

Supersingular isogeny graphs and endomorphism rings: reductions and solutions^{*}

Kirsten Eisenträger^{1**}, Sean Hallgren^{2***}, Kristin Lauter³, Travis Morrison^{1†},
and Christophe Petit⁴

¹ The Pennsylvania State University
Department of Mathematics

² The Pennsylvania State University
Department of Computer Science and Engineering

³ Microsoft Research

⁴ University of Birmingham

Abstract. In this paper, we study several related computational problems for supersingular elliptic curves, their isogeny graphs, and their endomorphism rings. We prove reductions between the problem of path finding in the ℓ -isogeny graph, computing maximal orders isomorphic to the endomorphism ring of a supersingular elliptic curve, and computing the endomorphism ring itself. We also give constructive versions of Deuring’s correspondence, which associates to a maximal order in a certain quaternion algebra an isomorphism class of supersingular elliptic curves. The reductions are based on heuristics regarding the distribution of norms of elements in quaternion algebras.

We show that conjugacy classes of maximal orders have a representative of polynomial size, and we define a way to represent endomorphism ring generators in a way that allows for efficient evaluation at points on the curve. We relate these problems to the security of the Charles-Goren-Lauter hash function. We provide a collision attack for special but natural parameters of the hash function and prove that for general parameters its preimage and collision resistance are also equivalent to the endomorphism ring computation problem.

1 Introduction

The recent search for new “post-quantum” cryptographic primitives and the ongoing international PQC competition sponsored by NIST has motivated a

^{*} This paper is the result of a merge of [EHM17] and [PL17].

^{**} The first author was partially supported by National Science Foundation awards DMS-1056703 and CNS-1617802, and by the National Security Agency (NSA) under Army Research Office (ARO) contract number W911NF-12-1-0541.

^{***} The second author was partially supported by National Science Foundation awards CNS-1617802 and CCF-1618287, and by the National Security Agency (NSA) under Army Research Office (ARO) contract number W911NF-12-1-0541.

[†] The fourth author was partially supported by National Science Foundation grants DMS-1056703 and CNS-1617802.

new era of research in the mathematics of cryptography. Ideas for cryptographic primitives based on hard mathematical problems are being actively proposed and examined. This paper focuses on supersingular isogeny-based cryptography, and in particular on the hardness of computing endomorphism rings of supersingular elliptic curves and its possible applications in cryptography.

In 2006, Charles, Goren, and Lauter [CGL06,CGL09] introduced the hardness of finding paths in Supersingular Isogeny Graphs into cryptography and used it for constructing cryptographic hash functions. In the CGL hash function, preimage resistance relies on the hardness of computing certain ℓ -power isogenies (for ℓ a small prime) between supersingular elliptic curves. Since then, this problem and related hard problems have been used as the basis for key exchange protocols [JDF11], signature schemes [YAJ⁺17,GPS17], and public key encryption [DFJP14]. There is also a submission [ACC⁺17] to the PQC standardization competition based on supersingular isogeny problems. While polynomial-time quantum algorithms are known for attacking widely deployed public key cryptosystems such as RSA and Elliptic Curve Cryptography (ECC), there are currently no known subexponential quantum attacks against these supersingular isogeny graph-based schemes.

In the supersingular case three problems have emerged as potential computational hardness assumptions related to the above systems. The first is computing isogenies between supersingular elliptic curves, the second one is computing the endomorphism ring of a supersingular elliptic curve, and the third is to compute a maximal order isomorphic to the endomorphism ring of a supersingular elliptic curve. In order to develop confidence that these new systems are secure against quantum computers, it is important to understand these problems, their relationships, and how they relate to the cryptosystems. The natural way to do this is to give polynomial-time reductions between the problems when possible, and there are heuristics for doing this [Koh96], [KLPT14]. However, one quickly runs into problems when attempting to find efficient reductions. For example, the main parameter for these problems is a large prime p , and it is not obvious that the endomorphism ring of an elliptic curve even has a basis with a representation size that is polynomial in $\log p$. The same problem exists for maximal orders.

The computational hardness assumption introduced in [CGL09] which underlies the security of Supersingular Isogeny Graph-based cryptography can be equivalently described as finding paths in the isogeny graph or as producing an ℓ -power isogeny (for ℓ a small prime) between two given supersingular elliptic curves. However, there exists another language to describe this problem, thanks to Deuring's correspondence [Deu41], which establishes (non-constructively) a one-to-one correspondence between supersingular j -invariants and maximal orders in a quaternion algebra, up to some equivalence relations. Following this correspondence, path-finding in the Supersingular Isogeny Graph can be translated, in theory, into a problem involving maximal orders in quaternion algebras which was solved in [KLPT14]. So this motivates the problem of finding explicit versions of Deuring's correspondence, namely constructive, efficient algorithms

to translate j -invariants into maximal orders in the quaternion algebra and conversely.

1.1 Contributions

Section 2 introduces preliminary material on supersingular elliptic curves and the arithmetic of quaternion algebras, and we recall some well-known facts from [Mes86, Piz80, Wat69], with an emphasis on explicit computations and representations. We state several problems for supersingular elliptic curves in Section 3. In Section 4, we show that an isomorphism class of maximal orders in a quaternion algebra has at least one representative of polynomial size. Since computing maximal orders is one of the central problems we consider, such a theorem is necessary to have meaningful polynomial-time reductions. The results in Section 4 are conditional on GRH but do not use any heuristics. In Section 6.4, we construct the quaternion algebra analogue of a factorization of an isogeny of ℓ -power degree into degree ℓ isogenies. The results in that section do not use any heuristics and are unconditional. The construction of Section 6.4 is used in our reductions between algorithms involving maximal orders and paths in the ℓ -isogeny graph in Sections 5 and 6.

Section 5 reduces three hard problems in supersingular graphs to each other: a constructive version of Deuring’s correspondence from j -invariants to maximal orders in $B_{p,\infty}$ (Problem 2); the endomorphism ring computation problem (Problem 3); and the preimage and collision resistance of the Charles-Goren-Lauter hash function, for a randomly chosen initial vertex. These reductions rely on various heuristic assumptions underlying the quaternion ℓ -isogeny algorithm of [KLPT14] and its powersmooth version described explicitly in [GPS17], along with new heuristics about using loops in the isogeny graph to generate endomorphism rings.

Section 6 shows that constructing paths in the ℓ -isogeny graph reduces to a different type of endomorphism ring computation. However, instead of just requiring an algorithm for computing the maximal order, one also needs to know how the generators of the order act on the ℓ -torsion of the curve. Thus this section contains a reduction to a harder problem. On the other hand, this section removes some of the heuristics used in Section 5. More precisely, the reductions in Section 5 use both the quaternion ℓ -isogeny algorithm and its powersmooth version, whereas the reductions in Section 6 only use the quaternion ℓ -isogeny algorithm [KLPT14].

Intuitively these heuristics say that numbers generated by the norm form of a quaternion algebra in the algorithm behave in the same way as random numbers of the same size, with respect to their factorization patterns.

Section 7 provides a (heuristic) probabilistic polynomial-time algorithm for computing the Deuring correspondence in one direction, and a partial attack on a special case of the Charles-Goren-Lauter hash function. In Section 8, we start by defining the notion of a compact representation of an endomorphism, which has as a requirement that it has size polynomial in $\log p$. We prove that every endomorphism ring has a basis specified by compact representations, and that

we can evaluate the endomorphism at points using the representation. We then show that the endomorphism problem reduces to computing a maximal order and the Action-on- ℓ -Torsion problem.

1.2 Related work

The endomorphism ring computation problem and constructive versions of Deuring’s correspondence have been studied in the past independently of their cryptographic applications, and all known algorithms for these problems have required exponential time. Computing the endomorphism ring of a supersingular elliptic curve was first studied by Kohel [Koh96, Theorem 75], who gave an approach for finding four linearly independent endomorphisms, generating a finite-index subring of $\text{End}(E)$. The algorithm was based on finding loops in the ℓ -isogeny graph of supersingular elliptic curves, and the running time of the probabilistic algorithm is $O(p^{1+\varepsilon})$. Another problem that has been considered is to list all isomorphism classes of supersingular elliptic curves together with a description of the maximal order in a quaternion algebra that is isomorphic to $\text{End}(E)$. This was done in [Cer04, LM04] and improved in [CG14, Section 5.2]. However, this approach is necessarily exponential in $\log p$ because there are roughly $\lfloor p/12 \rfloor$ isomorphism classes of supersingular elliptic curves.

The problem of computing isogenies between supersingular elliptic curves has also been studied, both in the classical setting [DG16, Section 4] where the complexity of the algorithm is $\tilde{O}(p^{1/2})$, and in the quantum setting [BJS14], where the complexity is $\tilde{O}(p^{1/4})$.

A signature scheme based on endomorphism ring computation is given in [GPS17, Section 4], where the secret key is a maximal order isomorphic to the endomorphism ring of a supersingular elliptic curve. While the scheme in [DFJP14] had to reveal auxiliary points, this is not necessary in this scheme.

Recently there have been several partial attacks on isogeny-based protocols (see [GPST16, Ti17, GW17]). These attacks target the key exchange protocol of Jao-De Feo [JDF11] in specific attack models, such as fault attacks, and are complementary to our work.

2 Preliminaries

2.1 Background on elliptic curves

Elliptic curves and isogenies By an elliptic curve E over a field k of characteristic $p > 3$ we mean a curve with equation $E : y^2 = x^3 + Ax + B$ for some $A, B \in k$ satisfying $4A^3 + 27B^2 \neq 0$. The points of E are the points (x, y) satisfying the curve equation, together with the point at infinity. These points form an abelian group. The j -invariant of an elliptic curve given as above is $j(E) = \frac{256 \cdot 27 \cdot A^3}{4A^3 + 27B^2}$. Two elliptic curves E, E' defined over a field k have the same j -invariant if and only if they are isomorphic over the algebraic closure of k . We

write $j(E)$ for the j -invariant of E . Given a j -invariant $j \neq 0, 1728$, we write $E(j)$ for the curve defined by the equation

$$y^2 + xy = x^3 - \frac{36}{j - 1728}x - \frac{1}{j - 1728}.$$

Such a curve can be put into a short Weierstrass equation $y^2 = x^3 + Ax + B$. We also write $E(0)$ and $E(1728)$ for the curves with equations $y^2 = x^3 + 1$ and $y^2 = x^3 + x$ respectively.

Let E_1 and E_2 be elliptic curves defined over a field k of positive characteristic p . An *isogeny* $\varphi : E_1 \rightarrow E_2$ defined over k is a non-constant rational map defined over k which is also a group homomorphism from $E_1(k)$ to $E_2(k)$ [Sil09, III.4]. The degree of an isogeny is its degree as a rational map. When the degree d of the isogeny φ is coprime to p , then φ is separable and the kernel of φ is a subgroup of the points on E_1 of size d . Every isogeny of degree n greater than one can be factored into a composition of isogenies of prime degrees such that the product of the degrees equals n . If $\psi : E_1 \rightarrow E_2$ is an isogeny of degree d , the *dual isogeny* of ψ is the unique isogeny $\hat{\psi} : E_2 \rightarrow E_1$ satisfying $\psi\hat{\psi} = [d]$, where $[d] : E_1 \rightarrow E_1$ is the multiplication-by- d map.

We can describe an isogeny via its kernel. Given an elliptic curve E and a finite subgroup H of E , there is, up to isomorphism a unique isogeny $\varphi : E \rightarrow E'$ having kernel H (see [Sil09, III.4.12]). Hence we can describe an isogeny of E to some other elliptic curve by giving its kernel. We can compute equations for the isogeny from its kernel by using Vélú's formula [Vél71].

Endomorphisms and supersingular versus ordinary curves An isogeny of an elliptic curve E to itself is called an endomorphism of E . If E is defined over some finite field \mathbb{F}_q , then an endomorphism of E will be defined over a finite extension of \mathbb{F}_q . The set of endomorphisms of E defined over \mathbb{F}_q together with the zero map form a ring under the operations addition and composition. It is called the endomorphism ring of E , and is denoted by $\text{End}(E)$. When E is defined over a finite field, then $\text{End}(E)$ is isomorphic either to an order in a quadratic imaginary field or to an order in a quaternion algebra. In the first case we call E an *ordinary elliptic curve*. An elliptic curve whose endomorphism ring is isomorphic to an order in a quaternion algebra is called a *supersingular elliptic curve*. Every supersingular elliptic curve over a field of characteristic p has a model that is defined over \mathbb{F}_{p^2} because the j -invariant of such a curve is in \mathbb{F}_{p^2} .

ℓ -power isogenies between supersingular elliptic curves Let E, E' be two supersingular elliptic curves defined over \mathbb{F}_{p^2} . It is a fact that for each prime $\ell \neq p$, E and E' are connected by a chain of isogenies of degree ℓ [Mes86]. By [Koh96, Theorem79], E and E' can be connected by m isogenies of degree ℓ , where $m = O(\log p)$. So any two supersingular elliptic curves can be connected by an isogeny of degree ℓ^m with $m = O(\log p)$. If $\ell = O(\log p)$ is a fixed prime,

then any ℓ -isogeny in the chain above can either be specified by rational maps or by giving the kernel of the isogeny, and both of these representations will have polynomial size in $\log p$. By Vélú's formula, and since $\ell = O(\log p)$, there is an efficient way to go back and forth between these two representations.

2.2 Quaternion algebras, $B_{p,\infty}$ and the Deuring correspondence

Quaternion algebras 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,

$$H(a, b) = \mathbb{Q} + \mathbb{Q}i + \mathbb{Q}j + \mathbb{Q}ij.$$

It is a fact that any quaternion algebra over \mathbb{Q} can be written in this form. Now let $B_{p,\infty}$ be the unique quaternion algebra over \mathbb{Q} that is ramified exactly at p and ∞ . Then $B_{p,\infty}$ is a definite quaternion algebra, so $B_{p,\infty} = H(a, b)$ for some $a, b \in \mathbb{Q}^\times$, and one can show a and b can be chosen to be negative integers. For example, when $p \equiv 3 \pmod{4}$, then $B_{p,\infty} = H(-p, -1)$.

There is a *canonical involution* on $B_{p,\infty}$ which sends an element $\alpha = a_1 + a_2i + a_3j + a_4ij$ to $\bar{\alpha} := a_1 - a_2i - a_3j - a_4ij$. Define the *reduced trace* of an element α as above to be

$$\text{Trd}(\alpha) = \alpha + \bar{\alpha} = 2a_1,$$

and the *reduced norm* to be

$$\text{Nrd}(\alpha) = \alpha\bar{\alpha} = a_1^2 - aa_2^2 - ba_3^2 + aba_4^2.$$

We say that Λ is a *lattice* in $B_{p,\infty}$ if $\Lambda = \mathbb{Z}x_1 + \cdots + \mathbb{Z}x_4$ and the elements x_1, \dots, x_4 are a vector space basis for $B_{p,\infty}$.

If $I \subseteq B_{p,\infty}$ is a lattice, the reduced norm of I , $\text{Nrd}(I)$, is the positive generator of the fractional \mathbb{Z} -ideal generated by $\{\text{Nrd}(\alpha) : \alpha \in I\}$. The quaternion algebra $B_{p,\infty}$ is an inner product space with respect to the bilinear form

$$\langle x, y \rangle = \frac{\text{Nrd}(x + y) - \text{Nrd}(x) - \text{Nrd}(y)}{2}.$$

The basis $\{1, i, j, ij\}$ is an orthogonal basis with respect to this inner product.

Orders in $B_{p,\infty}$ and representation of elements in $B_{p,\infty}$ An *order* \mathcal{O} of $B_{p,\infty}$ is a subring of $B_{p,\infty}$ which is also a lattice, and if \mathcal{O} is not properly contained in any other order, we call it a *maximal order*. For a lattice $I \subseteq B_{p,\infty}$ we define

$$\mathcal{O}_R(I) := \{x \in B_{p,\infty} : Ix \subseteq I\}$$

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 $B_{p,\infty}$ and $I \subseteq \mathcal{O}$ is a left ideal of \mathcal{O} , then $\mathcal{O}_R(I)$ is also a maximal order. Given any two maximal orders $\mathcal{O}, \mathcal{O}'$, there is a lattice

$I \subseteq B_{p,\infty}$ such that $\mathcal{O}_L(I) = \mathcal{O}$ and $\mathcal{O}_R(I) = \mathcal{O}'$; we say that I connects \mathcal{O} and \mathcal{O}' .

An element $\beta \in B_{p,\infty}$ is represented as a coefficient vector (a_1, a_2, a_3, a_4) in \mathbb{Q}^4 such that $\beta = a_1 + a_2i + a_3j + a_4ij$ in terms of the basis $\{1, i, j, ij\}$ for $B_{p,\infty}$. This will be used for specifying basis elements of maximal orders \mathcal{O} and elements of left ideals I of \mathcal{O} .

The Deuring correspondence and describing isogenies via kernel ideals

For a detailed overview of the information in this section, see Chapter 42 in [Voi]. Let E be a supersingular elliptic curve defined over \mathbb{F}_{p^2} . In [Deu41] Deuring proved that the endomorphism ring of E is isomorphic to a maximal order in $B_{p,\infty}$. Under this isomorphism, degrees and traces of endomorphisms correspond to norms and traces of quaternions. The correspondence between isomorphism classes of supersingular elliptic curves and maximal orders is often referred to as Deuring's correspondence.

Fix E , a supersingular elliptic curve over \mathbb{F}_{p^2} . We can associate to each pair (E', ϕ) with ϕ an isogeny $E \rightarrow E'$ of degree n a left $\text{End}(E)$ -ideal $I = \text{Hom}(E', E)\phi$ of norm n , and it was shown in [Koh96, Section 5.3] that every left $\text{End}(E)$ -ideal arises in this way. We now describe how to construct an isogeny from a left $\text{End}(E)$ -ideal.

Let I be a nonzero integral left ideal of $\text{End}(E)$. Define $E[I]$ to be the scheme-theoretic intersection

$$E[I] = \bigcap_{\alpha \in I} \ker(\alpha).$$

Thus to each left ideal I of $\text{End}(E)$ there is an associated isogeny $\phi_I : E \rightarrow E/E[I]$. If $\text{Nrd}(I)$ is coprime to p , then

$$E[I] = \{P \in E(\overline{\mathbb{F}}_{p^2}) : \alpha(P) = 0 \quad \forall \alpha \in I\}.$$

2.3 Supersingular isogeny graphs

For any prime $\ell \neq p$, one can construct a so-called ℓ -isogeny graph, where each vertex is associated to a supersingular j -invariant, and an edge between two vertices is associated to a degree ℓ isogeny between the corresponding curves. Isogeny graphs are regular with regularity degree $\ell + 1$; they are directed graphs (unless $p \equiv 1 \pmod{12}$). Isogeny graphs are Ramanujan, i.e. they are optimal *expander graphs*, with the consequence that random walks on the graph quickly reach the uniform distribution [HLW06].

2.4 The Charles-Goren-Lauter hash function

The first cryptographic construction based on supersingular isogeny problems is a hash function proposed by Charles, Goren and Lauter [CGL09]. The security of this construction relies on the hardness of computing some isogenies of special degrees between two supersingular elliptic curves.

More precisely, consider an ℓ -isogeny graph over \mathbb{F}_{p^2} , where p is a “large” prime and ℓ is a “small” prime. The authors suggest to take $p \equiv 1 \pmod{12}$ to avoid some annoying backtracking issues. The message is first mapped into $\{0, \dots, \ell - 1\}^*$, with some padding if necessary. At each vertex, a deterministic ordering of the edges is fixed (this can be done by sorting the j -invariants of the $\ell + 1$ neighbors). An initial vertex j_0 is also fixed, as well as an initial incoming direction.

Given a message $(m_1, m_2, \dots, m_N) \in \{0, \dots, \ell - 1\}^*$, an edge adjacent to j_0 (excluding the incoming edge) is first chosen according to the value of m_1 , and the corresponding neighbor E_1 is computed. Then an edge of j_1 (excluding the edge between j_0 and j_1) is chosen according to the value of m_2 , and the corresponding neighbor j_2 is computed, etc. The final invariant j_N reached by this computation is mapped to $\{0, 1\}^n$ in some deterministic way (here $n \approx \log p$) and the value obtained is returned as the output of the hash function.

Clearly the function is preimage resistant if and only if, given two supersingular j -invariants j_1 and j_2 , it is computationally hard to compute a positive integer e and an isogeny $\varphi : E(j_1) \rightarrow E(j_2)$ of degree ℓ^e .

In this paper we give two new results on the security of this construction. On the one hand (Section 5.5), we show that for a randomly chosen starting point j_0 the function is preimage and collision resistant if and only if the endomorphism ring computation problem is hard: loosely speaking this means computing some endomorphisms of $E(j)$ but not necessarily of the correct norms. The interest of this result lies in that computing endomorphisms of elliptic curves is a natural problem to consider from an algorithmic number theory point of view, and it has indeed been studied since Kohel’s thesis in 1996. On the other hand (Section 7.2), we also show that the collision resistance problem is easy for some particular starting points.

2.5 Isogeny-based cryptography

A few years after Charles, Goren and Lauter designed their hash function, Jao and De Feo proposed a variant of the Diffie-Hellman protocol based on supersingular isogeny problems, which is now known as the supersingular isogeny key exchange protocol [JDF11]. We briefly describe it here in a way to encompass both the original parameters and the generalization recently suggested by Petit [Pet17].

The parameters include a large prime p , a supersingular curve E , and two coprime integers N_A and N_B . Alice and Bob select cyclic subgroups of E of order N_A and N_B , respectively; they compute the corresponding isogenies and they exchange the values of the end vertices, which are E/G_A and E/G_B , respectively. The shared key is the value $j(E/\langle G_A, G_B \rangle)$. This shared key could a priori not be computed by any party from E/G_A , E/G_B and their respective secret keys only, so Alice (resp. Bob) additionally sends the images of a basis of $E[N_B]$ by ϕ_A (resp. a basis of $E[N_A]$ by ϕ_B).

Jao-De Feo suggested to use $N_A = 2^{e_B} \approx p^{1/2} \approx N_B = 3^{e_B}$ such that $(p - 1)/N_A N_B$ is a small integer for efficiency reasons; in [Pet17] Petit argued

that choosing $N_A \approx N_B \approx p^2$ both powersmooth numbers is a priori better from a security point of view while preserving polynomial-time complexity for the protocol execution. It was shown by Gабraith-Petit-Shani-Ti [GPST16] that computing the endomorphism ring of E and E_A is sufficient to break the key exchange for the parameters suggested by Jao-De Feo. The argument uses the fact that isogenies generated for Jao-De Feo’s parameters are of relatively small degree, and this does not seem to apply to Petit’s parameters.

The security of Jao-De Feo’s protocol relies on the hardness of computing isogenies of a given degree between two given curves, when provided in addition with the action of the isogeny on a large torsion group. This problem is not known to be equivalent to the endomorphism ring computation problem. Recent results by Petit [Pet17] show that revealing the action of isogenies on a torsion group does make some isogeny problems easier to solve, though at the moment his techniques do not apply to Jao-De Feo’s original parameters. We *believe* that the security of the key exchange protocol lies between these hard and easy problems, but leave its study to future work.

The interest in isogeny-based cryptography has recently increased in the context of NIST’s call for post-quantum cryptography algorithms [NIS16], and a submitted proposal was based on isogeny-based cryptography [ACC⁺17]. At the moment the best algorithms to solve supersingular isogeny problems all require exponential time in the security parameter, even when including quantum algorithms. Besides the hash function and the key exchange protocols, there are now constructions based on isogeny problems for public key encryption, identification protocols and signatures [DFJP14,YAJ⁺17,GPS17]. Constructions in the first two papers build on the key exchange protocol and rely on similar assumptions. The second signature scheme in [GPS17], however, only relies on the endomorphism computation problem.

3 Problem statements and heuristics

3.1 The Deuring Correspondence

The Deuring correspondence states that

$$\{\mathcal{O} \subseteq B_{p,\infty} \text{ maximal}\} / \simeq \leftrightarrow \{j \in \mathbb{F}_{p^2} : E(j) \text{ supersingular}\} / \text{Gal}(\mathbb{F}_{p^2}/\mathbb{F}_p)$$

is a bijective correspondence, given by associating a supersingular j -invariant to a maximal order in $B_{p,\infty}$ isomorphic to $\text{End}(E(j))$.

In this paper we will be interested in *constructing* Deuring’s correspondence for arbitrary maximal orders and supersingular j -invariants. This could a priori have different meanings, given by Problems 1 and 2 below.

Problem 1 (Constructive Deuring Correspondence.) *Given a maximal order $\mathcal{O} \subset B_{p,\infty}$, return a supersingular j -invariant such that the endomorphism ring of $E(j)$ is isomorphic to \mathcal{O} .*

We refer to the problem of computing a maximal order isomorphic to $\text{End}(E(j))$ for given a supersingular j -invariant as Problem MaxOrder or the “Inverse Deuring Correspondence.”

Problem 2 (MaxOrder) *Given p , the standard basis for $B_{p,\infty}$, and a supersingular elliptic curve E defined over \mathbb{F}_{p^2} , output vectors $\beta_1, \beta_2, \beta_3, \beta_4 \in B_{p,\infty}$ that form a \mathbb{Z} -basis of a maximal order \mathcal{O} in $B_{p,\infty}$ such that $\text{End}(E) \cong \mathcal{O}$. In addition, the output basis is required to have representation size polynomial in $\log p$.*

The j -invariant is naturally represented as an element of \mathbb{F}_{p^2} , and it is unique up to Galois conjugation. The maximal order is unique up to conjugation by an invertible quaternion element, and it can be described by a \mathbb{Z} -basis, namely four elements $1, \omega_2, \omega_3, \omega_4 \in B_{p,\infty}$ such that $\mathcal{O} = \mathbb{Z} + \omega_2\mathbb{Z} + \omega_3\mathbb{Z} + \omega_4\mathbb{Z}$. Choosing a Hermite basis makes this description unique.

In this paper we will provide a polynomial-time algorithm for Problem 1 (Section 7.1). We will also provide explicit connections between Problem 1 and the endomorphism ring computation problem, where instead of a maximal order in $B_{p,\infty}$ one needs to output a basis for $\text{End}(E(j))$.

3.2 The endomorphism ring computation problem

Given an elliptic curve, it is natural to ask to compute its endomorphism ring.

Problem 3 (Endomorphism ring computation problem.) *Given p and a supersingular j -invariant j , compute the endomorphism ring of $E(j)$.*

The endomorphism ring can be returned as four rational maps that form a \mathbb{Z} -basis with respect to scalar multiplication (in fact 3 maps, since one of these maps can always be chosen equal to the identity map). The maps themselves can usually not be returned in their canonical expression as rational maps, as in general this representation will require a space larger than the degree, and the degrees can be as big as p .

Various representations of the maps are a priori possible. We believe that any valid representation should be *concise* and *useful*, in the sense that it must require a space polynomial in $\log p$ to store, and it must allow the evaluation of the maps at arbitrary elliptic curve points in a time polynomial in both $\log p$ and the space required to store those points. To the best of our knowledge these two conditions are sufficient for all potential applications of Problem 3. When its degree is a smooth number, an endomorphism can be efficiently represented as a composition of small degree isogenies. In Section 5.1 we will consider a more general representation.

A first approximation to a solution to Problem 3 was provided by Kohel in his PhD thesis [Koh96], and later improved by Galbraith [Gal99] using a birthday argument. The resulting algorithm explores a tree in an ℓ -isogeny graph (for some small integer ℓ) until a collision is found, corresponding to an endomorphism. The

expected cost of this procedure is $O(\sqrt{p})$ times a polynomial in $\log p$. Repeating this procedure a few times, possibly with different values of ℓ , we obtain a set of endomorphisms which generate a subring of the whole endomorphism ring. The endomorphism ring computation problem was also considered in [DG16] for curves defined over \mathbb{F}_p . The identification protocol and signature schemes developed in [GPS17] explicitly rely on its potential hardness for security.

We observe that Problems 2 and 3 take the same input, and their outputs are also “equal” in the sense they are isomorphic. For this reason the two problems have sometimes been referred to interchangeably. In particular, a solution to Problem 2 does not a priori provide a useful description of the endomorphism ring so that one can evaluate endomorphisms at given points. Similarly, a solution to Problem 3 does not a priori provide a \mathbb{Z} -basis for an order in $B_{p,\infty}$, and this is necessary to apply the algorithms of [KLPT14].

It turns out that the two problems are equivalent: in Sections 5.1 and 5.4, we provide efficient algorithms to go from a representation of the endomorphism ring as a \mathbb{Z} basis over \mathbb{Q} to a representation as rational maps and conversely.

In Sections 6 and 8, our reductions will involve the following problem.

Problem 4 (Action-on- ℓ -Torsion) *Given p , a supersingular elliptic curve E defined over \mathbb{F}_{p^2} , and four elements $\{\beta_1, \beta_2, \beta_3, \beta_4\}$ in a maximal order \mathcal{O} of $B_{p,\infty}$ such that there exists an isomorphism $\iota : \text{End}(E) \rightarrow \mathcal{O}$, output eight pairs of points on E , $(P_1, Q_{1r}), (P_2, Q_{2r})$ ($r = 1, \dots, 4$) such that P_1, P_2 form a basis for the ℓ -torsion $E[\ell]$ of E , and such that $Q_{1r} = \iota^{-1}(\beta_r)(P_1)$ and $Q_{2r} = \iota^{-1}(\beta_r)(P_2)$ for $r = 1, \dots, 4$.*


The combination of this problem with Problem MaxOrder is, intuitively, to ask for both the algebraic structure of $\text{End}(E)$ (by asking for generators in $B_{p,\infty}$ for a maximal order $\mathcal{O} \simeq \text{End}(E)$), along with a small amount of geometric information, meaning asking for how those generators act as endomorphisms on $E[\ell]$.

Finally, we will be relating these various endomorphism ring problems to pathfinding in the ℓ -isogeny graph, which we often refer to as preimage resistance for the Charles-Goren-Lauter hash function or Problem ℓ -PowerIsogeny.

Problem 5 (ℓ -PowerIsogeny) *Given a prime p , along with two supersingular elliptic curves E and E' over \mathbb{F}_{p^2} , output an isogeny from E to E' represented as a chain of k isogenies whose degrees are ℓ .*

Since E is given as $y^2 = x^3 + ax + b$ with $a, b \in \mathbb{F}_{p^2}$, the input size for this problem is $O(\log p)$. By Section 2.1, the representation size of the output is also polynomial in $\log p$, if $\ell \in O(\log p)$ and the isogenies are represented by rational maps.

Below we map out the various reductions in this paper. An arrow represents the reduction from one problem to another, and its label indicates the algorithm or theorem giving that reduction.



3.3 Heuristics

Our reductions require several heuristics related to the distribution of numbers represented by certain quadratic forms and on isogeny graphs. When we refer to plausible heuristic assumptions, we mean one or more of the following:

1. We assume the heuristics used in [KLPT14], which can be summarized as saying that the distribution of outputs of quadratic forms arising from the norm form of a maximal order in $B_{p,\infty}$ is approximately like the uniform distribution on numbers of the same size.
2. We also assume the heuristics used in [GPS17] on representing powersmooth numbers by these quadratic forms.
3. We assume that the endomorphism ring of an elliptic curve can be generated by endomorphisms arising from loops in the ℓ -isogeny graph. In particular, we assume that given a suborder \mathcal{O}' of a maximal order \mathcal{O} such that \mathcal{O}' is generated by loops in an ℓ -isogeny graph, the probability that a randomly generated loop in the graph is in \mathcal{O}' is inversely proportional to $[\mathcal{O} : \mathcal{O}']$.

4 Efficient computations with maximal orders and their ideals

One of the main problems we consider in this paper is computing a maximal order associated to an elliptic curve E . The following sections will show that computing isogenies and computing endomorphisms reduces to computing maximal orders, together with a problem about ℓ -torsion action. In this section we show that maximal orders have polynomial-representation size, so that the reductions are meaningful. We will also show that the representation size of ideals inside these orders is related to their norms. Maximal orders are inside the algebra $B_{p,\infty}$, so we start with that.

Let p be a prime. In Proposition 5.1 of [Piz80] it is shown that $B_{p,\infty} = H(-1, -1)$ if $p = 2$, $B_{p,\infty} = H(-1, -p)$ if $p \equiv 3 \pmod{4}$, $B_{p,\infty} = H(-2, -p)$ if $p \equiv 5 \pmod{8}$, and $B_{p,\infty} = H(-q, -p)$ if $p \equiv 1 \pmod{8}$, where $q \equiv 3 \pmod{4}$ is prime and p is not a square modulo q .

So given p , we choose a and b as above (depending on the congruence class of p) such that $B_{p,\infty} = H(a, b)$. We obtain a basis $1, i, j, ij$ for $B_{p,\infty}$ such that $i^2 = a$ and $j^2 = b$. We refer to this as the *standard basis* of $B_{p,\infty}$. As stated in

Section 2.2, we represent elements of $B_{p,\infty}$ as their coefficient vectors in \mathbb{Q}^4 with respect to the standard basis.

To reduce problems to Problem MaxOrder in polynomial time, one requirement is that in every conjugacy class there is a maximal order that has a basis with representation size that is polynomial in $\log p$. Since a prime p is given, and E is given as $y^2 = x^3 + ax + b$ with $a, b \in \mathbb{F}_{p^2}$, the input size for this problem is $O(\log p)$.

To show that there is a maximal order that has a polynomial representation size, we first show this is true for a special maximal order \mathcal{O}_0 and then express all other classes of maximal orders as right orders $\mathcal{O}_R(I)$ for a left ideal I of \mathcal{O}_0 . Since every left ideal class of \mathcal{O}_0 contains an ideal whose reduced norm is $O(p^2)$, it will follow that in each conjugacy class of maximal orders, there is one with polynomial representation size.

As mentioned above, Pizer [Piz80] gave the following explicit description of $B_{p,\infty}$ for all p along with a basis for one maximal order.

Proposition 1. *Let $p > 2$ be a prime. Then we can define $B_{p,\infty}$ and a maximal order \mathcal{O}_0 as follows:*

p	(a, b)	\mathcal{O}_0
$3 \pmod{4}$	$(-p, -1)$	$\langle 1, j, \frac{j+k}{2}, \frac{1+i}{2} \rangle$
$5 \pmod{8}$	$(-p, -2)$	$\langle 1, j, \frac{2-j+k}{4}, \frac{-1+i+j}{2} \rangle$
$1 \pmod{8}$	$(-p, -q)$	$\langle \frac{1+j}{2}, \frac{i+k}{2}, \frac{j+ck}{2}, k \rangle$

where in the last row $q \equiv 3 \pmod{4}$, $(p/q) = -1$ and c is some integer with $q|c^2p+1$. Assuming that the generalized Riemann hypothesis is true, there exists $q = O(\log^2 p)$ satisfying these conditions.

Proof. The information in the table follows from [Piz80, p 368–369]. The only thing we need to prove is the statement that when $p \equiv 1 \pmod{8}$ there exists a prime $q \equiv 3 \pmod{4}$ such that $\left(\frac{p}{q}\right) = -1$. Equivalently, we require that q be an unramified prime which does not split in either $K_1 = \mathbb{Q}(\sqrt{p})$ or $K_2 = \mathbb{Q}(\sqrt{-1})$.

This is equivalent to the condition that the Frobenius of q in $\text{Gal}(K_1K_2/\mathbb{Q})$ is the unique automorphism which restricts to the nontrivial automorphisms of $\text{Gal}(K_1/\mathbb{Q})$ and $\text{Gal}(K_2/\mathbb{Q})$. By [LO77], there is a prime q of size $O((\log |D|)^2)$ whose Frobenius is this element, where D is the absolute discriminant of the compositum K_1K_2/\mathbb{Q} . The absolute discriminant of K_1/\mathbb{Q} is p since $p \equiv 1 \pmod{4}$, and the absolute discriminant of K_2/\mathbb{Q} is -4 . Because $(4, p) = 1$, we have that $\mathcal{O}_{K_1K_2} = \mathcal{O}_{K_1}\mathcal{O}_{K_2}$, and using this, a computation shows that $D = \text{Disc}(K_1K_2/\mathbb{Q}) = 4^2p^2$. Hence $q = O(\log^2 p)$, as desired. \square

We stress that in all cases the maximal orders \mathcal{O}_0 given by Proposition 1 contain $\langle 1, i, j, k \rangle$ as a small index subring.

For the remainder of this section, fix such an order \mathcal{O}_0 together with the small basis $\{b_1, \dots, b_4\}$ as in Proposition 1. We will now show that ideals of \mathcal{O}_0 of norm N have representations of size polynomial in $\log(N)$ in terms of the basis $\{b_1, \dots, b_4\}$.

Lemma 1. *Let I be a left ideal of \mathcal{O}_0 . Then there is a \mathbb{Z} -basis $\{\alpha_1, \dots, \alpha_4\}$ of I , consisting of elements $\alpha_i \in \mathcal{O}_0$, such that the coefficients of the α_i expressed, in terms of the basis $\{b_1, b_2, b_3, b_4\}$ of \mathcal{O}_0 , are bounded by $\text{Nrd}(I)^2$.*

Proof. Let $\{\gamma_1, \dots, \gamma_4\}$ be a \mathbb{Z} -basis of I and write γ_i as $\gamma_i = \sum_j a_{ij} b_j$. Let $A = (a_{ij})$ be the matrix whose rows are the coefficients of γ_i . Let $H = UA$ where H is the (row-)Hermite normal form of A and $U \in \text{SL}_4(\mathbb{Z})$. Then the rows of H correspond to elements of \mathcal{O}_0 which generate I as a \mathbb{Z} -basis. Additionally, H is upper triangular, its diagonal elements satisfy $0 < h_{ii}$, and $h_{ij} < h_{jj}$ for $i < j$. We have $\text{Nrd}(I)^2 = \det(A) = \prod h_{ii}$ and hence all $h_{ij} < \text{Nrd}(I)^2$. This gives us the desired basis $\{\alpha_1, \dots, \alpha_4\}$. \square

We will now prove that every conjugacy class of maximal orders has a representative whose basis has representation size $O(\log p)$ when written in terms of the standard basis $1, i, j, ij$ for $B_{p,\infty}$.

For this, we will show that the reduced norm Nrd is the Euclidean norm on $B_{p,\infty} = H(-q, -p)$ considered as a lattice in \mathbb{R}^4 . (Here $q = 1, 2$ or a prime $\equiv 3 \pmod{4}$ that is not a square modulo p , depending on the congruence class of p .) We can view orders \mathcal{O} in $B_{p,\infty}$ as lattices in \mathbb{R}^4 , and we will relate the covolume of a lattice to its discriminant. This is similar to the number field case. Together with Minkowski's Theorem, this will give us the desired result.

Note that $B_{p,\infty} \otimes \mathbb{R}$ is isomorphic to \mathbb{H} , the Hamiltonians. Let $1, i', j', i'j'$ be the basis of \mathbb{H} with $i'^2 = j'^2 = -1$. Let

$$f : B_{p,\infty} \otimes \mathbb{R} \xrightarrow{\sim} \mathbb{H},$$

and let the isomorphism be given by $i \mapsto \sqrt{q}i'$, $j \mapsto \sqrt{p}j'$. Then the norm on \mathbb{H} , which is the (square of) the standard Euclidean norm on \mathbb{R}^4 , is just the reduced norm on the image of $B_{p,\infty}$ in \mathbb{H} under the isomorphism f . Let $\Lambda \subseteq \mathbb{R}^n$ be a lattice. Define its *covolume*, denoted $\text{Covol}(\Lambda)$, to be $\sqrt{\det(L^T L)}$ for any matrix L consisting of a basis for Λ . If $\mathcal{O} \subseteq B_{p,\infty}$ is a lattice, define its covolume to be $\text{Covol}(f(\mathcal{O}))$.

If a lattice $\mathcal{O} \subseteq B_{p,\infty}$ has generators β_1, \dots, β_4 , its *discriminant*, denoted $\text{Disc}(\mathcal{O})$, is $\det((\text{Trd}(\beta_i \overline{\beta_j})))$. If a lattice \mathcal{O} is a maximal order in $B_{p,\infty}$, then $\text{Disc}(\mathcal{O}) = p^2$.

Proposition 2. *Let \mathcal{O} be a lattice in $B_{p,\infty}$. Then $\text{Covol}(\mathcal{O})^2 = \frac{1}{16} \text{Disc}(\mathcal{O})$.*

Proof. This is Equation 2.2 of [CG14]. \square

We need the notion of a Minkowski-reduced basis. A basis $\{v_1, \dots, v_n\}$ of a lattice $\Lambda \subseteq \mathbb{R}^n$ is *Minkowski-reduced* if for $1 \leq k \leq n$,

$$\|v_k\|_2 \leq \left| \sum_{i=1}^n x_i \|v_i\|_2 \right|,$$

whenever x_1, \dots, x_n are coprime integers. Here $\|\cdot\|_2$ denotes the Euclidean norm. Given a lattice Λ in \mathbb{R}^n , define the *i th successive minimum of Λ* , $\lambda_i(\Lambda)$, to be the

smallest nonnegative, real number r such that there are i linearly independent lattice vectors of Λ contained in the closed ball of radius r centered at the origin. So $\lambda_1(\Lambda)$ is the length of a shortest nonzero vector of Λ . For $n \leq 4$, there is a basis v_1, \dots, v_n of Λ such that $\|v_i\|_2 = \lambda_i(\Lambda)$; see [NS09]. Such a basis is Minkowski-reduced. When we refer to a Minkowski-reduced basis, we will always assume we choose such a basis.

Theorem 1 (Minkowski's second theorem). *Let V denote the volume of the n -dimensional unit ball of \mathbb{R}^n . Then*

$$\frac{2^n \text{Covol}(\Lambda)}{n! V} \leq \prod_{i=1}^n \lambda_i(\Lambda) \leq \frac{2^n}{V} \text{Covol}(\Lambda).$$

Corollary 1. *Let p be a prime, and let \mathcal{O}_0 be the maximal order of $B_{p,\infty}$ as above. Let $I \subseteq \mathcal{O}_0$ be a left ideal and let $\mathcal{O} := \mathcal{O}_R(I)$. Let $\alpha_1, \dots, \alpha_4$ be a basis of \mathcal{O} such that $\|\alpha_i\|_2 = \lambda_i(\mathcal{O})$ for $i = 1, \dots, 4$. Then*

$$\prod_{i=1}^4 \text{Nrd}(\alpha_i) \leq \text{Disc}(\mathcal{O}) = p^2.$$

Proof. We use Minkowski's second theorem applied to \mathcal{O} , and the fact that by Proposition 2, $\text{Covol}(\mathcal{O})^2 = \text{Disc}(\mathcal{O})/16$. These two facts, together with $\text{Nrd}(\alpha) = \|f(\alpha)\|_2^2$ give us that

$$\prod \text{Nrd}(\alpha_i) = \prod \lambda_i(\mathcal{O})^2 \leq \frac{16}{\pi^4/4} \text{Disc}(\mathcal{O}) \leq p^2.$$

□

Now we prove the main theorem on representation sizes of maximal orders:

Theorem 2. *Every conjugacy class of maximal orders in $B_{p,\infty}$ has a \mathbb{Z} -basis x_1, \dots, x_4 with $\text{Nrd}(x_i) \in O(p^2)$. If we express x_r (for $1 \leq r \leq 4$) as a coefficient vector in terms of $1, i, j, ij$, then the rational numbers appearing have numerators and denominators whose representation size are polynomial in $\log p$.*

Proof. The map $[I] \rightarrow [\mathcal{O}_R(I)]$ is a surjection from left ideal classes of \mathcal{O}_0 to isomorphism classes of maximal orders of $B_{p,\infty}$; see [Gro87], page 116. Every left ideal class of \mathcal{O}_0 contains an ideal I with $\text{Nrd}(I) \in O(p^2)$; see [Vig80, Proposition 17.5.6]. Set $\mathcal{O} = \mathcal{O}_R(I)$ and let $\langle 1, x_2, x_3, x_4 \rangle$ be a Minkowski-reduced \mathbb{Z} -basis of \mathcal{O} . By Corollary 1, $\text{Nrd}(x_i) \leq p^2$, since each x_i is integral. Since $\mathcal{O} = \mathcal{O}_R(I)$, it follows that $x_i \text{Nrd}(I) \in I$. This implies that if we express x_i as a \mathbb{Q} -linear combination of the elements $1, i, j, ij$, then the denominators of the coefficients are divisors of $\text{Nrd}(I) \cdot 4q$ where $q = \text{Nrd}(j)$. The numerator of each coefficient is then bounded by $8pq \text{Nrd}(I)$: indeed, if a/b is a coefficient of x_r , ($1 \leq r \leq 4$), then $(a/b)^2 \leq \text{Nrd}(x_r) \leq p^2$. Then

$$|a| \leq pb \leq 4pq \text{Nrd}(I).$$

□

5 Equivalent hard problems in supersingular isogeny graphs

In this section we consider the following problems:

- A constructive version of Deuring’s correspondence, from j -invariants to maximal orders in $B_{p,\infty}$ (Problem 2).
- The endomorphism ring computation problem (Problem 3).
- The preimage and collision resistance of the Charles-Goren-Lauter hash function, for a randomly chosen initial vertex.

We show that all these problems are heuristically equivalent, in the sense that there exist efficient reductions from one problem to another under plausible heuristics assumptions.

The first two problems have the same inputs and in a sense their outputs are also equal, so it is perhaps no surprise to the reader that they are equivalent. However, the two problems differ in the way the output should be represented: as a maximal order in $B_{p,\infty}$ for Problem 2, and as four rational maps for Problem 3. Sections 5.1 and 5.4 below clarify the steps from one representation to the other.

It should also be clear intuitively that (heuristically at least) an algorithm to find preimages or collisions for the hash function can be used to compute endomorphism rings. The other implication is perhaps not as intuitive, and our solution crucially requires the tools developed in [KLPT14]. These reductions are discussed in Section 5.5 below.

5.1 Endomorphism ring computation is not harder than Inverse Deuring Correspondence

When $p \equiv 3 \pmod{4}$ the curve $y^2 = x^3 + x$ is supersingular with invariant $j = 1728$. This curve corresponds to a maximal order \mathcal{O}_0 with \mathbb{Z} -basis $\{1, i, \frac{1+k}{2}, \frac{i+j}{2}\}$ under Deuring’s correspondence, and there is an isomorphism of quaternion algebras $\theta : B_{p,\infty} \rightarrow \text{End}(E_0) \otimes \mathbb{Q}$ sending $(1, i, j, k)$ to $(1, \phi, \pi, \pi\phi)$ where $\pi : (x, y) \rightarrow (x^p, y^p)$ is the Frobenius endomorphism, and $\phi : (x, y) \rightarrow (-x, \iota y)$ with $\iota^2 = -1$. More generally, it is easy to compute j -invariants corresponding to the maximal orders given by Proposition 1.

Proposition 3. *There is a polynomial-time algorithm that given a prime $p > 2$, computes a supersingular j -invariant $j_0 \in \mathbb{F}_p$ such that $\text{End}(E(j_0)) \cong \mathcal{O}_0$ (where \mathcal{O}_0 is as given by Proposition 1 together with a map $\phi \in \text{End}(E(j_0))$) such that $\theta : B_{p,\infty} \rightarrow \text{End}(E(j_0)) \otimes \mathbb{Q} : (1, i, j, k) \rightarrow (1, \phi, \pi, \pi\phi)$ is an isomorphism of quaternion algebras.*

Proof. Let q be chosen such that $B_{p,\infty} = H(-q, -p)$ as in Proposition 1 and let R be the ring of integers of $\mathbb{Q}(\sqrt{-q})$. Consider Algorithm 3 below. Step 1 can be executed in time polynomial in $\log p$ using a modification of Bröker’s Algorithm 2.4 in [Brö09]: the cardinality of $\mathcal{J} := \{j \in \mathbb{F}_{p^2} : R \subseteq \text{End}(E(j))\}$ is equal to the class number h_{-q} of R , and this is bounded by q . To see this requires

a surjectivity and injectivity argument. Suppose $j \in \mathbb{F}_{p^2}$ is a supersingular j -invariant such that R embeds into $\text{End}(E(j))$. Then if $R = \mathbb{Z}[\alpha]$, by Deuring's Lifting Theorem [Lan87, Theorem 14, page 184] applied to $E(j)$ and α , there is an elliptic curve \tilde{E}/\mathbb{C} such that $\text{End}(\tilde{E}) \simeq R$ and a prime \mathfrak{p} of R dividing p such that $\tilde{E} \pmod{\mathfrak{p}} = E(j)$. Since \tilde{E} has complex multiplication by R , $j(\tilde{E})$ is a root of the Hilbert class polynomial of $\mathbb{Q}(\sqrt{-q})$. Because $E(j)$ is supersingular, p is inert in R and $\mathfrak{p} = pR$. We see that the map is injective because principal prime ideals of R split completely in H , and so the Hilbert class polynomial will have h_{-q} distinct roots modulo p . To compute ϕ in Step 3 one can simply compute all isogenies of degree q using Vélú's formulae and identify the one corresponding to an endomorphism. The map ϕ defines an isomorphism of quaternion algebras $\theta : B_{p,\infty} \rightarrow \text{End}(E(j_0)) \otimes \mathbb{Q} : (1, i, j, k) \rightarrow (1, \phi, \pi, \pi\phi)$. To perform the check in Step 4, one applies θ to the numerators of \mathcal{O}_0 basis elements, and check whether the resulting maps annihilate the D torsion, where D is the denominator. \square

Algorithm 3 *Computing the Deuring correspondence for special orders*

Input: A prime p .

Output: A supersingular j -invariant $j_0 \in \mathbb{F}_p$ such that $\mathcal{O}_0 \cong \text{End}(E(j_0))$, and an endomorphism $\phi \in \text{End}(E(j_0))$ such that $\text{Nrd}(\phi) = q$ and $\text{Trd}(\phi) = 0$.

1. Compute \mathcal{J} , a set of supersingular j -invariants such that for $j \in \mathcal{J}$, R_{-q} embeds into $\text{End}(E(j))$, where R_{-q} is the integer ring of $\mathbb{Q}(\sqrt{-q})$.
2. For $j \in \mathcal{J}$:
 - (a) Compute ϕ , an endomorphism of degree q of $E(j)$.
 - (b) If $\text{End}(E(j)) \cong \mathcal{O}_0$:
 - i. Return j and ϕ .

5.2 Quaternion ℓ -isogeny algorithm

The quaternion ℓ -isogeny problem was introduced and solved in [KLPT14] as a step forward in the cryptanalysis of the Charles-Goren-Lauter hash function.

We refer to [KLPT14, GPS17] for a full description of the algorithm and its powersmooth version as well as their analysis. For our purposes the following proposition will be sufficient.

Lemma 2. [KLPT14, GPS17] *Under various heuristic assumptions, there exist two polynomial-time algorithms that given I a left ideal of \mathcal{O}_0 , returns J another left ideal of \mathcal{O}_0 in the same class as I of norm N such that $N \approx p^{7/2}$. Moreover for the first algorithm we have $N = \prod p_i^{e_i}$ with $p_i^{e_i} < \log p$ and for the second algorithm we have $N = \ell^e$ for some integer e and some small prime ℓ .*

Interestingly, [GPS17] also proves that (after a minor tweak) the outputs of these algorithms only depend on the ideal class of their inputs and not on the particular ideal class representative.

Many of our algorithms and reductions below will use these algorithms as black boxes. Their correctness will therefore rely on the same heuristics, and possibly some more.

5.3 Translating \mathcal{O}_0 -ideals to isogenies

Let \mathcal{O}_0 be the maximal order given by Proposition 1, let E_0 be a corresponding supersingular elliptic curve, and let I be a left \mathcal{O}_0 -ideal of norm N such that I is not contained in $\mathcal{O}_0 m$ for any $m \in \mathbb{N}$. This ideal corresponds to an isogeny $\phi : E_0 \rightarrow E_1$ of degree N . This isogeny is uniquely defined by its kernel, which is a cyclic subgroup of order N in E_0 by Proposition 10. Following Waterhouse [Wat69] one can identify the correct subgroup by evaluating the maps corresponding to an \mathcal{O}_0 -basis at a generator of each subgroup. Moreover when N is composite, the kernel can be represented more efficiently as a product of cyclic subgroups whose orders are powers of primes, and similarly the isogenies are represented more efficiently as a composition of prime degree isogenies. The details of such an algorithm can be found in [GPS17], which also analyzes its complexity. The following proposition will be sufficient for our purposes.

Proposition 4. *There exists an algorithm which, given an \mathcal{O}_0 left ideal I of norm $N = \prod_i p_i^{e_i}$, returns an isogeny $\phi : E_0 \rightarrow E_1$ corresponding to this ideal through Deuring's correspondence. Moreover the complexity of this algorithm is polynomial in $\max_i p_i^{e_i}$.*

We stress that this translation algorithm requires us to know the endomorphism ring of E_0 , and that it is only efficient when $\max_i p_i^{e_i}$ is small.

Let us first assume that we have an efficient algorithm for Problem 2, returning a \mathbb{Z} basis for a maximal order as discussed above. Algorithm 4 below uses this algorithm to solve Problem 3.

Algorithm 4 *Reduction from Problem 3 to Problem 2*

Input: A supersingular j -invariant j .

Output: Four maps that generate $\text{End}(E(j))$.

1. Use an algorithm for Problem 2 to obtain a maximal order $\mathcal{O} \simeq \text{End}(E(j))$.
2. Compute an ideal I connecting \mathcal{O}_0 and \mathcal{O} .
3. Compute an ideal J with powersmooth norm in the same class as I .
4. Translate the ideal J into an isogeny $\varphi : E_0 \rightarrow E$.
5. Let N be the norm of J .
6. Let $1, \phi_2, \phi_3, \phi_4$ generate $\text{End}(E(j_0))$.
7. Let $1, \omega_2, \omega_3, \omega_4$ generate \mathcal{O} , and let $1, \omega_{2,0}, \omega_{3,0}, \omega_{4,0} \in \mathcal{O}_0$ correspond to $1, \phi_2, \phi_3, \phi_4$.
8. Find integers c_{ij} such that $\omega_i = \frac{\sum_j c_{ij} \omega_{j,0}}{N}$.
9. Return N, φ, c_{ij} implicitly representing the maps $\frac{\sum_{i=1}^4 c_{ij} \widehat{\varphi} \phi_i \varphi}{N}$ for each i .

The maps returned by Algorithm 4 are of the form $\phi = \frac{\sum_{i=1}^4 c_{ij} \widehat{\varphi} \phi_i \varphi}{N}$ where N is a smooth number, $c_{ij} \in \mathbb{Z}$, $\{\phi_i\}_{i=1,2,3,4}$ form a basis for the endomorphism ring of a special curve E_0 , and $\varphi : E_0 \rightarrow E(j)$ is an isogeny of degree N , given as a composition of isogenies of low degree. In Section 8 we define compact representations of endomorphisms, and the data given by Algorithm 4 define four compact representations. This is arguably not the most natural representation

of endomorphisms, but it still allows to efficiently evaluate them at arbitrary points, as shown by Algorithm 5 and Lemma 3 below. See Section 8 for a detailed definition of how to represent the output of this algorithm.

Algorithm 5 *Endomorphism evaluation*

Input: A curve E , an isogeny $\varphi : E_0 \rightarrow E$ with powersmooth degree N , and integers a, b, c, d defining an endomorphism $\phi = \frac{\varphi(a+b\phi_2+c\phi_3+d\phi_4)\widehat{\varphi}}{N} \in \text{End}(E)$.

Input: A point $P \in E$.

Output: $\phi(P)$.

1. Let $N = \prod_i p_i^{e_i}$ and let $m_i = N/p_i^{e_i}$.
2. For all i :
 - (a) Compute Q_i such that $p_i^{e_i} Q_i = P$.
 - (b) Compute $S_i = \varphi(a + b\phi_2 + c\phi_3 + d\phi_4)\widehat{\varphi}(Q_i)$
3. Compute S such that $S_i = m_i S$ for all i .
4. Return S .

Lemma 3. *Let $P \in E(K)$ with K an extension of \mathbb{F}_{p^2} . Assume that $\log N$ and $\max_i p_i^{e_i}$ are polynomial in $\log p$. Then Algorithm 5 computes $\phi(P)$ and can be implemented to run in time polynomial in $\log |K|$.*

Proof. We will first prove the correctness of the above algorithm. Let $\gamma := \varphi(a + b\phi_2 + c\phi_3 + d\phi_4)\widehat{\varphi}$, so $[N] \circ \phi = \gamma$. While the choice of Q_i in Step 2a is not unique, in Step 2b the point S_i is independent of the choice of Q_i , because of the calculation

$$S_i = \gamma(Q_i) = ([N] \circ \phi)(Q_i) = ([m_i] \circ \phi)(P).$$

We now show that the S in Step 3 exists, is unique, and equals $\phi(P)$. The above calculation showed $\phi(P)$ satisfies $m_i \phi(P) = S_i$. On the other hand, the point S also satisfies $m_i S = S_i$ for all i , so $\phi(P) - S \in E[m_i]$ for all i . Since $\gcd(\{m_1, \dots, m_k\}) = 1$, we have $\bigcap_{i=1}^k E[m_i] = \{0\}$. This implies that $S = \phi(P)$.

We can efficiently compute S in Step 3 as follows. Since the greatest common divisor of $\{m_1, \dots, m_k\}$ is 1, there are integers a_1, \dots, a_k such that $\sum_{j=1}^k a_j m_j = 1$. These integers can be efficiently computed with the extended Euclidean algorithm since $k = O(\log p)$. Define $S := \sum_{i=1}^k a_i S_i$. Observe that for $i \neq j$, we have

$$m_i S_j = \frac{N}{p_i^{e_i} p_j^{e_j}} p_j^{e_j} S_j = \frac{N}{p_i^{e_i} p_j^{e_j}} p_j^{e_j} \gamma(Q_j) = \frac{N}{p_i^{e_i} p_j^{e_j}} \gamma(P) = \frac{N}{p_i^{e_i} p_j^{e_j}} \gamma(p_i^{e_i} Q_i) = m_j S.$$

This implies that $m_i S_j = m_j S_i$. Now we calculate

$$m_i S = m_i \sum_{j=1}^k a_j S_j = S_i - \left(\sum_{j \neq i} a_j m_j S_i \right) + \sum_{j \neq i} m_i a_j S_j = S_i.$$

Although Q may lie in a very large extension of \mathbb{F}_{p^2} , each of the Q_i lies in a reasonably small extension, namely the extension degree is polynomial in

$\log p$. Note that S lies in an extension of K of degree at most 6 by Theorem 4.1 of [Wat69], so Step 3 is efficient. Step 2a involves some univariate polynomial factorization, a task that is polynomial in both the degree of the polynomial and the logarithm of the field size. In Step 2b the isogeny φ and its dual can be evaluated stepwise, and evaluating the map $a + b\phi_2 + c\phi_3 + d\phi_4$ at an arbitrary point involves 4 scalar multiplications, three additions and the evaluation of the maps $\phi_i \in \text{End}(E(j_0))$ at certain points. \square

Proposition 5. *Under plausible heuristic assumptions, the reduction in Algorithm 4 from Problem 3 to Problem 2 can be implemented to run in time polynomial in $\log p$.*

Proof. By Theorem 2, we may assume that the maximal order isomorphic to $\text{End}(E(j))$ has size polynomial in $\log p$. In Step 2, the ideal I can be computed with Algorithm 3.5 of [KV10]. This can be done in time polynomial in $\log p$ since \mathcal{O}_0 and \mathcal{O} have size polynomial in $\log p$. By Lemma 2 the output of Step 3 is an ideal of norm $N = \prod p_i^{e_i}$ such that $S = \max_i p_i^{e_i} = O(\log p)$. The translation algorithm runs in a time polynomial in S , hence in $\log p$. The other steps also run in polynomial time. \square

5.4 Inverse Deuring Correspondence is not harder than endomorphism ring computation

Let us now assume that we have an efficient algorithm for Problem 3, returning four maps generating the endomorphism ring, in some format that allows efficient evaluation of the maps at arbitrary points. Algorithm 6 below uses this algorithm and then constructs a sequence of linear transformations that map $1, \alpha, \beta, \gamma$ to four orthogonal maps $1, \iota, \lambda, \iota\lambda$ corresponding to $1, i, j, k \in B_{p,\infty}$. Composing the inverses of these maps then gives a \mathbb{Z} -basis for \mathcal{O} .

Algorithm 6 *Reduction from Problem 2 to Problem 3*

Input: A supersingular j -invariant j .

Output: A maximal order $\mathcal{O} \subset B_{p,\infty}$ such that $\text{End}(E(j)) \simeq \mathcal{O}$.

1. Use an algorithm for Problem 3 to obtain four maps $1, \alpha, \beta, \gamma$ which generate $\text{End}(E(j))$, in a format that allows efficient evaluation at elliptic curve points.
2. Compute the Gram matrix associated to the sequence $(1, \alpha, \beta, \gamma)$.
3. Find a rational invertible linear transformation sending $(1, \alpha, \beta, \gamma)$ to some $(1, \alpha', \beta', \alpha'\beta')$, where $1, \alpha', \beta', \alpha'\beta'$ generate an orthogonal basis for $B_{p,\infty}$ over \mathbb{Q} .
4. If the numerators and denominators of $\text{Nrd}(\alpha')$ and $\text{Nrd}(\beta')$ are not easy to factor:
 - (a) Apply a random invertible linear transformation to (α, β, γ) .
 - (b) Go to Step 3.
5. Find $a, b, c \in \mathbb{Q}$ such that $\text{Nrd}(\iota) = q$, where $\iota = a\alpha' + b\beta' + c\alpha'\beta'$.
6. Find a rational invertible linear transformation sending $(1, \alpha', \beta', \alpha'\beta')$ to $(1, \iota, \delta, \iota\delta)$ for some $\delta \in B_{p,\infty}$ where $1, \iota, \delta, \iota\delta$ generate an orthogonal basis for $B_{p,\infty}$ over \mathbb{Q} .

7. If the numerator and denominator of $\text{Nrd}(\delta)$ is not easy to factor:
 - (a) Apply a random invertible linear transformation to (α, β, γ) .
 - (b) Go to Step 3.
8. Find $a, b \in \mathbb{Q}$ such that $\text{Nrd}(\delta)(a^2 + b^2q) = p$. Let $\lambda = a\delta + b\iota\delta$.
9. Compute a rational invertible linear transformation sending $(1, \iota, \delta, \iota\delta)$ to $(1, \iota, \lambda, \iota\lambda)$.
10. Invert and compose all linear transformations to express $1, \alpha, \beta, \gamma$ in the basis $(1, \iota, \lambda, \iota\lambda)$, and deduce a basis of \mathcal{O} in $B_{p,\infty}$.
11. Return the basis of \mathcal{O} .

Let B be a bound on the degrees of the maps α, β, γ returned in Step 1 of Algorithm 6. We analyze the complexity of the algorithm through the following lemmas and proposition.

Lemma 4. *There exists an algorithm for Step 2 that runs in time polynomial in $\log p$ and $\log B$.*

Proof. Given two endomorphisms α, β , one can compute their inner product $\langle \alpha, \beta \rangle = \alpha\beta + \beta\bar{\alpha} \in \mathbb{Z}$ by evaluating it on an appropriate set of torsion points of small prime order, and then applying the Chinese Remainder Theorem, following a strategy similar to Schoof's point counting algorithm (see [Koh96, Theorem 81]). Applying this algorithm to every pair of maps from $(1, \alpha, \beta, \gamma)$ gives the result. \square

Lemma 5. *There exists an algorithm for Steps 3 and 6 that runs in time polynomial in $\log p$ and $\log B$.*

Proof. We focus on Step 3, and Step 6 is similar. Given the Gram matrix one can apply the Gram-Schmidt orthogonalization process to obtain a new basis $(1, \alpha', \beta', \gamma')$. It remains to show that $\alpha'\beta'$ is a scalar multiple of γ' so that we can normalize γ' to obtain the result. It suffices to show that $\alpha'\beta'$ is orthogonal to $1, \alpha'$ and β' . Indeed we have $\langle \alpha'\beta', 1 \rangle = \alpha'\beta' + \bar{\beta}'\bar{\alpha}' = \langle \alpha', \bar{\beta}' \rangle = -\langle \alpha', \bar{\beta}' \rangle = 0$; we have $\langle \alpha'\beta', \alpha' \rangle = \alpha'\beta'\bar{\alpha}' + \alpha'\bar{\beta}'\bar{\alpha}' = \text{Nrd}(\alpha') \text{Trd}(\beta') = 0$; and similarly $\langle \alpha'\beta', \beta' \rangle = \alpha'\beta'\bar{\beta}' + \beta'\bar{\beta}'\bar{\alpha}' = \text{Nrd}(\beta') \text{Trd}(\alpha') = 0$. \square

Lemma 6. *Given the factorizations of the numerators and denominators of both $\text{Nrd}(\alpha')$ and $\text{Nrd}(\beta')$, there exists an algorithm for Step 5 that runs in time polynomial in $\log p$ and $\log B$.*

Proof. Finding such $a, b, c \in \mathbb{Q}$ satisfying the condition amounts to finding $a', b', c', d \in \mathbb{Z}$ such that $a'^2 \text{Nrd}(\alpha') + b'^2 \text{Nrd}(\beta') + c'^2 \text{Nrd}(\alpha') \text{Nrd}(\beta') = d^2 q$. According to Simon [Sim05, Section 8] there is an algorithm to solve this Diophantine equation in polynomial time. \square

Lemma 7. *Given the factorizations of the numerator and of the denominator of $\text{Nrd}(\delta)$, there exists an algorithm for Step 8 that runs in time polynomial in $\log p$ and $\log B$.*

Proof. Note that $\langle \delta, \iota\delta \rangle$ is by construction the orthogonal space of $\langle 1, \iota \rangle$, and this space must contain an element of norm p , so the equation has a solution. Given factorizations for both the numerator and the denominator of δ one can use Cornacchia's algorithm [Cor08] to solve Step 8. \square

Proposition 6. *Under plausible heuristic assumptions, the reduction provided by Algorithm 6 can be implemented to run in polynomial time.*

Proof. In Steps 4 and 7 the algorithm requires that some numbers are easy to factor. In Step 4 we may expect these numbers to behave like random numbers of the same sizes. In Step 7, p must divide the numerator of $\text{Nrd}(\delta)$. We may expect that both the numerator and the denominator factor like random numbers of the same size. One can require all those numbers to be large primes, or a product of large primes and small cofactors, two properties that will be satisfied with a probability inversely proportional to a polynomial function of $\log p$. Steps 4a and 7a randomize α, β, γ so that we expect the conditions to be satisfied after a number of steps that is polynomial in $\log p$. By the four lemmas before we then expect that the whole reduction runs in a time polynomial in $\log p$. \square

The reduction provided by Algorithm 6 and its runtime analysis relies on several heuristics, namely the probability to obtain suitable norms in Steps 4 and 7 as discussed in the above proposition, and the runtime assumption of Simon's algorithm for Step 5.

5.5 Preimage and collision resistance of the CGL hash function

In this section we show that the hardness of the endomorphism ring computation problem is equivalent to the security of the Charles-Goren-Lauter hash function.

Proposition 7. *Assume there exists an efficient algorithm for the endomorphism ring computation problem. Then there is an efficient algorithm to solve the preimage and collision problems for the Charles-Goren-Lauter hash function.*

Proof. By standard arguments on hash functions it is enough to focus on preimage resistance. Our reduction of this problem to the endomorphism ring computation problem is given in Algorithm 7. Besides two black box calls to an algorithm for the endomorphism ring computation problem, it uses other efficient algorithms described in this paper, including Algorithm 4 to translate a description of an endomorphism ring as rational maps into a description of a maximal order in $B_{p,\infty}$, both the ℓ -power and the powersmooth versions of the quaternion isogeny algorithm, and the translation algorithm from ideals to isogenies. All these routines are efficient by the lemmas and propositions of this paper. By the results in Section 6.4, the algorithm is correct. \square

Algorithm 7 *Reduction from preimage resistance to endomorphism ring computation*

Input: Two supersingular j -invariants $j_s, j_t \in \mathbb{F}_{p^2}$.

Output: A sequence of j -invariants $j_s = j_0, j_1, \dots, j_e = j_t$ such that for any i there exists an isogeny of degree ℓ from $E(j_i)$ to $E(j_{i+1})$.

1. Compute $\text{End}(E(j_s))$ and $\text{End}(E(j_t))$.
2. Compute $\mathcal{O}_s \simeq \text{End}(E(j_s))$ and $\mathcal{O}_t \simeq \text{End}(E(j_t))$ with Algorithm 4.
3. Compute ideals I_s and I_t connecting \mathcal{O}_0 respectively to \mathcal{O}_s and \mathcal{O}_t .
4. Compute ideals $J_s = \mathcal{O}_0\alpha_s + \mathcal{O}_0\ell^{e_s}$ and $J_t = \mathcal{O}_0\alpha_t + \mathcal{O}_0\ell^{e_t}$ with norm ℓ^{e_s}, ℓ^{e_t} for some e_s, e_t , in the same classes as I_s and I_t respectively.
5. For $r = s, t$ and corresponding $E = E(j_r)$:
 - (a) Compute a sequence of ideals $J_{r,i} = \mathcal{O}_0\alpha_r + \mathcal{O}_0\ell^i$ for $i = 0, \dots, e_r$
 - (b) For $0 \leq i \leq e_r$:
 - (c) Compute $K_{r,i}$ with powersmooth norm in the same class as $J_{r,i}$.
 - (d) Translate $K_{r,i}$ into an isogeny $\varphi_{r,i} : E_0 \rightarrow E_{r,i}$.
 - (e) Deduce a sequence $(j_0, j(E_{r,1}), j(E_{r,2}), \dots, j(E_{r,e}) = j(E))$.
6. Return $(j(E_s), \dots, j_0, \dots, j(E_t))$ the concatenation of both paths.

The reverse direction may a priori look easier. By standard arguments on hash functions it is sufficient to prove the claim with respect to a collision algorithm. A collision for the Charles-Goren-Lauter hash function gives a non-scalar endomorphism of the curve; four linearly independent endomorphisms give a full rank subring of the endomorphism ring; and heuristically one expects that a few such maps will be sufficient to generate the whole ring. To compute the endomorphism ring one would therefore call the collision finding algorithms multiple times until the resulting maps generate the full endomorphism ring. This strategy, however, has a potential caveat: the collision algorithm might be such that it always returns the same endomorphism. In Algorithm 8 we get around this problem by performing a random walk from the input invariant j , calling the collision algorithm on the end-vertex of the random walk, and concatenating paths to form endomorphisms of $E(j)$.

Proposition 8. *Assume there exists an efficient preimage or collision algorithm for the Charles-Goren-Lauter hash function. Then under plausible heuristic assumptions there is an efficient algorithm to solve the endomorphism ring computation problem.*

Proof. The reduction algorithm for collision resistance is given by Algorithm 8 below. Note that in Step 7 the discriminant can be computed from the Gram matrix, which by Lemma 4 can be efficiently computed. Heuristically, one expects that the loop will be executed at most $O(\log p)$ times. Indeed let us assume that after adding some elements to the subring we have a subring of index N . Then we can heuristically expect any new randomly generated endomorphism to lie in this subring with probability only $1/N$. Moreover when it does not lie in the subring, the element will decrease the index by a non trivial integer factor of N . \square

Algorithm 8 *Reduction from endomorphism ring computation to collision resistance*

Input: A supersingular j -invariant $j \in \mathbb{F}_{p^2}$.

Output: The endomorphism ring of $E(j)$.

1. Let $\mathcal{R} = \langle 1 \rangle \subset \text{End}(E(j))$.
2. While $\text{disc}(\mathcal{R}) \neq 4p^2$:
 - (a) Perform a random walk in the graph, leading to a new vertex j' .
 - (b) Apply a collision finding algorithm on j' , leading to an endomorphism of $E(j')$.
 - (c) Deduce an endomorphism ϕ of $E(j)$ by concatenating paths.
 - (d) Set $\mathcal{R} \leftarrow \langle \mathcal{R}, \phi \rangle$.
 - (e) Compute the discriminant of \mathcal{R} .
3. Return a \mathbb{Z} -basis for \mathcal{R} .

6 ℓ -PowerIsogeny Reduces to MaxOrder and Action-on- ℓ -Torsion


In this section we show that computing an ℓ -isogeny between two supersingular elliptic curves reduces to computing maximal orders of elliptic curves and solving the Action-on- ℓ -Torsion Problem.


6.1 Outline of reduction

Given two supersingular elliptic curves E, E' over \mathbb{F}_{p^2} , and oracles for the problems Action-on- ℓ -Torsion and MaxOrder, we will construct an ℓ -power isogeny $E \rightarrow E'$ by constructing a chain of ℓ -isogenies through intermediate curves. First, the oracle will give us two maximal orders $\mathcal{O}, \mathcal{O}' \subseteq B_{p,\infty}$ with $\mathcal{O} \simeq \text{End}(E)$ and $\mathcal{O}' \simeq \text{End}(E')$. We then compute a connecting ideal, meaning a left ideal of \mathcal{O} , whose left order is \mathcal{O} and right order is \mathcal{O}' . Next we use the main algorithm of [KLPT14] to compute an equivalent ideal I whose norm is ℓ^e for some $e = O(\log p)$. The isogeny $\phi_I : E \rightarrow E'$ corresponding to I has degree ℓ^e , so the representation size of the isogeny is exponential. To remedy this we will, given I , compute a chain of ℓ -isogenies ψ_1, \dots, ψ_e such that $\phi_I = \psi_e \circ \dots \circ \psi_1$. Since ψ_1, \dots, ψ_e have degree ℓ , they are of polynomial representation size as rational maps. To obtain the ψ_i we will first show that there is a factorization of the ideal I . The proper notion here is that of a *filtration* of ideals, namely a sequence

$$I = I_e \subseteq I_{e-1} \subseteq \dots \subseteq I_1 \subseteq I_0 = \mathcal{O}$$

such that the isogeny corresponding to I_k is a map ϕ_k from E to some intermediate curve E_k . The factorization of ϕ_I gives us a path starting at E and ending at E' of length e in the graph of isogenies of degree ℓ , and the filtration of I leads to a corresponding “path” between maximal orders in $B_{p,\infty}$. The maximal orders that appear in this path are $\mathcal{O}_R(I_k)$ and the ideal connecting $\mathcal{O}_R(I_k)$ to $\mathcal{O}_R(I_{k+1})$ is $J_k := I_{k-1}^{-1} I_k$. These paths are given in the following diagrams:





For each k , the isogeny $\phi_k : E_0 \rightarrow E_k$ has degree ℓ^k , and so corresponds to a left \mathcal{O} -ideal I_k of norm ℓ^k . We will show that $I_k = I + \mathcal{O}\ell^k$ is the desired ideal. As k grows, these ideals will have norms which are too big to find the corresponding isogenies, so we will compute the maps $\psi_k : E_{k-1} \rightarrow E_k$ which correspond to left ideals J_k of $\mathcal{O}_R(I_{k-1})$ of norm ℓ . Suppose we have computed ψ_k , the curve E_k , and J_{k+1} as above. We can use the oracle for MaxOrder to identify generators of J_{k+1} with endomorphisms of E_k . On the other hand, J_{k+1} corresponds to the isogeny ψ_{k+1} , whose kernel we compute using the information from the oracle Action-on- ℓ -Torsion. Using Vélú's formula, we can compute ψ_{k+1} from its kernel. This procedure iteratively computes the desired maps $\psi_1, \psi_2, \dots, \psi_e$.

6.2 Reduction from ℓ -PowerIsogeny to MaxOrder and Action-on- ℓ -Torsion

In this section, we give the reduction from ℓ -Power Isogeny to the problems MaxOrder and Action-on- ℓ -Torsion.

Algorithm 9 *Reduction from ℓ -PowerIsogeny to MaxOrder and Action-on- ℓ -Torsion*

Input: E, E' supersingular elliptic curves over \mathbb{F}_{p^2} , a prime $\ell \neq p$.

Output: a chain of ℓ -isogenies connecting E and E' .

1. Compute a basis $\langle 1, i, j, ij \rangle$ for $B_{p,\infty}$.
2. Call oracle MaxOrder on $p, \langle 1, i, j, ij \rangle, E$, resulting in $\alpha_1, \alpha_2, \alpha_3, \alpha_4$ where $\text{End}(E) \simeq \mathcal{O} := \langle \alpha_1, \alpha_2, \alpha_3, \alpha_4 \rangle \subseteq B_{p,\infty}$.
3. Call oracle MaxOrder on $p, \langle 1, i, j, ij \rangle, E'$, resulting in $\alpha'_1, \alpha'_2, \alpha'_3, \alpha'_4$ where $\text{End}(E') \simeq \mathcal{O}' := \langle \alpha'_1, \alpha'_2, \alpha'_3, \alpha'_4 \rangle \subseteq B_{p,\infty}$.
4. Compute connecting ideal: use $\alpha_1, \dots, \alpha_4$ and $\alpha'_1, \dots, \alpha'_4$ to compute a left ideal I of \mathcal{O} such that $\mathcal{O}_R(I) = \mathcal{O}'$ and $\text{Nrd}(I) = \ell^e$ with $e = O(\log p)$. Adjust I so that $I \not\subseteq \ell^k \cdot \mathcal{O}$ for any positive integer k .
5. For $0 \leq k \leq e$:
 - (a) Compute $I_k := I + \mathcal{O}\ell^k$. This is a left ideal of \mathcal{O} of norm ℓ^k . Also compute its right order $\mathcal{O}_R(I_k)$.
 - (b) Compute a \mathbb{Z} -basis $\gamma_1, \gamma_2, \gamma_3, \gamma_4$ for the ideal $J_{k+1} := I_k^{-1}I_{k+1}$ of $\mathcal{O}_R(I_k)$.
6. Set $E_0 := E$.
7. For $0 \leq k \leq e - 1$:
 - (a) Compute a basis $\{P_1, P_2\}$ for $E_k[\ell]$.
 - (b) Call oracle MaxOrder with $p, \langle 1, i, j, ij \rangle, E_k$, resulting in $\beta_1, \beta_2, \beta_3, \beta_4$ that generate $\mathcal{O}_k \subseteq B_{p,\infty}$.
 - (c) Call oracle Action-on- ℓ -Torsion with parameters $p, P_1, P_2, \langle 1, i, j, ij \rangle, E_k, \beta_1, \beta_2, \beta_3, \beta_4$ resulting in $Q_{st} = \iota_k^{-1}(\beta_s)(P_t)$ for $s = 1, \dots, 4, t = 1, 2$. Here, $\iota_k : \text{End}(E_k) \rightarrow \langle \beta_1, \dots, \beta_4 \rangle$ is an isomorphism.

- (d) Compute $v \in B_{p,\infty}$ such that $v\mathcal{O}_R(I_k)v^{-1} = \mathcal{O}_k$.
 - (e) Compute c_{rs} such that $v\gamma_r v^{-1} = \sum_s c_{rs}\beta_s$.
 - (f) Find $x, y \in \mathbb{Z}/\ell\mathbb{Z}$, not both 0, such that $\sum_s c_{rs}(xQ_{s1} + yQ_{s2}) = 0$ for $r = 1, \dots, 4$.
 - (g) Compute ψ_{k+1} and its image E_{k+1} corresponding to the kernel subgroup $\langle xP_1 + yP_2 \rangle = E_k[\iota_k^{-1}(J_{k+1})]$ using Vélú's formula
8. Return $\psi_1, \psi_2, \dots, \psi_e$.

Theorem 10. *ℓ -PowerIsogeny efficiently reduces to MaxOrder and Action-on- ℓ -Torsion. In particular, given a prime p , a prime $\ell \neq p$, and supersingular elliptic curves E, E' over \mathbb{F}_{p^2} , Algorithm 9 returns isogenies ψ_1, \dots, ψ_e of degree ℓ whose composition is an isogeny $\psi := \psi_e \circ \dots \circ \psi_1$ of degree ℓ^e from E to E' . Assuming ℓ is of size $O(\log p)$, Algorithm 9 runs in time polynomial in $\log p$ and makes $O(\log p)$ queries of MaxOrder and Action-on- ℓ -Torsion.*

Proof. By Theorem 2, the oracle returns a basis for \mathcal{O} and for \mathcal{O}' of polynomial size. To do Step 4, we first compute an arbitrary connecting ideal for \mathcal{O} and \mathcal{O}' in polynomial time using Algorithm 3.5 of [KV10]. An equivalent connecting ideal of norm ℓ^e , where $e = O(\log p)$, can be computed in polynomial time as claimed in [KLPT14].

Define $E_k := E/E[I_k]$ (here by $E[I_k]$ we mean the subgroup $E[\iota^{-1}(I_k)]$, where $\iota : \text{End}(E) \rightarrow \mathcal{O}$ is an isomorphism). We need to show that I_k has norm ℓ^k and that the left $\mathcal{O}_R(I_k)$ -ideal J_{k+1} corresponds to the isogeny $\psi_{k+1} : E_k \rightarrow E_{k+1}$ in the factorization $\phi_k = \psi_k \circ \phi_{k-1}$; this is proved in Theorem 11. Right orders and products of ideals can be computed efficiently with linear algebra over \mathbb{Z} , hence Step 4 is efficient; see [Rón92], Theorem 3.2 for the statement on right orders. Inverses can be computed from the formula $I^{-1} = \frac{1}{\text{Nrd}(I)} \bar{I}$. We make e calls to the oracle for generators of $\text{End}(E_k)$ and their action on ℓ -torsion. If $\mathcal{O} \simeq \mathcal{O}_k$, we can compute v such that $v\mathcal{O}_k v^{-1} = \mathcal{O}$ in polynomial time by Lemma 2.5, Corollary 3.6, and Proposition 6.9 of [KV10]. By Theorem 11, the isogeny corresponding to I factors as the product of the isogenies corresponding to J_k , $k = 1, \dots, e$, all of which have degree ℓ . Now compute the kernel of ψ_k using J_k and the action of $\text{End}(E_{k-1})$ on the ℓ -torsion of E_{k-1} ; see Proposition 9. Since ℓ is $O(\log p)$, rational maps for ψ_k from its kernel can be efficiently computed. \square

6.3 Going from an ideal of norm ℓ to a corresponding subgroup of order ℓ

At the beginning of Step 7 of the algorithm, we have an isogeny $E_{k-1} \rightarrow E_k$ represented by a left $\mathcal{O}_R(I_{k-1})$ -ideal J_k . We wish to specify the subgroup of E_{k-1} which is the kernel of this isogeny. If $\tilde{J}_k \subseteq \text{End}(E_{k-1})$ is the ideal isomorphic to J_k , recall from Section 2.2 that

$$E_{k-1}[\tilde{J}_k] = \bigcap_{\gamma \in \tilde{J}_k} \ker(\gamma_k),$$

and it suffices to compute $\ker(\gamma_1) \cap \cdots \cap \ker(\gamma_4)$, where $\gamma_1, \dots, \gamma_4$ are a \mathbb{Z} -basis of \widetilde{J}_k . Once we have $E_{k-1}[\widetilde{J}_k]$, we can use Vélú's formula to compute ψ_k .

Step 7 in our algorithm computes $E_{k-1}[\widetilde{J}_k]$ and is similar to Algorithm 2 in [GPS17]. In our version, we are working with ideals in consecutive endomorphism rings, rather than in the endomorphism ring of the starting curve, and we give proofs of correctness along with analysis of input size of left ideals of a maximal order.

Proposition 9. *Let E be a supersingular elliptic curve over \mathbb{F}_{p^2} , and assume $\iota : \text{End}(E) \rightarrow \mathcal{O} \subseteq B_{p,\infty}$ is an isomorphism, where \mathcal{O} has a basis of size polynomial in $\log p$. Let $I \subseteq \mathcal{O}$ be an ideal of norm ℓ^e for a prime $\ell \neq p$ with $\ell = O(\log p)$. For $k = 1, \dots, e$, define $I_k := I + \mathcal{O} \cdot \ell^k$ and $J_k = I_{k-1}^{-1} I_k \subseteq \mathcal{O}_R(I_{k-1})$ and $E_k := E/E[\iota^{-1}(I_k)]$ as in Theorem 11. Then if we are given $\iota_{k-1}(\text{End}(E_{k-1}))$ in $B_{p,\infty}$ where $\iota_{k-1} : \text{End}(E_{k-1}) \otimes \mathbb{Q} \rightarrow B_{p,\infty}$ is an isomorphism of quaternion algebras, along with the action of $\text{End}(E_{k-1})$ on $E_{k-1}[\ell]$, we can compute the kernel of the isogeny corresponding to $\iota_{k-1}^{-1}(J_k)$ in time polynomial in $\log p$.*

Proof. We wish to determine $E_{k-1}[\iota_{k-1}^{-1}(J_k)]$ so that we can compute the corresponding isogeny $\psi_k : E_{k-1} \rightarrow E_k$. If J_k has a \mathbb{Z} -basis $\gamma_1, \dots, \gamma_4 \in \mathcal{O}_R(I_{k-1})$, we need to understand how the γ_i act as endomorphisms of E_{k-1} . Suppose we are given the action of generators ϕ_1, \dots, ϕ_4 of $\text{End}(E_{k-1})$ on $E_{k-1}[\ell]$ and the image of an embedding $\iota_{k-1} : \text{End}(E_{k-1}) \rightarrow B_{p,\infty}$. Set $\mathcal{O}_{k-1} := \iota_{k-1}(\text{End}(E_{k-1}))$; then we can compute $v \in B_{p,\infty}^\times$ such that $\mathcal{O}_{k-1} = v\mathcal{O}_R(I_{k-1})v^{-1}$ in polynomial time by [KV10]. By expressing $v\gamma_i v^{-1}$ in terms of $\iota_{k-1}(\phi_j)$, say

$$v\gamma_r v^{-1} = \sum_s c_{rs} \iota_{k-1}(\phi_s),$$

we discern the kernel of the isogeny corresponding to J_k as follows. We require a nonzero point $P \in E_{k-1}[\ell]$ such that for all $r = 1, \dots, 4$,

$$\sum_s c_{rs} \phi_s(P) = 0.$$

Because we assume that we are given $\phi_s(P)$ for $s = 1, \dots, 4$ and $P \in E_{k-1}[\ell]$, we can find such a P by just calculating the sum for all $r = 1, \dots, 4$ and $P \neq 0 \in E_{k-1}[\ell]$. \square

6.4 Isogeny paths and corresponding filtrations of left ideals

Let $E, E'/\mathbb{F}_{p^2}$ be supersingular elliptic curves. We now prove the correctness of our earlier claims on how an ℓ -isogeny path between E and E' corresponds to a sequence of ideals of norm ℓ in $\text{End}(E) \otimes \mathbb{Q}$. In particular, suppose $\phi : E \rightarrow E'$ has degree ℓ^e for some prime $\ell \neq p$. Then the kernel ideal I of ϕ in $\text{End}(E)$ has degree ℓ^e . There is a factorization $\phi = \psi_e \circ \cdots \circ \psi_1$ with $\deg(\psi_k) = \ell$, and by setting $\phi_k := \psi_k \circ \cdots \circ \psi_1$, there is a corresponding ideal I_k of $\text{End}(E)$ of norm ℓ^k . Additionally, there is an ideal J_k of $\mathcal{O}_R(I_{k-1})$ which corresponds to the

factorization of the isogeny $\phi_k = \psi_k \circ \psi_{k-1}$; in this section, we construct I_k and J_k from I . Let I be a left ideal of $\text{End}(E)$ of norm ℓ^e such that $I \not\subseteq \text{End}(E) \cdot \ell^m$ for any positive integer m . In this section, we prove that for $k = 0, \dots, e$, $I_k = I + \text{End}(E) \cdot \ell^k$ is an ideal of norm ℓ^k and that

$$I = I_e \subseteq I_{e-1} \subseteq \dots \subseteq I_1 \subseteq I_0 = \text{End}(E).$$

We first establish when an ideal corresponds to an isogeny with cyclic kernel.

Proposition 10. *Suppose $I \subseteq \text{End}(E)$ is a left ideal with $\text{Nrd}(I)$ coprime to p . Then I is not contained in $\text{End}(E) \cdot m$ for any $m \in \mathbb{N}$ if and only if $E[I]$ is cyclic.*

Proof. Suppose that $I \subseteq \text{End}(E) \cdot m$. Then $E[I] \supseteq E[\text{End}(E) \cdot m] = E[m]$ and thus $m \mid \deg(\phi_I)$. Since p does not divide $\deg(\phi_I)$, it also does not divide m , so $E[m] \neq 0$ and has rank two as a $\mathbb{Z}/m\mathbb{Z}$ -module. Hence $E[I]$ is not cyclic. For the other direction, suppose that $E[I]$ is not cyclic. Then, by the structure theorem of abelian groups,

$$E[I] \simeq \bigoplus_{i=1}^j \mathbb{Z}/k_i \mathbb{Z}$$

and we can choose the k_i uniquely such that $k_i \mid k_{i+1}$. Since $E[I]$ is not cyclic, $j \neq 1$ and hence $E[I]$ has two elements of order k_1 which are linearly independent. Thus $E[k_1] \subseteq E[I]$ and hence $I \supseteq \text{End}(E) \cdot k_1$. \square

Proposition 11. *Suppose $I \subseteq \text{End}(E)$ and $N := \text{Nrd}(I)$ is coprime to p . Also suppose $M \mid N$, and that I is not contained in $\text{End}(E) \cdot m$ for any $m \in \mathbb{N}$. Then $I + \text{End}(E) \cdot M$ has norm M .*

Proof. We claim that

$$E[I + M\mathcal{O}] = E[I] \cap E[M].$$

Indeed, for an arbitrary left ideal J of $\text{End}(E)$ with $\text{Nrd}(J)$ coprime to p , $E[J]$ is the intersection of the kernels of a generating set of J , and for two left $\text{End}(E)$ -ideals J, J' , $J + J'$ is generated by $J \cup J'$. Since $E[I]$ is cyclic by Proposition 10, there is some $Q \in E[N]$ so that $E[I] = \langle Q \rangle$. Then $E[I] \cap E[M] = \langle [N/M]Q \rangle$, a group of order M as desired. \square

6.5 Matching up a filtration of an ideal with a factorization of an isogeny

In this section, we show that the definition of J_k in Algorithm 9 gives us the ideal which corresponds to the isogeny $E_{k-1} \rightarrow E_k$ of degree ℓ . To do this, it suffices to understand the horizontal isogeny and corresponding ideal in the following diagram:

$$\begin{array}{ccc} E & & \\ I_{k-1} \downarrow & \searrow I_k & \\ E_{k-1} := E/E[I_{k-1}] & \xrightarrow{J_k} & E_k := E/E[I_k] \end{array}$$

We will describe the relationship between the horizontal isogeny and its kernel ideal for two arbitrary left ideals I, I' of $\text{End}(E)$ satisfying $I' \subseteq I$, so in the above picture, we replace I_{k-1} with I and I_k with I' . The goal is to find, given $I' \subseteq I$, the horizontal isogeny $E_I \rightarrow E_{I'}$ by first computing its corresponding ideal \tilde{J} in the following diagram:

$$\begin{array}{ccc} E & & \\ \downarrow I & \searrow I' & \\ E_I := E/E[I] & \xrightarrow{\tilde{J}} & E_{I'} := E/E[I'] \end{array}$$

Let $\phi_I : E \rightarrow E_I := E/E[I]$ and $\phi_{I'} : E \rightarrow E_{I'} := E/E[I']$ be the corresponding isogenies; then $E[I] \subseteq E[I']$ and hence $\phi_{I'}$ factors as $\phi_{I'} = \psi \phi_I$ for some isogeny $\psi : E_I \rightarrow E_{I'}$. We wish to view the kernel of ψ as $E_I[\tilde{J}]$ for some left ideal \tilde{J} of $\text{End}(E_I)$. We make this idea precise in the following proposition.

Proposition 12. *Let $I' \subseteq I$ be two left $\text{End}(E)$ -ideals whose norms are coprime to p . Then there exists a separable isogeny $\psi : E_I \rightarrow E_{I'}$ such that $\phi_I = \psi \circ \phi_{I'}$, and a left ideal \tilde{J} of $\text{End}(E_I)$ with $E_I[\tilde{J}] = \ker(\psi)$ such that $J = \iota(\tilde{J}) = I^{-1}I'$, where $\iota : \text{End}(E_I) \rightarrow \text{End}(E) \otimes \mathbb{Q}$ is the map in Lemma 9 below.*

To prove this, we need the following three lemmas:

Lemma 8. *For a left ideal I of $\text{End}(E)$, the map*

$$\begin{aligned} \phi_I^* : \text{Hom}(E_I, E) &\rightarrow I \\ \psi &\mapsto \psi \phi_I \end{aligned}$$

is an isomorphism of left $\text{End}(E)$ -modules.

Proof. This is Lemma 42.2.6 of [Voi]. It also follows from Proposition 48 of [Koh96]. \square

Lemma 9. *Set $B = \text{End}(E) \otimes \mathbb{Q}$. The map*

$$\begin{aligned} \iota : \text{End}(E_I) &\rightarrow B \\ \beta &\mapsto \frac{1}{\deg(\phi_I)} \widehat{\phi_I} \beta \phi_I \end{aligned}$$

is injective, and its image is $\mathcal{O}_R(I)$.

Proof. This is Lemma 42.2.8 of [Voi] or Proposition 3.9 of [Wat69]. \square

Lemma 10. *We have a bijection*

$$\begin{aligned} g : \text{Hom}(E_{I'}, E_I) &\rightarrow I^{-1}I' \\ \psi &\mapsto \frac{1}{\deg(\phi_I)} \widehat{\phi_I} \psi \phi_{I'}. \end{aligned}$$

Proof. This is Lemma 42.2.19 of [Voi]. □

Now we can prove the proposition.

Proof (Proof of Proposition 12). We have that $I^{-1} = \frac{1}{\text{Nrd}(I)} \bar{I}$. Consider an element $x \in I^{-1}I'$ of the form

$$x = \frac{1}{\deg(\phi_I)} \widehat{\alpha'} \beta',$$

where $\alpha' \in I$, $\beta' \in I'$. Then by Lemma 8, there exists $\alpha \in \text{Hom}(E_I, E)$ and $\beta \in \text{Hom}(E_{I'}, E)$ with

$$\alpha' = \alpha \phi_I, \beta' = \beta \phi_{I'}.$$

Thus

$$x = \frac{1}{\deg(\phi_I)} \widehat{\phi_I \alpha \beta \phi_{I'}} = g(\widehat{\alpha \beta}),$$

where $g : \text{Hom}(E_{I'}, E_I) \rightarrow I^{-1}I'$ is the map in Lemma 10. Since $E[I] \subseteq E[I']$, and $\phi_I, \phi_{I'}$ are separable, by Corollary III.4.11 of [Sil09] there exists a unique separable isogeny $\psi : E_I \rightarrow E_{I'}$ such that $\phi_{I'} = \psi \circ \phi_I$. Then define

$$\tilde{J} := \{\alpha \in \text{End}(E_1) : \alpha(P) = 0 \quad \forall P \in \ker(\psi)\}.$$

Now map $g^{-1}(x) = \widehat{\alpha \beta} \in \text{Hom}(E_{I'}, E_I)$ to an element of \tilde{J} using $\psi^* : \widehat{\alpha \beta} \psi = \psi^*(\widehat{\alpha \beta}) \in \tilde{J}$. Finally, compute

$$\begin{aligned} x &= \frac{1}{\deg(\phi_I)} \widehat{\phi_I \alpha \beta \phi_{I'}} \\ &= \frac{1}{\deg(\phi_I)} \widehat{\phi_I \alpha \beta \psi \phi_I} \\ &= \iota(\widehat{\alpha \beta \psi}) \\ &= \iota(\psi^*(\widehat{\alpha \beta})) \\ &= (\iota \circ \psi^* \circ g^{-1})(x). \end{aligned}$$

In other words, we have

$$g = \iota \circ \psi^*.$$

From this, we conclude that the left ideal of $\mathcal{O}_R(I_1)$ corresponding to \tilde{J} indeed is $I^{-1}I'$. □

Combining the above results, we have our main theorem on matching up filtrations of ideals with factorizations of isogenies:

Theorem 11. *Suppose that $I \subseteq \text{End}(E)$ satisfies $\text{Nrd}(I) = \ell^e$ where $\ell \neq p$ is a prime and $I \not\subseteq \text{End}(E) \cdot \ell^k$ for any $k \in \mathbb{N}$. Then there exists a filtration*

$$I = I_e \subsetneq I_{e-1} \subsetneq \dots \subsetneq I_1 \subsetneq I_0 = \text{End}(E)$$

and a chain of isogenies

$$E = E_0 \xrightarrow{\psi_1} E_1 \xrightarrow{\psi_2} \dots \xrightarrow{\psi_{e-2}} E_{e-1} \xrightarrow{\psi_e} E_e = E'$$

such that if we set $\phi_k : E \rightarrow E/E[I_k]$, then $\phi_{k+1} = \psi_k \phi_k$. Moreover, for $k = 0, \dots, e-1$, the map $\psi_{k+1} : E_k \rightarrow E_{k+1}$ has degree ℓ , and its kernel ideal in $\text{End}(E_k)$ is isomorphic to $I_k^{-1} I_{k+1} \subseteq \mathcal{O}_R(I_k)$ under the map

$$\begin{aligned} \iota_k : \text{End}(E_k) &\rightarrow \mathcal{O}_R(I_k) \\ \rho &\mapsto \frac{1}{\deg(\phi_k)} \hat{\phi}_k \rho \phi. \end{aligned}$$

Proof. For $k = 0, 1, \dots, e$, define $I_k := I + \text{End}(E) \cdot \ell^k$. By Proposition 11, $\text{Nrd}(I_k) = \ell^k$. Let $\phi_I : E \rightarrow E_e := E/E[I_e] = E/E[I]$ be the isogeny corresponding to $I = I_e$. Set $\mathcal{O}_k := \mathcal{O}_R(I_k) \subseteq \text{End}(E) \otimes \mathbb{Q}$, and $J_k := I_{k-1}^{-1} I_k$. Then $\text{Nrd}(J_k) = \ell$. Let $E_k := E/E[I_k]$. From the ideals J_k , we have isogenies $\psi_k : E_{k-1} \rightarrow E_k$ such that

$$\phi = \psi_e \circ \dots \circ \psi_1$$

by Proposition 12 applied inductively to the ideals $I_{k+1} \subsetneq I_k$. \square

7 Some easy problems in supersingular isogeny graphs

The previous sections relied heavily on the quaternion ℓ -isogeny algorithm of [KLPT14] to derive the computational equivalence of several problems. In this section, we provide two additional applications of this algorithm. First, we give an algorithm for constructing the Deuring correspondence from maximal orders in $B_{p,\infty}$ to supersingular j -invariants. Second, we give a polynomial-time collision algorithm against the Charles-Goren-Lauter hash function when a special curve is chosen as the initial point.

7.1 Constructive Deuring correspondence, from quaternion orders to j -invariants

In this section we provide an efficient algorithm to solve Problem 1. Algorithm 12 first computes an ideal connecting \mathcal{O}_0 to \mathcal{O} . Then it uses the quaternion ℓ -isogeny algorithm from [KLPT14] (or rather, its powersmooth version) to compute another ideal in the same class but with a norm $N = \prod p_i^{e_i}$ such that $\max_i p_i^{e_i}$ is small. It finally translates that ideal into an isogeny $\phi : E_0 \rightarrow E_1$ that corresponds to it via Deuring's correspondence.

Algorithm 12 *Constructive Deuring correspondence, from maximal orders to j -invariants.*

Input: Maximal order $\mathcal{O} \subset B_{p,\infty}$.

Output: Supersingular j -invariant j such that $\text{End}(E(j)) \simeq \mathcal{O}$.

1. Compute an ideal I that is a left ideal of \mathcal{O}_0 and a right ideal of \mathcal{O} .
2. Compute an ideal J in the same class as I but with powersmooth norm.
3. Compute an isogeny $\phi : E_0 \rightarrow E_I$ that corresponds to J via Deuring's correspondence.
4. Return $j(E_I)$.

Let $\langle 1, \omega_2, \omega_2, \omega_3 \rangle$ be a basis for \mathcal{O} , and let $M \in GL(4, \mathbb{Q})$ be such that $(1, \omega_2, \omega_2, \omega_3) = M(1, i, j, k)$. Let B be a bound on the numerators and denominators of all the coefficients of M .

Proposition 13 (Constructive Deuring Correspondence.). *Under plausible heuristic assumptions, Algorithm 12 can be implemented to run in time polynomial in both $\log B$ and $\log p$.*

Proof. The analysis is similar to the proof of Proposition 5. \square

We remark that this algorithm is implicitly used in the recent identification protocol of Galbraith, Silva and Petit [GPS17].

7.2 An attack on the CGL hash function

It was shown in [CGL09] that computing collisions or preimages for the Charles-Goren-Lauter hash function amounts to computing large ℓ -power degree isogenies between two (possibly isomorphic) elliptic curves. The hardness arguments for these problems then essentially relied on the following arguments:

1. In general, these isogenies must have a degree so large that they cannot be efficiently computed with current algorithms.
2. The best known algorithms for these problems were variants that used birthday arguments, with an exponential complexity in the parameter's size [Gal99].

Paradoxically, the quaternion ℓ -isogeny algorithm [KLPT14] leads to both the security arguments of Section 5.5 and to a partial attack against the hash function. More precisely, in this section we present a collision attack for the hash function when the initial point used in the random walk is the special elliptic curve E_0 as constructed in Algorithm 3.

Our attack is summarized by Algorithm 13 below. We first compute $\alpha \in \langle 1, i, j, k \rangle \subset \mathcal{O}_0$ with $\text{Nrd}(\alpha) = \ell^e$ for some e , which defines a sequence of ideals I_i corresponding to a loop starting and ending at \mathcal{O}_0 . To ensure there is no backtracking in the loop (and moreover, that $\alpha \neq \ell^{e/2}$), we require that for any natural number k , $\ell^{-k}\alpha \notin \mathcal{O}_0$. Applying the translation algorithm directly to this sequence of ideals would have a prohibitive cost because ℓ^e is larger than p . As in Algorithm 7, we first replace each ideal in the sequence by another ideal in the same class but with powersmooth norm, and we apply the translation algorithm to each of them individually to obtain corresponding isogenies. The end vertices of these isogenies form a sequence of j -invariants that define a collision for the original elliptic curve version of the Charles-Goren-Lauter hash function.

Algorithm 13 *Collision attack on CGL hash function for special initial points*

Input: Special j_0 and \mathcal{O}_0 from Algorithm 3.

Output: A sequence of j -invariants $j_0, j_1, \dots, j_e = j_0$ such that for any i there exists an isogeny of degree ℓ from $E(j_i)$ to $E(j_{i+1})$.

1. Compute $e \in \mathbb{N}$ and $\alpha \in \langle 1, i, j, k \rangle \subset \mathcal{O}_0$ with $\text{Nrd}(\alpha) = \ell^e$.
2. Compute a sequence of ideals $I_i = \mathcal{O}_0 q + \mathcal{O}_0 \ell^i$.
3. For all i :
 - (a) Compute J_i with powersmooth norm in the same class as I_i .
 - (b) Translate J_i into an isogeny $\varphi_i : E_0 \rightarrow E_i$.
4. Return $(j_0, j(E_1), j(E_2), \dots, j(E_e) = j_0)$.

To obtain an element whose norm is a power of ℓ in Step 1, we fix e large enough, then pick random values of y and z until the equation $w^2 + qx^2 = \ell^e - p(y^2 + qz^2)$ can be solved with Cornacchia's algorithm. This solution is described in Algorithm 14.

Algorithm 14 ℓ -power norm element in \mathcal{O}_0

Input: Maximal order $\mathcal{O}_0 \subset B_{p,\infty}$ as defined in Proposition 1.

Output: $e \in \mathbb{N}$ and $\alpha \in \mathcal{O}_0$ with $\text{Nrd}(\alpha) = \ell^e$.

1. Let $e = \lceil 2 \log p \rceil$.
2. Choose random y, z smaller than $\sqrt{p/q}$.
3. Let $N \leftarrow \ell^e - p(y^2 + qz^2)$.
4. Find $w, x \in \mathbb{Z}$ such that $w^2 + qx^2 = N$ if there are some, otherwise go to Step 2.
5. Return $\alpha = w + xi + yj + zk$.

Proposition 14. *There exists an algorithm that computes a collision for the Charles-Goren-Lauter hash function when the initial vertex is a special curve in time polynomial in $\log p$.*

Proof. In Algorithm 14 we expect that the equation in Step 4 will have a solution for $1/2q \log p$ of the random choices (y, z) , so we expect this algorithm to run in time polynomial in $\log p$. Note that $e = \lceil 2 \log p \rceil$, and that Steps 4 and 5 in Algorithm 13 both run in time polynomial in $\log p$. We conclude that the runtime of Algorithm 13 is also polynomial in $\log p$. To ensure there is no backtracking in the loop in the isogeny graph, we require that the ideal $\mathcal{O}_0 \alpha$ satisfies $\mathcal{O}_0 \alpha \not\subset \mathcal{O}_0 \ell^k$ for any k . \square

We remark that we described our attack only for the maximal orders \mathcal{O}_0 defined in Proposition 1, but it can be extended to other maximal orders as long as the corresponding curve is known or can be computed, and as long as elements of norm a power of ℓ can be found in the order. This is the case for “special” orders, as defined in [KLPT14].

The attack provided by Algorithm 13 can be extended into a “backdoor attack” where an entity in charge of deciding the initial vertex for the hash function plays the role of the attacker. This entity could take a random walk

from j_0 to another curve E and publish this $j(E)$ as the initial vertex for the hash function. Due to the random walk the vertex $j(E)$ will be uniformly distributed, hence the function will be collision resistant based on the assumption that the endomorphism ring computation problem is hard (see Proposition 8). However, the entity can concatenate the path from j_0 to j and the collision which begins and ends at j_0 to obtain a collision which begins and ends at j .

To the best of our knowledge, there exists no efficient algorithm to sample supersingular j -invariants that does not involve this random walk procedure, so the backdoor attack cannot really be avoided. On the other hand, by inspecting such a collision, it is easy to recover a path to \mathcal{O}_0 and that will reveal that a backdoor was inserted. In that sense, the backdoor mechanism may not be too much of an issue in practice.

8 The EndomorphismRing Problem

In this section we provide an alternative study of the computational hardness of computing endomorphism rings of supersingular elliptic curves. The inputs are p and the curve, and so the running time must be polynomial in $\log p$. This brings up two important questions: 1) Does the endomorphism ring of an elliptic curve have a polynomial representation size? And 2) If it does, can the endomorphisms be evaluated in polynomial time? To have any meaningful efficient reduction, or to analyze how hard it is to compute the endomorphism ring, we need to know what the representation size of an endomorphism ring is. In particular, we need to discuss what we mean by *computing the endomorphism ring*.

We will define a compact representation of endomorphisms which has polynomial size, and show that the endomorphism ring of any supersingular elliptic curve has a basis of such representations. This answers question 1. We also show that these representations can be evaluated efficiently at arbitrary points, answering question 2. We then define the problem EndomorphismRing in terms of this new definition, and show that it efficiently reduces to MaximalOrder and Action-on- ℓ -Torsion for $\ell = 2, 3$. Our definition of compact representations is implicitly used in Algorithm 4. We also identify another problem that it reduces to, which is related to computing isogenies.

8.1 Representation size of endomorphism rings

There are two typical ways to represent the endomorphism ring of E . The first is to give rational functions $F_1(x, y), \dots, F_4(x, y)$ and $G_1(x, y), \dots, G_4(x, y)$ such that $\phi_i : (x, y) \mapsto (F_i(x, y), G_i(x, y))$ ($i = 1, \dots, 4$) are endomorphisms of E that form a basis for $\text{End}(E)$. The second is to give the kernel of the maps ϕ_i , which in general is not good enough for computations. However, it is not known if a basis for $\text{End}(E)$ exists in either representation that is of polynomial size. For example, the basis may contain an endomorphism of exponential degree, where exponentially many coefficients would be needed to describe it in general. For the case of using the kernel, the generators may lie in a finite field of exponential

degree over the base field, and there will be exponentially many points in the kernel.

8.2 Compact representations of endomorphisms

We will now show that the endomorphism ring $\text{End}(E)$ of any supersingular elliptic curve E/\mathbb{F}_{p^2} has compact representations if $p \equiv 3 \pmod{4}$. The proof will require a special curve E_0 for which a basis of the endomorphism ring is known; such a curve exists if $p \not\equiv 1 \pmod{12}$.

For simplicity, we will focus on the case where $p \equiv 3 \pmod{4}$ is a prime and let $E_0 : y^2 = x^3 + x$. Let $\pi : E_0 \rightarrow E_0$ denote the Frobenius map, and let $\phi : E_0 \rightarrow E_0$ be the map $(x, y) \mapsto (-x, \sqrt{-1}y)$. The maps $1 + \phi\pi$ and $\phi + \pi$ both have kernels containing $E[2]$, so they factor through the map $[2] : E_0 \rightarrow E_0$. Let $(1 + \phi\pi)/2$ and $(\phi + \pi)/2$ represent the maps in these factorizations. It can be shown that $1, \phi, (1 + \phi\pi)/2, (\phi + \pi)/2$ form a basis for $\text{End}(E_0)$, see [GPS17]. As rational maps, the size of this basis may not be polynomial in $\log p$, but the description as rational linear combinations of $1, \phi, \pi, \phi\pi$ uniquely identifies them, and so it is enough that ϕ and π have polynomial size. This representation allows for efficient evaluation at points P of E_0 by writing $P = [2]Q$ and then evaluating linear combinations of $1, \phi, \pi, \phi\pi$ at Q . Define $[\beta_1, \beta_2, \beta_3, \beta_4] := [1, \phi, (1 + \phi\pi)/2, (\phi + \pi)/2]$. We will use $\beta_1, \beta_2, \beta_3, \beta_4$ in our definition of compact representatives of endomorphisms for all other supersingular elliptic curves E/\mathbb{F}_{p^2} .

Definition 1 (Compact representation of an endomorphism).

Let $p \equiv 3 \pmod{4}$ be a prime, let $E_0 : y^2 = x^3 + x$, and $\beta_1, \dots, \beta_4 := 1, \phi, (1 + \phi\pi)/2, (\phi + \pi)/2$ be the endomorphisms of E_0 as above. Let E/\mathbb{F}_{p^2} be another supersingular elliptic curve, and let $\rho \in \text{End}(E)$. Define a compact representation of ρ to be a list

$$[d, [c_1, \dots, c_4], [\phi_1, \dots, \phi_m], [\widehat{\phi}_1, \dots, \widehat{\phi}_m]],$$

where $c_1, \dots, c_4, d \in \mathbb{Z}$, ϕ_i are isogenies on a path from E_0 to E , the total size of the list

$$\log(|d|) + \log(|c_1|) + \dots + \log(|c_4|) + \sum_{i=1}^m \log(\deg(\phi_m))$$

is at most polynomial in $\log p$, and

$$\rho = \frac{1}{d} \left(\phi_m \circ \dots \circ \phi_1 \circ \left(\sum_{i=1}^4 c_i \beta_i \right) \circ \widehat{\phi}_1 \circ \dots \circ \widehat{\phi}_m \right).$$

Theorem 15. Let $p \equiv 3 \pmod{4}$ and let E/\mathbb{F}_{p^2} be a supersingular elliptic curve. Then there exist two lists of four compact representatives of endomorphisms of E , such that each list represents a \mathbb{Z} -basis of $\text{End}(E)$.

Moreover, assume $\rho \in \text{End}(E)$ is a linear combination of the endomorphisms corresponding to one such basis, and assume that its coefficient vector in terms of this basis is of size polynomial in $\log p$. Using the two lists, we can evaluate ρ at arbitrary points of E in time polynomial in $\log p$ and the size of the point P .

Proof. Let \mathcal{O}_0 be the maximal order in $B_{p,\infty}$ with basis

$$b_1, \dots, b_4 := 1, i, (1 + ij)/2, (i + j)/2.$$

Then $\mathcal{O}_0 \cong \text{End}(E_0)$ and b_1, \dots, b_4 correspond to β_1, \dots, β_4 under an isomorphism. There exist chains of isogenies ϕ_1, \dots, ϕ_m and ψ_1, \dots, ψ_n between E_0 and E with $\deg(\phi_k) = 2$ and $\deg(\psi_k) = 3$, and with $m, n = O(\log p)$. Set $\phi = \phi_m \circ \dots \circ \phi_1$ and $\psi = \psi_n \circ \dots \circ \psi_1$. Let $I \subseteq \mathcal{O}_0$ and $J \subseteq \mathcal{O}_0$ be the left \mathcal{O}_0 -ideals corresponding to ϕ and ψ respectively.

There exist rational numbers c_{rs}^I whose denominators are divisors of $2 \text{Nrd}(I)$ and rational numbers c_{rs}^J whose denominators are divisors of $2 \text{Nrd}(J)$ such that

$$\gamma_r^I := \sum_s c_{rs}^I b_s, 1 \leq r \leq 4$$

is a Minkowski-reduced basis of $\mathcal{O}_R(I)$, and

$$\gamma_r^J := \sum_s c_{rs}^J b_s, 1 \leq r \leq 4$$

is a Minkowski-reduced basis of $\mathcal{O}_R(J)$. This follows from Theorem 2 and its proof. We can also efficiently find $v \in B_{p,\infty}$ such that $v\mathcal{O}_R(I)v^{-1} = \mathcal{O}_R(J)$, see [KV10].

Then $\rho_r^J := \frac{1}{2^m} \phi \gamma_r^I \widehat{\phi}$ and $\rho_r^I := \frac{1}{3^n} \psi \gamma_r^J \widehat{\psi}$ ($r = 1, \dots, 4$) each form a basis for $\text{End}(E)$. Then our compact representations are, for $r = 1, \dots, 4$,

$$\begin{aligned} &[\text{Nrd}(I), c_{r1}^I, \dots, c_{r4}^I, [\phi_1, \dots, \phi_m], [\widehat{\phi}_1, \dots, \widehat{\phi}_m]], \\ &[\text{Nrd}(J), c_{r1}^J, \dots, c_{r4}^J, [\psi_1, \dots, \psi_n], [\widehat{\psi}_1, \dots, \widehat{\psi}_n]]. \end{aligned}$$

Observe that we can efficiently evaluate ρ_r^J at any point P of E whose order is coprime to 2. This is because $[2^m]\rho_r^I$ can be evaluated at P as it is a composition of the $\widehat{\phi}_k$, an integer linear combination of the β_k and then ϕ_k , all of which we can efficiently evaluate in terms of the size of P . Set $Q = [2^m]\rho_r^I(P)$. Let N be the inverse of 2^m modulo the order of P . Then $[N]Q = \rho_r^I(P)$.

If we want to evaluate ρ_r^I at a point P with $P \in E[2^f]$, we will instead express $v\rho_r^I v^{-1}$ as an integral linear combination of $\rho_1^J, \dots, \rho_4^J$. We can evaluate each $\rho_1^J, \dots, \rho_4^J$ at any point of order coprime to 3 by the same argument.

Thus we can evaluate at arbitrary points P : if P has order $2^f M$ with $(2, M) = 1$, then we can write P as a sum of a point P_2 of order 2^f and P_M of order M . We can then evaluate at P by evaluating it at each summand with the two above strategies. \square

Computing compact representations of endomorphisms which can be evaluated at points of E and which generate $\text{End}(E)$ is a natural interpretation of the problem of computing endomorphism rings, so we formally state it here before relating it to other isogeny problems.

Problem 6 (EndomorphismRing) *Given a prime p and a supersingular elliptic curve E/\mathbb{F}_{p^2} , find a list of total length bounded by $O(\log p)$ of compact representations of endomorphisms of E such that using this list, we can evaluate the corresponding endomorphisms at points of E , and such that the corresponding endomorphisms generate $\text{End}(E)$ as a \mathbb{Z} -module.*

In the next section, we will discuss two reductions from EndomorphismRing.

8.3 EndomorphismRing reduces to MaxOrder and Action-on-2-Torsion and Action-on-3-Torsion

In Algorithm 9, we used embeddings of endomorphism rings in $B_{p,\infty}$, together with their action on ℓ -torsion, to construct an ℓ -isogeny.

Theorem 16. *If $p \equiv 3 \pmod{4}$, EndomorphismRing reduces to MaxOrder and Action-on- ℓ -Torsion for $\ell = 2$ and 3.*

Proof. Let E be a supersingular elliptic curve. Let E_0 be the curve $y^2 = x^3 + x$ and let \mathcal{O}_0 be the order isomorphic to $\text{End}(E_0)$. By Theorem 15, the necessary data to give compact representations of generators of $\text{End}(E)$ is a 2-power and 3-power isogeny from E_0 to E , and a basis for the right orders of the ideals which correspond to these isogenies in $B_{p,\infty}$. In the proof of Theorem 10, note that all of this data is constructed using the oracles for MaxOrder, and Problems Action-on-2-Torsion and Action-on-3-Torsion. \square

8.4 EndomorphismRing reduces to an isogeny problem

We can also reduce the problem EndomorphismRing to a variant of the ℓ -Isogeny Problem, where we require the ℓ -power isogeny to be represented both by a chain of ℓ -isogenies and by a left ideal in a maximal order.

Problem 7 (FindKernelIdeal) *Given a prime p and a sequence of supersingular elliptic curves E_0, \dots, E_{m-1} and ℓ -isogenies $\phi_k : E_{k-1} \rightarrow E_k$, $k = 1, \dots, m$, with $m = O(\log p)$, along with a maximal order $\mathcal{O}_0 \subseteq B_{p,\infty}$ isomorphic to $\text{End}(E_0)$, compute the ideal I of $\mathcal{O}_0 \subseteq B_{p,\infty}$ corresponding to $\phi_m \circ \dots \circ \phi_1 : E_0 \rightarrow E_m$.*

Theorem 17. *Problem EndomorphismRing reduces in polynomial time to Problems ℓ -PowerIsogeny and FindKernelIdeal.*

Proof. Let E be a supersingular elliptic curve. Assume we are given ϕ_1, \dots, ϕ_m and ψ_1, \dots, ψ_n whose compositions are 2^m - and 3^n -isogenies $E_0 \rightarrow E$ and m, n are $O(\log p)$. Also assume we are given ideals A and B of \mathcal{O}_0 such that A is the kernel ideal of $\phi := \phi_m \circ \dots \circ \phi_1 : E_0 \rightarrow E$ and B is the kernel ideal of $\psi := \psi_m \circ \dots \circ \psi_1$. Then we can compute \mathbb{Z} -bases of $\mathcal{O}_R(A)$ and $\mathcal{O}_R(B)$. The sequences $\{\phi_r\}$ and $\{\psi_s\}$ for $r = 1, \dots, m$ and $s = 1, \dots, n$, along with \mathbb{Z} -bases of $\mathcal{O}_R(A)$ and $\mathcal{O}_R(B)$, give us the compact representations of generators of $\text{End}(E)$ constructed in the proof of Theorem 15. \square

Acknowledgments

We thank John Voight for many helpful discussions regarding orders in quaternion algebras and their connection with supersingular elliptic curves. We would also like to thank the anonymous referees for their helpful suggestions and corrections.

References

- [ACC⁺17] 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. <https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/Round-1-Submissions>.
- [BJS14] Jean-François Biasse, David Jao, and Anirudh Sankar. A quantum algorithm for computing isogenies between supersingular elliptic curves. In *Progress in cryptography—INDOCRYPT 2014*, volume 8885 of *Lecture Notes in Comput. Sci.*, pages 428–442. Springer, Cham, 2014.
- [Brö09] Reinier Bröker. Constructing supersingular elliptic curves. *J. Comb. Number Theory*, 1(3):269–273, 2009.
- [Cer04] J. M. Cerviño. Supersingular elliptic curves and maximal quaternionic orders. In *Mathematisches Institut, Georg-August-Universität Göttingen: Seminars Summer Term 2004*, pages 53–60. Universitätsdrucke Göttingen, Göttingen, 2004.
- [CG14] Ilya Chevyrev and Steven D. Galbraith. Constructing supersingular elliptic curves with a given endomorphism ring. *LMS J. Comput. Math.*, 17(suppl. A):71–91, 2014.
- [CGL06] Denis Charles, Eyal Goren, and Kristin Lauter. Cryptographic hash functions from expander graphs. Cryptology ePrint Archive, Report 2006/021, 2006. <https://eprint.iacr.org/2006/021>.
- [CGL09] Denis X. Charles, Eyal Z. Goren, and Kristin Lauter. Cryptographic hash functions from expander graphs. *J. Cryptology*, 22(1):93–113, 2009.
- [Cor08] G. Cornacchia. Su di un metodo per la risoluzione in numeri interi dell’equazione $\sum_{h=0}^n c_h x^{n-h} y^h = p$. *Giornale di Matematiche di Battaglini*, 46:33–90, 1908.
- [Deu41] Max Deuring. Die Typen der Multiplikatorenringe elliptischer Funktionenkörper. *Abh. Math. Sem. Univ. Hamburg*, 14(1):197–272, 1941.
- [DFJP14] Luca De Feo, David Jao, and Jérôme Plüt. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. *J. Math. Cryptol.*, 3(3):209–247, 2014.
- [DG16] Christina Delfs and Steven D. Galbraith. Computing isogenies between supersingular elliptic curves over \mathbb{F}_p . *Des. Codes Cryptogr.*, 78(2):425–440, 2016.
- [EHM17] Kirsten Eisenträger, Sean Hallgren, and Travis Morrison. On the hardness of computing endomorphism rings of supersingular elliptic curves. Cryptology ePrint Archive, Report 2017/986, 2017. <https://eprint.iacr.org/2017/986>.

- [Gal99] Steven D. Galbraith. Constructing isogenies between elliptic curves over finite fields. *LMS J. Comput. Math.*, 2:118–138, 1999.
- [GPS17] Steven D. Galbraith, Christophe Petit, and Javier Silva. Identification protocols and signature schemes based on supersingular isogeny problems. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology – ASIACRYPT 2017*, pages 3–33, Cham, 2017. Springer International Publishing.
- [GPST16] Steven D. Galbraith, Christophe Petit, Barak Shani, and Yan Bo Ti. On the security of supersingular isogeny cryptosystems. In *Advances in cryptology—ASIACRYPT 2016. Part I*, volume 10031 of *Lecture Notes in Comput. Sci.*, pages 63–91. Springer, Berlin, 2016.
- [Gro87] Benedict H. Gross. Heights and the special values of L -series. In *Number theory (Montreal, Que., 1985)*, volume 7 of *CMS Conf. Proc.*, pages 115–187. Amer. Math. Soc., Providence, RI, 1987.
- [GW17] Alexandre Gélín and Benjamin Wesolowski. Loop-abort faults on supersingular isogeny cryptosystems. In *Post-Quantum Cryptography*, Tanja Lange and Tsuyoshi Takagi, editors, pages 93–106, Cham, 2017. Springer International Publishing.
- [HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc. (N.S.)*, 43(4):439–561, 2006.
- [JDF11] 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.
- [KLPT14] 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.
- [Koh96] David Kohel. *Endomorphism rings of elliptic curves over finite fields*. PhD thesis, University of California, Berkeley, 1996.
- [KV10] Markus Kirschmer and John Voight. Algorithmic enumeration of ideal classes for quaternion orders. *SIAM J. Comput.*, 39(5):1714–1747, 2010.
- [Lan87] Serge Lang. *Elliptic functions*, volume 112 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1987. With an appendix by J. Tate.
- [LM04] Kristin Lauter and Ken McMurdy. Explicit generators of endomorphism rings of supersingular elliptic curves. Preprint, 2004.
- [LO77] J. C. Lagarias and A. M. Odlyzko. Effective versions of the Chebotarev density theorem. In *Algebraic number fields: L -functions and Galois properties (Proc. Sympos., Univ. Durham, Durham, 1975)*, pages 409–464. Academic Press, London, 1977.
- [Mes86] J.-F. Mestre. La méthode des graphes. Exemples et applications. In *Proceedings of the international conference on class numbers and fundamental units of algebraic number fields (Katata, 1986)*, pages 217–242. Nagoya Univ., Nagoya, 1986.
- [NIS16] NIST. Post-quantum cryptography, 2016. csrc.nist.gov/Projects/Post-Quantum-Cryptography; accessed 30-September-2017.
- [NS09] Phong Q. Nguyen and Damien Stehlé. Low-dimensional lattice basis reduction revisited. *ACM Trans. Algorithms*, 5(4):Art. 46, 48, 2009.
- [Pet17] Christophe Petit. Faster algorithms for isogeny problems using torsion point images. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology – ASIACRYPT 2017*, pages 330–353, Cham, 2017. Springer International Publishing.

- [Piz80] Arnold Pizer. An algorithm for computing modular forms on $\Gamma_0(N)$. *J. Algebra*, 64(2):340–390, 1980.
- [PL17] Christophe Petit and Kristin Lauter. Hard and easy problems for supersingular isogeny graphs. Cryptology ePrint Archive, Report 2017/962, 2017. <https://eprint.iacr.org/2017/962>.
- [Rón92] Lajos Rónyai. Algorithmic properties of maximal orders in simple algebras over \mathbf{Q} . *Comput. Complexity*, 2(3):225–243, 1992.
- [Sil09] Joseph H. Silverman. *The arithmetic of elliptic curves*. Springer, New York, 2009.
- [Sim05] Simon. Quadratic equations in dimensions 4, 5 and more. Preprint, 2005.
- [Ti17] Yan Bo Ti. Fault attack on supersingular isogeny cryptosystems. In *Post-quantum cryptography*, volume 10346 of *Lecture Notes in Comput. Sci.*, pages 107–122. Springer, Cham, 2017.
- [Vél71] Jacques Vélú. Isogénies entre courbes elliptiques. *C. R. Acad. Sci. Paris Sér. A-B*, 273:A238–A241, 1971.
- [Vig80] Marie-France Vignéras. *Arithmétique des algèbres de quaternions*, volume 800 of *Lecture Notes in Mathematics*. Springer, Berlin, 1980.
- [Voi] John Voight. *Quaternion Algebras*. Version v0.9.7, September 3, 2017.
- [Wat69] William C. Waterhouse. Abelian varieties over finite fields. *Ann. Sci. École Norm. Sup. (4)*, 2:521–560, 1969.
- [YAJ⁺17] Youngho Yoo, Reza Azarderakhsh, Amir Jalali, David Jao, and Vladimir Soukharev. A post-quantum digital signature scheme based on supersingular isogenies. In Aggelos Kiayias, editor, *Financial Cryptography and Data Security - 21st International Conference, FC 2017, Sliema, Malta, April 3-7, 2017, Revised Selected Papers*, volume 10322 of *Lecture Notes in Computer Science*, pages 163–181. Springer, 2017.