Fast Re-Randomization and Galois-Derived Multi-Challenge Proofs for Module-LIP

Anonymous authors

Anonymous submission

Abstract. Re-randomizing a module-LIP instance requires a fresh point in the same orbit together with a transporter. Two obstacles prevent a compact multi-challenge use of this operation. First, fast NTRU completion applies to norm-coprime pairs, whereas the self-reduction samples only ideal-primitive columns; rejecting the remaining columns biases the joint endpoint–transporter law. We give a bounded public completion procedure that preserves the sampled column and is exact on success. Its statistical distance from formal re-randomization is bounded by the probability of its explicit public-budget exhaustion event. It supports exact batch accounting and a zero-failure deterministic completion algorithm running in $\widetilde{O}(n^2(B + \log n))$ bit operations under the Pornin–Prest size condition.
Second, deriving many challenges from one form by Galois action creates correlated module-LIP instances. We determine the information revealed by a transition between two derived endpoints, exclude the degenerate low-order quotients, and formulate the corresponding pairwise search assumption. Combining the two results yields a compact public-coin protocol with statistical honest-verifier zero knowledge under endpoint-uniform tail bounds, pairwise extraction from two openings, and full-witness extraction from all challenges. Experiments evaluate the completion component and the first-hit model.

Keywords: Module-LIP · re-randomization · NTRU completion · group actions · Sigma protocols

1 Introduction

The lattice isomorphism problem (LIP) and its module variant support structured group actions in post-quantum cryptography. HAWK uses rank-two module-LIP for compact signatures, and subsequent work develops self-reductions, automorphism algorithms, and commitment schemes for the same action [1,4,3,5]. A recurring operation is *re-randomization*: from a public Hermitian form, sample a fresh form in the same orbit together with a transporter relating the two.

This paper addresses two obstacles to using re-randomization in a compact multi-challenge proof. The first is computational. The HAWK self-reduction samples a primitive column and completes it to a determinant-one matrix. Pornin–Prest field-norm solvers make this step fast for norm-coprime inputs [8], but norm

coprimality is strictly stronger than ideal primitivity. Resampling until it holds changes the column distribution and can therefore change the joint endpoint–transporter law. The problem is to retain the fast solver without conditioning away valid primitive columns and while keeping the public work bounded.

The second obstacle is structural. Publishing an independent module-LIP instance for every challenge makes the statement grow with the challenge set. Galois action can derive all challenge endpoints from one form, but the derived instances are correlated. A transition extracted between two challenges is therefore not immediately covered by ordinary module-LIP: one must first identify the algebraic information forced by the correlation and exclude the degenerate challenge quotients.

Our results.

1. **Distribution-preserving fast re-randomization.** We make field-norm completion compatible with the primitive-column law used by module-LIP self-reduction. The resulting bounded procedure preserves the formal endpoint–transporter distribution up to an explicit public abort event, admits exact batch accounting, and has a deterministic zero-failure completion algorithm with complexity $\widetilde{O}(n^2(B + \log n))$ under the Pornin–Prest size condition.
2. **One-form multi-challenge instances with explicit correlated hardness.** We compress the entire challenge family to Galois images of one public form and determine exactly what a transition between two derived endpoints reveals. This yields an admissible challenge family and a precise correlated pairwise module-LIP search problem, rather than an unsupported appeal to ordinary module-LIP.
3. **Application to public-coin proofs.** Combining the two results gives a compact multi-challenge protocol. Under endpoint-uniform completion tails and the correlated pairwise assumption, it has statistical honest-verifier zero knowledge, pairwise extraction from two openings, full-witness extraction from all challenges, and parallel soundness $q^{-\mu} + \sqrt{\delta_{\mathrm{pair}}}$.

Techniques. For completion, write the sampled primitive column as $(f, g)^{\mathsf{T}}$ and probe the public shears $f + X^i g$. Whenever $\gcd(N(f + X^i g), N(g)) = 1$, one field-norm solve completes the sheared pair; a triangular correction converts this into a completion of the *original* column. The inherited HAWK post-processing therefore sees the same first column, and a direct coupling bounds the statistical loss by public-budget exhaustion. Convolving the supplied first-hit laws gives exact multi-equation quotas. For zero-failure comparison, n fixed shears have collective norm gcd one. A gcd support split and the Chinese remainder theorem combine them into a single norm-coprime shear, after which one Pornin–Prest solve suffices.

For the challenge family, let $\mathbf{Q}_0 = \mathbf{B}^*\mathbf{B}$ and $\mathbf{Q}_i = \sigma_i(\mathbf{Q}_0)$. Pulling a transition from $\mathbf{Q}_i$ to $\mathbf{Q}_j$ back to the basepoint gives $\widetilde{D} = \mathbf{B}^{-1} A \eta(\mathbf{B})$, where $\eta = \sigma_i^{-1}\sigma_j$ and A is integral unitary. Thus the hidden inverse basis lies in one of finitely many integral semilinear slices. Their dimensions identify the exceptional quotients

of order at most four and define the admissible challenge set. The same first message can then be opened with respect to any derived endpoint: two openings reveal a pairwise transporter, whereas openings for all challenges recover the full witness tuple.

Scope and experiments. The algebraic completion, distributional replacement, and protocol reductions do not rely on the occupancy model. We use that model to obtain auditable first-hit laws and public quotas, test its observable small-tail predictions on independent discrete-Gaussian cohorts, and measure the completion component. A concrete protocol instantiation additionally requires a uniform tail bound for every real and simulated endpoint law.

Organization. Section 2 fixes the algebraic and re-randomization interfaces. Sections 3 and 4 prove the completion and Galois-derived pairwise results, and Section 5 gives the protocol. Detailed proofs, model calculations, and experimental methodology appear in the appendices.

2 Notation and Preliminaries

2.1 Cyclotomic and module-LIP setting

Let $n = 2^k$ with $k \geq 3$ and

$$\mathbb{L} = \mathbb{Q}(\zeta_{2n}) \simeq \mathbb{Q}[X]/(X^n + 1), \qquad \mathcal{O}_{\mathbb{L}} = \mathbb{Z}[\zeta_{2n}] \simeq \mathbb{Z}[X]/(X^n + 1).$$

We identify X with ζ_{2n} and abbreviate the absolute field norm and trace by N and Tr. Complex conjugation is denoted by a bar, and $A^* = \bar{A}^{\mathsf{T}}$. For $x = \sum_{j=0}^{n-1} x_j X^j$, write $\mathrm{vec}(x) = (x_0, \ldots, x_{n-1})^{\mathsf{T}}$ and $\|x\|_1 = \sum_j |x_j|$. The normalized trace pairing is the coefficient Euclidean pairing:

$$\frac{1}{n} \mathrm{Tr}(\bar{x}y) = \langle \mathrm{vec}(x), \mathrm{vec}(y) \rangle_2. \tag{1}$$

We use $[m] = \{1, \ldots, m\}$, Δ_{SD} for statistical distance, $\perp$ for abort, and $\widetilde{O}$ for soft-O bit complexity. Independent randomized invocations use independent coins unless stated otherwise.

The Galois group $\Gamma = \mathrm{Gal}(\mathbb{L}/\mathbb{Q}) \simeq (\mathbb{Z}/2n\mathbb{Z})^\times$ acts entrywise on vectors and matrices. It is abelian, preserves $\mathcal{O}_{\mathbb{L}}$, and commutes with the adjoint:

$$\sigma(A^*) = \sigma(A)^* \qquad (\sigma \in \Gamma). \tag{2}$$

For $d \geq 1$, let $\Gamma[d] = \{\gamma \in \Gamma : \gamma^d = \mathrm{id}\}$.

A positive-definite Hermitian form is a matrix $\mathbf{Q} = \mathbf{Q}^* \in M_2(\mathbb{L})$ that is positive definite at every complex embedding. It induces

$$\langle x, y \rangle_{\mathbf{Q}} = \frac{1}{n} \mathrm{Tr}(x^* \mathbf{Q} y), \qquad \|x\|_{\mathbf{Q}}^2 = \langle x, x \rangle_{\mathbf{Q}}.$$

The right action of $\mathrm{SL}_2(\mathcal{O}_{\mathbb{L}})$ is

$$\mathbf{Q} \cdot U = U^* \mathbf{Q} U. \tag{3}$$

For forms in the same orbit, define the directed transporter set

$$\mathsf{Trans}(\mathbf{Q}_0, \mathbf{Q}_1) = \{U \in \mathrm{SL}_2(\mathcal{O}_{\mathbb{L}}) : U^* \mathbf{Q}_0 U = \mathbf{Q}_1\}.$$

Definition 1 (Rank-two determinant-one module-LIP). *Given positive-definite Hermitian forms $\mathbf{Q}_0, \mathbf{Q}_1$ with $\mathbf{Q}_1 \in [\mathbf{Q}_0]_{\mathrm{SL}}$, find $U \in \mathsf{Trans}(\mathbf{Q}_0, \mathbf{Q}_1)$.*

This is the rank-two special-linear form used by HAWK and related work [1,5,6]. In particular, $\mathbf{B} \in \mathrm{SL}_2(\mathcal{O}_{\mathbb{L}})$ defines $\mathbf{Q} = \mathbf{B}^* \mathbf{B}$ and transports I to $\mathbf{Q}$.

2.2 Protocol terminology

We use the standard three-move public-coin syntax (a, c, z) for a Σ-protocol. For a fixed challenge c, let $\mathsf{Real}_{\Pi}(\mathsf{st}, \mathsf{wit}; c)$ and $\mathsf{Sim}_{\Pi}(\mathsf{st}; c)$ denote the honest and simulated transcript laws. Statistical honest-verifier zero knowledge with error ϵ means that their statistical distance is at most ϵ for every valid statement–witness pair and every challenge; if an execution may abort, $\perp$ is part of the output space.

Our extraction statement uses *relaxed special soundness* [2, Definition 2.8]: two accepting transcripts with the same first message and different challenges yield a witness for an explicitly stated extraction relation, which may be weaker than the full statement relation. Perfect completeness below is conditioned on non-abort, with the abort probability accounted for separately.

2.3 Formal re-randomization

For $\sigma > 0$, let $D_{\mathbf{Q},\sigma}$ be the discrete Gaussian on $\mathcal{O}_{\mathbb{L}}^2$ with mass proportional to

$$\rho_{\mathbf{Q},\sigma}(x) = \exp\left(-\frac{\|x\|_{\mathbf{Q}}^2}{2\sigma^2}\right).$$

Write

$$\mathcal{M}(n, \sigma, \mathbf{Q}) = \mathcal{L}\big[(f, g)^{\mathsf{T}} \leftarrow D_{\mathbf{Q},\sigma} \;\big|\; f\mathcal{O}_{\mathbb{L}} + g\mathcal{O}_{\mathbb{L}} = \mathcal{O}_{\mathbb{L}}\big]$$

for the accepted primitive first-column law.

We use the HAWK self-reduction only through its endpoint–transporter interface [1, Section 6]; the same covariance is used by the module-LIP commitment re-randomizer [5, Section 4.2].

Algorithm 2 ($\mathsf{ReRand}_{\mathrm{form}}(\mathbf{Q}, \sigma)$). Sample $y_1 \leftarrow \mathcal{M}(n, \sigma, \mathbf{Q})$, complete y_1 by any exact matrix in $\mathrm{SL}_2(\mathcal{O}_{\mathbb{L}})$, apply the inherited HAWK post-completion map, and return its endpoint–transporter pair $(\mathbf{T}, W)$.

Lemma 3 (HAWK post-completion and self-reduction law). *The output of Algorithm 2 satisfies $W \in \mathrm{SL}_2(\mathcal{O}_{\mathbb{L}})$ and $\mathbf{T} = W^* \mathbf{Q} W$. Conditioned on the primitive first column, its joint law is independent of the chosen exact complement. Moreover, if $D \in \mathsf{Trans}(\mathbf{Q}, \mathbf{Q}')$ and $(\mathbf{T}, W) \leftarrow \mathsf{ReRand}_{\mathrm{form}}(\mathbf{Q}, \sigma)$, then $(\mathbf{T}, D^{-1}W)$ is distributed as $\mathsf{ReRand}_{\mathrm{form}}(\mathbf{Q}', \sigma)$.*

2.4 Primitive completion and field-norm solving

A column $(f, g)^{\mathsf{T}} \in \mathcal{O}_{\mathbb{L}}^2$ is *primitive* when $f\mathcal{O}_{\mathbb{L}} + g\mathcal{O}_{\mathbb{L}} = \mathcal{O}_{\mathbb{L}}$. Completing it to $\mathrm{SL}_2(\mathcal{O}_{\mathbb{L}})$ is equivalent to finding $(F, G) \in \mathcal{O}_{\mathbb{L}}^2$ such that

$$fG - gF = 1. \tag{4}$$

Lemma 4 (Primitive completion). *A pair (f, g) is primitive if and only if it admits a solution to (4).*

Proof. A solution expresses 1 in the ideal (f, g). Conversely, from $af + bg = 1$ take $(F, G) = (-b, a)$. □

Lemma 5 (Norm-coprime sufficiency). *If $\gcd(|N(f)|, |N(g)|) = 1$, then (f, g) is primitive.*

Proof. A proper ideal containing f and g lies below a rational prime dividing both principal-ideal norms. □

Pornin and Prest solve the target NTRU equation $fG - gF = q$ by repeated field norms along the power-of-two cyclotomic tower [8, Section 4, Algorithms 4–5]. Their solver succeeds when $\gcd(|N(f)|, |N(g)|)$ divides q; this condition is sufficient for the algorithm but not necessary for solvability.

Definition 6 (TowerSolve). *For an integer q, $\mathsf{TowerSolve}_q(f, g)$ denotes the exact Pornin–Prest target solver on inputs satisfying $\gcd(|N(f)|, |N(g)|) \mid q$. We abbreviate $\mathsf{TowerSolve} = \mathsf{TowerSolve}_1$.*

Norm coprimality therefore enables a fast right-hand-side-one completion, while section 3 removes it as a condition on the sampled primitive column.

Field-norm cost convention. Call a TowerSolve invocation *size-controlled* when the reduced-lift quantity in Pornin–Prest [8, Eq. (30)] satisfies the hypothesis of their Lemma 3. On degree-n inputs whose Euclidean norms have logarithms at most b, Lemmas 3–4 then give $\widetilde{O}(nb)$ bit operations. We use this condition only for the single final solve in the deterministic comparison; norm computation and standard integer or negacyclic-polynomial arithmetic are charged at their asymptotically fast bit complexities. The condition affects only the running-time claim: all completion identities and divisibility statements are unconditional.

3 Fast Primitive Completion

We give an exact primitive-completion interface with a public retry bound; its only statistical loss is budget exhaustion.

Proposition 7 (Obstruction to direct TowerSolve re-randomization). *Let*

$$\mathcal{A}_{\mathrm{TS}} = \{(f, g)^{\mathsf{T}} \in \mathcal{O}_{\mathbb{L}}^2 : \gcd(|N(f)|, |N(g)|) = 1\}.$$

Replacing the completion in formal re-randomization by the direct $q = 1$ TowerSolve path and restarting until the first column lies in $\mathcal{A}_{\mathrm{TS}}$ does not, in general, preserve basepoint covariance. For $D \in \mathsf{Trans}(\mathbf{Q}, \mathbf{Q}')$, the first columns of the transported real response and of a direct simulation from $\mathbf{Q}'$ are respectively distributed as $D_{\mathbf{Q}',\sigma}$ conditioned on $D^{-1}\mathcal{A}_{\mathrm{TS}}$ and on $\mathcal{A}_{\mathrm{TS}}$. Because the response transporter is public, membership in $\mathcal{A}_{\mathrm{TS}}$ is an efficient distinguishing test. Whenever the resulting mismatch probability is non-negligible, the canonical direct simulator is efficiently distinguishable from the real transcript. Thus hardness of distinguishing endpoint forms alone cannot establish honest-verifier zero knowledge for this replacement. The exact gap and the failure of invariance are proved in section D.1; the same obstruction to the formal re-randomization guarantee under this substitution is noted at the end of [5, Section 4].

Accordingly, the construction below keeps the ideal primitive-column law as its reference distribution and retains bounded-retry exhaustion as $\perp$; it does not renormalize re-randomization to the direct-TowerSolve success set.

3.1 Fixed Public Shears and Exact Completion

For $1 \le K \le 2n$, use the ordered shears $1, X, \ldots, X^{K-1}$. Multiplication by X^i is a negacyclic signed rotation, so

$$\|f + X^i g\|_1 \le \|f\|_1 + \|g\|_1. \tag{5}$$

Algorithm 8 (BoundedRetryComplete(f, g, K)). Require a primitive pair $(f, g) \in \mathcal{O}_{\mathbb{L}}^2$.

1. If $g = 0$, execute K public dummy probe slots and return $(0, f^{-1})$ without calling TowerSolve.
2. Set $M = |N(g)|$. For every $i = 0, \ldots, K-1$, compute
$$\delta_i = \gcd\big(N(f + X^i g) \bmod M, M\big).$$
3. Let $i^\star = \min\{i < K : \delta_i = 1\}$. If this set is empty, return $\perp$.
4. Compute $(F', G') \leftarrow \mathsf{TowerSolve}(f + X^{i^\star} g, g)$ and return $(F' - X^{i^\star} G', G')$.

Thus the algorithm performs all K public probes but invokes TowerSolve only once, on the first successful shear. The displayed algorithm is not by itself a complete constant-time design.

Theorem 9 (Exact fixed-shear completion). *The algorithm handles $g = 0$ correctly in K public dummy slots and makes no solver call on that branch. For $g \neq 0$, every non-abort output satisfies $fG - gF = 1$, uses exactly K norm/gcd probes and one call to TowerSolve, and obeys (5) before the solve.*

Proof. Primitivity of $(f, 0)$ makes f a unit. Otherwise, for a successful i,

$$fG' - g(F' - X^i G') = (f + X^i g)G' - gF' = 1.$$

The remaining statements follow from the public loop and the signed-rotation property. □

3.2 Bounded-Retry Distributional Replacement

Define the success event

$$S_K(f,g) = \big[\exists i < K : \gcd(N(f + X^i g), N(g)) = 1\big]. \tag{6}$$

The formal re-randomization interface samples its first column from $\mathcal{M}(n,\sigma,\mathbf{Q})$ as defined in section 2. For this primitive-column law, define

$$\varepsilon_{\mathbf{Q}}(K) = \Pr_{(f,g)\leftarrow\mathcal{M}(n,\sigma,\mathbf{Q})}[\neg S_K(f,g)]. \tag{7}$$

The budget K is public, and exhaustion produces $\perp$.

Theorem 10 (Statistical cost of bounded retry). *Include $\perp$ in the output space. Couple formal re-randomization and the bounded-retry experiment by the same primitive column and downstream coins. Their endpoint/transporter laws have statistical distance at most $\varepsilon_{\mathbf{Q}}(K)$. For calls with tails $\varepsilon_1,\ldots,\varepsilon_m$, the loss is at most their sum; real and simulator call counts are charged separately.*

Proof. On S_K, Theorem 9 gives a valid complement of the same sampled column. Complement independence of the inherited post-completion map, lemma 3, therefore gives the same joint output. The coupled executions may differ only on $\neg S_K$, so the first claim follows from the coupling lemma and the multi-call claim from a hybrid triangle inequality. □

3.3 Deterministic Worst-Case Completion

Bounded retry preserves the target distribution but may abort. For a zero-failure comparison, use the fixed set

$$A_{\mathrm{hit}} = \{0, 1, X, \ldots, X^{n-2}\}. \tag{8}$$

For every primitive (f,g), the integers $N(g)$ and $\{N(f+ag) : a \in A_{\mathrm{hit}}\}$ have collective gcd one. Without factoring $N(g)$, gcd support splitting assigns each prime-power factor of $|N(g)|$ to a candidate that is nonvanishing modulo its underlying prime. Coefficientwise CRT then produces one $h \in \mathcal{O}_{\mathbb{L}}$ such that

$$\gcd\big(|N(f+hg)|, |N(g)|\big) = 1.$$

One field-norm solve on $(f+hg, g)$, followed by undoing the shear, completes the original column.

Theorem 11 (Worst-case primitive completion). *For every primitive $(f,g) \in \mathcal{O}_{\mathbb{L}}^2$ with B-bit coefficients, the deterministic algorithm of section A computes (F,G) satisfying $fG - gF = 1$. Assuming that its final TowerSolve invocation is size-controlled in the sense of section 2, its running time is*

$$\widetilde{O}\big(n^2(B + \log n)\big) \tag{9}$$

bit operations.

The algorithm uses one norm computation for g, n sheared norm probes, one CRT-combined shear, and one norm-coprime TowerSolve call. It is a deterministic comparison result; the protocol uses bounded retry and accounts for its abort probability explicitly.

3.4 Single-Equation Laws

For a fixed endpoint law, define the first hit and survival function

$$\tau = \min\{k \geq 1 : S_k(f, g)\}, \qquad F_{\mathbf{Q}}(K) = \Pr[\tau > K]. \tag{10}$$

Right-censored or unresolved mass remains on the failure side. The results below distinguish an exact calculation conditional on a supplied law from the evidence used to obtain that law.

To provide a finite, auditable model input, fix an odd prime p, let $d_p = \mathrm{ord}_{2n}(p)$, $r_p = n/d_p$, and $q_p = p^{d_p}$. In each of the r_p residue-field components, let $\theta_{p,j}$ be the image of X and define the projective ratio

$$\rho_j = \begin{cases} -f_j/g_j, & g_j \neq 0, \\ \infty, & g_j = 0. \end{cases} \tag{11}$$

For fixed $K \leq 2n$, the values $1, \theta_{p,j}, \ldots, \theta_{p,j}^{K-1}$ are distinct in every component. We identify them componentwise with the common probe labels $1, \ldots, K$; this is only a relabeling of the projective line and preserves uniformity in Assumption 12.

Assumption 12 (Unit-root local occupancy heuristic). *For each retained prime, the r_p ratios are modeled as independent uniform labels in $\mathbf{P}^1(\mathbb{F}_{q_p})$, and retained primes are modeled as independent. This assumption defines a finite-cutoff heuristic model; it is not a theorem about, or a measurement of, a re-randomization sampler.*

For $r \in [K]$, let $C_p(K) \subseteq [K]$ contain r when some component has label ∞ and some component has the normalized probe label r (equivalently, its raw ratio is $\theta_{p,j}^{r-1}$). Algebraically, the first K probes all fail exactly when these contamination sets cover $[K]$:

$$\neg S_K(f, g) \quad \Longleftrightarrow \quad \bigcup_{p \mid N(g)} C_p(K) = [K]. \tag{12}$$

For a retained finite set $\mathcal{S}_B$ of primes, the occupancy assumption therefore gives the finite-cutoff evaluator. The reported parameter values use exactly the split-prime set $\mathcal{S}_B = \{p \leq B : p \equiv 1 \pmod{2n}\}$ with $B = 3 \cdot 10^7$:

$$F_B(K) = \sum_{j=0}^{K} (-1)^j \binom{K}{j} \prod_{p \in \mathcal{S}_B} \left[\left(1 - \frac{1}{q_p + 1}\right)^{r_p} + \left(1 - \frac{j}{q_p + 1}\right)^{r_p} - \left(1 - \frac{j+1}{q_p + 1}\right)^{r_p} \right]. \tag{13}$$

Equation (13) is a finite-cutoff heuristic prediction; Appendix B gives the DP, omitted-prime analysis, and precision checks.

Model-conditioned budgets. For a target tail δ and a conservatism factor $c \geq 1$ fixed independently of the parameter calculation, let

$$K_{\mathrm{raw}}(n, \delta) = \min\{K : F_B(K) \leq \delta\},$$
$$K_{\mathrm{cons}}(n, \delta) = \min\{K : \min(1, cF_B(K)) \leq \delta\},$$

These are public budgets derived from the supplied model, not empirical estimates of cryptographic tails.

For the parameter calculations below, calibration fixes $c = 2.68667279036$. Table 1 gives the resulting model-conditioned budgets. In particular, a conservative model tail of at most 2^{-128} requires $K = 40$ for $n = 512$ and $K = 44$ for $n = 1024$.

Table 1. Retry budgets from the split-prime occupancy model with cutoff $B = 3 \cdot 10^7$. "Raw" uses F_B and "cons." uses $\min(1, cF_B)$.

	$n = 512$		$n = 1024$	
target δ	raw	cons.	raw	cons.
10^{-6}	4	5	5	6
10^{-9}	7	7	8	9
2^{-32}	7	8	9	9
2^{-64}	17	17	20	20
2^{-128}	40	40	44	44
2^{-256}	92	92	97	98

Model check. Against disjoint calibration and holdout cohorts for $n \in \{16, 32, 64\}$, the experimental/model ratio lies in $[0.85, 1.17]$ for $K \leq 3$, and the conservative curve covers 89 of 90 cells through $K = 5$. The complete comparison, including the unique exception, is in section E. This supports the model only on the observed range, not at cryptographic tails.

Completion performance. At the conservative 2^{-128} model budgets, $K = 40$ for $n = 512$ and $K = 44$ for $n = 1024$. Across the three coefficient-variance profiles, bounded Completion has median latency 68.0–104.5 ms and 348.7–533.4 ms, respectively. Exact norm/mod probes account for 78.2–83.4% of these medians, whereas the single TowerSolve call accounts for 12.9–18.3%. An exact-width, reusable TowerSolve workspace reduces its retained high-water mark from 6.32–32.67 MiB to 0.17–0.53 MiB without a material timing penalty in the paired comparison. The complete process, including the shared exact-norm tables and arithmetic runtime, has a 13.85–13.88 MiB peak working set. The phase and memory accounting is reported in section E.

3.5 Batched Quotas

A shared schedule processes equations in public order, spending probes until each first hit and padding to exactly K' public slots. It returns $\perp$ if any equation remains unresolved and otherwise invokes TowerSolve once per equation on its recorded first clean shear. The padding and solver schedule are part of the public-work interface. In this shared schedule, a $g = 0$ equation has $\tau = 1$: it consumes one logical slot, performs no norm/gcd computation, and requires no solver call.

Definition 13 (Protocol batch profile). *A batch profile is*

$$\mathsf{BP} = (T, (F_j)_{j=1}^{T}, K'),$$

where the T equations use fresh independent coins and F_j is the supplied single-equation survival law for equation j, with $F_j(k) = \Pr[\tau_j > k]$ for $k \geq 0$ and $F_j(0) = 1$. The variable τ_j may take the value ∞; that mass remains on the failure side. Put

$$a_{j,k} = F_j(k-1) - F_j(k), \qquad A_j(z) = \sum_{k \geq 1} a_{j,k} z^k,$$

and define

$$\varepsilon_{T,K'} = 1 - \sum_{s=T}^{K'} [z^s] \prod_{j=1}^{T} A_j(z). \tag{14}$$

Theorem 14 (Rectangular and shared quotas). *For supplied monotone survival laws and independent coins, giving equation j at most K attempts has batch failure probability $1 - \prod_j (1 - F_j(K))$. A shared sequential budget K' is exhausted with probability exactly $\varepsilon_{T,K'}$. Censored mass is retained on the failure side and is never renormalized away.*

Proof. The rectangular event succeeds exactly when every equation hits by K. For the shared budget, independence makes $\prod_j A_j(z)$ the generating function of $\sum_j \tau_j$; summing its coefficients through K' gives the success probability in (14). □

For a common first-hit mass a_r, the equivalent recurrence is

$$P_0(0) = 1, \qquad P_t(s) = \sum_{r=1}^{s} P_{t-1}(s-r) a_r. \tag{15}$$

If values above $R < K'$ are censored, the resulting CDF is conservative and differs from the uncensored CDF by at most

$$\Pr[\exists j : \tau_j > R] \leq T \max_j F_j(R). \tag{16}$$

Thus convolution is exact for the supplied law without asserting that the law itself is correct.

4 Galois-Derived Pairwise Module-LIP

We derive many challenge endpoints from one module-LIP instance, determine what a transition between two derived endpoints reveals, and then isolate the remaining correlated search assumption.

4.1 Derived instances and the pairwise search problem

Use the field, order, action, and Galois group fixed in section 2.

Definition 15 (Quotient admissibility). *A challenge encoding Σ is admissible if*

$$\mathsf{Adm}_{>4}(\Sigma): \qquad \mathrm{ord}(\sigma_i^{-1}\sigma_j) > 4 \quad \textit{for every } i \neq j. \tag{17}$$

The condition is imposed on pairwise quotients: large individual orders do not prevent an order-two quotient.

Definition 16 (Derived setup). *Fix an efficiently samplable distribution family $\mathcal{D}_{\mathbf{B}}(1^\lambda)$ supported on $\mathrm{SL}_2(\mathcal{O}_{\mathbb{L}})$. Select $n = n(\lambda)$ and sample $\mathbf{B} \leftarrow \mathcal{D}_{\mathbf{B}}(1^\lambda)$. Put $\mathbf{Q}_0 = \mathbf{B}^*\mathbf{B}$. For an integer $q \geq 2$, let $\mathcal{I} = \{0, \ldots, q-1\}$ and deterministically encode a set $\Sigma = (\sigma_i)_{i\in\mathcal{I}} \subseteq \Gamma$ with $\sigma_0 = \mathrm{id}$ and $\mathsf{Adm}_{>4}(\Sigma)$. Evaluate $\mathbf{Q}_i = \sigma_i(\mathbf{Q}_0)$, and resample $\mathbf{B}$ if two endpoints collide. The setup family is required to make this acceptance probability inverse-polynomially bounded below, so expected setup time is polynomial. Publish $\mathsf{st} = (\mathbf{Q}_0, \Sigma)$ and retain $\mathsf{wit} = (D_i)_i$, where $D_i = \mathbf{B}^{-1}\sigma_i(\mathbf{B})$. Statement validation rechecks the ring and automorphism encodings, $q \geq 2$, $\sigma_0 = \mathrm{id}$, (17), and endpoint distinctness within the prescribed setup algorithm.*

The statement contains one form and the encoding of Σ; the verifier derives all $\mathbf{Q}_i$.

The derived pairwise search task is to output (i, j, D) such that

$$i \neq j, \qquad D \in \mathrm{SL}_2(\mathcal{O}_{\mathbb{L}}), \qquad D^*\mathbf{Q}_i D = \mathbf{Q}_j. \tag{18}$$

Collision rejection ensures that $i \neq j$ also means $\mathbf{Q}_i \neq \mathbf{Q}_j$, excluding the identity transition.

For an adversary $\mathcal{A}$ given the public output of Definition 16, define

$$\begin{aligned}\mathrm{Adv}^{\mathrm{pair}}_{\mathsf{Setup}}(\mathcal{A}) = \Pr\big[&(\mathsf{st}, \mathsf{wit}) \leftarrow \mathsf{Setup}(1^\lambda);\\ &(i, j, D) \leftarrow \mathcal{A}(\mathbf{Q}_0, \Sigma):\\ &i, j \in \mathcal{I},\ i \neq j,\ \mathbf{Q}_i \neq \mathbf{Q}_j, D \in \mathrm{SL}_2(\mathcal{O}_{\mathbb{L}}),\ D^*\mathbf{Q}_i D = \mathbf{Q}_j\big].\end{aligned}$$

The adversary chooses the endpoint pair after seeing the entire correlated family.

For fixed $i \neq j$, set

$$\eta = \sigma_i^{-1}\sigma_j, \qquad \widetilde{D} = \sigma_i^{-1}(D).$$

By (2), (18) is equivalent to

$$\widetilde{D}^*\mathbf{Q}_0\widetilde{D} = \eta(\mathbf{Q}_0). \tag{19}$$

4.2 Pairwise solution form and semilinear slices

The pairwise relation has a precise unitary ambiguity. Define only the group needed below,

$$\mathrm{SU}_2(\mathcal{O}_{\mathbb{L}}) = \{A \in \mathrm{SL}_2(\mathcal{O}_{\mathbb{L}}) : A^*A = I\}.$$

Proposition 17 (Pairwise solution form). *Fix* $\mathbf{Q}_0 = \mathbf{B}^*\mathbf{B}$. *The solutions to* (19) *are exactly*

$$\widetilde{D} = \mathbf{B}^{-1}A\eta(\mathbf{B}), \qquad A \in \mathrm{SU}_2(\mathcal{O}_{\mathbb{L}}).$$

Proof. Given a solution $\widetilde{D}$, put $A = \mathbf{B}\widetilde{D}\eta(\mathbf{B})^{-1}$. Then

$$\begin{aligned} A^*A &= \eta(\mathbf{B})^{-*}\widetilde{D}^*\mathbf{B}^*\mathbf{B}\widetilde{D}\eta(\mathbf{B})^{-1} \\ &= \eta(\mathbf{B})^{-*}\eta(\mathbf{Q}_0)\eta(\mathbf{B})^{-1} = I. \end{aligned}$$

The determinant is one, so $A \in \mathrm{SU}_2(\mathcal{O}_{\mathbb{L}})$; the converse follows by substitution. The hidden factorization gives the branch $A = I$, whereas an extracted transition may lie in any unitary branch. □

For the actual hidden $\mathbf{B}$, write the recovered transition as

$$\widetilde{D} = \mathbf{B}^{-1}A\eta(\mathbf{B})$$

using proposition 17, and define

$$\Lambda_{A,\eta} = \{P \in M_2(\mathcal{O}_{\mathbb{L}}) : PA = A\eta(P)\}, \qquad C_{A,\eta} = \{\mathrm{col}_1(P) : P \in \Lambda_{A,\eta}\}.$$

Theorem 18 (Semilinear slicing). *For every recovered transition written in the form above,*

$$\{Y \in M_2(\mathcal{O}_{\mathbb{L}}) : YA = \widetilde{D}\eta(Y)\} = \mathbf{B}^{-1}\Lambda_{A,\eta}.$$

Consequently, the first columns of all integral solutions form

$$\mathbf{B}^{-1}C_{A,\eta}.$$

In particular, $\mathbf{B}^{-1}$ *belongs to the matrix solution set.*

Proof. Substitute $Y = \mathbf{B}^{-1}P$. Since $\mathbf{B}$ is integral and unimodular, $Y \in M_2(\mathcal{O}_{\mathbb{L}})$ if and only if $P \in M_2(\mathcal{O}_{\mathbb{L}})$. Moreover,

$$YA = \widetilde{D}\eta(Y) \quad\Longleftrightarrow\quad \mathbf{B}^{-1}PA = \mathbf{B}^{-1}A\eta(P) \quad\Longleftrightarrow\quad PA = A\eta(P).$$

Taking first columns gives the second claim. Finally, $P = I$ satisfies $PA = A\eta(P)$. □

Thus extraction confines the hidden inverse basis to an explicitly enumerable family of integral semilinear slices.

4.3 First-column classification

Write $\mu(\mathbb{L}) = \langle\omega\rangle$ with $|\mu(\mathbb{L})| = M$, let $\eta(\omega) = \omega^s$, and set

$$m = \mathrm{ord}(\eta), \qquad F_\eta = \mathbb{L}^{\langle\eta\rangle}, \qquad F_{\eta^2} = \mathbb{L}^{\langle\eta^2\rangle},$$

with $\mathcal{O}_0 = \mathcal{O}_\mathbb{L} \cap F_\eta$ and $\mathcal{O}_2 = \mathcal{O}_\mathbb{L} \cap F_{\eta^2}$. The integral unitary classification for rank two over CM fields [6, Lemma 3.1] implies that the determinant-one subgroup is

$$\mathrm{SU}_2(\mathcal{O}_\mathbb{L}) = \{T_t, U_t : t \in \mathbb{Z}/M\mathbb{Z}\},$$

where

$$T_t = \begin{pmatrix} \omega^t & 0 \\ 0 & \omega^{-t} \end{pmatrix}, \qquad U_t = T_t \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}.$$

For an automorphism ν and integer r, put

$$\mathfrak{a}_{\nu,r} = \{x \in \mathcal{O}_\mathbb{L} : \nu(x) = \omega^r x\}.$$

Let $r_2 = m/\gcd(m, 2)$ and define the compatibility indicators

$$\epsilon_{\mathrm{diag}}(t, \eta) = 1 \iff M \mid 2t(1 + s + \cdots + s^{m-1}),$$
$$\epsilon_{\mathrm{anti}}(t, \eta) = 1 \iff M \mid 2t(1 - s)(1 + s^2 + \cdots + s^{2(r_2-1)}).$$

For a lattice $C \subseteq \mathbb{L}^2$, all dimensions below mean $\dim_{F_\eta} \mathrm{span}_{F_\eta}(C)$, not the rank or dimension of an underlying integer lattice.

Theorem 19 (First-column classification). *In the power-of-two cyclotomic setting above, the first-column slices and their F_η-span dimensions are as follows.*

A and condition	$C_{A,\eta}$	$\dim_{F_\eta} \mathrm{span}_{F_\eta} C_{A,\eta}$
T_t, $\epsilon_{\mathrm{diag}} = 1$	$\mathcal{O}_0 e_1 \oplus \mathfrak{a}_{\eta,2t} e_2$	2
T_t, $\epsilon_{\mathrm{diag}} = 0$	$\mathcal{O}_0 e_1$	1
U_t, m *odd*	$\mathcal{O}_0 e_1 \oplus \mathfrak{a}_{\eta^2,2t(1-s)} e_2$	2
U_t, m *even*, $\epsilon_{\mathrm{anti}} = 0$	$\mathcal{O}_2 e_1$	2
U_t, m *even*, $\epsilon_{\mathrm{anti}} = 1$	$\mathcal{O}_2 e_1 \oplus \mathfrak{a}_{\eta^2,2t(1-s)} e_2$	4

For the nonidentity quotients $\eta = \sigma_i^{-1}\sigma_j$ arising from distinct challenge automorphisms, the ambient space $\mathbb{L}^2$ has F_η-dimension $2m$. Hence the largest slice has dimension ratio $2/m$, and a branch has no dimension reduction only when

$$A = U_t, \qquad m = 2, \qquad M \mid 2t(1 - s).$$

The entrywise systems, the Hilbert–90 nonvanishing argument, and the exact enumeration procedure appear in Appendix C.

Remark 20 (Conjugation outside the admissible family). Outside the admissible challenge family, a complex-conjugation quotient gives the public transition

$$J = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \qquad D = J\mathbf{Q}_j, \qquad D^*\mathbf{Q}_i D = \mathbf{Q}_j.$$

Indeed, $\mathbf{Q}_j = \overline{\mathbf{Q}_i}$ and $J^*\mathbf{Q}_i J = \mathbf{Q}_j^{-1}$ for a determinant-one rank-two Hermitian form. Thus $D \in \mathrm{SL}_2(\mathcal{O}_\mathbb{L})$ and the displayed identity follows. This boundary counterexample motivates excluding conjugation through (17).

4.4 Admissible challenge sets and the pairwise assumption

Lemma 21 (Capacity of the strict challenge family). *Let $n = 2^k$ with $k \geq 4$. Choose one representative from each $\Gamma[4]$-coset, taking* id *for the subgroup coset. The resulting set is an admissible challenge set of size*

$$q = \frac{|\Gamma|}{|\Gamma[4]|} = \frac{n}{8}.$$

Proof. For the power-of-two cyclotomic field, $\Gamma \simeq (\mathbb{Z}/2n\mathbb{Z})^\times \simeq C_2 \times C_{n/2}$, so $|\Gamma| = n$ and $|\Gamma[4]| = 2 \cdot 4 = 8$. Distinct coset representatives have quotient outside $\Gamma[4]$; since Γ is a 2-group, that quotient has order greater than four. □

The strict quotient rule removes the classified order-two and order-four branches. Hardness of the remaining correlated family is the following assumption.

Assumption 22 (Galois-derived pairwise module-LIP). *For Definition 16, including $\mathcal{D}_{\mathbf{B}}$, $\mathsf{Adm}_{>4}(\Sigma)$, and collision rejection, the value $\mathrm{Adv}^{\mathrm{pair}}_{\mathsf{Setup}}(\mathcal{A})$ is negligible for every probabilistic polynomial-time adversary $\mathcal{A}$.*

This is a correlated-instance assumption, not a consequence of ordinary module-LIP; the slice theorem characterizes its leakage but does not prove hardness.

5 A Galois-Derived Multi-Challenge Protocol

The base protocol uses the derived endpoint family as its challenge space and one re-randomized form as its first message. Its parallel repetition uses one such form per coordinate and a declared shared completion quota; quota exhaustion returns $\perp$ before a valid transcript is formed.

5.1 Relation and Protocol

For the index set $\mathcal{I}$ of definition 16, the common input is a well-formed $\mathsf{st} = (\mathbf{Q}_0, \Sigma)$ with $\mathbf{Q}_i = \sigma_i(\mathbf{Q}_0)$. Define the full-witness relation $\mathcal{R}_{\mathrm{full}}$ by requiring $\mathsf{wit} = (D_i)_{i \in \mathcal{I}}$ to satisfy

$$D_0 = I, \qquad D_i \in \mathrm{SL}_2(\mathcal{O}_{\mathbb{L}}), \qquad D_i^* \mathbf{Q}_0 D_i = \mathbf{Q}_i. \tag{20}$$

Define the extraction relation $\mathcal{R}_{\mathrm{pair}}$ by

$$\big(\mathsf{st}, (i, j, D)\big) \in \mathcal{R}_{\mathrm{pair}} \iff i \neq j, \quad D \in \mathrm{SL}_2(\mathcal{O}_{\mathbb{L}}), \quad D^* \mathbf{Q}_i D = \mathbf{Q}_j.$$

Protocol 23 (Galois-derived multi-challenge protocol).

1. The prover obtains $(\mathbf{T}, W)$ by bounded-retry re-randomization from $\mathbf{Q}_0$. If the declared batch quota is exhausted, it returns $\perp$.
2. The verifier samples $i \leftarrow \mathcal{I}$ uniformly.

3. The prover returns $Z_i = D_i^{-1} W$.
4. The verifier accepts when $Z_i \in \mathrm{SL}_2(\mathcal{O}_{\mathbb{L}})$ and $Z_i^* \mathbf{Q}_i Z_i = \mathbf{T}$.

For $\mu \geq 1$, the parallel protocol Π^μ independently executes the first step in each coordinate, using a shared completion quota K'. It sends $(\mathbf{T}_\ell)_{\ell \in [\mu]}$, receives a uniform challenge vector $(i_\ell)_{\ell \in [\mu]} \leftarrow \mathcal{I}^\mu$, returns $(Z_{\ell, i_\ell})_{\ell \in [\mu]}$, and accepts exactly when every coordinate accepts. The completion batch therefore contains $T = \mu$ independent equations.

5.2 Completeness, HVZK, and Extraction

For Π^μ, the real completion batch samples every coordinate from $\mathcal{M}(n, \sigma, \mathbf{Q}_0)$; let its supplied profile have exhaustion probability $\varepsilon^{\mathrm{real}}_{\mu, K'}$. For a fixed challenge vector $\boldsymbol{i} = (i_1, \ldots, i_\mu)$, the simulator instead samples coordinate ℓ from $\mathcal{M}(n, \sigma, \mathbf{Q}_{i_\ell})$; let the corresponding profile have exhaustion probability $\varepsilon^{\mathrm{sim}}_{\boldsymbol{i}, K'}$. Both quantities are computed by (14), but their single-equation laws need not agree.

Theorem 24 (Completeness and honest-verifier zero knowledge). *Under the re-randomization law of lemma 3, verification conditioned on non-abort is perfectly complete. Overall completeness error is at most $\varepsilon^{\mathrm{real}}_{\mu, K'}$. A simulator that performs the same bounded-retry re-randomization independently from each challenged $\mathbf{Q}_{i_\ell}$ has statistical distance at most the sum of the real and simulator batch tails,*

$$\Delta_{\mathrm{SD}}(\mathsf{Real}, \mathsf{Sim}) \leq \varepsilon^{\mathrm{real}}_{\mu, K'} + \varepsilon^{\mathrm{sim}}_{\boldsymbol{i}, K'}. \tag{21}$$

Proof. For each coordinate of a non-abort transcript,

$$Z_i^* \mathbf{Q}_i Z_i = W^* D_i^{-*} \mathbf{Q}_i D_i^{-1} W = W^* \mathbf{Q}_0 W = \mathbf{T}.$$

The formal real and simulated joint laws agree by basepoint covariance. Applying theorem 10 to the real and simulated bounded-retry batches and using the triangle inequality gives the stated loss. □

Theorem 25 (Pairwise and full-witness extraction). *For accepting transcripts with the same first message,*

$$\begin{aligned} &\text{(pairwise)} \quad (\mathbf{T}, i, Z_i), (\mathbf{T}, j, Z_j),\ i \neq j \Longrightarrow D_{i \to j} = Z_i Z_j^{-1}, \\ &\text{(full witness)} \ \{(\mathbf{T}, i, Z_i) : i \in \mathcal{I}\} \Longrightarrow \widehat{D}_0 = I, \quad \widehat{D}_i = Z_0 Z_i^{-1}. \end{aligned}$$

The first output satisfies $D_{i \to j}^* \mathbf{Q}_i D_{i \to j} = \mathbf{Q}_j$; the second satisfies the full relation (20).

Proof. The verifier equations imply

$$(Z_i Z_j^{-1})^* \mathbf{Q}_i (Z_i Z_j^{-1}) = Z_j^{-*} \mathbf{T} Z_j^{-1} = \mathbf{Q}_j.$$

Taking the ordered pair $(0, i)$ gives $\widehat{D}_i = Z_0 Z_i^{-1}$ and hence the full tuple. □

Thus the protocol has relaxed special soundness from $\mathcal{R}_{\mathrm{full}}$ to $\mathcal{R}_{\mathrm{pair}}$ in the sense fixed in section 2. It does not have two-transcript special soundness for the full-witness relation.

5.3 Soundness

In the impersonation experiment for Π^{μ}, Setup gives only st to a malicious prover $\mathcal{P}^*$. The prover sends μ first messages, receives an independent uniform challenge vector in $\mathcal{I}^{\mu}$, and wins when every coordinate verifies. Write $\mathrm{Adv}_{\mathrm{imp}}(\mathcal{P}^*)$ for this acceptance probability. For $q = |\mathcal{I}|$, the two-fork argument in Appendix D gives the following bound.

Theorem 26 (Parallel impersonation bound). *If $\delta_{\mathrm{pair}}(t_{\mathcal{B}}, \lambda)$ bounds the pairwise advantage at the reduction time, then*

$$\mathrm{Adv}_{\mathrm{imp}} \le q^{-\mu} + \sqrt{\delta_{\mathrm{pair}}(t_{\mathcal{B}}, \lambda)}. \tag{22}$$

The bound separates an information-theoretic challenge-guessing term from the correlated-instance hardness term. Neither is estimated experimentally.

5.4 Conditional Parameter Selection

Let λ_{snd} and λ_{zk} be the target soundness and statistical-HVZK levels. Splitting each target between the two terms in (22) and between the real and simulated batches, it suffices that

$$\begin{aligned} q^{-\mu} &\le 2^{-(\lambda_{\mathrm{snd}}+1)}, & \delta_{\mathrm{pair}} &\le 2^{-2(\lambda_{\mathrm{snd}}+1)}, \\ \varepsilon^{\mathrm{real}}_{\mu,K'} &\le 2^{-(\lambda_{\mathrm{zk}}+1)}, & \sup_{\boldsymbol{i}\in\mathcal{I}^{\mu}} \varepsilon^{\mathrm{sim}}_{\boldsymbol{i},K'} &\le 2^{-(\lambda_{\mathrm{zk}}+1)}. \end{aligned} \tag{23}$$

The completion batch has $T = \mu$ equations. Using the strict challenge family of lemma 21, set $q = n/8$ and

$$\mu = \left\lceil \frac{\lambda_{\mathrm{snd}} + 1}{\log_2 q} \right\rceil. \tag{24}$$

Table 2 applies the conservative supplied first-hit law homogeneously to all coordinates and reports the minimal shared quota; in each row the tail at $K' - 1$ exceeds the target. The values become protocol HVZK parameters only if the same law uniformly dominates every real and simulated endpoint distribution.

Table 2. Conditional protocol accounting. The shared quota uses the homogeneous conservative first-hit law and requires endpoint-uniform validity.

n	q	security bits $\lambda_{\mathrm{snd}} = \lambda_{\mathrm{zk}}$	repetitions μ	batch target $\varepsilon^{\mathrm{side}}_{\mu,K'}$	shared quota K'
512	64	128	22	2^{-129}	65
1024	128	256	37	2^{-257}	140

At these parameters, the pairwise assumption must give advantage at most 2^{-258} and 2^{-514}, respectively. Applying the completion phase medians to the

full shared quota gives completion-only batch estimates of 0.316–0.530 seconds for $n = 512$ and 2.888–5.322 seconds for $n = 1024$; the full accounting is in section E.

6 Conclusion

We resolved the two obstacles that prevent compact multi-challenge use of module-LIP re-randomization. First, field-norm completion can be used without biasing the primitive-column law: the bounded construction preserves the formal joint distribution up to an explicit abort event and admits exact batch accounting, while deterministic completion is available as a zero-failure comparison. Second, a full challenge family can be derived from one form by Galois action once the correlation exposed by pairwise transitions is made explicit and the degenerate quotients are removed. Together these results yield the compact public-coin protocol of section 5.

A concrete instantiation still needs a reduction or independent analysis of the correlated pairwise module-LIP assumption and fixed-shear tail bounds uniform over all real and simulated endpoints. The experiments here evaluate the completion component and first-hit model rather than the full re-randomization pipeline.

References

1. Ducas, L., Postlethwaite, E.W., Pulles, L.N., van Woerden, W.: Hawk: Module LIP makes lattice signatures fast, compact and simple. In: Advances in Cryptology – ASIACRYPT 2022. Lecture Notes in Computer Science, vol. 13794, pp. 65–94. Springer (2022), iACR ePrint 2022/1155
2. Hanzlik, L., Lai, Y.F., Paracucchi, E., Persichetti, E.: Wombat: Post-quantum blind signature from standard group action assumptions and more. Cryptology ePrint Archive, Paper 2026/928 (2026)
3. Jiang, K., Wang, A., Luo, H., Liu, G., Tang, G., Pan, Y., Wang, X.: Re-randomize and extract: A novel commitment construction framework based on group actions. In: Advances in Cryptology – EUROCRYPT 2025. Springer (2025), iACR ePrint 2025/400
4. Jiang, K., Wang, A., Luo, H., Liu, G., Yu, Y., Wang, X.: Exploiting the symmetry of $\mathbb{Z}^n$: Randomization and the automorphism problem. In: Advances in Cryptology – ASIACRYPT 2023. pp. 167–200. Springer (2023), iACR ePrint 2023/1735
5. Luo, H., Jiang, K., Jin, R., Pan, Y., Wang, A.: Commitment schemes based on module-LIP (2025), cryptology ePrint Archive, Paper 2025/431
6. Luo, H., Jiang, K., Pan, Y., Wang, A.: Cryptanalysis of rank-2 module-LIP with symplectic automorphisms. Cryptology ePrint Archive, Paper 2024/1173 (2024)
7. Massart, P.: The tight constant in the Dvoretzky–Kiefer–Wolfowitz inequality. The Annals of Probability **18**(3), 1269–1283 (1990). https://doi.org/10.1214/aop/1176990746
8. Pornin, T., Prest, T.: More efficient algorithms for the NTRU key generation using the field norm. In: Public-Key Cryptography – PKC 2019. pp. 504–533. Springer International Publishing (2019), iACR ePrint 2019/015

A Proof of the Worst-Case Completion Theorem

This appendix proves Theorem 11. The construction is a zero-failure comparison rather than the distribution-preserving protocol API: it converts a collective norm condition into one norm-coprime shear and then uses a single field-norm solve.

A.1 The deterministic hitting set

We assume $n \geq 2$; the cryptographic parameters are far larger. Recall

$$A_{\mathrm{hit}} = \{0, 1, X, \ldots, X^{n-2}\}, \qquad |A_{\mathrm{hit}}| = n.$$

Theorem 27 (Deterministic monomial hitting set). *For every primitive pair $f, g \in \mathcal{O}_{\mathbb{L}}$,*

$$\gcd\Big(|N(g)|, \{|N(f + ag)| : a \in A_{\mathrm{hit}}\}\Big) = 1. \tag{25}$$

Equivalently, there is no rational prime p such that $p \mid N(g)$ and $p \mid N(f + ag)$ for every $a \in A_{\mathrm{hit}}$.

Proof. Fix first an odd prime p. With $d_p = \mathrm{ord}_{2n}(p)$ and $r_p = n/d_p$, reduction modulo p gives

$$\mathcal{O}_{\mathbb{L}}/p\mathcal{O}_{\mathbb{L}} \simeq \prod_{j=1}^{r_p} \mathbb{F}_{p^{d_p}}.$$

For $h \in \mathcal{O}_{\mathbb{L}}$, the residue of $N(h)$ modulo p is the determinant of multiplication by h on this product. Thus

$$p \mid N(h) \quad \Longleftrightarrow \quad h_j = 0 \text{ in at least one component } j. \tag{26}$$

Let θ_j be the image of X in component j; it has order $2n$. Primitivity ensures that $(f_j, g_j) \neq (0, 0)$. As in (11), write

$$\rho_j = \begin{cases} -f_j/g_j, & g_j \neq 0, \\ \infty, & g_j = 0. \end{cases}$$

A component with $g_j = 0$ can witness $p \mid N(g)$, but primitivity gives $f_j \neq 0$, so it cannot make $f_j + a_j g_j$ vanish for any a. When $g_j \neq 0$, the equality $f_j + a_j g_j = 0$ is equivalent to $a_j = \rho_j$.

The component values of the n elements of A_{hit} are

$$0, 1, \theta_j, \theta_j^2, \ldots, \theta_j^{n-2}.$$

They are distinct: zero is separate from every power, and the displayed powers have distinct exponents modulo the order $2n$. Consequently, a component with

$g_j \neq 0$ can cover at most one candidate a, where cover means that $f + ag$ vanishes in that component.

Suppose that p divided $N(g)$ and all $N(f + ag)$. At least one of the r_p components would be consumed by $g_j = 0$, and that component covers no candidate. The remaining at most $r_p - 1 \leq n - 1$ components could cover at most $n - 1$ of the n candidates, contradicting (26) for all $a \in A_{\text{hit}}$.

It remains to consider $p = 2$. Here

$$\mathcal{O}_{\mathbb{L}}/2\mathcal{O}_{\mathbb{L}} \simeq \mathbb{F}_2[X]/((X+1)^n)$$

is local. If $2 \mid N(g)$, then the residue of g is a nonunit. Primitivity forces the residue of f to be a unit. The nonunits form the maximal ideal, so ag is a nonunit and $f + ag$ is a unit for every a. Hence $2 \nmid N(f + ag)$ for every candidate. If $2 \nmid N(g)$, the prime clearly cannot divide the collective gcd. No prime divides every integer in (25), proving the claim. □

A.2 Combining the local choices into one shear

Assume first that $g \neq 0$ and put $M = |N(g)|$. Order $A_{\text{hit}} = \{a_1, \ldots, a_n\}$ publicly and set

$$d_i = \gcd\big(M, |N(f + a_i g)|\big). \tag{27}$$

By theorem 27,

$$\gcd(d_1, \ldots, d_n) = 1. \tag{28}$$

We next separate prime supports without factoring M.

Lemma 28 (Coprime-support split). *Let U, d be positive integers with $U > 1$, let $L_U = \lceil \log_2(U+1) \rceil$, and define*

$$z = \gcd(U, d), \qquad s = \gcd\big(U, z^{L_U}\big), \qquad m = U/s.$$

Then $U = ms$, $\gcd(m, s) = \gcd(m, d) = 1$, and a prime divides s exactly when it divides both U and d. The value s can be computed from $z^{L_U} \bmod U$, without factoring either input.

Proof. Fix $p^e \| U$. If $p \nmid d$, then $p \nmid z$, hence $p \nmid s$ and $p^e \mid m$. If $p \mid d$, then $p \mid z$ and $v_p(z^{L_U}) \geq L_U \geq e$, so $p^e \mid s$ and $p \nmid m$. Thus m and s contain complementary full prime-power factors of U, which proves every claim. Modular exponentiation gives the stated computation. □

Starting with $U_0 = M$, apply lemma 28 successively to (U_{i-1}, d_i) and write

$$U_{i-1} = m_i U_i, \qquad \gcd(m_i, U_i) = \gcd(m_i, d_i) = 1. \tag{29}$$

If $U_{i-1} = 1$, put $m_i = U_i = 1$. A prime survives in U_i exactly when it survived in U_{i-1} and divides d_i. Equation (28) therefore gives $U_n = 1$. Consequently, the nontrivial m_i are pairwise coprime and

$$M = \prod_{i=1}^{n} m_i. \tag{30}$$

For every i with $m_i > 1$, let $t_i \in [0, m_i)$ satisfy

$$t_i \frac{M}{m_i} \equiv 1 \pmod{m_i},$$

and define the usual CRT idempotent

$$e_i = \frac{M}{m_i} t_i, \qquad 0 \le e_i < M. \tag{31}$$

Set

$$h = \sum_{i:m_i>1} e_i a_i \in \mathcal{O}_{\mathbb{L}}. \tag{32}$$

Then $h \equiv a_i \pmod{m_i \mathcal{O}_{\mathbb{L}}}$ for every nontrivial m_i.

Lemma 29 (The combined shear is norm-coprime). *The element h of (32) satisfies*

$$\gcd\big(|N(f + hg)|, M\big) = 1.$$

Proof. Let $p \mid M$. By (30), there is a unique i with $p \mid m_i$. Since $\gcd(m_i, d_i) = 1$, we have $p \nmid N(f + a_i g)$. Moreover, $h \equiv a_i \pmod{p\mathcal{O}_{\mathbb{L}}}$, so multiplication by $f + hg$ and by $f + a_i g$ have congruent determinant modulo p. Hence

$$N(f + hg) \equiv N(f + a_i g) \not\equiv 0 \pmod{p}.$$

No prime divisor of M divides $N(f + hg)$. □

Run

$$(F', G') \leftarrow \mathsf{TowerSolve}(f + hg, g).$$

This call is admissible by lemma 29. Define

$$F = F' - hG', \qquad G = G'. \tag{33}$$

Then

$$fG - gF = (f + hg)G' - gF' = 1.$$

If $g = 0$, primitivity makes f a unit and a unit inverse gives the completion. If $M = 1$, the original pair is already admissible for TowerSolve. This proves the unconditional correctness part of theorem 11.

A.3 Bit complexity

Assume that every input coefficient has magnitude below 2^B and put $b = B + \lceil \log_2 n \rceil + 1$. The integer $M = |N(g)|$ has $O(nb)$ bits. Since each a_i is zero or a monomial, every $f + a_i g$ still has coefficient bit length $O(B)$; one norm probe therefore costs $\widetilde{O}(nb)$ bit operations. All n probes cost $\widetilde{O}(n^2 b)$.

Every integer in the support-allocation procedure has $O(nb)$ bits. A split uses two gcds and one modular exponentiation with an $O(\log(nb))$-bit exponent, so all

n splits cost $\widetilde{O}(n^2 b)$. The same bound covers the n CRT inversions and products. Because the nonzero candidates a_i are distinct monomials, each coefficient of h is either zero or one e_i; it is therefore bounded by M and has $O(nb)$ bits. Forming $u = f + hg$ by fast negacyclic multiplication costs $\widetilde{O}(n^2 b)$. The logarithm of the Euclidean norm of (u, g) is $O(nb)$. Under the Pornin–Prest size condition fixed in section 2, the single TowerSolve call therefore costs

$$\widetilde{O}(n \cdot nb) = \widetilde{O}(n^2 b)$$

by [8, Lemmas 3–4]. The multiplication hG' in (33) has the same bound. Thus the whole algorithm runs in

$$\widetilde{O}(n^2 b) = \widetilde{O}\big(n^2 (B + \log n)\big)$$

bit operations, as claimed.

B Local Occupancy Evaluator, Prime Tails, and Quota Computation

This appendix expands the finite-cutoff calculation behind Equation 13. Every probability in the first three subsections is conditional on Assumption 12; the algebraic coverage identity (12) itself is unconditional.

B.1 A fixed-prime occupancy DP

Fix an odd prime p, a probe count $K \le 2n$, and abbreviate

$$r = r_p, \qquad Q = q_p + 1.$$

Under the local model, r independent balls are thrown uniformly into the Q projective labels after the componentwise normalization specified before Assumption 12. We distinguish the label ∞, the K normalized probe labels, and all remaining labels.

Let $u_m(s)$ be the probability that, after m component ratios, ∞ has not appeared and exactly s of the K finite labels have appeared. Let $v_m(s)$ be the corresponding probability after ∞ has appeared. Initialize

$$u_0(0) = 1, \qquad u_0(s) = v_0(s) = 0 \quad (s \neq 0).$$

At each step initialize the next arrays to zero and, for $0 \le s \le K$, make the updates

$$u_{m+1}(s) \mathrel{+}= u_m(s)\frac{Q - K - 1 + s}{Q}, \quad u_{m+1}(s+1) \mathrel{+}= u_m(s)\frac{K - s}{Q}, \tag{34}$$

$$v_{m+1}(s) \mathrel{+}= u_m(s)\frac{1}{Q}, \qquad v_{m+1}(s) \mathrel{+}= v_m(s)\frac{Q - K + s}{Q}, \tag{35}$$

$$v_{m+1}(s+1) \mathrel{+}= v_m(s)\frac{K - s}{Q}. \tag{36}$$

Terms with index $K+1$ are omitted. The first line separates an unseen finite label from a label that is irrelevant or already occupied. The first term in (35) is the first hit at ∞; after that hit, the label ∞ is included among the labels that do not change s.

Let

$$W_p(s) = \Pr[|C_p(K)| = s].$$

The contamination set is empty when ∞ is absent, irrespective of the finite occupancy, and also when ∞ is present but no finite probe label is hit. Hence

$$W_p(0) = \sum_{s=0}^{K} u_r(s) + v_r(0), \qquad W_p(s) = v_r(s) \quad (1 \le s \le K). \tag{37}$$

By symmetry, conditioned on $|C_p(K)| = s$, the set $C_p(K)$ is uniform among the s-subsets of $[K]$.

For completeness, one exact coverage DP keeps a distribution $D_\ell(U)$ on subsets $U \subseteq [K]$ after the first ℓ primes. Starting from $D_0(\varnothing) = 1$, its update for the next prime p is

$$D_{\ell+1}(U) = \sum_{\substack{V,C \subseteq [K] \\ V \cup C = U}} D_\ell(V) \frac{W_p(|C|)}{\binom{K}{|C|}}. \tag{38}$$

After all primes in $\mathcal{S}$ have been processed, $D_{|\mathcal{S}|}([K]) = P_{\text{fail}}^{(\mathcal{S})}(K, n)$. This form is conceptually direct but exponential in K; the symmetric inclusion–exclusion formula is preferable for the reported parameters.

B.2 Inclusion–exclusion cross-check

For a fixed $J \subseteq [K]$ of size j, the event $C_p(K) \cap J = \varnothing$ is the union of two events: ∞ is absent, or all labels in J are absent. Inclusion–exclusion on these events gives

$$\begin{aligned} a_p(j) &= \Pr[C_p(K) \cap J = \varnothing] \\ &= \left(1 - \frac{1}{Q}\right)^r + \left(1 - \frac{j}{Q}\right)^r - \left(1 - \frac{j+1}{Q}\right)^r. \end{aligned} \tag{39}$$

Now let $H_i = [i \notin \bigcup_{p \in \mathcal{S}} C_p(K)]$ be the event that probe i is a hole after all primes in $\mathcal{S}$. Failure means that there are no holes. Inclusion–exclusion over the H_i, together with prime independence inside the finite model, yields

$$\begin{aligned} P_{\text{fail}}^{(\mathcal{S})}(K, n) &= \sum_{J \subseteq [K]} (-1)^{|J|} \Pr[H_i \text{ for every } i \in J] \\ &= \sum_{j=0}^{K} (-1)^j \binom{K}{j} \prod_{p \in \mathcal{S}} a_p(j), \end{aligned} \tag{40}$$

which is (13). At small K, agreement between (38) and (40) is an independent implementation check.

B.3 Prime classes and omitted-prime scale

For one fixed probe i, define the local contamination probability

$$\gamma_p = \Pr[i \in C_p(K)].$$

The labels ∞ and $\theta_{p,j}^{i-1}$ are distinct in component j, so

$$\gamma_p = 1 - 2\left(1 - \frac{1}{Q}\right)^r + \left(1 - \frac{2}{Q}\right)^r \le \frac{r(r-1)}{Q^2} \le \left(\frac{n}{d_p p^{d_p}}\right)^2. \tag{41}$$

The first inequality is the union bound over ordered choices of two distinct components, one hitting ∞ and one hitting the finite label.

More generally, for a fixed u-set $J \subseteq [K]$, let

$$\tau_{p,u} = \Pr[J \subseteq C_p(K)].$$

All u finite labels and ∞ must be occupied. Inclusion–exclusion on these $u+1$ labels gives

$$\tau_{p,u} = \sum_{a=0}^{1}\sum_{b=0}^{u}(-1)^{a+b}\binom{u}{b}\left(1 - \frac{a+b}{Q}\right)^r, \tag{42}$$

and the ordered-component bound gives

$$\tau_{p,u} \le \frac{(r)_{u+1}}{Q^{u+1}} \le \left(\frac{n}{d_p p^{d_p}}\right)^{u+1}, \tag{43}$$

where $(r)_{u+1} = r(r-1)\cdots(r-u)$.

Split primes. If $p \equiv 1 \pmod{2n}$, then $d_p = 1$, $r_p = n$, and for p on the scale $p \asymp n$,

$$\gamma_p \approx 1 - 2e^{-n/p} + e^{-2n/p} = (1 - e^{-n/p})^2. \tag{44}$$

Thus the first few split primes can contribute at constant scale and control the observable slopes.

Non-split primes. For odd non-split p, $d_p > 1$. Since $(\mathbb{Z}/2n\mathbb{Z})^\times$ is a 2-group, d_p is a power of two. The order-two residues modulo $2n$ have least positive representatives $n-1, n+1, 2n-1$, so

$$p^{d_p/2} \ge n - 1, \qquad p^{d_p} \ge (n-1)^2.$$

Substitution in (41) yields

$$\gamma_p \le \frac{n^2}{d_p^2(n-1)^4} = O\left(\frac{1}{d_p^2 n^2}\right). \tag{45}$$

Writing

$$A_{u,\mathrm{ns}} = \sum_{\substack{p \text{ odd} \\ d_p > 1}} \tau_{p,u},$$

the residue-class summation gives the conservative bound

$$A_{u,\mathrm{ns}} \le \frac{4}{2^u - 1} \frac{n^{u+1}}{(n-1)^{2u+2}} = O(n^{-u-1}). \tag{46}$$

This explains why the finite evaluator focuses on the split-prime head.

The ramified prime. The prime 2 covers no probes for a primitive pair. If $2 \mid N(g)$, the image of g in the local ring $\mathbb{F}_2[X]/((X+1)^n)$ is a nonunit, while primitivity forces f to be a unit. Hence every $f + X^i g$ is a unit and $2 \nmid N(f + X^i g)$.

Large-prime estimates. For omitted split primes beyond B, the finite-cutoff analysis uses the tail-scale estimates

$$A_{1,\mathrm{split}}(B) = \sum_{\substack{p > B \\ p \equiv 1 \ (2n)}} \gamma_p \approx \frac{n}{B \log B}, \tag{47}$$

$$A_{u,\mathrm{split}}(B) = \sum_{\substack{p > B \\ p \equiv 1 \ (2n)}} \tau_{p,u} \approx \frac{n^u}{u B^u \log B}. \tag{48}$$

For large non-split primes, $d_p \ge 2$ and

$$A_{1,\text{large ns}}(B) \le \frac{n^2}{12 B^3 \log B}, \tag{49}$$

which is smaller than (47) by roughly $n/(12B^2)$. Equations (47) and (48) are tail-scale estimates, not a replacement for the finite-cutoff label on a reported probability.

There is also a necessary finite-integer caveat. The local approximation $\Pr[p \mid N(g)] \approx n/p$ for a split prime has a divergent sum over all primes, but the integer $N(g)$ has finitely many prime divisors. If $L = \log_2 |N(g)|$, then every such divisor is at most 2^L. More importantly, fast-path contamination requires a *double* hit—division of both $N(g)$ and one shifted norm—whose scale n^2/p^2 is summable. The tail analysis must therefore be attached to the coverage event, not to an imagined infinite independent list of divisors of $N(g)$.

B.4 Head-DP correction

Let $\mathcal{H}$ be the primes retained by the evaluator and let $\mathcal{T}$ be an omitted set. Write

$$P_{\mathcal{H}} = \Pr\left[\bigcup_{p \in \mathcal{H}} C_p(K) = [K] \right],$$

and let $D_{\mathcal{H}}(K-h)$ be the probability that the head union has size $K-h$, i.e., leaves exactly h holes. By exchangeability, let T_h be the probability that the

omitted primes cover one fixed h-set. Under the head/tail independence of the evaluator,

$$P_{\mathcal{H}\cup\mathcal{T}} = P_{\mathcal{H}} + \sum_{h=1}^{K} D_{\mathcal{H}}(K-h)T_h. \tag{50}$$

For

$$A_u(\mathcal{T}) = \sum_{p\in\mathcal{T}} \tau_{p,u},$$

a set-partition union bound gives

$$T_h \le \sum_{\pi\in\Pi_h} \prod_{S\in\pi} A_{|S|}(\mathcal{T}), \tag{51}$$

where Π_h is the set of partitions of a fixed h-element hole set. A block represents holes assigned to one contaminating prime; summing over partitions and primes can only overcount.

The resulting logarithmic correction satisfies

$$\log_2 P_{\mathcal{H}\cup\mathcal{T}} - \log_2 P_{\mathcal{H}} \le \log_2\left(1 + \frac{\sum_{h=1}^{K} D_{\mathcal{H}}(K-h)\sum_{\pi\in\Pi_h}\prod_{S\in\pi} A_{|S|}(\mathcal{T})}{P_{\mathcal{H}}}\right). \tag{52}$$

Its first-order term is

$$\frac{1}{\ln 2}\frac{D_{\mathcal{H}}(K-1)}{P_{\mathcal{H}}}A_1(\mathcal{T}), \tag{53}$$

with (47) supplying the leading omitted-prime scale.

B.5 From model tails to a quota without silent truncation

For every evaluated k, set

$$\mathcal{F}^{(\le B)}(k) = P_{\text{fail}}^{(\le B)}(k,n), \qquad a_r^{(\le B)} = \mathcal{F}^{(\le B)}(r-1) - \mathcal{F}^{(\le B)}(r).$$

The first-hit mass must be checked for nonnegativity, and the tail must be checked for monotonicity. These are numerical consistency tests on the evaluator, not additional assumptions.

There are two distinct notions of "full tail." Mathematically, the fixed unit-root law has support in $\{1,\dots,2n\}\cup\{\infty\}$, and evaluating through $2n$ would identify the remaining atom as $\mathcal{F}^{(\le B)}(2n)$. A particular quota calculation need not evaluate that far. To compute

$$\Pr[S_T \le K'] = \sum_{s=T}^{K'} [z^s]A(z)^T,$$

one needs only those a_r that can occur in a coefficient of degree at most K' (in particular, values through $r = K'$ are sufficient). All uncomputed finite-tail mass

and any atom at ∞ remain automatically in $1 - \Pr[S_T \le K']$. Thus the CDF at a reported K' is exact for the specified finite-cutoff input model whenever the evaluator supplies every needed degree; it is not necessary to normalize the known a_r to sum to one or to claim that the infinite atom was numerically evaluated.

If instead an implementation deliberately censors all values above some $R < K'$ as ∞, the censored CDF is conservative and its distance from the uncensored CDF is at most

$$\Pr[\exists j : \tau_j > R] \le T\mathcal{F}(R), \tag{54}$$

as stated in (16). This residual must accompany any quota whose needed coefficient range exceeds the computed tail.

Finally, inclusion–exclusion near 2^{-256} contains severe cancellation. The evaluator must therefore use arbitrary precision (or a verified log-domain equivalent), and the convolution must retain comparable precision. Useful checks are: agreement of subset DP and inclusion–exclusion at small K; monotonicity of $\mathcal{F}$; nonnegativity of a_r; agreement of direct polynomial multiplication and the recurrence (15); and the minimality check that the returned K' meets the target while $K' - 1$ does not. Passing these checks establishes the numerical transformation of the model input, not the empirical validity of Assumption 12.

C Full Cyclotomic SL_2 Slice Classification

This appendix proves the matrix calculations summarized in theorem 19. The monomial description of $\mathrm{SU}_2(\mathcal{O}_\mathbb{L})$ follows from the rank-two integral unitary classification [6, Lemma 3.1]; this appendix derives the remaining slice equations and dimensions in the power-of-two cyclotomic setting.

Retain the notation

$$\mu(\mathbb{L}) = \langle\omega\rangle, \quad |\mu(\mathbb{L})| = M, \quad \eta(\omega) = \omega^s, \quad m = \mathrm{ord}(\eta),$$

and

$$F_\eta = \mathbb{L}^{\langle\eta\rangle}, \quad F_{\eta^2} = \mathbb{L}^{\langle\eta^2\rangle}, \quad \mathcal{O}_0 = \mathcal{O}_\mathbb{L} \cap F_\eta, \quad \mathcal{O}_2 = \mathcal{O}_\mathbb{L} \cap F_{\eta^2}.$$

For $\nu \in \langle\eta\rangle$ and $r \in \mathbb{Z}$, recall

$$\mathfrak{a}_{\nu,r} = \{x \in \mathcal{O}_\mathbb{L} : \nu(x) = \omega^r x\}.$$

Whenever this eigenspace lattice is nonzero, its field span is one-dimensional over the fixed field of ν: if nonzero x and y obey the same equation, then $\nu(x/y) = x/y$. We use the following explicit cyclic Hilbert–90 form for the converse.

Lemma 30 (Cyclic semilinear eigenspace). *Let ν have order r on $\mathbb{L}$ and let $u \in \mathbb{L}^\times$. The equation $\nu(x) = ux$ has a nonzero solution $x \in \mathbb{L}$ if and only if $\prod_{j=0}^{r-1} \nu^j(u) = 1$. When a solution exists, its solution space is one-dimensional over $\mathbb{L}^{\langle\nu\rangle}$ and contains a nonzero element of $\mathcal{O}_\mathbb{L}$.*

Proof. Iteration proves necessity. Conversely, put $c_0 = 1$ and $c_i = \prod_{j=0}^{i-1} \nu^j(u)$, so $c_r = 1$. For a suitable $y \in \mathbb{L}$, the nonzero element

$$x = \sum_{i=0}^{r-1} c_i^{-1} \nu^i(y)$$

satisfies $\nu(x) = ux$ after reindexing; such a y exists because distinct field automorphisms are linearly independent. The quotient of any two nonzero solutions is fixed by ν, proving the dimension statement. Clearing a rational integer denominator places a nonzero solution in $\mathcal{O}_{\mathbb{L}}$. □

The compatibility conditions below are precisely the product-one conditions of lemma 30.

C.1 Diagonal monomial branch

Let

$$A = T_t = \begin{pmatrix} \xi & 0 \\ 0 & \xi^{-1} \end{pmatrix}, \qquad \xi = \omega^t, \qquad P = \begin{pmatrix} a & b \\ c & d \end{pmatrix}.$$

Expanding $PA = A\eta(P)$ gives

$$a = \eta(a), \qquad d = \eta(d), \qquad \eta(b) = \omega^{-2t} b, \qquad \eta(c) = \omega^{2t} c. \tag{55}$$

Consequently,

$$\Lambda_{T_t,\eta} = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} : a, d \in \mathcal{O}_0,\ b \in \mathfrak{a}_{\eta,-2t},\ c \in \mathfrak{a}_{\eta,2t} \right\},$$

and

$$C_{T_t,\eta} = \mathcal{O}_0 e_1 \oplus \mathfrak{a}_{\eta,2t} e_2. \tag{56}$$

Iterating the last equation in (55) around the η-orbit yields

$$c = \eta^m(c) = \omega^{2t(1+s+\cdots+s^{m-1})} c.$$

Thus a nonzero second summand exists exactly when

$$M \mid 2t(1 + s + \cdots + s^{m-1}).$$

This is precisely $\epsilon_{\mathrm{diag}}(t, \eta) = 1$. If it holds, the two summands in (56) each have one-dimensional F_η-span; otherwise only $\mathcal{O}_0 e_1$ remains. The corresponding dimensions are therefore 2 and 1.

C.2 Anti-diagonal monomial branch

Now let

$$A = U_t = \begin{pmatrix} 0 & \xi \\ -\xi^{-1} & 0 \end{pmatrix}.$$

The equation $PA = A\eta(P)$ is equivalent to

$$d = \eta(a), \qquad a = \eta(d), \qquad b = -\omega^{2t}\eta(c), \qquad \eta^2(c) = \omega^{2t(1-s)}c. \tag{57}$$

Hence $a \in \mathcal{O}_2$, and every solution has the form

$$P(a,c) = \begin{pmatrix} a & -\omega^{2t}\eta(c) \\ c & \eta(a) \end{pmatrix}, \quad a \in \mathcal{O}_2, \quad c \in \mathfrak{a}_{\eta^2,2t(1-s)}.$$

In particular,

$$C_{U_t,\eta} = \mathcal{O}_2 e_1 \oplus \mathfrak{a}_{\eta^2,2t(1-s)} e_2, \tag{58}$$

where a zero eigenspace means that the second summand is absent.

Put $g = \gcd(m,2)$ and $r_2 = m/g$, so η^2 has order r_2. Iterating the last equation of (57) gives the compatibility condition

$$M \mid 2t(1-s)(1 + s^2 + \cdots + s^{2(r_2-1)}).$$

When m is odd, the 2-group property of Γ forces $m = 1$. Consequently $\eta = \mathrm{id}$, $s = 1$, and the eigenspace in (58) is nonzero; both summands have one-dimensional F_η-span. When m is even, $[F_{\eta^2} : F_\eta] = 2$. The first summand then has F_η-dimension 2, and lemma 30 gives a second summand of F_η-dimension 2 exactly when $\epsilon_{\mathrm{anti}}(t,\eta) = 1$. This gives dimensions 2 and 4 in the two even-order branches.

C.3 Dimension comparison and quotient orders

Combining the two entry systems gives the complete table

branch	$C_{A,\eta}$	$\dim_{F_\eta} \mathrm{span}_{F_\eta}(C_{A,\eta})$
$T_t,\ \epsilon_{\mathrm{diag}} = 1$	$\mathcal{O}_0 e_1 \oplus \mathfrak{a}_{\eta,2t} e_2$	2
$T_t,\ \epsilon_{\mathrm{diag}} = 0$	$\mathcal{O}_0 e_1$	1
$U_t,\ m$ odd	$\mathcal{O}_0 e_1 \oplus \mathfrak{a}_{\eta^2,2t(1-s)} e_2$	2
$U_t,\ m$ even, $\epsilon_{\mathrm{anti}} = 0$	$\mathcal{O}_2 e_1$	2
$U_t,\ m$ even, $\epsilon_{\mathrm{anti}} = 1$	$\mathcal{O}_2 e_1 \oplus \mathfrak{a}_{\eta^2,2t(1-s)} e_2$	4.

These are dimensions of field spans; no dimension is assigned to an integer lattice itself.

For nonidentity quotients η arising from distinct challenge automorphisms, $[\mathbb{L} : F_\eta] = m$ and the ambient first-column space $\mathbb{L}^2$ has F_η-dimension $2m$. The largest classified slice has dimension 4, so its ratio to the ambient space is $2/m$. Equality with the full ambient dimension forces $m = 2$ and the nonzero anti-diagonal eigenspace condition. For $m = 2$ one has $r_2 = 1$, so that condition reduces to

$$M \mid 2t(1-s).$$

This proves the final assertion of theorem 19.

The admissibility predicate $\mathsf{Adm}_{>4}$ removes both the $m = 2$ branch and the $m = 4$ exceptional branch, whose F_η-dimension 4 is one half of the ambient dimension 8. Equivalently, Setup chooses one representative from each $\Gamma[4]$-coset. This is a classification-based exclusion only; it supplies no hardness statement for the surviving quotients.

C.4 Exact linear-system enumeration

The enumeration can be made fully explicit. Fix an integral basis $(r_1, \ldots, r_d)$ of $\mathcal{O}_\mathbb{L}$. Represent each of the four entries of X by d integer coordinates, represent η by its $d \times d$ integral action matrix, and represent multiplication by the entries of A and $\widetilde{D}$ in the same basis. The four entry equations in

$$XA - \widetilde{D}\eta(X) = 0$$

then form a homogeneous integer linear system in $4d$ variables. Its integral kernel and its first-column projection describe the candidate slice for that A. Repeating this for the $2M$ monomial candidates is polynomial in M, d, and the coefficient bit lengths under the exact-arithmetic conditions stated above. The output is a union of candidate lattices, not a recovered factorization.

D Full Protocol Proofs and Parameter Accounting

This appendix gives the direct-TowerSolve obstruction and expands the formal re-randomization hybrids, extraction calculation, and rewind loss used in section 5.

D.1 The direct-TowerSolve obstruction

Let $\mathcal{A}_{\mathrm{TS}}$ be the norm-coprime set in proposition 7. Consider the direct replacement that samples $x \leftarrow D_{\mathbf{Q},\sigma}$ until $x \in \mathcal{A}_{\mathrm{TS}}$, calls $\mathsf{TowerSolve}$ on x, applies the inherited HAWK post-completion map, and returns the endpoint together with its transporter. Complement independence in lemma 3 shows that the result does not depend on which determinant-one complement the solver returns. Thus the only change from formal re-randomization is the accepted first-column set.

We first record why this set is not an orbit-invariant replacement for primitivity. For a rational prime p, put

$$R_p = \mathcal{O}_\mathbb{L}/p\mathcal{O}_\mathbb{L}, \qquad S_p = (R_p^\times \times R_p) \cup (R_p \times R_p^\times).$$

The field norm is the determinant of multiplication on the underlying $\mathbb{Z}$-module, and hence

$$p \nmid N(a) \iff a \bmod p\mathcal{O}_\mathbb{L} \in R_p^\times. \tag{59}$$

Consequently,

$$x \in \mathcal{A}_{\mathrm{TS}} \iff x \bmod p\mathcal{O}_\mathbb{L} \in S_p \quad \text{for every rational prime } p. \tag{60}$$

If R_p is local with maximal ideal $\mathfrak{m}$, then $S_p = R_p^2 \setminus \mathfrak{m}^2$ is the set of unimodular columns and is preserved by $\mathrm{GL}_2(R_p)$. If R_p is not local, its finite Artinian decomposition contains a nontrivial idempotent e. Lift e to $\widetilde{e} \in \mathcal{O}_{\mathbb{L}}$ and set

$$E = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix} \in \mathrm{SL}_2(\mathcal{O}_{\mathbb{L}}), \qquad x = (1, \widetilde{e})^{\mathsf{T}}.$$

Then $x \in \mathcal{A}_{\mathrm{TS}}$ because its first coordinate is one, whereas

$$Ex = (1 - \widetilde{e}, \widetilde{e})^{\mathsf{T}} \notin \mathcal{A}_{\mathrm{TS}}.$$

Indeed, both e and $1 - e$ are nonunits in R_p, so p divides the norms of both coordinates of Ex by (59). Thus split rational primes can destroy invariance; the local factors cannot.

We now compute the transcript-level gap. Fix $D \in \mathsf{Trans}(\mathbf{Q}, \mathbf{Q}')$ and write $D^{-1}\mathcal{A}_{\mathrm{TS}} = \{D^{-1}x : x \in \mathcal{A}_{\mathrm{TS}}\}$. Let P_{real} and P_{sim} denote the joint endpoint/response laws in the real and simulated experiments. In a real execution of the direct replacement, let x be the accepted first column and let $z = D^{-1}x$ be the first column of the response transporter. Since D^{-1} is a bijection of $\mathcal{O}_{\mathbb{L}}^2$ and

$$z^* \mathbf{Q}' z = x^* \mathbf{Q} x,$$

the first-column marginal of the transported real response is

$$P_{\mathrm{real}}^{(1)} = D_{\mathbf{Q}',\sigma} \mid D^{-1}\mathcal{A}_{\mathrm{TS}}. \tag{61}$$

A simulator that applies the same direct replacement at $\mathbf{Q}'$ instead has

$$P_{\mathrm{sim}}^{(1)} = D_{\mathbf{Q}',\sigma} \mid \mathcal{A}_{\mathrm{TS}}. \tag{62}$$

Write $\mu' = D_{\mathbf{Q}',\sigma}$ and define the exact mismatch mass

$$\beta_D(\mathbf{Q}', \sigma) = \frac{\mu'\big(D^{-1}\mathcal{A}_{\mathrm{TS}} \setminus \mathcal{A}_{\mathrm{TS}}\big)}{\mu'(D^{-1}\mathcal{A}_{\mathrm{TS}})}. \tag{63}$$

The efficient predicate that accepts exactly when the public response's first column lies in $\mathcal{A}_{\mathrm{TS}}$ accepts the simulated law with probability one and the real law with probability $1 - \beta_D(\mathbf{Q}', \sigma)$. Therefore

$$\Delta_{\mathrm{SD}}(P_{\mathrm{real}}, P_{\mathrm{sim}}) \geq \beta_D(\mathbf{Q}', \sigma), \tag{64}$$

both for the response laws and for the full transcript laws. This proves proposition 7.

This conclusion uses the fact that a protocol transcript exposes the response transporter. If only the endpoint form were revealed, the many-to-one map $W \mapsto W^* \mathbf{Q} W$ could contract statistical distance, and the first-column test would not by itself give an endpoint-form distinguisher. Endpoint indistinguishability therefore neither repairs nor contradicts the transcript-level gap above; it simply addresses a different marginal.

The proposition is a barrier to a generic zero-knowledge proof, not a claim that (63) is non-negligible for every derived transporter. Establishing negligibility would be an additional distributional theorem not implied by hardness of distinguishing endpoint forms. Our bounded construction instead compares each execution with the primitive-column formal law, retains exhaustion as $\perp$, and charges the explicit coupling loss of theorem 10; it never renormalizes the protocol to the direct-TowerSolve conditional law.

D.2 Formal re-randomization and completeness

Write

$$(\mathbf{T}, W) \leftarrow \mathsf{ReRand}_{\mathrm{form}}(\mathbf{Q}), \qquad \mathbf{T} = W^{*}\mathbf{Q}W.$$

If $D^{*}\mathbf{Q}D = \mathbf{Q}'$, basepoint covariance states that

$$\begin{aligned}(\mathbf{T}, D^{-1}W), \quad (\mathbf{T}, W) &\leftarrow \mathsf{ReRand}_{\mathrm{form}}(\mathbf{Q}) \\ \equiv (\mathbf{T}, Z), \quad (\mathbf{T}, Z) &\leftarrow \mathsf{ReRand}_{\mathrm{form}}(\mathbf{Q}').\end{aligned} \tag{65}$$

This equality concerns the joint endpoint/transporter law. Representative independence permits the exact complement from theorem 9 to replace another completion of the same sampled column before the inherited post-completion map.

For the setup witness $D_i = \mathbf{B}^{-1}\sigma_i(\mathbf{B})$,

$$D_i^{*}\mathbf{Q}_0 D_i = \sigma_i(\mathbf{B})^{*}\sigma_i(\mathbf{B}) = \sigma_i(\mathbf{B}^{*}\mathbf{B}) = \mathbf{Q}_i,$$

where the last identity uses (2). For a non-abort response $Z_i = D_i^{-1}W$, one then has $Z_i^{*}\mathbf{Q}_i Z_i = W^{*}\mathbf{Q}_0 W = \mathbf{T}$.

D.3 Batched HVZK hybrid

For fixed i, a formal real transcript contains $(\mathbf{T}, i, D_i^{-1}W)$, where

$$(\mathbf{T}, W) \leftarrow \mathsf{ReRand}_{\mathrm{form}}(\mathbf{Q}_0).$$

The formal simulator instead samples

$$(\mathbf{T}, Z) \leftarrow \mathsf{ReRand}_{\mathrm{form}}(\mathbf{Q}_i)$$

and returns $(\mathbf{T}, i, Z)$. Equation (65) makes these laws identical.

For T independent completion equations, the supplied first-hit masses and their generating functions are those of definition 13. For any one specified batch profile, a shared quota K' succeeds with probability

$$\sum_{s=T}^{K'} [z^{s}] \prod_{j=1}^{T} A_j(z).$$

All uncomputed and censored mass remains on the failure side. The real profile uses T copies of the law at $\mathbf{Q}_0$, whereas for challenge vector $\boldsymbol{i}$ the simulated profile uses the laws at $\mathbf{Q}_{i_1}, \ldots, \mathbf{Q}_{i_T}$. Applying theorem 10 separately to these batches gives distance at most $\varepsilon^{\mathrm{real}}_{T,K'} + \varepsilon^{\mathrm{sim}}_{\boldsymbol{i},K'}$. Equality of these two tails is neither needed nor implied by formal re-randomization covariance, because the fixed-shear success event is not invariant under a general transporter. Thus a uniform protocol bound must be supplied for every endpoint profile; neither convolution nor the protocol proof strengthens its model or empirical provenance.

D.4 Extraction from endpoint collisions

Suppose two accepting transcripts share $\mathbf{T}$ and have challenges $i \neq j$. Then

$$Z_i^* \mathbf{Q}_i Z_i = \mathbf{T} = Z_j^* \mathbf{Q}_j Z_j.$$

For $D = Z_i Z_j^{-1}$,

$$D^* \mathbf{Q}_i D = Z_j^{-*} Z_i^* \mathbf{Q}_i Z_i Z_j^{-1} = Z_j^{-*} \mathbf{T} Z_j^{-1} = \mathbf{Q}_j.$$

This fixes both the multiplication order and the transition direction. Duplicate endpoints are rejected during setup, so the output is a nontrivial pairwise instance. It does not by itself recover a factorization of $\mathbf{Q}_0$. Responses for challenge zero and every i instead give $D_i = Z_0 Z_i^{-1}$, which is the full witness tuple.

D.5 Exact two-fork reduction

Let a possibly cheating prover face μ parallel coordinates and put $N = q^\mu$. Fix its state after the vector of first messages. If s of the N challenge vectors lead to acceptance, write $p = s/N$. Two independent challenges produce two accepting, distinct vectors with probability

$$\frac{s(s-1)}{N^2} = p^2 - \frac{p}{N}.$$

Taking expectations, if $\epsilon = \mathbb{E}[p]$ is the prover's acceptance probability and α is the reduction's pairwise success probability, then

$$\alpha = \mathbb{E}[p^2] - \frac{\mathbb{E}[p]}{N} \geq \epsilon^2 - \frac{\epsilon}{N}.$$

Solving this quadratic gives

$$\epsilon \leq \frac{N^{-1} + \sqrt{N^{-2} + 4\alpha}}{2} \leq N^{-1} + \sqrt{\alpha},$$

which proves theorem 26.

D.6 Operation counts

One T-equation execution with a shared quota performs at most K' separate useful norm probes, pads the public probe schedule to K' slots, and makes at most T calls to TowerSolve on the successful path. Matrix transports and verification add $O(T)$ ring-matrix operations. Each $g = 0$ equation consumes one logical slot without a norm/gcd computation or a TowerSolve call. A leakage-aware implementation must separately specify padding for the public probe schedule and the solver schedule. No batch-norm speedup is implied by these counts.

E Empirical Methodology and Reproducibility

The empirical study has two roles: assessing whether the finite-cutoff occupancy model describes observed first-hit behavior over the reported range, and characterizing the cost of the bounded Completion implementation. It supplies no evidence for the algebraic identities, the re-randomization interface, or the pairwise module-LIP assumption, and no experiment is used in a security proof.

E.1 Sampling design

We sample primitive pairs with centered integer discrete-Gaussian coefficients under the three variance profiles n, $n^{3/2}$, and n^2. Calibration and holdout use disjoint seed and candidate domains. Each record stores the first successful public shear up to its declared horizon. If no success is observed, the record is right-censored and remains on the failure side at every available cutoff.

The supplementary $K \le 5$ study uses the fixed, previously recorded $n \in \{16, 32, 64\}$ cohorts. It neither refits the model nor selects a stopping rule from the observed outcomes. Its current reuse-only manifest references the immutable R1 sample authority and reuses 1,650,000 accepted observations; it generates no new samples and invokes no sampling validator. The accompanying cellwise CSV records the sample and failure counts underlying every displayed probability.

E.2 Model comparison

The evaluator computes equation (13) in arbitrary precision and is frozen before comparison with empirical data. At $n = 512, 1024$, the reported evaluator retains exactly the split primes $p \equiv 1 \pmod{2n}$ up to $3 \cdot 10^7$; omitted primes are treated only by the separate tail-scale discussion in section B. For each cell and cutoff it reports

$$\widehat{F}(K) = \frac{\#\{\tau > K\}}{N}, \qquad F_{\mathrm{model}}(K), \qquad F_{\mathrm{cons}}(K) = cF_{\mathrm{model}}(K),$$

where $c = 2.68667279036$. The multiplier is selected on calibration data over the primary range

$$n \in \{16, 32, 64, 128, 256\}$$

and the three variance profiles n, $n^{3/2}$, and n^2. These five dimensions and three profiles form fifteen calibration cells. For one cell with N independent samples, the one-sided Dvoretzky–Kiefer–Wolfowitz bound gives [7]

$$\Pr\left[\exists K:\ F(K) > \widehat{F}(K) + \epsilon\right] \le \exp(-2N\epsilon^2).$$

Only this upper deviation is needed: the calibration constructs an upper confidence band for the failure probability, not a two-sided confidence band. To obtain simultaneous coverage with family-wise error $\alpha = 0.05$, we assign equal error probability to the fifteen cells and set

$$\alpha_{\text{cell}} = \frac{0.05}{15}, \qquad \epsilon_N = \sqrt{\frac{\log(1/\alpha_{\text{cell}})}{2N}}.$$

Indeed, substituting ϵ_N makes the right-hand side of the one-sided bound equal to α_{cell}; a union bound over the fifteen cells then gives simultaneous coverage at least 0.95 over their recorded cutoffs. Independence among cells is not required for this union-bound step.

The calibration multiplier is the smallest common factor covering these upper bands:

$$c_{\text{cal}} = \max_{\substack{\text{primary calibration cells} \\ \text{recorded cutoffs } K}} \frac{\min\{1, \widehat{F}_{\text{cal}}(K) + \epsilon_N\}}{F_{\text{model}}(K)}.$$

The maximum occurs for $n = 128$, variance profile n, and $K = 5$, where

$$\begin{aligned} \widehat{F}_{\text{cal}}(5) &= 0.00456, \\ \epsilon_{50000} &= 0.00755233902486918, \\ F_{\text{model}}(5) &= 0.00450830449778397. \end{aligned}$$

Thus $c_{\text{cal}} = 2.68667279036341\ldots$, recorded above as $c = 2.68667279036$. This value is frozen before holdout evaluation. The raw model is treated as a center prediction, whereas the multiplied curve is a calibration-derived control used for model-conditioned parameter selection. It is not a distribution-free upper bound outside the calibrated dimensions and cutoffs. Higher-dimensional cells remain diagnostics and do not enlarge the calibrated range.

In the $K \le 5$ supplement, comparison is performed separately for every $(n, \text{profile}, \text{cohort}, K)$ cell even though the main table pools profiles and cohorts for readability. Agreement over this finite grid supports use of the model there; it does not turn Assumption 12 into a theorem or justify extrapolation beyond the displayed dimensions and cutoffs.

The empirical survival probability decreases sharply over the extended window. Across the six profile/cohort cells at each dimension, its range changes from $[6.00, 7.12] \cdot 10^{-4}$ at $K = 3$ to $[0, 4.00] \cdot 10^{-5}$ at $K = 5$ for $n = 16$; from $[9.40, 11.4] \cdot 10^{-4}$ to $[2.00, 7.00] \cdot 10^{-5}$ for $n = 32$; and from $[2.98, 3.24] \cdot 10^{-3}$ to $[1.00, 2.20] \cdot 10^{-4}$ for $n = 64$. The complete cellwise comparison reports the experimental and model probabilities, absolute and relative errors, ratios, and conservative-control indicators for every $K = 1, \ldots, 5$.

Table 3. Short-window model comparison. Each entry gives the range of $\widehat{F}(K)/F_{\mathrm{model}}(K)$ over the three profiles and two cohorts, followed by the number of cells below the conservative model. Zero in the first range at $K = 5$ is an observed zero-failure cell, not an estimate of zero tail probability.

K	$n = 16$	$n = 32$	$n = 64$
1	1.002–1.059; 6/6	0.982–1.027; 6/6	0.969–1.027; 6/6
2	0.978–1.100; 6/6	0.968–1.003; 6/6	0.965–1.057; 6/6
3	0.986–1.170; 6/6	0.850–1.030; 6/6	0.974–1.059; 6/6
4	0.735–1.379; 6/6	0.459–1.070; 6/6	0.806–1.196; 6/6
5	0–3.112; 5/6	0.530–1.853; 6/6	0.574–1.262; 6/6

The conservative curve covers 89 of the 90 cells. The unique exception is the $n = 16$, variance-$n^{3/2}$ calibration cell at $K = 5$, where $\widehat{F}(5) = 4 \cdot 10^{-5}$ and $F_{\mathrm{cons}}(5) = 3.45306056 \cdot 10^{-5}$.

E.3 Conditional batch calculation

The shared quotas convolve the supplied first-hit masses as in equation (15). Each reported K' is checked against $K' - 1$, with censored mass retained as failure, establishing minimality conditional on the model law. The empirical horizon is not treated as a cryptographic batch tail.

E.4 Performance measurements

The bounded Completion benchmark used one pinned logical processor of an Intel Core i9-13900, GCC 16.1.0 with `-O3 -DNDEBUG -march=native`, and single-threaded arithmetic. Each cell contains 50 calls after 10 warm-ups, with nearest-rank quantiles. The timed interval comprises one $N(g)$ computation, all K exact norm/mod and gcd probes, and one TowerSolve call; it excludes input preparation and the final identity check. Memory is reported separately for the reusable exact-norm cache, TowerSolve-owned retained storage, and process peak working set. The cache is included in the process measurement, but it is allocated once and is not charged anew to each probe or solver call.

The implementation computes $N(g)$ once, evaluates the fixed probes by exact NTT/CRT arithmetic, and calls TowerSolve only for the first clean shear. With 31-bit limbs, each lift uses $llen = \lceil lift\_bits/31 \rceil + 2$ limbs. The compact implementation keeps only the input and descent state in its static arena and reuses exact-width solution and lift banks whose capacities grow to the largest level reached by the current parameter cell. Its independent exact identity check streams the coefficients of $fG - gF$, retaining two GMP integers instead of an n-coefficient product polynomial. Consequently the reported TowerSolve high-water mark is the storage retained by these solver objects, whereas peak WS also includes the norm evaluator, input polynomials, FLINT/GMP, and the

Table 4. Bounded Completion at the conservative 2^{-128} model budgets. Times are milliseconds. The last columns give the shares of exact norm/mod probes (N/M), gcds, and TowerSolve (TS).

n	coefficient variance	K	median	p90	p99	N/M (%)	gcd (%)	TS (%)
512	n	40	68.01	73.62	89.39	83.44	1.57	12.89
512	$n^{3/2}$	40	85.99	90.71	97.78	82.17	1.64	13.37
512	n^2	40	104.51	110.22	133.81	80.77	1.72	15.12
1024	n	44	348.69	421.76	495.90	81.48	0.99	13.54
1024	$n^{3/2}$	44	431.05	536.07	668.15	80.77	1.09	15.14
1024	n^2	44	533.45	706.09	897.51	78.20	1.16	18.34

process runtime. Neither quantity is an asymptotic auxiliary-space bound for the Pornin–Prest solver [8].

Exact-norm initialization allocates one table of $521n$ 31-bit NTT roots. It is independent of the sampled polynomials and is reused by $N(g)$ and all K probe norms, rather than replicated per probe. The evaluator's two private residue arrays contain $2n$ words; the remaining private state consists of a few variable-size CRT integers.

Table 5. Memory accounting by dimension, in MiB. The norm cache is allocated once per process. $\mathrm{TS}_{\mathrm{base}}$ is the maximum-reservation baseline and $\mathrm{TS}_{\mathrm{compact}}$ is the compact retained high-water mark; the reduction and compact peak-WS columns give ranges over the three coefficient profiles. The columns are not additive because peak WS already contains the preceding components.

n	norm cache	$\mathrm{TS}_{\mathrm{base}}$	$\mathrm{TS}_{\mathrm{compact}}$	reduction (%)	peak WS
512	1.02	6.32–7.62	0.174–0.199	97.25–97.39	13.85–13.87
1024	2.04	26.21–32.67	0.417–0.531	98.38–98.45	13.85–13.88

The paired implementation comparison changes only the TowerSolve storage policy. Across all six cells, the ratio of compact to former TowerSolve median latency is 0.953–1.048 and the corresponding p90 ratio is 0.919–1.026. Thus the reduction in retained memory does not trade for a systematic timing regression at the resolution of this benchmark.

For the protocol accounting, let $t_{N(g)}$, $t_{\mathrm{probe}}(K)$, and t_{TS} be the corresponding phase medians. Charging K' probes and one solve for each of the μ equations gives

$$\widehat{t}_{\mathrm{batch}} = \mu t_{N(g)} + \frac{K'}{K} t_{\mathrm{probe}}(K) + \mu t_{\mathrm{TS}}. \tag{66}$$

Table 6. Estimated Completion contribution for one protocol batch. The estimate charges the full shared quota and rescales phase medians; it is not an end-to-end batch measurement.

n	coefficient variance	$\widehat{t}_{\mathrm{batch}}$ (s)	$N(g)$ (%)	probes (%)	TS (%)
512	n	0.316	9.1	29.8	61.1
512	$n^{3/2}$	0.405	8.5	29.0	62.5
512	n^2	0.530	8.0	26.5	65.6
1024	n	2.888	7.8	31.7	60.5
1024	$n^{3/2}$	3.815	7.3	29.4	63.3
1024	n^2	5.322	6.6	25.3	68.0

These batch values are estimates obtained from (66); they exclude discrete-Gaussian sampling, the inherited HAWK post-completion map, the form update, and all other end-to-end costs.