Solving the supersingular isogeny problem in time $p^{2/5+o(1)}$ using bivariate multipoint evaluation*

Aleksei Udovenko

SnT, University of Luxembourg
aleksei@affine.group

August 1, 2026

Abstract. This note presents a new unconditional attack on the supersingular isogeny problem, with time and memory complexity $p^{2/5+o(1)}$. It builds on the approach by Eisenträger-Hallgren-Leonardi-Morrison-Park (2020) and Fuselier-Iezzi-Kozek-Morrison-Namoijam (2025), and is related to the recent heuristic attack with complexity $p^{1/3+o(1)}$ by Wesolowski (ePrint 2026/1486): all of these search for a separable isogeny from a curve to its Galois conjugate to form a non-scalar endomorphism.
Our attack is based on highly theoretical multivariate multipoint evaluation algorithms from Kedlaya-Umans (2008, 2011), Bhargava-Ghosh-Guo-Kumar-Umans (2022), and Ghosh-Harsha-Herdade-Kumar-Saptharishi (2023), and therefore does not threaten isogeny cryptosystems in practice; it is of theoretical interest.

Keywords: Isogenies · Elliptic curves · OneEnd · Cryptanalysis

1 Introduction

In this note, we present a new attack on the supersingular isogeny problem, with expected time complexity $p^{2/5+o(1)}$. Its high-level approach closely follows [EHL+20] and its follow-up [FIK+25] (see also [CS22]). Both algorithms perform a walk in the supersingular isogeny graph aiming to find a curve with a cyclic d-isogeny to its Galois conjugate, for a (small) chosen degree d. Repeating this a small number of times is sufficient to solve the problem, in various ways.

Our approach is to choose a larger degree $d = \tilde{\Theta}(p^{1/5})$ and use fast multivariate multipoint evaluation (MME, Kedlaya-Umans [KU08, KU11] and recent improvements) for testing the curves. Since MME is a highly theoretical algorithm which in particular assumes negligible memory access cost, our attack in its current version is unlikely to threaten isogeny-based cryptosystems in practice and is mainly of theoretical interest.

The idea to use a larger degree is inspired by the recent result by Aubry-Oyono-Vincent [AOV26] and the subsequent attack on the supersingular isogeny problem by Wesolowski [Wes26]. The former work showed that *every* supersingular elliptic curve has a relatively low-degree separable isogeny to its conjugate (degree at most $(p/2)^{1/3}$). The latter work suggests finding such an isogeny by meet-in-the-middle enumerating all degrees up to some bound, hoping that the degree is smooth, which results in an attack with complexity $p^{1/3+o(1)}$ based on more practical components than our attack.

Due to Page-Wesolowski [PW24] and Herlédan Le Merdy-Wesolowski [MW25], the supersingular isogeny problem is equivalent to the full endomorphism problem, which in

*This work is funded by the Luxembourg National Research Fund's (FNR) project PQSeal (CORE/24/IS/18978392).

turn reduces to the one endomorphism problem (OneEnd), under probabilistic polynomial time reductions. We also remark that [FIK+25] studied in detail how to reconstruct the full endomorphism ring from such "inseparable reflections". We focus on the OneEnd problem, and propose the following unconditional result:

Theorem 1. *There exists a probabilistic algorithm that, given as input a supersingular elliptic curve $E/\mathbb{F}_{p^2}$, outputs a non-scalar endomorphism of E in efficient representation with expected time complexity $p^{2/5+o(1)}$, and the same memory complexity.*

The main idea of the attack is simple, here is a sketch. Sample $M = p^{2/5+o(1)}$ random walks from the starting curve E and denote the destinations by E_i. Choose $N = \Theta(p^{1/5})$ and compute the classical modular polynomial $\Phi_N(X, Y)$ modulo p. Then, use the bivariate multipoint evaluation algorithm to evaluate $\Phi_N(j(E_i), j(E_i)^{(p)})$ for all i. Applying MME after conversion to $\mathbb{F}_p$ yields complexity $p^{2/5+o(1)}$. We expect $\widetilde{\Theta}(\sqrt{Np}) = p^{3/5-o(1)}$ supersingular j-invariants N-isogenous to their conjugates. Therefore, with M samples we expect to hit one such E_i, and obtain a non-scalar endomorphism of E, given by the concatenation of: the walk from E to E_i, the discovered N-isogeny, the p-Frobenius $E_i^{(p)}$ to E_i, the dual walk from E_i to E. This note resolves the technical details and proves Theorem 1.

Conventions We assume a word RAM model with word size $\Theta(\log p)$; basic word/field operations and memory accesses have unit cost. This idealized model is unrealistic in practice but is appropriate for the theoretical setting of this work.

2 (Extended) Preliminaries

Define the Dedekind psi function $\Psi(d) = d \cdot \prod_{\text{prime } \ell \mid d} \left(1 + \frac{1}{\ell}\right) = O(d \log\log d)$.

2.1 Class numbers

Theorem 2 (Siegel-Tatuzawa [Sie35], [Tat51, Thm.1], see also [RL07])**.** *For any $\varepsilon > 0$ and for any imaginary quadratic number field K with absolute discriminant D_K, its class number $h_K = h(D_K)$ satisfies*

$$h_K \geq \frac{\varepsilon}{10\pi} D_K^{1/2-\varepsilon}$$

with at most one exceptional field K.

Corollary 1. *For a parameter p, set $\varepsilon = 1/\log_2 p$, so that*

$$h_K \geq \frac{\sqrt{D_K}}{10\pi \log_2 p \cdot 2^{\log_p D_K}} \geq \frac{\sqrt{D_K}}{40\pi \log_2 p} = \Omega\left(\frac{\sqrt{D_K}}{\log p}\right)$$

for all $D_K < p^2$ with at most one exceptional K.

2.2 Counting curves with separable isogenies to conjugates

Strong lower bounds on $\mathcal{O}$-orientable supersingular curves (for some quadratic order $\mathcal{O}$) were given by Leroux in [Ler24]. However, some of those results disallow p dividing the order's discriminant. Here, we propose an unconditional, slightly weak bound specific to orders of $\mathbb{Q}(\sqrt{-pd})$ for a squarefree $d = p^{1/3-o(1)}$. They are obtained by combining Chenu-Smith results [CS22] with a resultant-based multiplicity counting inspired by [dHKM+25]. We note that [EHL+20, Thm.3.9] has a closely related bound for a *constant* degree d; the proof actually reveals that the dependence on d is incompatible with our purposes.

Theorem 3 ([CS22, Cor.1]). *Let $p > 3$ be prime and let d be a squarefree integer prime to p. A pair of an elliptic curve $E/\mathbb{F}_{p^2}$ and a cyclic d-isogeny $\psi : E \to E^{(p)}$ is called a (d, ϵ)-structure if $\widehat{\psi} = \epsilon\psi^{(p)}$, $\epsilon = \pm 1$. The number of supersingular (d, ϵ)-structures (for fixed d, ϵ) over $\mathbb{F}_{p^2}$ up to $\mathbb{F}_{p^2}$-isomorphism is given by*

$$\alpha_d = \begin{cases} 2h_K & \text{if } -dp \equiv 1 \pmod 8, \\ 4h_K & \text{if } -dp \equiv 5 \pmod 8, \\ h_K & \text{otherwise.} \end{cases}$$

where $K = Q(\sqrt{-dp})$.

We note that a similar statement for a prime d can be found in [Arp24, Prop.5.3].

Corollary 2. *Let $\varepsilon > 0$, let $p > 3$ be prime, and let $d = O(p^{1-\varepsilon})$ be squarefree. The total number of isomorphism classes of pairs (E, G) with $E/\mathbb{F}_{p^2}$ a supersingular elliptic curve and $G \subseteq E[d]$ a cyclic subgroup of order d such that $E/G \simeq E^{(p)}$ is at least*

$$\frac{\alpha_d - 10\Psi(d)}{2} = \Omega\left(\sqrt{dp}/\log(dp)\right),$$

with at most one exceptional d (dividing by two corresponds to forgetting the sign of ψ).

Proof. Note that [CS22, Prop.2] shows that the trace of Frobenius of a supersingular (d, ϵ)-structure is $-2\epsilon p$, thus only two twists can support such structures, and ϵ fixes the choice. Therefore, we do not have to worry about the curve model.

The isomorphism definition of a (d, ϵ)-structure constrains the two curve isomorphisms to be conjugate. Thus, two (d, ϵ)-structures may be distinguished while differing only by an automorphism of the codomain, which the representation (E, G) does not distinguish. Every curve has the two automorphisms ± 1, while the one with $j = 0$ has six and the one with $j = 1728$ has four (when $p > 3$). These two curves are together involved in at most $(4 + 6)\Psi(d)$ structures, which we remove from counting. Each other pair (E, G), $\psi : E \to E/G$, is counted at most twice as $\pm\psi$, thus we halve the remaining count. Note that

$$\frac{\alpha_d}{\Psi(d)} \geq \frac{\left(\sqrt{pd}^{1+o(1)}\right)}{d^{1+o(1)}} = p^{\varepsilon/2-o(1)} \to \infty,$$

and thus α_d dominates. □

In the following, the number of edges between $j(E)$ and $j(E')$ in the isogeny graph is defined as the multiplicity of $Y = j(E')$ as a root of the univariate polynomial $\Phi_d|_{X=j(E)}(Y)$ (see e.g. [Gal12, §25.2]); this multiplicity is denoted by $\nu(j(E), j(E'))$. It is equal to the number of cyclic subgroups G of order d on (any model of) E such that $j(E/G) = j(E')$. Note that it distinguishes (counts) distinct subgroups related by an automorphism of E, while a (d, ϵ)-structure may not distinguish them.

Proposition 1. *Let $d > 1$ be an integer, and let $p > \Psi(d)$ be prime. It is possible to remove $O(\Psi(d)^2)$ edges from the supersingular d-isogeny graph so that the remaining graph has no loops or multiple edges.*

Proof. Let $a \in \overline{\mathbb{F}}_p$. A repeated root $Y = y$ of $\Phi_d(a, Y) = 0$ of multiplicity $m_y = \nu(a, y)$ must also satisfy $\partial_Y\Phi_d(a, Y) = 0$ with multiplicity $m_y - 1$. Let

$$R(X) = \mathrm{Res}_Y(\Phi_d(X, Y), \partial_Y\Phi_d(X, Y))$$

and let $S_Y(X) \in \mathbb{F}_p[X]^{D\times D}$ be the corresponding Sylvester matrix ($D = 2\Psi(d) - 1$), such that $R(X) = \det S_Y(X)$. Note that $\Phi_d(X, Y)$ has good reduction modulo $p \nmid d$ and serves

as a singular model of the irreducible algebraic curve $X_0(d)$ (see e.g. [Mil17, Thm.8.6]); thus it is irreducible over $\overline{\mathbb{F}}_p$. It is monic in Y, so it is primitive; by Gauss' lemma $\Phi_d(X,Y)$ is irreducible in $\overline{\mathbb{F}}_p(X)[Y]$ and thus the resultant is nonzero (note that $\partial_Y\Phi_d(X,Y)$ has degree exactly $\Psi(d)-1$ since the degree of Φ_d in Y is less than the characteristic p).

Under specialization to $X=a$, the two arguments share a common root y and the resultant $R(a)$ is equal to zero. The nullity of the corresponding Sylvester matrix $S_Y(a)\in\overline{\mathbb{F}}_p^{D\times D}$ is equal to the degree of the GCD of $\Phi_d(a,Y)$ and $\partial_Y\Phi_d(a,Y)$ which is equal to the sum v_a of $\nu(a,b)-1$ for all b with $\Phi_d(a,b)=0$. Let $A\in\mathrm{GL}_D(\overline{\mathbb{F}}_p)$ be such that the bottom v_a rows of $A\cdot S_Y(a)$ are equal to zero. Then, the bottom v_a rows of $A\cdot S_Y(X)$ will have every entry divisible by $(X-a)$. Therefore,

$$R(X)=\det S_Y(X)=\det A^{-1}\cdot\det\left(A\cdot S_Y(X)\right)$$

is divisible by $(X-a)^{v_a}$. Collecting these factors over all $a\in\overline{\mathbb{F}}_p$ (note that they are pairwise coprime), we conclude that the degree of $R(X)$ upper bounds the sum of $\nu(a,b)-1$ over all zeros (a,b) of Φ_d, which is the number of redundant multiple edges:

$$\sum_{a,b\in\overline{\mathbb{F}}_p:\,\nu(a,b)>1}(\nu(a,b)-1)\le\deg R.$$

The standard resultant degree bound is

$$\begin{aligned}\deg_X(R(X))&\le\deg_Y\partial_Y\Phi_d\cdot\deg_X\Phi_d+\deg_X\partial_Y\Phi_d\cdot\deg_Y\Phi_d\\&\le(\Psi(d)-1)\Psi(d)+\Psi(d)\Psi(d)\\&=2\Psi(d)^2-\Psi(d).\end{aligned}$$

After repeated edges are removed, there may remain at most one loop for each j-invariant. Such j-invariants Y are roots of $\Phi_d(Y,Y)=0$ (which is nonzero since $X-Y$ does not divide $\Phi_d(X,Y)$) and so their number is upper bounded by $2\Psi(d)$. We conclude that removing $2\Psi(d)^2+\Psi(d)$ edges is sufficient. □

For prime d, one could also use [BS25, Corollary 16] which shows that triple d-isogeny edges away from $j=0$ are possible only when $p<4d^3$, or a weaker bound $p<4d^4$ from the already published [dHKM+25, Remark 4] (which is an interpretation of Kaneko's result [Kan89]). Another option is to use [Gha24, Cor.3.3] which bounds the number of redundant edges by $d^2+d/4$ (again, only for prime d), improving the bound from Proposition 1.

Definition 1. Let p be a prime, and $d\ge1$ an integer. Define

$$S_d=\left\{j(E)\mid\exists\,\text{supersingular curve } E \text{ with a cyclic } \varphi:E\to E^{(p)} \text{ of degree } d\right\}.$$

Corollary 3. *Fix $\varepsilon>0$, let $p>3$ be prime and let $d=O(p^{1/3-\varepsilon}),d>1$, be a positive squarefree integer. Then, the cardinality of S_d is in $\Omega(\sqrt{pd}/\log(pd))$, with at most one exceptional value of d.*

Proof. Combining Corollary 2 with Corollary 1, we get that the number of d-isogenies to conjugates is at least $\Omega(\sqrt{pd}/\log(pd))$, with at most one exceptional value of d. This quantifies the number of cyclic d-isogenies connecting conjugate curves. By Proposition 1, we can remove $O(d^{2+o(1)})$ d-isogenies to obtain a set of

$$\Omega(\sqrt{pd}/\log(pd))-O(d^{2+o(1)})$$

isogenies with unique domain curve (since it determines the codomain by conjugation). The ratio of the two quantities is $p^{\frac{3}{2}\varepsilon-o(1)}/\log(pd)$ which tends to infinity with $p\to\infty$. Therefore, the first term dominates and we are left with $\Omega(\sqrt{pd}/\log(pd))$ unique supersingular curves, d-isogenous to their conjugates. □

2.3 Computing modular polynomials

The following is an unconditional variant of the classic method from [BLS11].

Theorem 4 ([KR24, Thm.1.2]). *There exists an unconditional CRT algorithm which computes the modular polynomial $\Phi_\ell(X,Y) \in \mathbb{Z}[X,Y]$ in quasi-linear time $O(\ell^3 \log^3 \ell \log\log \ell)$.*

Remark 1. Reduction of Φ_ℓ modulo a prime p can be done within the same time complexity.

The authors of [KR24] remark that the extension to composite N is not hard. We propose a simple way to efficiently combine classical modular polynomials of different degrees to achieve this.

Proposition 2. *Let p be a prime, and let a, b be two coprime integers such that $12\Psi(ab) + 11 < p$. Then, given classical modular polynomials $\Phi_a(X,Y), \Phi_b(X,Y) \in \mathbb{F}_p[X,Y]$, and a supersingular elliptic curve $E/\mathbb{F}_{p^2}$, one can compute the classical modular polynomial $\Phi_{ab}(X,Y) \in \mathbb{F}_p[X,Y]$ in expected time $\widetilde{O}((ab)^2 \log p)$.*

Proof. Let $A = \Psi(a), B = \Psi(b), D = \Psi(ab) = O(ab \log\log ab)$. Use a random 2-isogeny walk from E to compute D distinct supersingular j-invariants $j_1, \ldots, j_D$. If $12\Psi(ab)$ is close to p, we can enumerate all supersingular j-invariants by a depth-first search. Since there are at least $\left\lfloor \frac{p}{12} \right\rfloor$ supersingular j-invariants, we will find distinct D of them:

$$\left\lfloor \frac{p}{12} \right\rfloor \geq \frac{p-11}{12} \geq D.$$

Using multipoint evaluation of coefficients of Z^i in $\Phi_a(X,Z)$, compute all $F_i(Z) = \Phi_a(j_i, Z) \in \mathbb{F}_{p^2}[Z]$ in time $\widetilde{O}(DA)$ (a similar optimization was used in [IJ10, Sec.5.1]). All $F_i(Z)$ split into linear terms since all involved j-invariants are supersingular.

Factor each $F_i(Z) = \prod_{j=1}^{A}(Z - z_{i,j})$ (with possible repetitions of roots). This can be done in expected time $\widetilde{O}(A \log p)$ for each F_i, since all F_i split.

Since a is coprime to b, we have

$$\Phi_{ab}(X,Y) = \mathrm{Res}_Z(\Phi_a(X,Z), \Phi_b(Z,Y))$$

and it also holds for any fixed $X = j_i$. Note that

$$\begin{aligned}\Phi_{ab}(j_i,Y) &= \mathrm{Res}_Z\left(\Phi_a(j_i,Z), \Phi_b(Z,Y)\right) \\ &= \mathrm{Res}_Z\left(\prod_{j=1}^{A}(Z - z_{i,j}), \Phi_b(Z,Y)\right) = \prod_{j=1}^{A}\Phi_b(z_{i,j},Y).\end{aligned}$$

For each of the DA points $z_{i,j}$, we compute $G_{i,j}(Y) = \Phi_b(z_{i,j}, Y)$ using multipoint evaluation in total time $\widetilde{O}(DAB) = \widetilde{O}(D^2)$. This allows us to compute $\Phi_{ab}(j_i, Y)$ using the above expression via a product tree (product of A degree-B polynomials, for each of the D points j_i) in time $\widetilde{O}(D^2)$.

Write

$$\Phi_{ab}(X,Y) = X^D + Y^D + \sum_{k=0}^{D-1} g_k(X) Y^k$$

with $g_k \in \mathbb{F}_p[X]$ of degree at most $D-1$. Matching this with the recovered $\Phi_{ab}(j_i, Y)$, we obtain D known input-output pairs $(j_i, g_k(j_i))$ for g_k, and we can interpolate each polynomial g_k in quasi-linear time $\widetilde{O}(D)$, resulting in total expected time $\widetilde{O}(D^2 \log p)$. □

We will use the following instantiation for N-isogenous conjugation testing. A more direct approach for computing $\mathbb{F}_N(x,y)$ is given in Appendix A.

Lemma 1. *Let $p > 2$ be a prime not dividing an integer $N > 1$. Define $\mathbb{F}_{p^2} = \mathbb{F}_p[\beta]/(\beta^2-\delta)$ for a square non-residue $\delta \in \mathbb{F}_p$. Define*

$$F_N(x, y) = \Phi_N(x + \beta y, x - \beta y).$$

Note that $F_N(x, y) \in \mathbb{F}_p[x, y]$, with $\max(\deg_x F_N, \deg_y F_N) \leq 2\Psi(N)$. It can be computed efficiently from $\Phi_N(X, Y) \in \mathbb{F}_p[X, Y]$ in time $\widetilde{O}(\Psi(N)^2 \log p)$.

Proof. First, $F_N(x, y) \in \mathbb{F}_p[x, y]$ because it is stable under the conjugation $\beta \mapsto -\beta$, due to Φ_N being symmetric. Now, split Φ_N into homogeneous components as

$$\Phi_N(X, Y) = \sum_{k=0}^{2\Psi(N)} H_k(X, Y)$$

for homogeneous H_k of degree $k \leq 2\Psi(N)$. Observe that F_N can be expressed as a linear change of variables $X = x + \beta y, Y = x - \beta y$ applied to $\Phi_N(X, Y)$. By [vdHL21, Prop.2.15], each such homogeneous component can be processed in time $\widetilde{O}(k \log p)$. Summing for k from 0 up to $2\Psi(N)$ gives $\widetilde{O}(\Psi(N)^2 \log p)$. □

2.4 Multivariate multipoint evaluation

The following result from [GHH+23a] extends the fundamental work of Kedlaya-Umans [KU08, KU11], building upon ideas from [BGG+22]. Note that earlier works have an extra $\text{poly}(m, d, \log |\mathbb{F}|)$ factor, meaning that the complexity tends to $((d^m+Nm)\cdot\log p)^{1+o(1)}$ only with increasing number of variables. Despite the work's focus on numerical computations over rationals, reals, and complex numbers, it contains the following result for the case of the prime field.

Theorem 5 ([GHH+23b, Thm.3.2])**.** *There is a deterministic algorithm that, when given as input the coefficient vector of an m-variate polynomial f of degree less than d in each variable over a prime field $\mathbb{F}_p$, and N points $a^{(1)}, \ldots, a^{(N)} \in \mathbb{F}_p^m$, outputs $f(a^{(1)}), \ldots, f(a^{(N)})$ in time*

$$((d^m + N \cdot m) \cdot \log p)^{1+o(1)}$$

for all $m \in \mathbb{Z}_{\geq 1}$ and sufficiently large $d \in \mathbb{Z}_{\geq 1}$.

3 Solving the OneEnd problem

We now describe and analyze the detailed algorithm for proving Theorem 1. Its pseudocode is given in Algorithm 1.

Theorem 6. *Algorithm 1 is correct and finishes in expected time $p^{2/5+o(1)}$.*

Proof. The first seven steps are self-explanatory and have clear complexities described in the algorithm. We first show that the condition in step 8 holds with overwhelming probability. By Corollary 3, (note that hitting an exceptional quadratic field $\mathbb{Q}(\sqrt{-pN})$ has negligible probability since there exists at most one)

$$|S_N| = \Omega(\sqrt{pN}/\log(pN)) = \Omega(p^{3/5}/\log(pN)),$$

and every $j(E_i)$ hits S_N with probability $\Omega(p^{-2/5}/\log(pN))$ for an appropriately chosen walk length in $\Theta(\log p)$, so that the total-variation error is in $o(p^{-2/5}/logp)$ (see [PW24,

Algorithm 1 One endomorphism

Inputs: a supersingular elliptic curve $E/\mathbb{F}_{p^2}$, $p > 3$, $\mathbb{F}_{p^2} \simeq \mathbb{F}_p[\beta]/(\beta^2 - \delta)$
Output: a non-scalar endomorphism of E in efficient representation

1: choose primes $\ell_1, \ell_2 = \Theta(p^{1/10})$ at random, $\ell_1 \neq \ell_2$, and let $N = \ell_1\ell_2$ ▷ $N = \Theta(p^{1/5})$
2: compute $\Phi_{\ell_1}, \Phi_{\ell_2} \in \mathbb{F}_p[X, Y]$ using Theorem 4 ▷ $\widetilde{O}(p^{3/10})$
3: compute $\Phi_N(X, Y) \in \mathbb{F}_p[X, Y]$ using Proposition 2 ▷ $\widetilde{O}(p^{2/5})$
4: compute $F_N(x, y) = \Phi_N(x + \beta y, x - \beta y) \in \mathbb{F}_p[x, y]$ using Lemma 1 ▷ $\widetilde{O}(p^{2/5})$
5: let $M = \Theta(N^2 \log^2(pN))$ ▷ $M = \widetilde{\Theta}(p^{2/5})$
6: compute M random 2-isogeny walks $\varphi_i : E \to E_i$, $1 \leq i \leq M$, of length $\Theta(\log p)$ ▷ $\widetilde{O}(p^{2/5})$
7: compute $r_i = F_N(x_i, y_i) \in \mathbb{F}_p$ for all $1 \leq i \leq M$, where $j(E_i) = x_i + \beta y_i$, $x_i, y_i \in \mathbb{F}_p$, using Theorem 5 (MME of degree $2\Psi(N) + 1$ in each variable on M points) ▷ $p^{2/5+o(1)}$
8: **if** $r_i = 0$ for some i **then**
9: find any root j_i' of $\gcd\left(\Phi_{\ell_1}(j(E_i), X), \Phi_{\ell_2}(X, j(E_i^{(p)}))\right)$
10: recover the ℓ_1-isogeny $\psi_1 : E_i \to E_i'$, $j(E_i') = j_i'$, and the ℓ_2-isogeny $\psi_2 : E_i' \to E_i^{(p)}$ using the Elkies algorithm [Gal12, Alg.28] or [MM24] at singular points ▷ $\widetilde{O}(p^{1/5})$
11: compute efficient representation of $\psi = \psi_2 \circ \psi_1$ using [Rob22] ▷ $\widetilde{O}(p^{1/5})$
12: let $\pi : E_i^{(p)} \to E_i$ be the p-Frobenius isogeny
13: **return** $\phi = \widehat{\varphi}_i \circ \pi \circ \psi \circ \varphi_i : E \to E$
14: **else**
15: go to step 1
16: **end if**

Prop.2.7], [BCC+23, Thm.7,11]). Therefore, the probability that all endpoints of the $M = \Theta(p^{2/5} \log^2 pN)$ walks miss S_N is at most $\exp(-\Omega(\log(pN))) = p^{-\Omega(1)}$, i.e., negligible.

Clearly, $r_i = F_N(x_i, y_i) = 0$ if and only if $j(E_i)$ is N-isogenous to its conjugate. By the analysis above, we expect to find and detect such a curve; otherwise, we restart with fresh randomness. The probability analysis may fail due to exceptional N for Corollary 3; restarting solves this issue too.

In steps 9-10, the N-isogeny is factored as a composition of an ℓ_2- and an ℓ_1-isogeny. This can be done by computing the gcd of the two specialized modular polynomials (univariate) and choosing an arbitrary root. Both isogenies can then be recovered using standard techniques, namely:

1. The Elkies algorithm for recovering the degree-$(\ell_n - 1)/2$ kernel polynomial $D(x)$ [Gal12, Alg.28] and finding the normalized isogeny; complexity $O(\ell_n^2)$ field operations.

2. The Mahaney-Morrison algorithm [MM24] applicable in the case $(j(E_i), j(E_i'))$ and/or $(j(E_i'), j(E_i^{(p)}))$ are singular points of the curves $\Phi_{\ell_1} = 0, \Phi_{\ell_2} = 0$ respectively, except for endpoints with j-invariants $0, 1728$ (for which we can skip this part and choose a known endomorphism; if this happens for the intermediate E_i' then we skip it, which removes at most $2\Psi(\ell_1)$ candidates for E_i); complexity $O(m\ell_n^2)$, where m is the multiplicity of the singularity, $m \leq 2$ by [dHKM+25, Remark 4] since $4\ell_n^4 < p$ for a sufficiently large p.

3. The BMSS algorithms [BMSS08] for computing the normalized isogenies in quasi-linear time.

These algorithms require the classical modular polynomials $\Phi_{\ell_n}(X, Y), n = 1, 2$, which we have computed. One may need to compute an isomorphism from the constructed

codomain to $E_i^{(p)}$, which may require a constant-degree field extension. Finally, the separable isogeny $\psi_2 \circ \psi_1$ can be compressed to an efficient representation (needed for the subsequent EndRing reduction), using [Rob22], [ACD+24, Alg.2.1], in polylogarithmic time in the degree $N = O(p^{1/5})$ and a logarithmic number of evaluations of the isogeny on points from extensions of logarithmic degree.

The resulting isogeny ϕ is not a scalar multiplication on E since its inseparable degree is p (and not p^2). This is the same argument as in [Wes26].

The memory complexity is upper bounded by the time complexity, and is not lower than $\Theta(p^{2/5})$ field elements needed to store $\mathbb{F}_N(x, y)$ and the points for MME (which can be done in $\Theta(\log^2 pN)$ batches of N^2 points). □

Remark 2. Choosing a prime N slightly simplifies the algorithm, but optimal algorithms for computing Φ_N modulo p (e.g. [KR24]) are conditional under GRH (Generalized Riemann Hypothesis), which we otherwise avoid. Using a semiprime N with a resultant method solves the problem. However, it still requires computing an efficient representation of the isogeny. A better option is to use smooth, squarefree, and coprime ℓ_1, ℓ_2; this allows us to compute all steps in quasi-linear time, except MME, and to avoid the need for efficient representation. However, it slightly complicates the analysis due to the introduction of the smoothness bound for the choice of ℓ_1, ℓ_2.

Acknowledgements Relevant results from the literature and several proof ideas were suggested by an LLM (OpenAI GPT 5.6 Sol High). The main attack algorithm and bound strategy/ideas are from the author. All of the text is written by the author, with the LLM used for proofreading.

References

[ACD+24] Sarah Arpin, James Clements, Pierrick Dartois, Jonathan Komada Eriksen, Péter Kutas, and Benjamin Wesolowski. Finding orientations of supersingular elliptic curves and quaternion orders. *Designs, Codes and Cryptography*, 92(11):3447–3493, Nov 2024. doi:10.1007/s10623-024-01435-5. 8

[AOV26] Yves Aubry, Roger Oyono, and Christelle Vincent. Minimal degree of an isogeny between a supersingular elliptic curve and its conjugate, 2026. URL: https://arxiv.org/abs/2607.14624, arXiv:2607.14624. 1

[Arp24] Sarah Arpin. Adding Level Structure to Supersingular Elliptic Curve Isogeny Graphs. *Journal de théorie des nombres de Bordeaux*, 36(2):405–443, 2024. doi:10.5802/jtnb.1283. 3

[BBD+24] Jeremy Booher, Ross Bowden, Javad Doliskani, Tako Boris Fouotsa, Steven D. Galbraith, Sabrina Kunzweiler, Simon-Philipp Merz, Christophe Petit, Benjamin Smith, Katherine E. Stange, et al. Failing to hash into supersingular isogeny graphs. *The Computer Journal*, 67(8):2702–2719, 2024. doi:10.1093/comjnl/bxae038. 11

[BCC+23] Andrea Basso, Giulio Codogni, Deirdre Connolly, Luca De Feo, Tako Boris Fouotsa, Guido Maria Lido, Travis Morrison, Lorenz Panny, Sikhar Patranabis, and Benjamin Wesolowski. Supersingular curves you can trust. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part II*, volume 14005 of *LNCS*, pages 405–437. Springer, Cham, April 2023. doi:10.1007/978-3-031-30617-4_14. 7

[BGG+22] Vishwas Bhargava, Sumanta Ghosh, Zeyu Guo, Mrinal Kumar, and Chris Umans. Fast multivariate multipoint evaluation over all finite fields. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 221–232, 2022. doi:10.1109/FOCS54457.2022.00028. 6

[BLS11] Reinier Bröker, Kristin Lauter, and Andrew Sutherland. Modular polynomials via isogeny volcanoes. *Mathematics of Computation*, 81(278):1201–1231, 2011. URL: http://dx.doi.org/10.1090/S0025-5718-2011-02508-1, doi:10.1090/s0025-5718-2011-02508-1. 5

[BMSS08] A. Bostan, F. Morain, B. Salvy, and É. Schost. Fast algorithms for computing isogenies between elliptic curves. *Mathematics of Computation*, 77(263):1755–1778, January 2008. doi:10.1090/s0025-5718-08-02066-8. 7

[BS25] Sebastiano Boscardin and Sebastian A. Spindler. Exploring Kaneko's bound: On multi-edges, loops and the diameter of the supersingular ℓ-isogeny graph. Cryptology ePrint Archive, Paper 2025/1361, 2025. URL: https://eprint.iacr.org/2025/1361. 4

[CK91] David G. Cantor and Erich Kaltofen. On fast multiplication of polynomials over arbitrary algebras. *Acta Informatica*, 28(7):693–701, Jul 1991. doi:10.1007/BF01178683. 11

[CS22] Mathilde Chenu and Benjamin Smith. Higher-degree supersingular group actions. *Transactions on Mathematical Cryptology*, 1(2):85–101, Mar. 2022. URL: https://journals.flvc.org/mathcryptology/article/view/130604. 1, 2, 3

[dHKM+25] Thomas den Hollander, Sören Kleine, Marzio Mula, Daniel Slamanig, and Sebastian A. Spindler. More efficient isogeny proofs of knowledge via canonical modular polynomials. In Yael Tauman Kalai and Seny F. Kamara, editors, *Advances in Cryptology – CRYPTO 2025*, pages 131–166, Cham, 2025. Springer Nature Switzerland. 2, 4, 7

[EHL+20] Kirsten Eisenträger, Sean Hallgren, Chris Leonardi, Travis Morrison, and Jennifer Park. Computing endomorphism rings of supersingular elliptic curves and connections to path-finding in isogeny graphs. *Open Book Series*, 4(1):215–232, 2020. doi:10.2140/obs.2020.4.215. 1, 2

[FIK+25] Jenny Fuselier, Annamaria Iezzi, Mark Kozek, Travis Morrison, and Changningphaabi Namoijam. Computing supersingular endomorphism rings using inseparable endomorphisms. *Journal of Algebra*, 668:145–189, April 2025. doi:10.1016/j.jalgebra.2025.01.012. 1, 2

[Gal12] Steven D. Galbraith. *Mathematics of public key cryptography*. Cambridge University Press, 2012. 3, 7

[Gha24] Wissam Ghantous. Loops, multi-edges and collisions in supersingular isogeny graphs. *Advances in Mathematics of Communications*, 18(4):935–955, 2024. URL: https://www.aimsciences.org/article/id/62a94b882d80b75d2b6d8238, doi:10.3934/amc.2022038. 4

[GHH+23a] Sumanta Ghosh, Prahladh Harsha, Simão Herdade, Mrinal Kumar, and Ramprasad Saptharishi. Fast numerical multivariate multipoint evaluation. In *64th Annual Symposium on Foundations of Computer Science*, pages 1426–1439, Santa Cruz, CA, USA, November 6–9, 2023. IEEE Computer Society Press. doi:10.1109/FOCS57990.2023.00088. 6

[GHH+23b] Sumanta Ghosh, Prahladh Harsha, Simão Herdade, Mrinal Kumar, and Ramprasad Saptharishi. Fast numerical multivariate multipoint evaluation. *CoRR*, abs/2304.01191, 2023. arXiv:2304.01191, doi:10.48550/arXiv.2304.01191. 6

[IJ10] Sorina Ionica and Antoine Joux. Pairing the volcano. In *ANTS*, volume 6197 of *Lecture Notes in Computer Science*, pages 201–218. Springer, 2010. doi:10.1007/978-3-642-14518-6_18. 5

[Kan89] Masanobu Kaneko. Supersingular j-invariants as singular moduli mod p. *Osaka J. Math.*, 26(4):849–855, 1989. 4

[KR24] Sabrina Kunzweiler and Damien Robert. Computing modular polynomials by deformation. *Research in Number Theory*, 11(1):10, Dec 2024. doi:10.1007/s40993-024-00596-5. 5, 8

[KU08] Kiran S. Kedlaya and Christopher Umans. Fast modular composition in any characteristic. In *49th FOCS*, pages 146–155. IEEE Computer Society Press, October 2008. doi:10.1109/FOCS.2008.13. 1, 6

[KU11] Kiran S. Kedlaya and Christopher Umans. Fast polynomial factorization and modular composition. *SIAM Journal on Computing*, 40(6):1767–1802, 2011. doi:10.1137/08073408X. 1, 6

[Ler24] Antonin Leroux. An effective lower bound on the number of orientable supersingular elliptic curves. In Benjamin Smith and Huapeng Wu, editors, *SAC 2022*, volume 13742 of *LNCS*, pages 263–281. Springer, Cham, August 2024. doi:10.1007/978-3-031-58411-4_12. 2

[Mil17] James S. Milne. Modular functions and modular forms (v1.31), 2017. Available at www.jmilne.org/math/. 4

[MM24] William E. Mahaney and Travis Morrison. Computing isogenies at singular points of the modular polynomial, 2024. URL: https://arxiv.org/abs/2402.02038, arXiv:2402.02038. 7

[MW25] Arthur Herlédan Le Merdy and Benjamin Wesolowski. Unconditional foundations for supersingular isogeny-based cryptography. In Benny Applebaum and Huijia (Rachel) Lin, editors, *TCC 2025, Part III*, volume 16270 of *LNCS*, pages 266–297. Springer, Cham, December 2025. doi:10.1007/978-3-032-12296-4_9. 1

[PW24] Aurel Page and Benjamin Wesolowski. The supersingular endomorphism ring and one endomorphism problems are equivalent. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VI*, volume 14656 of *LNCS*, pages 388–417. Springer, Cham, May 2024. doi:10.1007/978-3-031-58751-1_14. 1, 6

[RL07] Stéphane R. Louboutin. Simple proofs of the Siegel-Tatuzawa and Brauer-Siegel theorems. *Colloquium Mathematicae*, 108(2):277–283, 2007. URL: http://eudml.org/doc/283504. 2

[Rob22] Damien Robert. Evaluating isogenies in polylogarithmic time. Cryptology ePrint Archive, Paper 2022/1068, 2022. URL: https://eprint.iacr.org/2022/1068. 7, 8

[Sie35] Carl Ludwig Siegel. Über die Klassenzahl quadratischer Zahlkörper. *Acta Arith.*, 1(1):83–86, 1935. 2

[Tat51] Tikao Tatuzawa. On a theorem of Siegel. *Japanese Journal of Mathematics: Transactions and Abstracts*, 21:163–178, 1951. The Mathematical Society of Japan. 2

[vdHL21] Joris van der Hoeven and Grégoire Lecerf. On the complexity exponent of polynomial system solving. *Foundations of Computational Mathematics*, 21(1):1–57, Feb 2021. doi:10.1007/s10208-020-09453-0. 6

[Wes26] Benjamin Wesolowski. The supersingular isogeny problem in time and memory $p^{1/3+o(1)}$. Cryptology ePrint Archive, Paper 2026/1486, 2026. URL: https://eprint.iacr.org/2026/1486. 1, 8

A Computing F_N directly

This approach is partially inspired by [BBD+24, Sect.4.2]. Split Φ_N into homogeneous components as

$$\Phi_N(X,Y) = \sum_{k=0}^{2\Psi(N)} H_k(X,Y) = \sum_{k=0}^{2\Psi(N)} Y^k \cdot H_k(X/Y,1) = \sum_{k=0}^{2\Psi(N)} Y^k \cdot h_k(X/Y)$$

for homogeneous H_k of degree $k \le 2\Psi(N)$, and univariate $h_k(z) = H_k(z,1)$. We can write

$$F_N(x,y) = \Phi_N(x+\beta y, x-\beta y) = \sum_{k=0}^{2\Psi(N)} (x-\beta y)^k \cdot h_k\left(\frac{x+\beta y}{x-\beta y}\right).$$

Define $z = x/y$, so that

$$F_N(x,y) = \sum_{k=0}^{2\Psi(N)} y^k \cdot (z-\beta)^k \cdot h_k\left(\frac{z+\beta}{z-\beta}\right) = \sum_{k=0}^{2\Psi(N)} y^k \cdot (z-\beta)^k \cdot h_k\left(1+\frac{2\beta}{z-\beta}\right).$$

Note that a term of the form $t_k(z) = (z-\beta)^k \cdot h_k\left(1+\frac{2\beta}{z-\beta}\right)$ can be computed using a sequence of additive shifts $g(x) = f(x+1)$ (cost equivalent to multiplying degree-k polynomials), scaling variables $g(x) = f(c \cdot x), c \in \mathbb{F}_{p^2}$ ($\mathcal{O}(k)$), and the reciprocal transformation $g(x) = x^k \cdot f(1/x)$ (reversal of the list of coefficients, $\mathcal{O}(k)$). Multiplying degree-k polynomials can be done in time $\mathcal{O}(k \log k \log\log k)$ [CK91]. Finally, the computation $y^k t_k(z) = y^k t_k(x/y)$ is implicit (identifying z^i with $x^i y^{k-i}$). Summing the costs over all $2\Psi(N)+1$ homogeneous components yields the result. Note that all the computations are over $\mathbb{F}_{p^2}$, while the result turns out to be over $\mathbb{F}_p$.