Perturbation of Hankel moment singular values and supersingular endomorphism rings via CVP

a p-adic super-resolution law and a fully computed pipeline

Radmir Isyanov
Bauman Moscow State Technical University
isyanovrd@gmail.com

August 3, 2026

Abstract

We present two self-contained but methodologically related results in post-quantum algebra, together with a p-adic strengthening and a complete, reproducible computational realization. **Part I** (spectral) proves an explicit *sufficient* noise bound under which the signal subspace of the Hankel matrix of the power moments of supersingular j-invariants stays strictly separated from the noise subspace, so that the roots of a separable polynomial are recovered deterministically; the propagation constant is written explicitly through the Vandermonde condition number via Weyl and Bauer–Fike (we do not claim tightness against actual failure). Part I is sharpened *non-archimedeanly*: the Teichmüller lift makes the Vandermonde matrix unimodular ($\mathrm{cond}_p = 1$ independently of the number of nodes L), and an *exact* p-adic super-resolution law gives the precision loss as $2\sum_{i<j} v_p(x_i - x_j) + \sum_i v_p(c_i)$, a sharp analogue of Moitra's condition-number bound. **Part II** (arithmetic) reduces the computation of a $\mathbb{Z}$-basis of $\mathrm{End}(E)$ to a rank-4 closest-vector problem and shows that a Shor-based quantum oracle for Weil-pairing discrete logarithms recovers the *exact* Gram matrix of the norm form in $\mathcal{O}(\mathrm{poly}(\log p))$ time, after which LLL in fixed dimension yields a canonical basis; a λ_1-criterion reads the node type from the shortest non-scalar endomorphism. Both parts are joined by an end-to-end theorem and *fully computed on real numbers*: node recovery, real Vélu chains with directly measured degree, deterministic KLPT construction of an endomorphism of prescribed smooth degree together with the ideal-to-isogeny step, reading the torsion action in $\mathbb{F}_{p^4}$ (Tonelli–Shanks) and via Weil pairing + Pohlig–Hellman without an $O(N)$ table, true LLL+Fincke–Pohst, and the simultaneous classification of all three supersingular nodes of $B_{23,\infty}$. Companion code (23 modules) verifies every numerical claim.

Status. A working preprint, not peer reviewed. Rigorous theorems are separated from heuristic components (construction of effective endomorphisms; the availability of two data channels), which are flagged; small-scale examples ($p \le 10^{12}$ for the number-theoretic core, $p \le 1019$ for the pairing route) illustrate but do not replace asymptotic guarantees.

Keywords: Hankel matrix, Weyl inequality, Wedin theorem, Prony/ESPRIT, Vandermonde conditioning, supersingular elliptic curves, endomorphism ring, quaternion order, Weil pairing, Shor's algorithm, CVP, LLL, KLPT, p-adic super-resolution.
MSC 2020: 11G05, 11R52, 15A18, 15A42, 65F35, 94A12, 81P94, 11Y16.

Contents

1 Introduction **2**
1.1 Motivation 2

1.2 Contributions 2
1.3 Related work and status 3

2 Preliminaries: quaternion orders and the Deuring correspondence 3

3 Part I: perturbation of Hankel moment singular values 3
3.1 Setup 3
3.2 Deterministic recovery 4
3.3 p-adic reformulation and super-resolution 5

4 Part II: reduction of $\mathrm{End}(E)$ to CVP 6
4.1 Node criterion 8

5 End-to-end theorem 9

6 Computational realization 9

7 Discussion, limitations, and honest assessment 12

8 Conclusion 13

1 Introduction

1.1 Motivation

Identifying a node of the supersingular isogeny graph and computing the endomorphism ring $\mathrm{End}(E)$ are central problems in isogeny-based post-quantum cryptography [5, 20]. We treat them in two complementary settings. *Part I* (spectral): given noisy power *moments* $m_k = \sum_i c_i x_i^k$ of a set of j-invariants x_i, what is an explicit sufficient noise bound for *deterministic* recovery of the nodes? *Part II* (arithmetic): from the action of endomorphisms on the torsion $E[N]$, recover a $\mathbb{Z}$-basis of $\mathrm{End}(E)$ and the node type. Both lead to classical objects of computational linear algebra and lattice theory: Hankel/Vandermonde matrices, a CVP in a quaternion order, and lattice reduction.

1.2 Contributions

All numerical statements are reproducible (companion code, §6).

1. **Explicit sufficient noise bound (Part I).** Theorem 3.2: an explicit sufficient condition on the noise $\eta = \|\mathbf{E}\|_2$ with propagation constant $\Lambda_{\mathrm{prop}} = 2\kappa_2(\mathbf{V}_\square)(1+\rho)/\sigma_L(\mathbf{H}^{(0)})$, $\rho := \kappa_2(\mathbf{V}_\square)\max_i|x_i|$, via Weyl and Bauer–Fike; clustering asymptotics (Prop. 3.3).
2. **p-adic reformulation and super-resolution.** The Teichmüller lift makes the Vandermonde matrix unimodular ($\mathrm{cond}_p = 1$ for all L; Prop. 3.4, Thm. 3.5); a confluent version (Thm. 3.6) and an *exact* super-resolution law (Thm. 3.7): the p-adic digit loss equals $2\sum_{i<j} v_p(x_i - x_j) + \sum_i v_p(c_i)$.
3. **Reduction to CVP (Part II).** Theorems 4.1, 4.4, 4.6: computing $\mathbf{G}$ and a $\mathbb{Z}$-basis of $\mathrm{End}(E)$ in $\mathrm{poly}(\log p)$; the exact number of CVP calls via the Bruhat–Tits tree (Prop. 4.2).
4. **Node criterion.** The λ_1-criterion (Prop. 4.7) reads the node type (j=0, 1728, generic) from the shortest non-scalar endomorphism; exact orientation count via Eichler's embedding numbers (Prop. 4.8).

5. **End-to-end theorem and full computation.** Theorem 5.1 joins both parts; §6 gives a *fully computed* real-number pipeline: node recovery, real Vélu chains (degree measured by kernel count), deterministic KLPT construction of a prescribed smooth degree with the *ideal-to-isogeny* step, reading the action in $\mathbb{F}_{p^4}$ (Tonelli–Shanks) and via Weil pairing + Pohlig–Hellman (no $O(N)$ table), true LLL+Fincke–Pohst, and the simultaneous classification of all three nodes of $B_{23,\infty}$.

1.3 Related work and status

Conditioning of real Vandermonde/Hankel matrices: Beckermann [2], Gautschi [6], Moitra [15]. Subspace perturbation: Weyl [21], Wedin [19], Bauer–Fike [1]. Endomorphism rings and isogeny graphs: Deuring [4], Voight [18], [5, 20]; orientations: Onuki [16], Colò–Kohel [3], Gross [11]; smooth-degree paths: KLPT [20]. Lattices: LLL [13], Kannan [12], Micciancio–Voulgaris [14]; quantum part: Shor [17]. *Status.* A working preprint; rigorous results are separated from heuristic components (§7), and small-scale examples illustrate but do not replace asymptotics.

2 Preliminaries: quaternion orders and the Deuring correspondence

We collect the arithmetic background used in Part II; see Voight [18] for proofs.

Let p be a prime and $B_{p,\infty}$ the quaternion algebra over $\mathbb{Q}$ ramified exactly at p and ∞; it carries the reduced trace Trd and reduced norm Nrd, and $\bar{\alpha} = \mathrm{Trd}(\alpha) - \alpha$ is the conjugate. An *order* $\mathcal{O} \subset B_{p,\infty}$ is a rank-4 $\mathbb{Z}$-lattice that is also a subring with 1; it is *maximal* if not properly contained in another order. On any order the pairing $\langle \alpha, \beta \rangle = \mathrm{Trd}(\alpha\bar{\beta})$ is symmetric positive definite, so $(\mathcal{O}, \mathrm{Nrd})$ is a positive-definite rank-4 $\mathbb{Z}$-lattice; its Gram matrix $\mathbf{G} = (\langle \theta_i, \theta_j \rangle)$ has $\det \mathbf{G} = \mathrm{disc}(\mathcal{O})^2$, and for a maximal order $\mathrm{disc}(\mathcal{O}) = p$, so $\det \mathbf{G} = p^2$. Two orders are *of the same type* if conjugate in $B_{p,\infty}$; the number of types is the class number $\approx p/12$.

The *Deuring correspondence* [4] makes this arithmetic geometric: the isomorphism classes of supersingular elliptic curves over $\overline{\mathbb{F}}_p$ are in bijection with the types of maximal orders of $B_{p,\infty}$, via $E \mapsto \mathrm{End}(E)$. Isogenies $\phi\colon E \to E'$ correspond to (left) ideals connecting the two orders, with $\deg \phi = \mathrm{Nrd}$ of the ideal; automorphisms of E are the units $\mathcal{O}^\times$, so $|\mathcal{O}^\times|$ distinguishes the special curves $j = 0$ ($|\mathcal{O}^\times| = 6$) and $j = 1728$ ($|\mathcal{O}^\times| = 4$) from generic nodes ($|\mathcal{O}^\times| = 2$). Locally at a prime $\ell \neq p$ one has $B_{p,\infty} \otimes \mathbb{Q}_\ell \cong M_2(\mathbb{Q}_\ell)$, and the maximal orders are the vertices of the $(\ell+1)$-regular Bruhat–Tits tree of $\mathrm{PGL}_2(\mathbb{Q}_\ell)$; ℓ-isogenies are edges. Finally, the total "mass" of the type set is rigidly fixed by the *Eichler mass formula* $\sum_i 1/\left|\mathcal{O}_i^\times\right| = (p-1)/24$, and the number of optimal embeddings of a quadratic order $\mathcal{O}_d$ (discriminant d) into the type set is $\sum_i m_i(d) = h(d)\,(1 - \{d/p\})$ with $\{d/p\} \in \{+1, -1, 0\}$ the split/inert/ramified symbol. These two identities underlie the node-counting of §4.

3 Part I: perturbation of Hankel moment singular values

3.1 Setup

Let $P(x) = \prod_{i=1}^{L}(x - x_i) \in \mathbb{F}_q[x]$ be separable (x_i distinct); in the application the x_i are supersingular j-invariants. Singular values are defined over $\mathbb{R}$ or $\mathbb{C}$, so moments over $\mathbb{F}_q$ are read through a fixed lift $\iota\colon \mathbb{F}_q \hookrightarrow \mathbb{C}$ (Teichmüller, or a p-adic lift below). For nonzero weights $c_i \in \mathbb{C}^\times$ put $m_k = \sum_i c_i x_i^k$ ($0 \le k \le 2N-2$, $N \ge L+1$) and form the Hankel matrix $\mathbf{H} = (m_{i+j})_{0 \le i,j \le N-1}$.

Lemma 3.1 (Vandermonde factorization and exact rank)**.** *With* $\mathbf{V} = (x_i^j)$ *and* $D = \mathrm{diag}(c_i)$ *one has* $\mathbf{H} = \mathbf{V}D\mathbf{V}^\top$*, and for distinct* x_i*, all* $c_i \neq 0$*,* $\operatorname{rank}\mathbf{H} = L$*; hence* $\sigma_1 \geq \cdots \geq \sigma_L > 0 = \sigma_{L+1} = \cdots = \sigma_N$*. The smallest* nonzero *singular value is* $\sigma_{\min}^+ := \sigma_L(\mathbf{H})$*.*

Proof. $\mathbf{H}_{ab} = m_{a+b} = \sum_i c_i x_i^{a+b} = \sum_i (x_i^a) c_i (x_i^b) = (\mathbf{V}D\mathbf{V}^\top)_{ab}$, giving the factorization. $\mathbf{V}$ is an $N \times L$ Vandermonde with distinct nodes, hence of full column rank L, and D is invertible; therefore $\operatorname{rank}\mathbf{H} = \operatorname{rank}(\mathbf{V}D\mathbf{V}^\top) = L$. A real/complex symmetric matrix of rank L has exactly L nonzero singular values, so $\sigma_{L+1} = \cdots = \sigma_N = 0$. □

3.2 Deterministic recovery

We give an *explicit sufficient condition* on the noise below which recovery is deterministic; the proof chains Weyl's inequality (spectral separation), invertibility of the ESPRIT block, and the explicit Bauer–Fike bound on the recovery matrix $A = \mathbf{H}^{(0)-1}\mathbf{H}^{(1)}$ (Prop. 3.3 gives the constant's scaling). We do not claim tightness against actual failure.

Theorem 3.2 (deterministic recovery: explicit sufficient condition)**.** *Let* $\Delta = \min_{i \neq j} |x_i - x_j|$*, let* $\eta := \|\mathbf{E}\|_2$ *be the spectral norm of the moment-noise Hankel,* $\Lambda_{\mathrm{prop}} := 2\,\kappa_2(\mathbf{V}_\square)(1+\rho)/\sigma_L(\mathbf{H}^{(0)})$ *with* $\rho := \kappa_2(\mathbf{V}_\square) \cdot \max_i |x_i|$*, and* $\sigma_{\min}^+ = \sigma_L(\mathbf{H})$ *the smallest nonzero singular value of the full* $N \times N$ *Hankel. If*

$$\eta \;<\; \min\Big\{ \tfrac{1}{2}\sigma_{\min}^+,\; \tfrac{1}{2}\sigma_L(\mathbf{H}^{(0)}),\; \tfrac{\Delta}{2\Lambda_{\mathrm{prop}}} \Big\},$$

then (i) the signal subspace is strictly separated from the noise subspace; (ii) the L *recovered nodes are in* bijection *with the true ones with* $\max_i |\hat{x}_i - x_i| \leq \Lambda_{\mathrm{prop}}\eta < \Delta/2$*; (iii) rounding* $\hat{x}_i$ *to the nearest lifted node returns the exact* x_i*, hence* $P(x) = \prod_i (x - x_i)$ *over* $\mathbb{F}_q$*. In model (B) with entrywise zero-mean noise* $|\varepsilon_k| \leq B$ *one has* $\eta \leq NB$*, so the condition holds whenever* $B < \frac{1}{N}\min\{\frac{1}{2}\sigma_{\min}^+, \frac{1}{2}\sigma_L(\mathbf{H}^{(0)}), \Delta/(2\Lambda_{\mathrm{prop}})\}$*. The propagation constant is* explicit*,* $\Lambda_{\mathrm{prop}} = 2\,\kappa_2(\mathbf{V}_\square)(1+\rho)/\sigma_L(\mathbf{H}^{(0)})$*, with* $\rho := \kappa_2(\mathbf{V}_\square) \cdot \max_i |x_i|$ *(the a-priori bound on* $\big\|\mathbf{H}^{(0)-1}\mathbf{H}^{(1)}\big\|$*; both* $\kappa_2(\mathbf{V}_\square)$ *and* ρ *are finite and computable from the data).*

Proof. The perturbation is Hankel with $\eta = \|\mathbf{E}\|_2$; in model (B) with $|\varepsilon_k| \leq B$ one has $\eta \leq \|\mathbf{E}\|_F \leq NB$. *(i)* Let $\tilde{\mathbf{H}} = \mathbf{H} + \mathbf{E}$ be the noisy full $N \times N$ Hankel ($\mathbf{H}$ the clean one, $\sigma_{\min}^+ = \sigma_L(\mathbf{H})$, $\sigma_{L+1}(\mathbf{H}) = 0$ by Lemma 3.1). By Weyl's inequality $\big|\sigma_k(\tilde{\mathbf{H}}) - \sigma_k(\mathbf{H})\big| \leq \|\mathbf{E}\|_2 = \eta$; since $\eta < \frac{1}{2}\sigma_{\min}^+$, $\sigma_L(\tilde{\mathbf{H}}) \geq \sigma_{\min}^+ - \eta > \frac{1}{2}\sigma_{\min}^+ > \eta \geq \sigma_{L+1}(\tilde{\mathbf{H}})$, i.e. a strict gap between the L-dim signal space and the noise space. *(ii)* We make the node perturbation explicit through the ESPRIT recovery matrix [7, 8]. With the square Vandermonde $\mathbf{V}_\square = (x_i^j)$, $D = \mathrm{diag}(c_i)$, $X = \mathrm{diag}(x_i)$, the two shifted Hankel blocks factor as $\mathbf{H}^{(0)} = \mathbf{V}_\square^\top D\,\mathbf{V}_\square$ and $\mathbf{H}^{(1)} = (m_{r+s+1}) = \mathbf{V}_\square^\top D X\,\mathbf{V}_\square$, so the *recovery matrix*

$$A := \mathbf{H}^{(0)-1}\mathbf{H}^{(1)} = \mathbf{V}_\square^{-1} X\,\mathbf{V}_\square$$

is a genuine diagonalizable matrix (not a singular pencil): its eigenvalues are exactly the nodes x_i and its eigenvector matrix is $S = \mathbf{V}_\square^{-1}$, with $\kappa_2(S) = \kappa_2(\mathbf{V}_\square) = \kappa_2(\mathbf{V}_\square)$. Under noise $\mathbf{H}^{(0)} \to \mathbf{H}^{(0)} + \mathbf{E}_0$, $\mathbf{H}^{(1)} \to \mathbf{H}^{(1)} + \mathbf{E}_1$ ($\|\mathbf{E}_0\|, \|\mathbf{E}_1\| \leq \|\mathbf{E}\|$) the algebraic identity $\tilde{A} - A = (\mathbf{H}^{(0)} + \mathbf{E}_0)^{-1}\big(\mathbf{E}_1 - \mathbf{E}_0 A\big)$, together with Weyl ($\big\|(\mathbf{H}^{(0)} + \mathbf{E}_0)^{-1}\big\| \leq 2/\sigma_L(\mathbf{H}^{(0)})$ once $\|\mathbf{E}_0\| \leq \eta \leq \frac{1}{2}\sigma_L(\mathbf{H}^{(0)})$, which the *second* term of the condition enforces — note $\mathbf{E}_0, \mathbf{E}_1$ are submatrices of $\mathbf{E}$, so $\|\mathbf{E}_0\|, \|\mathbf{E}_1\| \leq \eta$) and $\|A\| = \big\|\mathbf{V}_\square^{-1} X \mathbf{V}_\square\big\| \leq \kappa_2(\mathbf{V}_\square) \max_i |x_i| = \rho$, gives

$$\big\|\tilde{A} - A\big\| \;\leq\; \frac{2}{\sigma_L(\mathbf{H}^{(0)})}\big(\|\mathbf{E}_1\| + \|\mathbf{E}_0\|\,\|A\|\big) \;\leq\; \frac{2\eta}{\sigma_L(\mathbf{H}^{(0)})}\,(1+\rho).$$

Bauer–Fike for the *diagonalizable* A then yields $\min_i |\mu - x_i| \le \kappa_2(S) \left\|\tilde{A} - A\right\| \le \Lambda_{\text{prop}}\eta$ for every eigenvalue μ of $\tilde{A}$, with $\Lambda_{\text{prop}} = 2\kappa_2(\mathbf{V}_\square)(1+\rho)/\sigma_L(\mathbf{H}^{(0)})$; here $\rho := \kappa_2(\mathbf{V}_\square)\max_i |x_i|$ is the concrete value of the earlier "pencil-shift factor", now defined and bounded. The third term of the condition gives $\Lambda_{\text{prop}}\eta < \Delta/2$, so every eigenvalue of $\tilde{A}$ lies in exactly one of the disjoint radius-$\Delta/2$ balls $B(x_i, \Delta/2)$. *Bijection:* since $\eta < \frac{1}{2}\sigma_L(\mathbf{H}^{(0)})$, $\tilde{\mathbf{H}}^{(0)} = \mathbf{H}^{(0)} + \mathbf{E}_0$ is invertible and $\tilde{A}$ has exactly L eigenvalues; a homotopy $\mathbf{E}_0(s), \mathbf{E}_1(s) = s\mathbf{E}_0, s\mathbf{E}_1$, $s \in [0,1]$, keeps $\tilde{\mathbf{H}}^{(0)}(s)$ invertible (Weyl) and moves the eigenvalues continuously from $\{x_i\}$ without leaving their balls (each stays within $\Lambda_{\text{prop}}\eta < \Delta/2$), so exactly one eigenvalue ends in each ball — the map $\hat{x}_i \mapsto x_i$ is a bijection and $\max_i |\hat{x}_i - x_i| \le \Lambda_{\text{prop}}\eta < \Delta/2$. *(iii)* Rounding each $\hat{x}_i$ to the nearest lifted node is then exact; applying ι^{-1} returns the $x_i \in \mathbb{F}_q$ and $P = \prod_i (x - x_i)$. All inequalities hold surely under the stated condition, hence the recovery is deterministic. □

Proposition 3.3 (clustering asymptotics)**.** *If exactly one pair of nodes has distance $\Delta \to 0$ and the rest are fixed, then $\sigma_L(\mathbf{V}_\square) = \Theta(\Delta)$, $\kappa_2(\mathbf{V}_\square) = \Theta(\Delta^{-1})$, $\sigma_L(\mathbf{H}^{(0)}) = \Theta(\Delta^2)$, $\rho = \kappa_2(\mathbf{V}_\square)\max_i |x_i| = \Theta(\Delta^{-1})$, hence $\Lambda_{\text{prop}} = \Theta(\Delta^{-4})$; the noise-norm threshold $\eta < \Delta/(2\Lambda_{\text{prop}})$ is $\Theta(\Delta^5)$ and, equivalently, the noise-*variance *threshold $\sigma^2 \le B^2$ ($\eta \le NB$) is $\Theta(\Delta^{10})$. Under full confluence $\sigma_L(\mathbf{V}_\square) = \Theta(\Delta^{L-1})$.*

Proof. Write the two colliding nodes as x_1, x_2 with $x_2 - x_1 = \Delta \to 0$ and the remaining $x_3, \dots, x_L$ fixed and distinct. In the square Vandermonde $\mathbf{V}_\square$ the determinant $\prod_{i<i'}(x_{i'} - x_i)$ has a single factor $\to 0$ linearly in Δ, and all cofactors stay bounded away from 0; hence $\sigma_L(\mathbf{V}_\square) = \Theta(\Delta)$ while $\sigma_1 = \Theta(1)$, so $\kappa_2(\mathbf{V}_\square) = \sigma_1/\sigma_L = \Theta(\Delta^{-1})$. From $\mathbf{H}^{(0)} = \mathbf{V}_\square^\top D\,\mathbf{V}_\square$ (as in the proof of Thm. 3.2) with $D = \Theta(1)$, $\sigma_L(\mathbf{H}^{(0)}) = \Theta(\sigma_L(\mathbf{V}_\square)^2) = \Theta(\Delta^2)$. Now $\rho = \kappa_2(\mathbf{V}_\square)\max_i |x_i| = \Theta(\Delta^{-1})$ since $\max_i |x_i| = \Theta(1)$, so $1 + \rho = \Theta(\Delta^{-1})$; substituting into $\Lambda_{\text{prop}} = 2\kappa_2(\mathbf{V}_\square)(1+\rho)/\sigma_L(\mathbf{H}^{(0)})$ gives $\Lambda_{\text{prop}} = \Theta(\Delta^{-1} \cdot \Delta^{-1}/\Delta^2) = \Theta(\Delta^{-4})$, so the noise-norm threshold $\Delta/(2\Lambda_{\text{prop}}) = \Theta(\Delta^5)$ and the variance threshold $\sigma^2 \propto (\Delta/(N\Lambda_{\text{prop}}))^2 = \Theta(\Delta^{10})$ (this corrects the naive Δ^{-3}/Δ^8, which treated ρ as $O(1)$; the honest derivation of Λ_{prop} shows $\rho = \Theta(\Delta^{-1})$, verified in `recovery_bound.py`). Under full confluence of all L nodes the Vandermonde determinant vanishes to order $\binom{L}{2}$ and its smallest singular value scales as the shortest edge, $\Theta(\Delta^{L-1})$. □

Over $\mathbb{R}$ the threshold is exponentially small in L (Beckermann [2]); this archimedean defect is removed p-adically.

3.3 p-adic reformulation and super-resolution

Let $W(\mathbb{F}_q)$ be the Witt ring (integers of the unramified extension $\mathbb{Q}_{p^f}$) and $x_i = [\overline{x_i}]$ Teichmüller lifts.

Proposition 3.4 (unimodular Vandermonde)**.** *If $\overline{x_i} \in \mathbb{F}_q$ are distinct then $\det \mathbf{V}_\square = \prod_{i<i'}(x_{i'} - x_i) \in W(\mathbb{F}_q)^\times$, so $\mathbf{V}_\square \in \mathrm{GL}_L(W(\mathbb{F}_q))$ and $\mathrm{cond}_p(\mathbf{V}_\square) = 1$, independently of L and of the node configuration.*

Theorem 3.5 (p-adic Prony; threshold independent of L)**.** *For distinct residues and unit weights, $\mathbf{H}_L^{(0)} \in \mathrm{GL}_L(W(\mathbb{F}_q))$, and Prony recovery of $P = \prod_i (x - x_i)$ is* exact *over $W(\mathbb{F}_q)$ under p-adically small noise $\varepsilon_k \in p^t W$, $t \ge 1$; no digit is lost and the threshold does not depend on L.*

Proof. By Proposition 3.4 the square Vandermonde $\mathbf{V}_\square$ lies in $\mathrm{GL}_L(W)$, and with unit weights $D \in \mathrm{GL}_L(W)$; hence $\mathbf{H}_L^{(0)} = \mathbf{V}_\square D \mathbf{V}_\square^\top \in \mathrm{GL}_L(W)$ and $\det \mathbf{H}_L^{(0)} \in W^\times$, i.e. $v_p(\det \mathbf{H}_L^{(0)}) = 0$. The Prony system $\mathbf{H}_L^{(0)}\mathbf{a} = -\mathbf{m}$ for the coefficient vector of P then has $(\mathbf{H}_L^{(0)})^{-1} \in M_L(W)$, so Cramer/inverse is entrywise integral. Under noise $\boldsymbol{\varepsilon} \in p^t W$ the perturbed solution satisfies $\mathbf{a} - \mathbf{a}^{\text{true}} = (\mathbf{H}_L)^{-1}\boldsymbol{\varepsilon}$;

since $\mathbf{H}_L \equiv \mathbf{H}_L^{(0)} \pmod{p^t}$ is still in $\mathrm{GL}_L(W)$ and its inverse is integral, the difference lies in $p^t W$. Thus every coefficient is correct modulo p^t, in particular exactly for the integer coefficients of P once $t \geq 1$ exceeds their (zero) valuation of failure; no bound involves L. □

Theorem 3.6 (confluent p-adic Prony)**.** *For distinct nodes x_i with multiplicities μ_i, unit leading weights, the confluent Hankel matrix is unimodular over $W(\mathbb{F}_q)$ iff $\max_i \mu_i \leq p$ (the superfactorial $\prod_r r!$ loses unit status once $\mu_i > p$); recovery of $\prod_i (x - x_i)^{\mu_i}$ is then exact with L-independent threshold.*

Proof. The confluent Hankel matrix factors as $\mathbf{H}^{(0)} = \mathbf{V}_{\mathrm{c}} D \mathbf{V}_{\mathrm{c}}^{\top}$ with $\mathbf{V}_{\mathrm{c}}$ the confluent (Hermite) Vandermonde whose columns are the successive *ordinary* derivatives $\partial^r (x^k)|_{x=x_i} = k^{\underline{r}} x_i^{k-r}$, $0 \leq r < \mu_i$ (the natural basis of the confluent moment/Prony problem; the falling factorials $k^{\underline{r}}$ are what introduce the superfactorial below — the divided/Hasse basis would cancel it and is *not* the moment basis here). Its determinant is the confluent Vandermonde determinant $\prod_{i<i'} (x_{i'} - x_i)^{\mu_i \mu_{i'}}$ times a superfactorial $\prod_i \prod_{r<\mu_i} r!$. Over W the difference factors are units (distinct residues), so $\det \mathbf{V}_{\mathrm{c}} \in W^{\times}$ iff the superfactorial $\prod_i \prod_{r<\mu_i} r!$ is a unit, i.e. iff no factor $r!$ with $r < \max_i \mu_i$ is divisible by p, which holds exactly when $\max_i \mu_i \leq p$ (the first p in $r!$ appears at $r = p$). When the matrix is unimodular the argument of Theorem 3.5 applies verbatim, giving exact recovery of $\prod_i (x - x_i)^{\mu_i}$ with L-independent threshold. □

Theorem 3.7 (p-adic super-resolution: exact digit-loss law)**.** *For distinct $x_i \in W(\mathbb{F}_q)$ and weights $c_i \in W \setminus \{0\}$,*

$$v_p\big(\det \mathbf{H}_L^{(0)}\big) \;=\; 2 \sum_{1 \leq i < i' \leq L} v_p(x_{i'} - x_i) \;+\; \sum_{i=1}^{L} v_p(c_i),$$

and p-adic Prony loses exactly $v_p(\det \mathbf{H}_L^{(0)})$ digits: under noise $\varepsilon \in p^t W$ the recovered coefficients agree with the true ones modulo $p^{\,t - v_p(\det \mathbf{H}_L^{(0)})}$. For distinct residues the loss is zero; a pair of nodes agreeing to p-adic depth v costs exactly $2v$ digits.

Proof. $\det \mathbf{H}_L^{(0)} = (\det \mathbf{V}_{\square})^2 \prod_i c_i$ by Lemma 3.1, and $\det \mathbf{V}_{\square} = \prod_{i<i'} (x_{i'} - x_i)$; take v_p. Cramer's rule gives $\mathbf{a} - \mathbf{a}^{\mathrm{true}} = (\mathbf{H}_L^{(0)})^{-1} \boldsymbol{\delta}$ with $\boldsymbol{\delta} \in p^t W$, whose valuation is $\geq t - v_p(\det)$. This is the guaranteed (worst-case) retained precision; the loss is *strictly* less than $v_p(\det)$ only for $\boldsymbol{\delta}$ in a proper p-adic subscheme (Haar-measure zero), so "$t - v_p(\det)$ digits retained" is the correct universal statement, and the digit loss cannot be reduced by taking more moments (floor: supplement, Thm. 1.2). The word *exact* refers to the determinant identity, not to equality in the recovery bound for every $\boldsymbol{\delta}$. □

Remark 3.8 (sharper than Moitra). Archimedean super-resolution (Moitra [15]) needs a minimal separation $\gtrsim 1/M$ and yields only a *polynomial* condition-number bound. The non-archimedean law is an *exact identity*: the loss is the sum of difference-valuations, vanishing for distinct residues and growing by 2 per shared p-adic digit, independently of L.

4 Part II: reduction of $\mathrm{End}(E)$ to CVP

Let $E/\mathbb{F}_{p^2}$ be supersingular; $\mathcal{O} = \mathrm{End}(E)$ is a maximal order in $B_{p,\infty}$ (Deuring [4]). With the reduced trace form $\langle \alpha, \beta \rangle = \mathrm{Trd}(\alpha \bar{\beta})$, $(\mathcal{O}, \mathrm{Nrd})$ is a positive-definite rank-4 lattice with $\det \mathbf{G} = p^2$, covolume p (Voight [18, Ch. 15]).

Theorem 4.1 (reduction to rank-4 CVP)**.** *Given a full-rank effective set $\Lambda = \sum_i \mathbb{Z}\theta_i \subseteq \mathcal{O}$, computing a $\mathbb{Z}$-basis of $\mathcal{O} \supseteq \Lambda$ reduces to finitely many CVP instances in the rank-4 lattice (Λ, Nrd); each is solved exactly in polynomial time (fixed dimension, Kannan [12]).*

Proof. $\Lambda \subseteq \mathcal{O}$ are both full rank-4 lattices in $B_{p,\infty}$, so the index $m = [\mathcal{O} : \Lambda]$ is finite and $m^2 = \det \mathbf{G}_\Lambda / \det \mathbf{G}_\mathcal{O} = \det \mathbf{G}_\Lambda / p^2$; thus m is read off the computable Gram matrix $\mathbf{G}_\Lambda$. Saturation proceeds prime by prime: for each $\ell \mid m$ the overorder $\Lambda \subseteq \Lambda' \subseteq \frac{1}{\ell}\Lambda$ is spanned by Λ together with those $\alpha \in \frac{1}{\ell}\Lambda$ that are integral ($\operatorname{Trd}\alpha, \operatorname{Nrd}\alpha \in \mathbb{Z}$) and closed under multiplication; the candidate α are the vectors of $\frac{1}{\ell}\Lambda/\Lambda \cong (\mathbb{Z}/\ell)^4$ of bounded norm, i.e. the lattice points of $(\Lambda, \operatorname{Nrd})$ nearest to the finitely many coset representatives — a CVP instance. Each CVP in the fixed rank 4 is solved exactly by Kannan enumeration in $\operatorname{poly}(\log p)$ bit operations. Iterating until the Gram determinant reaches p^2 certifies $\mathcal{O}$; the number of iterations is $\Omega(m) = O(\log p)$ (Prop. 4.2). □

Proposition 4.2 (number of CVP calls; Bruhat–Tits)**.** *ℓ-saturation reaches $\mathcal{O}$ using $O(\log p)$ exact rank-4 CVP calls. For an* Eichler *local order of level ℓ^e there are exactly 2 minimal overorders, giving $\#\mathrm{CVP} \le 2\,\Omega(m)$; in general the branching is $\le \ell + 1$, so $\#\mathrm{CVP} \le \sum_{\ell \mid m} (\ell+1) e_\ell$.*

Proof. Locally at ℓ the maximal orders of $B_{p,\infty} \otimes \mathbb{Q}_\ell \cong M_2(\mathbb{Q}_\ell)$ are the vertices of the Bruhat–Tits tree of $\mathrm{PGL}_2(\mathbb{Q}_\ell)$, a regular tree of degree $\ell + 1$; the minimal overorders of a given order are its neighbours, of which there are at most $\ell + 1$. An Eichler order of level ℓ^e sits on the geodesic between two maximal vertices and has exactly 2 neighbours that decrease the level, giving the $2\,\Omega(m)$ headline. For the generic (non-Eichler) local order the full degree $\ell + 1$ may be attained, so each of the $e_\ell = v_\ell(m)$ saturation steps at ℓ costs $\le \ell + 1$ CVP calls, and summing over $\ell \mid m$ gives $\#\mathrm{CVP} \le \sum_{\ell \mid m} (\ell+1) e_\ell$. Since $m \le \sqrt{\det \mathbf{G}_\Lambda}/p$ has $O(\log p)$ prime factors with multiplicity for the inputs considered, this is $O(\log p)$ calls. □

Remark 4.3 (which budget is operative). The order $\Lambda = \sum_i \mathbb{Z}\theta_i$ from the trace form need not be locally Eichler; the *operative* bound is thus the general $\sum_{\ell \mid m} (\ell+1) e_\ell = O(\ell\,\Omega(m))$, i.e. $O(\log p)$ CVP calls with per-step branching $\le \ell + 1$ (the $2\,\Omega(m)$ headline holds only in the Eichler case).

Theorem 4.4 (exact traces via Weil pairing and Shor)**.** *Let N be coprime to p, smooth, $N > 4\sqrt{D}$ with $D = \max_{i,j} \deg(\phi_i \bar{\phi}_j)$. There is a quantum algorithm computing every $\operatorname{Trd}(\phi_i \bar{\phi}_j) \in \mathbb{Z}$ exactly in $\operatorname{poly}(\log p, \log N)$ time: the action of $\psi = \phi_i \bar{\phi}_j$ on $E[N]$ is read via Weil pairings e_N (Miller [9]), the entries of its matrix are Shor discrete logs in μ_N, and* Trd *lifts uniquely because $|\operatorname{Trd}\psi| \le 2\sqrt{D} < N/2$.* The Gram matrix uses only traces *(bounded by $2\sqrt{D}$); reading a reduced norm separately would require $N > 2D$.*

Proof. Fix $\psi = \phi_i \bar{\phi}_j$ (effectively represented: $\bar{\phi}_j$ is the dual isogeny). *Step 1.* ψ acts on $E[N] \cong (\mathbb{Z}/N)^2$ by a matrix $\mathbf{M}_\psi = \left(\begin{smallmatrix} a & c \\ b & d \end{smallmatrix}\right)$. *Step 2.* With $\zeta = e_N(P, Q)$ a primitive N-th root, bilinearity and antisymmetry of the Weil pairing give $e_N(\psi P, Q) = \zeta^a$, $e_N(\psi P, P) = \zeta^{-b}$, $e_N(\psi Q, Q) = \zeta^c$, $e_N(\psi Q, P) = \zeta^{-d}$; the points $\psi P, \psi Q$ and the pairings are computed classically (Miller, Vélu) in $\operatorname{poly}(\log p, \log N)$, and each of $a, b, c, d \in \mathbb{Z}/N$ is a discrete log in μ_N, obtained by Shor in $\operatorname{poly}(\log N)$. *Step 3.* $\operatorname{Trd}(\psi) \equiv a + d$ and $\operatorname{Nrd}(\psi) \equiv ad - bc \pmod N$. *Step 4 (exact lift).* Cauchy–Schwarz for the positive-definite Nrd gives $|\operatorname{Trd}(\psi)| = |\langle \psi, 1 \rangle| \le 2\sqrt{\deg \psi} \le 2\sqrt{D}$; since $N > 4\sqrt{D} \ge 2\,|\operatorname{Trd}\psi|$, the unique representative of $\operatorname{Trd}(\psi) \bmod N$ in $(-N/2, N/2]$ is the exact integer. The Gram entries $G_{ij} = \operatorname{Trd}(\phi_i \bar{\phi}_j)$ are thus recovered exactly; only traces (bound $2\sqrt{D}$) are used, so $N > 4\sqrt{D}$ suffices. □

Remark 4.5 (Shor is only needed for non-smooth N; the demonstrated regime is classical). In every computation of this paper N is chosen *smooth* (so that $E[N]$ is rational in a small extension of $\mathbb{F}_{p^2}$); for smooth N the discrete logs in μ_N are computed *classically* in $\operatorname{poly}(\log N)$ by Pohlig–Hellman (`crypto.py`), with no quantum computer. Shor (Step 2) is genuinely required only for a non-smooth N. Thus the pipeline as run is fully classical and poly-time; the quantum statement is about the general N case, for which a resource estimate (logical qubits $\approx 4 \log_2 N + \log_2 p^k$, Toffolis $\approx 2 \log_2 N \cdot (\log_2 p^k)^2$) is given in the supplement.

Theorem 4.6 (LLL in fixed dimension)**.** *From the exact* $\mathbf{G} \in M_4(\mathbb{Z})$ *(*$\det \mathbf{G} = p^2$*), LLL reduction runs in* $\mathrm{poly}(\log p)$*, and in dimension* 4 *finite enumeration solves SVP/CVP exactly, yielding a canonical* $\mathbb{Z}$*-basis of* $\mathrm{End}(E)$ *with multiplication table; maximality is certified by* $\det \mathbf{G} = p^2$*.*

Proof. Given the exact integral Gram matrix $\mathbf{G}$, a Cholesky factor $\mathbf{G} = R^\top R$ (over $\mathbb{Q}$, exact rationals) presents the lattice; LLL on R terminates in $O(4^2 \log(\max_{ij} |G_{ij}|)) = \mathrm{poly}(\log p)$ arithmetic steps and returns a reduced basis whose Gram–Schmidt norms are within a $2^{O(1)}$ factor of the successive minima (dimension 4 fixed). A reduced basis bounds every coordinate of a shortest/closest vector, so Fincke–Pohst enumeration over the finite box solves SVP and CVP exactly. Multiplying the reduced basis vectors as quaternions and reducing back onto the basis produces the exact structure constants (multiplication table). Finally $\det \mathbf{G} = p^2$ is the discriminant of a maximal order in $B_{p,\infty}$, so it certifies that the computed order is $\mathrm{End}(E)$ itself, not a proper suborder. □

4.1 Node criterion

Proposition 4.7 (λ_1-criterion)**.** *For* $p > 3$*, let* $\mu(\mathcal{O}) = \min\{\mathrm{Nrd}(\alpha) : \alpha \in \mathcal{O} \setminus \mathbb{Z}\}$*. Then* $\mu = 1 \iff \mathrm{Aut}(E) \supsetneq \{\pm 1\} \iff j \in \{0, 1728\}$*; moreover* $\mu = 1, \mathrm{Trd} = 0 \iff j = 1728$ *and* $\mu = 1, |\mathrm{Trd}| = 1 \iff j = 0$*; otherwise* $\mu \geq 2$ *(generic). In characteristic* $2, 3$ *the number of shortest vectors equals* $|\mathrm{Aut}(E)| \in \{24, 12\}$*.*

Proof. A non-scalar $\alpha \in \mathcal{O}$ with $\mathrm{Nrd}(\alpha) = 1$ is a unit of $\mathcal{O}$ outside $\mathbb{Z}$, i.e. an automorphism of E beyond $\{\pm 1\}$; conversely each extra automorphism is such an α. Over $\overline{\mathbb{F}}_p$ with $p > 3$ the only supersingular curves with $\mathrm{Aut}(E) \supsetneq \{\pm 1\}$ are $j = 0$ ($\mathrm{Aut} \cong \mathbb{Z}/6$, ζ_6 with $\mathrm{Trd} = 1, \mathrm{Nrd} = 1$) and $j = 1728$ ($\mathrm{Aut} \cong \mathbb{Z}/4$, i with $\mathrm{Trd} = 0, \mathrm{Nrd} = 1$); this gives $\mu = 1$ with the stated trace, and for every other node the shortest non-scalar element has $\mathrm{Nrd} \geq 2$. The count of shortest vectors is the number of units of $\mathrm{Nrd} = 1$ outside $\{\pm 1\}$, which equals $|\mathrm{Aut}(E)|$; in characteristic $2, 3$ the exceptional curve has $|\mathrm{Aut}| = 24, 12$ respectively. □

Proposition 4.8 (exact orientation count)**.** *For a quadratic order* $\mathcal{O}_K$ *of discriminant* d *and* $p \geq 5$*, with the Eichler symbol* $\{d/p\} \in \{+1, -1, 0\}$ *(split/inert/ramified), the number of primitively* $\mathcal{O}_K$*-oriented supersingular curves is* $h(d)$ *if* $\{d/p\} \leq 0$ *and* 0 *if split (a* $\mathrm{Cl}(\mathcal{O}_K)$*-torsor, Onuki [16]); locally,* $\sum_i m(\mathcal{O}_K, \mathcal{O}_i) = h(d)(1 - \{d/p\})$ *(Eichler [18, Thm. 30.4.7]), counting embeddings up to* $\mathcal{O}_i^\times$*-conjugacy.*

Theorem 4.9 (Eichler mass and total embedding count)**.** *With* $w_i = \left|\mathcal{O}_i^\times\right|/2$ *and* $m_i(d)$ *the number of optimal embeddings* $\mathcal{O}_d \hookrightarrow \mathcal{O}_i$ *up to* $\mathcal{O}_i^\times$*-conjugacy,* $\sum_i 1/\left|\mathcal{O}_i^\times\right| = (p-1)/24$ *and* $\sum_i m_i(d) = h(d)(1 - \{d/p\})$*. Here* $m_i(d)$ *is the number of* conjugacy classes *(orbits) of* $\{\alpha : \mathrm{Trd}\,\alpha = t_0, \mathrm{Nrd}\,\alpha = n_0\}$ *under* $\alpha \mapsto u\alpha u^{-1}$*,* $u \in \mathcal{O}_i^\times$ — not *the element count, which over-counts at the special nodes* $j = 0, 1728$ *(proof). Verified across all three classes of* $B_{23,\infty}$ *for* $d \in \{-3, \ldots, -24\}$ *(companion* `eichler_mass.py`*). This resolves the normalization ambiguity flagged in review (item M2).*

Proof. The Eichler mass formula for the maximal order type set $\{\mathcal{O}_i\}$ of $B_{p,\infty}$ states $\sum_i 1/\left|\mathcal{O}_i^\times\right| = (p-1)/24$; equivalently, with $w_i = \left|\mathcal{O}_i^\times\right|/2$, $\sum_i 1/(2w_i) = (p-1)/24$. For $p = 23$ the three classes have $|\mathrm{Aut}| = (2, 4, 6)$, i.e. $\left|\mathcal{O}_i^\times\right| = (2, 4, 6)$, so $\sum_i 1/\left|\mathcal{O}_i^\times\right| = \frac{1}{2} + \frac{1}{4} + \frac{1}{6} = \frac{11}{12} = (23-1)/24$, confirming the normalization. The optimal-embedding count is the trace formula of Eichler [18, Thm. 30.4.7]: $\sum_i m_i(d) = h(d)\,(1 - \{d/p\})$, where the local factor at p is $1 - \{d/p\} \in \{0, 1, 2\}$ (respectively split, ramified, inert, since $\{d/p\} = +1, 0, -1$) and $m_i(d)$ counts optimal embeddings of the quadratic order $\mathcal{O}_d$ into $\mathcal{O}_i$ *up to* $\mathcal{O}_i^\times$*-conjugacy.* We take this class-level identity as the definition of $m_i(d)$; the M2 concern was that our earlier code counted embedding *elements* $r_i(d) = \#\{\alpha \in \mathcal{O}_i : \mathrm{Trd}\,\alpha = t_0, \mathrm{Nrd}\,\alpha = n_0\}$, which over-counts by the orbit size $\left|\mathcal{O}_i^\times\right| / \left|\mathcal{O}_i^\times \cap \mathbb{Q}(\alpha)^\times\right|$ whenever $\mathcal{O}_i^\times$ has units outside the embedded field $\mathbb{Q}(\alpha)$ — precisely at the special nodes $j = 0, 1728$. The fix is to

divide the element count by that orbit size (equivalently, count conjugacy classes directly), after which $\sum_i m_i(d) = h(d)(1 - \{d/p\})$; we do *not* claim orbits are singletons. Direct enumeration over $d \in \{-3, \ldots, -24\}$ in `eichler_mass.py` confirms the class-level identity in every case. □

5 End-to-end theorem

Theorem 5.1 (moments → node)**.** *Assume* (H1) *bounded noise below the threshold of Thm. 3.2 (or, p-adically, any $t \geq 1$, Thm. 3.5, 3.7);* (H2) *a concrete effective input (B-smooth isogeny cycles of length K with $K \log B = O(\log p)$, giving explicit $D \leq B^{2K}$);* (H3) *smooth $N > 4\sqrt{D}$,* $\gcd(N, p) = 1$*. Then a quantum-classical algorithm deterministically outputs, in* $\mathrm{poly}(\log p)$ *(quantum part): (a) the exact j-invariants; (b) the exact* $\mathbf{G}$ *and an LLL-reduced* $\mathbb{Z}$*-basis of* $\mathrm{End}(E)$*; (c) the type of each node. Budget: one size-N SVD/ESPRIT +* $O(\log p)$ *Shor logs +* $O(\log p)$ *rank-4 CVP + fixed-dimension LLL.*

Proof. The three claims chain the preceding results. (a) Under (H1) the recovery theorem (Thm. 3.2, or its exact p-adic form Thms. 3.5, 3.7) returns $P(x) = \prod_i (x - x_i)$ from the moments, hence the exact j-invariants x_i, using one size-N SVD/ESPRIT. (b) Under (H2) an effective B-smooth generating set of $\mathrm{End}(E)$ is available with $D \leq B^{2K}$; under (H3), $N > 4\sqrt{D}$ smooth and prime to p, Theorem 4.4 recovers the exact integral Gram matrix $\mathbf{G}$ in $\mathrm{poly}(\log p, \log N) = \mathrm{poly}(\log p)$ quantum time ($O(\log p)$ Shor discrete logs), and Theorems 4.1, 4.6 produce an LLL-reduced $\mathbb{Z}$-basis of $\mathrm{End}(E)$ with $O(\log p)$ rank-4 CVP calls and fixed-dimension LLL. (c) The λ_1-criterion (Prop. 4.7) reads the node type — generic, $j = 0$, or $j = 1728$ — from the shortest non-scalar endomorphism of the reduced basis. Summing the budgets gives the stated cost, with the quantum part confined to the $O(\log p)$ Shor logs. □

Proposition 5.2 (the "two channels" are two problems)**.** *The Part II channel (action on $E[N]$) already determines the node: it yields* $\mathrm{End}(E)$ *(Thm. 4.1, 4.6), whose isomorphism class as a maximal order determines the class of E (Deuring [4]), hence $j(E)$. Thus* one *channel identifies a* single *node. Part I solves the different problem of separating* several *nodes from aggregate noisy moments $m_k = \sum_i c_i x_i^k$, in which the individual x_i are not observed. Theorem 5.1 joins two* distinct problems*, not two channels for the same data.*

Proposition 5.3 ((H2) is unconditional in torsion-point protocols)**.** *SIDH/SIKE-type protocols publish the images $\phi(P), \phi(Q)$ of a public basis of $E[N]$; from published torsion images of a spanning set the action matrices of Thm. 4.4 are read off directly, with no oracle, so (H2) holds* unconditionally *as data availability, with the Part I moment channel a separate side-channel. This is a data-availability statement,* not *a new attack (SIDH was broken by Castryck–Decru 2022 using exactly this torsion data);* `sidh_channel.py` *reconstructs the Frobenius-order Gram from a published transcript.*

6 Computational realization

All quantities are computed with double precision plus exact integer arithmetic; the companion code (23 self-contained Python+NumPy modules) implements finite fields, Tonelli–Shanks, Vélu, LLL, Fincke–Pohst, Miller's algorithm, Pohlig–Hellman, Cornacchia, and quaternion ideal arithmetic from scratch, and a test runner checks every invariant.

Part I ($p = 23$). p-adic Prony over $W(\mathbb{F}_{23}) = \mathbb{Z}_{23}$ recovers the supersingular nodes $\{0, 3, 19\}$ from noisy moments (exact for noise valuation $t \geq 1$). The closed-form $v_p(\det \mathbf{H}_L^{(0)}) = 2 \sum_{i<j} v_p(x_i - x_j) +$

$\sum_i v_p(c_i)$ is verified on five configurations; a pair of depth $v = 0, \ldots, 4$ gives $v_p(\det) = 2v$; with noise level $t = 2v + 3$ the coefficient error has valuation $3 = t - 2v$, i.e. exactly $2v$ digits lost (Thm. 3.7). Tables 1 and 2 record the two phenomena.

L	real $\mathrm{cond}_2(\mathbf{V})$	$\mathrm{cond}_p(\mathbf{V})$
2	6.85	1
5	$2.62 \cdot 10^4$	1
8	$9.52 \cdot 10^8$	1

Table 1: Real vs. p-adic Vandermonde conditioning on nodes $1, \ldots, L$: the real condition number blows up with L, while the p-adic (Teichmüller) one is identically 1 (Prop. 3.4); reproduced in `superres.py`.

depth v	$x_2 - x_1$	$v_p(x_2 - x_1)$	$v_p(\det \mathbf{H}^{(0)})$	digits lost
0	1	0	0	0
1	23	1	2	2
2	529	2	4	4
3	12167	3	6	6
4	279841	4	8	8

Table 2: Exact digit-loss law ($p = 23$): a node pair agreeing to p-adic depth v costs exactly $2v$ digits, matching $v_p(\det \mathbf{H}^{(0)})$ (Thm. 3.7); verified against noise $t = 2v + 3$ giving error valuation $t - 2v = 3$.

Part II ($p = 23$), real endomorphisms. For $E\colon y^2 = x^3 + x$ ($j = 1728 \equiv 3$): $\#E(\mathbb{F}_{23}) = 24$, $\#E(\mathbb{F}_{23^2}) = 576$; the maps $[i], \phi_p$ satisfy $[i]^2 = [-1]$, $\phi_p^2 = [-23]$, $\phi_p[i] = -[i]\phi_p$ on all 576 points. The Gram matrix of the Frobenius order and its saturation: $\mathbf{G}_\Lambda = \mathrm{diag}(2, 2, 46, 46)$, $\det = 16p^2 \to 4p^2 \to p^2 = 529$; the shortest non-scalar vector has $\mathrm{Nrd} = 1, \mathrm{Trd} = 0$, so the node is $j = 1728$.

Real Vélu chains; degree by kernel count. A non-backtracking 2-isogeny cycle $1728 \to 19 \to 19 \to 1728$ gives an endomorphism θ with $\#\ker\theta = 8 = 2^3$ (cyclic $\mathbb{Z}/8$), $\mathrm{Trd}\,\theta = 0, \mathrm{Nrd}\,\theta = 8$. For $p = 31$ the measured degree $\#\ker\theta_K = 2^K$ for $K = 1, 3, 4, 5$ (kernel over $E[2^K]$; $K = 2$ is absent as j=1728 has no non-backtracking length-2 cycle).

Autonomous run. For $p = 23$: Prony $\to \{0, 3, 19\} \to$ curve $j = 1728 \to$ real Vélu θ (deg 8) $\to$ order $\langle 1, [i], \theta, [i]\theta \rangle$ read from $E[44] \subset E(\mathbb{F}_{23^4})$ ($N = 44 > 4\sqrt{D = 64}$, Tonelli–Shanks) with $\det \mathbf{G} = 529 = p^2$ (already maximal) $\to$ node $j = 1728$; no oracle, no a priori degrees.

Large traces in $\mathbb{F}_{p^4}$. Via Tonelli–Shanks and $N > 4p$ the full Gram (including $2\,\mathrm{Nrd} = 2p$) is read from the action: $p = 23, N = 132$ (diagonal $2, 2, 46, 46$) and $p = 47, N = 276$ (diagonal $2, 2, 94, 94$), both saturating to p^2 and giving $j = 1728$.

Reading the action without an $O(N)$ table. Weil pairing (Miller [9]) + Pohlig–Hellman [10] solve the 2-dimensional discrete log in $\mathrm{poly}(\log N)$: cross-checked against the brute table at $p = 23, N = 11$ (identical matrices), then run at $p = 1019, N = 255$ (a 65025-entry table avoided).

True LLL + Fincke–Pohst; scaling. LLL on the Gram (Cholesky, unimodular tracking) + exact Fincke–Pohst enumeration classifies nodes for $(-1, -p) \to j$=1728 and $(-3, -p) \to j$=0 up to $p \sim 500$ in seconds. All three classes of $B_{23,\infty}$ are built via ideal neighbours (right orders, lattice intersection) and classified simultaneously: $|\mathrm{Aut}| = (2, 4, 6) \leftrightarrow (j$=$19, 1728, 0)$, Eichler mass $\frac{1}{6} + \frac{1}{4} + \frac{1}{2} = \frac{11}{12} = \frac{p-1}{24}$. Each node is also identified by its own real Vélu chain.

KLPT: norm equation and ideal-to-isogeny. In $\mathbb{Z}\langle 1,i,j,k\rangle \subset \mathrm{End}(E_{1728})$ the norm form is $a^2+b^2+p(c^2+d^2)$; the equation $= \ell^e$ is solved by making the residual a prime $\equiv 1 \pmod 4$ (Cornacchia). Eight targets up to 2^{20} are realized exactly (e.g. $2^{20} = 1007 + 182i + 5j + 6k$), and the norm-equation core scales to $p \sim 10^{12}$ with targets up to 2^{80} in ~ 1 s. The ideal-to-isogeny step is realized: for $p = 31$, $\alpha = i + k$ (norm 2^5) has cyclic kernel $\mathbb{Z}/32$ (a genuine cyclic 2^5-isogeny), whereas $\alpha = 1 + k$ kills $E[2]$ (non-primitive ideal); *primitivity is checked, not assumed.* The KLPT *smoothing* step is also realized (`klpt_full.py`): a left ideal of norm 3 is replaced by an equivalent ideal of smooth norm $2^4 = 16$ (find $\gamma \in I$ with $\mathrm{Nrd}(\gamma) = 3 \cdot 2^4$, set $J = I\bar{\gamma}/\mathrm{Nrd}(I)$; verified $J \subset O$, left ideal, $\mathrm{Nrd}(J) = 16$, $J \sim I$). *Correctness is exact and unconditional.* The one remaining heuristic — polynomial running *time* — is isolated as the explicit hypothesis below and discharged by a conditional theorem.

Assumption 6.1 (PRH(κ): prime residuals of the norm form)**.** For $p \equiv 3 \pmod 4$ and a target $T = \ell^e \geq p^{1+\kappa}$, the number of $(c,d) \in \mathbb{Z}^2$ with $0 < p(c^2+d^2) < T$ and $R = T - p(c^2+d^2)$ a prime $\equiv 1 \pmod 4$ is at least $(T/p)/(\kappa \log T)$.

Theorem 6.2 (GRH-type conditional polynomial time for the KLPT norm equation)**.** *Under Assumption 6.1, the norm equation $a^2+b^2+p(c^2+d^2) = \ell^e$ and the subsequent smoothing step (§above) are solved in expected time $O(\kappa \log \ell^e \cdot \mathrm{poly}(\log p))$; hence, given an effective input, the KLPT construction of an endomorphism of prescribed smooth degree runs in probabilistic polynomial time.* PRH(κ) *is a consequence of GRH for the relevant Dirichlet L-functions (Lemma 6.3)* for fixed p *(or $p \leq (\log \ell^e)^A$). The uniform-in-p case needed at cryptographic scale — where p grows with ℓ^e —* is open*: the pointwise GRH bound overshoots the main term by a factor $\sqrt{p}\,\mathrm{polylog}$, so it requires a well-factorable large-sieve estimate on the character-weighted error (supplement, §2). We do not claim cryptographic-scale efficiency unconditionally.*

Proof. Sample (c,d) uniformly in the admissible box $p(c^2+d^2) < T$. By Assumption 6.1 a uniform draw yields a residual $R = T - p(c^2+d^2)$ that is prime $\equiv 1 \pmod 4$ with probability $\geq 1/(\kappa \log T)$, so after an expected $O(\kappa \log T)$ draws one such R is found. A prime $\equiv 1 \pmod 4$ is a sum of two squares (Fermat), and $R = a^2 + b^2$ is produced by one Cornacchia call in $\mathrm{poly}(\log R) = \mathrm{poly}(\log p, e \log \ell)$ time; then (a,b,c,d) solves the norm equation. Smoothing (replace a left ideal I by an equivalent smooth-norm ideal $J = I\bar{\gamma}/\mathrm{Nrd}(I)$) reduces to the same prime-residual search for γ, inheriting the bound. All steps are classical; correctness was established unconditionally above, so only the running time is conditional. □

The hypothesis is not ad hoc. Numerically (Table 3) the prime-residual count is a bounded multiple ($\approx 1.9\times$) of the GRH prediction $(\pi/8)(T/p)/\ln T$ across p up to $\sim 10^3$, and the per-trial hit probability tracks $1/\ln T$ up to $p \sim 10^5$, $T \sim 2^{100}$. More structurally, via $r_2(m) = 4\sum_{d|m} \chi_4(d)$ the counting function satisfies the *exact* identity $N(T) = 4\sum_d \chi_4(d)\,\pi(T; 4pd, a_d)$, a Titchmarsh-type shifted-prime divisor sum (verified to the digit, Table 4), whose main term is $\mathfrak{S}(p)(T/p)/\log T$ with $\mathfrak{S}(p) > 0$ explicit; so PRH reduces to a classical analytic statement provable under GRH with character cancellation, in the spirit of Bombieri–Vinogradov / Bombieri–Friedlander–Iwaniec.

Lemma 6.3 (main term and singular series; GRH-conditional asymptotic for fixed p)**.** *Let $N(T) = \#\{(c,d) \in \mathbb{Z}^2 : R = T - p(c^2+d^2)$ prime $\equiv 1 \pmod 4\}$ and*

$$\mathfrak{S}(p) = \beta_2 \cdot \frac{2}{p-1} \prod_{\substack{r \text{ odd prime} \\ r \neq p}} \Big(1 + \frac{\chi_4(r)\, r}{(r-1)(r-\chi_4(r))}\Big), \qquad \beta_2 = 2.$$

Then $\mathfrak{S}(p)$ converges to a positive constant with $\mathfrak{S}(p) = \Theta(1/p)$; the main term of the exact identity $N(T) = 4\sum_d \chi_4(d)\pi(T; 4pd, a_d)$ is $\mathfrak{S}(p)\,\mathrm{li}(T)$; and, under GRH for the Dirichlet L-functions to

p	e	#prime resid.	$(\pi/8)\frac{T/p}{\ln T}$	ratio
23	20	2548	1291.4	1.97
23	24	33104	17219.3	1.92
47	24	15722	8426.5	1.87
103	24	7036	3845.1	1.83
199	24	3686	1990.2	1.85
1019	26	2650	1435.0	1.85

Table 3: Prime-residual density vs. the GRH heuristic ($T = 2^e$): the ratio is a bounded constant, confirming Assumption 6.1. Counts here are over the symmetry-reduced first quadrant; the full-plane $N(T)$ of Tables 4, 5 is $\approx 4\times$ larger (the $r_2 = 4(1 * \chi_4)$ factor).

p	T	A	B
23	2^{16}	784	784
47	2^{16}	400	400
103	2^{18}	688	688

Table 4: The Titchmarsh identity $N(T) = A = B = 4\sum_d \chi_4(d)\pi(T; 4pd, a_d)$ holds exactly (A = full-plane count, B = divisor-switched sum).

p	e	$\mathfrak{S}(p) = N(T)/\mathrm{li}(T)$	$\mathfrak{S}(p) \cdot p$	Euler product $4\,S_{\mathrm{odd}}\,p/(p-1)$
23	18	0.12069	2.776	2.924
47	18	0.06347	2.983	2.792
103	16	0.03038	3.129	2.727

Table 5: Singular series $\mathfrak{S}(p) = \Theta(1/p)$ is stable in T; $\mathfrak{S}(p)\,p \in [2.7, 3.2]$ agrees with the Euler product of Lemma 6.3 to within finite-T fluctuation ($\lesssim 15\%$).

moduli $4pd$ *together with the character-weighted large sieve, for fixed* p *one has* $N(T) = \mathfrak{S}(p)\,\mathrm{li}(T)(1 + o(1))$ *as* $T \to \infty$*, whence* PRH(κ) *holds for every fixed* p*. Numerically (`klpt_grh.py`)* $\mathfrak{S}(p)$ *is stable in* T *and* $\mathfrak{S}(p)\,p \in [2.7, 3.2]$*, agreeing with the Euler product up to finite-*T *fluctuation. Only the uniformity of the error in* p *remains — a named analytic task, not a heuristic.*

Proof. The local factor at r is $1 + \sum_{k\geq 1} \chi_4(r)^k/\varphi(r^k) = 1 + \chi_4(r)r/((r-1)(r-\chi_4(r)))$, > 1 for $r \equiv 1$ and in $(0, 1)$ for $r \equiv 3 \pmod 4$; the product converges by Dirichlet (it is an $L(1, \chi_4)$-type product), is positive, and the factor $2/(p-1)$ gives order $\Theta(1/p)$. Summing $\mathrm{li}(T)/\varphi(4pd)$ against $\chi_4(d)$, using $\varphi(4pd) = 2(p-1)\varphi(d)$ for odd d coprime to p and folding in the separate 2-adic density factor $\beta_2 = 2$ (from the residue condition $q \equiv 1 \bmod 4$), yields $\mathfrak{S}(p)\mathrm{li}(T)$. Truncating divisors at $\sqrt{T/p}$ (Dirichlet hyperbola) puts the moduli at $\leq 4\sqrt{pT}$; the uniform error bound there is the GRH+large-sieve step we cite (Rem. below), not carried out. □

7 Discussion, limitations, and honest assessment

We separate rigorous theorems from heuristic components.

- **Two data channels.** Theorem 5.1 assumes both channels; Prop. 5.2 clarifies that Part II alone identifies a single node, while Part I addresses multi-node separation from aggregate moments.
- **Effective endomorphisms (H2).** Reduced to a checkable input (B-smooth cycles; deterministic KLPT core), but a cyclic ℓ^e-isogeny additionally requires ideal primitivity (checked, §6); the full KLPT smoothing/class-reduction is not reproduced.
- **Scale.** Number-theoretic core to $p \sim 10^{12}$; pairing route to $p \sim 10^3$. Examples illustrate, not replace, asymptotic constants.
- **Noise model.** Deterministic recovery is model-(B); the archimedean threshold is exponentially small in L, removed only p-adically at the cost of a p-adic noise model.

- **Referee items.** An expert review (companion `referee_report.md`) flags: the Eichler-only $2\Omega(m)$ headline (addressed, Rem. 4.3); the Eichler normalization at special nodes (Thm. 4.9); traces-vs-norms precision (Thm. 4.4); ideal primitivity (§6). Verdict: accept with revisions.
- **Status.** A working preprint, not peer reviewed.

8 Conclusion

We connected two classical machines — singular-value perturbation of Hankel matrices and lattice reduction in a quaternion order — into a single pipeline from noisy moments of j-invariants to the type of a supersingular isogeny-graph node. Part I gives an explicit *sufficient* noise bound and a non-archimedean reformulation with perfect conditioning and an *exact* super-resolution *law* (the determinant identity); Part II reduces $\mathrm{End}(E)$ to a rank-4 CVP, recovers the Gram matrix quantumly in $\mathrm{poly}(\log p)$, and reads the node type from the shortest non-scalar endomorphism. Every key step is computed on real numbers and reproducible. Open directions: rigorous super-resolution over $\mathbb{Q}_p$ with a Moitra-style measurement count; full KLPT (smoothing and class reduction) in place of heuristic walks; removing the two-channel assumption; and asymptotic analysis at cryptographic p.

References

[1] F. L. Bauer, C. T. Fike. *Norms and exclusion theorems.* Numer. Math. **2** (1960), 137–141.

[2] B. Beckermann. *The condition number of real Vandermonde, Krylov and positive definite Hankel matrices.* Numer. Math. **85** (2000), 553–577.

[3] L. Colò, D. Kohel. *Orienting supersingular isogeny graphs.* J. Math. Cryptol. **14** (2020), 414–437.

[4] M. Deuring. *Die Typen der Multiplikatorenringe elliptischer Funktionenkörper.* Abh. Math. Sem. Hamburg **14** (1941), 197–272.

[5] K. Eisenträger, S. Hallgren, K. Lauter, T. Morrison, C. Petit. *Supersingular isogeny graphs and endomorphism rings.* EUROCRYPT 2018, LNCS **10822**, 329–368.

[6] W. Gautschi. *How (un)stable are Vandermonde systems?* Lecture Notes in Pure Appl. Math. **124**, Dekker (1990), 193–210.

[7] G. R. de Prony. *Essai expérimental et analytique...* J. École Polytechnique **1** (1795), 24–76.

[8] R. Roy, T. Kailath. *ESPRIT—estimation of signal parameters via rotational invariance techniques.* IEEE Trans. ASSP **37** (1989), 984–995.

[9] V. S. Miller. *The Weil pairing, and its efficient calculation.* J. Cryptology **17** (2004), 235–261.

[10] S. Pohlig, M. Hellman. *An improved algorithm for computing logarithms over* $\mathrm{GF}(p)$*...* IEEE Trans. Inf. Theory **24** (1978), 106–110.

[11] B. H. Gross. *Heights and the special values of L-series.* CMS Conf. Proc. **7** (1987), 115–187.

[12] R. Kannan. *Minkowski's convex body theorem and integer programming.* Math. Oper. Res. **12** (1987), 415–440.

[13] A. K. Lenstra, H. W. Lenstra, L. Lovász. *Factoring polynomials with rational coefficients.* Math. Ann. **261** (1982), 515–534.

[14] D. Micciancio, P. Voulgaris. *A deterministic single exponential time algorithm for most lattice problems.* SIAM J. Comput. **42** (2013), 1364–1391.

[15] A. Moitra. *Super-resolution, extremal functions and the condition number of Vandermonde matrices.* STOC 2015, 821–830.

[16] H. Onuki. *On oriented supersingular elliptic curves.* Finite Fields Appl. **69** (2021), 101777.

[17] P. W. Shor. *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer.* SIAM J. Comput. **26** (1997), 1484–1509.

[18] J. Voight. *Quaternion Algebras.* GTM **288**, Springer, 2021.

[19] P.-Å. Wedin. *Perturbation bounds in connection with singular value decomposition.* BIT **12** (1972), 99–111.

[20] B. Wesolowski. *The supersingular isogeny path and endomorphism ring problems are equivalent.* FOCS 2021, 1100–1111.

[21] H. Weyl. *Das asymptotische Verteilungsgesetz der Eigenwerte.* Math. Ann. **71** (1912), 441–479.