ECC RESEARCH PROGRAM · P3.3

The Norm Gap in Ideal-to-Isogeny Translation: Certified Lattice Optima versus Prescribed Power Shapes

Exact SVP baselines through 28-bit p, growing pure-power penalties, and the unexplained remainder of the KLPT output scale

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

with GPT 5.6 Sol · Claude Fable 5

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

ABSTRACT

The ideal-to-isogeny translation step behind the Deuring correspondence requires, given a left ideal I of a maximal order $\mathcal{O} \subset B_{p,\infty}$, an equivalent ideal $J \sim I$ of small norm — ideally $N(J) \leq p^{1/2+o(1)}$, the scale of the lattice-theoretic target line, whereas the complete basic KLPT algorithm's ℓ-power output is estimated near $p^{7/2}$ under its heuristics. Problem P3.3 asks whether that gap is structural (short representatives of the prescribed shape fail to exist) or algorithmic (the search does not find them). We attack the question with exact, certified computation below the asymptotic regime: seeded prime-norm neighbor ideals in the explicit maximal order for $p \equiv 3 (\mathrm{mod}\, 4)$, exact norm-aware LLL on their rank-4 normalized norm lattices, true shortest vectors certified by inverse-Gram-bounded exhaustive enumeration, and exact minimum representatives constrained to powers of 2, powers of 3, or 5-smooth values. On 108 near-p ideals across balanced 12/20/28-bit bands, LLL equals exact SVP on every row and the unconstrained optimum sits at a stable constant fraction of $\sqrt{p}$ (per-band mean $N(J)/\sqrt{p}$ between 0.35809 and 0.37189; overall mean exponent 0.40731). Exact pure-power constraints are costly and increasingly so — overall median penalties 222.64 (2^e) and 168.23 (3^e), rising to roughly 2500 at 28 bits — refuting a small-p impression that shape is nearly free; yet every measured shaped optimum stays below $p^{1.013}$, far under the $p^{7/2}$ scale. Within the tested range the verdict is therefore algorithmic/search-structural rather than structural non-existence. All findings are empirical at toy parameters ($\log_2 p \leq 28$); no KLPT implementation was run, and the sampler is not proven uniform in the ideal class group. We also correct a fatally class-biased first baseline (a fixed small input norm forces $N(J) \leq \ell \leq 31$ for every p) and document the certification machinery, including a quadratic coordinate-elimination branch that keeps billion-tuple certified boxes exhaustive rather than censored.

Keywords. quaternion algebras · ideal lattices · KLPT · Deuring correspondence · lattice reduction · supersingular isogenies · exact SVP

1 Introduction

Supersingular-isogeny cryptography rests on the Deuring correspondence [1]: left ideals of a maximal order $\mathcal{O}$ in the quaternion algebra $B_{p,\infty}$ correspond to isogenies from the curve attached to $\mathcal{O}$, with ideal norm equal to isogeny degree. Protocols and security reductions constantly need the **translation** direction — replace an arbitrary ideal I by an equivalent ideal $J \sim I$ whose norm is small, or small **and** of prescribed shape (a power of a fixed small prime ℓ, so that the isogeny can be evaluated as a chain of cheap ℓ-isogenies). The reference algorithm is KLPT [2]. Its lattice-side target is well understood: the normalized norm form of an ideal is a rank-4 positive-definite quadratic form, and Kohel–Lauter–Petit–Tignol bound the product of its successive minima by a quantity of scale p^2, expecting the individual minima — hence the least equivalent-ideal norm — near $\tilde{O}(\sqrt{p})$ for generic ideals ([2], Section 3.1). Against this stands the output of the complete basic KLPT ℓ-power algorithm, estimated near $p^{7/2}$ under the paper's optimistic heuristics ([2], Theorem 7). Between $p^{1/2}$ and $p^{7/2}$ lies a factor of p^3 whose **cause** is the subject of problem P3.3.

The formal target of P3.3 is to find, in polynomial time, an equivalent ideal $J \sim I$ with $N(J) \leq p^{1/2+o(1)}$ — or to define a precise KLPT-style strategy class and prove no algorithm in that class reaches this bound. Neither half of that alternative is resolved here, and we do not claim otherwise. What the present work delivers is the experimental program that the problem statement requests as its first deliverables: an exact, certified, KLPT-independent measurement of **which part of the gap is structural**. The design principle throughout is exactness — no floating-point reduction, no sampling of the search space, no uncertified "shortest" vectors. Every optimum reported in this paper is accompanied by a finite exhaustion certificate, and every claim carries the epistemic tag it carried in the research log.

Main findings (bounded to the measured ranges). On 108 seeded near-p prime-neighbor left ideals in balanced 12/20/28-bit bands ($2203 \leq p \leq 245\,000\,047$):

(i) EMPIRICAL: 108 ideals, 12/20/28-bit p Exact norm-aware LLL attains the certified exhaustive-SVP optimum on 108/108 rows; the least equivalent-ideal norm sits at a constant fraction of $\sqrt{p}$ (per-band mean $N(J)/\sqrt{p} \in [0.35809, 0.37189]$), with mean exponent $\log_p N(J) = 0.40731$ and maximum 0.47016.

(ii) EMPIRICAL: same rows, targets through 4p Forcing the norm to be an exact power of 2 or 3 is a material and growing structural cost: overall median penalties over unconstrained SVP are 222.64 and 168.23, and the 28-bit medians are 2476.82 and 2622.49. A small-p spectrum experiment that suggested penalties rarely exceed 20 was a finite-size effect.

(iii) EMPIRICAL: same rows Nevertheless, every measured shaped optimum but one is at most p, and the single exception has exponent 1.01280. Short pure-power representatives **exist** in every sampled class far below the $p^{7/2}$ KLPT output scale CITED — shaped existence alone does not explain the basic KLPT gap in the tested range. The bounded verdict is algorithmic/search-structural.

1.1 Contributions and honest scope

We contribute (i) a validated exact-arithmetic implementation of the explicit maximal order, its ideals, norm-aware LLL, and three layers of certified exhaustive search (SVP, full normalized-norm spectrum, sparse increasing-target scan) (§3–§4); (ii) a corrected unconstrained baseline through 28 bits with a dependence audit (§5); (iii) exact shape-penalty measurements at two scales, including the reversal of the small-p conclusion (§6); (iv) a documented correction of a class-biased first experiment, kept as a cautionary result (§3.3); and (v) a bounded structural-versus-algorithmic verdict with the explicit list of what remains unmeasured (§7–§8).

The scope limits are strict and worth stating up front. All experiments are at $\log_2 p \leq 28$, far below cryptographic size and below the repository ceiling $\log_2 p \leq 60$. The ideal sampler is a toy sampler: it is **not** proven uniform in the ideal class group and is **not** protocol-derived HEURISTIC. No KLPT implementation was run, so nothing here measures KLPT's own lifting, congruence, or norm-equation steps PROVED. Every conclusion below is tagged and bounded accordingly.

2 Setting and notation

2.1 The quaternion algebra and the explicit maximal order

Fix a prime $p \equiv 3 \pmod 4$. Let $B = \mathbb{Q}\langle i, j\rangle$ be the quaternion algebra with

$$i^2 = -1, \quad j^2 = -p, \quad ij = -ji, \tag{1}$$

which is ramified exactly at p and ∞; write $k = ij$. Among the maximal orders containing $\mathbb{Z}\langle i, j\rangle$ we use the explicit one generated by $i, (1+j)/2, (1+k)/2$ ([2], Lemma 2) CITED, with $\mathbb{Z}$-basis

$$e_0 = 1, \quad e_1 = i, \quad e_2 = \frac{1+j}{2}, \quad e_3 = \frac{i+ij}{2}. \tag{2}$$

Proposition 1 (reduced norm on the order basis). PROVED For $x = \sum_r z_r e_r$ with $z_r \in \mathbb{Z}$,

$$\mathrm{nrd}(x) = z_0^2 + z_0 z_2 + \frac{p+1}{4} z_2^2 + z_1^2 + z_1 z_3 + \frac{p+1}{4} z_3^2. \tag{3}$$

Proof. Substitute $x = (z_0 + z_2/2) + (z_1 + z_3/2)i + (z_2/2)j + (z_3/2)ij$ into $\mathrm{nrd}(a + bi + cj + dij) = a^2 + b^2 + p(c^2 + d^2)$ and expand; the cross terms collect into the two binary blocks displayed. Integrality of the coefficients uses $p \equiv 3 \pmod 4$, so $(p+1)/4 \in \mathbb{Z}$. □

The associated bilinear form has a Gram matrix of determinant scale p^2 (the trace discriminant of the order), matching the successive-minima budget cited above; the repository verifies the order relations, the trace discriminant p^2, and the norm identity by unit tests at $p \in \{11, 19, 43\}$ EMPIRICAL: unit tests, p in {11,19,43}.

2.2 Ideals, equivalence, and the normalized norm

For an integral left $\mathcal{O}$-ideal I, the reduced norm $N(I)$ satisfies $[\mathcal{O} : I] = N(I)^2$. Two left ideals are equivalent, $J \sim I$, if $J = I\gamma$ for some $\gamma \in B^\times$. The translation primitive we measure is the one used by KLPT:

Proposition 2 (equivalent-ideal construction, cited). CITED Let I be an integral left $\mathcal{O}$-ideal and $x \in I$ nonzero. Then $J = I\overline{x}/N(I)$ is an integral left $\mathcal{O}$-ideal, equivalent to I, with

$$N(J) = \frac{\mathrm{nrd}(x)}{N(I)} =: q_I(x). \tag{4}$$

([2], Section 2.4 and Lemma 5.)

Thus minimizing the equivalent-ideal norm over the class of I is exactly minimizing the **normalized norm form** q_I over nonzero $x \in I$: a rank-4 positive-definite integral quadratic-form minimization. Throughout, “exponent” means $\log_p N(J)$ and the Minkowski-scale target line is exponent $1/2$. Our invariant convention — every plot and table reports $q_I(x)$, never $\mathrm{nrd}(x)$ — follows the research log’s recorded rules.

2.3 The KLPT scale, and a literature correction

CITED Under the heuristics of [2] (Theorem 7), the **complete** basic ℓ-power algorithm outputs an ideal of norm approximately $p^{7/2}$ when its prime-norm input is near $p^{1/2}$; the intermediate lifted quaternion has norm scale p^3. PROVED The P3.3 problem prompt’s phrase “KLPT produces ideals around p^3” conflates these two quantities: p^3 is the lifted element, $p^{7/2}$ the final ideal estimate. We measure against the correct $p^{7/2}$ figure. CITED The same paper reports a separate prime-norm representative procedure whose outputs were only slightly larger than $p^{1/2}$ ([2], Sections 3.1, 4.5, Appendix A.1) — consistent with our unconstrained measurements below — while the recent parametrization of maximal orders along ℓ-isogeny paths by Amorós–Clements–Martindale describes the KLPT-generated path as far from optimally short and leaves exploiting their norm forms for short paths as future work [3].

3 Samplers and measured quantities

3.1 Prime-neighbor ideals and the near-p policy

Definition 3 (prime-neighbor ideal). For an odd prime $\ell \notin \{2, p\}$, sample a residue vector $(z_0, z_1, z_2, z_3) \bmod \ell$, not all zero, with $\mathrm{nrd}(\alpha) \equiv 0 (\mathrm{mod}\, \ell)$ for $\alpha = \sum_r z_r e_r$, and set

$$I = \mathcal{O}\ell + \mathcal{O}\alpha. \tag{5}$$

Every sampled I is checked to satisfy $[\mathcal{O} : I] = \ell^2$ (so $N(I) = \ell$) and left-$\mathcal{O}$-closure; rows failing either check abort the run.

Proposition 4 (rejection-free sampling at large ℓ). PROVED Isotropic residues mod ℓ can be sampled without $O(\ell)$ rejection: for random (z_1, z_2, z_3), the condition $\mathrm{nrd}(\alpha) \equiv 0$ is the monic quadratic

$$z_0^2 + z_2 z_0 + C(z_1, z_2, z_3) \equiv 0 (\mathrm{mod}\, \ell), \quad C = z_1^2 + z_1 z_3 + \frac{p+1}{4}(z_2^2 + z_3^2), \tag{6}$$

in z_0, solvable exactly when the discriminant $z_2^2 - 4C$ is a square modulo ℓ; an exact Tonelli–Shanks square root then yields z_0. Sampled triples are redrawn until the discriminant is a square. The resulting ideal still passes the index-ℓ^2 and closure checks.

The **near-p policy** draws ℓ uniformly from the eight primes $\ell_s = \mathrm{nextprime}(p + 2 + s \cdot \max(16, \lfloor p/100 \rfloor))$, $s = 0, ..., 7$, i.e. from the window $(p, \approx 1.08p]$. Why this matters is the content of §3.3.

Remark 5 (what the sampler is not). HEURISTIC Prime-norm neighbor ideals from uniformly sampled nonzero isotropic residues are a useful toy sampler, but they are **not** assumed uniformly distributed in the ideal class group, and they are not protocol-generated. The heuristic would be refuted as a protocol-representative model by a comparison showing a material distributional difference from protocol-generated ideals. All distributional statements below inherit this caveat.

3.2 Measured quantities

Definition 6 (norm-gap quantities). For a sampled ideal I with $N(I) = \ell$: the **unconstrained optimum** is $N(J^*) = \min_{x \in I, x \neq 0} q_I(x)$, reported as the exponent $\log_p N(J^*)$ and the ratio $N(J^*)/\sqrt{p}$; the **LLL value** is $q_I(x_{\text{lll}})$ for the best vector of an exact norm-aware LLL [4] reduction of I (rational arithmetic, Lovász parameter $3/4$, using the reduced-norm bilinear form), and the **LLL approximation factor** is $q_I(x_{\text{lll}})/N(J^*)$; for a shape class $S \subset \mathbb{Z}_{>0}$ (powers of 2, powers of 3, 5-smooth numbers), the **shaped optimum** is $\min\{q_I(x) : x \in I, q_I(x) \in S\}$ and the **shape penalty** is its ratio to $N(J^*)$.

Each optimum is converted back into the actual equivalent ideal $J = I\overline{x}/N(I)$, whose norm and left closure are re-verified independently; the norm identity $N(J) = q_I(x)$ failing on any row aborts the run. A finite **theta prefix** $\Theta_I(n) = \#\{x \in I : q_I(x) = n\}$, $n \leq 2\lceil\sqrt{p}\rceil$, is recorded per row.

Proposition 7 (theta prefixes as class-distinguishing fingerprints). PROVED The finite normalized theta prefix is invariant under ideal equivalence: right multiplication by γ is a bijection $I \to I\gamma$ scaling nrd by $\mathrm{nrd}(\gamma)$, and normalization by the ideal norm removes the scale, so equivalent ideals have identical prefixes. Distinct recorded fingerprints therefore certify distinct normalized norm lattices. Equal finite prefixes may still come from different ideal classes, so fingerprints separate but never merge classes.

3.3 A fatally biased first baseline, corrected

The first campaign (attempt A001; 140 ideals, $7 \leq p \leq 223$, input norms $\ell \leq 31$) produced a seductive picture: LLL exact on 140/140 rows, median $N(J)/\sqrt{p} = 0.35921$, maximum exponent 0.42357, and a ratio to $\sqrt{p}$ that **decreased** with p. The decrease was an artifact.

Proposition 8 (small fixed input norms trivialize the target). PROVED Let I be an integral left ideal with $N(I) = \ell$ prime. Then the central element $\ell \in I$ satisfies

$$q_I(\ell) = \frac{\mathrm{nrd}(\ell)}{N(I)} = \frac{\ell^2}{\ell} = \ell. \tag{7}$$

Hence every class sampled with $\ell \leq 31$ has an equivalent representative of norm at most 31 **independently of p**, and any apparent convergence of $N(J)/\sqrt{p}$ toward 0 is built into the input distribution.

Proof. $\ell = \ell \cdot e_0 \in \mathcal{O}\ell \subseteq I$, and $\mathrm{nrd}(\ell) = \ell^2$ by the norm form of Proposition 1 with $z = (\ell, 0, 0, 0)$. □

The A001 dataset is therefore retained **only** as an arithmetic regression suite PROVED every distributional claim in this paper uses the near-p sampler, which removes the trivial cap because $q_I(\ell) = \ell > p$ is then no better than exponent 1. This correction — found by the Session 2 rerun when $N(J)/\sqrt{p}$ collapsed in the 20-bit band — is, in our view, the single most transferable methodological lesson of the problem: in equivalent-ideal-norm experiments, the input-norm policy **is** part of the class distribution.

4 Exact certification machinery

All searches share one certificate schema, instantiated three times: exact SVP, full normalized-norm spectrum, and sparse increasing-target scan.

4.1 The inverse-Gram box

Proposition 9 (coefficient bounds certify exhaustion). PROVED Let H be the (doubled-norm) Gram matrix of the reduced basis of I in coefficient space, and let R be any known upper bound on the minimal norm value under search. Every integer coefficient vector c with $c^T Hc \leq 2R$ satisfies

$$c_i^2 \leq 2R(H^{-1})_{ii}, \quad i = 0, ..., 3. \tag{8}$$

Scanning the finite box these bounds define, and evaluating the exact integer form on every tuple, is therefore an exhaustive search below $2R$.

Proof. H is positive definite. By Cauchy–Schwarz in the H-inner product, $c_i = (H^{-1}e_i)^T Hc$ obeys $c_i^2 \leq \left((H^{-1}e_i)^T H(H^{-1}e_i)\right)(c^T Hc) = (H^{-1})_{ii}(c^T Hc) \leq 2R(H^{-1})_{ii}$. □

For SVP the bound R is the exact LLL value, so the box is tiny in practice — the maximum across all 108 corrected rows was 80 nonzero tuples EMPIRICAL: 108 rows, 12/20/28-bit p. For the spectrum and target searches, R is the cutoff or target itself and the box can be large; the spectrum implementation evaluates the integer form in overflow-checked NumPy int64 blocks, re-evaluating every stored witness with the exact Python norm form afterwards PROVED. (The first spectrum implementation materialized every lattice point as exact rational objects and needed 104 seconds for three tiny tests; the block rewrite reduced the same tests to 0.37 seconds without changing any certified value.)

4.2 Sparse targets by quadratic coordinate elimination

Scanning pure powers $T = 2^e$ or 3^e in increasing order certifies the first represented power without enumerating the norms below it. At 28 bits, however, a single target box can exceed 10^9 tuples. The solver then switches branch rather than censoring:

Proposition 10 (exact coordinate elimination). PROVED Fix a target T and enumerate three of the four bounded coefficients. The remaining equation is the quadratic

$$at^2 + 2ht + k = 2TN(I), \tag{9}$$

where $a > 0$ is the corresponding diagonal Gram entry and h, k are integers determined by the fixed coefficients. An integral solution exists iff $h^2 + a(2TN(I) - k)$ is a perfect square whose root gives an integral, in-range t via the divisibility test $a \mid (-h \pm \sqrt{\cdot})$. Both directions are exact, so the branch is exhaustive, not heuristic.

On the recorded 28-bit run this branch was exercised on exactly two rows; the largest certified full box held 5,254,365,169 tuples, of which the successful elimination tested only 1,241,059 triples before finding its witness EMPIRICAL: 108 rows, targets through 4p. The box threshold (10^9 tuples) is thereby a strategy switch, not a censoring limit — the final dataset has zero censored rows, where a first version of the run had treated the same guard as a hard censor and left two rows undetermined.

5 The unconstrained norm gap through 28 bits

5.1 Design

The corrected main experiment (attempt A002) uses 18 primes in three balanced bit bands — 12-bit (2203…3719), 20-bit (560083…960031), 28-bit (145000043…245000047) — six trials per prime, 36 rows per band, each band split evenly between $p \equiv 3$ and $7 \pmod 8$, with near-p input norms and master seed 33032028. Every row records the LLL value, the certified exact optimum, both witnesses, the reconstructed equivalent ideals, and the theta fingerprint.

5.2 Results

Table 1: Corrected unconstrained experiment (A002), 108 near-p prime-neighbor ideals, seed 33032028. LLL's approximation factor is exactly 1 on every row. Data: `measure_norm_gap_p2203-245000047x18_t6_s33032028_enearp_20260701_{raw,summary}.csv`.

p bits	n	LLL exact hits	Mean $\log_p N(J)$	Mean $N(J)/\sqrt{p}$	Max nonzero box tuples
12	36	36	0.36378	0.36186	80
20	36	36	0.41647	0.37189	80
28	36	36	0.44167	0.35809	80

EMPIRICAL: 108 near-p ideals, 12/20/28-bit p Across all 108 rows, the LLL exact-hit rate is 100%, the median $N(J)/\sqrt{p}$ is 0.37974, the mean is 0.36395 (normal 95% half-width 0.02398), and the maximum observed exponent is 0.47016 — every row below the $1/2$ line. The mean exponent rises across bands ($0.36378 \longrightarrow 0.41647 \longrightarrow 0.44167$) while the mean **ratio** to $\sqrt{p}$ stays flat within $[0.35809, 0.37189]$: exactly the signature of a $\Theta(\sqrt{p})$-scale optimum whose exponent approaches $1/2$ from below as the constant factor is amortized over more bits (Figure 1).

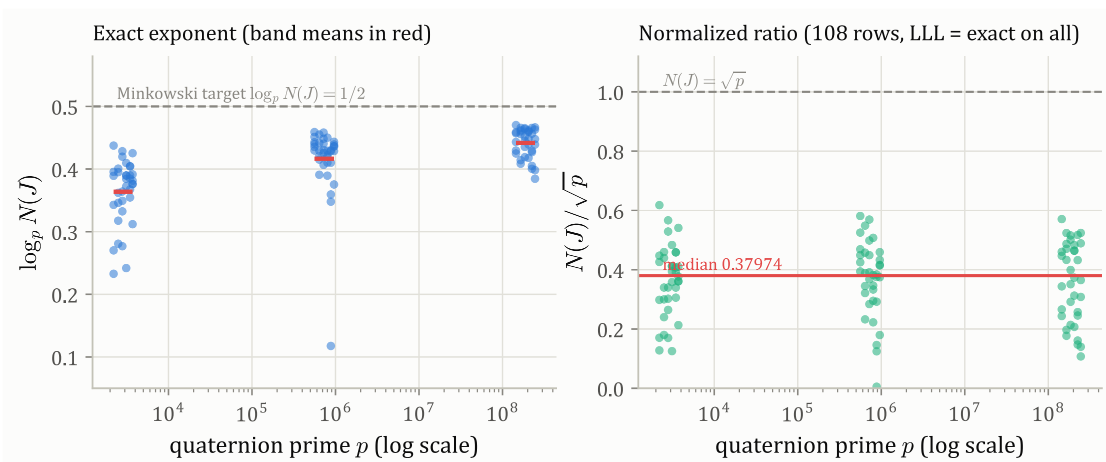

Figure 1: EMPIRICAL: 108 ideals, seed 33032028 The corrected unconstrained baseline. Left: certified exact exponents $\log_p N(J)$ against p (log scale) with per-band means (red) and the Minkowski-scale line $1/2$ (dashed); the three clusters are the 12/20/28-bit bands. Right: the same rows as ratios $N(J)/\sqrt{p}$ with the overall median 0.37974. Data: `measure_norm_gap_p2203-245000047x18_t6_s33032028_enearp_20260701_raw.csv`.

Remark 11 (no reduction gap at rank 4). EMPIRICAL: 108 rows through 28 bits The prediction registered before the Session 2 run allowed LLL to miss on up to 5% of rows; in fact no row in any band shows a gap between exact norm-aware LLL and the certified exhaustive optimum. At this rank and size, the interesting obstruction is **not** lattice reduction quality — it is the norm shape constraint of §6. The exact certificates are also cheap: at most 80 nonzero tuples per row against a pre-registered refutation threshold of 10,000.

5.3 Dependence audit

EMPIRICAL: same 108 rows Mean exponents split by residue are statistically indistinguishable at this resolution: 0.40468 ($p \equiv 3 (\operatorname{mod} 8)$, 95% half-width 0.01726) versus 0.40993 ($p \equiv 7 (\operatorname{mod} 8)$, half-width 0.01256), with mean ratios 0.36511 versus 0.36279. Splitting the input norm by $\ell \equiv 1$ or $3 (\operatorname{mod} 4)$ gives 0.40937 versus 0.40539, again with overlapping intervals. We flag a design limit rather than promote a null: exact input norms form 77 groups across 108 rows and 107 of 108 theta fingerprints are unique, so input-norm and ideal-class effects cannot be separately estimated from this frame; the grouped rows remain in the summary CSV without a dependence claim.

5.4 Cost of the exact baseline

EMPIRICAL: same run Mean exact-search times per band are 0.00281, 0.00464, and 0.00615 seconds; a three-point fit of $\log_2(\text{seconds})$ against $\log_2 p$ has slope 0.07064 with residuals $(-0.05264, 0.10528, -0.05264)$. This is stored for reproducibility and is **not** an asymptotic complexity claim — three points spanning 16 bits cannot support one. The full 108-row run, including validation and plotting, took 13.17 seconds on the recorded environment.

6 The price of shaped norms

The isogeny-side application does not want the **shortest** equivalent ideal; it wants a short ideal whose norm is ℓ-power (or at least smooth), because translation to an isogeny walk requires factoring the norm into cheap prime degrees. The gap question is therefore: what does the shape constraint cost, **structurally**, before any algorithm is blamed?

6.1 Exact spectrum at small p: shape looks nearly free

Attempt A003 enumerates, exhaustively, **every** represented normalized norm up to cutoff p (doubling adaptively to $16p$ if a shape class is empty — never needed) on 70 near-p ideals over the 14 balanced primes $7 \le p \le 223$, seed 33032030, and reads off the least power of 2, least power of 3, and least 5-smooth value.

Table 2: Exact shape penalties at small p (A003), 70 near-p ideals, $7 \le p \le 223$. Every row represented all three shapes by cutoff p. Data: `measure_shape_gap_p7-223x14_t5_m16_b5_s33032030_20260713_{raw,summary}.csv`.

Shape	Mean $\log_p N(J)$	Median $\log_p N(J)$	Median penalty	Max penalty
Unconstrained	0.22152	0.24283	1.00	1.00
2^e	0.31608	0.30642	1.00	21.33
3^e	0.38109	0.40635	1.50	20.25
5-smooth	0.22488	0.24283	1.00	1.43

EMPIRICAL: 70 ideals, 7 <= p <= 223 The medians say shape is nearly free (median penalty 1 for 2^e and 5-smooth, 1.5 for 3^e) with long tails to only about 20; every class contained a power of 2, a power of 3, and a 5-smooth value at normalized norm at most p (Figure 2). Fifteen of the 70 rows are principal ($N(J^*) = 1$), so this sampler certainly overrepresents easy classes relative to an unknown uniform-class baseline — one more reason not to extrapolate. And indeed the extrapolation fails.

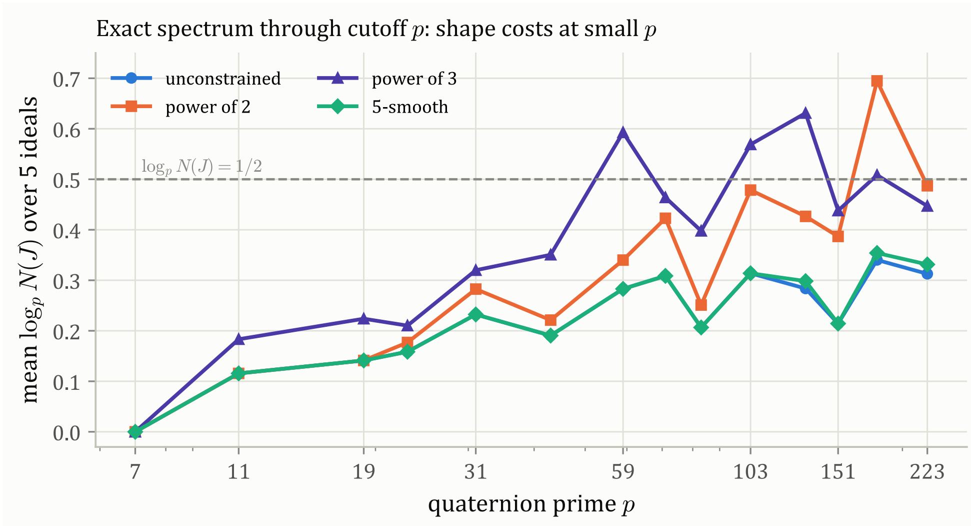

Figure 2: EMPIRICAL: 70 ideals, seed 33032030 Mean exponents per prime for the four shape classes in the exact small-p spectrum (A003). The 5-smooth series hugs the unconstrained optimum; the pure-power series drift upward but stay far below exponent 1. Data: `measure_shape_gap_p7-223x14_t5_m16_b5_s33032030_20260713_summary.csv`.

6.2 Exact sparse targets through 28 bits: shape is a real and growing cost

Attempt A004 replays the **same** 108 ideals as A002 (same seed, same grid) and certifies the least represented 2^e and 3^e normalized norm per row by the increasing-target scan of §4.2, with targets through $4p$.

Table 3: Exact sparse pure-power targets (A004) on the A002 ideals, targets through $4p$, zero censored rows. Data: `measure_power_targets_p2203-245000047x18_t6_m4_s33032028_20260713_{raw,summary}.csv`.

p bits	n	Unc. mean exp.	2^e mean exp.	2^e med. penalty	3^e mean exp.	3^e med. penalty
12	36	0.36378	0.68955	18.31	0.67265	11.86
20	36	0.41647	0.82582	255.28	0.79410	166.13
28	36	0.44167	0.84601	2476.82	0.86253	2622.49
All	108	0.40731	0.78713	222.64	0.77643	168.23

EMPIRICAL: 108 ideals, targets through 4p All 216 pure-power optima were determined. Mean exponents are 0.78713 (2^e) and 0.77643 (3^e) against 0.40731 unconstrained; overall median penalties are 222.64 and 168.23, with maxima 209,715.2 and 43,554.86. The per-band medians in Table 3 are the paper's central quantitative finding: the pure-power penalty grows from tens at 12 bits through hundreds at 20 bits to roughly 2500 at 28 bits (Figure 4). The small-p tail bound of about 20 from Table 2 was a finite-size effect, full stop; the earlier conclusion did not survive one order of magnitude more bits.

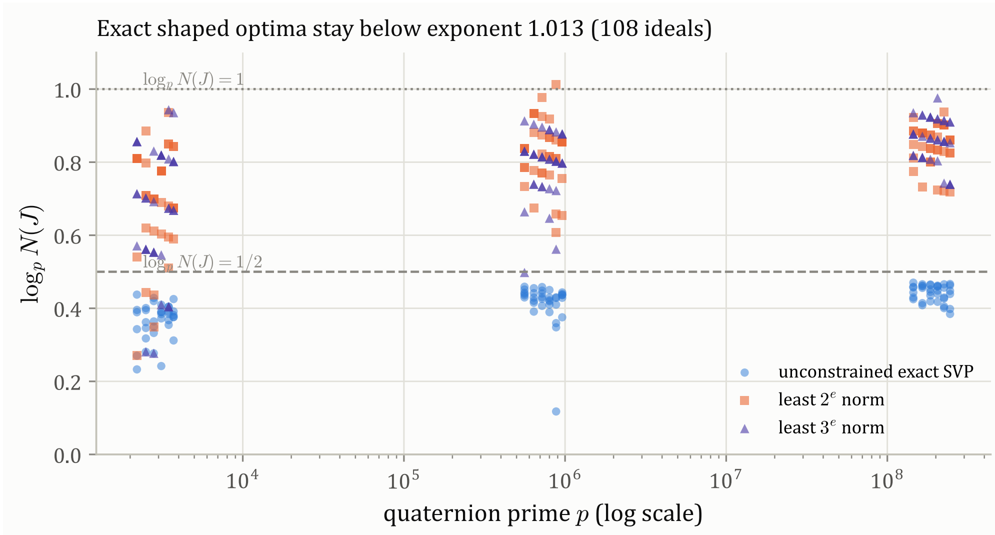

Figure 3: EMPIRICAL: 108 ideals, seed 33032028 Certified exponents of the unconstrained optimum and of the least 2^e and 3^e normalized norms per row (A004). Exactly one point exceeds exponent 1 (a 2^e optimum at 1.01280); the maximum 3^e exponent is 0.97585. Data: `measure_power_targets_p2203-245000047x18_t6_m4_s33032028_20260713_raw.csv`.

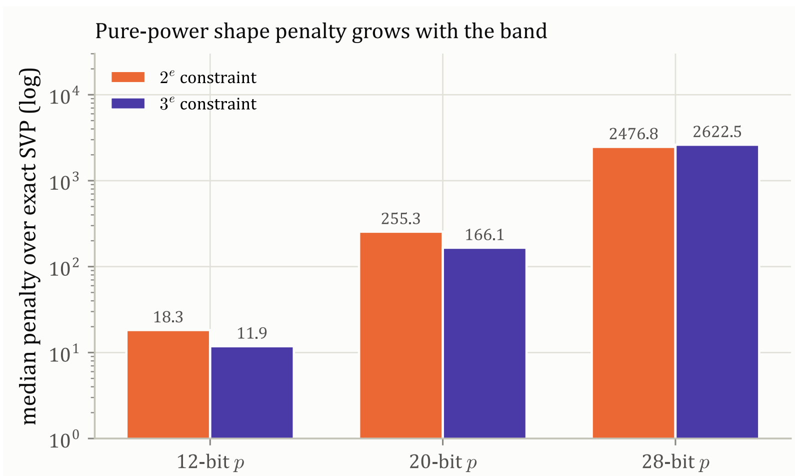

Figure 4: EMPIRICAL: same run Median shape penalty over the unconstrained certified optimum per bit band, log scale (A004 summary). The growth by band — not the small-p medians of Table 2 — is the structurally meaningful trend. Data: `measure_power_targets_p2203-245000047x18_t6_m4_s33032028_20260713_summary.csv`.

> **Remark 12** (yet existence stays cheap on the KLPT scale). EMPIRICAL: same run Even at 28 bits, every shaped optimum except one sits at normalized norm at most p — exponent at most 1 — and the single exception is $p^{1.01280}$. On the scale that matters for the KLPT comparison ($p^{7/2}$ CITED), exact shaped representatives are **three orders of magnitude of exponent** below the algorithm's estimated output. Shape constrains, but it does not remotely explain.

7 The remaining gap to basic KLPT

Putting §5 and §6 together yields the structural-versus-algorithmic verdict the problem requested, in its bounded form.

Bounded verdict. EMPIRICAL: 108 rows through 28 bits (a) No unconstrained lattice-reduction gap exists in the measured range: norm-aware rank-4 LLL always equals certified exact SVP, and the optimum tracks $\approx 0.36\sqrt{p}$. (b) Prescribed pure-power shape is a material structural cost — median penalties grow from tens to thousands across the bands — so structural shape effects are real and were invisible at small p. (c) Every measured shaped optimum is below $p^{1.013}$, so shaped **existence** accounts for almost none of the distance from $p^{1/2}$ to the cited $p^{7/2}$ basic-KLPT output scale. PROVED Because no KLPT implementation was run here, the remaining gap cannot be allocated among KLPT's lifting, congruence, and binary norm-equation steps; within the tested ranges it is algorithmic/search-structural, not an absence of short shaped representatives.

Three points sharpen what this does and does not say. First, the verdict is consistent with the literature's own internal evidence: [2] report their prime-norm representative step returning norms only slightly above $p^{1/2}$, and [3] describe the KLPT path as far from optimally short — both statements attribute the length to the prescribed-shape **search**, not to non-existence, exactly as our exact measurements do at toy scale. Second, the exponent trend in Table 3 (pure-power mean exponents $0.69 \longrightarrow 0.83 \longrightarrow 0.85$, still rising at the measurement boundary) leaves open where the shaped existence exponent saturates; nothing measured here excludes it drifting toward 1 or slightly above as p grows, and the single 1.01280 row is a hint in that direction. Third, the comparison target itself is heuristic: $p^{7/2}$ is an estimate under the 2014 paper's optimistic randomness assumptions CITED, not a measured quantity, and a matched measurement is precisely the open experiment Q099 (§8).

8 Limitations, obstructions, and open questions

8.1 Limitations of record

(L1) Toy parameter range. All distributional results live at $\log_2 p \leq 28$; the repository ceiling is $\log_2 p \leq 60$, and cryptographic sizes are far beyond both. The A003-to-A004 reversal is a concrete demonstration that small-p conclusions in this problem family can fail one order of magnitude up — the same caution applies to our 28-bit conclusions.

(L2) Sampler bias. HEURISTIC The near-p prime-neighbor sampler is not proven class-group uniform and is not protocol-derived; 15 of 70 small-p rows were principal. The dependence audit (§5.3) additionally shows the design cannot separate input-norm from ideal-class effects (77 input-norm groups, 107 of 108 unique fingerprints).

(L3) No KLPT measurement. PROVED The repository contained no quaternion, isogeny, or KLPT module when this problem started; everything used here was built and validated within it, and basic KLPT itself was deliberately left out of scope after the exact-existence questions were closed through 28 bits. The end-to-end gap attribution (lifting versus congruence versus norm-equation costs) is therefore unmeasured.

(L4) A protocol miss, recorded. No numerical prediction was preregistered in the log before the A004 run (the two prior campaigns had preregistered refutation thresholds and met them). The operational expectation — that coordinate elimination would uncensor the two guarded rows — was stated during implementation and confirmed, but the omission is recorded rather than back-filled, and the next matched-KLPT experiment carries preregistration as an explicit requirement.

8.2 Open questions

(Q1 — the matched comparison; logged as Q099.) Implement and validate basic KLPT on the **same** seeded A002/A004 ideals; record output norm, runtime, and success/censoring; report the ratio of KLPT's output to A004′s exact matching $2^e/3^e$ optimum per row. This is the highest-value next experiment: the exact-existence side of the comparison is already closed through 28 bits, so every observed factor above the A004 line is attributable to the algorithm.

(Q2 — distributional realism.) Replace the neighbor sampler by protocol-derived ideals (or prove/refute its class-group near-uniformity at computable sizes via the theta-fingerprint census). Any material difference refutes the sampler heuristic and re-opens the distributional claims.

(Q3 — where does shaped existence saturate?) Extend A004 toward the repository ceiling ($\log_2 p \leq 60$) to test whether pure-power exponents stabilize near 1, and whether the penalty growth in Figure 4 is polynomial in $\log p$ or in p. The certified machinery already switches to coordinate elimination precisely where this extension will need it.

(Q4 — protocol-relevant shapes.) KLPT-style constraints are not bare pure-powers: they are ℓ-powers **with congruence side conditions** from the lifting steps. Measuring the exact cost of those composite constraints — shape plus congruence — on the same ideals would interpolate between A004 and a full KLPT run, and may localize the gap more finely than Q1 alone.

9 Conclusion

Problem P3.3 asks whether the ideal-to-isogeny translation gap — Minkowski says $p^{1/2}$, basic KLPT is estimated at $p^{7/2}$ — is a fact about lattices or a fact about algorithms. At every parameter this project could reach with certified exactness, it is a fact about algorithms. The unconstrained side is as easy as it could possibly be: exact norm-aware LLL is never beaten by exhaustive search, and the optimum is a flat $\approx 0.36\sqrt{p}$ across 16 bits of growth (§5). The shaped side is genuinely constrained — the exact price of a pure-power norm grows by band from tens to thousands, and we document how a small-p experiment concealed this (§6) — but shaped representatives still exist at exponent ≤ 1.013 on every measured row, nowhere near the $p^{7/2}$ output scale of the algorithm whose behavior motivated the question (§7). The honest form of the verdict is bounded: through 28 bits, with a toy sampler, and against a heuristic literature estimate, the basic-KLPT gap is not explained by the non-existence of short shaped representatives. Closing the question — in either direction — now runs through exactly one experiment: a validated KLPT implementation on these same seeded ideals, with the exact shaped optima of this paper as its per-row baseline.

Reproducibility

All experiments run from the repository with Python 3.13.4, SymPy 1.14.0, and Matplotlib 3.11.1 (no SageMath; Windows 11); the shared and problem-local test suites (53 shared + 3 problem tests at Session 2 start) pass before every campaign. Order arithmetic, ideals, norm-aware LLL, exhaustive SVP, spectrum, theta prefixes, and the sparse-target solver live in `lib/quaternion.py`; the experiment drivers are `code/measure_norm_gap.py` (A002; 108 rows, 13.17 s), `code/measure_shape_gap.py` (A003; 70 rows, 21.20 s), and `code/measure_power_targets.py` (A004; 108 rows, 43.53 s, `--max-box 10^9`, targets through $4p$). Master seeds: 33032028 (A002/A004 grid), 33032030 (A003); superseded datasets (seed 33032026 small-ℓ regression baseline, seed 33032027 confounded first scaling run, and a 90-row `t5` near-p precursor) are retained under `problems/P3.3-ideal-to-isogeny-norms/data/`. Recorded CSVs used here: `measure_norm_gap_p2203-245000047x18_t6_s33032028_enearp_20260701_{raw,summary,scaling}.csv`, `measure_shape_gap_p7-223x14_t5_m16_b5_s33032030_20260713_{raw,summary}.csv`, and `measure_power_targets_p2203-245000047x18_t6_m4_s33032028_20260713_{raw,summary}.csv`. All four figures are regenerated from those CSVs by `papers/figures/P3.3/make.py`. Epistemic tags — PROVED, CITED, EMPIRICAL: range, HEURISTIC — are carried verbatim from the research log; untagged sentences are exposition, not claims.

References

[1] M. Deuring, “Die Typen der Multiplikatorenringe elliptischer Funktionenkörper,” *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, vol. 14, pp. 197–272, 1941.

[2] D. Kohel, K. Lauter, C. Petit, and J.-P. Tignol, “On the quaternion ℓ-isogeny path problem,” *LMS Journal of Computation and Mathematics*, vol. 17, no. A, pp. 418–432, 2014, doi: 10.1112/S1461157014000151.

[3] L. Amorós, J. Clements, and C. Martindale, “Parametrizing maximal orders along supersingular ℓ-isogeny paths,” *IACR Cryptology ePrint Archive*, no. 2025/033, 2025.

[4] A. K. Lenstra, H. W. Lenstra, and L. Lovász, “Factoring polynomials with rational coefficients,” *Mathematische Annalen*, vol. 261, no. 4, pp. 515–534, 1982.