THE OPEN BOOK SERIES 4

ANTS XIV
Proceedings of the Fourteenth Algorithmic Number Theory Symposium

Computing endomorphism rings of supersingular elliptic curves and connections to path-finding in isogeny graphs

Kirsten Eisenträger, Sean Hallgren, Chris Leonardi, Travis Morrison, and Jennifer Park

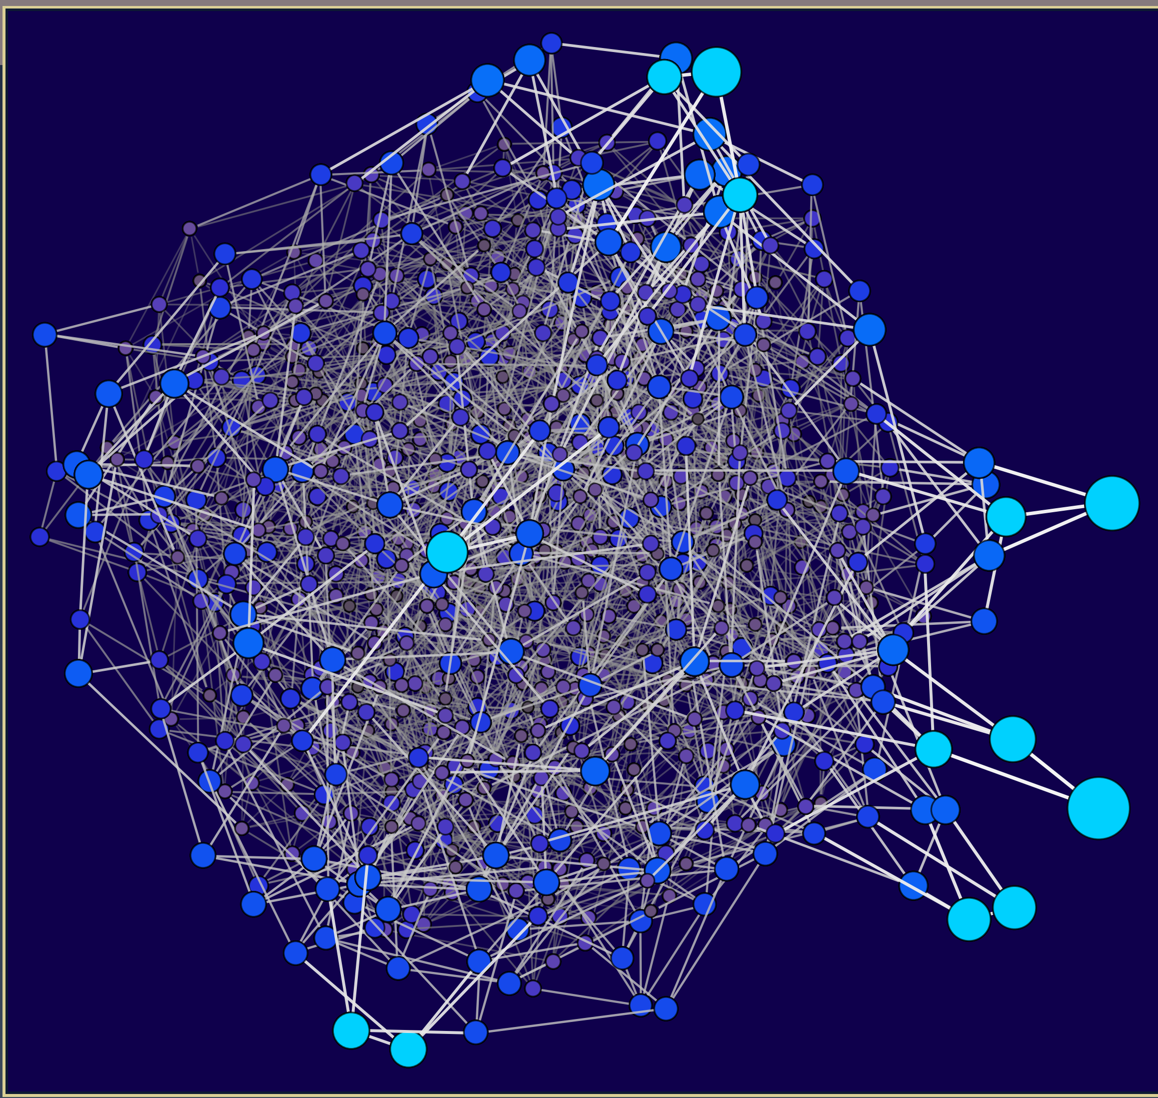

msp

THE OPEN BOOK SERIES 4 (2020)
Fourteenth Algorithmic Number Theory Symposium
https://doi.org/10.2140/obs.2020.4.215

Computing endomorphism rings of supersingular elliptic curves and connections to path-finding in isogeny graphs

Kirsten Eisenträger, Sean Hallgren, Chris Leonardi, Travis Morrison, and Jennifer Park

Computing endomorphism rings of supersingular elliptic curves is an important problem in computational number theory, and it is also closely connected to the security of some of the recently proposed isogeny-based cryptosystems. We give a new algorithm for computing the endomorphism ring of a supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$ that runs, under certain heuristics, in time $O((\log p)^2 p^{1/2})$. The algorithm works by first finding two cycles of a certain form in the supersingular ℓ-isogeny graph $G(p, \ell)$, generating an order $\Lambda \subseteq \operatorname{End}(E)$. Then all maximal orders containing Λ are computed, extending work of Voight (2013). The final step is to determine which of these maximal orders is the endomorphism ring. As part of the cycle-finding algorithm, we give a lower bound on the set of all j-invariants j that are adjacent to j^p in $G(p, \ell)$, answering a question of Arpin et al. (2019).

We also give a polynomial-time reduction from computing $\operatorname{End}(E)$ to path-finding in the ℓ-isogeny graph which is simpler in several ways than previous ones. We show that this reduction leads to another algorithm for computing endomorphism rings which runs in time $\tilde{O}(p^{1/2})$. This allows us to break the second preimage resistance of a hash function in the family constructed by Charles, Goren and Lauter.

1. Introduction

Computing the endomorphism ring of an elliptic curve defined over a finite field is a fundamental problem in computational arithmetic geometry. For ordinary elliptic curves the fastest algorithm is due to Bisson and Sutherland [5] who gave a subexponential time algorithm to solve this problem. No subexponential time algorithm is known for general supersingular elliptic curves.

Computing endomorphism rings of supersingular elliptic curves has emerged as a central problem for isogeny-based cryptography. The first cryptographic application of isogenies between supersingular

---

Eisenträger was partially supported by National Science Foundation award CNS-1617802 and a Vannevar Bush Faculty Fellowship from the US Department of Defense. Hallgren was partially supported by National Science Foundation awards CCF-1618287, CNS-1617802, and a Vannevar Bush Faculty Fellowship from the US Department of Defense. Morrison was supported by the Natural Sciences and Engineering Research Council of Canada, the Canada First Research Excellence Fund, CryptoWorks21, Public Works and Government Services Canada, and the Royal Bank of Canada. Park was partially supported by National Science Foundation award DMS-1902199. This work was done in part while Eisenträger and Hallgren were visiting the Simons Institute for the Theory of Computing.
MSC2010: 11T71, 14G50, 14H52, 14Q05.
Keywords: supersingular elliptic curves, endomorphism ring, isogeny-based cryptography, quaternion algebras, isogeny graph.

elliptic curves was the hash function in [9]. An efficient algorithm for computing the endomorphism ring of a supersingular elliptic curve would, under certain assumptions, completely break this hash function and also SIKE [18; 2]. It would also have a major impact on the security of CSIDH [7].

Computing the endomorphism ring of a supersingular elliptic curve E was first studied by Kohel [20, Theorem 75], who gave an approach for generating a subring of finite index of the endomorphism ring $\mathrm{End}(E)$. The algorithm was based on finding cycles in the ℓ-isogeny graph of supersingular elliptic curves in characteristic p, and the running time of the probabilistic algorithm was $O(p^{1+\varepsilon})$. In this paper we complete Kohel's approach by showing how to compute $\mathrm{End}(E)$ from a suborder when the order is Bass. In a different direction, in [14] it is argued that heuristically one expects $O(\log p)$ calls to a cycle-finding algorithm until the cycles generate $\mathrm{End}(E)$. An algorithm for computing powersmooth endomorphisms with complexity $\tilde{O}(p^{1/2})$ and polynomial storage is given by Delfs and Galbraith [11].

One can also compute $\mathrm{End}(E)$ using an isogeny $\phi : \tilde{E} \to E$, where $\tilde{E}$ is an elliptic curve with known endomorphism ring. McMurdy was the first to compute $\mathrm{End}(E)$ via such an approach [24], but did not determine its complexity. In [14] a polynomial-time reduction from computing $\mathrm{End}(E)$ to finding an isogeny ϕ of powersmooth degree was given assuming some heuristics, while [10] used an isogeny ϕ of ℓ-power degree.

In this paper we give a new algorithm for computing the endomorphism ring of a supersingular elliptic curve E: first we compute two cycles through E in the supersingular ℓ-isogeny graph that generate an order Λ in $\mathrm{End}(E)$. We show that this order will be a Bass order with constant probability, assuming that the discriminants of the two cycles are random in a certain way. Then we compute all maximal orders that contain the Bass order Λ by first solving the problem locally, showing how to efficiently compute all maximal superorders of Λ when Λ is local and Bass. This extends work of Voight [29, Theorem 7.14]. The main property of local Bass orders used here is that there are at most $e+1$ maximal orders containing a local Bass order $\Lambda \otimes \mathbb{Z}_q$, where $e = v_q(\mathrm{discrd}(\Lambda))$ is the valuation of the reduced discriminant of Λ (see [6]). To solve the global case, we use the local data and a local-global principle for quaternionic orders. To bound the running time in this step, we prove that the number of maximal global orders containing Λ is $O(p^{\varepsilon})$ for any $\varepsilon > 0$ when the size of Λ is polynomial in $\log p$ and $\mathrm{discrd}(\Lambda)$ is square-free. We conjecture that this bound also holds when $\mathrm{discrd}(\Lambda)$ is not square-free. Finally, as we compute each global maximal order, we check if it is isomorphic to $\mathrm{End}(E)$. As part of the analysis of the cycle-finding algorithm, we give a lower bound on the size of the set of all j-invariants j that are adjacent to j^p in $G(p, \ell)$, answering the lower-bound part of Question 3 in [1].

Our overall algorithm is still exponential: the two cycles are found in time $O((\log p)^2 p^{1/2})$, and the overall algorithm has the same running time, assuming several heuristics. This saves at least a factor of $\log p$ versus the previous approach in [14] that finds cycles in $G(p, \ell)$ until they generate all of $\mathrm{End}(E)$. This is because with that approach one expects to compute $O(\log p)$ cycles, while our algorithm for the endomorphism ring computes just one pair of cycles and succeeds with constant probability, assuming that the above heuristic about the discriminants of cycles holds. Also, our cycle-finding algorithm requires only polynomial storage, while a generic collision-finding algorithm that relies on the birthday paradox has the same running time as our algorithm but requires exponential storage.

In the last section of the paper we give a new polynomial-time reduction from computing $\mathrm{End}(E)$ to path-finding in the ℓ-isogeny graph which is simpler in several ways than previous ones. For this we need to assume GRH and the heuristics of [14]. We use this to break the second preimage resistance of a hash function in the family constructed in [9].

The paper is organized as follows. Section 2 gives some necessary background. In Section 3 we give an algorithm for computing cycles in the ℓ-isogeny graph $G(p,\ell)$ so that the corresponding endomorphisms generate an order in the endomorphism ring of the associated elliptic curve. In Section 4 we show how to compute all maximal local orders containing a given $\mathbb{Z}_q$-order Λ. In Section 5 we construct global orders from these local orders and compute $\mathrm{End}(E)$. In Section 6 we give a reduction from the endomorphism ring problem to the problem of computing ℓ-power isogenies in $G(p,\ell)$ that is then used to attack the second preimage resistance of the hash function in [9].

2. Background on elliptic curves and quaternion algebras

For the definition of an elliptic curve, its j-invariant, isogenies of elliptic curves, their degrees, and the dual isogeny see [26].

2A. ***Endomorphism rings, supersingular curves, ℓ-power isogenies.*** Let E be an elliptic curve defined over a finite field $\mathbb{F}_q$. An isogeny of E to itself is called an *endomorphism* of E. The set of endomorphisms of E defined over $\overline{\mathbb{F}}_q$ together with the zero map is called the endomorphism ring of E, and is denoted by $\mathrm{End}(E)$.

If the endomorphism ring of E is noncommutative, E is called a *supersingular elliptic curve*. Otherwise we call E *ordinary*. Every supersingular elliptic curve over a field of characteristic p has a model that is defined over $\mathbb{F}_{p^2}$.

Let E, E' be two supersingular elliptic curves defined over $\mathbb{F}_{p^2}$. For each prime $\ell \neq p$, E and E' are connected by a chain of isogenies of degree ℓ. By [20, Theorem 79], E and E' can be connected by m isogenies of degree ℓ, where $m = O(\log p)$. For ℓ a prime different from p, the *supersingular ℓ-isogeny graph in characteristic p* is the multigraph $G(p,\ell)$ whose vertex set is

$$V = V(G(p,\ell)) = \{j \in \mathbb{F}_{p^2} : j = j(E) \text{ for } E \text{ supersingular}\},$$

and the number of directed edges from j to j' is equal to the multiplicity of j' as a root of $\Phi_\ell(j, Y)$. Here, given a prime ℓ, $\Phi_\ell(X,Y) \in \mathbb{Z}[X,Y]$ is the *modular polynomial*. This polynomial has the property that $\Phi_\ell(j, j') = 0$ for $j, j' \in \mathbb{F}_q$ and $q = p^r$ if and only if there exist elliptic curves $E(j), E(j')$ defined over $\mathbb{F}_q$ with j-invariants j, j' such that there is a separable ℓ-isogeny from $E(j)$ to $E(j')$.

2B. ***Quaternion algebras, orders and sizes of orders.*** For $a, b \in \mathbb{Q}^\times$, let $H(a,b)$ denote the quaternion algebra over $\mathbb{Q}$ with basis $1, i, j, ij$ such that $i^2 = a$, $j^2 = b$ and $ij = -ji$. That is, $H(a,b) = \mathbb{Q} + \mathbb{Q}i + \mathbb{Q}j + \mathbb{Q}ij$. Any quaternion algebra over $\mathbb{Q}$ can be written in this form. There is a *canonical* *involution* on $H(a,b)$ which sends an element $\alpha = a_1 + a_2 i + a_3 j + a_4 ij$ to $\overline{\alpha} := a_1 - a_2 i - a_3 j - a_4 ij$. Define the *reduced trace* of an element α as above to be $\operatorname{Trd}(\alpha) = \alpha + \overline{\alpha} = 2a_1$, and the *reduced norm* to be $\operatorname{Nrd}(\alpha) = \alpha\overline{\alpha} = a_1^2 - aa_2^2 - ba_3^2 + aba_4^2$.

A subset $I \subseteq H(a,b)$ is a *lattice* if I is finitely generated as a $\mathbb{Z}$-module and $I \otimes \mathbb{Q} \simeq H(a,b)$. If $I \subseteq H(a,b)$ is a lattice, the *reduced norm of I*, denoted $\operatorname{Nrd}(I)$, is the positive generator of the fractional $\mathbb{Z}$-ideal generated by $\{\operatorname{Nrd}(\alpha) : \alpha \in I\}$. An *order* $\mathcal{O}$ of $H(a,b)$ is a subring of $H(a,b)$ which is also a lattice, and if $\mathcal{O}$ is not properly contained in any other order, we call it a *maximal order*. We call an order $\mathcal{O} \subseteq H(a,b)$ *q-maximal* if $\mathcal{O} \otimes \mathbb{Z}_q$ is a maximal order in $H(a,b) \otimes \mathbb{Z}_q$.

We define $\mathcal{O}_R(I) := \{x \in H(a,b) : Ix \subseteq I\}$ to be the *right order of the lattice I*, and we similarly define its *left order* $\mathcal{O}_L(I)$. If $\mathcal{O}$ is a maximal order in $H(a,b)$ and $I \subseteq \mathcal{O}$ is a left ideal of $\mathcal{O}$, then $\mathcal{O}_R(I)$ is also a maximal order. Here a *left ideal of $\mathcal{O}$* is an additive subgroup of $\mathcal{O}$ that is closed under scalar multiplication on the left. In our setting, a lattice or an order is always specified by a basis. The *size* of a lattice or an order Λ specified by a basis $\mathcal{B}$ in a quaternion algebra B is the number of bits needed to write down the coefficients of the basis $\mathcal{B}$ plus the size of B, which is specified by a basis and a multiplication table. In the following we write $\operatorname{size}(\Lambda)$ for simplicity even though the size depends on the basis chosen to represent Λ. If $\{a_1, a_2, a_3, a_4\}$ is a basis of Λ, the *Gram matrix* of this basis is the 4×4 matrix whose ij-th entry is $\operatorname{Trd}(a_i a_j)$. We denote by $B_{p,\infty}$ the unique quaternion algebra over $\mathbb{Q}$ that is ramified exactly at p and ∞, and this algebra has a standard basis [25, Proposition 5.1]. The endomorphism ring of a supersingular elliptic curve is isomorphic to a maximal order in $B_{p,\infty}$.

2C. ***Bass, Eichler, and Gorenstein orders in quaternion algebras; discriminants and reduced discriminants.*** Let B be a quaternion algebra over $\mathbb{Q}$. We define the *discriminant of B*, denoted $\operatorname{disc} B$, to be the product of primes that ramify in B; then $\operatorname{disc} B$ is a squarefree positive integer. If $\mathcal{O} \subset B$ is an order, we define the *discriminant of $\mathcal{O}$* to be $\operatorname{disc}(\mathcal{O}) := |\det(\operatorname{Trd}(\alpha_i\alpha_j))_{i,j}| \in \mathbb{Z} > 0$, where $\alpha_1, \ldots, \alpha_4$ is a $\mathbb{Z}$-basis for $\mathcal{O}$ [28, §15.2].

The discriminant of an order is always a square, and the *reduced discriminant* $\operatorname{discrd}(\mathcal{O})$ is the positive integer square root so that $\operatorname{discrd}(\mathcal{O})^2 = \operatorname{disc}(\mathcal{O})$ [28, §15.4]. The discriminant of an order measures how far the order is from being a maximal order. The order $\mathcal{O}$ is maximal if and only if $\operatorname{discrd}(\mathcal{O}) = \operatorname{disc} B$ [28, Theorem 23.2.9]. Associated to a quaternion algebra B over $\mathbb{Q}$ there is a *discriminant form* $\Delta : B \to \mathbb{Q}$, defined by $\Delta(\alpha) = \operatorname{Trd}(\alpha)^2 - 4\operatorname{Nrd}(\alpha)$, and we refer to $\Delta(\alpha)$ as the *discriminant* of α. Now let $\mathcal{O} \subset B$ be a $\mathbb{Z}$-order. We say that $\mathcal{O}$ is an *Eichler order* if $\mathcal{O} \subseteq B$ is the intersection of two (not necessarily distinct) maximal orders. The *codifferent* of an order is defined as $\operatorname{codiff}(\mathcal{O}) = \{\alpha \in B : \operatorname{Trd}(\alpha\mathcal{O}) \subseteq \mathbb{Z}\}$. Following [28, Definition 24.2.1], we say that $\mathcal{O}$ is *Gorenstein* if the lattice $\operatorname{codiff}(\mathcal{O})$ is invertible as a lattice as in [28, Definition 16.5.1]. An order $\mathcal{O}$ is *Bass* if every superorder $\mathcal{O}' \supseteq \mathcal{O}$ is Gorenstein. An order is *basic* if it contains a commutative, quadratic subalgebra R such that R is integrally closed in $\mathbb{Q}R$ [28, §24.5]. Given an order Λ, its *radical idealizer* $\Lambda^\natural$ is defined as $\Lambda^\natural = \mathcal{O}_R(\operatorname{rad}\Lambda)$, where $\operatorname{rad}\Lambda$ is the Jacobson radical of the ring Λ. When B is a quaternion algebra over $\mathbb{Q}_p$ and $\mathcal{O}$ is a $\mathbb{Z}_p$-order in B, we similarly define lattices, ideals, and orders in B.

3. Computing an order in the endomorphism ring of a supersingular elliptic curve

3A. *Computing cycles in $G(p, \ell)$.* Fix a supersingular elliptic curve E_0 defined over $\mathbb{F}_{p^2}$ with j-invariant j_0. In this section we describe and analyze an algorithm for computing two cycles through j_0 in $G(p, \ell)$ that generate an order in $\mathrm{End}(E_0)$.

We will first show how to construct two distinct paths from j_0 to j_0^p. Given two such paths P and P', then first traversing through P and then traversing through P' in reverse gives a cycle through j_0. This uses the fact that if j is adjacent to j', then j^p is adjacent to $(j')^p$.

Let P_1 be a path of length k from j_0 to j_k in $G(p, \ell)$. Denote the not necessarily distinct vertices on the path by $j_0, j_1, \ldots, j_k$ and assume that j_k is adjacent to j_k^p in $G(p, \ell)$. Let

$$P_1^p = [j_k, j_k^p, j_{k-1}^p, \ldots, j_1^p, j_0^p].$$

The concatenation $P := P_1 P_1^p$ is a path from j_0 to j_0^p. Such paths were also considered in [9, Section 7].

If $j_0 = j_0^p$, then P is already a cycle. Otherwise, we repeat this process to find another path $P' := P_2 P_2^p$ that passes through at least one vertex not in P. Concatenating P and P' (in reverse order) gives a cycle starting and ending at j_0; this corresponds to an endomorphism of E. We will need the notion of a path/cycle with no *backtracking* and *trimming a path/cycle* to remove backtracking.

Definition 3.1. Suppose $e_j, e_{j'}$ are edges in $G(p, \ell)$ that correspond to ℓ-isogenies

$$\phi_j : E(j) \to E(j') \quad \text{and} \quad \phi_{j'} : E(j') \to E(j)$$

between curves $E(j)$ and $E(j')$ with j-invariants j, j'. We say that e_j is *dual* to $e_{j'}$ if up to isomorphism $\phi_{j'}$ equals the dual isogeny $\hat{\phi}_j$ of ϕ_j. That is $\phi_{j'} = \alpha\hat{\phi}_j$, where $\alpha \in \mathrm{Aut}(E(j))$. We say that a path or cycle with a specified start vertex j_0, following edges $(e_1, \ldots, e_k)$ and ending at vertex j_k *has no backtracking* if e_{i+1} is not dual to e_i for $i = 1, \ldots, k-1$.

In our definition, a cycle has a specified start vertex j_0. According to our definition, if the first edge e_1 and the last edge e_k in such a cycle are dual to each other, it is not considered backtracking.

Definition 3.2. Given a path $(e_1, \ldots, e_k)$ from j_0 to j_k (with $j_0 \neq j_k$) or a cycle with specified start vertex $j_0 = j_k$, define *trimming* as the process of iteratively removing pairs of adjacent dual edges until none are left.

One can show that given a path P from j_0 to j_k with $j_0 \neq j_k$, or a cycle C with start vertex $j_0 = j_k$, the trimmed versions $\tilde{P}$ or $\tilde{C}$ may result in a smaller set of vertices. The vertices j_0 and j_k will still be there in $\tilde{P}$, and the only way that j_0 and j_k may disappear from $\tilde{C}$ is if the whole cycle gets removed.

Definition 3.3. Given a path P in $G_{p,\ell}$ from j_0 to j_k, we define P^R to be the path P traversed in reverse order, from j_k to j_0, using the dual isogenies.

Let

$$S^p := \{j \in \mathbb{F}_{p^2} : j \text{ is supersingular and } j \text{ is adjacent to } j^p \text{ in } G(p, \ell)\}.$$

We can now give the algorithm to find cycle pairs:

Algorithm 3.4. Find cycle pairs for prime ℓ.

Input: A prime $p \neq \ell$ and a supersingular j-invariant $j_0 \in \mathbb{F}_{p^2}$.

Output: Two cycles in $G(p, \ell)$ through j_0.

(1) Perform $N = \Theta(\sqrt{p} \log p \log\log p)$ random walks of length $k = \Theta(\log(p^{3/4}(\log\log p)^{1/2}))$ starting at j_0 and select a walk that hits a vertex $j_k \in S^p$, i.e., such that j_k is ℓ-isogenous to j_k^p; let P_1 denote the path from j_0 to j_k.

(2) Let P_1^p be the path given by $j_k, j_k^p, j_{k-1}^p, \ldots, j_0^p$.

(3) Let P denote the path from j_0 to j_0^p given as the concatenation of P_1 and P_1^p. Remove any self-dual self-loops and trim $P_1 P_1^p$.

(4) If $j_0 \in \mathbb{F}_p$ then $P_1 P_1^p$ is a cycle through j_0.

(5) If $j_0 \in \mathbb{F}_{p^2} - \mathbb{F}_p$ repeat Steps (1)–(3) again to find another path $P' = P_2 P_2^p$ from j_0 to j_0^p; then $P(P')^R$ is a cycle. Remove any self-dual self-loops and trim the cycle.

(6) Repeat Steps (1)–(5) a second time to get a second cycle.

Remark 3.5. Instead of searching for a vertex j in Step (1) such that j is adjacent to j^p, one could also search for a vertex $j \in \mathbb{F}_p$, i.e., j with $j = j^p$, or a vertex j whose distance from j^p in the graph is bounded by some fixed integer B. Our algorithm that searches for a vertex j such that j is adjacent to j^p was easier to analyze because there were fewer cases to consider.

To analyze the running time of Algorithm 3.4, we will use the mixing properties in the Ramanujan graph $G(p, \ell)$. This is captured in the following proposition, which is an extension of [19, Lemma 2.1] in the case that $G(p, \ell)$ is not regular or undirected (that is, when $p \not\equiv 1 \pmod{12}$).

Proposition 3.6. *Let $p > 3$ be prime, and let $\ell \neq p$ also be a prime. Let S be any subset of the vertices of $G(p, \ell)$ not containing* 0 *or* 1728*. Then a random walk of length at least*

$$t = \frac{\log\left(\frac{p}{6|S|^{1/2}}\right)}{\log\left(\frac{\ell+1}{2\sqrt{\ell}}\right)}$$

will land in S with probability at least $6|S|/p$.

One can prove this since the eigenvalues for the adjacency matrix of $G(p, \ell)$ satisfy the Ramanujan bound. This allows us to prove the following theorem.

Theorem 3.7. *Let ℓ, p be primes such that $\ell < p/4$. Under GRH, Algorithm 3.4 computes two cycles in $G(p, \ell)$ through j_0 that generate an order in the endomorphism ring of E_0 in time $O(\sqrt{p}\,(\log p)^2)$, as long as the two cycles do not pass through the vertices* 0 *or* 1728, *with probability* $1 - O(\log p/p)$*. The algorithm requires* poly$(\log p)$ *space.*

Remark 3.8. In Section 5 we use this proposition to compute endomorphism rings, and from this point there is no problem with excluding paths through 0 or 1728. This is because the endomorphism rings of the curves with j-invariants 0 and 1728 are known, and a path of length $\log P$, starting at j_0 going through 0 or 1728 lets us compute $\mathrm{End}(E_0)$ via the reduction in Section 6.

Proof. We implement Step (1) by letting j_{i+1} be a random root of $\Phi_\ell(j_i, Y)$. To test if $j \in S^p$ we check if $\Phi_\ell(j, j^p) = 0$. Assuming GRH, Theorem 3.9 implies that $|S^p| = \Omega(\sqrt{p}/\log\log p)$ (treating ℓ as a constant). Proposition 3.6 implies that the endpoint j_k of a random path found in Step (1) is in S^p with probability $\Omega(1/(\sqrt{p}\log\log p))$. The probability that none of the $N+1$ paths land in S^p is at most $(1 - C/(\sqrt{p}\log\log p))^{N+1} \leq (1 + C/(\sqrt{p}\log\log p))^{-(N+1)} \leq e^{-c\log p/C} = O(1/p)$ if $c = C$, where C is from Theorem 3.9 and c is the constant used in the choice of N.

Now we show that with high probability the two cycles C_0, C_1 returned by the algorithm are linearly independent. We will use Corollary 4.12 of [3]. This corollary states that two cycles C_0 and C_1 with no self-loops generate an order inside $\mathrm{End}(E_0)$ if they

(1) do not go through 0 or 1728,

(2) have no backtracking, and

(3) have the property that one cycle contains a vertex that the other does not contain.

By construction, the cycles C_0 and C_1 returned by our algorithm do not have any self-loops or backtracking. To prove that condition (3) holds, we first claim that with high probability, the end vertex $j_k \in S^p$ in the path P_1 from j_0 to j_k will not get removed when the path $P_1 P_1^p$ is trimmed in Step (3). Then we show it's also still there in the trimmed cycle after Step (5). Observe that if the path P_1 were to be trimmed to obtain a path $\tilde{P}_1$ with no backtracking, then $\tilde{P}_1$ is still a nontrivial path that starts at j_0 and ends at j_k as long as j_0 and j_k are different which occurs with probability $1 - O(1/p)$. After concatenating $\tilde{P}_1$ with its corresponding path $\tilde{P}_1^p$, the path $\tilde{P}_1\tilde{P}_1^p$ has backtracking only if the last edge of $\tilde{P}_1$ is dual to the first edge in $\tilde{P}_1^p$, i.e., if $j_{k-1} = j_k^p$. If that is the case, remove the last edge from $\tilde{P}_1$ and the first edge from $\tilde{P}_1^p$, and call the remaining path $\hat{P}_1$. The new path $\hat{P}_1$ still has the property that it ends in a vertex $j = j_k^p$ that is ℓ-isogenous to its conjugate $(j_k^p)^p = j_k$. After concatenating $\hat{P}_1$ with its corresponding $\hat{P}_1^p$, this still gives a path from j_0 to j_0^p. Again, the concatenation of these two paths has no backtracking unless the last edge in $\hat{P}_1$ is the first edge in $\hat{P}_1^p$, i.e., if the last edge in $\hat{P}_1$ is an edge from j_k to j_k^p. But this cannot happen, because otherwise the trimmed path $\tilde{P}_1$ would have backtracking because it would go from j_k to j_k^p and back to j_k, contradicting the definition of a trimmed cycle. (With negligible probability, the vertex j_k has multiple edges, so we exclude this case here.) Hence the trimmed version of $P_1 P_1^p$ is $\hat{P}_1\hat{P}_1^p$, and this path still contains the vertex j_k, since $\hat{P}_1^p$ contains the vertex j_k. Now we can finish the argument by considering two cases:

Case 1: $j_0 \in \mathbb{F}_p$. The above argument about trimming shows that if the vertex j_k appearing in the second cycle C_1 is different from all the vertices appearing in C_0 and their conjugates, which happens with probability $1 - O(\log p/p)$, then that vertex j_k will appear in the trimmed cycle $\tilde{C}_1$, but not in $\tilde{C}_0$. (This is because in this case the trimmed path $P_1 P_1^p$ is already a cycle.) Hence by [3, Corollary 4.12], $\tilde{C}_0$ and $\tilde{C}_1$ are linearly independent.

Case 2: $j_0 \in \mathbb{F}_{p^2} - \mathbb{F}_p$. Here, with probability $1 - O(\log(p)/p)$, the endpoint j_k of P_2 is a vertex such that neither it nor its conjugate appear as a vertex in P_1. The concatenation of the two paths $P = P_1 P_1^p$ and $P' = P_2 P_2^p$ in reverse is a cycle C_0 through j_0. When we trim it, it is still a cycle through j_0 in which the endpoint j_k from P_2 appears because neither that j_k nor its conjugate appeared in P_1. Similarly, Algorithm 3.4 finds a second cycle C_1 with probability $1 - \log(p)/p$ that contains a random vertex that was not on the first cycle C_0. This means that by Corollary 4.12 of [3], $\tilde{C}_0$ and $\tilde{C}_1$ and hence C_0 and C_1 are linearly independent.

The running time is $O(\sqrt{p}\,(\log p)^2)$ because we are considering $O(\sqrt{p})$ paths of length $O(\log p)$, going from one vertex to the next takes time polynomial in $\ell \log p$, and we are assuming that ℓ is fixed. The storage is polynomial in $\log p$ because we only have to store the paths P_1, P_2 that land in S^p. □

3B. ***Determining the size of S^p.*** We will now determine a lower bound for the size of the set

$$S^p := \{j \in \mathbb{F}_{p^2} : j \text{ is supersingular and } j \text{ is adjacent to } j^p \text{ in } G(p, \ell)\}.$$

In [9, Section 7], an upper bound is given for S^p, but in order to estimate the chance that a path lands in S^p we need a lower bound for this set.

Let ℓ, p be primes such that $\ell < p/4$. Let $\mathcal{O}_K$ be the ring of integers of $K := \mathbb{Q}(\sqrt{-\ell p})$. We use the terminology and notation in [13; 4]. Let $\mathrm{Emb}_{\mathcal{O}_K}(\mathbb{F}_{p^2})$ be the collection of pairs (E, f) such that E is an elliptic curve over $\mathbb{F}_{p^2}$ and $f : \mathcal{O}_K \hookrightarrow \mathrm{End}(E)$ is a normalized embedding, taken up to isomorphism. We say $f : \mathcal{O}_K \hookrightarrow \mathrm{End}(E)$ is *normalized* if each $\alpha \in \mathcal{O}_K$ induces multiplication by its image in $\mathbb{F}_{p^2}$ on the tangent space of E, and (E, f) is isomorphic to (E', f') if there exists an isomorphism $g : E \to E'$ such that $f(\alpha)' = g f(\alpha) g^{-1}$ for all $\alpha \in \mathcal{O}_K$.

Theorem 3.9. *Let ℓ be a prime and assume that $\ell < p/4$. Let*

$$S^p = \{j \in \mathbb{F}_{p^2} : j \text{ is supersingular and } \Phi_\ell(j, j^p) = 0\}.$$

Under GRH there is a constant $C > 0$ (depending on ℓ) such that $|S^p| > C\sqrt{p}/\log\log(p)$.

Proof. First, if E is a supersingular elliptic curve defined over $\mathbb{F}_{p^2}$ with j-invariant j and $E^{(p)}$ is a curve with j-invariant j^p and $\ell < p/4$ is also a prime, then E is ℓ-isogenous to $E^{(p)}$ if and only if $\mathbb{Z}[\sqrt{-\ell p}]$ embeds into $\mathrm{End}(E)$ [9, Lemma 6].

For any element $(E, f) \in \mathrm{Emb}_{\mathcal{O}_K}(\mathbb{F}_{p^2})$, E is supersingular, since p ramifies in $\mathbb{Q}(\sqrt{-\ell p})$. Moreover $j(E) \in S^p$ by the above fact. Thus the map $\rho : \mathrm{Emb}_{\mathcal{O}_K}(\mathbb{F}_{p^2}) \to S^p$ that sends (E, f) to $\rho(E, f) = j(E)$ is well-defined.

To get a lower bound for S^p we will show that for $j \in S^p$, the size of $\rho^{-1}(j)$ is bounded by $(\ell + 1) \cdot 6$ and that $|\mathrm{Emb}_{\mathcal{O}_K}(\mathbb{F}_{p^2})| \gg \sqrt{\ell p}/\log\log(\ell p)$. These two facts imply

$$|S^p| \geq |\mathrm{Emb}_{\mathcal{O}_K}(\mathbb{F}_{p^2})|/((\ell+1)\cdot 6) > \frac{1}{(\ell+1)\cdot 6} \cdot \frac{\sqrt{\ell p}}{\log\log(\ell p)}.$$

To get a lower bound for $|\mathrm{Emb}_{\mathcal{O}_K}(\mathbb{F}_{p^2})|$ we can use [15, Proposition 2.7] to show that $\mathrm{Emb}_{\mathcal{O}_K}(\mathbb{F}_{p^2})$ is in bijection with $\mathrm{Ell}_{\mathcal{O}_K}(\hat{L}_{\mathfrak{P}})$, where $\hat{L}_{\mathfrak{P}}$ is the algebraic closure of the completion of the ring class field $H_{\mathcal{O}_K}$ at a prime $\mathfrak{P}$ above p, and $\mathrm{Ell}_{\mathcal{O}_K}(\hat{L}_{\mathfrak{P}})$ is the set of isomorphism classes of elliptic curves over $\hat{L}_{\mathfrak{P}}$ with endomorphism ring $\mathcal{O}_K$. Hence $|\mathrm{Emb}_{\mathcal{O}_K}(\mathbb{F}_{p^2})| = |\mathrm{Ell}_{\mathcal{O}_K}(\hat{L}_{\mathfrak{P}})|$, whose order equals $|\mathrm{Cl}(\mathcal{O}_K)|$. Class group estimates from [23] give

$$|\mathrm{Cl}(\mathcal{O}_K)| = h(-\ell p) \gg \sqrt{\ell p}/\log\log(\ell p).$$

It remains to bound the size of $\rho^{-1}(j)$. We claim that an equivalence class of pairs (E, f) determines an edge in $G(p, \ell)$. Let $[(E, f)] \in \mathrm{Emb}_{\mathcal{O}_K}(\mathbb{F}_{p^2})$ be given by some representative curve E. First assume that $j(E) \neq 0, 1728$. Then $(E, f) \simeq (E, g)$ implies that $f = g$, since $\mathrm{Aut}(E) = \pm 1$. Thus we may identify $[(E, f)]$ with the edge in $G(p, \ell)$ corresponding to the kernel of $f(\sqrt{-\ell p})$. When $j(E) = 0$ or 1728, we may assume that E is defined over $\mathbb{F}_p$. Then let $[(E, f)] \in \mathrm{Emb}_{\mathcal{O}_K}(\mathbb{F}_{p^2})$ and suppose (E, f) is equivalent to (E, g). We can factor $f(\sqrt{-\ell p}) = \pi \circ \phi$ and $g(\sqrt{-\ell p}) = \pi \circ \phi'$, where ϕ, ϕ' are degree ℓ endomorphisms of E and π is the Frobenius endomorphism of E. Additionally, $\pi\phi = u\pi\phi' u^{-1}$. We claim that u and ϕ commute. If not, then they generate an order Λ such that the following formula holds (see [22]):

$$\mathrm{discrd}(\Lambda) = \tfrac{1}{4}(\Delta(u)\Delta(\phi) - (\mathrm{Trd}(u)\,\mathrm{Trd}(\phi) - 2\,\mathrm{Trd}(u\hat{\phi}))^2) \leq \tfrac{1}{4}\Delta(u)\Delta(\phi). \tag{3-1}$$

One can show that this contradicts our assumption that $p/4 > \ell$. Thus u and ϕ commute, and we see that $f(\sqrt{-\ell p})$ and $g(\sqrt{-\ell p})$ have the same kernel and thus determine the same edge in $G(p, \ell)$.

We now count how many elements of $\mathrm{Emb}_{\mathcal{O}_K}(\mathbb{F}_{p^2})$ determine the same edge in $G(p, \ell)$. Suppose that $[(E, f)], [(E, g)] \in \mathrm{Emb}_{\mathcal{O}_K}(\mathbb{F}_{p^2})$ and that $\ker(f(\sqrt{-\ell p})) = \ker(g(\sqrt{-\ell p}))$. Writing $f(\sqrt{-\ell p}) = \phi \circ \pi$ and $g(\sqrt{-\ell p}) = \phi' \circ \pi$ we see that ϕ and ϕ' must have the same kernel. Thus $\phi' = u\phi$ for some $u \in \mathrm{Aut}(E)$. Because $p > 4\ell > 3$, $\mathrm{Aut}(E) \leq 6$ and we conclude that there are at most 6 classes $[(E, f)]$ determining the same edge emanating from $j(E)$ in $G(p, \ell)$. Thus

$$|\rho^{-1}(j)| \leq (\ell + 1) \cdot 6. \qquad \square$$

Assuming GRH, this result settles the lower-bound portion of Question 3 in [1]. See Lemma 6 of [9] for the upper-bound.

4. Enumerating maximal superorders: the local case

Let q be a prime. In this section, we give an algorithm for the following problem:

Problem. Given a $\mathbb{Z}_q$-order $\Lambda \subseteq M_2(\mathbb{Q}_q)$, find all maximal orders containing Λ.

For general Λ there might be an exponential number of maximal orders containing it, so the algorithm for enumerating them would also be exponential time. However, we will show that the above problem can be solved efficiently when Λ is Bass. The main property of local Bass orders Λ we use is that there are at most $e + 1$ maximal orders containing Λ, where $e = v_q(\mathrm{discrd}(\Lambda))$ [6, Corollaries 2.5, 3.2, 4.3 and Proposition 3.1].

We use the Bruhat–Tits tree $\mathcal{T}$ [28, §23.5] to compute the maximal superorders of Λ. The vertices of $\mathcal{T}$ are in bijection with maximal orders in $M_2(\mathbb{Q}_q)$.

A homothety class of lattices $[L] \subseteq \mathbb{Q}_q^2$ corresponds to a maximal order via

$$L \mapsto \operatorname{End}_{\mathbb{Z}_q}(L) = \{x \in M_2(\mathbb{Q}_q) : xL \subseteq L\} \subseteq M_2(\mathbb{Q}_q) \tag{4-1}$$

for every choice of $L \in [L]$. Two maximal orders $\mathcal{O}$ and $\mathcal{O}'$ are adjacent in $\mathcal{T}$ if there exist lattices L and L' for $\mathcal{O}$ and $\mathcal{O}'$ such that $qL \subsetneq L' \subsetneq L$. Hence the neighbors of $\mathcal{O}$ in $\mathcal{T}$ correspond to the one-dimensional subspaces of $L/qL \cong \mathbb{F}_q \times \mathbb{F}_q$.

A division quaternion algebra B over $\mathbb{Q}_q$ has only one maximal order, which can be found using the algorithm in [29]. The split case is solved by Algorithm 4.1, and also relies on the algorithm in [29].

Algorithm 4.1. Enumerate all maximal orders containing a local order.

Input: A $\mathbb{Z}_q$-order $\Lambda \subseteq M_2(\mathbb{Q}_q)$, represented by a $\mathbb{Z}_q$-basis.

Output: The maximal orders in $M_2(\mathbb{Q}_q)$ containing Λ, each specified by a $\mathbb{Z}_q$-basis.

(1) Compute a maximal order $\tilde{\mathcal{O}} \supseteq \Lambda$ with [29, Algorithm 7.10] and a lattice $\tilde{L}$ in $\mathbb{Q}_q \times \mathbb{Q}_q$ such that $\tilde{\mathcal{O}} = \operatorname{End}_{\mathbb{Z}_q}(\tilde{L})$.

(2) Let $A = \{\tilde{\mathcal{O}}\}$ and $B = \{\tilde{L}\}$.

(3) While $B \neq \varnothing$:

   (a) Remove L from B, and label it as discovered. Set $\mathcal{O} = \operatorname{End}_{\mathbb{Z}_q}(L)$.

   (b) Compute the set of neighbors $\mathcal{N}_{\mathcal{O}}$ of $\mathcal{O}$ that contain Λ.

   (c) For each $\mathcal{O}' \in \mathcal{N}_{\mathcal{O}}$ not labeled as discovered, add $\mathcal{O}'$ to A and its corresponding lattice to B.

(4) Return A.

Now we show that Algorithm 4.1 is efficient when the input lattice Λ is Bass.

Proposition 4.2. *Let $\Lambda \subseteq M_2(\mathbb{Q}_q)$ be a Bass $\mathbb{Z}_q$-order, and $e := v_q(\operatorname{discrd}(\Lambda))$. Algorithm 4.1 computes $A := \{\mathcal{O} \supseteq \Lambda : \mathcal{O}$ is maximal$\}$, and $|A| \le e+1$. The runtime is polynomial in* $\log q \cdot \operatorname{size}(\Lambda)$.

Proof. To prove correctness we first show that the maximal orders containing an arbitrary order Λ' in $M_2(\mathbb{Q}_q)$ form a subtree of $\mathcal{T}$. If $\mathcal{O}, \mathcal{O}'$ are two maximal orders containing Λ', then the maximal orders containing $\mathcal{O} \cap \mathcal{O}'$ are precisely the vertices in the path between $\mathcal{O}$ and $\mathcal{O}'$ in $\mathcal{T}$ [28, §23.5.15]. Each order on this path also contains Λ', so the maximal orders containing Λ' form a connected subset of $\mathcal{T}$. The above algorithm explores this subtree.

If Λ is Bass and Eichler, i.e., $\Lambda = \mathcal{O} \cap \mathcal{O}'$ for maximal orders $\mathcal{O}, \mathcal{O}'$, then there are $e+1$ maximal orders containing Λ [6, Corollary 2.5], and they are exactly the vertices on the path from $\mathcal{O}$ to $\mathcal{O}'$. If Λ is Bass but not Eichler, then there are either 1 or 2 maximal orders containing Λ by [6, Proposition 3.1 and Corollaries 3.2 and 4.3]. Since they form a tree, they must also form a path. In either case, $|A| \le e+1$, and the vertices in A form a path.

As for the running time, in Step (1) we run [29, Algorithm 7.10], which is polynomial in $\log q \cdot \mathrm{size}(\Lambda)$. Let L be a lattice such that $\mathcal{O} = \mathrm{End}_{\mathbb{Z}_q}(L)$ contains Λ. The neighbors of $\mathcal{O}$ containing Λ are in bijection with the lines in L/qL fixed by the action of the image of Λ in $\mathcal{O}/q\mathcal{O} \simeq M_2(\mathbb{F}_q)$. For each such line, let $\bar{v} \in L/qL$ be a nonzero vector, and let v be a lift to L. Let $w \in L$ be such that $\{v, w\}$ is a $\mathbb{Z}_q$-basis of L. Then $L' := \mathrm{span}\{v, qw\}$ is a $\mathbb{Z}_q$-lattice such that $\mathcal{O}' := \mathrm{End}_{\mathbb{Z}_q}(L')$ contains Λ. So we can efficiently compute the lattices L' corresponding to the neighbors of $\mathcal{O}$ which contain Λ. Given such an L', let $x \in M_2(\mathbb{Q}_q)$ be the base change matrix from L to L'. If $\mathcal{B}$ is a basis for $\mathcal{O}$, then $\mathcal{B}' := x\mathcal{B}x^{-1}$ is a basis for $\mathcal{O}'$. The size of $\mathcal{B}'$ is $c(\log q) + \mathrm{size}(\mathcal{O})$ for some constant c, so each neighbor of $\mathcal{O}$ containing Λ can be computed in time polynomial in $\log q \cdot \mathrm{size}(\mathcal{O})$.

Since the length of the path explored in the algorithm is at most e, where $e = v_q(\mathrm{discrd}(\Lambda))$ is polynomial in $\mathrm{size}(\Lambda)$, and the size of the starting order $\tilde{\mathcal{O}}$ is polynomial in $\log q \cdot \mathrm{size}(\Lambda)$ we obtain that the size of any maximal order containing Λ is polynomial in $\mathrm{size}(\Lambda) \cdot \log q$. Each step takes time polynomial in $\log q \cdot \mathrm{size}(\Lambda)$, so the whole algorithm is polynomial in $\log q \cdot \mathrm{size}(\Lambda)$. □

Later we will need to enumerate the q-maximal $\mathbb{Z}$-orders containing a Bass $\mathbb{Z}$-order Λ. The algorithm below uses Algorithm 4.1 to accomplish this.

Algorithm 4.3. Enumerate the q-maximal $\mathbb{Z}$-orders $\mathcal{O}$ containing Λ.

Input: A $\mathbb{Z}$-order Λ, specified by a $\mathbb{Z}$-basis, and prime q such that $\Lambda \otimes \mathbb{Z}_q$ is Bass.

Output: All $\mathbb{Z}$-orders $\mathcal{O} \supseteq \Lambda$ such that $\mathcal{O}$ is q-maximal and $\mathcal{O} \otimes \mathbb{Z}_{q'} = \Lambda \otimes \mathbb{Z}_{q'}$ for all primes $q \neq q'$.

(1) Compute an embedding $f : \Lambda \otimes \mathbb{Q} \hookrightarrow M_2(\mathbb{Q}_q)$ such that $f(\Lambda) \subseteq M_2(\mathbb{Z}_q)$.

(2) Let A be the output of Algorithm 4.1 on input $f(\Lambda)$.

(3) Return $\{f^{-1}(\mathcal{O}) + \Lambda : \mathcal{O} \in A\}$.

Lemma 4.4. *Algorithm 4.3 is correct. The run time is polynomial in* $\log q \cdot \mathrm{size}(\Lambda)$.

Proof. Step (1) can be accomplished with Algorithms 3.12, 7.9, and 7.10 in [29], which run in time polynomial in $\log q \cdot \mathrm{size}\,\Lambda$. For each maximal $\mathbb{Z}_q$-order $\mathcal{O} \supseteq f(\Lambda)$, we then compute a corresponding $\mathbb{Z}$-lattice $\mathcal{O}' \supseteq \Lambda$, whose generators are $\mathbb{Z}[q^{-1}]$-linear combinations of generators of Λ. The denominator of these coefficients is at most q^e where $e := v_q(\mathrm{discrd}(\Lambda))$. By Proposition 4.2, there are at most $e+1$ maximal orders containing $f(\Lambda)$ if $\Lambda \otimes \mathbb{Z}_q$ is Bass. It is straightforward to check that the lattice $\Lambda + \mathcal{O}'$ is actually a $\mathbb{Z}$-order and has the desired completions. Moreover, these are all such orders by the local-global principle for orders, [28, Theorem 9.5.1]. □

Remark 4.5 (global case). Algorithm 4.3 can be used to enumerate all maximal orders $\mathcal{O}$ of a quaternion algebra B over $\mathbb{Q}$ that contain a $\mathbb{Z}$-order Λ which is Bass, given Λ and the factorization of $\mathrm{discrd}(\Lambda)$ as $\mathrm{discrd}(\Lambda) = \prod_{i=1}^m q_i^{e_i}$:

We run Algorithm 4.3 m times, namely on $(\Lambda, q_1), \ldots, (\Lambda, q_m)$. Let $\{X_1, \ldots, X_m\}$ be the output, where $X_i = \{\mathcal{O}_{i1}, \ldots, \mathcal{O}_{in_i}\}$. The global orders containing Λ are in bijection with $\prod_i X_i$, by associating to $(\mathcal{O}_{1j_1}, \ldots, \mathcal{O}_{mj_m}) \in \prod X_i$ the order $\sum_i \mathcal{O}_{ij_i}$. In particular, the number of such orders is at most $\prod_i (e_i + 1)$.

The correctness of this follows from the local-global principle for maximal orders [28, Lemma 10.4.2]. The above results show that each order in the enumeration can be computed in time polynomial in the size of Λ. However, for an arbitrary order Λ, there might be an exponential number of orders containing it.

5. Computing $\mathrm{End}(E)$

Now we describe our algorithm to compute the endomorphism ring of E. By computing $\mathrm{End}(E)$ we mean computing a basis for an order $\mathcal{O}$ in $B_{p,\infty}$ that is isomorphic to $\mathrm{End}(E)$, and that we can evaluate the basis at all points of E via an isomorphism $B_{p,\infty} \to \mathrm{End}(E) \otimes \mathbb{Q}$. First we give an algorithm that uses Algorithm 3.4 to generate a Bass suborder of $\mathrm{End}(E)$. A heuristic about the distribution of discriminants of cycles is used to show that just one call to Algorithm 3.4 generates a Bass order with constant probability. Then we give an algorithm which recovers $\mathrm{End}(E)$ from a Bass suborder. The key property used here is that Bass orders Λ (whose basis is of size polynomial in $\log p$ and whose discriminant is $O(p^k)$) only have $O(p^\varepsilon)$ maximal orders containing them for any $\varepsilon > 0$. This is proved in Proposition 5.5 when the reduced discriminant is square-free. Based on our numerical evidence, we conjecture that this holds for general Bass orders as well.

5A. *Computing a Bass order.*

Algorithm 5.1. Compute a Bass suborder $\Lambda \subseteq \mathrm{End}(E)$.

Input: A supersingular elliptic curve E.

Output: A Bass order $\Lambda \subseteq \mathrm{End}(E)$ and the factorization of $\mathrm{discrd}(\Lambda)$, or "false".

(1) Compute two cycles in $G(p, \ell)$ through $j(E)$ using Algorithm 3.4.

(2) Let α, β be the endomorphisms corresponding to the cycles from Step (1). Compute the Gram matrix for $\{1, \alpha, \beta, \alpha\beta\}$ and from it an abstract representation for $\Lambda = \langle 1, \alpha, \beta, \alpha\beta \rangle$.

(3) Factor $\mathrm{discrd}(\Lambda) = \prod_{i=1}^n q_i^{e_i}$.

(4) If Λ is Bass return Λ and the factorization of $\mathrm{discrd}(\Lambda)$, else return "false".

To analyze the algorithm we introduce a new heuristic:

Heuristic 5.2. The probability that the discriminants of the two endomorphisms corresponding to the cycles produced by Algorithm 3.4 are coprime is at least μ for some constant $\mu > 0$ not depending on p.

This heuristic is based on our numerical experiments. Intuitively, we are assuming that the endomorphisms we compute with Algorithm 3.4 have discriminants which are distributed like random integers that satisfy the congruency conditions to be the discriminant of an order in a quadratic imaginary field in which p is inert and ℓ splits. Two random integers are coprime with probability $6/\pi^2$. We are assuming that the discriminants of our cycles are coprime with constant probability.

Theorem 5.3. *Assume GRH and Heuristic 5.2. Then with probability at least μ, Algorithm 5.1 computes a Bass order $\Lambda \subseteq End(E)$, and the runtime is $O(\sqrt{p}(\log p)^2)$.*

Proof. In Step (2), the Gram matrix for Λ, whose entries are the reduced traces of pairwise products of the basis elements, is computed. This uses a generalization of Schoof's algorithm (see Theorem A.6 of [3]), which runs in time polynomial in $\log p$ and log of the norm of α, β. Since α and β arise from cycles of length at most $c\lceil \log p \rceil$, for some constant c which is independent of p, the norms of α and β are at most p^c. From the Gram matrix we can efficiently compute $\operatorname{discrd}(\Lambda)$.

To check that Λ is Bass, it is enough to check that Λ is Bass at each q dividing $\operatorname{discrd}(\Lambda)$ [8, Theorem 1.2]. To check that Λ is Bass at q it is enough to check that $\Lambda \otimes \mathbb{Z}_q$ and $(\Lambda \otimes \mathbb{Z}_q)^\natural$ are Gorenstein [8, Corollary 1.3]. An order is Gorenstein if and only if its ternary quadratic form is primitive [28, Corollary 24.2.10], and this can be checked efficiently. Thus, given a factorization of $\operatorname{discrd}(\Lambda)$, we can efficiently decide if Λ is Bass.

Finally, we compute the probability that the order returned by Algorithm 3.4 is Bass. By [8, Theorem 1.2], an order is Bass if and only if it is basic, and being basic is a local property. It follows that the order Λ is Bass whenever the conductors of $\mathbb{Z}[\alpha]$ and $\mathbb{Z}[\beta]$ are coprime. A sufficient condition for this is that the discriminants of α and β are coprime which will happen with probability at least μ by the above heuristic. This sufficient condition also covers the case when the cycle for α or β goes through 0 or 1728 even though Theorem 3.7 does not apply here. □

5B. ***Computing End(E) from a Bass order.*** In this section we compute $\operatorname{End}(E)$ from a given Bass suborder Λ. For this we enumerate the maximal orders containing Λ by taking sums of the q-maximal orders returned by Algorithm 4.3. As we enumerate the orders, we check each one to see if it is isomorphic to $\operatorname{End}(E)$.

Algorithm 5.4. Compute $\operatorname{End}(E)$ from a Bass order.

Input: A Bass order $\Lambda \subseteq \operatorname{End}(E)$ with factored reduced discriminant $\prod_{i=1}^{n} q_i^{e_i}$.

Output: A compact representation of $\operatorname{End}(E)$, as defined in [12, Section 8.2].

(1) For each $i = 1$ to n:
   (a) Compute all orders $\{\mathcal{O}_{i,1}, \dots, \mathcal{O}_{i,m_i}\}$ which are maximal at q_i and equal to Λ at primes $q' \neq q_i$ by running Algorithm 4.3 with input Λ and prime q_i.

(2) Compute $f : \Lambda \otimes \mathbb{Q} \to B_{p,\infty}$.

(3) For each choice of indices $(i_1, \dots, i_n) \in [m_1] \times \cdots \times [m_n]$:
   (a) Set $\mathcal{O} := \mathcal{O}_{1,i_1} + \cdots + \mathcal{O}_{n,i_n}$.
   (b) Compute $E'/\mathbb{F}_{p^2}$ such that $\operatorname{End}(E') \simeq f(\mathcal{O})$ along with a compact representation of $\operatorname{End}(E')$.
   (c) If $j(E') = j(E)$ or $j(E') = j(E)^p$, return $f(\mathcal{O})$ and the compact representation of $\operatorname{End}(E')$.

Proposition 5.5. *Fix a positive integer k, and let Λ be a Bass order whose size is polynomial in* $\log p$ *and whose reduced discriminant is square-free and of size $O(p^k)$. Assume that the factorization of the reduced discriminant is given. There are $O(p^\varepsilon)$ maximal orders containing Λ and Algorithm 5.4 terminates in time $\tilde{O}(p^\varepsilon)$ for any $\varepsilon > 0$, assuming that the heuristics in* [14; 12] *hold.*

p	orders	Bass orders	average $N(\Lambda)$
70,001	92	76	122.21
90,001	80	67	322.04
100,003	81	75	337.59

Table 1. Results from computing 100 pairs of cycles in $G(p, 2)$ at random $j \in \mathbb{F}_{p^2} - \mathbb{F}_p$.

Proof. Computing the isomorphism $f : \Lambda \otimes \mathbb{Q} \simeq B_{p,\infty}$ requires one call to an algorithm for factoring integers (and poly$(\log p)$ calls to algorithms for factoring polynomials over $\mathbb{F}_p$, see [17]). Let

$$\operatorname{discrd}(\Lambda) = p \cdot \prod_{i=1}^{m} q_i$$

with $q_1, \ldots, q_m$ distinct and different from p. By the local-global principle for maximal orders there is one maximal order corresponding to each collection of q_i-maximal orders $\{\mathcal{O}_i\}$ with $\mathcal{O}_i \supseteq \Lambda \otimes \mathbb{Z}_{q_i}$. We loop through these orders in Step (3). The size of the index set in that loop and hence the number of distinct maximal orders containing Λ is at most $2^{\omega(\operatorname{discrd}(\Lambda))-1}$, where $\omega(n)$ denotes the number of distinct prime factors of an integer n. Fix $\varepsilon > 0$. Since $\omega(n) = O(\log n / \log\log n)$ [16, Chapter 22, §10], for p large enough, the number of maximal orders $\mathcal{O} \supseteq \Lambda$ is at most

$$2^{c' \frac{\log c \cdot p^k}{\log\log c \cdot p^k}} = (c \cdot p^k)^{\frac{c'}{\log\log c \cdot p^k}}$$

for some $c, c' > 0$, which is $O(p^{\varepsilon})$.

As we loop through the maximal orders $\mathcal{O}$ containing Λ, we check each one to see if it is isomorphic to $\operatorname{End}(E)$: after constructing such an order in Step (3)(a), we compute in Step (3)(b) a curve E' whose endomorphism ring is isomorphic to $\mathcal{O}$. This can be solved efficiently with the algorithms in [14]: one computes a connecting ideal I between $\mathcal{O}$ and a special order $\mathcal{O}'$ and then applies Algorithm 2 of [14] (see also Algorithm 12 of [12]). Then, in Step (3)(c), we compare j-invariants. Checking each order takes time polynomial in $\log p$ (assuming the heuristics in [14; 12]), so the total running time of the algorithm is $\tilde{O}(p^{\varepsilon})$ for any $\varepsilon > 0$. □

Our computational data from Section 5C suggests that we will get the same running time when the reduced discriminant of Λ is not square-free. This motivates the following conjecture:

Conjecture 5.6. *Fix an integer $k \geq 0$ and assume that $\Lambda \subseteq \operatorname{End}(E)$ is a Bass order of size polynomial in* $\log p$ *and with* $\operatorname{discrd}(\Lambda) = O(p^k)$. *Then for any $\varepsilon > 0$, the number of maximal orders containing Λ is $O(p^{\varepsilon})$.*

Theorem 5.7. *Assume GRH, Conjecture 5.6, Heuristic 5.2, and the heuristics in* [14]. *Let E be a supersingular elliptic curve. Then the algorithm which combines Algorithm 5.1 and Algorithm 5.4 computes $\operatorname{End}(E)$ with probability at least μ, in time $O((\log p)^2 \sqrt{p})$.*

Proof. By the proof of Theorem 5.3, the norms of the endomorphisms α_1, α_2 computed by Algorithm 3.4 are bounded by p^c for some constant c independent of p, so their discriminants satisfy $|\Delta(\alpha_i)| < 4p^c$.

Hence by (3-1), they generate an order Λ whose reduced discriminant satisfies $\operatorname{discrd}(\Lambda) = O(p^{2c})$. This means we can apply Conjecture 5.6, so the theorem follows from Theorem 5.3. □

5C. ***Computational data.*** We implemented a cycle-finding algorithm in Sage along with an algorithm for computing traces of cycles in $G(p, \ell)$, which is based on the implementation of Schoof's algorithm available in [27]. For each p in Table 1, and for 100 iterations, we computed a pair of cycles in $G(p, 2)$. We then tested whether they generated a Bass order by testing whether the two quadratic orders had coprime conductors and computed the discriminant of the order that they generated. We also computed an upper bound on the number of maximal orders containing Λ when Λ was Bass: suppose $\operatorname{discrd}(\Lambda) = p \prod_i q_i^{e_i}$, then there are at most $N(\Lambda) := \prod_i (e_i + 1)$ maximal orders containing Λ. We report how often the two cycles generated an order, how many of those orders were Bass, and the average value of $N(\Lambda)$. The cycle-finding algorithm we implemented is the variant discussed in Remark 3.5: it searches for $j \in \mathbb{F}_p$ to construct the cycles using walks of length $\lceil \log p \rceil$. We also did not avoid a second cycle which may commute with the first since even without that more than 80% of cases were orders. We also only computed cycles at $j \in \mathbb{F}_{p^2} - \mathbb{F}_p$ because this is the case of interest as there are no obvious noninteger endomorphisms.

6. Computing End(E) via path-finding in the ℓ-isogeny graph

In this section, we give a reduction from the endomorphism ring problem to the problem of computing ℓ-power isogenies in $G(p, \ell)$, using ideas from [21], [14], and [12]. This reduction is simpler than the one in [12], and uses only one call to a path-finding oracle (rather than poly($\log p$) calls to an oracle for finding cycles in $G(p, \ell)$, as in [12]). We apply this reduction in two ways, noting that it gives an algorithm for computing the endomorphism ring, and that it breaks second preimage resistance of the variable-length version of the hash function in [9].

6A. ***Reduction from computing End(E) to path-finding in the ℓ-isogeny graph.*** We first define the path-finding problem in $G(p, \ell)$:

Problem (ℓ-PowerIsogeny). Given a prime p and supersingular elliptic curves E and E' over $\mathbb{F}_{p^2}$, output a chain of ℓ-isogenies of length $O(\log p)$ from E to E'.

Computing the endomorphism ring of a supersingular elliptic curve via an oracle for ℓ-PowerIsogeny proceeds as follows. On input p, Algorithm 3 of [12] returns a supersingular elliptic curve $\tilde{E}$ defined over $\mathbb{F}_{p^2}$ and a maximal order $\tilde{\mathcal{O}} \subseteq B_{p,\infty}$ with an explicit $\mathbb{Z}$-basis $\{x_1, \ldots, x_4\}$. Proposition 3 of [12] gives an isomorphism $g : \tilde{\mathcal{O}} \to \operatorname{End}(\tilde{E})$ such that we can efficiently evaluate $g(x_i)$ at points of E_0. From this, the endomorphism ring of any supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$ can be computed, given a path in $G(p, \ell)$ from $\tilde{E}$ to E, with $\ell \neq p$ a small fixed prime, for example $\ell = 2$ or 3.

The following algorithm gives a polynomial time reduction from computing endomorphism rings to the path-finding problem, which uses only one call to the path-finding oracle. It assumes the heuristics of [14] and GRH (to compute $\tilde{E}$). A similar algorithm appeared in [10].

Algorithm 6.1. Reduction from computing $\mathrm{End}(E)$ to ℓ-PowerIsogeny.

Input: A prime p, and $E/\mathbb{F}_{p^2}$ supersingular.

Output: A maximal order $\mathcal{O} \simeq \mathrm{End}(E)$, whose elements can be evaluated at any point of E, and a powersmooth isogeny $\psi_e : \tilde{E} \to E$, with $\tilde{E}$ as above.

(1) Compute $\tilde{E}, \tilde{\mathcal{O}}$ with Algorithm 3 in [12].

(2) Run the oracle for path-finding on $\tilde{E}$, E to obtain an ℓ-power isogeny $\phi = \phi_e \circ \cdots \circ \phi_1 : \tilde{E} \to E$ of degree ℓ^e.

(3) Let $J_0 := \tilde{\mathcal{O}}$, $P_0 := \tilde{\mathcal{O}}$, $\mathcal{O}_0 := \tilde{\mathcal{O}}$.

(4) For $k := 1, \dots, e$:

   (a) Compute $I_k \subseteq \mathcal{O}_{k-1}$, the kernel ideal of ϕ_k.
   (b) Compute $J_k := J_{k-1} I_k$.
   (c) Compute P_k, an ideal equivalent to J_k of powersmooth norm.
   (d) Compute an isogeny $\psi_k : \tilde{E} \to E_k$ corresponding to P_k.
   (e) Set $\mathcal{O}_k := \mathcal{O}_R(P_k)$.

(5) Return $\mathcal{O}_R(P_e)$, ψ_e.

Orders and ideals appearing in the above algorithm are represented by a $\mathbb{Z}$-basis, and we can compute right orders of ideals using linear algebra over $\mathbb{Z}$, as in [12]. The ideal I_k, which is the ideal of $\mathcal{O}_{k-1}$ of norm ℓ corresponding to ϕ_k, can be computed efficiently because we can evaluate endomorphisms efficiently using Proposition 3 of [12]. The algorithm is correct because $\mathcal{O}_R(P_e) = \mathcal{O}_R(J_e) = \mathrm{End}(E_e) = \mathrm{End}(E)$.

6B. ***Using Algorithm 6.1 to compute endomorphism rings and break the second preimage of the CGL hash.*** Algorithm 6.1 can be used to give an algorithm for computing the endomorphism ring of a supersingular elliptic curve E by combining it with algorithms from [11; 14; 12]. This yields a $O((\log p)^2 p^{1/2})$ time algorithm with polynomial storage, assuming the relevant heuristics in [14; 12].

We now consider the hash function in [9] constructed from Pizer's graphs $G(p, 2)$. For each supersingular elliptic curve $\tilde{E}$, there is an associated hash function. An input $s \in \{0, 1\}^*$ to the hash function determines a walk in $G(p, 2)$ from $\tilde{E}$ to another curve E, and the output of the hash function is $j(E)$. The following is an improvement over [12], which gave a collision attack for this specific hash function.

Proposition 6.2. *Let $\tilde{E}$ be the elliptic curve computed in Step* (1) *of Algorithm 6.1. For the hash function associated to $\tilde{E}$, Algorithm 6.1 gives a second preimage attack* (*and hence, also a collision attack*) *that runs in time polynomial in* $\log p$.

Proof. The attack works as follows: Given a path from $\tilde{E}$ to E, use Algorithm 6.1 to compute $\mathrm{End}(E)$. Then use Algorithm 7 of [12] to compute new paths from $\tilde{E}$ to E. □

Acknowledgements

We would like to thank Ben Diamond, Daniel Smertnig and John Voight for several helpful discussions and suggestions. We would like to thank an anonymous reviewer of an earlier version of this paper whose suggestions greatly simplified Section 4.

References

[1] Sarah Arpin, Catalina Camacho-Navarro, Kristin Lauter, Joelle Lim, Kristina Nelson, Travis Scholl, and Jana Sotáková. Adventures in supersingularland. Preprint, 2019. arXiv:1909.07779.

[2] Reza Azarderakhsh, Matthew Campagna, Craig Costello, Luca De Feo, Basil Hess, Amir Jalali, David Jao, Brian Koziel, Brian LaMacchia, Patrick Longa, Michael Naehrig, Joost Renes, Vladimir Soukharev, and David Urbanik. Supersingular isogeny key encapsulation. Submission to the NIST Post-Quantum Standardization project, 2017. csrc.nist.gov/Projects/Post-Quantum-Cryptography/Round-1-Submissions.

[3] Efrat Bank, Catalina Camacho-Navarro, Kirsten Eisenträger, Travis Morrison, and Jennifer Park. Cycles in the supersingular ℓ-isogeny graph and corresponding endomorphisms. *Proceedings of the Women in Numbers 4 Conference*, To appear in WIN 4 proceedings, 2019. arXiv:1804.04063.

[4] Juliana Belding, Reinier Bröker, Andreas Enge, and Kristin Lauter. Computing Hilbert class polynomials. In *Algorithmic number theory*, volume 5011 of *Lecture Notes in Comput. Sci.*, pages 282–295. Springer, Berlin, 2008.

[5] Gaetan Bisson and Andrew V. Sutherland. Computing the endomorphism ring of an ordinary elliptic curve over a finite field. *J. Number Theory*, 131(5):815–831, 2011.

[6] Juliusz. Brzeziński. On orders in quaternion algebras. *Comm. Algebra*, 11(5):501–522, 1983.

[7] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: an efficient post-quantum commutative group action. In *Advances in cryptology—ASIACRYPT 2018. Part III*, volume 11274 of *Lecture Notes in Comput. Sci.*, pages 395–427. Springer, Cham, 2018.

[8] Sara Chari, Daniel Smertnig, and John Voight. On basic and Bass quaternion orders. Preprint, 2019. arXiv:1903.00560.

[9] Denis X. Charles, Eyal Z. Goren, and Kristin Lauter. Cryptographic hash functions from expander graphs. *J. Cryptology*, 22(1):93–113, 2009.

[10] Luca De Feo, Simon Masson, Christophe Petit, and Antonio Sanso. Verifiable delay functions from supersingular isogenies and pairings. In *Advances in Cryptology – ASIACRYPT 2019*, volume 11921 of *Lecture Notes in Comput. Sci.*, pages 248–277. Springer, 2019.

[11] Christina Delfs and Steven D. Galbraith. Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$. *Des. Codes Cryptography*, 78(2):425–440, February 2016.

[12] Kirsten Eisenträger, Sean Hallgren, Kristin Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: reductions and solutions. *Eurocrypt 2018, LNCS 10822*, pages 329–368, 2018.

[13] Noam D. Elkies. Supersingular primes for elliptic curves over real number fields. *Compositio Math.*, 72(2):165–172, 1989.

[14] Steven D. Galbraith, Christophe Petit, and Javier Silva. Identification protocols and signature schemes based on supersingular isogeny problems. In *Advances in cryptology—ASIACRYPT 2017. Part I*, volume 10624 of *Lecture Notes in Comput. Sci.*, pages 3–33. Springer, 2017.

[15] Benedict Gross and Don Zagier. On singular moduli. *J. Reine Angew. Math.*, 355:191–220, 1985.

[16] Godfrey H. Hardy and Edward M. Wright. *An introduction to the theory of numbers*. Oxford University Press, Oxford, sixth edition, 2008.

[17] Gábor Ivanyos, Lajos Rónyai, and Josef Schicho. Splitting full matrix algebras over algebraic number fields. *J. Algebra*, 354:211–223, 2012.

[18] David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In *Post-quantum cryptography*, volume 7071 of *Lecture Notes in Comput. Sci.*, pages 19–34. Springer, Heidelberg, 2011.

[19] David Jao, Stephen D. Miller, and Ramarathnam Venkatesan. Expander graphs based on GRH with an application to elliptic curve cryptography. *J. Number Theory*, 129(6):1491–1504, 2009.

[20] David Kohel. *Endomorphism rings of elliptic curves over finite fields*. PhD thesis, University of California, Berkeley, 1996.

[21] David Kohel, Kristin Lauter, Christophe Petit, and Jean-Pierre Tignol. On the quaternion ℓ-isogeny path problem. *LMS Journal of Computation and Mathematics*, 17:418–432, 2014.

[22] Kristin Lauter and Bianca Viray. An arithmetic intersection formula for denominators of Igusa class polynomials. *Amer. J. Math.*, 137(2):497–533, 2015.

[23] John E. Littlewood. On the class-number of the corpus $P(\sqrt{-k})$. *Proc. London Math. Soc. (2)*, 27(5):358–372, 1928.

[24] Ken McMurdy. Explicit representations of the endomorphism rings of supersingular elliptic curves. 2014.

[25] Arnold Pizer. An algorithm for computing modular forms on $\Gamma_0(N)$. *J. Algebra*, 64(2):340–390, 1980.

[26] Joseph H. Silverman. *The arithmetic of elliptic curves*. Springer, New York, 2009.

[27] Andrew Sutherland. 18.783 Elliptic Curves. Massachusetts Institute of Technology: MIT OpenCourseWare, ocw.mit.edu License: Creative Commons BY-NC-SA.

[28] John Voight. *Quaternion algebras*. Version v.0.9.14, July 7, 2018.

[29] John Voight. Identifying the matrix ring: algorithms for quaternion algebras and quadratic forms. *Developments in Mathematics*, 31:255–298, 2013.

Received 27 Feb 2020. Revised 31 Jul 2020.

KIRSTEN EISENTRÄGER: eisentra@math.psu.edu
Department of Mathematics, The Pennsylvania State University, University Park, PA, United States

SEAN HALLGREN: hallgren@cse.psu.edu
Department of Computer Science and Engineering, Penn State University, University Park, PA, United States

CHRIS LEONARDI: cfoleona@uwaterloo.ca
Department of Combinatorics and Optimization, The University of Waterloo, Waterloo, ON, Canada

TRAVIS MORRISON: travis.morrison@uwaterloo.ca
Institute for Quantum Computing, The University of Waterloo, Waterloo, ON, Canada

JENNIFER PARK: park.2720@osu.edu
Department of Mathematics, The Ohio State University, Columbus, OH, United States

VOLUME EDITORS

Stephen D. Galbraith
Mathematics Department
University of Auckland
New Zealand
https://orcid.org/0000-0001-7114-8377

The cover image is based on an illustration from the article "Supersingular curves with small noninteger endomorphisms", by Jonathan Love and Dan Boneh (see p. 9).

Electronic copies can be obtained free of charge from http://msp.org/obs/4 and printed copies can be ordered from MSP (contact@msp.org).

The Open Book Series is a trademark of Mathematical Sciences Publishers.

ISSN: 2329-9061 (print), 2329-907X (electronic)

ISBN: 978-1-935107-07-1 (print), 978-1-935107-08-8 (electronic)

First published 2020.

---

msp

MATHEMATICAL SCIENCES PUBLISHERS

798 Evans Hall #3840, c/o University of California, Berkeley CA 94720-3840

contact@msp.org **http://msp.org**

THE OPEN BOOK SERIES 4

Fourteenth Algorithmic Number Theory Symposium

The Algorithmic Number Theory Symposium (ANTS), held biennially since 1994, is the premier international forum for research in computational and algorithmic number theory. ANTS is devoted to algorithmic aspects of number theory, including elementary, algebraic, and analytic number theory, the geometry of numbers, arithmetic algebraic geometry, the theory of finite fields, and cryptography.

This volume is the proceedings of the fourteenth ANTS meeting, which took place 29 June to 4 July 2020 via video conference, the plans for holding it at the University of Auckland, New Zealand, having been disrupted by the COVID-19 pandemic. The volume contains revised and edited versions of 24 refereed papers and one invited paper presented at the conference.

TABLE OF CONTENTS