The Pennsylvania State University
The Graduate School

HORIZONTAL ISOGENIES AND ENDOMORPHISM RINGS OF SUPERSINGULAR ELLIPTIC CURVES

A Dissertation in
Mathematics
by
Gabrielle Scullard

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

August 2024

The dissertation of Gabrielle Scullard was reviewed and approved by the following:

Kirsten Eisenträger
Professor of Mathematics
Dissertation Advisor
Chair of Committee

Mihran Papikian
Professor of Mathematics
Committee Member

Wen-Ching Winnie Li
Professor of Mathematics
Committee Member

Antonio Blanca
Assistant Professor of Computer Science
Outside Committee Member

Pierre-Emmanuel Jabin
Professor of Mathematics
Graduate Program Head

Abstract

This dissertation considers two problems regarding the structure of isogenies between supersingular elliptic curves which are motivated by isogeny-based cryptography. The first problem is inspired by the setup of OSIDH, a recent proposal for an isogeny-based protocol whose security relies on the hardness of finding a "horizontal" isogeny between two "oriented" supersingular elliptic curves. A basic question is to ask if it is harder to find horizontal isogenies between oriented curves than it is to find isogenies between oriented curves. Under certain conditions, the answer is no. We give conditions for which all or most isogenies of fixed degree between oriented supersingular elliptic curves are horizontal, and we classify the exceptions. Our work can be applied to extend an attack on OSIDH by Dartois and de Feo.

The second problem is to compute the endomorphism ring of a supersingular elliptic curve. Most problems in isogeny-based cryptography have polynomial-time reductions to computing the endomorphism ring of a supersingular elliptic curve, so an efficient algorithm for computing endomorphism rings of supersingular elliptic curves would have consequences for isogeny-based protocols. We give a deterministic polynomial time algorithm which computes the endomorphism ring of a supersingular elliptic curve from the input of two noncommuting isogenies. Our algorithm uses techniques of higher-dimensional isogenies to navigate towards the local endomorphism ring in the Bruhat-Tits tree.

Contents

List of Figures

Acknowledgments

I am immensely grateful to have been surrounded by and supported by mentors, collaborators, and friends throughout my time in graduate school. First and foremost, I would like to thank my advisor, Kirsten Eisenträger, for her encouragement, guidance, and support. She has been an incredible teacher, collaborator, and advocate. I also thank Yuri Zarhin for fruitful discussion about higher dimensional abelian varieties; Damien Robert for helpful conversations about Algorithm 6.1.9; and Travis Morrison, who provided the proof of Proposition 3.3.6 in an Email correspondence.

I am deeply indebted to my academic siblings, Caleb Springer and Travis Morrison, for their kindness and mentorship as I entered the world of number theory and arithmetic geometry. I'm also grateful for the personal support I've received from my fellow graduate students and friends: Jeff Katen, Jacob Bradd, Matt Bernstein, Steve White, Gözde Sert, Sebastian Calvo, Ufuoma Asarhasa, William Noland, Alex Safsten, Dominic Veconi, Ana Cristina Chavez Caliz, Sergio Zamora Barrera, Caitlin Lienkamper, Jonathan Love, Ben Dees, Xiaoling Chen, Niru Murali, and so many others. Finally, I'm grateful for my the support of my family: my mother, Nhi; my father, Tim; my stepmother, Marilou; my siblings, Abigail, Alexandra, Michael, and Vanessa; and my many aunts, uncles, and cousins.

I was supported in part by the Maryam Mirzakhani Science Achievement Graduate Scholarship in Mathematics and National Science Foundation Grant CNS-2001470. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

Material from Chapters 5 and 6 was produced in collaboration with Kirsten Eisenträger.

Dedication

To my mom, who worked so hard for my entire childhood and still helped me study when I failed my first timed multiplication test.

Chapter 1
Introduction

1.1 Motivation: Isogeny-based cryptography

The problems considered in this dissertation are motivated by isogeny-based cryptography. Today, all widely-used public key cryptosystems are based on factoring or discrete log problems which can be solved efficiently by quantum algorithms [46]. This makes it necessary to investigate post-quantum alternatives. In 2016, NIST began a process to solicit, assess, and standardize public-key cryptographic protocols which would be secure against attacks by quantum computers. At the end of its fourth round in 2022, four cryptographic schemes were chosen for standardization. Three of the schemes chosen for standardization are from lattice-based cryptography and are based on similar hard problems, so it is important to continue studying diverse alternatives. The only isogeny-based proposal was SIDH (Supersingular Isogeny Diffie-Hellman) [5], which had comparatively smaller key sizes. It was selected as an alternate candidate for further evaluation at the end of the fourth round. However, SIDH was broken later that year by a series of groundbreaking attacks [10, 33, 42] and could no longer be considered for standardization.

Despite the break of SIDH, isogenies are still of cryptographic interest. The fundamental hard problem for isogeny-based cryptography is the ℓ-isogeny path-finding problem: Given two supersingular elliptic curves, compute an ℓ-power degree isogeny between them. This problem can be phrased as finding a path in the ℓ-isogeny graph, whose vertices are supersingular elliptic curves and whose edges represent ℓ-isogenies between supersingular elliptic curves. The ℓ-isogeny graph is an expander graph [39]. In SIDH, one is also told the image of certain points under the desired isogeny; this is the key information exploited by the attacks against SIDH. This means that the fundamental hard problem for isogeny-based cryptography is not broken and is not made easier by the techniques of the SIDH attacks, leaving many isogeny-based protocols unaffected. A related problem, which is equivalent to the path-finding

problem under polynomial time reductions [57], is the problem of computing endomorphism rings of supersingular elliptic curves. An algorithm to solve either problem efficiently would affect all isogeny-based cryptosystems [20, 57, 58], but the current best algorithms have exponential complexity.

This dissertation highlights the interplay between these two problems. First, we will consider a variant of the path-finding problem which underlies the security of OSIDH (Oriented Supersingular Isogeny Diffie-Hellman). In this setting, we only consider curves which are "oriented" by a fixed quadratic imaginary ring $\mathcal{O}$. This gives information about the structure of the endomorphism ring, which we use to characterize the structure of isogenies between $\mathcal{O}$-oriented curves. While we are motivated by the setting of OSIDH, our results apply more generally. In the second part of our dissertation, we will give a deterministic algorithm for computing the endomorphism ring of a supersingular elliptic curve from the input of a full-rank subring. We use techniques similar to those used to break SIDH in order to navigate to the endomorphism ring locally in the Bruhat-Tits tree.

1.2 Oriented elliptic curves and horizontal isogenies

In 2020, Colò and Kohel introduced OSIDH [14], a public key exchange protocol based on the action of the class group $\mathcal{CL}(\mathcal{O})$ on the set of supersingular elliptic curves which are "oriented" by a fixed quadratic imaginary ring $\mathcal{O}$. These are exactly the elliptic curves obtained by reducing CM elliptic curves modulo an appropriate prime, and the action of the class group is inherited from the action on CM elliptic curves. Isogenies which are in the image of the reduction map, corresponding to action by an invertible ideal, are called "horizontal." A more precise definition of both terms will be given in Chapter 3. Many other isogeny-based cryptosystems are based on this group action for different orders $\mathcal{O}$. The first isogeny-based cryptosystem, CRS [16, 45, 49], used ordinary elliptic curves, which are oriented by their endomorphism rings; and CSIDH [11] uses a set of $\mathbb{F}_p$-rational supersingular elliptic curves, which are oriented by $\mathbb{Z}[\sqrt{-p}]$. OSIDH was the first public key exchange to consider this structure in a more general setting, motivated by replacing the ring $\mathbb{Z}[\sqrt{-p}]$ in CSIDH by a more general quadratic imaginary order $\mathcal{O}$ in the hopes of increasing the number of curves in the keyspace.

In 2021, Dartois and De Feo [18] described an attack on OSIDH which makes use of the public information about the group action to construct a horizontal endomorphism. However, they rely on heuristics to argue that they can construct a horizontal endomorphism from the public group action alone [18, Section 4]. If β cannot be constructed this way, then one can construct a horizontal isogeny using the public group action and supplement by

using a meet-in-the-middle attack, which *a priori* does not produce a horizontal isogeny, only an isogeny of the correct degree. Even under their heuristic assumptions, the attack has exponential complexity and relies on solving a shortest vector problem in a particular lattice, so OSIDH is still of interest from a security perspective. In addition to the attack, Dartois and De Feo propose a set of countermeasures which they suggest could even make OSIDH secure against subexponential complexity quantum algorithms for solving hidden shift problems, which are known to apply to CRS and CSIDH. We will give conditions for which an isogeny of degree N is likely to be horizontal, which extends Dartois and De Feo's attack in the event that the heuristics do not hold.

In addition to being useful for constructing cryptographic schemes (in addition to those mentioned, see [29] and [22]), orientations have been used to compute endomorphism rings of supersingular elliptic curves [3, 24, 38, 58]. Most work in this direction assumes knowledge of the orientation. However, a recent preprint shows that under certain conditions, knowing that a curve is orientable (without the full data of an orientation) can be used to compute the orientation [4].

In Chapter 3, we state and give new proofs of basic facts about $\mathcal{O}$-oriented supersingular curves and the class group action. While they are "basic" facts, they were only proved recently and highlight the differences between the ordinary case and the supersingular case. The main results of this chapter are not new, but we give original proofs. We show that any supersingular elliptic curve is oriented by infinitely many quadratic imaginary orders $\mathcal{O}$, and we show that there are quadratic imaginary orders $\mathcal{O}$ such that every supersingualar elliptic curve is $\mathcal{O}$-oriented; a proof of this statement was given in [29, Proposition 5.10], but it does not work for all primes p, so we give a corrected proof. It was proved in [2] that the group action is free and that the number of orbits (one or two) depends on the splitting behavior of p in $\mathcal{O}$. We re-prove this statement, along the way proving some lemmas which appear to be new and may be of independent interest.

In Chapter 4, we restrict our attention to the setup of OSIDH, although our techniques can be applied more generally. The security of OSIDH depends on the hardness of the following problem: Given E, $\mathfrak{b} * E$, and information about how to compute the group action (see [14, Section 5.2, page 23]), recover the ideal class $\mathfrak{b}$. Onuki observed (and Dartois and De Feo made more explicit) that one can recover $\mathfrak{b}$ by constructing a non-integer endomorphism β which generates the image of $\mathcal{O}$ in $\mathrm{End}(\mathfrak{b} * E)$ ([37, Section 6.3] and [18, Section 4]). In the specific setting of OSIDH, this condition on β is equivalent to β being horizontal.

We study the structure of isogenies, not necessarily horizontal, between $\mathcal{O}$-oriented elliptic curves and give conditions for which any or most isogenies of the correct degree are likely to be horizontal. We use an explicit description of endomorphism rings of $\mathcal{O}$-oriented curves

given by Lauter and Viray [31]. As an application, we extend the attack of Dartois and De Feo in the event that the heuristics in Dartois and De Feo's attack do not hold.

We prove the following, which is our main tool for understanding horizontal N-isogenies.

Theorem 1.2.1. *Let (E, ι) be an $\mathcal{O}$-oriented supersingular elliptic curve over $\overline{\mathbb{F}}_p$, where $\mathcal{O}$ is a quadratic imaginary order of discriminant D, and let K be the field of fractions of $\mathcal{O}$. Let $\mathfrak{b}$ be an invertible ideal in $\mathcal{O}$ with norm coprime to p, and let E' such that $\mathfrak{b} * (E, \iota) = (E', \iota')$. Let N be a positive integer. There are primitive positive definite binary quadratic forms Q and Q' of discriminant D such that the following hold.*

1. *If p is inert in K, there is an injective map*

$$\{N\text{-isogenies } \phi : E \to E'\} \to \{(w, x, y, z) \in \mathbb{Z}^4 : Q(w, x) + pQ'(y, z) = -DN\}.$$

2. *If p ramifies in K, there is an injective map*

$$\{N\text{-isogenies } \phi : E \to E'\} \to \{(w, x, y, z) : Q(w, x) + Q'(y, z) = -DN/p\}.$$

In the setting of OSIDH, horizontal isogenies correspond to the solutions with $y = z = 0$, and it is straightforward to describe all ways that non-horizontal N-isogenies may arise. Moreover, in the special case that p is inert in $\mathcal{O}$ and $-DN < 2p$, this map extends to a bijection up to a certain equivalence, and in this case, we give a more complete classification of N-isogenies between $\mathcal{O}$-oriented supersingular elliptic curves.

1.3 Computing endomorphism rings of supersingular elliptic curves

Outside of cryptographic interest, computing the endomorphism ring of an elliptic curve is a fundamental problem in computational arithmetic geometry. Faster algorithms are known in the ordinary case. Bisson and Sutherland [9] gave a subexponential time algorithm to compute the endomorphism ring of an ordinary curve under certain heuristics, later improved to rely only on GRH [8]. Recently, Robert outlined an algorithm to compute the endomorphism ring of an ordinary elliptic curve in polynomial time, given the factorization of its discriminant [43, Theorem 4.2].

For supersingular elliptic curves, the problem is more complex. In the ordinary case, there is a convenient subring of finite index which is generated by the Frobenius endomorphism, which is no longer true for supersingular elliptic curves. One approach to endomorphism ring

computation is to first compute a subring of finite index, then to compute the endomophism ring from the subring. The first step was first accomplished by Kohel, who gave an algorithm for generating a subring of finite index by finding cycles in the ℓ-isogeny graph of supersingular elliptic curves in characteristic p [27, Theorem 75], running in time $O(p^{1+\varepsilon})$. In [23], it was shown that a Bass suborder of $\mathrm{End}(E)$ can be computed in $O(p^{1/2+\varepsilon})$ time, assuming GRH. Our algorithm focuses on the second step, computing the endomorphism ring from the input of a subring of finite index. We give a deterministic polynomial time algorithm that computes the endomorphism ring of a supersingular elliptic curve, given two non-commuting endomorphisms and a factorization of the reduced discriminant of the order they generate. We build on [21] which gave a subexponential algorithm, under certain heuristics, if the input suborder was Bass.

Our algorithm is an improvement and a generalization of the main result in [21], which we summarize briefly here. Given a Bass suborder $\mathcal{O}_0$ of the endomorphism ring $\mathrm{End}(E)$ and a factorization of the reduced discriminant, they compute all local maximal orders containing $\mathcal{O}_0 \otimes \mathbb{Z}_q$ at each prime q dividing the reduced discriminant. They then combine local maximal orders to obtain all global maximal orders containing $\mathcal{O}_0$ and check each one until they find $\mathrm{End}(E)$. In the general case, $\mathcal{O}_0$ may be contained in exponentially many maximal orders, so the Bass restriction is needed to reasonably bound the number of maximal orders.

We also approach the problem locally. However, we are able to compute the local maximal order $\mathrm{End}(E) \otimes \mathbb{Z}_q$ without constructing all global candidates for $\mathrm{End}(E)$. We also do not require the Bass restriction, although we give a more efficient algorithm in the Bass case.

There are two key tools. The first is a polynomial-time algorithm which, when given an endomorphism β and an integer n, determines if $\frac{\beta}{n}$ is an endomorphism. This algorithm is implicit in Robert's algorithm for computing an endomorphism ring of an ordinary elliptic curve [43, Section 4]. A detailed proof and runtime analysis are given in [24, Section 4]. This technique is also used in [38], which gives a probabilistic polynomial time reduction to compute the endomorphism ring from an oracle which computes a non-scalar endomorphism. In contrast, our deterministic algorithm applies this algorithm to test if certain local orders are contained in the endomorphism ring.

The second tool is a theorem by Tu [51, Theorem] which expresses an intersection of finitely many maximal orders in $M_2(\mathbb{Q}_q)$ as an intersection of at most three maximal orders, which can be constructed explicitly [51, Theorem 8]. This theorem allows us to relate the placement of the local endomorphism ring in the Bruhat-Tits tree to its containment of certain local orders, which allows us to use the containment testing to find the local endomorphism ring in the Bruhat-Tits tree.

In Chapter 6, we describe the algorithm and give a proof of correctness and complexity

analysis.

Theorem 1.3.1. *There exists an algorithm that computes the endomorphism ring of a supersingular elliptic curve E when given E, two noncommuting endomorphisms α and γ which can be evaluated efficiently on powersmooth torsion points, and a factorization of the reduced discriminant Δ of the order generated by α and γ. The algorithm runs in polynomial time in $\log p$, the size of α and γ, $\log \Delta$, and is linear in the number of primes dividing Δ/p and the largest prime dividing Δ/p.*

Our algorithm enlarges the order $\mathcal{O}_0$ generated by α and γ by finding the appropriate local maximal order at each prime q dividing Δ/p. The complexity of this step is linear in q. At the cost of a higher degree polynomial complexity in q, one could instead enlarge $\mathcal{O}_0$ at q by searching for an endomorphism in $\mathcal{O}_0$ which is divisible by q and adjoining its division by q.

Chapter 2
Preliminaries

In this chapter, we briefly review background on supersingular elliptic curves and their endomorphism rings which will be used throughout.

2.1 Elliptic Curves

2.1.1 Elliptic curves

Definition 2.1.1. An elliptic curve is a smooth, projective curve of genus one, together with a distinguished point.

In this text, we will mainly be concerned with elliptic curves defined over a finite field of large characteristic, so we consider a more concrete definition: An elliptic curve defined over a field K of characteristic not equal to 2 or 3 is the set of solutions (x, y) to an equation in Weierstrass form

$$E : y^2 = x^3 + Ax + B$$

with $A, B \in K$ such that the function $x^3 + Ax + B$ has no repeated roots, together with a point at infinity denoted O_E. There is a group law on the points of E.

For an elliptic curve $E : y^2 = x^3 + Ax + B$, we define the **j-invariant** of E to be $j(E) := 4A^3 + 27B^2$. If E is defined over K, then it is clear that the j-invariant is defined over K. The converse is true as well: If the j-invariant is an element of K, then there is a model for E with coefficients in K. The j-invariant classifies elliptic curves up to isomorphism over the algebraic closure of K.

2.1.2 Isogenies and endomorphisms

One of the fundamental properties of elliptic curves is that the points of an elliptic curve form an abelian group under a geometric group law, where O_E is the identity. For a positive

integer n and a point P, we let $[n]P := \underbrace{P + P + \ldots + P}_{\text{n times}}$ and $[-n]P := [n](-P)$, where $-P$ is the inverse of P. The multiplication-by-n map,

$$[n] : P \mapsto [n]P,$$

is an example of an endomorphism.

Definition 2.1.2. Let E and E' be elliptic curves. An **isogeny** $\phi : E \to E'$ is a morphism such that $\phi(O_E) = O_{E'}$. When $E = E'$, we call ϕ an **endomorphism**. If there exists an isogeny between E and E', we say E and E' are **isogenous**.

In addition to being geometric morphisms, isogenies are group homomorphisms. If ϕ is nonconstant, then ϕ is surjective and has finite kernel. An isogeny can be computed explicitly from its kernel using Vélu's formula.

For a separable isogeny $\phi : E \to E'$, the degree $\deg(\phi)$ is the size of the kernel. Over a field of characteristic zero, all isogenies are separable. Over a field of characteristic $p > 0$, every isogeny $\phi : E \to E'$ factors as $\phi = \phi_{\text{sep}} \circ \phi_{\text{insep}}$, where ϕ_{insep} is the p^r-th power Frobenius endomorphism and ϕ_{sep} is separable (see [48, Chapter II, Corollary 2.12]). In this case, $\deg(\phi) = p^r \deg(\phi_{\text{sep}})$.

If $\deg(\phi) = N$, we call ϕ an N-isogeny. Each isogeny $\phi : E \to E'$ comes with a dual $\hat{\phi} : E' \to E$. The dual $\hat{\phi}$ is the map such that $\phi\hat{\phi} = [\deg(\phi)]$.

2.1.3 Endomorphism rings of elliptic curves

A fundamental object of interest is the endomorphism ring.

Definition 2.1.3. Let E be an elliptic curve. The **endomorphism ring** of E, denoted $\text{End}(E)$, is the ring consisting of all endomorphisms of E defined over an algebraic closure.

One can restrict to endomorphisms defined over a finite field and potentially obtain different structures, so we emphasize that in this dissertation, we consider all endomorphisms.

For any elliptic curve E, there is an injection $\mathbb{Z} \to \text{End}(E)$ given by $n \mapsto [n]$. When E is defined over a finite field with p^k elements, $\text{End}(E)$ contains the **Frobenius endomorphism** $\pi_E : (x, y) \mapsto (x^{p^k}, y^{p^k})$.

The endomorphism ring is isomorphic to one of $\mathbb{Z}$, a quadratic imaginary order, or a maximal order in a quaternion algebra [48, Chapter III, Corollary 9.4].

2.1.4 Supersingular elliptic curves

For most of this dissertation, we will be concerned with elliptic curves over finite fields. For these elliptic curves, the endomorphism ring is larger than $\mathbb{Z}$. We can classify curves based on the kinds of endomorphism rings they have, in the following way:

Definition 2.1.4. An elliptic curve is called **ordinary** if $\mathrm{End}(E)$ is isomorphic to a quadratic imaginary order and **supersingular** if $\mathrm{End}(E)$ is a maximal order in a quaternion algebra.

One can show that all supersingular elliptic curves over an algebraic closure of a finite field of characteristics p are defined over $\mathbb{F}_{p^2}$.

2.2 Isogeny graphs

Definition 2.2.1. The supersingular ℓ-isogeny graph over $\overline{\mathbb{F}}_p$ is the directed graph whose vertices are isomorphism classes of supersingular elliptic curves and edges between vertices correspond to ℓ-isogenies between the corresponding curves, up to equivalence. The isogenies $\phi : E_0 \to E_1$ and $\phi' : E_0 \to E_1$ are considered equivalent if and only if there is an isomorphism $\psi : E_1 \to E_1$ such that $\psi\phi = \phi'$. We denote this graph by $G_\ell(p)$.

There are approximately $\frac{p-1}{12}$ vertices (see [48, Theorem 4.1(c)]). In the setting of Theorem 1.2.1, we consider a subset of the vertices, and when $N = \ell^k$, we are considering walks of length k which begin and end at a vertex in this subset.

When $p \equiv 1 \pmod{12}$, we can identify isogenies with their duals and view $G_\ell(p)$ as an undirected graph. In this case, Pizer showed that the graph is Ramanujan [39]. This means that $G_\ell(p)$ is $(\ell+1)$-regular and has largest possible spectral gap, and as a consequence, the ℓ-isogeny graph is an expander with rapid mixing properties. When $p \not\equiv 1 \pmod{12}$, we can no longer identify duals due to vertices with extra automorphisms (corresponding to curves with j-invariants 0 and 1728). However, the ℓ-isogeny graph still has out-degree $\ell+1$ and has rapid mixing properties (see [7, Theorem 3] with $d = 1$ for example).

Starting at a vertex, it is easy to compute neighbors via Vélu's formula. The fundamental hard problem for isogeny-based cryptography is the ℓ-isogeny pathfinding problem: find a path between two given vertices in the ℓ-isogeny graph.

We can compare the supersingular ℓ-isogeny graph with the analagous graph for ordinary curves defined over a finite field $\mathbb{F}_{p^r}$. The connected components of the ordinary ℓ-isogeny graph form a volcano (see [50]).

In Chapters 3 and 4, we will consider the structure of isogenies of fixed degree between *oriented* supersingular elliptic curves. For ℓ-isogenies, this can be viewed as considering a

subgraph of the ℓ-isogeny graph. We will show when all such isogenies are *horizontal*, which is precisely when this subgraph represents the ordinary case.

2.3 Quaternion algebras

In this section, we describe background on quaternion algebras.

Definition 2.3.1. Let R be a domain with field of fractions F, and let B be a finite-dimensional F-algebra. A subset $M \subseteq V$ is an *R-lattice* if M is finitely generated as an R-module and $MF = B$. An **R-order** $\mathcal{O} \subseteq B$ is an R-lattice that is also a subring of B. An order is **maximal** if it is not properly contained in another order.

Definition 2.3.2. Let F be $\mathbb{Q}$ or $\mathbb{Q}_q$ (with $q > 2$). For $a, b \in F^\times$, let $H(a, b)$ denote the quaternion algebra over F with basis $1, i, j, ij$ such that $i^2 = a$, $j^2 = b$ and $ij = -ji$. That is, $H(a, b) = \mathbb{Q} + \mathbb{Q}\,i + \mathbb{Q}\,j + \mathbb{Q}\,ij$.

Any quaternion algebra over F can be written in this form. For the case of $F = \mathbb{Q}_2$, see [54, Section 6.2].

There is a **standard involution** on $H(a, b)$ which maps $\alpha = a_1 + a_2 i + a_3 j + a_4 ij$ to its **conjugate** $\overline{\alpha} := a_1 - a_2 i - a_3 j - a_4 ij$. The **reduced trace** of such an element α is defined as $\mathrm{Trd}(\alpha) = \alpha + \overline{\alpha} = 2a_1,$. The **reduced norm** is $\mathrm{Nrd}(\alpha) = \alpha\overline{\alpha} = a_1^2 - aa_2^2 - ba_3^2 + aba_4^2$.

We say that a quaternion algebra B over $\mathbb{Q}$ **ramifies** at a prime q (respectively ∞) if $B \otimes \mathbb{Q}_q$ (respectively $B \otimes \mathbb{R}$) is a division algebra. Otherwise, B is said to be **split** at q. In this case, $B \otimes \mathbb{Q}_q \cong M_2(\mathbb{Q}_q)$.

The **discriminant** of a quaternion algebra B, denoted $\mathrm{disc}\, B$, is the product of the finite primes ramifying in B. For an order $\mathcal{O} \subset B$ with $\mathbb{Z}$-basis $\{\alpha_1, \alpha_2, \alpha_3, \alpha_4\}$, the **discriminant** of $\mathcal{O}$ is defined to be $\mathrm{disc}(\mathcal{O}) \coloneqq |\det(\mathrm{Trd}(\alpha_i\alpha_j))_{i,j}| \in \mathbb{Z} > 0$ [54, p. 242]. The discriminant of an order is always a square. The **reduced discriminant** $\mathrm{discrd}(\mathcal{O})$ is the positive integer square root of $\mathrm{disc}(\mathcal{O})$ [54, p. 242].

For a supersingular elliptic curve E defined over a finite field of characteristic p, the endomorphism ring $\mathrm{End}(E)$ is isomorphic to a maximal order of the quaternion algebra $B_{p,\infty}$, the unique (up to isomorphism) quaternion algebra ramified at the primes p and ∞.

Chapter 3
Properties of supersingular oriented elliptic curves

For this chapter, $\mathcal{O}$ will always refer to a quadratic imaginary order, with K its field of fractions. Unless otherwise stated, all elliptic curves in this chapter are supersingular.

We prove some basic facts about oriented supersingular elliptic curves which highlight the difference between the ordinary case and the supersingular case. The main results of this chapter are not new, but we give different proofs from those appearing in the literature.

In Section 3.1, we give background for quadratic imaginary orders and their class groups, which play a central role in the discussion of oriented elliptic curves. In Section 3.2, we give background on oriented curves and the action of the class group on them.

In Section 3.3, we show that $\mathrm{End}(E)$ admits optimal embeddings of infinitely many orders when E is supersingular. In fact, for any quadratic imaginary field K in which p does not split, $\mathrm{End}(E)$ admits optimal embeddings of infinitely many orders contained in K. This is in contrast to the ordinary case, where the order is uniquely determined by E.

In Section 3.4, we show that the action of the class group defined in Chapter 3 is transitive when p is ramified in K and has two orbits when p is inert in K. A proof of this fact appears in [2, Proposition 4.2, Proposition 4.3]. We give a different proof, along the way proving new statements about when optimal embeddings which are complex-conjugate or Galois-conjugate are K-isomorphic or appear in the same orbit.

3.1 Quadratic imaginary orders

In this section, we define the class group of a (not necessarily maximal) quadratic imaginary order and describe a bijection between ideals in the class group and binary quadratic forms. The theory of binary quadratic forms is classical and was developed in large part by Gauss.

Many details are omitted for simplicity of exposition; see [17, Chapter 3] for a complete treatment of the topic.

3.1.1 Definitions and Notation

An order in a quadratic imaginary field K is a subring $\mathcal{O} \subset K$ such that $\mathcal{O}$ is a finitely-generated $\mathbb{Z}$-module and $\mathcal{O} \otimes \mathbb{Q} = K$. The maximal order $\mathcal{O}_K$ is the set of integral elements of K.

Definition 3.1.1. Let $\mathcal{O}$ be a quadratic imaginary order. The **conductor** of $\mathcal{O}$ is its index $f = [\mathcal{O}_K : \mathcal{O}]$. If $\mathcal{O}$ has basis $\{\alpha_1, \alpha_2\}$ as a $\mathbb{Z}$-module, the **discriminant** of $\mathcal{O}$ is

$$\mathrm{disc}(\mathcal{O}) = \left(\det \begin{pmatrix} \alpha_1 & \alpha_2 \\ \overline{\alpha_1} & \overline{\alpha_2} \end{pmatrix} \right)^2,$$

where the bar denotes complex conjugation.

It is well-known that every order $\mathcal{O} \subset K$ is of the form $\mathcal{O} = \mathbb{Z} + f\mathcal{O}_K$, where f is the conductor. If $\mathcal{O}$ has conductor f, it is not hard to see that $\mathrm{disc}(\mathcal{O}) = f^2 D_K$, where $D_K = \mathrm{disc}(\mathcal{O}_K)$, and D is a negative integer with $D \equiv 0, 1 \pmod 4$. (See [17, Lemma 7.2] and the discussion immediately following the proof for details.)

The main subject of interest will be actions of ideals on a set of elliptic curves whose endomorphism rings contain a copy of $\mathcal{O}$ in an "optimal" way, which we will make precise later.

For an ideal $\mathfrak{a} \subset \mathcal{O}$, the quotient $\mathcal{O}/\mathfrak{a}$ is finite. We define the **norm** of an ideal, denoted $N(\mathfrak{a})$, to be $|O/\mathfrak{a}|$.

Definition 3.1.2. Let $\mathcal{O}$ be an order. A **fractional ideal of** $\mathcal{O}$ is a subset of K which is a nonzero finitely-generated $\mathcal{O}$-module. A fractional $\mathcal{O}$-ideal $\mathfrak{a}$ is **principal** if $\mathfrak{a} = \alpha\mathcal{O}$ for some $\alpha \in K^*$, and $P(\mathcal{O})$ denotes the set of principal fractional ideals of $\mathcal{O}$. A fractional ideal is **invertible** if there exists a fractional $\mathcal{O}$-ideal $\mathfrak{b}$ such that $\mathfrak{a}\mathfrak{b} = \mathcal{O}$, and $I(\mathcal{O})$ denotes the set of invertible ideals of $\mathcal{O}$.

One can show that any fractional ideal is of the form $\alpha\mathfrak{a}$ for $\alpha \in K^*$ and an integral ideal $\mathfrak{a}$. We extend the definition of the norm in the natural way, to define $N(\alpha\mathfrak{a}) = N_{\mathbb{Q}}^K(\alpha)N(\mathfrak{a})$.

Definition 3.1.3. The **class group** of $\mathcal{O}$ is the quotient $\mathcal{CL}(\mathcal{O}) := I(\mathcal{O})/P(\mathcal{O})$. An element of $\mathcal{CL}(\mathcal{O})$ is called an **ideal class**.

3.1.2 Properties of class groups

We list some basic properties of the class group which will be used later. It is well-known that $\mathcal{CL}(\mathcal{O})$ is an abelian group under ideal multiplication.

For a non-maximal order, it is no longer true that all ideals are invertible.

Proposition 3.1.4. *[17, Lemma 7.18, Proposition 7.4] Let f be the conductor of $\mathcal{O}$. Every $\mathcal{O}$-ideal $\mathfrak{a}$ with $N(\mathfrak{a})$ coprime to f is invertible.*

In the other direction, starting with an ideal class, we can always choose a representative of norm coprime to the conductor.

Proposition 3.1.5. *[17, Proposition 7.19] Let f be the conductor of $\mathcal{O}$. Let $I(\mathcal{O}, f)$ denote the subgroup of $I(\mathcal{O})$ generated by ideals of norm coprime to f. Let $P(\mathcal{O}, f)$ denote the subgroup of $I(\mathcal{O}, f)$ generated by principal ideals $\alpha\mathcal{O}$ where $\alpha \in \mathcal{O}$ has norm $N(\alpha)$ coprime to f. Then the inclusion map $I(\mathcal{O}, f) \to I(\mathcal{O})$ induces an isomorphism:*

$$I(\mathcal{O}, f)/P(\mathcal{O}, f) \cong I(\mathcal{O})/P(\mathcal{O}) = \mathcal{CL}(\mathcal{O}).$$

In fact, we may choose a representative to avoid any finite number of primes.

Proposition 3.1.6. *[15, Theorem 5.2] Let $[\mathfrak{a}] \in \mathcal{CL}(\mathcal{O})$, and let $\mathfrak{n}$ be a nonzero integral ideal in $\mathcal{O}$. A representative $\mathfrak{a}$ for the ideal class can be chosen so that $\mathfrak{a}$ is integral and $\mathfrak{a}$ is coprime to $\mathfrak{n}$.*

3.2 Oriented elliptic curves

3.2.1 Definitions

Definition 3.2.1. An **$\mathcal{O}$-oriented elliptic curve** (E, ι) is an elliptic curve E defined over $\overline{\mathbb{F}}_p$ with an embedding $\iota : K \to \mathrm{End}(E) \otimes \mathbb{Q}$ such that $\iota(K) \cap \mathrm{End}(E) = \iota(\mathcal{O})$. The embedding ι is called an **optimal embedding** of $\mathcal{O}$ into $\mathrm{End}(E)$.

Remark 3.2.2. In other papers, what we call an "optimal embedding" is often referred to as a *primitive orientation*, and what we call "$\mathcal{O}$-oriented" is often called "*primitively $\mathcal{O}$-oriented*."

Definition 3.2.3. Let (E, ι) be an $\mathcal{O}$-oriented elliptic curve, and let $\phi : E \to E'$ be an isogeny. The **pushforward** of ι by ϕ is denoted by

$$\phi_* \iota = \frac{1}{\deg(\phi)} \phi \iota \hat{\phi}.$$

If $\phi_*\iota$ is also an optimal embedding of $\mathcal{O}$ into $\mathrm{End}(E')$, i.e. $\phi_*\iota(K) \cap \mathrm{End}(E') = \phi_*\iota(\mathcal{O})$, then ϕ is **horizontal**. Otherwise, $\phi_*\iota$ is optimal for a different order. It is called **ascending** if $\phi_*\iota(K) \cap \mathrm{End}(E') \supsetneq \phi_*\iota(\mathcal{O})$ or **descending** if $\phi_*\iota(K) \cap \mathrm{End}(E') \subsetneq \phi_*\iota(\mathcal{O})$.

Remark 3.2.4. If $\deg(\phi) = \ell$ is prime, then ϕ can only be horizontal, ascending, or descending, and $\phi_*\iota$ is optimal for an order with relative index 1 or ℓ in $\mathcal{O}$ (see [2, Proposition 2.15]). The order for which $\phi_*\iota$ is optimal may not be comparable to $\mathcal{O}$ if $\deg(\phi)$ has more than one distinct prime factor (for example, ϕ may factor into coprime descending and ascending isogenies).

Definition 3.2.5. We say (E, ι) and (E', ι') are **K-isomorphic** if there is an isomorphism $\psi : E \to E'$ such that $\iota' = \psi_*\iota$. Here, ψ may be defined over any field. In this case, we use the notation $(E, \iota) \cong_K (E, \iota')$. If ι and ι' are optimal embeddings of $\mathcal{O}$ into $\mathrm{End}(E)$, we will say ι and ι' are K-isomorphic if $(E, \iota) \cong_K (E, \iota')$.

For ordinary elliptic curves, an optimal embedding is an isomorphism with the endomorphism ring, and there are exactly two optimal embeddings which are complex conjugates of each other. By contrast, the situation for supersingular elliptic curves is more interesting, owing to the rich structure of the endomorphism ring. A supersingular elliptic curve E admits optimal embeddings of infinitely many quadratic imaginary orders into $\mathrm{End}(E)$. Even for a fixed quadratic imaginary order $\mathcal{O}$, there may be many optimal embeddings into $\mathrm{End}(E)$ which are not K-isomorphic.

3.2.2 Action of ideals on oriented curves

In [37, Theorem 3.4], Onuki shows that the class group $\mathcal{CL}(\mathcal{O})$ acts on the set of $\mathcal{O}$-oriented curves up to K-isomorphism. This action is defined as follows. Suppose $\mathfrak{a} \subset \mathcal{O}$ is an invertible ideal of norm prime to p. Then $E[\mathfrak{a}] = \cap_{\alpha \in \mathfrak{a}} \ker(\iota(\alpha))$ is a subgroup of E of size equal to $N(\mathfrak{a})$. Let $\phi_\mathfrak{a}$ be the isogeny with kernel equal to $E[\mathfrak{a}]$. We define $\mathfrak{a} * E$ to be the codomain of $\phi_\mathfrak{a}$, and we define $\mathfrak{a} * \iota$ to be the embedding $(\phi_\mathfrak{a})_*\iota$.

Onuki also shows that the action is free and has one or two orbits. It was shown in [2, Proposition 4.2, Proposition 4.3] that the action is transitive when p is ramified in K and has two orbits when p is inert in K. We will give a different proof of this fact.

3.3 Infinitely many optimal embeddings

In this section, we show that $\mathrm{End}(E)$ admits optimal embeddings of infinitely many quadratic imaginary orders when E is supersingular. By the Deuring correspondence, this is a statement purely about maximal orders in a quaternion algebra.

For this section, let $B_{p,\infty}$ denote the quaternion algebra ramified at p and ∞.

3.3.1 Restricting and extending to optimal embeddings.

We start with two straightforward propositions which show that embeddings are always optimal for some order.

Proposition 3.3.1. *Let K be a quadratic imaginary field. Let R be a maximal order of $B_{p,\infty}$. Any embedding $\iota : K \to B_{p,\infty}$ restricts to an optimal embedding of some order $\mathcal{O} \subset K$ into R.*

Proof. We take $\mathcal{O} = \iota^{-1}(\iota(K) \cap R)$. □

Proposition 3.3.2. *Any embedding of $\mathcal{O}$ into R extends to an optimal embedding of $\mathcal{O}'$ into R for some $\mathcal{O}' \supset \mathcal{O}$.*

Proof. Let $\iota : \mathcal{O} \to R$ be an embedding. Let ι' denote the extension of ι to K by $\mathbb{Q}$-scalars. By Proposition 3.3.1, ι' restricts to an optimal embedding of some order $\mathcal{O}' \subset K$. In particular, $\mathcal{O}' = \iota'^{-1}(\iota'(K) \cap R)$. Since $\iota'(\mathcal{O}) \subset R$, it is clear that $\mathcal{O} \subset \mathcal{O}'$. □

3.3.2 Quadratic imaginary fields embedding into quaternion algebras

We give a splitting condition for quadratic imaginary fields embedding into $B_{p,\infty}$.

Proposition 3.3.3. *Let K be a quadratic imaginary field. There is an embedding of K into $B_{p,\infty}$ if and only if p does not split in K.*

Proof. As $B_{p,\infty}$ ramifies at p, $B_{p,\infty} \otimes \mathbb{Q}_p$ is a division ring and therefore has no zero divisors. But if K is a number field in which p splits, say $p\mathcal{O}_K = \mathfrak{p}_1\mathfrak{p}_2$, then $K \otimes \mathbb{Q}_p \cong K_{\mathfrak{p}_1} \oplus K_{\mathfrak{p}_2}$, which contains zero divisors. Therefore, $B_{p,\infty}$ does not contain any field in which p splits.

If K is a number field in which p does not split, there is an elliptic curve E defined over the Hilbert class field of K such that E has CM by the maximal order $\mathcal{O}_K$ [47, Theorem 4.1 and Remark 4.2]. As $j(E)$ is algebraic, E has potential good reduction [47, Theorem 6.1], so there is a finite extension of H over which E is defined and has good reduction at a prime over p.

By [30, Theorem 12 in Chapter 13], since p does not split in K and E has good reduction at a prime over p, the reduction $\widetilde{E}$ is a supersingular elliptic curve over $\overline{\mathbb{F}}_p$. Reduction induces an embedding $\iota : \mathrm{End}(E) \to \mathrm{End}(\widetilde{E})$ [47, Proposition 4.4]. Composing with isomorphisms gives an embedding of $\mathcal{O}_K$ into a maximal order of $B_{p,\infty}$. Extending scalars to $\mathbb{Q}$ gives the desired embedding into $B_{p,\infty}$. □

Remark 3.3.4. The result of Proposition 3.3.3 is part of [54, Proposition 14.6.7]; our proof uses elliptic curves and theory of lifting and reduction.

Proposition 3.3.5. *Fix a supersingular elliptic curve $E/\mathbb{F}_{p^2}$. If K is a quadratic imaginary field such that p does not split in K, then there is an order $\mathcal{O} \subset K$ such that $\mathcal{O}$ embeds optimally into End(E).*

Proof. Fix E and K as in the statement. Let $R \subset B_{p,\infty}$ be a maximal order such that $\mathrm{End}(E) \cong R$.

By Proposition 3.3.3, there is a supersingular elliptic curve E_K which admits an optimal embedding of $\mathcal{O}_K \to \mathrm{End}(E_K)$.

Let $R_K \subset B_{p,\infty}$ be a maximal order such that $\mathrm{End}(E_K) \cong R_K$. Since R and R_K are full-rank $\mathbb{Z}$-modules, $R \cap R_K$ is a full-rank submodule of R and thus the index $[R : R \cap R_K]$ is finite. Let $n = [R : R \cap R_K]$. Then $nR_K \subset R$. As there is an embedding of $\mathcal{O}_K$ into R_K, this shows that there is an embedding of $n\mathcal{O}_K$ into R. Since $\mathbb{Z} \subset R$, we have an embedding $\mathbb{Z} \oplus n\mathcal{O}_K \hookrightarrow R$. By Proposition 3.3.2, this embedding extends to an optimal embedding of an order containing $\mathbb{Z} \oplus n\mathcal{O}_K$. □

3.3.3 Optimal embeddings into all supersingular elliptic curves

Fix a supersingular elliptic curve E. We have just shown that every quadratic imaginary field K in which p does not split contains an order $\mathcal{O}$ which embeds optimally into $\mathrm{End}(E)$. One can ask if $\mathcal{O}$ can be chosen to optimally embed into the endomorphism ring of every supersingular elliptic curve.

It was observed in [29, Proposition 5.10] that the answer is yes. However, the proof given requires a slight modification for p for which $j = 0$ and $j = 1728$ are supersingular. They use the result that andom nonbacktracking walks in regular, undirected expander graphs converge to the uniform distribution (see [1]) to prove the next proposition. However, when $p \not\equiv 1 \pmod{12}$, the ℓ-isogeny graph is no longer a regular, undirected graph. A small modification is needed to obtain the result for all primes p.

Proposition 3.3.6. *Let $\ell \neq p$ be a prime. There exists an integer N such that for all $k > N$ and for any two supersingular elliptic curves E and E', there exists a cyclic ℓ^N-isogeny $\phi : E \to E'$.*

Proof. Random nonbacktracking walks in the ℓ-isogeny graph converge (in total variation) to the stationary distribution by [7, Theorem 11]. The stationary distribution weights the vertex corresponding to the elliptic curve E_i by $\frac{1}{w_i}$, where w_i is half the size of the automorphism group of E_i. For the curve with j-invariant 1728, $w_i = 2$, and for the curve with $j = 0$,

$w_i = 3$; for all other curves, $w_i = 1$. More precisely, the total variation between the stationary distribution and the probability distribution obtained after a random nonbacktracking walk of length k is bounded above by a function in k which decreases and approaches 0. For sufficiently large k, this guarantees that there exists a nonbacktracking walk of length k between any two vertices. □

The rest of the proof proceeds as in [29, Proposition 5.10]:

Proposition 3.3.7. *Let K be a quadratic imaginary field in which p does not split, and fix any quadratic imaginary order $\mathcal{O}_0 \subset K$ with conductor coprime to p. Let ℓ be a prime which is inert in K, coprime to the conductor of $\mathcal{O}_0$, and is not equal to p. Let $n > N$ in Proposition 3.3.6, and let $\mathcal{O} \subset \mathcal{O}_0$ have index $[\mathcal{O}_0 : \mathcal{O}] = \ell^n$. Every supersingular elliptic curve E admits an optimal embedding $\mathcal{O} \to End(E)$.*

Proof. Let (E_0, ι_0) be an $\mathcal{O}_0$-oriented elliptic curve, and let E be any other supersingular elliptic curve. By Proposition 3.3.6, there is a cyclic ℓ^n-isogeny $\phi : E_0 \to E$. By decomposing ϕ as a sequence of n ℓ-isogenies such that no isogeny is followed by its dual, we see that each step must be descending: since ℓ does not divide the conductor of $\mathcal{O}_0$, there are no ascending ℓ-isogenies, and since ℓ is inert in K, there are no horizontal ℓ-isogenies. Every ℓ-isogeny starting at (E_0, ι_0) induces an optimal embedding of the index ℓ subring of $\mathcal{O}_0$ into the endomorphism ring of the image; at each subsequent step, ℓ divides the conductor, so there are no horizontal ℓ-isogenies, and the only ascending ℓ-isogeny is dual to the previous descending isogeny. Hence, ϕ induces an optimal embedding of $\mathcal{O}$, the ℓ^n-index subring of $\mathcal{O}_0$, into the endomorphism ring of the target curve, E. Since E could be taken to be any supersingular elliptic curve, $\mathcal{O}$ embeds optimally into $\mathrm{End}(E)$ for all supersingular elliptic curves E. □

In fact, we can get arbitrarily many optimal embeddings into every supersingular elliptic curve by starting with $\mathcal{O}_0$ with a large class number.

Corollary 3.3.8. *Keeping K, ℓ, $\mathcal{O}_0$, $\mathcal{O}$, and n the same as in Proposition 3.3.7, every supersingular elliptic curve E admits at least $h(\mathcal{O}_0)$ optimal embeddings of $\mathcal{O} \to End(E)$. If p is inert in K, there are at least $2h(\mathcal{O}_0)$ optimal embeddings of $\mathcal{O} \to End(E)$.*

Proof. We repeat the proof of Proposition 3.3.7 for all choices of (E_0, ι_0). □

It would be interesting to see if the splitting condition on ℓ can be dropped. This would show that "most" orders optimally embed into all endomorphism rings with as many optimal embeddings as desired.

3.4 Orbits of the group action

The main result of this section is that the action of $\mathcal{CL}(\mathcal{O})$ on $\mathcal{O}$-oriented curves is transitive when p ramifies in K, and the action has two orbits when p is inert in K. This result appears in [2, Proposition 4.2, Proposition 4.3] with a different proof. In particular, while they also consider K-isomorphisms among conjugate embeddings, we show more precise statements about when conjugate embeddings may be K-isomorphic.

3.4.1 Conjugate embeddings

Fix a quadratic imaginary field K in which p does not split. Fix an order $\mathcal{O} \subset K$ such that p does not divide the conductor. Let $\iota : \mathcal{O} \to \mathrm{End}(E)$ be an optimal embedding.

There are two "conjugate" embeddings related to ι, which we define as follows:

Definition 3.4.1. The **complex conjugate** of ι is the embedding $\bar{\iota} : \mathcal{O} \to \mathrm{End}(E)$ defined by $\bar{\iota}(\alpha) = \iota(\overline{\alpha})$, where the overline denotes complex conjugation.

The **Galois conjugate** of ι is the embedding $\iota^{(p)} : \mathcal{O} \to \mathrm{End}(E^{(p)})$, which is obtained by action of the Frobenius. If $\pi_p : E \to E^{(p)}$ is the Frobenius endomorphism, then $\iota^{(p)}(\alpha) = \frac{1}{p}\pi_p \iota(\alpha)\hat{\pi}_p$.

Later, we will show that when the class group action has two orbits, one orbit will contain (E, ι) and $(E^{(p)}, \bar{\iota}^{(p)})$, and the other will contain $(E, \bar{\iota})$ and $(E^{(p)}, \iota^{(p)})$. To prove this, we will first classify when there are K-isomorphisms among these four embeddings.

First, we consider (E, ι) and $(E, \bar{\iota})$.

Lemma 3.4.2. *Let $p > 3$. Let E be supersingular with an optimal embedding $\iota : \mathcal{O} \to End(E)$. If (E, ι) and $(E, \bar{\iota})$ are K-isomorphic, then $j(E) = 1728$ and p ramifies in K.*

Proof. Let ω be a nonzero element of $\mathcal{O}$ with trace zero. Note that ω may not be a generator of $\mathcal{O}$, but since $K = \mathbb{Q}(\omega)$, the embedding ι is determined by its action on ω. Thus, (E, ι) and $(E, \bar{\iota})$ are K-isomorphic if and only if there is an automorphism $\psi : E \to E$ such that $\psi\iota(\omega) = \bar{\iota}(\omega)\psi = \iota(\overline{\omega})\psi$. Since the trace of ω is zero, $\overline{\omega} = -\omega$, so (E, ι) and $(E, \bar{\iota})$ are K-isomorphic if and only if there is an automorphism ψ such that $\psi\iota(\omega) = -\iota(\omega)\psi$.

Thus, (E, ι) and $(E, \bar{\iota})$ are K-isomorphic if and only if there is an automorphism ψ such that ψ and $\iota(\omega)$ can be embedded as anticommuting elements of $B_{p,\infty}$.

Anticommuting elements of $B_{p,\infty}$ have zero real part and therefore have zero trace. The automorphisms of elliptic curves with zero trace arise only for $j(E) = 1728$, which are the automorphisms of order 4, so we must have $j(E) = 1728$.

To show that p ramifies in K, we use that the endomorphism ring for E can be wrtten explictily, from which we can express K explicitly. For $j(E) = 1728$, E is isomorphic to the curve $y^2 = x^3 + x$. Let i denote the automorphism $(x, y) \mapsto (-x, ay)$ where $a^2 = -1$, so that $i^2 = -1$. Then $\mathrm{End}(E)$ has basis $\{1, i, \frac{1-\pi_p}{2}, \frac{1-i\pi_p}{2}\}$ [28, Lemma 2].

One can check that the set of elements of $\mathrm{End}(E)$ anticommuting with i are $\mathbb{Q}$-linear combinations of π_p and $i\pi_p$. Replacing $\iota(\omega)$ by an integer multiple if necessary, we may write $\iota(\omega) = (a + bi)\pi_p$ for some $a, b \in \mathbb{Z}$. Therefore the order $\mathbb{Z}[\omega]$ has discriminant $4(a^2 + b^2)p$. Since E with $j(E) = 1728$ is supersingular if and only if $p \equiv 3 \pmod 4$, $a^2 + b^2$ can contribute only even powers of p. In particular, this implies $4(a^2 + b^2)p$ is exactly divisible by an odd power of p, and therefore p divides the fundamental discriminant. Thus, p ramifies in $K = \mathbb{Q}(\omega)$. □

Now, we consider when (E, ι) and $(E, \iota^{(p)})$ are K-isomorphic. We require the following lemma.

Lemma 3.4.3. *Suppose $\phi : E \to E^{(p)}$ has degree n. If ϕ satisfies $\frac{1}{n}\phi\iota(\alpha)\hat{\phi} = \iota^{(p)}(\alpha)$ for all $\alpha \in \mathcal{O}$, then $\pi_p\phi \in \iota(\mathcal{O})$*

Proof. The condition $\frac{1}{n}\phi\iota(\alpha)\hat{\phi} = \iota^{(p)}(\alpha)$ is equivalent to the condition that

$$\phi\iota(\alpha) = \iota^{(p)}(\alpha)\phi,$$

by composing on the right with ϕ.

We will show that $\pi_p\phi$ commutes with $\iota(\alpha)$. By hypothesis, $\pi_p\phi\iota(\alpha) = \pi_p\iota^{(p)}(\alpha)\phi$. By definition of the action of Frobenius, $\pi_p\iota^{(p)} = \iota\pi_p$, so $\pi_p\iota^{(p)}(\alpha)\phi = \iota(\alpha)\pi_p\phi$. Hence $\pi_p\phi$ commutes with all elements of $\iota(\alpha)$, and therefore $\pi_p\phi \in \iota(\mathcal{O})$. □

Lemma 3.4.4. *Let E be defined over $\mathbb{F}_p$ and $p > 3$. If (E, ι) and $(E, \iota^{(p)})$ are K-isomorphic, then $K \cong \mathbb{Q}(\sqrt{-p})$ and $\mathcal{O} \supset \mathbb{Z}[\sqrt{-p}]$.*

Proof. If (E, ι) and $(E, \iota^{(p)})$ are K-isomorphic, there is an automorphism $\psi : E \to E'$ such that $\psi\iota\psi^{-1} = \iota^{(p)}$. By Lemma 3.4.3, this implies that $\pi_p\psi \in \iota(\mathcal{O})$. Since ψ is an automorphism, this implies $\mathcal{O}$ contains an element of norm p. Thus, there is an ideal of norm p in $\mathcal{O}$ (and also in $\mathcal{O}_K$), so either p splits in K or p ramifies in K. By Proposition 3.3.3, p does not split, so p must ramify in K and divide the discriminant of $\mathcal{O}_K$.

Let $\alpha \in \mathcal{O}$ such that $\iota(\alpha) = \pi_p\psi$. Then α and $\overline{\alpha}$ generate the same ideal of norm p in $\mathcal{O}_K$, so α and $\overline{\alpha}$ must differ multiplicatively by a unit of $\mathcal{O}_K$. Since $p > 3$ divides the discriminant of K, the only units of $\mathcal{O}_K$ are ± 1, so $\overline{\alpha} = -\alpha$. Equivalently, α is an element of trace zero and norm p. So, $\mathcal{O}$ contains $\pm\sqrt{p}$. □

3.4.2 Orbits of the class group

Now, we will show that the orbit of the class group action depends on the splitting behavior of p in K.

Proposition 3.4.5. *Let K be a quadratic imaginary field in which p does not split, and fix an order $\mathcal{O} \subset K$. Let $S_{\mathcal{O}}$ be the set of $\mathcal{O}$-oriented supersingular elliptic curves considered up to K-isomorphism. Then $\mathcal{CL}(\mathcal{O})$ acts on $S_{\mathcal{O}}$. If p ramifies in K, the action is free and transitive. If p is inert in K, the action partitions $S_{\mathcal{O}}$ into two orbits, which are related via the bijections $(E, \iota) \mapsto (E, \bar{\iota})$ or $(E, \iota) \mapsto (E, \bar{\iota}^{(p)})$.*

The proof is not much more than highlighting details in the proof of [37, Proposition 3.3].

Proof. Let $\mathcal{ELL}_{\mathcal{O}}$ be the set of roots of the Hilbert class polynomial of $\mathcal{O}$; these are precisely the j-invariants of elliptic curves over $\mathbb{C}$ with CM by $\mathcal{O}$.

There are a number field L and a prime $\mathfrak{p}$ over p such that for every $j \in \mathcal{ELL}_{\mathcal{O}}$, there is an elliptic curve E_j defined over L with j-invariant j and with good reduction at $\mathfrak{p}$. By CM theory, the j-invariants generate the Hilbert Class Field, H [47, Theorem 4.1], so we may assume $H \subset L$. Onuki shows that $\mathrm{Cl}(\mathcal{O})$ acts freely and transitively on the set $\rho(\{E_j : j \in L_{\mathcal{O}}\})$, where $\rho(E_j) = (\tilde{E}_j, \theta) \in S_{\mathcal{O}}$: $\tilde{E}_j$ is the reduction modulo $\mathfrak{p}$ and θ is the embedding induced by the reduction of the normalized embedding $[\cdot]_{E_j} : \mathcal{O}$ into $\mathrm{End}(E_j)$.

We will show that the image of ρ depends on the splitting behavior of p.

By the Deuring Lifting Theorem (see [30, Chapter 13, Section 5, Theorem 14]), any elliptic curve E in characteristic p and endomorphism $\alpha \in \mathrm{End}(E)$ can be lifted to an elliptic curve defined over a number field and an endomorphism, whose reductions modulo a prime over p are E and α respectively. In particular, for every $(E, \iota) \in S_{\mathcal{O}}$, we can lift E and the image of a generator of $\mathcal{O}$ under ι to an elliptic curve E' defined over a number field L' and with good reduction at a prime $\mathfrak{p}'$, such that the reduction $\tilde{E}' \cong E$ and $\mathrm{End}(E')$ is isomorphic to $\mathcal{O}$. By taking L' larger if necessary, we may assume L' is Galois and contains K. Thus, $\rho(E')$ is one of (E, ι) or $(E, \bar{\iota})$. For ease of notation, we relabel the image by (E, ι).

From (E, ι), we will determine conjugate elements of $S_{\mathcal{O}}$, which are elements of the form $\rho(E'^{\sigma})$ for $\sigma \in \mathrm{Gal}(L'/\mathbb{Q})$ which fix the prime $\mathfrak{p}'$. Any such curves must be conjugate to E as they are related by action of $\sigma \pmod{\mathfrak{p}'} \in \mathrm{Gal}(\mathbb{F}_{p^f}/\mathbb{F}_p)$. The induced embeddings of $\mathcal{O}$ are obtained by reducing the normalized embedding $[\cdot]_{E'^{\sigma}}$ modulo $\mathfrak{p}$. As $[\cdot]_{E'^{\sigma}} = ([\sigma(\cdot)]_{E'})^{\sigma}$ (see [47, Theorem 2.2]), the resulting embedding is determined by both the action of σ restricted to K (which is the identity or complex conjugation, as K is quadratic imaginary), as well as the action of σ modulo $\mathfrak{p}'$ (a power of Frobenius). We get four possibilities for $\rho(E'^{\sigma})$, corresponding to the conjugate embeddings: $(E, \iota), (E, \bar{\iota}), (E^{(p)}, \iota^{(p)}), (E^{(p)}, \bar{\iota}^{(p)})$.

Let $G = \mathrm{Gal}(L'/\mathbb{Q})$, and let $H = \mathrm{Gal}(L'/K)$. Define the decomposition group

$$D := \{\sigma \in G : \sigma\mathfrak{p}' = \mathfrak{p}'\}$$

and the inertia group

$$I := \{\sigma \in G : \sigma\alpha \equiv \alpha \mathrm{mod}\mathfrak{p}' \text{ for all } \alpha \in L'\}.$$

We want to first consider the action of D on E'.

Since p does not split in K, K is not contained in the fixed field of D, L_D. The fixed field of $D \cap H$ is KL_D, so $D \cap H$ is an index 2 subgroup of D.

Case 1: p ramifies in K.

When p ramifies in K, we have $D/I \cong (D \cap H)/(I \cap H) \cong \mathrm{Gal}(\mathbb{F}_p^f/\mathbb{F}_p)$ (where f is the inertial degree). Reaching E or $E^(p)$ depends on if σ modulo $\mathfrak{p}'$ acts as an even or odd power of Frobenius. The isomorphism of $(D \cap H)/(I \cap H)$ with D/I implies that each such choice can be lifted to elements $\sigma \in D$ either belonging to $D \cap H$, which acts trivially on K, or to $\sigma \in D \setminus H$, which acts as conjugation on K. As all choices can be made in this case, each of the four possibilities listed are realized as the reduction of $(E')^\sigma$ for some σ.

Since ρ is reduction mod $\mathfrak{p}$, we apply an automorphism moving $\mathfrak{p}'$ to $\mathfrak{p}$. Let M be Galois so that $LL' \subset M$, and let P' be a prime of M over $\mathfrak{p}'$ and P a prime of M over $\mathfrak{p}$. Since there is only one prime over p in K, we can choose $\tau \in \mathrm{Gal}(M/K)$ with $\tau(P') = P$.

Then $(E'^\sigma)^\tau$ has good reduction at P. If j is the j-invariant of (E'^σ) and $g \in K[x]$ reduces to the minimal polynomial of j over $\mathbb{F}_p$, then $g(j) \in P$, so $g(\tau(j)) = \tau(g(j)) \in \tau(P)$, hence $(E'^\sigma)^\tau$ reduces mod $\mathfrak{p}$ to one of E or $E^{(p)}$. Furthermore, $[\cdot]_{(E')^{\tau\sigma}} = [\cdot]^\tau_{(E')^\sigma}$, since τ fixes K. If p ramifies in K, this ranges over all conjugate embeddings as σ ranges through elements of D, and therefore we must be in the transitive case.

Case 2: p is inert in K.

When p is inert in K, we have $(D \cap H)/(I \cap H)$ is a subgroup of D/I of index 2. To reach E or $E^{(p)}$ depends on if σ modulo $\mathfrak{p}'$ acts as the p-power Frobenius, all of whose lifts are necessarily nontrivial on K, or the p^2-power Frobenius, all of whose lifts are necessarily trivial on K. Thus, the only pairs that can be reached via an automorphism in D followed by reduction mod $\mathfrak{p}'$ are (E, ι) and $(E^{(p)}, \bar{\iota}^p)$.

If $j(E) \notin \mathbb{F}_p$, then (E, ι) and $(E^{(p)}, \bar{\iota}^{(p)})$ are not K-isomorphic (as, in particular, E and $E^{(p)}$ are not isomorphic). By Lemma 3.4.2, (E, ι) is also not K-isomorphic to $(E, \bar{\iota})$. Hence, there are four distinct K-isomorphism classes among these four possibilities, and two of them can be reached via action of D.

If $j(E) \in \mathbb{F}_p$, then E and $E^{(p)}$ are isomorphic. But by Lemma 3.4.4, as p is inert in K (and in particular $K \ncong \mathbb{Q}(\sqrt{-p})$), we have that (E, ι) and $(E, \iota^{(p)})$ are not K-isomorphic. Either there are 4 distinct K-isomorphism classes among these four possibilities, or if there is an identification, then $(E, \iota) \cong (E^{(p)}, \bar{\iota}^{(p)})$, in which case there are only two possibilities (E, ι) and $(E, \bar{\iota})$, and action by D only reaches one of them. In any case, this gives half of the possibilities from action by D: $\{(E, \iota), (E^{(p)}, \bar{\iota}^{(p)})\}$.

Since ρ is reduction mod $\mathfrak{p}$, we apply an automorphism moving $\mathfrak{p}'$ to $\mathfrak{p}$. As in the ramified case, we obtain an elliptic curve $(E'^{\sigma})^{\tau}$ which reduces to E or $E^{(p)}$. Since p is inert, the two possibilities we obtain are $\{(E, \iota), (E^{(p)}, \bar{\iota}^{(p)})\}$ or $\{(E^{(p)}, \iota^{(p)}), (E, \bar{\iota})\}$, depending on the action of τ mod P.

When p is inert in K, we have shown that half of the conjugate embeddings appear in the image of ρ. We need to rule out the possibility that the rest of the conjugate embeddings could appear in the image of ρ.

So, assume p is inert in K. Suppose (E, ι) and $(E^{(p)}, \iota^{(p)})$ appear in the same orbit, so they are given by reductions of elliptic curves E', E'' over a number field L with good reduction at $\mathfrak{p}$ and with CM by $\mathcal{O}$. Using the class group action on curves with CM by $\mathcal{O}$, there is an ideal $\mathfrak{a} \subset \mathcal{O}$ such that $\mathfrak{a} * E' = E''$. We can choose $\mathfrak{a}$ to be coprime to the conductor of $\mathcal{O}$ and p.

Correspondingly, there is an isogeny $\phi : E' \to E''$ of degree $N = N(\mathfrak{a})$. Note that ϕ is necessarily separable, as N has degree prime to p. Normalized embeddings satisfy the commuting property $\phi[\cdot]_{E'} = [\cdot]_{E''}\phi$ [47, Chapter II, Cor 1.1.1], so by reducing mod $\mathfrak{p}$, we get $\phi\iota = \iota^{(p)}\phi$.

This means in particular that $\pi_p\phi$ commutes with the image of ι, so that $\pi_p\phi = \iota(\alpha)$ for some $\alpha \in \mathcal{O}$ by Lemma 3.4.3. The endomorphism $\pi_p\phi$ has degree $N \cdot p$. Since ϕ is separable, the reduction of ϕ mod $\mathfrak{p}$ has the same degree, $N = N(\mathfrak{a})$. So α has reduced norm Np. Thus, the ideal $\alpha\mathcal{O}$ is divisible by a prime of K over p. The only prime of K over p is $p\mathcal{O}$, which has norm p^2. Hence if $p \mid N(\alpha)$, necessarily $p^2 \mid N(\alpha)$. This contradicts that $N(\alpha) = N \cdot p$ such that $p \nmid N$.

Hence, there is no such $\mathfrak{a}$ such that $\mathfrak{a} * (E, \iota) = (E^{(p)}, \iota^{(p)})$. □

Chapter 4
Horizontal isogenies between oriented elliptic curves

In this chapter, we fix a prime p, a quadratic imaginary field K in which p does not split, and an order $\mathcal{O} \subset K$ of discriminant D and conductor coprime to p. In this chapter, we describe the structure of isogenies of fixed degree between $\mathcal{O}$-oriented supersingular elliptic curves.

We are motivated by the setting of OSIDH and will usually assume that $\mathcal{O}$ is chosen so that any supersingular elliptic curve E has at most one optimal embedding of $\mathcal{O}$ into $\mathrm{End}(E)$ up to conjugation, as in the setting of OSIDH. We give conditions under which all or most N-isogenies between $\mathcal{O}$-oriented curves are horizontal, and we classify the exceptions. As an application, we extend an attack on OSIDH by Dartois and De Feo.

4.1 Orders in a quadratic imaginary order

4.1.1 Binary quadratic forms

A **binary quadratic form** is a function of the form $f(x, y) = ax^2 + bxy + cy^2$, where $a, b, c \in \mathbb{Z}$ and a and c are nonzero. A form is **primitive** if a, b, and c are relatively prime. The **discriminant** of such a form is $b^2 - 4ac$. A form is called **positive definite** if the discriminant is negative. Positive definite forms are precisely those such that $f(x, y)$ takes only positive values when $(x, y) \neq (0, 0)$.

For this dissertation, all quadratic forms considered will be positive definite, primitive, binary quadratic forms.

Two forms f and g are **properly equivalent** if there are integers p, q, r, and s such that $f(x, y) = g(px + qy, rx + sy)$ and $ps - qr = 1$. One can consider g to be obtained from an action of a matrix $\begin{pmatrix} p & q \\ r & s \end{pmatrix} \in SL_2(\mathbb{Z})$ on f. It is a calculation to see that proper equivalence

fixes the discriminant. Proper equivalence defines an equivalence relation on the set of binary quadratic forms of some fixed discriminant.

Let $C(D)$ be the set of primitive, positive definite, binary quadratic forms of discriminant D considered up to proper equivalence. Gauss defined a group law on $C(D)$ (see [17, Theorem 3.9] for details), under which $C(D)$ is isomorphic to the ideal class group $\mathcal{CL}(\mathcal{O})$, where $\mathcal{O}$ is a quadratic imaginary order of discriminant D.

Theorem 4.1.1. *[17, Theorem 7.7, Exercise 7.17] Let $\mathcal{O}$ be the order of discriminant D. Then the map*

$$ax^2 + bxy + cy^2 \mapsto \mathbb{Z}\ a + \mathbb{Z}\ \frac{-b + \sqrt{D}}{2}$$

induces an isomorphism betwen $C(D)$ and $\mathcal{CL}(\mathcal{O})$. If $\mathfrak{a} = \mathbb{Z}\alpha + \mathbb{Z}\beta$ is an invertible $\mathcal{O}$-ideal, then the map

$$\mathfrak{a} \mapsto f(x, y) = \frac{N(x\alpha + y\beta)}{N(\mathfrak{a})}$$

is the inverse map.

4.2 Classifying N-Isogenies Using Positive Definite Quadratic Forms

Our main tool is a correspondence between N-isogenies and a solution to an equation involving quadratic forms of discriminant D.

Theorem 1.2.1. *Let (E, ι) be an $\mathcal{O}$-oriented supersingular elliptic curve over $\overline{\mathbb{F}}_p$. Let $\mathfrak{b}$ be an invertible ideal in $\mathcal{O}$ with norm coprime to p, and let E' such that $\mathfrak{b} * (E, \iota) = (E', \iota')$. Let N be a positive integer. The following hold:*

1. *If p is inert in K, there is an injective map*

$$\{N\text{-isogenies } \phi : E \to E'\} \to \{(w, x, y, z) \in \mathbb{Z}^4 : Q(w, x) + pQ'(y, z) = -DN\}.$$

2. *If p ramifies in K, there is an injective map*

$$\{N\text{-isogenies } \phi : E \to E'\} \to \{(w, x, y, z) : Q(w, x) + Q'(y, z) = -DN/p\}.$$

Here, Q depends only on $\mathfrak{b}$, and Q' depends on both E and $\mathfrak{b}$.

To prove this theorem, we give an explicit description of $\mathrm{Hom}(E', E)$ and use the correspondence between invertible ideals and quadratic forms.

Lemma 4.2.1. *With notation as in Theorem 1.2.1,* $\mathrm{Hom}(E', E) \cong \mathit{End}(E)\iota(\mathfrak{b})$*, where* $\cong$ *is an isomorphism of End(E)-modules. Under the isomorphism, an isogeny of degree N corresponds to an endomorphism of degree* $N * N^K(\mathfrak{b})$.

Proof. From the action of $\mathcal{CL}(\mathcal{O})$ on $\mathcal{O}$-oriented curves, there is an isogeny $\phi_{\mathfrak{b}} : E \to \mathfrak{b} * E$ with $\deg(\phi_{\mathfrak{b}}) = N^K(\mathfrak{b})$ and $\ker(\phi_{\mathfrak{b}}) = \cap_{\alpha \in \mathfrak{b}} \ker(\theta(\alpha))$. We consider the $\mathrm{End}(E)$-module homomorphism

$$\mathrm{Hom}(\mathfrak{b} * E, E) \to \mathrm{End}(E)\iota(\mathfrak{b})$$

$$\phi \mapsto \phi\phi_{\mathfrak{b}}.$$

It is clear that the degree is multiplied by $\deg(\phi_{\mathfrak{b}}) = N^K(\mathfrak{b})$. It remains to show that the image is $\mathrm{End}(E)\iota(\mathfrak{b})$.

By [56, Theorem 3.15], the left ideal $I := \mathrm{End}(E)\iota(\mathfrak{b})$ is a kernel ideal in the maximal order $\mathrm{End}(E)$. In other words, $I = \{\alpha \in \mathrm{End}(E) : \alpha H(I) = 0\}$, where $H(I) = \cap_{\alpha \in I} \ker(\alpha)$. Since $\cap_{\alpha \in I} \ker(\alpha) = \cap_{\alpha \in \mathfrak{b}} \ker(\iota(\alpha)) = \ker(\phi_{\mathfrak{b}})$, we have

$$\mathrm{End}(E)\iota(\mathfrak{b}) = \{\alpha \in \mathrm{End}(E) : \alpha(\ker(\phi_{\mathfrak{b}})) = 0\}.$$

This shows that the image of the map $\phi \mapsto \phi\phi_{\mathfrak{b}}$ is contained in $\mathrm{End}(E)\iota(\mathfrak{b})$. Moreover, for any $\alpha \in \mathrm{End}(E)\iota(\mathfrak{b})$, we know that $\alpha(\ker(\phi_{\mathfrak{b}})) = 0$, and by [48, Corollary 4.11], this implies $\alpha = \phi\phi_{\mathfrak{b}}$ for some $\phi \in \mathrm{Hom}(E', E)$, hence α is in the image of the map. Thus, this mapping is an $\mathrm{End}(E)$-module isomorphism. □

We combine this isomorphism with the description of maximal orders with optimal embeddings by a quadratic imaginary order given by Lauter and Viray in [31]. First, we set notation which we will use for the remainder of this chapter. More details are given in [31, Section 5, Section 6.1].

Fix a prime p, a quadratic imaginary order $\mathcal{O}$ of discriminant D, and a prime q satisfying certain congruence conditions mod $|D|$. In particular, for an integer k depending on D, there is a map $\Psi : (\mathbb{Z}/|D|\mathbb{Z})^\times \to \{\pm 1\}^k$ such that a positive integer $m \in \mathrm{Ker}(\Psi)$ if and only if there exists an ideal $\mathfrak{a} \in \mathcal{CL}(\mathcal{O})^2$ with $N(\mathfrak{a}) = m$ [31, Theorem 5.1]. If p is inert in $\mathcal{O}$, q is chosen so that $-pq \in \mathrm{Ker}(\Psi)$. If p ramifies in $\mathcal{O}$, q is chosen so that $-q \in \mathrm{Ker}(\Psi)$. Finally, a prime $\mathfrak{q}$ lying over q is chosen.

With this notation, for any fixed embedding $\iota : K \to B_{p,\infty}$, $B_{p,\infty}$ can be written as $B_{p,\infty} = \iota(K) \oplus \iota(K)j$ for an element $j \in B_{p,\infty}$, and elements of $B_{p,\infty}$ can be expressed as

matrices in the following way:

$$B_{p,\infty} = \{[\alpha, \beta] \colon = \begin{pmatrix} \alpha & \beta \\ j^2\overline{\beta} & \overline{\alpha} \end{pmatrix}\}, \alpha, \beta \in K$$

where $\iota(\alpha) = [\alpha, 0]$. Here, $j \in B_{p,\infty}$ is chosen to satisfy $j^2 = -pq$ if p is inert in $\mathcal{O}$ and $j^2 = -q$ if p ramifies in $\mathcal{O}$.

The maximal orders of $B_{p,\infty}$ into which $\mathcal{O}$ embeds optimally are described explicitly by the following proposition.

Proposition 4.2.2. *[31, Section 6.2, Section 6.3] For each invertible ideal $\mathfrak{a} \subset \mathcal{O}$ coprime to the conductor, there exists $\lambda_\mathfrak{a} \in \mathcal{O}$ and a ring $R(\mathfrak{a}, \lambda_\mathfrak{a}) \subset B_{p,\infty}$ satisfying the following properties:*

1. $\lambda_\mathfrak{a}\mathfrak{q}^{-1}\overline{\mathfrak{a}}\mathfrak{a}^{-1} \subset \mathcal{O}$.
2. $N(\lambda_\mathfrak{a}) \equiv -pq \pmod{D}$ *if p is inert in $\mathcal{O}$.*
3. $N(\lambda_\mathfrak{a}) \equiv -q \pmod{D/p}$ *if p is ramified in $\mathcal{O}$.*
4. $R(\mathfrak{a}, \lambda_\mathfrak{a})$ *is a maximal order.*
5. $R(\mathfrak{a}, \lambda_\mathfrak{a}) \cap K = \mathcal{O}$.
6. *The optimal embedding $\iota : \mathcal{O} \hookrightarrow R(\mathfrak{a}, \lambda_\mathfrak{a})$ is isomorphic to the embedding $End(E(\mathfrak{a})) \hookrightarrow \widetilde{End(E(\mathfrak{a}))}$, where $E(\mathfrak{a})$ is the elliptic curve with CM by $\mathcal{O}$ corresponding to $\mathfrak{a}$.*
7. *For any invertible ideal $\mathfrak{b} \subset \mathcal{O}$ coprime to the conductor, $R(\mathfrak{a}, \lambda_\mathfrak{a})\mathfrak{b} = \mathfrak{b}R(\mathfrak{a}\mathfrak{b}, \lambda_{\mathfrak{a}\mathfrak{b}})$.*

When p is inert in K,

$$R(\mathfrak{a}, \lambda_\mathfrak{a}) : \{[\alpha, \beta] : \alpha \in \sqrt{D}^{-1}\mathcal{O}, \beta \in \mathfrak{q}^{-1}\sqrt{D}^{-1}\overline{\mathfrak{a}}\mathfrak{a}^{-1}, \alpha - \lambda_\mathfrak{a}\beta \in \mathcal{O}\}.$$

When p ramifies in K,

$$R(\mathfrak{a}, \lambda_\mathfrak{a}) : \{[\alpha, \beta] : \alpha \in \sqrt{D}^{-1}\mathfrak{p}\mathcal{O}, \beta \in \mathfrak{q}^{-1}\sqrt{D}^{-1}\overline{\mathfrak{a}}\mathfrak{a}^{-1}\mathfrak{p}, \alpha - \lambda_\mathfrak{a}\beta \in \mathcal{O}\}.$$

Lauter and Viray construct $\lambda_\mathfrak{a}$ explicitly for each ideal so that the above properties hold, but we omit details of the construction here beyond the properties which are summarized above. We will usually refer to $R(\mathfrak{a}, \lambda_\mathfrak{a})$ by $R(\mathfrak{a})$.

Corollary 4.2.3. *Let E and E' be $\mathcal{O}$-oriented curves with $D = \mathrm{disc}(\mathcal{O})$. For integral ideals $\mathfrak{a}$ and $\mathfrak{b}$ depending on E and E', $\mathrm{Hom}(E', E)$ is isomorphic to one of the following ideals depending on the splitting behavior of p in K:*

If p is inert in K, then

$$\mathrm{Hom}(E', E) \cong \{[\alpha, \beta] : \alpha \in (\sqrt{D}\mathcal{O})^{-1}\mathfrak{b}, \beta \in (\sqrt{D}\mathcal{O})^{-1}\mathfrak{q}^{-1}\bar{\mathfrak{a}}\bar{\mathfrak{b}}\mathfrak{a}^{-1}, \alpha - \lambda_{\mathfrak{a}\mathfrak{b}}\beta \in \mathfrak{b}\}.$$

If p is ramified in K, then

$$\mathrm{Hom}(E', E) \cong \{[\alpha, \beta] : \alpha \in (\sqrt{D}\mathcal{O})^{-1}\mathfrak{b}\mathfrak{p}, \beta \in (\sqrt{D}\mathcal{O})^{-1}\mathfrak{q}^{-1}\bar{\mathfrak{a}}\bar{\mathfrak{b}}\mathfrak{a}^{-1}\mathfrak{p}, \alpha - \lambda_{\mathfrak{a}\mathfrak{b}}\beta \in \mathfrak{b}\}$$

Proof. By Lemma 4.2.1, $\mathrm{Hom}(E', E) \cong \mathrm{End}(E)\iota(\mathfrak{b})$. Under the isomorphism $\mathrm{End}(E) \to R(\mathfrak{a})$, we have $\mathrm{End}(E)\iota(\mathfrak{b}) \cong R(\mathfrak{a})\mathfrak{b}$. By [31, Lemma 6.5, Lemma 6.9], $R(\mathfrak{a})\mathfrak{b} = \mathfrak{b}R(\mathfrak{a}\mathfrak{b})$. We will show that $\mathfrak{b}R(\mathfrak{a}\mathfrak{b})$ can be written as described in the corollary. There are two cases depending on the splitting behavior of p. We will show the inert case and note that the ramified case is similar.

First, we show that the set

$$H := \{[\alpha, \beta] : \alpha \in (\sqrt{D}\mathcal{O})^{-1}\mathfrak{b}, \beta \in (\sqrt{D}\mathcal{O})^{-1}\mathfrak{q}^{-1}\bar{\mathfrak{a}}\bar{\mathfrak{b}}\mathfrak{a}^{-1}, \alpha - \lambda_{\mathfrak{a}\mathfrak{b}}\beta \in \mathfrak{b}\}$$

is a right $R(\mathfrak{a}\mathfrak{b})$-ideal.

Let $[\alpha', \beta'] \in H$ and let $[\alpha, \beta] \in R(\mathfrak{a}\mathfrak{b})$. We need to show that $[\alpha', \beta'][\alpha, \beta] \in H$. For ease of notation, let $\lambda = \lambda_{\mathfrak{a}\mathfrak{b}}$, $a = \sqrt{D}\alpha$, $a' = \sqrt{D}\alpha'$, $b = \sqrt{D}\beta$, and $b' = \sqrt{D}\beta'$. We need to show:

(1) $a'a + pqb\overline{b'} \in \sqrt{D}\mathfrak{b}$,

(2) $a'b + \overline{a}b' \in \sqrt{D}\mathfrak{q}^{-1}\bar{\mathfrak{a}}\bar{\mathfrak{b}}\mathfrak{a}^{-1}$, and

(3) $a'a + pqb'b - \lambda(a'b + \overline{a}b') \in D\mathfrak{b}$.

The proofs of these three statements are completely analagous to the proof of [31, Lemma 6.3].

To prove (1), rewrite $a'a + pqb'b$ as

$$(a' - \lambda b')a + \lambda b'(a - \lambda b) + \lambda b'(\lambda b - \overline{\lambda b}) + (N(\lambda) + pq)b'\overline{b}.$$

Using Proposition 4.2.2 properties (1) and (2), the definition of H, and the fact that for any $c \in \mathcal{O}$, $c - \overline{c} \in \sqrt{D}\mathcal{O}$, it is clear that each term is in $\sqrt{D}\mathfrak{b}$.

To prove (2), rewrite $a'b + \overline{a}b'$ as $(a' - \lambda b')b - (\overline{a} - \overline{\lambda b})b' + b(\lambda b' - \overline{\lambda b'})$.

To prove (3), let c and c' such that $a = \lambda b + \sqrt{D}c$ and $a' = \lambda b' + \sqrt{D}c'$. Note that

$c \in \sqrt{D}\mathcal{O}$ and $c' \in \sqrt{D}\mathfrak{b}$. We can rewrite $a'a + pqb'b - \lambda(a'b + \overline{a}b')$ as

$$Dcc' + (N(\lambda) + pq)b'\overline{b} + \lambda\sqrt{D}b'(c' - \overline{c'}) \in D\mathcal{O}.$$

This shows that H is a right $R(\mathfrak{ab})$-ideal. Furthermore, the generators of $\mathfrak{b}R(\mathfrak{ab})$ are of the form $[\gamma, 0][\alpha, \beta]$, and multiplying through the corresponding matrices, we obtain the matrix $[\alpha\gamma, \beta\gamma]$ with $\gamma \in \mathfrak{b}$ and $[\alpha, \beta] \in R(\mathfrak{ab})$, and it is clear that $[\alpha\gamma, \beta\gamma] \in H$. Thus, H contains $\mathfrak{b}R(\mathfrak{ab})$.

Finally, to see that $H = \mathfrak{b}R(\mathfrak{ab})$, we show that every element of H has norm divisible by $N(\mathfrak{b})$. Let $[\alpha, \beta] \in H$. Then $N([\alpha, \beta]) = N(\alpha) + pqN(\beta)$. As $\sqrt{D}\alpha \in \mathfrak{b}$ and $\sqrt{D}\beta \in \mathfrak{q}^{-1}\overline{\mathfrak{a}}\mathfrak{a}^{-1}\overline{\mathfrak{b}}$, it follows that $N(\sqrt{D}\alpha)$ and $qN(\sqrt{D}\beta)$ have norm divisible by $N(\mathfrak{b})$. We have $N(\sqrt{D}\alpha) + pqN(\sqrt{D}\beta)$ is divisible by D, since $N(\sqrt{D}\alpha - \lambda\sqrt{D}\beta) \equiv N(\sqrt{D}\alpha) + pqN(\sqrt{D}\beta) \pmod{D}$. Since $\sqrt{D}\alpha - \sqrt{D}\lambda\beta \in \sqrt{D}\mathfrak{b}$ by construction, the norm must be divisible by D, and therefore $N(\alpha) + pqN(\beta)$ is an integer. Since D and $\mathfrak{b}$ are coprime, $N(\alpha) + pqN(\beta)$ is divisible by $N(\mathfrak{b})$. This shows that H has norm divisible by $N(\mathfrak{b})$, and since it contains an ideal of norm $N(\mathfrak{b})$, they must be equal.

This gives the result when p is inert. When p is ramified, an identical argument for the corresponding set H gives the result. □

Now, we prove Theorem 1.2.1 using the explicit description from Corollary 4.2.3 and the correspondence between ideals and binary quadratic forms described in Section 4.1.1.

Proof of Theorem 1.2.1. By Lemma 4.2.1, if ϕ corresponds to $[\alpha, \beta]$ under the described isomorphism, then $\deg(\phi)N^K(\mathfrak{b}) = \det([\alpha, \beta]) = (N^K(\alpha) - j^2N^K(\beta))$. Write $\mathfrak{b} = \mathbb{Z}a + \mathbb{Z}(\frac{b+\sqrt{D}}{2})$, where $a, b \in \mathbb{Z}$.

We have two cases.

Case: p is inert in K

Write $\alpha = \frac{\alpha'}{\sqrt{D}}$ such that $\alpha' \in \mathfrak{b}$. Then all possible values of α' range over $wa + x(\frac{b+\sqrt{D}}{2})$ for all integers $w, x \in \mathbb{Z}$. Letting $Q_{\mathfrak{b}}(w, x) = N^K(wa + x(\frac{b+\sqrt{D}}{2}))/N^K(\mathfrak{b})$ denote the corresponding binary quadratic form, we have $N^K(\alpha) = \frac{N^K(\mathfrak{b})Q_{\mathfrak{b}}(w,x)}{-D}$.

Similarly, write $\beta = \frac{\beta'}{\sqrt{D}}$ such that $\beta' \in \mathfrak{q}^{-1}\bar{\mathfrak{a}}\bar{\mathfrak{b}}\mathfrak{a}^{-1}$. We can rewrite this ideal as $\frac{1}{qN^K(\mathfrak{a})}\bar{\mathfrak{a}}^2\bar{\mathfrak{b}}\bar{\mathfrak{q}}$. Then β' ranges over $\frac{1}{N^K(\mathfrak{a})q}(ya' + z(\frac{b'+\sqrt{D}}{2}))$ such that $\mathbb{Z}\ a' + \mathbb{Z}\frac{b'+\sqrt{D}}{2} = \bar{\mathfrak{a}}^2\bar{\mathfrak{q}}\bar{\mathfrak{b}}$ for $y, z \in \mathbb{Z}$. Letting $Q_{\bar{\mathfrak{a}}^2\bar{\mathfrak{q}}\bar{\mathfrak{b}}}(y, z) = N^K(ya' + z(\frac{b'+\sqrt{D}}{2}))/(qN^K(\mathfrak{a})^2N^K(\mathfrak{b}))$ denote the corresponding binary quadratic form, we have

$$N^K(\beta) = \frac{N^K(\mathfrak{b})Q_{\bar{\mathfrak{a}}^2\bar{\mathfrak{q}}\bar{\mathfrak{b}}}(y, z)}{-Dq}.$$

Hence

$$\deg(\phi) = \frac{Q_{\mathfrak{b}}(w,x) + pQ_{\bar{\mathfrak{a}}^2\bar{\mathfrak{q}}\bar{\mathfrak{b}}}(y,z)}{-D}.$$

Case: p is ramified in K.

In this case, we have $\alpha = \frac{\alpha'}{\sqrt{D}}$ such that $\alpha' \in \mathfrak{bp}$, and hence, replacing $\mathfrak{b}$ with $\mathfrak{bp}$ in the preceding section, we have

$$N^K(\alpha) = \frac{pN^K(\mathfrak{b})Q_{\mathfrak{pb}}(w,x)}{-D}.$$

Similarly, $\beta = \frac{\beta'}{\sqrt{D}}$ such that $\beta' \in \frac{1}{qN^K(\mathfrak{a})}\bar{\mathfrak{a}}^2\bar{\mathfrak{b}}\mathfrak{qp}$. Then

$$N^K(\beta) = \frac{N^K(\mathfrak{b})pQ_{\bar{\mathfrak{a}}^2\bar{\mathfrak{b}}\mathfrak{qp}}(y,z)}{-Dq}.$$

Hence

$$\deg(\phi) = \frac{p(Q_{\mathfrak{pb}}(w,x) + Q_{\bar{\mathfrak{a}}^2\bar{\mathfrak{b}}\mathfrak{qp}}(y,z))}{-D}.\square$$

4.3 Classifying Horizontal and Non-Horizontal Isogenies

Fix an integer N coprime to p. In this section, we use Theorem 1.2.1 to count horizontal N-isogenies, assuming that $\mathcal{O}$ is chosen so that embeddings are unique up to conjugation. If θ and θ' are both optimal embeddings of E into $\mathrm{End}(E)$, then $\theta(\mathcal{O}) = \theta'(\mathcal{O})$. This holds whenever $-D < p$ by [25].

Lemma 4.3.1. *Let (E,θ) be an $\mathcal{O}$-oriented elliptic curve and further assume θ is unique up to conjugation. Let $\mathfrak{b}$ be an integral ideal of $\mathcal{O}$ of norm coprime to D. Let $\phi : \mathfrak{b} * E \to E$ be an isogeny, and let $\alpha, \beta \in K$ so that ϕ corresponds to $[\alpha,\beta] \in End(E)\mathfrak{b}$. Then ϕ is horizontal (with respect to the orientation on $\mathfrak{b} * E$ induced by the action of $\mathfrak{b}$ on θ) if and only if $\beta = 0$ (in which case the induced orientation is θ) or $\alpha = 0$ (in which case the induced orientation is $\bar{\theta}$).*

Proof. The optimal embedding induced by the action of $\mathfrak{b}$ on (E,θ)) is

$$\theta' := \frac{1}{N^K(\mathfrak{b})}\phi_{\mathfrak{b}}\theta\hat{\phi}_{\mathfrak{b}},$$

and therefore the embedding of $K \hookrightarrow \mathrm{End}(E)$ induced by ϕ is

$$\theta" := \frac{1}{\deg(\phi)N^K(\mathfrak{b})}\phi\phi_{\mathfrak{b}}\theta\phi\hat{\phi}_{\mathfrak{b}}.$$

By hypothesis, θ'' is an optimal embedding of $\mathcal{O} \hookrightarrow \mathrm{End}(E)$ if and only if $\theta'' = \theta$ or $\theta'' = \bar{\theta}$.

Write $\phi\phi_{\mathfrak{b}} = [\alpha, \beta] = \alpha + \beta j$ and note that the dual $\phi\hat{\phi}_{\mathfrak{b}} = [\bar{\alpha}, -\beta]$. Then

$$\theta'' = \frac{1}{N^K(\alpha) - j^2 N^K(\beta)} \begin{pmatrix} \alpha & \beta \\ j^2\bar{\beta} & \bar{\alpha} \end{pmatrix} \begin{pmatrix} \theta & 0 \\ 0 & \bar{\theta} \end{pmatrix} \begin{pmatrix} \bar{\alpha} & -\beta \\ -j^2\bar{\beta} & \alpha \end{pmatrix}$$

$$= \frac{1}{N^K(\alpha) - j^2 N^K(\beta)} \begin{pmatrix} N^K(\alpha)\theta - j^2 N^K(\beta)\bar{\theta} & -\alpha\beta(\theta - \bar{\theta}) \\ j^2\bar{\alpha}\bar{\beta}(\theta - \bar{\theta}) & N^K(\alpha)\bar{\theta} - j^2 N^K(\beta)\theta \end{pmatrix}.$$

θ'' is an optimal embedding of $\mathcal{O}$ into $\mathrm{End}(E)$ if and only if

$$\frac{1}{N^K(\alpha) - j^2 N^K(\beta)} \begin{pmatrix} N^K(\alpha)\theta - j^2 N^K(\beta)\bar{\theta} & -\alpha\beta(\theta - \bar{\theta}) \\ j^2\bar{\alpha}\bar{\beta}(\theta - \bar{\theta}) & N^K(\alpha)\bar{\theta} - j^2 N^K(\beta)\theta \end{pmatrix} = \begin{pmatrix} \theta & 0 \\ 0 & \bar{\theta} \end{pmatrix}$$

or

$$\frac{1}{N^K(\alpha) - j^2 N^K(\beta)} \begin{pmatrix} N^K(\alpha)\theta - j^2 N^K(\beta)\bar{\theta} & -\alpha\beta(\theta - \bar{\theta}) \\ j^2\bar{\alpha}\bar{\beta}(\theta - \bar{\theta}) & N^K(\alpha)\bar{\theta} - j^2 N^K(\beta)\theta \end{pmatrix} = \begin{pmatrix} \bar{\theta} & 0 \\ 0 & \theta \end{pmatrix}.$$

As $\theta - \bar{\theta}$ is nonzero on $K \setminus \mathbb{Q}$, we have $\theta'' = \theta$ if and only if $\beta = 0$ and $\theta'' = \bar{\theta}$ if and only if $\alpha = 0$.

□

In the case that p is inert in K and ϕ is separable, the asymmetry in the expression for $\deg(\phi)$ restricts horizontal isogenies further.

Corollary 4.3.2. *In addition to the hypotheses of Lemma 4.3.1, assume p is inert in K and that ϕ is separable. Then ϕ is horizontal if and only if $\beta = 0$, i.e. the induced embedding is θ.*

Proof. Assume the hypotheses of Lemma 4.3.1 and suppose ϕ is a horizontal isogeny. Let $[\alpha, \beta]$ be the corresponding matrix. By Lemma 4.3.1, ϕ is horizontal if and only if $\alpha = 0$ or $\beta = 0$. We will show that if $\alpha = 0$, then ϕ is not separable.

If $\alpha = 0$ and $\phi_*\theta = \bar{\theta}$, the associated tuple (w, x, y, z) has $w = x = 0$, so $\deg(\phi) = \frac{pQ'(y,z)}{-D}$. Since p is inert in K and does not divide the conductor of $\mathcal{O}$, it follows that $p \mid \deg(\phi)$. This implies ϕ is not separable. □

Proposition 4.3.3. *Assume p is inert in k and let $N < \frac{p}{-D}$. If E, E' are curves with $\mathcal{O}$-orientation, then all isogenies $\phi : E \to E'$ with $\deg(\phi) = N$ are horizontal.*

Proof. Every isogeny of degree N can be expressed with a tuple (w, x, y, z) such that

$$\frac{Q_{\mathfrak{b}}(w,x) + pQ_{\bar{\mathfrak{q}}\bar{\mathfrak{a}}^2\bar{\mathfrak{b}}}(y,z)}{-D} = N.$$

Since $Q_{\mathfrak{b}}$ and $Q_{\bar{\mathfrak{q}}\bar{\mathfrak{a}}^2\bar{\mathfrak{b}}}$ are positive definite forms, $N < \frac{p}{-D}$ implies that $Q_{\bar{\mathfrak{q}}\bar{\mathfrak{a}}^2\bar{\mathfrak{b}}}(w,z) = 0$, hence $w = z = 0$ and thus $\beta = 0$. Thus $\phi\phi_{\mathfrak{b}}$ corresponds to $[\alpha, 0]$ and is in the image of K in $\mathrm{End}(E) \otimes \mathbb{Q}$. □

Corollary 4.3.4. *With N, E, and E' as in Proposition 4.3.3, let $\mathfrak{b}$ such that $E = \mathfrak{b} * E'$.*

(1) If there is a non-integer element of $\mathfrak{b}$ of norm $N^K(\mathfrak{b})N$, then optimal embeddings $\theta : \mathcal{O} \hookrightarrow End(E)$ are unique (up to conjugation and K-isomorphism).

(2) If there is no non-integer element of $\mathfrak{b}$ of norm $N^K(\mathfrak{b})N$, then there are no isogenies $\phi : E \to E'$ with $\deg(\phi) = N$.

Proof. (1) If there is a non-integer element α of $\mathfrak{b}$ of norm $N^K(\mathfrak{b})N$, then there is an isogeny $\phi : E \to E'$ (corresponding to $[\alpha, 0]$) of degree N. Suppose θ and θ' are optimal embeddings of $\mathcal{O} \hookrightarrow \mathrm{End}(E)$. Then $\phi\phi_{\mathfrak{b}}$ is in the image of θ and in the image of θ', hence the images of θ and θ' must agree.

(2) In this case, there is no α such that ϕ corresponds to $[\alpha, 0]$. There are therefore no horizontal isogenies of degree N, and thus there are no isogenies of degree N.

□

Proposition 4.3.5. *Suppose $p \mid N$ and $-DN < p^2$. If E is $\mathcal{O}$-oriented and $\phi : \mathfrak{b} * E \to E$ has $\deg(\phi) = N$, then there is an element of $\beta \in \bar{\mathfrak{q}}\bar{\mathfrak{a}}^2\bar{\mathfrak{b}}\frac{1}{qN^K(\mathfrak{a})}$ such that $N^K(\beta) = \frac{N}{pq}N^K(\mathfrak{b})$*

Proof. As $-DN < p^2$, we must have $p \nmid D$, so we are in the case that p is inert in K. Thus ϕ corresponds to a matrix $[\alpha, \beta]$ which corresponds to a tuple (x, y, w, z) such that

$$Q_{\mathfrak{b}}(x,y) + pQ_{\bar{\mathfrak{q}}\bar{\mathfrak{a}}^2\bar{\mathfrak{b}}}(w,z) = -DN.$$

An integer M is represented by $Q_{\mathfrak{b}}(x, y)$ if and only if there is an element of $\mathfrak{b}$ of norm $MN^K(\mathfrak{b})$. As p is inert in $\mathcal{O}$, any element of $\mathcal{O}$ of norm divisible by p must be divisible by an even power of p. As $-DN < p^2$, we must have $Q_{\mathfrak{b}}(x, y) = 0$ (hence $x = y = 0$ and thus $\alpha = 0$) and thus $pQ_{\bar{\mathfrak{q}}\bar{\mathfrak{a}}^2\bar{\mathfrak{b}}}(w, z) = -DN$.

In particular, the corresponding $[\alpha, \beta] = [0, \beta]$ satisfies $pqN^K(\beta) = NN^K(\mathfrak{b})$. □

Corollary 4.3.6. *Suppose $-D < p$. Then the number of $\mathbb{F}_p$-rational curves E with an optimal embedding of $\mathcal{O}$ into $End(E)$ is $\#\mathcal{CL}(\mathcal{O})[2]$ if and only if there is an element of $\mathcal{CL}(\mathcal{O})^2$ of norm q, where q is the prime defined at the beginning of Section 4.2. Otherwise, there are no $\mathbb{F}_p$-rational curves.*

Proof. E is $\mathbb{F}_p$-rational if and only if $\mathrm{End}(E)$ contains an element of degree p. So, we apply the above Proposition with $\mathfrak{b} = (1)$ and $N = p$.

If $\mathrm{End}(E) \cong R(\mathfrak{a})$ contains an element of degree p, then there is an element $\beta \in \bar{\mathfrak{q}}\bar{\mathfrak{a}}^2 \frac{1}{qN^K(\mathfrak{a})}$ of norm $\frac{1}{q}$ and such that $\lambda_{\mathfrak{a}}\beta \in \mathcal{O}$.

But $N^K(\bar{\mathfrak{q}}\bar{\mathfrak{a}}^2 \frac{1}{qN^K(\mathfrak{a})}) = \frac{1}{q}$, hence contains such a β if and only if $\bar{\mathfrak{q}}\bar{\mathfrak{a}}^2$ is principal (and generated by β), which is true if and only if $[\mathfrak{q}] = [\mathfrak{c}]^2$ and $[\mathfrak{a}] \in [\mathfrak{c}]\mathcal{CL}(\mathcal{O})[2]$. Thus, E is $\mathbb{F}_p$-rational implies that $[\mathfrak{q}]$ is a square, and there are at most $\mathcal{CL}(\mathcal{O})$ possibilities for the corresponding ideal classes $[\mathfrak{a}]$.

Now, we see that every choice of $[\mathfrak{a}] \in [\mathfrak{c}]\mathcal{CL}(\mathcal{O})[2]$ yields an endomorphism, i.e. that $\lambda_{\mathfrak{a}}\beta \in \mathcal{O}$. As $\lambda_{\mathfrak{a}}$ is chosen so that $\lambda_{\mathfrak{a}}\bar{\mathfrak{q}}\bar{\mathfrak{a}}^2 \frac{1}{qN^K(\mathfrak{a}} \subset \mathcal{O}$, and we choose β to generate $\bar{\mathfrak{q}}\bar{\mathfrak{a}}^2 \frac{1}{qN^K(\mathfrak{a}}$, we have $\lambda_{\mathfrak{a}}\beta \in \mathcal{O}$.

Hence we have two cases: Either $\mathfrak{q}$ is not a square in $\mathcal{CL}(\mathcal{O})^2$ and there are no $\mathbb{F}_p$-rational $\mathcal{O}$-oriented curves, or $\mathfrak{q} \in \mathcal{CL}(\mathcal{O})^2$ and we have exactly $\#\mathcal{CL}(\mathcal{O})[2]$ $\mathbb{F}_p$-rational $\mathcal{O}$-oriented curves.

□

Remark 4.3.7. Li, Li, and Ouyang also consider factorization of the Hilbert Class Polynomial mod p, including counting the number of $\mathbb{F}_p$-rational roots [32, Theorem 4.1]. They obtain the same number ($\#\mathcal{CL}(\mathcal{O})[2]$ or 0), under slightly different assumptions.

4.4 Special Case: $-DN < 2p$

In general, the injection in Theorem 1.2.1 is not a bijection. The expression with quadratic forms does not account for the condition depending on λ. Moreover, the quadratic forms only depend on the ideals $\mathfrak{b}$ and $\mathfrak{a}^2$, rather than $\mathfrak{b}$ and $\mathfrak{a}$. In other words, there are $\mathcal{CL}(\mathcal{O})[2]$ $\mathcal{O}$-oriented curves (E, θ) which give rise to the same binary quadratic forms.

In the special case that $-DN < 2p$, the injection in Theorem 1.2.1 can be extended to a bijection (up to a certain equivalence), and we give a more complete characterization of how non-horizontal isogenies arise.

Corollary 4.4.1. *Assume $-DN$ is coprime to p. Further assume that $\mathcal{O}$-orientations are unique up to conjugation. Then we have the following:*

1. *If $-DN < p$, all N-isogenies between $\mathcal{O}$-oriented curves are horizontal.*
2. *If $p < -DN < 2p$, write $-DN = p + r$ with $0 < r < p$. For each integral invertible ideal $\mathfrak{b}$ with $N(\mathfrak{b}) = r$, there is an $\mathcal{O}$-oriented curve E such that there is a non-horizontal N-isogeny $\mathfrak{b} * E \to E$. If ϕ is a non-horizontal isogeny between oriented curves,*

$\phi : \mathfrak{b} * E \to E$, then $\mathfrak{b}$ is in the same ideal class as an integral invertible ideal of norm r.

Proof. Since $-DN$ is coprime to p, p is inert in K and any N-isogenies are separable.

In the case that $-DN < p$, all N-isogenies correspond to a solution (w, x, y, z) of $Q(w, x) + pQ'(y, z) = -DN$. Since the forms Q and Q' are positive definite, any solution must have $y = z = 0$. By Lemma 4.3.1, the corresponding isogeny is horizontal.

Now suppose $-DN = p + r$ with $0 < r < p$.Let $\mathfrak{b}$ be an invertible ideal of norm r. First, we show that for some $\mathfrak{a}$, there exists is a solution (w, x, y, z) of the equation $Q(w, x)+pQ'(y, z) = -DN$ such that $(y, z) \neq 0$. Here Q is the form associated to $\mathfrak{b}$ and Q' is the form associated to $\bar{\mathfrak{q}}\bar{\mathfrak{b}}\bar{\mathfrak{a}}^2$. In this case, it is clear that the only solutions with $(y, z) \neq (0, 0)$ satisfy $Q'(y, z) = 1$ and $Q(w, x) = r$.

There is a solution to $Q'(y, z) = 1$ if and only if the ideal $\bar{\mathfrak{q}}\bar{\mathfrak{b}}\bar{\mathfrak{a}}^2$ is principal. We have that $\bar{\mathfrak{q}}\bar{\mathfrak{b}}\bar{\mathfrak{a}}^2$ is principal if and only if $[\mathfrak{a}]^2 = \bar{\mathfrak{q}}\bar{\mathfrak{b}}$. Since $N(\mathfrak{b}\bar{\mathfrak{q}}) = rq \equiv -pq \pmod{D}$, it follows from construction of q that $\mathfrak{b}\bar{\mathfrak{q}} \in \mathcal{CL}(\mathcal{O})^2$, so $\mathfrak{a}$ can be chosen so that $Q'(y, z) = 1$ has a solution.

Since $N(\mathfrak{b}) = r$, the integer r is an element of $\mathfrak{b}$ and $\frac{N(r)}{N(\mathfrak{b})} = r$. Therefore, there is a solution (w, x) to $Q(w, x) = r$.

We obtain a solution (w, x, y, z) for any $\mathfrak{a}$ such that $\mathfrak{a}^2 = \bar{\mathfrak{b}}\mathfrak{q}$. Note that if (w, x, y, z) is a solution, then so is $(w, x, -y, -z)$. We will show that exactly one of the two corresponds to an isogeny.

From (w, x) and (y, z), we construct the corresponding $\alpha \in \mathfrak{b}$ and $\beta \in \mathfrak{q}^{-1}\bar{\mathfrak{b}}\bar{\mathfrak{a}}\mathfrak{a}^{-1}$ such that $N(\alpha) + pqN(\beta) = -DNN(\mathfrak{b})$.

We will show that there is exactly one choice of $\mathfrak{a}$ such that one of $\frac{\alpha}{\sqrt{D}} \pm \lambda_{\mathfrak{a}\mathfrak{b}}\frac{\beta}{\sqrt{D}} \in \mathfrak{b}$.

Fix any choice of $\mathfrak{a}$ such that $\mathfrak{a}^2 = \bar{\mathfrak{q}}\mathfrak{b}$, and let $\lambda_{\mathfrak{a}\mathfrak{b}}$ be as described in Section 4.2. By construction, $\lambda_{\mathfrak{a}\mathfrak{b}}\mathfrak{q}^{-1}\bar{\mathfrak{a}}\mathfrak{a}^{-1}\bar{\mathfrak{b}}\mathfrak{b}^{-1} \subset \mathcal{O}$, so $\lambda_{\mathfrak{a}\mathfrak{b}}\beta \in \mathfrak{b}$.

By construction, $N(\lambda_{\mathfrak{a}\mathfrak{b}}) \equiv -pq \pmod{D}$. Thus, $N(\alpha) \equiv N(\lambda_{\mathfrak{a}\mathfrak{b}}\beta) \pmod{D}$. Since p is coprime to D, we also have $N(\alpha) = r^2$ is coprime to D. By [31, Lemma 7.8], there exists $c \in \mathcal{O}$ such that $\alpha - c\lambda_{\mathfrak{a}\mathfrak{b}}\beta \in \sqrt{D}\mathcal{O}$ and $N(c) \equiv 1 \pmod{D}$. Replacing β with $-\beta$ replaces c with $-c$. By [31, Lemma 7.7], there is a unique ideal $\mathfrak{c} \in \mathcal{CL}(\mathcal{O})^2$ such that $\lambda_{\mathfrak{a}\mathfrak{c}\mathfrak{b}}$ is one of $\pm c\lambda_{\mathfrak{a}\mathfrak{b}}$.

If E is the elliptic curve associated to $\mathfrak{a}\mathfrak{c}$, then we have shown that one of $[\alpha/\sqrt{D}, \pm\beta/\sqrt{D}]$ is an element of $\mathrm{Hom}(\mathfrak{b} * E, E)$. Since $\beta \neq 0$, it corresponds to an isogeny which is not horizontal.

To see that every non-horizontal N-isogeny between $\mathcal{O}$-oriented curves arises this way, suppose $\phi : \mathfrak{b} * E \to E$ is an isogeny which is not horizontal. Under the isomorphism in Corollary 4.2.3, there is a corresponding $\alpha \in \mathfrak{b}$ and $\beta \in \mathfrak{q}^{-1}\bar{\mathfrak{b}}\bar{\mathfrak{a}}\mathfrak{a}^{-1}$ such that $N(\alpha)+pqN(\beta) = -DNN(\mathfrak{b})$. Furthermore, because ϕ is not horizontal, it follows that $N(\beta) \neq 0$ and $N(\alpha) \neq 0$.

Since $\beta\mathfrak{q}\mathfrak{a} \subset \bar{\mathfrak{a}}\bar{\mathfrak{b}}$, it follows that $N(\beta)q$ is an integer. As $-DN = p + r$, the only possibility with β and α nonzero is that $N(\beta)/N(\mathfrak{b}) = \frac{1}{q}$ and $N(\alpha)/N(\mathfrak{b}) = r$.

Now, we show that the ideal class of $\mathfrak{b}$ contains an ideal of norm r. Since $\alpha \in \mathfrak{b}$, the conjugate $\bar{\alpha} \in \bar{\mathfrak{b}}$, and $\bar{\alpha}\bar{\mathfrak{b}}^{-1}$ is an integral ideal of norm r. Since $\bar{\mathfrak{b}}$ is in the same ideal class as $\mathfrak{b}^{-1}$, this shows that $\bar{\alpha}\bar{\mathfrak{b}}^{-1}$ is in the same ideal class as $\mathfrak{b}$ and has the desired norm.

□

When $-DN < 2p$, the existence of non-horizontal N-isogenies depends only on the existence of ideals of norm $-DN - p$, and every solution to the norm form equation gives rise to an isogeny. In the more general case that $-DN = mp + r$ where $0 < r < p$, we cannot classify the situation as completely.

Corollary 4.4.2. *Assume $-DN$ is coprime to p. Further assume that $\mathcal{O}$-orientations are unique up to conjugation. Then the following hold:*

1. *All pairs of $\mathcal{O}$-oriented curves connected by a non-horizontal N-isogeny are of the form $(E, \mathfrak{b} * E)$ where $N(\mathfrak{b}) \in \{r, p + r, 2p + r, \ldots, (m - 1)p + r\}$.*

2. *Suppose $(m - 1)p + r$ is prime to the conductor of $\mathcal{O}$. For each integral invertible ideal with $N(\mathfrak{b}) = (m - 1)p + r$, there exists an $\mathcal{O}$-oriented curve E and a non-horizontal N-isogeny between E and $\mathfrak{b} * E$.*

Proof. To see (1), every non-horizontal N-isogeny arises from a solution to $-DN = Q(w, x) + pQ'(y, z)$, where Q is the binary quadratic form associated to an ideal $\mathfrak{b}$. The possibilities for $Q(w, x) = n$ satisfying this equation and corresponding to non-horizontal isogenies are the integers $n \in \{r, p + r, 2p + r, \ldots, (m - 1)p + r\}$. As in the last step of Corollary 4.4.1, any such solution corresponds to an element $\alpha \in \mathfrak{b}$ with norm $N(\alpha) = N(\mathfrak{b})n$, and $\bar{\alpha}\bar{\mathfrak{b}}^{-1}$ is an integral ideal in the ideal class of $\mathfrak{b}$ of norm n.

To see (2), we repeat the first part of the proof of Corollary 4.4.1: our choice of q guarantees a solution to $Q'(y, z) = 1$ for some choices of $\mathfrak{a}$ (corresponding to $\#\mathcal{CL}(\mathcal{O})[2]$ choices of E), and if there exists an integral invertible ideal with $N(\mathfrak{b}) = (m - 1)p + r$, there is a solution to $Q(w, x) = (m - 1)p + r$ (corresponding to $\mathfrak{b}$). Since $N(\mathfrak{b})$ is coprime to D, we can use [31, Lemma 7.7] to show there is a unique E such that one of (w, x, y, z) or $(w, x, -y, -z)$ corresponds to an isogeny in $\mathrm{Hom}(E, \mathfrak{b} * E)$. □

4.5 Multiple Embeddings

We have been assuming that $\mathcal{O}$-oriented curves have optimal embeddings which are unique up to conjugation, which allows us to give a simple classification of which tuples correspond

to horizontal isogenies. However, our proof of Theorem 1.2.1 does not require uniqueness of optimal embeddings. In this section, we use Theorem 1.2.1 to give sufficient conditions for embeddings to be unique up to conjugation. When p is inert in K, this is equivalent to the statement that the Hilbert Class Polynomial has no repeated roots.

Let $\mathfrak{b}$ be non-principal (so that E and $\mathfrak{b} * E$ correspond to different oriented curves). To check if E and $\mathfrak{b} * E$ are isomorphic, we want to check if $\mathrm{Hom}(\mathfrak{b} * E, E)$ contains any elements of degree 1. In other words, we consider solutions to the following over all pairs of ideal classes $(\mathfrak{a}, \mathfrak{b})$:

$$Q(w, x) + pQ'(y, z) = -D$$

We make the following easy observation.

Proposition 4.5.1. *Let E be an elliptic curve over $\bar{\mathbb{F}}_p$. If $-D < p$, then there is at most one optimal embedding (up to conjugation) of $\mathcal{O}$ into End(E).*

Proof. Since Q and Q' are positive definite forms, the only solutions to this equation occur when $Q(w, x) = -D$. There is a corresponding $\alpha \in \mathfrak{b}$ such that $N(\alpha) = -DN(\mathfrak{b})$ and $\alpha \in \sqrt{D}\mathfrak{b}$. This implies that α generates $\sqrt{D}\mathfrak{b}$, and therefore $\mathfrak{b}$ is principal and generated by $\frac{\alpha}{\sqrt{D}}$. Since $\mathfrak{b}$ is principal, (E, θ) and $\mathfrak{b} * (E, \theta)$ are K-isomorphic, and the embedding induced by $\mathfrak{b}$ is equal to θ. □

The above also follows from work of Kaneko, who gives sufficient conditions for two (not necessarily distinct) quadratic imaginary orders to simultaneously embed into a maximal order in $B_{p,\infty}$ [25, Theorem 2']. We can extend the work of Kaneko using our techniques to the following situations.

Proposition 4.5.2. *Suppose $-D = p + r$ for $0 < r < p$. If there are no non-principal ideals of norm r in $\mathcal{O}$, then there is at most one optimal embedding of $\mathcal{O}$ into End(E) (up to K-isomorphism) for any elliptic curve E over $\bar{\mathbb{F}}_p$.*

Proof. Note that the statement is only interesting for supersingular E (optimal embeddings for ordinary curves are necessarily isomorphisms), so we assume p does not split in K. Also note that since $p \nmid p + r$, we are in the case that p is inert in $K = \mathbb{Q}(\sqrt{D})$.

Let $\mathfrak{b}$ be an ideal which is not principal, so that (E, ι) and $\mathfrak{b} * (E, \iota)$ are not K-isomorphic. An isomorphism (on the level of elliptic curves) $\mathfrak{b} * E \to E$ corresponds to a solution:

$$Q_{\mathfrak{b}}(w, x) + pQ_{\bar{\mathfrak{q}}\bar{\mathfrak{a}}^2\bar{\mathfrak{b}}}(y, z) = -D.$$

The forms are positive definite, so we have exactly two cases:

(1) $(y, z) = (0, 0)$ and $Q_{\mathfrak{b}}(w, x) = p + r$. Equivalently, the corresponding $[\alpha, \beta]$ has $\beta = 0$, so $\alpha - \lambda\beta \in \mathfrak{b}$, and $N(\sqrt{D}\alpha)/N(\mathfrak{b}) = p + r = -D$. But this implies $N(\alpha) = N(\mathfrak{b})$, so $\mathfrak{b}$ is a principal ideal.

(2) $Q_{\mathfrak{b}}(x, y) = Q_{\bar{\mathfrak{q}}\bar{\mathfrak{a}}^2\bar{\mathfrak{b}}} = r$. But by hypothesis, $Q_{\mathfrak{b}}(x, y) = r$ implies $\mathfrak{b}$ is principal.

In either case, we obtain a contradiction. □

Proposition 4.5.2 extends Kaneko's result, for example, to the case that $-D = p + 1$ or $-D = p + \ell$ where ℓ is inert in $\mathcal{O}$.

4.6 Extending the Attack on OSIDH

In this section, we will show how these results can be applied to extend Dartois and De Feo's attack to generate horizontal endomorphisms of higher degrees.

4.6.1 Dartois and De Feo's Original Attack.

In the setup of OSIDH, K is chosen to be a quadratic imaginary field of class number 1, and $\mathcal{O} \subset K$ is chosen to be an order of conductor ℓ^n for a prime ℓ (usually $\ell = 2$) and a large exponent n. A set of primes $q_1, \ldots, q_t$ splitting in K and an integer r are chosen so that (heuristically) each ideal class of $\mathcal{CL}(\mathcal{O})$ can be represented by an ideal of norm $N := \prod_{i=1}^{t} q_i^{e_i}$ for $|e_i| \leq r$, without "too much" overlap in the ideal classes.

Dartois and De Feo's algorithm to construct a horizontal endomorphism β has two steps: First, compute an integer $N = \prod_{i=1}^{t} q_i^{e_i}$ which is the norm of a principal ideal in $\mathcal{O}$; they do this by computing the shortest vector in the rank t relation lattice associated to $\mathcal{CL}(\mathcal{O})$. Then, construct β as a composition of isogenies corresponding to prime ideals over q_i. If $|e_i| \leq 2r$, the corresponding isogenies are public knowledge and β can be computed and guaranteed to be horizontal.

Heuristically, the integer N output by the first step has sufficiently small exponents e_i. Our application examines the case that for some e_i, we have $|e_i| > 2r$, in which case the best-known approach is meet-in-the-middle. This case would be relevant in the event that the heuristics fail. Even under the heuristics, this analysis would be relevant for a version of the attack which does not use the shortest vector (which would be desirable, as this step is

exponential time in the number of primes q_i), as well as for Dartois and De Feo's proposed countermeasure, for which the corresponding exponents e_i would be much larger.

4.6.2 Meet-in-the-Middle Extension

We describe how to supplement the approach of Dartois and De Feo via a meet-in-the-middle attack. Dartois and De Feo, following Onuki, propose this extension in their paper; however, they have no guarantee that the resulting isogeny is horizontal as desired. We describe parameters which guarantee that the resulting isogeny is horizontal, as well as parameters where it is easy to analyze the exceptional cases.

Let $N = \prod_{i=1}^{t} q_i^{e_i}$ be the norm of a principal ideal in $\mathcal{O}$, which can be computed by finding the shortest vector of the relation lattice associated to $\mathcal{CL}(\mathcal{O})$. Let E_0 be an $\mathcal{O}$-oriented curve. N is the degree of the horizontal endomorphism β we would like to produce in $\mathrm{End}(E_0)$.

Write $N = N'N''$ where $N' = \prod_{i=1}^{t} q_i^{d_i}$ with $|d_i| \leq 2r$, and assume $N'' \neq 1$. N' is the degree of the isogeny that we can produce via the public information of OSIDH, and N'' is the degree of the isogeny that we must supplement. We obtain an N'-isogeny $\beta' : E \to \mathfrak{b} * E$ which factors through E_0 and decomposes as a sequence of isogenies corresponding to the prime ideals $\mathfrak{q}_i$.

Next, we use a meet-in-the-middle attack to generate an N''-isogeny $\beta'' : E \to \mathfrak{b} * E$.

We examine the probability that β'' is horizontal in the two simplest cases below.

Proposition 4.6.1. *Let N'' be an integer coprime to p. Assume $-DN'' < 2p$. Fix an invertible ideal $\mathfrak{b} \subset \mathcal{O}$.*

1. *Suppose $-DN'' < p$. Then for all $\mathcal{O}$-oriented curves E, all N''-isogenies between E and $\mathfrak{b} * E$ are horizontal.*

2. *Suppose $-DN'' = p + r$ with $0 < r < p$. At most $\#\mathcal{CL}(\mathcal{O})[2]$ curves E admit non-horizontal N''-isogenies $E \to \mathfrak{b} * E$.*

3. *Let k be the number of invertible integral ideals of norm r in the ideal class of $\mathfrak{b}$. As E ranges over all $\mathcal{O}$-oriented curves, there are exactly k non-horizontal N''-isogenies starting at E and ending at $\mathfrak{b} * E$.*

As an easy example, let $-DN'' = p + r$ where r is prime. If r is inert in $\mathcal{O}$, there are no ideals of norm r, so in particular, 1(c) implies that all N''-isogenies are horizontal. If r splits in $\mathcal{O}$, there are two ideals $\mathfrak{b}$ and $\overline{\mathfrak{b}}$ of norm r. Then there are two non-horizontal N''-isogenies, given by the N''-isogeny $E \to \mathfrak{b} * E$ and its dual $\mathfrak{b} * E \to \overline{\mathfrak{b}} * \mathfrak{b} * E \cong E$, for a unique $\mathcal{O}$-oriented curve E.

We can detect if the meet-in-the-middle attack fails by computing the discriminant of β and checking if it is equal to $-D$. If not, then β is not horizontal, and we search for another N''-isogeny until we succeed. Without computing relations in the class group, we can bound the number of isogenies to try by the number of ideals of norm r in $\mathcal{O}$, which is determined only by the splitting behavior of prime factors of r in $\mathcal{O}$.

In more general parameter ranges, the situation becomes more complicated; one expects that horizontal isogenies become less likely as $\lfloor \frac{-DN}{p} \rfloor$, but it depends heavily on the structure of the class group.

4.6.3 Specialization to OSIDH

We consider the specific parameters of OSIDH. When $\ell = 2$ and $K = \mathbb{Q}(i)$ or $K = \mathbb{Q}(\sqrt{-3})$, $\#\mathcal{CL}(\mathcal{O})[2]$ is 2 or 4 (respectively) and $\#\mathcal{CL}(\mathcal{O}) \cong 2^{n-1}$, so the probability that there exist any non-horizontal N''-isogenies between E and $\mathfrak{b} * E$ is $\frac{1}{2^{n-2}}$ or $\frac{1}{2^{n-3}}$, where $n = 256$. As n grows, the probability decreases (keeping in mind that the N'' covered is also changing as D increases).

If the ring class polynomial associated to $\mathcal{O}$ can be computed efficiently, we can, under the parameters of OSIDH, modify the meet-in-the-middle algorithm so that we will always output a horizontal isogeny, regardless of the size of $-DN''$. This is because the only known way to make the group action computation described by Colo and Kohel well-defined is to ensure $-Dq_i < p$ for all primes q_i considered [37, Section 6.2].

Proposition 4.6.2. *Let $\phi_i(x, y)$ denote the q_i-th modular polynomial mod p, whose roots (j, j') are j-invariants of q_i-isogenous elliptic curves over $\overline{\mathbb{F}}_p$. Let $H(x)$ denote the ring class polynomial associated to $\mathcal{O}$, whose roots are j-invariants of $\mathcal{O}$-oriented elliptic curves. If $-Dq_i < p$, we have the following:*

For each root j of H, there are two roots of $\gcd(\phi_i(j, x), H(x))$, corresponding precisely to the action by $\mathfrak{q}_i$ or $\overline{\mathfrak{q}}_i$ on j. The corresponding q_i-isogenies are horizontal.

In the case that we can compute the required polynomials and gcds, we can modify the meet-in-the-middle algorithm: We construct isogeny paths of j-invariants $j_1, j_2, \ldots, j_m$ by selecting j_k from roots of $\gcd(\phi_i(j_{k-1}, x), H(x))$. Every step is necessarily horizontal, and any two paths meeting at a j-invariant j_m can be composed to form a horizontal isogeny.

4.6.4 Modification to OSIDH

One implication of our work is that N-isogenies between $\mathcal{O}$-oriented curves are less likely to be horizontal as $\lfloor \frac{-DN}{p} \rfloor$ grows, although a more precise statement would depend on the

structure of $\mathcal{CL}(\mathcal{O})$. Dartois and De Feo's countermeasure to their attack increases D as well as (heuristically) increasing N by considering the same set of primes, chosen to only cover a portion of the class group [18, Section 5.2]. An additional countermeasure to explore is to decrease the prime p, which would need to be done in a way to make the group action still possible to compute via modular polynomials. The group action computation can be reduced to an assumption about when two isogenies induce the same orientation [37, Assumption 5.1], which we can analyze via the same techniques that we introduce here.

Chapter 5
Notation and preliminaries for local orders and the Bruhat-Tits tree

In this chapter and the next, all orders $\mathcal{O}$ will be orders in a quaternion algebra. We consider the problem of computing the endomorphism ring of a supersingular elliptic curve E from a given order $\mathcal{O}_0 \subset \mathrm{End}(E)$ of finite index. Our strategy is to work locally by enlarging $\mathcal{O}_0$ at each prime dividing the reduced discriminant and finding a path in the Bruhat-Tits tree.

5.1 Notation and preliminaries for orders in a quaternion algebra

One key fact that we will use is the local-global principle for orders in a quaternion algebra. This states that an order in a global quaternion algebra is determined by its completions at each prime, see [54, Theorem 9.4.9, Lemma 9.5.3]. Maximality is a local property, i.e. an order $\mathcal{O} \subseteq B_{p,\infty}$ is maximal if and only if for all primes q, $\mathcal{O} \otimes \mathbb{Z}_q$ is a maximal order. Being maximal can also be expressed in terms of the reduced discriminant: an order $\mathcal{O}$ in $B_{p,\infty}$ is maximal if and only if the reduced discriminant $\mathrm{discrd}(\mathcal{O})$ is equal to p [54, p. 375]. The primes at which $\mathcal{O}$ fails to be maximal are exactly those primes dividing $\mathrm{discrd}(\mathcal{O})/p$.

Thus, the local-global principle reduces finding $\mathrm{End}(E)$ to finding $\mathrm{End}(E) \otimes \mathbb{Z}_q$ at each prime q dividing $\mathrm{discrd}(\mathcal{O})/p$. When $q = p$, there is a unique maximal order in the division algebra $B_{p,\infty} \otimes \mathbb{Z}_p$. In the case that $q \neq p$, the local order $\mathrm{End}(E) \otimes \mathbb{Z}_q$ is a maximal order of $B_{p,\infty} \otimes \mathbb{Q}_q \cong M_2(\mathbb{Q}_q)$.

Our strategy will be to enlarge $\mathcal{O}_0$ locally. We give the following definitions for the orders we obtain in this way.

Definition 5.1.1. Let $\mathcal{O}_0$ be an order. We say that an order $\mathcal{O}$ is a *q-enlargement* of $\mathcal{O}_0$ if $\mathcal{O}_0 \subset \mathcal{O}$ and $\mathcal{O} \otimes \mathbb{Z}_{q'} = \mathcal{O}_0 \otimes \mathbb{Z}_{q'}$ for all $q' \neq q$. We say that $\mathcal{O}$ is a *q-maximal q-enlargement*

if $\mathcal{O}$ is a q-enlargement such that $\mathcal{O} \otimes \mathbb{Z}_q$ is maximal.

Let $\mathcal{O} \subset B$ be a $\mathbb{Z}$-order We call $\mathcal{O}$ an **Eichler order** if $\mathcal{O} \subseteq B$ is the intersection of two (not necessarily distinct) maximal orders. The **codifferent** of an order is $\text{codiff}(\mathcal{O}) = \{\alpha \in B : \text{Trd}(\alpha\mathcal{O}) \subseteq \mathbb{Z}\}$. We say that $\mathcal{O}$ is *Gorenstein* if the lattice $\text{codiff}(\mathcal{O})$ is invertible as a lattice [54, 24.1.1]. We call $\mathcal{O}$ **Bass** if every superorder $\mathcal{O}' \supseteq \mathcal{O}$ is Gorenstein.

5.2 Local orders and the Bruhat-Tits tree

Given an order $\mathcal{O}_0$ of finite index in $\text{End}(E)$, we will compute $\text{End}(E)$ from $\mathcal{O}_0$ by enlarging it so that locally at a prime q it is maximal and equal to $\text{End}(E) \otimes \mathbb{Z}_q$. When $q = p$ this step follows from work of [53] since $B_{p,\infty} \otimes \mathbb{Z}_p$ is a division algebra and has a unique maximal order. When $q \neq p$ we will first compute some maximal order containing $\mathcal{O}_0 \otimes \mathbb{Z}_q$ and then find a path from that maximal order to $\text{End}(E) \otimes \mathbb{Z}_q$. Here we view both orders as vertices in the Bruhat-Tits tree for $GL_2(\mathbb{Q}_q)$.

Remark 5.2.1. Throughout this paper, a path in the Bruhat-Tits tree always refers to a *nonbacktracking* path.

For the remainder of this section, fix a prime $q \neq p$. We use the labelling conventions described by Tu [51].

Definition 5.2.2. The *Bruhat-Tits tree* is the graph whose vertices are rank 2 $\mathbb{Z}_q$-lattices up to homothety. Two lattice classes $[L]$ and $[L']$ are connected by an edge if and only if there are representatives L and L' such that $qL' \subsetneq L \subsetneq L'$.

Equivalently, one can consider the vertices of the Bruhat-Tits tree as maximal orders of $M_2(\mathbb{Q}_q)$, via the correspondence $[L] \mapsto \text{End}(L)$. In this case, two maximal orders Λ and Λ' are neighbors if and only if $[\Lambda : \Lambda \cap \Lambda'] = [\Lambda' : \Lambda \cap \Lambda'] = q$.

Fixing a basis for a lattice L_0 and identifying $\text{End}(L_0)$ with $M_2(\mathbb{Z}_q)$, we associate to each basis defining a lattice L a 2×2 matrix T that transforms the basis of L into that of L_0.

Then $\text{End}(L) = T^{-1}M_2(\mathbb{Z}_q)T$. Given L_0 and L, the matrix T is well-defined as an element of $\mathbb{Q}_q^* \,\text{GL}_2(\mathbb{Z}_q) \backslash \,\text{GL}_2(\mathbb{Q}_q)$, so we may assume that T is of the form $T = \begin{pmatrix} q^a & c \\ 0 & q^b \end{pmatrix}$ with $a, b \geq 0$, $c \in \mathbb{Z}/q^b\mathbb{Z}$, and $v_q(c) = 0$ if both a and b are positive.

The Bruhat-Tits tree is a $(q+1)$-regular tree. Each neighbor of a lattice L corresponds to a choice of cyclic sublattice of index q, which corresponds to a choice of matrices of the form $\begin{pmatrix} q & c \\ 0 & 1 \end{pmatrix}$ and $\begin{pmatrix} 1 & 0 \\ 0 & q \end{pmatrix}$. More generally, a path of length n starting at the root $M_2(\mathbb{Z}_q)$

(labelled by the 2×2 identity matrix) corresponds to a product of such matrices. We make the following definition.

Definition 5.2.3. For each c such that $0 \leq c \leq q-1$, let $\gamma_c := \begin{pmatrix} 1 & c \\ 0 & q \end{pmatrix}$. Let $\gamma_\infty = \begin{pmatrix} q & 0 \\ 0 & 1 \end{pmatrix}$. Let $\Sigma = \{\gamma_c : 0 \leq c \leq q-1\} \cup \{\gamma_\infty\}$. We call a finite sequence of matrices $\{c_i\}_{i=1}^n$ a *matrix path* if each $c_i \in \Sigma$ and $c_{i+1}c_i \notin qM_2(\mathbb{Z}_q)$. The *length* of the matrix path $\{c_i\}_{i=1}^n$ is n. Each path of length n starting at the root corresponds to a matrix path. If $\{c_i\}_{i=1}^n$ is a matrix path, we call the product $T = c_n c_{n-1} \dots c_1$ the *associated matrix*.

There is a bijection between paths of length n in the Bruhat-Tits tree starting at $M_2(\mathbb{Z}_q)$ and matrix paths of length n. The endpoint of the path corresponding to the matrix path $\{c_i\}_{i=1}^n$ is the order $T^{-1}M_2(\mathbb{Z}_q)T$, where T is the associated matrix. The vertices of the path are

$$M_2(\mathbb{Z}_q), c_1^{-1}M_2(\mathbb{Z}_q)c_1, c_1^{-1}c_2^{-1}M_2(\mathbb{Z}_q)c_2c_1, \dots, (c_1^{-1}c_2^{-1}\cdots c_n^{-1})M_2(\mathbb{Z}_q)(c_n \cdots c_2c_1).$$

Depending on the context, we may represent vertices in the Bruhat-Tits tree as maximal orders in $M_2(\mathbb{Z}_q)$, lattices, or 2×2 matrices T as above.

5.2.1 Distance

We have the usual notion of distance in the Bruhat-Tits tree.

Definition 5.2.4. The *distance* between two vertices in the Bruhat-Tits tree v and v', denoted $d(v, v')$, is the length of the unique path between v and v'. We denote the distance between v and v' by $d(v, v')$. Here, v and v' may be represented by homothety classes of lattices, maximal orders in $M_2(\mathbb{Z}_q)$, or the matrices associated to the matrix path.

Definition 5.2.5. Let ℓ be a postive integer and v a vertex in the Bruhat-Tits tree. The *ℓ-neighborhood of v* is the set

$$N_\ell(v) \colon = \{v' : d(v', v) \leq \ell\}.$$

We also have the analogous notion of distance to a path and neighborhood of a path.

Definition 5.2.6. Let P be the set of vertices along a path in the Bruhat-Tits tree. The *distance* between a vertex v and P is $\min\{d(v, v') \colon v' \in P\}$. The distance between v and P is denoted $d(v, P)$.

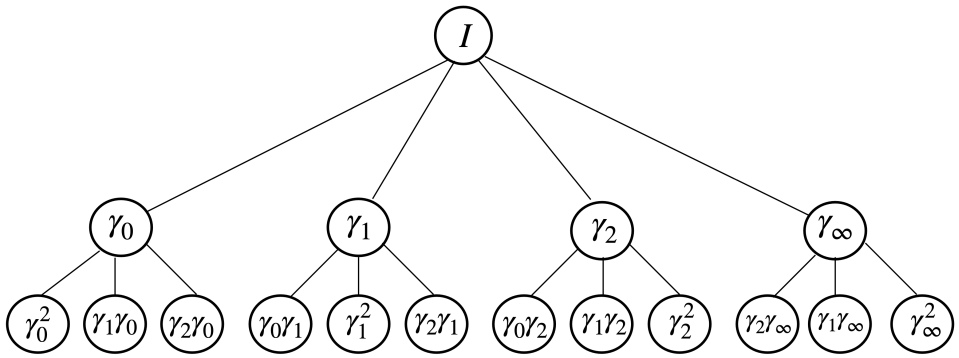

Figure 5.1. The (truncated) Bruhat-Tits tree for $q = 3$, with vertices labelled by the corresponding matrices. The root of the tree, labelled I, corresponds to $M_2(\mathbb{Z}_q)$. The vertex labelled with matrix T corresponds to the order $T^{-1}M_2(\mathbb{Z}_q)T$.

Definition 5.2.7. Let P be the set of vertices along a path in the Bruhat-Tits tree and let ℓ be a nonnegative integer. The *ℓ-neighborhood of P* is the set

$$N_\ell(P) \colon = \{v' : d(v', P) \leq \ell\}.$$

5.2.2 Distance and matrix labelling

This section relates the distance between two vertices in the Bruhat-Tits tree to the matrix labelling just described. We will also get a bound on the distance in terms of the reduced discriminant of the intersection of the two maximal orders.

Proposition 5.2.8. *Let $T = b_r \cdots b_2 b_1 a_k \cdots a_2 a_1$ and $T' = c_s \cdots c_1 a_k \cdots a_2 a_1$, such that $a_i, b_i, c_i \in \Sigma$, the product of two consecutive matrices is not in $M_2(q\mathbb{Z}_q)$, and $b_1 \neq c_1$. Then*

$$d(T^{-1}M_2(\mathbb{Z}_q)T, T'^{-1}M_2(\mathbb{Z}_q)T') = r + s.$$

Proof. Let $\gamma = a_k \cdots a_1$. The unique path from $T^{-1}M_2(\mathbb{Z}_q)T$ to $M_2(\mathbb{Z}_q)$ and the unique path from $T'^{-1}M_2(\mathbb{Z}_q)T'$ intersect exactly in the path from $\gamma^{-1}M_2(\mathbb{Z}_q)\gamma$ to $M_2(\mathbb{Z}_q)$. Thus, we obtain a (nonbacktracking) path from $T^{-1}M_2(\mathbb{Z}_q)T$ to $T'^{-1}M_2(\mathbb{Z}_q)T'$ by first taking the path of length r from $T^{-1}M_2(\mathbb{Z}_q)T$ to $\gamma^{-1}M_2(\mathbb{Z}_q)\gamma$, and concatenating it with the path of length s from $\gamma^{-1}M_2(\mathbb{Z}_q)\gamma$ to $T'^{-1}M_2(\mathbb{Z}_q)T'$. □

We can relate the reduced discriminant of an order to the distance between maximal orders containing it.

Lemma 5.2.9. *Let Λ be a $\mathbb{Z}_q$-order in $M_2(\mathbb{Q}_q)$ such that $\Lambda \subset \Lambda_1 \cap \Lambda_2$ for maximal orders Λ_1, Λ_2. Then $v_q(\mathrm{discrd}(\Lambda)) \geq v_q(\mathrm{discrd}(\Lambda_1 \cap \Lambda_2))$.*

Proof. By [54, Lemma 15.2.15], $\operatorname{discrd}(\Lambda) = [\Lambda_1 \cap \Lambda_2 : \Lambda] \operatorname{discrd}(\Lambda_1 \cap \Lambda_2)$. The index is an integral ideal, so $v_q(\operatorname{discrd}(\Lambda)) \geq v_q(\operatorname{discrd}(\Lambda_1 \cap \Lambda_2))$. □

From the lemma, we obtain the two following useful corollaries.

Corollary 5.2.10. *Suppose $\Lambda \subset M_2(\mathbb{Q}_q)$ is a $\mathbb{Z}_q$-order. Let $e = v_q(\operatorname{discrd}(\Lambda))$. If Λ is contained in two maximal orders Λ_1 and Λ_2, then $d(\Lambda_1, \Lambda_2) \leq e$.*

Proof. When Λ_1 and Λ_2 are maximal orders, the distance between them is $v_q(\operatorname{discrd}(\Lambda_1 \cap \Lambda_2))$. The result then follows from Lemma 5.2.9. □

Corollary 5.2.11. *Let Λ be a $\mathbb{Z}_q$-order of finite index in $M_2(\mathbb{Q}_q)$. Then Λ is contained in finitely many maximal orders.*

Proof. This is immediate from Corollary 5.2.10. If Λ_1 is a maximal order containing Λ, then all maximal orders containing Λ are at most $v_q(\operatorname{discrd}(\Lambda))$ steps from Λ_1. □

Remark 5.2.12. Suppose $\Lambda \subset \operatorname{End}(E) \otimes \mathbb{Z}_q$. If we can construct a maximal order Λ_1 which contains Λ, the preceding corollaries give us a starting point for how to locate $\operatorname{End}(E) \otimes \mathbb{Z}_q$ in the Bruhat-Tits tree. A naive approach would be to check all orders within $e = v_q(\operatorname{discrd}(\Lambda))$ steps from Λ_1 in the Bruhat-Tits tree. However, when $e \geq 1$, there are $1 + (q+1)\frac{q^e - 1}{q-1}$ maximal orders at most e steps from Λ_1. Working with each of these orders is computationally infeasible for general Λ.

5.3 Finite intersections of maximal orders

In this section, we review Tu's results on finite intersections of maximal orders in $M_2(\mathbb{Q}_q)$. As an application, for each path P and $\ell \geq 0$, we construct an order $\tilde{\Lambda}$ which is contained in a maximal order Λ' if and only if $\Lambda' \in N_\ell(P)$. This construction allows us to work with many maximal orders at once.

Definition 5.3.1. [51, Notation 7] Let S be a finite set of maximal orders. We define

$$d_3(S) := \max\{d(\Lambda_1, \Lambda_2) + d(\Lambda_2, \Lambda_3) + d(\Lambda_3, \Lambda_1)\},$$

where the maximum is taken over all choices of $\Lambda_1, \Lambda_2, \Lambda_3 \in S$. The orders Λ_i need not be distinct.

We restate Tu's main theorem, specialized to our case $K = \mathbb{Q}_q$.

Theorem 5.3.2. *[51, Theorem 8] Let S be a finite set of maximal orders in $M_2(\mathbb{Q}_q)$. Let $\Lambda_1, \Lambda_2, \Lambda_3 \in S$ be such that $d_3(\{\Lambda_1, \Lambda_2, \Lambda_3\}) = d_3(S)$. Then $\bigcap_{\Lambda \in S} \Lambda = \Lambda_1 \cap \Lambda_2 \cap \Lambda_3$. The orders $\Lambda_1, \Lambda_2, \Lambda_3$ need not be distinct.*

Our first lemma relates $d_3(S)$ to the reduced discriminant $\mathrm{discrd}(\bigcap_{\Lambda \in S} \Lambda)$.

Lemma 5.3.3. *Let S be a finite set of maximal orders in $M_2(\mathbb{Q}_q)$. Then*

$$v_q(\mathrm{discrd}(\bigcap_{\Lambda \in S} \Lambda)) = d_3(S)/2.$$

Proof. If S consists of a single maximal order Λ_1, then $d_3(S) = 0$. Furthermore, $\cap_{\Lambda \in S} \Lambda = \Lambda_1$ is conjugate to $M_2(\mathbb{Z}_q)$, and therefore $v_q(\mathrm{discrd}(\Lambda_1)) = v_q(\mathrm{discrd}(M_2(\mathbb{Z}_q))) = 0 = d_3(S)/2$.

Suppose $\cap_{\Lambda \in S} \Lambda$ is Eichler, say $\cap_{\Lambda \in S} = \Lambda_1 \cap \Lambda_2$ with $\Lambda_1, \Lambda_2 \in S$ and such that $d_3(S) = d_3(\{\Lambda_1, \Lambda_2\})$. Then $v_q(\mathrm{discrd}(\Lambda_1 \cap \Lambda_2)) = d(\Lambda_1, \Lambda_2)$. Writing $\Lambda_1 \cap \Lambda_2 = \Lambda_1 \cap \Lambda_2 \cap \Lambda_2$, we have $d_3(S) = d(\Lambda_1, \Lambda_2) + d(\Lambda_2, \Lambda_2) + d(\Lambda_2, \Lambda_1) = 2d(\Lambda_1, \Lambda_2)$. Hence $v_q(\mathrm{discrd}(\cap_{\Lambda \in S} \Lambda)) = d_3(S)/2$.

If $\cap_{\Lambda \in S} \Lambda = \Lambda_1 \cap \Lambda_2 \cap \Lambda_3$ is not Eichler, there is an order Λ_0 which lies on the path between any two of the Λ_i. Letting $m = d(\Lambda_1, \Lambda_0)$, $n = d(\Lambda_2, \Lambda_0)$, and $\ell = d(\Lambda_3, \Lambda_0)$, we have $d_3(S) = d(\Lambda_1, \Lambda_2) + d(\Lambda_2, \Lambda_3) + d(\Lambda_3, \Lambda_1) = (m+n) + (n+\ell) + (\ell+m) = 2m + 2n + 2\ell$.

Furthermore, the intersection $\Lambda_1 \cap \Lambda_2 \cap \Lambda_3$ is conjugate to the order with basis

$$\left\{ \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 0 & q^n \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ q^m & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 0 & q^\ell \end{pmatrix} \right\}.$$

(See [51, proof of Theorem 2] for details.) As $\mathrm{disc}(\{\alpha_1, \alpha_2, \alpha_3, \alpha_4\}) = \det(\mathrm{Trd}(\alpha_i \alpha_j))$, a computation shows

$$\mathrm{disc}(\Lambda_1 \cap \Lambda_2 \cap \Lambda_3) = -q^{2m+2n+2\ell} = -q^{d_3(S)}.$$

Hence $v_q(\mathrm{discrd}(\Lambda_1 \cap \Lambda_2 \cap \Lambda_3)) = d_3(S)/2$. □

We'll also use the following lemma which is key to the proof of Tu's Theorem 8.

Lemma 5.3.4. *[51, Lemma 12] Let $S = \{\Lambda_1, \Lambda_2, \Lambda_3\}$ be a set of maximal orders, and let Λ_4 be a maximal order such that $d_3(S \cup \{\Lambda_4\}) = d_3(S)$. Then $\Lambda_4 \supset \Lambda_1 \cap \Lambda_2 \cap \Lambda_3$.*

The converse is also true, which is shown in the next Lemma.

Lemma 5.3.5. *Let $S = \{\Lambda_1, \Lambda_2, \Lambda_3\}$ be a set of maximal orders. Suppose $\Lambda_4 \supset \Lambda_1 \cap \Lambda_2 \cap \Lambda_3$. Then $d_3(S) = d_3(S \cup \{\Lambda_4\})$.*

Proof. We have $v_q(\mathrm{discrd}((\bigcap_{\Lambda \in S} \Lambda) \cap \Lambda_4)) = d_3(S \cup \{\Lambda_4\})/2$, and $v_q(\mathrm{discrd}(\bigcap_{\Lambda \in S} \Lambda)) = d_3(S)/2$. But $\bigcap_{\Lambda \in S} \Lambda \subset \Lambda_4$ implies that $(\bigcap_{\Lambda \in S} \Lambda) \cap \Lambda_4 = \cap_{\Lambda_\in S} \Lambda$, hence the reduced discriminants are equal. By Lemma 5.3.3, $d_3(S \cup \{\Lambda_4\}) = d_3(S)$. □

Corollary 5.3.6. *Let P be the set of maximal orders along a path, and let $\ell \geq 0$. Let $\tilde{\Lambda} = \bigcap_{\Lambda \in N_\ell(P)} \Lambda$. Then the set of maximal orders containing $\tilde{\Lambda}$ is $N_\ell(P)$. Moreover, $v_q(\mathrm{discrd}(\tilde{\Lambda})) = 3\ell + \mathrm{card}(P) - 1$.*

Proof. We will describe $\tilde{\Lambda}$ as an intersection of at most 3 maximal orders.

If $\ell = 0$ and P is a single point, then $N_\ell(P)$ consists of a single order, which is equal to $\tilde{\Lambda}$. In this case, as $\tilde{\Lambda}$ is maximal, we have $v_q(\mathrm{discrd}(\tilde{\Lambda})) = 0 = \mathrm{card}(P) - 1$.

If $\ell = 0$ and $\mathrm{card}(P) > 1$, then $N_\ell(P) = P$. In this case, $\tilde{\Lambda} = \Lambda_1 \cap \Lambda_2$ where Λ_1 and Λ_2 are the endpoints of P. The only orders containing $\tilde{\Lambda}$ are those in P. We have $v_q(\mathrm{discrd}(\Lambda_1 \cap \Lambda_2)) = d(\Lambda_1, \Lambda_2) = \mathrm{card}(P) - 1$.

Now, assume $\ell > 0$. Let $\tilde{\Lambda}_1, \tilde{\Lambda}_2, \tilde{\Lambda}_3$ denote any choice of three orders in $N_\ell(P)$, and let $\tilde{\Lambda}'_i$ denote the order on the path P which is closest to $\tilde{\Lambda}_i$.

By the triangle inequality, we have $d(\tilde{\Lambda}_i, \tilde{\Lambda}_j) \leq d(\tilde{\Lambda}_i, \tilde{\Lambda}'_i) + d(\tilde{\Lambda}'_i, \tilde{\Lambda}'_j) + d(\tilde{\Lambda}'_j, \tilde{\Lambda}_j) \leq 2\ell + d(\tilde{\Lambda}'_i, \tilde{\Lambda}'_j)$. Hence $d_3(\{\tilde{\Lambda}_1, \tilde{\Lambda}_2, \tilde{\Lambda}_3\}) \leq 6\ell + d(\tilde{\Lambda}'_1, \tilde{\Lambda}'_2) + d(\tilde{\Lambda}'_2, \tilde{\Lambda}'_3) + d(\tilde{\Lambda}'_3, \tilde{\Lambda}'_1)$. Since $\tilde{\Lambda}'_1, \tilde{\Lambda}'_2, \tilde{\Lambda}'_3$ lie along the same path P, this sum is at most $2(\mathrm{card}(P) - 1)$. Hence $d_3(N_\ell(P)) \leq 6\ell + 2(\mathrm{card}(P) - 1)$.

Now, choose orders $\Lambda_i \in N_\ell(P)$ in the following way: Choose a path of length ℓ starting at an endpoint Λ'_1 of P which is otherwise disjoint from P; the end of this path will be Λ_1. To construct Λ_2, choose a path of length ℓ which starts at the opposite endpoint Λ'_2 of P and is otherwise disjoint from both P and the path from Λ_1 to Λ'_1. As the Bruhat-Tits tree is $(q+1)$-regular and $q \geq 2$, this can be done. By construction, $d(\Lambda_1, \Lambda_2) = 2\ell + \mathrm{card}(P) - 1$.

We then construct Λ_3 as follows. Choose a path of length ℓ starting at any point of P which is otherwise disjoint from the paths P, the path from Λ_1 to Λ'_1, and the path from Λ_2 to Λ'_2. The path from Λ_1 to Λ'_1 and Λ_2 to Λ'_2 are automatically disjoint unless P is a single point. Thus, this disjointness restriction can be accomplished if and only if we can choose a path starting at any point of P to avoid two adjacent edges. As the Bruhat-Tits tree is $(q+1)$-regular and $q \geq 2$, we can choose Λ_1, Λ_2, and Λ_3 as specified. By construction, $d(\Lambda_1, \Lambda_3) + d(\Lambda_2, \Lambda_3) = d(\Lambda_1, \Lambda_2) + 2\ell = 4\ell + 2(\mathrm{card}(P) - 1)$.

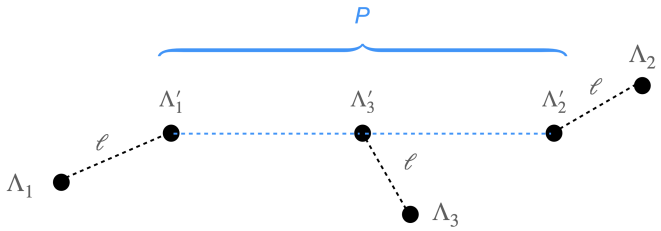

Figure 5.2. Constructing $\Lambda_1, \Lambda_2, \Lambda_3$ such that $\bigcap_{\Lambda \in N_\ell(P)} \Lambda = \Lambda_1 \cap \Lambda_2 \cap \Lambda_3$.

For this choice of $\Lambda_1, \Lambda_2, \Lambda_3$, we have $d_3(\{\Lambda_1, \Lambda_2, \Lambda_3\}) = 6\ell + 2(\mathrm{card}(P) - 1)$. Thus, $d_3(N_\ell(P)) = 6\ell + 2(\mathrm{card}(P) - 1)$. By Theorem 5.3.2, we can write $\tilde{\Lambda} = \Lambda_1 \cap \Lambda_2 \cap \Lambda_3$, and by

Lemma 5.3.3, we have $v_q(\mathrm{discrd}(\tilde{\Lambda})) = 3\ell + \mathrm{card}(P) - 1$.

Now, we want to show that the set $S := \{\Lambda \text{ maximal}\colon \tilde{\Lambda} \subset \Lambda\}$ is equal to $N_\ell(P)$. It is clear that $N_\ell(P) \subset S$ by construction. Suppose $\Lambda_4 \notin N_\ell(P)$, so that $d(\Lambda_4, P) > \ell$. We will show that $d_3(\{\Lambda_1, \Lambda_2, \Lambda_3, \Lambda_4\}) > 6\ell + 2(\mathrm{card}(P) - 1)$ and hence $\Lambda_4 \notin S$ by Lemma 5.3.5.

Let Λ_4' be the point of P which is closest to Λ_4. By construction, the paths Λ_i to Λ_i' for $i \leq 3$ are pairwise disjoint except possibly at Λ_i'. Thus, there is at most one k such that the paths Λ_k to Λ_k' and Λ_4 to Λ_4' intersect in more than one point.

Case 1: There is no such k, or Λ_4' is not one of the endpoints of P.

Consider $d_3(\{\Lambda_1, \Lambda_2, \Lambda_4\})$. For distinct $i, j \in \{1, 2, 4\}$, the path between Λ_i and Λ_j passes through Λ_i' and Λ_j'. We have $d(\Lambda_i, \Lambda_j) = d(\Lambda_i, \Lambda_i') + d(\Lambda_i', \Lambda_j') + d(\Lambda_j, \Lambda_j')$. By construction, $d(\Lambda_i, \Lambda_i') = \ell$ if $i \in \{1, 2, 3\}$, and since Λ_1' and Λ_2' are the endpoints of P, we have $d(\Lambda_1', \Lambda_2') + d(\Lambda_2', \Lambda_4') + d(\Lambda_4', \Lambda_1') = \mathrm{card}(P) - 1$. Hence $d_3(\{\Lambda_1, \Lambda_2, \Lambda_4\}) = 4\ell + 2(\mathrm{card}(P) - 1) + 2d(\Lambda_4, \Lambda_4') > d_3(S)$ if $d(\Lambda_4, \Lambda_4') > \ell$. Hence $\Lambda_4 \notin S$.

Case 2: Either $\Lambda_4' = \Lambda_1'$ or $\Lambda_4' = \Lambda_2'$.

Say $\Lambda_4' = \Lambda_1'$. Then consider $d_3(\{\Lambda_2, \Lambda_3, \Lambda_4\})$. Arguing as in the previous case, and noting that Λ_4' and Λ_2' are endpoints of P, we similarly get $d_3(\{\Lambda_2, \Lambda_3, \Lambda_4\}) = 4\ell + 2(\mathrm{card}(P) - 1) + 2d(\Lambda_4, \Lambda_4') > d_3(S)$ if $d(\Lambda_4, \Lambda_4') > \ell$. Hence $\Lambda_4 \notin S$. □

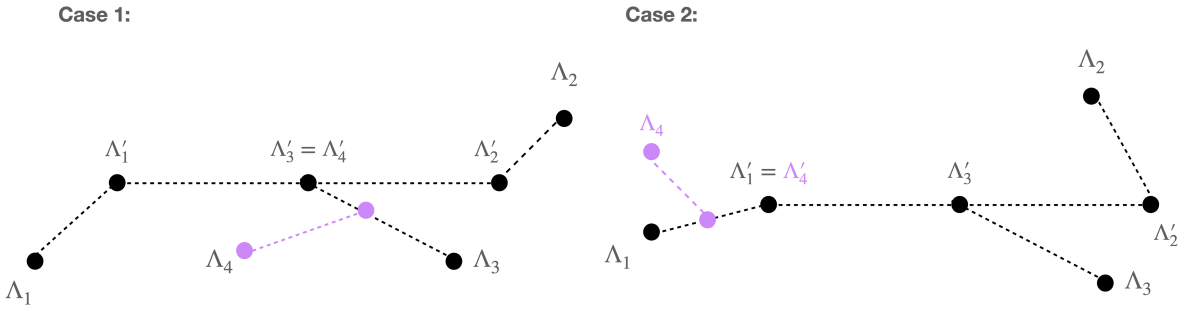

Figure 5.3. Case 1 and Case 2 in the proof of Corollary 5.3.6.

Chapter 6
Connecting Kani's Lemma and the Bruhat-Tits tree to compute supersingular endomorphism rings

In this chapter, we describe our algorithm for pathfinding in the Bruhat-Tits tree from a given order $\mathcal{O}_0 \subset \mathrm{End}(E)$. In Section 6.1, we give background on higher-dimensional isogenies and describe an Endomorphism-Testing Algorithm which determines if a rational multiple of an endomorphism is still an endomorphism. In Section 6.2, we give an algorithm which computes the distance between a q-maximal enlargment of $\mathcal{O}_0$ and $\mathrm{End}(E) \otimes \mathbb{Z}_q$ by testing if $\mathrm{End}(E) \otimes \mathbb{Z}_q$ contains certain orders. In Section 6.3, we give an algorithm to compute an explicit embedding of $\mathcal{O}_0$ into $M_2(\mathbb{Z}_q)$ and for testing containment of a local order in $\mathrm{End}(E) \otimes \mathbb{Z}_q$. In Section 6.4, we give an algorithm which constructs the path from $M_2(\mathbb{Z}_q)$ to the endomorphism ring in the Bruhat-Tits tree. In Section 6.5, we give a more efficient algorithm when the input order $\mathcal{O}_0$ is locally Bass. In Section 6.6, we describe the full algorithm to compute the endomorphism ring $\mathrm{End}(E)$ from input $\mathcal{O}_0 \subset \mathrm{End}(E)$. In Section 6.7, we describe a more efficient algorithm if more is known about the subgraph of orders containing $\mathcal{O}_0$.

6.1 Using Higher-Dimensional Isogenies for Endomorphism-Testing

One of the key tools we will use is an algorithm which determines if a rational multiple of an endomorphism is an endomorphism. This algorithm first appeared in [43, Section 4], and a detailed algorithm with correctness proof and complexity analysis were given in [24]. The content of Section 6.1 was written before [24] was posted and is formulated slightly differently.

Proposition 6.1.1. *[Divide algorithm] There exists an algorithm which takes as input an elliptic curve E defined over $\mathbb{F}_{p^k}$, an endomorphism $\beta \in End(E)$, and an integer n, and outputs TRUE if $\frac{\beta}{n} \in End(E)$ and FALSE if $\frac{\beta}{n} \notin End(E)$. The algorithm runs in polynomial-time in $\log(p^k)$ and $\log(\deg(\beta))$.*

This algorithm was first outlined by Robert in a special case to compute endomorphism rings of ordinary curves [43, Section 4]. The main idea is to use Kani's Lemma to translate the problem into a higher dimension, where there is enough flexibility to impose powersmoothness. A proof of Proposition 6.1.1 is given in [24, Section 4]. Before [24] was posted, we had written down the details for the algorithm, proof of correctness, and run-time analysis, which we include here. We will refer to this algorithm by our original name for it, the Endomorphism-Testing Algorithm.

We also note that higher-dimensional isogenies will only be used in this section, so one could take Proposition 6.1.1 as a blackbox and skip to Section 6.2 if desired.

6.1.1 Isogenies between polarized abelian varieties and their degrees

Definition 6.1.2. [35, p. 126] A *polarization* of an abelian variety X defined over a field k is an isogeny $\lambda : X \to X^\vee$ to the dual variety $X^\vee$ so that $\lambda_{\overline{k}} = \phi_{\mathcal{L}}$ for some ample invertible sheaf $\mathcal{L}$ on $X_{\overline{k}}$. Here $\phi_{\mathcal{L}} : A(k) \to \mathrm{Pic}(A)$ is the map given by $a \mapsto t_a^*\mathcal{L} \otimes \mathcal{L}^{-1}$ with t_a the translation-by-a map.

Notation: Given an isogeny $\Phi : A \to B$ between abelian varieties we denote by $\Phi^\vee : B^\vee \to A^\vee$ the dual isogeny (see [36, p. 143]).

Definition 6.1.3. Given a positive integer N, an *N-isogeny* $\Phi : (A, \lambda_A) \to (B, \lambda_B)$ between principally polarized abelian varieties (A, λ_A) and (B, λ_B) is an isogeny such that $\Phi^\vee \circ \lambda_B \circ \Phi = N\lambda_A$. An (N, N)-isogeny $\Phi : (A, \lambda_A) \to (B, \lambda_B)$ of abelian varieties of dimension g is an N-isogeny whose kernel is isomorphic to $(\mathbb{Z}/N\mathbb{Z})^g$.

Let A be an abelian variety with a polarization λ. Since λ is an isogeny $A \to \hat{A}$, it has an inverse in $\mathrm{Hom}(\hat{A}, A) \otimes \mathbb{Q}$. The *Rosati involution* on $\mathrm{Hom}(\hat{A}, A) \otimes \mathbb{Q}$ corresponding to λ is

$$a \mapsto a^\dagger = \lambda^{-1} \circ \hat{\alpha} \circ \lambda.$$

In this paper we will consider endomorphisms of products of elliptic curves and abelian varieties. Given an abelian variety A, an integer $r > 1$ and isogenies $\phi_{i,j} : A \to A$ for $1 \leq i, j \leq r$, the $r \times r$ matrix $M = (\phi_{i,j})_{1 \leq i,j \leq r}$ represents the isogeny

$\Phi : A^r \to A^r$ sending

$$(P_1, \ldots, P_r) \text{ to } (\phi_{1,1}(P_1) + \cdots + \phi_{1,r}(P_r), \ldots, \phi_{r,1}(P_1) + \cdots + \phi_{r,r}(P_r)) .$$

We refer to this as the *matrix form of* Φ.

Definition 6.1.4. Let A be a principally polarized abelian variety. Consider $\Phi : A^r \to A^r$ with matrix form $M = (\phi_{i,j})_{1 \le i,j \le r}$ as above. Let $\phi_{i,j}^\dagger : A \to A$ be the Rosati involution of $\phi_{i,j}$. Define $\hat{\Phi} : A^r \to A^r$ as the endomorphism represented by the matrix $\hat{M} = (\phi_{j,i}^\dagger)_{1 \le i,j \le r}$.

Definition 6.1.3 can also be rephrased as follows, see [42, Section 3.1].

Proposition 6.1.5. *Let A be principally polarized, and let $\Phi : A^r \to A^r$ be an isogeny with matrix form M. Then $\hat{M} \cdot M = N \cdot Id_r$ if and only if Φ is an N-isogeny with respect to the product polarization.*

Proposition 6.1.6. *Let E be an elliptic curve. Let $\Phi : E^r \to E^r$ be an N-isogeny of principally-polarized abelian varieties whose matrix form is $M = (\phi_{i,j})_{1 \le i,j \le r}$. Then the degrees of the isogenies $\phi_{i,j} : E \to E$ are bounded above by N.*

Proof. By the previous proposition, $\hat{M} \cdot M = N \cdot \mathrm{Id}_r$. In particular, the i-th diagonal entry of $\hat{M} \cdot M$ is given by $\sum_{j=1}^r \phi_{j,i}^\dagger \phi_{j,i} = N$. For elliptic curves, $\phi_{j,i}^\dagger$ is the dual isogeny of $\phi_{j,i}$, so we have $\sum_{j=1}^r \deg(\phi_{j,i}) = N$ (by convention, the degree of the 0 map is 0). As the degree of an isogeny is nonnegative, this implies that $\deg(\phi_{j,i}) \le N$ for $1 \le i, j \le r$. □

6.1.2 Isogeny Diamonds and Kani's Lemma

We now give the definition of an isogeny diamond in the setting of abelian varieties. This was first introduced by Kani [26] for elliptic curves and generalized in [42] to principally polarized abelian varieties.

Definition 6.1.7. A (d_1, d_2)-*isogeny factorization configuration* is a $d_1 \cdot d_2$-isogeny $f : A \to B$ between principally polarized abelian varieties of dimension g which has two factorizations $f = f_1' \circ f_1 = f_2' \circ f_2$ with f_1 a d_1-isogeny, f_2 a d_2-isogeny. If, in addition, d_1 and d_2 are relatively prime we call this configuration a (d_1, d_2)-*isogeny diamond configuration.*

$$\begin{array}{ccc} A & \xrightarrow[f_1]{} & A_1 \\ {\scriptstyle f_2}\downarrow & & \downarrow{\scriptstyle f_1'} \\ A_2 & \xrightarrow[f_2']{} & B \end{array}$$

Lemma 6.1.8 (Kani's Lemma)**.** *Let $f = f_1' \circ f_1 = f_2' \circ f_2$ be a (d_1, d_2)-isogeny diamond configuration. Then $F = \begin{pmatrix} f_1 & \tilde{f}_1' \\ -f_2 & f_2' \end{pmatrix}$ is d-isogeny $F : A \times B \to A_1 \times A_2$ with $d = d_1 + d_2$ and kernel* $\operatorname{Ker} F = \{(\tilde{f}_1(P), f_1'(P)) : P \in A_1[d]\}$.

Proof. This is Lemma 6 in [42], which generalizes Theorem 2.3 in [26]. □

6.1.3 Endomorphism-Testing Algorithm

The following gives more details for the algorithm described in [43, Section 4], which we will use repeatedly. This algorithm also appears as the Divide algorithm in [24, Section 4]. We will refer to this algorithm by our original name for it, the Endomorphism-Testing Algorithm.

Algorithm 6.1.9. Endomorphism-Testing Algorithm

Input: Elliptic curve E defined over $\mathbb{F}_{p^k}$; $\beta \in \operatorname{End}(E)$ which is written as a sum $\beta = b_1\beta_1 + b_2\beta_2 + b_3\beta_3 + b_4\beta_4$ where β_i are linearly independent endomorphisms which can be evaluated efficiently at powersmooth points of E and $b_i \in \mathbb{Z}$; n a positive integer; Q the norm form such that $Q(x_1, x_2, x_3, x_4) = \deg(\sum_{i=1}^4 x_i\beta_i)$

Output: TRUE if $\frac{\beta}{n}$ is an endomorphism of E and FALSE if $\frac{\beta}{n}$ is not an endomorphism.

1. Compute $\deg(\beta)$. If $n^2 \nmid \deg(\beta)$, conclude that $\frac{\beta}{n}$ is not an endomorphism and output FALSE. Otherwise, set $N := \deg(\beta)/n^2$.

2. Choose $a \in \mathbb{Z}$ such that $N' := N + a$ is powersmooth and $\gcd(N', n) = 1$.

3. Compute integers $a_1, a_2, a_3, a_4 \in \mathbb{Z}$ such that $a_1^2 + a_2^2 + a_3^2 + a_4^2 = a$. Let $\alpha \in \operatorname{End}(E^4)$ be the a-isogeny given by the matrix

$$\begin{pmatrix} a_1 & -a_2 & -a_3 & -a_4 \\ a_2 & a_1 & a_4 & -a_3 \\ a_3 & -a_4 & a_1 & a_2 \\ a_4 & a_3 & -a_2 & a_1 \end{pmatrix}.$$

4. Compute $K := \{(\frac{\widehat{\beta}}{n} \cdot \mathrm{Id}_4(P), \alpha(P)) : P \in E^4[N + a]\}$. Note that K can be computed even if $\frac{\beta}{n}$ is not an endomorphism: we can compute $\widehat{\beta}$ on $E[N + a]$, and by choice of a, n is invertible mod $N + a$.

5. Determine if $F : E^8 \to E^8/K$ is an endomorphism of principally polarized abelian varieties. (We do so by computing an appropriate theta structure for E^8/K and checking

that the projective theta constant of E^8 is the same as the projective theta constant of E^8/K.) If not, then terminate and conclude that $\frac{\beta}{n}$ is not an endomorphism.

6. Choose $M > \sqrt{\deg(\beta)} + \sqrt{n^2(N+a)}$ which is powersmooth. We check if $F_{ij}|_{E[M]} = \psi\frac{\beta}{n}|_{E[M]}$ for some $\psi \in \mathrm{Aut}(E)$, by evaluating the composition $E \xrightarrow{\iota_i} E^8 \xrightarrow{F} E^8 \xrightarrow{\pi_j} E$ on $E[M]$. If for some F_{ij} we have $F_{ij}|_{E[M]} = \psi\frac{\beta}{n}|_{E[M]}$, then we terminate and output TRUE. If no entry F_{ij} satisfies $F_{ij} = \psi\frac{\beta}{n}$, then terminate and output FALSE.

Proposition 6.1.10. *Algorithm 6.1.9 is correct and runs in time polynomial in* $\log(p^k)$ *and* $\log(\deg(\beta))$.

The proof of Proposition 6.1.10 follows from Lemmas 6.1.12, 6.1.15, and 6.1.16 below.

Lemma 6.1.11. *Let $\psi \in \mathrm{Aut}(E^n, \lambda)$, with λ the product polarization. Suppose ψ is given by its matrix form $M = (\psi_{i,j})_{1\le i,j\le n}$ as in Section 6.1.1. Then M has exactly one nonzero entry in each row and each column. Whenever ψ_{ij} is nonzero, ψ_{ij} is an automorphism of E.*

Proof. As ψ preserves the polarization λ on E^n, $\lambda = \psi^\vee\lambda\psi$. Therefore $\psi^\dagger\psi = 1$, with $\psi^\dagger$ the image of ψ under the Rosati involution. By [42, Lemma 3], the matrix form of $\psi^\dagger$ is given by the matrix $(\psi_{ij}^\dagger)_{i,j}$, with $\psi_{i,j}^\dagger$ the Rosati involution of $\psi_{i,j}$, which for elliptic curves equals the dual isogeny $\widehat{\psi_{i,j}}$. Call this matrix $\hat{M}$. Since $\psi^\dagger\psi = 1$, it follows that $\hat{M}M = \mathrm{Id}_n$.

Fix $1 \le i \le n$. We have $\sum_{k=1}^n \widehat{\psi_{ik}}\psi_{ik} = \sum_{k=1}^n \deg(\psi_{ik}) = 1$. As $\deg(\psi_{ik})$ is a positive integer whenever ψ_{ik} is nonzero, $\deg(\psi_{ik})$ is nonzero for exactly one k, and for this k, $\deg(\psi_{ik}) = 1$.

For $j \ne i$, we have $\sum_{k=1}^n \widehat{\psi_{ik}}\psi_{jk} = 0$. By the above argument, $\psi_{ik} = 0$ for all but one k. For this k, the fact that $\widehat{\psi_{ik}}\psi_{jk} = 0$ implies that $\psi_{jk} = 0$.

This shows that there is a unique nonzero entry in the i-th row, and that it is the only nonzero entry in its column. As there are n rows and n columns, this shows that there is a unique nonzero entry in each column, which is necessarily an automorphism. □

Lemma 6.1.12. *Let $\beta \in End(E)$ and n a positive integer. If $\frac{\beta}{n}$ is an endomorphism, then Algorithm 6.1.9 outputs True.*

Proof. Let $\phi = \frac{\beta}{n} \in \mathrm{End}(E)$. Then $\deg(\phi) = \frac{\deg(\beta)}{n^2} = N$. Since α is built out of scalar multiplications, we have the following commutative diagram, which is an (N, a)-isogeny diamond configuration.

$$\begin{array}{ccc} E^4 & \xrightarrow{\phi\cdot\mathrm{Id}_4} & E^4 \\ {\scriptstyle\alpha}\downarrow & & \downarrow{\scriptstyle\alpha} \\ E^4 & \xrightarrow{\phi\cdot\mathrm{Id}_4} & E^4 \end{array}$$

By Kani's Lemma, there is an (N+a)-endomorphism $G : (E^8, \lambda) \to (E^8, \lambda)$, where λ is the product polarization, such that G is given by the matrix $\begin{pmatrix} \phi \cdot \mathrm{Id}_4 & \alpha^\dagger \\ -\alpha & \widehat{\phi} \cdot \mathrm{Id}_4 \end{pmatrix}$. Moreover, as a was chosen such that $(N, a) = 1$, we can write $\ker(G) = \{\widehat{\phi} \cdot \mathrm{Id}_4(P), \alpha(P) : P \in E^4[N+a]\}$, which is the subgroup K constructed in Step 4.

If F is an isogeny with $\ker(F) = K$, then F is an $(N+a)$-endomorphism of principally polarized abelian varieties and the computed theta constants are equal. Therefore, we proceed to Step 6.

By [26, Proposition 1.1], there is an automorphism $\psi : E^8 \to E^8$ which preserves the product polarization and such that $F = \psi G$. By Lemma 6.1.11 each row and each column of the matrix form of ψ has exactly one nonzero entry, which is an automorphism of E. Thus, the entries of the matrix form of F are precisely the entries of the matrix form of G, composed with an automorphism of E. In particular, four of the nonzero entries of F will be given by $\psi_{ij}\phi$ for some automorphism $\psi_{ij} \in \mathrm{End}(E)$. □

Lemma 6.1.13. *The subgroup K in Step 4 of Algorithm 6.1.9 is a maximally isotropic subgroup of $E^8[N+a]$ (whether or not $\frac{\beta}{n}$ is an endomorphism). Thus, K is the kernel of an $(N+a)$-isogeny with respect to some polarization on E^8.*

Proof. Let K denote the subgroup in Step 4 of Algorithm 6.1.9, which is precisely the image of $F^\dagger = \begin{pmatrix} \frac{1}{n}\widehat{\beta} \cdot \mathrm{Id}_4 & -\alpha^\dagger \\ \alpha & \frac{1}{n}\beta \cdot \mathrm{Id}_4 \end{pmatrix}$ on $(E^4 \times E^4)[N+a]$

Let $m \in \mathbb{Z}$ such that $mn \equiv 1 \pmod{N+a}$. Consider the following isogeny factorization configuration:

$$\begin{array}{ccc} E^4 & \xrightarrow{m\beta \cdot \mathrm{Id}_4} & E^4 \\ {\scriptstyle mn\alpha}\downarrow & & \downarrow{\scriptstyle mn\alpha} \\ E^4 & \xrightarrow{m\beta \cdot \mathrm{Id}_4} & E^4 \end{array}$$

By Kani's Lemma, there is an $m^2n^2(N+a)$-endomorphism of E^8 with respect to the product polarization, given by $F' = \begin{pmatrix} m\beta \cdot \mathrm{Id}_4 & mn\alpha^\dagger \\ -mn\alpha & m\widehat{\beta} \cdot \mathrm{Id}_4 \end{pmatrix}$ and with kernel equal to the image of $F'^\dagger = \begin{pmatrix} m\widehat{\beta} \cdot \mathrm{Id}_4 & -mn\alpha^\dagger \\ mn\alpha & m\beta \cdot \mathrm{Id}_4 \end{pmatrix}$ on $(E^4 \times E^4)[m^2n^2(N+a)]$. Let $K' = F'^\dagger(E^4 \times E^4)[m^2n^2(N+a)]$. By Kani's Lemma, K' is a maximal isotropic subgroup of $E^8[m^2n^2(N+a)]$.

First, $K' \cap E^8[N+a]$ is a maximal isotropic subgroup of $E^8[N+a]$. Let $e_{m^2n^2(N+a)}$ be the Weil pairing on $E^8[m^2n^2(N+a)]$ and $P, Q \in E^8[N+a] \cap K'$. By compatibility of the Weil pairing, $1 = e_{m^2n^2(N+a)}(P, Q) = e_{N+a}(mnP, mnQ)$. By choice of m, we have

$e_{N+a}(mnP, mnQ) = e_{N+a}(P, Q)$. Thus, $K' \cap E^8[N+a]$ is an isotropic subgroup of $E^8[N+a]$. Since K' is a maximal isotropic subgroup of $E^8[m^2n^2(N+a)]$, and $(m^2n^2, N+a) = 1$, we have $K' \cap E^8[N+a]$ has order $(N+a)^8$ and is therefore a maximal isotropic subgroup of $E^8[N+a]$.

Finally, we have $K = K' \cap E^8[N+a]$. It is clear that $K \subset K' \cap E^8[N+a]$, since $F^\dagger = F'^\dagger$ on $E^8[N+a]$. Moreover, by the description of K as $\{(\frac{\widehat{\beta}}{n} \cdot \mathrm{Id}_4(P), \alpha(P)) : P \in E^4[N+a]\}$, where β and α have degrees coprime to $N+a$, it is clear that the order of $\#K = (N+a)^8 = \#(K' \cap E^8[N+a])$. Thus, K is a maximal isotropic subgroup of $E^8[N+a]$.

By [26, Proposition 1.1], K is therefore the kernel of an $N+a$-isogeny with respect to some polarization. □

The following lemma shows that an endomorphism is uniquely determined by its degree and its action on M-torsion, for suitably large M (depending on the degree).

Lemma 6.1.14. *Let E be an elliptic curve and $\phi, \psi \in End(E)$. Let $M > \sqrt{\deg(\phi)} + \sqrt{\deg(\psi)}$. If $\psi|_{E[M]} = \phi|_{E[M]}$, then $\psi = \phi$.*

Proof. For contradiction, assume the hypotheses of the lemma and that $\phi - \psi$ is nonzero. Since $\psi|_{E[M]} = \phi|_{E[M]}$, $E[M] \subset \ker(\phi - \psi)$. Since $\phi - \psi$ is nonzero, we must have $\phi - \psi = M\gamma$ for some nonzero $\gamma \in \mathrm{End}(E)$. Thus, $\deg(\phi - \psi) = M^2 \deg(\gamma)$. By the Cauchy-Schwartz inequality, $\deg(\phi - \psi) \leq (\sqrt{\deg(\phi)} + \sqrt{\deg(\psi)})^2$. Hence $M^2 \leq M^2 \deg(\gamma) \leq (\sqrt{\deg(\phi)} + \sqrt{\deg(\psi)})^2$, which is a contradiction. □

Lemma 6.1.15. *If $\frac{\beta}{n}$ is not an endomorphism, Algorithm 6.1.9 outputs False.*

Proof. Assume $F : E^8 \to E^8$ respects the product polarization and has kernel K as defined in Step 4. Let F_{ij} be an entry in the matrix form of F. Then $\deg(F_{ij}) \leq (N+a)$. If $F_{ij}|_{E[M]} = \frac{\psi\beta}{n}|_{E[M]}$ for some $M > \sqrt{\deg(\beta)} + \sqrt{n^2(N+a)}$ and an automorphism ψ, then $nF_{ij}|_{E[M]} = \psi\beta|_{E[M]}$. As we know $\psi\beta, nF_{ij}$ are endomorphisms, and $M > \sqrt{\deg(\beta)} + \sqrt{n^2(N+a)} > \sqrt{\deg(\psi\beta)} + \sqrt{n \deg(F_{ij})}$, Lemma 6.1.14 implies that $\frac{\beta}{n} = \psi^{-1}F_{ij} \in \mathrm{End}(E)$. □

Lemma 6.1.16. *Algorithm 6.1.9 runs in time polynomial in $\log(p^k)$ and $\log(\deg(\beta))$.*

Proof. Let B be a powersmoothness bound for $N+a$ (as in Step 2), and let C be a powersmoothness bound for M (as in Step 6). Given Q, computing the degree $\deg(\beta)$ amounts to evaluating Q at (b_1, b_2, b_3, b_4). The complexity of computing a_1, a_2, a_3, a_4 is $O((\log(a))^2(\log\log(a))^{-1})$, see [40, 41].

Computing a basis for K means first computing a basis for $E[N+a]$; decomposing into at most $\log(N+a)$ prime power parts, this can be done in $O(B^2 \log(p^k)^2 \log(N+a))$

operations [42, Lemma 7]. Evaluating $\widehat{\beta}$ on a basis for $E[N+a]$ and α on the induced basis for $E^4[N+a]$ can be done efficiently by our assumption on β and powersmoothness of $N+a$.

For Step 5, we need to check that F is truly an endomorphism. We place the additional data of a symmetric theta structure of level 2 on E^8, by taking an appropriate symplectic basis of $E[4]$ if $N+a$ is odd, or $E[2^{m+2}]$ where 2^m is the largest power of 2 dividing $N+a$ otherwise. (See Proposition C.2.6 of [19] and the preceding remark about how to choose a basis which is compatible with K in different cases.) Decomposing K into prime components, we can compute the theta null point of E^8/K with the induced theta structure in $O(\ell_{N+a}^8 \log(N+a))$ operations, where ℓ_{N+a} is the largest prime dividing $N+a$. (See Theorem C.2.2 and Theorem C.2.5 of [19].) Finally, as F may not preserve the product theta structure even if it is the desired endomorphism, we need to act on the theta null point by a polarization-preserving matrix in order to directly compare theta null points. When $N+a$ is odd, this matrix is computed explicitly [19, Proposition C.2.4] from the action of F on $E[4]$, which can also be evaluated in $O(\ell_{N+a}^8 \log(N+a))$ operations. This gives $O(B^8 \log(N+a))$ operations for this step.

In Step 6, computing a basis for the prime-power parts of $E[M]$ takes $O(C^2 \log(p^k)^2 \log(M))$ operations. If F is an endomorphism, then having already computed theta coordinates for E^8 and E^8/K in the previous step, we can evaluate F in terms of theta coordinates [19, Theorem C.2.2, Theorem C.2.5] and translate back to Weierstrass coordinates to check the equality. Note that there are only finitely many, and usually two, automorphisms to consider. Each evaluation costs $O(\ell_{N+a}^8 \log(N+a))$ operations where ℓ_{N+a} is the largest prime dividing $N+a$. There are 64 entries F_{ij} to check, by checking the equality on at most $2\log(M)$ points. Thus, this step requires at most $O(C^2 \log(p^k)^2 \log(M) + B^8 \log(N+a)\log(M))$ operations.

Now, we show that B and C can be taken polynomially sized in $\deg(\beta)$, that $N+a$ is $\tilde{O}(\deg(\beta))$, and that M is $\tilde{O}(1+n)\sqrt{\log(\deg(\beta))\deg(\beta)}$. Here, $\tilde{O}$ ignores logarithmic factors.

M and C are easier to analyze, as we have no restrictions on the primes which can divide M. When $k \geq 6$, we have that the k-th prime p_k satisfies $k\log(k) < p_k < k(\log(k)+\log\log(k))$ [44, Corollary of Theorem 3]. Therefore, we can take M to be a product of the first k primes where k is at most $\log(\sqrt{\deg(\beta)} + \sqrt{n^2(N+a)})$ and $C = \tilde{O}(\log(\sqrt{\deg(\beta)} + \sqrt{n^2(N+a)})$. Such a product is bounded by $\tilde{O}((\log(\sqrt{\deg(\beta)} + \sqrt{n^2(N+a)})(\sqrt{\deg(\beta)} + \sqrt{n^2(N+a)}))$.

We can bound $N+a$ and B similarly. However, $N+a$ is chosen to be coprime to Nn (equivalently, coprime to $\deg(\beta)$), so we instead take $N+a$ to be the product of the first at most $\log(N)$ primes which are coprime to Nn. Then we can take $B = \tilde{O}(\log(\deg(\beta)))$, noting that Nn has at most $\log(\deg(\beta))$ prime factors, so the largest prime we use is the k-th prime for $k \leq 2\log(\deg(\beta))$. The smallest such product which is larger than N is at most $\tilde{O}(\log(\deg(\beta))N)$. Thus, we have $N+a = \tilde{O}(\deg(\beta))$.

Returning to M and C, we get $\sqrt{\deg(\beta)}+\sqrt{n^2(N+a)} \leq \tilde{O}((1+n)\sqrt{\log(\deg(\beta))\deg(\beta)})$. Hence $M = \tilde{O}((1+n)\sqrt{\log(\deg(\beta))\deg\beta})$ and $C = \tilde{O}(\log((1+n)\sqrt{\log(\deg(\beta)\deg(\beta))}))$. $\square$

One can get speedups by replacing E^8 by E^4 and tweaking parameters as discussed by Robert in [42, Section 6]; for simplicity and for a proven complexity we don't go into those details here.

6.2 Computing the Distance From the Root

Let $\tilde{O}$ be a q-maximal q-enlargement of a suborder of $\mathrm{End}(E)$. In this section, we show how to compute the distance between $\mathrm{End}(E)\otimes\mathbb{Z}_q$ and $\tilde{O}\otimes\mathbb{Z}_q$, viewed as vertices on the Bruhat-Tits tree.

In the following proposition, we give a convenient expression for $\bigcap_{\Lambda\in N_r(M_2(\mathbb{Z}_q))}\Lambda$.

Proposition 6.2.1. *Fix an integer $r\geq 0$. Let $\tilde{\Lambda} = \cap_{\Lambda\in N_r(M_2(\mathbb{Z}_q))}\Lambda$. Then $\tilde{\Lambda} = \mathbb{Z}_q + q^r M_2(\mathbb{Z}_q)$.*

Proof. We first give an explicit basis for $\tilde{\Lambda}$. For any $\gamma_c \in \Sigma$ as in the notation of Definition 5.2.3, we have $d(\gamma_c^r, \mathrm{Id}) = r$ by Proposition 5.2.8. Hence $\gamma_c^{-r} M_2(\mathbb{Z}_q)\gamma_c^r \in N_r(M_2(\mathbb{Z}_q))$. It also follows that for $c\neq c'$, $d(\gamma_c^r, \gamma_{c'}^r) = 2r$. In particular,

$$d_3(\{\gamma_\infty^{-r} M_2(\mathbb{Z}_q)\gamma_\infty^r, \gamma_0^{-r} M_2(\mathbb{Z}_q)\gamma_0^r, \gamma_1^{-r} M_2(\mathbb{Z}_q)\gamma_1^r\}) = 6r.$$

By Corollary 5.3.6 and Lemma 5.3.3, we have $d_3(N_r(M_2(\mathbb{Z}_q))) = 6r$. It follows from Theorem 5.3.2 that

$$\tilde{\Lambda} = \gamma_\infty^{-r} M_2(\mathbb{Z}_q)\gamma_\infty^r \cap \gamma_0^{-r} M_2(\mathbb{Z}_q)\gamma_0^r \cap \gamma_1^{-r} M_2(\mathbb{Z}_q)\gamma_1^r.$$

A basis of $\gamma_\infty^{-r} M_2(\mathbb{Z}_q)\gamma_\infty^r \cap \gamma_0^{-r} M_2(\mathbb{Z}_q)\gamma_0^r$ is

$$\left\{ \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, q^r \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, q^r \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \right\}.$$

An element of $\gamma_1^{-r} M_2(\mathbb{Z}_q)\gamma_1^r$ can be written as

$$a \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + bq^r \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} + c \begin{pmatrix} -Aq^{-r} & q^{-r} \\ q^{-r} & 0 \end{pmatrix} + d \begin{pmatrix} 0 & -A \\ 0 & 1 \end{pmatrix}$$

where $A \equiv 1 \pmod{q^r}$. Hence an element is in the triple intersection if and only if $c \equiv 0$

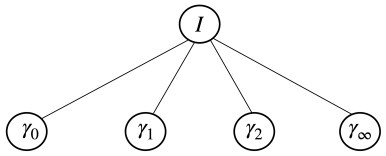

Figure 6.1. The maximal orders containing $\bigcap_{\Lambda' \subset N_1(M_2(\mathbb{Z}_q))} \Lambda'$ when $q = 3$.

$\pmod{q^{2r}}$ and $d \equiv 0 \pmod{q^r}$. In other words, a basis for the triple intersection is given by

$$\left\{ \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, q^r \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, q^r \begin{pmatrix} -A & 1 \\ 1 & 0 \end{pmatrix}, q^r \begin{pmatrix} 0 & -A \\ 0 & 1 \end{pmatrix} \right\},$$

which can be rewritten as

$$\left\{ \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, q^r \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, q^r \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, q^r \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \right\}.$$

□

We will use Proposition 6.2.1 to compute the distance between $\tilde{O}$ and $\mathrm{End}(E) \otimes \mathbb{Z}_q$ under an appropriate embedding into $M_2(\mathbb{Q}_q)$. Namely, if $\tilde{O} \otimes \mathbb{Z}_q$ is mapped to the root of the Bruhat-Tits tree, the distance will be the least r for which $\mathbb{Z}_q + q^r M_2(\mathbb{Z}_q) \subset \mathrm{End}(E) \otimes \mathbb{Z}_q$. We compute r by finding the least r for which $q^r \tilde{O} \subset \mathrm{End}(E)$. We show that such an r exists and is bounded in terms of the reduced discriminant of the input order.

Proposition 6.2.2. *Let $\mathcal{O}_0$ be an order, and let $\mathcal{O}$ be a q-enlargement of $\mathcal{O}_0$ for a prime q. If $k \geq v_q([\mathcal{O} : \mathcal{O}_0])$, then $q^k \mathcal{O} \subset \mathcal{O}_0$.*

Proof. For any prime q', the power of q' exactly dividing the global index $[\mathcal{O} : \mathcal{O}_0]$ is a generator for the local index at q', $[\mathcal{O} \otimes \mathbb{Z}_{(q')} : \mathcal{O}_0 \otimes \mathbb{Z}_{(q')}]$, by [54, Lemma 9.6.7]. As $\mathcal{O}$ is a q-enlargement of $\mathcal{O}_0$, we have $\mathcal{O} \otimes \mathbb{Z}_{q'} = \mathcal{O}_0 \otimes \mathbb{Z}_{q'}$, and thus $\mathcal{O}_0 \otimes \mathbb{Z}_{(q')} = \mathcal{O} \otimes \mathbb{Z}_{(q')}$ [54, Lemma 9.5.3]. Hence $[\mathcal{O} \otimes \mathbb{Z}_{(q')} : \mathcal{O}_0 \otimes \mathbb{Z}_{(q')}]$ is generated by a unit of $\mathbb{Z}_{(q')}$ whenever $q' \neq q$. Thus, the global index $[\mathcal{O} : \mathcal{O}_0]$ is a power of q. If $e = v_q([\mathcal{O} : \mathcal{O}_0])$, then $q^e \mathcal{O} \subset \mathcal{O}_0$. If $k \geq e$, then $q^k \mathcal{O} \subset q^e \mathcal{O} \subset \mathcal{O}_0$. □

Corollary 6.2.3. *Let $\mathcal{O}_0 \subset End(E)$ be an order, and let $\tilde{\mathcal{O}}$ be a q-maximal q-enlargement of $\mathcal{O}_0$ for $q \neq p$. Let $e = v_q(\mathrm{discrd}(\mathcal{O}_0))$. Then $q^e \tilde{\mathcal{O}} \subset End(E)$.*

Proof. We have $\mathrm{discrd}(\mathcal{O}_0) = [\tilde{\mathcal{O}} : \mathcal{O}_0] \, \mathrm{discrd}(\tilde{\mathcal{O}})$ by [54, Lemma 15.2.15]. As $\tilde{\mathcal{O}}$ is maximal at q, and $q \neq p$, we have $v_q(\mathrm{discrd}(\tilde{\mathcal{O}})) = 0$. Thus $e = v_q(\mathrm{discrd}(\mathcal{O}_0)) = v_q([\tilde{\mathcal{O}} : \mathcal{O}_0])$. It now

follows from Proposition 6.2.2 that $q^e\tilde{\mathcal{O}} \subset \mathcal{O}_0$, and since $\mathcal{O}_0 \subset \mathrm{End}(E)$ by hypothesis, we have $q^e\tilde{\mathcal{O}} \subset \mathrm{End}(E)$. □

The next proposition shows that the distance can be computed without reference to an embedding into $M_2(\mathbb{Q}_q)$.

Proposition 6.2.4. *Let $\mathcal{O}_0$ be a suborder of End(E) and let $\tilde{\mathcal{O}}$ be a q-maximal q-enlargement of $\mathcal{O}_0$. Let f be any isomorphism $f : \tilde{\mathcal{O}} \otimes \mathbb{Q}_q \to M_2(\mathbb{Q}_q)$ such that $f(\tilde{\mathcal{O}} \otimes \mathbb{Z}_q) = M_2(\mathbb{Z}_q)$. Let $\Lambda_E = f(End(E) \otimes \mathbb{Z}_q)$. Let r be the least integer such that $q^r\tilde{\mathcal{O}} \subset End(E)$. Then $r = d(M_2(\mathbb{Z}_q), \Lambda_E)$.*

Proof. Let f be any isomorphism satisfying the hypotheses of the proposition, and let $\Lambda_E = f(\mathrm{End}(E) \otimes \mathbb{Z}_q)$. Let k be a nonnegative integer, and let $\tilde{\Lambda} = \cap_{\Lambda \in N_k(M_2(\mathbb{Z}_q))}\Lambda$. By Corollary 5.3.6, $d(M_2(\mathbb{Z}_q), \Lambda_E) \leq k$ if and only if $\tilde{\Lambda} \subset \Lambda_E$. By Proposition 6.2.1, we have $\tilde{\Lambda} = \mathbb{Z}_q + q^k M_2(\mathbb{Z}_q)$. This shows that $d(M_2(\mathbb{Z}_q), \Lambda_E)$ is the least r for which $\mathbb{Z}_q + q^r M_2(\mathbb{Z}_q) \subset \Lambda_E$. We will now show that $\mathbb{Z}_q + q^k M_2(\mathbb{Z}_q) \subset \Lambda_E$ if and only if $q^k\tilde{\mathcal{O}} \subset \mathrm{End}(E)$.

As $\mathbb{Z}_q$ is contained in every $\mathbb{Z}_q$-order, we have $\mathbb{Z}_q + q^k M_2(\mathbb{Z}_q) \subset \Lambda_E$ if and only if $q^k M_2(\mathbb{Z}_q) \subset \Lambda_E$. By hypothesis, $q^k M_2(\mathbb{Z}_q) = f(q^k\tilde{O} \otimes \mathbb{Z}_q)$. It follows that $q^k M_2(\mathbb{Z}_q) \subset \Lambda_E$ if and only if $q^k\tilde{O} \subset \mathrm{End}(E) \otimes \mathbb{Z}_q$. Since $\tilde{O}$ is a q-maximal order, at primes $q' \neq q$ we have $q^k\tilde{\mathcal{O}} \otimes \mathbb{Z}_{q'} = \tilde{\mathcal{O}} \otimes \mathbb{Z}_{q'} = \mathcal{O}_0 \otimes \mathbb{Z}_{q'} \subset \mathrm{End}(E) \otimes \mathbb{Z}_{q'}$. It follows from the local-global principle that $q^k\tilde{\mathcal{O}} \subset \mathrm{End}(E)$ if and only if $q^k\tilde{\mathcal{O}} \otimes \mathbb{Z}_q \subset \mathrm{End}(E) \otimes \mathbb{Z}_q$. This shows that $d(M_2(\mathbb{Z}_q), \Lambda_E)$ is the least r for which $q^r\tilde{O} \subset \mathrm{End}(E)$. □

We now give an algorithm to compute the distance $d(M_2(\mathbb{Z}_q), \Lambda_E)$, where Λ_E is the image of $\mathrm{End}(E) \otimes \mathbb{Z}_q$ under any isomorphism satisfying the hypotheses of Proposition 6.2.4. We will construct both $\tilde{\mathcal{O}}$ and f in Section 6.3.

Algorithm 6.2.5. Computing the distance $d(M_2(\mathbb{Z}_q), \Lambda_E)$

Input: $E/\mathbb{F}_{p^2}$ supersingular; a prime $q \neq p$; a basis B for a q-maximal q-enlargement of $\mathcal{O}_0 \subset \mathrm{End}(E)$ with elements of B expressed as elements of $\mathcal{O}_0 \otimes \mathbb{Z}_{(q)}$; $e := v_q(\mathrm{discrd}(\mathcal{O}_0))$

Output: $r = d(M_2(\mathbb{Z}_q), \Lambda_E)$

1. Set $i := e - 1$.
2. While $0 \leq i \leq e - 1$:
   (a) Use Proposition 6.1.1 to determine if $q^i b \in \mathrm{End}(E)$ for each $b \in B$.
   (b) If for any $b \in B$, $q^i b \notin \mathrm{End}(E)$, output $i + 1$. Otherwise, set $i := i - 1$.
3. Output 0.

Proposition 6.2.6. *Algorithm 6.2.5 is correct and uses at most* $4(e+1)$ *applications of Proposition 6.1.1, with input* $n = q$ *and* $\beta \in End(E)$*. The run time is polynomial in* $\log(p^2)$ *and* $e \log(q) \max\{\log(\mathrm{Nrd}(b)) : b \in B\}$.

Proof. Let $\tilde{O}$ be the order generated by B. By Corollary 6.2.3, we have $q^e\tilde{\mathcal{O}} \subset \mathrm{End}(E)$, so the least r for which $q^r\tilde{\mathcal{O}} \subset \mathrm{End}(E)$ is at most e. This shows that the output is the least r for which $q^r\tilde{\mathcal{O}} \subset \mathrm{End}(E)$. By Proposition 6.2.4, this is the distance $d(M_2(\mathbb{Z}_q), \Lambda_E)$. In each iteration of the while loop, we apply Proposition 6.1.1 at most 4 times, and there are at most $e+1$ iterations of the while loop. At the i-th stage, we can express each candidate endomorphism as $\frac{q^{i+1}b}{q}$, where $q^{i+1}b \in \mathrm{End}(E)$ was verified in the $(i-1)$-th stage. We have $\deg(q^{i+1}b) = q^{2(i+1)}\,\mathrm{Nrd}(b)$, so the run-time follows from Proposition 6.1.1. □

6.3 Using global containment to test local containment

In this section, we show how to translate between computations in $\mathrm{End}(E) \otimes \mathbb{Q}$ and computations in the Bruhat-Tits tree. In the former, we have Proposition 6.1.1, and we would like to use this algorithm to deduce information about $\mathrm{End}(E) \otimes \mathbb{Z}_q$. Using work of Voight, we construct a q-maximal q-enlargement $\tilde{O}$ of our input order and an isomorphism $f : \tilde{O} \otimes \mathbb{Z}_q \to M_2(\mathbb{Z}_q)$. This maps the global order $\tilde{O}$ to the root of our Bruhat-Tits tree, so that $\tilde{O} \subset \mathrm{End}(E)$ if and only if $M_2(\mathbb{Z}_q) = f(\mathrm{End}(E) \otimes \mathbb{Z}_q)$. The main result of this section is Corollary 6.3.7, which shows that for any finite intersection of orders Λ in $M_2(\mathbb{Q}_q)$, we can construct a global order O such that $O \subset \mathrm{End}(E)$ if and only if $\Lambda \subset f(\mathrm{End}(E) \otimes \mathbb{Z}_q)$. This will allow us to test many candidates for $\mathrm{End}(E) \otimes \mathbb{Z}_q$ at once.

Our first task is to construct an explicit embedding of a given order $\mathcal{O}_0$ into $M_2(\mathbb{Z}_q)$. We first construct a q-maximal q-enlargement of $\mathcal{O}_0$.

Proposition 6.3.1. *Suppose an order* $\mathcal{O}_0 \subset End(E)$ *is given by a basis and a multiplication table, and let* q *be a prime. Then there is an algorithm which computes a* q*-maximal* q*-enlargement* $\tilde{O}$ *of* $\mathcal{O}_0$*. The run time is polynomial in the size of the basis and multiplication table. The basis elements which are output are of the form* $\frac{\beta}{q^k}$*, for an endomorphism* $\beta \in \mathcal{O}_0$ *and* $k \leq e = v_q(\mathrm{discrd}(\mathcal{O}_0))$*. Furthermore,* $\log(\deg(\beta))$ *is polynomial in the pairwise reduced traces of the basis elements of* $\mathcal{O}_0$*.*

Proof. On input $\mathcal{O}_0$, specified by the multiplication table and Q, we compute a q-maximal q-enlargement of $\mathcal{O}_0$, denoted $\tilde{\mathcal{O}}$ [53, Algorithms 3.12, 7.9, 7.10]. More specifically, Algorithm 3.12 produces a basis for $\mathcal{O}_0 \otimes \mathbb{Z}_q$ such that the norm form is normalized. Algorithm 7.9 gives a basis for a potentially larger "q-saturated" order, whose elements are of the form $\frac{x}{q^k}$. Here, x has coefficients in terms of the original basis at most $\max(\mathrm{Trd}(\beta_i\hat{\beta}_j))^4$, where β_i and β_j

range over basis elements of the original basis. The power k in the denominator is at most $\lfloor j/2 \rfloor$ where j is the valuation of the atomic form corresponding to the basis element, and hence $k \leq e = v_q(\mathrm{discrd}(\mathcal{O}_0))$.

Since $|\mathrm{Trd}(\beta_i \hat{\beta}_j)| \leq 2\sqrt{\deg(\beta_i)\deg(\beta_j)}$, the coefficients are polynomial in the original basis. Applying Algorithm 7.10 of [53] adjoins a zero divisor mod q, which is of the form $\frac{x}{q}$; here, x is expressed as linear combinations of the original basis with polynomially-sized coefficients. Thus, the basis which is output for $\tilde{\mathcal{O}}$ has coefficients which are polynomially-sized in the degrees of the original basis elements and q. Therefore, a basis element $\frac{\beta}{q^k}$ satisfies $\log(\deg(\beta))$ is at most polynomially-sized in $\log(\deg(\beta_i))$, where β_i ranges over the original basis elements, and $\log(q)$. □

Once we have obtained a q-maximal q-enlargment $\tilde{O}$, we compute an explicit isomorphism of $\tilde{O} \otimes \mathbb{Z}_q$ with the matrix ring $M_2(\mathbb{Z}_q)$, specified modulo q^{r+1}. In the context of our larger algorithm, r will be the distance computed by Algorithm 6.2.5.

Proposition 6.3.2. *Given a basis and multiplication table for a q-maximal order $\tilde{\mathcal{O}}$, and an integer r, there is an algorithm which computes a zero divisor $x \in \tilde{\mathcal{O}} \otimes \mathbb{Z}_q$ mod q^{r+1}. In other words, there is an algorithm to compute an element $x \in \tilde{\mathcal{O}} \otimes \mathbb{Z}_{(q)}$ such that there exists a zero divisor $x' \in \tilde{\mathcal{O}} \otimes \mathbb{Z}_q$ with $v_q(x - x') \geq r + 1$. The element x is expressed as a linear combination of the given basis such that coefficients are polynomially-sized in q^{r+1} and $\deg(\beta_i)\deg(\beta_j)$, where β_i and β_j range over elements of the given basis. The runtime is polynomial in $\log(q^{r+1})$ and the size of $\tilde{\mathcal{O}}$.*

Proof. First, use [53, Algorithm 3.12] on $\tilde{\mathcal{O}} \otimes \mathbb{Z}_q$ to obtain a normalized basis $\{f_1, f_2, f_3, f_4\}$ for $\tilde{\mathcal{O}} \otimes \mathbb{Z}_q$. By clearing denominators by units in $\mathbb{Z}_q$ if necessary, we can ensure $f_i \in \mathcal{O}_0 \otimes \mathbb{Z}_{(q)}$.

As $\tilde{\mathcal{O}}$ is q-maximal, the output basis being normalized means that the reduced norm form $Q(x_1, x_2, x_3, x_4) = \mathrm{Nrd}(\sum_{i=1}^4 x_i f_i)$ is given by a sum of atomic forms.

When q is odd, this means that $Q(x_1, x_2, x_3, x_4) = \sum_{i=1}^4 a_i x_i^2$ where $a_i \in (\mathbb{Z}_q)^\times$ and $\mathrm{Trd}(f_i \hat{f}_j) = 0$ when $i \neq j$. When $q = 2$, atomic forms are of one of the two following types: (i) ax^2 for $a \in (\mathbb{Z}_q)^\times$ or (ii) $a_i x_i^2 + a_{ij} x_i x_j + a_j x_i^2$ such that $v_2(a_{ij}) \leq v_2(a_i) \leq v_2(a_j)$ and $v_2(a_i) v_2(a_{ij}) = 0$. Up to reordering basis elements if necessary, we may therefore write $Q(x_1, x_2, x_3, x_4) = A_{12}(x_1, x_2) + A_{34}(x_3, x_4)$, where A_{ij} is either atomic of type (ii) or a sum of atomic forms of type (i).

We split up rest of the proof into the case that q is odd and $q = 2$: We first produce a nonzero element $x \in (\mathbb{Z}/q\mathbb{Z})^4$ such that $Q(x) \equiv 0 \pmod{q}$. Then, we show that there exists a lift x' in $\tilde{\mathcal{O}} \otimes \mathbb{Z}_q$, and we compute and output a lift of x in $\tilde{\mathcal{O}} \otimes \mathbb{Q}$ up to our desired precision q^r. In each case, the coefficients (in terms of the f_i) x_1, x_2, x_3, x_4 will be chosen mod q^r, so

the resulting output coefficients (in terms of the input basis) is polynomially-sized in q^r and $\deg(\beta_i)\deg(\beta_j)$.

Case 1: q is odd. In this case, the resulting reduced norm form is $Q(x_1,x_2,x_3,x_4) = \mathrm{Nrd}(\sum_{i=1}^4 x_i f_i) = \sum_{i=1}^4 a_i x_i^2$. The coefficients a_i may be rational, but $v_q(a_i) = 0$, so we may replace a_i by an integer mod q^r. Then there is a nonzero solution $(x_1,x_2,x_3) \in (\mathbb{F}_q)^3$ to the equation $\sum_{i=1}^3 a_i x_i^2 \equiv 0$, which can be found by a deterministic algorithm running in polynomial time in $\log(q)$ [52]. Reindexing the basis elements f_i and the corresponding a_i as necessary, we can assume $x_1 \neq 0$, so that the quadratic polynomial $Q_1(x) = Q(x,x_2,x_3,0)$ has a nonzero solution, x_1, mod q. Furthermore, $Q_1'(x_1) = 2a_1x_1$, which is nonzero mod q. Thus, by Hensel's Lemma, x can be lifted to a solution to $Q_1(x) = 0$ over $\mathbb{Z}_q$. A solution mod q^{r+1} can be recovered in (at most) r Hensel lifts, each running in polynomial time in $\log(q)$ (see [55, Algorithm 15.10 and Theorem 15.11] or [13, Theorem 3.5.3]).

Case 2: $q = 2$. In this case, the resulting reduced norm form is given by the normalized form $Q(x_1,x_2,x_3,x_4) = A_{1,2}(x_1,x_2) + A_{3,4}(x_3,x_4)$. Here $A_{i,j}(x_i,x_j) = a_i x_i^2 + a_{i,j}x_ix_j + a_j x_j^2$. The discriminant of Q, and therefore of $\tilde{\mathcal{O}} \otimes \mathbb{Z}_q$, is $(4a_1a_2 - a_{1,2}^2)(4a_3a_4 - a_{ij})^2$. As $\tilde{\mathcal{O}} \otimes \mathbb{Z}_q$ is 2-maximal, $a_{i,2}$ and $a_{3,4}$ are necessarily nonzero (mod 2).

Let $A(y,z)$ be an atomic form of type (ii), say $A(y,z) = ay^2 + byz + cz^2$ such that $v_2(b) \leq v_2(a) \leq v_2(c)$. Further assume $v_2(b) = 0$. We show that we can choose $y_0, z_0 \in \mathbb{Z}/2\mathbb{Z}$ such that $A(y_0,z_0) \equiv 1 \pmod 2$ and at least one of y_0 or z_0 is odd. If $v_2(a) \geq 1$ (and therefore $v_2(c) \geq 1$ as well), or if $v_2(a) = v_2(c) = 0$, we can set $y_0 \equiv z_0 \equiv 1 \pmod 2$. Otherwise, in the case that $v_2(a) = 0$ and $v_2(c) > 0$, we can set $y_0 \equiv 1 \pmod 2$ and $z_0 \equiv 0 \pmod 2$.

The quadratic form $Q(x_1,x_2,x_3,x_4)$ is the sum of two atomic quadratic forms $A_{1,2}$ and $A_{3,4}$ as above. We obtain a solution mod 2 by choosing x_1,x_2,x_3,x_4 mod 2 as just described. If x_1 and x_2 are both odd, i.e. in the case that a_1 and a_2 are of the same parity, we lift x_2,x_3,x_4 to $\mathbb{Z}/q^r\mathbb{Z}$ to obtain a quadratic polynomial $Q_1(x) = Q(x,x_2,x_3,x_4)$ with a solution mod 2 at $x \equiv 1 \pmod 2$. Then the derivative $Q_1'(1) = 2a_1 + a_{1,2}x_2$ is a unit in $\mathbb{Z}_2$. Otherwise, in the case that x_1 is odd and x_2 is even, we fix integers $x_1,x_3,x_4 \in \mathbb{Z}/q^r\mathbb{Z}$ to obtain a quadratic polynomial $Q_2(x) = Q(x_1,x,x_3,x_4)$ with a solution mod 2 at $x \equiv 0 \pmod 2$. Then the derivative $Q_2'(0) = a_{1,2}x_1$ is a unit in $\mathbb{Z}_2$. In either case, we obtain a solution to $Q = 0$ mod 2 which can be lifted to a solution in $\mathbb{Z}_q^4$ via Hensel's Lemma. As in the case that q is odd, a solution mod q^{r+1} can be recovered in r lifts, running in polynomial time in $\log(q)$. □

Finally, we construct the isomorphism.

Proposition 6.3.3. *Let $q \neq p$. Given a q-maximal order $\tilde{\mathcal{O}} \subset End(E) \otimes \mathbb{Q}$ and a nonnegative integer r, there is an algorithm which computes an isomorphism $f : \tilde{\mathcal{O}} \otimes \mathbb{Z}_q \to M_2(\mathbb{Z}_q)$ modulo q^{r+1}. This isomorphism is specified by giving the inverse image of standard basis elements i'*

and j' determined mod q^{r+1} in $\tilde{\mathcal{O}}$, such that

$$j' \mapsto \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

and

$$i' \mapsto \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \text{ if } q \neq 2, \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \text{ otherwise.}$$

The run time is polynomial in $\log(q^r)$ and the size of the basis and multiplication table for $\tilde{\mathcal{O}}$. In terms of the basis for $\tilde{\mathcal{O}}$, the representatives i' and j' are expressed with coefficients which are determined mod q^{r+1}

Proof. By Proposition 6.3.2, there is an algorithm to compute $x \in \tilde{\mathcal{O}} \otimes \mathbb{Z}_{(q)}$ such that $\mathrm{Nrd}(x) \equiv 0 \pmod{q^{r+1}}$. We first use x as input for [53, Algorithm 4.2]to compute nonzero $e \in \tilde{\mathcal{O}} \otimes \mathbb{Z}_q$ such that $e^2 = 0$. As in the proof of Proposition 6.3.2, we only specify e up to precision q^{r+1} and can therefore approximate e with an element of $\tilde{\mathcal{O}} \otimes \mathbb{Z}$. Furthermore, we can choose $e = \sum_{i=1}^{4} e_i f_i$ such that for some i, $q \nmid e_i$.

Then, on input e, we use [53, Algorithm 4.3] to compute i' and j' as a $\mathbb{Z}$-linear combination of $\frac{1}{s}e$ and $\frac{1}{s}f_i e$, for a basis element f_i such that $s = \mathrm{Trd}(f_i e)$ is nonzero.

In fact, we will modify the algorithm by choosing f_i such that $\mathrm{Trd}(f_i e)$ is nonzero mod q. If no such i exists, then $\mathrm{Trd}(ye) = 0$ for all $y \in \tilde{\mathcal{O}}$, so we show this cannot happen. Write $y = \sum_{j=1}^{4} y_j f_j$ and $e = \sum_{i=1}^{4} e_i f_i$, and consider the expression for $\mathrm{Trd}(ye) = -\mathrm{Trd}(y\bar{e})$ given by $\sum_{j=1}^{4}\sum_{i=1}^{4} -y_i e_j \mathrm{Trd}(f_i \hat{f}_j)$. As $\{f_1, f_2, f_3, f_4\}$ is a normalized basis, the equation simplifies in the following ways, depending on if q is even or odd.

If q is odd, then the expression simplifies to $\sum_{i=1}^{4} -e_i \mathrm{Trd}(f_i \hat{f}_i) y_i$. This is identically 0 mod q if and only if q divides $e_i \mathrm{Trd}(f_i \hat{f}_i)$ for all i. In the notation of the proof of Proposition 6.3.2, $\mathrm{Trd}(f_i \hat{f}_i)$ is exactly $2a_i$ and hence is not divisible by q by q-maximality. Hence, this expression is identically 0 mod q if and only if q divides e_i for all i, and we chose e such that this does not happen.

If $q = 2$, we have that $\mathrm{Trd}(f_i \hat{f}_i) = 2\,\mathrm{Nrd}(f_i) \equiv 0 \pmod{q}$ for all i, so the only nonzero terms are $-e_1 \mathrm{Trd}(f_2 \hat{f}_1), -e_2 \mathrm{Trd}(f_1 \hat{f}_2), -e_3 \mathrm{Trd}(f_4 \hat{f}_3), -e_4 \mathrm{Trd}(f_3 \hat{f}_4)$. We have $\mathrm{Trd}(f_1 \hat{f}_2) = Trd(f_2 \hat{f}_1) = a_{1,2}$ and $\mathrm{Trd}(f_3 \hat{f}_4) = \mathrm{Trd}(f_4 \hat{f}_3) = a_{3,4}$, which are not divisible by q as we showed in the proof of Proposition 6.3.2. Hence this expression is identically 0 mod q if and only if q divides e_i for all i, but we chose e such that this does not happen.

This shows that $v_q(\mathrm{Trd}(ef_i)) = 0$ for some i, so that $\frac{1}{s} \in \mathbb{Z}_q$, and the elements i' and j' output by Algorithm 4.3 of [53] (with this modification) are elements of $\tilde{\mathcal{O}} \otimes \mathbb{Z}_q$ and furnish an isomorphism of $\tilde{\mathcal{O}} \otimes \mathbb{Z}_q \to M_2(\mathbb{Z}_q)$. To get i' and j' in $\tilde{\mathcal{O}}$ rather than in $\tilde{\mathcal{O}} \otimes \mathbb{Q}$, replace $\frac{1}{s}$

by an integer $m \equiv s^{-1} \pmod{q^r}$. $\square$

The other maximal orders that we will work with in the Bruhat-Tits tree are of the form $T^{-1}M_2(\mathbb{Z}_q)T$, where T is a matrix associated to a matrix path, as described in Definition 5.2.3. The next lemma shows that T can be replaced with T' such that $T \equiv T' \pmod{q^{r+1}}$, where $r > v_q(\det(T))$.

Lemma 6.3.4. *Let* $T = \begin{pmatrix} q^a & c \\ 0 & q^b \end{pmatrix}$, *with* $a, b \geq 0$, $c \in \mathbb{Z}/q^b\mathbb{Z}$, *and* $v_q(c) = 0$ *if both* a *and* b *are positive. Let* $T' = T + q^e M$ *where* $M \in M_2(\mathbb{Z}_q)$ *and* $e > a + b$. *Then there is* $x \in \mathbb{Q}_q^*$ *and* $C \in \mathrm{GL}_2(\mathbb{Z}_q)$ *such that* $T = xCT'$. *In particular,* $T^{-1}M_2(\mathbb{Z}_q)T = T'^{-1}M_2(\mathbb{Z}_q)T'$.

Proof. We write $T' = \begin{pmatrix} q^a + dq^e & c + fq^e \\ gq^e & q^b + hq^e \end{pmatrix}$ with $d, f, g, h \in \mathbb{Z}_q$.

Multiplying on the left by $\begin{pmatrix} (1 + dq^{e-a})^{-1} & 0 \\ -gq^{e-a}(1 + dq^{e-a})^{-1} & 1 \end{pmatrix} \in \mathrm{GL}_2(\mathbb{Z}_q)$ gives $\begin{pmatrix} q^a & \alpha \\ 0 & q^b\beta \end{pmatrix}$, where $\alpha = (c + fq^e)(1 + dq^{e-a})^{-1}$ and $\beta = 1 + hq^{e-a} - gq^{e-a-b}\alpha$.

For any $e > a + b$, we have $v_q(\beta) = 0$ and $v_q(\alpha) \geq 0$. We also have $c = \alpha + q^b(dq^{e-a-b}\alpha - fq^{e-b})$. Multiply on the left by $\begin{pmatrix} 1 & \beta^{-1}(dq^{e-a-b}\alpha - fq^{e-b}) \\ 0 & \beta^{-1} \end{pmatrix} \in \mathrm{GL}_2(\mathbb{Z}_q)$ to arrive at T. All operations are invertible over $\mathbb{Z}_q$ provided $e > a + b$. $\square$

To check that a finite intersection of local maximal orders is contained in $\mathrm{End}(E) \otimes \mathbb{Z}_q$, we will check that the intersection of related maximal orders is contained in $\mathrm{End}(E)$. The following lemma allows us to compute the basis of an intersection of $\mathbb{Z}$-orders in polynomial time.

Lemma 6.3.5. *Let* $L_1, L_2 \subseteq \mathbb{Z}^4$ *be two* $\mathbb{Z}$*-lattices of full rank, specified by a possibly dependent set of generators. A basis for each lattice, a lattice basis for the sum* $L_1 + L_2$ *and for the intersection* $L_1 \cap L_2$ *can be computed in polynomial time.*

Proof. A basis for each lattice from a set of m generators can be computed by writing the generators as the columns of a matrix and the compute the Hermite Normal Form (HNF) of the matrix (see [34, p. 149] for the definition). The HNF of this $4 \times m$ matrix can be computed in time polynomial in m, and the bit length of the matrix entries, see [**?**, **?**]. Computing a basis for the sum of two lattices immediately reduces to the problem of computing a basis of a lattice from a set of generators. To compute the intersection of two lattices L_1, L_2 each specified by a 4×4 basis in matrix form B_1, B_2, we first compute $(B_1^T)^{-1}$ and $(B_2^T)^{-1}$. Here B_i^T denotes the transpose of B_i. The matrix $(B_i^T)^{-1}$ is a basis for the dual $\hat{L}_i$ of L_i [34, p. 19]. By Cramer's rule, the inverse of this 4×4 matrix can be computed efficiently. Since the dual

of a lattice L consists of all vectors y in $L \otimes \mathbb{R}$ whose real inner product $\langle x, y \rangle$ is an integer for every $x \in L$, it follows easily that the dual of $L_1 \cap L_2$ is $\hat{L}_1 + \hat{L}_2$, i.e. the smallest lattice containing both $\hat{L_1}$ and $\hat{L_2}$. So a basis for the intersection is obtained by computing a basis for the lattice $\tilde{L} := \hat{L}_1$ and $\hat{L}_2$ and then computing the dual of $\tilde{L}$. By the above argument, this can be computed in polynomial time. □

Corollary 6.3.6. *Let $\mathcal{O}(i), i = 1, \ldots, 3$ be orders in $B_{p,\infty}$ such that $(\mathrm{disc}(\mathcal{O}_0))\mathcal{O}_0 \subseteq \mathcal{O}(i)$. A basis for $\cap\mathcal{O}(i)$ in which each basis vector is given as a $\mathbb{Q}$-linear combination of the basis vectors for $\mathcal{O}_0$ can be computed in polynomial time.*

Proof. We can reduce this to matrix computations with 4×4 integer matrices and use the previous lemma. We identify our starting global order $\mathcal{O}_0$ with $\mathbb{Z}^4$, whose Hermite Normal Form is just the 4×4 identity matrix. Since $(\mathrm{disc}(\mathcal{O}_0))\mathcal{O}_0 \subseteq \mathcal{O}(i)$, we can scale the matrices representing the orders $\mathcal{O}(i)$ by $\mathrm{disc}\,\mathcal{O}_0$ and work with integer matrices. See [13, page 73f.] for generalizing the computation of the HNF to matrices with bounded rational coefficients. By the previous lemma, $\cap\mathcal{O}(i)$ can be computed in polynomial time. □

The following is the main result of this section, which shows that we can check local containment by checking global containment.

Corollary 6.3.7. *Let $\mathcal{O}_0 \subset End(E)$. Let Λ be a finite intersection of maximal orders of $M_2(\mathbb{Q}_q)$, and let $r \leq v_q(\mathrm{discrd}(\mathcal{O}_0))$ be an integer such that $\Lambda \subset \cap_{\Lambda' \in N_r(M_2(\mathbb{Z}_q))}\Lambda'$. Let $\tilde{\mathcal{O}}$ be a q-maximal q-enlargement of $\mathcal{O}_0$ and let f be the isomorphism computed in Proposition 6.3.3. Then there exists a global order $\mathcal{O}$ such that $f(\mathcal{O} \otimes \mathbb{Z}_q) = \Lambda$ and $\mathcal{O} \subset End(E)$ if and only if $\Lambda \subset f(End(E) \otimes \mathbb{Z}_q)$. The basis of $\mathcal{O}$ can be computed in polynomial time in the size of $\mathcal{O}_0$ and $\log(q)$, and the basis elements of $\mathcal{O}$ have degree polynomial in the size of $\mathcal{O}_0$ and q^r.*

Proof. We will show that $\mathcal{O}$ can be computed such that $f(\mathcal{O} \otimes \mathbb{Z}_q) = \Lambda$ and $\mathcal{O} \otimes \mathbb{Z}_{q'} \subset \mathcal{O}_0 \otimes \mathbb{Z}_{q'}$ for all primes $q' \neq q$.

Suppose that Λ is maximal. Write $\Lambda = T^{-1}M_2(\mathbb{Z}_q)T$ where T is the matrix associated to a matrix path of length at most r, written $T = \begin{pmatrix} q^a & c \\ 0 & q^b \end{pmatrix}$ where $a + b \leq r$.

We construct an element $t \in \tilde{\mathcal{O}}$ such that $f(t) \equiv T \pmod{q^{r+1}}$. Let $i', j' \in \tilde{\mathcal{O}}$ denote the inverse image of the standard basis elements of $M_2(\mathbb{Z}_q)$ modulo q^{r+1}, as in Proposition 6.3.3. If q is odd, then with $f(i') \equiv \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \pmod{q^{r+1}}$ and $f(j') \equiv \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \pmod{q^{r+1}}$, we can take

$$t = \frac{q^a + q^b}{2} + \frac{q^a - q^b}{2}i' + \frac{c}{2}j' + \frac{c}{2}i'j'.$$

As written, t is an element of $\tilde{\mathcal{O}} \otimes \mathbb{Q}$, but we can replace division by 2 with multiplication by an integer $m \equiv 2^{-1} \pmod{q^{r+1}}$ to ensure $t \in \tilde{\mathcal{O}}$ and $f(t) \equiv T \pmod{q^{r+1}}$. If $q = 2$, then with $f(i') \equiv \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix} \pmod{q^{r+1}}$ and $f(j') \equiv \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \pmod{q^{r+1}}$, we can take

$$t = (q^a + c) + (q^b - q^a)i' + (c - q^b + q^a)j' + (-c)i'j'.$$

Let $\mathcal{O} = \frac{1}{q^{a+b}}\hat{t}\tilde{\mathcal{O}}t$, where $\hat{}$ denotes the dual isogeny. At $q' \neq q$, we have $\mathcal{O} \otimes \mathbb{Z}_{q'} \subset \mathcal{O}_0 \otimes \mathbb{Z}_{q'} \subset \mathrm{End}(E) \otimes \mathbb{Z}_{q'}$. This is because $\tilde{\mathcal{O}} \otimes \mathbb{Z}_{q'} = \mathcal{O}_0 \otimes \mathbb{Z}_{q'}$, and $\hat{t}\tilde{\mathcal{O}}t \subset \tilde{\mathcal{O}}$. At q, we have $f(\mathcal{O} \otimes \mathbb{Z}_q) = T^{-1}M_2(\mathbb{Z}_q)T \supset f(\mathcal{O}_0 \otimes \mathbb{Z}_q)$. By the local-global principle, we get $f(\mathcal{O} \otimes \mathbb{Z}_q) \subset f(\mathrm{End}(E) \otimes \mathbb{Z}_q)$ if and only if $\mathcal{O} \subset \mathrm{End}(E)$.

In the general case, let $\Lambda_1, \Lambda_2, \Lambda_3$ be maximal orders in $M_2(\mathbb{Q}_q)$ such that $\Lambda = \cap_{i=1}^3 \Lambda_i$. We can choose three orders Λ_i for which this is true by Theorem 5.3.2. Let $\mathcal{O}(i)$ denote a global order such that $\mathcal{O}(i) \otimes \mathbb{Z}_{q'} \subset \mathcal{O}_0 \otimes \mathbb{Z}_{q'}$ for all $q' \neq q$ and $f(\mathcal{O}(i) \otimes \mathbb{Z}_q) = \Lambda_i$, as constructed in the previous paragraph.

Let $\mathcal{O} = \cap_{i=1}^3 \mathcal{O}(i)$. By construction, $q^{a+b}\mathcal{O}(i) \subset \tilde{\mathcal{O}}$, and by Corollary 6.2.3, we have $q^{v_q(\mathrm{discrd}(\mathcal{O}_0))}\tilde{\mathcal{O}} \subset \mathcal{O}_0$. As $a + b \leq r \leq v_q(\mathrm{discrd}(\mathcal{O}_0))$, we have $q^{2v_q(\mathrm{discrd}(\mathcal{O}_0))}\mathcal{O}(i) \subset \mathcal{O}_0$, and thus $\mathrm{disc}(\mathcal{O}_0)\mathcal{O}(i) \subset \mathcal{O}_0$. By Corollary 6.3.6, the basis of $\mathcal{O}$ can be computed in polynomial time.

Tensoring by $\mathbb{Z}_{q'}$ for any prime q' commutes with taking intersections, as $\mathbb{Z}_{q'}$ is a flat $\mathbb{Z}$-module. Hence $\mathcal{O} \otimes \mathbb{Z}_{q'} = \cap_{i=1}^3 \mathcal{O}(i) \otimes \mathbb{Z}_{q'} \subset \mathcal{O}_0 \otimes \mathbb{Z}_{q'}$ for all $q' \neq q$, and $f(\mathcal{O} \otimes \mathbb{Z}_q) = \cap_{i=1}^3 f(\mathcal{O}(i) \otimes \mathbb{Z}_q) = \cap_{i=1}^3 \Lambda_i = \Lambda$.

We have that $\mathcal{O} \otimes \mathbb{Z}_{q'} \subset \mathcal{O}_0 \otimes \mathbb{Z}_{q'} \subset \mathrm{End}(E) \otimes \mathbb{Z}_{q'}$ for all $q' \neq q$, so $\mathcal{O} \subset \mathrm{End}(E)$ if and only if $\mathcal{O} \otimes \mathbb{Z}_q \subset \mathrm{End}(E) \otimes \mathbb{Z}_q$. As f is an isomorphism, this is equivalent to $\Lambda \subset f(\mathrm{End}(E) \otimes \mathbb{Z}_q)$, as desired. □

6.4 Finding Λ_E in the Bruhat-Tits Tree

Let $\Lambda_E = f(\mathrm{End}(E) \otimes \mathbb{Z}_q)$ and $r = d(\Lambda_E, M_2(\mathbb{Z}_q))$. Once we have computed r, we know that Λ_E is of the form $T^{-1}M_2(\mathbb{Z}_q)T$, where T is a matrix associated to a matrix path of length r. In this section, we show how to recover the matrix path one step at a time, which will allow us to compute Λ_E.

For each matrix path $\{c_i\}_{i=1}^k$ of length at most r, the next two propositions define an order which is contained in Λ_E if and only if the corresponding path begins with $\{c_i\}_{i=1}^k$.

Proposition 6.4.1. *Let $q > 2$. Let γ be the matrix associated to a matrix path $\{c_i\}_{i=1}^k$, where*

$1 \leq k \leq r$. *Let*

$$S := N_r(M_2(\mathbb{Z}_q)) \cap \{\Lambda' : \Lambda' \text{ maximal corresponding to a matrix path starting with } \{c_i\}_{i=1}^k\}.$$

Let $\Lambda_{\gamma,r} := \bigcap_{\Lambda \in S} \Lambda$ *and* $P = \{\gamma^{-1} M_2(\mathbb{Z}_q)\gamma\}$. *Then:*

1. $\Lambda_{\gamma,r} = \bigcap_{\Lambda \in N_{r-k}(P)} \Lambda$.
2. $v_q(\mathrm{discrd}(\Lambda_{\gamma,r})) = 3r - 3k$.
3. *Suppose* $\Lambda_{\gamma,r} \subset \Lambda$ *for a maximal order* $\Lambda \in N_r(M_2(\mathbb{Z}_q))$. *Then* Λ *corresponds to a matrix path starting with* $\{c_i\}_{i=1}^k$.

Proof. Any order in S corresponds to a matrix path $\{c_1, \ldots, c_k, b_1, b_2, \ldots, b_{r-k}\}$. Letting $\delta = b_{r-k} b_{r-k-1} \cdots b_1$, the associated matrix is $\delta\gamma$. By Proposition 5.2.8, we have $d(\gamma, \delta\gamma) = r - k$. This shows that $S \subset N_{r-k}(P)$. To show $\Lambda_{\gamma,r} = \bigcap_{\Lambda \in N_{r-k}(P)} \Lambda$, we will choose $\Lambda_1 \Lambda_2, \Lambda_3$ in S which maximize d_3 among the orders of $N_{r-k}(P)$ and therefore in S. Choose three distinct matrices $\delta_i \in \Sigma$ such that $\delta_i c_k \notin M_2(q\mathbb{Z}_q)$. There are $q \geq 3$ possibilities for δ_i. Let $\Lambda_i = (\delta_i^{r-k}\gamma)^{-1} M_2(\mathbb{Z}_q) \delta_i^{r-k}\gamma$. If $i \neq j$, we have $d(\delta_i^{r-k}\gamma, \delta_j^{r-k}\gamma) = 2r - 2k$ by Proposition 5.2.8.

As $\delta_i c_k \notin M_2(q\mathbb{Z}_q)$, the matrices $\delta_i^{r-k}\gamma$ correspond to matrix paths of length r starting with $\{c_i\}_{i=1}^k$, so $\Lambda_i \in S$. Therefore, $d_3(\{\Lambda_1, \Lambda_2, \Lambda_3\}) = 6r - 6k \leq d_3(S)$.

By Corollary 5.3.6, we have $d_3(N_{r-k}(P)) = 6r - 6k$, so $d_3(\Lambda_1, \Lambda_2, \Lambda_3) = d_3(N_{r-k}(P)) = d_3(S)$. By Theorem 5.3.2, this shows $\Lambda_{\gamma,r} = \Lambda_1 \cap \Lambda_2 \cap \Lambda_3 = \cap_{\Lambda \subset N_{r-k}(P)} \Lambda$, which gives (1) and (2).

It follows from (1) and Corollary 5.3.6 that the set of maximal orders containing $\Lambda_{\gamma,r}$ is $N_{r-k}(P)$. To show (3), we must show that $N_r(M_2(\mathbb{Z}_q)) \cap N_{r-k}(\gamma^{-1} M_2(\mathbb{Z}_q)\gamma)$ consists of those orders corresponding to a matrix path of length r starting with $\{c_i\}_{i=1}^k$.

Consider a matrix path $\{g_i\}_{i=1}^r$ and the associated matrix $g = g_r g_{r-1} \cdots g_1$. The order $g^{-1} M_2(\mathbb{Z}_q) g$ is in the intersection $N_r(M_2(\mathbb{Z}_q)) \cap N_{r-k}(\gamma^{-1} M_2(\mathbb{Z}_q)\gamma)$ if and only if $d(g, \gamma) \leq r - k$. But by Lemma 5.2.8, $d(g, \gamma) \leq r - k$ if and only if the sequences agree in the first k matrices, which are exactly $\{c_i\}_{i=1}^k$. This shows (3). □

Proposition 6.4.2. *Let* $q = 2$. *Let* γ *be the matrix associated to a matrix path* $\{c_i\}_{i=1}^k$, *where* $1 \leq k \leq r$. *Let*

$$S := N_r(M_2(\mathbb{Z}_q)) \cap \{\Lambda' : \Lambda' \text{ maximal corresponding to a matrix path starting with } \{c_i\}_{i=1}^k\}.$$

Let $\Lambda_{\gamma,r} := \cap_{\Lambda \in S} \Lambda$ *and* $\delta_1, \delta_2 \in \Sigma$ *be such that* $\delta_1 c_k \notin M_2(q\mathbb{Z}_q)$ *and* $\delta_2 c_k \notin M_2(q\mathbb{Z}_q)$. *Let* $P := \{(\delta_1\gamma)^{-1} M_2(\mathbb{Z}_q)\delta_1\gamma, \gamma^{-1} M_2(\mathbb{Z}_q)\gamma, (\delta_2\gamma)^{-1} M_2(\mathbb{Z}_q)\delta_2\gamma\}$. *Then:*

1. $\Lambda_{\gamma,r} = \bigcap_{\Lambda \in N_{r-k-1}(P)} \Lambda$.

2. $v_q(\mathrm{discrd}(\Lambda_{\gamma,r})) = 3r - 3k - 1$.

3. *Suppose* $\Lambda_\gamma \subset \Lambda$ *for a maximal order* Λ *such that* $d(\Lambda, M_2(\mathbb{Z}_q)) = r$. *Then* Λ *corresponds to a matrix path starting with* $\{c_i\}_{i=1}^k$.

Proof. Any order in S corresponds to a matrix path $\{c_1, \ldots, c_k, \delta_i, b_1, b_2, \ldots, b_{r-k-1}\}$ for one of $i = 1, 2$. Letting $\delta = b_{r-k-1} b_{r-k} \cdots b_1$, the associated matrix is $\delta\delta_i\gamma$. By Proposition 5.2.8, $d(\delta_i\gamma, \delta\delta_i\gamma) = r - k - 1$. This shows $S \subset N_{r-k}(P)$. To show $\Lambda_{\gamma,r} = \bigcap_{\Lambda \in N_{r-k}(P)} \Lambda$, we will choose $\Lambda_1\Lambda_2, \Lambda_3$ in S which maximize d_3 among the orders of $N_{r-k-1}(P)$ and hence in S.

Let $\Lambda_1 = (\delta_1^{r-k}\gamma)^{-1} M_2(\mathbb{Z}_q) \delta_1^{r-k}\gamma$, and $\Lambda_2 = (\delta_2^{r-k}\gamma)^{-1} M_2(\mathbb{Z}_q)(\delta_2^{r-k}\gamma)$. Choose $\delta_3 \in \Sigma$ such that $\delta_3\delta_1 \notin M_2(q\mathbb{Z}_q)$ and $\delta_3 \neq \delta_1$. Let $\Lambda_3 = (\delta_3^{r-k-1}\delta_1\gamma)^{-1} M_2(\mathbb{Z}_q)(\delta_3^{r-k-1}\delta_1\gamma)$. It is easy to see that $\Lambda_i \in S$ for each i, as the corresponding matrix paths have length r and begin with $\{c_i\}_{=1}^k$. By Proposition 5.2.8, we compute $d_3(\{\Lambda_1, \Lambda_2, \Lambda_3\}) = 2(r-k) + 2(r-k) + 2(r-k-1) = 6r - 6k - 2$.

By Corollary 5.3.6, $d_3(N_{r-k-1}(P)) = 6r - 6k - 2$, so $d_3(\Lambda_1, \Lambda_2, \Lambda_3) = d_3(N_{r-k-1}(P)) = d_3(S)$. By Theorem 5.3.2, this shows $\Lambda_{\gamma,r} = \Lambda_1 \cap \Lambda_2 \cap \Lambda_3 = \cap_{\Lambda \subset N_{r-k}(P)} \Lambda$, which gives (1) and (2).

It follows from (1) and Corollary 5.3.6 that the set of maximal orders containing $\Lambda_{\gamma,r}$ is $N_{r-k-1}(P)$. To show (3), we must show that $N_r(M_2(\mathbb{Z}_q)) \cap N_{r-k-1}(P)$ consists of those orders corresponding to a matrix path of length r starting with $\{c_i\}_{i=1}^k$.

Consider a matrix path $\{g_i\}_{i=1}^r$ and the associated matrix $g = g_r g_{r-1} \cdots g_1$. The order $g^{-1} M_2(\mathbb{Z}_q) g$ is in the intersection $N_r(M_2(\mathbb{Z}_q)) \cap N_{r-k-1}(P)$ if and only if one of $d(g, \delta_i\gamma) \leq r - k - 1$ or $d(g, \gamma) \leq r - k - 1$. Since γ corresponds to a matrix path of length k and g corresponds to a matrix path of length r, we have $d(g, \gamma) \geq r - k$, so the condition $d(g, \gamma) \leq r - k - 1$ is impossible in this setup. We also have $d(g, \delta_i\gamma) \leq r - k - 1$ if and only if $g_i = c_i$ for all $1 \leq i \leq k$ and $c_{k+1} = \delta_i$. Since δ_1 and δ_2 are the only two choices for the $(k+1)$-th entry in a matrix path starting with $\{c_i\}_{i=1}^k$, we have $d(g, P) \leq r - k - 1$ if and only if $g_i = c_i$ for all $1 \leq i \leq k$. This shows (3). □

We obtain the following algorithm to recover the matrix path $\{d_i\}_{i=1}^r$ corresponding to Λ_E. For each γ_c as in Definition 5.2.3, we check if $\Lambda_{\gamma_c,r} \subset \Lambda_E$. By Propositions 6.4.1 and 6.4.2, we have $\Lambda_{\gamma_c,r} \subset \Lambda_E$ if and only if $d_1 = \gamma_c$, so we recover the first matrix in the path in $q + 1$ checks. Once we have recovered the first $k - 1$ matrices, we test each of the q possibilities for d_k, continuing until we have recovered the full matrix path.

Algorithm 6.4.3. Computing the path from $M_2(\mathbb{Z}_q)$ to Λ_E

Input: An order $\mathcal{O}_0 \subset \mathrm{End}(E)$; a prime $q \neq p$; q-maximal q-enlargement $\tilde{\mathcal{O}}$ of $\mathcal{O}_0$; an isomorphism $f : \mathcal{O}_0 \otimes \mathbb{Q}_q \to M_2(\mathbb{Q}_q)$ such that $f(\tilde{\mathcal{O}} \otimes \mathbb{Z}_q) = M_2(\mathbb{Z}_q)$ computed mod q^{r+1}; $r := d(M_2(\mathbb{Z}_q), \Lambda_E)$

Output: γ such that $\Lambda_E = \gamma^{-1} M_2(\mathbb{Z}_q)\gamma$

1. Set $k := 1$, $\gamma' := \mathrm{Id}$, $d_0 := \mathrm{Id}$.
2. While $k \leq r$:
   (a) For each $\gamma_c \in \Sigma$ as in Definition 5.2.3 such that $\gamma_c d_{k-1} \notin qM_2(\mathbb{Z}_q)$:
      i. Set $\gamma := \gamma_c \gamma'$.
      ii. Compute a basis B_γ for an order $\mathcal{O}_\gamma$ such that $\mathcal{O}_\gamma \subset \mathrm{End}(E)$ if and only if $\Lambda_{\gamma,r} \subset \Lambda_E$. (See Corollary 6.3.7.)
      iii. Apply Proposition 6.1.1 to each $b \in B_\gamma$ to decide if $b \in \mathrm{End}(E)$.
      iv. If for all $b \in B_\gamma$, we have $b \in \mathrm{End}(E)$: Set $\gamma' := \gamma$, $d_k := \gamma_c$, $k := k + 1$, and return to Step 2.
3. Output γ.

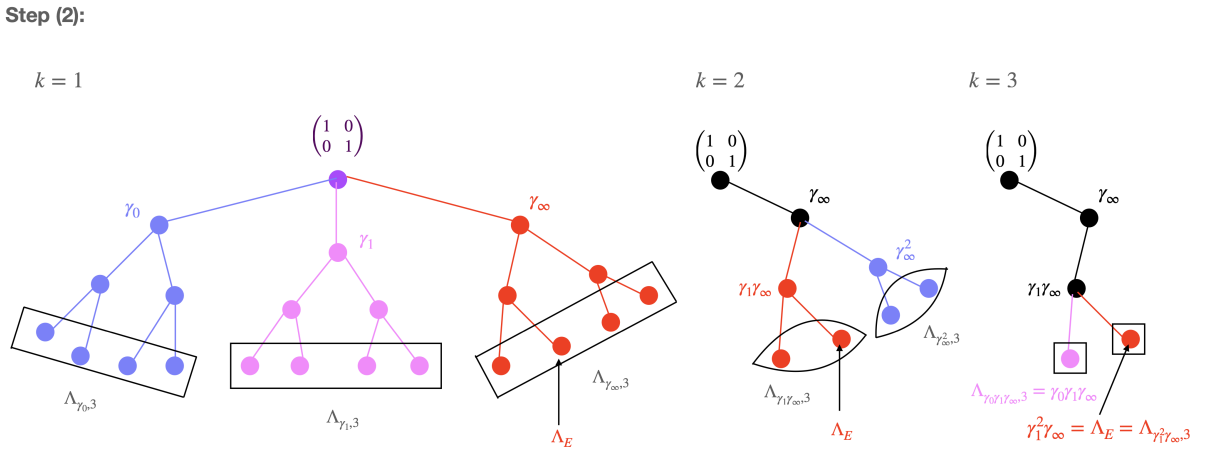

Figure 6.2. Algorithm 6.4.3 Step 2 with $q = 2$ and $d(\Lambda_E, M_2(\mathbb{Z}_q))) = 3$. Black edges correspond to edges of the path to Λ_E revealed in previous steps.

Proposition 6.4.4. *Algorithm 6.4.3 is correct and requires at most* $4(rq+1)$ *applications of Proposition 6.1.1. In each application, the input for the algorithm in Proposition 6.1.1 has* $n = q^3$ *and* β *with* $\log(\deg(\beta))$ *polynomially-sized in* $\log(q^r)$ *and the size of* $\mathcal{O}_0$.

Proof. We know that for some $T_E = d_r d_{r-1} \cdots d_1$, we have $\Lambda_E = T_E^{-1} M_2(\mathbb{Z}_q) T_E$. We show that T_E is the output of the algorithm.

Let $\gamma = c_k \cdots c_1$. By the construction in Corollary 6.3.7, we have $\mathcal{O}_\gamma \subset \mathrm{End}(E)$ if and only if $\Lambda_\gamma \subset \Lambda_E$. By Proposition 6.4.1(3) and 6.4.2(3), we have $\Lambda_{c_k \cdots c_2 c_1} \subset \Lambda_E$ if and only if $c_i = d_i$ for each $i \leq k$. Therefore, Step 2(a)(iv) has $b \in \mathrm{End}(E)$ for all $b \in B_\gamma$ if and only if $c_i = d_i$ for each $i \leq k$. The output after the $k = r$ step must be $\gamma = d_r d_{r-1} \cdots d_1 = T_E$, as desired.

When $k = 1$, there are $q + 1$ choices for γ_c. For $1 < k \leq r$, there are q choices for γ_c. Thus, we test $\mathcal{O}_\gamma \subset \mathrm{End}(E)$ for at most $rq + 1$ values of γ.

If $\gamma = \gamma_c \gamma'$, then a basis for $\mathcal{O}_\gamma$ can be written as a $\mathbb{Z}_{(q)}$-linear combination of elements of $\mathcal{O}_{\gamma'} + \mathcal{O}_0$. The exponent of q in the denominator is at most 3, since $v_q([\Lambda_\gamma : \Lambda_{\gamma'}]) = v_q \operatorname{discrd}(\Lambda_{\gamma'}) - v_q \operatorname{discrd}(\Lambda_\gamma)) = 3$ by Propositions 6.4.1(2) and 6.4.2(2), and $\mathcal{O}_\gamma \otimes \mathbb{Z}_{q'} \subset \mathcal{O}_0 \otimes \mathbb{Z}_{q'}$ when $q' \neq q$ by construction. As the previous step verifies that $\mathcal{O}_{\gamma'} \subset \mathrm{End}(E)$, this shows that elements of B_γ can be expressed as $\frac{\beta}{n}$ for $\beta \in \mathrm{End}(E)$ and $n = q^3$. By Corollary 6.3.7, we have that $\log(\deg(\beta))$ is polynomially-sized in $\log(q^r)$ and the size of $\mathcal{O}_0$. □

6.5 Special case: Bass orders

Let $\Lambda_0 = f(\mathcal{O}_0 \otimes \mathbb{Z}_q)$ and $\Lambda_E = f(\mathrm{End}(E) \otimes \mathbb{Z}_q)$. If Λ_0 is Bass, the subgraph of maximal orders containing Λ_0 forms a path which can be recovered efficiently. In this case, we can give a simpler algorithm which performs a binary search along the path to find Λ_E.

Algorithm 6.5.1. Finding Λ_E When Λ_0 is Bass

Input: An order $\mathcal{O}_0 \subset \mathrm{End}(E)$ which is Bass at q; $e = v_q(\operatorname{discrd}(\mathcal{O}_0))$

Output: γ such that $\Lambda_E = \gamma^{-1} M_2(\mathbb{Z}_q) \gamma$

1. Compute a q-maximal q-enlargement $\tilde{\mathcal{O}} \supset \mathcal{O}_0$ and an isomorphism $f : \tilde{\mathcal{O}} \otimes \mathbb{Q}_q \to M_2(\mathbb{Q}_q)$ such that $f(\tilde{\mathcal{O}} \otimes \mathbb{Z}_q) = M_2(\mathbb{Z}_q)$ up to precision q^{e+1}. Set $\Lambda_0 := f(\mathcal{O}_0 \otimes \mathbb{Z}_q)$.

2. Compute a list L of matrices T_i associated to matrix paths such that $\Lambda_0 \subset T_i^{-1} M_2(\mathbb{Z}_q) T_i$. Index the matrices T_i, starting with $i = 1$, such that T_i and T_{i+1} are adjacent in the Bruhat-Tits tree.

3. While $|L| > 1$:

   (a) Set $m := \lfloor \frac{|L|}{2} \rfloor$, $\Lambda_{\text{start}} := T_1^{-1} M_2(\mathbb{Z}_q) T_1$, and $\Lambda_{\text{mid}} := T_m^{-1} M_2(\mathbb{Z}_q) T_m$.

   (b) Compute a basis B for an order $\mathcal{O}$ such that $f(\mathcal{O} \otimes \mathbb{Z}_q) = \Lambda_{\text{start}} \cap \Lambda_{\text{mid}}$ and $\mathcal{O} \otimes \mathbb{Z}_{q'} \subset \mathcal{O}_0 \otimes \mathbb{Z}_{q'}$ for all $q' \neq q$. (Corollary 6.3.7)

   (c) For each $b \in B$, use Proposition 6.1.1 to determine if $b \in \mathrm{End}(E)$.

(d) If $b \in \mathrm{End}(E)$ for all b, set $L := \{T_i : 1 \leq i \leq m\}$. Otherwise, set $L := \{T_i : m < i \leq |L|\}$, and reindex the matrices, by replacing the index $i + m$ with the index i for $1 \leq i \leq |L| - m$.

4. Output the single element of L.

Proposition 6.5.2. *Algorithm 6.5.1 is correct, requires at most* $4\log_2(e+1)$ *applications of Proposition 6.1.1, and runs in polynomial time in* $\log(q)$ *and the size of* $\mathcal{O}_0$.

Proof. Step 1 can be done in polynomial-time by Propositions 6.3.1 and 6.3.3. The list L has size at most $e + 1$, the orders of L form a path, and L can be computed in polynomial time in $\log(q)$ and the size of Λ_0 by [21, Algorithm 4.1, Proposition 4.2]. As Λ_E contains Λ_0, the order Λ_E must be one of the orders in the list L. Each iteration of Step 3 tests which half of the path contains Λ_E and discards the other half. After $\log_2(e+1)$ loops, there is only one order remaining in L, which must be Λ_E. □

6.6 Computing the Endomorphism Ring

Now we can give the full algorithm to compute the endomorphism ring of $\mathrm{End}(E)$ on input of two noncommuting endomorphisms, α and γ, which generate a subring $\mathcal{O}_0$ of $\mathrm{End}(E)$. At every prime q for which $\mathcal{O}_0$ is not maximal, we find the path from $M_2(\mathbb{Z}_q)$ to $\mathrm{End}(E) \otimes \mathbb{Z}_q$ in the Bruhat-Tits tree. We emphasize that the only tool we have to distinguish $\mathrm{End}(E) \otimes \mathbb{Z}_q$ from the other orders of the Bruhat-Tits tree is the existence of an algorithm which determines if a local order Λ is contained in $\mathrm{End}(E) \otimes \mathbb{Z}_q$.

Algorithm 6.6.1. Computing the Endomorphism Ring

Input: A supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$; a suborder $\mathcal{O}_0$ of $\mathrm{End}(E)$ represented by a basis $\{1, \alpha, \gamma, \alpha\gamma\}$, such that α and γ can be evaluated efficiently on powersmooth torsion points of E; a factorization of $\mathrm{discrd}(\mathcal{O}_0)$

Output: A basis for $\mathrm{End}(E)$

1. For each prime $q \mid (\mathrm{discrd}(\mathcal{O}_0)/p)$:

   (a) Test if $\mathcal{O}_0 \otimes \mathbb{Z}_q$ is Bass. If so, use Algorithm 6.5.1 to compute $\mathcal{O}_q = \mathrm{End}(E) \otimes \mathbb{Z}_q$.

   (b) Compute a q-maximal q-enlargement $\tilde{\mathcal{O}}$ of $\mathcal{O}_0$. If $q = p$, output $\tilde{\mathcal{O}}$. Otherwise, proceed to Step 2. [53, Algorithm 3.12, 7.9, 7.10]

   (c) Compute the distance r between $\tilde{\mathcal{O}} \otimes \mathbb{Z}_q$ and $\mathrm{End}(E) \otimes \mathbb{Z}_q$, considered as vertices in the Bruhat-Tits tree. [Algorithm 6.2.5]

(d) Compute an isomorphism $f : \tilde{\mathcal{O}} \otimes \mathbb{Z}_q \to M_2(\mathbb{Z}_q)$ given modulo q^{r+1}, which extends to an isomorphism $f : \tilde{\mathcal{O}} \otimes \mathbb{Q}_q \to M_2(\mathbb{Q}_q)$. [Proposition 6.3.3]

(e) Compute the matrix γ such that $f(\mathrm{End}(E) \otimes \mathbb{Z}_q) = \gamma^{-1} M_2(\mathbb{Z}_q)\gamma$. [Algorithm 6.4.3]

(f) Compute a basis for a global order $\mathcal{O}_q$ such that $f(\mathcal{O}_q \otimes \mathbb{Z}_q) = \gamma^{-1} M_2(\mathbb{Z}_q)\gamma$ and $\mathcal{O}_q \otimes \mathbb{Z}_{q'} \subset \mathcal{O}_0 \otimes \mathbb{Z}_{q'}$ for all $q' \neq q$. [Corollary 6.3.7]

2. Return $\mathcal{O}_0 + \sum_{q|\mathrm{discrd}(\mathcal{O}_0)/p} \mathcal{O}_q$.

We now prove Theorem 1.3.1.

Proof. The order $\mathcal{O}_0 \otimes \mathbb{Z}_q$ is Bass if and only if $\mathcal{O}_0 \otimes \mathbb{Z}_q$ and the radical idealizer $(\mathcal{O}_0 \otimes \mathbb{Z}_q)^\flat$ are Gorenstein [12, Corollary 1.3]. An order is Gorenstein if and only if the associated ternary quadratic form is primitive [54, Theorem 24.2.10], which can be checked efficiently.

For Steps 1b and 1d, we must compute a multiplication table and reduced norm form Q for $\mathcal{O}_0$. Coefficients are given by the reduced traces of pairwise products of the basis, which can be evaluated efficiently using a modified Schoof's Algorithm by computing the trace on powersmooth torsion points (see [6, Theorem 6.10]).

Each substep of Step 1 has polynomial runtime. In the worst case, Step 1 requires $\sum_{i=1}^{m} 4(e_i q_i + 2)$ applications of Proposition 6.1.1, where $\mathrm{discrd}(\mathcal{O}_0) = \prod_{i=1}^{m} q_i^{e_i}$.

By construction, $q^{2e}\mathcal{O}_q \subset \mathcal{O}_0$, and hence $\mathrm{disc}(\mathcal{O}_0)\mathcal{O}_q \subset \mathcal{O}_0$. By Corollary 6.3.6, a basis for the sum $\mathcal{O}_0 + \sum_{q|\mathrm{discrd}(\mathcal{O}_0)/p} \mathcal{O}_q$ can be computed in polynomial-time. By construction, the output has completion $\mathcal{O}_q \otimes \mathbb{Z}_q = \mathrm{End}(E) \otimes \mathbb{Z}_q$ at every prime $q \mid \mathrm{discrd}(\mathcal{O}_0)$ and also still at all other primes. By the local-global principle, the output is $\mathrm{End}(E)$. □

6.7 Subgraphs specified by intersection

Suppose we are given $\mathcal{O}_0 \subset \mathrm{End}(E)$ and, under an isomorphism $f : \mathcal{O}_0 \otimes \mathbb{Q}_q \to M_2(\mathbb{Q}_q)$, we set $\Lambda_0 = f(\mathcal{O}_0 \otimes \mathbb{Z}_q)$. Our paper describes how to find $\Lambda_E = f(\mathrm{End}(E) \otimes \mathbb{Z}_q)$ without recovering the tree of orders containing Λ_0. However, in the case that the graph of orders containing Λ_0 is a path, we have have seen that Λ_E can be recovered more efficiently, as in Algorithm 6.5.1.

In this section, we describe how our algorithm can be made more efficient if we already know the subgraph of maximal orders containing Λ_0, specified by the intersection of three maximal orders.

6.7.1 Possible subgraphs

We can use Tu's results to describe the subgraph of maximal orders containing any order in $M_2(\mathbb{Q}_q)$. We summarize the possible subgraphs in the following corollary.

Corollary 6.7.1. *Suppose Λ is an order in $M_2(\mathbb{Q}_q)$. Let $S = \{\Lambda' \text{ maximal} : \Lambda \subset \Lambda'\}$. Then there exists a path P and an integer $\ell \geq 0$ such that $S = N_\ell(P)$.*

Proof. By Lemma 5.2.9, Λ is contained in only finitely many maximal orders even when Λ is not a finite intersection of maximal orders. Hence the set S of maximal orders containing Λ is a finite set. Let $\Lambda'' = \cap_{\Lambda' \in S} \Lambda'$. The set of maximal orders containing Λ is precisely the set of maximal orders containing Λ''. Thus, it suffices to prove the statement in the case that Λ is equal to a finite intersection of maximal orders.

For the rest of the proof, assume Λ is a finite intersection of maximal orders. By Theorem 5.3.2, we can choose $\Lambda_1, \Lambda_2, \Lambda_3 \in S$ such that $\Lambda = \Lambda_1 \cap \Lambda_2 \cap \Lambda_3$. We need to construct a path P and an integer $\ell \geq 0$ such that for a maximal order Λ', we have $\Lambda_1 \cap \Lambda_2 \cap \Lambda_3 \subset \Lambda'$ if and only if $\Lambda' \in N_\ell(P)$.

By reindexing if necessary, assume $d(\Lambda_1, \Lambda_2)$ is maximal among $d(\Lambda_i, \Lambda_j)$. Let P' denote the path from Λ_1 to Λ_2, and let $\ell = d(\Lambda_3, P')$.

By maximality of $d(\Lambda_1, \Lambda_2)$, the path P' has length at least 2ℓ. For $i = 1, 2$, let Λ_i' denote the vertex on the path P' such that $d(\Lambda_i', \Lambda_i) = \ell$. Let P be the path of vertices from Λ_1' to Λ_2'. We will show that $S = N_\ell(P)$.

First, note that $\ell = d(\Lambda_3, P)$. Suppose not. Then the closest vertex v' of P' to Λ_3 lies between Λ_i' and Λ_i for some i, and $v' \neq \Lambda_i'$. Then for $j \neq i, j \neq 3$, we have $d(\Lambda_j, \Lambda_i) = d(\Lambda_j, \Lambda_i') + d(\Lambda_i', \Lambda_i) = d(\Lambda_j, \Lambda_i') + d(\Lambda_i', v') + d(v', \Lambda_i)$. Since $d(v', \Lambda_i) < \ell = d(v', \Lambda_3)$, we have $d(\Lambda_j, \Lambda_i) \leq d(\Lambda_j, v') + d(v', \Lambda_3)$. But $d(\Lambda_j, v') + d(v', \Lambda_3) = d(\Lambda_j, \Lambda_3)$. This contradicts maximality of $d(\Lambda_1, \Lambda_2)$.

Let Λ_4 be a maximal order, and let $m = d(\Lambda_4, P)$. We need to show that $\Lambda_4 \in S$ if and only if $m \leq \ell$. By Lemmas 5.3.4 and 5.3.5, this is the same as showing that $d_3(S) = d_3(S \cup \{\Lambda_4\})$ if and only if $m \leq \ell$.

We have $d(\Lambda_1, \Lambda_2) + d(\Lambda_2, \Lambda_4) + d(\Lambda_4, \Lambda_1) = 2d(\Lambda_1, \Lambda_2) + 2m$. We will show that $d_3(S \cup \{\Lambda_4\}) = 2d(\Lambda_1, \Lambda_2) + 2\max\{\ell, m\}$.

If $i = 1$ or $i = 2$, let P_i denote the path between Λ_3 and Λ_i, and let $n_i = d(\Lambda_4, P_i)$.. We have $d(\Lambda_i, \Lambda_4) + d(\Lambda_4, \Lambda_3) + d(\Lambda_3, \Lambda_i) = 2d(\Lambda_i, \Lambda_3) + 2n_i$. If $n_i \leq m$, then this is clearly at most $2d(\Lambda_1, \Lambda_2) + 2m$. Let v_i denote the vertex of P_i which is closest to Λ_4, so $d(\Lambda_4, v_i) = n_i$ If $n_i > m$, the path P_i does not contain the closest vertex v on P to Λ_4 and $d(v, v_i) = n_i - m$. In this case, it follows that v_i lies on the path P, as otherwise $v_i = v$ and $n_i = m$, and that v_i is the closest vertex of P to Λ_3. Thus, $d(\Lambda_1, \Lambda_2) = d(\Lambda_i, v_i) + n_i - m + d(v, \Lambda_j)$,

where $j \neq i, 3, 4$. We also have $d(\Lambda_i, \Lambda_3) = d(\Lambda_i, v_i) + d(v_i, \Lambda_3)$. Thus $2d(\Lambda_i, \Lambda_3) + 2n_i = 2(d(\Lambda_i, v_i) + n_i - m) + 2m \leq 2d(\Lambda_1, \Lambda_2) + 2m$.

We have shown that $d_3(S \cup \{\Lambda_4\}) = 2d(\Lambda_1, \Lambda_2) + 2\max\{\ell, m\}$. This is equal to $d_3(S)$ if and only if $m \leq \ell$, and hence Λ_4 is in S if and only if $\Lambda_4 \in N_\ell(P)$. □

6.7.2 Special case: known subgraph of the Bruhat-Tits tree

In the general case, Algorithm 1.3.1 works by finding the distance between Λ_E and $M_2(\mathbb{Z}_q)$ and then finding the path between them. Finding the path is the most costly step. In the worst case, if $e = v_q(\mathrm{discrd}(\mathcal{O}_0))$ and $e = d(M_2(\mathbb{Z}_q), \Lambda_E)$, the algorithm tests $eq + 1$ steps.

If we can describe the set of maximal orders containing Λ_0 as $N_\ell(P)$ for a path P and an integer $\ell \geq 0$, we can obtain Λ_E more efficiently, by replacing $M_2(\mathbb{Z}_q)$ with a closer vertex. First, we compute the distance r of Λ_E from the path P; next, we compute the order Λ' in P which is closest to Λ_E; finally, we recover the path from Λ' to Λ_E, one step at a time. This last step is the most costly, but in the worst case, we only need to recover ℓ steps, where ℓ is at most $\mathrm{discrd}(\Lambda_0)/3$.

Algorithm 6.7.2. Finding Λ_E When the Subgraph is Known

Input: An order $\mathcal{O}_0 \subset \mathrm{End}(E)$; $e = v_q(\mathrm{discrd}(\mathcal{O}_0))$; a q-maximal q-enlargement $\tilde{O}$ of $\mathcal{O}_0$; an isomorphism $f : \mathcal{O}_0 \otimes \mathbb{Q}_q \to M_2(\mathbb{Q}_q)$ such that $f(\tilde{O} \otimes \mathbb{Z}_q) = M_2(\mathbb{Z}_q)$, given up to precision q^{e+1}; matrices T_1' and T_2' corresponding to endpoints of a path P and an integer $\ell \geq 0$ such that $\cap_{\Lambda \supset f(\mathcal{O}_0 \otimes \mathbb{Z}_q)} \Lambda = \cap_{\Lambda \in N_\ell(P)} \Lambda$.

Output: γ such that $\Lambda_E = \gamma^{-1} M_2(\mathbb{Z}_q) \gamma$

1. Compute the least $r \leq \ell$ such that $\cap_{\Lambda \in N_r(P)} \Lambda \subset \Lambda_E$.

2. Similar to Algorithm 6.5.1, partition P into two disjoint paths P_0 and P_1 of equal length, and check if P_0 satisfies $\Lambda_E \in N_r(P_0)$. Set $P' = P_0$ if $\Lambda_E \in N_r(P_0)$ and $P' = P_1$ otherwise. Then replace P by P' and continue until P consists of a single order $T^{-1} M_2(\mathbb{Z}_q) T$.

3. Similar to Algorithm 6.4.3, recover the matrix path $d_1, d_2, \ldots, d_r$ of length r from $T^{-1} M_2(\mathbb{Z}_q) T$ to Λ_E, so that $\Lambda_E = (d_r \cdots d_2 d_1 T)^{-1} M_2(\mathbb{Z}_q) d_r \cdots d_2 d_1 T$.
   Output $d_r \cdots d_2 d_1 T$.

Proposition 6.7.3. *Algorithm 6.7.2 requires at most* $4(\ell + \log(\mathrm{card}(P)) + \ell q + 1)$ *applications of Algorithm 6.1.9.*

Proof. Step 1 is a generalization of Algorithm 6.2.5, replacing $N_k(M_2(\mathbb{Z}_q))$ by $N_k(P)$ and testing locally rather than globally. Step 2 is a generalization of Algorithm 6.5.1, replacing

the path from Λ_{start} to Λ_{mid} with the r-neighborhood of that path, where $r = d(\Lambda_E, P)$. Step 3 is a generalization of Algorithm 6.4.3, replacing $M_2(\mathbb{Z}_q)$ by $T^{-1}M_2(\mathbb{Z}_q)T$. □

The extra information about the graph structure allows us to replace $M_2(\mathbb{Z}_q)$ with an order whose which is close to Λ_E, thus minimizing the most costly step (recovering the path step-by-step). However, we stress that it is not clear how to efficiently obtain Λ_1, Λ_2, and Λ_3 from Λ_0.

Bibliography

[1] Noga Alon, Itai Benjamini, Eyal Lubetzky, and Sasha Sodin. Non-backtracking random walks mix faster. *Commun Contemp Math*, 9, 10 2006.

[2] Sarah Arpin, Mingjie Chen, Kristin E. Lauter, Renate Scheidler, Katherine E. Stange, and Ha T. N. Tran. Orientations and cycles in supersingular isogeny graphs, 2022.

[3] Sarah Arpin, Mingjie Chen, Kristin E. Lauter, Renate Scheidler, Katherine E. Stange, and Ha Thanh Nguyen Tran. Orienteering with one endomorphism. *La Matematica*, 2:523 – 582, 2022.

[4] Sarah Arpin, James Clements, Pierrick Dartois, Jonathan Komada Eriksen, Péter Kutas, and Benjamin Wesolowski. Finding orientations of supersingular elliptic curves and quaternion orders. Cryptology ePrint Archive, Paper 2023/1268, 2023. `https://eprint.iacr.org/2023/1268`.

[5] 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`.

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

[7] Andrea Basso, Giulio Codogni, Deirdre Connolly, Luca De Feo, Tako Boris Fouotsa, Guido Maria Lido, Travis Morrison, Lorenz Panny, Sikhar Patranabis, and Benjamin Wesolowski. Supersingular curves you can trust. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023*, pages 405–437, Cham, 2023. Springer Nature Switzerland.

[8] Gaetan Bisson. Computing endomorphism rings of elliptic curves under the grh. *Journal of Mathematical Cryptology*, 5(2):101–114, 2012.

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

[10] Wouter Castryck and Thomas Decru. An efficient key recovery attack on SIDH. In *Advances in cryptology—EUROCRYPT 2023. Part V*, volume 14008 of *Lecture Notes in Comput. Sci.*, pages 423–447. Springer, Cham, 2023.

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

[12] Sara Chari, Daniel Smertnig, and John Voight. On basic and Bass quaternion orders. *Proc. Amer. Math. Soc. Ser. B*, 8:11–26, 2021.

[13] Henri Cohen. *A course in computational algebraic number theory*, volume 138 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 1993.

[14] Leonardo Colò and David Kohel. Orienting supersingular isogeny graphs. Cryptology ePrint Archive, Paper 2020/985, 2020. `https://eprint.iacr.org/2020/985`.

[15] Keith Conrad. The conductor ideal of an order. `https://kconrad.math.uconn.edu/blurbs/gradnumthy/conductor.pdf`.

[16] Jean-Marc Couveignes. Hard homogeneous spaces. Cryptology ePrint Archive, Paper 2006/291, 2006. `https://eprint.iacr.org/2006/291`.

[17] David A. Cox. *Primes of the form $x^2 + ny^2$*. Pure and Applied Mathematics (Hoboken). John Wiley & Sons, Inc., Hoboken, NJ, second edition, 2013. Fermat, class field theory, and complex multiplication.

[18] Pierrick Dartois and Luca De Feo. On the security of osidh. In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *Public-Key Cryptography – PKC 2022*, pages 52–81, Cham, 2022. Springer International Publishing.

[19] Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski. SQISignHD: New dimensions in cryptography. Cryptology ePrint Archive, Paper 2023/436, 2023. `https://eprint.iacr.org/2023/436`.

[20] Kirsten Eisenträger, Sean Hallgren, Kristin Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: reductions and solutions. In *Advances in cryptology—EUROCRYPT 2018. Part III*, volume 10822 of *Lecture Notes in Comput. Sci.*, pages 329–368. Springer, Cham, 2018.

[21] Kirsten Eisenträger, Sean Hallgren, Chris Leonardi, Travis Morrison, and Jennifer Park. Computing endomorphism rings of supersingular elliptic curves and connections to path-finding in isogeny graphs. In *ANTS XIV—Proceedings of the Fourteenth Algorithmic Number Theory Symposium*, volume 4 of *Open Book Ser.*, pages 215–232. Math. Sci. Publ., Berkeley, CA, 2020.

[22] Luca De Feo, Tako Boris Fouotsa, Péter Kutas, Antonin Leroux, Simon-Philipp Merz, Lorenz Panny, and Benjamin Wesolowski. Scallop: Scaling the csi-fish. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *Public-Key Cryptography – PKC 2023*, pages 345–375, Cham, 2023. Springer Nature Switzerland.

[23] Jenny Fuselier, Annamaria Iezzi, Mark Kozek, Travis Morrison, and Changningphaabi Namoijam. Computing supersingular endomorphism rings using inseparable endomorphisms, Preprint, 2023. `https://arxiv.org/pdf/2306.03051`.

[24] Arthur Herlédan Le Merdy and Benjamin Wesolowski. The supersingular endomorphism ring problem given one endomorphism. Cryptology ePrint Archive, Paper 2023/1448, 2023. `https://eprint.iacr.org/2023/1448`.

[25] Masanobu Kaneko. Supersingular j-invariants as singular moduli mod p. *Osaka Journal of Mathematics*, 26:849–855, 1989.

[26] Ernst Kani. The number of curves of genus two with elliptic differentials. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 1997:122–93, 1997.

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

[28] 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.

[29] Péter Kutas, Christophe Petit, Luca de Feo, Antonin Leroux, Tako Boris Fouotsa, Benjamin Wesolowski, Cyprien Delpech de Saint Guilhem, and Javier Silva. Séta: Supersingular encryption from torsion attacks. Springer-Verlag, 2021.

[30] Serge Lang. *Elliptic functions*. Graduate Texts in Mathematics. Springer New York, NY, 2012.

[31] Kristin E. Lauter and Bianca Viray. On singular moduli for arbitrary discriminants. *International Mathematics Research Notices*, 2015:9206–9250, 2012.

[32] Jianing Li, Songsong Li, and Yi Ouyang. Factorization of hilbert class polynomials over prime fields. 2021.

[33] Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, and Benjamin Wesolowski. A direct key recovery attack on SIDH. In *Advances in cryptology—EUROCRYPT 2023. Part V*, volume 14008 of *Lecture Notes in Comput. Sci.*, pages 448–471. Springer, Cham, [2023] ©2023.

[34] Daniele Micciancio and Shafi Goldwasser. *Complexity of lattice problems*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, MA, 2002. A cryptographic perspective.

[35] J. S. Milne. Abelian varieties. In *Arithmetic geometry (Storrs, Conn., 1984)*, pages 103–150. Springer, New York, 1986.

[36] David Mumford. *Abelian varieties*. Published for the Tata Institute of Fundamental Research, Bombay by Oxford University Press, London, 1970.

[37] Hiroshi Onuki. On oriented supersingular elliptic curves. *Finite Fields and Their Applications*, 69:101777, 2021.

[38] Aurel Page and Benjamin Wesolowski. The supersingular endomorphism ring and one endomorphism problems are equivalent. Cryptology ePrint Archive, Paper 2023/1399, 2023. `https://eprint.iacr.org/2023/1399`.

[39] Arnold K. Pizer. Ramanujan graphs and hecke operators. *Bulletin of the American Mathematical Society*, 23:127–137, 1990.

[40] Paul Pollack and Enrique Treviño. Finding the four squares in Lagrange's theorem. *Integers*, 18A:A15, 2018.

[41] Michael O. Rabin and Jeffery O. Shallit. Randomized algorithms in number theory. *Communications on Pure and Applied Mathematics*, 39(S1):S239–S256, 1986.

[42] Damien Robert. Breaking SIDH in polynomial time. In *Advances in cryptology—EUROCRYPT 2023. Part V*, volume 14008 of *Lecture Notes in Comput. Sci.*, pages 472–503. Springer, Cham, [2023] ©2023.

[43] Damien Robert. Some applications of higher dimensional isogenies to elliptic curves, Preprint, 2023.

[44] J. Barkley Rosser and Lowell Schoenfeld. Approximate formulas for some functions of prime numbers. *Illinois Journal of Mathematics*, 6:64–94, 1962.

[45] Alexander Rostovtsev and Anton Stolbunov. Public-key cryptosystem based on isogenies. Cryptology ePrint Archive, Paper 2006/145, 2006. `https://eprint.iacr.org/2006/145`.

[46] P.W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, 1994.

[47] Joseph H. Silverman. *Advanced topics in the arithmetic of elliptic curves*. Graduate Texts in Mathematics. Springer New York, NY, first edition, 1994.

[48] Joseph H. Silverman. *The arithmetic of elliptic curves*. Graduate Texts in Mathematics. Springer-Verlag, New York, second edition, 2009.

[49] Anton Stolbunov. Constructing public-key cryptographic schemes based on class group action on a set of isogenous elliptic curves. *Advances in Mathematics of Communications*, 4(2):215–235, 2010.

[50] Andrew Sutherland. Isogeny volcanoes. *The Open Book Series*, 1(1):507–530, November 2013.

[51] Fang-Ting Tu. On orders of $M(2, K)$ over a non-Archimedean local field. *Int. J. Number Theory*, 7(5):1137–1149, 2011.

[52] Christiaan E. van de Woestijne. Deterministic equation solving over finite fields. In *International Symposium on Symbolic and Algebraic Computation*, 2005.

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

[54] John Voight. *Quaternion algebras*, volume 288 of *Graduate Texts in Mathematics*. Springer, Cham, [2021] ©2021.

[55] Joachim von zur Gathen and Jürgen Gerhard. *Modern computer algebra*. Cambridge University Press, Cambridge, third edition, 2013.

[56] William C. Waterhouse. Abelian varieties over finite fields. *Annales Scientifiques De l'École Normale Superieure*, 2:521–560, 1969.

[57] Benjamin Wesolowski. The supersingular isogeny path and endomorphism ring problems are equivalent. *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1100–1111, 2021.

[58] Benjamin Wesolowski. Orientations and the supersingular endomorphism ring problem. In *Advances in cryptology—EUROCRYPT 2022. Part III*, volume 13277 of *Lecture Notes in Comput. Sci.*, pages 345–371. Springer, Cham, [2022] ©2022.

Vita

Gabrielle Scullard

Education

May 2018 *Rochester, NY*	Honors B.S. in Mathematics, **University of Rochester** • GPA: 3.99
May 2024 *State College, PA*	Ph.D. in Mathematics, **The Pennsylvania State University** • Advisor: Kirsten Eisenträger

Awards

Fall 2018	Maryam Mirzhakani Graduate Scholarship in Mathematics
Spring 2017	Elected to Phi Beta Kappa, Iota Chapter of New York as a junior

Papers

Preprint in progress	Horizontal isogenies and applications to OSIDH *Gabrielle Scullard*
Submitted 2024	Computing endomorphism rings of supersingular elliptic curves using higher-dimensional isogenies *Kirsten Eisenträger and Gabrielle Scullard*
Submitted 2024	The probability of non-isomorphic group structures of isogenous elliptic curves in finite field extensions, II *John Cullinan, Shanna Dobson, Jorge de Mello, Linda Frey, Asimina Hamakiotes, Roberto Hernandez, Nathan Kaplan, Gabrielle Scullard*