Master's thesis

NTNU
Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Mathematical Sciences

Jørgen Juel

Constructing Isogenies Between Elliptic Curves

Master's thesis in Mathematical Sciences
Supervisor: Kristian Gjøsteen
June 2021

NTNU
Norwegian University of
Science and Technology

Jørgen Juel

Constructing Isogenies Between Elliptic Curves

Master's thesis in Mathematical Sciences
Supervisor: Kristian Gjøsteen
June 2021

Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Mathematical Sciences

Abstract. In this thesis, we will explore methods to compute isogenies between elliptic curves in the ordinary and supersingular case. This is an attempt to understand the security foundations of isogeny-based cryptography. We will explore topics in algebraic number theory, quaternion algebra, lattices and some specific algorithms for dealing with such objects. Next, we will explore how to compute a connecting isogeny given only the starting and ending curve in both the ordinary case and the supersingular case. Finally we will look at two applications for this theory in terms of the concrete isogeny-based cryptosystems SIDH and SQISign.

Sammendrag. I denne oppgaven utforsker vi metoder for å beregne isogenier mellom elliptiske kurver i både det ordinære og supersingulære tilfellet. Dette er et forsøk på å forstå sikkerhetsgrunnlaget til isogenibasert kryptografi. Gjennom oppgaven kommer vi til å utforske algebraisk tallteori, kvaternionalgebra, gitre og noen spesifikke algoritmer for å jobbe med slike objekter. Videre kommer vi til å forklare hvordan vi konstruerer en isogeni mellom en gitt start- og sluttkurve i det ordinære og supersingulære tilfellet. Til slutt skal vi se på to anvendelser av teorien mot de konkrete isogenibaserte kryptosystemene SIDH og SQISign.

Preface

This thesis concludes my studies toward a Master of Science in Mathematics at the Norwegian University of Science and Technology (NTNU). It has been five challenging and rewarding years, where Matteland has proven its true value. Especially during the last year of master writing, where study room 395B and all its students have pressured me to work.

I would like to thank my supervisor Kristian Gjøsteen for helpful feedback and useful discussions during the year, and previously through my bachelor thesis. It has been very interesting to be able to explore the world of isogenies once again.

Although the ongoing pandemic could have made this year extremely difficult, my stay at Mørlenda has made this a rather enjoyable year. I would like to thank the Chinese gang, you know who you are, for our frequent get-togethers and weekly dinner competition. I would also like to thank Julie and Sindre for hosting the biweekly Wednesday dinners.

Jørgen Juel
Trondheim, 2021

Contents

CHAPTER 1

Introduction

While physicists and engineers tackle the challenge of constructing quantum computers, mathematicians and cryptographers work on building post-quantum cryptosystems able to resist the potential computing power in these machines. The National Institute of Standards and Technology (NIST) currently have a running project on establishing post-quantum cryptography (NIST-PQC)[31] which is expected to provide a standards draft in the period 2022-2024. One of the candidates rely on the hardness of finding isogenies between supersingular elliptic curves. Unfortunately, on the 22^{th} of July 2020, when Round 3 candidates where selected, the isogeny based cryptosystem (SIKE) was removed from the list of primary candidates [27].

SIKE was characterised as being the only cryptosystem based on isogenies and with exceptionally small key sizes. It does however require longer running time (roughly an order magnitude longer than its competitors) and its security assumption is much less studied.

> NIST sees SIKE as a strong candidate for future standardization with continued improvements and accordingly selected SIKE to move into the third round as an alternate candidate. There are applications which would benefit from SIKE's small key and ciphertext sizes and which may be able to accept the performance impact. Further research in isogeny-based cryptography is encouraged. [27]

Further research is exactly what motivated this thesis. Although we do not provide any new results, this thesis will hopefully make the current advances in the field more available. We will look into how one might go about constructing isogenies by understanding the mathematics underlying them. Overall, we are interested in looking further into the following problem which, roughly speaking, is the underlying security assumption of all isogeny-based cryptography.

PROBLEM. Given two elliptic curves E, E', find an isogeny $\phi : E \to E'$

We will begin by looking into the mathematical foundations required for studying the endomorphism ring of ordinary and supersingular elliptic curves in Chapter 2. There we will introduce some concepts from algebraic number theory, the notion of lattices in a number field and the complex numbers, what quaternions are and finally recall some facts about elliptic curves.

Next, in Chapter 3, we will look at how isogenies between ordinary elliptic curves behave. We will describe an algorithm by Galbraith at how one can construct isogenies between arbitrary

ordinary elliptic curves and look at its running time. This section is meant as a way to see the clear distinction between ordinary and supersingular elliptic curves.

In Chapter 4 we will look at supersingular elliptic curves and how one can always construct an l-power isogeny connecting two such curves. This is perhaps the chapter of most interest to the reader who would like to learn more about recent isogeny-based research. We will discuss the so-called KLPT algorithm[22] in detail. An algorithm which has found applications in recent signature schemes. We will also explore the difficulty of computing the endomorphism ring and see how there even exists some problems with mapping a well understood abstract endomorphism ring to actual endomorphisms on the elliptic curve.

Then, in Chapter 5 we will provide the reader with two applications in the supersingular case. We will explore how one can break the SIDH[14] cryptosystem (the basis of SIKE) and finally discuss how one can build a signature scheme, SQISign [15], based on secret knowledge of the endomorphism ring using a modified KLPT algorithm.

Finally, in Chapter 6 we will discuss the differences between constructing isogenies within ordinary curves and supersingular curves.

We conclude the thesis with Appendix A, describing how one can implement some of the algorithms from Chapter 4 in SageMath.

CHAPTER 2

Mathematical Foundations

In this chapter we will describe the mathematical foundations required for reading this thesis. We will look at algebraic number theory, lattices, quaternions and finally elliptic curves. The contents of algebraic number theory and elliptic curves is considered well-known and only provided to make the notation clear and provide results for later references. The sections on lattices and quaternions are treated in more detail as their content will be important for us in the upcoming chapters.

1. Algebraic Number Theory

In this section we will recall some facts from number theory focusing on imaginary quadratic fields. Then we will introduce quadratic forms before we end this section with a useful algorithm.

1.1. Imaginary number fields. This section is mainly based on [9, Chapter 2] and [33].

A **quadratic field** K is simply a field of the form $K = \mathbb{Q}(\sqrt{N})$ where $N \in \mathbb{Z} \setminus \{0, 1\}$ is square-free. Whenever $N < 0$ we say the the number field is **imaginary** as it somehow contains the traditional imaginary element i satisfying $i^2 = -1$. For any quadratic field there is an invariant called its **discriminant**, Δ_K. It is defined as

$$\Delta_K = \begin{cases} N & \text{if } N \equiv 1 \mod 4 \\ 4N & \text{otherwise} \end{cases}$$

We note that we can always embed our quadratic number field K in $\mathbb{C}$ for the simple reason that the square roots of N exists in $\mathbb{C}$. Therefore the numbers $a + b\sqrt{N} \in \mathbb{Q}$ are just elements $z \in \mathbb{C}$. That being said there are two solutions to the equation $\sqrt{N}$ in $\mathbb{C}$, thus there are two ways of embedding K in $\mathbb{C}$. Mapping from one to the other is what we refer to as the **nontrivial automorphism** of K. Notice that elements of $\mathbb{Q}$ necessarily remain fixed under this automorphism. Another useful map is **conjugation**, denoted $\overline{a}$ which simply is the map $\overline{a + b\sqrt{N}} = a - b\sqrt{N}$.

Once we have a number field we can define its **ring of integers**, O_K. It serves the same purpose as $\mathbb{Z}$ does in $\mathbb{Q}$. It can be described explicitly as

$$O_K = \begin{cases} \mathbb{Z}[\sqrt{N}] \text{ if } N \not\equiv 1 \mod 4 \\ \mathbb{Z}\left[\frac{1+\sqrt{N}}{2}\right] \text{ otherwise} \end{cases}$$

or once we know Δ_K it can be written as $O_K = \mathbb{Z}\left[\frac{\Delta_K+\sqrt{\Delta_K}}{2}\right]$

The ring of integers is sometimes referred to as an **order**, meaning that it is a $\mathbb{Z}$-submodule of K which is closed under multiplication. There can be many other orders in a number field, but for quadratic number fields there is always a unique maximal order, which is the ring of integers. A useful object related to orders are fractional ideals of an order. That is, a **fractional ideal** of O_K is simply a finitely generated O_K-submodule of K.

Sometimes we are interested in how a number a behaves with respect to a prime p, one such interesting measure is the **Kronecker symbol** denoted (a/p) or $\left(\frac{a}{p}\right)$. For a prime $p \neq 2$ it is just the Legendre symbol which is defined as 0 if p divides a, 1 if it is a quadratic residue modulo p and -1 otherwise. When $p = 2$ we define the Kronecker symbol to be 0 if a is even, 1 if $a \equiv \pm 3 \mod 8$ and -1 if $a \equiv \pm 1 \mod 8$.

When dealing with O_K we are often talking about prime ideals $\mathfrak{p}$ **laying above** p, this is just a prime ideal $\mathfrak{p} \subseteq O_K$ that satisfies $p \in \mathfrak{p}$. It can be shown that these prime ideals satisfy $\mathfrak{p} \cap \mathbb{Z} = (p)$, the ideal generated by p in $\mathbb{Z}$. We are also interested in what pO_K is. This will clearly be an ideal, but how it decomposes into prime ideals is what is called the ramification of p in O_K or sometimes just p in K. We use the following results to define it.

PROPOSITION 2.1. *Let K be a quadratic field of discriminant Δ_K, let the nontrivial automorphism of K be denoted as $\alpha \mapsto \alpha'$, and p be a prime in $\mathbb{Z}$. Then we have the following description of pO_K.*

(1) if $(\Delta_K/p) = 0$, then $pO_K = \mathfrak{p}^2$ for some prime ideal $\mathfrak{p} \subset O_K$
(2) if $(\Delta_K/p) = 1$, then $pO_K = \mathfrak{p}\mathfrak{p}'$ for distinct prime ideals $\mathfrak{p}, \mathfrak{p}' \subset O_K$
(3) if $(\Delta_K/p) = -1$, then $pO_K = \mathfrak{p}$ is a prime ideal in O_K

Proof: See [9, Proposition 5.16]

Thus we say that p **ramifies** in K if p divides Δ_K, that p **splits completely** in K if $(\Delta_K/p) = 1$ and that p is **inert** in K if $(\Delta_K/p) = -1$.

Moreover, the ring of integers, O_K is in fact a Dedekind domain [9, Theorem 5.5], which amongst other things means that every nonzero prime ideal of O_K is maximal. This allows us to look at the quotient $O_K/\mathfrak{p}$ for a prime $\mathfrak{p}$ of O_K which is a field is a field of finite size by [9, Corollary 5.4].

We define the **norm** on elements $\alpha \in K$ as $\mathrm{Nr}(\alpha) = \alpha\alpha'$ where α' is the element obtained from taking the nontrivial automorphism of α. This gives us that for $a \in \mathbb{Q}$, $\mathrm{Nr}(a) = a^2$. For ideals $\mathfrak{p}$, the norm is defined to be the size of $O_K/\mathfrak{p}$, that is $\mathrm{Nr}(\mathfrak{p}) = |O_K/\mathfrak{p}|$.

LEMMA 2.2. *Let O be an order in an imaginary quadratic field, then*

(1) $\mathrm{Nr}(pO) = \mathrm{Nr}(p)$
(2) $\mathrm{Nr}(\mathfrak{a}\mathfrak{b}) = \mathrm{Nr}(\mathfrak{a})\,\mathrm{Nr}(\mathfrak{b})$
(3) $\mathfrak{a}\overline{\mathfrak{a}} = \mathrm{Nr}(\mathfrak{a})O$

Proof: See [9, Lemma 7.14]

Thus for a prime p, with $\mathfrak{p}$ laying above p we have $O_K/\mathfrak{p} \cong \mathbb{F}_q$ for $q = p^2$ when p is inert, and $q = p$ if it is ramified or splits completely.

Finally we would like to describe the class number. Whenever we have an order O, we can talk about fractional O-ideals. If we take the set of such fractional ideals, modulo principal fractional O-ideals we get what is known as the **ideal class group** of O, often denoted $\mathcal{CL}(O)$. In other words, $\mathfrak{a}$ and $\mathfrak{b}$ are in the same class if there exists some α such that $\mathfrak{a} = \alpha\mathfrak{b}$. There are some details to be worked out when dealing with orders properly contained in O_K instead of the entire ring of integers, but this is left for later.

The **class number**, $h(O_K)$, is defined to be the number of elements of $\mathcal{CL}(O)$. It will sometimes be denoted h_K to say that it is the class number of the ring of integers of K, but we will stick to $h(O_K)$. It is in general considered difficult to compute it, but in Cohen [8, Exercise 5.27 (b)] it is given a bound based on the discriminant.

$$h(\mathcal{O}_K) \leq \frac{1}{\pi}\sqrt{|\Delta_K|}\ln|\Delta_K|$$

To show it one uses the fact that for imaginary quadratic fields of discriminant $\Delta_K < -4$ the relation $\sum_{n\geq 1}\left(\frac{\Delta_K}{n}\right)/n = \pi h(\mathcal{O}_K)/\sqrt{|\Delta_K|}$ [8, Proposition 5.3.12]. The sum can be shown to be bounded by $\ln(|\Delta_K|)$ giving the desired result.

1.2. Forms. In addition to number fields, we would also like the notion of forms. These are just specific kinds of polynomials which will make our lives easier when dealing with supersingular isogenies later on.

Definition 2.3. A **quadratic form** in two variables over a ring R is a polynomial

$$f(x, y) = ax^2 + bxy + cy^2 \text{ with } a, b, c \in R$$

We say that it is **primitive** if the numbers a, b, c are relatively prime. The **discriminant** of a quadratic form f is simply $b^2 - 4ac$ and is sometimes denoted $\text{disc}(f)$. In our case the ring R will be the integers $\mathbb{Z}$ and we look for certain primitive forms defined by their discriminant.

Example 2.4. A primitive quadratic form over $\mathbb{Z}$ of discriminant -4 is just the polynomial $x^2 + y^2$.

1.3. Cornaccia's algorithm. Finally we will end this section on Algebraic Number theory with an important algorithm known as Cornaccia's algorithm. The algorithm by Cornaccia, described in [3], solves the equation $x^2 + dy^2 = m$ whenever d and m are coprime. This will be useful when we attempt to represent the integer m as the reduced norm in a number fields and when we would like to find solutions to primitive quadratic forms of discriminant D.

Proposition 2.5. *Algorithm 1 returns failure or the correct solution in time approximately* $\log_{10}(m/2)$

Proof. For the correctness we note that if $r_k \leq \sqrt{m}$ we have

$$r_k^2 + d\sqrt{\frac{m - r_k^2}{d}}^2 = r_k^2 + d\frac{m - r_k^2}{d} = m$$

so the solution is correct.

Algorithm 1: Cornaccia(m, d)

Input: Integers m, d
Output: A tuple (x, y) satisfying $x^2 + dy^2 = m$

Find r_0 such that $r_0^2 \equiv -d \pmod{m}$;
if $r_0 > m/2$ **then**
  $r_0 \leftarrow m - r_0$;
end
$k \leftarrow 1$;
$r_1 \leftarrow m \mod r_0$;
while $r_k^2 > m$ **do**
  $k \leftarrow k + 1$;
  $r_k \leftarrow r_{k-2} \mod r_{k-1}$;
end
$s \leftarrow \sqrt{\frac{m - r_k^2}{d}}$;
if *s is not integer* **then**
  Select another r_0 and repeat. If second iteration fails, then return failure ;
end
return (r_k, s)

Furthermore, for the running time, we notice that the algorithm we are performing is really just the Euclidian algorithm starting at (r_0, m), except that we end earlier. Since $r_0 < m/2$ and the running time of the Euclidian algorithm is bounded by the number of digits in base 10 we have that the running time of this algorithm is bounded by $\lceil \log_{10}(m/2) \rceil$ plus the time required to compute square roots modulo m and potentially running the algorithm twice if the first root was unsuccessful. □

Thus in general, this algorithm is considered to be of low computational cost.

2. Lattices

Much of the theory related to elliptic curves is related to lattices. We begin by describing them algebraically for vector fields along some useful results like how we can view them locally instead of globally. Then we explain a method for representing them with a useful basis before we finally look at lattices in the complex plane $\mathbb{C}$.

Let R be a domain with field of fractions F and V be a finitely dimensional vector space over F. For our purposes we will mostly have $R = \mathbb{Z}$, $F = \mathbb{Q}$ and V be of dimension 2 or 4 over $\mathbb{Q}$.

Definition 2.6. An R-**Lattice** in V is a finitely generated R-submodule M of V.

If in addition $MF = V$ we say that the R-lattice is **full**. When $R = \mathbb{Z}$ we will simply call it a lattice instead of saying $\mathbb{Z}$-lattice. It is clear that lattices are fractional ideals of $\mathbb{Z}$.

REMARK. Some texts define a lattice to be a discrete subgroup of $\mathbb{R}^n$. To make equivalent definitions we would need to modify our definition to let the lattice M be a finitely generated $\mathbb{Z}$-submodule of V satisfying $M\mathbb{R} = V$.

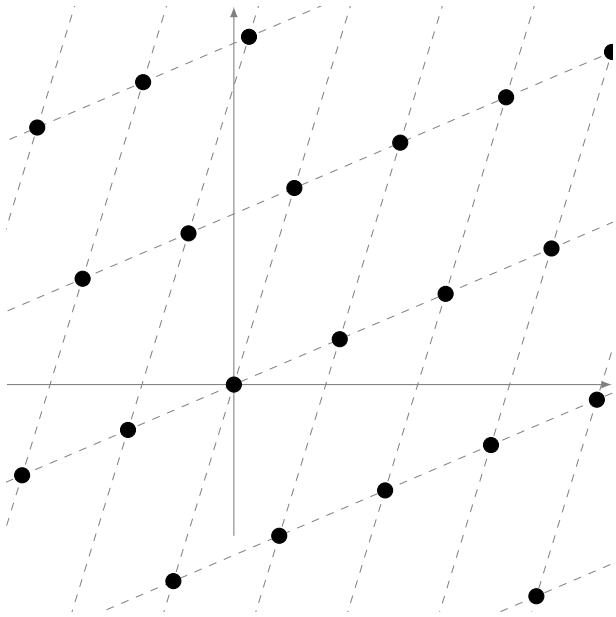

FIGURE 1. Lattice generated by $1.4 + 0.6i$ and $0.8 + 2.6i$ in $\mathbb{Q}(i)$

EXAMPLE 2.7. Let V be $\mathbb{Q}(i) = \mathbb{Q}(\sqrt{-1})$. Then the subset of V consisting of integer linear combinations of $1.4 + 0.6i$ and $0.8 + 2.6i$ is a full lattice in V as illustrated by the black dots in Figure 1. This follows since it, by definition, is of the form $\mathbb{Z}x_1 + \mathbb{Z}x_2$ and the generators are linearly independent, thus spanning V when coefficients from $\mathbb{Q}$ are allowed.

LEMMA 2.8. *Let $M \subseteq V$ be a full lattice and J be a finitely generated $\mathbb{Z}$-submodule of B, then the following statements hold:*

(1) For all $x \in B$ there exists a nonzero $r \in \mathbb{Z}$ such that $rx \in M$
(2) There exists a nonzero $r \in \mathbb{Z}$ such that $rJ \subseteq M$
(3) J is a full lattice if and only if there exists a nonzero $r \in \mathbb{Z}$ such that $rM \subseteq J \subseteq r^{-1}M$

PROOF. This proof is inspired by [39, Lemma 9.3.5]

(1) Writing x in the basis emitted by M we have $x = ax_1 + bx_2 + cx_3 + dx_4$ with $a, b, c, d \in \mathbb{Q}$. Letting z be the least common denominator of a, b, c, d we have $zx = a'x_1 + b'x_2 + c'x_3 + d'x_4$ where each $a', b', c', d' \in \mathbb{Z}$ so, since $\mathbb{Z}M = B$, we have $zx \in M$.

(2) Let $y_1, \dots y_n$ be the generators of J. For each generator we have an $r_i \in \mathbb{Z}$ such that $r_i y_i \in M$. Simply writing $r = \prod r_i$ gives us an element of $\mathbb{Z}$ satisfying $rJ \subseteq M$.

(3) By 2 we already have $rJ \subseteq M$, giving us $J \subseteq r^{-1}M$. Using the exact same argument we can find $r' \in \mathbb{Z}$ such that $r'M \subseteq J$. Clearly $r'M \subseteq M$ so we get

$$rr'M \subseteq r'M \subseteq J \subseteq r^{-1}M \subseteq (rr')^{-1}M$$

□

One rather powerful feature of lattices is the local-global principle. To understand them we need to introduce localizations.

Definition 2.9. A **Localization** of $\mathbb{Z}$ away from a prime p is

$$\mathbb{Z}_{(p)} := \{a/b \in \mathbb{Q} \mid p \nmid b\}$$

We can extend this to lattices through tensoring, where we define

$$M_{(p)} := M \otimes_{\mathbb{Z}} \mathbb{Z}_{(p)} \subseteq B$$

as the localization of M away from p where we identify B as $M \otimes_{\mathbb{Z}} \mathbb{Q}$. This is indeed an example of a full $\mathbb{Z}_{(p)}$-lattice in B.

Lemma 2.10. *Let M be an $\mathbb{Z}$-lattice in V, then we have*

$$M = \bigcap_p M_{(p)}$$

Proof. We give a proof based on [39, Lemma 9.4.6]. Since $M \subseteq M_{(p)}$ for any p by definition we have $M \subseteq \bigcap_p M_{(p)}$.

For the other direction, suppose $\alpha \in B$ such that $\alpha \in \bigcap_p M_{(p)}$. Then the set $S = \{r \in \mathbb{Z} \mid r\alpha \in M\}$ is an ideal of $\mathbb{Z}$. This is easily shown since if $r\alpha \in M$ then $\mathbb{Z}r\alpha \in \mathbb{Z}M \subseteq M$ as M is a $\mathbb{Z}$-module. Furthermore, given $r, r' \in S$ then $(r - r')\alpha = r\alpha - r'\alpha \in M$ since M is closed under addition. Now, since $x \in M_{(p)}$, we can clear its denominators by multiplying with some $r_p \in \mathbb{Z} \setminus (p)$. That is $r_p x \in M$, so $r_p \in S$. This applies to every prime ideal (p) which are all maximal in $\mathbb{Z}$, so in particular S is not properly contained in any maximal ideal giving us that $S = \mathbb{Z}$. Since $1 \in \mathbb{Z}$ we have $x \in M$. □

Corollary 2.11. *Let M, N be lattices in B. Then the following are equivalent*

(1) $M \subseteq N$
(2) $M_{(p)} \subseteq N_{(p)}$ for all primes $p \in \mathbb{Z}$

Proof. This follows immediately from Lemma 2.10. □

THEOREM 2.12 (Local Global dictionary of lattices). *Let B be a finite-dimensional $\mathbb{Q}$-vector space and let $M \subseteq B$ be a full lattice, then the map $N \mapsto (N_{(p)})_p$ is a bijection between full lattices N and collections of lattices $(N_{(p)})$ indexed by p where $M_{(p)} = N_{(p)}$ for all but finitely many primes p.*

PROOF. Inspired by [39, Theorem 9.4.9] Given the two lattices M, N we know that there exists some $r \in \mathbb{Z}$ such that $rM \subseteq N \subseteq \frac{1}{r}M$. Factoring r into primes we see that there are only finitely many prime ideals of $\mathbb{Z}$ where r is contained. Everywhere else we get $(rM)_{(p)} = M_{(p)}$ and similarly $(\frac{1}{r}M)_{(p)} = M_{(p)}$ so we have $M_{(p)} = N_{(p)}$ for all but finitely many primes p.

Next suppose we have a collection of lattices $(N_{(p)})_p$ (a priori not related to some lattice N) that satisfy $N_{(p)} = M_{(p)}$ for all but finitely many primes p, then we define N to be the $\mathbb{Z}$-submodule $N = \bigcap_p N_{(p)}$ of V. To show that N is a full lattice we need to find $r \in \mathbb{Z}$ such that $rM \subseteq N \subseteq r^{-1}M$. At every p where $M_{(p)} \neq N_{(p)}$ we have $r_p M_{(p)} \subseteq N_{(p)} \subseteq r_p^{-1} M_{(p)}$. Thus the integer $r = \prod r_p$ is a good candidate. Indeed, at the primes of inequality we still have

$$rM_{(p)} = r_p M_{(p)} \subseteq N_{(p)} \subseteq r_p^{-1} M_{(p)} = r^{-1} M_{(p)}$$

whilst at the primes of equality we already have $M_{(p)} = N_{(p)}$ and in particular

$$rM_{(p)} = M_{(p)} = N_{(p)} = M_{(p)} r^{-1} M_{(p)}$$

Thus by the above corollary $rM \subseteq N \subseteq r^{-1}M$ and N is a full lattice. We have just shown that given a collection of lattices that equal $M_{(p)}$ at almost every prime p, then we have a corresponding full lattice.

To show that the sets are bijective, let us start with a full lattice N, then the map $N \to (N_{(p)})_p$ maps N to a collection of lattices with inverse map $(N_{(p)})_p \mapsto \bigcap_p N_{(p)} = N$ as in the above Lemma. Conversely, given the collection $(N_{(p)})_p$, the maps are inverses since

$$\left(\bigcap_p N_{(p)}\right)_{(q)} = \left(\bigcap_p N \otimes_{\mathbb{Z}} \mathbb{Z}_{(p)}\right) \otimes_{\mathbb{Z}} \mathbb{Z}_{(q)} = \left(\bigcap_{p \neq q} N \otimes_{\mathbb{Z}} \mathbb{Q}\right) \bigcap (N \otimes_{\mathbb{Z}} \mathbb{Z}_{(q)}) = (N_{(q)})$$

□

DEFINITION 2.13. The **index** of N in M, for lattices N, M, is written $[M : N]$ and is the $\mathbb{Z}$-submodule of $\mathbb{Q}$ generated by the set

$$\{\det(\delta) \mid \delta \in \mathrm{End}_F(B) \text{ and } \delta(M) \subseteq N\}$$

Since $\mathbb{Z}$ is a PID this submodule is generated by a single integer which we will sometimes identify as the index. It will correspond to the index of abelian groups, namely $\#(M/N)$.

LEMMA 2.14. *If M and N are free as $\mathbb{Z}$-modules, then $[M : N]$ is a free $\mathbb{Z}$-module generated by the determinant of any $\delta \in \mathrm{End}_{\mathbb{Q}}(V)$ giving a change of basis from M to N*

Proof: See [39, Lemma 9.6.4]

EXAMPLE 2.15. Let M the the lattice generated by $1, i, (1+ij)/2$ and $(i+j)/2$ and N be the lattice generated by $1, i, j, ij$. We would like to compute the index of N inside M ($[M : N]$). We do this using the lemma above by finding a change of basis map. This follows easily when we write the generators of M as vectors in V with basis $1, i, j, ij$. We get the endomorphism

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & -1 & -1 \\ 0 & 0 & 0 & 2 \\ 0 & 0 & 2 & 0 \end{bmatrix} : M \to N$$

Which clearly has determinant -4 (swap the two last rows and look at the diagonal).

2.1. Bases. As lattices are finitely generated $\mathbb{Z}$-submodules we can look further into how to represent them. There are several bases which are useful, and the orthonormal one is perhaps the most known. This section is mainly about introducing the Minkowski reduced basis which we shall use later. The definitions are taken from [5].

DEFINITION 2.16. Let M be a lattice of rank m. For $i \in \{1, \ldots, m\}$ we define the **ith successive minimum** as:

$$\lambda_i(M) = \inf\{r \mid \dim(\text{span}(M \cap B(0, r])) \geq i\}$$

Where $B(0, r]$ is the closed ball of radius r centered at 0. In other words, $\lambda_i(M)$ is the smallest radius such that a ball centered at the origin contains at least i linearly independent elements of M. If $i = 1$ we have $\lambda_i(M)$ equal to the norm of the shortest vectors of M.

DEFINITION 2.17. An ordered basis $(b_1, \ldots, b_m)$ is a **Minkowski reduced basis** of M if it is a basis of M and for every $i \in \{1, \ldots, m\}$, there exists no element b_i' of norm less than b_i such that $\{b_1, \ldots, b_i'\}$ form a linear independent set.

What remains is to show how we can compute the Minkowski reduced basis. To do this we also need to recall what a Gram matrix is.

DEFINITION 2.18. The **Gram matrix** of a set of vectors $b_1, \ldots b_m$ is the matrix G with entries $G_{ij} := \langle b_i, b_j \rangle$ - that is the inner product of the vectors.

One can compute the Gram matrix in time $O(\log^2 ||b_m||)$ for fixed m [28, Section 3.2].

THEOREM 2.19. *Let $1 \leq d \leq 4$. Given as input an ordered basis $(b_1, \ldots, b_d)$ of the lattice L, Algorithm 2 outputs a minkowski reduced basis in*

$$O\left(\log(||b_d||)[1 + \log(||b_d||) - \log(\lambda_1(L))]\right)$$

bit operations

Proof: See [28, Theorem 6]

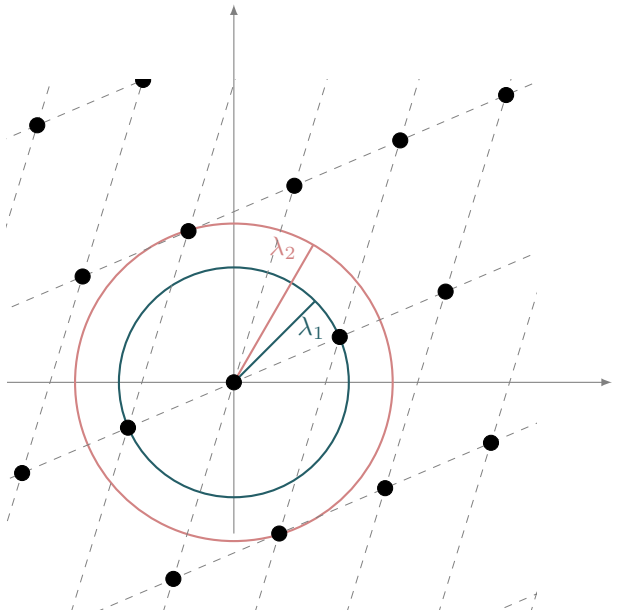

FIGURE 2. Ith successive minima of lattice defiend by $[(1, 0.4), (0.3, 1.3)]$.

2.2. Complex lattices. In this section we will describe complex lattices similarly to how we defined R-lattices, define what homothetic means, describe the Wierstrass function and finally find some invariants which will be used to describe these lattices. This section is based on [36, Chapter VI].

Although it isn't entirely compatibly with our definition we say that a **complex lattice** is a full $\mathbb{Z}$-lattice in $\mathbb{C}$. That is a finitely generated $\mathbb{Z}$-submodule Λ of $\mathbb{C}$ such that $\Lambda\mathbb{R} = \mathbb{C}$. We will use the notation $\Lambda \subseteq \mathbb{C}$ to denote such lattices. There is no problem with this inconsistency as we will not use the results from the previous section on lattices. Throughout this section we will simply refer to complex lattices as lattices.

We can view the lattices Λ as an additive subgroup of $\mathbb{C}$ and thus look at the quotient $\mathbb{C}/\Lambda$ where we identify z with z' if $z = z' + \omega$ for some $\omega \in \Lambda$. This looks like a torus giving us a nice way of thinking of these spaces which we will later connect to ordinary elliptic curves.

Two lattices Λ_1, Λ_2 are said to be **homothetic** if there exists some $0 \neq \alpha \in \mathbb{C}$ such that $\Lambda_1 = \alpha\Lambda_2$. Often we are only interested in lattices up to homothethy and this simply means that two lattices are considered equivalent if they are homothetic.

For each lattice Λ we define its **Wierstrass** $\wp$-function as

$$\wp(z) = \frac{1}{z^2} + \sum_{0 \neq \omega \in \Lambda} \left(\frac{1}{(z-\omega)^2} - \frac{1}{\omega^2} \right)$$

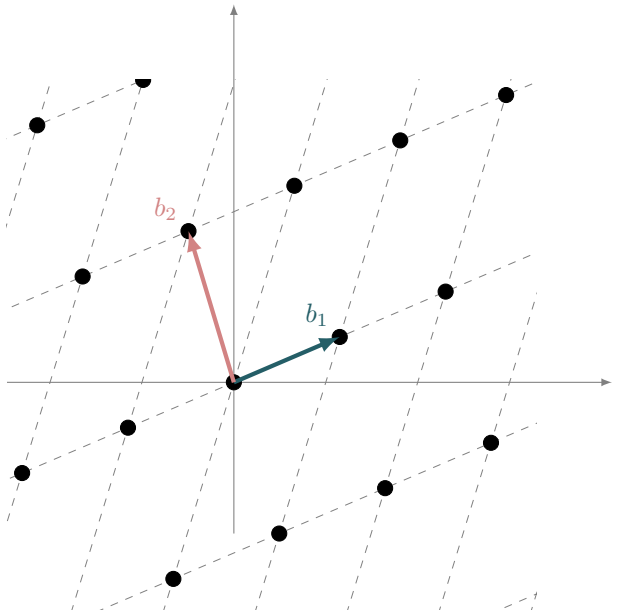

FIGURE 3. A Minkowski reduced basis $\{b_1, b_2\}$ of lattice $[(1, 0.4), (0.3, 1.3)]$

Algorithm 2: Greedy$(b_1, \ldots, b_m)$

Input: A basis $[b_1, \ldots b_m]$ and its Gram matrix
Output: A greedy-reduced basis $(b_1, \ldots, b_m)$

if $d = 1$ **then**
  return b_1
end
repeat
  Order $(b_1, \ldots, b_m)$ by increasing length and update Gram matrix ;
  $(b_1, \ldots, b_{m-1}) \leftarrow$Greedy$(b_1, \ldots, b_{m-1})$;
  Compute vector c closest to $b_d \in \{\mathbb{Z}b_1 + \ldots + \mathbb{Z}b_{m-1}\}$;
  $b_m \leftarrow b_m - c$ and update Gram matrix.
until $||b_m|| \geq ||b_{m-1}||$;
return $(b_1, \ldots, b_m)$

Furthermore we have the **Eisenstein series of weight** $2k$

$$G_{2k}(\Lambda) = \sum_{0 \neq \omega \in \Lambda} \omega^{-2k}$$

By performing the Laurent series expansion of $\wp$ we get the relation

$$\wp'(z)^2 = 4\wp(z)^3 - 60G_4\wp(z) + 140G_6$$

Where the quantities $60G_4$ and $140G_6$ are denoted g_2 or $g_2(\Lambda)$ and g_3 or $g_3(\Lambda)$ respectively. These values turn out to describe the lattice in a nice way since given a lattice Λ, we can create its Wierstrass function, and subsequently find the values g_2 and g_3 from the above equation. Furthermore, the Uniformization theorem [35, Chapter I, Corollary 4.3] states that given two values $A, B \in \mathbb{C}$ there is a unique lattice Λ which satisfy $g_2(\Lambda) = A$ and $g_3(\Lambda) = B$.

3. Quaternion Algebra

This section is based on John Voight's Quaternion Algebras [39, Part I, II] where only the most relevant things are extracted. I would highly recommend reading the book as it gives an excellent introduction to the subject. Throughout this section we will define a quaternion algebra, talk about conjugation or the standard involution of elements, define the trace and norm in the language of quaternions, introduce quaternion orders and ideals, describe some invariants for quaternion algebras and orders, quickly look at ideal classes, and introduce orthogonality. Finally we will provide an algorithm for constructing specific kinds of ideals. The concepts introduced and explained here will be particularly interesting when we later look into supersingular elliptic curves in Chapter 4.

The rings we are using will always contain the multiplicative identity 1, so homomorphisms necessarily preserve 1. If R is a ring we will write $R^\times$ for the group of multiplicative of units of R. When we say **algebra** over some field F we mean a ring B that has a homomorphism $f : F \to B$ such that $f(F) \subseteq Z(B)$ where $Z(B)$ is the center of B - the set of all elements of $\beta \in B$ such that $\beta\alpha = \alpha\beta$ for every $\alpha \in B$. Instead of saying that B is an algebra over F we often just say that B is an F-algebra.

A **central** algebra is one where $f(F) = Z(B)$. Furthermore the map f is injective as $f(1) = 1$, so F can be identified with $f(F)$, giving rise to a F-vector space structure of B, allowing us to talk about the dimension of B as a F-vector space. We will write $\dim_F(B)$ for this dimension.

Example 2.20. Let $F = \mathbb{Q}$ be the field of rational numbers, then $B = M_2(\mathbb{Q})$ is an $\mathbb{Q}$-algebra of dimension 4.

Let $\phi : B \to B'$ be a ring ring homomorphism of F-algebras B, B'. We can then define the following maps. ϕ is an F-algebra **homomorphism** if $\phi|_F = \mathrm{id}_F$. An F-algebra **endomorphism** is an homomorphism where $B = B'$. An F-algebra **isomorphism** is an homomorphism that is also invertible. And finaly, an F-algebra **automorphism** is an endomorphism that is also invertible.

The write $\mathrm{End}_F(B)$ for the set of all F-algebra endomorphisms of B. This is in fact a ring where we define the operations as the usual function compositions $(\phi + \psi)(\alpha) = \phi(\alpha) + \psi(\alpha)$ and $(\phi \circ \psi)(\alpha) = \phi(\psi(\alpha))$.

Recall that a **division ring** is a ring where every non-zero element has a two-sided inverse, therefore we define a **division algebra** to be an algebra that is also a division ring.

DEFINITION 2.21. A **quaternion algebra** B is an F-algebra where there exists $i, j \in B$ such that $1, i, j, ij$ forms an F-basis for B which satisfy

$$i^2 = a, j^2 = b, \text{ and } ij = -ji$$

with $a, b \in F^\times$

For simplicity, we sometimes write $\left(\frac{a,b}{F}\right)$ or just $(a, b \mid F)$ for the quaternion algebra as defined above.

The ring $M_2(F)$ is a the quaternion algebra $(1, 1 \mid F)$ where we identify i with $\left[\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right]$ and j with $\left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right]$. In fact, we can always view a quaternion algebra as a subalgebra of matrices:

PROPOSITION 2.22. *Let $B = (a, b \mid F)$ and $F(\sqrt{a})$ be the splitting field of $x^2 - a$ over F. Then we have a map*

$$\begin{aligned} \lambda : B &\to M_2(F(\sqrt{a})) \\ i, j &\mapsto \begin{bmatrix} \sqrt{a} & 0 \\ 0 & -\sqrt{a} \end{bmatrix}, \begin{bmatrix} 0 & b \\ 1 & 0 \end{bmatrix} \\ t + xi + yj + zij &\mapsto \begin{bmatrix} t + x\sqrt{a} & b(y + z\sqrt{a}) \\ y - z\sqrt{a} & t - x\sqrt{a} \end{bmatrix} \end{aligned}$$

that is an injective F-algebra homomorphism which is isomorphic onto its image.

Proof: See [39, Proposition 2.3.1]

Finally we state the two first results of the main theorem of quaternion algebras.

THEOREM 2.23 (Main theorem of quaternion algebras). *Let $B = (a, b \mid F)$ be a quaternion algebra with $\operatorname{char}(F) \neq 2$. Then the following are equivalent*

(1) $B \cong (1, 1 \mid F) \cong M_2(F)$
(2) B is not a division ring

Proof: See [39, Theorem 5.4.4]

3.1. Conjugation/Involution. Just like complex conjugation has several applications when dealing with complex numbers, the conjugation of quaternion elements turns out the be very useful for us. In this section we will describe the uniqueness of the standard involution of quaternion algebras and how it can be used to defined trace and norm of quaternion elements. Later we shall see that the dual map of an isogeny satisfies the properties of being a standard involution which, by the uniqueness, forces it to be the standard involution.

We begin with the definition of an involution

DEFINITION 2.24. A **standard involution** $f : B \to B$ of an F-algebra B is an F-linear map that satisfies the following properties:

(1) $f(1) = 1$
(2) $f(f(\alpha)) = \alpha$ for all $\alpha \in B$
(3) $f(\alpha\beta) = f(\beta)f(\alpha)$ for all $\alpha, \beta \in B$
(4) $\alpha f(\alpha) \in F$ for all $\alpha \in B$

REMARK. If we discard the last property we have the definition of an involution.

EXAMPLE 2.25. In the complex case, $\mathbb{C}$, we can simply take the conjugation map $f(\alpha) := \overline{\alpha}$. Simple verification shows that the properties always satisfied.

PROPOSITION 2.26. *Let $B = (a, b|F)$ be a quaternion algebra with generators $1, i, j, k$. Then the map*

$$\bar{\cdot} : \alpha = t + xi + yj + zk \mapsto t - xi - yj - zk = \overline{\alpha}$$

is a standard involution on B

PROOF. The properties are verified directly using elementary algebra. We start by checking F-linearity by taking $\lambda \in F$ and $\alpha = t + xi + yj + zk \in B$

$$\begin{aligned}\overline{\lambda\alpha} &= \overline{\lambda t + \lambda xi + \lambda yj + \lambda zk} = \lambda t - \lambda xi - \lambda yj - \lambda zk \\ &= \lambda(t - zi - yj - zk) = \lambda\overline{\alpha}\end{aligned}$$

Similarly, with $\beta = t' + x'i + y'j + z'k \in B$

$$\overline{\alpha + \beta} = t + t' - (x + x')i - (y + y')j - (z + z')k = \overline{\alpha} + \overline{\beta}$$

Next we see that $\overline{1} = 1$. Then we look at the repeated conjugation

$$\overline{\overline{\alpha}} = \overline{t - xi - yj - zk} = t + xi + yj + zk = \alpha$$

The third and fourth property requires a slightly more cumbersone computation:

$$\begin{aligned}\alpha\beta = (tt' + axx' + yy'b - abzz') + (tx' + xt' - byz' + bzy')i \\ + (ty' + axz' + yt' - azx')j + (tz' + xy' - yx' + zt')k\end{aligned}$$

and similarly

$$\begin{aligned}\overline{\beta}\overline{\alpha} = (tt' + axx' + byy' - abzz') + (-xt' - tx' - bzy' + byz')i \\ + (-ty' - axz' - yt' + azx')j + (-tz' - xy' + yx' - zt')k\end{aligned}$$

Which gives us that $\overline{\alpha\beta} = \overline{\beta}\overline{\alpha}$. Furthermore, if $\beta = \overline{\alpha}$, we do in particular have that $t = t'$, $x' = -x$, $y' = -y$ and $z' = -z$, which gives us that $\alpha\beta = t^2 + ax^2 + by^2 - abz^2$, an element of F since $a, b, t, x, y, z \in F$. □

Once we have a standard involution f, we can define the reduced trace and reduced norm on B. A priori, the definitions might seem to depend on the particular choice of involution, but as we will see shortly there is in fact no such choice.

The **reduced trace** $\mathrm{trd} : B \to F$ of a standard involution f is the map

$$\alpha \mapsto \alpha + f(\alpha)$$

Similarly we define the **reduced norm** $\mathrm{nrd} : B \to F$ of a standard involution f is the map

$$\alpha \mapsto \alpha f(\alpha)$$

LEMMA 2.27. *If B is a nonzero F-algebra with standard involution f, then $\alpha \in B$ is a unit if and only if* $\mathrm{nrd}\,\alpha \neq 0$

Proof, see[39, Part I, Exercise 3.5]

REMARK. The reason for calling them reduced is their relation to the algebra trace and norm where we have $\mathrm{Tr}(\alpha) = 2\,\mathrm{trd}(\alpha)$ and $\mathrm{Nr}(\alpha) = \mathrm{nrd}(\alpha)^2$. Intuitively this makes sense as the quaternion algebra is of dimension 4 giving us four automorphisms. The algebra norm would then be the product of four elements while our definition is the product of two. This is however not important for us, so we will not show this fact.

To show uniqueness we start off with a lemma

LEMMA 2.28. *Let K be an F-algebra of dimension 2 over F. Then K is commutative and has a unique standard basis.*

PROOF. We follow [39, Lemma 3.4.2]. First let $\alpha \in K \setminus F$ be an arbitrary element. This exists because K has dimension 2. We decompose K to $K = F \oplus F\alpha$. To show that K is commutative we take any two elements $(a, b), (c, d) \in F \oplus F\alpha$ and verify that we have $(a, b)(c, d) = (ac, bd) = (ca, db) = (c, d)(a, b)$ since F is a field.

Next we can always write $\alpha^2 = t\alpha - n$ for unique $t, n \in F$ since $1, \alpha$ forms a basis of K and $\alpha^2 \in K$. Now Let $\bar{\cdot} : K \to K$ be any standard involution. Noticing that

$$\alpha^2 - (\alpha + \overline{\alpha})\alpha + \alpha\overline{\alpha} = 0$$

and keeping in mind that $\alpha\overline{\alpha} \in F$ we have our unique decomposition with $t = \mathrm{trd}(\alpha)$ and $n = \mathrm{nrd}(\alpha)$. Thus in particular, the involution must satisfy $t = \alpha + \overline{\alpha}$ so we have $\overline{\alpha} = t - \alpha$.

Finally we need to show that this requirement, that $\overline{\alpha} = t - \alpha$, is enough to determine the involution on any element of K. Let $\beta \in K$, then $\beta = a + b\alpha$ with unique $a, b \in F$. Giving us

$$\overline{\beta} = \overline{a + b\alpha} = \overline{a} + b\overline{\alpha} = a + b(t - \alpha)$$

Where we have only used the facts that $\overline{F} = F$ and that the involution is F-linear, and that $\overline{\alpha} = t - \alpha$. □

COROLLARY 2.29. *If an F-algebra B has a standard involution, then this involution is unique.*

PROOF. Let $\alpha \in B \setminus F$. Then $\alpha^2 - \mathrm{trd}(\alpha)\alpha + \mathrm{nrd}(\alpha) = 0$, so we have that the F-dimension of $F[\alpha]$ is 2. Restricting the standard involution on B to $F[\alpha]$ gives us a unique involution by the above lemma. This happens for every $\alpha \in B \setminus F$ and for $\alpha \in F$ we have $\overline{\alpha} = \alpha$, so the involution must be unique. □

As a consequence, the standard involution we defined: $t + xi + yj + zk \mapsto t - xi - yj - zk$ is the standard involution of a quaternion algebra with basis i, j, k. The definitions of the reduced trace and reduced norm becomes well-defined and not dependent on the choice of involution.

3.2. Orders. In this and the following sections we will be more concrete by considering quaternion algebras over $\mathbb{Q}$ instead of some arbitrary field. This allows us to simplify some results and move faster through the theory. We will in particular use the fact that the ring $\mathbb{Z}$ is PID (and thus noetherian). This is nevertheless no problem for us since we are only interested in quaternions over $\mathbb{Q}$ and the $\mathbb{Z}$-ideals and $\mathbb{Z}$-orders that arise.

Since we have already dedicated an entire section for lattices we will not repeat the information here. Instead we note that lattices of a quaternion algebra B are simply $\mathbb{Z}$-submodules of B when we consider B as a vector space over $\mathbb{Q}$.

EXAMPLE 2.30. Let M be the lattice generated by 1, $2i$, $3j$ and ij with $i^2 = -1$ and $j^2 = -p$ for some prime p. Then this is a full lattice in some quaternion algebra, but it is not closed under multiplication. Take for example $2i \cdot ij = -2j$. This cannot be written as a linear combination of the generators 1, $2i$, $3j$ and ij. Therefore this is not an order.

As lattices are submodules they only require closure under addition. Sometimes we are however interested in this multiplicative structure as well. Therefore we introduce the notion of orders, which are lattices closed under multiplication.

DEFINITION 2.31. An **order** $O \subseteq B$ is a full lattice O that is also a subring of B.

EXAMPLE 2.32. The lattice generated by the set $1, i, j, ij$ is an order.

Given a lattice I we can create its **left order**

$$O_L(I) := \{\alpha \in B \mid \alpha I \subseteq I\}$$

and similarly its **right order**

$$O_R(I) := \{\alpha \in B \mid I\alpha \subseteq I\}$$

Since B is non-commutative, these orders need not be the same.

LEMMA 2.33. *$O_L(I)$ is an order*

PROOF. Inspired by [39, Lemma 10.2.7].

We need to show two things. That $O_L(I)$ is a full lattice and that it is a subring of B. To show the subring property all we need is that $\alpha - \beta \in O_L(I)$ and $\alpha\beta \in O_L(I)$ for all $\alpha, \beta \in O_L(I)$. This follows from easy inspections. We have $(\alpha\beta)I = \alpha(\beta I) \subseteq \alpha I \subseteq I$ and

$$(\alpha - \beta)I = \{\alpha\gamma - \beta\gamma \mid \gamma \in I\} = \{\alpha' - \beta' \mid \alpha', \beta' \in I\} \subseteq I$$

Where the last equality follows from $\alpha\gamma, \beta\gamma \in I$, so $O_L(I)$ is a subring of B.

To show that it is a full lattice we will use the properties of Lemma 2.8 quite a few times. We need it to be finitely generated and satisfy $\mathbb{Q}O_L(I) = B$. First, for every $\alpha \in B$, αI is still finitely generated. By the lemma we have a nonzero $r \in \mathbb{Z}$ satisfying $r(\alpha I) \subseteq I$. Since $r \in \mathbb{Z}$ it commutes with α and we have $\alpha(rI) \subseteq I$ giving us $O_L(I)\mathbb{Z} = I\mathbb{Z}$ and thus $O_L(I)\mathbb{Q} = I\mathbb{Q} = B$.

Second, to show that it is finitely generated we use the fact that $1 \in B$, so we must have a nonzero $r \in \mathbb{Z}$ such that $r = r \cdot 1 \in I$. Taking an arbitrary $\alpha \in O_L(I)$ we already have that $\alpha I \subseteq I$, but since $r \in I$ we get $\alpha r \subseteq I$ so we have $O_L(I)s \subseteq I$ giving us $O_L(I) \subseteq s^{-1}I$ (a nonzero integer has inverse in $\mathbb{Q}$). Since $s^{-1}I$ is a finitely generated $\mathbb{Z}$-module, and $O_L(I)$ is a submodule of $s^{-1}I$ it is also finitely generated since $\mathbb{Z}$ is noetherian. □

We say that an element $\alpha \in B$ is integral if it satisfies a monic polynomial with coefficients in $\mathbb{Z}$. This could be defined in any of the following ways.

LEMMA 2.34. *Let $\alpha \in B$, then the following are equivalent.*

(1) α is integral
(2) $\mathbb{Z}[\alpha]$ is a finitely generated $\mathbb{Z}$-module
(3) α is contained in a subring A that is finitely genereated as a $\mathbb{Z}$-module

Proof: See [39, Lemma 10.3.2]

COROLLARY 2.35. *If O is an order of B, and $\alpha \in O$, then α is integral*

PROOF. Follows immediately from the above lemma as O is a finitely generated $\mathbb{Z}$ module of B that is also a subring. □

DEFINITION 2.36. A **maximal** order O is an order that is not properly contained in any other order.

LEMMA 2.37. *An order O of B is maximal if and only if $O_{(p)}$ is a maximal $\mathbb{Z}_{(p)}$ order for every prime p*

Proof: See [39, Lemma 10.4.3]

3.3. Ideals. In general we are not interested in ideals of B, but rather ideals of some order O. This is just the classical notion of ideals from algebra. We must however separate left ideals from right ideals as O (and B) is non-commutative.

Given two lattices I, J, we say that I is **compatible** with J if the right order of I is equal to the left order of J. That is $O_R(I) = O_L(J)$.

Furthermore, we say that I is **right invertible** if there exists some lattice I' such that $II' = O_L(I)$. The lattice I' is called the **right inverse**. Similarly we say that I is **invertible** if it has both a left and a right inverse. That is

$$II' = O_L(I) = O_R(I') \text{ and } I'I = O_L(I') = O_R(I)$$

If a lattice has an inverse then $I^{-1} := \{\alpha \in B \mid I\alpha I \subseteq I\}$ is the unique inverse.

For an order O, we say that a **Left fractional O-ideal** is a lattice $I \subseteq B$ such that $O \subseteq O_L(I)$. Similarly a **Right fractional O-ideal** is a lattice such that $O \subseteq O_R(I)$.

REMARK. Note that O already has a ring structure, so we are still able to talk about ideals in the usual sense. We will for certain have $O \subseteq O_L(I)$ for any ordinary ideal I of O, so I is also a fractional ideal.

PROPOSITION 2.38. *Let O be a maximal order, then every left or right fractional O-ideal is invertible.*

PROOF. Let I be a left fraction ideal (the proof is the same for right fractional ideals). Notice that I^{-1} is always defined and $II^{-1} \subseteq O_L(I)$. The last part is easily seen by taking $\alpha \in I$ and $\beta \in I^{-1}$, then $\alpha\beta I \subseteq I$ by the definition of I^{-1} so $\alpha\beta \in O_L(I)$.

We have $O \subseteq II^{-1}$ as $1 \in I^{-1}$ and every $\alpha \in O$ satisfy $\alpha I \subseteq I$ since $O \subseteq O_L(I)$. Furthermore, since O is maximal we have $O = II{-}1 = O_L(I) = O$, so I is invertible. □

A lattice I is **principal** if there exists $\alpha \in B$ such that

$$I = O_L(I)\alpha = \alpha O_R(I)$$

and we say that I is **generated** by α.

DEFINITION 2.39. A lattice I is **integral** if $I^2 \subseteq I$

LEMMA 2.40. *$O_L(I) = O_R(\overline{I})$ and $O_R(I) = O_L(\overline{I})$*

PROOF. [39, Lemma 16.6.7]

This follows from the definitions as $\alpha \in O_L(I)$ if and only if $\alpha I \subseteq I$ if and only if $\overline{\alpha I} = \overline{I}\overline{\alpha} \subseteq \overline{I}$ if and only if $\overline{\alpha} \in O_R(\overline{I})$. But since Orders are closed $\overline{\alpha} \in O_R(\overline{I})$ implies that $\alpha \in O_R(\overline{I})$ □

LEMMA 2.41. *Let I be a lattice. Then the following are equivalent:*

(1) I is integral
(2) For all $\alpha, \beta \in I$, we have $\alpha\beta \in I$
(3) $I \subseteq O_L(I)$
(4) $I \subseteq O_R(I)$
(5) $I \subseteq O_L(I) \cap O_R(I)$

PROOF. Taken from [39, Lemma 16.2.8].

(1) and (2) are equivalent by the definition of an integral lattice. To show that (1) and (3) are equivalent we first notice that since $II \subseteq I$, then $I \subseteq O_L(I)$. Similarly, if $I \subseteq O_L(I)$, then $II \subseteq I$. The same argument goes for the equivalence of (1) and (4). Then (3) and (4) are equivalent to (5). □

REMARK. Notice that by (3) I is a left $O_L(I)$-ideal in the usual sense. It is necessarily closed under left multiplication by $O_L(I)$ and already have the additive structure from being a lattice. Furthermore, since every element of O is integral, then our slightly different notion of integrality of lattices coincides with the standard notion.

Just like we defined the reduced norm of elements we define the reduced norm of lattices.

DEFINITION 2.42. The **reduced norm** $\mathrm{nrd}(I)$ of I is the $\mathbb{Z}$-submodule of $\mathbb{Q}$ generated by the set $\{\mathrm{nrd}(\alpha) \mid \alpha \in I\}$

LEMMA 2.43. *The reduced norm $\mathrm{nrd}(I)$ is a fractional ideal of $\mathbb{Q}$ - that is, it is finitely generated as an $\mathbb{Z}$-module*

Proof: See [39, Lemma 16.3.2]

In other words, we can write $\mathrm{nrd}(I) = aJ$ for $a \in \mathbb{Q}$ and $J \subset \mathbb{Z}$ an ideal. Since $\mathbb{Z}$ is a PID this is in turn generated by a single element $J = (b)$ for some $b \in \mathbb{Z}$, thus we can always write $\mathrm{nrd}(I) = a(b) = (c)$ for some $c \in \mathbb{Q}$. As a consequence we will rarely talk about the fractional ideal $\mathrm{nrd}(I) \subseteq \mathbb{Q}$, but rather the element $\mathrm{nrd}(I) \in \mathbb{Q}$.

Similarly, when I is integral, the reduced norm of its elements, $\mathrm{nrd}(\alpha)$, lie in $\mathbb{Z}$, so the reduced norm of I is actually just generated by the greatest common divisor of its elements. That is

$$\mathrm{nrd}(I) = \gcd(\{\mathrm{nrd}(\alpha) \mid \alpha \in I\})$$

If I is principal, generated by α, we have $\mathrm{nrd}(I) = \mathrm{nrd}(\alpha)\mathbb{Z}$. If $\alpha \in B$ we have $\mathrm{nrd}(\alpha I) = \mathrm{nrd}(\alpha)\,\mathrm{nrd}(I)$ by the multiplicative structure of nrd. That is

$$\mathrm{nrd}(\alpha I) = \{\mathrm{nrd}(\alpha\gamma) \mid \gamma \in I\} = \mathrm{nrd}(\alpha)\{\mathrm{nrd}(\gamma) \mid \gamma \in I\} = \mathrm{nrd}(\alpha)\,\mathrm{nrd}(I)$$

PROPOSITION 2.44. *Let O be a maximal order and I an left O-ideal, then $I\overline{I} = \mathrm{nrd}(I)O$.*

Proof: See [39, Section 16.6.14] adding that $O \subseteq O_L(I)$ implies $O = O_L(I)$ when O is maximal.

We finish this section with the main theorem of quaternion ideals.

THEOREM 2.45 (Main theorem of quaternion ideals). *Let B be a quaternion algebra over $\mathbb{Q}$ and $I \subseteq B$ be a lattice. Then the following are equivalent*

(1) I is locally principal ($I_{(p)} = I \otimes_{\mathbb{Z}} \mathbb{Z}_{(p)}$ is principal for all primes p)
(2) I is invertible
(3) I is right invertible
(4) I is left invertible
(5) $\mathrm{nrd}(I)^2 = [O_R(I) : I]$ and
(6) $\mathrm{nrd}(I)^2 = [O_L(I) : I]$

Proof: See [39, Theorem 16.7.7]

3.4. Ramification. One way to classify quaternion algebras is based on their ramification which we will use for describing the endomorphism ring of some elliptic curves later on.

We define the set of places to be the set of primes and ∞. These correspond to embeddings of $\mathbb{Q}$ into $\mathbb{Q}_p$ (the p-adic numbers) for the primes and $\mathbb{R}$ for ∞.

DEFINITION 2.46. A quaternion algebra B is **ramified** at a place v if $B_v := B \otimes_{\mathbb{Q}} \mathbb{Q}_v$ is a division ring. Otherwise B is **split** (unramified).

We write the set of ramified places as $\mathrm{Ram}(B)$ and by convention we represent our places with the primes and ∞.

DEFINITION 2.47. The **discriminant** of a quaternion algebra B is the integer

$$\mathrm{disc}(B) := \prod_{p \in \mathrm{Ram}'(B)} p \in \mathbb{Z}$$

Where $\mathrm{Ram}'(B)$ is the set of ramified places except ∞ if that is inside the set.

Note, in general one defines the discriminant to be an ideal, but this corresponds to a generating integer in $\mathbb{Z}$ as every ideal is principal

EXAMPLE 2.48. Let B be the quaternion algebra ramified at p and ∞, then $\mathrm{disc}(B) = p$.

The usefulness of ramification comes from the following proposition where we see that it is enough to look at the local fields $\mathbb{Q}_v$ instead of over $\mathbb{Q}$ when checking if two quaternion algebras are isomorphic.

PROPOSITION 2.49. *Let B, B' be quaternion algebras over $\mathbb{Q}$, then the following are equivalent*

(1) $B \cong B'$
(2) $\operatorname{Ram} B = \operatorname{Ram} B'$
(3) $B_v \cong B'_v$ *for all places v and —item $B_v \cong B'_v$ for all but one place v*

Proof: See [39, Proposition 14.3.1]

3.5. Discriminants. Similar to how we defined the discriminant of a quaternion algebra, we can look at the discriminant of lattices. Here the discriminant is a measure of volume and we will see why we used a similar name for the discriminant of a quaternion algebra shortly.

Suppose I is a lattice in B generated by $\{\alpha_1, \ldots, \alpha_n\}$, then we say that the **discriminant** of I is

$$\operatorname{disc}(I) = |\det(\operatorname{trd}(\alpha_i \alpha_j))|_{i,j}$$

Clearly for orders O, since they are integral, every $\alpha_i \in O$ has reduced trace in $\mathbb{Z}$, so the discriminant of O is in $\mathbb{Z}$.

LEMMA 2.50. *Let $I, J \in B$ be lattices. Then*

$$\operatorname{disc}(I) = [J : I]^2 \operatorname{disc}(J)$$

Proof: See [39, Lemma 15.2.15]

DEFINITION 2.51. Let I be a lattice, then the **reduced discriminant**, $\operatorname{discrd}(I)$ is defined to be the square root of $\operatorname{disc}(I)$, that is

$$\operatorname{disc}(I) = \operatorname{discrd}(I)^2$$

THEOREM 2.52. *An order $O \subseteq B$ is maximal if and only if discrd$(O) = \operatorname{disc}(B)$*

Proof: See [39, Theorem 15.5.5]

Thus, thinking of the reduced discriminant as a measure of volume we see that a smaller order has smaller discriminant, and the maximal orders have the largest possible discriminant, equal to the entire quaternion algebra.

3.6. Ideal Classes. Unlike number fields where there simply was one ideal class group, when dealing with non-commutative rings one can deal with either left ideals, right ideals or two sided ideals. We make a choice of dealing with right ideal classes and follow the book by John Voight [39].

DEFINITION 2.53. Let I, J be lattices. We say that I, J are **in the same right class** and write $I \sim_R J$ if there exists some $\alpha \in B^\times$ such that $\alpha I = J$.

PROPOSITION 2.54. *The relation $\sim_R$ is an equivalence relation*

PROOF. Reflexivity: $I \sim_R I$ since $1 \in B^\times$ and $1I = I$. Symmetry: if $I \sim_R J$ with $\alpha I = J$ for $\alpha \in B^\times$ we have $I = \alpha^{-1}J$ so $J \sim_R I$. Transitivity: if $\alpha I = J$ and $\beta J = K$ we have $\alpha\beta I = K$ and $\alpha\beta \in B^\times$. □

We denote the right equivalence class of the lattice I as $[I]_R$. If I is invertible then αI is still invertible, so every element of the class is invertible. We then say that $[I]_R$ is an invertible class.

DEFINITION 2.55. The **right class set** of the order O is

$$\mathrm{CL}_{\mathrm{S}\,R}(O) := \{[I]_R \mid I \text{ is an invertible right fractional } O\text{-ideal}\}$$

For simpler notation we write just $\mathrm{CL}_{\mathrm{S}}(O)$ instead of $\mathrm{CL}_{\mathrm{S}\,R}(O)$. The element $[O]_R$ is always in $\mathrm{CL}_{\mathrm{S}}(O)$ and

We say that two orders, O and O', are **of the same type** if there exists some $\alpha \in B^\times$ such that $O' = \alpha^{-1}O\alpha$

LEMMA 2.56. *The orders O, O' are of the same type if and only if they are isomorphic as $\mathbb{Z}$-algebras.*

Proof: See [39, Lemma 17.4.2]

DEFINITION 2.57. Two orders O, O' are **connected** if there exists a locally principal fractional O, O'-ideal $J \subset B$ called a **connecting ideal**

LEMMA 2.58. *If O, O' are maximal orders, then OO' is a O, O'-connecting ideal.*

Proof: See [39, Lemma 17.4.7]

PROPOSITION 2.59. *Let O, O' be maximal orders. Then there is a unique integral O, O'-connecting ideal I of minimal reduced norm. Furthermore we have*

$$\mathrm{nrd}(I) = [O : O \cap O']$$

Proof: See [39, Exercise 17.4]

We say that B is definite if it is ramified at ∞. The notation $\mathrm{discrd}(O)$ is the reduced discriminant of O, which equals the discriminant of B for maximal orders. We use the notation $N(I) := \mathrm{nrd}(I)^2$ for the algebra norm instead of the reduced norm we introduced.

PROPOSITION 2.60. *Let B be a definite quaternion algebra over $\mathbb{Q}$ and let $O \subseteq B$ be an order. Then $O^\times = O^1$ is a finite group, and every right ideal class in $\mathrm{CL}_{\mathrm{S}}(O)$ is represented by an integral right O-ideal with*

$$N(I) \leq \frac{8}{\pi^2}\mathit{discrd}(O)$$

And the right class set $\mathrm{CL}_{\mathrm{S}}(O)$ is finite.

Proof: See [39, Proposition 17.5.6]

3.7. Bilinear forms and orthogonality. In this section we will introduce bilinear forms and define orthogonal complements in the quaternion algebra B.

Let $B = (a, b \mid \mathbb{Q})$ be a quaternion algebra, then the reduced norm map

$$\mathrm{nrd} : B \to \mathbb{Q} \qquad \mathrm{nrd}(t + xi + yj + zij) = t^2 - ax^2 - by^2 + abz^2)$$

can be thought of as a quadratic form - that is a homogenous polynomial of degree 2 in $\mathbb{Q}[t, x, y, z]$, where the accompanying bilinear map is

$$T : B \times B \to \mathbb{Q} \qquad (\alpha, \beta) \mapsto \mathrm{nrd}(\alpha + \beta) - \mathrm{nrd}(\alpha) - \mathrm{nrd}(\beta)$$

We say that two elements $\alpha, \beta \in B$ are **orthogonal** if $T(\alpha, \beta) = 0$, where T is the map above. A subspace V of B is the **orthogonal complement** of another subspace W of B, denoted $V = W^\perp$ if

$$V = \{\alpha \in B \mid T(\alpha, \beta) = 0 \text{ for every } \beta \in W\}$$

PROPOSITION 2.61. *Let $B = (a, b \mid \mathbb{Q})$ and $1, i, j, ij$ be the standard basis ($i^2 = a$ and $j^2 = b$), then $\mathbb{Q}(i)^\perp = \mathbb{Q}(i)j$ in B.*

PROOF. We show this in the usual way by showing containment. First let $\alpha \in \mathbb{Q}(i)^\perp$, then we have $T(\alpha, t' + x'i) = 0$ for any $t', x' \in \mathbb{Q}$. We want to show that $\alpha \in \mathbb{Q}(i)j$, so we write it as $\alpha = t + xi + yj + zij$ and simply compute the bilinear map

$$T(\alpha, t' + x'i) = 2tt' + 2xx'a = 0 \iff tt' = -xx'a$$

Since that equation should hold for every x', t' it must in particular hold for $t' = 1$ and when $x' = 1$ and $x' = 2$. That is it must satisfy $t = -ax$ and $t = -2x$. Thus the only possibility is $t = x = 0$, so $\alpha = yj + zij \in \mathbb{Q}(i)j$.

Second, let $\alpha \in \mathbb{Q}(i)j$, then $\alpha = (y + zi)j = yj + zij$. If $\beta \in \mathbb{Q}(i)$ we have $\beta = t + xi$, computing

$$T(\alpha, \beta) = (t^2 - ax^2 - by^2 - abz^2) - (-by^2 - abz^2) - (t^2 - ax^2) = 0$$

shows us that $\alpha \in \mathbb{Q}(i)^\perp$ □

3.8. Creating prime ideals. We will end the section on quaternions with an algorithm that allows us to turn an ideal into a prime ideal. This is a result of [22], expanded by other authors. When we deal with supersingular curves in Chapter 4 we will make use of this algorithm.

Whenever we have a maximal order O and some left O-ideal I, we want to find another left O-ideal J of prime norm. To accomplish this, we use Lemma 4.11, and simply find an element of I of reduced norm p/N for some prime p.

PROPOSITION 2.62. *Algorithm 3 returns a prime ideal and runs in expected time bounded by* $\log^4(p)$

PROOF IDEA. If we assume that $\mathrm{nrd}(\alpha)/N$ behave like a random number when α is sampled as in the algorithm, what we need to ensure is that there are sufficiently many numbers in $[-m, m]^4$ such that a prime can be found in reasonable time. First we notice that since $\mathrm{nrd}(I) = N$ and $\alpha \in I$, then we must have $N \mid \mathrm{nrd}(\alpha)$, so $\mathrm{nrd}(\alpha)/N \in \mathbb{Z}$.

Algorithm 3: MakeIdealPrime(I)

Input: A left O-ideal I represented by Minkowski reduced basis $(\alpha_1, \alpha_2, \alpha_3, \alpha_4)$
Output: An isomorphic left O-ideal J of prime norm, isomorphic to I

```
N ← nrd(I) ;
repeat
    α ← Σ_i x_i α_i for x_i ∈ [−m, m]^4
until nrd(α)/N is prime;
return I(ᾱ/N)
```

A rather tight bound for the generators α_i can be shown to be

$$p^2 \leq 16 \operatorname{nrd}(\alpha_1) \operatorname{nrd}(\alpha_2) \operatorname{nrd}(\alpha_3) \operatorname{nrd}(\alpha_4)/N^4 < 4p^2$$

This is accomplished by viewing the ideal as a lattice and mapping it to the Hamiltonian quaternion algebra $(-1, -1 \mid \mathbb{R})$ which works nicely since B is ramified at ∞. Next using results from [6], [11] and [17, Section 16] we can get the above bounds. It will however be a massive detour to only provide a slightly more detailed description on the bounds.

Using the above bounds, the values of $\operatorname{nrd}(\alpha_i)/N$ are roughly $\sqrt{p}$. We can think of the possible values of $\operatorname{nrd}(\alpha)/N$ as an arithmetic progression with difference $\operatorname{nrd}(\alpha_i)/N$. In [22, Section 3.1] it is claimed that it is expected to be enough prime candidates when m is roughly $\log(p)$. Thus giving a running time bounded by $\log^4(p)$ and output ideal of prime norm roughly $\log^2(p)\sqrt{p}$.

The exception is when I is not an arbitrary ideal, for then the values $\operatorname{nrd}(\alpha_i)/N$ might not be distributed evenly. In particular $\operatorname{nrd}(\alpha_1)/N$ might be a lot smaller than $\sqrt{p}$, while $\operatorname{nrd}(\alpha_4)/N$ might be a lot bigger than $\sqrt{p}$. Galbraith, Petit and Silva solves this in [18] by then simply using linear combinations of α_1 and α_2, leaving out the large α_3 and α_4 parts. □

4. Elliptic Curves

In this section we will recall some well known facts about elliptic curves. We will define them, show how the group operation works, define K-rationality, and look a bit further into isogenies. This section is considered well-known and included only to provide clear notation and state results to be used for further references later.

In general an **elliptic curve** is a smooth projective curve of genus 1 with a distinguished point $\mathcal{O}$ (usually the point at infinity) that give it a group variety structure. This allows for points to be added together, having the identity element $\mathcal{O}$. In general such curves can be described with a homogeneous polynomial in three variables with coefficients in $\overline{K}$, the closure of whatever field we are working over (In our case K will be $\mathbb{C}$ or $\mathbb{F}_{p^n}$ for some n). It can be shown that if $\operatorname{char} K \notin \{2, 3\}$, then the defining polynomial can be dehomogenized and written as

$$y^2 = x^3 + ax + b$$

for $a, b \in \overline{K}$. This is called the Weierstrass equation of an elliptic curve. The smoothness requirement is satisfied whenever the **discriminant**

$$\Delta := -16(4a^3 + 27b^2)$$

is nonzero. We will denote this $E : y^2 = x^3 + ax + b$ to say that E is the elliptic curve defined by the given polynomial. The point at infinity is obtained from the embedding in the projective plane $\mathbb{P}^2(\overline{K})$, where our polynomial is homogenized to $ZY^2 - X^3 - aZ^2X - bZ^3$, giving us the point $\mathcal{O} := [0 : 1 : 0]$. We will in general use the notation $[X : Y : Z]$ for projective coordinates and (x, y) for affine coordinates. If this is unfamiliar concepts to the reader, feel free to always think of the elliptic curve points as (x, y), solutions to $y^2 = x^3 + ax + b$, with the additional point $(0, 1)$. See Figure 4 for an illustration

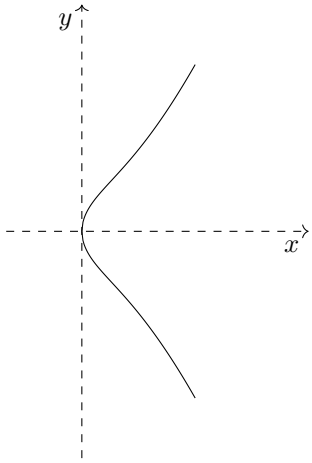

(A) Over $\mathbb{R}$

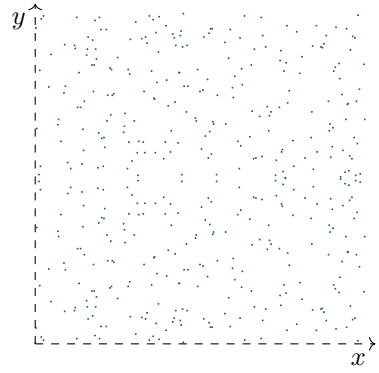

(B) Over $\mathbb{F}_{439}$

FIGURE 4. The elliptic curve $E : y^2 = x^3 + x$ over two fields

The group structure is what makes these objects interesting in cryptography. With the designated identity element $\mathcal{O}$ we can obtain an abelian group where we write $P + Q$ for the operation on two points. Geometrically we can define this operation by looking at intersection points. Letting P, Q be two arbitrary points on E, then we can construct a line through P and Q (it will be the tangent line at P if $Q = P$), which by Bezout's theorem must intersect E at a third (not necessarily distinct) point $R' = [X : Y : Z]$. If $R' \neq \mathcal{O}$ we flip the y coordinate to obtain $R = [X : -Y : Z]$, otherwise $R = R' = \mathcal{O}$ (technically we take the third intersection point of the line through R' and $\mathcal{O}$). Then we simply define $P + Q := R$, see Figure 5. To simplify notation we will write $[2]P$ for $P + P$ and extend this to $[m]$ for multiplication by m (brackets are there to make this operation distinct from scalar multiplication of the coordinates).

(A) P and Q (B) Intersection point R' and $R := P + Q$

FIGURE 5. Illustration of group operation on $y^2 = x^3 + 1$

Sometimes we are interested in curves defined over some non-closed field, this is especially true in cryptography. We say that E **is defined over** K, and write E/K, if the coefficients a, b are in K (not the algebraic closure of K). Sometimes it might be necessary to take a small extension field of K to obtain the coefficients. Similarly we are sometimes only interested in the K-**rational points** on E, namely the points $(x, y) \in E$ with $x, y \in K$ and the point $\mathcal{O}$, which we will denote $E(K)$.

Another important subset of points is the **l-torsion** group, $E[l]$, which is simply defined as the points of order dividing l

$$E[l] := \{P \in E \mid [l]P = \mathcal{O}\}$$

Sometimes we are even just interested in the K-rational l-torsion points, namely $E(K)[l]$, that is the subset of $E[l]$ whose points are defined over K.

As with any other useful algebraic object we can define maps between elliptic curves, this is what will be our primary interest in this thesis. We say that an **isogeny** $\phi : E \to E'$ is a map of elliptic curves if ϕ is a morphism of varieties that preserves $\mathcal{O}$, that is $\phi(\mathcal{O}_E) = \phi(\mathcal{O}_{E'})$. The second condition turns ϕ into a group homomorphism of the underlying group structure, while the first condition preserves the variety structure. It is common to exclude the constant isogeny $\phi : E \to \mathcal{O}$, and in this case all isogenies are surjective - a result that follows from being a morphism of projective curves. We denote the set of isogenies between E and E' as $\mathrm{Hom}(E, E')$.

We say that two elliptic curves E, E' are **isogenous** if there exists an isogeny $\phi : E \to E'$. Similarly we say that E, E' are **isomorphic**, and write $E \cong E'$, if there exists two isogenies $\phi : E \to E'$ and $\psi : E' \to E$ such that $\phi \circ \psi = \mathrm{id}_{E'}$ and $\psi \circ \phi = \mathrm{id}_E$. Whenever we have an isogeny $\phi : E \to E$ we call it an **endomorphism** and under isogeny composition

the set of all endomorphisms, $\mathrm{End}(E)$, form a ring. That is $(\phi+\psi)(P) = \phi(P)+\psi(P)$ and $(\phi\circ\psi)(P) = \phi(\psi(P))$.

A useful invariant for elliptic curves is the **j-invariant** which is defined in terms of the coefficients a, b as

$$j(E) = 1728\frac{4a^3}{16(4a^3+27b^2)}$$

In fact we have that $j(E) = j(E')$ if and only if $E \cong E'$.

4.1. Isogenies. To be more explicit about the isogenies we need the notion of function fields. In general this can be done for the projective case, but it is more enlightening to think of it in the affine case where we view E as the zeros of $f(x,y) = y^2 - x^3 - ax - b$ (and the point $\mathcal{O}$). In this case the **coordinate ring** of E, $\overline{K}[E]$ is defined as the integral domain $\overline{K}[x,y]/(f)$. If we take a nonzero polynomial in the coordinate ring we know that it will not vanish on every point of E. This is then extended to the **function field** of E, $\overline{K}(E)$, by taking the field of fractions of $\overline{K}[E]$. In this case every isogeny $\phi : E \to E'$ (as long as we exclude the constant isogeny) can be extended to provide an injection of function fields

$$\phi^* : \overline{K}(E') \to \overline{K}(E)$$

where we simply take $\phi^*(f) = f\circ\phi$. This allows us to create the most abstract definition for the degree of an isogeny. We say that the **degree** of ϕ is the degree of the field extension $\overline{K}(E)/\phi^*\overline{K}(E')$.

This extension can be split into an purely inseparable extension $L \supseteq \phi^*\overline{K}(E')$ and a separable extension $\overline{K}(E) \supseteq L$, where we denote the respective degrees as $\deg_i(\phi)$ and $\deg_s(\phi)$. Note however that we are usually interested in separable isogenies. In those cases the degree corresponds to the number of points in the kernel which can be thought of as the number of points that gets mapped to the same point under the isogeny.

Example 2.63. Let $E_1 : y^2 = x^3 + x$ and $E_2 : y^2 = x^3 + 13$ be elliptic curves over $\mathbb{F}_{23}$. Then the isogeny $\phi : E_1 \to E_2$, given by rational maps (of $\overline{FF}_{23}(E_1)$)

$$\phi(x,y) = \left(\frac{x^3+10x^2+7x+10}{x^2+10x+2}, y\frac{x^3-8x^2-8x+9}{x^3-8x^2+6x+10}\right)$$

is an separable isogeny of degree 3. This can be illustrated by looking at how many points of E_1 are mapped to the same point of E_2. What we can see is that the points $(1,-5)$, $(13,-5)$ and $(16,8)$ of E_1 are all mapped to the point $(17,21) = (17,-2)$ on E_2 through this isogeny, thus justifying the degree. This is illustrated in Figure 6.

Suppose K is the field $\mathbb{F}_q := \mathbb{F}_{p^n}$ for some n, and let $E/\mathbb{F}_q$ be an elliptic curve. Then the **q-th power Frobenius map**

$$\pi : (x,y) \mapsto (x^q, y^q)$$

is a purely inseparable endomorphism of degree q. This will be an endomorphism of E as it acts like the identity on $E(\mathbb{F}_q)$. Note however that although it has a positive degree it has trivial kernel as it acts like the identity.

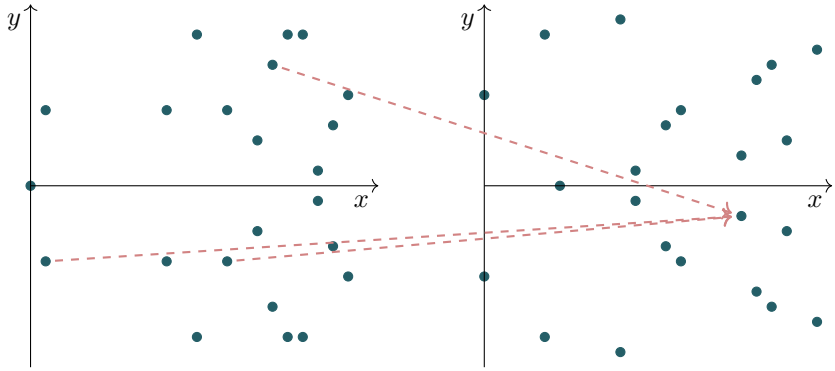

FIGURE 6. Illustration of the isogeny $\phi : E_1 \to E_2$ with $E_1 : y^2 = x^3 + x$ and $E_2 : y^2 = x^3 + 13$, both defined over $\mathbb{F}_{23}$

The Frobenius map gives us a nice way to compute the number of points on an elliptic curve.

THEOREM 2.64. *Let E be an elliptic curve defined over $\mathbb{F}_q$. Let $a = q + 1 - \#E(\mathbb{F}_q)$, then*

$$\pi^2 - [a] \circ \pi + [1] = [0] \equiv \mathcal{O}$$

where π is the qth power Frobenius, furthermore a is the unique integer that satisfies this equation.

Proof: See Washington [40, Theorem 4.10 Section 4.2]

REMARK. The integer a in the above equation is in fact the trace of π. This is more easily seen when viewing π in the abstract imaginary quadratic field or quaternion algebra. To see this in the case of elliptic curve endomorphisms we need to view it as an action on the m-torsion subgroups for various integers m. Then we can represent it as a 2 by 2 matrix which we can take the trace of. Therefore we often see it denoted by t and called the Frobenius trace

THEOREM 2.65 (Hasse). *Let E be an elliptic curve over $\mathbb{F}_q$. Then the order of $E(\mathbb{F}_q)$ satisfies*

$$|q + 1 - \#E(\mathbb{F}_q)| \leq 2\sqrt{q}$$

Proof: See Washington [40, Theorem 4.2 Section 4.2]

Which again allows us to verify whether two elliptic curves are isogenous or not.

THEOREM 2.66 (Tate). *Let $E/\mathbb{F}_q$ and $E'/\mathbb{F}_q$ be two elliptic curves. Then E and E' are $\mathbb{F}_q$-isogenous if and only if $\#E(\mathbb{F}_q) = \#E'(\mathbb{F}_q)$.*

Proof: See Tate [38, Theorem C4.1].

Separable isogenies on the other hand are much easier to deal with and they are the main interest for us. They are uniquely defined (up to isomorphism) by their kernel, $\ker(\phi)$, and

similarly the number of points in the kernel will equal the degree of the isogeny. Furthermore we have a useful result on kernel containment as depicted in Figure 7.

COROLLARY 2.67. *Let $\phi : E_1 \to E_2$ and $\psi : E_1 \to E_3$ be non-constant isogenies and assume that ϕ is separable. If $\ker \phi \subset \ker \psi$ then there is a unique isogeny $\lambda : E_2 \to E_3$ satisfying $\psi = \lambda \circ \phi$*

Proof: See Silverman [36, Corollary III 4.11]

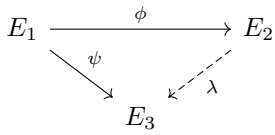

FIGURE 7. Kernel containment as in Corollary 2.67

With every isogeny $\phi : E \to E'$, there exists a unique **dual isogeny** $\widehat{\phi} : E' \to$ that satisfy $\widehat{\phi} \circ \phi = [\deg(m)]$. See Silverman [36, Theorem III 6.1(a)]. This dual isogeny satisfy many useful properties.

THEOREM 2.68. *Let $\phi : E \to E'$ be an isogeny, then the following holds true*

(1) Let $m = \deg(\phi)$, then $\widehat{\phi} \circ \phi = [m]$ on E and $\phi \circ \widehat{\phi} = [m]$ on E'

(2) Let $\psi : E' \to E''$ be another isogeny, then

$$\widehat{\psi \circ \phi} = \widehat{\phi} \circ \widehat{\psi}$$

(3) Let $\psi : E \to E'$ be another isogeny, then

$$\widehat{\psi + \phi} = \widehat{\psi} + \widehat{\phi}$$

(4) For all $m \in \mathbb{Z}$ we have $\widehat{[m]} = [m]$ and $\deg([m]) = m^2$

(5) $\deg(\widehat{\phi}) = \deg(\phi)$

(6) $\widehat{\widehat{\phi}} = \phi$

Proof: See Silverman [36, Theorem III 6.2] Notice how the dual corresponds nicely with the conjugate/standard involution defined in the section on quaternions. This is a property we will explore more in Chapter 4.

4.2. Endomorphism ring and classifictaion of elliptic curves. We briefly discuss the endomorphism ring before ending this section on elliptic curves.

COROLLARY 2.69. *Let E be an elliptic curve. Then the endomorphism ring is either $\mathbb{Z}$, an order in an imaginary quadratic field, or an order in a quaternion algebra.*

Proof: see Silverman [36, Corollary III.9.4].

Notice that since multiplication by m is always an endomorphism, so the identification $\mathbb{Z} \subseteq \mathrm{End}(E)$ is clear. Furthermore, if we are working over a finite field we always have the

Frobenius endomorphism which is not a multiplication, so $\mathbb{Z} \subsetneq \mathrm{End}(E)$ and we have that the endomorphism ring is either an order in an imaginary quadratic field or a quaternion algebra. Similarly if K has characteristics 0, then $\mathrm{End}(E)$ cannot be an order in a quaternion algebra.

We say that E is **supersingular** if $\mathrm{End}(E)$ is an order in a quaternion algebra and **ordinary** otherwise.

THEOREM 2.70. *Let $E/\mathbb{F}_q$ be a supersingular elliptic curve, then $B = \mathrm{End}(E) \otimes \mathbb{Q}$ is a quaternion algebra over $\mathbb{Q}$ ramified at $p = \mathrm{char}\, \mathbb{F}_q$ and ∞. Furthermore, $\mathrm{End}(E)$ is a maximal order in B.*

Proof: See Voight [39, Theorem 42.1.9]

An interesting fact about supersingular elliptic curves is that the j invariant is defined over $\mathbb{F}_{p^2}$ implying that every supersingular elliptic curve defined over $E/\overline{\mathbb{F}_p}$ has an isomorphic elliptic curve defined over $\mathbb{F}_{p^2}$ [36, Theorem V 3.1]. This tells us that there are only finitely many supersingular elliptic curves up to isomorphism, and in fact this number only depends on the characteristics of K.

THEOREM 2.71. *Let K be a field of characteristics $p > 3$, then the number of supersingular elliptic curves defined over $\overline{K}$ up to isomorphism is*

$$\left[\frac{p}{12}\right] + \begin{cases} 0 & \text{if } p \equiv 1 \pmod{12} \\ 1 & \text{if } p \equiv 5 \pmod{12} \\ 1 & \text{if } p \equiv 7 \pmod{12} \\ 2 & \text{if } p \equiv 11 \pmod{12} \end{cases}$$

Proof: See Silverman [36, Theorem V 4.1(c)]

Finally we mention the automorphism group of an elliptic curve. By **automorphism** we mean an endomorphism which is also an isomorphism. Suppose that the characteristics of K is neither 2 nor 3. Then the automorphism group of E has order 4 if $j(E) = 1728$, order 6 if $j(E) = 0$, and order 2 if $j(E) \notin \{0, 1728\}$. See silverman [36, Theorem III 10.1]. In particular for most j-invariants, there can be no automorphism of order greater than 2. That is, if ϕ is an automorphism then $\phi = \mathrm{id}$ or $\phi^2 = \mathrm{id}$.

CHAPTER 3

Ordinary Elliptic Curves and Isogenies

This chapter is about constructing isogenies between ordinary elliptic curves. We will begin by building up the motivation for an algorithm by Galbraith. Then, in the first section, we will discuss the endomorphism ring further by looking at what the possibilities are for ordinary elliptic curves. We will describe an algorithm for determining the endomorphism ring of such curves, and provide an algorithm for going from an elliptic curve with an arbitrary endomorphism ring to one that has a more useful endomorphism ring.

In the second section we will discuss why this is more useful by first examining how ordinary elliptic curves over $\mathbb{F}_p$ can be lifted to elliptic curves over $\mathbb{C}$. This allows us to relate the notion of isogenies with holomorphic maps between lattice quotients $\mathbb{C}/\Lambda$, which again lets us relate them to ideals of orders in number fields. Since orders are well studied this gives us plenty of results which we can use to determine the size of the graph and tell us which isogenies we should look at.

In the third section we will put everything together to a complete algorithm for computing isogenies between arbitrary ordinary elliptic curves. This section will conclude this chapter with a short discussion on the running time of the algorithm. Overall, our goal is to solve the following problem.

PROBLEM. Given two ordinary elliptic curves E_1, E_2, defined over some finite field $\mathbb{F}_q$, find an isogeny $\phi : E_1 \to E_2$, also defined over $\mathbb{F}_q$.

A natural place to start is to think of this more graph-theoretically. We let elliptic curves be nodes and isogenies be edges. If we allow for edges of arbitrary degrees we would have edges between every node, therefore we instead limit our graphs to isogenies of some fixed degree l, as depicted in Figure 8.

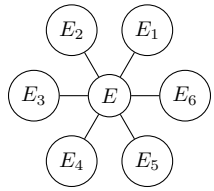

FIGURE 8. Graph of elliptic curves connected to E through 5-isogenies

This is a quite useful representation, however we are interested in computing explicit isogenies between specific nodes of the graph. Firstly the question of how to compute the isogenies and the connected elliptic curves arise. One could compute the l-torsion, find the $l+1$ subgroups, compute their corresponding isogenies and then find the elliptic curve of the codomain.

An alternative approach would be to use the **modular polynomial**, $\Phi_l(X, Y) \in \mathbb{Z}[X, Y]$. It is a polynomial of degree $l + 1$ which allows us to retrieve the j-invariants connected to some curve with a single l-isogeny rather easily.

REMARK. The actual description of the modular polynomial would takes us on a path far away from what we are really interested in. More details on the polynomial can be found in MIT Lecture Notes on Elliptic Curves [37, Lecture 20] and Silverman's Advanced Topics in the Arithmetic of Elliptic Curves [35, Section II.6] where it is denoted F_n instead of Φ_N.

Using the following theorem we see that the roots of $\Phi_l(j_1, Y)$ (with respect to Y) are exactly the j-invariants corresponding to curves that are connected to j_1 in the l-isogeny graph.

THEOREM 3.1. *Let $l > 1$ be an integer and let $\mathbb{F}_q$ be a field of characteristic p, such that $p \nmid l$. For any $j_1, j_2 \in \mathbb{F}_q$, we have $\Phi_l(j_1, j_2) = 0$ if and only if j_1 and j_2 are j-invariants of elliptic curves over $\mathbb{F}_q$ that are connected by a cyclic isogeny of degree l defined over $\mathbb{F}_q$.*

Proof: For Characteristics 0, see [37, Chapter 21], while for prime characteristic, see [20].

We can then construct a similar graph where the nodes are the different j-invariants, and edges are l-isogenies as depicted in Figure 9.

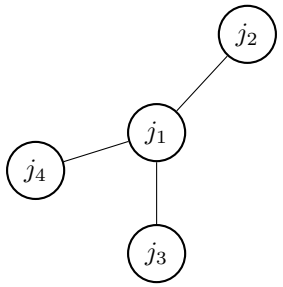

FIGURE 9. Graph of j-invariants

This helps us quite a lot along the way and we could begin a standard graph-search algorithm starting at $j(E_1)$ and search until we found $j(E_2)$ by changing the degree l every once in a while. Then, when a path was found, we could compose it with an isomorphism to get the actual isogeny $\phi : E_1 \to E_2$. Such an approach would work, however we can be much more efficient by exploiting the structure of the endomorphism rings even further.

1. Computing the endomorphism ring

In this section we will look further into the endomorphism ring. We will first describe what sorts of endomorphisms rings we can find, then an algorithm for computing the actual

endomorphism ring. Finally we will use this knowledge to build an algorithm to move our known endomorphism ring to another type which we will later show to be more useful.

Recall that for ordinary elliptic curves, by Corollary 2.69, the endomorphism ring is either $\mathbb{Z}$ or an order in an imaginary quadratic field. Furthermore, over finite fields, we have that the Frobenius map $\pi : E \to E$ is an endomorphism, so $\mathbb{Z} \subsetneq \mathbb{Z}[\pi] \subseteq \mathrm{End}(E)$. giving us that $\mathbb{Z} \neq \mathrm{End}(E)$. Let us write K for this imaginary quadratic field. In fact, this π allows us to describe K entirely by Theorem 2.65.

Furthermore by Theorem 2.64 we know that π has the characteristical polynomial $\pi^2 - t\pi + q = 0$ with $t = q + 1 - \#E(\mathbb{F}_q)$. Thus it is an element of degree 2 in $K \setminus \mathbb{Q}$. This gives us that

$$K = \mathbb{Q}(\sqrt{t^2 - 4q})$$

a standard result in algebraic number theory which can be seen in Pierre Samuel [33, Section 2.5].

Example 3.2. Let $E : y^2 = x^3 + 2x + 3$ be an elliptic curve defined over $\mathbb{F}_{101}$. The trace of the Frobenius, $t = 6$, so the endomorphism ring of E is an order in

$$K = \mathbb{Q}(\sqrt{t^2 - 4q}) = \mathbb{Q}(4\sqrt{-23}) = \mathbb{Q}(\sqrt{-23})$$

This gives us that the maximal order $\mathcal{O}_K$ is equal to $\mathbb{Z}[\frac{\sqrt{-23}+1}{2}]$. The discriminant of K is simply $\Delta_K = -23$. Since the roots of π are $3 \pm 2\sqrt{-23}$, we get that $\mathbb{Z}[\pi] = \mathbb{Z}[3 + 2\sqrt{-23}] = \mathbb{Z}[2\sqrt{-23}]$. Notice that we cannot remove 2 here since $2^{-1} \notin \mathbb{Z}$.

This would not be very useful if these number fields were different for various elliptic curves. However from Tate's theorem 2.66, we get that two isogenous curves have the same number of points, so their Frobenius trace t are the same and their endomorphism rings are orders of the same imaginary quadratic field $\mathbb{Q}(\sqrt{t^2 - 4q})$. This motivates us to dive deeper into algebraic number theory and study their endomorphism rings further.

Let $\mathcal{O}_K$ be the ring of integers of K, we then have that $\mathcal{O} := \mathrm{End}(E)$ lies somewhere between $\mathbb{Z}[\pi]$ and $\mathcal{O}_K$. That is

$$\mathbb{Z}[\pi] \subseteq \mathcal{O} \subseteq \mathcal{O}_K$$

Kohel [21] has studied these endomorphism rings quite extensively, so we have a couple of useful results which we can use to navigate the graph of isogenous elliptic curves.

Proposition 3.3. *Let E be an ordinary elliptic curve over $\mathbb{F}_q = \mathbb{F}_{p^n}$ and $\phi : E \to E'$ an isogeny of prime degree l such that $l \neq p$, then $l \mid [\mathrm{End}(E) : \mathrm{End}(E')]$ or $l \mid [\mathrm{End}(E') : \mathrm{End}(E)]$*

Proof: See Kohel [21, Proposition 21].

The above proposition gives us some information about the isogenies connecting the endomorphism rings. Let $\mathcal{O}_1$ and $\mathcal{O}_2$ be the endomorphism rings of E_1 and E_2. Then if $l \mid [\mathcal{O}_1 : \mathcal{O}_2]$, we know that the isogeny $\phi : E_1 \to E_2$ must be composed of an isogeny of degree l, so $l \mid \deg(\phi)$.

We use this to classify the different endomorphism rings which can occur. We know that whenever we have an order, $\mathcal{O}$, in $\mathcal{O}_K$ it is just a subring and it is of the form $\mathcal{O} = \mathbb{Z} + c\mathcal{O}_K$ (See Theorem 3.16). So we are left with a finite number of orders of the form $\mathcal{O} = \mathbb{Z} + c\mathcal{O}_k$ where c divides $[\mathcal{O}_K : \mathbb{Z}[\pi]]$.

EXAMPLE 3.4. Extending Example 3.2, we have

$$\mathbb{Z}\left[2\sqrt{-23}\right] \subseteq \mathcal{O} \subseteq \mathbb{Z}\left[\frac{\sqrt{-23}+1}{2}\right]$$

Notice that $\mathbb{Z}[2\sqrt{-23}] = \mathbb{Z} + 4\mathcal{O}_K$, so the number $c = 4$ is the conductor of $\mathbb{Z}[\pi]$, thus the order $\mathcal{O}$ can be any of $\mathbb{Z} + 4\mathcal{O}_K$, $\mathbb{Z} + 2\mathcal{O}_K$ or $\mathcal{O}_K$.

The way we classify the endomorphism rings is using the notion of a volcano where we say that an isogeny $\phi : E \to E'$ of degree l is

Ascending: If $\text{End}(E') : \text{End}(E)] = l$ (that is the $\text{End}(E')$ contains $\text{End}(E)$).
Descending: If $[\text{End}(E) : \text{End}(E')] = l$ ($\text{End}(E')$ is contained in $\text{End}(E)$).
Horizontal: If $\text{End}(E) \cong \text{End}(E')$

As illustrated in Figure 10.

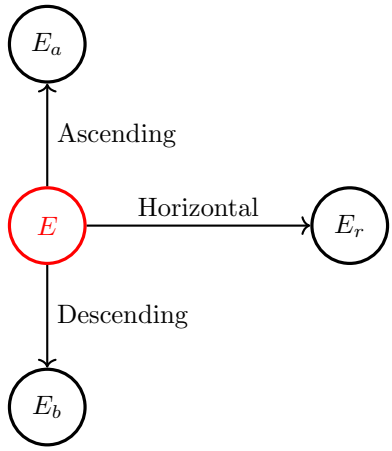

FIGURE 10. Illustration of isogenies from E where $\text{End}(E_b) \subsetneq \text{End}(E) \subsetneq \text{End}(E_a)$ and $\text{End}(E) \cong \text{End}(E_r)$

We call the very top of this volcano (the place where no more ascending isogenies exists) for the surface, while the very bottom (where no descending isogenies exists) the floor. We call the different orders for the levels of the volcano. That is, if $\phi : E \to E'$ is a horizontal isogeny, then E and E' are on the same level. Naturally we have that if $\text{End}(E) \cong \mathbb{Z}[\pi]$ then it is on the floor of the l-volcano and if $\text{End}(E) = \mathcal{O}_K$ it is on the surface. It may however newer reach those end-orders using only l-isogenies, so the surface or floor of an l-volcano may be different from $\mathcal{O}_K$ and $\mathbb{Z}[\pi]$ given some starting curve E.

We can describe the edges of this volcano even further using another one of Kohel's propositions.

PROPOSITION 3.5. *Let E be an ordinary elliptic curve over $\mathbb{F}_q$ and $\mathcal{O}$ be its endomorphism ring. Let $c = [\mathcal{O}_K : \mathcal{O}]$, $t = \mathrm{Tr}(\pi)$, and l be a prime. Then every l-isogeny $\phi : E \to E'$ arises from the following cases.*

(1) If $l \nmid c$ then there are exactly $\left(1 + \left(\frac{t^2-4q}{l}\right)\right)$[1] horizontal elliptic curves from E.
(2) If $l \mid c$ then there are no horizontal l-isogenies starting at E.
(3) If $l \mid c$ there is exactly one ascending l-isogeny starting at E
(4) If $l \mid [\mathcal{O} : \mathbb{Z}[\pi]$, then there are $l + 1$ isogenies of degree l going to different elliptic curves where the horizontal and ascending isogenies are as above and the remaining are descending.
(5) If $l \nmid [\mathcal{O} : \mathbb{Z}[\pi]]$ then there is no descending l-isogeny.

Proof: See Kohel [21, Proposition 23].

From this proposition, we see that we can actually make a volcano-looking structure from the isogeny graph we created before as depicted in Figure 11.

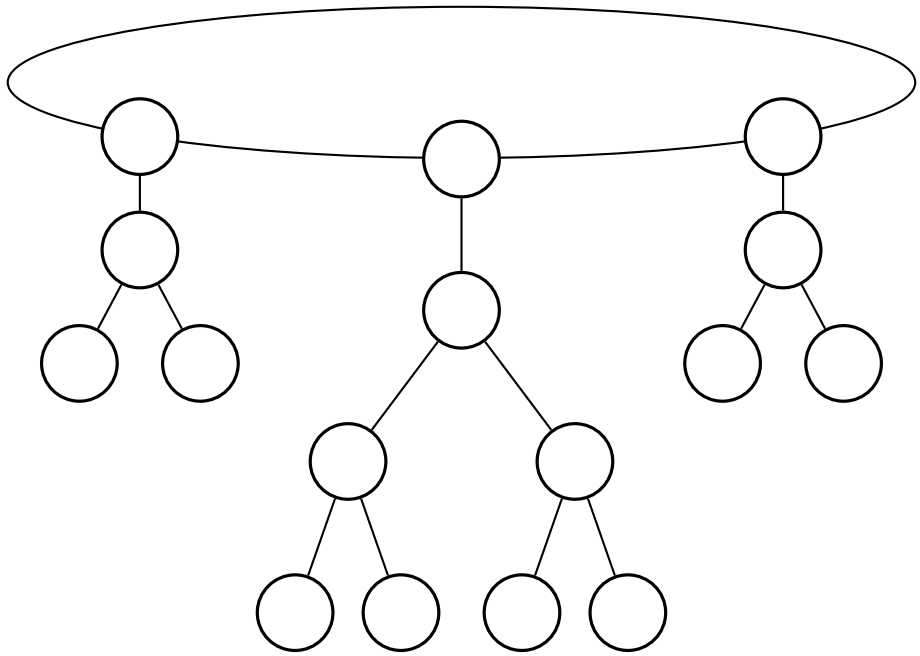

FIGURE 11. Isogeny Volcano of degree 2 isogenies with a of total 4 levels

This volcano along with Proposition 3.5 gives us a lot of information about how to reduce the search among the isogeny graph. Firstly we notice that it is only when $\mathrm{End}(E_1) \cong \mathrm{End}(E_2)$ that

[1]Where $\left(\frac{a}{l}\right)$ is the Legendre symbol

we would want to go through horizontal isogenies to find a match. Thus the first step should be to move $\mathrm{End}(E_1)$ and $\mathrm{End}(E_2)$ to the same level [2]. If we could compute the actual endomorphism ring of E_1 and E_2, this would be trivial, then we could simply compute $c_1 = [\mathcal{O}_K : \mathcal{O}_1]$ and $c_2 = [\mathcal{O}_K, : \mathcal{O}_2]$, factor them, and look at each prime to determine if one needs to go up, down or stay at the same level with respect to that prime. Kohel has constructed an algorithm for determining how far it is to the bottom of am l-isogeny graph, and thus telling us to what degree l divides c. The algorithm requires being able to compute $\mathcal{O}_K$ and $\mathbb{Z}[\pi]$ which only requires computing $\#E(\mathbb{F}_q)$ which can be done in $O(\log(q)^8)$ steps by Schoof's Algorithm [34, Section 5] explained in [36, Section XI.3].

1.1. Find the floor. What one notices is that if l divides $[\mathcal{O}_K : \mathbb{Z}[\pi]]$, then $l \nmid [\mathcal{O} : \mathbb{Z}[\pi]]$ if and only if there is a single l-isogeny over $\mathbb{F}_q$ starting from E. That is the modular polynomial $\Phi_l(j(E), Y)$ has exactly one root. Similarly we have a nice result related to the dual isogeny.

PROPOSITION 3.6. *Let ϕ be a horizontal (resp ascending, descending), then $\hat{\phi}$ is horizontal (resp descending, ascending)*

PROOF. This proposition is immediate once we recall that $\phi : E \to E'$ has dual $\widehat{\phi} : E' \to E$. □

Which implies that if we start with a descending isogeny ϕ and select the next isogeny as something different from its dual, we must continue a descending path. To ensure that we choose a non-backtracking isogeny ψ after having chosen $\phi : E \to E'$ can simply quotient out the j-invariant of E from the modular polynomial of E', that is

$$\Phi_l(j(E'), Y)/(Y - j(E))$$

Now, let us assume that $l^m \mid c$ but $l^{m+1} \nmid c$ for some integer m. Kohel's algorithm is then Algorithm 4 taking the value l^m and E as input and returning the level, n, at which E is.

The algorithm is depicted by Figure 12 where the orange nodes are those that are visited and the green node corresponds to the terminal node. This algorithm terminates correctly because if one starts on the floor, it returns 0 immediately. If one is in the middle of the volcano, then there are at most m steps to the floor, and taking two paths gives you either a path upward or downward the volcano which must terminate after at most m steps. If you move upward and hit the surface you simply start moving around with horizontal isogenies and possibly going downward again after some steps. If you start on the surface there are exactly m steps to the bottom, but you need not have selected any descending isogenies, so you may essentially just move around the surface, therefore the last if-statement ensures that the algorithm terminates even in this case.

Now, moving horizontally is only possible when $l \nmid c$ so it makes sense to move E_1 and E_2 to the surface of the volcano (making $[\mathcal{O}_K : \mathcal{O}] = 1$), before trying to connect them. Doing so would allow us to not only use primes not dividing $[\mathcal{O}_K : \mathbb{Z}[\pi]]$ but also those that do divide that

[2] We shall shortly see that it is beneficial to move both curves to the surface instead of the same level, but the approach is more or less the same.

Algorithm 4: Kohel's Find the floor

Input: l, m, E, such that $l^m \mid [\mathcal{O}_K : \mathrm{End}(E)]$
Output: Largest n such that $l^n \mid [\mathrm{End}(E) : \mathbb{Z}[\pi]]$

$n \leftarrow 0$;
$j_1 \leftarrow j(E)$;
if $\Phi_l(j_1, Y)$ *has exactly one root* **then**
  return 0 ; // E is already on the floor
end
$j_1, j_2 \leftarrow \Phi_l(j_1, Y)$; // Any two distinct roots of the polynomial
while $\Phi_l(j_1, Y)$ *and* $\Phi_l(j_2, Y)$ *has more than one root* **do**
  $j_1 \leftarrow \Phi_l(j_1, Y)/(Y - j_1)$; // Any root
  $j_2 \leftarrow \Phi_l(j_2, Y)/(Y - j_2)$; // Any root
  $n \leftarrow n + 1$;
  if $n = m$ **then**
    return m ; // E is on the surface
  end
end
return n

index [3]. We can go to the surface by modifying Algorithm 4 slightly. At each step, starting at E_1 one computes all the j-invariants using $\Phi_l(j(E_1), Y)$. Then one computes the level of each j-invariant, and keeps the one corresponding to the curve above E_1. This process is repeated until we have the j-invariant of an elliptic curve at the surface, it must necessarily be repeated for each $l \mid [\mathcal{O}_K : \mathbb{Z}[\pi]$.

Example 3.7. Continuing Example 3.2 we have that $c = 4$, so it is only necessary to look at the degree 2 isogenies, and the floor should be reached within 2 steps. We know that the j invariant of E is 74. Computing $\Phi_2(74, Y)$ gives us only one root, 98 (over $\mathbb{F}_{101}$), so we are already on the floor. Going one step up, solving $\Phi_2(98, Y)$ gives us $74, 22, 30$ as the three roots, satisfying Kohel's proposition. We know 74 is below, so we need to expand 22 and 30. $\Phi_2(22, Y)$ has only one root, 98, so it is below 98. Thus 30 is above 98 and it must necessarily be on the surface, since it is two levels above 74.

Now that we have found the j-invariants at the surface, we are ready to attempt to connect them. In order to do so we need more information about how the graph expands, and we use the ideal class group to explain this.

[3] just verifying that we are not choosing an elliptic curve below the surface when walking.

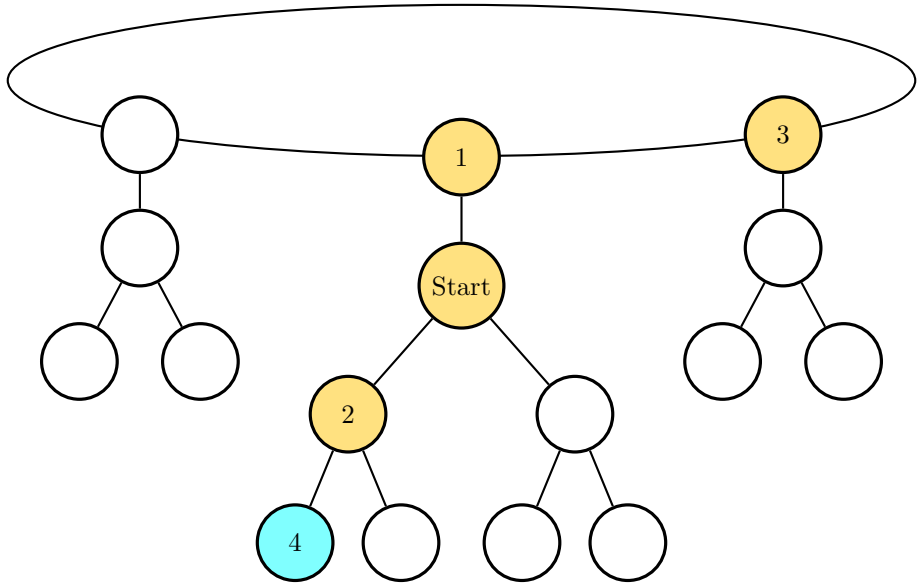

FIGURE 12. Kohel's algorithm: Find the floor, starting at "Start" and traversing nodes in order described

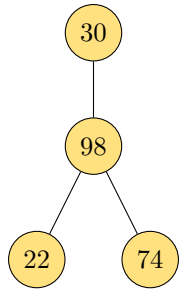

FIGURE 13. Example, going to the surface of $E : y^2 = x^3 + 2x + 3$ over $\mathbb{F}_{101}$, with $j(E) = 74$

2. Traversing theory: Ideal classes and lattices

In this section we will look further into how we should traverse the surface graph of the volcano described earlier. We begin by introducin Deuring's lifting theorem to move our elliptic curves to the complex case. Then we will show the connection between isogenies and maps

between complex lattices. Finally we will connect this information to the theory of ideals and orders which will allow us to get some restrictions as to which isogenies we should look at.

2.1. Lifting to C. Using the Deuring lifting theorem we can move our endomorphism ring of $E_1/\mathbb{F}_q$ to some $E/\mathbb{C}$ which allows us to characterize the behaviour of endomorphisms more clearly.

THEOREM 3.8 (Deuring). *Let $E_0/\mathbb{F}_q$ be an elliptic curve in characteristics p, with an endomorphism α_0 that is non trivial. Then there exists some elliptic curve E defined over some number field K, an endomorphism α of E and a non-degenerate reduction of A at place $\mathfrak{p}$ above p such that E_0 is isomorphic to $\bar{E}$ and α_0 corresponds to $\bar{\alpha}$ under the isomorphism.*

Proof: See [24, Theorem 14, Chapter 13].

Now let us explore what this theorem is really saying in regards to E/K. Firstly, we notice that the coefficients describing the elliptic curve E all lie in the ring of integers, O_K, of K. From algebraic number theory we know that $\mathcal{O}_K/\mathfrak{p} \cong \mathbb{F}_q$ where $q = \mathrm{Nr}(\mathfrak{p})$, so to get the elliptic curve over $\mathbb{F}_q$ it makes sense to look the map $x \mapsto \tilde{x} := x \pmod{\mathfrak{p}}$ for some prime ideal $\mathfrak{p}$ of norm q above p.

This is exactly what the theorem is doing. Writing $E : y^2 = x^3 + ax + b$ and then taking the reduction modulo $\mathfrak{p}$ yielding $a \mapsto \tilde{a}$, $b \mapsto \tilde{b}$ making $\tilde{E}$ defined over $\mathbb{F}_q$. The theorem then states that $\tilde{E}$ is isomorphic to E_0 in such a way that the endomorphisms α_0 of E_0 lifts to some endomorphism α of E that then reduces to another isomorphic endomorphism $\tilde{\alpha}$ of $\tilde{E}$, see Figure 14.

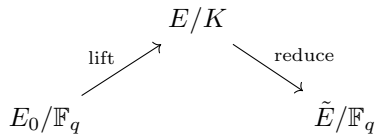

FIGURE 14. Deuring's lifting theorem

Now since the number field K has characteristics 0, we can use what is known as the *Lefschets principle*, which says that whenever we do algebraic geometry over some algebraically closed field of characteristics 0, it is the same as doing algebraic geometry over $\mathbb{C}$, see [36, Section VI.6].

2.2. Connection to lattices. Now that we know that we can lift our elliptic curves to something equivalent to an elliptic curve defined over $\mathbb{C}$, we can study the behaviour of elliptic curves there in order to figure out what happens when composing isogenies together. We will begin by describing the connection between complex elliptic curves and lattices in $\mathbb{C}$.

PROPOSITION 3.9. *Let $E/\mathbb{C}$ be the elliptic curve $E : y^2 = 4x^3 - g_2x - g_3$ where g_2 and g_3 are the quantities associated to the lattice $\Lambda \subset \mathbb{C}$. Then the map*

$$\phi : \mathbb{C}/\Lambda \to E(\mathbb{C}) \qquad z \mapsto [\wp(z), \wp'(z), 1]$$

is a complex analytic isomorphism

Proof: See Silverman [36, Proposition 3.6].

In other words, whenever we have a lattice $\Lambda \subset \mathbb{C}$, there is an elliptic curve $E/\mathbb{C}$ such that $\mathbb{C}/\Lambda \cong E(\mathbb{C})$.

COROLLARY 3.10. *Let $E/\mathbb{C}$ be an elliptic curve, then there exists a lattice $\Lambda \subset \mathbb{C}$, unique up to homothety, and a complex analytic isomorphism*

$$\phi : \mathbb{C}/\Lambda \to E(\mathbb{C}) \qquad z \mapsto [\wp(z), \wp'(z), 1]$$

Proof: See Silverman [36, Corollary 5.1.1]

That is, not only is there a complex elliptic curve for every lattice, but every complex elliptic curve is isomorphic to a unique lattice up to homothety. Furthermore, as the next theorem shows, there is a deeper connection between maps of elliptic curves and lattices. For any $\alpha \in \mathbb{C}$ satisfying $\alpha\Lambda_1 \subseteq \Lambda_2$ we can construct the holomorphic map $\phi_\alpha : \mathbb{C}/\Lambda_1 \to \mathbb{C}/\Lambda_2$ as $\phi_a(z) = \alpha z$.

THEOREM 3.11. *Using the map ϕ_a as above we get the following results*

(1) The following map is a bijection:

$$\{\alpha \in \mathbb{C} \mid \alpha\Lambda_1 \subseteq \Lambda_2\} \to \left\{ \begin{array}{c} \text{Holomorphic maps} \\ \phi : \mathbb{C}/\Lambda_1 \to \mathbb{C}/\Lambda_2 \\ \text{such that } \phi(0) = 0 \end{array} \right\}$$
$$\alpha \mapsto \phi_\alpha$$

(2) Let E_1, E_2 be elliptic curves corresponding to lattices Λ_1 and Λ_2, then the following inclusion is a bijection:

$$\{\text{Isogenies } \phi : E_1 \to E_2\} \to \left\{ \begin{array}{c} \text{Holomorphic maps} \\ \phi : \mathbb{C}/\Lambda_1 \to \mathbb{C}/\Lambda_2 \\ \text{such that } \phi(0) = 0 \end{array} \right\}$$

Proof: See Silverman [36, Theorem VI.4.1].

This allows us to think of maps between complex elliptic curves as maps between lattices. To make things even more clear we have the following corollary.

COROLLARY 3.12. *Let $E_1/\mathbb{C}$ and $E_2/\mathbb{C}$ be two elliptic curves corresponding to lattices Λ_1 and Λ_2. Then E_1 and E_2 are isomorphic if and only if Λ_1 and Λ_2 are homothetic.*

PROOF. E_1 and E_2 are isomorphic if and only if there are isogenies $\phi : E_1 \to E_2$ and $\psi : E_2 \to E_1$ such that $\phi \circ \psi : E_2 \to E_1 = id_{E_2}$ and $\psi \circ \phi = id_{E_1}$. From the theorem above ϕ and ψ corresponds to α_1, α_2 such that $\alpha_1\Lambda_1 \subseteq \Lambda_2$ and $\alpha_2\Lambda_2 \subseteq \Lambda_1$. Now let us prove the two directions.

$\boxed{\implies}$ Assuming E_1 and E_2 are isomorphic, then $\alpha_1\alpha_2 = 1$ since they correspond to the identity map. Looking at the equations we get

$$\begin{aligned} & \alpha_1\Lambda_1 \subseteq \Lambda_2 \\ \implies & \alpha_2\Lambda_1 \subseteq \alpha_2\Lambda_2 \subseteq \Lambda_1 \\ \iff & \Lambda_1 \subseteq \alpha_2\Lambda_2 \subseteq \Lambda_1 \end{aligned}$$

So $\alpha_2\Lambda_2 = \Lambda_1$, similarly $\alpha_1\Lambda_1 = \Lambda_2$, which is exactly the definition of Λ_1, Λ_2 being homothetic.

$\boxed{\impliedby :}$ Assuming Λ_1 and Λ_2 are homothetic, there exists an $\alpha \in \mathbb{C}^*$ such that $\alpha\Lambda_1 = \Lambda_2$. Setting $\alpha_1 = \alpha$ and $\alpha_2 = \alpha^{-1}$ gives us the corresponding maps $\phi_{\alpha_1} : \mathbb{C}/\Lambda_1 \to \mathbb{C}/\Lambda_2$ and $\phi_{\alpha_2} : \mathbb{C}/\Lambda_2 \to \mathbb{C}/\Lambda_1$ which corresponds to isogenies $\phi : E_1 \to E_2$ and $\psi : E_2 \to E_1$ under the bijection of the above theorem. Furthermore, since the composition $\phi \circ \psi$ corresponds to $\phi_{\alpha_1\alpha_2} = \phi_1$ it is the identity, and similarly for $\psi \circ \phi = \phi_1$, so we must have $E_1 \cong E_2$. □

Therefore we see that the endomorphism ring of E can be represented as a set of scalars in $\mathbb{C}$:

$$\mathrm{End}(E) \cong \{\alpha \in \mathbb{C} \mid \alpha\Lambda \subseteq \Lambda\} \subseteq \mathbb{C}$$

2.3. Lattices, ideals and orders. Now that we can think of elliptic curves as lattices, we shall show that there are finitely many elliptic curves with the same endomorphism ring and that they can be reached using rather few small-prime isogenies composed together. The results in this section comes from [36, Chapter VI], [35, Chapter II], and [24, Chapter 8].

For simplicity we define the set that we are looking at, namely isomorphism classes of elliptic curves. That is those elliptic curves that are isomorphic are treated as the same.

Definition 3.13. Let $\mathcal{O}$ be an order in an imaginary quadratic field K, then we say that $\mathcal{ELL}(\mathcal{O})$ is the set of **isomorphism classes of elliptic curves** over $\mathbb{C}$ with endomorphism ring isomorphic to $\mathcal{O}$. That is

$$\mathcal{ELL}(\mathcal{O}) = \frac{\{\text{elliptic curves } E/\mathbb{C} \text{ with } \mathrm{End}(E) \cong \mathcal{O}\}}{\text{isomorphism over } \mathbb{C}}$$

Which by Theorem 3.11 can be rewritten as

$$\mathcal{ELL}(\mathcal{O}) = \frac{\{\text{Lattices } \Lambda \text{ with } \mathrm{End}(E_\Lambda) \cong \mathcal{O}\}}{\text{homothety}}$$

Noting that any number field K can be embedded in $\mathbb{C}$ we shall look at lattices in K and follow Lang's Elliptic Functions [24, Chapter 8] to describe the generalized class group. To simplify the notation we assume that $K = \mathbb{Q}(\tau)$ is an imaginary quadratic number field, which is what is useful for our purposes.

Lattices in this chapter will be what we previously defined as full lattices of K. They correspond to free additive subgroups of K of dimension 2 over $\mathbb{Z}$. Similarly orders will be subrings of O_K with dimension 2 over $\mathbb{Z}$. The order of a lattice is equivalent to the left order of a lattice is just the set $\{\alpha \in K \mid \alpha\Lambda \subseteq \Lambda\}$ which corresponds to being a proper subset of $\mathcal{O}_K$

and an order, just like how the left orders of a lattice in the quaternions corresponds to being an order.

Since a lattice has degree 2 we can write it as $\Lambda = [w_0, w_1]$ for $w_0, w_1 \in K$. Furthermore as we are mainly interested in lattices up to homothety, so we can write it as $[1, w] := [1, w_1/w_0]$.

DEFINITION 3.14. Given an order $\mathcal{O} \subseteq K$, we say that a lattice Λ is a **proper** $\mathcal{O}$-lattice if

$$\mathcal{O} = \{\alpha \in K \mid \alpha\Lambda \subseteq \Lambda\}$$

From the above discussion we see that Λ is a proper $\mathcal{O}$ lattice if and only if it corresponds to an elliptic curve with endomorphism ring $\mathcal{O}$. Similarly we define the $\mathcal{O}$-ideals to be the ideals $\mathfrak{a} \subseteq \mathcal{O}$ which is also a lattice. This allows us to talk about lattices and ideals of orders interchangeably. That is if we look at an isomorphism class of elliptic curves $[E]$ over $\mathbb{C}$, they correspond to a class of homothetic lattices $[\Lambda] \subseteq \mathbb{C}$. Taking the left order of Λ necessarily gives us the order O isomorphic to the endomorphism ring of E for any representative $\Lambda \in [\Lambda]$. Thus Λ is a proper O-lattice. Since any $\alpha \in O$ satisfy $\alpha\Lambda \subseteq \Lambda$ (as its left order is O), and α is already an additive group it is a left O-ideal. Since O_K is commutative it is necessarily a right O ideal as well, so we Λ is an O-ideal.

COROLLARY 3.15. *Let $\mathcal{O}$ be an order in K. Every proper $\mathcal{O}$-lattice is $\mathcal{O}$-invertible and conversely any lattice which is $\mathcal{O}$-invertible is a proper $\mathcal{O}$-lattice. Furthermore the set of proper $\mathcal{O}$-lattices is a multiplicative group.*

Proof: Lang [24, Corollary, Chapter 8, p 91]

The corollary above tells us that whenever we have a lattice Λ that has an inverse Λ^{-1} in $\mathcal{O}$, ie $\Lambda^{-1}\Lambda = \mathcal{O}$, then Λ corresponds to an elliptic curve with endomorphism ring $\mathcal{O}$.

THEOREM 3.16. *Let $\mathcal{O}$ be an order in K, and $\mathcal{O}_K := [1, z]$. Then there exists a unique positive integer c such that*

$$\mathcal{O} = [1, cz] = \mathbb{Z} + c\mathcal{O}_K$$

Proof: See Lang [24, Theorem 3, Chapter 8].

The integer c in the above theorem is called the **conductor** of $\mathcal{O}$. If we have an $\mathcal{O}$-ideal $\mathfrak{a}$, we say that $\mathfrak{a}$ is prime to c if $\mathfrak{a} + c\mathcal{O} = \mathcal{O}$ or equivalently $\mathfrak{a} + c\mathcal{O}_K = \mathcal{O}$. Let $I_K(c)$ be the set of all $\mathcal{O}_K$-ideals that are prime to c and similarly $I_{\mathcal{O}}(c)$ be the set of all $\mathcal{O}$-ideals prime to c. These sets coincide by the following theorem.

THEOREM 3.17. *There is a multiplicative bijection between $I_K(c)$ and $I_{\mathcal{O}}(C)$ given by*

$$\mathfrak{a} \mapsto \mathfrak{a} \cap \mathcal{O}$$

and

$$\mathfrak{a} \mapsto \mathfrak{a}\mathcal{O}_K$$

Furthermore, any ideal of $\mathcal{O}$ prime to c is a proper $\mathcal{O}$-ideal

Proof: See Lang [24, Theorem 4, Chapter 8].

Meaning that we can just as easily look at I_K as $I_{\mathcal{O}}$.

DEFINITION 3.18. Let $\mathcal{O}$ be an order in K, the group of **proper $\mathcal{O}$-ideal classes**, $\mathcal{CL}(\mathcal{O})$ is the quotient

$$\mathcal{CL}(\mathcal{O}) := I_{\mathcal{O}}/P_{\mathcal{O}}$$

where $I_{\mathcal{O}}$ are the proper $\mathcal{O}$ ideals and $P_{\mathcal{O}}$ are the principal $\mathcal{O}$ ideals (those on the form $a\mathcal{O}$).

Notice that for $\mathcal{O} = \mathcal{O}_K$ this is the ideal class group, justifying the notation.

Since $I_{\mathcal{O}}$ are the proper $\mathcal{O}$-lattices, they correspond to elliptic curves with endomorphism rings equal isomorphic to $\mathcal{O}$. Two lattices Λ_1, Λ_2 are homothetic if $\alpha\Lambda_1 = \Lambda_2$ implying that they differ by exactly a principal ideal $\alpha\mathcal{O}$. This gives us that

$$\mathcal{ELL}(\mathcal{O}) \cong \mathcal{CL}(\mathcal{O})$$

When $O = \mathcal{O}_K$, this class group is well understood, and we see immediately that the class number h tells us how many isogenous elliptic curves there are. To understand what happens when $\mathcal{O} \subsetneq \mathcal{O}_K$ we will therefore look a bit further into the generalized class group for a (potentially) smaller $\mathcal{O}$.

THEOREM 3.19. *Let Λ be a proper $\mathcal{O}$-lattice and m a positive integer. Then there exists an element $\alpha \in K$ such that $\alpha\Lambda \subseteq \mathcal{O}$ and*

$$\alpha\Lambda + m\mathcal{O} = \mathcal{O}$$

In other words, in the equivalence class of Λ, there is a lattice prime to m.

See Lang [24, Theorem 5, Chapter 8].

We define the sets $I_{\mathcal{O}}(c)$ and $P_{\mathcal{O}}(c)$ as for $I_{\mathcal{O}}$ and $P_{\mathcal{O}}$, except that we also require that the proper ideals are prime to c. Thus from the theorem above, we have that

$$\mathcal{CL}(\mathcal{O}) \cong I_{\mathcal{O}}(c)/P_{\mathcal{O}}(c)$$

Next we want to relate this to $\mathcal{O}_K$. As above, we let $\mathcal{O}_K = [1, z]$ and define $P_z(c)$ as the set of $\mathcal{O}_K$-ideals that are principal and of the form $\mathfrak{a} = \mathcal{O}_K\alpha$ where $\alpha \equiv a \pmod{c\mathcal{O}_K}$, where c is the conductor of $\mathcal{O}$ and a is just some integer relatively prime to c.

LEMMA 3.20. *Let $\mathfrak{a} \in P_z(c)$ as above, then*

$$\mathfrak{a} \cap \mathcal{O} = \mathcal{O}\alpha$$

Proof: See Lang [24, Lemma 1, Chapter 8].

THEOREM 3.21. *Consider the homomorphism $I_K(c) \to I_{\mathcal{O}}(c)$ defined as $\mathfrak{a} \mapsto \mathfrak{a} \cap \mathcal{O}$, then the inerse image of $P_{\mathcal{O}}(c)$ is $P_z(\mathcal{O})$*

Proof: See Lang [24, Theorem 6, Chapter 8].

From the above lemma and theorem we get that we can rewrite the class group as

$$\mathcal{CL}(\mathcal{O}) \cong I_k(c)/P_z(c)$$

Which allows us to compute the order of $\mathcal{CL}(\mathcal{O})$ based on $h(O_K)$, the class number of K.

THEOREM 3.22. *Let $\mathcal{O}$ be the order of conductor c in an imaginary quadratic field K of discriminant Δ_K, then we have*

$$h_{\mathcal{O}} = \frac{h(\mathcal{O}_K)c}{[\mathcal{O}_K^* : \mathcal{O}^*]} \prod_{p|c} \left(1 - \left(\frac{\Delta_K}{p}\right)\frac{1}{p}\right)$$

Where $\mathcal{O}_K^*$ and $\mathcal{O}^*$ are the groups of units in $\mathcal{O}_K$ and $\mathcal{O}$ respectively, and $\left(\frac{\Delta_K}{p}\right)$ is the Legendre symbol for p odd and whenever $p = 2$ it equals 0 if $2 \mid \Delta_K$, 1 if $\Delta_K \mod 8 = 1$ or 2 if $\Delta_K \mod 8 = 5$.

Furthermore, $h(\mathcal{O})$ is always an integer multiple of $h(\mathcal{O}_K)$

See Lang [9, Theorem 7.24].

This theorem is interesting because it tells us that we have nothing to gain from searching among elliptic curves at a level below $\mathcal{O}_K$ as opposed to looking at $\mathcal{O}_K$. This comes from the fact that $c \geq [\mathcal{O}_K^* : \mathcal{O}^*]$ by definition, and that $\prod_{p|c}\left(1 - \left(\frac{\Delta_K}{p}\right)\frac{1}{p}\right) \geq 0$ (0 if and only if every $p \mid c$ is prime in K).

Furthermore, it is well known that the class number $h(\mathcal{O}_K)$ is finite, meaning that $h(\mathcal{O})$ is also finite by the above theorem.

Finally, Eric Bach proved some bounds for the representatives within the class group of a quadratic order in [2].

PROPOSITION 3.23. *Let $\mathcal{O}_K$ be the ring of integers in a quadratic number field K of absolute discriminant Δ_K. Then the class group $\mathcal{CL}(\mathcal{O}_K)$ is generated by prime ideals of norm less than $6\ln^2|\Delta_K|$.*

Proof: See Eric Bach [2, Application of Lemma 7.2]

Now suppose we have two elliptic curves E_1, E_2 whose endomorphism rings are isomorphic to $\mathcal{O}_K$, such that they are connected by some isogeny. Then we know that this isogeny corresponds to an representative in some class of $\mathcal{CL}(\mathcal{O}_K)$. Since the class is generated by small prime ideals, we only need to compute the isogenies corresponding to $\mathfrak{p}$ where $\mathfrak{p}$ is a prime $\mathcal{O}_K$-ideal of norm less than $6\ln^2|\Delta_K|$. Then we will eventually end up in the right class of $\mathcal{CL}(\mathcal{O}_K)$ and find our isogeny.

3. Galbraith's algorithm

We now have enough information to describe the algorithm by Galbraith. This section will conclude the chapter on ordinary elliptic curves by providing an algorithm for connecting two arbitrary elliptic curves using all the theory we have discussed so far. The algorithm is split into three steps, where the last one is simply a computational step. Recall that the input consists of two elliptic curves E_1, E_2, and the output is supposed to be an isogeny $\phi : E_1 \to E_2$

Step 1 Find j_1, j_2 such that j_i is at the surface of E_i for $i \in \{1, 2\}$ using the modified version of Kohel's find the floor.

Step 2 Traverse the isogeny graphs of j_1 and j_2 using random prime degree isogenies until a collision between the graphs is found.

Step 3 Compose the isogeny $\phi = \lambda \circ \phi_3 \circ \phi_2 \circ \phi_1 : E_1 \to E_2$, where $\phi_1 : E_1 \to E_1'$ and $\phi_3 : E_2' \to E_2$ are the isogenies obtained from the first step, $\phi_2 : E_1' \to E_2'$ is the isogeny from the second step, and λ is just an automorphism used to ensure that we end up with the correct end curve instead of just a curve isomorphic to E_2.

Step 1 follow from the discussion after Algorithm 4, step 3 is obvious, so let us discuss the second step further.

3.1. Traversing the isogeny graph - step 2. In this step one starts out with the j-invariants of two elliptic curves whose endomorphism rings are isomorphic to the maximal order of a imaginary quadratic number field $K := \mathbb{Q}(\sqrt{D})$.

Since a class of the class group of $\mathcal{O}_K$ corresponds to an isogeny and thus a new elliptic curve up to isomorphism we shall use the structure of $\mathcal{CL}(\mathcal{O}_K)$ to decide how to traverse this graph. Firstly we recall $\mathcal{CL}(\mathcal{O}_K)$ is generated by prime ideals of norm less than $6\ln|\Delta_K|^2$. This means that any fractional ideal $\mathfrak{a}$ is a product of prime fractional ideals $\mathfrak{p}_1, \ldots, \mathfrak{p}_t$ where the norm of $\mathfrak{p}_i < 6\ln|\Delta_K|^2$. Given the first condition, that $\left(\frac{\Delta_K}{l}\right) \neq -1$ we know that $l\mathcal{O}_K$ is ramified or splits, that is $l\mathcal{O}_K = \mathfrak{l}_1\mathfrak{l}_2$ or $l\mathcal{O}_K = \mathfrak{l}_1{}^2$.

Secondly, by Kohel, Proposition 3.5, there are $1 + \left(\frac{\Delta_K}{l}\right)$ horizontal isogenies from E, so we should limit our primes to those where $\left(\frac{\Delta_K}{l}\right) \neq -1$.

Thus we make a set of useful primes:

$$\mathcal{L} := \{\text{primes } l \mid \left(\frac{\Delta_K}{l}\right) \neq -1 \text{ and } l < 6\ln|\Delta_K|^2\}$$

Then we start by creating two graphs $\Gamma_0 = \{j_1\}, \Gamma_1 = \{j_2\}$, and iterate which one we are expanding, starting at Γ_0. We select a random prime $l \in \mathcal{L}$, compute the connected j-invariants of every $j' \in \Gamma_0$ by finding the roots of $\Phi_l(j', Y)$, and add them to Γ_0 with proper edges corresponding to the isogenies. Then do the same form Γ_1 and then alternate between expanding Γ_0 and Γ_1 picking a new random prime $l \in \mathcal{L}$ every time. This process stops as soon as $\Gamma_0 \cap \Gamma_1 \neq \emptyset$, at which point we have a path from j_1 to j_2.

Example 3.24. Continuing Example 3.2, we have an elliptic curve E with surface curve of j invariant 30. We want to connect this to another elliptic curve, say $E_2 : y^2 = x^3 + 4x + 62$. $j(E_2) = 28$ and it is on the surface. The set of primes that we are trying is

$$\mathcal{L} = \left\{l \mid \left(\frac{-23}{l}\right) \neq -1 \text{ and } l < 6\ln 23^2\right\} = \{2, 3, 13, 23, 29, 31, 41, 47\}$$

Let us sample some random primes, for simplicity we remove 2 so we don't have to verify that the image curve is on the surface. If we were to compute the 2-isogenies, we would only include the j-invariants of curves on the surface.

Initially we have $\Gamma_1 = \{30\}$ and $\Gamma_2 = \{28\}$ with no nodes connecting them.

We pick a prime $l \in \mathcal{L}$, say $l = 23$, and compute the roots of $\Phi_{23}(30, Y)$ giving us only one root (since $\left(\frac{-23}{3}\right) = 0$ and $3 \nmid 4 = [\mathcal{O}_K : \mathbb{Z}[\pi]]$): 30. Since 30 is already in Γ_1, nothing happens.

Next we pick another prime, say 13 and compute the roots of $\Phi_{13}(28, Y)$, this gives us two roots: 65 and 30. Now we are done.

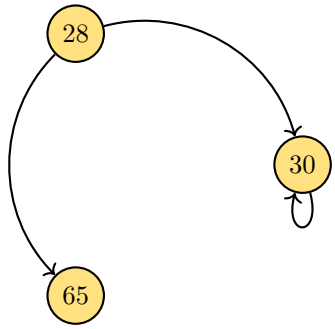

FIGURE 15. Graph of step 2

Note that the class number of $\Delta_k = -23$ is 3, so there are only three nodes on the surface.

3.2. Running time. In this section we shall compute the running time of the algorithm. We will threat each step individually and the compose the results together in the end.

Computing Φ_l. Firstly we note that our analysis depends on the class number. We have the equation

$$h(O_K) \leq \frac{1}{\pi}\sqrt{|\Delta_K|}\ln|\Delta_K|$$

Giving us a worst case bound of $h(O_K) \approx \sqrt{p}\ln p$. We let c be the conductor of $\mathbb{Z}[\pi]$ and assume that it can be computed efficiently.

We are going to compute the modular polynomial $\Phi_l(X, Y)$ for most of the primes l that are less than our bound $6(\ln|\Delta_K|^2) < 6(\ln 4p)^2$.

THEOREM 3.25 (Prime Number theorem). *Let $\pi(n)$ be the number of primes less than or equal to n, then we have*

$$\pi(n) \cong \frac{n}{\ln n}$$

Using the prime number theorem we have roughly $\frac{(\ln p)^2}{(\ln(\ln p)^2)} = \frac{(\ln p)^2}{(\ln\ln p)}$ primes in $\mathcal{L}$.

Whenever we have a prime, we need to compute the polynomial $\Phi_l(X, Y)$, which takes $O(l^3)$ operations and $O(l^2)$ storage to compute (By Elkies [13, Section 3]).

Thus computing $\Phi_l(X, Y)$ for $l \in \mathcal{L}$ is bounded by $O(\ln p^{2*3}) = O(\ln p^6)$ with space $O(\ln p^4)$. Multiplying this with the number of elements in $\mathcal{L}$ gives us a total complexity of $O(\frac{(\ln p)^8}{(\ln\ln p)})$ using $O(\frac{(\ln p)^6}{(\ln\ln p)})$ space.

We also need to compute $\Phi_l(X, Y)$ for each prime dividing c, giving us a total complexity of $O(c^3)$ with time $O(c^2)$

The final step of the computation of Φ requires finding its roots. Since Φ_l is a degree l polynomial finding the roots modulo p can be performed in probabilistic time $O(l^2 \ln p)$.

Step 1, going to the surface. For every prime l that divides the conductor $c := [\mathcal{O}_K : \mathbb{Z}[\pi]$, assume a is the largest integer such that $l^a \mid c$. Then the process of going to the surface of the l-volcano requires computing the roots of up to l^{a-1} j-values before we are certain that we are at the surface. This can be simplified to be $O(c)$. Since every l is bounded by c we have total running time of complexity $O(c^3 \ln p)$ field operations. The length of the chain from E_1 to the surface will be $O(\ln c)$.

Step 2, connecting the surface nodes. Both trees are expected to be of size $O(\sqrt{h(O_k)})$ (by the birthday problem)

LEMMA 3.26. *Step 2 is expected to terminate after $\mathcal{O}(\ln h(O_K))$ iterations.*

Proof: See Galbraith [16, Lemma 2].

Thus, by lemma above, we need roughly $\ln h(O_K)$ iterations. At each iteration we find the roots $\Phi_l(j, Y)$ for every j already in the tree, this takes time $O(l^2 \ln p) = O((\ln p)^5$. The trees are bounded by $\sqrt{h(O_K)}$ (heuristically), giving us a running time of

$$O\left(\ln h(O_K)\sqrt{h(O_K)}(\ln p^5 + \ln h)\right)$$

using space roughly $\sqrt{h(O_K)}$ for the tree and $\ln h(O_K)$ for the chain of isogenies connecting j_1 and j_2.

Step 3, composing the chains. First finding the chain connecting the j-invariants, takes time $O(\sqrt{h})$ (where $h := h(O_K)$) and provides a chain of length roughly $\ln h$, combining them with the chain of Step 1 yields a chain of length $\ln h + \ln c$. For every l-isogeny in the chain we compute the isogeny using Elkies' and Vélu's methods requiring time $O(c^3)$ for the primes $l \mid c$ and $O(\ln p^6)$ for the remaining primes $l \in \mathcal{L}$. Giving us a running total running time $O\left(\sqrt{h} + \ln h(\ln p^6 + (\ln c)c^3)\right)$ and requires storage space roughly $O\left(\ln h(\ln p^4 + \ln cc^2)\right)$.

Conclusion. In total we get an expected running time that is

$$O\left(\left(\frac{\ln p^8}{\ln \ln p} + c^3 + c^3(\ln p) + \ln h\sqrt{h}((\ln p)^5 + (\ln h)\right) + \sqrt{h} + \ln h(\ln p)^6 + (\ln c)c^3\right)$$

And requires expected space

$$O\left(\frac{(\ln p)^6}{\ln \ln p} + c^2 + \ln c + \sqrt{h} + \ln h(\ln p)^4 + (\ln c)c^2\right)$$

So in the worst case $c \approx p^{1/2}$ and h could be $p^{1/2} \ln p$, we get that the expected running time would be $O\left(p^{3/2}(\ln p)\right)$ requiring space $O\left(p(\ln p)\right)$.

If, on the other hand, the conductor is $\ln p$-smooth, the terms featuring c becomes polynomial so the algorithm will have expected running time of

$$O\left(p^{1/4}(\ln p)^{13/2}\right)$$

Also, in the case when the class number $h(O_K)$ is small the algorithm becomes polynomial.

CHAPTER 4

Supersingular Elliptic Curves and Isogenies

In this chapter we will look at supersingular elliptic curves and how computing their endomorphism rings allows us to find isogenies connecting them. We begin by describing two methods of computing the endomorphism ring given an elliptic curve. This is perhaps the most striking difference between the ordinary and the supersingular case. As we will see, there is no simple method of doing this.

In the second section we will introduce some more theory. Like we did in the chapter on ordinary curves we will now connect the isogenies of supersingular curves to lattices, but in this setting they are lattices of quaternion algebras which correspond to ideals of maximal orders. We end this section with an algorithm for computing the corresponding ideal given an isogeny of elliptic curves.

In the third section we will describe a way of computing a connecting ideal between two elliptic curves once the endomorphism ring is known for both curves. This is the Kohel-Lauter-Petit-Tignol (KLPT) algorithm [22]. We will introduce a special kind of maximal order which is very useful for sampling elements of the order before we give a high level overview of the algorithm. Next we will describe two slightly intricate steps in more detail before we wrap up the algorithm, remove the condition initially set regarding the special maximal orders, and then finally end the section with a run time analysis.

In the last section we will describe a way of going from the connecting ideal I to an actual, usable, isogeny. As we will see, this is not necessarily as straight forward as it was in the case of ordinary elliptic curves.

We recall that for a supersingular elliptic curve E, its endomorphism ring $\mathrm{End}(E)$ is isomorphic to a maximal order in a quaternion algebra over $\mathbb{Q}$ ramified at exactly p and ∞. Overall we are interested in solving the following problem.

PROBLEM. Given two supersingular elliptic curves E_0, E_1 defined over $\mathbb{F}_q$ that are connected by some isogeny, find an isogeny $\phi : E_0 \to E_1$.

REMARK. Actually we might be interested in putting some restrictions on the isogeny. For the security assumption of SIDH [14] we are looking for isogenies of a given degree l^e, while for the CGL-hash function [4] we are looking for isogenies of prime power degree. Other applications simply require the degree to be l-power-smooth, meaning that ϕ can be decomposed into isogenies of degree $l_i^{e_i}$ for small primes l_i. Nevertheless, the we will describe how to compute

an l-power isogeny, and discuss how to use this for the SIDH case later in Chapter 5 when we describe some applications.

To make things more concrete, we will explore an example with the starting curve $E_0 : y^2 = x^3 + x$ over $\mathbb{F}_q$ with $q = 439^2$ and use it throughout the chapter.

PROPOSITION 4.1. *If $p \equiv 3 \mod 4$ then the elliptic curve $E : y^2 = x^3 + x$ is supersingular over $\mathbb{F}_p$.*

Proof: See [40, Proposition 4.37]

1. Computing the endomorphism ring

In this section we would like to show how one can compute the endomorphism ring of a supersingular elliptic curve. Unlike the ordinary case where there is a clear algorithm for doing so, there does not appear to exist a deterministic and efficient algorithm for the supersingular case. As we will see, there are some elliptic curves whose endomorphism ring is known. This allows us to use the knowledge of an isogeny from any of these curves to construct the endomorphism ring of its co-domain curve, something which we will discuss in the first subsection. Then we will look at the ideas behind a generic algorithm for computing the endomorphism ring of an arbitrary supersingular curve.

We let $E/\mathbb{F}_{p^2}$ be the supersingular elliptic curve we are working over, and assume that $B = (a, b \mid \mathbb{Q})$ is the quaternion algebra isomorphic to $\mathrm{End}(E) \otimes \mathbb{Q}$ ramified at p and ∞. Since $\mathrm{End}(E)$ is a maximal order in B it is in particular generated by 4 elements. Our goal is thus to come up with 4 isogenies $\phi_i : E \to E$ that are linearly independent, generating $\mathrm{End}(E)$. Notice that multiplication by one, $[1]$, is always an endomorphism so we may choose $\phi_1 := [1]$. This leaves us with the problem of finding three endomorphisms.

Pizer[30, Proposition 5.1 and 5.2] showed us that the quaternion algebra $\mathrm{End}(E) \otimes \mathbb{Q}$ is isomorphic to $(a, b \mid \mathbb{Q})$ where a, b only depend on the characteristics p of the field we are working over . He further gave examples of a known maximal order in the various quaternion algebras. See Table 1.

Characteristics	Quaternion algebra	Maximal order
2	$(-1, -1 \mid \mathbb{Q})$	$\frac{1+i+j+k}{2}, i, j, ij$
$3 \mod 4$	$(-1, -p \mid \mathbb{Q})$	$\frac{1+j}{2}, \frac{i+ij}{2}, j, ij$
$5 \mod 8$	$(-2, -p \mid \mathbb{Q})$	$\frac{1+j+ij}{2}, \frac{1+2j+ij}{4}, j, ij$
$1 \mod 8$	$(-q, -p \mid \mathbb{Q})$	$\frac{1+j}{2}, \frac{i+ij}{2}, \frac{1}{q(j+aij)}, ij$

TABLE 1. Possible quaternion algebras for elliptic curves defined over $\mathbb{F}_{p^n}$, where $\gcd(p, q) = 1$ and a is some integer satisfying $q \mid (a^2 p + 1)$

In other words, there is an injective map $\iota_E^{-1} : \mathrm{End}(E) \to (a, b \mid \mathbb{Q})$ taking us from the language of isogenies to the language of quaternions. The isogeny corresponding to j is exactly

what one would expect, namely the p-th power Frobenius π, which can be represented on any curve. This follows easily since $j^2 = -p$, and we have that $\pi^2 = [-p]$. The isogeny corresponding to i depends on whether $i^2 = -1, -2$ or $-q$ and what curve we are looking at. For example, when $i^2 = -1$ we have $i^4 = 1$, thus i is an automorphism with inverse i^3. Furthermore i^2 is an automorphism with inverse i^2. 1 is always an automorphism, and when $j(E) \notin \{0, 1728\}$ (and $p \neq 2$), there are only 2 automorphisms so there cannot be an endomorphism corresponding to i.

EXAMPLE 4.2. We will focus on elliptic curves defined over $\mathbb{F}_{p^2} = \mathbb{F}_{439^2}$. Here we have that $q \equiv 3 \pmod 4$, so the endomorphism ring of our elliptic curves are maximal orders in the quaternion algebra $(-1, -p \mid \mathbb{Q})$. We will look at the isogenies corresponding to i when we describe the endomorphism ring of a certain elliptic curve.

Next we would like to sketch some approaches to actually compute the endomorphism ring of a given elliptic curve E.

1.1. Using a known endomorphism ring. If we assume that we already know the endomorphism ring of some starting curve E_0 and we have knowledge of an isogeny $\phi : E_0 \to E$, then we can compute the endomorphism ring of E. We begin by explicitly showing the endomorphism ring of some known supersingular curves.

PROPOSITION 4.3. *Let $E_0 : y^2 = x^3 + x$ be an elliptic curve defined over $\mathbb{F}_q$ with characteristic $p \equiv 3 \pmod 4$. Then $\iota_{E_0} : (-1, -p \mid \mathbb{Q}) \to \mathrm{End}(E_0) \otimes \mathbb{Q}$ is an isomorphism of quaternion algebras which maps the endomorphism ring of E_0 to the maximal order as in Table 1 - that is*

$$\iota_{E_0}^{-1}(\mathrm{End}(E_0)) = \left\langle \frac{1+j}{2}, \frac{i+ij}{2}, j, ij \right\rangle$$

Furthermore, the map sends the Frobenius $\pi : [x : y : z] \mapsto [x^p : y^p : z^p] \in \mathrm{End}(E_0)$ to j and $\phi_i : [x : y : z] \mapsto [-x, y\sqrt{-1}, z]$ to i.

Proof, see [1, Proposition 3.1]

REMARK. Some authors prefer to work over the maximal order generated $1, \iota, (1 + \iota \circ \pi)/2$ and $(\iota + \pi)/2$ instead of the one used above. These orders are in fact equivalent which can be shown by writing out the generator matrices for the two maximal orders (using the basis $1, i, j, ij$), and showing that the basis change matrix has determinant -1.

$$\begin{bmatrix} 1 & 0 & 1/2 & 0 \\ 0 & 1 & 0 & 1/2 \\ 0 & 0 & 0 & 1/2 \\ 0 & 0 & 1/2 & 0 \end{bmatrix} \begin{bmatrix} 1/2 & -1/2 & 0 & -1 \\ -1/2 & 1/2 & -1 & 0 \\ 0 & 1 & 0 & 2 \\ 1 & 0 & 2 & 0 \end{bmatrix} = \begin{bmatrix} 1/2 & 0 & 0 & 0 \\ 0 & 1/2 & 0 & 0 \\ 1/2 & 0 & 1 & 0 \\ 0 & 1/2 & 0 & 1 \end{bmatrix}$$

Ken McMurdy [1] has shown similar results for the remaining characteristics, but this case is sufficient for our purposes.

Next we would like to turn the isogeny ϕ and our knowledge of $\mathrm{End}(E_0)$ into knowledge of $\mathrm{End}(E)$. What we do is we compute the quaternion ideal $I \subseteq B$ corresponding to the isogeny ϕ.

This will be a left O_0-ideal and right O-ideal where $O_0 = \iota_{E_0}^{-1}(\mathrm{End}(E_0))$ and $\iota_E(O) = \mathrm{End}(E)$. Since I is a right O-ideal it will satisfy $IO \subseteq I$ and we have

$$O \subseteq O_R(I) = \{\alpha \in B \mid I\alpha \subseteq I\}$$

Since O is an order corresponding to an endomorphism ring it is maximal and thus $O = O_R(I)$. Intuitively this can be seen when thinking of elements of B as isogenies, then the only isogenies we can post-compose with our ideal are those who correspond to endomorphisms. Therefore we can simply compute $O_R(I)$ to retrieve the endomorphism ring of E.

This gives us an abstract description of our endomorphism ring as an order in a quaternion algebra. To evaluate the endomorphisms we need to use our explicit isogeny to map the endomorphisms back to our original curve, evaluate them there, and map them back again.

Remark. The reason for computing the endomorphism ring is that we want a simpler way of computing isogenies between curves. Therefore this method is useless when attempting to construct isogenies between random curves. It should however be noted the importance of this method when dealing with trapdoors and potential backdoors embedded in cryptosystems. Systems like the CGL Hash function [4] are vulnerable to attacks if the starting curve, E, is chosen in a manner which allows someone to know an isogeny from E_0 to E [29, Section 3.3]. On the other hand, signatures like SQI-Sign [15] rely on being able to compute the endomorphism through a secret isogeny from E_0 to E.

Example 4.4. We perform this computation more explicitly. Assume that we start with the elliptic curve E_0 with endomorphism ring

$$\mathrm{End}(E_0) = \left\langle \frac{[1] + \pi}{[2]}, \frac{\iota + \iota \circ \pi}{[2]}, \pi, \iota \circ \pi \right\rangle$$

and we have the isogeny ϕ of degree 5 with kernel generated by $(125, 82)$, mapping E_0 to $E : y^2 = x^3 + 163x + 400$ with $j(E) = 288$. We can embed the endomorphism ring through $\iota_{E_0}^{-1}$ as the maximal order $O_0 = \langle \frac{1+j}{2}, \frac{i+ij}{2}, j, ij \rangle$. Furthermore, using an algorithm we will describe later we get the kernel ideal I_ϕ:

$$I_\phi = \left\langle \frac{1 + 9j}{2}, \frac{i + 9ij}{2}, 5j, 5ij \right\rangle$$

This computation can be found in Example 4.15. Then the explicit description of the right order allows us to compute $O_R(I_\phi)$:

$$O_R(I_\phi) = \left\langle \frac{1 + j}{2}, \frac{i + 29ij}{10}, j, 5ij \right\rangle$$

Thus we have an abstract definition of $\mathrm{End}(E)$. Furthermore, notice that i is not in this endomorphism ring, which is exactly what one would expect since the j invariant is neither 0 nor 1728, so we cannot have automorphisms other than -1 and 1.

The main issue with our computations of the endomorphism ring is due to the fact that although π corresponds nicely with j, we don't have a representation of the endomorphism corresponding to i.

However, knowing $\phi : E_0 \to E$ and $\mathrm{End}(E_0)$ allows us to pull the points of E back to E_0 through $\widehat{\phi}$, then perform the endomorphism computations there, and then push the point forward to E through ϕ. We only need to take care of the added degree of our map, dividing out points as needed.

Waterhouse [41, Section 3.1], proposes a composition for evaluating elements of the quaternion algebra B on points of the elliptic curve by mapping them to a curve with a known embedding in B. Let $\alpha \in O = \iota_E^{-1}(\mathrm{End}(E)) \subseteq B$, and $\phi : E_0 \to E$ where we can evaluate ι_{E_0} on any $\alpha' \in O_0$. Now we can turn α into an endomorphism through the map

$$\iota_E(\alpha) := \frac{\phi \circ \iota_{E_0}(\alpha) \circ \widehat{\phi}}{[\deg \phi]}$$

When dividing P by the degree of ϕ we simply mean taking any element Q satisfying $[\deg \phi]Q = P$

Remark. The method proposed by Waterhouse is only valid for endomorphisms $\alpha \in O \cap O_0$, but since we are only interested in evaluating $\alpha \in I$ for the connecting ideal I, this requirement is satisfied. This will become clear when we discuss these ideals further later.

1.2. Using suborders. Eisenträger, Hallgreen, Leonardi, Morrison and Park [12] gives a new algorithm for computing the endomorphism ring O of a supersingular elliptic curve over $\mathbb{F}_{p^2}$ in time $O\left((\log p)^2 p^{1/2}\right)$. Explaining the entire algorithm is beyond the scope of this thesis but we shall give an overview.

As Kohel did in 1996 [21], the algorithm is based on finding a suborder of O and then using this information to construct O. The suborder is found in a specific way, giving rise to a particular family of orders which behave nicely. The suborder is what is known as a Bass order and it is contained in rather few maximal orders, enabling a brute-force search of the possible maximal orders until the correct endomorphism ring is found.

Finding the cycles that generate the Bass order is based on the work of finding elliptic curves that are l-isogenous to their p-th power Frobenius (note that we are not using the natural choice of the q-th power Frobenius). That is, we want to find an elliptic curve with j invariant j_k such that it is connected to an elliptic curve of j invariant j_k^p by a single l-isogeny. The reason for going through the Frobenius is that every isogeny can be decomposed into a separable and purely inseparable isogeny. Including every endomorphism will thus involve the inseparable Frobenius isogeny π. The interesting part is that if j and j' are adjacent, that is $\phi : j \to j'$ and $\deg \phi = l$, then j^p and j'^p are adjacent as well (for the simple reason that $\pi \circ \phi = \phi \circ \pi$).

Eisenträger et al. constructs cycles through composing l-isogenies $\phi_i : E_{i-1} \to E_i$ together starting at E_0 with $j(E_0) = j_0$ until finding a curve whose j-invariant j_k has adjacent Frobenius curve j_k^p. Then this path is reversed, going from j_k^p to j_{k-1}^p and all the way back to j_0^p. When the initial value j_0^p is found, the path is stored as P. If $j_0^p = j_0$, then this is a cycle. Otherwise

another path $P' : j_0 \to j_0^p$ is found and the composition $\widehat{P'} \circ P : j_0 \to j_0$ is a cycle. See Figure 16.

After finding two such cycles, they are highly likely to be linearly independent ([12, Theorem 3.7]). Furthermore, under the generalized Riemann hypothesis, there is a positive constant C such that there are more than $C\frac{\sqrt{p}}{\log\log p}$ j-invariants whose Frobenius j^p is adjacent to j in the l-isogeny graph ([12, Theorem 3.9]). Since there are only roughly $p/12$ supersingular curves in the graph this justifies the high probability of finding a curve with adjacent j^p node.

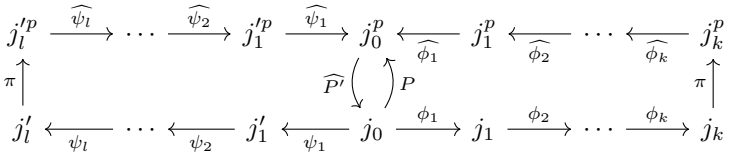

FIGURE 16. Computing j_0 to j_0^p cycles

2. Traversing theory: Ideals and orders

In this section we will make the connection between isogenies between supersingular curves and ideals of quaternion orders. This is the supersingular analogue of lattices in $\mathbb{C}$ which now becomes lattices in maximal orders O of quaternion algebras. We begin by describing the endomorphism ring slightly before we make the actual connection in the first subsection. There we will explain how to go between ideals and isogenis, how the dual isogeny corresponds to the standard involution and some ways of constructing connecting ideals between maximal orders. Finally we will look at an algorithm for computing the ideal corresponding to an isogeny once the isogeny is explicit.

Overall, the results of this section allows us to use the language of quaternions to construct ideals corresponding to nice isogenies in the next section.

PROPOSITION 4.5. *Let E_0, E_1 be isogenous supersingular elliptic curves defined over $\mathbb{F}_q$ and $l \neq \mathrm{char}(\mathbb{F}_q)$ a prime. Then there exists an isogeny of degree l^e connecting E_0 to E_1 where e is some positive integer.*

Proof: see Mestre [25, Section 2.4].

PROPOSITION 4.6. *Let E be a supersingular elliptic curve defined over $\overline{\mathbb{F}_p}$, then $j(E) \in \mathbb{F}_{p^2}$.*

In other words, if we want to sketch the graph of supersingular elliptic curves connected by some l isogeny then every node can be represented in $\mathbb{F}_{p^2}$.

EXAMPLE 4.7. To construct an example, let us use $p = 439$ which is congruent to 3 (mod 4). Then we have that $E_0 : y^2 = x^3 + x$ is supersingular with j-invariant 411 (1728 mod 439) and there are a total of $\lfloor p/12 \rfloor + 1 = 37$ supersingular elliptic curves up to isomorphism. In Figure 17 we have depicted the 7-isogeny graph starting at E_0, expanding the j-invariants of one child

at each level. Since the j-invariant is defined over $\mathbb{F}_{p^2}$ we require ω to be able to represent some of them. That is, $\mathbb{F}_{p^2} = \langle 1, \omega \rangle$.

We notice a few things right away. First, there are only 4 nodes connected to E_0. This is because curves of j-invariant 1728 have the extra automorphism ι, mapping $(x, y) \mapsto (-x, y\sqrt{-1})$, thus every isogeny $\phi : E_0 \to E$ of degree 7 also has an isomorphic (but not equal) isogeny $\phi \circ \iota$. This is not the case for $E(157)$ and $E(126)$, where there simply exists fewer non-isomorphic elliptic curves since some distinct 7-isogenies have isomorphic codomain.

Furthermore we notice that the j-invariant 411 is repeated on level 5 when choosing a "random" path. In fact, if we where to expand every node we would find 411 repeated after performing 5 7-isogenies at the earliest.

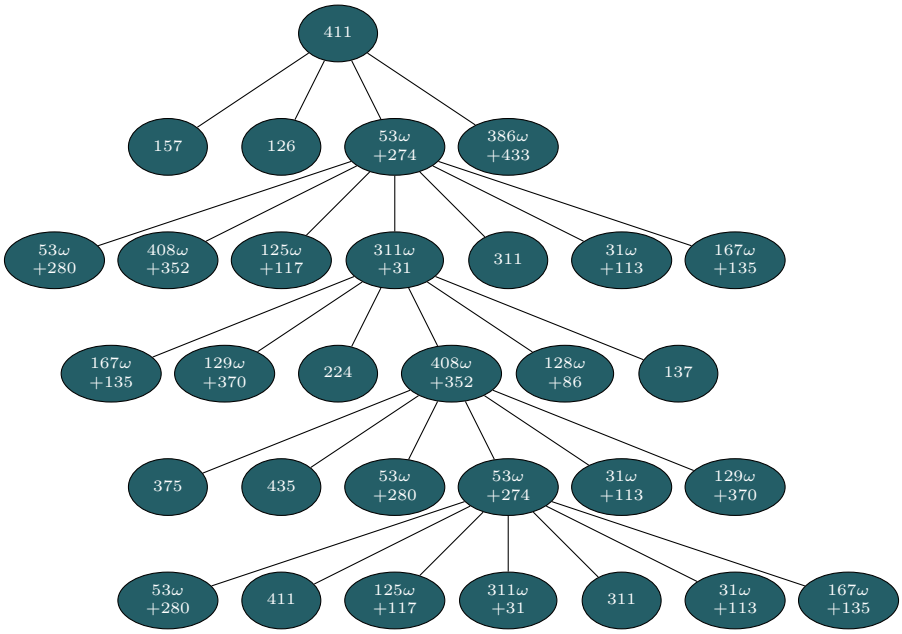

FIGURE 17. Graph of supersingular j-invariants, starting at 411 with degree 7 isogenies

2.1. Ideals and quaternion relation. In this section we will look further into the connection between ideals in maximal orders of quaternions and isogenies between supersingular curves. We will describe how the dual isogeny corresponds to the standard involution on quaternions and two useful propositions.

We already know that $\mathrm{End}(E)$ is a maximal order in the quaternion algebra $B = \mathrm{End}(E)\otimes\mathbb{Q}$ which is is ramified at p and ∞. We follow Voight [39, Chapter 42] when describing some useful results.

Suppose $\phi : E \to E'$ is a separable isogeny with finite kernel H. Furthermore let $O = \mathrm{End}(E)$ and $O' = \mathrm{End}(E')$ be the maximal orders corresponding to the endomorphism rings of E, E' in B. Then we can define the left O-ideal corresponding to H as follows:

$$I(H) := \{\alpha \in O \mid \alpha(P) = 0 \text{ for all } P \in H\}$$

We will prove that it is in fact an O, O' connecting ideal shortly.

Just like we can go from isogenies to ideals, we can go from ideals to isogenies. Let I be a left O-ideal, then we define the kernel group

$$E[I] := \cap_{\alpha\in I} E[\alpha] = \{P \in E \mid \alpha(P) = 0 \text{ for all } \alpha \in I\}$$

which will correspond to the isogeny $\phi_I : E \to E/E[I]$.

Remark. The above construction only holds when $\mathrm{nrd}(I)$ is coprime to p. This is hidden in the fact that our first intersection of $E[\alpha]$-values is not defined as the points of E which are mapped to 0 through α, but rather the group scheme kernel of α. These definitions coincide when dealing with separable isogenies, but for inseparable isogenies there is another story. Kernels of these isogenies are empty when viewed as group variety isogenies, while they are non-empty when viewed as group scheme isogenies. To get to the second definition of kernel points on E we simply split the ideal I into $I = P^a I'$ where P is the unique ideal of reduced norm p (this exists since B is ramified at p so $pO = P^2$). Then we can use the above definition on I', and map P^a to π^a.

We are nevertheless only interested in expressing separable isogenies so this is not something we need to worry to much about.

We first show that the two definitions are compatible. That is if we start with a left O-ideal I, then $I(\ker(\phi_I)) = I(E[I]) = I$.

Proposition 4.8. *The following statements hold:*

(1) $\deg\phi_I = \mathrm{nrd}(I)$
(2) $I(E[I]) = I$

Proof: See [39, Proposition 42.2.16]

Corollary 4.9. *For every isogeny $\phi : E \to E'$, there exists a left O-ideal I and an isomorphism $\rho : E_I \to E'$ such that $\phi = \rho\phi_I$. Moreover, for every maximal order $O' \subseteq B$ there exists a supersingular elliptic curve E' such that $O' \cong \mathrm{End}(E')$*

Proof: See [39, Corollary 42.2.21]

Thus we have the correspondence between ideals and isogenies. An interesting result is an application Proposition 2.60 where we saw that given a left O-ideal I, there exists an equivalent left O-ideal J satisfying

$$\mathrm{nrd}(J) \leq \sqrt{\frac{8}{\pi^2}\mathrm{discrd}(O)} = \sqrt{\frac{8p}{\pi^2}}$$

This is however not that useful to us as it only tells us that elliptic curves can be connected by an isogeny of degree smaller than $\sqrt{p}$. The isogenies may be of degree $p_1^{e_1} \dots p_n^{e_n}$ which is not what we are interested in for now. Nevertheless, the proposition does not tell us anything about how to obtain such an ideal.

Next we would like to show that the dual isogeny corresponds to the standard involution on B. That is whenever we take the conjugate in the abstract world of quaternion algebras we are really taking the dual in the concrete world of elliptic curve isogenies. Since there is unique standard involution on B, verifying that

$$\begin{aligned} \widehat{\cdot} : \mathrm{End}(E) &\to \mathrm{End}(E) \\ \phi &\mapsto \widehat{\phi} \end{aligned}$$

satisfies the requirements for a standard involution would imply that extending it to $B = \mathrm{End}(E) \otimes \mathbb{Q}$ would correspond to the conjugate on B.

The requirements are easily satisfied by our discussion on the dual isogeny. $\widehat{[1]} = [1]$, $\widehat{\widehat{\phi}} = \phi$, $\widehat{\phi \circ \psi} = \widehat{\psi} \circ \widehat{\phi}$ and $\phi \circ \widehat{\phi} = [\deg \phi] \in \mathrm{End}(E)$. Furthermore it is $\mathbb{Z}$-linear as $\widehat{\phi + \psi} = \widehat{\phi} + \widehat{\psi}$ and $\widehat{[m]} = [m]$, so $\widehat{[m] \circ \phi} = \widehat{\phi}[m] = [m]\widehat{\phi}$. Extending this to $\mathrm{End}(E) \otimes \mathbb{Q}$ gives us the standard involution on B.

Suppose now that $\phi : E \to E'$ is a separable isogeny. Then the ideal $I = I(\ker(\phi))$ is a left O ideal. Furthermore we have the isogeny $\widehat{\phi} : E' \to E$ which necessarily corresponds to $\overline{I}$. Now, $\overline{I}$ is a left O'-ideal and since conjugation on ideals simply exchange their left and right orders ($O_R(I) = O_L(\overline{I})$) we see that $O_R(I) = O'$, so I is an O, O' connecting ideal.

In particular, what we see is that any element $\alpha \in I$ is in fact in both the endomorphism ring of E and E'. This inspires the following result which gives us a way of finding a connecting ideal of two endomorphism rings. Later, during the last section of the KLPT algorithm, we will describe yet another method of computing a O, O'-connecting ideal.

LEMMA 4.10. *If O, O' are maximal orders, then OO' is a O, O'-connecting ideal.*

Proof: See [39, Lemma 17.4.7]

Thus finding a connecting ideal given endomorphism rings is easy. What remains is to turn this into a useful ideal. That is one which corresponds to an isogeny that we want. To do this we use an important, but simple, lemma.

LEMMA 4.11. *Let O be a maximal order and I be a left O-ideal of reduced norm N and $\alpha \in I$ an element of I. Then $I\frac{\overline{\alpha}}{N}$ is a left O-ideal of norm* $\mathrm{nrd}(\alpha)/N$.

PROOF. First, let us show that $J := I\frac{\overline{\alpha}}{N}$ is an O-ideal. All we need is that J is a lattice contained in O. Since I is an ideal and thus a fractional left O-ideal or rather a lattice contained in O we know that J is still a lattice. What remains is to show that it is contained in O. Since $\overline{\alpha} \in \overline{I}$ we have

$$I\overline{\alpha} \subseteq I\overline{I} = \mathrm{nrd}(I)O = NO$$

Yielding the desired result that $J \subseteq O$.

Second, to show that it has the required norm, we simply compute it directly:

$$\mathrm{nrd}\left(I\frac{\overline{\alpha}}{N}\right) = \mathrm{nrd}(I)\frac{\mathrm{nrd}(\overline{\alpha})}{N^2} = \frac{\mathrm{nrd}(\overline{\alpha})}{N} = \frac{\mathrm{nrd}(\alpha)}{N}$$

□

In particular, assuming I has norm N, if I contains an element of norm NM we can turn I into an isomorphic ideal of norm M.

Finally, we will conclude this subsection with two results that will be useful later.

PROPOSITION 4.12. *Let O be a maximal order and I a left O-ideal such that $\mathrm{nrd}(I) = N$ is prime, then $I = ON + O\alpha$ with $\gcd(N^2, \mathrm{nrd}(\alpha)) = N$.*

PROOF. By Proposition 2.39 the left fractional O-ideal I is invertible. Then by the main theorem of quaternion ideals it is locally principle, that is $I_{(p)} := I \otimes_{\mathbb{Z}} \mathbb{Z}_{(p)}$ is principal. Since I has reduced norm N we have $I\overline{I} = ON$, so in particular we have $ON \subseteq I$. Thus, locally we still have $(ON)_{(p)} \subseteq I_{(p)}$. However, since N is prime, at any $p \neq N$ we have

$$ON_{(p)} = O_{(p)} \subseteq I_{(p)} \subseteq O_{(p)}$$

where the last inclusion comes from the fact that I is a left O-ideal. In other words, locally at any $p \neq N$, we have $I_{(p)} = O_{(p)}$. Locally at N we have $I_{(N)} = O_{(N)}\alpha$ for some $\alpha \in B$. But since $I = \cap_p I_{(p)}$, we have

$$\mathrm{nrd}(I) = \bigcap_p \mathrm{nrd}(I_{(p)})$$

when we view the reduced norm as $a\mathbb{Z}$ (as opposed to just a). Then, for any $p \neq N$, we have $\mathrm{nrd}(I_{(p)}) = \mathrm{nrd}(O_{(p)}) = \mathbb{Z}$, thus must have $\mathrm{nrd}(I_{(N)}) = \mathrm{nrd}(O_{(N)}\alpha) = \mathrm{nrd}(\alpha) = N\mathbb{Z}$. So in our standard notation we have $\mathrm{nrd}(\alpha) = N$ when viewed locally away from N.

What remains to show is that $I = ON + O\alpha$. By the Local-global dictionary for lattices we see that it is enough to view the lattices locally. At any $p \neq N$, since $(ON)_{(p)} = O_{(p)}$, we have

$$I_{(p)} = O_{(p)} = (ON + O\alpha)_{(p)}$$

Then at N we have $I_{(N)} = O_{(N)}\alpha$, while

$$(ON + O\alpha)_{(N)} = ON_{(N)} + O_{(N)}\alpha = ON_{(N)} + I_{(N)} = I_{(N)}$$

since we already have that $ON_{(N)} \subseteq I_{(N)}$.

Furthermore since $\mathrm{nrd}(I) = \{\gcd(\mathrm{nrd}(\beta)) \mid \beta \in I\} = N$, we must necessarily have $\gcd(\mathrm{nrd}(N), \mathrm{nrd}(\alpha))$ N as we already know that $N \mid \mathrm{nrd}(\alpha)$ but $N^2 \nmid \mathrm{nrd}(\alpha)$ for then its reduced norm would be at least N^2 locally away from N.

□

PROPOSITION 4.13. *Let O_1 be a maximal order, O_2 be any order, I be an O_1, O_2 ideal, and J is an O_2, O_1 ideal. Then IJ is left O_1 ideal.*

PROOF. First we see that $O_L(I) \subseteq O_L(IJ)$ since for any $\alpha \in O_L(I)$ we have $\alpha I \subseteq I$, so in particular $\alpha IJ \subseteq IJ$ so $\alpha \in O_L(IJ)$.

Second, since O_1 is maximal we have $O_1 \subseteq O_L(I) \subseteq O_L(IJ)$ where $O_L(IJ)$ is also an order, so we must have equality. $O_1 = O_L(IJ)$ □

2.2. Isogeny to Ideal. As promised in a Example 4.4, we shall show how to create the ideal corresponding to an isogeny given that we know the endomorphism ring of the starting curve. We use an algorithm taken from [18, section 4.4.2] where we simplify the algorithm slightly.

Recall that we would like to construct an ideal of reduced norm N corresponding to an isogeny of degree N. We make use of the fact that any ideal of reduced norm N can be written as $I = ON + O\alpha$ for some $\alpha \in O$ since $\text{nrd}(I) = \gcd(\text{nrd}(N), \text{nrd}(\alpha))$. Thus we only need to ensure that N divide $\text{nrd}(\alpha)$. To accomplish this, we sample random α-values by sampling random elements of $\mathbb{Z}$ and creating linear combinations of the basis elements of O. To verify that it kills off the kernel of ϕ we then check that $\alpha(P) = 0$ for a kernel generator P of ϕ.

Algorithm 5: FindIdealGenerator(I, O)

Input: The kernel generator P, $b_1, \ldots, b_4$ a basis of O, and the isogeny degree N
Output: $\alpha \in O$, the element satisfying $I_\phi = ON + O\alpha$

$Q_i \leftarrow \iota_{E_0}(b_i)(P)$ for $i \in \{1, 2, 3, 4\}$;
repeat
  $a_i \leftarrow^{\$} \{1, \ldots, p-1\}$;
  $\alpha \leftarrow [a_1]b_1 + [a_2]b_2 + [a_3]b_3 + [a_4]b_4$;
until $N \mid \text{nrd}(\alpha)$ *and* $[a_2]Q_1 + [a_2]Q_2 + [a_3]Q_3 + [a_4]Q_4 = \mathcal{O}$;
return α

LEMMA 4.14. *Assuming $N < \log(p)$, Algorithm 5 runs in expected time $\tilde{O}(\log^4(p))$*[1]

Proof: See [18, Lemma 6]

EXAMPLE 4.15. Let $E_0 : y^2 = x^3 + x/\mathbb{F}_{439^2}$ and $\phi : E_0 \to E$ be an isogeny with kernel generating point $P = (125, 82)$. We begin the algorithm by computing the image points of P under $\iota_{E_0}(b_1) = \frac{[1]+\pi}{[2]}$, $\iota_{E_0}(b_2) = \frac{\iota+\iota\circ\pi}{[2]}$, $\iota_{E_0}(b_3) = \pi$ and $\iota_{E_0}(b_4) = \iota \circ \pi$ and get that they are $Q_1 = (125, 82)$, $Q_2 = (295, 416\omega + 254)$, $Q_3 = (125, 82)$ and $Q_4 = (314, 273\omega + 249)$. Next we sample random values for a_i until we find a match. The combination $a_1 = 414$, $a_2 = 24$, $a_3 = 326$, $a_4 = 191$ gives us

$$\alpha = a_1b_1 + a_2b_2 + a_3b_3 + a_4b_4 = 207 + 12i + 533j + 203ij$$

which satisfies $\text{nrd}(\alpha) = 142848815$ which clearly is divisible by 5, furthermore $\iota_{E_0}(\alpha)(P) = \mathcal{O}$.

[1] $\tilde{O}$ is the so-called soft-O notation where $\tilde{O}(f(x)) = O(f(x)\log^a(f(x)))$ for some a

Computing the actual generators of $I_\phi = O[5] + O\alpha$ can be done by performing basis reduction on the set $\bigcup_{i\in\{1,2,3,4\}}\{5b_i\} \cup \{b_i \circ \alpha\}$ This gives us that

$$I_\phi = \left\langle \frac{1+9\circ j}{2}, \frac{i+9ij}{2}, 5j, 5ij \right\rangle$$

Which has reduced norm 5, a fact which can be seen after taking the reduced norm of each generator and then their greatest common divisor.

3. KLPT Algorithm

With our knowledge of the endomorphism rings, let us now see how we can use this information to compute a connecting ideal of some prime power norm. This section will describe an algorithm by Kohel, Lauter, Petit and Tignol[22] which can turn maximal orders into an ideal. They show that if one has representatives for two maximal orders O_1, O_2, then one can find an ideal I whose left order is O_1 and right order O_2 with reduced norm l^e for a given prime l. The algorithm may also be modified to produce an ideal of smooth prime norm (ie $\text{nrd}(I) = \prod l_i^{e_i}$), which is done in [18, 15], but we will stick to the most basic notion in this section.

We begin this section by describing another kind of maximal order that we will use throughout this section for easier computations. Then we will show how we can use the structure of these orders to obtain elements of given reduced norm. Once we have dealt with this background quite thoroughly we will give an overview of the algorithm and explain the steps in more detail in the following subsections. We will end this section by wrapping up the algorithm, and explaining why the initial restriction to these special maximal orders is of no issue. Finally we will conclude with a short run time analysis of the KLPT algorithm.

The KLPT algorithm works by first creating the connecting ideal, J, of O_1 and O_2, then finding an element $\alpha \in J$ that satisfy $\text{nrd}(\alpha) = l^e/N$ where $N = \text{nrd}(J)$. This allows us to use Lemma 4.11 to make the ideal $I := J\overline{\alpha}/N$ of reduced norm l^e. Finding α is where the KLPT algorithm comes into play, where it finds special elements of O_1 using the structure of a nice suborder of O_1 to compute $\alpha \in J$ of the desired norm.

To efficiently find special elements of O_1, the algorithm requires that O_1 is a special p-extremal order. This allows us to find elements of O_1 rather easily using a special norm form that we will introduce shortly. Later we shall, as in [22], show it is not really required that O_1 nor O_2 are special p-extremal, but rather that we can find one such order O_*. Then one can simply perform the algorithm twice, going from O_1 to O_* and then O_* to O_2. We will explain this in the end of this section. Therefore we do not need to worry to much about this fact except to note that these orders exist.

For simplicity, throughout this section we will assume that $p \equiv 3 \pmod 4$. This gives us a nice representation of a maximal order for our starting curve $y^2 = x^3 + x$, and it simplifies some explanations. Note however that this is not a requirement of the algorithm.

Now, let us explain what a special p-extremal order is. Let S be the set of all maximal orders in $B = (-1, -p \mid \mathbb{Q})$ that corresponds to elliptic curves over $\mathbb{F}_{p^2}$. They all contain the Frobenius endomorphism $\pi \in O$ which satisfies $\pi^2 = -p$, so we call them p-extremal orders.

Being a special p-extremal order is related to a subring structure inside O. We define the value $d(O)$ as the discriminant of the smallest quadratic subring R that can be embedded in O, that is

$$d(O) = \min\{|\operatorname{disc}(R)| \mid \mathbb{Z} \neq R \subsetneq O\}$$

Being special p-extremal simply means being amongst the orders O that have the smallest value $d(O)$ achievable in the set S.

Definition 4.16. A maximal order O is **special p-extremal order** if it satisfies $d(O) \leq d(O')$ for any $O' \in S$

Example 4.17. For our purposes the quadratic ring which we want to embed is $\mathbb{Z}[\sqrt{-1}]$. This is the case since the maximal order $O_0 = \langle (1+j)/2, (i+ij)/2, j, ij\rangle = \langle b_1, b_2, b_3, b_4\rangle$ contains this ring of integers as $1 = 2b_1 - b_3$ and $i = 2b_2 - b_4$. Furthermore, $\mathbb{Z}[i]$ has the smallest possible absolute discriminant 4. We clearly have $\mathbb{Z} \neq \mathbb{Z}[\sqrt{-1}] \subsetneq O_0$.

Lemma 4.18. *Let O be a maximal order in B containing a subring $\mathbb{Z}\langle i, j\rangle$ with $i^2 = -q$ and $j^2 = -p$ for q coprime to p. Set $R = O \cap \mathbb{Q}(i)$ and let D be its discriminant.*

If R is the ring of integers of $\mathbb{Q}(i)$, then $R^\perp = Rj$ and $R + Rj$ is a suborder of index $|D|$ in O.

Furthermore, we have

$$\operatorname{nrd}(x_1 + y_1\omega + (x_2 + y_2\omega)j) = f(x_1, y_1) + pf(x_2, y_2)$$

where $f(x, y)$ is a principal quadratic form of discriminant D and ω is a nontrivial generator of R.

Example 4.19. Before we make a proof, let us show this with an example. $\mathbb{Z}[\sqrt{-1}] = O_0 \cap \mathbb{Q}(i)$, having discriminant -4. It is the ring of integers of $\mathbb{Q}(i)$ since $-1 \equiv 3 \mod 4$. Using the basis b_1, b_2, b_3, b_4 of O_0, any such element $aj + bij \in Rj$ can be written as $a(2b_4 - b_2) + b(2b_3 - b_1) \in Rj$, so we have $Rj = \mathbb{Z}[\sqrt{-1}]j = \mathbb{Z}j + \mathbb{Z}ij \subseteq O_0$. The principal quadratic form is $f(x, y) = x^2 + y^2$, having discriminant -4 with $\omega = i$, and the reduced norm for equation becomes obvious:

$$\operatorname{nrd}(x_1 + y_1 i + (x_2 + y_2 i)j) = x_1^2 + y_1^2 + px_2^2 + py_2^2$$

Notice however that we cannot evaluate every element of O_0 in this manner, simply because we cannot represent every $\alpha \in O_0$ in this way (for example the element b_4). Furthermore, by Example 2.15 we see that $\mathbb{Z} + \mathbb{Z}i + \mathbb{Z}j + \mathbb{Z}ij$ has index 4 in O_0.

Proof. To show the lemma we will first show that $R^\perp = Rj$, then we will show that $R + Rj$ forms a suborder of O, then we will show the norm form $\operatorname{nrd}(\beta_1 + \beta_2 j) = f(\beta_1) + pf(\beta_2)$, and finally we will show that $[O : R + Rj] = |D|$.

$\boxed{R^\perp = Rj:}$ By Proposition 2.61 we have that $\mathbb{Q}(i)^\perp = \mathbb{Q}(i)j$ in B.

Thus, by taking the orthogonal complement of R inside O we get that $R^\perp := \{\alpha \in O \mid \langle \alpha, \beta\rangle = 0 \ \forall \beta \in R\}$ is a subset of $\mathbb{Q}(i)j$.

It is a $\mathbb{Z}$-submodule of $\mathbb{Q}(i)j$ since for any $\beta, \gamma \in R^\perp$, $\alpha \in R$ and $a \in \mathbb{Z}$, we have

$$\langle \alpha, a\beta + \gamma\rangle = a^2\langle \alpha, \beta\rangle + \langle \alpha, \gamma\rangle = 0 + 0 = 0$$

It is finitely generated since $\mathbb{Z}$ is noetherian and O is finitely generated. Since $\mathbb{Q}(i)j = \mathrm{Frac}(\mathbb{Z}[i]j) = \mathrm{Frac}(Rj)$, we can find an element $\alpha \in Rj$ that clears the denominators of the generators so $R^\perp$ is a Rj fractional ideal. We can factor this ideal and write it as $\mathfrak{a}j$ where $\mathfrak{a}$ is a fractional R ideal.

Next we show that Rj is in fact contained in $R^\perp$. Let $\alpha \in Rj$, then $\alpha = aj + bij$, then for any $\beta = c + di \in R$ we have $\langle c + di, aj + bij\rangle = 0$, so $Rj \subseteq R^\perp$. Then since $R^\perp = \mathfrak{a}j$ we must have $R \subseteq \mathfrak{a}$.

What we would like to show is that $\mathfrak{a} = R$ and thus $R^\perp = Rj$, to do this we use the ramification of R together with the integrality of the order and look at the localization at p. By assumption p is ramified in B, implying that $pO_p = P^2$ where P is the maximal ideal of B_p. Thus intersecting with the image of $\mathbb{Q}(i)$ in B_p gives us that $pR = P'^2$ or $pR = P'$ for some prime ideal P'. However by Proposition 2.1 p is not ramified in R as it is coprime with q and the discriminant of R is either $-q$ or $-4q$. Thus pR remains prime. To conclude, since $\mathfrak{a}j$ is a subset of O it is in particular integral, so $\mathrm{nrd}(\mathfrak{a}j) \in \mathbb{Z}$. But since $\mathrm{nrd}(j) = p$ and p is inert, $\mathrm{nrd}(\mathfrak{a})$ must also be integral. An integral fractional ideal of R, containing R, must equal R as R is the ring of integers, so $\mathfrak{a} = R$ and we have $R^\perp = \mathfrak{a}j = Rj$.

$\boxed{R + Rj \text{ is a suborder of } O:}$ Next we would like to show that $R + Rj$ is a suborder. It is clearly a sublattice of O as both R and Rj are sublattices. What remains is to show that it is closed under multiplication. Taking an the product we see that $(R + Rj)(R + Rj) = RR + RRj + RjR + RjRj$. Clearly $RR \subseteq R$ and $RjRj \subseteq Rj$ as they are both orders themselves, but RjR is not necessarily in $R + Rj$. If we can show that $jR = Rj$, then $RjR = RRj \subseteq Rj$, so $RjR \subseteq R + Rj$ and we have that $R + Rj$ is closed under multiplication. To show this we let $\beta \in R$ be an arbitrary element, clearly $\overline{\beta} \in R$ as well. Writing it out as $\beta = t + xi$, we see that that $j\beta = \overline{\beta}j$ as $j\beta = tj + xji$ and $\overline{\beta}j = tj - xij = tj + xji$. Similarly $\beta j = j\overline{\beta}$. Thus $jR \subseteq Rj$ and $Rj \subseteq jR$ so $jR = Rj$.

$\boxed{\mathrm{nrd}(x_1 + y_1\omega + (x_2 + y_2\omega)j) = f(x, y) + pf(x, y):}$ Let $\omega = \frac{D+\sqrt{D}}{2}$, then ω is a nontrivial generator of R. Let $\beta_1, \beta_2 \in R$ be two elements, each of the form $x_i + y_i\omega$ for $x_i, y_i \in \mathbb{Z}$. Then since R and Rj are orthogonal, $\langle \beta_1, \beta_2 j\rangle = 0$ so we get

$$\mathrm{nrd}(\beta_1 + \beta_2 j) = \mathrm{nrd}(\beta_1) + \mathrm{nrd}(\beta_2 j)$$

but $\mathrm{nrd}(\beta_2 j) = p\,\mathrm{nrd}(\beta_2)$ and $\mathrm{nrd}(\beta_i) = x_i^2 + qy_i^2$. So it is a principal form of discriminant $-4q$ and we can rewrite our expression as

$$\mathrm{nrd}(x_1 + y_1\omega + (x_2 + y_2\omega)j) = f(x, y) + pf(x, y)$$

where f is the principal form $f(x, y) = x^2 + qy^2$.

$\boxed{[O : R + Rj] = |D|:}$ To conclude our lemma we need for $R + Rj$ to be of index $|D|$ in O. But we already know that it is an order and that O is maximal, thus $\mathrm{disc}(O)$ is p^2 and we have $\mathrm{disc}(R + Rj) = [O : R + Rj]^2 \,\mathrm{disc}(O)$. Therefore, the discriminant of $R + Rj$ is A^2p^2, so $[O : R + Rj]$ is necessarily $|A|$ for some A, we only need to show that $A = D$. We already know that $R + Rj$ is generated by the set $\{1, \omega, j, \omega j\}$, furthermore we use the fact that $\mathrm{trd}(\alpha\beta) = \mathrm{nrd}(\alpha + \beta) - \mathrm{nrd}(\alpha) - \mathrm{nrd}(\beta)$, allowing us to simplify a lot when comparing the

elements $1, \omega \in R$ with $j, \omega j \in Rj$ as they are perpendicular, giving us results like $\mathrm{trd}(\omega j) = 0$. We compute the discriminant of $R + Rj$ directly:

$$\begin{aligned}
\mathrm{disc}(R + Rj) &= \det(\mathrm{trd}(\alpha_i \alpha_j))_{i,j} \\
&= \det \begin{bmatrix} \mathrm{trd}(11) & \mathrm{trd}(1\omega) & \mathrm{trd}(1j) & \mathrm{trd}(1\omega j) \\ \mathrm{trd}(\omega 1) & \mathrm{trd}(\omega\omega) & \mathrm{trd}(\omega j) & \mathrm{trd}(\omega\omega j) \\ \mathrm{trd}(j1) & \mathrm{trd}(j\omega) & \mathrm{trd}(jj) & \mathrm{trd}(j\omega j) \\ \mathrm{trd}(\omega j 1) & \mathrm{trd}(\omega j\omega) & \mathrm{trd}(\omega jj) & \mathrm{trd}(\omega j\omega j) \end{bmatrix} \\
&= \det \begin{bmatrix} \mathrm{trd}(11) & \mathrm{trd}(1\omega) & 0 & 0 \\ \mathrm{trd}(\omega 1) & \mathrm{trd}(\omega\omega) & 0 & 0 \\ 0 & 0 & \mathrm{trd}(jj) & \mathrm{trd}(j\omega j) \\ 0 & 0 & \mathrm{trd}(\omega jj) & \mathrm{trd}(\omega j\omega j) \end{bmatrix} \\
&= \det \begin{bmatrix} \mathrm{trd}(11) & \mathrm{trd}(1\omega) \\ \mathrm{trd}(\omega 1) & \mathrm{trd}(\omega\omega) \end{bmatrix} \det \begin{bmatrix} \mathrm{trd}(jj) & \mathrm{trd}(j\omega j) \\ \mathrm{trd}(\omega jj) & \mathrm{trd}(\omega j\omega j) \end{bmatrix} \\
&= \det \begin{bmatrix} 2 & D \\ D & \frac{D^2 - D}{2} \end{bmatrix} \det \begin{bmatrix} -2p & -Dp \\ -Dp & \frac{-D^2 p + Dp}{2} \end{bmatrix} = p^2 D^2
\end{aligned}$$

Which gives us $\mathrm{disc}(R + Rj) = p^2 D^2$, and the index $[O : R + Rj]$ is $|D|$. □

If we have a special p-extremal order, we fix $\mathbb{Z}[i] \subseteq R$ such that $i^2 = -q, j^2 = -p$ and the discriminant of R is $-d(O)$. Since R is the quadratic ring of minimal discriminant it will be the ring of integers of $\mathbb{Q}(i)$. Furthermore, for our setting, when $i^2 = -1$ we have $R = \mathbb{Z}[i]$.

3.1. Representing integers by special orders. The usefulness of these special p-extremal orders is exactly the Lemma above. The special norm form and the integer ring structure of R allows us to construct elements of a given integer norm quite easily. This section will provide an algorithm for finding such numbers before we move on with the KLPT algorithm in the subsequent sections.

The first algorithm which we will look into lets us find an element of O which has a given integer norm M. It requires that O is special p-extremal and makes use of the suborder $R + Rj$ which in our case has discriminant $D = -4$. Furthermore the principal quadratic form of discriminant -4, $f(x, y)$, has the simple form $f(x, y) = x^2 + y^2$.

Remark. As before we simplify the algorithm by assuming that $i^2 = -1$, in the original paper this assumption is not made and the results are still valid with slight modifications.

The paper makes the rather abstract choice of having a monotone function $\Phi(x)$. It is defined by the requirement that the interval $[x, x + \Phi(x)]$ has sufficiently many primes. The word sufficient is rather non-explicit, however for the algorithm to terminate we are expected to test $2h(R)$ primes, so sufficient needs to be greater than this number (where $h(R)$ is the class number of R). Furthermore, to be able to represent M we require that $M \geq p\Phi(M)$. In other words we are not able to represent any integer element with this algorithm, but if we want an element of the form ab^c where we can choose c ourselves, we can just increase it until the

inequality is satisfied. This is indeed what the KLPT algorithm does so Algorithrm 6 is suitable for our purposes.

Algorithm 6: RepresentInteger(M)

Input: An integer M
Output: An element $\alpha + \beta j \in O$ of norm M

$m \leftarrow \lfloor \sqrt{\frac{\Phi(M)}{4}} \rfloor$;
repeat
  $(x_2, y_2) \leftarrow^r [-m, m]^2$;
  $r \leftarrow M - pf(x_2, y_2)$;
  $\mathfrak{r} \leftarrow Cornaccia(r, -4)$; `// Computes prime above` r
until r *is prime,* r *splits in* R *and* $\mathfrak{r}$ *is principal*;
$\beta \leftarrow x_2 + y_2\omega$;
$\alpha \leftarrow x_1 + y_1\omega = \mathfrak{r}$; `// Generator of` $\mathfrak{r}$
return $\alpha + \beta j$

PROPOSITION 4.20. *Algorithm 6 outputs the correct norm and runs in expected time* $2h(R)\log(M)$.

PROOF. For correctness, we notice that if we set $m = \sqrt{\Phi(M)/4}$ we have $f(x, y) < \Phi(M)$. This comes from the fact that $f(x, y) = x^2 + y^2$, so

$$f(x, y) \leq \Phi(M)/4 + \phi(M)/4 < M/2 + M/2 = M$$

Similarly we see that $\mathrm{nrd}(\beta j) = pf(x, y) < p\Phi(M) \leq M$ using the assumption on M. It is clear that the output is of desired norm since $\mathrm{nrd}(\alpha) = M - pf(x_2, y_2)$, so $\mathrm{nrd}(\alpha + \beta j) = M$.

The numbers we are testing for primality are of the form $M - pa$ where $a \in [0, \Phi(M)] \subset \mathbb{Z}$. We assume that the primes have density $1/\log(M)$ within this set. Furthermore we assume that such primes are equally likely to be split and non-split in R. Whenever a prime splits in R we assume that it is equally likely to be in any of the $h(R)$ ideal classes of R. The last thing we require is that one can sample random x, y and compute $r = M - p\,\mathrm{nrd}(x, y)$ instead of sampling r at random from the set while still maintaining the same properties.

If this holds true we are expected to test a total of $2h(R)\log(M)$ different $\beta := (x_2 + y_2\omega)$ values before successfully terminating. That is, after $\log(M)$ tries one is expected to find a prime, every second prime will be split, and every $h(R)$ split-prime will be in the trivial ideal class - those that are principal. □

EXAMPLE 4.21. We attempt to represent the integer $M = 5 * 7^5 = 84035$ with $p = 439$ with our standard maximal order O_0. Using $m = 7$ we get the result $(x_2, y_2) = (7, 3)$. Which gives us the prime $r = M - p(x_2^2 + y_2^2) = 58573$. Next we look at the ideal generated by r inside $\mathbb{Q}(i)$, and see that it splits into $(-242 + 3i)(-242 - 3i) = (58573)$. Therefore r splits in R.

Next, we represent r with $(242, 3)$ (That is $242^2 + 3^2 = 58573$). Viewing this as an ideal $(242+3i)$ in $\mathbb{Q}(i)$, we see that it is principal, ie generated by $242+3i$. Finally we simply return the element $242+3\omega+(7+3\omega)j$ which in our case is the quaternion element $242+3i+7j+3ij \in O_0$. Notice how this is indeed in O_0 since it can be written as $484b_1+6b_2-245b_3$ where $b_1 = (1+j)/2$, $b_2 = (i+ij)/2$, $b_3 = j$ and $b_4 = k$ is a basis of O_0.

3.2. Algorithm overview. Going back to the KLPT algorithm we recall that we have a O_1, O_2-connecting ideal I of reduced prime norm N. We wish to compute an equivalent ideal J of reduced norm l^e for a given prime l and some integer e (chosen by the algorithm). Assuming O_1 is special p-extremal, we can use Algorithm 6 to find elements of O_1 of a given norm. We find one such element γ of reduced norm Nl^{e_1} where we select e_1 to be the smallest possible that still satisfies $Nl^{e_1} \geq p\Phi(Nl^{e_1})$.

Example 4.22. In Example 4.21 we found an element of reduced norm $5*7^5$ with $p = 439$. With our initial attempts at finding en element with a smaller e (smaller than 5), we failed as there where not enough primes r of the form $M - pf(x, y)$. Thus we have our $\gamma = 242 + 3i + 7j + 3ij$.

We are really interested in having an element $\alpha \in I$ with a norm of similar form (Nl^e). This would allow us to use Lemma 4.11 to create the equivalent ideal $J := I\overline{\alpha}/N$ of reduced norm l^e. To solve this problem, the KLPT algorithm makes use of the fact that any left O-ideal of norm N is of the form $O\beta + ON$, in particular our ideal I is of this form. Finding β does not help us that much as it is only in O, but looking at elements in the quotient ring O/NO lets us find a $\mu \in O$ such that $\gamma\mu \equiv \beta \pmod{NO}$, thus necessarily $\gamma\mu \in I$. To find this μ, the KLPT algorithm first finds a quotient element $[\mu] \in O/NO$ that satisfies $(O\gamma/NO)[\mu] = I/NO$. This element is not a priori of the suitable norm l^{e_2}, but since we can choose the representative ourselves when lifting it to O we have just enough choice to find one that we want. Having lifted $[\mu]$ to μ with $\mathrm{nrd}(\mu) = l^{e_2}$ we can compute $\gamma\mu$ which has reduced norm $Nl^{e_1+e_2}$, satisfying $\gamma\mu \in I$.

We summarize the algorithm in the following steps:

Step 1 Ensure that $I = O\beta + ON$ has prime norm N (Algorithm 3)
Step 2 Sample $\gamma \in O$ of reduced norm Nl^{e_1} giving us the proper ideal $O\gamma$ (Algorithm 6)
Step 3 Find $[\mu] \in O/NO$ such that $\gamma\mu \equiv \beta \pmod{NO}$
Step 4 Lift to $\mu \in O$ such that $\mathrm{nrd}(\mu) = l^{e_2}$ for some e_2
Step 5 Return $I\overline{\gamma\mu}/N$, an ideal of reduced norm $l^{e_1+e_2}$ (Lemma 4.11)

Remark. For the algorithm to work, we require that N (the reduced norm of I) is relatively prime to p, D and l. Furthermore, l needs to be a quadratic non-residue modulo N. The connecting ideal I that we start with might not satisfy these requirements, but running Algorithm 3 we can turn I into one of prime norm. Since p and l are prime, the norm of I is unlikely to be the same. Furthermore $D = -4$ and I is unlikely to be of norm 2. Therefore we can assume that these requirements are satisfied without much trouble.

Step 1, 2, and 5 are already discussed. The interesting parts are step 3 and 4 which is what we will discuss next.

3.3. Finding $[\mu_0]$. Our first task is to find $[\mu_0] \in O_1/NO$ where $[\mu_0]$ is simply the notation of an element of that quotient. We want to find $[\mu_0]$ such that $(O\gamma/NO)[\mu_0] = I/NO$ and $[\mu_0]$ is a unit in Rj/NO. By Proposition 4.12 we have $I = O_1\alpha + O_1N$, thus in the quotient we have $I/NO_1 = O_1\alpha/NO_1$. We are then left with solving $O_1\gamma\mu_0 \equiv O_1\alpha \pmod{NO_1}$. Furthermore, we need to have $\mu_0 \in Rj = \mathbb{Z}j + \mathbb{Z}ij$, which is not always possible but as we will see it usually is. If it turns out that it is not, we simply choose another γ in the earlier step.

Our proofs are inspired by Dimtrij Ray's Master's thesis [10, Chapter 3].

PROPOSITION 4.23. *Let O be a special p-extremal order with distinguished subring $R = \mathbb{Z}[i]$, such that $R + Rj$ forms a suborder of O and $\gcd(p, N) = 1$, then we have the following isomorphisms*

$$O/NO \cong (R + Rj)/N(R + Rj) \cong M_2(\mathbb{Z}/N\mathbb{Z})$$

PROOF. Since O is a ring, then O/NO is also a ring. Furthermore, since $1 \in O$ we can create a homomorphism $f : \mathbb{Z}/N\mathbb{Z} \to O/NO$ by sending 1 to its residue class. Therefore O/NO is a $\mathbb{Z}/N\mathbb{Z}$ algebra. To show that it is a quaternion algebra we need for $i, j \in O/NO$ such that $i^2 = a$, $j^2 = b$ and $ij = -ji$ for $a, b \in (\mathbb{Z}/N\mathbb{Z})^\times = \mathbb{Z}/N\mathbb{Z} \setminus \{0\}$. However since $R = \mathbb{Z} + \mathbb{Z}i$ we already have the elements $1, i, j$ and ij inside our order O. In the quotient we send $i^2 = -1$ to $-1 \in (\mathbb{Z}/N\mathbb{Z})^\times$ and since p and N are relatively prime we also have $j^2 = -p$ that can be sent to $-p \in (\mathbb{Z}/N\mathbb{Z})^\times$. The non-commutative structure $ij = -ji$ follow from O.

Using the Main Theorem of quaternion algebras 2.23 we have that $O/NO \cong M_2(\mathbb{Z}/N\mathbb{Z})$ if and only if O/NO is not a division ring. However, using Algorithm 6, we can always find an element $\alpha \in O$ such that $\mathrm{nrd}(\alpha) = Na$ for some (large) integer a. So we have a nonzero element in O/NO with a $\mathrm{nrd}([\alpha]) = 0$, thus it cannot be inverted so in particular O/NO is not division and thus it is isomorphic to $M_2(\mathbb{Z}/N\mathbb{Z})$.

To show the first isomorphism we note that since O/NO is a quaternion algebra over $\mathbb{Z}/N\mathbb{Z}$ it is trivially isomorphic to

$$(R + Rj)/N(R + Rj) = \mathbb{Z}/N\mathbb{Z} + \mathbb{Z}/N\mathbb{Z}i + \mathbb{Z}/N\mathbb{Z}j + \mathbb{Z}/N\mathbb{Z}ij$$

□

EXAMPLE 4.24. We can make the ismorphism $\phi(R + Rj)/N(R + Rj) \to M_2(\mathbb{Z}/N\mathbb{Z})$ explicit in our case when $N = 5$ and $p \equiv -1 \pmod 5$ by sending the generators

$$[1] \mapsto \left[\begin{smallmatrix} 1 & 0 \\ 0 & 1 \end{smallmatrix}\right], [i] \mapsto \left[\begin{smallmatrix} 2 & 0 \\ 0 & 3 \end{smallmatrix}\right], [j] \mapsto \left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right] \text{ and } [ij] \mapsto \left[\begin{smallmatrix} 0 & 2 \\ 3 & 0 \end{smallmatrix}\right]$$

Since $O\gamma$ and I are (proper) left O-ideals we have that $O\gamma/NO$ and I/NO are proper ideals. They are both nonzero since $NO \subsetneq O\gamma$ and $NO \subsetneq I$. Using the explicit isomorphism ϕ defined in the example above, we can solve the equation $(O\gamma/NO)[\mu_0] = I/NO$ for $[\mu_0]$ by mapping the principal ideals $O\gamma/NO$ and I/NO to matrices $C, D \in M_2(\mathbb{Z}/N\mathbb{Z})$ and then solving the equation $CX = D$ for X. Then mapping the solution back using $\phi^{-1}(X)$.

EXAMPLE 4.25. Continuing with our $\gamma = 242 + 3i + 7j + 3ij$, we construct the ideal $O\gamma$ which then is generated by $\{\frac{1+74071j+120620ij}{2}, \frac{i+47450j+74071ij}{2}, 84035j, 84035ij\}$. Then we have $I = O\alpha + ON$, where $\alpha = 207 + 12i + 533j + 203ij$. Writing these with respect to the generators

of O gives us that $O\gamma = \{b_1 + 37035b_3 + 60310b_4, b_2 + 23725b_3 + 37035b_4\}$ and $\alpha = 414b_1 + 24b_2 + 326b_3 + 191b_4$ where b_i are the generators of O.

We first map these to the quotient modulo $5O$ which gives us that

$$O\gamma/5O = \langle b_1, b_2\rangle = \langle 6b_1, 6b_2\rangle = \langle 3 + 3j, 3i + 3ij\rangle$$

and

$$\alpha/5O = 4b_1 + 4b_2 + 1b_3 + 1b_4 = 2 + 2i + 3j + 3ij$$

Choosing our representatives carefully as to be within $R + Rj = \mathbb{Z} + \mathbb{Z}i + \mathbb{Z}j + \mathbb{Z}ij$

Next we map these generators to $M_2(\mathbb{Z}/5\mathbb{Z})$ using the ϕ map from the above example.

$$\phi(O\gamma/5O) = \left\langle \begin{bmatrix} 3 & 3 \\ 3 & 3 \end{bmatrix}, \begin{bmatrix} 1 & 1 \\ 4 & 4 \end{bmatrix} \right\rangle = \left\langle \begin{bmatrix} 2 & 2 \\ 4 & 4 \end{bmatrix} \right\rangle \qquad \phi(\alpha/5O) = \begin{bmatrix} 1 & 4 \\ 2 & 3 \end{bmatrix}$$

Finally we need to find $[\mu_0]$ such that $O\gamma/5O[\mu_0] = \alpha/5O$. To do this we simply solve for $\phi([\mu_0])$ and find the solution $\left[\begin{smallmatrix} 0 & 1 \\ 4 & 0 \end{smallmatrix}\right]$. This can be mapped back to $3ij$ in O/5O.

Now that we have found a $[\mu_0]$ we are almost done. We only need to make sure that we can find a $[\mu_0]$ which a unit in $Rj/(NO \cap Rj)$. To show this we make use of some linear algebra and group theory. The following results are only necessary to prove the existence of $[\mu_0]$ and gives little insights into how we can actually find it. The reader may therefore safely move ahead until the next subsection if he can believe in the existence of such a $[\mu_0]$.

Henryk Minc describes in his lecture notes [26, Theorem 1 & 2] that the ring $M_2(\mathbb{Z}/N\mathbb{Z})$ is a principal left-ideal ring. Furthermore, during the proof of Theorem 2 he shows that the rank of the generating matrices of nonzero proper left ideals is 1.

LEMMA 4.26. *The ring $M_2(\mathbb{Z}/N\mathbb{Z})$ has $N + 1$ nontrivial proper left ideals*

PROOF. Since the proper left ideal generators of $M_2(\mathbb{Z}/N\mathbb{Z})$ have rank 1, if $A' = \left[\begin{smallmatrix} 0 & 0 \\ a & b \end{smallmatrix}\right]$ is a generator of the ideal, then $A = \left[\begin{smallmatrix} a & b \\ 0 & 0 \end{smallmatrix}\right] = \left[\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right] A'$ is also a generator. Therefore we can assume that the ideal is generated by $A = \left[\begin{smallmatrix} a & b \\ 0 & 0 \end{smallmatrix}\right]$ where at least one of a, b is nonzero. Thus a total of $N^2 - 1$ generator matrices exist.

Furthermore, let $B = \left[\begin{smallmatrix} b_1 & b_2 \\ b_3 & b_4 \end{smallmatrix}\right] \in M_2(\mathbb{Z}/N\mathbb{Z})$ be an arbitrary matrix. Then $BA = \left[\begin{smallmatrix} ab_1 & bb_1 \\ ab_3 & bb_3 \end{smallmatrix}\right]$. In particular any nonzero scalar multiple of A generates the same left ideal. That is $\left[\begin{smallmatrix} a & b \\ 0 & 0 \end{smallmatrix}\right]$ and $\left[\begin{smallmatrix} \lambda a & \lambda b \\ 0 & 0 \end{smallmatrix}\right]$ generate the same ideal if $\lambda \in (\mathbb{Z}/N\mathbb{Z})^\times$. There are $N - 1$ such scalars, so we end up with $(N^2 - 1)/(N - 1) = N + 1$ possible left ideals. □

COROLLARY 4.27. *The ring O/NO has $N + 1$ nontrivial left ideals*

REMARK. Notice how an ideal of O/NO corresponds to a generator $[\alpha] \in O/NO$ which can be lifted to $O\alpha + ON$ which then corresponds to an isogeny of degree N. We already know that there are $N + 1$ isogenies of degree N, so this corollary should come as no surprise.

Next we will describe an action on the left O/NO ideals, which will allow us to determine if there is a solution to the equation

$$(O\gamma)\mu_0 = O\alpha \pmod{NO}$$

with respect to μ_0.

LEMMA 4.28. *Let R be the distinguished subring of O and $\mathcal{L}$ be the set of all nontrivial proper left O/NO ideals. Then the map*

$$\rho : \mathcal{L} \times (R/NR)^\times \to \mathcal{L}$$
$$(I, \beta) \to I\beta$$

defines a group action with kernel $(\mathbb{Z}/N\mathbb{Z})^\times$. Furthermore

(1) If N is split in R, the group action has an orbit of size $N-1$ and two fixed points
(2) If N is inert in R, the group action has only orbit.

PROOF. $\boxed{\rho \text{ is a group action:}}$ Let $I \in \mathcal{L}$ be a left ideal generated by the element α, and $\beta \in (R/NR)^\times$ be any element. Then

$$\rho(I, \beta) = I\beta = (O\alpha)\beta = (O\alpha\beta) \pmod{NO}$$

Which is also a nontrivial proper left O/NO ideal unless $\alpha\beta = 0$, but since $\beta \neq 0$ and N is prime this cannot happen. So the map is clearly well defined. To show that it is a group action we need to ensure that $\rho(I, 1) = I$ and that $\rho(\rho(I, \beta), \gamma) = \rho(I, \beta\gamma)$. The first property follows immediately. The second can similarly be seen quite easily:

$$\rho(\rho(I, \beta), \gamma) = \rho((I\beta), \gamma) = (I\beta)\gamma) = I(\beta\gamma) = \rho(I, \beta\gamma)$$

$\boxed{\ker(\rho) = (\mathbb{Z}/N\mathbb{Z})^\times:}$ Clearly $(\mathbb{Z}/N\mathbb{Z})^\times$ is in the kernel since $Ia = aI = I$ as I is a left O/NO-ideal and $\mathbb{Z}/N\mathbb{Z}$ is in the center of the quaternion algebra O/NO. What remains to show is that no element of the form $b_1 + b_2 i$ with $b_i \in \mathbb{Z}/N\mathbb{Z}$ and $b_2 \neq 0$ is in the kernel. Suppose I is generated by $\alpha = a_1 + a_2 i + a_3 j + a_4 ij$ and $b_2 i$ for $b_2 \in (\mathbb{Z}/N\mathbb{Z})^\times$

$$\alpha b_2 i = -a_2 b_2 + a_1 b_2 i + a_4 b_2 j - a_3 b_2 ij = b_2(-a_2 + a_1 i + a_4 j - a_3 ij)$$

Since $b_2 \in (\mathbb{Z}/N\mathbb{Z})^\times$ it is invertible, so the ideal generated by $\alpha b_2 i$ is also generated by $-a_2 + a_1 i + a_4 j - a_3 ij$. Writing them as matrices with the basis $1, i, j, ij$ we look for a solution matrix C to the equation

$$C \begin{bmatrix} -a_2 & 0 & 0 & 0 \\ 0 & a_1 & 0 & 0 \\ 0 & 0 & a_4 & 0 \\ 0 & 0 & 0 & -a_3 \end{bmatrix} = \begin{bmatrix} a_1 & 0 & 0 & 0 \\ 0 & a_2 & 0 & 0 \\ 0 & 0 & a_3 & 0 \\ 0 & 0 & 0 & a_4 \end{bmatrix}$$

Solving the equations we see that the solution matrix is

$$C = \begin{bmatrix} -a_2^{-1}a_1 & 0 & 0 & 0 \\ 0 & a_1^{-1}a_2 & 0 & 0 \\ 0 & 0 & a_4^{-1}a_3 & 0 \\ 0 & 0 & 0 & -a_3^{-1}a_4 \end{bmatrix}$$

Thus in particular, if only one of a_1 and a_2 is zero or only one of a_3 and a_4 is zero, then the ideal generated by $\alpha b_2 i$ is not the same as the ideal generated by α, so $b_2 i \notin \ker(\rho)$. Furthermore if we assume I is one of those ideals that satisfy $Ib_2 i = J \neq I$, then we have

$$I(b_1 + b_2 i) = Ib_1 + Ib_2 i = I + J \neq I$$

In particular, there exists ideals, I, where for any $\beta = b_1 + b_2 i$ with $b_2 \neq 0$ $\rho(I, \beta) \neq I$. Thus the kernel must be $(\mathbb{Z}/N\mathbb{Z})^\times$

$\boxed{N \text{ is split in } R:}$ Assuming that N splits we can write it as $N = (a_1 + a_2 i)(a_1 - a_2 i)$ for $a_1, a_2 \in \mathbb{Z} \setminus \{0\}$. Thus both a_1 and a_2 are invertible modulo $N\mathbb{Z}$ and $(R/NR)^\times$ acts trivially on the left ideals $(a_1 + a_2 i)$ and $(a_1 - a_2 i)$ by the above discussion. Thus these are fixed points of ρ. Let $\alpha \in (R/NR)$ be an element that is not a multiple of $a_1 + a_2 i$ nor $a_1 - a_2 i$, then $\gcd(\text{nrd}(\alpha), N) = 1$, so we can find its inverse: $\overline{\alpha}/\text{nrd}(\alpha)$. Thus α is a unit in R/NR, and we have $|(R/NR)^\times| = N^2 - 1 - 2(N-1) = (N-1)^2$. Since there are ideals where the only stabilizer is $(\mathbb{Z}/N\mathbb{Z})^\times$ by the orbit stabilizer theorem their orbit is

$$[(R/NR)^\times : (\mathbb{Z}/N\mathbb{Z})^\times] = \frac{|(R/NR)^\times|}{|(\mathbb{Z}/N\mathbb{Z})^\times|} = \frac{(N-1)^2}{N-1} = N - 1$$

Furthermore since the size of $\mathcal{L}$ is $N+1$ and we already have two fixed points, we conclude that there are exactly two fixed points and one orbit of size $N-1$.

$\boxed{N \text{ is inert in } R:}$ Since N is inert, we have $|(R/NR)^\times| = N^2 - 1$ as every nonzero element is a unit. Again, using the orbit stabilizer theorem and the fact that there exists an ideal whose only stabilizer is $(\mathbb{Z}/N\mathbb{Z})^\times$ we have

$$[(R/NR)^\times : (\mathbb{Z}/N\mathbb{Z})^\times] = \frac{|(R/NR)^\times|}{|(\mathbb{Z}/N\mathbb{Z})^\times|} = \frac{N^2 - 1}{N-1} = N + 1$$

Again by the size of $\mathcal{L}$ we conclude that this is the only orbit. □

Using this group action we can show the solvability of our earlier equation.

PROPOSITION 4.29. *Let $\alpha \in I$, $\gamma \in O$. Then the equation*

$$(O\gamma)\mu_0 = O\alpha \pmod{NO}$$

is always solvable for $\mu_0 \in Rj$ whenever N is inert, and solvable with probability $(N^2 - 2N + 3)/(N+1)^2$ assuming that α and γ are chosen at random.

PROOF. Since j is a unit, the lemma above also holds for the action of $(R/NR)^\times[j]$ where $[j]$ is a representative of j in O/NO.

$\boxed{N \text{ is inert:}}$ By the lemma above both ideals $O\gamma/NO$ and $O\alpha/NO$ lie in the same orbit. In particular, there exists some $\mu_0 \in (Rj/NRj)^\times$ such that $\rho(O\gamma/NO, \mu_0) = O\alpha/NO$

$\boxed{N \text{ is split:}}$ Again, by the lemma above, if $O\gamma/NO$ and $O\alpha/NO$ are in the orbit of size $N-1$ there exists a solution, so among the $(N+1)^2$ possible ideal-combinations that we can get from $O\gamma/NO$ and $O\alpha/NO$, $(N-1)^2/(N+1)^2$ have a solution. j can either swap the fixed

points or keep them in place. Nevertheless there are only 2 out of $(N+1)^2$ ideal-combinations. Thus giving us the total probability of

$$\left(\frac{N-1}{N+1}\right)^2 + \frac{2}{(N+1)^2} = \frac{N^2-2N+3}{(N+1)^2}$$

□

In other words we are highly likely to be able to find a μ_0 such that $\mu_0 \in Rj$ which solves the equation $(O\gamma/NO)[\mu_0] = I/NO$. If we are not able to find this μ_0 we backtrack and sample another γ instead.

3.4. Lifting $[\mu_0]$ to its $R + Rj$ representative μ. Recall that we want to lift $[\mu_0] \in (R/NR)^*[j]$ to some $\mu \in O$ satisfying $\mathrm{nrd}(\mu) = l^{e_1}$ for some positive e_1 given a prime l. As before, we make use of the norm form on our special p-extremal order O

$$\mathrm{nrd}(x_1 + y_1\omega + (z_1 + w_1\omega)j) = f(x_1, y_1) + pf(z_1, y_1)$$

We first lift $[\mu_0]$ to a representative $\mu_0 \in R + Rj$ (noting that μ_0 will be entirely in Rj).

Example 4.30. With our $[\mu_0] = ij \in Rj/NRj$ this is trivially lifted to $\mu_0 = ij \in R + Rj$.

The canonical lift of $[\mu_0]$ is unlikely to be of the correct norm. To remedy this we make use of the freedom we are given with respect to the requirements of $[\mu_0]$. We can replace it with an element of the form $\mu = \lambda\mu_0 + N\mu_1$ for $\lambda \in \{1, \ldots, N-1\}$ and $\mu_1 \in R + Rj$. The only requirement of $[\mu_0]$ is that $(O\gamma/NO)[\mu_0] = I/NO$, which is also satisfied for our new element:

$$(O\gamma/NO)[\lambda\mu_0 + N\mu_1] = (O\gamma/NO)[\lambda\mu_0] = \lambda(O\gamma/NO)[\mu] = \lambda I/NO = I/NO$$

Writing $\mu_0 = (z_0 + w_0\omega)j$ and $\mu_1 = x_1 + y_1 + (z_1 + w_1\omega)j$ and then $\mu = \lambda\mu_0 + N\mu_1$ we get the simplified norm equation

$$\mathrm{nrd}(\mu) = f(Nx_1, Ny_1) + pf(\lambda z_0 + Nz_1, \lambda w_0 + Nw_1) = l^{e_1} \tag{1}$$

We first lift $[\mu_0]$ to its canonical representative μ_0 and look at (1) modulo N. This gives us

$$\lambda^2 pf(z_0, w_0) = l^{e_1} \mod N$$

Since l is a quadratic residue mod N we choose the parity of the exponent e_1 depending on whether $pf(z_0, w_0)$ is a quadratic residue or not. Next we solve the square root equation $\lambda^2 \equiv l^{e_0}/pf(z_0, w_0) \pmod N$ and choose representative of λ within the range $1, \ldots, N-1$.

Example 4.31. Continuing our example we have $\mu_0 = (z_0 + w_0 i)j = (0 + 3i)j$ and $p = 439 \equiv 4 \pmod 5$ giving us

$$pf(z_0, w_0) = 439 * (3^2) \equiv 1 \pmod 5$$

Since 1 is a quadratic residue modulo 5, we have that e_1 needs to be even. $1^{-1} \equiv 1 \pmod 5$ so $\lambda^2 \equiv 7^{e_1} \pmod 5$. The smallest possible e_1 is 2, however if we choose this value we need to repeat with a larger one. Therefore we cheat and choose 8 right away, giving us $\lambda^2 \equiv 7^8 \equiv 1 \pmod 5$. Solving for λ gives us the possibilities $1, 4$. We select an arbitrary value, for example $\lambda = 1$.

Rewriting equation (1), using the equation $\operatorname{nrd}(\alpha+\beta)-\operatorname{nrd}(\alpha)-\operatorname{nrd}(\beta)=\langle\alpha,\beta\rangle$ we get

$$\langle\lambda\mu_0, N\mu_1\rangle = l^{e_1} - pf(z_0,w_0) \bmod N$$

We can simplify the inner product as follows:

$$\begin{aligned}\langle\lambda\mu_0, N\mu_1\rangle &= \operatorname{nrd}(\lambda\mu_0+N\mu_1)-\operatorname{nrd}(\lambda\mu_0)-\operatorname{nrd}(N\mu_1)\\ &= \operatorname{nrd}(\lambda(z_0+w_0\omega)j)+N(z_1+w_1\omega)j)\\ &\quad-\lambda^2\operatorname{nrd}((z_0+w_0\omega)j)-N^2\operatorname{nrd}((z_1+w_1\omega)j)\\ &= p\lambda N(2z_0z_1+z_0w_1(\overline{\omega}+\omega)+z_1w_0(\omega+\overline{\omega}+2w_0w_1\omega\overline{\omega})\\ &= p\lambda N(2z_0z_1+\operatorname{trd}(\omega)(z_0w_1+z_1w_0)+2\operatorname{nrd}(\omega)w_0w_1)\end{aligned}$$

Where the second line follows from the fact $N\mu_1$ is the only part of the sum that contains components from the first two dimensions. They will therefore necessarily subtract to zero when subtracting $\operatorname{nrd}(N\mu_1)$. The simplification follows from expanding the reduced norm equations, replacing $\operatorname{nrd}(\alpha)$ with $\alpha\overline{\alpha}$, recalling that $\overline{\omega j}=\overline{j}\overline{\omega}=-j\overline{\omega}$, and that $z_0,z_1,w_0,w_1,\lambda,N$ and p commute with ω and j. Thus giving us the equation

$$p\lambda\left(2z_0z_1+\operatorname{trd}(\omega)(z_0w_1+z_1w_0)+2\operatorname{nrd}(\omega)w_0w_1\right)=\frac{l^{e_1}-pf(z_0,w_0)}{N} \mod N$$

Since N is coprime to $w_0, z_0, |D|$ and p, this equation has exactly N solutions (z_1,w_1). We choose a random one among them that satisfy

$$|\lambda z_0+Nz_1|<N^2 \quad \text{and} \quad |\lambda w_0+Nw_1|<N^2$$

Example 4.32. We continue our example and note that since our $\omega=i$ we get that $\operatorname{trd}(\omega)=i+\overline{i}=i-i=0$ and $\operatorname{nrd}(\omega)=i(-i)=-i^2=1$. Furthermore we still have $pf(z_0,w_0)=3951$ and $p\equiv 4 \pmod 5$ and our chosen $\lambda=1$. Thus the expression becomes $4(2*0*z_1+2*3*w_1)=\frac{7^8-3951}{5}=1152170\equiv 0 \pmod 5$ which can be simplified to

$$w_0\equiv 0 \pmod 5$$

That is we have the 5 solutions $(0,0),(1,0),(2,0),(3,0)$ and $(4,0)$ for (z_1,w_1). We would like for z_1 to satisfy $|\lambda z_0+Nz_1|=|5*z_1|<5^2=25$. This gives us the options $z_1\in\{0,1,2,3,4\}$. The second requirement is immediately satisfied as $w_0=0$. Thus every solution for z_1,w_1 satisfy the requirements. We pick an arbitrary tuple, $(z_1,w_1)=(0,0)$.

The final part of finding μ is to solve the equation when not working modulo N. That is we need to solve

$$f(x_1,y_1)=\frac{l^{e_1}-pf(\lambda z_0+Nz_1,\lambda w_0+Nw_1)}{N^2}$$

This can be solved with Cornaccia's algorithm, Algorithm 1, since we know f has discriminant D. If the fraction is negative, we simply need to increase our e_1 chosen earlier and repeat the steps. Conraccia's algorithm will either be able to solve this equation or determine that no solution exist. If no such solution exist we chose a new value for (z_1,w_1).

EXAMPLE 4.33. Continuing our example we have

$$f(x_1, y_1) = \frac{7^4 - 439f(3*4+5*1, 3*4+5*1)}{5*5} = 230434$$

Using Cornaccia's algorithm we get that $f(455, 153)$ satisfy this expression. Thus we have found our element $\mu_0 = 3ij$ and $\mu_1 = 455 + 153i$, giving us the lift

$$\mu = \lambda\mu_0 + N\mu_1 = 2275 + 765i + 3ij$$

For the running time of this algorithm, we assume that the values

$$\frac{l^{e_1} - pf(\lambda z_0 + Nz_1, \lambda w_0 + Nw_1)}{N^2}$$

behave like random numbers close to $N^4|D|p$, thus we expect to choose $\log(N^4|D|p)h(D)$ values before finding a solution.

The authors expect the value e_1 to be of size $\log_l(N^4|D|p) \approx 3\log_l(p)$, and note that it should be selected as the minimal value satisfying the parity requirement, incrementing it if needed. We notice that this approximation fits well with our running example since $\log_7(5^4 * 4 * 439) \approx 7.1$ while we required to choose $e_1 = 7$ to find a solution.

3.5. Wrapping up the algorithm. Now we have found $\gamma, \mu \in O$ with the properties $\text{nrd}(\gamma) = Nl^{e_0}$ and $\text{nrd}(\mu) = l^{e_1}$, thus our element $\beta := \gamma\mu \in I$ has reduced norm $\text{nrd}(\beta) = Nl^e = Nl^{e_0+e_1}$. Multiplying with I gives us the equivalent left ideal $J := I\overline{\beta}/N$ of prime power norm l^e.

EXAMPLE 4.34. Continuing our example we have $\gamma = 247+90i+3j+5ij$ and $\mu = 2275+765i+3ij$ giving us that

$$\beta = \gamma\mu = 544304 + 201174i + 18211j + 2196ij$$

Computing its reduced norm we see that $\text{nrd}(\beta) = 5 * 7^{13}$, which is what we would expect since $e_0 = 5$ and $e_1 = 8$. Finally scaling the ideal

$$I = \langle 1/2 + 9/2j, 1/2i + 9/2ij, 5j, 5ij \rangle$$

with $\overline{\beta}/N$ gives us the ideal $J = \langle \alpha_1, \alpha_2, \alpha_3, \alpha_4 \rangle$ of reduced norm 7^{13} with

$$\begin{aligned}
\alpha_1 &= \frac{14499193 - 1775514i + 976105j + 361674ij}{2} \\
\alpha_2 &= \frac{1775514 + 14499193i - 361674j + 976105ij}{10} \\
\alpha_3 &= 7994629 - 964044i + 544304j + 201174ij \\
\alpha_4 &= 964044 + 7994629i - 201174j + 544304ij
\end{aligned}$$

We can also verify that the right order of J is the same as the right order of I which is what we would expect ($O_R(I) = O_R(J)$).

Under the assumption that O is a special p-extremal order we conclude the KLPT algorithm with the following theorem.

THEOREM 4.35. *Let O be a special p-extremal maximal order in quaternion algebra B, and let l be a small prime. Then there exists a probabilistic algorithm which takes as input a left O-ideal I and outputs an isomorphic left ideal O-ideal J satisfying* $\mathrm{nrd}(J) = l^e$ *for some positive integer e. Furthermore the exponent e is of size*

$$e \approx \log_l(Np\Phi(p)|D|) + \log_l(N^4|D|p) - \log_l(N^2)$$

Proof: See [22, Theorem 7]

We can simplify the equation for the exponent approximation by assuming that $\log_l(N) \approx \frac{1}{2}\log_l(p)$ and that $\Phi(p) \approx \log(p)^n$ a factor completely negligible in this case, thus giving us $e \approx \frac{7}{2}\log_l(p)$

EXAMPLE 4.36. Our example produce an isogeny with $e = 13$, while the theorem gives us the approximation $\log_7(439) * 7/2 \approx 11$ which is not that far away.

3.6. Removing the need for special p-extremal orders. Going back to our initial claim that we don't need for O_1 to be a special p-extremal maximal order, we will now prove that we can use the algorithm even if this is not the case. Actually, all we need is to run the algorithm twice to get our desired result.

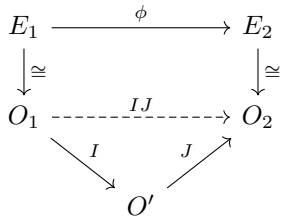

FIGURE 18. Computing the ideal connecting O_1 to O_2

That is, we find the connecting ideals I and J satisfying $O_L(I) = O_1$, $O_R(I) = O'$, $O_L(J) = O'$ and $O_R(J) = O_2$. Clearly these are compatible, and their product IJ satisfy $O_L(IJ) = O_1$ and $O_R(IJ) = O_2$. This is easily seen since $O_1 I \subseteq I$, so $O_1 \subseteq O_L(IJ)$, similarly $O_2 \subseteq O_R(J) \subseteq O_R(IJ)$. Since O_1 and O_2 are maximal and O_L and O_R are orders, these inclusions must be equalities.

The only remaining thing to show is that one can in fact find this $I : O_1 \to O'$ when O' is the special p-extremal order. The KLPT algorithm will a priori only construct it in the other direction.

LEMMA 4.37. *Let O_1, O_2 be maximal orders of a quaternion algebra B. Then the Eichler order $O_1 \cap O_2$ satisfy $[O_1 : O_1 \cap O_2] = M = [O_2 : O_1 \cap O_2]$ and the set*

$$I(O_1, O_2) := \{\alpha \in B \mid \alpha O_2 \overline{\alpha} \subseteq MO_1\}$$

is a left O_1-ideal and a right O_2-ideal of reduced norm M. Conversely if I is a left O_1 ideal with right order O_2 such that $I \not\subseteq nO_1$ for any $n > 1$, then $I = I(O_1, O_2)$.

Proof: See [22, Lemma 8].

Using the lemma above we can easily find a connecting ideal between maximal orders. This along with the Theorem below allows us to construct the ideal I in the opposite direction.

THEOREM 4.38. *Let O_1 and O_2 be maximal orders in a quaternion algebra B, and let l be prime. If there exists an algorithm which takes as input a left O_1 ideal and computes an equivalent left O_1 ideal of l-power norm, there exists an algorithm which takes as input a left O_2 ideal and computes an equivalent left O_2 ideal of l-power norm.*

PROOF. Let A be an algorithm that turns a left O_1 ideal I into an element $\gamma \in I$ such that $I\overline{\gamma}/\operatorname{nrd}(I)$ has l power norm.

Given left O_2 ideal J, set $I = I(O_1, O_2)$. Note that I is an O_1, O_2 ideal, so by Proposition 4.13 IJ is a left O_1 ideal. Using our algorithm we get the elements $\gamma_1 \leftarrow \mathsf{A}(I)$ and $\gamma_2 \leftarrow \mathsf{A}(IJ)$. such that $I_1 = I\overline{\gamma_1}/\operatorname{nrd}(I)$ for $\gamma_1 \in I$ and $I_2 = IJ\overline{\gamma_2}/\operatorname{nrd}(IJ)$ with $\gamma_2 \in IJ$. These elements satisfy reduced norms $\operatorname{nrd}(\gamma_1) = \operatorname{nrd}(I)l^{e_1}$ and $\operatorname{nrd}(\gamma_2) = \operatorname{nrd}(IJ)l^{e_2}$. Thus the element $\gamma := \overline{\gamma_1}\gamma_2/\operatorname{nrd}(I)$ has reduced norm $\operatorname{nrd}(J)l^{e_1+e_2}$.

What remains to show is that the element $\overline{\gamma_1}\gamma_2$ is in fact in J. First we make use of the fact that $\overline{I}I = \operatorname{nrd}(I)O_2$. Writing γ_2 as $\alpha\beta$ with $\alpha \in I$ and $\beta \in J$ we therefore get that $\gamma' := \overline{\gamma_1}\alpha/\operatorname{nrd}(I) \in O_2$. Second, since $O_2 = O_L(J)$, we have $\gamma := \gamma'\beta \in J$.

Therefore, the element γ, can be used to produce the ideal

$$J' := J\frac{\overline{\gamma}}{\operatorname{nrd}(J)}$$

which is equivalent to J, and of reduced norm $l^{e_1+e_2}$ □

Thus to be able to find the ideal $I : O_1 \to O'$, one instead performs the algorithm on the two ideals $J := I(O', O_1)$ and JI and use the outputs to find our ideal J' of reduced l-power norm.

3.7. Run time analysis. For the run time analysis we make the same estimations that are done in [22]. That is, we assume that $N \approx \sqrt{p}$, that $\Phi(p) \approx \log(p)^n$ for some integer n and that the discriminant D is so small that it does not contribute to the running time. As we saw in the example, when $p \equiv 3 \pmod 4$, $D = -4$, and is thus a small constant. However, unlike [22], we give a concrete estimate for the running time based on their assumptions.

We quickly recall the five steps:

Step 1 Turn I into an ideal of prime norm N
Step 2 Sample $\gamma \in O$ of reduced norm Nl^{e_0} giving us the proper ideal $O\gamma$ ()
Step 3 Find $[\mu] \in O/NO$ such that $\gamma\mu \equiv \beta \pmod{NO}$
Step 4 Lift to $\mu \in O$ such that $\operatorname{nrd}(\mu) = l^{e_1}$ for some e_1
Step 5 Return $I\overline{\gamma\mu}/N$, an ideal of reduced norm $l^{e_0+e_1}$ (Lemma 4.11)

Step 1 uses Algorithm 3 which runs in time bounded by $\log^4(p)$, where p is the ramification prime of $B = (-1, -p \mid \mathbb{Q})$.

Step 2 uses Algorithm 6 which runs in time $2h(R)\log(M)$ where M is Nl^{e_0}. We have already estimated e_0 to be roughly $\log_l(Np\Phi(p)|D|)$, so Nl^{e_0} is roughly $N^2p\Phi(p)|D|$ which is $O(p^2)$ under our heuristics. We also know that the class number $h(R)$ is bounded by its discriminant, that is

$$h(R) \leq \frac{1}{\pi}\sqrt{|D|}\ln(|D|)$$

so in particular we have

$$O(2h(R)\log(M)) = O(\log(p^2)) = O(\log(p))$$

Step 3 only consists of a relatively few number of steps which involves solving linear algebra over $M_2(\mathbb{Z}/N\mathbb{Z})$, implicitly depending on the factor N. For our run-time analysis we will omit this step.

Step 4 only consists of a few number of steps, but it involves calling Cornaccia's algorithm by trying to represent numbers close to l^{e_1} where l^{e_1} is expected to be $N^4p|D|$, which under our estimations become $O(p^2p) = O(p^3)$, thus giving us a running time bounded by $O(\log_{10}(p^{3/2})) = O(\log_{10}(p))$.

Step 5 is also just a straight forward computation.

Thus the algorithm runs in time $O(\log(p))$ making this suitable for virtually any prime p. In other words, finding an l-power isogeny (represented as a connecting ideal) given the endomorphism rings of elliptic curves is considered easy.

Remark. In practice this is no problem since the security assumption relies on computing the endomorphism rings, which we have shown to be a difficult task.

4. Computing isogenies from ideals

In this section we will look at how we can turn the ideals we have just created into isogenies. We begin with a short introduction to the problem before we describe a way to compute the isogenies if our ideal is of prime norm l, then we describe how we can turn an ideal of prime power norm l^e into e ideals of norm l, and finally we will put this together to create a complete algorithm.

At this point we have an ideal I of reduced norm l^e that connects O_1 to O_2, the orders in B corresponding to the endomorphism rings of the elliptic curves E_1 and E_2. We will describe an approach for computing those isogenies using algorithms based on Eisenträger, Hallgren, Lauter, Morrison and Petit's [11, Section 6].

We begin with the simplest observation of how to compute an isogeny. Assuming I is an ideal of $\mathrm{End}(E)$, we have

$$E[I] := \cap_{\alpha\in I}\ker(\iota_E(\alpha)) = \ker(\iota_E(\gamma_1)) \cap \ker(\iota_E(\gamma_2)) \cap \ker(\iota_E(\gamma_3)) \cap \ker(\iota_E(\gamma_4))$$

Where $\gamma_1, \ldots, \gamma_4$ is a $\mathbb{Z}$-basis of I. Since we assume $\mathrm{nrd}(I)$ and p are coprime, the isogeny will be separable and thus correspond to its kernel. After computing $E[I]$ we can thus use the explicit kernel to compute the isogeny using Vélu's formulas.

That being said, the kernel has size l^e. For the algorithm to be usable we require it to be polynomial in $\log(p)$ (or better). However, as explained in the previous section, e is expected

to be roughly $\log_l(p)$, so l^e is roughly p - not polynomial in $\log(p)$. To remedy this we use the classical approach of splitting l^e into e ideals of size l which gives us a solution of size el instead of l^e. This has some issues which we will discuss shortly, but first let us simply show how to compute the ideal corresponding to I given that we can represent $E[\mathrm{nrd}(I)]$.

4.1. Computing the isoegnies of l-normed ideals. We begin with the algorithm for turning an ideal I of reduced prime norm l into an isogeny of degree l. We assume we can represent $E[l]$, and that we have explicit ways of computing the action of $\alpha \in I$ on points of $E[l]$. We shall stick to our usual notation ι_E for the map between $(-1, -p \mid \mathbb{Q})$ and $\mathrm{End}(E) \otimes \mathbb{Q}$, but we stress that it is not necessary to have this map as long as we can evaluate $\iota_E(\alpha)$ on every $P \in E[l]$.

Algorithm 7: IdealToIsogeny(I, E)

Input: An ideal I of reduced norm l and elliptic curve E. Assuming I is given as explicit $\mathbb{Z}$-basis $\beta_1, \beta_2, \beta_3, \beta_4$ with $\beta_i \in O_L(I)$

Output: An isogeny $\phi_I : E \to E'$ corresponding to I

```
Compute the basis {P_1, P_2} of E[l] ;
Compute the map ι_E that can be evaluated on β_i ;
for s in {1, ..., 4} do
    Compute Q_st := ι_E(β_s)(P_t) for s ∈ {1, 2, 3, 4} and t ∈ {1, 2} ;
    Initialize L to empty list ;
    for (x, y) ∈ {(0, 1), (1, 0), (1, 1), (1, 2), ..., (0, l − 1)} do
        append (x, y) to L if [x]Q_s1 + [y]Q_s2 = 0 ;
    end
    if Size of L is exactly 1 then
        Extract x, y from L[0] ;
        Compute ψ : E → E/⟨xP_1 + yP_2⟩ ;
        return ψ
    end
end
return ⊥
```

Proposition 4.39. *Algorithm 7 is correct and runs in time roughly* $\mathrm{nrd}(I) + T$ *where* T *is the time required to find and evaluate* ι_E *on the generators of* I.

Proof. Correctness: First, assuming that an isogeny is returned, we show that it is the correct one. Since l is prime, $[x]P_1 + [y]P_2$ is an element of order l, and the isogeny ψ has degree l. Thus ψ can correspond to I or some other ideal of norm l. Since $[x]\iota_E(\beta_s)(P_1) + [y]\iota_E(\beta_s)(P_2) = 0$ we have $\iota_E(\beta_s)([x]P_1 + [y]P_2) = 0$, so $[x]P_1 + [y]P_2 \in \ker(\iota_E(\beta_s))$. Furthermore, since the size of L is 1, we have $\langle [x]P_1 + [y]P_2 \rangle = \ker(\iota_E(\beta_s)) \cap E[l]$. If any other generator β_i has a different $E[l]$

kernel subgroup (and not the entire $E[l]$ torsion), then $\ker(\iota_E\beta_1) \cap \ker(\iota_E\beta_s) \cap E[l] = \emptyset$ which cannot happen because we already know that $\emptyset \neq E[I] \subseteq E[l]$.

Second, assuming I is an ideal of prime norm l, we want to show that it returns an isogeny (not $\perp$). Since it corresponds to an isogeny of prime degree l the kernel $E[I]$ must correspond to a subgroup S of $E[l]$ of size l. Therefore the generators β_i has kernels $\ker(\iota_E\beta_i) \cap E[l] \in \{S, E[l]\}$. At least one of them must equal S, otherwise $E[I]$ would contain $E[l]$ and thus I would correspond to an isogeny of degree at least l^2. If $\ker(\iota_E\beta_i) \cap E[l] = E[l]$, then the size of L would be $l+1$ so the isogeny will not return. If $\ker(\iota_E\beta_i) \cap E[l] = S$, then $S = \langle [x]P_1 + [y]P_2 \rangle$ for exactly one of the pairs (x, y), so the size of L is 1 and ψ is returned.

Running time: The first step is a classical operation of finding elements of a subgroup of given order which is considered easy. A simple approach is to simply sample random elements of $E(\mathbb{F}_q)$, multiply them with $|E(\mathbb{F}_q)|/l$ and then check if they have order l by multiplying with $[l]$ (if l is not prime we must also verify the prime powers dividing l). Assuming l is prime, $E[l]$ contains l^2 elements, sampling two at random will give two independent elements except with probability $1/l$, so it is likely to succeed after very few attempts.

The second step is considered difficult for a random supersingular curve. Therefore we shall simply denote the cost of computing this map and the evaluation $\iota_E(\beta_i)$ as T.

Step 4 through 8 are simple evaluations and elliptic curve arithmetic, where step 7 is repeated $l+1$ times.

Step 11 can be performed using Vélus formulas which runs in time linear to the size of the kernel which is l.

Thus the algorithm runs in time l plus some factor T depending on how ι_E is evaluated. □

Remark. This algorithm can be modified to return an l^e isogeny if we during line 6 select (x, y) from the set $(0, 1) \cup S$ where

$$S = \{(1, y) \mid \gcd(y, l) = 1 \text{ and } y \in \{0, \ldots, l^e - 1\}$$

Thus we sample $l^e - l^{e-1}$ elements corresponding to all the l^e subgroups of $E[l^e]$.

Example 4.40. Let us perform the reverse of our isogeny to ideal computation from Example 4.15, where we ended up with the ideal

$$I_\phi = \left\langle \frac{1 + 9 \circ j}{2}, \frac{i + 9ij}{2}, 5j, 5ij \right\rangle$$

Which corresponded to the isogeny $\phi : E_0 \to E$ of degree 5 with kernel point $[125 : 82 : 1]$ going from $j(E_0) = 1728$ to $j(E) = 288$. Now we would like to use the above algorithm to verify that it works. Since we already know the endomorphism of O_0 by explicitly that ι_{E_0} sends the map i to $[x : y : z] \mapsto [-x : y\sqrt{-1} : z]$ and j to the Frobenius π, we can perform the algorithm easily.

First we compute a basis of $E[5]$ with ω the nontrivial generator of $\mathbb{F}_q$

$$P_1 = [27 : 430\omega + 233 : 1] \qquad P_2 = [320\omega + 129 : 356\omega + 120 : 1]$$

We start with the first generator of I_ϕ and compute its image of the generators. The explicit endomorphism corresponding to β_1 is $([1] + [9] \circ \pi)/[2]$ which gives us the points

$$Q_{11} = [27 : 430\omega + 233 : 1] \qquad Q_{12} = [314 : 273\omega + 249 : 1]$$

Then we see which tuples (x, y) give us $[x]Q_{11} + [y]Q_{21} = 0$, and see that it is exactly $(1, 3)$. Since the size of L is 1 we compute the isogeny

$$\psi : E_0 \to E_0/\langle P_1 + 3P_2 \rangle = E_0/\langle [125 : 357 : 1] \rangle = E$$

We see immediately that the kernel point $P_1 + 3P_2$ correspond to our initial isogeny as it has the same coordinates except $-y$ instead of y.

4.2. Getting to the l-normed ideals. As our ideal I connecting E to E' is of reduced norm l^e, with $e \approx 7/2 \log_l(p)$, Algorithm 7 runs in time polynomial in $l^e \approx p^{7/2}$, which would be quite intensive for large values of p. Furthermore $E[l^e]$ will possibly be in some large extension field of E when e is this large. Thus, instead of computing the large degree isogeny directly we make the classical approach of splitting I into e l-normed ideals corresponding to the l isogenies connecting E and E'.

Let us explain this more clearly. Initially we have I connecting E to E', thus implicitly we have an isogeny $\phi : E \to E'$ of degree l^e. This can be decomposed into e isogenies of degree l

$$\phi = \psi_e \circ \cdots \circ \psi_1$$

This allows us to create the partially composed isogenies $\phi_k := \psi_k \circ \cdots \circ \psi_1$ as illustrated in Figure 19.

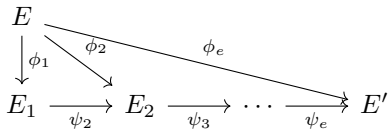

FIGURE 19. Illustration of isogeny decomposition

Just like the kernels of ϕ_k form a filtration $(\ker(\phi_i) \subseteq \ker(\phi_{i+1}))$ we can make a similar filtration for the ideal I:

$$I = I_e \subseteq I_{e-1} \subseteq \cdots \subseteq I_2 \subseteq I_1 \subseteq I_0 = \mathrm{End}(E)$$

we will show how this is done shortly, but for now we will just assume that it can be done and that each ideal I_k has reduced norm l^k. Then at each intermediate step we would like to construct the ideal J_k being equivalent to the ψ_k isogenies in Figure 19.

From Figure 20 it might seem clear that our ideal J should be of the form $J = I_k^{-1} I_{k+1}$. The following proposition gives us the desired result as we assume that $I_{k+1} \subseteq I_k$.

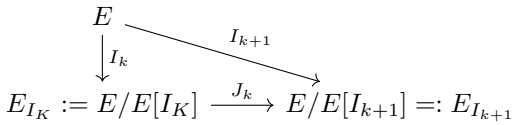

FIGURE 20. Illustration of ideal decomposition

PROPOSITION 4.41. *Let $I' \subseteq I$ be two left $\mathrm{End}(E)$-ideals with reduced norms coprime to p. Then there exists a separable isogeny $\psi : E_I \to E_{I'}$ such that $\phi_I = \psi \circ \phi_{I'}$ and a left $\mathrm{End}(E_I)$-ideal $\tilde{J}$ satisfying $E_I[\tilde{J}] = \ker(\psi)$ such that $J := \iota(\tilde{J}) = I^{-1}I'$ where $\iota : \mathrm{End}(E_I) \to \mathrm{End}(E) \otimes \mathbb{Q}$ is the injective map defined earlier.*

PROOF. Proof based on [11, Proposition 10, 11 and 12]

By Corollary 2.67 and the fact that $E[I] \subseteq E[I']$ we have that there exists a unique isogeny $\psi : E_I \to E_{I'}$ such that $\phi_{I'} = \psi \circ \phi_I$. Furthermore, since $\phi_{I'}$ is separable, so is ψ.

Let $\tilde{J}$ be the left $\mathrm{End}(E_I)$-ideal corresponding to ψ. What remains to be shown is that $\iota(\tilde{J}) = I^{-1}I'$. Let $x \in I^{-1}I'$. Since $I^{-1} = \frac{1}{\deg \phi_I}\overline{I}$ we have $x = \frac{1}{\deg \phi_I}\widehat{\alpha}\beta$ for $\alpha' \in I$ and $\beta' \in I'$. Using the push-forwards ϕ_I^* and $\phi_{I'}^*$ and their isomoprhism result we can find $\alpha \in \mathrm{Hom}(E_I, E)$ and $\beta \in \mathrm{Hom}(E_{I'}, E)$ satisfying $\alpha = \alpha'\phi_I$ and $\beta = \beta'\phi_{I'}$.

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

Furthermore, since $\widehat{\alpha}\beta \in \mathrm{Hom}(I', I)$, using the bijective map $g : \mathrm{Hom}(E_{I'}, E_I) \to I^{-1}I'$ we see that

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

giving us

$$x = (\iota \circ \psi^* \circ g^{-1})(x)$$

or rather that $g = \iota \circ \psi^*$. Which tells us that $I^{-1}I$ is equal to $\iota(\tilde{J})$ which we simply call J. □

Now we can make the ideal decomposition diagram equivalent of the earlier isogeny decomposition:

What remains to be shown is that we can decompose the ideal I into the filtration described earlier. This is the main theorem of [11, Section 6], which creates this filtration and collects the previous results about decomposing the ideal I into a sequence of l normed ideals J_k.

THEOREM 4.42. *Let I be a left $\mathrm{End}(E)$-ideal that satisfies $\mathrm{nrd}(I) = l^e$ with $l \neq p$ a prime and $I \not\subset \mathrm{End}(E)l^k$ for any integer $k > 0$. Then there exists a filtration*

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

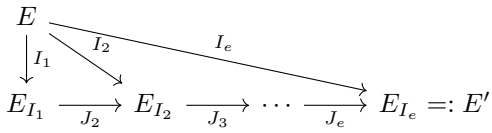

FIGURE 21. Illustration of complete isogeny ideal decomposition

and a chain of isogenies

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

Such that if we set $\phi_k : E \to E/E[I_K]$ *we get* $\phi_{k+1} = \psi_k\phi_k$ *and for* $k = 0, \ldots, e-1$ *the map* $\psi_{k+1} : E_k \to E_{k+1}$ *has degree* l *and kernel ideal isomorphic to* $J_{k+1} := I_k^{-1}I_{k+1}$.

PROOF. Proof based on [11, Theorem 11].

We first show that $I_k := I + \mathrm{End}(E)l^k$ has reduced norm l^k whenever I is not contained in any $\mathrm{End}(E)l^m$. Since I has reduced norm l^k and is not contained in any $\mathrm{End}(E)l^m$ then $E[I]$ is cyclic, ie generated by a single point P. Suppose that $E[I]$ is non-cyclic. Then $E[I] \cong \oplus_{i=1}^{n}\mathbb{Z}/k_i\mathbb{Z}$ by the structure theorem of abelian groups, such that $k_i \mid k_{i+1}$. Since $E[I]$ is non-cyclic $n > 1$ and hence we have two linearly independent elements of order k_1, so $E[k_1] \subseteq E[I]$. Thus $\mathrm{End}(E)k_1 \subseteq I$ and we have in particular that $k_1 \mid l^k$ so there exists an m such that $I \subseteq \mathrm{End}(E)l^m$.

So if we assume that no m exists such that $I \subseteq \mathrm{End}(E)l^m$, then $E[I]$ is cyclic and in particular generated by a single point P. We want to show that $I_k := I + \mathrm{End}(E)l^k$ is of norm l^k. Since P has order l^e and $E[I_k]$ is generated by $E[I] \cap E[\mathrm{End}(E)l^k] = E[I] \cap E[l^k]$ we get that $E[I_k] = \langle [l^e/l^k]P \rangle$ is a group of order l^k. Thus the isogeny corresponding to $E[I_k]$ has degree l^k and thus the ideal I_k must be of reduced norm l^k.

Clearly $I_k \subsetneq I_{k-1}$ by the description of the kernel ideal, so our filtration is also immediate.

From Proposition 4.41 we get that $J_k := I_{k-1}^{-1}I_k$ is a valid candidate for the horizontal isogenies. It will certainly satisfy Figure 21. What remains to show is that J_k has reduced norm l. But this follows from simple computations:

$$\mathrm{nrd}(J_k) = \mathrm{nrd}(I_{k-1}^{-1}I_k) = \mathrm{nrd}\left(\frac{1}{\mathrm{nrd}(I_{k-1})}\overline{I_{k-1}}I_k\right) = \left(\frac{1}{l^{k-1}}\right)^2 l^{k-1}l^k = l$$

□

4.3. The complete ideal to isogeny algorithm. Putting the previous two sections together we can make an algorithm for computing the isogeny corresponding to the ideal I.

This is the slightly modified version of [11, Algorithm 9].

THEOREM 4.43. *Excluding the time needed to find basis of* B*, computing the maximal order of* E *and furthermore assuming that* l *is of size polynomial in* $\log p$ *and that Algorithm 7 runs in time polynomial in* $\mathrm{nrd}(J_k)$*. Then Algorithm 8 runs in time polynomial in* $\log p$.

Proof: See [11, Theorem 10]

Algorithm 8: IdealToIsogenyChain(I, E, E')

Input: An ideal I of reduced norm l^e connecting E to E'
Output: An chain of l-isogenies composing to ϕ_I

Compute the basis $\langle 1, i, j, ij\rangle$ for B ;
Compute the maximal order of E in B - giving us the basis $\langle \beta_1, \beta_2, \beta_3, \beta_4\rangle$ of $O \subset B$;
for $0 \leq k \leq e$ **do**
    Compute $I_k := I + Ol^k$ and its right order $O_R(I_k)$;
    Compute a $\mathbb{Z}$-basis of the $O_R(I_k)$-ideal $J_{k+1} := I_k^{-1} I_{k+1}$;
end
Set $E_0 := E$;
for $0 \leq k \leq e-1$ **do**
    $\phi_{k+1} \leftarrow$ IdealToIsogeny(J_k, E_k) ;
    $E_{k+1} = \text{Image}(\phi_k)$
end
return $(\phi_1, \ldots, \phi_e)$

In other words, if we are able to compute the endomorphism ring of an elliptic curve, then we can construct the connecting ideal I using the KLPT algorithm and then compute the corresponding chain of isogenies $\phi_e \cdots \phi_1$ that provide an representable isogeny.

Example 4.44. We continue our example. At this point we have an ideal I of reduced norm 7^{13}, we would like to compute the isogeny corresponding to this ideal. Unfortunately for us, $7^2 \nmid |E(\mathbb{F}_{p^2})|$ so wee need to look at an extension field to find our 7-torsion points. It turns out that $|E(\mathbb{F}_{p^6})| = 2^6 * 5^2 * 7^2 * 11^2 * 13^2 * 2113^2$ which is exactly what we need.

Our ideal I has generators

$$\begin{aligned}
\beta_1 &= 1/10 + 180794141123/5j + 64278645487/2ij \\
\beta_2 &= 1/10i + 74134226523/2j + 180794141123/5ij \\
\beta_3 &= 69206436005j \\
\beta_4 &= 69206436005ij
\end{aligned}$$

When performing the algorithm we get the J_k ideals generated by β_i as depicted in Figure 22.

Finally we would like to turn this sequence of ideals into isogenies like we did in the previous section. This can be easily accomplished using the first algorithm if we know how to evaluate the generators of J_k on the 7-torsion of the various elliptic curves. However we are only able to compute the endomorphism ring of the starting curve E_0. This allows us to construct J_1, which corresponds to the isogeny ϕ with kernel

$$(77\omega^5 + 4\omega^4 + 106\omega^3 + 417\omega^2 + 6\omega + 414, 185\omega^5 + 431\omega^4 + 437\omega^3 + 160\omega^2 + 435\omega + 122)$$

Corresponding to an isogeny going from E_0 to $E_1 : y^2 = x^3 + Ax + B$ with

$$A = 132\omega^5 + 303\omega^4 + 135\omega^3 + 271\omega^2 + 23\omega + 59$$
$$B = 243\omega^5 + 29\omega^4 + 39\omega^3 + 10\omega^2 + 192\omega + 141$$

The next step would either require computing the map ι_{E_1} or finding a larger extension field where we can find the 7^2 torsion. If we have the generators $P, Q \in E_1[7]$ we can use the earlier composition by Washington to evaluate β_i on P through $\frac{\phi \circ \iota_{E_0}(\beta_i)\widehat{\phi}}{[7]}(P)$. Noticing that if we start dividing P and Q with 7 we get points of order 7^2 which will survive the $\phi \circ \widehat{\phi}$ composition. Unfortunately our example requires very large extension fields to find the 7^2 torsion, and even longer to find the entire 7^{13} torsion required to complete the entire isogeny composition. Our attempts at finding a more suitable field for such computations did not hold through, so we unfortunately end this section without a complete example.

This illustrates how the difficulty of computing the endomorphism ring underlies the hardness assumption on the isogeny problem of supersingular elliptic curves. Given explicit maps ι_E it would have been easy to construct an isogeny of a prime degree l. Similarly if the entire l^e-torsion was $\mathbb{F}_{p^n}$-rational for some small n it would be possible to compute it. Otherwise there is no clear way to retrieve the isogeny from its ideal.

Ideal J_1	
β_1	$\frac{1+11j+4ij}{2}$
β_2	$\frac{1i+10j+11ij}{2}$
β_3	$\frac{14j}{2}$
β_4	$\frac{14ij}{2}$

Ideal J_2	
β_1	$\frac{7+77j+518ij}{14}$
β_2	$\frac{1i+66j+165ij}{14}$
β_3	$\frac{98j+98ij}{14}$
β_4	$\frac{686ij}{14}$

Ideal J_3	
β_1	$\frac{49+539j+8428ij}{98}$
β_2	$\frac{1i+458j+14277ij}{98}$
β_3	$\frac{686j+19894ij}{98}$
β_4	$\frac{33614ij}{98}$

Ideal J_4	
β_1	$\frac{343+3773j+1000188ij}{686}$
β_2	$\frac{1i+1830j+1062485ij}{686}$
β_3	$\frac{4802j+374556ij}{686}$
β_4	$\frac{1647086ij}{686}$

Ideal J_5	
β_1	$\frac{2401+26411j+64649326ij}{4802}$
β_2	$\frac{1i+16236j+3833239ij}{4802}$
β_3	$\frac{33614j+60269902ij}{4802}$
β_4	$\frac{80707214ij}{4802}$

Ideal J_6	
β_1	$\frac{16807+184877j+3277297772ij}{33614}$
β_2	$\frac{1i+16236j+3554950655ij}{33614}$
β_3	$\frac{235298j+3811592302ij}{33614}$
β_4	$\frac{3954653486ij}{33614}$

Ideal J_7	
β_1	$\frac{117649+1294139j+1613539564l4ij}{235298}$
β_2	$\frac{1i+1192726j+22612912165ij}{235298}$
β_3	$\frac{1647086j+820462994918ij}{235298}$
β_4	$\frac{193778020814ij}{235298}$

Ideal J_8	
β_1	$\frac{823543+9058973j+384236998294ij}{1647086}$
β_2	$\frac{1i+9428156j+4889738865477ij}{1647086}$
β_3	$\frac{11529602j+574324064426ij}{1647086}$
β_4	$\frac{949512301988 6ij}{1647086}$

Ideal J_9	
β_1	$\frac{5764801+63412811j+3592258956000068ij}{11529602}$
β_2	$\frac{1i+9428156j+4226751517404 61ij}{11529602}$
β_3	$\frac{80707214j+20341785186858 8ij}{11529602}$
β_4	$\frac{465261027974414ij}{11529602}$

Ideal J_{10}	
β_1	$\frac{40353607+443889677j+122850628566631 70ij}{80707214}$
β_2	$\frac{1i+4129642 26j+7953418802725197ij}{80707214}$
β_3	$\frac{564950498j+1770806094218 4606ij}{80707214}$
β_4	$\frac{22797790370746286ij}{80707214}$

Ideal J_{11}	
β_1	$\frac{282475249+3107227739j+1043502635567986202ij}{564950498}$
β_2	$\frac{1i+3237716716j+8488208057482 75665ij}{564950498}$
β_3	$\frac{3954653486j+9218790895714 12252ij}{564950498}$
β_4	$\frac{1117091728166568014ij}{564950498}$

Ideal J_{12}	
β_1	$\frac{1977326743+21750594173j+3858308683763980 7806ij}{3954653486}$
β_2	$\frac{1i+7192370202j+4004883351652823 945ij}{3954653486}$
β_3	$\frac{27682574402j+1427279572416586 1862ij}{3954653486}$
β_4	$\frac{5473749468016183268 6ij}{3954653486}$

Ideal J_{13}	
β_1	$\frac{13841287201+152254159211j+2700816078634786 54642ij}{27682574402}$
β_2	$\frac{1i+117922667810j+2524283326855598742263ij}{27682574402}$
β_3	$\frac{193778020814j+9990957006916 1033034ij}{27682574402}$
β_4	$\frac{2682137239327929801614ij}{27682574402}$

FIGURE 22. Table of J_k ideals with their generators $J_k = \langle \beta_1, \ldots, \beta_4 \rangle$

CHAPTER 5

Applications

Now that we have discussed how to construct isogenies between arbitrary elliptic curves, this chapter will explain two applications. We will look into the supersingular cryptosystem SIDH [14] and see how knowing the endomorphism ring allows us to break the system completely. Next we will look into how the KLPT algorithm can be turned into a post-quantum signature scheme as is done in SQISign [15].

We will sketch the cryptosystems by going into just enough details to get a feel for how our knowledge of isogeny construction can be applied. This chapter is by no means meant to introduce the cryptosystems nor be precise with their inner workings.

1. SIDH

In this section we will describe the Supseringular Isogeny Diffie-Hellman key exchange (SIDH) [14] proposed by De Feo, Jao, and Plût in 2011. We will show how it can be broken if one is able to compute the endomorphism ring of either E_A or E_B.

1.1. Description of the cryptosystem. We choose the prime $p = l_A^{e_A} l_B^{e_B} f - 1$, and have our two participants Alice and Bob. The protocol is initialized by creating the starting curve $E : y^2 = x^3 + x$ over $\mathbb{F}_{p^2}$ and finding kernel generators $\langle P_A, Q_A \rangle = E[l_A^{e_A}]$ and $\langle P_B, Q_B \rangle = E[l_B^{e_B}]$. The protocol is depicted in Figure 23. Alice constructs a secret isogeny ϕ_A by choosing random integers n_A and m_A which gives rise to its kernel $\langle [n_A]P_A + [m_A]Q_A \rangle$. Bob does the same by sampling n_B and m_B, creating $\langle [n_B]P_B + [m_B]Q_B \rangle$ and constructing its isogeny ϕ_B. Next Alice transfers $E_A := \text{Image}(\phi_A)$, $P'_B := \phi_A(P_B)$ and $Q'_B := \phi_A(Q_B)$ to Bob. Similarly Bob replies with $E_B := \text{Image}(\phi_B)$, $P'_A := \phi_B(P_A)$ and $Q'_A := \phi_B(Q_A)$. Now they are both able to compute their shared secret target curve E_{AB} as $E_B/\langle [n_A]P'_A + [m_A]P'_B \rangle$ or equivalently $E_A/\langle [n_B]P'_B + [m_B]Q'_B \rangle$. Simply computing the j-invariant of E_{AB} gives Alice and Bob a shared secret that can be used for other cryptographic purposes.

The hardness assumption of SIDH is essentially the following problem.

Problem. As in the setting above, given E, E_A, $l_A^{e_A}$, P_A, Q_A, P_B, Q_B, $\phi_A(P_B)$, $\phi_A(Q_B)$, $\phi_B(P_A)$ and $\phi_B(Q_A)$. Find an isogeny $\phi : E \to E_A$ of degree $l_A^{e_A}$.

That is, if one can solve this problem, then one can write the kernel generator of ϕ with respect to P_A, Q_A and create the isogeny ψ_A which gives rise to the secret curve E_{AB}. Finding E_{AB} breaks the system as it allows an adversary to obtain the shared secret.

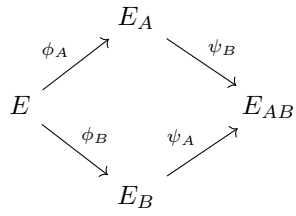

FIGURE 23. SIDH Cryptosystem

This problem seems reasonable, but let us explain the necessity for ϕ to have degree $l_A^{e_A}$, and not just be any isogeny connecting E to E_A. This is based on [19, Section 4.1].

We first recall that the honest participant Alice creates her isogeny ϕ_A from the generator point $G_A := [n_A]P_A + [m_a]Q_A$. To construct the secret key E_{AB}, she takes the public curve E_B and the kernel points $\phi_B(P_A), \phi_B(P_B)$ and creates the new equivalent kernel generator point $\phi_B(G_A) := [n_A]\phi_B(P_A) + [m_A]\phi_B(Q_A)$. This gives rise to an isogeny $\psi_A : E_B \to E_{AB}$.

For the adversary, attempting to recover the secret E_{AB} he can approach the problem like Alice by attempting to create her secret $\phi_A : E \to E_A$. Let us assume that he is able to construct an isogeny $\phi : E \to E_A$ of l_A-power norm. Then he is only able to map the equivalent isogeny to E_B if the kernel generator G' is in $\langle P_A, Q_A \rangle$. Otherwise the image points $\phi_B(P_A)$ and $\phi_B(Q_A)$ does not provide enough information to recreate the isogeny at the curve E_B.

Furthermore, we argue that it is unlikely that another isogeny $\phi : E \to E_A$ of l_A-power norm has a generator point inside $E[l_A^{e_A}] = \langle P_A, Q_A \rangle$. There are only $l_A^{e_A} + 1$ subgroups of order dividing $l_A^{e_A}$ (thus laying in $\langle P_A, Q_A \rangle$). let S be the set of all elliptic curves obtained from taking the image curve of the isogenies corresponding to these subgroups. Since there are roughly $p/12$ supersingular curves in total, the probability that a random curve is in this set is roughly $1/l_B^{E_B}$. Thus if one has found an isogeny of l_A-power degree it will most likely be larger than $l_A^{e_A}$.

1.2. Breaking the security assumption. We follow the paper by Galbraith, Petit, Shani and Ti [19, Section 4] on how to break the security of the SIDH cryptosystem. In order to break the system they assume that the adversary is in possession of the endomorphism ring of the starting curve E_A (or equivalently E_B). With the knowledge of $\mathrm{End}(E)$ and $\mathrm{End}(E_A)$ they provide an algorithm which is highly likely to produce the correct isogeny connecting E and E_A.

It is tempting to just apply the previously described KLPT algorithm to provide us with an isogeny of degree l_A^e for some e. However this algorithm produces an isogeny ϕ with $\deg \phi \approx l_A^{7/2 \log_{l_a}(p)} = p^{7/2}$ while the SIDH cryptosystem produces ϕ with $\deg(\phi) \approx p^{1/2}$.

By our previous discussion on supersingular elliptic curves we know that I is an $\mathrm{End}(E), \mathrm{End}(E_A)$-ideal corresponding to an isogeny $\phi : \mathrm{End}(E) \to \mathrm{End}(E_A)$. Computing a Minkowski reduced basis is possible, and easy, since the ideal has dimension 4 (or smaller). Once the Minkowski reduced basis is found, one can easily compute the smallest non-zero element α of I. Assuming this has $\mathrm{nrd}(\alpha) = \mathrm{nrd}(I) l_A^{e_A}$ we can compute I' by scaling I as in step 7 which gives us I' with

Algorithm 9: Computing the smallest degree isogeny of given degree [19, Algorithm 2]

Input: l_A, e_A, E, E_A, $\mathrm{End}(E)$, $\mathrm{End}(E_A)$ where $\phi : E \to E_A$ exists with $\deg(\phi) = l_A^{e_A}$
Output: An isogeny $\phi : E \to E_A$ of degree $l_A^{e_A}$ or Failure

Compute $I := \mathrm{End}(E) \cap \mathrm{End}(E_A)$;
Compute a Minkowski reduced basis of I ;
Let α be the non-zero element of I of minimal norm ;
if $\mathrm{nrd}(\alpha) \neq \mathrm{nrd}(I) l_A^{e_A}$ **then**
  return *Failure*
end
Compute $I' := I\overline{\alpha}/\mathrm{nrd}(I)$;
Compute isogeny ϕ_A corresponding to I' ;
return ϕ_A

$\mathrm{nrd}(I') = l_A^{e_A}$. Since we know the endomorphism rings of $\mathrm{End}(E)$ we can compute the ideal corresponding to ϕ_A by splitting I' into e_A ideals of reduced norm l_A and compute the chain of isogenies one by one. Therefore Algorithm 9 runs in polynomial time and succeeds assuming that α is this correct element.

Remark. The issue of computing the isogeny from the ideal as described in Example 4.44 is not a problem for us in this case as we already know that we can represent the entire $E[l_A^{e_A}]$ torsion subgroup by the way the cryptosystem is constructed.

The existence of α. Firstly let us remark that there exists an $\alpha \in I$ of reduced norm $\mathrm{nrd}(I) l_A^{e_A}$ by the construction of the SIDH cryptosystem. This follows since we know there exists an $\mathrm{End}(E), \mathrm{End}(E_A)$-ideal J of reduced norm $l_A^{e_A}$ which is equivalent to I. Furthermore, since $J \subseteq \mathrm{End}(E)$ and $J \subseteq \mathrm{End}(E_A)$ we have

$$J' := J \cap I = J \cap \mathrm{End}(E) \cap \mathrm{End}(E_A) \neq \emptyset$$

This intersection is still an $\mathrm{End}(E), \mathrm{End}(E_A)$-ideal and its reduced norm is necessarily $\mathrm{nrd}(I)\,\mathrm{nrd}(J) = \mathrm{nrd}(I) l_A^{e_A}$. Thus taking any $\alpha \in I \cap J$ of smallest reduced norm would be sufficient.

To show that α is the smallest element of I, Galbraith et al. argues that two random curves are unlikely to be connected by isogenies of degrees significantly smaller than $\sqrt{p}$. Using analytic number theory they show that there are at most

$$\frac{15}{2\pi^2} D^2$$

supersingular elliptic curves connected to E with an isogeny of degree at most D. setting $D = \sqrt{p}$ we get at most $\frac{15}{2\pi^2} p \approx 0.75p$. Since there are roughly $p/12 \approx 0.08p$ supersingular curves we indeed barely have the possibility to reach every curve when $D = \sqrt{p}$. Therefore, if the degree is significantly smaller than $\sqrt{p}$, it is unlikely that the specific curve E_A is among this set.

However, for the parameters of the cryptosystem, one could choose $p = 2^{e_A}3^{e_B} - 1$ with $2^{e_A} \approx 3^{e_B}$. Thus creating the entire 2^{e_A} isogeny this estimate would not hold as $\sqrt{p} \approx 2^{e_A} \approx 3^{e_B}$. Instead one can perform a small number of degree 2 isogenies from E (the article suggest 10), and then perform the algorithm on each candidate. If we choose to perform 10 isogenies of degree 2, one would need to test $2^{10} = 1024$ elliptic curves. Then one would have at most $\frac{15}{2\pi^2}2^{(e_A-10)2}$ target curves and $p/12 \approx 2^{e_A}3^{e_B}/12$ curves in total. Assuming that E_A is a random curve, the likelihood that it is among this set would be bounded by

$$\frac{2^{(e_A-10)2}15 * 12}{2\pi^2 2^{e_A}3^{e_B}} \approx 9\frac{2^{2e_A-20}}{2^{e_A}3^{e_B}} \approx 9\frac{2^{2e_A}2^{-20}}{2^{e_A}2^{e_A}} \approx \frac{9}{2^{20}} < 2^{-16}$$

Giving us a really large success probability.

2. SQISign

In 2020, De Feo, Kohel, Leroux, Petit and Wesolowski proposed a post-quantum signature scheme named SQISign [15]. It generalizes the KLPT algorithm to create a connecting ideal of smooth norm (that is $\text{nrd}(I) = \prod_j p_j^{e_j}$ where p_j are small primes) and uses this to prove knowledge of the endomorphism ring of some elliptic curve. Since it is considered difficult to compute the endomorphism ring of random supersingular elliptic curves (as we have seen) this then gives rise to a signature scheme. The interesting part of the scheme is that it, like other isogeny-based systems, provide really short keys which is unheard of in the post-quantum signature world. For NIST Level 1 security[1], the secret key is 16 bytes, the public key is 64 bytes and the signature is 204 bytes for one concrete instantiation of the protocol [15, Section 8.7].

The signature scheme requires a function, $\Phi_{D_c}(E, s)$ that maps integers $s \in [1, f(D_c)]$ to a non-backtracking isogeny of degree D_c from E, and a hash function $H : \{0,1\}^* \to [1, f(D_c)]$ which is cryptographically secure. The function $f(D_c)$ is the map $\prod_j l_j^{e_j} \mapsto \prod_j l_j^{e_j-1}(l_j + 1)$ where the left product is the prime factorization of D_c and $l_j \neq l_i$ whenever $i \neq j$.

Setup: Pick prime p and supersingular elliptic curve E_0 with known endomorphism ring O_0. Select odd smooth number D_c and let $D = 2^e$ be a power of 2 where e is greater than the diameter of the 2-isogeny graph of supersingular elliptic curves.

Keygen: Pick random isogeny $\tau : E_0 \to E_A$. The public key is E_A and the secret key is τ.

Sign: Pick a random (secret) isogeny $\psi : E_0 \to E_1$, let $s = H(j(E_1), m)$ be the hash of the j-invariant of the curve E_1 along with the message we wish to sign. Then map this hash value to an isogeny $\phi = \Phi_{D_c}(E, s) : E_1 \to E_2$. Using knowledge of τ, construct an isogeny $\sigma : E_A \to E_2$ of degree D, such that $\widehat{\phi} \circ \sigma$ is cyclic. The signature is (E_1, σ)

Verify: Create $\phi = \Phi_{D_c}(E, H(E_1, m)) : E_1 \to E_2$. Verify that σ is an isogeny from E_A to E_2 of degree D and that $\widehat{\phi} \circ \sigma$ is cyclic. If so, then the signature is valid.

This is illustrated in Figure 24

[1] security equivalent to a 128-bit block cipher

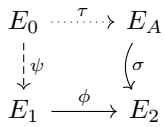

FIGURE 24. SQISign protocol

The signer knows the endomorphism ring of E_A through the isogeny τ as well as the endomorphism ring of E_2 through $\phi \circ \psi$. He is able to construct τ of degree D in a manner similar to the KLPT algorithm which we will describe next. Furthermore since τ and ψ are kept secret to the verifier, computing $\mathrm{End}(E_2)$ and $\mathrm{End}(E_A)$ is considered difficult so he is unable to compute σ like the signer. Finally the choice of sampling σ such that $\widehat{\phi} \circ \sigma$ is cyclic ensure that the signer does not reveal τ, thus providing the zero-knowledge property. We do not have the required background material to go through the proof of this, but the SQISign article [15] explains this well.

2.1. Modified KLPT Algorithm. We will describe the modified KLPT algorithm with the notation taken from [15], rather than using our own. Next we will describe it line by line to ensure we understand what is going on. This approach allows us to read the paper more easily later on.

In this scenario we already know the isogeny τ and we have computed its ideal I_τ. Our prime l is 2, and we denote the maximal orders of E_0, E_A, E_2 by O_0, O and O_2 respectively. We call the algorithm with a connecting O, O_2-ideal I for example the intersection of O and O_2. We depict the various ideals in Figure 25.

Algorithm 10: SigningKLPT(I, I_τ)

Input: I, a left O-ideal, I_τ a O_0, O-connecting ideal of norm N_τ
Output: $J \sim I$ of norm l^e with fixed e

1 Compute $K =$**EquivalentRandomEichlerIdeal**(I, N_τ) ;
2 Compute $K' = [I_\tau]^* K$ and set $L =$**EquivalentPrimeIdeal**(K') where $L = \chi_{K'}(\delta)$ for $\delta \in K'$ with $N = \mathrm{nrd}(L)$. Then we set $e_0 = e_0(N)$ and $e_1 = e - e_0$;
3 Compute $\gamma =$**RepresentInteger**$_{O_0}(Nl^{e_0})$;
4 Compute $(C_0 : D_0) =$ **IdealModConstraint**(L, γ) ;
5 Compute $(C_1 : D_1) =$ **EichlerModConstraint**$(\mathbb{Z} + I_\tau, \gamma, \delta)$;
6 Compute $C =$**CRT**$_{N,N_\tau}(C_0, C_1)$ and $D =$**CRT**$_{N,N_\tau}(D_0, D_1)$. If $l^e p(C^2 + D^2)$ is not a quadratic residue, repeat from step 3 ;
7 Compute $\mu =$**StrongApproximation**$({}_l(NN_\tau, C, D))$ of norm l^{e_1} ;
8 Set $\beta = \gamma\mu$;
9 **return** $J := [I_\tau]_* \chi_L(\beta)$

We first note that the function $\chi_I(\alpha)$ simply returns $I\overline{\alpha}/\operatorname{nrd}(I)$, that is it constructs the equivalent ideal of a different norm.

In step 1 we create a random element K which is equivalent to I and has norm relatively prime to N_τ. It is selected this way to ensure that we can pull it back to O_0 without trouble. Furthermore, choosing a random K affects the distribution of L later on, which in turn underlies the entire security of this scheme. By sampling K at random, ensures that L is uniformly distributed, which in turn allows us to reveal σ corresponding to J without revealing the isogeny τ.

The method works by finding an element $\omega \in O$ such that N_τ is inert in $\mathbb{Z}[\omega]$, this allows us to use a result similar to Lemma 4.28 to ensure that the ideals in $O/N_\tau O$ are in the same orbit. Thus sampling a random element of $\mathbb{P}^1(\mathbb{Z}/N_\tau\mathbb{Z})$ gives us the possibility to map the ideal I to a random equivalent ideal. To ensure that the output ideal has norm relatively prime to N_τ we sample random elements γ until one is found which satisfies $\gcd(\operatorname{nrd}(\gamma)/\operatorname{nrd}(I), N_\tau) = 1$. We lift the random value of $\mathbb{P}^1(\mathbb{Z}/N_\tau\mathbb{Z})$ to μ and return $\chi_I(\mu\gamma)$, which will be a random ideal equivalent to I with norm relatively prime to N_τ.

In step 2 we use the notion $[I_\tau]^*$ which is equivalent of the pull-back of I_τ acting on whatever ideal comes next. This is really a notion from the isogeny-maps, but it is extended to ideals by mapping them to isogenies, performing the pull-back (or push-forward) there and then mapping back to ideals afterwards. For isogenies, the pullback is just the push-forward of the dual isogeny. When the norm of I and J are relatively prime we can create the maps explicitly as

$$[I]^*J = IJ + \operatorname{nrd}(J)O$$

where I is a left O-ideal and I, J are compatible. In other words, K' is an O_0, O_2-connecting ideal. We set L to be an ideal equivalent to K' but of prime norm. This is just Algorithm 3, so we know how to do this, with the added functionality that it also returns δ for use later. The definition of e_0 as $e_0(N)$ is just to indicate that we select the integer e_0 depending on the value of N, more details of this is given in [15, Section 6.4].

In step 3 we do the same as in KLPT where we simply find a suitable element $\gamma \in O_0$ which has reduced norm Nl^{e_0}.

In step 4 we do the same as in KLPT where we find values $(C_0 : D_0)$ which we previously described as elements in Rj that would solve $L \equiv O_0\gamma\mu \pmod{NO}$ for $\mu = (C_0 + D_0\omega)j \in L$.

In step 5, we do the same as in step 4, except with the added constraint that $\mu \in O_0 \cap O_2$ - the Eichler order of O_0 and O_2. More concretely it searches for $\mu = (C_1 + \omega D_1)j$ such that $\gamma\mu\delta \in \mathbb{Z} + I_\tau$.

In step 6 we use the Chinese Remainder Theorem (CRT) to find a solution to the congruence $C \equiv C_1 \pmod{N}$ and $C \equiv C_2 \pmod{N_\tau}$ and similar for D. Ensuring that $l^e p(C^2 + D^2)$ is a quadratic residue allows one to use the strong approximation theorem in the next step.

In step 7 we use the so called StrongApproximation algorithm, which is the equivalent of our method for lifting $[\mu] = (C + \omega D)j$ to μ of a suitable norm. Notice however that there is a slight distinction from the algorithm we described. In our case the integer N was prime which is the one we used for modular arithmetic. Furthermore we were allowed to choose e ourselves when

lifting to an element of norm l^e. This method performs the arithmetic modulo NN_τ for distinct primes N and N_τ, furthermore the requirements of e is fixed to be e_1. This has a probability of failing, forcing us to go back to step 3 and choose another γ.

In step 8 we simply compute our β as in the KLPT algorithm, scale the ideal L by β, giving us an ideal of reduced norm l^e, then we use the push-forward to move this ideal which originally is an O_0, O_2-connecting ideal to an O, O_2-connecting ideal.

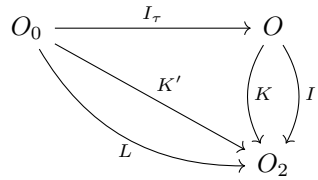

FIGURE 25. An illustration of where the various ideals are located with respect to the maximal orders

Or in shorter, less precise terminology, we take the combined ideal I, turn it into an equivalent ideal $K : O \to O_2$, then pull it back to the ideal $K' : O_1 \to O_2$. There we create a random but equivalent ideal $L : O_1 \to O_2$. This ideal is then turned to an equivalent ideal of reduced norm D which is pushed forward to $J : O \to O_2$, converted to an isogeny and returned as the signature of the message. With the exception of J, this is depicted in Figure 25.

CHAPTER 6

Concluding Remarks

In this thesis we have explored the world of isogenies between supersingular curves and between ordinary elliptic curves. Although the methods for creating the connecting isogeny have been rather different we have shown ways of doing this for both cases. Despite this, the overall process is quite similar. In this chapter we will recall the differences between the two methods and explain why we prefer the supersingular elliptic curves.

1. Differences in isogeny computation

In general, the isogenies are computing in the following way

(1) Compute the endomorphism rings
(2) Connecting the elliptic curves
(3) Map the result back to an actual isogeny

Next we will explore the differences between the ordinary and supersingular case a bit further.

1.1. Endomorphism rings. In the ordinary case, the endomorphisms rings are orders in an imaginary quadratic number field. There are only a finitely number of possible endomorphism rings and they are all of the form $\mathbb{Z} + cO_K$, where c divides the conductor $f = [O_K : \mathbb{Z}[\pi]]$. Kohel's "Find the floor" algorithm allows us to find the exact endomorphism ring by simply factoring the value f and figuring out at which level one is with respect to each prime. Thus there is a rather simple way of computing the exact abstract endomorphism ring of any elliptic curve. Notice the the conductor f is in the worst case of size $O(\sqrt{p})$, but allegedly much smaller in practice. In fact, for the large f conductors, Galbraith [16] suggests to simply discard this step and try to connect the elliptic curves directly (without moving their endomorphisms to the surface of the volcano). In total, this step takes time $O(f^3 \log(p))$

In the supersingular case on the other hand, the endomorphism rings are maximal orders in a quaternion algebra. For some curves (like $y^2 = x^3 + x$) we know what the maximal order is, but for most curves we do not. Furthermore we have no way to compute it efficiently. The perhaps most promising and recent development is using suborders, as described in [12] which runs in time $O((\log p)^2\sqrt{p})$ which can easily be defeated by choosing the prime p large enough.

1.2. Connecting the elliptic curves. In the ordinary case, the trick is to move our curve to the surface. There one uses probabilistic arguments to determine when one is expected to

find a connecting isogeny by sampling random elements from a rather small set of candidates. A rather simple and generic approach of using a breadth-first search from both ends eventually give a collision resulting in a connecting isogeny. This step takes time $O(\log(h)\sqrt{h}(\log(p)^5 + \log(h)))$, where h is the class number of $\mathrm{End}(E) \otimes \mathbb{Q}$. This depends heavily on which curve is selected. It can be as high as $O(p^{1/3}\log(p))$ which would make this run in time roughly $O(p^{1/4})$.

For the supersingular curves we have a more concrete approach where the only non-deterministic approach is equivalent of searching for primes. Thereafter the entire process is about solving equations which in the end provide an ideal of the proper norm. Unlike the ordinary case, this approach is rather explicit giving us a clear ideal in the end with the form we want. Furthermore, as we saw in the run time analysis, this method runs in time roughly $O(\log(p))$.

Thus it is in fact much faster to connect the elliptic curves in the supersingular case than in the ordinary case.

1.3. Computing the isogeny. Once we have connected the two elliptic curves we need to get a connecting isogeny. For the ordinary case, this is easy as we throughout the algorithm keep working with isogenies (or ratherthe actual j-invariants). The theory of lattices in the complex plane is only used to argue for termination and isogeny candidates. Thus this final step simply involves computing the dual of some isogenies so the resulting isogeny goes in the correct direction.

For the supersingular case the story is quite different. Here we are working with the more abstract object of ideals in quaternion orders directly. We use the structure of certain orders to find specific elements and the output is an ideal of a given norm. This then needs to be translated back to an isogeny between elliptic curves. To do so, we are required to evaluate elements of the endomorphism on actual points of the elliptic curve. Since the endomorphisms are given as elements of an abstract Q-algebra this is not a trivial task. Essentially we are required to know the endomorphism ring explicitly or at least have a way of evaluating it. However, assuming we can evaluate abstract endomorphism rings easily, this process can be achieved in time polynomial in $\log p$ (assuming that the isogeny is $(\log p)$-power smooth).

1.4. Remarks. From the above discussion it might seem counter-intuitive to use supseringular elliptic curves. Choosing an ordinary elliptic curve with a really large conductor would provide a presumably better security than supersingular curves of the same size. However, we are interested in post-quantum security, where the commutative endomorphism ring structure of ordinary curves turn out to be broken easily on a quantum computer. Using Kuperberg's hidden shift algorithm [23], Childs Jao and Soukharev [7] provide a quantum algorithm that constructs isogenies between ordinary elliptic curves in subexponential time. Unfortunately we did not have time to explore such quantum algorithms in this thesis.

One interesting fact about the supersingular case is that as soon as one has knowledge of the endomorphism ring, everything runs fast. Thus one can build schemes relying on proof of knowledge, just like we saw in the case of SQISign application.

Bibliography

[1] K. M. August. *Explicit Representation of the Endomorphism Rings of Supersingular Elliptic Curves*. https://pages.ramapo.edu/~kmcmurdy/research/McMurdy-ssEndoRings.pdf. 2014.

[2] E. Bach and J. Sorenson. "Explicit bounds for primes in residue classes". In: *Mathematics of Computation* 65 (216 1996). DOI: 10.1090/s0025-5718-96-00763-6.

[3] J. M. Basilla and H. Wada. "On the solution of $x^2 - dy^2 = \pm m$". In: *Proceedings of the Japan Academy Series A: Mathematical Sciences* 81 (8 2005). DOI: 10.3792/pjaa.81.137.

[4] D. X. Charles, K. E. Lauter, and E. Z. Goren. "Cryptographic hash functions from expander graphs". In: *Journal of Cryptology* 22 (1 2009). DOI: 10.1007/s00145-007-9002-x.

[5] Y. Chen et al. "Relations between minkowski-reduced basis and θ-orthogonal basis of lattice". In: vol. 9219. 2015. DOI: 10.1007/978-3-319-21969-1_15.

[6] I. Chevyrev and S. D. Galbraith. "Constructing supersingular elliptic curves with a given endomorphism ring". In: *LMS Journal of Computation and Mathematics* 17 (Special Issue A 2014). DOI: 10.1112/S1461157014000254.

[7] A. Childs, D. Jao, and V. Soukharev. "Constructing elliptic curve isogenies in quantum subexponential time". In: *Journal of Mathematical Cryptology* 8.1 (2014), pp. 1–29. DOI: doi:10.1515/jmc-2012-0016.

[8] H. Cohen. *A Course in Computational Algebraic Number Theory*. Graduate Texts in Mathematics. Springer Berlin Heidelberg, 2013. DOI: 10.1007/978-3-662-02945-9.

[9] D. A. Cox. *Primes of the form $p = x^2 + ny^2$: Fermat, class field theory, and complex multiplication*. Second edition. Hoboken, New Jersey: John Wiley & Sons, Inc, 2013. DOI: 10.1002/9781118400722.

[10] R. Dimitrij. "Constructing the Deuring correspondence with applications to supersingular isogeny-based cryptography". https://www.martindale.info/research/Dimitrij-thesis.pdf. MA thesis. Technische Universiteit Eindhoven, 2018.

[11] K. Eisenträger et al. "Supersingular isogeny graphs and endomorphism rings: reductions and solutions". In: vol. 10822 LNCS. 2018. DOI: 10.1007/978-3-319-78372-7_11.

[12] K. Eisenträger et al. "Computing endomorphism rings of supersingular elliptic curves and connections to path-finding in isogeny graphs". In: *Open Book Series* 4 (1 2020). DOI: 10.2140/obs.2020.4.215.

[13] N. Elkies. *Elliptic and modular curves over finite fields and related computational issues*. 1997. DOI: 10.1090/amsip/007/03.

[14] L. D. Feo, D. Jao, and J. Plût. "Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies". In: *Journal of Mathematical Cryptology* (2014). DOI: 10.1515/jmc-2012-0015.

[15] L. D. Feo et al. "SQISign: Compact Post-quantum Signatures from Quaternions and Isogenies". In: vol. 12491 LNCS. 2020. DOI: 10.1007/978-3-030-64837-4_3.

[16] S. D. Galbraith. "Constructing Isogenies between Elliptic Curves Over Finite Fields". In: *LMS Journal of Computation and Mathematics* 2 (1999). DOI: 10.1112/s1461157000000097.

[17] S. D. Galbraith. *Mathematics of public key cryptography*. 2012. DOI: 10.1017/CBO9781139012843.

[18] S. D. Galbraith, C. Petit, and J. Silva. "Identification Protocols and Signature Schemes Based on Supersingular Isogeny Problems". In: *Journal of Cryptology* 33 (1 2020). DOI: 10.1007/s00145-019-09316-0.

[19] S. D. Galbraith et al. "On the security of supersingular isogeny cryptosystems". In: vol. 10031 LNCS. 2016. DOI: 10.1007/978-3-662-53887-6_3.

[20] J.-I. Igusa. "Kroneckerian Model of Fields of Elliptic Modular Functions". In: *American Journal of Mathematics* 81 (3 1959). DOI: 10.2307/2372914.

[21] D. R. Kohel. "Endomorphism rings of elliptic curves over finite fields". http://iml.univ-mrs.fr/~kohel/pub/thesis.pdf. PhD thesis. University of California, Berkeley, 1996.

[22] D. Kohel et al. "On the quaternion l-isogeny path problem". In: *LMS Journal of Computation and Mathematics* 17 (Special Issue A 2014), pp. 418–432. DOI: 10.1112/S1461157014000151.

[23] G. Kuperberg. "A subexponential-time quantum algorithm for the dihedral hidden subgroup problem". In: *SIAM Journal on Computing* 35 (1 2006). DOI: 10.1137/S0097539703436345.

[24] S. Lang. *Elliptic Functions*. Graduate texts in mathematics. Springer, 1987. DOI: 10.1007/978-1-4612-4752-4.

[25] J.-F. Mestre. "Le Méthode des graphes. Exemples et applications". In: *Preceedings of the International Conference on Class Numbers and Fundamental Units of Algebraic Number Fields* (1986), pp. 217–242.

[26] H. Minc. "Left and Right Ideals in the Ring of 2×2 Matrices". In: *The American Mathematical Monthly* 71.1 (1964), pp. 72–75. DOI: 10.2307/2311311.

[27] D. Moody et al. *Status Report on the Second Round of the NIST Post-Quantum Cryptography Standardization Process*. en. July 2020. DOI: https://doi.org/10.6028/NIST.IR.8309.

[28] P. Q. Nguyen and D. Stehlé. "Low-Dimensional Lattice Basis Reduction Revisited". In: *Algorithmic Number Theory*. Ed. by D. Buell. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 338–357. DOI: 10.1007/978-3-540-24847-7_26.

[29] C. Petit and K. Lauter. *Hard and Easy Problems for Supersingular Isogeny Graphs*. Cryptology ePrint Archive, Report 2017/962. https://eprint.iacr.org/2017/962. 2017.

[30] A. Pizer. "An algorithm for computing modular forms on $\Gamma_0(N)$". In: *Journal of Algebra* 64 (2 1980). DOI: 10.1016/0021-8693(80)90151-9.

[31] *Post-Quantum Cryptography — CSRC*. URL: https://csrc.nist.gov/Projects/post-quantum-cryptography/. (accessed: 16.05.2021).

[32] W. Stein et al. *Sage Mathematics Software (Version 9.2)*. http://www.sagemath.org. The Sage Development Team. 2021.

[33] P. Samuel and A. Silberger. *Algebraic Theory of Numbers*. Dover Books on Mathematics Series. Dover Publications, 2008.

[34] R. Schoof. "Counting points on elliptic curves over finite fields". In: *Journal de Théorie des Nombres de Bordeaux* 7 (1 1995). DOI: 10.5802/jtnb.142.

[35] J. Silverman. *Advanced Topics in the Arithmetic of Elliptic Curves*. Graduate Texts in Mathematics. Springer New York, 1994. DOI: 10.1007/978-1-4612-0851-8.

[36] J. Silverman. *The Arithmetic of Elliptic Curves*. Graduate Texts in Mathematics. Springer New York, 2009. DOI: 10.1007/978-0-387-09494-6.

[37] A. Sutherland. *Elliptic Curves. 18.783*. Massachusetts Institute of Technology: MIT OpenCourseWare. 2019. URL: https://ocw.mit.edu. License: Creative Commons BY-NC-SA.

[38] J. Tate. "Endomorphisms of abelian varieties over finite fields". In: *Inventiones Mathematicae* 2 (2 1966). DOI: 10.1007/BF01404549.

[39] J. Voight. *Quaternion Algebras*. Graduate texts in mathematics. Springer, 2021. DOI: 10.1007/978-3-030-56694-4.

[40] L. C. Washington. *Elliptic Curves: Number Theory and Cryptography, Second Edition*. 2nd ed. Chapman & Hall/CRC, 2008. DOI: 10.1201/9781420071474.

[41] W. C. Waterhouse. "Abelian varieties over finite fields". en. In: *Annales scientifiques de l'École Normale Supérieure* Ser. 4, 2.4 (1969), pp. 521–560. DOI: 10.24033/asens.1183.

APPENDIX A

Performing computations using SageMath

In this appendix we will describe the various computations we performed using SageMath[32]. There is already implemented a lot of functionality related to elliptic curves. However the documentation is slim and all the required functionality is not implemented yet. Therefore, by explaining how we created the algorithms, we hope that we have simplified the task for others.

We begin with the simplest case of creating isogenies from an elliptic curve. Then we show how we can evaluate endomorphisms on elliptic curve points when the maps are viewed abstractly as elements of $(-1, -p \mid \mathbb{Q})$. Then we show how we can create an ideal based on the kernel point of an isogeny and knowing how to evaluate the endomorphisms. Next we give a method to represent integers in orders when they are special p-extremal, and show some tricks related to performing linear algebra in Sage. Finally we show how to turn an ideal into an isogeny through first creating the chain of ideals and then an algorithm for computing the ideal, given that we know the embedding of the endomorphism ring into the quaternion algebra.

In the code snippets we present, we will sometimes leave out repetitive code. Therefore if you read a line of code using something that appears to be undefined, take a look at the previous code and see if it is defined there. Furthermore for the readers of the digital version of this thesis, we have provided links for the documentation throughout this appendix. The analogue reader may search for the same titles in an appropriate search engine to obtain the same documentation.

1. Isogeny graphs

We begin with the easiest task, which is more or less implemented by default, namely computing separable isogenies. To perform this task we used the methods described in the documentation for Elliptic curves over finite fields and Isogenies. The methods we are using are already implemented, so the challenge was only related to finding the proper combination yielding the desired result. In essence our goal is to compute isogenous elliptic curves when restricting ourselves to isogenies of a given prime degree, and using a rather small finite field. We choose the prime $p = 439$ and starting curve $E : y^2 = x^3 + x$.

```
p = 439
field = GF(p*p) # Creates a galois field of p*p elements
E = EllipticCurve(field, [1,0]) #Creates elliptic curve E with A = 1 and B = 0
l = 5 # The degree for our isogenies
```

```
isogenies = E.isogenies_prime_degree(l)
for isogeny in isogenies:
    print("Elliptic curve:", isogeny.codomain())
    print("Isogeny:", isogeny)
    print("J invariant:", isogeny.codomain().j_invariant())
```

The above code can easily be extended to a breadth first search to expand all elliptic curves by appending the codomain to the list for further expansion. This is in particular how we create the image in Figure 17.

2. Evaluating the endomorphism generators on points

We assume that we are working with a curve of known endomorphism ring generated by $(1+\pi)/2$, $(\iota+\iota\circ\pi)/2$, π and $\iota\circ\pi$, where $\iota : (x, y) \mapsto (-x, \sqrt{-1}y)$ and $\pi : (x, y) \mapsto (x^p, y^p)$. Our approach for evaluating them on the elliptic curve E consists of performing each of the ι and π maps separately on the coordinates of the points, then mapping the points to the elliptic curve.

We create the methods `Ec1`, `EcI`, `EcJ` and `EcK` corresponding to the generators $1, i, j, ij$ of the quatenion algebra. These methods will be useful later on for more general quaternion element evaluations. Notice that we cannot simply return the tuple $(-P[0], \sqrt{-1}P[1], P[2])$ as this would not have been an elliptic curve point so we cannot add it together with another point later.

Furthermore, since the quaternion elements are fractions we are often in the need to divide the points by some integer. To accomplish this we have written a helper method, `point_divide(P,n)`, which takes as input a point and the integer to divide it by. It uses the built in method of elliptic curve points to evaluate if it is divisible by a certain number. Then `P.division_points(n)` returns a list of points, Q, of E which satisfy $[n]Q = P$.

Remark. The method division_points provide points of different orders. For other applications it might be necessary to choose a particular division point, instead of simply the first like we do here.

Finally we create the methods g_i which corresponds to the generators of the maximal order.

```
p = 439
field = GF(439**2)
sqrt = field(-1).sqrt() # creates the square root of -1

E = EllipticCurve(field, [1,0])

def point_divide(P, n):
    if P.is_divisible_by(n):
        return P.division_points(n)[0]
    else:
        # Handle exception where point is not divisible by n

def Ec1(P):
  return P
```

```
def EcI(P):
  return E(-1 *P[0], sqrt* P[1])
def EcJ(P):
  return E(P[0]^p, P[1]^p, P[2]^p)
def EcK(P):
  return E(-1 *P[0]^p, sqrt* P[1]^p, P[2]^p)

## First generator: (1 + j)/2
def g1(P):
  return point_divide(EcI(P) + EcJ(P), 2)

## Second generator: (i + ij)/2
def g2(P):
  return point_divide(EcI(P) + EcK(P), 2)

## Third generator: j
def g3(P):
  return EcJ(P)

## Fourth generator: ij
def g4(P):
  return EcK(P)
```

3. Isogeny to ideal

In this section we introduce the quaternions in order to implement Algorithm 5 to turn an isogeny and known endomorphism ring into an ideal in a quaternion algebra. The documentation for Sage can be found in Quaternion Algebra and Quaternion Algebra Element.

We would like to construct the ideal corresponding to an isogeny, and we assume that the endomorphism ring of the elliptic curve is known. That is, we use the maps as defined in the previous section for evaluating endomorphisms on elliptic curve points. Furthermore we let E be the curve as defined before and n be the degree of our isogeny. We assume that the point $(125, 82)$ generates the kernel of the isogeny.

To make this algorithm work we need to introduce the QuaternionAlgebra class, which we instantiate in line 3. Furthermore we make use of the default basis, and fortunate for us the function `Q.maximal_order()` returns exactly the maximal order we are interested in.

```
import random

Q.<i,j,k> = QuaternionAlgebra(QQ, -1, -p) #Initialize a quaternion algebra
    ramified at -1 and -p with variables i, j and k := ij
O0 = Q.maximal_order()
b1 = O0.basis()[0]
b2 = O0.basis()[1]
b3 = O0.basis()[2]
```

```
8  b4 = O0.basis()[3]
9
10 isogeny_ker_gen = E((125, 82, 1)) #Specify kernel generator of an isogeny
11
12 # Compute the elliptic curve point associated to the linear
13 # combination of a,b,c,d of the image of the kernel point
14 def alphaComputer(a,b,c,d, gen1, gen2, gen3, gen4):
15     return a*gen1 + b*gen2 + c*gen3 + d * gen4
16
17 def simpleIsogenyToIdeal(kernel_point):
18     gen1 = g1(kernel_point)
19     gen2 = g2(kernel_point)
20     gen3 = g3(kernel_point)
21     gen4 = g4(kernel_point)
22     while True:
23         a = randint(1,p)
24         b = randint(1,p)
25         c = randint(1,p)
26         d = randint(1,p)
27         alpha = a * b1 + b * b2 + c*b3 + d*b4
28         if( gcd(alpha.reduced_norm(), 5) == 5):
29             if( alphaComputer(a,b,c,d,gen1,gen2,gen3,gen4) == 0):
30                 return (a,b,c,d)
31
32 (a,b,c,d) = simpleIsogenyToIdeal(isogeny_ker_gen)
33 alpha = a * b1 + b * b2 + c*b3 + d*b4
34
35 # Want I = On + O alpha
36 On_generators = [b1 * n, b2 * n, b3*n, b4*n]
37 Oalpha_generators = [b1 * alpha, b2 * alpha, b3 * alpha, b4 * alpha]
38 I = O0.left_ideal(ideal_On + ideal_Oalpha)
```

To perform the norm computation we simply create an element of the quaternion algebra by making a linear combination of the basis elements of `O0`. We get those elements by calling `O0.basis()` which returns them as a list, next we compute the value alpha as a linear combination of these elements. Elements of the quaternion algebra have the associated method `reduced_norm()` as shown on line 28, which computes the reduced norm of that element. Next on line 29 we compare the linear combination of the image points of the kernel element to the element 0. Sage is smart enough to understand that we really mean the element $\mathcal{O}$, when writing 0, so this works great.

To compute the ideal $I = On + O\alpha$ we make two lists of the generators of On and $O\alpha$ respectively. Then we combine those generators to create a left O_0-ideal using the built in function `left_ideal()` which works on any order, and takes as input a list of generators. In

the end we should get an ideal of reduced norm n, which we can verify by calling `ideal.norm()`. Furthermore we can verify that its left order is `O0` by calling `ideal.left_order()`.

3.1. Computing the forward endomorphism ring. Using the above code we can successfully transform an isogeny to its ideal. This ideal has the wonderful method which gives its right order: `I.right_order()`. Thus computing $\phi : E_0 \to E$ from the known starting curve $E_0 : y^2 = x^3 + x$, then finding its kernel generator, and then performing the above code gives us the connecting ideal I and right order corresponding the the endomorphism ring of E quite easily.

4. KLPT Algorithm

In this section we will look at the main subalgorithm of the KLPT algorithm, namely how to represent an integer using the quadratic form. The remaining computations of the KLPT algorithm involves the explicit isomorphism to the matrix ring $M_2(\mathbb{Z}/N\mathbb{Z})$ and some elementary methods which we will only discuss briefly in the following subsection.

4.1. Represent integer in order. This section is about performing Algorithm 6. The Sage documentation can be found in Number Fields and Binary Quadratic Forms with Integer Coefficients.

```
Fx.<x>=QQ[] # Initialize variable x
NF.<i2> = NumberField(x^2 +1) # Q(sqrt(i))

def representInteger(M, m):
    f = BinaryQF([1,0,1])
    while True:
        x_2 = randint(0,m)
        y_2 = randint(0,m)
        r = M - p*f(x_2, y_2)
        if not is_prime(r):
            continue # r is not prime
        Ir = NF.ideal(r) # Ideal above r
        if len(Ir.factor()) < 2:
            continue # r does not split
        t = f.solve_integer(r) # Equivalent to Cornaccia's algorithm
        It = NF.ideal(t[0] + t[1] * NF.gen()) # Ideal generated by t
        if len(It.factor()) == 1:
            break # It is principal ideal
    return (t[0], t[1], x_2, y_2)
```

We instantiate the number field $\mathbb{Q}(\sqrt{-1})$ with the variable `i2` since `i` is taken by the previously declared quaternion algebra `Q`. Next we need to use binary quadratic forms which is available as the class `BinaryQf([a,b,c])` which creates the form $f(x, y) = ax^2 + bxy + cy^2$. Instead of implementing Cornaccia's algorithm as in Algorithm 1, we use the built-in `solve_integer(r)` method of the `BinaryQF` class which returns a tuple (x, y) such that $f(x, y) = r$.

Notice how we create the ideals corresponding generated by given elements by calling `NF.ideal(r)` much like we created the ideals in the quaternion algebra. Here we also have the method `factor` which will factor the ideal and return generators of it. If the length of this list is 1 then the ideal is principal, otherwise it is split.

4.2. Linear algebra. The core of the KLPT algorithm depends on the explicit embedding in $M_2(\mathbb{Z}/N\mathbb{Z})$ and selecting certain outputs. We therefore did not spend time writing a generic algorithm for this part of the process. Instead we will describe some useful functions when performing the calculations in Sage. More documentation on this can be found in Linear Algebra

First of all we can perform matrix computations in $M_2(\mathbb{Z}/N\mathbb{Z})$. We do this by initializing the number field with `ZN = Integers(N)`. Next we can create the matrices by calling `matrix`. For example we can construct the identity like this `M1 = matrix(ZN, [[1, 0], [0,1]])`. This is quite useful when performing matrix multiplication. It also allows us to call the method `A.solve_right(B)` if we want to find the matrix X that satisfy $AX = B$.

The remainder of the algorithm is by working with basis elements in the quotient, and solving some modular calculations. Performing inversions can be done with `inverse_mod(a,N)` which returns $a^{-1} \mod N$. There is also a method for solving modular expressions `solve_mod()`, but we did not explore this method as our numbers where already so small that they could be easily computed by hand.

5. Ideal to isogeny

Going from an ideal to an isogeny is a two step process. First we create the chain of ideals, each of prime norm l. Then we compute the isogeny corresponding to each prime normed ideal.

5.1. Compute ideal chain. To turn an ideal of l power norm we use part of Algorithm 8. Here we simply assume that we have the quaternion algebra Q available from earlier.

```
O0 = Q.maximal_order()
OI = Q.ideal(O0.basis())

def idealToIdealChain(I, l, e):
  Ik = []
  for a in range(e + 1):
    Ik.append(O0.left_ideal(I.basis() + OI.scale(7**a).basis()))
  Jk = []
  for a in range(e):
    Jk.append(Ik[a].conjugate().scale(1/Ik[a].norm()) * Ik[a + 1])
  return Jk
```

The method is takes as input an ideal I to be converted to a chain of ideals J_k, and the integers l, e satisfying $\mathrm{nrd}(I) = l^e$. The method is quite straight forward, essentially just providing what the algorithm is supposed to do. It ends by returning a list of J_k ideals corresponding to the chain of ideals, each will have reduced norm l.

First we notice that when we append to `Ik` we use `OI.scale()` instead of `OO.scale` where `OI` is a quaternion ideal while `OO` is a quaternion order. The reason for this is that `scale` is only available for ideals, so this is just a simplification of our code. We could just as easily have taken the basis elements of `OO` and multiplied them with `7**a`.

Second when we append to `Jk` we make take the product of two ideals which is already defined in the package for quaternion algebra. Here the inverse of `Ik[a]` is computed as `Ik[a].conjugate().scale(1/Ik[a].norm())` which is just $\overline{I_a}/\operatorname{nrd}(I_a)$, exactly what we want.

5.2. Compute isogeny from ideal. This is an implementation of Algorithm 7 where we assume that we know the map $\iota_{E_0}^{-1}$ of the basis elements of O_0 given by the explicit methods `Ec1` `EcI` `EcJ` and `EcK` as defined earlier. Furthermore we make use of the `point_divide` method from earlier.

```
def findLTorsion(E,l):
  while True:
    R = E.random_point()
    if gcd(R.order(), l) == l:
      # l divides order => we can divide l
      R = (R.order() // l) * R
      if R.order() == l:
        return R

def isLinearlyDependent(P, Q, n):
  for i in range(n):
    if P == n*Q:
      return True
  return False

def findLTorsionGenerators(E, l):
  P = findLTorsion(E,l)
  Q = findLTorsion(E,l)

  while isLinearlyDependent(P, Q, l):
    Q = findLTorsion(E,l)
  return (P,Q)

def endomorphismEvaluator(endomorphism, P):
  knownEmbeddings = [Ec1, EcI, EcJ, EcK]
  (d,x,y,z,w) = endomorphism.denominator_and_integer_coefficient_tuple()
  numerator = x*knownEmbeddings[0](P) + y * knownEmbeddings[1](P) + z*
    knownEmbeddings[2](P) + w*knownEmbeddings[3](P)
  if d != 1:
      return point_divide(numerator, d)
  return numerator

```

```
def idealToIsogeny(ideal, EC):
  l = int(ideal.norm()[0]) # Integer ie numerator
  (R,S) = findLTorsionGenerators(E,l)
  for basisElement in ideal.basis():
    R1 = endomorphismEvaluator(basisElement, R)
    S1 = endomorphismEvaluator(basisElement, S)
    Q = []
    if S2 == 0:
        Q.append(S)
    for b in range(l):
      if R1 + b*S1 == 0:
          Q.append(R + b*S)
    if len(Q) == 1:
        return EC.isogeny(Q[0])
```

The first three methods are naive methods for computing some l-torsion point generators. The first method takes an elliptic curve and an integer l. It samples random elements and divides out l if their order allows it. If it finds an element of order l it is simply returned. The second method is a brute-force attempt at verifying whether two points are linearly independent or not. The third method creates two generators for the l-torsion by calling the previous two helper-methods.

Next we have the `endomorphismEvaluator` method which takes in an endomorphism and a point P in the elliptic curve. It has hard-coded the methods of evaluating the endomorphism generators on points of the curve as described earlier. Next it calls `endomorphism.denominator_and_integer_coefficient()` which returns an $5-tuple$ containing the integers d, x, y, z, w satisfying $endomorphism = \frac{x+yi+zj+wij}{d}$. Then we compute the numerator as of this fraction as one would guess and finally divide out if $d \neq 1$.

Finally we have the `idealToIsogeny` method which takes the ideal and an elliptic curve to be mapped to an isogeny. It first creates the l-torsion generators. Notice that we need to turn the ideal norm into an integer, otherwise Sage will be unhappy with the code for our findLTorsionGenerators. Next we simply iterate the basis of the ideal. For each basis we see where they map the l-torsion generators and return if they are mapped to a non-trivial subgroup of $E[l]$ as described in the algorithm.

REMARK. It is possible to tweak this code to work with an arbitrary curve E having a known isogeny ϕ from E_0 to E. To do this we would extend `endomorphismEvaluator` to include an isogeny `phi`. It would then begin by computing `point_divide` on P with the degree of `phi`, next it would compute `P = phi.dual()(P)` and perform the rest of the method. Before returning it would map the point, Q, back to E by calling `phi(Q)`

Jørgen Juel

Constructing isogenies

NTNU
Norwegian University of
Science and Technology