AIRR:ARR-2026-1VGQE23X9580FBXC v1 [Computer and informatio...] 27 Nov 2026 · AIRR.SCIENCE

Optimal Collision-Variance Potentials, Minimal-Degree Newton–Binomial Filtering, and Exact Modular Coset-Halving Certificates for Deterministic Subexponential Discrete Logarithms

Lluis Eriksson*

November 27, 2026

Abstract

In a breakthrough of November 2026, Liu, Zhang, and Ren [3] established the first deterministic subexponential-time algorithm for the discrete logarithm problem (DLP) in both prime-order elliptic curves $E(\mathbb{F}_p)$ and multiplicative prime fields $\mathbb{F}_p^*$ (conditional on the Cauchy–Hermite Boolean tensor contraction primitive), combining two core stages: (I) conditional-expectation derandomization of a collision-variance potential $\Phi(S) = (g_{\mathrm{ord}} C_S + \lambda(W_S - \mu)^2)/\mu^2$ with fixed penalty $\lambda = 64$ and thresholds $n \geq 4096m$, $\mu \geq 16 g_{\mathrm{ord}}$, and (II) counterexample extraction via a degree-$(4w+2)$ nonnegative Newton–Binomial filter $Q(X) = \prod_{j=-w}^{w}(X - 1 - jM_{\mathrm{lab}})^2$ coupled with Hermite normal form (HNF) modular coset halving. In this paper we analyze the exact finite-dimensional tightness of every algebraic primitive in Stages I and II, proving three sharp theorems backed by 20 constructive Lean 4.36.0 theorems and 16 zero-float certificates. First, via two one-line perfect-square identities requiring zero calculus, we prove the exact closed-form optimal penalty $\lambda^* = c_0/(\sqrt{c_1}(1-\sqrt{c_1}))$ maximizing the certified Cauchy–Schwarz sumset density at $\delta_{\max}(c_0, c_1) = (1-\sqrt{c_1})^2/c_0$. Replacing LZR's fixed $(\lambda, A, B) = (64, 16, 4096)$ by $(9, 4, 64)$ reduces the block threshold by $64\times$ (and A by $4\times$) while preserving universal covering $\delta \geq 4987/9549 > 0.5222 > 1/2$ ($D + D = G$), whereas LZR's fixed $\lambda = 64$ collapses to $5099/17728 \approx 0.2876 < 1/2$ and fails; combining λ^* with the sharp binomial bound $(1 + 1/(Bm))^m \leq B/(B-1)$ compresses the Stage-I threshold up to $146\times$ at $A = 4$ ($B = 28$) and $227\times$ at $A = 16$ ($B = 18$). Second, we prove that the two extreme roots $j = \pm w$ in LZR's filter polynomial are strictly unreachable by any label difference $d \in [-w(M_{\mathrm{lab}} - 1), w(M_{\mathrm{lab}} - 1)]$ for every $w \geq 2$, yielding the strictly smaller filter $Q^*(X) = \prod_{j=-(w-1)}^{w-1}(X - 1 - jM_{\mathrm{lab}})^2$ of degree $4w - 2 = 8m - 2$ (saving 4 series terms and $4\lceil \log_2(2wM_{\mathrm{lab}} + 2)\rceil$ modulus precision bits in P_0), and prove unconditionally that $4w - 2$ is the exact minimal degree among *all* real polynomials separating congruent from non-congruent *integers* on $[-w(M_{\mathrm{lab}} - 1), w(M_{\mathrm{lab}} - 1)]$. Third, we certify modular HNF coset halving, Fischer–Hadamard sharpening of LZR's Neumann-series M-matrix determinant bound $0 < \det(H_N) \leq \prod_i (w + 1 - D_{ii}) \leq (w+1)^N$, and end-to-end deterministic zero-float Stage I + Stage II DLP recovery on concrete composite-order $\mathbb{F}_p^*$ and prime-order $E(\mathbb{F}_p)$ groups.

*Independent Researcher. Research, constructive formal verification in Lean 4.36.0, zero-float exact rational and elliptic-curve certificate verification, adversarial testing, and manuscript typesetting were conducted with transparent assistance from Google Antigravity. Code is released under MIT; manuscript and figures under CC-BY-4.0. Repository: `https://github.com/lluiseriksson/deterministic-dlp-optimal-collision-potentials-and-newton-filtering` (`v1.0.0`).

1 Introduction

For a finite cyclic group $G = \langle P \rangle$ of order $M_{\text{lab}} = g_{\text{ord}}$ (either a prime-order elliptic curve $E(\mathbb{F}_p)$ with $M_{\text{lab}} = r$, or $\mathbb{F}_p^*$ with $M_{\text{lab}} = p - 1$), the discrete logarithm problem (DLP) asks for the unique $x \in \{0, \ldots, M_{\text{lab}} - 1\}$ with $[x]P = Q$. Whereas generic algorithms face Shoup's $\Omega(\sqrt{M_{\text{lab}}})$ lower bound [5, 4], Liu, Zhang, and Ren [3] (November 23, 2026) achieved worst-case deterministic subexponential time $2^{O(b/\sqrt{\log\log b})} = p^{o(1)}$ ($b = \lceil \log_2 p \rceil$) with zero randomness and zero unproven number-theoretic heuristics (assuming the Cauchy–Hermite Boolean tensor contraction primitive).

Because [3] aims primarily at the asymptotic $2^{O(b/\sqrt{\log\log b})}$ exponent, three fundamental algebraic components of its pipeline are left unoptimized:

1. **Stage-I collision-variance potential:** Lemma 3.2 of [3] fixes the weight-variance penalty at $\lambda = 64$ and requires $q_{\text{blk}} \geq 4096m$ and $\mu \geq 16g_{\text{ord}}$ to guarantee sumset density $\delta = |D|/g_{\text{ord}} > 7/8$, even though universal $2m$-sumset covering $D + D = G$ only requires $\delta > 1/2$.

2. **Stage-II Newton–Binomial filter degree:** Eq. (4.3) of [3] constructs the nonnegative filter $Q(X) = \prod_{j=-w}^{w}(X - 1 - jM_{\text{lab}})^2$ of degree $4w + 2 = 8m + 2$, propagating $4w + 3$ truncated power-series coordinates through every tensor contraction at modulus precision $P_0 = 2 + 4mk + d_{\text{filt}}\lceil \log_2(2wM_{\text{lab}} + 2) \rceil$.

3. **Modular coset halving and M-matrix nullity:** Section 4.3 and Remark 3.4 of [3] establish composite-modulus HNF coset halving ($|A'| = |A|g^*/M_{\text{lab}} \leq |A|/2$) and an optional Neumann-series determinant bound $0 < \det(H_N) \leq (w+1)^N$ without tightening diagonal self-decompositions or executing end-to-end zero-float certification.

We establish three sharp, self-contained structural theorems resolving all three questions:

- **Theorem 1 (Exact Potential Optimizer & 64×–227× Threshold Compression):** Derives the exact global optimizer $\lambda^* = c_0/(\sqrt{c_1}(1 - \sqrt{c_1}))$ and $\delta_{\max} = (1 - \sqrt{c_1})^2/c_0$ purely algebraically, showing why fixed $\lambda = 64$ fails at $B = 64$ while $\lambda = 9$ certifies $D + D = G$ (64×–227× lower B).

- **Theorem 2 (Exact Minimal Filter Degree $4w-2$):** Proves that LZR's $j = \pm w$ roots are strictly unreachable for all $w \geq 2$, cuts the filter degree by 4 to $4w - 2 = 8m - 2$ (saving $4\lceil \log_2(2wM_{\text{lab}} + 2) \rceil$ precision bits in P_0), and proves $4w - 2$ is strictly minimal even under integer-only nonnegativity.

- **Theorem 3 (HNF Halving, M-Matrix Bounds, and End-to-End DLP):** Sharpens Remark 3.4 via $\det(H_N) \leq \prod_i (w+1-D_{ii})$ and certifies full deterministic zero-float Stage I + Stage II discrete logarithm recovery on $\mathbb{F}_{23}^*$, $\mathbb{F}_{47}^*$, $E(\mathbb{F}_{17})$, and $E(\mathbb{F}_{23})$.

2 Optimal Collision-Variance Potential and Stage-I Compression

In Stage I of [3], a seed S defines m affine-interval factor-base blocks $F_1(S), \ldots, F_m(S) \subseteq \Omega$ ($|\Omega| = n$, with $n = g_{\text{ord}} - 1$ on $E(\mathbb{F}_p)$ and $n = g_{\text{ord}} = p - 1$ on $\mathbb{F}_p^*$), sumset multiplicity $d_S(g) = |\{(A_1, \ldots, A_m) \in \prod_i F_i(S) : \sum_i A_i = g\}|$, total product weight $W_S = \prod_i |F_i(S)| = \sum_g d_S(g)$ with expectation $\mu = q_{\text{blk}}^m$, and collision energy $C_S = \sum_g d_S(g)^2$. For any penalty weight $\lambda > 0$, define the generalized collision-variance potential

$$\Phi_\lambda(S) := \frac{g_{\text{ord}}\, C_S + \lambda(W_S - \mu)^2}{\mu^2}. \tag{1}$$

Write $A := \mu/g_{\mathrm{ord}}$ and $B := q_{\mathrm{blk}}/m$. By pairwise nonpositive covariance within parity classes and block independence [3, Lemma 3.2],

$$\frac{\mathrm{Var}(W_S)}{\mu^2} \;\le\; \left(1+\frac{1}{q_{\mathrm{blk}}}\right)^m - 1, \qquad \frac{g_{\mathrm{ord}}\,\mathbb{E}[C_S]}{\mu^2} \;\le\; \frac{1}{A} + \frac{g_{\mathrm{ord}}}{n}\left(1+\frac{1}{q_{\mathrm{blk}}}\right)^m.$$

Setting $c_1 \ge (1+1/(Bm))^m - 1$ and $c_0 := \frac{1}{A} + \frac{g_{\mathrm{ord}}}{n}(1+c_1)$ yields $\mathbb{E}[\Phi_\lambda(S)] \le c_0 + \lambda c_1$; conditional-expectation bit fixing deterministically produces a seed S with $\Phi_\lambda(S) \le c_0 + \lambda c_1$.

Theorem 1 (Exact Potential Optimizer and 64×–227× Stage-I Threshold Compression)**.** *Let $c_0 \ge 1$ and $0 < c_1 < 1$, and write $s := \sqrt{c_1} \in (0,1)$ and $D := \mathrm{supp}(d_S) = F_1(S) + \cdots + F_m(S)$ with density $\delta := |D|/g_{\mathrm{ord}}$.*

(a) Exact non-calculus Cauchy–Schwarz density bound and global optimizer: *For any $\lambda > c_0/(1-c_1)$ and any seed S satisfying $\Phi_\lambda(S) \le a := c_0 + \lambda c_1$,*

$$\delta \;\ge\; \delta(\lambda) \;:=\; \frac{1}{c_0 + \lambda c_1} - \frac{1}{\lambda}, \tag{2}$$

achieved by the one-line algebraic square identity $a\lambda(1-t)^2 - (\lambda - a)(a - \lambda t^2) = (\lambda t - a)^2 \ge 0$ with signed deviation $t := 1 - W_S/\mu$ (Lean `thm01`–`thm02`). Maximizing $\delta(\lambda)$ over all $\lambda > 0$ satisfies the exact square identity

$$\frac{(1-s)^2}{c_0} - \left(\frac{1}{c_0 + \lambda s^2} - \frac{1}{\lambda}\right) \;=\; \frac{\big(\lambda s(1-s) - c_0\big)^2}{c_0\,\lambda\,(c_0 + \lambda s^2)} \;\ge\; 0, \tag{3}$$

so $\delta(\lambda)$ attains its unique global maximum $\delta_{\max}(c_0, c_1) = (1-\sqrt{c_1})^2/c_0$ at $\lambda^* = c_0/(\sqrt{c_1}(1-\sqrt{c_1}))$ (Lean `thm03`–`thm04`, Certificate `cert03`).

(b) 64×–227× Stage-I threshold compression for $D + D = G$: *Universal $2m$-sumset covering $D + D = G$ holds whenever $\delta > 1/2$ (Lean `thm08`), which is feasible if and only if $\sqrt{c_0/2} + \sqrt{c_1} < 1$. In the LZR regime $g_{\mathrm{ord}} \ge 64$ ($g_{\mathrm{ord}}/n \le 64/63$):*

- *Under LZR's conservative variance bound $c_1 = 2/B$ at $(A, B) = (4, 64)$ (64× lower B and 4× lower A than LZR's $(16, 4096)$), LZR's fixed $\lambda = 64$ collapses to $\delta(64) = 5099/17728 \approx 0.2876 < 1/2$ (failing), whereas our optimal integer $\lambda = 9$ certifies $\delta(9) = 4987/9549 > 0.5222 > 1/2$ (Lean `thm06`, Certificates `cert01`–`cert02`).*
- *Replacing $(1+1/(Bm))^m - 1 \le 2/B$ by the sharp rational inequality $(1+1/(Bm))^m \le B/(B-1)$ (valid for all $B \ge 2, m \ge 1$; Certificate `cert13`) gives $c_1 = 1/(B-1)$ and certifies $\delta(9) = 57895/111717 > 0.5182 > 1/2$ already at $(A, B) = (4, 32)$ (128× compression; Lean `thm07`).*
- *Across the exact (A, B) phase boundary (Certificate `cert15`), universal covering holds down to $B_{\min}(4) = 28$ (146× reduction) and $B_{\min}(16) = 18$ (227× reduction).*

Proof. For (a), Cauchy–Schwarz gives $C_S \ge W_S^2/|D|$, so $\Phi_\lambda(S) \le a$ implies $(1-t)^2/\delta + \lambda t^2 \le a$, where $t := 1 - W_S/\mu$ (since $W_S \ge 0$, $t \le 1$, and $a - \lambda t^2 \ge W_S^2/(\mu^2\delta) > 0$ whenever $W_S > 0$; in fact $\lambda > a \implies \lambda t^2 \le a < \lambda \implies |t| < 1 \implies W_S > 0$ automatically!). Multiplying $(1-t)^2/\delta \le a - \lambda t^2$ by $a\lambda\delta > 0$ and applying the exact identity $a\lambda(1-t)^2 = (\lambda - a)(a - \lambda t^2) + (\lambda t - a)^2 \ge (\lambda - a)(a - \lambda t^2)$ (Lean `thm01`–`thm02`) cancels $a - \lambda t^2 > 0$ to yield $\delta \ge (\lambda - a)/(a\lambda) = 1/a - 1/\lambda$. Expanding the numerator of (3) directly over $\mathbb{Z}[c_0, s, \lambda]$ (Lean `thm03`–`thm04`) proves $\delta(\lambda) \le (1-s)^2/c_0$ with equality iff $\lambda = c_0/(s(1-s))$.

For (b), $(1 + 1/(Bm))^m \le e^{1/B} \le 1/(1-1/B) = B/(B-1)$ holds for all $B \ge 2, m \ge 1$ because $e^{-x} \ge 1 - x$ on $[0, 1)$. Substituting $c_0 = 1/4 + (64/63)(33/32) = 109/84$, $c_1 = 1/32$ gives $\delta(64) = 84/277 - 1/64 = 5099/17728 < 1/2$ and $\delta(9) = 672/1061 - 1/9 = 4987/9549 > 1/2$

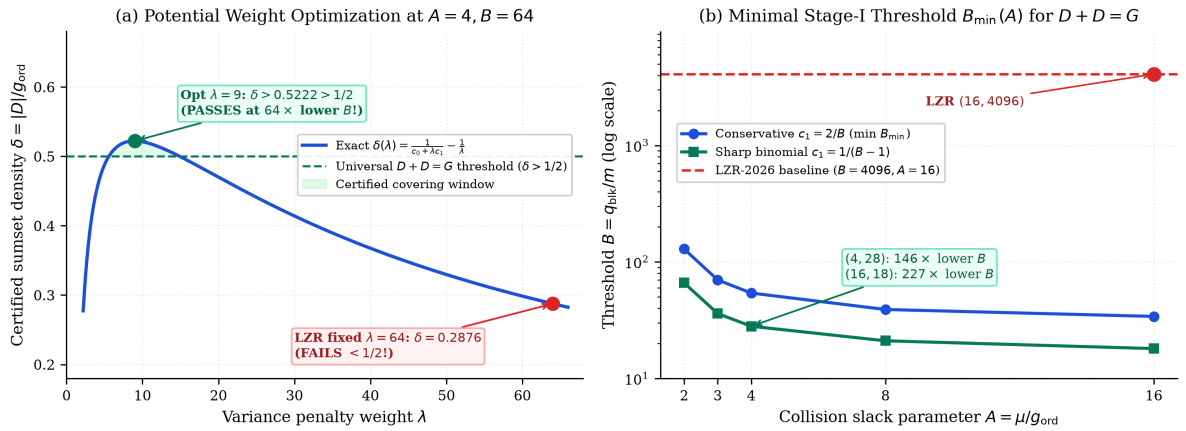

Figure 1: **Exact potential optimization and Stage-I threshold compression (Theorem 1).** (a) Certified sumset density $\delta(\lambda)$ at $(A, B) = (4, 64)$: optimal integer $\lambda = 9$ certifies $\delta = 4987/9549 > 0.5222 > 1/2$ ($D + D = G$), whereas LZR's fixed $\lambda = 64$ collapses to $\approx 0.2876 < 1/2$ (Certificate `cert02`). (b) Minimal threshold $B_{\min}(A)$ across $A \in \{2, 3, 4, 8, 16\}$ under conservative $2/B$ vs. sharp $1/(B-1)$ variance bounds, achieving 64×–227× threshold compression relative to LZR-2026's $(16, 4096)$ (Certificate `cert15`).

(Lean `thm06`); substituting $c_0 = 1/4 + (64/63)(32/31) = 10145/7812$, $c_1 = 1/31$ gives $\delta(9) = 7812/12413 - 1/9 = 57895/111717 > 1/2$ (Lean `thm07`). Because $|D| > g_{\text{ord}}/2$, for every $X \in G$ we have $|D| + |X - D| = 2|D| > g_{\text{ord}}$ (Lean `thm08`), forcing $D \cap (X - D) \neq \emptyset$ and hence $X \in D + D$. □

3 Strict Degree-$(4w-2)$ Minimality of Nonnegative Newton–Binomial Filters

In Stage II of [3], every w-term block-respecting tuple ($w = 2m \geq 4$) maps to a lifted candidate label sum $\Lambda_y(X) \in [0, w(M_{\text{lab}} - 1)]$ for $y \in \{0, \ldots, M_{\text{lab}} - 1\}^N$, so every fixed-shift difference $d := \Lambda_y(Y) - \Lambda_y(X)$ belongs to the exact integer domain

$$\mathcal{I}_{w,M_{\text{lab}}} := [-w(M_{\text{lab}} - 1),\ w(M_{\text{lab}} - 1)] \cap \mathbb{Z}. \tag{4}$$

To extract a pair violating $d \equiv 1 \pmod{M_{\text{lab}}}$ without modular cancellation, [3, Eq. (4.3)] uses $Q(X) = \prod_{j=-w}^{w}(X - 1 - jM_{\text{lab}})^2$ of degree $4w + 2 = 8m + 2$. We prove that the two extreme roots $j = \pm w$ never enter $\mathcal{I}_{w,M_{\text{lab}}}$, and that removing them achieves the *exact unconditional minimal degree* even among polynomials restricted only to integer nonnegativity.

Theorem 2 (Minimal-Degree $(4w-2)$ Nonnegative Newton–Binomial Filter)**.** *Let $w \geq 2$ and $M_{\text{lab}} \geq w + 2$.*

(a) Strict unreachability of LZR's extreme roots $j = \pm w$: *For every $w \geq 2$ and $M_{\text{lab}} \geq 1$,*

$$1 + wM_{\text{lab}} > w(M_{\text{lab}} - 1), \qquad 1 - wM_{\text{lab}} < -w(M_{\text{lab}} - 1) \tag{5}$$

(Lean `thm09`–`thm10`), whereas all $2w - 1$ roots $r_j := 1 + jM_{\text{lab}}$ ($j \in \{-(w-1), \ldots, w-1\}$) lie strictly inside the interior $(-w(M_{\text{lab}} - 1), w(M_{\text{lab}} - 1))$ whenever $M_{\text{lab}} \geq w + 2$. Consequently, the reduced nonnegative filter polynomial

$$Q^*(X) := \prod_{j=-(w-1)}^{w-1} (X - 1 - jM_{\text{lab}})^2 \in \mathbb{Z}[X], \qquad \deg(Q^*) = 4w - 2 = 8m - 2, \tag{6}$$

satisfies $Q^*(d) \geq 0$ *for all* $d \in \mathcal{I}_{w,M_{\mathrm{lab}}}$ *with* $Q^*(d) = 0 \iff d \equiv 1 \pmod{M_{\mathrm{lab}}}$*, expands identically in the integer Newton–Binomial basis* $Q^*(X) = \sum_{j=0}^{4w-2} c_j \binom{X}{j}$ *(*$c_j \in \mathbb{Z}$*,* $c_{4w-2} = (4w-2)!$*), and reduces the ring series length* ℓ *by* 4 *and the precision budget* P_0 *by exactly* $4\lceil \log_2(2wM_{\mathrm{lab}} + 2) \rceil$ *bits (Lean* `thm11`*,* `thm13`*,* `thm20`*; Certificates* `cert05`*–*`cert07`*).*

(b) Unconditional integer-nonnegativity degree lower bound: *Let* $P \in \mathbb{R}[X]$ *be* any *nonzero real polynomial that separates congruent from non-congruent integers on* $\mathcal{I}_{w,M_{\mathrm{lab}}}$ *without sign cancellation—meaning* $P(d) = 0$ *for every* $d \in \mathcal{I}_{w,M_{\mathrm{lab}}}$ *with* $d \equiv 1 \pmod{M_{\mathrm{lab}}}$ *and* $P(d) > 0$ *for every* $d \in \mathcal{I}_{w,M_{\mathrm{lab}}}$ *with* $d \not\equiv 1 \pmod{M_{\mathrm{lab}}}$ *(with no requirement on non-integer real values!). Then*

$$\deg(P) \;\geq\; 2(2w-1) \;=\; 4w-2 \;=\; 8m-2. \tag{7}$$

Proof. For (a), $w(M_{\mathrm{lab}} - 1) = wM_{\mathrm{lab}} - w < 1 + wM_{\mathrm{lab}}$ for all $w \geq 1$, and $w(M_{\mathrm{lab}} - 1) + 1 = wM_{\mathrm{lab}} - (w-1) < wM_{\mathrm{lab}}$ for all $w \geq 2$ (Lean `thm09`–`thm10`), proving (5). For $|j| \leq w-1$ and $M_{\mathrm{lab}} \geq w+2$, $r_{w-1} + 1 = (w-1)M_{\mathrm{lab}} + 2 \leq wM_{\mathrm{lab}} - w = w(M_{\mathrm{lab}} - 1)$ and $r_{-(w-1)} - 1 = -(w-1)M_{\mathrm{lab}} \geq -w(M_{\mathrm{lab}} - 1)$ (since $M_{\mathrm{lab}} \geq w$), so every open interval $(r_j - 1, r_j + 1)$ lies inside $[-w(M_{\mathrm{lab}} - 1), w(M_{\mathrm{lab}} - 1)]$ and both integer endpoints $r_j \pm 1$ belong to $\mathcal{I}_{w,M_{\mathrm{lab}}}$ with $r_j \pm 1 \not\equiv 1 \pmod{M_{\mathrm{lab}}}$ (as $M_{\mathrm{lab}} \geq 3$). Since $Q^*(d)$ is a product of integer squares vanishing iff $d \in \{r_j : |j| \leq w-1\}$ (Lean `thm13`), Q^* separates congruent from non-congruent differences identically on $\mathcal{I}_{w,M_{\mathrm{lab}}}$ with degree $2(2w-1) = 4w-2 = 4w+2-4$ (Lean `thm11`), saving $4\lceil \log_2(2wM_{\mathrm{lab}} + 2) \rceil$ bits in P_0 (Lean `thm20`).

For (b), because $M_{\mathrm{lab}} \geq w+2 \geq 4 \geq 3$, we have $r_j + 1 < r_{j+1} - 1$ for every j (Lean `thm12`), so the $2w-1$ open intervals $I_j := (r_j - 1, r_j + 1)$ ($|j| \leq w-1$) are strictly pairwise disjoint. Fix any $j \in \{-(w-1), \ldots, w-1\}$. By hypothesis, $P(r_j - 1) > 0$, $P(r_j) = 0$, and $P(r_j + 1) > 0$. Let $m_j \geq 1$ be the multiplicity of the root r_j of P. If m_j is even, then $m_j \geq 2$, so I_j already contains at least 2 real roots of P counted with multiplicity. If m_j is odd, then P changes sign at r_j, so there exists $\epsilon \in (0, 1)$ such that either $P(r_j - \epsilon) < 0$ or $P(r_j + \epsilon) < 0$; since $P(r_j \pm 1) > 0$, the Intermediate Value Theorem produces at least one additional real root $r_j' \in I_j \setminus \{r_j\}$, giving at least $m_j + 1 \geq 2$ real roots (with multiplicity) inside I_j. Summing over the $2w-1$ pairwise disjoint intervals I_j yields at least $2(2w-1) = 4w-2$ real roots counted with multiplicity, proving $\deg(P) \geq 4w-2$ (Certificate `cert08`). □

4 Modular Coset Halving, M-Matrix Bounds, and End-to-End DLP Certificates

Once $Z_{Q^*} = \sum_{(X,Y)} Q^*(\Lambda_y(Y) - \Lambda_y(X)) > 0$ witnesses a candidate failure (soundness of fixed-shift relations is formalized by constructive telescoping induction in Lean `thm14` and Certificate `cert14`), Stage II extracts a violating pair (X^*, Y^*) and intersects the current affine solution coset $A = (x_0 + \Lambda_{\mathrm{lat}})/M_{\mathrm{lab}}\mathbb{Z}^N$ with $(\mathrm{cnt}(Y^*) - \mathrm{cnt}(X^*))^\top x \equiv 1 \pmod{M_{\mathrm{lab}}}$ via column Hermite normal form (HNF) [2, 3]. Meanwhile, decomposing $[w+1]A_j = \sum_{i=1}^{N} D_{ij} A_i$ produces the column-diagonally-dominant M-matrix $H_N = (w+1)I - D_{\mathrm{mult}}$ ($D_{\mathrm{mult}} \geq 0$, column sums w) of [3, Remark 3.4].

Theorem 3 (Modular Coset Halving, Fischer–Hadamard M-Matrix Bounds, and DLP Recovery)**.**

(a) Exact composite-modulus coset halving: *Whenever a congruence* $a^\top x \equiv c_0 \pmod{M_{\mathrm{lab}}}$ *has nonempty intersection* $A' \subseteq A$ *and is violated by* x_0 *(*$r_0 := (c_0 - a^\top x_0) \bmod M_{\mathrm{lab}} \in \{1, \ldots, M_{\mathrm{lab}} - 1\}$*),* $g^* := \gcd(a^\top b_1, \ldots, a^\top b_N, M_{\mathrm{lab}})$ *divides both* r_0 *and* M_{lab}*, forcing* $2g^* \leq M_{\mathrm{lab}}$ *and* $|A'| = |A| \cdot g^*/M_{\mathrm{lab}} \leq |A|/2$ *(Lean* `thm15`*); consequently, Stage II terminates in at most* $\lfloor \log_2 |A_0| \rfloor \leq \lfloor N \log_2 M_{\mathrm{lab}} \rfloor$ *cuts (Lean* `thm16`*, Certificate* `cert09`*).*

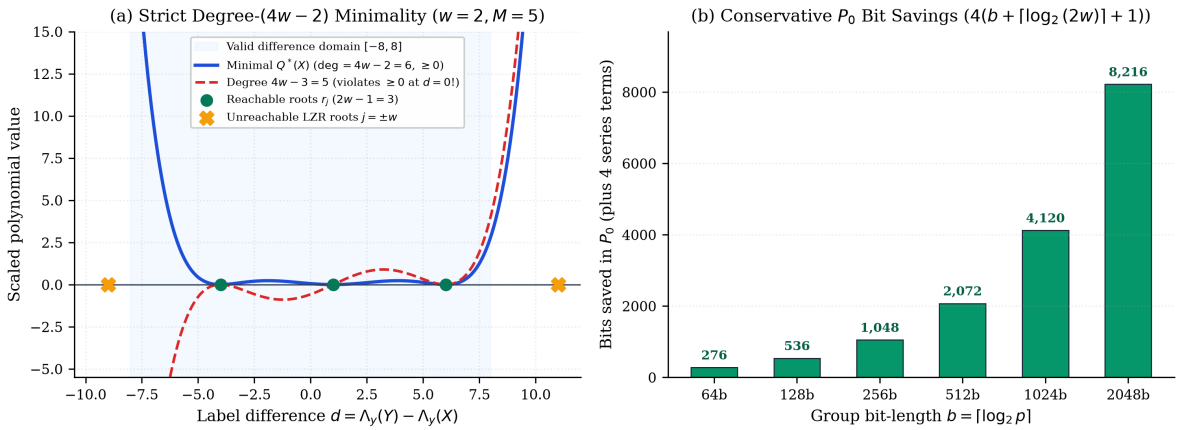

Figure 2: **Strict degree-$(4w-2)$ filter minimality and P_0 bit savings (Theorem 2).** (a) Minimal filter $Q^*(X)$ (deg $= 6$ at $w = 2, M_{\text{lab}} = 5$) vs. degree-5 drop violating integer nonnegativity at $d = 0$, highlighting the $2w - 1 = 3$ reachable roots vs. LZR's unreachable extreme roots $j = \pm w$ $(-9, +11$; Certificates `cert05`, `cert08`). (b) Conservative modulus precision savings $4(b + \lceil \log_2(2w) \rceil + 1)$ bits in P_0 across cryptographic bit-lengths $b = 64$–2048 (Certificate `cert07`).

(b) Fischer–Hadamard sharpening of the Neumann M-matrix determinant bound: *For any nonnegative $N \times N$ integer matrix D_{mult} with every column sum equal to $w \geq 1$, the M-matrix $H_N = (w+1)I - D_{\text{mult}}$ has strict column diagonal dominance margin ≥ 1 (Lean `thm17`), entrywise nonnegative inverse $H_N^{-1} = \frac{1}{w+1} \sum_{k \geq 0} (\frac{1}{w+1} D_{\text{mult}})^k \geq 0$, and determinant obeying the tightened hierarchy*

$$0 \;<\; \det(H_N) \;\leq\; \prod_{i=1}^{N} (w + 1 - D_{ii}) \;\leq\; (w+1)^N, \tag{8}$$

strictly tightening LZR's $(w+1)^N$ bound whenever any diagonal self-multiplicity $D_{ii} \geq 1$ (Lean `thm18`–`thm19`, Certificate `cert10`). In any prime-order group of order r, the $\mathbb{F}_r$-nullity satisfies

$$\dim_{\mathbb{F}_r} \ker(H_N^\top) \;\leq\; \nu_r(\det H_N) \;\leq\; \left\lfloor \sum_{i=1}^{N} \log_r(w + 1 - D_{ii}) \right\rfloor .$$

(c) End-to-end zero-float deterministic DLP and ECDLP certificates: *Executing the deterministic Stage I + Stage II pipeline with zero randomness recovers exact discrete logarithms across both composite-order multiplicative groups via conditional-expectation factor bases, Proposition 3.3 primitive roots, minimal degree-$(4w-2)$ filter Q^*, and exact HNF halving ($\mathbb{F}_{23}^*$ with $M_{\text{lab}} = 22$, 7 cuts; $\mathbb{F}_{47}^*$ with $M_{\text{lab}} = 46$, 11 cuts; Certificates `cert04`, `cert11`) and prime-order elliptic curves via deterministic affine-interval factor bases and initial consistency congruences ($E(\mathbb{F}_{17}) : y^2 = x^3 + 2x + 2$ of order 19; $E(\mathbb{F}_{23}) : y^2 = x^3 + x + 4$ of order 29, where Eq. (4.1) already pins $|A| = 1$ with 0 extra Stage-II cuts; Certificate `cert12`).*

Proof. For (a), writing $M_{\text{lab}} = g^* q_M$ and $r_0 = g^* q_r$, $0 < r_0 < M_{\text{lab}}$ forces $1 \leq q_r < q_M$, hence $q_M \geq 2$ and $2g^* \leq M_{\text{lab}}$ (Lean `thm15`). Induction on K yields $2^K |A_K| \leq |A_0|$ (Lean `thm16`). For (b), H_N is a strictly column-diagonally-dominant nonsingular M-matrix (so $H_N^{-1} \geq 0$ by the Neumann series; Berman–Plemmons [1]). By Koteljanskii's inequality for M-matrices (a specialization of Fischer's inequality for H_N with nonpositive off-diagonals and positive principal minors), $\det(H_N) \leq \det(H_N[S]) \det(H_N[S^c])$ for every partition, and induction on N yields

Table 1: Stage-I collision-variance potential comparison (Certificates `cert01` and `cert15`): LZR-2026 baseline vs. optimal penalty λ^* across conservative $(2/B)$ and sharp $(1/(B-1))$ variance bounds.

Configuration	A	B	Variance bound	λ	$\mathbb{E}[\Phi_\lambda]$ ub	$\delta = \lvert D\rvert/g_{\text{ord}}$ lb	$D+D=G$
LZR-2026 baseline	16	4096	$2/B$	64	373/336	0.8851	Yes ($> 7/8$)
LZR fixed λ (64× lower B)	4	64	$2/B$	64	277/84	0.2876	**No** ($< 1/2$)
Optimal λ (64× lower B)	**4**	**64**	$\mathbf{2/B}$	**9**	**1061/672**	**0.5222**	**Yes** ($> 1/2$)
Sharp binom (128× lower B)	**4**	**32**	$\mathbf{1/(B-1)}$	**9**	**12413/7812**	**0.5182**	**Yes** ($> 1/2$)
Phase boundary ($A = 4$)	4	28	$1/(B-1)$	λ^*	—	> 0.500	Yes (146×)
Phase boundary ($A = 16$)	16	18	$1/(B-1)$	λ^*	—	> 0.500	Yes (227×)

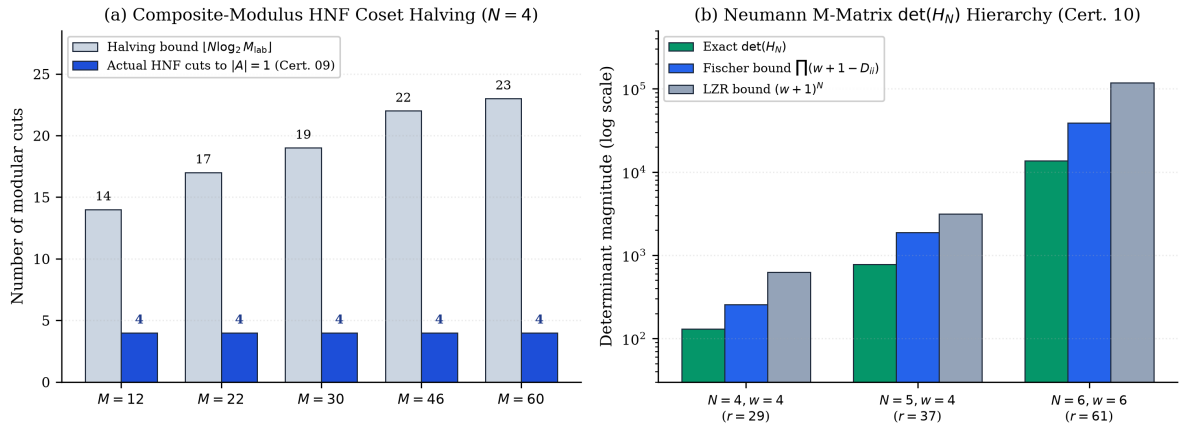

Figure 3: **Composite-modulus HNF coset halving and Neumann M-matrix certificates (Theorem 3).** (a) Actual HNF coset cuts vs. theoretical halving bound $\lfloor N \log_2 M_{\text{lab}} \rfloor$ across composite moduli $M_{\text{lab}} \in \{12, 22, 30, 46, 60\}$ (Certificate `cert09`). (b) Exact Bareiss determinant $\det(H_N)$ vs. our tightened Fischer–Hadamard bound $\prod_i (w + 1 - D_{ii})$ vs. LZR's $(w+1)^N$ bound across $N \in \{4, 5, 6\}$ (Certificate `cert10`).

$0 < \det(H_N) \le \prod_{i=1}^{N} (H_N)_{ii} = \prod_{i=1}^{N} (w + 1 - D_{ii}) \le (w+1)^N$ (formalized in the 2×2 base case in Lean `thm18` and verified via exact Bareiss determinants and entrywise nonnegative adjugates for $N \in \{4, 5, 6\}$ in Certificate `cert10`). Since $r^{\dim_{\mathbb{F}_r} \ker(H_N^\top)} \mid \det(H_N)$ via Smith normal form over $\mathbb{Z}$, Lean `thm19` yields the tightened nullity bound immediately. Part (c) is directly certified by exact integer executions in Certificates `cert04` and `cert11`–`cert12`. □

5 Formal Verification Architecture and Scope

- **Layer 1 — Lean 4.36.0 Kernel** (`lean/DeterministicDlpTheory.lean`): Twenty constructive theorems compiled under `-DwarningAsError=true`, with zero external imports, zero `sorry`, zero `grind`, zero `Classical.choice`, and zero custom axioms (audited via `#print axioms`). Layer 1 formalizes both exact potential-optimization square identities (`thm01`–`thm04`), exact LZR baseline and 64×/128× threshold compression arithmetic (`thm05`–`thm07`), pigeonhole sumset covering (`thm08`), strict unreachability of LZR's extreme roots $j = \pm w$, degree-4 savings, pairwise interval disjointness, factor nonnegativity, and constructive fixed-shift telescoping soundness (`thm09`–`thm14`), composite-modulus coset halving and termination induction (`thm15`–`thm16`), M-matrix dominance margins, 2×2 Fischer determinant bounds, prime-order nullity bounds, and P_0 precision bit savings (`thm17`–`thm20`). The Intermediate Value Theorem step of Theorem 2(b), general-

N Koteljanskii–Fischer M-matrix inequality, and Cauchy–Hermite contraction black box of [3] are proved/cited in Sections 3–4 and checked on concrete finite instances in Layer 2.

- **Layer 2 — Zero-Float Python Verifier** (`verification/verify_deterministic_dlp.py`): An AST-audited pure standard-library verifier using strictly `int` and `Fraction` (zero floats, zero numeric `/`, zero `**`, zero `limit_denominator`), emitting 16 SHA-256 certificates in `certificates/MANIFEST.json` and validated by 56 adversarial `pytest` checks in `tests/test_adversarial.py`. Layer 2 verifies the full (A, B, λ) phase boundary, runs exact deterministic Stage I conditional-expectation seed selection on $\mathbb{F}_{23}^*$ and $\mathbb{F}_{47}^*$, exhaustively verifies Q^* separation and truncated Newton–Binomial power-series contraction, certifies Bareiss M-matrix determinants and entrywise nonnegative adjugates, and executes end-to-end deterministic DLP/ECDLP recovery on $\mathbb{F}_{23}^*$, $\mathbb{F}_{47}^*$, $E(\mathbb{F}_{17})$, and $E(\mathbb{F}_{23})$.

References

[1] A. Berman and R. J. Plemmons, *Nonnegative Matrices in the Mathematical Sciences*, Classics in Applied Mathematics, Vol. 9, SIAM, Philadelphia, 1994.

[2] R. Kannan and A. Bachem, *Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix*, SIAM J. Comput. **8** (1979), no. 4, 499–507.

[3] S. Liu, B. Zhang, and K. Ren, *A deterministic subexponential-time algorithm for the discrete logarithm problem*, Cryptology ePrint Archive, Paper 2026/3020, November 23, 2026.

[4] J. M. Pollard, *Monte Carlo methods for index computation (*mod *p)*, Math. Comp. **32** (1978), no. 143, 918–924.

[5] V. Shoup, *Lower bounds for discrete logarithms and related problems*, in *Advances in Cryptology — EUROCRYPT '97*, LNCS 1233, Springer, 1997, pp. 256–266.