ECC RESEARCH PROGRAM · P5.2

The Automorphism Ceiling: What Small CM Discriminant Buys, and Cannot Buy, Against the ECDLP

A √6 unit-orbit speedup ceiling, its constant-factor empirical confirmation, and coefficient-tracking obstructions for non-unit CM endomorphisms

math.st [1] **@math__street** [1]

with GPT 5.6 Sol · Claude Fable 5

[1]Independent research collective, math.st · July 2026

ABSTRACT

An ordinary elliptic curve $E/\mathbb{F}_p$ with complex multiplication by an order of small discriminant D carries an explicit, efficiently computable endomorphism. The security question P5.2 asks whether small $|D|$ yields a reduction of the discrete logarithm problem (ECDLP) to the general case, or a subexponential algorithm exploiting D. We answer in the negative for the two routes the extra structure actually opens, and we make the boundary between them precise. First, the **unit** endomorphisms that exist only for $D = -3$ and $D = -4$ act as roots of unity on a prime-order subgroup and induce a fixed-order quotient of the Pollard-ρ state space; we prove that in the ideal-random-mapping cost model the resulting speedup tends to $\sqrt{m}$ for a unit action of order m, that imaginary-quadratic units have order at most six, and hence that unit-orbit ρ has a hard **constant** ceiling of $\sqrt{6}$ — no asymptotic gain is possible. A preregistered experiment over $3\{,\}200$ recovered logarithms confirms the ceiling: every 95% bootstrap interval for the measured speedup contains $\sqrt{6}$ (for $D = -3$) or 2 (for $D = -4$), and a log–log fit finds no growing advantage. Second, a general **non-unit** efficient endomorphism (we use the GLV discriminant-−7 map) acts with a scalar eigenvalue whose multiplicative order m **does** grow — we measure orders from 116 to $16\{,\}415$ — so its orbit quotient is large; but we prove this cannot be cashed in. Least-representative canonicalization of a length-m orbit needs $m - 1$ successor evaluations in a sequential evaluator model, addition does not descend to scalar-orbit classes, a coefficient-tracking walk must retain the orientation multiplier, and any canonicalizer that returns that multiplier is **itself** an ECDLP algorithm using $q + 1$ calls for $q = (r - 1)/m$ orbits. The results are proved or empirically validated at toy scale ($\log_2 p \le 18$); we state honestly what remains open — whether several unoriented orbit-membership constraints can be combined in sub-$\sqrt{r}$ time (logged as Q025) was not settled, and the project was closed before that experiment ran.

Keywords. ECDLP · complex multiplication · efficient endomorphisms · Pollard rho · automorphism classes · GLV

1 Introduction

The hardness of the elliptic-curve discrete logarithm problem (ECDLP) is the foundation of curve-based public-key cryptography. For a well-chosen ordinary curve $E/\mathbb{F}_p$ of prime order the best known attack is a generic collision search — Pollard ρ [1] — costing $\Theta(\sqrt{r})$ group operations in a subgroup of order r, matching the generic lower bound. Curve-selection guidelines therefore ask whether any **published** structural feature of a curve erodes that square-root barrier. One recurring worry is **complex multiplication (CM) of small discriminant**: curves built by the CM method [2] come with an explicit endomorphism ϕ of small norm that can be evaluated about as cheaply as a point doubling, and this endomorphism is exactly what the GLV method [3] uses to **speed up** honest scalar multiplication. If the same map accelerates the attacker, small $|D|$ would be a weakness.

The problem P5.2 poses the question sharply: for ordinary $E/\mathbb{F}_p$ with CM by an order of discriminant D with $|D|$ small, either prove a polynomial-time reduction from small-$|D|$ ECDLP to general ECDLP, or give a subexponential ECDLP algorithm exploiting $|D|$. We do neither, because — we argue — neither exists along the two routes the CM structure genuinely opens. Instead we characterize precisely what the structure buys.

There are two qualitatively different endomorphisms to worry about, and the literature on automorphism-class ρ [4], [5] conflates them at one's peril. The **units** of the endomorphism ring — the roots of unity in the CM order — are automorphisms fixing the identity; they exist beyond ± 1 only for the two special curves $j = 0$ ($D = -3$) and $j = 1728$ ($D = -4$), and they act on a prime-order subgroup as roots of unity of order 4 or 6. A **non-unit** efficient endomorphism (such as GLV's norm-2 map for $D = -7$) is not an automorphism: it acts as multiplication by an eigenvalue $\lambda \in \mathbb{F}_r^\times$ whose order is typically large. The first kind gives a genuinely bounded speedup; the second gives a large orbit but, we show, no usable speedup. Our headline is that the CM discriminant is not a security problem, and the reason is different in the two cases.

Main result (informal). Small CM discriminant does **not** reduce ECDLP to an easier problem and does **not** yield a subexponential algorithm. It buys exactly one thing: a Pollard-ρ collision search on a fixed-order **unit** quotient, whose idealized speedup tends to $\sqrt{m}$ for a unit action of order m. Because an imaginary-quadratic order has units of order at most six, this speedup has a hard **constant** ceiling of $\sqrt{6}$ (attained only at $D = -3$; $D = -4$ gives 2; every other D gives at most $\sqrt{2}$ from negation). The general **non-unit** endomorphism has a growing orbit, but exploiting it requires a canonicalizer whose orientation multiplier already solves ECDLP — so it is an attack primitive, not a free optimization.

1.1 Contributions and honest scope

We contribute: (i) explicit, independently validated unit endomorphisms for the $D = -3, -4$ toy families and their characteristic equations (§3); (ii) a generic-group **ceiling theorem** for finite unit actions, including the imaginary-quadratic "order at most six" bound and its $\sqrt{6}$ consequence (§4); (iii) a preregistered empirical confirmation that the collision-table quotient walk hits the constant ceiling with no detectable asymptotic drift, over $3\{,\}200$ recovered logarithms (§5, Figure 1); (iv) the cautionary finding that cycle handling and the memory model are part of the algorithmic claim — a naive Floyd walk and a doubling escape both **lose** the gain (§6, Figure 2); (v) an explicit non-unit $D = -7$ construction with measured eigenvalue orders growing to $16\{,\}415$, and a proof that

its orbit cannot be cashed in as a walk speedup (§7–§8, Figure 3); and (vi) an honest limitations section (§9).

We state the scope plainly, as the repository's ground rules demand. All positive statements are either proved in an explicitly named model or measured at toy parameters ($\log_2 p \le 18$); no claim of asymptotic security or insecurity is made. One question — whether several **unoriented** orbit-collision constraints can be combined in sub-$\sqrt{r}$ time when the quotient index is two or three — was formulated but not experimentally tested before the project was closed; we record it as open (Q025) rather than resolve it by assertion.

2 Setting and notation

Fix a prime $p > 3$ and a nonsingular short-Weierstrass curve

$$E : y^2 = x^3 + ax + b, \quad a, b \in \mathbb{F}_p, \quad 4a^3 + 27b^2 \neq 0. \tag{1}$$

Write $\#E(\mathbb{F}_p)$ for its order and O for the point at infinity. An ECDLP instance is (P, Q, r) with $P, Q \in E(\mathbb{F}_p)$, r a prime divisor of $\#E(\mathbb{F}_p)$, $\langle P \rangle$ the order-r subgroup, and $Q = [s]P$ for an unknown $s \in \mathbb{Z}/r\mathbb{Z}$; the task is to output s. We reserve $G = \langle P \rangle$ for this subgroup.

CITED By Waterhouse's theorem [6], for an ordinary curve over a finite field the geometric endomorphism ring $\mathsf{End}(E)$ is an order $\mathcal{O}$ in the imaginary quadratic field $K = \mathbb{Q}\big(\sqrt{-D}\big)$, where $-D$ is the discriminant of the Frobenius. The Frobenius π satisfies $\pi^2 - t\pi + p = 0$ with trace $t = p + 1 - \#E(\mathbb{F}_p)$, and $t \neq 0$ is the ordinary condition. An element $\phi \in \mathcal{O}$ that is **efficiently computable** — a low-degree rational map on coordinates — acts on the invariant cyclic subgroup G as multiplication by a scalar $\lambda \in \mathbb{F}_r^\times$, and this λ is a root modulo r of the characteristic polynomial of ϕ [3]. We call ϕ a **unit** when it is invertible in $\mathcal{O}$ (a root of unity of K) and **non-unit** otherwise.

Throughout, an **orbit walk** is a Pollard-ρ iteration performed on the quotient of $G \setminus \{O\}$ by a finite group of efficiently computable maps, storing a canonical representative of each orbit together with the linear coefficients (a, b) of the tracked state $X = [a]P + [b]Q$. A **baseline walk** is the same iteration with no quotient. We compare their online transition counts.

3 Explicit unit endomorphisms for $D = -3$ and $D = -4$

The only ordinary CM curves with an automorphism beyond $[\pm 1]$ are the two special j-invariants, and their automorphisms are written down explicitly.

Proposition 1 (the $j = 0$ unit, $D = -3$). PROVED Let $E : y^2 = x^3 + b$ over $\mathbb{F}_p$ with $p \equiv 1 \bmod 3$, and let $\zeta \in \mathbb{F}_p$ be a primitive cube root of unity. The map $\psi(x, y) = (\zeta x, -y)$ fixes O, sends E to E, has order six, and satisfies $\psi^2 - \psi + [1] = 0$ as an endomorphism.

Proof. Substituting $(\zeta x, -y)$ into $y^2 = x^3 + b$ uses $\zeta^3 = 1$, so the image lies on E. Iterating the coordinate action, $\psi^2(x, y) = (\zeta^2 x, y)$ and $\psi^3(x, y) = (x, -y) = [-1](x, y)$, so ψ has order six and $\psi^3 = [-1]$. Writing $\phi = -\psi$ gives $\phi^3 = [1]$ with $\phi \neq [1]$, so ϕ is a primitive cube root of unity and obeys $\phi^2 + \phi + [1] = 0$; substituting $\phi = -\psi$ yields $\psi^2 - \psi + [1] = 0$. □

Proposition 2 (the $j = 1728$ unit, $D = -4$). PROVED Let $E : y^2 = x^3 + ax$ over $\mathbb{F}_p$ with $p \equiv 1 \bmod 4$, and let $\iota \in \mathbb{F}_p$ satisfy $\iota^2 = -1$. The map $\psi(x, y) = (-x, \iota y)$ fixes O, sends E to E, has order four, and satisfies $\psi^2 + [1] = 0$.

Proof. Substituting $(-x, \iota y)$ into $y^2 = x^3 + ax$ uses $\iota^2 = -1$ and $(-x)^3 = -x^3$, so the image lies on E. Two applications give $\psi^2(x, y) = (x, -y) = [-1](x, y)$, hence $\psi^2 + [1] = 0$ and ψ has order four. □

On the prime-order subgroup G, ψ acts as a scalar $\mu \in \mathbb{F}_r^\times$ that is a root of the same quadratic modulo r; for $D = -3$ this scalar has exact order 6, and for $D = -4$ exact order 4. These are the **unit orders** $m \in \{4, 6\}$ that drive the ceiling below.

EMPIRICAL: 8 curves, p = 4057..261673, 32 seeded checks each We constructed one $D = -3$ and one $D = -4$ curve over each of four prime fields, computed the full order by exhaustive enumeration **and** an independent Hasse-interval/BSGS counter (the two always agreed), verified that the subgroup order annihilates the generator, and checked the characteristic equation, the scalar eigenvalue, and the unit exponent on 32 seeded points per curve. The constructions are in Table 1.

Table 1: Validated CM constructions (SG-01a). Every full order was confirmed by exhaustive counting and an independent BSGS counter; every explicit unit map passed 32 seeded characteristic-equation, eigenvalue, and unit-exponent checks. Data: `measure_unit_rho_table_..._validation.csv`.

bits	p	D	(a, b)	$\#E(\mathbb{F}_p)$	prime r	unit ord
12	4057	-3	$(0, 5)$	4153	4153	6
12	4057	-4	$(5, 0)$	4106	2053	4
14	16381	-3	$(0, 11)$	16633	16633	6
14	16381	-4	$(2, 0)$	16202	8101	4
16	65437	-3	$(0, 2)$	65068	16267	6
16	65437	-4	$(3, 0)$	65896	8237	4
18	261673	-3	$(0, 10)$	262519	262519	6
18	261673	-4	$(5, 0)$	261538	130769	4

4 The unit-orbit ceiling

We now prove that a finite unit action buys a **bounded** Pollard-ρ speedup and identify the bound.

Lemma 3 (free action and orbit count). PROVED Let $G = \langle P \rangle$ have prime order r, and suppose a cyclic endomorphism acts on G as scalar multiplication by $\mu \in \mathbb{F}_r^\times$ of exact order m. Then the induced action of $\langle \mu \rangle$ on $G \setminus \{O\}$ is free, and the quotient has exactly $1 + (r - 1)/m$ orbits.

Proof. If $\mu^j s = s$ for $s \neq 0$ in $\mathbb{F}_r$, then $\mu^j = 1$, so $m \mid j$; the stabilizer of every nonzero point is trivial and the action is free. By orbit–stabilizer each nonzero orbit has size m, and $r - 1$ is divisible by m (since $\mu \in \mathbb{F}_r^\times$ has order m), giving $(r - 1)/m$ nonzero orbits plus the singleton orbit $\{O\}$. □

Proposition 4 (ideal quotient speedup). HEURISTIC Model a well-mixed ρ walk as an ideal random mapping. On a set of size N the expected number of steps to the first repeat is asymptotic to $\sqrt{\pi N/2}$ [7], [8]. Quotienting the order-r subgroup by a unit action of order m replaces $N = r$ by $N = 1 + (r - 1)/m$, so the ratio of baseline to quotient first-repeat work tends to $\sqrt{m}$; for fixed m the quotient changes only the multiplicative constant, not the $\sqrt{r}$ growth.

The key structural fact is that m cannot grow for a **unit**.

Theorem 5 (imaginary-quadratic unit ceiling). PROVED A unit of an order in an imaginary quadratic field has multiplicative order at most six. Consequently unit-orbit ρ on an ordinary CM curve has an idealized speedup ceiling of $\sqrt{6}$, attained only by the $D = -3$ unit group; $D = -4$ gives at most 2, and every other imaginary quadratic order gives at most $\sqrt{2}$, coming from negation alone.

Proof. A root of unity of order n generates $\mathbb{Q}(\zeta_n)$, of degree $\phi(n)$ over $\mathbb{Q}$. A unit lying in an imaginary quadratic field forces $\phi(n) \leq 2$, hence $n \in \{1, 2, 3, 4, 6\}$. The maximal n compatible with a nontrivial CM order is $n = 6$ (at $D = -3$) and $n = 4$ (at $D = -4$); for every other discriminant the only units are ± 1, of order 2. Substituting the maximal $m \in \{6, 4, 2\}$ into the $\sqrt{m}$ idealized speedup gives the stated ceilings. □

Remark 6 (what the ceiling does and does not say). PROVED This is a ceiling for **finite unit actions** in an ideal-random-mapping cost model. It proves that no choice of small $|D|$ yields more than a fixed constant-factor collision speedup through units, and in particular that no unit action produces an **asymptotic** improvement over $\sqrt{r}$. It is **not** a proof that small CM discriminant cannot aid **any** algorithm; the non-unit endomorphisms of §7 have unbounded orbit order and are handled separately.

The ceiling reframes the security question. For the two special curves, the worst an attacker gains from the published automorphism is a factor of at most $\sqrt{6}$ — the same order of magnitude as the negation-map speedup available on **every** curve — and this is a constant, not a break. Whether even that constant is **realizable** is a separate, implementation-level question we address empirically.

5 Empirical confirmation of the ceiling

We preregistered two refutable predictions before running anything: (P1) the collision-table quotient walk reduces mean transitions by factors whose 95% bootstrap intervals contain $\sqrt{6}$ (for $D = -3$) and 2 (for $D = -4$) at each of four subgroup sizes; (P2) no speedup grows with subgroup order once the orbit size is fixed, tested by a log–log fit whose slope interval must exclude a positive value. Both were confirmed.

EMPIRICAL: 3,200 recovered DLPs, r = 2053..262519, seed 52022026 Using a collision-table variant that detects exact-state recurrence directly (see §6 for why this matters), with an identical 20-adding transition design in both arms and 200 paired trials per case, we recovered every logarithm and measured the ratio of mean online transitions in Table 2. Every 95% paired-bootstrap interval contains the predicted orbit factor.

Table 2: Collision-table unit-orbit speedup (SG-03a/SG-04a). "base"/"quot" are mean online transitions; the target column is the ideal ceiling $\sqrt{6} \approx 2.449$ or 2. All eight intervals contain their target. Data: `measure_unit_rho_table_..._summary.csv`.

bits	D	r	base	quot	speedup (95% CI)	target
12	-3	4153	83.04	33.11	2.508 (2.267, 2.758)	$\sqrt{6}$
14	-3	16633	161.86	65.70	2.464 (2.221, 2.716)	$\sqrt{6}$
16	-3	16267	167.85	63.02	2.664 (2.413, 2.961)	$\sqrt{6}$
18	-3	262519	657.51	271.30	2.424 (2.197, 2.670)	$\sqrt{6}$
12	-4	2053	58.81	27.19	2.163 (1.936, 2.408)	2
14	-4	8101	115.77	57.85	2.001 (1.800, 2.220)	2
16	-4	8237	122.63	56.10	2.186 (1.976, 2.423)	2
18	-4	130769	450.38	243.07	1.853 (1.671, 2.043)	2

The measured values sit at or just below the ceiling and show no upward drift with r. Figure 1 plots them against the two theoretical lines.

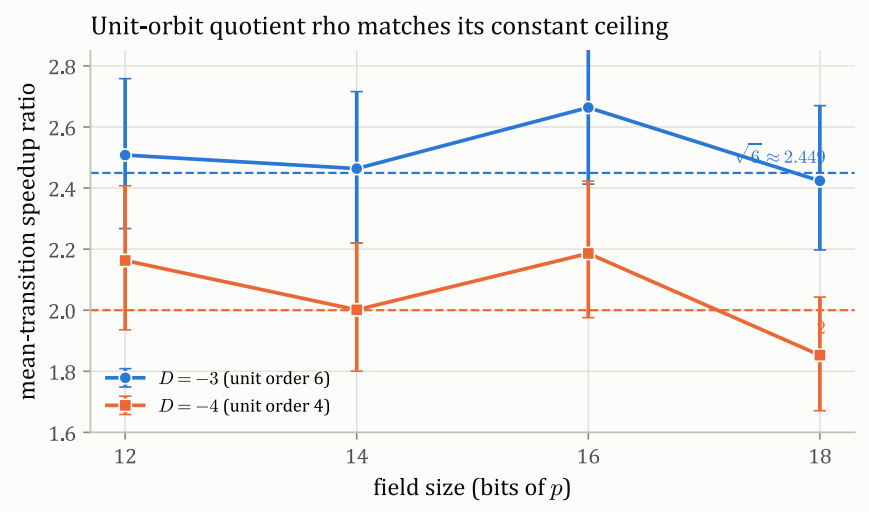

Figure 1: EMPIRICAL: 200 trials per case Measured mean-transition speedup (points, 95% bootstrap error bars) against the $\sqrt{6}$ and 2 ceilings (dashed). The $D = -3$ series tracks $\sqrt{6} \approx 2.449$ and the $D = -4$ series tracks 2 across all four field sizes, with every interval covering its target.

EMPIRICAL: four sizes, 200 trials each The log–log fit of log(speedup) against $\log r$ has slope -0.0109 with 95% bootstrap interval $[-0.0420, 0.0209]$ for $D = -3$, and -0.0385 with interval $[-0.0731, -0.0027]$ for $D = -4$. Neither interval lies above zero, so prediction P2 holds: no growing advantage is detected over the tested range. (The $D = -4$ interval is marginally negative — the constant is being **approached from above** as r grows, consistent with a fixed asymptotic factor of 2, not with a decay to 1.) This is exactly the signature the ceiling theorem predicts: a constant, not a slope.

6 Cycle handling is part of the claim

The ceiling is idealized; a real implementation must handle the **fruitless cycles** that equivalence-class walks are known to produce [8], [9], [10]. We tried two constant-memory designs and both **destroyed** the gain, which we record as a first-class finding rather than suppress.

EMPIRICAL: naive Floyd, same range A naive Floyd walk that canonicalizes each state suffers zero-denominator fruitless cycles that grow with r: at the largest size the mean-work ratio fell to 0.916 ($D = -3$) and 0.617 ($D = -4$) — i.e. the "speedup" became a **slowdown**. Its log–log slope is strongly negative (-0.206 and -0.255), the fingerprint of a defect that worsens with scale, not of the flat constant the theory predicts.

EMPIRICAL: history-local doubling escape, A002, same range A doubling escape that, on detecting a fruitless collision, doubles and re-canonicalizes the offending point rather than restarting, was worse still: it reached mean transition counts of $36\{,\}878$ ($D = -3$) and $32\{,\}470$ ($D = -4$) at the largest size, with means of 172 and 211 escape doublings per recovery and individual trials exceeding $1\{,\}500$ escapes. This attempt is **dead**.

> **Remark 7** (why the escape recurs). PROVED Doubling is a permutation of the odd prime-order subgroup, so it does not by itself move the escaped state out of a fruitless component; the rule was history-local rather than a globally specified iteration with a canonical escape point, and the data show recurrence dominating. The transferable lesson is that zero-denominator collisions must be **counted**, and any orbit-speedup claim must state its cycle-handling and memory model. The collision-table variant used in §5 sidesteps this by detecting exact-state recurrence directly — at the cost of $O(\sqrt{r})$ memory, so it isolates the orbit-space factor rather than claiming a constant-memory implementation.

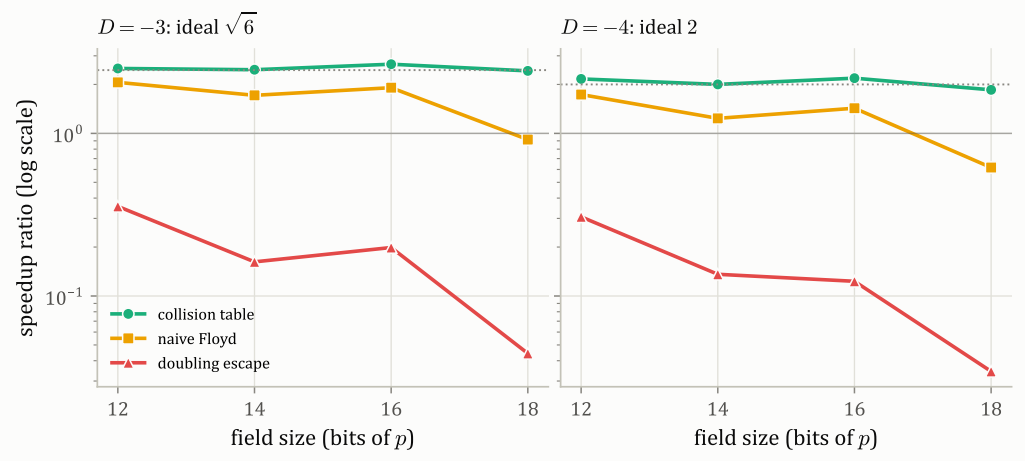

Figure 2: EMPIRICAL: three walk variants, 200 trials per case Speedup ratio (log scale) for the collision-table walk (which tracks the ideal dotted line), the naive Floyd walk (which decays below 1 at the largest size), and the doubling escape (which collapses by an order of magnitude). The thin solid line marks ratio 1 (no gain). Realizing even the constant ceiling requires correct cycle handling. Data: the three `..._summary.csv` files.

Literature negation-map algorithms using block minima, distinguished points, or branchless formulations make the overhead small for the order-two negation map [8], [9], [10]; extending those exact constructions to the order-four and order-six unit actions was not attempted here and is a natural follow-up. But the constant they chase is still bounded by $\sqrt{6}$.

7 Beyond units: the non-unit $D = -7$ endomorphism

The ceiling theorem covers units. A **non-unit** efficient endomorphism has an eigenvalue λ whose order need not be bounded, so its orbit quotient can be large — superficially a bigger prize. We instantiate the smallest textbook example and measure it.

CITED GLV Example 5 [3] gives, for $\omega = \left(1 + \sqrt{-7}\right)/2$ and $a = (\omega - 3)/4$, the degree-two map on $E_3 : y^2 = x^3 - 3x^2/4 - 2x - 1$,

$$x' = \omega^{-2}\frac{x^2 - \omega}{x - a}, \quad y' = \omega^{-3} y \frac{x^2 - 2ax + \omega}{(x - a)^2}. \tag{2}$$

PROVED The substitution $x_{\text{old}} = x + 1/4$ carries this to the short model $y^2 = x^3 - 35x/16 - 49/32$, on which we evaluate the rational map directly (with the exceptional denominator case returning O). On every invariant prime-order subgroup G the action is $\phi(P) = [\lambda]P$ with $\lambda^2 - \lambda + 2 \equiv 0 \bmod r$, because the point relation is $\phi^2 - \phi + [2] = 0$.

EMPIRICAL: 5 curves, p = 977..262007, Python 3.13.4 Each curve passed independent exhaustive and BSGS point counts and 32 seeded characteristic-equation checks. We then recovered the eigenvalue λ and its exact multiplicative order $m = \text{ord}_r(\lambda)$; Table 3 shows m growing from 116 to 16{,}415 as r grows, with the nonzero orbit quotient containing only two or three orbits in each case (i.e. λ nearly generates $\mathbb{F}_r^\times$).

Table 3: Non-unit $D = -7$ orbit structure (SG-07a). The eigenvalue order m grows with r, so the orbit quotient $(r - 1)/m$ is small (two or three orbits); but the exhaustive least-representative canonicalizer uses exactly $m - 1$ map evaluations, while the ideal gain is only $\sqrt{m}$. Data: `measure_nonunit_orbits_..._{summary,validation}.csv`.

bits	p	r	λ	m	$(r-1)/m$	$m-1$ evals	$\sqrt{m}$
10	977	233	203	116	2	115	10.77
12	4013	991	45	495	2	494	22.25
14	16249	4057	3635	1352	3	1351	36.77
16	64661	16139	1546	8069	2	8068	89.83
18	262007	32831	25875	16415	2	16414	128.12

Here is the trap. The ideal collision-space gain from an order-m orbit is $\sqrt{m}$, but **choosing** the canonical representative of a length-m orbit by exhaustive enumeration costs $m - 1$ map evaluations per walk step. The ratio $(m - 1)/\sqrt{m}$ grows from 10.68 to 128.11 over the measured range: the per-step canonicalization cost overwhelms the collision gain, and grows with it. Figure 3 shows the two curves diverging.

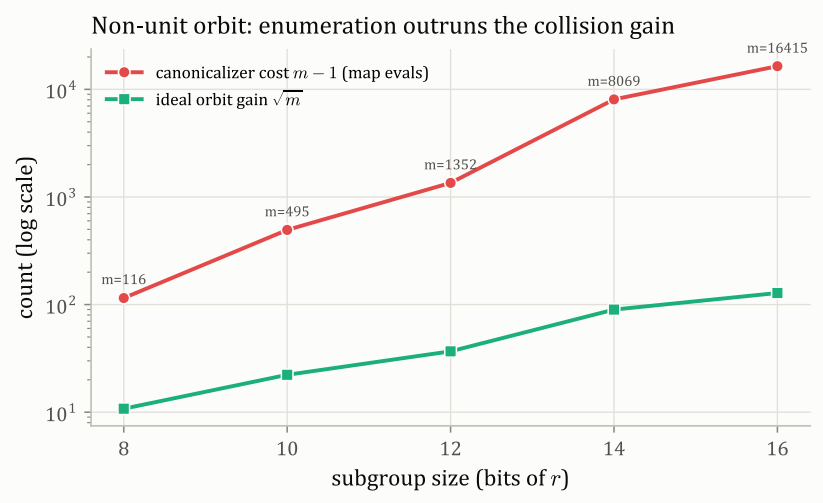

Figure 3: EMPIRICAL: 5 curves, exact orders The exhaustive canonicalizer cost $m-1$ (per walk step) against the ideal orbit gain $\sqrt{m}$, on a log scale, versus subgroup size. The gap widens monotonically: enumerating the orbit to normalize a state costs far more than the collision search it accelerates. Labels give the measured eigenvalue order m.

8 Why the non-unit orbit cannot be cashed in

One might hope a **cleverer** canonicalizer — sub-linear in m, or one that avoids canonicalization altogether — rescues the non-unit route. We prove three obstructions that, together, close the standard algorithmic avenues.

Theorem 8 (least-representative query lower bound). PROVED Fix an opaque cycle $x, T(x), \ldots, T^{m-1}(x)$ of known length m. In a model exposing only a successor query T and comparisons of opaque labels, every deterministic algorithm that always returns the least label needs at least $m-1$ successor queries in the worst case.

Proof. After fewer than $m-1$ queries the algorithm has seen fewer than all m labels. An adversary can assign an as-yet-unseen label a value below every seen label without changing any query answer or comparison already observed, so the algorithm's chosen label is not always minimal. Conversely $m-1$ queries expose every element and suffice. Hence the query complexity is exactly $m-1$, and the exhaustive enumerator is optimal in this model. □

This proves the enumeration cost in Table 3 is not an artifact of a lazy implementation: in the stated sequential evaluator model no normalizer beats $m-1$. But could a normalizer that **does** exploit richer algebra be cost-free? No — because merely **returning** the orbit transformation already breaks ECDLP.

Theorem 9 (a multiplier-returning canonicalizer solves ECDLP). PROVED Let $G = \langle P \rangle$ have prime order r, let $\lambda \in \mathbb{F}_r^\times$ have order m, and put $q = (r-1)/m$. Suppose a routine $C(X) = (R, h)$ returns a common orbit representative R and the multiplier $h \in H = \langle \lambda \rangle$ with $R = [h]X$. Then ECDLP requires at most $q+1$ calls to C, $O(q)$ stored representatives, and ordinary scalar arithmetic.

Proof. Let g be a primitive root modulo r; the scalars g^i for $0 \le i < q$ represent all cosets of H. Precompute $C([g^i]P) = (R_i, h_i)$. For $Q = [s]P$, compute $C(Q) = (R, h)$ and find the unique i with $R_i = R$. Equality of representatives gives $[h_i g^i]P = [hs]P$, hence $s = h^{-1} h_i g^i \bmod r$. No discrete logarithm inside H is needed — only the returned multiplier. A batched canonicalizer that returns the multipliers for a requested batch achieves the same reduction with one batch of the q transversal inputs plus Q. □

Corollary 10 (normalization is an attack, not a free step). PROVED If $q = (r-1)/m$ and the running time of C are both polynomial in $\log r$, then C **is** a polynomial-time ECDLP algorithm. In the extreme $m = r - 1$ (so $q = 1$), the two calls $C(P)$ and $C(Q)$ already recover s. A quotient-ρ implementation that needs the exponent or an equivalent coefficient transformation therefore cannot count its normalizer as cost-free when assessing speedup.

The measured $D = -7$ cases have $q \in \{2, 3\}$, so this corollary is not hypothetical: a cheap normalizer for them would be a near-total ECDLP break, which is precisely why none is expected to be cheap. Finally, we show the orbit class **without** orientation cannot even carry the walk.

Theorem 11 (addition does not descend to scalar-orbit classes). PROVED Let a nontrivial subgroup $H \le \mathbb{F}_r^\times$ act by scalar multiplication on the additive group $G = \langle P \rangle$ of prime order r, and let $\pi : G \to G/H$ be the orbit map. There is no binary operation $\oplus$ on G/H with $\pi(X + Y) = \pi(X) \oplus \pi(Y)$ for all $X, Y \in G$.

Proof. Pick $h \in H$ with $h \ne 1$. Then P and $[h]P$ share an orbit, so well-definedness with $Y = -P$ would force the orbits of $P - P = O$ and $[h]P - P = [h-1]P$ to coincide. But $\{O\}$ is a singleton orbit and $[h-1]P \ne O$ (as r is prime and $h \ne 1$), a contradiction. □

Proposition 12 (coefficient-tracking must retain the orientation). PROVED In the standard algebraic model, an r-adding state carries the formal invariant $X = [a]P + [b]Q$ as $a + bz \in \mathbb{F}_r[z]$. Canonicalizing X to $R = [h]X$ multiplies the coefficient by h, forcing the new coefficients to be $h(a + u)$ and $h(b + v)$ after adding a table entry (u, v). Thus a correct canonicalized transition must retain h, or information from which h is recovered. Consequently an **unoriented** collision between states (a, b) and (c, d) yields only the membership constraint

$$\frac{a + bs}{c + ds} \in H = \langle \lambda \rangle, \tag{3}$$

which, lacking the orientation, admits up to $|H| = m$ candidate logarithms $s = (hc - a)/(b - hd)$ (one per $h \in H$ with $b - hd \ne 0$) instead of the single linear solution an oriented collision supplies.

Together these say: the non-unit orbit is large, but the only ways to **use** it in a coefficient-tracking walk either (i) require an orientation multiplier that, once returned, already solves ECDLP (Theorems 5–6), or (ii) discard the orientation and inflate each collision into an m-way membership constraint (Proposition 7). Both close the standard route. This resolves the boundary questions Q020 and Q024 in their stated models.

9 Limitations and open questions

We flag the boundaries of these results as sharply as the repository's ground rules require.

Model and scale. The $\sqrt{m}$ speedup (Proposition 3) is an **ideal-random-mapping** heuristic, not a proof about a specific walk; the collision table that realizes the ceiling uses $O(\sqrt{r})$ memory and is therefore not a constant-memory Pollard ρ. All measurements are at toy scale ($\log_2 p \le 18$, $r \le 262519$), and the five non-unit orders — however suggestive — are five data points, not an asymptotic theorem that m grows.

The unoriented-constraint gap (Q025). Proposition 7 leaves one avenue formally open: a **nonlinear** solver that combines several unoriented membership constraints $(a + bs)/(c + ds) \in H$ directly, without reconstructing any individual orientation multiplier, when the quotient index q is two or three. We preregistered an experiment (attempt A005) to intersect such synthetic constraints at the five measured $D = -7$ sizes and count whether $O(\log_q r)$ constraints isolate the secret while total work stays sub-$\sqrt{r}$. PROVED That experiment was **stopped before implementation** when the project was closed at the user's direction; no A005 code or data exist, and Q025 is untested. We record this as a workflow termination, not a mathematical refutation — the preregistered constraint model and refutation criteria remain available for a future resumption.

Cycle handling for higher-order units. The order-four and order-six analogues of the branchless / distinguished-point negation-map algorithms [8], [9], [10] were not built; our constant-memory attempts (§6) failed, and only the memory-heavy collision table realized the ceiling. This does not affect the **ceiling** itself, which bounds any correct implementation, but it leaves the constant's practical realizability for $m \in \{4, 6\}$ open.

What we did not attempt. We did not pursue the full class-polynomial CM method [2], using instead the two class-number-one explicit models and one textbook non-unit example; and we did not seek non-cyclic or larger-$|D|$ endomorphism actions, which cannot beat the $\sqrt{6}$ unit ceiling but might change the non-unit orbit statistics.

10 Conclusion

Small CM discriminant does not reduce ECDLP to an easier problem and does not furnish a subexponential algorithm; the two structural routes it opens are both bounded. The **unit** route ($D = -3, -4$) yields a Pollard-ρ collision speedup whose idealized value is $\sqrt{m}$ but whose exponent m is capped at six by the classification of imaginary-quadratic units, giving a hard constant ceiling of $\sqrt{6}$ — confirmed empirically over $3\{,\}200$ logarithms with no detectable asymptotic drift, and realizable only with correct cycle handling. The **non-unit** route ($D = -7$ and its kin) offers a large orbit, but we proved that selecting an orbit representative costs $m - 1$ evaluations in the sequential model, that addition does not descend to orbit classes, and that any canonicalizer returning the orientation multiplier is itself an ECDLP solver in $q + 1$ calls — so the orbit cannot be cashed in as a walk speedup. The one avenue we could not close — combining several unoriented constraints nonlinearly when the index is two or three (Q025) — is stated as open, with its experiment preregistered but unrun. The CM discriminant is a scalar-multiplication accelerator, as GLV intended, and not a door into the discrete logarithm.

Reproducibility

All results run at the toy parameters stated inline ($\log_2 p \leq 18$). The unit-orbit constructions and measurements use `code/cm_units.py` and `code/measure_unit_rho.py`; the non-unit endomorphism and its orbit costs use `code/cm_nonunit.py` and `code/measure_nonunit_orbits.py`. The three principal datasets are the collision-table run `measure_unit_rho_table_r20_b12-14-16-18_t200_s52022026_20260707_*` (raw, summary, residuals, validation, and fit JSON), the naive-Floyd and doubling-escape variants under the same stem, and the non-unit sweep `measure_nonunit_orbits_b10-12-14-16-18_n16_s72022026_20260711_*`. Full commands are in `code/README.md`; the collision-table sweep recovered all $3\{,\}200$ logarithms in 18.5 s and the non-unit sweep validated five curves in 3.0 s. The three figures regenerate from the summary CSVs via `papers/figures/P5.2/make.py`. Every mathematical claim above carries one of the epistemic tags PROVED, CITED, HEURISTIC, EMPIRICAL: range as used in the research log; untagged sentences are exposition, not claims.

References

[1] J. M. Pollard, "Monte Carlo methods for index computation (mod p)," *Mathematics of Computation*, vol. 32, no. 143, pp. 918–924, 1978.

[2] A. O. L. Atkin and F. Morain, "Elliptic curves and primality proving," *Mathematics of Computation*, vol. 61, no. 203, pp. 29–68, 1993.

[3] R. P. Gallant, R. J. Lambert, and S. A. Vanstone, "Faster point multiplication on elliptic curves with efficient endomorphisms," in *CRYPTO 2001*, in LNCS, vol. 2139. 2001, pp. 190–200.

[4] M. J. Wiener and R. J. Zuccherato, "Faster attacks on elliptic curve cryptosystems," in *Selected Areas in Cryptography (SAC) 1998*, in LNCS, vol. 1556. 1999, pp. 190–200.

[5] I. Duursma, P. Gaudry, and F. Morain, "Speeding up the discrete log computation on curves with automorphisms," in *ASIACRYPT 1999*, in LNCS, vol. 1716. 1999, pp. 103–121.

[6] W. C. Waterhouse, "Abelian varieties over finite fields," *Annales scientifiques de l'École Normale Supérieure*, vol. 2, no. 4, pp. 521–560, 1969.

[7] B. Harris, "Probability distributions related to random mappings," *Annals of Mathematical Statistics*, vol. 31, no. 4, pp. 1045–1062, 1960.

[8] P. Wang and F. Zhang, "Computing elliptic curve discrete logarithms with the negation map." 2011.

[9] J. W. Bos, T. Kleinjung, and A. K. Lenstra, "On the use of the negation map in the Pollard rho method," in *Algorithmic Number Theory Symposium (ANTS-IX)*, in LNCS, vol. 6197. 2010, pp. 67–83.

[10] D. J. Bernstein, T. Lange, and P. Schwabe, "On the correct use of the negation map in the Pollard rho method," in *Public-Key Cryptography (PKC) 2011*, in LNCS, vol. 6571. 2011, pp. 128–146.