LEASH: When LaBRADOR's Guarantees Go Astray

Xiang Wang
Guizhou University

Ji Xiang
Chongqing University

***Abstract*—Succinct arguments enable efficient verification of large computations, while lattice-based constructions offer a promising path toward post-quantum security. LaBRADOR advances this direction through recursive reductions of knowledge that achieve compact proofs. However, these reductions require careful control of witness norms, relation preservation, and knowledge extraction to maintain their claimed guarantees.**

In this paper, we revisit the correctness and security analysis of LaBRADOR. We construct explicit counterexamples showing that its concrete parameter-selection strategy can fail to preserve completeness, even for valid binary R1CS witnesses. We further identify inconsistencies in recursive relation transformations and inaccuracies in the probability bounds underlying its knowledge extraction analysis. We develop targeted repairs and provide corresponding mathematical analyses, examining how changes to norm bounds, decomposition parameters, and challenge distributions affect the resulting security and efficiency guarantees. Our results highlight the importance of rigorous parameterization and compositional reasoning in lattice-based succinct proof systems.[1]

1. Introduction

Succinct arguments enable the verification of large computations using proofs much smaller than the underlying witnesses. Lattice-based constructions offer a promising path toward post-quantum succinct arguments, and substantial progress has been made from linear-size proofs to sublinear constructions. Nevertheless, achieving proof sizes competitive with hash-based systems remains challenging. To address this challenge, LaBRADOR introduced recursive reductions of knowledge for rank-one constraint systems (R1CS) under Module-SIS assumptions, reporting a proof size of only 58 KB for an arithmetic R1CS with 2^{20} constraints [1], [2]. Subsequent work has built on LaBRADOR to develop influential constructions for polynomial commitments, signature aggregation, practical proof frameworks, and cryptographic applications, including Greyhound, Aggregating Falcon Signatures with LaBRADOR, and LaZer, among others [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21].

However, recursive compression must preserve more than witness size. Each reduction must transform a valid witness into one satisfying both the next relation and its prescribed norm bound. This is nontrivial because the resulting witness includes folded vectors, decomposed commitment openings, and quadratic auxiliary terms. Knowledge extraction further relies on probability bounds and commitment-binding conditions that must remain valid across successive reductions. LaBRADOR's concrete parameter selection estimates the sizes of intermediate witness components using heuristic distributional models. While useful for optimizing proof size, these estimates do not by themselves establish uniform completeness guarantees for all valid witnesses.

In this paper, we revisit the correctness and security analysis of LaBRADOR. We construct an explicit binary R1CS instance with 448 variables and a valid witness for which the decomposed Gram terms alone necessarily exceed the entire next-layer norm budget prescribed by a fixed application of the parameter-selection rule. The violation is deterministic, independent of the folding challenges, and persists regardless of the admissible two-digit decomposition. We further identify inconsistencies in source and recursive relation transformations and inaccuracies in the discrete probability bounds underlying the knowledge-extraction analysis.

We develop corrected source and recursive relations, derive deterministic bounds for the complete next-layer witness, and establish replacement probability and knowledge-extraction analyses under explicit assumptions. We quantify the resulting changes to parameter selection and proof efficiency, and provide exact arithmetic certificates and reproducible code.

2. Background and Preliminaries

2.1. Algebraic Setting and Relations

LaBRADOR operates over the polynomial ring $R = \mathbb{Z}[X]/(X^d + 1)$ and its quotient $R_q = R/qR$, where d is a power of two and q is an odd prime. Norms are Euclidean norms of centered integer coefficients, extended to polynomial vectors. We write $\mathrm{ct}(f)$ for the constant coefficient of f and $\overline{f} = f(X^{-1})$ for conjugation. The latter

1. This full version includes the counterexample calculations, complete proofs of the repaired bounds and extraction results, and detailed parameter and communication accounting. Code and experimental artifacts are available at https://github.com/cxcx52/LEASH. The repository provides a Rust binary-R1CS frontend and recursive lattice backend, library and command-line interfaces, proof serialization, a runnable completeness example, benchmark configurations and results, and a register of downstream papers and implementation dependencies.

allows coefficient inner products to be represented through constant terms of polynomial products.

The central object in LaBRADOR is its principal relation. A witness consists of r polynomial vectors $s = (s_1, \ldots, s_r)$, with $s_i \in R_q^n$. It must satisfy quadratic dot-product equations and a global norm bound:

$$\begin{aligned} f(s) &= 0 \qquad (f \in \mathcal{F}), \\ \mathrm{ct}(f'(s)) &= 0 \qquad (f' \in \mathcal{F}'), \\ \sum_{i=1}^{r} \|s_i\|_2^2 &\le \beta^2. \end{aligned}$$

The functions in $\mathcal{F}$ must vanish as full polynomials, whereas those in $\mathcal{F}'$ are constrained only at their constant coefficients.

LaBRADOR supports two reductions from rank-one constraint systems (R1CS). The binary route packs witnesses over $\mathbb{F}_2$, together with auxiliary vectors and their conjugates, into polynomial blocks. The arithmetic route handles R1CS over $\mathbb{Z}/(2^d+1)\mathbb{Z}$ using short polynomial encodings and evaluation at $X = 2$. This evaluation is defined over the integer ring R, while the proof operates over R_q. Recovering the intended arithmetic from residues therefore requires an appropriate integer lift.

2.2. LaBRADOR's Recursive Reduction

LaBRADOR achieves compact proofs through a sequence of recursive reductions. Each reduction transforms a principal-relation instance into another instance with a smaller witness representation. The resulting relation is then passed to the next round, rather than checking the entire intermediate witness directly.

Given witness blocks $s_1, \ldots, s_r$, the prover computes inner commitments t_i and Gram polynomials $g_{ij} = \langle s_i, s_j \rangle$. Their decomposition digits are bound by a first outer commitment. The verifier then samples a random projection to check witness shortness, and the protocol aggregates the original constraints. A second outer commitment binds the remaining linear auxiliary terms h_{ij}.

The verifier subsequently samples challenge polynomials $c_1, \ldots, c_r$ from a prescribed set $\mathcal{C}$. The prover forms the folded vector $z = \sum_i c_i s_i$, which combines the original witness blocks into one. The auxiliary terms allow the verifier to check the corresponding commitment and quadratic relations without revealing every original block.

Folding reduces the witness dimension but increases coefficient magnitudes. LaBRADOR therefore decomposes $z = z^{(0)} + bz^{(1)}$ using a chosen base b. The next witness also contains decomposition digits of the commitments and auxiliary polynomials. Writing these collections as $\widetilde{t}, \widetilde{g}, \widetilde{h}$, its norm must satisfy

$$\|z^{(0)}\|_2^2 + \|z^{(1)}\|_2^2 + \|\widetilde{t}\|_2^2 + \|\widetilde{g}\|_2^2 + \|\widetilde{h}\|_2^2 \le (\beta')^2.$$

The witness components are then repartitioned and padded to form the next principal-relation instance. Hence, the norm budget must cover the complete transformed witness, not just the folded vector.

2.3. Guarantees and Parameter Selection

We briefly review LaBRADOR's underlying relations, recursive reduction, and parameter-selection strategy. Our primary reference is the 34-page version of ePrint 2022/1341 [2].

The security of recursive proofs depends on the guarantees established at each reduction. Completeness requires valid witnesses to survive the reduction with the prescribed probability. If the next-layer norm bound is too small, even an honest prover may be unable to complete the proof.

Knowledge soundness concerns what an accepted proof establishes. The verifier checks the final relation rather than every intermediate witness. Each reduction must therefore preserve the constraints needed to recover a witness for the original statement. Otherwise, a valid final witness may not establish knowledge of a valid source witness.

The extraction analysis also relies on probability bounds. Random projections are used to detect witnesses that exceed the allowed norm, while challenge polynomials support the recovery of committed witness blocks. If these probabilities are underestimated, the claimed extraction error may no longer follow. Moreover, extraction may produce a witness longer than an honest witness. This larger norm affects the Module-SIS assumptions needed for commitment binding.

These requirements are closely connected. Changing a norm bound may restore honest completeness, but it can also change the commitment parameters and proof size. Changing a challenge distribution affects both folding and extraction. LaBRADOR uses heuristic estimates of intermediate coefficient sizes to select practical parameters. It also considers restarts and parameter adjustments when those estimates are exceeded.

Consequently, a recursive proof cannot be evaluated through its compression ratio alone. Its guarantees depend on whether valid witnesses remain within the prescribed bounds, the transformed relations preserve the original statement, and the extraction analysis holds for the actual parameters. These conditions jointly determine the completeness, knowledge soundness, and efficiency of the resulting proof system.

3. Where the Guarantees Fall Short

We examine LaBRADOR by following a witness through its reductions. We first apply the parameter-selection rule to an honest binary witness and check whether it fits the next-layer norm budget. We then trace the arithmetic and constraints passed between relations. Finally, we examine the probability bounds used to recover a source witness from an accepting proof.

3.1. A Valid Witness Exceeds the Recursive Norm Budget

LaBRADOR's parameter-selection rule estimates the sizes of intermediate coefficients to choose decomposition

bases and norm bounds. Its Gram estimates model sums of products of witness coefficients using Gaussian distributions. A valid witness, however, need not have the distribution assumed by this model. Applying the rule to a simple all-one binary witness reveals a failure: the Gram decomposition alone requires more squared norm than the budget allocated to the entire next witness [2].

Consider a binary R1CS with 448 variables and 448 constraints, where $A = B = C = I_{448}$. The all-one assignment is valid, and the auxiliary vectors $a = Aw$, $b = Bw$, and $c = Cw$ are also all one. At ring degree $d = 64$, each group of 64 coefficients is represented by $J = \sum_{j=0}^{63} X^j$. Including the conjugate copies produces eight blocks of length seven: four containing J and four containing $\overline{J}$. Thus, $r = 8$, $n = 7$, and the squared witness norm is 3,584.

For the ordinary-round calculation, we fix $q = 4294967197$ and inner commitment rank $\kappa = 24$. We use the public binary norm bound $\beta^2 = 2(448+3\cdot 448) = 3584$, rather than the looser bound displayed in the source protocol. With the original recipe's challenge-energy convention $\tau = 71$, Section 5.4 of LaBRADOR [2] gives $b = 9$, $(t_1, b_1) = (10, 10)$, and $(t_2, b_2) = (2, 10)$. The source-size and modular-projection conditions hold for these choices. The accompanying certificate checks these conditions and the parameter rounding using integer arithmetic.

Proposition 3.1 (Fixed-profile completeness failure)**.** *For this witness and parameter configuration, every two-digit modular decomposition of the upper-triangular Gram entries exceeds the prescribed total next-layer norm budget.*

Proof. Negacyclic multiplication gives $[X^k]J^2 = 2k+2-64$ for $0 \le k < 64$. The polynomials J^2, $J\overline{J}$, and $\overline{J}^2$ each have squared coefficient norm 87,424. Every Gram entry is seven times one of these polynomials. There are 36 upper-triangular entries, so their total squared norm is $E_G = 36 \cdot 7^2 \cdot 87424 = 154215936$. Each coefficient has magnitude at most 448, so reduction modulo q does not change these values.

Now consider any two digits $u_0, u_1 \in R_q$ satisfying $g = u_0 + 10u_1$ modulo q, and let U_0, U_1 be their centered integer lifts. Since centering minimizes coefficient magnitudes, Cauchy–Schwarz gives

$$\|g\|_2^2 \le \|U_0 + 10U_1\|_2^2 \le 101\big(\|U_0\|_2^2 + \|U_1\|_2^2\big).$$

Summing over the Gram entries yields a lower bound on their decomposition digits. Substituting the selected parameters into LaBRADOR's target budget gives

$$\|\widetilde{g}\|_2^2 \ge \frac{154215936}{101} > \frac{102115328}{81} = (\beta')^2.$$

All other components of the next witness contribute nonnegative squared norm. The complete witness therefore cannot satisfy the target bound. This argument allows arbitrary high digits and does not depend on a particular decomposition algorithm. □

The Gram-digit lower bound exceeds the entire budget by approximately 21.1%. These Gram entries are fixed before the verifier samples any challenges. Resampling challenges therefore cannot resolve the failure while the witness packing and parameters remain unchanged. LaBRADOR explicitly allows restarts or parameter adjustments when its estimates are exceeded. Our example distinguishes these options: a restart with the same fixed profile cannot help, whereas changing the parameters requires a new calculation of the norm bounds, commitment requirements, and proof size. The result concerns this fixed profile, not every possible parameter choice or an execution of the original complete implementation [2].

3.2. Relation Preservation Across Reductions

The preceding example satisfies the source constraints but fails the next-layer norm check. We next examine the constraints themselves. In particular, we follow the arithmetic source reduction to determine whether its ring equations retain the meaning of the original R1CS statement.

This reduction encodes arithmetic modulo $N = 2^d + 1$ as polynomials and performs proof computations modulo q. To recover the source meaning, a polynomial residue must be lifted back to an integer polynomial before evaluation at $X = 2$. A coefficient that crosses the centered interval boundary changes by a multiple of q. That change need not vanish modulo N.

Let L_q denote centered coefficient lifting and φ evaluation at $X = 2$ modulo N. Theorem 6.3 of LaBRADOR [2] uses the condition $S = \sqrt{Md/2}\,\beta + \beta^2/2 < q$, where $M = n + (3+\ell)k$, to justify the lifting step. We construct a valid witness satisfying this condition for which the step changes a zero source residual into a nonzero one. Take $d = 4$, $N = 17$, $q = 13$, and $n = k = \ell = 1$. With $A = 15$, $B = 11$, $C = 14$, and $w = 4$, the source values are $(a, b, c) = (9, 10, 5)$, satisfying $ab = c$ modulo 17. An allowed choice of aggregation challenges produces the integer residual $F = 2 - X - 7X^2 - 5X^3$.

The encoded witness has squared norm $8 < \beta^2 = 801/100$, and the stated condition $S < q$ holds. Nevertheless, centered reduction changes the coefficient -7 to 6. The residual's evaluation changes from -68 to -16, giving

$$\varphi(F) = 0, \qquad \varphi(L_{13}(F \bmod 13)) = 1 \quad \text{in } \mathbb{Z}_{17}.$$

The encodings and aggregation challenges are given in Appendix A.2. This example isolates the lifting premise of Theorem 6.3. It does not satisfy the separate restriction $Md < 0.3q$ used for composition with the main protocol, and is therefore not a counterexample to that complete composed configuration.

Following the same residuals into the target relation reveals a separate mismatch. The prover sends $g_j = \widetilde{f}_j(s)$, and the source test allows g_j to be a nonzero polynomial whose lifted evaluation vanishes. Figure 5 of LaBRADOR [2] nevertheless lists $\widetilde{f}_j$ among the functions required to vanish in the principal relation. The target must instead contain

$\widetilde{f}_j - g_j$. It must also include both source commitment-opening equations. These equations tie the final witness to the values committed before the corresponding challenges. An additional example in Appendix A.2 separates the missing residual constant from the lifting failure: lifting is unambiguous, but the printed target still rejects the honest residual.

The recursive norm check has a related consistency issue. The packed Gram and linear auxiliary vectors store only upper-triangular entries, as do their commitments. Figure 3, line 14, of LaBRADOR [2] instead sums their squared norms over all ordered pairs. Off-diagonal entries are consequently counted twice in that check but once in the packed witness. The parameter calculation and verifier must use the same convention. The counterexample in Section 3.1 already uses the more favorable single-count convention, so its failure does not depend on this discrepancy.

3.3. Probability Bounds in Knowledge Extraction

After checking the transformed witnesses and relations, we turn to the extraction analysis. Its random checks must rule out invalid witnesses with the stated probabilities. LaBRADOR presents approximate numerical projection bounds, then uses them as concrete error budgets in the extraction proof. We evaluate these probabilities directly for the discrete distribution used by the protocol [2].

Each projection entry is zero with probability $1/2$ and takes either sign with probability $1/4$. Consider a 128-dimensional vector w whose first coordinate is $a = 2^{21}$, followed by 100 ones and 27 zeros. For the same modulus $q = 4294967197$, we have $\|w\|_2 > a$ and $125a < q$, meeting the relevant premises of the modular projection lemma. Each projected coordinate has magnitude at most $a + 100 < q/2$, so modular reduction has no effect.

Write a projected coordinate as $aA_j + B_j$, where A_j is the ternary coefficient multiplying the large coordinate and B_j sums the 100 remaining ternary terms. Let $K = \sum_j A_j^2$ and $H = \sum_j A_j B_j$. The projected squared norm is $a^2K + 2aH + \sum_j B_j^2$. Since $|B_j| \le 100$, either $K \le 29$, or $K = 30$ together with $H \le -1$, forces this norm below $30a^2$.

Here K follows $\mathrm{Bin}(256, 1/2)$. Conditional on $K = 30$, H follows $\mathrm{Bin}(6000, 1/2) - 3000$. Counting the two disjoint events therefore gives

$$\begin{aligned}
&\Pr\left[\|\Pi w \bmod q\|_2^2 < 30a^2\right] \\
&\quad \ge 2^{-256} \sum_{k=0}^{29} \binom{256}{k} \\
&\qquad + 2^{-257} \binom{256}{30} \left(1 - 2^{-6000} \binom{6000}{3000}\right) \\
&\quad > 2^{-127} > 2^{-128}.
\end{aligned}$$

The displayed lower bound is approximately $2^{-126.89}$. Its comparison with 2^{-127} is checked using exact integers, not floating-point approximations. The mechanism is simple: the large coordinate places discrete probability mass near the threshold, and the smaller coordinates move part of that mass below it. The Gaussian estimate does not account for this effect.

A second exact calculation shows a discrepancy in the single-row upper tail. For a vector of 4,096 ones, the probability of exceeding $9.5\|w\|_2$ is approximately $2^{-134.98}$, about 64.7 times the stated 2^{-141} estimate. Appendix A.3 gives the corresponding binomial sum.

The challenge-set calculation also requires the distribution after sampling is complete. LaBRADOR starts with polynomials having 23 zero coefficients, 31 coefficients of magnitude one, and 10 of magnitude two. This candidate set has size $N_{\mathrm{raw}} = 64!\,2^{41}/(23!\,31!\,10!)$, with $\log_2 N_{\mathrm{raw}} \approx 128.088$. The sampler then rejects candidates whose operator norm exceeds its threshold. Root-counting terms such as $2/|\mathcal{C}|$ must use the retained set, not the raw candidate count. In a sample of one million candidates, 179,602 passed the filter, yielding an estimated retained cardinality of roughly $2^{125.611}$. This sample indicates the size of the correction, but does not replace the certified cardinality bound needed by the proof [2].

These findings invalidate specific numerical premises, rather than demonstrate a complete forgery. In particular, the projection event exceeding 2^{-128} does not by itself establish that the protocol's overall 2^{-125} knowledge-error bound is violated. It shows that the local probability cannot be charged at the stated value. The projection and challenge bounds must be corrected before they are combined with the remaining extraction and recursion terms [2].

4. Reconstructing the Recursive Guarantees

A valid recursive reduction must preserve the underlying relation, control the norm of the next witness, and provide reliable probability bounds for its random checks. These requirements are closely connected. Increasing a norm bound may restore completeness, but it also changes the conditions required for commitment binding and knowledge extraction. We therefore reconstruct the relations and derive compatible deterministic and probabilistic bounds.

4.1. Corrected Source and Target Relations

We distinguish the honest relation $\mathcal{R}_\beta$ from the extraction relation $\mathcal{R}_{\beta_{\mathrm{ext}}}$. Both impose the same algebraic constraints, but the latter permits a larger witness norm. Conditions used to decode an extracted witness must therefore hold at β_{ext}, rather than only at the honest bound β.

Arithmetic R1CS. We retain LaBRADOR's short polynomial encodings and two source commitments. The corrected target includes both commitment-opening equations and the residual equations $\widetilde{f}_j(s) - g_j = 0$, where g_j is the value previously sent by the prover. It also fixes the constant wire and public inputs. These constraints preserve the connection between the final witness and the values committed before the corresponding challenges.

The source arithmetic takes place modulo $N = 2^d + 1$, whereas the proof operates over R_q. Let L_q denote centered

coefficient lifting and φ denote evaluation at $X = 2$ modulo N. The verifier checks $\varphi(L_q(g_j)) = 0$. To justify this test, the linear coefficients are first combined in the source ring and then encoded.

The resulting lifting condition follows from Cauchy–Schwarz. With $M = n_{\mathrm{src}} + (3 + \ell)k$ encoded witness elements, the linear contribution to each residual coefficient is bounded by $\sqrt{Md/2}\,\beta_{\mathrm{ext}}$, while the quadratic contribution is at most $\beta_{\mathrm{ext}}^2/2$. Thus,

$$\sqrt{\frac{Md}{2}}\,\beta_{\mathrm{ext}} + \frac{\beta_{\mathrm{ext}}^2}{2} < \frac{q}{2}$$

ensures that centered lifting recovers the intended integer residual. Evaluation at $X = 2$ then yields the source equation, without assuming that evaluation is a homomorphism on R_q.

Binary R1CS. The binary route instead uses coefficient-level constraints. We retain the source commitment, conjugacy equations, bitness and product checks, scalar aggregation equations, and explicit padding constraints. The bitness identities use the conjugate of $J = 1 + X + \cdots + X^{d-1}$. Public inputs and the constant wire are pinned at their individual coefficient positions, rather than as entire ring elements.

The conditions $\beta_{\mathrm{ext}}^2 < q$ and $6k < q$ ensure that the bitness and product equations recover their intended integer meaning. Scalar replies are represented as even integers in $[-n_{\mathrm{src}}, 3k]$. When $n_{\mathrm{src}} + 3k < q$, the target equations uniquely determine these integers from their residues, allowing the verifier to check their parity correctly.

Recursive relations. The ordinary target retains the inner opening equation, both outer commitment-opening equations, and all quadratic consistency checks. Symmetric auxiliary entries are stored once using upper-triangular indexing, while off-diagonal terms retain their required algebraic multiplicities. Unused padding coordinates are constrained to zero.

These corrections preserve the intended algebraic relations in honest executions. Knowledge extraction additionally requires commitment binding and valid aggregation arguments.

4.2. Deterministic Bounds for the Next-Layer Witness

We replace the heuristic coefficient estimates with deterministic bounds on the complete target witness. The key difficulty is digit decomposition. Although the low digits are bounded by construction, the final quotient can be larger and must be included in the norm calculation.

For centered base-a decomposition into $t \geq 2$ digits, the first $t-1$ digits have magnitude at most $a/2$. The final quotient is bounded by $|x|/a^{t-1} + C(a,t)$, where $C(a,t) = \frac{a}{2(a-1)}(1 - a^{-(t-1)})$ accounts for accumulated carries.

This yields a general decomposition bound. A coefficient vector of length D and norm at most L has decomposed squared norm bounded by

$$F_{a,t}(D, L) = \frac{D(t-1)a^2}{4} + \left(\frac{L}{a^{t-1}} + \sqrt{D}\,C(a,t)\right)^2.$$

We apply this bound to the commitment and auxiliary terms. Write $G = r(r+1)/2$, $D_z = nd$, $D_u = d(r\kappa + G)$, and $D_g = dG$, where κ is the inner commitment rank. Let b be the decomposition base for the folded vector, and let (a_1, t_1) and (a_2, t_2) describe the commitment/linear and Gram decompositions.

The folded vector contributes at most $E_z = D_z b^2/4 + (\Gamma/b + \sqrt{D_z}/2)^2$ when its norm is bounded by Γ. The commitment and linear digits contribute at most $E_u = F_{a_1,t_1}(D_u, q\sqrt{D_u}/2)$. For the Gram terms, Young's inequality gives $\sum_{i \leq j} \|g_{ij}\|_2^2 \leq d\beta^4$, leading to $E_g = F_{a_2,t_2}(D_g, \sqrt{d}\,\beta^2)$.

Proposition 4.1 (Complete Next-Layer Norm Bound)**.** *The repaired ordinary reduction admits a deterministic bound for its complete target witness. Every valid input witness $s = (s_1, \ldots, s_r)$ with $\sum_i \|s_i\|_2^2 \leq \beta^2$ and folded response satisfying $\|z\|_2 \leq \Gamma$ produces a target witness w' such that*

$$\|w'\|_2^2 \leq E_z + E_u + E_g.$$

Hence, $\beta' = \lceil \sqrt{E_z + E_u + E_g} \rceil$ is a sufficient target norm bound, including all decomposition digits and final quotients.

Proof. Decomposing $z = z^{(0)} + bz^{(1)}$ bounds its low part by $b\sqrt{D_z}/2$ and its high part by $\Gamma/b + \sqrt{D_z}/2$, giving E_z. The commitment and linear auxiliary coefficients are centered residues of magnitude at most $q/2$, so the general decomposition bound gives E_u.

For the Gram terms, $\|g_{ij}\|_2 \leq \sqrt{d}\,\|s_i\|_2\|s_j\|_2$. Summing over the upper triangle gives the stated Gram bound. Centering modulo q cannot increase the norm, and applying the decomposition bound yields E_g. The three groups occupy disjoint target coordinates, so their squared norms add. □

The proposition does not assume independent or Gaussian witness coefficients. It therefore covers the correlated Gram terms responsible for the counterexample in Section 3.1.

The target layout must also preserve the original dot-product structure. Its unpadded size is $W' = 2n + rt_1\kappa + (t_1 + t_2)G$. Writing $M_{\mathrm{aux}} = W' - 2n$, a split into $r' = 2\nu + \mu$ blocks uses common rank $n' = \max\{\lceil n/\nu \rceil, \lceil M_{\mathrm{aux}}/\mu \rceil\}$. This keeps the two folded components aligned. Zero-padding changes the dimension but not the norm.

Unlike the original heuristic estimates, the resulting budget covers every component of the next witness, including the Gram terms, without assumptions about their coefficient distributions.

4.3. Uniform Modular Projection Bounds

The deterministic decomposition bound does not replace LaBRADOR's probabilistic shortness check. To support extraction, the random projection must reject oversized witnesses except with sufficiently small probability. The counterexample in Section 3.3 shows that the original numerical estimates do not provide the required uniform bound.

We establish a replacement for the actual discrete distribution. Let χ take values $0, 1, -1$ with probabilities $1/2, 1/4, 1/4$, respectively.

Theorem 4.2 (Uniform Modular Projection)**.** *The 1,024-row ternary projection satisfies a uniform lower-tail bound for arbitrary centered witnesses. Every $w \in \mathbb{Z}^D$ with $\|w\|_2 \geq u > 0$ and $250u \leq q$ satisfies*

$$\Pr_{\Pi \leftarrow \chi^{1024 \times D}} \left[\|\Pi w \bmod q\|_2^2 \leq 120u^2 \right] < 2^{-164}.$$

The bound includes the closed acceptance event used by the verifier.

Proof sketch. We distinguish three cases. For small-norm vectors, exponential moments control the lower tail and modular wraparound. If one coordinate is large, conditioning on the remaining coordinates gives a modular small-ball bound. Otherwise, we isolate a subvector with sufficiently spread-out norm and apply Berry–Esseen, retaining its approximation error [22]. The complete inequalities and numerical certificates appear in Appendix B.3.

Honest completeness requires an upper-tail bound as well. For a witness of norm at most β, the squared projection norm has expectation at most $512\beta^2$ and variance at most $512\beta^4$. Cantelli's inequality bounds rejection at threshold $640\beta^2$ by $1/33$.

The same threshold gives extraction slack $\beta_{\text{ext}} = \sqrt{16/3}\,\beta$, since $120\beta_{\text{ext}}^2 = 640\beta^2$. Under $250\beta_{\text{ext}} \leq q$, Theorem 4.2 bounds the probability that an oversized fixed witness passes the projection check by 2^{-164}. Thus, the honest rejection probability and extracted norm are calibrated together.

4.4. Certified Challenges and Folding Bounds

The final requirement is to control the folded witness while retaining enough challenges for extraction. A larger challenge set reduces challenge-related error terms, but larger challenge norms increase the requirements on commitment binding. Both properties must therefore be considered when selecting the distribution.

We construct a bounded-energy family $\mathcal{C}_{d,E,T}$ consisting of coefficient vectors over $\{-2, -1, 0, 1, 2\}^d$ with squared norm at most E and multiplication operator norm strictly below T. Sampling uniformly from the bounded-energy set and applying the operator-norm filter gives the intended challenge distribution.

Proposition 4.3 (Certified Challenge Family)**.** *The family $\mathcal{C}_{64,104,30}$ retains more than two thirds of its bounded-energy candidates and contains more than 2^{143} challenges. Every retained element has squared coefficient norm at most 104 and operator norm below 30. For prime $q > 32$ with $q \equiv 5 \pmod 8$, differences of distinct challenges are invertible in R_q [23].*

The cardinality follows from an exact recurrence counting coefficient vectors by energy. An exponential-moment argument bounds the fraction removed by the operator-norm filter. All decisive comparisons use integer or rational arithmetic. The same counts support uniform sampling without enumerating the challenge set.

Filtering changes the coefficient distribution. In particular, the energy restriction creates dependencies even before the operator-norm test. We therefore derive the folding bound using conditional moments rather than assuming independent coefficients.

Proposition 4.4 (Folding Concentration)**.** *Independent challenges sampled uniformly from $\mathcal{C}_{64,104,30}$ satisfy the following bound for any witness blocks with $\sum_i \|s_i\|_2^2 \leq \beta^2$:*

$$\Pr\left[\left\| \sum_{i=1}^{r} c_i s_i \right\|_2^2 > 9d\beta^2 \right] < \frac{1}{16}.$$

The bound holds for integer polynomial products and remains valid after centered reduction modulo q.

Proof sketch. Symmetries of the filtered distribution give zero mean and scalar covariance, despite coefficient dependencies. Exact fourth moments control the conditional variance of the folded norm. The paired eigenvalue structure of negacyclic multiplication bounds the cross terms between independent challenges. Cantelli's inequality then gives the stated result. The moment calculations and certificates are given in Appendices B.5 and D.1.

At $d = 64$, Proposition 4.4 permits the fold cap $\Gamma = 24\beta$. Substituting this value into Proposition 4.1 gives a complete next-layer norm budget with folding rejection probability below $1/16$.

The revised distribution uses operator threshold 30, rather than LaBRADOR's original threshold 15. The larger operator-norm allowance must therefore be reflected in the commitment-binding parameters.

Corollary 4.5 (Honest Completeness)**.** *Under the corrected relations and certified parameter conditions, each honest ordinary reduction succeeds with probability greater than $479/528$. A compatible seven-layer configuration consequently achieves honest acceptance probability greater than $1/2$, without local challenge retries.*

Proof. The projection and folding checks reject an honest witness with probabilities at most $1/33$ and less than $1/16$, respectively. When both pass, Proposition 4.1 ensures that the complete target witness meets its norm bound, while the corrected relations hold by construction. A union bound gives per-layer acceptance greater than $1 - 1/33 - 1/16 = 479/528$. Since this guarantee holds uniformly for valid incoming witnesses, it applies at each successive layer. The

certified seven-layer parameter path satisfies $(479/528)^7 > 1/2$. □

5. Knowledge Extraction and Recursive Composition

The repaired norm bounds explain why an honest prover can complete the reductions. They do not yet show how to recover a source witness from a prover that convinces the verifier. Rewinding provides the responses needed for recovery, but their consistency and the cost of obtaining them must also be established.

5.1. Extracted Witnesses and Commitment Binding

We consider an efficient initial-state generator that receives independently uniform commitment matrices and produces a statement x and a classical prover. The prover has a polynomial per-execution time bound, and the extractor can restore its state and original random tape. Statements may depend on the public parameters, but additional setup-dependent information must be permitted by the stated Module-SIS model. The extraction guarantee covers all efficiently generated statements, including false ones.

For a single layer, we first analyze the bare reduction, whose final target witness is available and explicitly checked. The extractor must recover a witness in $\mathcal{R}_{\beta_{\text{ext}}}(x)$. This includes every source equation and the extracted norm bound, not just the commitment-opening equations.

Two accepting responses under different challenges can reveal a witness block. However, changing a challenge may also change the opening supplied by the prover. Computational binding does not make every opening in the entire rewinding tree unique. It bounds the success of algorithms that find inconsistent openings. We therefore account for these inconsistencies through explicit Module-SIS collision experiments.

Write B for the actual target-witness norm cap, b for the folded-vector decomposition base, and T for the challenge operator-norm bound. The relevant collision radii are

$$L_{\text{out}} = 2B, \qquad L_{\text{weak}} = 8T(b+1)B,$$
$$L_{\text{Gram}} = 2(b+1)B + 4T\beta_{\text{ext}}.$$

The first radius applies to either outer commitment. The other two arise from inner linear-opening and Gram-consistency checks. Each radius belongs to a particular matrix, with its own rank and dimension. The corresponding hardness assumption must account for the full collision-finding algorithm, including initialization and any nested extraction calls. A bound established only at the honest witness norm is insufficient.

5.2. An Acceptance-Gated Extractor

Even when two responses determine a witness block algebraically, obtaining them efficiently is not automatic. Acceptance can be rare after a particular prefix. Repeating until another response accepts may therefore require substantial work, and an extractor parameterized by a known acceptance gap does not resolve the case where that gap is unknown.

Algorithm 1: Acceptance-gated extraction

1) Run the bare reduction with fresh verifier randomness. If an early check or the final target-witness check fails, return $\perp$.
2) Save the original accepted transcript, its first outer opening, its final challenge vector c^0, and its reconstructed folded response z^0.
3) For each coordinate i, restore the original prefix before the final challenges. Keep the other verifier challenges unchanged, independently resample c_i, and run the suffix.
4) If the suffix accepts and $c_i \neq c_i^0$, compare its first outer opening with the saved opening. Different openings give a verified outer Module-SIS collision; record it and return $\perp$. Otherwise, recover $S_i = (c_i - c_i^0)^{-1}(z - z^0)$ and advance to the next coordinate.
5) After every unsuccessful trial, including acceptance with $c_i = c_i^0$, independently abort with probability $1/N$. Otherwise, retry the same coordinate.
6) Return $S = (S_1, \ldots, S_r)$ only if it satisfies every equation and the norm bound of $\mathcal{R}_{\beta_{\text{ext}}}(x)$. Otherwise, return $\perp$.

Our extractor begins with one ordinary execution. It proceeds only if this execution accepts, and never replaces that original transcript with a later successful fork. After an unsuccessful recovery attempt, an independent abort coin limits further work.

We use an exactly uniform challenge set $\mathcal{C}$ of size $N > 2$, with invertible differences between distinct challenges. The ring has the two-field decomposition $R_q \simeq \mathbb{F}_{q^{d/2}} \times \mathbb{F}_{q^{d/2}}$. All verifier challenges follow the prescribed message order.

For the challenge family constructed above, the abort coin can be implemented by sampling an independent challenge and testing whether it is zero. This has probability exactly $1/N$, without requiring the sampler to know the filtered cardinality.

Let η_0, η_h denote the success probabilities of the two outer-opening collision experiments. Let η_y correspond to the inner linear-opening experiment on a uniformly chosen coordinate, and η_g to the experiment using a selected inconsistent Gram entry. Finally, η_E accounts for collisions found during candidate generation. These are unconditional probabilities for complete experiments, including failed candidate generation. Additional test executions use fresh randomness sampled after the candidate has been fixed.

Theorem 5.1 (Single-Layer Extraction)**.** *Algorithm 1 extracts a verified witness for the repaired ordinary reduction without an acceptance-probability estimate. Under the stated relation, setup, chronology, projection, and challenge*

conditions, its success probability satisfies

$$\Pr[S \in \mathcal{R}_{\beta_{\mathrm{ext}}}(x)] \geq \bar{\varepsilon} - \delta - \nu,$$
$$\delta = \delta_{\mathrm{JL}} + q^{-m} + q^{-d/2} + \frac{r+2}{N},$$
$$\nu = \sqrt{\eta_0} + (2\eta_h + 2r\eta_y + \eta_g)^{1/3} + \eta_E.$$

Here $\bar{\varepsilon}$ is ordinary acceptance probability over the setup and initial-state experiment, m is the number of independent scalar aggregations, and δ_{JL} is the modular projection error. The expected number of ordinary or suffix executions is at most $1+r/(1-1/N)$, uniformly over permitted initial states. Thus, polynomial local work and expected-polynomial challenge generation give an expected-polynomial-time extractor.

Proof sketch. Fix the original prover tape and all verifier coins except one challenge coordinate. Let p be the fraction of challenges producing acceptance, and write $\alpha = 1/N$. A row reached by the initial gate has $p \geq \alpha$. A fresh trial obtains a distinct accepting challenge with probability $p - \alpha$. Including the abort coin, its stopping probability is $\rho = p(1-\alpha) + \alpha^2$.

The gate reaches this row with probability p. Its contribution to expected recovery work is therefore $p/\rho \leq 1/(1-\alpha)$, even when recovery conditioned on reaching the row is expensive. The gate-weighted abort probability is at most α. Summing over coordinates gives the stated work bound and an additional extraction loss of at most r/N.

To bound invalid candidates, the analysis selects a most likely accepting opening at each relevant prefix. This is an analytical device, not an opening the extractor must compute. The probability mass of inconsistent openings is bounded through the collision experiments. The remaining terms come from projection, scalar aggregation, ring aggregation, and degree-two challenge checks. Fresh tests are essential: the candidate being tested is fixed before their randomness is sampled.

Retaining the invalid-candidate probability inside the Cauchy–Schwarz argument gives the cubic-root term in ν. Discarding that weight earlier gives a weaker fourth-root bound. The refinement changes the reduction loss, not the extractor or its collision radii.

Under the corresponding Module-SIS assumptions, the computational terms are negligible asymptotically. Expected-time reductions are handled by truncating any hypothetically nonnegligible solver at a suitable polynomial cutoff. Concrete security bounds must instead charge the actual reduction time and any chosen truncation loss; a small statistical error alone does not establish a numerical knowledge-security level.

5.3. Recursive Composition and Complete Restarts

The single-layer extractor expects a checked target witness. In a recursive proof, that witness is supplied by another extractor. This substitution changes both the available norm bound and the randomness used to generate the final reply.

First, the inner extractor must return a witness satisfying every equation and the norm cap required by the preceding layer. If layer $i+1$ extracts norm at most $\beta_{\mathrm{ext},i+1}$, layer i must use $B_i \geq \beta_{\mathrm{ext},i+1}$ in its target check and collision bounds. The smaller honest bound cannot replace it. At the terminal layer, B_i is the cap of the explicitly checked final witness.

Second, each inner invocation restores the original continuation state but uses fresh extraction coins. It cannot reuse a previously extracted witness or a mutable extraction cache. Its expected cost must be bounded uniformly over all permitted continuation states, including those reached by an outer rewind. Early verifier rejection always forces failure.

The gate analysis extends to these randomized final replies. In particular, it handles prefixes whose acceptance probability is below $1/N$, a case absent when the bare prover's tape is fixed. The same expected-work and abort bounds remain valid.

Theorem 5.2 (Recursive Composition)**.** *A sequence of D repaired ordinary reductions ending in an explicitly checked target witness admits a composite extractor under the relation-matching, fresh-randomness, and setup conditions above. Its success probability is*

$$\Pr[S \in \mathcal{R}_{\mathrm{out}}(x)] \geq \varepsilon_{\mathrm{total}} - \sum_{i=1}^{D} (\delta_i + \nu_i),$$

where δ_i and ν_i are the single-layer losses for the complete recursively initialized experiments. If τ_i bounds local work and C_{i+1} bounds the inner extractor's expected cost uniformly over continuation states, then

$$C_i \leq c_i(\tau_i + C_{i+1}), \qquad c_i = 1 + \frac{r_i}{1 - 1/N_i}.$$

The extractor runs in expected polynomial time when the depth, local costs, dimensions, and product $\prod_i c_i$ are polynomially bounded.

Proof sketch. Construct the extractor from the terminal witness upward. Replace each layer's final reply by a fresh invocation of the next extractor, returning its witness only after verification. Apply the inner theorem once to the ordinary distribution of outer prefixes. Then apply the outer theorem to the resulting randomized prover.

The extraction errors consequently add once per layer, rather than once per rewind. Repeated inner calls still contribute to running time and to the computational advantages in ν_i. Conditioning before each fresh call gives the runtime recurrence. A polynomial number of layers alone does not bound its multiplicative cost. Under uniformly quantified Module-SIS hardness at these complete reduction runtimes, the summed computational loss is negligible.

The repaired source reductions fit the same construction. With ℓ independent aggregation rows, the binary source adds $2^{-\ell} + \sqrt{\eta_{\mathrm{src}}}$ to the extraction error. The arithmetic source adds $2p^{-\ell} + \sqrt{\eta_{\mathrm{src},0}} + \sqrt{\eta_{\mathrm{src},1}}$, where p is the smallest prime factor of $2^d + 1$. These source collision experiments include fresh core-extractor calls and their full costs. The source lifting conditions are checked at the actual extracted norm.

Complete restarts. Honest completeness can be amplified by repeating an entire rejected execution. This changes the protocol experiment and must also be reflected in extraction. We fix a maximum of Q attempts before interaction, keep the statement and setup unchanged, and stop at the first real acceptance. The prover may adapt to earlier rejection histories, but no local challenge retries are introduced.

The composed extractor preserves an initial execution of the real complete protocol: its first nested calls reach the terminal check before any extraction fork begins. If that execution rejects, the original post-rejection state is retained for the next attempt. If it accepts, extraction finishes and stops, even if witness recovery fails. Continuing after an accepted execution merely because extraction failed would analyze a different protocol.

Corollary 5.3 (Bounded Complete Restarts)**.** *Under the single-attempt hypotheses for each efficiently generated rejection history, the bounded-restart protocol admits an extractor satisfying*

$$\Pr[S \in \mathcal{R}_{\mathrm{out}}(x)] \geq \varepsilon_{\mathrm{restart}} - Q\delta_{\mathrm{tot}} - \sum_{j=1}^{Q} \nu_j,$$

where δ_{tot} includes the source and all recursive statistical terms. Each ν_j includes initialization by the preceding ordinary attempts. A uniform single-attempt expected extraction bound C gives total expected cost at most QC.

Proof sketch. For attempt j, simulate the preceding ordinary attempts and install a rejecting continuation if any earlier attempt accepted. Its acceptance probability is exactly the probability that the real protocol first accepts at j. These first-acceptance events are disjoint. Because the extractor returns a witness only after its original real execution accepts, its output events can be coupled to the same disjoint branches. Summing the single-attempt bounds proves the result. The uniform state-wise cost bound applies before each reached attempt.

If each reached honest attempt accepts with conditional probability at least $a > 0$, final rejection is at most $(1-a)^Q$, and the expected number of attempts is at most $\min\{Q, 1/a\}$. For the seven-layer configuration with $a > 1/2$, taking $Q = 128$ gives honest rejection below 2^{-128} and fewer than two attempts in expectation. The knowledge-error bound still includes the statistical cost of all attempts and their computational losses.

These guarantees concern the specified interactive ordinary reductions and classical reset access. They do not establish security for the optimized terminal round or the implementation's Fiat–Shamir compilation.

6. Implementation and Quantitative Evaluation

6.1. Evaluation Questions and Configurations

Our evaluation examines whether the counterexamples can be checked independently, whether the corrected construction can prove the same valid source instance, and what costs arise from the revised parameter conditions. We separate exact certificates, executable measurements, and analytical parameter calculations because they establish different results.

The Rust benchmarks use one ordinary reduction and the full five-symbol challenge cube. The seven-layer calculations use filtered challenges and separately selected decomposition parameters. Their analytical encoding bounds are therefore reported separately from measured proof sizes.

6.2. Exact Certificates and Regression Checks

We begin with the 448-variable binary identity instance from the completeness counterexample. The certificate checks the all-one assignment, reconstructs its packed witness, and computes all 36 upper-triangular Gram entries by negacyclic multiplication. It then compares the necessary Gram-digit energy with the prescribed total budget after clearing denominators. Parameter rounding is also checked using integer inequalities. The result does not depend on floating-point tolerances or the behavior of a particular decomposition algorithm. The projection certificates similarly evaluate finite binomial sums and compare their integer numerators against the claimed probability thresholds.

The budget contradiction does not prevent the same source witness from being proved with different parameters. The completeness example therefore passes the same instance and assignment to the repaired Rust implementation. It generates a proof, serializes and decodes it, checks that reencoding preserves the bytes, and verifies the decoded proof. One recorded execution produced an 85,147-byte proof. The source instance and supplied witness remain unchanged, while the implementation uses its own parameter planner and corrected target relation. This is a counterexample-to-repair check, not a comparison between executions of the original and repaired libraries.

The regression suite also tests constraints that an honest round trip alone would not exercise. Source tests modify public pins, commitment values, and scalar replies, then check the resulting target relation against the original witness. Ordinary-round tests insert nonzero padding, alter aggregate constants and outer commitments, and exercise centered decomposition boundaries. A further test feeds one ordinary reduction's output into a second reduction and checks that the verifier independently reconstructs the same target at each step.

TABLE 2. EVALUATION CONFIGURATIONS AND EVIDENCE SCOPE.

Configuration	Method	Evidence provided
Fixed-profile counterexample	Exact integer and rational calculations	Budget contradiction for the 448-variable witness; no execution of the original prover
Rust prototype	Proof generation, serialization, and verification	End-to-end execution on small binary instances
Seven-layer parameter paths	Norm, error, and encoding calculations	Feasibility and communication bounds for specified interactive configurations

TABLE 3. END-TO-END RUST MEASUREMENTS. SIZES AND TIMES ARE MEDIANS OF THREE RUNS. RETRIES COUNT ADDITIONAL COMPLETE ATTEMPTS ACROSS THOSE THREE RUNS.

Variables	Proof bytes	Prove (ms)	Verify (ms)	Retries
64	84,059	400.807	160.807	0
256	84,715	477.447	226.623	0
1,024	84,891	864.888	486.945	1

6.3. End-to-End Measurements

The Rust library exposes setup, parameter planning, proving, verification, and serialization. Its binary frontend supports sparse R1CS inputs and public pins. The backend implements both outer commitments, projection checks, scalar and ring aggregation, folding, recursive target reconstruction, and a terminal opening. Commitment and auxiliary residues use centered base-16 decomposition with eight digits. Public matrices are expanded from seeds, and transcript challenges are derived through SHAKE256. This executable noninteractive prototype is distinct from the interactive distributions covered by the extraction theorems.

We measured synthetic identity R1CS instances with all-one assignments at 64, 256, and 1,024 variables, with equally many constraints. Each size used three fresh setups. Experiments ran sequentially on one thread using a release build of revision 508a5b3, Rust 1.98.0, an Intel Core i9-13980HX, and Windows 11. The profile uses degree 64, modulus 4,294,967,197, source rank 6, inner rank 24, outer rank 9, 160 source aggregation rows, and six scalar aggregations. Each benchmark performs one ordinary reduction with full-cube challenges. Its 1,024-row projection uses threshold $552\beta^2$, rather than the $640\beta^2$ threshold of the seven-layer parameter analysis.

Proof sizes include the header, source messages, recursive messages, and terminal opening. The public parameters and circuit are separate inputs. Setup, encoding, and decoding are timed separately; proving time includes complete-session retries.

All nine recorded proofs verified after serialization and decoding. Across the tested sizes, median proof size increased by 832 bytes, while proving and verification times increased more substantially. One 1,024-variable run required a second complete attempt and took 1,274.292 ms; the other two took 840.440 and 864.888 ms. Median encoding times were 0.643–0.713 ms, and median decoding times were 2.815–3.152 ms. Small differences in serialized size reflect the per-polynomial integer encoding. These measurements establish end-to-end execution for the tested configurations, rather than large-instance scaling or a reliable estimate of retry probability.

6.4. The Cost of Restoring the Guarantees

The small executable instances do not determine the cost of maintaining the repaired conditions across a larger recursion. We therefore evaluate seven-layer parameter paths, recalculating the target dimensions, honest and extracted norm caps, decomposition widths, and collision radii at every layer.

Controlled parameter comparison. We compare two repaired challenge profiles within the same finite search domain. Both use degree 64, modulus 4,294,967,197, 1,024 projection rows, threshold $640\beta^2$, inner rank 24, and outer ranks 16. They share the multiplicity schedule $(112, 33, 15, 8, 7, 5, 4)$, initial block length 37,450, initial norm bound 16,384, and the same collision-radius limits. Only the challenge family and its proved folding bounds change.

The reference profile filters the full five-symbol cube at operator threshold 34 and uses fold cap $(21/2)d\beta^2$. The bounded-energy profile restricts challenge energy to 104, filters at threshold 30, and uses fold cap $9d\beta^2$. The reference is itself a repaired configuration, not LaBRADOR's original parameter selection.

We separate the effect of this distribution change from subsequent parameter optimization. The optimized filtered-cube profile has a core encoding bound of 185.734 KiB. Reusing its decomposition choices with the bounded-energy profile reduces the bound to 182.023 KiB, saving 30,400 bits. Reoptimizing the decomposition path reduces it further to 179.648 KiB, saving another 19,456 bits. The total reduction is 6,232 bytes, or approximately 3.28%. The search computes the exact minimum within the specified finite domain; it is not a global optimum over protocol designs. All three figures include ordinary prover messages and the final opening, but exclude source messages, verifier messages, setup, framing, and Fiat–Shamir material.

The savings have a direct cause. With the old decomposition choices fixed, the smaller fold cap reduces the next-layer norm bounds, while the lower operator threshold reduces the corresponding inner collision bounds. Reoptimization then adjusts the digit path to exploit these changes. The improvement is not uniform across every measure:

after 128 complete attempts, the calculated statistical-error bounds are approximately $2^{-133.79}$ for the filtered cube and $2^{-128.69}$ for the bounded-energy family. Both meet the selected statistical target, but the latter retains less margin. Neither figure includes the computational extraction losses.

Complete communication accounting. A separate parameter ledger includes the native binary source and seven ordinary layers for $n_{\mathrm{src}} = k = 2^{25}$. This configuration uses source rank 6, inner ranks 24, and outer ranks 9. It also replaces independent aggregation rows with polynomial fingerprinting [24], so its statistical accounting includes the corresponding degree-dependent error terms. These changes distinguish it from the controlled core-only comparison above.

For one complete attempt, this ledger gives 1,288,313 prover-message bits, rounded to 161,040 bytes, or 157.266 KiB. The figure includes source and ordinary-core messages but excludes setup, the statement, framing, rejected attempts, and Fiat–Shamir material. Explicitly representing the verifier challenges requires approximately 67.54 GiB, dominated by the dense projection matrices. This is a communication cost of the specified interactive representation, not a noninteractive proof size or a lower bound for alternative constructions. The smaller prover-message figure reflects different commitment ranks and aggregation choices, not an additional saving attributable solely to the challenge change.

These results cannot be converted into an overhead ratio against LaBRADOR's reported 58 KB proof. That figure concerns arithmetic R1CS modulo $2^{64}+1$ with 2^{20} constraints, whereas the evaluated paths use binary sources and ordinary terminal openings. The source relation, commitment parameters, encoding, and terminal protocol must match before such a comparison is meaningful. Recovering the original size claim requires a corrected instantiation of that same construction.

7. Implications for Downstream Constructions

We trace the probability bound, extracted-witness norm, or terminal guarantee used by each construction. The accompanying register records 19 downstream works and six optional or proposed configurations, with source locations for the relevant claims. The following cases show how these dependencies enter concrete security and parameter calculations.

7.1. Inherited Bounds and Parameters

RoKoko explicitly imports LaBRADOR's modular projection bound with 256 ternary rows, lower threshold $\sqrt{30}$, and failure probability 2^{-128} [6, Lemma 5]. Its premises require $0 < \theta \le q/125$, a centered witness, and $\|w\|_2 \ge \theta$. The example in Section 3.3 satisfies these conditions with $\theta = a$, while its strict lower-tail event already has probability exceeding 2^{-127}. This event is contained in the lemma's non-strict lower-tail event, so the stated upper bound fails. Lemma 6 applies that bound blockwise, and Lemma 7 carries the resulting error into the knowledge-soundness analysis of the committed coarse projection. Those concrete projection and extraction budgets require a corrected lower-tail bound.

The parameter dependence is explicit in Albrecht and Gur's lattice VOPRF [15]. Its client-request proof uses LaBRADOR compression, and Appendix B.4 fixes `LABRADOR_SLACK` to $\sqrt{128/30}$. In Appendix B.3, this factor scales an extracted-response bound that feeds the Module-SIS solution bound and the search for the commitment dimension. The selected dimensions and decomposition parameters then enter the bandwidth calculation. The direct LaBRADOR estimator likewise uses the slack to choose inner and outer commitment ranks before computing proof size. Thus, revising the extraction analysis requires revalidating the parameter search and the reported client-proof bandwidth. The size of any resulting change must be determined by that calculation.

7.2. Backend and Instantiation Dependencies

Greyhound reduces polynomial evaluation to a final proof of a short linear relation and instantiates this terminal argument with LaBRADOR [3, Sections 4.2–4.3]. Its PCS theorem inherits the terminal completeness error and subtracts the terminal knowledge error from the extraction guarantee; the composition proof invokes LaBRADOR's Lemma 3.7. The end-to-end guarantee therefore needs valid terminal norm bounds and a compatible recursive extraction argument. Correcting those bounds also requires checking the terminal parameters used to estimate the complete evaluation-proof size.

The implementation must be assessed under its actual distribution. Greyhound's frontend performs no JL projection, while the joint C implementation changes the LaBRADOR backend's projection coefficients to uniform signs [3, Sections 5–6]. Its projection tails require calibration for that Rademacher distribution. The ternary event in Section 3.3 does not evaluate that probability; the terminal norm and composition requirements still determine which backend guarantees the PCS can use.

Optional and proposed compression configurations assign a further proof guarantee or cost to the selected LaBRADOR backend [25], [26], [27], [28], [29], [30]. For LatticeFold, this dependency arises when the final IVC or PCD statements are compressed using LaBRADOR. That compressed proof requires a valid terminal instantiation; the folding protocol's sumcheck-based norm analysis supplies a separate guarantee. Likewise, measurements taken before an optional compression step establish the cost of that measured stage. The additional compression cost must be calculated for the corrected backend configuration.

8. Discussion and Outlook

When do parameter failures matter? A failed norm check can lead to a restart or a change of parameters, but these responses are not interchangeable. In our completeness

example, the excess comes from Gram terms fixed before the folding challenges. Resampling those challenges cannot make the same witness fit the original budget. Changing the decomposition or target bound can restore feasibility, as demonstrated by the repaired execution on the same source instance. The new parameters, however, require their own norm, binding, and cost calculations.

The example establishes a failure of the fixed configuration, not its frequency in application workloads. Successful tests on typical witnesses cannot establish a guarantee for all valid witnesses; conversely, a counterexample to completeness is not an accepted proof of a false statement. Implementations that permit retries should also specify which failures trigger them, whether parameters may change, and how rejected attempts are charged. A bounded complete-restart policy has a different analysis from locally resampling an unfavorable challenge.

Making parameter selection verifiable. Heuristics remain useful for finding compact candidate configurations. A practical safeguard is to separate this search from validation. Each candidate should carry a record of the complete witness layout, decomposition bounds, challenge distribution, and norm limits required by extraction. The witness budget must include auxiliary terms and final decomposition quotients, rather than only the folded vector. At recursive boundaries, validation must use the norm that the next extractor can return, not merely the norm expected during honest execution.

An independent checker could verify these arithmetic conditions and reject configurations outside their certified range before proving begins. Changes to the challenge sampler or decomposition convention should invalidate the corresponding certificate until its conditions are checked again. Regression tests should also target constraints that honest executions rarely expose, such as public-input pins and padding zeros. Such tooling would reduce the risk of silently carrying obsolete bounds into an optimized implementation. It would not replace the mathematical proofs or establish the hardness of the resulting Module-SIS instances.

From repaired reductions to practical succinct proofs. Our reconstruction provides guarantees for a specified interactive construction and an executable prototype with a different challenge-generation mechanism. The interactive extraction theorem does not automatically establish security of the prototype's Fiat–Shamir compilation. Concrete knowledge security also requires accounting for the computational reduction losses at the full extraction runtime, including recursive calls and complete restarts. These obligations remain distinct from meeting a statistical-error target.

There is nevertheless room to improve efficiency within explicit constraints. Our controlled parameter comparison shows that changing the challenge distribution and reoptimizing decomposition choices can reduce encoding costs while preserving the selected feasibility conditions. The resulting costs are not lower bounds for other repaired constructions. Structured projections and alternative terminal protocols may offer further savings, but their probability bounds and extraction interfaces must be established for the modified protocol. Compactness should ultimately be reported for the same construction and parameter profile covered by the security analysis.

References

[1] W. Beullens and G. Seiler, "LaBRADOR: Compact proofs for R1CS from module-SIS," in *CRYPTO 2023, Part V*, ser. Lecture Notes in Computer Science, vol. 14085. Springer, 2023, pp. 518–548.

[2] ——, "LaBRADOR: Compact proofs for R1CS from module-SIS," Cryptology ePrint Archive, Paper 2022/1341, 2022. [Online]. Available: https://eprint.iacr.org/2022/1341

[3] N. K. Nguyen and G. Seiler, "Greyhound: Fast polynomial commitments from lattices," in *CRYPTO 2024*, 2024, full version: Cryptology ePrint Archive, Paper 2024/1293. [Online]. Available: https://eprint.iacr.org/2024/1293

[4] M. A. Aardal, D. F. Aranha, K. Boudgoust, S. Kolby, and A. Takahashi, "Aggregating falcon signatures with LaBRADOR," in *CRYPTO 2024*, 2024, full version: Cryptology ePrint Archive, Paper 2024/311. [Online]. Available: https://eprint.iacr.org/2024/311

[5] V. Lyubashevsky, G. Seiler, and P. Steuer, "The LaZer library: Lattice-based zero knowledge and succinct proofs for quantum-safe privacy," in *ACM CCS 2024*, 2024, full version: Cryptology ePrint Archive, Paper 2024/1846. [Online]. Available: https://eprint.iacr.org/2024/1846

[6] M. Klooss, R. W. F. Lai, N. K. Nguyen, M. Osadnik, and L. Tucci, "RoKoko: Lattice-based succinct arguments, a committed refinement," in *ASIACRYPT 2026*, 2026, full version: Cryptology ePrint Archive, Paper 2026/575. [Online]. Available: https://eprint.iacr.org/2026/575

[7] M. Bolboceanu, J. Bootle, V. Lyubashevsky, A. Merino-Gallardo, and G. Seiler, "Orthus: Practical sublinear batch-verification of lattice relations from standard assumptions," in *CRYPTO 2026*, 2026, full version: Cryptology ePrint Archive, Paper 2026/398. [Online]. Available: https://eprint.iacr.org/2026/398

[8] A. Kallesoe and H. Khoshakhlagh, "Grand danois: Succinct multilinear polynomial commitments over lattices," Cryptology ePrint Archive, Paper 2026/1196, 2026. [Online]. Available: https://eprint.iacr.org/2026/1196

[9] B. Chen, "Symphony: Scalable SNARKs in the random oracle model from lattice-based high-arity folding," in *ASIACRYPT 2026*, 2026, full version: Cryptology ePrint Archive, Paper 2025/1905. [Online]. Available: https://eprint.iacr.org/2025/1905

[10] B. Biasioli, M. Bolboceanu, V. Lyubashevsky, A. Merino-Gallardo, M. Osadnik, G. Seiler, and P. Steuer, "A toolkit for succinct lattice-based zero knowledge proofs," in *ACM CCS 2026*, 2026, full version: Cryptology ePrint Archive, Paper 2026/1289. [Online]. Available: https://eprint.iacr.org/2026/1289

[11] Y. Lin, H. Wang, M. F. Esgin, A. Sakzad, and R. Steinfeld, "Lemur+: Compact post-quantum synchronized multi-signatures and multi-hop aggregation," Cryptology ePrint Archive, Paper 2026/2054, 2026. [Online]. Available: https://eprint.iacr.org/2026/2054

[12] T. Tomita and J. Shikata, "Compact aggregate signature from module-lattices," Cryptology ePrint Archive, Paper 2023/471, 2023. [Online]. Available: https://eprint.iacr.org/2023/471

[13] S. Jeon, C. A. Haidar, and M. Tibouchi, "LastRings: Lattice-based scalable threshold ring signatures," in *ISC 2025*, 2025, full version: Cryptology ePrint Archive, Paper 2025/1633. [Online]. Available: https://eprint.iacr.org/2025/1633

[14] M. R. Albrecht, A. Davidson, A. Deo, and D. Gardham, "Crypto dark matter on the torus: Oblivious PRFs from shallow PRFs and FHE," in *EUROCRYPT 2024*, 2024, full version: Cryptology ePrint Archive, Paper 2023/232. [Online]. Available: https://eprint.iacr.org/2023/232

[15] M. R. Albrecht and K. D. Gur, "Verifiable oblivious pseudorandom functions from lattices: Practical-ish and thresholdisable," in *ASIACRYPT 2024*, 2024, full version: Cryptology ePrint Archive, Paper 2024/1459. [Online]. Available: https://eprint.iacr.org/2024/1459

[16] A. Faonio and L. Tong, "Non-malleable reductions of knowledge," in *ASIACRYPT 2026*, 2026, full version: Cryptology ePrint Archive, Paper 2026/1822. [Online]. Available: https://eprint.iacr.org/2026/1822

[17] N. K. Nguyen, G. O'Rourke, and J. Zhang, "Hachi: Efficient lattice-based multilinear polynomial commitments over extension fields," Cryptology ePrint Archive, Paper 2026/156, 2026. [Online]. Available: https://eprint.iacr.org/2026/156

[18] R. del Pino, S. Katsumata, G. Niot, M. Reichle, and K. Takemure, "Unmasking TRaccoon: A lattice-based threshold signature with an efficient identifiable abort protocol," in *CRYPTO 2025*, 2025, full version: Cryptology ePrint Archive, Paper 2025/849. [Online]. Available: https://eprint.iacr.org/2025/849

[19] M. Osadnik and G. Seiler, "LaBinius – fast lattice-based binary polynomial commitment scheme," Cryptology ePrint Archive, Paper 2026/2103, 2026. [Online]. Available: https://eprint.iacr.org/2026/2103

[20] L. Hanzlik, M. Manulis, M. Mula, A. Pulval-Dady, T. Silde, and D. Slamanig, "BRaccoon: Concurrently secure blind lattice signatures from raccoon," Cryptology ePrint Archive, Paper 2026/1084, 2026. [Online]. Available: https://eprint.iacr.org/2026/1084

[21] M. F. Esgin, R. Steinfeld, E. Tairi, and J. Xu, "LeOPaRd: Towards practical post-quantum oblivious PRFs via 2HashDH paradigm," in *CRYPTO 2026*, 2026, full version: Cryptology ePrint Archive, Paper 2024/1615. [Online]. Available: https://eprint.iacr.org/2024/1615

[22] I. G. Shevtsova, "On the absolute constants in the Berry–Esseen inequality and its structural and nonuniform improvements," *Informatika i Ee Primeneniya*, vol. 7, no. 1, pp. 124–125, 2013. [Online]. Available: https://www.mathnet.ru/eng/ia252

[23] V. Lyubashevsky and G. Seiler, "Short, invertible elements in partially splitting cyclotomic rings and applications to lattice-based zero-knowledge proofs," in *EUROCRYPT 2018, Part I*, ser. Lecture Notes in Computer Science, vol. 10820. Springer, 2018, pp. 204–224.

[24] J. Thaler, "Proofs, arguments, and zero-knowledge," Author manuscript, 2023. [Online]. Available: https://people.cs.georgetown.edu/jthaler/ProofsArgsAndZK.pdf

[25] M. Klooß, R. W. F. Lai, N. K. Nguyen, and M. Osadnik, "RoK, paper, SISsors – toolkit for lattice-based succinct arguments," in *ASIACRYPT 2024*, 2024, full version: Cryptology ePrint Archive, Paper 2024/1972. [Online]. Available: https://eprint.iacr.org/2024/1972

[26] D. Boneh and B. Chen, "LatticeFold: A lattice-based folding scheme and its applications to succinct proof systems," in *ASIACRYPT 2025*, 2025, full version: Cryptology ePrint Archive, Paper 2024/257. [Online]. Available: https://eprint.iacr.org/2024/257

[27] L. Zhang, S. S. M. Chow, S. Gao, and B. Xiao, "Serval: Slack-free ℓ_2-sound polynomial commitments from lattices," Cryptology ePrint Archive, Paper 2025/1903, 2025. [Online]. Available: https://eprint.iacr.org/2025/1903

[28] S. Faller, G. Niot, and M. Reichle, "Lattice-based threshold blind signatures," in *IEEE Symposium on Security and Privacy 2026*, 2026, full version: Cryptology ePrint Archive, Paper 2025/1566. [Online]. Available: https://eprint.iacr.org/2025/1566

[29] J. Bootle, V. Lyubashevsky, and A. Merino-Gallardo, "Efficient verifiable mixnets from lattices, revisited," in *PKC 2025*, 2025, full version: Cryptology ePrint Archive, Paper 2025/658. [Online]. Available: https://eprint.iacr.org/2025/658

[30] W. Beullens, V. Lyubashevsky, N. K. Nguyen, and G. Seiler, "Lattice-based blind signatures: Short, efficient, and round-optimal," in *ACM CCS 2023*, 2023, full version: Cryptology ePrint Archive, Paper 2023/077. [Online]. Available: https://eprint.iacr.org/2023/077

Appendix A. Exact Counterexample Calculations

A.1. Parameter Certificate for Proposition 3.1

The profile in Section 3.1 gives $b = 9$, $(t_1, b_1) = (10, 10)$, and $(t_2, b_2) = (2, 10)$. The rounding checks are

$$\begin{aligned} &17^4 < 16 \cdot 6816 < 19^4, \quad && 9^{19} < q^2 < 9^{21}, \\ &9^{10} < q \le 10^{10}, && 9^3 < 10752 < 9^5, \\ &19^4 < 16 \cdot 10752 < 21^4. \end{aligned}$$

These certify b, t_1, b_1, t_2, b_2, respectively, with no ties. The source conditions follow from $1792 < 15q/128$ and $2688 < q$. The projection condition is $3584 \cdot 128 \cdot 125^2 < 30q^2$. Substitution gives $(\beta')^2 = 508928/81 + 1024000 + 230400 = 102115328/81$. The Gram energy computed in the proposition then gives the exact excess $2177842688/8181 > 0$.

A.2. Arithmetic Relation Counterexamples

A.2.1. Lifting changes an honest residual. The instance in Section 3.2 uses $d = 4$, $N = 17$, $q = 13$, $n = k = \ell = 1$, and $\beta^2 = 801/100$. Signed non-adjacent binary expansions of representatives in $[-8, 8]$ encode every source residue with degree below four and squared norm at most two.

With $(A, B, C, w) = (15, 11, 14, 4)$, the source values are $(a, b, c) = (9, 10, 5)$ and satisfy $ab = c$. The challenge $\phi = 3$ gives $e = 10$. The second challenges $(\alpha, \rho, \gamma, \delta) = (12, 8, 2, 7)$ test

$$f = \alpha(Aw - a) + \rho(Bw - b) + \gamma(Cw - c) + \delta(\phi a - e) + eb - \phi c.$$

In the order (a, b, c, w, e), the encoded witness and combined linear coefficients are

$$\begin{aligned} s &= (-X^3, 1 - X^3, 1 + X^2, X^2, 1 - X^3), \\ \lambda &= (9, 9, 12, 7, 10) \pmod{17}. \end{aligned}$$

Encoding λ and adding the quadratic term yields $F = 2 - X - 7X^2 - 5X^3$, with $\|s\|_2^2 = 8$. Its evaluation is $F(2) = -68 = 0 \pmod{17}$. Centered reduction replaces -7 by 6, so the evaluation becomes $-16 = 1 \pmod{17}$.

The lifting premise holds because $13 - 801/200 > 0$ and

$$\left(13 - \frac{801}{200}\right)^2 - \frac{801}{10} = \frac{32401}{40000} > 0.$$

Also, $X^4 + 1 = (X^2 + 5)(X^2 + 8)$ has two irreducible factors modulo 13. The separate composition condition fails: $Md = 20 > 0.3q$.

A.2.2. A missing residual constant. With $d = 4$, $N = 17$, $q = 101$, $(A, B, C, w) = (1, 1, 3, 3)$, and $\phi = 1$, the source values are $a = b = e = 3$ and $c = 9$. All second-stage linear coefficients are zero. The encodings $E(3) = X^2 - 1$ and $E(9) = -X^3$ yield

$$\widetilde{f} = (X^2 - 1)^2 + X^3 = X^2(X - 2) \neq 0 \quad \text{in } R_{101}.$$

Its evaluation vanishes. The witness has squared norm nine; $\beta^2 = 10$ gives lifting bound $15 < 101/2$. Thus the honest reply $g = \widetilde{f}$ passes the source test. The target must impose $\widetilde{f} - g = 0$ to retain this witness.

A.3. Exact Projection Probabilities

The lower-tail event in Section 3.3 has probability at least the displayed binomial sum. Multiplying that sum by 2^{6257} reduces its comparison with 2^{-127} to

$$2^{6001} \sum_{k=0}^{29} \binom{256}{k} + \binom{256}{30} \left(2^{6000} - \binom{6000}{3000} \right) > 2^{6130}.$$

The event conditions follow from $29a^2 + 5800a + 2560000 < 30a^2$ and $2560000 < 2a$, with $a = 2^{21}$.

For $w = 1^{4096}$, the row inner product is $\mathrm{Bin}(8192, 1/2) - 4096$. The strict upper tail starts at 609, since $9.5\|w\|_2 = 608$. Hence

$$\Pr[|\langle \pi, w \rangle| > 9.5\|w\|_2] = 2^{-8191} \sum_{j=4705}^{8192} \binom{8192}{j} > 2^{-135}.$$

The integer numerator exceeds 2^{8056}, which also certifies the violation of the 2^{-141} bound.

Appendix B. Proofs of the Repaired Relations and Bounds

B.1. Source Relations

The signed non-adjacent expansion $E(v)$ of a centered representative in $S = \mathbb{Z}/(2^d + 1)\mathbb{Z}$ has degree below d. Its nonzero digits are separated, so $\|E(v)\|_2^2 \le d/2$. It preserves evaluation: $\varphi(E(v)) = v$. After combining the linear coefficients in the source ring, the residual of Section 4.1 is $\widetilde{f}_i(s) = \langle s_{e_i}, s_b \rangle + \sum_v \mathrm{Enc}(\lambda_{i,v}) s_v$.

Lemma B.1 (Centered lifting)**.** *A target witness of norm at most β_{ext} satisfies $L_q(\widetilde{f}_i(s)) = \widetilde{f}_i^{\mathbb{Z}}(L_q(s))$ if $\sqrt{Md/2}\,\beta_{\mathrm{ext}} + \beta_{\mathrm{ext}}^2/2 < q/2$, where $M = n_{\mathrm{src}} + (3 + \ell)k$. Here $\widetilde{f}_i^{\mathbb{Z}}$ uses $E(\lambda_{i,v})$. Thus $\varphi(L_q(g_i)) = 0$ implies the decoded source equation $f_i = 0$.*

Proof. Each coefficient of a negacyclic product is a signed inner product. Cauchy–Schwarz bounds the linear contribution by $\sqrt{Md/2}\,\beta_{\mathrm{ext}}$. The distinct blocks s_{e_i}, s_b contribute at most $\|s_{e_i}\|_2 \|s_b\|_2 \le (\|s_{e_i}\|_2^2 + \|s_b\|_2^2)/2 \le \beta_{\mathrm{ext}}^2/2$. Every coefficient lies strictly within the centered interval. Lifting recovers the integer residual, and $\varphi : R \to S$ recovers its source value. □

The binary equations admit a direct integer argument. Conjugation commutes with centered lifting. The bitness residual is $\sum_j u_j(u_j - 1)$, and $0 \le z(z-1) \le 2z^2$ for integers z. The witness includes conjugate copies, so this residual lies in $[0, \beta_{\mathrm{ext}}^2]$. Its zero congruence and $\beta_{\mathrm{ext}}^2 < q$ force every coefficient to be a bit. Then $v_j = a_j + b_j - 2c_j$ lies in $\{-2, -1, 0, 1, 2\}$. Each $v_j(v_j - 1)$ lies in $[0, 6]$. The product equation and $6k < q$ force every term to vanish, equivalently $a_j b_j = c_j$.

The aggregation coefficients satisfy $\delta_i = A^\top \alpha_i + B^\top \rho_i + C^\top \gamma_i + H^\top \eta_i$ over $\mathbb{F}_2$. The integer reply $G_i = \langle \alpha_i, a \rangle + \langle \rho_i, b \rangle + \langle \gamma_i, c \rangle - \langle \delta_i, w \rangle$ and the transmitted even reply g_i both lie in $[-n_{\mathrm{src}}, 3k]$. Since this interval has width below q, their congruence gives $G_i = g_i$. Its parity recovers the binary check. The pins and zero-padding constraints are those of Section 4.1.

B.2. Decomposition Carries

Lemma B.2 (Centered decomposition)**.** *Centered base-a decomposition into $t \ge 2$ digits satisfies $\|\mathrm{Dec}_{a,t}(x)\|_2^2 \le F_{a,t}(D, L)$ for $x \in \mathbb{Z}^D$ with $\|x\|_2 \le L$, using $F_{a,t}$ and $C(a,t)$ from Section 4.2.*

Proof. The quotient recurrence $x^{[j+1]} = (x^{[j]} - x_j)/a$, with $x^{[0]} = x$ and $|x_j| \le a/2$, gives

$$|x^{[t-1]}| \le \frac{|x|}{a^{t-1}} + \frac{1}{2} \sum_{j=0}^{t-2} a^{-j} = \frac{|x|}{a^{t-1}} + C(a,t).$$

The low digits contribute at most $D(t-1)a^2/4$. The triangle inequality bounds the quotient vector by $L/a^{t-1} + \sqrt{D}\,C(a,t)$. Their squared norms sum to $F_{a,t}(D, L)$ under either fixed centered tie convention. □

This supplies the carry bound used in the proof of Proposition 4.1. The component bounds and paired partition are established in Section 4.2.

B.3. Proof of Theorem 4.2

Proof. The centered vector w is fixed before Π is sampled. A normalized row product $Y = \langle \pi, v \rangle / \|v\|_2$ has $\mathbb{E}Y^2 = 1/2$ and $\mathbb{E}Y^4 = 3/4 - \sum_i v_i^4/(4\|v\|_2^4) \le 3/4$. Hence $e^{-x} \le 1 - x + x^2/2$ gives $\mathbb{E}e^{-Y^2/2} \le 27/32$. The representation $\chi = (\epsilon + \epsilon')/2$ by independent uniform signs also gives

$$\Pr[|\langle \pi, v \rangle| \ge t] \le 2e^{-t^2/\|v\|_2^2}.$$

The acceptance threshold is $t_0 = \sqrt{120}\,u$.

Case 1: $\|w\|_2 < q/12$. Acceptance requires either the unreduced event or a row of unreduced magnitude at least $q - t_0$. Independence and a union bound give

$$\Pr[\|\Pi w \bmod q\|_2^2 \le 120u^2] \le e^{60}(27/32)^{1024} + 2048e^{-144(1-\sqrt{120}/250)^2}.$$

Since $163/60 < e < 87/32$, $\sqrt{120} < 11$, and $144(239/250)^2 > 131$, this is strictly below

$$(87/32)^{60}(27/32)^{1024} + 2048(60/163)^{131} < 2^{-164}.$$

The last comparison follows by clearing denominators.

Case 2: $\|w\|_\infty \ge q/72$. Conditioning on the other coordinates leaves $a, a + w_i, a - w_i$ modulo q, with masses $1/2, 1/4, 1/4$. An open arc of length $q/72$ cannot contain a

and either shifted point. A row therefore has modular magnitude at least $q/144$ with probability at least $1/2$, including boundary atoms. Forty such rows exceed $120q^2/250^2$. Acceptance has probability at most $2^{-1024}\sum_{j=0}^{39}\binom{1024}{j} < 2^{-788}$.

Case 3: $\|w\|_2 \ge q/12$ *and* $\|w\|_\infty < q/72$. Deleting coordinates until the remaining vector v first has norm below $q/12$ gives $\sqrt{35}\,q/72 < \|v\|_2 < q/12$. Berry–Esseen [22] applied to $v_i\epsilon_i/2, v_i\epsilon_i'/2$ bounds the distance from $X = \langle \pi, v\rangle$ to its matching normal law by

$$\Delta \le \frac{0.56\sum_i |v_i|^3}{\sqrt{2}\,\|v\|_2^3} < \frac{0.56}{\sqrt{70}}.$$

Conditioning on the independent remainder leaves one interval or two boundary intervals of total length $2t_0$ in the centered interval. The normal density is at most $1/(\sqrt{\pi}\|v\|_2)$. In the two-interval case, taking the complement uses only two distribution-function errors. Actual and normal tails outside the centered interval each cost at most $2e^{-36}$. Thus a row is small with probability below

$$\frac{144\sqrt{120}}{250\sqrt{35\pi}} + \frac{1.12}{\sqrt{70}} + 4e^{-36} < \frac{736}{1000}.$$

Acceptance requires every row to be small, giving $(736/1000)^{1024} < 2^{-452}$. The three cases exhaust all centered w. □

The honest bound follows from the unreduced energy $S = \sum_{j=1}^{J}\langle \pi_j, w\rangle^2$. Its mean is $(J/2)\|w\|_2^2$ and its variance is $J(\|w\|_2^4/2 - \sum_i w_i^4/4)$. At $J = 1024$, Cantelli bounds rejection above $640\beta^2$ by $1/33$. Centering can only reduce the energy; the zero witness always passes.

B.4. Proof of Proposition 4.3

Proof. The count $a_j(e)$ of length-j words with energy e obeys

$$a_j(e) = a_{j-1}(e) + 2a_{j-1}(e-1) + 2a_{j-1}(e-4),$$

with $a_0(0) = 1$ and out-of-range counts zero. Hence $N_B = |\mathcal{B}_{64,104}| = \sum_{e=0}^{104} a_{64}(e)$.

Under the auxiliary product law proportional to t^{c^2}, where $t = 4/5$, normalization is $Z_t = 2137/625$ and variance is $v_t = 3048/2137$. The fourth moment is at most $3v_t^2$. Later even-moment ratios are at most four, below the Gaussian ratios $5v_t, 7v_t, \ldots$. Termwise comparison gives $\mathbb{E}_t e^{uC} \le e^{v_t u^2/2}$.

At an odd 128th root, each directional real projection has squared coefficient sum 32. A 12-direction net detects at least $\cos(\pi/12)$ of any complex magnitude. Since $\cos^2(\pi/12) > 933/1000$, the 32 conjugate root representatives give rejection below $384e^{-x}$, with $x = (933/1000)900/(64v_t)$. Every word in $\mathcal{B}$ has auxiliary mass at least t^{104}/Z_t^{64}, so

$$\frac{|\mathcal{B}\setminus\mathcal{C}|}{N_B} < \frac{Z_t^{64}t^{-104}}{N_B}\,384e^{-x} < \frac{1}{3}.$$

Appendix D.1 certifies the last comparison and $\lfloor 2N_B/3\rfloor + 1 > 2^{143}$.

The two irreducible factors of $X^{64}+1$ are $X^{32} \pm \omega$, where $\omega^2 = -1$. A challenge difference $u + X^{32}v$ has coefficients in $[-4, 4]$. Vanishing in either factor gives $u_j \pm \omega v_j = 0$, hence $u_j^2 + v_j^2 = 0$ modulo q. Since $u_j^2 + v_j^2 \le 32 < q$, the difference is zero. Distinct challenges therefore differ by a unit. □

The recurrence supplies completion counts for uniform sampling from $\mathcal{B}$. Filtering by the exact operator predicate then samples $\mathcal{C}$ uniformly, using fewer than $3/2$ candidates in expectation.

B.5. Proof of Proposition 4.4

Lemma B.3 (Conditional folding bound)**.** *Let $d \ge 2$ be a power of two. A vector $X \in \mathbb{R}^d$ has finite fourth moments and a law invariant under permutations and independent sign changes. Conditioning on an event F of probability $p > a > 0$ gives $C = X \mid F$, with $\mathbb{E}C = 0$ and $\mathrm{Cov}(C) = vI$. The raw moments $v_0 = \mathbb{E}X_0^2$, $m_4 = \mathbb{E}X_0^4$, and $m_{22} = \mathbb{E}X_0^2X_1^2$ determine*

$$V_U = dm_4 + d(d-1)m_{22} - d^2v_0^2,$$
$$g = 2m_{22} - v_0^2 + \frac{m_4 - 3m_{22}}{d}, \qquad \rho = \sqrt{\frac{V_U(1-a)}{ad^2}}.$$

With $f \in (0,1)$ and

$$G = \max\left\{ v_0 + \sqrt{\frac{g}{af}},\ (v_0+\rho)\left(1 + \sqrt{\frac{1-f}{f}}\right)\right\},$$

independent copies C_i and fixed ring vectors s_i satisfy $\Pr[\|\sum_i C_i s_i\|_2^2 > Gd\beta^2] \le f$ whenever $\sum_i \|s_i\|_2^2 \le \beta^2$.

Proof. The stacked negacyclic multiplication matrix M_s has Gram matrix $H = M_s^\top M_s$ with diagonal $t = \|s\|_2^2$ and trace dt. Multiplication by $X^{d/2}$ commutes with these maps and gives an orthogonal complex structure. Eigenvalues of H and the output covariance occur in pairs, so $\mathrm{tr}(H^2) \le (\mathrm{tr}\,H)^2/2$.

Raw sign and permutation symmetry give

$$\mathbb{E}(X^\top HX)^2 = m_{22}\big((\mathrm{tr}\,H)^2 + 2\,\mathrm{tr}(H^2)\big) + (m_4 - 3m_{22})\sum_j H_{jj}^2.$$

The raw variance is at most gd^2t^2. Centering at the raw mean before conditioning yields

$$\mathrm{Var}(\|M_sC\|_2^2) \le \left(\frac{g}{p} - (v-v_0)^2\right)d^2t^2.$$

Independent $Y_i = M_{s_i} C_i$ have zero means. Their covariances obey $\operatorname{tr}(\Sigma_i \Sigma_j) \le \operatorname{tr}(\Sigma_i)\operatorname{tr}(\Sigma_j)/2$ by the paired eigenvalues. Expanding the variance therefore gives

$$\operatorname{Var}\left(\left\|\sum_i C_i s_i\right\|_2^2\right) \le d^2 \left(\sum_i \|s_i\|_2^2\right)^2 \max\left\{\frac{g}{p} - (v - v_0)^2,\ v^2\right\}.$$

Also, $\operatorname{Cov}(\|X\|_2^2, 1_F) = pd(v - v_0)$, so Cauchy–Schwarz gives $|v - v_0| \le \rho$.

Cantelli uses $k = (1-f)/f$. With $u = v - v_0$ and $h = g/p$, the first variance branch requires at most $v_0 + \sqrt{g/(af)}$, since $u + \sqrt{k(h-u^2)} \le \sqrt{(1+k)h}$. The second requires $v(1+\sqrt{k}) \le (v_0+\rho)(1+\sqrt{k})$. Both are bounded by G. □

Proof of Proposition 4.4. The raw law is uniform on $\mathcal{B}_{64,104}$ and $F = \{\|X\|_{\mathrm{op}} < 30\}$ has probability above $2/3$. Global negation gives conditional mean zero; signed cyclic shifts equalize variances. Each off-diagonal covariance reduces to $\mathbb{E} C_0 C_\delta$. Writing $\delta = 2^e u$ with u odd, the automorphism $X \mapsto X^{1+64/2^e}$ fixes coefficient zero and negates coefficient δ. It preserves both filters, so this covariance vanishes.

The exact moments in Appendix D.1 satisfy $(9 - v_0)^2 > 24g$, $v_0 < 25/16$, and $V_U < 32$. At $a = 2/3$, $f = 1/16$, these give $\rho < 1/16$. Both branches of G are below nine: the second is below $(26/16)(39/8) < 9$ because $\sqrt{15} < 31/8$. Thus $9d\beta^2$ strictly exceeds the Cantelli threshold. The rejection bound is strictly below $1/16$, and centering preserves it. When $\beta = 0$, the folded vector is zero. □

Appendix C. Knowledge Extraction and Composition Proofs

C.1. Collision Experiments

An *anchor* records Algorithm 1's completed candidate, recovery certificates, and gate transcript before the final validity check. Failure or collision returns $\perp$. All experiments use the setup and reset model of Section 5.1. Their probabilities include initialization and anchor generation; subsequent tests use fresh randomness.

The response bound is $Z = (b+1)B$. The η_0 experiment compares the anchor's first opening with an independent accepted execution from the original state. The η_h experiment forks the final challenges twice after a fresh prefix through the second commitment. Distinct openings give collisions at radius $2B$ in the respective outer matrices.

The η_y experiment chooses i uniformly. Its anchor certificate contains a unit d_i and an integer numerator $v_i = d_i S_i \bmod q$ with $\|v_i\|_2 \le 2Z$. On a fresh slice, $y = z - c_i S_i$ is the complement opening. Its scaled value $d_i y$ has representative $d_i z - c_i v_i$ of norm at most $4TZ$. Since d_i is a unit, distinct openings give a nonzero collision at radius $8TZ$.

The η_g experiment selects the first incorrect Gram entry of a valid-norm candidate before drawing test challenges. Its complement opening subtracts $c_i S_i + c_j S_j$ from z, or only $c_i S_i$ on the diagonal. Its norm is at most $Z + 2T\beta_{\mathrm{ext}}$, giving radius $2Z + 4T\beta_{\mathrm{ext}}$.

Except for the η_0 comparison, tests retain only accepted branches with the anchor's first opening. The remaining probability η_E covers verified collisions during anchor generation. These are unconditional probabilities over complete experiments.

C.2. Gate Cost and Loss

Lemma C.1 (Acceptance gate)**.** *From every permitted initial state, Algorithm 1 uses at most $1 + r/(1 - 1/N)$ ordinary or suffix executions in expectation. It produces a complete candidate or a verified collision with probability at least $\varepsilon - r/N$, where ε is ordinary acceptance probability.*

Proof. Condition on the prover tape and all verifier coins except coordinate i. A fraction p of challenges accepts, and the gate reaches this row with probability p. With $\alpha = 1/N$, a reached row has $p \ge \alpha$. A fresh trial accepts with a distinct challenge with probability $p - \alpha$. Including abort, its stopping probability is $\rho = p(1-\alpha) + \alpha^2$. Thus

$$\frac{p}{\rho} \le \frac{1}{1-\alpha}, \qquad \frac{p\alpha(1-p+\alpha)}{\rho} \le \alpha.$$

The second inequality follows from $\rho - p(1-p+\alpha) = (p-\alpha)^2$. Earlier loops can only reduce the mass reaching coordinate i. Summing proves the work and loss bounds. Positive stopping probability gives almost-sure termination. □

C.3. Proof of Theorem 5.1

Proof. We bound invalid candidates using two opening-mass inequalities.

At a fixed prefix, let p_v be the unnormalized mass of anchor-consistent accepted suffixes with opening v, and $a_{\mathrm{acc}} = \sum_v p_v$. The nonmodal mass $e = a_{\mathrm{acc}} - \max_v p_v$ and collision probability $c = a_{\mathrm{acc}}^2 - \sum_v p_v^2$ satisfy $c \ge a_{\mathrm{acc}} e \ge e^2$. Cauchy–Schwarz therefore gives, for any surrounding event D_0,

$$\left(\mathbb{E}[1_{D_0} e]\right)^2 \le \Pr[D_0]\,\mathbb{E}[c].$$

Modal openings are used only in the analysis.

The second inequality retains the invalid-candidate mass in this bound. Write $\kappa_0 = \delta_{\mathrm{JL}} + q^{-m} + q^{-d/2} + 2/N$. Invalid anchors with valid norm but incorrect Gram matrix have mass b_G; the remainder have mass b_N. We will show

$$b_*^2 - \kappa_0 b_* \le \eta_0 + K\sqrt{b_*}, \qquad b_* = b_N + b_G, \qquad K^2 = (\sqrt{\eta_h} + \sqrt{r\eta_y})^2 + \eta_g.$$

Each candidate is fixed before the fresh tests. An oversized candidate passes the true projection test with probability at

most δ_{JL}. Any accepted, different claimed projection leaves a nonzero scalar residual, missed with probability at most q^{-m}. For a valid-norm candidate with correct Gram matrix but an invalid source equation, scalar or ring aggregation misses the residual with probability at most q^{-m} or $q^{-d/2}$. The ring bound uses a CRT component where the residual is nonzero.

Outside these events, the aggregated equation forces some modal h_{ii} to differ from $\langle \phi_i, S_i \rangle$. Fix the modal second opening at its prefix and the first such coordinate. A modal complement $z - c_i S_i$ then makes the linear consistency check a nonzero polynomial of degree at most two in c_i. An incorrect Gram entry similarly gives a degree-two polynomial in the selected challenges after fixing its modal complement. Unit challenge differences make the challenge set injective in each CRT field. Each polynomial passes with probability at most $2/N$.

The first inequality charges second-opening inconsistencies by $\sqrt{b_N \eta_h}$ and weak-opening inconsistencies by $\sqrt{r b_N \eta_y}$. Uniform coordinate selection accounts for r. The Gram pair is selected from the anchor, giving $\sqrt{b_G \eta_g}$. First-opening mismatches contribute η_0. Hence

$$\mathbb{E}[1_{\mathrm{bad}}(\varepsilon_\theta - \kappa_0)_+] \le \eta_0 + \sqrt{b_N}(\sqrt{\eta_h} + \sqrt{r\eta_y}) + \sqrt{b_G \eta_g} \le \eta_0 + K\sqrt{b_*}.$$

Here θ includes setup and initial state. The gate gives $b_\theta = \Pr[\mathrm{bad} \mid \theta] \le \varepsilon_\theta$. The left side is therefore at least $\mathbb{E}[b_\theta^2] - \kappa_0 b_* \ge b_*^2 - \kappa_0 b_*$, proving the second inequality.

If $b_* > \kappa_0 + \sqrt{\eta_0} + K^{2/3}$, then $\eta_0 + K\sqrt{b_*} < (\sqrt{\eta_0} + K^{2/3}) b_*$ contradicts that inequality. Thus $b_* \le \kappa_0 + \sqrt{\eta_0} + K^{2/3}$. Since $K^2 \le 2\eta_h + 2r\eta_y + \eta_g$, subtracting this mass and η_E from the candidate-generation bound in Lemma C.1 proves the claimed error. That lemma also gives the invocation count. Polynomial local work and expected-polynomial sampling give the runtime bound. □

A collision solver with expected time $P(\lambda)$ and success at least $1/p(\lambda)$ on infinitely many parameters can be truncated at $2P(\lambda)p(\lambda)$. Markov's inequality leaves success at least $1/(2p(\lambda))$, contradicting polynomial-time Module-SIS hardness. A concrete cutoff $T_{\max}$ instead incurs the explicit loss $P/T_{\max}$.

C.4. Proof of Theorem 5.2

Lemma C.2 (Fresh final replies)**.** *Lemma C.1 and the single-layer analysis hold when a kernel $K(\mathrm{prefix}; \xi)$ generates the final witness. Each invocation must restore the original prover state and use fresh independent coins, with no mutable extraction cache. Expected cost must be bounded uniformly over permitted prefixes.*

Proof. Condition on all verifier coins except coordinate i. Write a_c for kernel acceptance at challenge c, $p = N^{-1}\sum_c a_c$, and $\alpha = 1/N$. A gate accepting at c_0 has mass αa_{c_0}. A new trial accepts distinctly with probability $s_{c_0} = p - \alpha a_{c_0}$.

When $p \ge \alpha$, we have $s_{c_0} \ge p - \alpha$. Expected trials and abort probability decrease with s_{c_0}. Summing the gate masses gives the bounds $1/(1-\alpha)$ and α from Lemma C.1. When $p < \alpha$, stopping probability is at least α. Gate-weighted work is at most $p/\alpha < 1$, and abort mass is at most $p < \alpha$.

Opening masses include the kernel coins. Root counting remains valid because kernel acceptance at any fixed root is at most one. Immediately before each call, its prefix is fixed and its coins are fresh. The uniform conditional cost bound and the tower property therefore charge every reached call. □

Proof of Theorem 5.2. Starting from the checked terminal witness, replace each layer's final reply by a fresh call to the next extractor. Return its witness only after checking the entire target relation. This requires $B_i \ge \beta_{\mathrm{ext},i+1}$ in both the target check and collision bounds; the last layer uses the actual terminal cap.

If the suffix extractor has total error Δ_{i+1}, apply its guarantee to the ordinary distribution of layer-i prefixes, with a rejecting continuation after early failure. The virtual prover accepts with probability at least $\varepsilon_i - \Delta_{i+1}$. Lemma C.2 then permits the single-layer theorem, giving success at least $\varepsilon_i - \Delta_{i+1} - \delta_i - \nu_i$. Induction adds the errors once per layer. Repeated inner calls enter the runtime and computational experiments in ν_i.

Conditioning before each fresh call gives the theorem's cost recurrence. With $c_i = 1 + r_i/(1 - 1/N_i)$, its expansion is

$$C_1 \le \sum_{i=1}^{D} \tau_i \prod_{j=1}^{i} c_j + C_{D+1} \prod_{j=1}^{D} c_j.$$

The hypotheses bound this sum polynomially. Computational losses use Module-SIS hardness at the complete reduction runtimes, uniformly over the layer family. □

C.5. Source Composition and Bounded Restarts

Source composition. For the binary source, distinct original openings from two accepted suffixes yield a collision. Bitness bounds their difference in $\{-1, 0, 1\}^{D_{\mathrm{src}}}$, where $D_{\mathrm{src}} = n_{\mathrm{src}} + 3k$ counts the committed coefficients without conjugate copies. The squared collision norm is at most D_{src}. The opening-mass inequality charges nonmodal mass by $\sqrt{\eta_{\mathrm{src}}}$. On modal branches, an invalid opening has a fixed nonzero residual before the source rows. Each independent binary row misses it with probability $1/2$, giving $2^{-\ell}$.

In the arithmetic source, a nonzero residual over S has an additive image of size at least p, the smallest prime divisor of $2^d + 1$. Each uniform linear combination hits a fixed target with probability at most $1/p$. Fix the first opening before ϕ and the auxiliary opening before the second-stage rows. A nonzero linear residual $U = (Aw - a, Bw - b, Cw - c, Hw)$ survives every auxiliary choice, giving error at most $p^{-\ell}$. If $U = 0$ but the source is invalid, the first-stage fingerprints of $a \circ b - c$ all vanish with probability at most $p^{-\ell}$. Otherwise,

TABLE D.1. SELECTED SEVEN-LAYER PATH. THE TERMINAL OPENING HAS NO ADDITIONAL EXTRACTION SLACK.

Layer	n_i	r_i	β_i	b_i	$t_{1,i}$	$t_{2,i}$	β_{i+1}	B_i
1	37,450	112	16,384	16	8	6	35,098	81,056
2	5,795	33	35,098	32	7	5	32,627	75,349
3	1,754	15	32,627	64	6	5	21,402	49,426
4	877	8	21,402	64	5	5	20,834	48,115
5	440	7	20,834	64	6	5	12,763	29,475
6	440	5	12,763	64	6	5	10,040	23,187
7	443	4	10,040	4	15	8	77,874	77,874

correct auxiliaries force rejection, while incorrect auxiliaries leave a nonzero second-stage residual. This adds at most $p^{-\ell}$.

The two arithmetic commitments contribute $\sqrt{\eta_{\mathrm{src},0}} + \sqrt{\eta_{\mathrm{src},1}}$ at radius $2\beta_{\mathrm{actual}}$, where β_{actual} bounds the returned source witness. Each source experiment includes fresh core extraction and its full cost. Applying Theorem 5.2, with decoding at the returned norm, gives the source bounds in Section 5.3. □

Proof of Corollary 5.3. The nested initial calls reach the terminal check before any extraction fork. Real rejection therefore leaves the original post-rejection state intact. Real acceptance ends the protocol, even if recovery later fails.

For attempt j, an initializer simulates the preceding $j-1$ ordinary attempts. It supplies their rejection state, or a rejecting continuation if an earlier attempt accepted. Its acceptance probability is $\Pr[A_j]$, where A_j denotes first real acceptance at j.

Couple these experiments and the restart extractor through the same setup, prover tape, and rejection history. On A_j, use the same fresh extraction coins. Outside A_j, the gate prevents the j-th extractor from returning a witness. These output events are disjoint, so their single-attempt guarantees sum to

$$\Pr[E_{\mathrm{restart}} \text{ returns a valid witness}] \geq \varepsilon_{\mathrm{restart}} - \sum_{j=1}^{Q} \delta_j - \sum_{j=1}^{Q} \nu_j.$$

A common statistical bound gives $Q\delta_{\mathrm{tot}}$. Each ν_j includes the preceding rejection history and nested extraction. Conditional expected cost before each reached attempt gives $\mathbb{E}[\mathrm{time}] \leq \sum_j \Pr[\mathrm{reach}_j] C_j \leq Q \max_j C_j$.

If every reached honest attempt accepts with probability at least a, then $\Pr[R_j] \leq (1-a)\Pr[R_{j-1}]$. Thus final rejection is at most $(1-a)^Q$, and expected attempts are at most $[1-(1-a)^Q]/a$. The seven-layer path has $a > (479/528)^7 > 1/2$, so $Q = 128$ gives rejection below 2^{-128} and fewer than two expected attempts. □

Appendix D. Exact Certificates and Selected Parameters

D.1. Challenge and Moment Certificates

The recurrence in Appendix B.4 gives N_B exactly. With $v_t = 3048/2137$, $x = (933/1000)900/(64v_t)$, and $L_{80}(x) = \sum_{j=0}^{80} x^j/j!$, the rational certificate is

$$\frac{(2137/625)^{64}(5/4)^{104}}{N_B} \frac{384}{L_{80}(x)} < \frac{1}{3}.$$

Since $L_{80}(x) < e^x$, the retained count is at least $N_{\min} = \lfloor 2N_B/3 \rfloor + 1 > 2^{143}$.

The same recurrence gives the moments in Appendix B.5. Writing $B_j(E) = \sum_{e \leq E} a_j(e)$, their numerators are

$$N_B v_0 = \sum_x x^2 B_{63}(104 - x^2),$$
$$N_B m_4 = \sum_x x^4 B_{63}(104 - x^2),$$
$$N_B m_{22} = \sum_{x,y} x^2 y^2 B_{62}(104 - x^2 - y^2).$$

The sums range over $\{-2,-1,0,1,2\}$, with $B_j(E) = 0$ for $E < 0$. Exact rational comparisons certify $(9 - v_0)^2 > 24g$, $16v_0 < 25$, and $V_U < 32$. The artifact retains their integer forms and outputs.

D.2. Selected Parameter Path

Table D.1 instantiates the complete ledger of Section 6.4. Both auxiliary bases equal b_i. Challenges are uniform in $\mathcal{C}_{64,104,30}$. The projection threshold is $640\beta_i^2$, and $\Gamma_i = 24\beta_i$. Each nonterminal cap is $B_i = \lceil \sqrt{16/3}\, \beta_{i+1} \rceil$; the terminal uses $B_7 = \beta_8$. Upward rounding in Proposition 4.1 gives the listed outgoing norms. Every layer satisfies $250\beta_{\mathrm{ext},i} \leq q$, and all collision radii are below q.

D.3. Aggregation and Statistical Accounting

The source fingerprints its $3k$ residual entries over $\mathbb{F}_{2^{187}}$. Each ordinary layer fingerprints M_i scalar equations over $\mathbb{F}_{q^6}$ and the ring equations over the two CRT fields. Every residual is fixed before its point is sampled. Including all

padding equations gives $\sum_i (M_i - 1) = 134225082$ and $\sum_i (K_i + 5) = 9135$. With 128 complete attempts, the statistical bound is

$$128\Big[\frac{198}{N_{\min}} + 7 \cdot 2^{-164} + \frac{134225082}{q^6} + \frac{9135}{q^{32}} + \frac{100663295}{2^{187}}\Big] < 2^{-128}.$$

Here $198 = \sum_i (r_i + 2)$ and $100663295 = 3k - 1$. The exact comparison gives exponent approximately -128.69244; computational losses are accounted for separately in Section 5.3.

D.4. Communication Ledger

One complete attempt sends 490,496 bits in early core messages, 780,480 in the terminal opening, 12,288 in the source commitment, and $187 \cdot 27$ in source replies. The total is 1,288,313 bits, or 161,040 bytes after final padding. Explicit verifier messages total 580,126,427,131 bits, approximately 67.54 GiB. These use the communication scope specified in Section 6.4. The artifact supplies coefficient widths, padding counts, field certificates, and the component arithmetic.