(v_b) &= t_0 + \sum_{a=1}^m \mathbf{t}_a s_a^{(b)}, \\ \mathcal{A}(\mathbf{x}_a^{(b)}) &= \mathbf{t}_a D_a^{(b)}. \end{aligned} \tag{23}$$

Here $s_a^{(b)} \in \mathcal{S}^{r_a}$ is a column and $\varepsilon_a = s_a^{(0)} - s_a^{(1)}$; equivalently, one may use $s_a^{(1)}, \varepsilon_a$ as the independent degree-zero center variables. After localizing at $\det D_a^{(b)}$, write

$$\mathbf{u}_a^{(b)} = \mathbf{x}_a^{(b)} (D_a^{(b)})^{-1}, \quad u_0^{(b)} = v_b - \sum_a \mathbf{u}_a^{(b)} s_a^{(b)}. \tag{24}$$

Every polynomial module-linear certificate in these formal generators has the form

$$K = p(v_0 - v_1) + \sum_{a=1}^m (\mathbf{x}_a^{(0)} c_a^{(0)} + \mathbf{x}_a^{(1)} c_a^{(1)}). \quad (25)$$

Comparing the independent target rows in Eq. (23) shows that K is a universal kernel element if and only if, for every block,

$$p\varepsilon_a + D_a^{(0)} c_a^{(0)} + D_a^{(1)} c_a^{(1)} = 0. \quad (26)$$

Its degree is the maximum denominator-variable degree of its nonzero coefficients. For local comparison put

$$G_a = [D_a^{(0)} \mid D_a^{(1)}] \in \mathcal{S}^{r_a \times 2r_a}. \quad (27)$$

A full direct comparer system must detect a nonzero normalized mismatch in every block and, after all normalized rows agree, a nonzero constant normalized mismatch for every residual active set J . Completeness is evaluated on the unit chart where the block determinants are invertible. The lower-bound identities hold in the universal polynomial ring, and localization certifies nonzeroness of the resulting comparer outputs.

Theorem 2 (Block comparison degree). *In the generic block-linear polynomial module-linear comparer model of Section 3.2, the exact direct degree for residual active set J is*

$$d_{\text{direct}}^*(J) = \max \left\{ \max_a r_a, \sum_{a \in J} r_a \right\}. \quad (28)$$

Proof. For the local upper bound, let $h_a = \det D_a^{(0)}$ and let $d_{a,j}^{(1)}$ denote column j of $D_a^{(1)}$. The block-adjugate comparers

$$K_{a,j} = x_{a,j}^{(1)} h_a - \mathbf{x}_a^{(0)} \text{adj}(D_a^{(0)}) d_{a,j}^{(1)} \quad (29)$$

are universal kernel elements of coefficient degree r_a . On the unit chart,

$$K_{a,j} = (\mathbf{u}_a^{(1)} - \mathbf{u}_a^{(0)}) d_{a,j}^{(1)} h_a. \quad (30)$$

If every $K_{a,j}$ vanishes, collecting the columns gives $(\mathbf{u}_a^{(1)} - \mathbf{u}_a^{(0)}) D_a^{(1)} h_a = 0$; invertibility then forces the normalized rows to agree. Hence these comparers detect every local mismatch at degree r_a .

For the baseline bridge, the generic determinantal identity

$$\text{Ann}_{\mathcal{S}}([\varepsilon_a] \in \text{coker } G_a) = I_{r_a}(G_a) \quad (31)$$

means that a maximal minor clears the mismatch in block a . The square submatrix $D_a^{(0)}$ gives the explicit lift $D_a^{(0)} \text{adj}(D_a^{(0)}) \varepsilon_a = h_a \varepsilon_a$. Thus, with $P_J = \prod_{a \in J} h_a$,

$$K_0 = (v_0 - v_1) P_J - \sum_{a \in J} \mathbf{x}_a^{(0)} \text{adj}(D_a^{(0)}) \varepsilon_a \prod_{c \in J \setminus \{a\}} h_c \quad (32)$$

is a universal kernel element of degree $\sum_{a \in J} r_a$. On the unit chart it factors as $K_0 = (u_0^{(0)} - u_0^{(1)})P_J$, so it detects the residual baseline mismatch. The two constructions establish the upper bound by their maximum.

For the lower bound, a local comparer restricts to a nonzero element of $\ker G_a$. The generic matrix has

$$\text{indeg } \ker G_a = r_a, \quad (33)$$

so every complete system has degree at least $\max_a r_a$. A pure-baseline bridge has a nonzero baseline coefficient in

$$\bigcap_{a \in J} I_{r_a}(G_a).$$

The block matrices use disjoint denominator variables, and the intersection equals the product. Its initial degree is $\sum_{a \in J} r_a$. The determinantal identities, adjugate lifts, and both reductions are proved in Appendix B, using the standard generic-matrix syzygy and determinantal-ideal theory [5,6].

The two mismatch cases are alternative, so the worst-case degree is their maximum.

3.3 Branch-separated comparison

The branch-separated architecture commits to an integral representation of each whole branch before the other branch is used. This forces every block determinant into each branch multiplier.

Theorem 3 (Branch-separated comparison degree). *Within the branch-separated polynomial integralize-then-compare architecture of Section 2.3, instantiated in the generic block model of Section 3.2, the exact denominator-variable degree of the two-branch common multiplier is*

$$d_{\text{branch}}^* = 2 \sum_{a=1}^m r_a. \quad (34)$$

Proof. In one branch, compare coefficients in the requirement that an integralizer has image $t_0 D_b$. For each block a , the resulting equation places $D_b s_a^{(b)}$ in the image of $D_a^{(b)}$. The generic-vector annihilator is the maximal-minor ideal $I_{r_a}(D_a^{(b)}) = (\det D_a^{(b)})$. Hence

$$\prod_{a=1}^m \det D_a^{(b)} \mid D_b,$$

and every branch-local integralizer has degree at least $\sum_a r_a$. The variable sets of the two branches are disjoint, so any nonzero common multiplier is divisible by both branch products and has degree at least $2 \sum_a r_a$.

For the converse, multiply each branch by its determinant product and use the block adjugates to clear its normalized terms. Cross multiplication of the two resulting t_0 -multiples has degree $2 \sum_a r_a$. Appendix B gives the coefficient calculation and the divisibility proof.

4 Conditional-Fiber Extraction

The comparer needs two successful executions that agree on a chosen part of the challenge state; conditional resampling supplies the second execution.

4.1 Conditional resampling

Fix the part of the challenge that the two executions must share. From that stored label, repeatedly sample the remaining challenge and continuation until the selected-success predicate holds. The retry count at a label is the reciprocal of its selected-success probability, but the label is entered with at most the corresponding probability mass. Averaging before conditioning cancels this reciprocal.

Let S have raw law μ on a standard Borel challenge space $\mathcal{S}$, and let

$$\pi : \mathcal{S} \longrightarrow \mathcal{Z}$$

be a measurable fiber label. Write $\lambda = \pi_{\#}\mu$ and let $\{\mu_z\}_{z \in \mathcal{Z}}$ be a regular conditional law of S given $\pi(S) = z$. A retry entered at the stored label z draws

$$S' \leftarrow \mu_z = \mu(\cdot \mid \pi(S') = z). \quad (35)$$

The checkpoint also restores the prover state and random tape fixed at that point, and every other continuation coin is sampled from its raw conditional law. This is the exact conditional-fiber experiment used below.

For a fixed label z , let u collect the raw challenge and continuation context sampled by one trial, with conditional law $\rho_z(du)$. Let $q(z, u)$ be the probability of the efficiently decidable selected-success event after u is fixed. A retry node is entered with a subdensity $w(z, u)$ satisfying

$$0 \leq w(z, u) \leq q(z, u).$$

Set

$$\bar{w}(z) = \int w(z, u) \rho_z(du), \quad \bar{q}(z) = \int q(z, u) \rho_z(du).$$

The ratio $\bar{w}(z)/\bar{q}(z)$ is defined as zero when $\bar{q}(z) = 0$; the domination $w \leq q$ then also gives $\bar{w}(z) = 0$.

Lemma 3 (Entry-weighted cancellation). *Under the setup above, the unconditional entry-weighted number of retry invocations satisfies*

$$\int_{\mathcal{Z}} \frac{\bar{w}(z)}{\bar{q}(z)} \lambda(dz) \leq 1. \quad (36)$$

For finite $\mathcal{Z}$, this is

$$\sum_z \lambda(z) \frac{\bar{w}(z)}{\bar{q}(z)} \leq 1.$$

This is an unconditional entry-weighted bound; individual success-selected labels may have larger conditional retry cost.

After the retry selects a successful trial, its entry-weighted density with respect to $\lambda(dz)\rho_z(du)$ is

$$\tilde{w}(z, u) = q(z, u) \frac{\bar{w}(z)}{\bar{q}(z)} \leq q(z, u). \quad (37)$$

This domination is preserved when the selected state is projected to a later checkpoint. Appendix C gives the measure-theoretic statement and the tower argument for nested retries.

4.2 Two-star extraction

Assume an outer checkpoint from which the final challenge is sampled, branch-local checkpoints for coordinate forks, exact conditional sampling on the retained fiber, and the affine rows and shortness bounds required by the deterministic comparer. Let **TgtSucc** be the target event and **Succ** the retry event containing it. For coordinate i , let κ_i bound the raw joint mass of selected fork success and a nonunit fork difference, uniformly in the fixed history and baseline coordinate. After the first normalized tuple is fixed, let κ_{local} bound failure of the fresh downstream validation. Appendix C gives the event inclusions and checkpoint typing used in the proof.

Theorem 4 (Conditional-fiber two-star extraction). *Assume the exact conditional-sampling and checkpoint interface above. Suppose that two selected stars with a common fiber label satisfy the deterministic comparison theorem of Sections 2–3: a mismatch returns a nonzero kernel vector within its stated comparison radius, while equality supplies one normalized tuple to the fresh downstream procedure. Suppose the downstream procedure returns source witnesses or a nonzero kernel vector outside its raw bad event. If*

$$\epsilon = \mathbb{P}[\text{TgtSucc}],$$

then an interactive black-box extractor succeeds with probability at least

$$\mathbb{P}[\text{extract}] \geq \epsilon - \left(2 \sum_{i=1}^L \kappa_i + \kappa_{\text{local}} \right). \quad (38)$$

If every raw invocation and exact conditional sampler has polynomial worst-case cost, the extractor is expected polynomial time. The coordinate two-star schedule uses at most

$$\mathbb{E}[\text{additional retry invocations}] \leq 2L + 1 \quad (39)$$

in unconditional expectation beyond the initial genuine invocation.

The count in Eq. (39) comprises L coordinate forks for each star and one conditional-fiber retry for the second baseline. Unit failures are charged through the raw bad-event terms, while the retry predicate remains selected success. The proof in Appendix C applies Eqs. (36) and (37) at these nodes.

4.3 Coordinate fibers

Same-root replay is the special case $\pi = \text{id}_S$. Its conditional law is the point mass at the stored numerical challenge, so the retry resamples only the continuation.

For a product challenge $S = (S_1, \dots, S_L) \leftarrow \mathcal{D}^L$ and a residual unsynchronized set $J \subseteq [L]$, define

$$\pi_J(S) = S_{[L] \setminus J}. \quad (40)$$

The exact conditional law fixes $S_{[L] \setminus J}$ and samples $S'_J \leftarrow \mathcal{D}^J$. Hence $J = \emptyset$ recovers same-root replay, while larger J exposes precisely the residual center mismatch handled by the different-center comparer.

5 Cyclo

For Cyclo, the fiber label specifies which final challenge coordinates are retained across the two outer branches.

5.1 Interface and partial synchronization

Let

$$R = \mathbb{Z}[X]/(X^{128} + 1), \quad R_q = R/qR, \quad q = 2^{50} - 351.$$

Ring elements use canonical centered representatives and the coefficient infinity norm. At the checkpoint H_{preRange} , the Ajtai matrix $\mathbf{R}$ and the affine commitment images $t_0, t_1, \dots, t_L$ are fixed. Range, batching, sum-check, and final folding challenges remain fresh. Each outer continuation reaches a branch-local checkpoint H_{final} before sampling its final vector challenge

$$S = (S_1, \dots, S_L) \leftarrow \mathcal{D}^L.$$

Appendix D verifies this checkpoint chronology and the corresponding protocol rows. The exact-strong benchmark uses the uniform fixed-weight support

$$\mathcal{D} = \mathcal{D}_{32} := \{s \in \{-1, 0, 1\}^{128} : \text{wt}(s) = 32\}.$$

For a selected terminal baseline V and its i th coordinate fork, the commitment rows give

$$\mathbf{R}V = t_0 + \sum_{j=1}^L t_j S_j, \quad \mathbf{R}X_i = t_i \Delta_i,$$

where X_i is the difference of two terminal openings and Δ_i is the corresponding challenge difference. Fixing L , abbreviate the terminal opening cap by β . The conservative interface bounds are

$$B_V = \beta, \quad B_X = 2\beta, \quad B_\Delta = 2\gamma, \quad B_\varepsilon = 2\gamma. \quad (41)$$

We use $\gamma = 128$ throughout this parameter line. The fixed-weight support has the smaller native operator norm 32, but the terminal budget and every dependent radius here retain the common conservative bound 128.

The exact-capacity Extension component has digit cap c and contributes a radius- $2c$ kernel branch. With $Q_{\text{fld}} = q^e$, the branch-local raw validation loss is

$$\kappa_{\text{local}} = \frac{\ell_0 + 2\ell_1}{Q_{\text{fld}}} + \frac{L\ell_1(2b+3)}{Q_{\text{fld}}} + \frac{L\ell_C}{|C|}. \quad (42)$$

The exact-capacity source encoding and raw Range and sum-check bounds are verified in Appendix D; the support arithmetic is in Appendix E.

Let $J \subseteq [L]$ be the coordinates that may differ between the two outer centers, and put $h = |J|$. The retained fiber label is

$$\pi_J(S) = S_{[L] \setminus J}. \quad (43)$$

Since the final coordinates are independent under $\mathcal{D}^L$, an exact fiber retry fixes $S_{[L] \setminus J}$ and draws $S'_J \leftarrow \mathcal{D}^J$. The checkpoint chronology therefore realizes the conditional law required by Theorem 4. For every coordinate fork, pairwise unit differences on $\mathcal{D}_{32}$ give

$$\kappa_i = |\mathcal{D}_{32}|^{-1}, \quad |\mathcal{D}_{32}| = \binom{128}{32} 2^{32}.$$

5.2 Radius and benchmark

Equation (19) gives the different-center comparison bound for h residual coordinates as

$$\max \{2B_X B_\Delta, 2B_V B_\Delta^h + hB_X B_\varepsilon B_\Delta^{h-1}\},$$

where the second summand involving h is zero at $h = 0$. Substitution of Eq. (41), followed by the downstream radius- $2c$ branch, yields

$$B_{\text{Cyclo}}(h) = \max \{8\beta\gamma, 2(h+1)\beta(2\gamma)^h, 2c\}. \quad (44)$$

Theorem 5 (Cyclo conditional-fiber extraction). *Assume the checkpoint and exact-capacity interfaces verified in Appendix D, independent final coordinates from $\mathcal{D}_{32}$, and the exact conditional sampler of Eq. (43). For target success probability ϵ , the interactive extractor outputs source witnesses or a nonzero $\mathbf{R}$ -kernel vector of coefficient radius at most $B_{\text{Cyclo}}(h)$ with probability at least*

$$\epsilon - \left(\frac{2L}{|\mathcal{D}_{32}|} + \kappa_{\text{local}} \right). \quad (45)$$

Its outer coordinate schedule uses at most $2L+1$ additional retry invocations in unconditional expectation.

The local coupling term already equals the defect term at $h = 1$. At $h = 0$, 2β is dominated by $8\beta\gamma$ on this parameter line. Hence

$$B_{\text{Cyclo}}(1) = B_{\text{Cyclo}}(0). \quad (46)$$

Synchronizing all but one coordinate is therefore radius-equivalent to exact same-root synchronization under the certified bound.

For the displayed benchmark, retain

$$\beta = \beta_L := B + 2TLb\gamma, \quad B = 2^{10}, \quad T = 64, \quad b = c = 1, \quad \gamma = 128. \quad (47)$$

Centered uniqueness is certified when $B_{\text{Cyclo}}(h) < (q - 1)/2$, whose base-two logarithm is 49 to the displayed precision. Under this conservative parameter line, $h = 0$ and $h = 1$ have the same radius. For $L = 2$, the bound remains below the centered threshold through $h = 2$; for $L = 4$ and $L = 8$, it remains below through $h = 3$. Appendix E gives the exact integer radii.

At $L = 2, h = 0$, the earlier same-root value is 25.0444 bits, while the centered threshold is 49; the branch-separated safe bound is 49.6294.

6 Vandermonde Two-Star Consistency

Vandermonde extraction gives a dense counterpart to the coordinate case of Sections 2–3. Its natural cost is challenge-difference degree.

6.1 Two-star setup

Let $k \geq 1$ and $n = k + 1$. Write

$$r(x) = (1, x, \dots, x^k), \quad V(x_0, \dots, x_k) = \begin{pmatrix} r(x_0) \\ \vdots \\ r(x_k) \end{pmatrix}.$$

The source interface is the $(k + 1)$ -special-sound template of Esgin, Steinfeld, Liu, and Liu [4, Protocol 1]. If $C = (C_0, \dots, C_k)^\top$ is the common commitment anchor and $O = (o_0, \dots, o_k)^\top$ is a column of accepted openings, its equations are

$$\mathcal{B}(o_i) = r(x_i)C, \quad \mathcal{B}^n(O) = VC. \quad (48)$$

Esgin et al. apply $\text{adj}(V)$ to one such system. The polynomial identity $\text{adj}(V)V = \det(V)I$ produces an integral relaxed opening before specialization [4, Sect. 3.2 and Facts 1–2].

We compare two independently extracted stars under the same anchor C . Their challenges are $x_0, \dots, x_k$ and $y_0, \dots, y_k$, and their opening columns are O_0 and O_1 . Put

$$V_0 = V(x_0, \dots, x_k), \quad V_1 = V(y_0, \dots, y_k), \quad \delta_b = \det(V_b).$$

Within each star the challenges are distinct, as required by special soundness. Cross-star challenge values may coincide.

6.2 Direct comparison

For every row i of the second star, define

$$K_i = \delta_0 o_i^{(1)} - r(y_i) \operatorname{adj}(V_0) O_0. \quad (49)$$

Linearity and Eq. (48) give

$$\mathcal{B}(K_i) = \delta_0 r(y_i) C - r(y_i) \operatorname{adj}(V_0) V_0 C = 0.$$

Each K_i is an integral kernel opening by a polynomial identity valid before localization.

The coefficients expose the natural grading. If

$$\Delta_{\hat{j}} = \prod_{\substack{0 \leq a < b \leq k \\ a, b \neq j}} (x_b - x_a),$$

then

$$\lambda_j(y) := [r(y) \operatorname{adj}(V_0)]_j = (-1)^{k+j} \Delta_{\hat{j}} \prod_{\ell \neq j} (y - x_\ell). \quad (50)$$

Both δ_0 and every nonzero $\lambda_j(y)$ contain

$$\kappa := \binom{k+1}{2}$$

challenge-difference factors. Appendix F proves the factorization as a polynomial identity and relates nonzeroness to the standard branchwise comparison.

Exactness here is within the anchor-universal polynomial-linear comparer model: coefficients are polynomials in independent challenge variables, the relation holds for every common anchor, and cost is the largest total challenge-difference degree of a raw-opening coefficient.

For two accepted stars, put

$$N_b = \operatorname{adj}(V_b) O_b, \quad W = \delta_0 N_1 - \delta_1 N_0, \quad W_k = e_k^\top W.$$

Here W_k is the same-factor witness comparison obtained after separately integralizing the two branches.

We call a nonempty family $\mathcal{K}$ complete for this two-star consistency problem if, for every specialization to two accepted stars under one anchor,

$$W_k \neq 0 \implies \text{some comparer in } \mathcal{K} \text{ outputs a nonzero kernel opening.}$$

The challenge-difference degree of $\mathcal{K}$ is the supremum of the degrees of its members.

Theorem 6 (Direct comparison degree). *Let $k \geq 1$. Every consistency-complete family in the model above has challenge-difference degree at least $\binom{k+1}{2}$. The family $\{K_0, \dots, K_k\}$ from Eq. (49) attains this bound, so*

$$d_{\text{Vand, direct}}^* = \binom{k+1}{2}.$$

Place the $2n$ challenge variables into the $n \times 2n$ evaluation matrix whose columns are $(1, z, \dots, z^{n-1})^\top$. Its maximal minors are Vandermonde determinants of degree $\binom{n}{2}$. The simultaneous rank-drop locus has codimension $n+1$, so the Buchsbaum–Rim complex generates the kernel by the $(n+1)$ -column circuit syzygies. The kernel therefore has no nonzero coefficient vector below degree $\binom{n}{2}$. Equation (49) gives $k+1$ such circuits, all of degree $\binom{n}{2}$. Appendix F.3 gives the height and generation argument, while Appendix F.5 proves

$$W_k \neq 0 \implies K \neq 0 \implies \exists i : K_i \neq 0.$$

The finite-support hypotheses below convert the comparer coefficients into safe norm upper bounds.

6.3 Branch-separated comparison

A branch-separated integralize-then-compare architecture chooses independently for each star polynomials $a^{(b)} = (a_0^{(b)}, \dots, a_k^{(b)})$ and $D_b \neq 0$ satisfying

$$(a^{(b)})^\top V_b = D_b e_k^\top. \quad (51)$$

Joint scalar elimination follows the two branch-local integralizers and chooses $c_0 D_0 = c_1 D_1 = M \neq 0$. The canonical choice outputs W_k up to sign, so the branch architecture is measured against the same consistency criterion as the direct comparer family.

Theorem 7 (Branch-separated comparison degree). *Every branch-local scalar integralizer has the form*

$$D_b = h_b \delta_b, \quad a_i^{(b)} = h_b \Gamma_i^{(b)}, \quad \Gamma^{(b)} = e_k^\top \text{adj}(V_b).$$

Consequently, within the branch-separated polynomial integralize-then-compare architecture, the minimum raw-response coefficient degree is

$$d_{\text{Vand,branch}}^* = k^2.$$

The two branch determinant factors are coprime because their variable sets are disjoint. Their common multiplier is therefore divisible by $\delta_0 \delta_1$, while each raw coefficient retains one adjugate-row factor of degree $k(k-1)/2$. The canonical adjugate integralizers attain the resulting degree $\binom{k+1}{2} + k(k-1)/2 = k^2$; Appendix F gives the full divisibility proof. For $k \geq 2$,

$$\binom{k+1}{2} < k^2. \quad (52)$$

For $k = 1$, both architectures have degree one. Their triangle-inequality prefactors can nevertheless differ, as the norm bounds below show.

6.4 Norm consequences

Let N_R denote the ring coefficient dimension called d by Esgin et al. Suppose every nonzero challenge difference belongs to their set $\Delta\mathcal{C}_{w,p}^{N_R}$, and let

$$T_f = \max_{b,i} \|f_i^{(b)}\|, \quad T_z = \max_{b,i} \|z_i^{(b)}\|, \quad T_H = \sqrt{T_z^2 + T_f^2}.$$

Their product bound and negacyclic multiplication inequality imply, for a product c of $s \geq 1$ challenge differences,

$$\|cu\| \leq N_R(2p)^s w^{s-1} \|u\|. \quad (53)$$

This scalar-times-vector inequality follows from Lemma 3 and the norm lemma in Appendix B.1 of Esgin et al. [4, Lemma 3, Lemma 8, and App. H].

Corollary 1 (Safe HMC and UMC comparison radii). *Let $\kappa = \binom{k+1}{2}$. Whenever the relevant comparison vector is nonzero, the direct and branch-separated routes give the safe bounds*

$$\begin{aligned} B_{\text{dir}}^{\text{HMC}} &= (k+2)N_R(2p)^\kappa w^{\kappa-1} T_H, \\ B_{\text{dir}}^{\text{UMC}} &= (k+2)N_R(2p)^\kappa w^{\kappa-1} T_z, \\ B_{\text{bw}}^{\text{HMC}} &= 2(k+1)N_R(2p)^{k^2} w^{k^2-1} T_H, \\ B_{\text{bw}}^{\text{UMC}} &= 2(k+1)N_R(2p)^{k^2} w^{k^2-1} T_z. \end{aligned}$$

For either commitment type,

$$\frac{B_{\text{bw}}}{B_{\text{dir}}} = \frac{2(k+1)}{k+2} (2pw)^{k(k-1)/2}.$$

For HMC, a nonzero full opening vector in the kernel of its commitment matrix is directly an M-SIS vector. For UMC, the top block of the opening equation places the randomness component in an M-SIS kernel, and nonzeroness of the full opening pair forces this randomness component to be nonzero. Appendix G records this typing and derives all four bounds.

The challenge-space example of Esgin et al. uses $(N_R, w, p) = (256, 60, 1)$ [4, Sect. 3.3]. Factoring out the common response cap T gives the following deterministic evaluations.

We apply the one-star extractor of Esgin et al. twice and measure the subsequent reconciliation step under a common commitment anchor.

Table 1. Direct versus branch-separated Vandermonde comparison for $(N_R, w, p) = (256, 60, 1)$. The logarithms are formula-derived norm multipliers of the common response cap T ; “gain” denotes $\log_2(B_{\text{bw}}/B_{\text{dir}})$.

k	direct	degree	branch	degree	$\log_2(B_{\text{dir}}/T)$	$\log_2(B_{\text{bw}}/T)$	gain
2		3		4	24.8138	32.3056	7.4919
3		6		9	45.8564	67.2551	21.3987
4		10		16	73.7470	115.9253	42.1783
5		15		25	108.5038	178.3503	69.8465

7 Related Work

The closest algebraic antecedents concern integral comparison and special soundness. Esgin, Steinfeld, Liu, and Liu replace inverse Vandermonde coefficients by $\text{adj}(V)V = \det(V)I$, obtaining a one-star relaxed opening whose factor is the Vandermonde product of challenge differences [4, Protocol 1, Facts 1–2, and Sect. 3.2]. Attema, Lyubashevsky, and Seiler instead cross-multiply two weak openings with relaxation factors to obtain a short Module-SIS binding vector [7, Lemma 4.3]. Coordinate-wise special soundness appears in the lattice polynomial commitments of Fenzi, Moghaddas, and Nguyen, whose recursive extractor combines two t -difference slacks and a fresh difference into a $(2t + 1)$ -difference slack [8, Lemmas 5.10–5.11]. Our Vandermonde application takes the Esgin et al. one-star interface as input and studies the integral comparison of two independently extracted stars.

Lattice folding systems also certify range or norm bounds under Module-SIS relations; LatticeFold, LatticeFold+, and Cyclo do so with different support and protocol choices [9,10,3]. Lova uses probabilistic norm bounds and Chernoff–Hoeffding-style completeness tradeoffs [11]. Other constructions control extraction-side shortness through sum-checks, projections, embeddings, or norm-reducing recursion [12,13,14] [15, Abstract and Sects. 2.2–2.3] [16, Abstract and Sect. 1.1] [17, Abstract and pp. 5–7] [18, Abstract, Fig. 2, and Sect. 7.2]. This line of work supplies the surrounding norm-certification problem; our bounds isolate the integral comparison step after the extractor has retained only part of the challenge.

Accumulation schemes and recursive folding determine where that extraction boundary lies. Accumulation reduces repeated verification to one accumulated claim, while Nova folds committed relaxed R1CS for incrementally verifiable computation [1, Abstract and Sect. 5] [2, Abstract and Sect. 1.1]. ProtoStar and ProtoGalaxy respectively build from special-sound protocols and fold multiple instances in a Lagrange-basis representation [19, Abstract and Sect. 3] [20, Sects. 1.1 and 4]. Reductions of knowledge formalize composable maps between knowledge relations [21, Abstract, Def. 8, and Sect. 2.2]; later designs move sum-check, relation reduction, or auxiliary commitments to different stages of the fold [22, Sects. 1.3 and 4] [23, Abstract and Sect. 1.1] [24, Abstract and Cor. 1]. In this paper these distinctions matter only insofar as they determine which public

state is fixed before a fork and which challenge coordinates remain available for resampling.

Multi-round Fiat–Shamir analyses and adaptive special soundness separate transcript selection from the later challenges used by an extractor [25,26]. Cyclo’s coordinate-wise forking lemma provides the baseline-and-fork chronology used in our coordinate application [3, Lemma 4 and proof]. Our retry step conditions on the stored fiber label and regenerates the downstream verifier randomness, so the comparison degree is tied to this extraction boundary rather than to the protocol name.

8 Conclusion

Comparing two relaxed extraction objects need not require clearing every denominator carried by either branch. In the generic block model, direct comparison pays the larger of the local block width and the denominator degree needed to bridge the remaining center mismatch. Clearing the two branches independently first pays the determinant product of both branches instead.

The conditional-fiber argument links these degree bounds to extraction by supplying successful executions with a controlled synchronization defect. Cyclo realizes the coordinate case, while the Esgin-style Vandermonde construction realizes the dense case. The result is a model-specific link between retained synchronization, comparison degree, and the certified SIS radii displayed in the two instantiations.

Declaration on the use of large language models. Large language models were used solely for language polishing, typesetting and grammar checks, and editorial review of the clarity, precision, and rigor of the exposition. They played no role in formulating or validating the paper’s core ideas, mathematical results, proofs, or conclusions. The authors reviewed all AI-assisted edits and take full responsibility for the content of this manuscript.

References

1. Bünz, B., Chiesa, A., Mishra, P., Spooner, N.: Recursive proof composition from accumulation schemes. In: Theory of Cryptography–TCC 2020, Part II. LNCS, vol. 12551, pp. 1–18. Springer (2020). https://doi.org/10.1007/978-3-030-64378-2_1, <https://eprint.iacr.org/2020/499>
2. Kothapalli, A., Setty, S., Tzalla, I.: Nova: Recursive zero-knowledge arguments from folding schemes. In: Advances in Cryptology–CRYPTO 2022, Part IV. LNCS, vol. 13510, pp. 359–388. Springer (2022). https://doi.org/10.1007/978-3-031-15985-5_13, <https://eprint.iacr.org/2021/370>
3. Garreta, A., Lipmaa, H., Luhaäär, U., Osadnik, M.: Cyclo: Lightweight lattice-based folding via partial range checks. In: Advances in Cryptology–EUROCRYPT 2026, Part VII. LNCS, vol. 16547, pp. 33–62. Springer (2026). https://doi.org/10.1007/978-3-032-25336-1_2

4. Esgin, M.F., Steinfeld, R., Liu, J.K., Liu, D.: Lattice-based zero-knowledge proofs: New techniques for shorter and faster constructions and applications. In: *Advances in Cryptology–CRYPTO 2019, Part I*. LNCS, vol. 11692, pp. 115–146. Springer (2019). https://doi.org/10.1007/978-3-030-26948-7_5
5. Eisenbud, D.: *Commutative Algebra with a View Toward Algebraic Geometry*, Graduate Texts in Mathematics, vol. 150. Springer (1995). <https://doi.org/10.1007/978-1-4612-5350-1>
6. Bruns, W., Vetter, U.: *Determinantal Rings*, Lecture Notes in Mathematics, vol. 1327. Springer (1988). <https://doi.org/10.1007/BFb0080378>
7. Attema, T., Lyubashevsky, V., Seiler, G.: Practical product proofs for lattice commitments. In: *Advances in Cryptology–CRYPTO 2020, Part II*. LNCS, vol. 12171, pp. 470–499. Springer (2020). https://doi.org/10.1007/978-3-030-56880-1_17
8. Fenzi, G., Moghaddas, H., Nguyen, N.K.: Lattice-based polynomial commitments: Towards asymptotic and concrete efficiency. *Journal of Cryptology* **37**(3), 31 (2024). <https://doi.org/10.1007/s00145-024-09511-8>
9. Boneh, D., Chen, B.: LatticeFold: A lattice-based folding scheme and its applications to succinct proof systems. In: *Advances in Cryptology–ASIACRYPT 2025, Part III*. LNCS, vol. 16247, pp. 330–362. Springer (2025). https://doi.org/10.1007/978-981-95-5099-9_11
10. Boneh, D., Chen, B.: LatticeFold+: Faster, simpler, shorter lattice-based folding for succinct proof systems. In: *Advances in Cryptology–CRYPTO 2025, Part VII*. LNCS, vol. 16006, pp. 327–361. Springer (2025). https://doi.org/10.1007/978-3-032-01907-3_11
11. Fenzi, G., Knabenhans, C., Nguyen, N.K., Pham, D.T.: Lova: Lattice-based folding scheme from unstructured lattices. In: *Advances in Cryptology–ASIACRYPT 2024, Part IV*. LNCS, vol. 15487, pp. 303–326. Springer (2024). https://doi.org/10.1007/978-981-96-0894-2_10
12. Kuriyama, S., Lai, R.W.F., Osadnik, M., Tucci, L.: SALSAA: Sumcheck-aided lattice-based succinct arguments and applications. *IACR Cryptology ePrint Archive*, Report 2025/2124 (2025)
13. Kloof, M., Lai, R.W.F., Nguyen, N.K., Osadnik, M.: RoK and Roll: Verifier-efficient random projection for $\tilde{O}(\lambda)$ -size lattice arguments. In: *Advances in Cryptology–ASIACRYPT 2025, Part III*. LNCS, vol. 16247, pp. 297–329. Springer (2025). https://doi.org/10.1007/978-981-95-5099-9_10
14. Beullens, W., Seiler, G.: LaBRADOR: Compact proofs for R1CS from module-SIS. In: *Advances in Cryptology–CRYPTO 2023, Part V*. LNCS, vol. 14085, pp. 518–548. Springer (2023). https://doi.org/10.1007/978-3-031-38554-4_17
15. Nguyen, W., Setty, S.: Neo and SuperNeo: Post-quantum folding with pay-per-bit costs over small fields. In: *Advances in Cryptology–CRYPTO 2026, Part IX*. LNCS, vol. 16808, pp. 492–523. Springer (2026). https://doi.org/10.1007/978-3-032-35424-2_17
16. Balbás, D., Nitulescu, A., Plançon, M.: ProtogaLattice: Constant-round lattice-based folding for general polynomial relations. *IACR Cryptology ePrint Archive*, Report 2026/1317 (2026), <https://eprint.iacr.org/2026/1317>
17. Kloof, M., Lai, R.W.F., Nguyen, N.K., Osadnik, M., Tucci, L.: RoKoko: Lattice-based succinct arguments, a committed refinement. *IACR Cryptology ePrint Archive*, Report 2026/575 (2026), <https://eprint.iacr.org/2026/575>
18. Osadnik, M.: PikkuFold: Efficient folding in a few kilobytes. *IACR Cryptology ePrint Archive*, Report 2026/1809 (2026), <https://eprint.iacr.org/2026/1809>

19. Bünz, B., Chen, B.: ProtoStar: Generic efficient accumulation/folding for special-sound protocols. In: Advances in Cryptology–ASIACRYPT 2023, Part II. LNCS, vol. 14439, pp. 77–110. Springer (2023). https://doi.org/10.1007/978-981-99-8724-5_3, <https://eprint.iacr.org/2023/620>
20. Eagen, L., Gabizon, A.: ProtoGalaxy: Efficient ProtoStar-style folding of multiple instances. IACR Cryptology ePrint Archive, Report 2023/1106 (2023), <https://eprint.iacr.org/2023/1106>
21. Kothapalli, A., Parno, B.: Algebraic reductions of knowledge. In: Advances in Cryptology–CRYPTO 2023, Part IV. LNCS, vol. 14084, pp. 669–701. Springer (2023). https://doi.org/10.1007/978-3-031-38551-3_21
22. Kothapalli, A., Setty, S.: HyperNova: Recursive arguments for customizable constraint systems. In: Advances in Cryptology–CRYPTO 2024, Part X. LNCS, vol. 14929, pp. 345–379. Springer (2024). https://doi.org/10.1007/978-3-031-68403-6_11, <https://eprint.iacr.org/2023/573>
23. Dimitriou, N., Garreta, A., Manzur, I., Vlasov, I.: Mova: Nova folding without committing to error terms. IACR Cryptology ePrint Archive, Report 2024/1220 (2024), <https://eprint.iacr.org/2024/1220>
24. Kothapalli, A., Setty, S.: NeutronNova: Group-based folding done right. In: Advances in Cryptology–CRYPTO 2026, Part IX. LNCS, vol. 16808, pp. 524–556. Springer (2026). https://doi.org/10.1007/978-3-032-35424-2_18, <https://eprint.iacr.org/2024/1606>
25. Attema, T., Fehr, S., Klooß, M.: Fiat–Shamir transformation of multi-round interactive proofs. In: Theory of Cryptography–TCC 2022, Part I. LNCS, vol. 13747, pp. 113–142. Springer (2022). https://doi.org/10.1007/978-3-031-22318-1_5
26. Attema, T., Klooß, M., Lai, R.W.F., Yatsyna, P.: Adaptive special soundness: Improved knowledge extraction by adaptive useful challenge sampling. IACR Cryptology ePrint Archive, Report 2024/2038 (2024), <https://eprint.iacr.org/2024/2038>
27. Lyubashevsky, V., Seiler, G.: Short, invertible elements in partially splitting cyclotomic rings and applications to lattice-based zero-knowledge proofs. In: Advances in Cryptology–EUROCRYPT 2018, Part I. LNCS, vol. 10820, pp. 204–224. Springer (2018). https://doi.org/10.1007/978-3-319-78381-9_8

Supplementary Material

A Coordinate Comparer Algebra

This appendix gives the coefficient calculations behind the different-center bridge and the two coordinate lower bounds. All degree statements use the generic polynomial module-linear model of Section 2.3.

A.1 The oriented bridge

Assume the setup of Section 2.1 and write $U_i = U_i^{(0)} = U_i^{(1)}$ for $i \geq 1$. Subtracting the two baseline equations gives

$$A(V^{(0)} - V^{(1)}) = \sum_{i \in J} t_i \varepsilon_i. \quad (54)$$

For an orientation $\sigma : J \rightarrow \{0, 1\}$, multiplication by $P_\sigma = \prod_{j \in J} \Delta_j^{(\sigma_j)}$ and use of $A(X_i^{(\sigma_i)}) = t_i \Delta_i^{(\sigma_i)}$ yields

$$\begin{aligned} A(K_{0,\sigma}) &= \sum_{i \in J} t_i \varepsilon_i P_\sigma - \sum_{i \in J} t_i \Delta_i^{(\sigma_i)} \varepsilon_i \prod_{j \in J \setminus \{i\}} \Delta_j^{(\sigma_j)} \\ &= 0. \end{aligned}$$

The normalized-coordinate identity contains the nonzeroness argument. From Eq. (9),

$$\begin{aligned} V^{(0)} - V^{(1)} &= (U_0^{(0)} - U_0^{(1)}) + \sum_{i=1}^L (U_i^{(0)} s_i^{(0)} - U_i^{(1)} s_i^{(1)}) \\ &= (U_0^{(0)} - U_0^{(1)}) + \sum_{i \in J} U_i \varepsilon_i. \end{aligned}$$

Moreover,

$$X_i^{(\sigma_i)} \varepsilon_i \prod_{j \in J \setminus \{i\}} \Delta_j^{(\sigma_j)} = U_i \varepsilon_i P_\sigma.$$

Substitution in Eq. (15) proves

$$K_{0,\sigma} = (U_0^{(0)} - U_0^{(1)}) P_\sigma.$$

Every factor of P_σ is a unit in the instantiated comparison, so this kernel vector is nonzero precisely when the two constant normalized coordinates differ.

For the heterogeneous norm bound, apply the triangle and action inequalities before using any normalized opening:

$$\begin{aligned} \|K_{0,\sigma}\| &\leq \|V^{(0)} - V^{(1)}\| \prod_{j \in J} \|\Delta_j^{(\sigma_j)}\|_{\text{op}} \\ &\quad + \sum_{i \in J} \|X_i^{(\sigma_i)}\| \|\varepsilon_i\|_{\text{op}} \prod_{j \in J \setminus \{i\}} \|\Delta_j^{(\sigma_j)}\|_{\text{op}}. \end{aligned}$$

Inserting the branch- and coordinate-dependent bounds gives Eq. (17). Under homogeneous bounds, the first line contributes $2B_V B_\Delta^h$, and the h terms in the second line contribute at most $hB_X B_\varepsilon B_\Delta^{h-1}$. When $h = 0$, the second sum is empty and the first product is 1.

A.2 The coordinate mismatch annihilator

We now pass to a polynomial ring $\mathcal{S}$ over a field. Give each denominator variable degree one and each center variable degree zero. For an active coordinate i , set

$$G_i = \begin{bmatrix} \delta_i^{(0)} & \delta_i^{(1)} \end{bmatrix}, \quad I_i = (\delta_i^{(0)}, \delta_i^{(1)}). \quad (55)$$

The center mismatch $\varepsilon_i = s_i^{(0)} - s_i^{(1)}$ is independent of the two denominator variables. Therefore

$$\text{Ann}_{\mathcal{S}}([\varepsilon_i] \in \text{coker } G_i) = I_i. \quad (56)$$

Indeed, the left side is the colon ideal $I_i : \varepsilon_i$. Modulo I_i , the polynomial ε_i is nonzero in a polynomial domain, so this colon ideal is I_i .

Take the direct sum of the maps G_i for $i \in J$ and let $\varepsilon = (\varepsilon_i)_{i \in J}$. An element annihilates the class of ε exactly when it annihilates every component. Hence

$$\text{Ann}_{\mathcal{S}}([\varepsilon] \in \text{coker } G_J) = \bigcap_{i \in J} I_i = \prod_{i \in J} I_i. \quad (57)$$

The second equality uses the disjoint denominator-variable pairs. Every nonzero element of this product has denominator degree at least h , and each oriented product

$$p_\sigma = \prod_{i \in J} \delta_i^{(\sigma_i)} \quad (58)$$

attains degree h .

Equation (57) also recovers the oriented lift directly. For $a = p_\sigma$, define $c_i = (c_i^{(0)}, c_i^{(1)})$ by

$$c_i = \begin{cases} (\varepsilon_i a / \delta_i^{(0)}, 0), & \sigma_i = 0, \\ (0, \varepsilon_i a / \delta_i^{(1)}), & \sigma_i = 1. \end{cases}$$

Then $G_i c_i = a \varepsilon_i$. Inserting these lifts into

$$(V^{(0)} - V^{(1)})a - \sum_{i \in J} (X_i^{(0)} c_i^{(0)} + X_i^{(1)} c_i^{(1)})$$

recovers Eq. (15).

A.3 The two coordinate lower bounds

Lemma 4 (Local coordinate lower bound). *In the generic coordinate-linear polynomial comparer model, every nonzero local comparer has degree at least one in the denominator variables.*

Proof. Suppress all generators except the two fork numerators in one coordinate. A local kernel expression has coefficient vector (p, q) satisfying

$$\delta^{(0)}p + \delta^{(1)}q = 0. \quad (59)$$

The two denominator variables are coprime. Thus $p = \delta^{(1)}g$ and $q = -\delta^{(0)}g$ for some polynomial g . Every nonzero solution has coefficient degree at least one, and $g = 1$ gives the local comparer of Eq. (11).

Lemma 5 (Coordinate synchronization-defect lower bound). *Let the residual active set be J . In the generic coordinate-linear polynomial comparer model, every pure-baseline bridge that detects an arbitrary nonzero constant-coordinate mismatch has denominator-variable degree at least $|J|$.*

Proof. Let a be the coefficient multiplying $v_0 - v_1$. After the nonconstant normalized coordinates have been identified, cancellation in coordinate i requires a solution of

$$G_i c_i = a \varepsilon_i.$$

Thus $a \in \text{Ann}([\varepsilon_i]) = I_i$ for every $i \in J$. The bridge must have $a \neq 0$ to detect a free nonzero constant-coordinate mismatch. Equation (57) places a in a product ideal of initial denominator degree $|J|$. Hence the degree of the bridge is at least $|J|$. The oriented construction attains this bound.

Lemmas 4 and 5 concern different possible mismatches. A complete direct comparer system must handle both, so its worst-case degree is their maximum.

B Generic Block Lower Bounds

This appendix proves the algebraic statements used in Theorems 2 and 3. Let k be a field. Every matrix called generic has algebraically independent entries over k , and different blocks and branches use disjoint variable sets. For a matrix H , write $I_r(H)$ for its ideal of $r \times r$ minors.

B.1 Local coupling by generic matrix syzygies

Fix $r \geq 1$ and two generic matrices $D^{(0)}, D^{(1)} \in S^{r \times r}$ over their joint polynomial ring. Put

$$G = [D^{(0)} \mid D^{(1)}] \in S^{r \times 2r}. \quad (60)$$

A vector $z = (p, q) \in S^{2r}$ lies in $\ker G$ exactly when $D^{(0)}p + D^{(1)}q = 0$. The associated module-linear comparer is

$$\mathbf{x}^{(0)}p + \mathbf{x}^{(1)}q. \quad (61)$$

Lemma 6 (Degree and attainment for local block coupling). *For the generic $r \times 2r$ matrix G ,*

$$\text{indeg } \ker G = r. \quad (62)$$

There is a degree- r family whose induced comparers detect every nonzero normalized row mismatch on the unit chart.

Proof. Choose the first r columns as the square submatrix $D^{(0)}$. If $d_j^{(1)}$ is column j of $D^{(1)}$, then

$$p_j = -\text{adj}(D^{(0)})d_j^{(1)}, \quad q_j = \det(D^{(0)})e_j \quad (63)$$

satisfy $D^{(0)}p_j + D^{(1)}q_j = 0$. Every entry of (p_j, q_j) is homogeneous of denominator degree r . If $h = \det D^{(0)}$, the induced comparer is

$$K_j = x_j^{(1)}h - \mathbf{x}^{(0)} \text{adj}(D^{(0)})d_j^{(1)} = (\mathbf{u}^{(1)} - \mathbf{u}^{(0)})d_j^{(1)}h. \quad (64)$$

If all K_j vanish, then $(\mathbf{u}^{(1)} - \mathbf{u}^{(0)})D^{(1)}h = 0$; on the unit chart this forces $\mathbf{u}^{(0)} = \mathbf{u}^{(1)}$. The family therefore detects every nonzero normalized mismatch at degree r .

For the lower bound, let $I = I_r(G)$. The generic height formula gives $\text{ht } I = 2r - r + 1 = r + 1$, hence the grade condition for the Buchsbaum–Rim complex is sharp. That complex is exact, and its first map onto $\ker G$ consists of the circuit relations on $r + 1$ columns. Each circuit coefficient is a signed maximal minor, hence is homogeneous of degree r . Exactness therefore says that $\ker G$ is generated in degree r ; since a circuit is nonzero, its initial degree is exactly r . See [5,6] for these standard generic determinantal and Buchsbaum–Rim inputs.

The construction in Eq. (63) is one convenient unit-minor chart. Taking all $r + 1$ column circuits gives a chart-independent polynomial family.

To transfer the initial-degree statement to comparison, use the following explicit quotient test for a chosen block a . Set all centers to zero, kill the target rows and source generators outside block a , put $G_a = [D_a^{(0)} \mid D_a^{(1)}]$ (so the preceding lemma is used with $r = r_a$), and let

$$N_{\text{test}} = \widehat{St}_0 \oplus \bigoplus_{j=1}^{r_a} \widehat{St}_{a,j}, \quad M_{\text{test}} = N_{\text{test}} \oplus \bigoplus_{j=1}^{r_a} SW_j. \quad (65)$$

Define $\mathcal{A}_{\text{test}} : M_{\text{test}} \rightarrow N_{\text{test}}$ as projection onto the first summand, and write $\widehat{\mathbf{t}}_a = (\widehat{t}_{a,1}, \dots, \widehat{t}_{a,r_a})$ and $\mathbf{W} = (W_1, \dots, W_{r_a})$. The assignments

$$v_0 = v_1 = \widehat{t}_0, \quad \mathbf{u}_a^{(0)} = \widehat{\mathbf{t}}_a, \quad \mathbf{u}_a^{(1)} = \widehat{\mathbf{t}}_a + \mathbf{W}, \quad \mathbf{x}_a^{(b)} = \mathbf{u}_a^{(b)} D_a^{(b)} \quad (66)$$

satisfy the quotient anchor equations, while the entries of $\mathbf{W}$ span a nonzero free submodule of $\ker \mathcal{A}_{\text{test}}$. In this test, the baseline term and all other block terms vanish. Equation (26) becomes

$$D_a^{(0)}c_a^{(0)} + D_a^{(1)}c_a^{(1)} = 0, \quad (67)$$

whereas the evaluated certificate is $\mathbf{W}D_a^{(1)}c_a^{(1)}$. If the certificate detects the displayed normalized mismatch, this value is nonzero; in particular $(c_a^{(0)}, c_a^{(1)})$ is a nonzero element of $\ker G_a$. The lemma then forces coefficient degree at least r_a .

B.2 Mismatch classes

Let S be a commutative ring, let $D \in S^{r \times m}$, and let $\varepsilon \in S^r$. Write

$$C_D = \text{coker } D, \quad \mathfrak{a}(D, \varepsilon) := \text{Ann}_S([\varepsilon] \in C_D). \quad (68)$$

By definition, $a \in \mathfrak{a}(D, \varepsilon)$ exactly when there is $c \in S^m$ such that

$$Dc = a\varepsilon. \quad (69)$$

Fix a target row $T \in N^{1 \times r}$ and elements $B \in M$ and $X \in M^{1 \times m}$ satisfying

$$A(B) = T\varepsilon, \quad A^{(m)}(X) = TD. \quad (70)$$

Any solution of Eq. (69) gives

$$K = Ba - Xc \in \ker A. \quad (71)$$

Indeed, $A(K) = T(a\varepsilon - Dc) = 0$. If moreover $X = UD$ and $B = W + U\varepsilon$ for a row $U \in M^{1 \times r}$ and $W \in \ker A$, then

$$K = Wa. \quad (72)$$

When a is a unit and $W \neq 0$, Eq. (72) also gives $K \neq 0$. The oriented coordinate bridge chooses one such coefficient explicitly.

B.3 A generic mismatch annihilator

Let $m \geq r$, let G now denote a generic $r \times m$ matrix, and let $\varepsilon \in S^r$ be a generic column independent of its entries, where

$$S = k[G, \varepsilon].$$

Lemma 7 (Generic mismatch-annihilator equality). *For $m \geq r$ and a generic $r \times m$ matrix G , the class of the independent generic column ε in $\text{coker } G$ has annihilator*

$$\text{Ann}_S([\varepsilon] \in \text{coker } G) = I_r(G). \quad (73)$$

We use this statement with $m = 2r$ for the two-branch block matrix $[D^{(0)} \mid D^{(1)}]$, and with $m = r$ for a single square block in the branch-separated argument.

Proof. Let H be any $r \times r$ column submatrix of G and let $\mu_H = \det H$. Embedding $\text{adj}(H)\varepsilon$ in the coordinates belonging to H gives

$$Gc_H = H \text{adj}(H)\varepsilon = \mu_H \varepsilon.$$

Thus every maximal minor annihilates $[\varepsilon]$, proving $I_r(G) \subseteq \text{Ann}_S([\varepsilon])$.

For the reverse inclusion, write $P = k[G]$, $I = I_r(G)$, and $A = S/IS \simeq (P/I)[\varepsilon]$. The generic maximal-minor ideal I is prime, so A is a domain. Write $\overline{G}$ and $\overline{\varepsilon}$ for reduction modulo IS . Then $\overline{G}$ has rank $r - 1$ over $\text{Frac}(A)$ [5,6].

The augmented matrix $[\overline{G} \mid \overline{\varepsilon}]$ has rank r over that fraction field. Indeed, take a determinant using ε and $r - 1$ columns of G . The corresponding polynomial is not in IS : specialize those $r - 1$ columns to the first $r - 1$ standard basis vectors, all other columns of G to zero, and ε to the last standard basis vector. Every maximal minor of G then vanishes, while the chosen augmented determinant becomes 1. Consequently the class of $\overline{\varepsilon}$ in $\text{coker}(\overline{G}) \otimes_A \text{Frac}(A)$ is nonzero.

If a nonzero element of A annihilated this class, it would become a unit over $\text{Frac}(A)$ and force that nonzero generic-fiber class to vanish. Thus the class has zero annihilator modulo IS . The image in A of any element of $\text{Ann}_S([\varepsilon])$ is therefore zero, which proves $\text{Ann}_S([\varepsilon]) \subseteq IS = I_r(G)S$. Together with the first inclusion, this proves Eq. (73).

B.4 The synchronization-defect degree

Return to blocks $a \in J$. For each block, let

$$G_a = [D_a^{(0)} \mid D_a^{(1)}], \quad \varepsilon_a = s_a^{(0)} - s_a^{(1)}.$$

Lemma 7 gives

$$\mathfrak{a}(G_a, \varepsilon_a) = I_{r_a}(G_a). \quad (74)$$

Choose an $r_a \times r_a$ column submatrix H_a of G_a and put $\mu_a = \det H_a$. Define

$$\alpha = \prod_{a \in J} \mu_a. \quad (75)$$

For each $a \in J$, embed

$$c_a = \left(\prod_{b \in J \setminus \{a\}} \mu_b \right) \text{adj}(H_a) \varepsilon_a \quad (76)$$

in the coordinates of H_a . Then

$$G_a c_a = \alpha \varepsilon_a. \quad (77)$$

The multiplier α has degree $\sum_{a \in J} r_a$, while every entry of c_a has degree one less. If $\mathbf{X}_a = [\mathbf{x}_a^{(0)} \mid \mathbf{x}_a^{(1)}]$ and the two normalized rows in block a agree, then $\mathbf{X}_a = \mathbf{u}_a G_a$. Consequently

$$\begin{aligned} K &= (v_0 - v_1) \alpha - \sum_{a \in J} \mathbf{X}_a c_a \\ &= (u_0^{(0)} - u_0^{(1)}) \alpha. \end{aligned} \quad (78)$$

This is the block adjugate bridge used for the direct upper bound.

We will also use the following elementary disjoint-variable calculation. Let $R = k[X]$ and $T = k[Y]$ have disjoint variable sets, let $I \subseteq R$ and $L \subseteq T$, and put $Q = k[X, Y]$. Choose k -vector-space decompositions $R = I \oplus R'$ and $T = L \oplus T'$. The induced decomposition of $Q = R \otimes_k T$ shows

$$IQ \cap LQ = (I \otimes_k L) = ILQ. \quad (79)$$

Induction gives intersection equals product for any finite collection of ideals in pairwise disjoint variable sets. Extending by the degree-zero center variables preserves the identity.

Lemma 8 (Generic synchronization-defect lower bound). *In the generic block-linear polynomial module-linear comparer model, a pure-baseline bridge for active set J has denominator-variable degree at least $\sum_{a \in J} r_a$.*

Proof. To force the baseline coefficient to be visible, realize the universal module as a free module with an independent kernel direction W_0 . Write $\mathbf{X}_a = \mathbf{U}_a G_a$ and

$$v_0 - v_1 = W_0 + \sum_{a \in J} \mathbf{U}_a \varepsilon_a, \quad W_0 \in \ker \mathcal{A}. \quad (80)$$

After substituting the universal certificate equations (26), every pure-baseline certificate reduces to $K = \alpha W_0$, where α is the coefficient of $v_0 - v_1$. Detection of a nonzero free W_0 therefore requires $\alpha \neq 0$.

Image cancellation in block a requires $G_a c_a = -\alpha \varepsilon_a$ for some c_a . The sign does not affect ideal membership, so Eq. (74) gives

$$\alpha \in \bigcap_{a \in J} I_{r_a}(G_a). \quad (81)$$

The ideals are homogeneous and use disjoint sets of positive-degree denominator variables. Applying Eq. (79) successively gives

$$\bigcap_{a \in J} I_{r_a}(G_a) = \prod_{a \in J} I_{r_a}(G_a). \quad (82)$$

Each factor is generated in degree r_a . The product therefore has initial degree $\sum_{a \in J} r_a$, proving the lower bound. The multiplier in Eq. (75) attains it.

Lemma 6 supplies the independent lower bound r_a for local coupling in block a . Since a complete direct system must handle both a local normalized mismatch and a pure constant-coordinate mismatch, its degree is at least

$$\max \left\{ \max_a r_a, \sum_{a \in J} r_a \right\}.$$

The two constructions above give equality.

B.5 Branch-separated integralization

We finally characterize the branch-local multiplier. Fix one branch and write D_a and s_a for its block denominator matrix and center. A branch-local polynomial integralizer has the form

$$n = qv + \sum_{a=1}^m \mathbf{x}_a c_a, \quad \mathcal{A}(n) = t_0 q, \quad (83)$$

with all coefficients in that branch's polynomial subring. Comparing the row $\mathbf{t}_a$ gives

$$q s_a + D_a c_a = 0. \quad (84)$$

Thus q annihilates the class of the generic vector s_a in coker D_a . Lemma 7, in the square-matrix specialization, yields

$$\text{Ann}([s_a] \in \text{coker } D_a) = I_{r_a}(D_a) = (\det D_a). \quad (85)$$

It follows that

$$P = \prod_{a=1}^m \det D_a \mid q, \quad \deg q \geq \sum_{a=1}^m r_a. \quad (86)$$

The lower bound is attained. Define

$$n_P = vP - \sum_{a=1}^m \mathbf{x}_a \text{adj}(D_a) s_a \prod_{c \neq a} \det D_c. \quad (87)$$

Since $D_a \text{adj}(D_a) = (\det D_a) I_{r_a}$,

$$\mathcal{A}(n_P) = t_0 P. \quad (88)$$

The baseline coefficient has degree $\sum_a r_a$ and every fork-numerator coefficient has degree $\sum_a r_a - 1$.

For two branches, let P_0, P_1 be the two determinant products. Their variable sets are disjoint. If branch-local integralizers with image $t_0 q_b$ are eliminated through a nonzero common multiplier

$$C_0 q_0 = C_1 q_1 = Q, \quad (89)$$

then $P_0 P_1 \mid Q$. Therefore

$$\deg Q \geq 2 \sum_{a=1}^m r_a. \quad (90)$$

Conversely,

$$K_{\text{branch}} = n_{P_0}^{(0)} P_1 - n_{P_1}^{(1)} P_0 \quad (91)$$

lies in $\ker \mathcal{A}$ and has common-multiplier degree $2 \sum_a r_a$. This proves the exact frontier within the branch-separated polynomial integralize-then-compare architecture.

Support specialization. The equalities above characterize universal polynomial module-linear comparers over generic denominator variables. They transfer to a concrete finite challenge support when evaluation is faithful on the relevant bounded-degree certificate relations. Additional support identities, structured-matrix syzygies, or postchallenge proof objects define a different grading or comparison boundary.

C Conditional-Fiber Probability Proofs

This appendix proves the probability and retry statements of Section 4. All bounds are taken in the raw experiment before rejection selection. This order is what allows a success-selected first branch to pay for later retries without a reciprocal loss in the extraction probability.

C.1 Conditional laws and zero-success fibers

Let $\lambda = \pi_{\#}\mu$. A regular conditional law μ_z satisfies

$$\mu(E) = \int_{\mathcal{Z}} \mu_z(E) \lambda(dz), \quad \mu_z(\pi^{-1}(z)) = 1$$

for every measurable $E \subseteq \mathcal{S}$, for λ -almost every z . The raw trial context u includes a draw from μ_z and every fresh continuation coin sampled before the final selected-success test. We denote its conditional law by ρ_z .

The functions w and q are measurable and obey $0 \leq w \leq q \leq 1$. Consequently $0 \leq \bar{w}(z) \leq \bar{q}(z) \leq 1$. On a fiber with $\bar{q}(z) = 0$, also $\bar{w}(z) = 0$, so the retry node is entered with zero raw mass. Every ratio below is set to zero on these fibers.

C.2 Entry-weighted cancellation and selection

Proof (Proof of Lemma 3). At a label with $\bar{q}(z) > 0$, independent raw trials until selected success have geometric mean $1/\bar{q}(z)$. The raw entry mass at that label is $\bar{w}(z)$. Therefore the unconditional entry-weighted number of trial invocations is

$$\int_{\{\bar{q}>0\}} \frac{\bar{w}(z)}{\bar{q}(z)} \lambda(dz).$$

Since $\bar{w}(z) \leq \bar{q}(z)$ pointwise, the integrand is at most one. Integration against the probability measure λ proves Eq. (36).

At fixed z with $\bar{q}(z) > 0$, the context selected by rejection sampling has density $q(z, u)/\bar{q}(z)$ with respect to $\rho_z(du)$. Multiplying by the raw entry mass $\bar{w}(z)$ gives

$$\tilde{w}(z, u) = q(z, u) \frac{\bar{w}(z)}{\bar{q}(z)}.$$

The inequality $\bar{w} \leq \bar{q}$ proves $\tilde{w}(z, u) \leq q(z, u)$ and hence Eq. (37).

Lemma 9 (Checkpoint pushforward domination). *Let $Y \leftarrow \nu$ be a raw state on a standard Borel space, and let the entered-state and raw selected-success submeasures have densities w_Y and q_Y with respect to ν , where $0 \leq w_Y \leq q_Y$. For a measurable checkpoint map $G = \psi(Y)$, put $\nu_G = \psi_{\#}\nu$ and let w_G, q_G be the densities of the two pushforward submeasures. If ν_g is a regular conditional law of Y given $G = g$, then*

$$w_G(g) = \int w_Y(y) \nu_g(dy) \leq \int q_Y(y) \nu_g(dy) = q_G(g) \quad \text{for } \nu_G\text{-a.e. } g. \quad (92)$$

Proof (Proof of Lemma 9). Put

$$W_Y(dy) = w_Y(y) \nu(dy), \quad Q_Y(dy) = q_Y(y) \nu(dy).$$

For every measurable A in the checkpoint space, pushforward monotonicity gives

$$(\psi_{\#}W_Y)(A) = \int_{\psi^{-1}(A)} w_Y d\nu \leq \int_{\psi^{-1}(A)} q_Y d\nu = (\psi_{\#}Q_Y)(A).$$

Both pushforwards are absolutely continuous with respect to $\nu_G = \psi_{\#}\nu$. Disintegrating ν over G identifies their Radon–Nikodym densities as

$$w_G(g) = \int w_Y(y) \nu_g(dy), \quad q_G(g) = \int q_Y(y) \nu_g(dy).$$

Conditional-integration monotonicity yields $w_G \leq q_G$ for ν_G -almost every g , proving Eq. (92).

Lemma 10 (Raw bad-mass preservation). *Let $b(z, u)$ be the raw joint probability density of selected success and a bad event after (z, u) is fixed, so $0 \leq b(z, u) \leq q(z, u)$. A retry entered with density w has selected bad mass at most the corresponding raw bad mass:*

$$\int_{\mathcal{Z}} \int b(z, u) \frac{\bar{w}(z)}{\bar{q}(z)} \rho_z(du) \lambda(dz) \leq \int_{\mathcal{Z}} \int b(z, u) \rho_z(du) \lambda(dz). \quad (93)$$

Proof. On $\{\bar{q} > 0\}$, the multiplier $\bar{w}/\bar{q}$ lies in $[0, 1]$. On $\{\bar{q} = 0\}$, both sides contributed by the retry are zero. Integrate the pointwise inequality.

Equation (93) applies equally to a collection of bad events; their raw bounds add by a union bound. In particular, conditioning on selected success is never applied to a unit-failure or validation estimate.

C.3 Nested retry layers

The selected density from one retry satisfies the same domination condition as the entry density of the next. More formally, Eq. (37) first gives domination on the complete parent state Y . If a child retains checkpoint data $G = \psi(Y)$, Lemma 9 gives $w_G \leq q_G$. Exact checkpoint restoration identifies the conditional law given G with the child’s raw trial law, so Lemma 3 charges at most one invocation at that child. Integrating over the parent history is the tower property.

Lemma 11 (Retry-tree accounting). *Consider a finite or polynomial-size adaptive retry tree. At each retry node, the loop-entry density is dominated by the raw selected-success density for the exact conditional trial law used at that node. Then every node contributes at most one unconditional expected trial invocation. If a trial together with its extractor-side processing has worst-case cost at most T_{trial} , the corresponding expected running-time contribution is at most T_{trial} .*

Proof. Condition on the complete reachable parent history. Lemma 3 gives a unit invocation charge. The tower property averages this charge over the parent distribution. Summing over the nodes proves the claim; multiplying each charge by its worst-case trial cost gives the running-time statement.

The domination premise matters at the raw selected-success level. If a retry were instead required to satisfy both selected success and a unit difference, entry mass q against success-and-unit probability g would leave the uncontrolled ratio q/g . The extraction schedule therefore retries on **Succ** alone and charges nonunits through Lemma 10.

C.4 Proof of the two-star theorem

Proof (Proof of Theorem 4). Run one genuine target invocation. On **TgtSucc**, the pointwise event inclusion supplies a selected first baseline, including its exact affine row and terminal shortness bound. Save its label $z = \pi(S)$ and its branch-local pre-final checkpoint.

For each coordinate i , restore the first branch-local checkpoint and sample the modified coordinate from its real raw coordinate law. Repeat only until **Succ**. For G_i equal to the retained branch-local state and coordinates other than i , pointwise inclusion $\text{TgtSucc} \subseteq \text{Succ}$ and Lemma 9 give $w_{G_i} \leq q_{G_i}$. The checkpoint interface identifies the conditional continuation with the raw coordinate-fork law. Lemma 3 therefore charges one invocation for each of the L forks. Lemma 10 charges their nonunit events by $\sum_i \kappa_i$.

Next restore the outer checkpoint. Sample the second center exactly from $\mu(\cdot \mid \pi(S') = z)$, sample all fresh continuation coins from their raw conditional law, and repeat until **Succ**. The genuine target event is pointwise contained in selected success, so its entry density obeys the $w \leq q$ premise. Lemma 3 charges this second baseline loop once. Equation (37) shows that the selected second branch is dominated on its complete outer state. Applying Lemma 9 to each retained G_i , then using the exact checkpoint law, supplies the coordinate-fork domination. Its L forks therefore contribute another L invocation charges and another $\sum_i \kappa_i$ of raw nonunit mass.

Outside the total raw nonunit mass $2\sum_i \kappa_i$, the two selected stars satisfy the deterministic comparison hypotheses. A mismatch returns its nonzero kernel vector. On equality, the normalized tuple determined by the first branch is fixed before the fresh validation coins of the selected second branch are sampled. Intersecting the raw validation event with tuple equality can only decrease its mass, and Lemma 10 preserves the bound through rejection selection. The all-equal bad mass is therefore at most κ_{local} .

Subtracting these raw bad events from the initial target mass ϵ gives Eq. (38). The retry nodes contribute

$$\underbrace{L}_{\text{first star}} + \underbrace{1}_{\text{conditional-fiber baseline}} + \underbrace{L}_{\text{second star}} = 2L + 1,$$

which proves Eq. (39). Lemma 11 and the assumed polynomial trial costs give expected polynomial time.

C.5 Coordinate-product specialization

Let $\mu = \mathcal{D}^L$ and $\pi_J(s) = s_{[L] \setminus J}$. For every supported retained coordinate vector z ,

$$\mu(ds' \mid \pi_J(S') = z) = \delta_z(ds'_{[L] \setminus J}) \otimes \mathcal{D}^J(ds'_J). \quad (94)$$

Indeed, independence factors the probability of every cylinder event into the fixed-coordinate mass and the product mass on J ; division by the positive fixed-coordinate mass gives Eq. (94). For $J = \emptyset$, the right-hand side is the point mass at the stored full challenge. This proves the two specializations stated in Section 4.

D Cyclo Interface Verification

This appendix verifies the protocol facts used by Theorem 5. The proof follows the message chronology: exact-capacity Extension first fixes the affine anchor, Range and the batched sum-check then create a branch-local validation context, and the final challenge is sampled only afterward.

D.1 Ring, relation, and checkpoints

Let

$$R = \mathbb{Z}[X]/(X^{128} + 1), \quad R_q = R/qR,$$

and identify ring elements with their canonical centered coefficient vectors. The norm $\|v\|_\infty$ is the largest absolute centered coefficient, and $\|s\|_{\text{op}}$ is the induced coefficient-infinity multiplication norm.

Cyclo’s principal relation stacks an Ajtai commitment row, matrix-evaluation rows, and the additional evaluation rows of its Range interface [3, Eq. (6), pp. 16–17]. In Figure 3 of Cyclo, every Step–1 Extension invocation terminates before Step–2 Range begins. Step 3 samples the batching challenge, Step 4 completes the shared sum-check, and Step 5 samples the final folding vector S before the affine output in Step 6 [3, p. 20]. We use two checkpoints:

- H_{preRange} , immediately after all Extension invocations and before the first Range verifier coin; and
- H_{final} , after Range, batching, and sum-check data are fixed and immediately before the final vector challenge.

At H_{preRange} , $\mathbf{R}$, the old image t_0 , the fresh images $t_1, \dots, t_L$, and the deterministic Extension rows are fixed. Range, batching, sum-check, and final-vector coins remain fresh. At H_{final} , the branch-local rows, evaluation points, and validation predicate are fixed, while $S \leftarrow \mathcal{D}^L$ remains fresh.

Lemma 12 (Cyclo checkpoint interface). *Restoring H_{preRange} reinstates the prover state and random tape at the common affine anchor and permits a fresh raw Range and Step-3/4 continuation followed by a verifier-chosen final vector challenge. Every selected continuation reaches its own H_{final} . Restoring that branch-local checkpoint keeps its context and S_{-i} fixed and permits a fresh draw of coordinate S_i from $\mathcal{D}$.*

Proof. The common data are fixed by the completed Extension invocations. The final vector is first sampled in Step 5, so restoring either checkpoint precedes that verifier message. At H_{preRange} the intervening Range and Step-3/4 coins are sampled from their real raw law. At H_{final} they are already fixed, and only the designated coordinate of the final vector is changed. This is the coordinate-wise fork used in Cyclo’s extraction [3, Lem. 4, p. 21].

For a fixed history h , final vector s , and continuation coins C , let $V(h, s; C)$ be the prescribed terminal output. The selected event is

$$\begin{aligned} \text{RelSucc}(h, s; C) &:= \text{Acc}(h, s; C) \\ &\wedge [(x_e(h, s; C), V(h, s; C)) \in \Xi_1] \wedge [\|V(h, s; C)\|_\infty \leq \beta]. \end{aligned} \quad (95)$$

Membership in the terminal target relation supplies

$$\mathbf{R}V(h, s; C) = t_0 + \sum_{i=1}^L t_i s_i, \quad \|V(h, s; C)\|_\infty \leq \beta. \quad (96)$$

The event is publicly decidable from the transcript, relation, and prescribed terminal output. Every baseline and fork used below is selected by its own instance of Eq. (95).

D.2 Exact-capacity Extension

Let $c, B \geq 1$, assume $c \mid B$ and $2B < q$, and choose positive integer weights $g = (g_0, \dots, g_{\ell-1})$ satisfying

$$g_0 = 1, \quad g_j \leq 2c \sum_{r < j} g_r + 1, \quad c \sum_{j=0}^{\ell-1} g_j = B. \quad (97)$$

Define

$$P_g = (g_0 I \mid \dots \mid g_{\ell-1} I), \quad P_g v = \sum_j g_j v_j.$$

Each deterministic source row L_0 is regenerated as $L_0 P_g$. The independent Ajtai commitment matrix and the primary witness dimensions remain unchanged.

Lemma 13 (Exact-capacity signed-digit image). *Under Eq. (97),*

$$\left\{ \sum_{j=0}^{\ell-1} g_j d_j : d_j \in [-c, c] \cap \mathbb{Z} \right\} = [-B, B] \cap \mathbb{Z}. \quad (98)$$

Hence every source coefficient of magnitude at most B has a legal radius- c digit representation, and every legal digit vector reconstructs inside the centered radius B .

Proof. Put $M_{-1} = 0$ and $M_j = c \sum_{r=0}^j g_r$. The one-digit sums form $[-M_0, M_0]$. Assuming that the first j digit positions fill $[-M_{j-1}, M_{j-1}]$, the intervals obtained from two adjacent values of the new digit meet without an integer gap exactly when

$$g_j \leq 2M_{j-1} + 1.$$

Their union is therefore $[-M_j, M_j]$. Induction and $M_{\ell-1} = B$ prove Eq. (98). The condition $2B < q$ identifies these integers with their unique centered representatives.

The benchmark has

$$c = b = 1, \quad B = 1024, \quad g = (1, 1, 2, 4, 8, 16, 32, 64, 128, 256, 512).$$

The weight sum is 1024, and every prefix satisfies Eq. (97).

Let $\mathcal{C}$ be a strong Extension challenge set: for a transcript fixed before $Z \leftarrow \mathcal{C}^{\ell_C}$, every nonzero compressed polynomial of total degree at most ℓ_C has root probability at most $\ell_C/|\mathcal{C}|$.

Theorem 8 (Exact-capacity Extension extraction). *The regenerated Extension relation returns an original radius- B source witness or a nonzero $\mathbf{R}$ -kernel vector of coefficient radius at most $2c$, except with raw mass*

$$\kappa_{\text{dec}} = \frac{\ell_C}{|\mathcal{C}|}. \quad (99)$$

Proof. Let v^* be the first selected opening, of radius c , and set $w^* = P_g v^*$. Lemma 13 gives $\|w^*\|_\infty \leq B$, and the regenerated rows recover every deterministic source equation. If the randomized source residual vanishes, w^* is an original witness. Otherwise rewind to the post-commitment, pre-challenge checkpoint and obtain a second selected opening v' . Unequal openings give $v^* - v' \in \ker \mathbf{R}$ with radius at most $2c$. Equal openings leave a fixed nonzero residual before the fresh Extension challenge. Cyclo's dual-basis compression represents a nonzero multilinear polynomial of total degree at most ℓ_C , so the strong challenge-set assumption bounds its raw root mass by Eq. (99). Entry-weighted cancellation preserves this raw bound through the retry.

D.3 Branch-local affine row reconstruction

Fix one completed Range and Step-3/4 context ρ . For every displayed final row r , let $\mathcal{L}_{\rho,r}$ be its branch-fixed linear map and write the Step-6 public image as

$$Y_{\rho,r}(s) = Y_{\rho,0,r} + \sum_{i=1}^L Y_{\rho,i,r} s_i.$$

Lemma 14 (Branch-local public-row reconstruction). *Suppose a baseline and all of its final-coordinate forks are selected by RelSucc within the fixed context ρ . Let*

$$X_i = V^{(i)} - V, \quad \Delta_i = S'_i - S_i, \quad U_i = X_i \Delta_i^{-1}, \quad U_0 = V - \sum_i U_i S_i,$$

with every Δ_i a unit. Then

$$\mathcal{L}_{\rho,r} U_i = Y_{\rho,i,r}, \quad \mathcal{L}_{\rho,r} U_0 = Y_{\rho,0,r}$$

for the commitment row, all regenerated Extension evaluation rows, the post-batching Range row, and the terminal/evaluation-claim rows. On validity of the intermediate Range leaves, the same tuple opens those rows as well.

Proof. For every displayed row, Eq. (96) and final target membership give the baseline and modified-challenge equations. Subtracting them yields

$$\mathcal{L}_{\rho,r} X_i = Y_{\rho,i,r} \Delta_i.$$

Multiplication by a unit is a module automorphism, giving the U_i equation; substitution into the definition of U_0 gives the remaining equation. The argument is applied separately in each branch-local context.

In particular, the common commitment row has

$$\mathbf{R} X_i = t_i \Delta_i, \quad \|X_i\|_\infty \leq 2\beta.$$

The challenge bound gives $\|\Delta_i\|_{\text{op}} \leq 2\gamma$.

D.4 Fresh Range and sum-check validation

Write $Q_{\text{fld}} = q^e$. For a fixed normalized tuple U , the Range polynomial for component i is

$$f_{i,U}(X) = \text{MLE}[\text{cf}(U_i)](X), \quad \widehat{f}_{i,U}(X; \eta_i) = \text{eq}(X; \eta_i) \prod_{j=-b}^b (f_{i,U}(X) - j). \quad (100)$$

The tuple and the complete reachable Range prefix are fixed before the fresh point η_i and its subsequent sum-check coins.

Lemma 15 (Raw Range validation). *Assume $q > 2b$. If $\|U_i\|_\infty > b$, then the raw joint probability that the actual Range transcript accepts and its candidate leaf row opens to U_i is at most*

$$\kappa_{\text{short}} = \frac{\ell_1(2b+3)}{Q_{\text{fld}}}. \quad (101)$$

Proof. If a centered coefficient is outside $[-b, b]$, the fixed polynomial in Eq. (100) is nonzero at a Boolean table point. The first Range stage tests its multilinear extension and contributes ℓ_1/Q_{fld} . On the complement, the fixed-candidate Range sum-check has individual degree $2b+2$ and contributes $\ell_1(2b+2)/Q_{\text{fld}}$. These disjoint cases give Eq. (101); this is Cyclo’s fixed-witness Range hybrid applied to the reconstructed public candidate [3, pp. 17–18].

The batch in Cyclo Steps 3 and 4 contains four typed claim families. After the NTT/CRT map, let α index the scalar slots and let $\chi_\alpha(d) = \text{eq}(\iota(\alpha); d)$ be the public selector. For a fixed tuple, the difference between its claimed slots and the public slots is a vector $(\delta_\alpha)_\alpha$ fixed before the fresh batching point d . The discrepancy polynomial

$$D_U(d) = \sum_{\alpha} \chi_\alpha(d) \delta_\alpha$$

has total degree at most ℓ_0 . When the vector is nonzero, faithful NTT scalarization and selector padding make D_U nonzero.

Lemma 16 (Raw fixed-tuple batching and sum-check validation). *Fix the reachable pre-Step-3 history and a normalized tuple U before the fresh Step-3/4 coins. If U violates an intermediate Range leaf or an intended pre-sum-check source row, the raw joint mass of acceptance together with the reconstructed final endpoint rows is at most*

$$\kappa_{\text{chk}} = \frac{\ell_0 + 2\ell_1}{Q_{\text{fld}}}. \quad (102)$$

Proof. The selector collision $D_U(d) = 0$ has probability at most ℓ_0/Q_{fld} . On its complement, the aggregate initial claim is false. The fixed-candidate Step-4 polynomial is a sum of products of multilinear extensions and equality polynomials, so its individual degree is at most two. The first-difference sum-check hybrid contributes $2\ell_1/Q_{\text{fld}}$. Adding the two disjoint contributions proves Eq. (102).

D.5 Fresh-validation handoff

Complete the first outer branch and all of its coordinate forks before restoring H_{preRange} . Its terminal-derived tuple $U^{(0)}$ and the retained fiber label are then fixed before the second branch samples Range or Step-3/4 randomness. On the all-equal comparison branch, Lemma 14 identifies the second branch’s reconstructed endpoint with this fixed tuple. The validation estimates are pointwise in the retained center: for every supported $z = \pi_J(S')$, conditioning on that value only fixes the later coordinates $S'_{[L] \setminus J}$; the coordinates S'_J keep their raw

product law. Every Range and Step-3/4 coin also keeps its raw law, and the proofs of Lemmas 15 and 16 are uniform in the fixed tuple and reachable prefix. Hence the same constants bound the raw joint bad mass for each z separately; no averaging over centers or reciprocal fiber-mass factor is used. Exposing the Range copies in protocol order gives total raw mass

$$L\kappa_{\text{short}} + \kappa_{\text{chk}}.$$

On its complement, all fresh coordinates have norm at most b . With the aligned cap $c = b$, the reconstructed tuple satisfies the public input of the exact-capacity Extension interface. Applying Theorem 8 to the L fresh components adds $L\kappa_{\text{dec}}$ and returns original source witnesses or the radius- $2c$ kernel branch.

Consequently the branch-local term used in Eq. (42) is

$$\kappa_{\text{local}} = \kappa_{\text{chk}} + L\kappa_{\text{short}} + L\kappa_{\text{dec}}.$$

Every term is a raw joint-event bound applied before conditional-fiber retry selection. Lemma 10 therefore transports the sum through the selected second branch without a reciprocal success loss.

E Cyclo Challenge and Support Certificates

This appendix supplies the finite support facts used by the Cyclo phase diagram: pairwise unit differences, an exact uniform sampler, exact conditional sampling for coordinate fibers, and the arithmetic behind the centered-gate classifications.

E.1 Pairwise unit differences

Let

$$q_{\text{exact}} = 2^{50} - 351, \quad R_{q_{\text{exact}}} = \mathbb{F}_{q_{\text{exact}}}[X]/(X^{128} + 1).$$

The integer q_{exact} is prime,

$$q_{\text{exact}} \equiv 161 \pmod{256}, \quad \text{ord}_{256}(q_{\text{exact}}) = 8,$$

so $X^{128} + 1$ has sixteen irreducible factors of degree eight over $\mathbb{F}_{q_{\text{exact}}}$. The power-of-two short-unit criterion [27], with splitting parameter 16, gives

$$\frac{q_{\text{exact}}^{1/16}}{4} = 2.181015\dots > 2. \tag{103}$$

Every nonzero element with coefficient infinity norm at most 2 is therefore a unit. In particular, every nonzero difference of two ternary coefficient vectors is a unit.

For $0 \leq \tau \leq 128$, define

$$\mathcal{D}_\tau = \{s \in \{-1, 0, 1\}^{128} : \text{wt}(s) = \tau\}, \quad N_\tau = \binom{128}{\tau} 2^\tau.$$

Theorem 9 (Exact-strong fixed-weight support). *For distinct $s, s' \in \mathcal{D}_\tau$, the difference $s' - s$ is a unit in $R_{q_{\text{exact}}}$. Uniform sampling from $\mathcal{D}_\tau$ has point mass N_τ^{-1} , and negacyclic multiplication satisfies*

$$\|M_s\|_{\infty \rightarrow \infty} = \sum_j |s_j| = \tau. \quad (104)$$

Consequently, for a fixed supported s and an independent $S' \leftarrow \mathcal{D}_\tau$,

$$\mathbb{P}[S' - s \notin R_{q_{\text{exact}}}^\times] = \mathbb{P}[S' = s] = N_\tau^{-1}. \quad (105)$$

Proof. The support size follows by choosing the nonzero positions and their signs. A distinct ternary difference is nonzero and has coefficient norm at most 2, so Eq. (103) makes it a unit. Every row of the negacyclic convolution matrix is a signed permutation of the coefficient vector of s ; its absolute row sum is therefore τ , proving Eq. (104). Equality is the only nonunit difference between two supported samples, which gives Eq. (105).

For $\tau = 32$,

$$\log_2 N_{32} = 132.221300637359\dots, \quad \|M_s\|_{\infty \rightarrow \infty} = 32.$$

The two-star pointwise unit term $2L/N_{32}$ has inverse-bit values

L	1	2	4	8
$-\log_2(2L/N_{32})$	131.2213	130.2213	129.2213	128.2213.

E.2 Exact sampling and coordinate conditioning

An exact uniform sampler for $\mathcal{D}_\tau$ first samples an ordered τ -tuple of distinct positions by partial Fisher–Yates sampling, using exact rejection before each modular reduction, and then samples τ independent sign bits. Every ordered position tuple has probability $1/(128)_\tau$. A fixed signed support has $\tau!$ orderings, hence probability

$$\frac{\tau!}{(128)_\tau} 2^{-\tau} = \frac{1}{\binom{128}{\tau} 2^\tau}.$$

Thus the sampler is exactly uniform on $\mathcal{D}_\tau$.

Lemma 17 (Exact coordinate-fiber sampler). *Let $S = (S_1, \dots, S_L) \leftarrow \mathcal{D}_\tau^L$, let $J \subseteq [L]$, and let z be a supported value of $\pi_J(S) = S_{[L] \setminus J}$. The following procedure samples exactly from the real conditional law $\mu(\cdot \mid \pi_J(S') = z)$:*

1. set $S'_{[L] \setminus J} = z$; and
2. sample the coordinates $S'_j \leftarrow \mathcal{D}_\tau$ independently for $j \in J$ using the exact sampler above.

Proof. For any s' with $s'_{[L]\setminus J} = z$, product independence gives

$$\mathbb{P}[S = s' \mid S_{[L]\setminus J} = z] = \prod_{j \in J} \mathbb{P}[S_j = s'_j].$$

The right-hand side is precisely the probability assigned by the stated procedure. Values outside the fiber receive probability zero.

The case $J = \emptyset$ is the point mass at the stored full challenge. The case $J = [L]$ is a fresh draw from the raw product law.

E.3 Maximal ideals and fixed-support minimax

In a finite commutative ring, an element is a unit exactly when it lies in no maximal ideal. Hence a support $\mathcal{T} \subseteq R_q$ has pairwise unit differences exactly when the quotient map

$$\mathcal{T} \longrightarrow R_q/\mathfrak{m}$$

is injective for every maximal ideal $\mathfrak{m}$. In particular, if $R_q/\mathfrak{m} \cong \mathbb{F}_{q^f}$, then

$$|\mathcal{T}| \leq q^f. \quad (106)$$

Lemma 18 (Support minimax). *Let $\mathcal{T}$ be a finite pairwise-unit-difference support and let P be any distribution on $\mathcal{T}$. Its worst-case pointwise nonunit loss is*

$$\sup_{s \in \mathcal{T}} \Pr_{S' \leftarrow P}[S' - s \notin R_q^\times] = \|P\|_\infty \geq \frac{1}{|\mathcal{T}|}.$$

Uniform sampling attains equality and therefore minimizes the worst-case pointwise loss among distributions with support $\mathcal{T}$.

Proof. Pairwise unit differences make equality the only nonunit event at every supported baseline. The loss is therefore $P(s)$ at baseline s , and its supremum is the largest point mass. The largest of $|\mathcal{T}|$ masses summing to one is at least $1/|\mathcal{T}|$, with equality for the uniform distribution.

For the ambient norm ball, let

$$\mathcal{B}_\Gamma = \{s \in \{-1, 0, 1\}^{128} : \text{wt}(s) \leq \Gamma\}.$$

By Eq. (104), every ternary distribution of worst-case multiplication norm at most Γ is supported in $\mathcal{B}_\Gamma$. For the two-star $L \leq 8$ ledger at a 2^{-128} pointwise-unit target, the largest point mass must be at most 2^{-132} . Direct counting gives

$$\log_2 |\mathcal{B}_{31}| = 129.867682452 \dots < 132, \quad \log_2 |\mathcal{B}_{32}| = 132.479101466 \dots > 132.$$

Thus norm 31 cannot meet this fixed ledger, while uniform sampling from $\mathcal{B}_{32}$ does. Within the uniform exact-weight family,

$$\log_2 N_{31} = 129.621387795 \dots < 132 < 132.221300637 \dots = \log_2 N_{32},$$

so weight 32 is the smallest exact weight meeting the same $L \leq 8$ target. These support optimizations are separate from the conservative $\gamma = 128$ terminal-budget convention used in Section 5.

Table 2. Formula-derived Cyclo radii on the conservative parameter line. An outside entry means that the displayed safe upper bound does not certify centered uniqueness.

L	h	β_L	$B_{\text{Cyclo}}(h)$	centered gate
2	0	33,792	34,603,008	inside
2	1	33,792	34,603,008	inside
2	2	33,792	13,287,555,072	inside
4	0	66,560	68,157,440	inside
4	1	66,560	68,157,440	inside
4	2	66,560	26,172,456,960	inside
4	3	66,560	8,933,531,975,680	inside
4	4	66,560	2,858,730,232,217,600	outside
8	0	132,096	135,266,304	inside
8	1	132,096	135,266,304	inside
8	2	132,096	51,942,260,736	inside
8	3	132,096	17,729,624,997,888	inside
8	4	132,096	5,673,479,999,324,160	outside
8	5	132,096	1,742,893,055,792,381,952	outside
8	6	132,096	520,544,059,329,991,409,664	outside
8	7	132,096	152,296,319,072,546,058,141,696	outside
8	8	132,096	43,861,339,892,893,264,744,808,448	outside

E.4 Phase-diagram arithmetic

The conservative parameter line has

$$\beta_L = 1024 + 16384L, \quad B_{\text{Cyclo}}(h) = \max\{1024\beta_L, 2(h+1)\beta_L 256^h, 2\}.$$

The centered threshold is

$$\frac{q_{\text{exact}} - 1}{2} = 562949953421136, \quad \log_2 \frac{q_{\text{exact}} - 1}{2} = 48.99999999999955 \dots$$

The exact integer radii are listed below.

At $L = 2$, the two formula-derived reference values are

$$\begin{aligned} \log_2 B_{\text{Cyclo}}(0) &= 25.04439411935845 \dots, \\ \log_2 B_{\text{branch}} &= 49.62935662007961 \dots \end{aligned}$$

They give the comparison $25.0444 < 49 < 49.6294$.

F Vandermonde Comparer Algebra

This appendix proves the degree statements of Theorems 6 and 7. All degree claims are universal polynomial claims. The lower-bound model uses independent challenge variables; the actual-support analysis in Appendix G instead derives safe upper bounds after specialization.

F.1 Polynomial comparer model

Fix a field $\mathbb{F}$, let $n = k + 1$, and put

$$S = \mathbb{F}[x_0, \dots, x_k, y_0, \dots, y_k].$$

Each difference of two challenge variables has degree one. The degree of a polynomial coefficient vector is the maximum total degree of its nonzero coordinates. Consider the evaluation matrix

$$E_n(z_1, \dots, z_{2n}) = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ z_1 & z_2 & \cdots & z_{2n} \\ \vdots & \vdots & & \vdots \\ z_1^{n-1} & z_2^{n-1} & \cdots & z_{2n}^{n-1} \end{pmatrix}. \quad (107)$$

The first and last n columns represent the two stars; changing the sign of the second group does not change kernel degrees.

An anchor-universal polynomial-linear comparer is a nonzero vector $a \in S^{2n}$ such that $E_n a = 0$. Indeed, applying a row-opening combination with coefficient vector a to Eq. (48) must vanish for every common anchor exactly when its coefficient vector annihilates all monomials $1, z, \dots, z^{n-1}$.

F.2 The direct circuit and its factorization

Let $V_0 = V(x_0, \dots, x_k)$ and $\delta_0 = \det(V_0)$. For an additional variable y , set

$$\lambda_j(y) = [r(y) \operatorname{adj}(V_0)]_j.$$

In the fraction field of S ,

$$r(y)V_0^{-1} = (L_0(y), \dots, L_k(y)), \quad L_j(y) = \prod_{\ell \neq j} \frac{y - x_\ell}{x_j - x_\ell}.$$

Multiplying by δ_0 and collecting the factors not incident to j gives

$$\lambda_j(y) = (-1)^{k+j} \left(\prod_{\substack{0 \leq a < b \leq k \\ a, b \neq j}} (x_b - x_a) \right) \left(\prod_{\ell \neq j} (y - x_\ell) \right). \quad (108)$$

Both sides are polynomials, so Eq. (108) is a polynomial identity. Its proof in the fraction field imposes no inverse or unit requirement on a later specialization.

The first product has $k(k-1)/2$ factors and the second has k , while δ_0 has $\binom{k+1}{2}$ factors by the Vandermonde determinant formula [4, Facts 1–2]. Therefore the coefficient vector

$$(-\lambda_0(y), \dots, -\lambda_k(y), \delta_0)$$

is an $(n+1)$ -column circuit of coefficient degree $\kappa = \binom{k+1}{2}$. Substitution of $y = y_i$ is precisely the direct comparer K_i in Eq. (49).

F.3 Universal evaluation-kernel lower bound

Let $I_n(E_n)$ be the ideal of maximal minors of E_n . For a subset $J \subseteq [2n]$ of size n , its minor is

$$\Delta_J = \prod_{\substack{a < b \\ a, b \in J}} (z_b - z_a), \quad \deg \Delta_J = \binom{n}{2}.$$

Over an algebraic closure of $\mathbb{F}$, all maximal minors vanish exactly when the $2n$ coordinates z_i assume at most $n - 1$ distinct values. The rank-drop locus is therefore a finite union of coordinate-equality subspaces. A component with $q \leq n - 1$ value classes has dimension q , and components of dimension $n - 1$ occur. Hence

$$\text{ht } I_n(E_n) = 2n - (n - 1) = n + 1. \quad (109)$$

The polynomial ring S is Cohen–Macaulay, so the height in Eq. (109) is also the required grade. The Buchsbaum–Rim exactness criterion applies at the expected maximal-minor height [5,6]. Its first syzygy map generates $\ker E_n$ by the circuit vectors supported on $n + 1$ columns; the nonzero coordinates of each such vector are signed maximal minors.

Give the codomain of E_n the row shifts that make Eq. (107) graded. Every circuit generator then has coefficient degree $\binom{n}{2}$, and polynomial multiples have no homogeneous coefficient component below this degree. Thus

$$\text{indeg } \ker E_n \geq \binom{n}{2}.$$

The direct circuit constructed above attains equality. This proves the lower-bound half of Theorem 6: every member of a consistency-complete family is a nonzero vector in $\ker E_n$ and therefore has coefficient degree at least $\binom{n}{2}$. The direct family’s operational completeness is proved below.

F.4 Branch-local integralizers

Work first in the branch ring $S_b = \mathbb{F}[x_0, \dots, x_k]$. Let $V = V(x_0, \dots, x_k)$, $\delta = \det(V)$, and $\Gamma = e_k^\top \text{adj}(V)$. Suppose

$$a^\top V = D e_k^\top, \quad D \neq 0. \quad (110)$$

In the fraction field,

$$a^\top = D e_k^\top V^{-1} = \frac{D}{\delta} \Gamma,$$

and therefore $\delta \mid D \Gamma_i$ for every i . Fix $u < v$ and let $\ell_{uv} = x_v - x_u$. This is an irreducible factor of δ . The coordinate Γ_u contains only differences whose two indices differ from u , so $\ell_{uv} \nmid \Gamma_u$. Unique factorization gives $\ell_{uv} \mid D$. Repeating this for every pair shows $\delta \mid D$. Consequently there is a unique nonzero $h \in S_b$ such that

$$D = h\delta, \quad a_i = h\Gamma_i. \quad (111)$$

This proves the integralizer characterization in Theorem 7.

Now restore the two disjoint variable sets. If branch-local integralizers are jointly eliminated through

$$c_0 D_0 = c_1 D_1 = M \neq 0,$$

then $\delta_0 \mid M$ and $\delta_1 \mid M$. Since the determinants use disjoint variables, they are coprime, and

$$\delta_0 \delta_1 \mid M.$$

The final coefficient of $o_i^{(0)}$ is

$$c_0 h_0 \Gamma_i^{(0)} = \frac{M}{\delta_0} \Gamma_i^{(0)}.$$

It is divisible by $\delta_1 \Gamma_i^{(0)}$ and hence has degree at least

$$\binom{k+1}{2} + \frac{k(k-1)}{2} = k^2.$$

The same argument applies to the other branch. Equality is attained by $D_b = \delta_b$, $a_i^{(b)} = \Gamma_i^{(b)}$, $c_0 = \delta_1$, and $c_1 = \delta_0$. This completes the branch-separated lower and upper bounds.

F.5 Transfer of nonzeroness

The direct identities always produce kernel openings; a cryptographic hardness output additionally requires one to be nonzero. Define

$$N_b = \text{adj}(V_b) O_b, \quad K = (K_0, \dots, K_k)^\top.$$

Equation (49) gives

$$\text{adj}(V_1) K = \delta_0 N_1 - \delta_1 N_0 =: W. \quad (112)$$

In particular,

$$W_k = e_k^\top W \neq 0 \implies W \neq 0 \implies K \neq 0 \implies \exists i : K_i \neq 0.$$

Each K_i is an anchor-universal kernel comparer by Eq. (49). Hence $\{K_0, \dots, K_k\}$ satisfies the completeness condition of Section 6.2. Together with its degree- $\binom{n}{2}$ circuit coefficients and the lower bound in Appendix F.3, this proves Theorem 6. No determinant cancellation is used in the nonzeroness implication.

G Vandermonde HMC and UMC Norm Derivations

We derive the safe radii in Corollary 1 from the norm conventions and product bounds of Esgin et al. The derivation is separate from the polynomial-degree lower bounds: it gives certified upper bounds under an actual challenge support, not norm optimality.

G.1 Source norm convention

Let

$$R = \mathbb{Z}[X]/(X^{N_R} + 1), \quad R_q = R/qR.$$

Esgin et al. write d for N_R . Their norm of a polynomial is the Euclidean norm of its coefficient vector, and for a polynomial vector $u = (u_0, \dots, u_{m-1})$,

$$\|u\|^2 = \sum_{j=0}^{m-1} \|u_j\|^2.$$

In particular,

$$\|(z, f)\| = \sqrt{\|z\|^2 + \|f\|^2}.$$

These conventions and the negacyclic multiplication inequality

$$\|fg\| \leq \sqrt{N_R} \|f\| \|g\| \quad (113)$$

appear in Appendix B.1 and Lemma 8 of Esgin et al. [4, App. B.1, Lemma 8].

For $s \geq 1$ and nonzero differences $d_1, \dots, d_s \in \Delta\mathcal{C}_{w,p}^{N_R}$, their Lemma 3 gives

$$\left\| \prod_{t=1}^s d_t \right\|_{\infty} \leq (2p)^s w^{s-1}, \quad \left\| \prod_{t=1}^s d_t \right\| \leq \sqrt{N_R} (2p)^s w^{s-1}. \quad (114)$$

The proof in their Appendix H uses $\|d_t\|_{\infty} \leq 2p$ and $\|d_t\|_1 \leq 2pw$ [4, Lemma 3 and App. H].

Lemma 19 (Scalar-times-vector bound). *If c is a product of $s \geq 1$ nonzero challenge differences and u is a polynomial vector, then*

$$\|cu\| \leq N_R (2p)^s w^{s-1} \|u\|.$$

Proof. Apply Eq. (113) componentwise and sum squares:

$$\|cu\|^2 = \sum_j \|cu_j\|^2 \leq N_R \|c\|^2 \sum_j \|u_j\|^2.$$

Taking square roots and substituting Eq. (114) yields the claim. The two factors $\sqrt{N_R}$ come respectively from the coefficient-product bound and negacyclic multiplication.

Esgin et al.'s Lemma 5 provides an independent check of this accounting: its adjugate row contains $k(k-1)/2$ differences, and applying the same two source inequalities gives their displayed factor N_R [4, Lemma 5 and App. H]. When $k = 1$, that adjugate row is $(-1, 1)$ and should be handled directly rather than by inserting $s = 0$ into Lemma 19.

G.2 Direct comparison

Let $\kappa = \binom{k+1}{2}$. By Eq. (50), every nonzero coefficient in

$$K_i = \delta_0 o_i^{(1)} - \sum_{j=0}^k \lambda_j(y_i) o_j^{(0)}$$

is a product of κ nonzero challenge differences. If a cross difference vanishes, the corresponding $\lambda_j(y_i)$ term is zero and is omitted. For any common cap $\|o_i^{(b)}\| \leq T$, the triangle inequality and Lemma 19 give

$$\|K_i\| \leq (k+2)N_R(2p)^\kappa w^{\kappa-1}T. \quad (115)$$

Write $o_i^{(b)} = (z_i^{(b)}, f_i^{(b)})$ and suppose

$$\|z_i^{(b)}\| \leq T_z, \quad \|f_i^{(b)}\| \leq T_f.$$

For HMC, the opening map is $\mathcal{B}(z, f) = G(z, f)$ and the verifier checks the Euclidean norm of the complete message–randomness pair [4, Sect. 2.4, Lemma 1].

Hence $\|(z_i^{(b)}, f_i^{(b)})\| \leq T_H$ for $T_H = \sqrt{T_z^2 + T_f^2}$, and a nonzero K_i satisfies

$$GK_i = 0, \quad 0 < \|K_i\| \leq B_{\text{dir}}^{\text{HMC}} = (k+2)N_R(2p)^\kappa w^{\kappa-1}T_H.$$

It is therefore an M-SIS vector in the full HMC opening module.

For UMC,

$$\mathcal{B}(z, f) = Gz + (0, f),$$

and only $\|z\|$ is checked [4, Sect. 2.4, Lemma 2]. If a full pair $K_i = (K_{i,z}, K_{i,f})$ is nonzero and $GK_{i,z} + (0, K_{i,f}) = 0$, the top rows give $G_1 K_{i,z} = 0$. Moreover, $K_{i,z} = 0$ would force $K_{i,f} = 0$, a contradiction. Thus the actual randomness component is a nonzero M-SIS vector and

$$0 < \|K_{i,z}\| \leq B_{\text{dir}}^{\text{UMC}} = (k+2)N_R(2p)^\kappa w^{\kappa-1}T_z.$$

G.3 Branch-separated comparison

For the canonical integralizer, write

$$n_b = \sum_{i=0}^k \Gamma_i^{(b)} o_i^{(b)}, \quad W_k = \delta_0 n_1 - \delta_1 n_0.$$

Each final raw-response coefficient $\delta_{1-b} \Gamma_i^{(b)}$ is a product of

$$\binom{k+1}{2} + \frac{k(k-1)}{2} = k^2$$

challenge differences. Applying Lemma 19 directly to this fully expanded coefficient avoids paying a second, artificial norm-conversion factor. Each common-factor branch opening has norm at most

$$(k+1)N_R(2p)^{k^2}w^{k^2-1}T,$$

and their difference has norm at most twice this quantity. Therefore

$$\begin{aligned} B_{\text{bw}}^{\text{HMC}} &= 2(k+1)N_R(2p)^{k^2}w^{k^2-1}T_H, \\ B_{\text{bw}}^{\text{UMC}} &= 2(k+1)N_R(2p)^{k^2}w^{k^2-1}T_z. \end{aligned}$$

The HMC difference is a full opening-kernel vector. In UMC, the same top-row argument used above shows that a nonzero full difference has a nonzero randomness component in the M-SIS kernel.

G.4 Ratio and numerical derivation

Let $T = T_H$ for HMC and $T = T_z$ for UMC. Dividing the preceding bounds and using $k^2 - \kappa = k(k-1)/2$ gives

$$\frac{B_{\text{bw}}}{B_{\text{dir}}} = \frac{2(k+1)}{k+2}(2pw)^{k(k-1)/2}. \quad (116)$$

For $k = 1$, the two coefficient degrees are both one and Eq. (116) equals $4/3$. This is a difference between the displayed safe triangle bounds, not a strict degree separation.

For reproducibility, the two logarithms used in Table 1 are

$$\begin{aligned} \log_2(B_{\text{dir}}/T) &= \log_2(k+2) + \log_2 N_R + \kappa \log_2(2p) + (\kappa-1) \log_2 w, \\ \log_2(B_{\text{bw}}/T) &= 1 + \log_2(k+1) + \log_2 N_R + k^2 \log_2(2p) + (k^2-1) \log_2 w. \end{aligned}$$

Substitution of the sourced challenge example $(N_R, w, p) = (256, 60, 1)$ gives, before rounding,

k	$\log_2(B_{\text{dir}}/T)$	$\log_2(B_{\text{bw}}/T)$	$\log_2(B_{\text{bw}}/B_{\text{dir}})$
2	24.8137811912	32.3056342875	7.4918530963
3	45.8563810729	67.2551247649	21.3987436919
4	73.7469778612	115.9252870290	42.1783091678
5	108.5038232606	178.3503367953	69.8465135347

These are deterministic formula evaluations. They are neither implementation measurements nor security estimates.