ISOGENY VOLCANOES

NICHOLAS GEORGE TRIANTAFILLOU

Abstract. Many modern cryptographic systems rely on having access to an elliptic curve with a prescribed number of points over some fixed finite field. The Complex Multiplication method, one leading approach to finding such curves, has been improved substantially within the last 20 years using the structure of "isogeny graphs". We will describe the structure of these graphs and how they can be used to speed computation.

Of course, long before these cryptographic application were known, the study of elliptic curves with complex multiplication was a major area of study in its own right, providing the first real progress towards Hilbert's twelfth problem. As we work towards the algorithmic applications of isogeny volcanoes, we will introduce complex multiplication and Hilbert class polynomials, a beautiful theory that Kronecker described as the "dearest dream of his youth."

Contents

1. Introduction

2. Orders in $\mathbb{Q}$-Algebras

Definition 2.1. Given a finitely-generated $\mathbb{Q}$-algebra A, an order $\mathcal{O}$ of A is a subring (containing 1) such that

(1) $\mathcal{O}$ is finitely-generated as a $\mathbb{Z}$-module and
(2) $\mathcal{O}$ contains a $\mathbb{Q}$-basis of A.

Example 2.2. Let $A = K$ be a number field with ring of integers $\mathcal{O}_K$. Suppose that $\alpha_1, \ldots, \alpha_n \in \mathcal{O}_K$ are such that $K = \mathbb{Q}(\alpha_1, \ldots, \alpha_n)$. Then, $\mathbb{Z}[\alpha_1, \ldots, \alpha_n]$ is an order of K. In particular, $\mathcal{O}_K$ is an order of K.

Remark 2.3. Indeed, these are the only examples of orders in a number field, since the finite generation property implies that every order $\mathcal{O}$ of K is a subset of $\mathcal{O}_K$ and the fact that $\mathcal{O}$ contains a $\mathbb{Q}$-basis of K implies that $\mathcal{O}$ has finite index in $\mathcal{O}_K$. For this reason, the ring of integers $\mathcal{O}_K$ is often called the maximal order of K. While maximal orders exist in more general $\mathbb{Q}$-algebras, they are typically not unique.

Definition 2.4. The quaternion algebra $Q_{a,b}$ is the algebra of the form

$$Q_{a,b} = \mathbb{Q} + \mathbb{Q}\alpha + \mathbb{Q}\beta + \mathbb{Q}\alpha\beta$$

with the relations $\alpha^2, \beta^2 \in \mathbb{Q}$, $\alpha^2 = a < 0$, $\beta^2 = b < 0$, $\beta\alpha = -\alpha\beta$.

Example 2.5. Let $A = Q_{a,b}$ where $a, b \in \mathbb{Z}$ and let $\mathcal{O} = \mathbb{Z}[\alpha, \beta]$. Then, $\mathcal{O}$ is clearly an order in a quaternion algebra.

2.1. **Orders In Number Fields.** Our treatment has been strongly influenced by the coverage in Section 7 of [1]. The main difference is that we state several of the results and definitions in slightly greater generality before specializing to the case of quadratic number fields in the next section.

We first recall the definition of a discriminant of (an order of) a number field.

Definition 2.6. Let K be a number field with $\mathrm{Gal}(K/\mathbb{Q}) = \{\sigma_1, \ldots, \sigma_n\}$ and let $\mathcal{O} = [\alpha_1, \ldots, \alpha_n]$ be an order in K. Then, the discriminant $D = D_{\mathcal{O}}$ of $\mathcal{O}$ is

$$D_{\mathcal{O}} = \left(\det \begin{pmatrix} \sigma_1(\alpha_1) & \cdots & \sigma_1(\alpha_n) \\ \vdots & \ddots & \vdots \\ \sigma_n(\alpha_1) & \cdots & \sigma_n(\alpha_n) \end{pmatrix} \right)^2$$

It is clear that if $\mathcal{O} = \mathcal{O}_K$, then, this gives the usual definition for the discriminant of a number field, i.e. $d_K = D_{\mathcal{O}_K}$.

We will typically be interested in studying the ideal structure of orders. As in the case of number fields, it will be more useful to consider *fractional ideals* (i.e. subsets of K which are non-zero finitely-generated $\mathcal{O}$-modules). The following proposition summarizes some important ways that the ideal structure of $\mathcal{O}$ is similar to that of $\mathcal{O}_K$.

Proposition 2.7. *Let $\mathcal{O}$ be an order of a number field K. Then,*

(1) *If $\mathfrak{a} \subset \mathcal{O}$ is a non-zero ideal, the* norm $N(\mathfrak{a}) = |\mathcal{O}/\mathfrak{a}|$ *is finite.*
(2) *Every prime ideal of $\mathcal{O}$ is maximal.*
(3) *$\mathcal{O}$ is Noetherian.*

(4) *Every fractional ideal of $\mathcal{O}$ can be written as $\alpha\mathfrak{a}$ for $\alpha \in K^*$ and $\mathfrak{a}$ an ideal of $\mathcal{O}$.*

Proof. First, note that any (non-zero) ideal $\mathfrak{a} \subset \mathcal{O}$ is a sublattice and every full-rank subgroup of $\mathbb{Z}^n$ has finite index. Then, for any non-zero prime $\mathfrak{p}$, $\mathcal{O}/\mathfrak{p}$ is a finite integral domain, whence a finite field, so every prime ideal is maximal. $\mathcal{O}$ is finitely generated as a $\mathbb{Z}$-module, so it is certainly Noetherian as a ring. Finally, since fractional ideals are finitely generated, multiplying by some large constant makes every generator an element of $\mathcal{O}_K$, and multiplying by the index of $\mathcal{O}$ in $\mathcal{O}_K$ makes every generator an element of $\mathcal{O}$. □

The main difference between a general order $\mathcal{O}$ and $\mathcal{O}_K$ is that $\mathcal{O}$ is not integrally closed in K. As a consequence $\mathcal{O}$ is not a Dedekind domain. Therefore, $\mathcal{O}$ typically does not have unique factorization of ideals and fractional ideals do not typically have inverses. To deal with this lack of structure, we make the following definitions:

Definition 2.8. Let $\mathfrak{a}$ be a (fractional) ideal of $\mathcal{O}$.

(1) $\mathfrak{a}$ is *proper* if $\mathcal{O} = \{\beta \in K : \beta\mathfrak{a} \subset \mathfrak{a}\}$.
(2) $\mathfrak{a}$ is *invertible* if there is some fractional $\mathcal{O}$-ideal $\mathfrak{b}$ with $\mathfrak{a}\mathfrak{b} = \mathcal{O}$.

It is easy to see that $\mathcal{O} \subset \{\beta \in K : \beta\mathfrak{a} \subset \mathfrak{a}\} \subset \mathcal{O}_K$, so if $\mathcal{O} = \mathcal{O}_K$ every (fractional) ideal is proper. Similarly, it is a standard fact that fractional $\mathcal{O}_K$ ideals are invertible, so these definitions both generalize the fractional ideals of $\mathcal{O}_K$. This allows us to generalize the ideal class group of $\mathcal{O}_K$.

Definition 2.9. Let $\mathcal{O}$ be an order in a number field K. Set $I(\mathcal{O})$ to be the set of invertible fractional ideals of $\mathcal{O}$ and set $P(\mathcal{O})$ to be the set of principal fractional ideals of $\mathcal{O}$. Clearly $P(\mathcal{O}) \subset I(\mathcal{O})$. Then, the quotient

$$C(\mathcal{O}) = P(\mathcal{O})/I(\mathcal{O})$$

is the *ideal class group* of the order $\mathcal{O}$.

Unfortunately, restricting to proper or invertible ideals is not sufficient to recover prime factorization. For that, we need the concept of the *conductor* of an order.

Definition 2.10. Let $\mathcal{O}$ be an order in a number field K. The *conductor* $\mathfrak{f} = \mathfrak{f}_{\mathcal{O}}$ of the order $\mathcal{O}$ (in $\mathcal{O}_K$) is the ideal

$$\mathfrak{f} = \{\alpha \in \mathcal{O}_K : \alpha\mathcal{O}_K \subset \mathcal{O}\}.$$

As the following theorem demonstrates, the conductor is essentially the only obstruction to unique factorization. So, it is reasonable to consider what happens if we "avoid" the conductor by considering only relatively prime ideals. It turns out that this does not affect the class group. In fact, it will provide an alternate definition of $C(\mathcal{O})$ that will be extremely useful when we discuss the Class Field Theory of $\mathcal{O}$ in Section 4.

Theorem 2.11. *Let $\mathcal{O}$ be an order in a number field K with conductor $\mathfrak{f}$. Let $I(\mathcal{O}, \mathfrak{f})$ be the group of invertible fractional ideals of $\mathcal{O}$ generated by the ideals that are relatively prime to $\mathfrak{f}$ and let $P(\mathcal{O}, \mathfrak{f})$ be the subgroup of principal ideals of $\mathcal{O}$ that are generated by the principal ideals relatively prime to $\mathfrak{f}$. Then,*

(1) *$I(\mathcal{O}, \mathfrak{f})$ has unique factorization into prime ideals*
(2) *$C(\mathcal{O}) \cong I(\mathcal{O}, \mathfrak{f})/P(\mathcal{O}, \mathfrak{f})$ by a canonical isomorphism.*

For a proof of (i) in the general case, see [5]. For a proof of (ii) in the general case, see Osserman's Notes [6]. Unfortunately, discussing the general proof would take us rather far afield from our goal of discussing complex multiplication. In theorem 2.15, we will give a slightly more specific characterization of the ideal class group in the case where K is an imaginary quadratic extension of $\mathbb{Q}$. The discussion following theorem 2.15 will provide some commentary on the proof of theorem 2.11 in this special case.

2.2. **Orders in Imaginary Quadratic Fields.** Having laid out the general properties of orders in general number fields, we now specialize to the case where K is an imaginary quadratic field extension, following the exposition of Cox in [1]. For $N \in \mathbb{Z}$, $N < 0$, square-free, an easy computation shows that the quadratic field $K = \mathbb{Q}(\sqrt{N})$ has discriminant

$$d_K = D_{\mathcal{O}_K} = \begin{cases} N & \text{if } N \equiv 1 \pmod 4 \\ 4N & \text{otherwise.} \end{cases}$$

Writing $\tau_K = \frac{d_K + \sqrt{d_K}}{2}$, it is simple to check that τ_K is integral in K and generates the ring of integers, so

$$\mathcal{O}_K = \mathbb{Z}[\tau_K] = \begin{cases} \mathbb{Z}\left[\frac{1+\sqrt{N}}{2}\right] & \text{if } N \equiv 1 \pmod 4 \\ \mathbb{Z}\left[\sqrt{N}\right] & \text{otherwise.} \end{cases} \tag{2.1}$$

Because we shall be working with lattices a great deal in the remainder of this paper, we note that $\mathbb{Z}[\tau_K] = [1, \tau_K] \subset \mathbb{C}$. By example 2.2, it is clear that $[1, f\tau_K]$ is an order of K for all $f \in \mathbb{Z}_{>0}$. Indeed, as the following lemma shows, these are the only orders of K.

Lemma 2.12. *Let $\mathcal{O}$ be an order in a quadratic field K with discriminant d_K. Set $\tau_K = \frac{d_K + \sqrt{d_K}}{2}$. Then, $\mathcal{O}$ has finite index in $\mathcal{O}_K$ and if $f = [\mathcal{O}_K : \mathcal{O}]$, then*

$$\mathcal{O} = \mathbb{Z} + f\mathcal{O}_K = [1, f\tau_K]$$

Remark 2.13. It is not hard to see that the conductor of $\mathbb{Z} + f\mathcal{O}K$ is precisely the principal ideal $\mathfrak{f} = f\mathcal{O}_K$. For this reason, we follow convention and abuse notation by referring to f as the conductor of $\mathcal{O}$.

Proof. Since $\mathcal{O}$ has index f in $\mathcal{O}_K$, it is clear that $f\mathcal{O}_K \subset \mathcal{O}$. Also, $\mathbb{Z} \subset \mathcal{O}_K$ by definition, and it is clear that $\mathbb{Z} + f\mathcal{O}_K = [1, f\tau_K]$ has index f in $\mathcal{O}_K = [1, \tau_K]$, which completes the proof. □

It is an easy consequence of Lemma 2.12 that the discriminant of $\mathcal{O}$ is $D_{\mathcal{O}} = f^2 d_K < 0$.

We have already seen that the conductor is much simpler in an imaginary quadratic field. Proper ideals (recall Definition 2.8) are also much easier to understand in this setting. In fact, we have:

Proposition 2.14. *Let $\mathcal{O}$ be an order in an imaginary quadratic field K, and let $\mathfrak{a}$ be a fractional ideal of $\mathcal{O}$. Then, $\mathfrak{a}$ is proper if and only if $\mathfrak{a}$ is invertible.*

Proof. We describe the proof of Proposition 7.4 from [1].

The fact that $\mathfrak{a}$ invertible implies $\mathfrak{a}$ proper is completely general. Suppose $\mathfrak{b}$ is a fractional ideal of $\mathcal{O}$ with $\mathfrak{ab} = \mathcal{O}$. Then, for any $\beta \in K$ with $\beta\mathfrak{a} \subset \mathfrak{a}$, $\beta\mathcal{O} = \beta\mathfrak{ab} \subset \mathfrak{ab} = \mathcal{O}$, so $\beta \in \mathcal{O}$.

The other direction relies on the fact that K is quadratic and involves an explicit computation of the inverse using the minimal polynomial for an element τ with $K = \mathbb{Q}(\tau)$. □

To conclude our treatment of orders, we make the following characterization of the ideal class group of an order in an imaginary quadratic field.

Theorem 2.15. *Let $\mathcal{O}$ be an order with conductor f in an imaginary quadratic field K. Let $I(\mathcal{O}, f)$ and $P(\mathcal{O}, f)$ be as in Theorem 2.11, let $I_K(f)$ be the group of ideals of $\mathcal{O}_K$ that are relatively prime to f and let $P_{K,\mathbb{Z}}(f)$ be the subgroup of principal ideals of $\mathcal{O}_K$ of the form $\alpha\mathcal{O}_K$, where $\alpha \in \mathcal{O}_K$ satisfies $\alpha \equiv a \mod f\mathcal{O}_K$ for some integer relatively prime to f. Then, there are natural isomorphisms*

$$C(\mathcal{O}) \cong I(\mathcal{O}, f)/P(\mathcal{O}, f) \cong I_K(f)/P_{K,\mathbb{Z}}(f)$$

Proof. We merely sketch a proof the proof given in [1] and refer the reader to Section 7.C of [1] for full details.

For the first isomorphism, the theory of quadratic forms implies that for any non-zero integer M, every ideal class in $C(\mathcal{O})$ contains an element with norm relatively prime to M. Using the structure theorem for finite abelian groups, it is not hard to see that an ideal of $\mathcal{O}$ is prime to f if and only if its norm is prime to f. Together, these facts imply that the map $I(\mathcal{O}, f) \to C(\mathcal{O})$ is surjective with kernel $I(\mathcal{O}, f) \cap P(\mathcal{O})$. Checking that $I(\mathcal{O}, f) \cap P(\mathcal{O}) \subset P(\mathcal{O}, f)$ is slightly subtle - one needs to express everything in terms of various ideals of $\mathcal{O}$ that are prime to f, making use of the norm - but is a more or less standard computation.

The second isomorphism follows from the observation that $\mathcal{O}_K$ ideals that are relatively prime to f correspond to $\mathcal{O}$ ideals that are relatively prime to f under the maps $\mathfrak{a} \mapsto \mathfrak{a} \cap \mathcal{O}$ and $\mathfrak{b} \mapsto \mathfrak{b}\mathcal{O}_K$. Moreover, these maps induce an isomorphism $I_K(f) \cong I(\mathcal{O}, f)$ that preserves the norm of the ideals. The fact that the kernel is as claimed follows from easily from the observation that $\alpha \in \mathcal{O}_K$ satisfies $\alpha \equiv a \pmod{f\mathcal{O}_K}$ for some $a \in \mathbb{Z}$ if and only if $\alpha \in \mathcal{O}$, in which case, $N(\alpha) \equiv a^2 \pmod{f}$. □

Remark 2.16. This correspondence between ideals prime to f in $\mathcal{O}$ and ideals prime to f in $\mathcal{O}_K$ gives an easy proof of (i) from 2.11 in the case where K is an imaginary quadratic field, or more generally when the conductor of $\mathcal{O}$ is principal. Essentially, if an ideal $\mathfrak{a}$ of $\mathcal{O}$ is prime to f, then every factorization into primes comes from a distinct factorization of the lifted ideal in $\mathcal{O}_K$. Then, the unique factorization in $\mathcal{O}_K$ implies that the factorization of $\mathfrak{a}$ is also unique.

2.3. **Relating Class Numbers.** Before we move on, we will use Theorem 2.15 to provide an important formula for the class number of an order $\mathcal{O}$ in terms of the class number of the maximal order $\mathcal{O}_K$ and the conductor f. This formula has very important computational implications as it will be critically important when we prove the structure of the ℓ-isogeny graph in Section 6 Our exposition is inspired by Section 7.D. of [1], which also serves as a reference providing greater detail of the results discussed here.

As anyone familiar with modern algebraic number theory knows, the sizes of various groups of units play an important role in class group computations. With this in mind, we first recall the following general lemma.

Lemma 2.17. *Let K be a number field and let $\mathfrak{a}$ be an $\mathcal{O}_K$ ideal. Then,*

$$|(\mathcal{O}_K/\mathfrak{a})^*| = N(\mathfrak{a}) \prod_{\mathfrak{p}|\mathfrak{a}} \left(1 - \frac{1}{N(\mathfrak{p})}\right),$$

where $\mathfrak{p}$ runs over prime ideals of $\mathcal{O}_K$.

Proof. This result is a generalization of the more famous result that

$$|(\mathbb{Z}/a\mathbb{Z})^*| = a \prod_{p|a} \left(1 - \frac{1}{p}\right),$$

and the proof is essentially the same. The case $\mathfrak{a} = \mathfrak{p}^t$, follows by an inductive argument. The key is to consider the exact sequence

$$1 \to \mathcal{O}_K/\mathfrak{p} \xrightarrow{\phi} (\mathcal{O}_K/\mathfrak{p}^n)^* \to (\mathcal{O}_K/\mathfrak{p}^{n-1})^* \to 1,$$

where ϕ is defined by fixing $u \in \mathfrak{p}^{n-1} \setminus \mathfrak{p}^n$ and then taking $\phi([\alpha]) = [1 + \alpha u]$. The general case follows by the Chinese Remainder Theorem. Exercises 7.28 and 7.29 of [1] give a more thorough walk-through. □

Before we prove our main result, we define a bit of notation.

Definition 2.18. Given an imaginary quadratic field K, and a prime $p \in \mathbb{Z}$, define

$$\alpha_{K,p} = \begin{cases} -1, & p \text{ is inert in } K \\ 0, & p \text{ is ramified in } K \\ 1, & p \text{ splits completely in } K. \end{cases}$$

Remark 2.19. In fact, $\alpha_{K,p}$ is the Kronecker symbol of the discriminant of the maximal order of K on p, but since we will be more concerned with how p splits in K, we use this alternate notation.

Theorem 2.20. *Let K be an imaginary quadratic field with $K \neq \mathbb{Q}(\sqrt{-3}), K \neq \mathbb{Q}(i)$ and let $\mathcal{O} = \mathbb{Z} + f\mathcal{O}_K$ be the order of conductor f in K. Then,*

$$\frac{|C(\mathcal{O})|}{|C(\mathcal{O}_K)|} = f \prod_{p|f} \left(1 - \alpha_{K,p} \frac{1}{p}\right)$$

Remark 2.21. We omit the cases $K = \mathbb{Q}(\sqrt{-3})$ and $K = \mathbb{Q}(i)$ because in these cases, $\mathcal{O}_K^* \neq \{\pm 1\}$, which complicates the proof. Besides, these will correspond to exceptional cases that we shall often omit when discussing applications to elliptic curves later on. For the general case, one needs to multiply the left-hand side by $[\mathcal{O}_K^* : \mathcal{O}^*]$.

Proof. For full details of the proof, we refer the reader to Section 7.D. of [1]. The key idea is to use the result of Theorem 2.15 that $C(\mathcal{O}) \cong I_K(f)/P_{K,\mathbb{Z}}(f)$. Then it is an easy exercise in commutative algebra to show that the sequences

$$1 \to (I_K(f) \cap P_K)/P_{K,\mathbb{Z}}(f) \to I_K(f)/P_{K,\mathbb{Z}}(f) \to I_K/P_K \to 1$$
$$1 \to (\mathbb{Z}/f\mathbb{Z})^* \to (\mathcal{O}_K/f\mathcal{O}_K)^* \to (I_K(f) \cap P_K)/P_{K,\mathbb{Z}}(f) \to 1$$

are exact, where the last map is given by $[\alpha] \mapsto [\alpha\mathcal{O}_K]$. It is clear from these sequences that

$$\frac{|C(\mathcal{O})|}{|C(\mathcal{O}_K)|} = \frac{|(\mathcal{O}_K/f\mathcal{O}_K)^*|}{|(\mathbb{Z}/f\mathbb{Z})^*|}$$

Now, K is imaginary quadratic so applying Lemma 2.17,

$$\begin{aligned}|(\mathcal{O}_K/f\mathcal{O}_K)^*| &= N(f)\prod_{\mathfrak{p}|f}\left(1-\frac{1}{N(\mathfrak{p})}\right) = f^2\prod_{p|f}\prod_{\mathfrak{p}|p}\left(1-\frac{1}{N(\mathfrak{p})}\right)\\ &= f^2\prod_{p|f}\begin{cases}\left(1-1/p^2\right), & p \text{ is inert in } K\\ (1-1/p) & p \text{ is ramified in } K\\ (1-1/p)^2 & p \text{ splits completely in } K.\end{cases}\end{aligned}$$

Substituting the corresponding formula for $|(\mathbb{Z}/f\mathbb{Z})^*|$ completes the proof. □

3. Endomorphism Rings of Elliptic Curves

Our treatment of the endomorphism rings of elliptic curves attempts to synthesize several sources, predominantly Silverman's books, [7, 8], Cox's book [1], and Sutherland's Lecture notes.

3.1. **General Restrictions on Endomorphism Rings.** To start, we shall study some general restrictions on the possible endomorphism rings of an elliptic curve. Our treatment is largely inspired by Sections III.7 and III.9 of [7], where all of the proofs that are omitted or abbreviated here can be found in full detail. Along the way, we will recall without proof several facts with which we expect the reader to be familiar, providing references when possible. Throughout this (sub)-section, K will denote a field, which may not be a quadratic number field (and indeed may not even have characteristic zero).

The first important observation is that the endomorphism ring is a chacteristic zero integral domain since the multiplication by m map is non-constant and every (non-zero) isogeny has finite kernel (c.f. III.4.2 and II.2.3 of [7].)

3.1.1. *The ℓ-adic Tate Module.* The next step is to limit the size of the endomorphism ring. Our main tool will be the ℓ-adic Tate module. Recall the notation $E[n]$ for the n-torsion points on E. We use the definition from III.7 of [7].

Definition 3.1. Let E be an elliptic curve and $\ell \in \mathbb{Z}$ a prime. The ℓ-adic Tate module of E is the group

$$T_\ell(E) = \varprojlim_n E[\ell^n],$$

where the inverse limit is taken with respect to the "multiplication-by-ℓ maps.

Now, the multiplication by ℓ map $E[\ell^{n+1}] \to E[\ell^n]$ is surjective and we know that

$$E[\ell^n] = \begin{cases}\mathbb{Z}/\ell^n\mathbb{Z} \times \mathbb{Z}/\ell^n\mathbb{Z}, & \text{if } \ell \neq \text{char}(K),\\ \mathbb{Z}/\ell^n\mathbb{Z}, & \text{if } \ell = \text{char}(K) \text{ and } E \text{ ordinary},\\ \{0\}, & \text{if } \ell = \text{char}(K) \text{ and } E \text{ supersingular}.\end{cases} \tag{3.1}$$

Then, letting $\mathbb{Z}_\ell$ denote the ℓ-adic integers, it is simple to verify:

Proposition 3.2. *$T_\ell(E)$ is a $\mathbb{Z}_\ell$ module with the following structure*

$$T_\ell(E) = \begin{cases} \mathbb{Z}_\ell \times \mathbb{Z}_\ell, & \text{if } \ell \neq \text{char}(K), \\ \mathbb{Z}_p, & \text{if } \ell = \text{char}(K) \text{ and } E \text{ ordinary,} \\ \{0\}, & \text{if } \ell = \text{char}(K) \text{ and } E \text{ supersingular.} \end{cases}$$

Remark 3.3. if $\phi : E_1 \to E_2$ is an isogeny, then $\phi(E_1[n]) \subset \phi(E_2[n])$ for all n, so there is an induced map $\phi : E_1[n] \to E_2[n]$. Moreover, taking $[n]$ to denote the multiplication-by-n map, $\phi \circ [n] = [n] \circ \phi$, so the induced maps are compatible with the direct limit structure and we have an induced map $\phi_\ell : T_\ell(E_1) \to T_\ell(E_2)$. In particular, if $\phi \in \text{End}(E)$ is an endomorphism, it induces an endomorphism $\phi_\ell \in T_\ell(E)$. Moreover, if $\phi_\ell(T_\ell(E)) = \{0\}$ and $\text{char}(K) \neq \ell$ or E is ordinary, the kernel of ϕ is infinite, and so ϕ is the zero isogeny. It is also clear that the map

$$\text{End}(E) \to \text{End}(T_\ell(E)), \phi \mapsto \phi_\ell$$

is a ring homomorphism (as it is given by restriction). Then, by Proposition 3.2, when $\ell \neq \text{char}(K)$ we can naturally view $\text{End}(E)$ as a subring of $\text{Mat}_2(\mathbb{Z}_\ell)$, the space of 2×2 matrices over $\mathbb{Z}_\ell$. Similarly, if $p = \text{char}(K)$ and E is ordinary, $\text{End}(E)$ injects into $\text{Mat}_1(\mathbb{Z}_p) = \mathbb{Z}_p$.

With relatively little work (beyond the admittedly non-trivial characterization of the torsion subgroups of E), we have already seen that the Tate module allows us to view the endomorphism ring of E inside of a $\mathbb{Z}_\ell$ module of dimension at most 4, providing a limit on the possible size of the endomorphism ring. In fact, continuing along this line, we can further restrict the size of the endomorphism ring. However, before we continue, we wish to provide some additional motivation as to why the Tate module was a natural object to look at in the first place and why we might hope that there is more that we can say.

One motivating idea is that up to composition with an isomorphism, a (non-zero) isogeny is uniquely determined by its (finite) kernel. Hence, it is natural to hope that finite subgroups of E contain enough data to severely restrict the number of possible automorphisms. Unfortunately, looking at a finite collection of finite subgroups (equivalently a single finite subgroup) of E is not enough to determine an isogeny uniquely, since both the zero isogeny and any isogeny containing that subgroup in its kernel be zero on that subgroup. However, we know that only the zero isogeny kills arbitrarily large subgroups of E and the ℓ-adic Tate module is one of the most natural ways to collect arbitrarily large finite subgroups of E into a single object.

An analogous (albeit simpler) situation arises when studying the endomorphism ring of the multiplicative group of an algebraically closed field K. Again (up to composition with an isomorphism), every endomorphism is determined by its finite kernel. The Tate module in this case corresponds to the group of all ℓth power roots of unity and is a one-dimensional $\mathbb{Z}_\ell$ module, so $\text{End}(K^*)$ lives inside of $\text{End}(\mathbb{Z}_\ell)$, a one-dimensional $\mathbb{Z}_\ell$ module. In this case, we are able to show that in fact, $\text{End}(K^*) \cong \mathbb{Z}$ is one-dimensional as a $\mathbb{Z}$-module. We might hope for a similar result for E. Indeed, the following theorem leads quickly to a proof that $\text{End}(E)$ is a $\mathbb{Z}$-module of dimension at most 4.

Theorem 3.4. *Let E_1 and E_2 be elliptic curves over K and let ℓ be a prime with $\ell \neq \mathrm{char}(K)$. Then, the map*

$$\mathrm{Hom}(E_1, E_2) \otimes \mathbb{Z}_\ell \to \mathrm{Hom}(T_\ell(E_1), T_\ell(E_2)), \qquad \phi \mapsto \phi_\ell$$

is injective.

Proof. Our proof synthesizes the arguments in Section III.7 of [7] or 13.1 of [4].

Let $\phi_1, \ldots, \phi_r$ be linearly independent in $\mathrm{Hom}(E_1, E_2)$ We claim that if $c_1, \ldots, c_r \in \mathbb{Z}_\ell$ satisfy

$$c_1\phi_1 + \cdots + c_r\phi_r = 0$$

in $\mathrm{Hom}(T_\ell(E_1), T_\ell(E_2))$, then $c_1 = \cdots = c_r = 0$.

The general approach is to show that $c_i \equiv 0 \pmod{\ell^n}$ for all $1 \leq i \leq r$. The proof requires two key ideas.

First, if $c_i = d_i + \ell^n m_i$, then if $P \in E_1[\ell^n] \subset T_\ell(E_1)$,

$$\sum_{i=1}^{r} [d_i]\phi_i(P) = \sum_{i=1}^{r} d_i\phi_i(P) + m_i\phi_i\ell(P) = \sum_{i=1}^{r} c_i\phi(P),$$

so

$$\psi = [d_1]\phi_1(P) + \cdots + [d_r]\phi_r$$

kills $E_1[\ell^n]$. Then, since ℓ is a prime not equal to the characteristic of K, $[\ell^n]$ is a separable isogeny and ψ factors as $\psi = [\ell^n] \circ \chi$ for some $\chi \in \mathrm{Hom}(E_1, E_2)$.

Now, if we knew that χ were in the span of the ϕ_i, we would be done. The second key idea is to choose the ϕ_i carefully so that this is the case. We will choose the ϕ_i to be minimal in the sense that every endomorphism that is a $\mathbb{Q}$-linear combination of the ϕ_i is a $\mathbb{Z}$-linear combination of the ϕ_i. In this case, it is clear that $\chi \in \mathrm{span}_{1 \leq i \leq r} \phi_i$ and we would be done.

To this end, note that $M = \bigoplus_{i=1}^{r} \mathbb{Z}\phi_i$ is a lattice in the r-dimensional vector space $\mathcal{M} = \bigoplus_{i=1}^{r} \mathbb{Q}\phi_i$. Extending the degree map from M to $\mathcal{M}$ and noting that $\deg(\phi) \geq 1$ for any non-zero isogeny, it is clear that $M' = \mathcal{M} \cap \mathrm{Hom}(E_1, E_2)$ is a discrete subgroup of $\mathcal{M}$ and hence a lattice containing M. Taking $\{\phi_i'\}_{i=1}^{r}$ to be a generating set for M', the previous discussion shows that the claim holds for the isogenies $\phi_1', \ldots, \phi_r'$. But the ϕ_i lie in the span of the ϕ_i', so this implies the claim for $\phi_1, \ldots, \phi_r$, completing the proof. □

With this result in hand, we are ready to prove the following corollary, which also appears in [7] and [4].

Corollary 3.5. *Let E_1 and E_2 be elliptic curves over an arbitrary field K. Then,* $\mathrm{Hom}(E_1, E_2)$ *is a free $\mathbb{Z}$-module of rank at most* 4. *In particular, taking $E = E_1 = E_2$, the additive group of* $\mathrm{End}(E)$ *is free abelian of rank at most* 4.

Proof. Let $\ell \neq \mathrm{char}(K)$. Now, $\mathrm{Hom}(E_1, E_2)$ is torsion free, so from commutative algebra, 3.4 and 3.2

$$\begin{aligned}
\mathrm{rank}_{\mathbb{Z}}(\mathrm{Hom}(E_1, E_2)) &= \mathrm{rank}_{\mathbb{Z}_\ell} \mathrm{Hom}(E_1, E_2) \\
&\leq \mathrm{rank}_{\mathbb{Z}_\ell} \mathrm{Hom}(T_\ell(E_1), T_\ell(E_2)) \\
&= \mathrm{rank}_{\mathbb{Z}_\ell} \mathrm{Mat}_2(\mathbb{Z}_\ell) \\
&= 4
\end{aligned}$$

□

Remark 3.6. When E/K is an ordinary elliptic curve over a field of characteristic p, we know that $T_p(E) = \mathbb{Z}_p$. At first glance, it might seem that this proof shows that $\mathrm{End}(E) = \mathbb{Z}$ in this case. However, the proof relies on the fact that the multiplication-by-ℓ map is separable. Since multiplication-by-p is inseparable in characteristic p, these arguments do not apply. As we shall see in Section 3.3, the Frobenius endomorphism π_E is never in $\mathbb{Z}$ in this case. In fact, $\mathrm{End}(E)$ will always be an order in an imaginary quadratic field.

Remark 3.7. It may seem rather surprising that we are able to constrain the size of $\mathrm{End}(E)$ so effectively by looking at the behaviour of the Tate module at a single prime. However, this is not the only place in the study of elliptic curves where such a phenomenon occurs. For instance, after developing the machinery of heights, the Mordell-Weil theorem follows from weak Mordell-Weil for a single prime. In fact, this is not the only parallel between these two proofs, as the proof of Mordell-Weil also relies on the Kummer pairing, which makes explicit the analogy between m-torsion points on an elliptic curve and mth roots of unity in a field.

Before we go on, we quickly remark on a partial converse to Theorem 3.4. We denote by $\mathrm{Hom}_K(E_1, E_2)$ the subset of isogenies from E_1 to E_2 that are defined over K. Then,

Theorem 3.8. *Let K be a finite field or a number field and $\ell \neq \mathrm{char}(K)$ a prime. Then, the natural map*

$$\mathrm{Hom}_K(E_1, E_2) \otimes \mathbb{Z}_\ell \to \mathrm{Hom}_K(T_\ell(E_1), T_\ell(E_2))$$

is an isomorphism.

The proof in the finite field case is due to Tate. The number field case was first published by Faltings in his proof of the Mordell Conjecture and requires most of the machinery developed there. Needless to say, this is beyond the scope of this paper.

3.1.2. *The Structure of the Endomorphism Ring.* Having restricted the size of $\mathrm{End}(E)$ through our discussion of the Tate module, we remind the reader of one additional piece of structure - the existence of a dual isogeny - that will limit the possible algebra structures for $\mathrm{End}(E)$.

We quickly recall a few important properties of the dual isogeny (see Theorems III.6.1 and III.6.2 of [7] for proofs).

Proposition 3.9. *Suppose that $\phi : E_1 \to E_2$ is a non-constant isogeny of degree m. Then, there exists a unique isogeny*

$$\widehat{\phi} : E_2 \to E_1$$

such that $\widehat{\phi} \circ \phi = [m]$. Moreover, if $E = E_1 = E_2$, then the map

$$\mathrm{End}(E) \to \mathrm{End}(E), \qquad \phi \mapsto \widehat{\phi}$$

is an anti-involution, i.e. for $\phi, \psi \in \mathrm{Hom}(E_1, E_2)$, $n \in \mathbb{Z}$

(1) $\widehat{\phi + \psi} = \widehat{\phi} + \widehat{\psi}$
(2) $\widehat{\phi\psi} = \widehat{\psi}\widehat{\phi}$
(3) $\hat{\widehat{\phi}} = \phi$
(4) $\widehat{[n]} = [n]$

With this in hand, the following theorem (9.3 of [7]) allows us to place a strong restriction on the possibilities for the the endomorphism ring of an elliptic curve.

Theorem 3.10. *Suppose R is a ring with the following properties:*

1. *The additive group of R is a free $\mathbb{Z}$-module of rank at most 4.*
2. *R has an anti-involution $\phi \mapsto \widehat{\phi}$.*
3. *For $\phi \in R$, $\widehat{\phi}\phi \in \mathbb{Z}_{\geq 0}$ and $\widehat{\phi}\phi = 0$ if and only if $\phi = 0$.*

In particular, we may take R to be the endomorphism ring of an elliptic curve. Then, one of the following three possibilities hold:

1. *$R \cong \mathbb{Z}$.*
2. *R is an order in an imaginary quadratic field.*
3. *R is an order in a quaternion algebra (as defined in 2.4).*

We briefly comment on the proof, which is a fairly straight-forward computation and refer the reader to III.9 of [7] for full details. The key idea is to extend the anti-involution to $\mathbb{Q}$ and define a norm and trace by $N(\phi) = \widehat{\phi}\phi$, $T(\phi) = \widehat{\phi} + \phi$, which corresponds to the usual norm and trace of an isogeny of elliptic curves. Using the observation that if $T(\phi) = 0$, then $\phi^2 \in \mathbb{Q}_{\leq 0}$, choosing generators and following a process reminiscent of Gram-Schmidt orthonormalization,the restriction on the rank of R quickly leads to the desired result.

In fact, it is possible to achieve the same result without an a priori bound on the rank of R. The trick is to showing that if $\phi \notin \mathbb{Q}$ and ϕ and ψ commute, then $\psi \in \mathbb{Q}(\phi)$. Then given three generators ϕ, ψ, χ appropriately normalized by linear tranformation, $\psi\chi$ commutes with ϕ and so $\chi \in \mathbb{Q}(\phi, \psi)$. See Lecture 14 of [9] for further detail. Despite the existence of this somewhat more elementary proof, we believe there are two very important reasons for introducing Tate modules.

First, Tate modules will be a useful tool when discussing how the ring of endomorphisms of different elliptic curves are related, whether these elliptic curves are related by an isogeny (as we discuss in Section 6) or by reduction modulo a prime ideal (as we discuss in Section 3.4 .)

Second, introducing the Tate module allows us to reiterate the major theme in the theory of elliptic curves that the torsion points on an elliptic curve play an analogous role to the roots of unity in a field. As we shall see when stating the theorems of Class Field Theory in Section 4, this analogy lies at the heart of the deepest and most profound results in the theory of Complex Multiplication.

3.2. **Endomorphism Rings of Elliptic Curves over $\mathbb{C}$.** Our next task is to study the endomorphisms rings of elliptic curves over $\mathbb{C}$. Our main reference for the classification of possible endomporphism rings over $\mathbb{C}$ is Section VI of [7]. Our main reference for the class group action in the complex multiplication case is Section II.1 of [8]. We will be quite brief, typically referring the reader to these texts, or their favorite textbook on compact Riemann surfaces or modular forms for the proofs.

The first major goal of the section is to discuss the following equivalence of categories

Theorem 3.11. *There is an equivalence of categories*

$$\{\textit{Objects: Elliptic curves Over } \mathbb{C} \textit{ up to isomorphism, Maps: Isogenies}\}$$
$$\leftrightarrow \{\textit{Objects: Elliptic curves Over } \mathbb{C} \textit{ up to isomorphism, Maps: Complex analytic maps taking } O_{E_1} \textit{ to } O_{E_2}\}$$
$$\leftrightarrow \{\textit{Objects: Lattices } \Lambda \subset \mathbb{C} \textit{ up to scaling, Maps: (from } \Lambda_1 \textit{ to } \Lambda_2\textit{)} \ \{\alpha : \alpha\Lambda_1 \subset \Lambda_2\} \ \}.$$

The first equivalence in Theorem 3.11 is just a bit of general theory about compact Riemann surfaces. In order to discuss the second equivalence, we first recall a few important definitions from Section 6 of [7].

Definition 3.12. Let $\Lambda \subset \mathbb{C}$ be a lattice. The Weierstrass $\wp$-function (for Λ) is given by the series

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

The Eisenstein series of weight $2k$ (for Λ) is the series

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

Finally, define

$$g_2 = g_2(\Lambda) = 60G_4(\Lambda) \qquad g_3 = g_3(\Lambda) = 140G_6(\Lambda).$$

We quickly recall a few properties relevant to our applications.

Proposition 3.13. *Let Λ be a lattice.*

1. *Every Λ-periodic meromorphic function on $\mathbb{C}$ can be written $P(\wp(z), \wp'(z))$ for some rational function P.*
2. *Every even Λ-periodic meromorphic function on $\mathbb{C}$ can be written $Q(\wp(z))$ for some rational function Q.*

The typical proof of (2) multiplies or divides by appropriate $(\wp(z) - \wp(z_0))$ to cancel the poles/zeros of the function f, using Liouville's Theorem to finish. The proof of (1) follows by writing $f = f_{odd} + f_{even}$ and multiplying the odd part by $\wp'(z)$. See VI.3.2 of [7] for further detail.

Proposition 3.14. (1) *$\wp$ satisfies the differential equation*

$$\wp'(z)^2 = 4\wp(z)^3 - g_2\wp(z) - g_3.$$

(2) *Let $E/\mathbb{C}$ be the elliptic curve*

$$E : y^2 = 4x^3 - g_2x - g_3.$$

Then, the map

$$\varphi : \mathbb{C}/\Lambda \to E \subset \mathbb{P}^2(\mathbb{C})$$
$$z \mapsto [\wp(z), \wp'(z), 1]$$

is an isomorphism of Riemann surfaces and of groups.

The proof of (1) is a simple computation using the Laurent expansion of $\wp$. For (2), standard computations verify that E is an elliptic curve and that φ is bijective. Since the pullback of the invariant differential dx/y under φ^* is just dz, the map is a local and hence a global isomophism of Riemann surfaces. The fact that the map

is a group isomorphism follows from looking at divisors and applying 3.13. Again, see Section VI.3 of [7] for further details.

The Uniformization Theorem (stated in VI.5 of [7] with numerous references and proved in I.4 of [8], albeit with a slightly different normalization so that the curve will be in Weierstrass form) implies that every elliptic curve of the form in (2) of Proposition 3.14 is the image of a unique lattice.

To complete our discussion of Theorem 3.11, we need to consider the maps between objects. If E_1 and E_2 correspond to lattices Λ_1 and Λ_2, a straightforward computation shows that holomorphic maps between E_1 and E_2 taking O_{E_1} to O_{E_2} correspond to holomorphic maps from $\mathbb{C}/\Lambda_1$ to $\mathbb{C}/\Lambda_2$ which fix 0. In turn, looking at the derivative of a lift to a map from $\mathbb{C}$ to $\mathbb{C}$, one sees that these maps are exactly the multiplication-by-α maps for $\alpha \in \mathbb{C}$ with $\alpha\Lambda_1 \subset \Lambda_2$. The easy consequence that two elliptic curves are isomorphic over $\mathbb{C}$ if and only if the corresponding lattices are scalar multiples of one another completes our discussion of Theorem 3.11. A full proof and precise statement of these facts can be found in Section VI.4 of [7].

Theorem 3.11 allows us to describe a large class of endomorphism rings of elliptic curves over $\mathbb{C}$.

Example 3.15. Let $\mathcal{O}$ be an order in an imaginary quadratic field and let $\mathfrak{a}$ be an proper fractional ideal of $\mathcal{O}$. Then, $\mathfrak{a}$ is a lattice in $\mathbb{C}$, so it corresponds to some elliptic curve $E_\mathfrak{a}$. Moreover, $\mathfrak{a}$ is a proper $\mathcal{O}$ ideal, so

$$\mathrm{End}(E_\mathfrak{a}) = \{\alpha : \alpha\Lambda_1 \subset \Lambda_2\} = \mathcal{O},$$

The following corollary shows that (together with $\mathbb{Z}$,) these are all of the possible endomorphism rings of elliptic curves over $\mathbb{C}$.

Corollary 3.16. *Let $E/\mathbb{C}$ be an elliptic curve and take $\Lambda = [1, \tau]$ to be a lattice associated to E under the equivalence of categories from Theorem 3.11. Let $K = \mathbb{Q}(\tau)$.*

1. *If K is not an imaginary quadratic field,* $\mathrm{End}(E) \cong \mathbb{Z}$.
2. *If K is an imaginary quadratic field, then* $\mathrm{End}(E) \cong \mathcal{O}$ *for some order $\mathcal{O}$ of K.*

Proof. By 3.11,

$$\mathrm{End}(E) \cong \{\alpha : \alpha\Lambda \subset \Lambda\}.$$

Suppose $\alpha\Lambda \subset \Lambda$ for $\alpha \notin \mathbb{Z}$. Then, there exist $a, b, c, d \in \mathbb{Z}$ with

$$\alpha \cdot 1 = a + b\tau,$$
$$\alpha \cdot \tau = c + d\tau.$$

Substituting for α in the second equation, we have that

$$b\tau^2 + (a - d)\tau - c = 0.$$

$\alpha \notin \mathbb{Z}$, so $b \neq 0$, whence $\mathbb{Q}(\tau)$ is a quadratic extension. Since $\tau \notin \mathbb{R}$, $\mathbb{Q}(\tau)$ is an imaginary quadratic extension.

Multiplying the second equation by b and substituting for $b\tau$ in the second equation, we have that

$$\alpha^2 - (a + d)\alpha + (ad - bc) = 0,$$

whence $\alpha \in [1, \tau]$ is algebraic over $\mathbb{Q}$. Then, $\mathrm{End}(E) \subset \mathcal{O}_K$.

Moreover, if $\mathbb{Q}(\tau)$ is quadratic imaginary and the minimal polynomial of τ is $a_2z^2+a_1z+a_0$ with $a_2, a_1, a_0 \in \mathbb{Z}$, then $a_2\tau^2, a_2\tau \in [1, \tau]$, so in this case, $\mathrm{End}(E) \subset \mathcal{O}_K$ contains a $\mathbb{R}$-basis for $\mathbb{C}$ and thus is an order in K. □

When the endomorphism ring of $E/\mathbb{C}$ is strictly larger than $\mathbb{Z}$ (i.e. case (2) of Corollary 3.16) we say that E has complex multiplication. We have seen in Corollary 3.16 that the possible endomorphism rings are orders in imaginary quadratic fields and ideal class of proper fractional ideals of $\mathcal{O}$ corresponds to an elliptic curve with endomorphism ring $\mathcal{O}$. We finish our discussion of the endomorphism rings of elliptic curves over $\mathbb{C}$ by showing that up to isomorphism, these are the only elliptic curves with endomorphism ring $\mathcal{O}$.

Proposition 3.17. *Suppose that $\mathcal{O}$ is an order of the imaginary quadratic field K and $\Lambda \subset \mathbb{C}$ is a lattice with $\mathrm{End}(E_\Lambda) = \mathcal{O}$. Then, $\Lambda = c\mathfrak{a}$ for some $c \in K^\times$ and $\mathfrak{a}$ a proper fractional ideal of $\mathcal{O}$.*

Proof. Suppose $\Lambda = [c_1, c_2]$ with $c \neq 0$ and consider the lattice $c_1^{-1}\Lambda = [1, c_1^{-1}c_2]$. By the proof of Corollary 3.16, we see that $c^{-1}c_2 \in K$, so $c_1^{-1}\Lambda \subset K$. Also, since $\{\alpha : \alpha c_1^{-1}\Lambda \subset c_1^{-1}\Lambda\} = \mathcal{O}\}$. Thus, $c_1^{-1}\Lambda$ is a finitely-generated $\mathcal{O}$-module, say $\mathfrak{a}$. Putting this together, we have that $\Lambda = c_1\mathfrak{a}$, as desired. □

Now, let $\mathcal{ELL}_{\mathbb{C}}(\mathcal{O})$ denote the set of (isomorphism classes of) elliptic curves over $\mathbb{C}$ with complex multiplication by $\mathcal{O}$.

Proposition 3.17 and Example 3.15 show that the map

$$C(\mathcal{O}) \to \mathcal{ELL}_{\mathbb{C}}(\mathcal{O}), \qquad [\mathfrak{a}] \mapsto [E_{\mathfrak{a}}]$$

is a bijection. This allows us to define an simply transitive action of $C(\mathcal{O})$ on $\mathcal{ELL}_{\mathbb{C}}(\mathcal{O})$ by $[\mathfrak{a}] \cdot [E_{\mathfrak{b}}] = E_{\mathfrak{a}^{-1}\mathfrak{b}}$. The choice to make $\mathfrak{a}$ act by multiplication by $\mathfrak{a}^{-1}$ will ease our notation later on when we discuss Hilbert Class polynomials in Section 5.3.

Before moving on to fields of positive characteristic, we briefly mention the *Lefschetz principle*, which Silverman describes as saying "roughly that algebraic geometry over an arbitrary algebraically closed field of characteristic zero is 'the same' as algebraic geometry over $\mathbb{C}$." in Section VI.6 of [7].

For example, one can show that for any elliptic curve over an algebraically closed field of characteristic 0, either $\mathrm{End}(E) \cong \mathbb{Z}$ or $\mathrm{End}(E) \cong \mathcal{O}$, an order in an imaginary quadratic field. Essentially, since $\mathrm{End}(E)$ is finitely generated, we can embed all of the coefficients of all of the rational functions defining isogenies in $\mathrm{End}(E)$ inside of $\mathbb{C}$, and apply the results presented in this section.

3.3. **Endomorphism Rings of Elliptic Curves in characteristic $p > 0$.** Having thoroughly described the $\mathrm{End}(E)$ for an elliptic curve $E/\mathbb{C}$, we now turn to elliptic curves defined over a field F where $\mathrm{char}(F) = p$.

We saw in Theorem 3.4 that the ℓ-adic Tate module $T_\ell(E)$ for $\ell \neq p$ provides a major restriction on the size of the endomorphism ring of E. While Remark 3.6 revealed that we cannot use this argument to restrict the size of $\mathrm{End}(E)$ further do to the inseparability of $[p]$, the p-adic Tate module does play a role in determining the endomorphism ring of E.

Since our main applications are to find elliptic curves over finite fields with a prescribed number of points (relatively prime to the order of the field) and to study the theory of complex multiplication, we will primarily be interested in the case

where $F = \mathbb{F}_{p^n}$ is a finite field and E/F is ordinary, closely following the treatment of Lecture 15 of [9]. We will also remark briefly on the case of supersingular curves and more general curves, referring the reader to V.III of Silverman's [7] for further details.

We first recall an alternate characterization for supersingular curves over a finite field (Theorem 15.1 of [9])

Proposition 3.18. *Let $E/\mathbb{F}_{p^n}$ be an elliptic curve over a finite field with Frobenius endomorphism π_E. Then E is supersingular if and only if $\operatorname{tr} \pi_E \equiv 0 \pmod{p}$.*

Proof. First, note that π_E is inseparable. Then, $\operatorname{tr} \pi_E \equiv 0 \pmod{p}$ if and only if $[\operatorname{tr} \pi_E]$ is inseparable. Since $[\operatorname{tr} \pi_E] = \pi_E + \widehat{\pi_E}$, this holds if and only if $\widehat{\pi_E}$ is inseparable. Since $\deg(\widehat{\pi_E}) = p^n$, and $E[\pi_E] = \{O_E\}$, this holds if and only if

$$|E[p^n]| = |E[\pi_E]| \cdot |E[\widehat{\pi_E}]| = |E[\widehat{\pi_E}]| < p^n.$$

By (3.1), which shows that $E[p^n] \cong \mathbb{Z}/p^n\mathbb{Z}$ or $E[p^n] = 0$ according to whether E is ordinary or supersingular, this occurs if and only if E is supersingular, which completes the proof. □

Sutherland gives a slightly different proof in Lecture 15 of [9], , writing $\pi_E = \pi^n$ where each π is the pth power Frobenius maps and looking at $[p] = \pi\widehat{\pi}$. (Note that each π is an isogeny, but not necessarily an endomorphism of E.) This approach has the advantage that it only requires us to know the structure of $E[p]$, rather than $E[p^n]$, but introduces additional subtleties in the form of additional elliptic curves.

We are now ready to prove the main result of this section.

Theorem 3.19. *Let $E/\mathbb{F}_{p^n}$ be an elliptic curve. Then,*

1. *If E is ordinary, then $\operatorname{End}(E)$ is an order $\mathcal{O}$ in an imaginary quadratic field with p not dividing the conductor f of $\mathcal{O}$.*
2. *If E is supersingular, then $\operatorname{End}(E)$ is an order $\mathcal{O}$ in a quaternion algebra $Q_{a,b}$.*

Proof. In the ordinary case, our proof is inspired by Section 13.2 of [4]. For the supersingular case, we outline the proof from V.III of [7] in the supersingular case, referring the reader to these sources for further details. First, suppose that E is ordinary. Then, for all $m \neq \pm 1$, $|E[m]| > 1$.

First, suppose that E is ordinary. By Remark 3.3, we know that $\operatorname{End}(E) \hookrightarrow \mathbb{Z}_p$, so it is commutative. Now, take $\pi_E \in \operatorname{End}(E)$ be the Frobenius endomorphism $(a, b) \mapsto (a^{p^n}, b^{p^n})$. Then, $\ker(\pi_E^k) = \{O_E\}$, so if $\pi_E^k \in \mathbb{Z}$, we must have $\pi_E^k = \pm 1$. But $\deg(\pi_E^k) = p^{nk}$, so this is ridiculous. Now, quaternion algebras are not commutative, so by Theorem 3.10, we must have that $\operatorname{End}(E)$ is an order in the imaginary quadratic number field $\mathbb{Q}(\pi_E)$.

Now, we have

$$\mathbb{Z}[\pi_E] \subset \operatorname{End}(E) \subset \mathcal{O}_{\mathbb{Q}(\pi_E)},$$

so to complete the proof of (1), it suffices check that $\pi_E \notin \mathbb{Z} + p\mathcal{O}_K$. Lang's treatment in 13.2 of [4] is very terse, so we provide a full argument.

Suppose $\pi_E \in \mathbb{Z} + p\mathcal{O}_K$. Then, we can write

$$\pi_E = a + p\alpha, \quad \widehat{\pi_E} = a + p\overline{\alpha}, \qquad \text{for } a \in \mathbb{Z}, \alpha \in \mathcal{O}_K$$

Since the Frobenius map has degree p^n, we have

$$p^{2n} = a^2 + pa(\alpha + \overline{\alpha}) + p^2\alpha\overline{\alpha},$$

so $p|a$ and we may write $\pi_E = p(a' + \alpha)$, $\pi'_E = p(a' + \overline{\alpha})$. Now, the action of $\mathrm{End}(E)$ on the Tate module gives us an embedding $\sigma : \mathcal{O} \hookrightarrow \mathbb{Z}_p$ which we can extend to an embedding $K \hookrightarrow \mathbb{Q}_p$. Now, the action of π_E on $T_\ell(E)$ has trivial kernel, so we must have $\sigma(\pi_E) \in \mathbb{Z}_p \smallsetminus p\mathbb{Z}_p$. Then, since $\pi_E\widehat{\pi_E} = p^{2n}$, $\widehat{\pi_E} \in p^{2n}\mathbb{Z}_p$. Hence, $\pi_E + \widehat{\pi_E} = p(2a' + \alpha + \overline{\alpha}) \notin p\mathbb{Z}_p$. But we know that $\alpha + \overline{\alpha} \in \mathbb{Z}$ and $2a' \in \mathbb{Z}$, so this is a contradiction, which completes the proof.

In the supersingular case $T_p(E)$ is trivial, and these methods do not even imply that $\pi_E \notin \mathbb{Z}$, so we need a different approach. For this case, we follow the proof from V.3 of [7] and refer the reader there for full details.

The proof is by contradiction, assuming that $\mathrm{End}(E) \otimes \mathbb{Q}$ is a number field for some supersingular curve E. There are three key ingredients.

The first is that $\mathrm{End}(E) \otimes \mathbb{Q}$ is preserved by isogenies, which we discuss further in Section 6. .

The second is that there are only finitely many (isomorphism classes of) supersingular elliptic curves (the j-invariant, which we discuss further in must lie in $\mathbb{F}_{p^2}$) and that isogenous elliptic curves are either both supersingular or both ordinary (which follows easily from the fact that isogenies have finite kernel, considering the structure of the p-adic Tate modules $T_p(E)$ and $T_p(E')$.)

The third idea is to use the fact that every finite subgroup of an elliptic curve to construct an endomorphism ϕ of some supersingular elliptic curve E' with cyclic kernel of order ℓ^n for ℓ a prime in $\mathrm{End}(E')$. This step uses finiteness twice - once to ensure that we can choose ℓ such that ℓ is prime in every $\mathrm{End}(E)$ for E supersingular and again to show that two isogenies from E with kernels $\Phi_1 \subset \Phi_2$ isomorphic to $\mathbb{Z}/\ell^{m_1}\mathbb{Z}$ and $\mathbb{Z}/\ell^{m_2}\mathbb{Z}$ respectively must map to the same elliptic curve E'. This induces an endomorphism of E' with cyclic kernel that must factor through $[\ell^{(m_2-m_1)/2}]$, which is impossible. See the proof of V.3 of [7] for full details. □

Remark 3.20. To complete our treatment of the elliptic curves over fields of characteristic p, we quickly remark that if the j-invariant is algebraic over $\mathbb{F}_p$, we may assume that E is defined over some finite field and our previous analysis applies. If instead, the j-invariant of E is transcendental over $\mathbb{F}_p$, then there is no Frobenius endomorphism and in fact, $\mathrm{End}(E) = \mathbb{Z}$.

Under the principal that this is the most generic possible behaviour, some authors refer to such elliptic curves as *ordinary* and refer to the curves that we have termed ordinary as *singular*. With this terminology, the phrasing *supersingular* is reflects the fact that these are the "rarest" elliptic curves counting by number of isomorphism classes over a sufficiently large field.

3.4. **Reduction and Some Results of Deuring.** Having characterized the possible endomorphism rings for elliptic curves over $\mathbb{C}$ and for elliptic curves over finite fields, we would like to be able to relate these two notions. In particular, if E is defined over a number field L and E has good reduction mod $\mathfrak{p}$, then $\overline{E}$ is an elliptic curve over the finite field $K/\mathfrak{p}$. One might hope that $\mathrm{End}(E) = \mathrm{End}(\overline{E})$, or at least that our knowledge of $\mathrm{End}(E)$ can be used to determine $\mathrm{End}(\overline{E})$. Conversely, we know that every ordinary curve $E'/\mathbb{F}_{p^n}$ has $\mathrm{End}(E') = \mathcal{O}$, where $\mathcal{O}$ is an order in some imaginary quadratic field. We also know from Proposition 3.17 that there

are exactly $C(\mathcal{O})$ isomorphism classes of elliptic curves over $\mathbb{C}$ with complex multiplication by $\mathcal{O}$. One might hope that a curve in one of these isomorphism classes reduces to E'.

Indeed, in 1941, Deuring proved that both of these hopes are often true. In the remainder of this section we will present a few of Deuring's main results, inspired by the treatment of Lang in Sections 13.4 and 13.5 [4], and remark on how they will be important for our computational applications.

We will always consider the case of good reduction. Note that if $\phi : E \to E'$ is an isogeny of elliptic curves over L and both E and E' have good reduction mod $\mathfrak{p}$, then ϕ descends to an isogeny $\overline{\phi} : \overline{E} \to \overline{E'}$, since it is defined algebraically. In fact, we have

Proposition 3.21. *Let E_1, E_2 be elliptic curves defined over a number field L with good reduction mod $\mathfrak{p}$. Let $L_\mathfrak{p}$ be the residue field. Then, the reduction map*

$$\mathrm{Hom}(E_1, E_2) \to \mathrm{Hom}(\overline{E_1}, \overline{E_2})$$

is injective and preserves degrees.

Proof. Choose a prime $\ell \in \mathbb{Z}$ with $\mathfrak{p} \nmid \ell$. For given $n \in \mathbb{Z}$, all of the ℓ^n-torsion points of E_1 are defined over some number field L'. If $\mathfrak{p}'$ is a prime lying over $\mathfrak{p}$, then

$$(\mathbb{Z}/\ell^n\mathbb{Z})^2 = E_1[\ell^n] = E_1(L')[\ell^n] \hookrightarrow \overline{E_1}(L'/\mathfrak{p}')[\ell^n] \subset \overline{E_1}[\ell^n] = (\mathbb{Z}/\ell^n\mathbb{Z})^2.$$

Hence, the reduction map induces an injection, and therefore an isomorphism $T_\ell(E_1) \to T_\ell(\overline{E_1})$.

Now, suppose $\phi \in \mathrm{Hom}(E_1, E_2)$ with $\overline{\phi} = 0$. Then, $\overline{\phi}(T_\ell(\overline{E_1})) = 0$, so $\phi(T_\ell(E_1)) = 0$. But then the kernel of ϕ is infinite, so $\phi = 0$, as desired, proving injectivity.

The proof that reduction preserves degrees is a similar computation involving the Weil pairing on the ℓ-adic Tate module. We refer the reader to Proposition 4.4 of [8] for details. □

Remark 3.22. When applied to the case $E = E_1 = E_2$, Proposition 3.21 shows that $\mathrm{End}(E) \hookrightarrow \mathrm{End}(\overline{E})$ is an injection under the reduction map. In particular, if $\mathrm{End}(E) = \mathcal{O}$ is an order $\mathcal{O}$ in an imaginary quadratic field K and $\overline{E}$ is ordinary, then implies that $\mathrm{End}(\overline{E})$, since $\mathrm{End}(\overline{E})$ is also an order in K.

In fact, the following result (Theorem 12 of Section 13 of [4]) allows us to determine $\mathrm{End}(\overline{E})$ precisely.

Theorem 3.23. *Let E be an elliptic curve over a number field L, with $\mathrm{End}(E) \cong \mathcal{O}$ an order in an imaginary quadratic field K. Let $\mathfrak{P}|p$ be a prime such that E has good reduction to $\overline{E}$ mod $\mathfrak{P}$. Then, $\overline{E}$ is ordinary if and only if p splits completely in K. In this case, let the conductor of $\mathcal{O}$ be $f_0 p^r$ where $p \nmid f_0$. Then,*

1. *$\mathrm{End}(\overline{E}) = \mathbb{Z} + f_0\mathcal{O}_K$ is the order of conductor f_0 in K.*
2. *If $r = 0$, then, $\phi \mapsto \overline{\phi}$ is an isomorphism of $\mathrm{End}(E)$ onto $\mathrm{End}(\overline{E})$.*

Proof. We follow the proof of 13.12 in [4], providing extra commentary on a few steps and refering the reader to [4] for a few technical points.

To prove that $\overline{E}$ is ordinary, it suffices to check that an isogenous curve is ordinary because only the zero isogeny can annihilate the p-adic Tate module $T_p(E)$. Moreover, the equivalence of categories from Theorem 3.11 makes it clear that we can find an isogenous curve E' with $\mathrm{End}(E') = \mathcal{O}_K$. We will see in Corollary 5.7 that the j-invariant of E' is algebraic, so we may assume that E' is defined over

some number field. Since isogenies descend under reduction, we have that $\overline{E}$ and $\overline{E'}$ are isogenous, as well. This reduction is quite useful, because it allows us to view elements of $\mathcal{O}_K$ as endomorphisms of E'. In particular, by the finiteness of the class group, if $p = \mathfrak{p}\mathfrak{p}'$, splits completely, some power of $\mathfrak{p}$ and $\mathfrak{p}'$ is principal. Without loss of generality, say $\mathfrak{p}^m = \alpha\mathcal{O}_K$ and $(\mathfrak{p}')^m = \alpha'\mathcal{O}_K$, where α and α' are conjugate. Then, $p^m = \alpha\alpha'$ is a product of endomorphisms. Then, the (reduction of) the endomorphism corresponding to α' is separable since E' has good reduction mod $\mathfrak{P}$ and $\deg(\alpha) = \deg(\alpha') = p^m$, so its reduction does as well and so $\overline{E'}$ has a p^m-torsion point, whence E' is singular.

The proof when p does not split completely requires a bit more machinery, including Hecke characters of ideles, so we omit the proof and refer the reader to [4].

For the proofs of (1) and (2), we follow the proof in 13.4 of [4], citing lemma 13.1 of [4], which states that for ℓ a prime not equal to the characteristic of the base field of E, the localization of $\mathrm{End}(E)$ at ℓ is determined uniquely by $\mathrm{End}(E) \otimes \mathbb{Q}$ and $T_\ell(E)$. Then, the localizations of $\mathrm{End}(E)$ and $\mathrm{End}(\overline{E})$ are the same at each $\ell \neq p$. In particular, this means that $\mathrm{End}(\overline{E}) = \mathbb{Z} + p^{r'} f_0 \mathcal{O}_K$ for some $r' \in \mathbb{Z}_{\geq 0}$. Quoting Theorem 3.19, which says that the conductor of $\mathrm{End}(\overline{E})$ is prime to p completes the proof.

□

Theorem 3.23 shows that when p splits completely in $\mathrm{End}(E) \otimes \mathbb{Q}$ and does not divide the conductor of $\mathrm{End}(E)$, then E descends to an elliptic curve over some finite field of characteristic p with $\mathrm{End}(\overline{E}) = \mathrm{End}(E)$.

We now state a sort of converse, known as the Deuring Lifting Theorem (see Theorem 13.14 of [4]).

Theorem 3.24. *Let $\tilde{E}/F$ be an elliptic curve over a field F of characteristic p and let $\tilde{\phi}$ be an endomorphism of $\tilde{E}$. Then, there is an elliptic curve E defined over a number field L, an endomorphism ϕ and a prime $\mathfrak{p}$ of L lying over p such that E has good reduction at $\mathfrak{p}$, $\tilde{E} \cong \overline{E}$, and $\overline{\phi}$ corresponds to $\tilde{\phi}$.*

Proof. We very briefly outline the proof in the case where $\tilde{A}$ is ordinary, referring the reader to [4] for full details. Shifting by a multiplication-by-n map and factoring out any integer multiples, one may assume that the kernel of $\tilde{\phi}$ is cyclic. As we shall see in Section 5.2 if we adjoin a variable j-invariant to $\mathbb{Q}$, the curves related by a cyclic isogeny of order n are all defined over some finite extension L' of $\mathbb{Q}(j)$ and in fact have j-invariants that are integral over $\mathbb{Z}[j]$. If the image curve of the isogeny ϕ corresponding to $\tilde{\phi}$ has j-invariant j', then reducing ϕ modulo an appropriately chosen prime containing $j - j'$ makes $\overline{\phi}$ an endomorphism of an elliptic curve defined over some number field, and reducing further, again by an appropriate prime, gives $\overline{\overline{\phi}} = \tilde{\phi}$, up to isomorphism. Verifying these claims requires a bit of dimension theory. This proves the theorem for curves with $\mathrm{Aut}(\tilde{E}) = \{\pm 1\}$, which is the case for ordinary elliptic curves. □

With the Deuring Lifting Theorem in hand, we can easily prove the following corollary.

Corollary 3.25. *Let $\tilde{E}/F$ be an elliptic curve over a field F of characteristic p with $\mathrm{End}(\tilde{E}) = \mathcal{O}$ where $\mathcal{O} = \mathbb{Z} + f\mathcal{O}_K$ where $p \nmid f$ is an order in the imaginary*

quadratic number field K. Also suppose that p splits completely in K. Then, there is an elliptic curve E defined over a number field L, and a prime $\mathfrak{p}$ of L lying over p such that E has good reduction at $\mathfrak{p}$, $\tilde{E} \cong \overline{E}$, and $\mathrm{End}(E) = \mathcal{O}$.

Proof. Choose $\tilde{\phi} \in \mathrm{End}(\tilde{E})$ such that $\mathbb{Z}[\tilde{\phi}] = \mathcal{O}$ and apply Theorem 3.24. Then, Theorem 3.23 implies that $\mathrm{End}(E) \cong \mathrm{End}(\tilde{E}) = \mathcal{O}$, as desired. □

We now state prove an important result, which underlies the efficient implentation of the complex multiplication method for constructing curves with a given number of points over a given finite field.

Proposition 3.26. *Let $\mathcal{ELL}_L(\mathcal{O}_K)$ denote the set of isomorphism classes of elliptic curves E over L with* $\mathrm{End}(E) \cong \mathcal{O}_K$. *Then, for any prime p that splits completely in K,*

$$|\mathcal{ELL}_{\overline{\mathbb{Q}}}(\mathcal{O}_K)| = |\mathcal{ELL}_{\overline{\mathbb{F}_p}}(\mathcal{O}_K)|.$$

The same holds replacing $\mathcal{O}_K$ with any order of K with conductor prime to p.

Proof. First, we claim that

$$|\mathcal{ELL}_{\overline{\mathbb{Q}}}(\mathcal{O}_K)| \leq |\mathcal{ELL}_{\overline{\mathbb{F}_p}}(\mathcal{O}_K)|.$$

We note first of all that $\mathcal{ELL}_{\overline{\mathbb{Q}}}(\mathcal{O}_K)$ is finite and so we can choose some number field L, a prime $\mathfrak{p}$ of L lying over p, and representatives E/L for the elements of $\mathcal{ELL}_{\overline{\mathbb{Q}}}(\mathcal{O}_K)$ such that every E has good reduction mod $\mathfrak{p}$.

The remainder of this direction of the proof follows the proof of Theorem 13.12 of [4] exactly, so we omit the proof and refer the reader there. The idea is that if E_1 and E_2 reduce to isomorphic curves, then we can construct a graph $E_1 \times E_1 \to E_1 \times E_2$ which reduces to the graph of an isomorphism. By reduction theory, the original graph must also be the graph of an isomorphism, so reduction mod $\mathfrak{p}$ is injective and so

$$|\mathcal{ELL}_{\overline{\mathbb{Q}}}(\mathcal{O}_K)| \leq |\mathcal{ELL}_{\overline{\mathbb{F}_p}}(\mathcal{O}_K)|.$$

To show that

$$|\mathcal{ELL}_{\overline{\mathbb{Q}}}(\mathcal{O}_K)| \geq |\mathcal{ELL}_{\overline{\mathbb{F}_p}}(\mathcal{O}_K)|,$$

we note that by slightly modifying the proof of Theorem 3.24 (and applying this as in Corollary 3.25,) given any finite subset of $\mathcal{E} \subset \mathcal{ELL}_{\mathbb{F}_p}(\mathcal{O}_K)$, we can lift each $\tilde{E} \in \mathcal{E}$ to the same field and prime. If $\tilde{E}_1$ and $\tilde{E}_2$ lift to isomorphic curves, the isomorphism descends and so $\tilde{E}_1 \cong \tilde{E}_2$. Hence, $|\mathcal{ELL}_{\overline{\mathbb{Q}}}(\mathcal{O}_K)| \geq \mathcal{E}$, and so

$$|\mathcal{ELL}_{\overline{\mathbb{Q}}}(\mathcal{O}_K)| \geq |\mathcal{ELL}_{\overline{\mathbb{F}_p}}(\mathcal{O}_K)|,$$

as desired. □

When we discuss the Hilbert Class polynomial in Section 5.3, we will see that it is the polynomial with the j-invariants of the elliptic curves in $\mathcal{ELL}_{\overline{\mathbb{Q}}}(\mathcal{O}_K)$ as roots is defined over $\mathbb{Z}$. Then, Proposition 3.26 shows that the roots of this polynomial in $\overline{\mathbb{F}_p}$ are exactly the j-invariants of the curves in $\mathcal{ELL}_{\overline{\mathbb{F}_p}}(\mathcal{O}_K)$. This will have the important consequence that we can compute the Hilbert Class polynomial by computing the j-invariants of elliptic curves with complex multiplication by $\mathcal{O}_K$ over p for small primes using the Chinese remainder theorem and bounds on the size of the coefficients. This technique makes the CM Method feasible for primes of cryptographic size and is discussed further in section 6.

4. A Quick Recap of Class Field Theory

The goal of (abelian) class field theory is to classify the abelian extensions of a number field K in terms of data innate to the field. In particular, we will relate the ideal class groups of K to the Galois theory of these extensions.

It would be foolhardy to attempt to provide a comprehensive treatment of this massive subject in a paper as short as this. Indeed, a through exposition on class field theory can and often does fill whole books. Instead, in an effort to be self-contained, we follow the approach in [1, 4, 8] and content ourselves to provide a few relevant definitions and major results almost entirely without proof. In an effort to keep the exposition as elementary as possible, we elect not to describe the idele theoretic formulation of class field theory. To simplify exposition further, we will often restrict to the case of quadratic imaginary fields, which are all we need for the remaining applications in this paper. For the reader interested in a more thorough treatment of Class Field Theory, there are many excellent resources available, including lecture notes by Milne and books by Lang, Tate, and Neukirch.

Following chapters 5 and 8 of [1], we begin our treatment of class field theory by defining the Artin map.

Proposition 4.1. *Suppose L/K is a Galois extension of number fields and $\mathfrak{p}$ is a prime of $\mathcal{O}_K$ which is unramified in L. Then, for any prime $\mathfrak{P}$ lying over $\mathfrak{p}$, there is a unique element $\sigma \in \mathrm{Gal}(L/K)$ such that*

$$\sigma(\alpha) \equiv \alpha^{N(\mathfrak{p})} \pmod{\mathfrak{P}}$$

for all $\alpha \in \mathcal{O}_L$, where $N(\mathfrak{p}) = |\mathcal{O}_K/\mathfrak{p}|$ is the norm of $\mathfrak{p}$.

The proof is a standard exercise in reduction theory since the Frobenius automorphism generates the Galois group of a finite extension of finite fields.

Definition 4.2. In the setup of Proposition 4.1, let the Artin symbol $\left(\frac{L/K}{\mathfrak{P}}\right)$ denote the unique $\sigma \in \mathrm{Gal}(L/K)$ such that

$$\left(\frac{L/K}{\mathfrak{P}}\right)(\alpha) \equiv \alpha^{N(\mathfrak{p})} \pmod{\mathfrak{P}}$$

From the uniqueness assertion in Definition 4.2, it follows almost immediately that for any $\sigma \in \mathrm{Gal}(L/K)$,

$$\left(\frac{L/K}{\sigma(\mathfrak{P})}\right) = \sigma \left(\frac{L/K}{\sigma(\mathfrak{P})}\right) \sigma^{-1}.$$

Now, the action of $\mathrm{Gal}(L/K)$ on the set of primes $\mathfrak{P}$ lying over a given prime $\mathfrak{p}$ is transitive, so if L/K is an abelian extension, the Artin symbol is defined uniquely by the prime $\mathfrak{p}$ and so we can write the Artin symbol as $\left(\frac{L/K}{\mathfrak{p}}\right)$.

Remark 4.3. In what proceeds, we assume that K is a totally imaginary field. If K has real embeddings, the use of $\mathcal{O}_K$-ideals below must be replaced by the notion of a modulus, which also captures ramification at the infinite places. Since we are considering behaviour at all of the places of K, it might seem that this is a natural application for the ideles. Indeed, Class Field Theory can naturally be expressed in terms of ideles. We will comment on some of the advantages of that approach later on. For now, however, we stick to this more concrete setup, since we have a concrete application in mind.

We now fix a bit of notation, following 8.A of [1]:

Definition 4.4. Suppose K is a totally imaginary field and $\mathfrak{a} \subset \mathcal{O}_K$ is an ideal.

(1) Let $I_K(\mathfrak{a})$ be the group of fractional ideals prime to $\mathfrak{a}$.
(2) Let $P_{K,1}(\mathfrak{a})$ be the subgroup of $I_K(\mathfrak{a})$ generated by the principal ideals $\alpha\mathcal{O}_K$ with $\alpha \equiv 1 \pmod{\mathfrak{a}}$.
(3) Call H a *congruence subgroup* for $\mathfrak{a}$ if $P_{K,1}(\mathfrak{a}) \subset H \subset I_K(\mathfrak{a})$.
(4) For a congruence subgroup H, call the quotient $I_K(\mathfrak{a})/H$ a *generalized ideal class group.*

As is pointed out in 8.A of [1], it is immediately clear from Theorem 2.15 that when K is an imaginary quadratic field,

$$P_{K,1} \subset P_{K,\mathbb{Z}}(f) \subset I_K(f) \subset I_K(f\mathcal{O}_K),$$

so the notion of a generalized ideal class group does indeed generalize the notion of an ideal class group that we have seen previously.

Remark 4.5. The notion of a congruence subgroup in the setup of class field theory may remind the reader of congruence subgroup in $\mathrm{SL}_2(\mathbb{Z})$ in the theory of modular forms. While this author was unable to determine a direct correspondence between the two notions, this parallel may serve as some motivation for why we might hope that some of the remarkable connections between complex multiplication and class field theory discussed in section 5.3 exist.

Remark 4.6. One drawback of our concrete approach to class field theory is that congruence subgroups may be subgroups of different groups, which is somewhat awkward. As is discussed in Section 8.C of [1], in the idelic formulation of class field theory, the role of congruence subgroups is played by closed subgroups of finite index in the idele class group, which is somewhat more satisfying.

With these definition in hand, we are ready to defined the Artin map, which connects (generalized) ideal class groups and Galois theory.

Definition 4.7. Suppose L/K is an abelian extension of totally imaginary number fields and $\mathfrak{a} \subset \mathcal{O}_K$ is an ideal such that if the prime $\mathfrak{p}$ ramifies in L, then $\mathfrak{p}|\mathfrak{a}$. Then, the Artin symbol $\left(\frac{L/K}{\mathfrak{p}}\right)$ is defined for all $\mathfrak{p} \in I_K(\mathfrak{a})$. This extends by multiplicativity to a homomorphism

$$\Phi_{\mathfrak{a}} = \Phi_{L/K,\mathfrak{a}} : I_K(\mathfrak{a}) \to \mathrm{Gal}(L/K),$$

called the *Artin map* for L/K and $\mathfrak{a}$.

We now recall three main theorems of Class Field Theory (essentially using the statements from [1]) before discussing a few applications:

Theorem 4.8 (The Artin Reciprocity Theorem)**.** *Let L/K and $\mathfrak{a}$ be as in Definition 4.7. Then,*

(1) *The Artin map $\Phi_{\mathfrak{a}}$ is surjective.*
(2) *There exists some ideal $\mathfrak{b}$ with $\mathfrak{a}|\mathfrak{b}$ such that $\ker(\Phi_{\mathfrak{b}})$ is a congruence subgroup for $\mathfrak{b}$.*

In particular, $\mathrm{Gal}(L/K)$ *is a generalized ideal class group for $\mathfrak{b}$.*

Theorem 4.9 (The Conductor Theorem)**.** *Let L/K be an abelian extension of totally imaginary number fields. Then, there is some $\mathcal{O}_K$-ideal $\mathfrak{a} = \mathfrak{a}(L/K)$ such that*

(1) *A prime* $\mathfrak{p}$ *of* $\mathcal{O}_K$ *ramifies in* L *if and only if* $\mathfrak{p}|\mathfrak{a}$.
(2) *If* $\mathfrak{b} \subset \mathcal{O}_K$ *is divisible by all primes of* $\mathcal{O}_K$ *that ramify in* L*, then* $\ker(\Phi_{\mathfrak{b}})$ *is a congruence subgroup for* $\mathfrak{b}$ *if and only if* $\mathfrak{a}|\mathfrak{b}$.

Theorem 4.10 (The Existence Theorem)**.** *If* K *is a totally imaginary number field and* $\mathfrak{a} \subset \mathcal{O}_K$ *is an ideal with a congruence subgroup* H*, then there is a unique abelian extension* $L \supset K$ *such that*

(1) *If* $\mathfrak{p} \subset \mathcal{O}_K$ *ramifies in* L*, then* $\mathfrak{p}|\mathfrak{a}$.
(2) H *is the kernel of the Artin map*

$$\Phi_{\mathfrak{a}} : I_K(\mathfrak{a}) \to \mathrm{Gal}(L/K).$$

Essentially, the Artin Reciprocity Theorem and the Existence Theorem say that the Galois groups of abelian extensions are exactly the generalized ideal class groups of K. The conductor theorem says that given L, the set of ideals for which $\mathrm{Gal}(L/K)$ is a generalized ideal class group has a maximal element with respect to inclusion.

The existence theorem (4.10) is particularly useful for our purposes as it allows us to generalize the ideal class group by associating a unique (abelian) field extension L to each order $\mathcal{O}$ of an imaginary quadratic field K.

Definition 4.11. Given an imaginary quadratic field K and an order $\mathcal{O} = \mathbb{Z} + f\mathcal{O}_K$, let $\mathfrak{a} = f\mathcal{O}_K$ and $H = P_{K,\mathbb{Z}}(f)$. Then, the unique extension $L = L_{\mathcal{O}}$ of K from Theorem 4.10 is called the *ring class field* of $\mathcal{O}$.

Remark 4.12. It is clear from Theorems 4.10 and 2.15 that $C(\mathcal{O}) \cong \mathrm{Gal}(L_{\mathcal{O}}/K)$, allowing us to realize the ideal class group of $\mathcal{O}$ as the Galois group of an appropriate extension.

Remark 4.13. If $\mathcal{O} = \mathcal{O}_K$ is the ring of integers in K, then setting $\mathfrak{a} = \mathcal{O}_K$ and $H = P_K = P_{K,1}(\mathcal{O}_K)$, the corresponding L is called the *Hilbert Class Field* of K. It is an easy consequence of the Conductor Theorem (4.9) that L is the maximal unramified abelian extension of K.

The previous remarks show that the ring class field of an order $\mathcal{O}$ provides one generalization of the Hilbert Class field. For completeness, we define one other generalization.

Definition 4.14. For $\mathfrak{a} \subset \mathcal{O}_K$ an ideal, let L be the unique field extension corresponding to the "minimal" congruence subgroup $P_{K,1}(\mathfrak{a})$ from the Existence Theorem 4.10. Then, L is called the *ray class field* of $\mathfrak{a}$.

The last bit of class field theory we will need are a few consequences of the Čebotarev Density Theorem. Roughly speaking, for a Galois extension L/K of number fields and a conjugacy class $\langle\sigma\rangle$ in $\mathrm{Gal}(L/K)$, the Density Theorem gives an measurement of the size of the set $\mathcal{S}$ of primes $\mathfrak{p}$ of $\mathcal{O}_K$ such that

(1) $\mathfrak{p}$ is unramified in L.
(2) $\left(\frac{L/K}{\mathfrak{p}}\right) = \langle\sigma\rangle$.

In particular, it shows that $\mathcal{S}$ has "Dirichlet density" $\frac{|\langle\sigma\rangle|}{[L:K]} > 0$ which implies that $\mathcal{S}$ is infinite.

This result is particularly easy to interpret when $|\langle\sigma\rangle| = 1$, for example when L/K is abelian or $\sigma = 1$.

Corollary 4.15. *If L/K is an abelian extension and $\sigma \in \mathrm{Gal}(L/K)$, there are infinitely many prime ideals $\mathfrak{p} \subset \mathcal{O}_K$ such that $\mathfrak{p}$ is unramified in L and $\left(\frac{L/K}{\mathfrak{p}}\right) = \sigma$.*

Corollary 4.16. *If L/K is a Galois extension of number fields, then infinitely many prime $\mathcal{O}_K$ ideals split completely in L.*

Proof. The Čebotarev Density Theorem says that there are infinitely many primes $\mathfrak{p}$ that are unramified in L such that $\left(\frac{L/K}{\mathfrak{p}}\right) = 1$. Now, if $\left(\frac{L/K}{\mathfrak{p}}\right) = 1$, then for any $\mathfrak{P}$ lying over $\mathfrak{p}$, the residue fields $L_{\mathfrak{P}}$ and $K_{\mathfrak{p}}$ are equal, so the inertial degree of $\mathfrak{P}/\mathfrak{p}$ is 1. Since $\mathfrak{p}$ is unramified, this implies that $\mathfrak{p}$ splits completely in L. □

In fact, with a bit more work, one can use the Čebotarev Density Theorem to show that a Galois extension L/K is completely determined by the primes of $\mathcal{O}_K$ that split completely.

We have the following proposition (8.20 of [1])

Proposition 4.17. *Suppose M and L are algebraic extensions of K, at least one of which is Galois. Let $\mathcal{S}_{M/K}$ denote the sets of primes of K that are unramified in K such that there is some prime $\mathfrak{P}$ of M lying over $\mathfrak{p}$ such that the residue fields $M_{\mathfrak{P}}$ and $K_{\mathfrak{p}}$ are equal. Define $\mathcal{S}_{L/K}$ similarly.*

Then, $L \subset M$ if and only if there is some finite set Σ such that $S_{M/K} \subset S_{L/K} \cup \Sigma$.

In particular, $M = L$ if and only if there exists a finite set Σ with $S_{M/K} \cup \Sigma = S_{L/K} \cup \Sigma$.

Note that if M/K (resp. L/K) is Galois, $S_{M/K}$ (resp. $S_{L/K}$) is the set of primes of K that split completely in K (resp. L).

For a proof of Proposition 4.17, see Section 8.B of [1].

At first glance, the fact that Galois extensions are determined by the primes that split completely may seem shocking. As Cox points out, this result is intimately related to the fact that if $\mathfrak{a}$ is an $\mathcal{O}_K$ ideal divisible by all of the prime ideals that ramify in L, then as in the proof of Corollary 4.16, up to a finite set, the prime ideals in the kernel of $\Phi_{\mathfrak{a}}$ are exactly the prime ideals that split completely in L. Since the existence theorem relates $\ker(\Phi_{\mathfrak{a}})$ to Galois extensions of K, the results of Proposition 4.17 should perhaps not be so surprising as they may seem at first.

Combining the power of the Čebatarov Density Theorem with the our discussion of the Artin map, we are able to prove one more important result that will allow us to assume that we are working with ideals of prime norm when choosing representatives of the ideal class group later on.

Proposition 4.18. *Let $\mathcal{O}$ be an order in the imaginary quadratic field K and let $\alpha \in C(\mathcal{O})$ be an arbitrary ideal class. Then, the set*

$$\{p \text{ prime } : \exists \mathfrak{p} \subset \mathcal{O}, N(\mathfrak{p}) = p, [\mathfrak{p}] = \alpha\}$$

of primes that are the norm of a $\mathcal{O}$-ideal in the ideal class $\mathfrak{a}$ is infinite.

Proof. Throughout the proof, we shall assume that p is relatively prime to the conductor of $\mathcal{O}$, which excludes only a finite set of primes. Then, when we consider ideals of order prime to p, we abuse notation slightly by using $\mathfrak{p}$ to refer both to the $\mathcal{O}$-ideal and the corresponding $\mathcal{O}_K$-ideal $\mathfrak{p}\mathcal{O}_K$ of the same norm from the isomorphism discussed in the proof of Theorem 2.15. It should be clear from context which meaning is intended.

Let L be the ring class field of K. Given any $\mathfrak{p}$ prime to the conductor of $\mathcal{O}$, the Artin symbol $\left(\frac{L/K}{\mathfrak{p}}\right) = \sigma_\alpha \in \mathrm{Gal}(L/K)$ depends only on the class α.

Now, we can also interpret σ_α as an element of $\mathrm{Gal}(L/\mathbb{Q})$. If $\mathfrak{p}$ has prime norm p, then the local fields of K at $\mathfrak{p}$ and $\mathbb{Q}$ at p are equal. Hence, if the prime $\mathfrak{P} \subset \mathcal{O}_L$ lies over $\mathfrak{p}$, taking some $\left(\frac{L/K}{\mathfrak{P}}\right) = \left(\frac{L/\mathbb{Q}}{\mathfrak{P}}\right)$ and so $\left(\frac{L/\mathbb{Q}}{p}\right)$ is the conjugacy class of σ_α in $\mathrm{Gal}(L/\mathbb{Q})$.

Similarly, if p is prime to the conductor of $\mathcal{O}$ and $\left(\frac{L/\mathbb{Q}}{p}\right)$ contains σ_α, there is some prime $\mathfrak{P}$ of $\mathcal{O}_L$ with $\left(\frac{L/\mathbb{Q}}{\mathfrak{P}}\right) = \sigma_\alpha$. Hence, $\sigma_\alpha = \left(\frac{L/K}{\mathfrak{P}}\right) = \left(\frac{L/K}{\mathfrak{P}\cap\mathcal{O}_K}\right)$. Thus, the local field at $\mathfrak{p} = \mathfrak{P} \cap \mathcal{O}_K$ must be $\mathbb{F}_p$ and so $N(\mathfrak{p}) = p$.

Thus, up to a finite set,

$$\{p \text{ prime } : \exists \mathfrak{p} \subset \mathcal{O}, N(\mathfrak{p}) = p, [\mathfrak{p}] = \alpha\}$$

is the set of primes such that $\left(\frac{L/\mathbb{Q}}{p}\right)$ is the conjugacy class of σ_α. This set is infinite by the Čebetarov Density Theorem, which completes the proof. □

This completes our review of Class Field Theory.

5. The Modular Equation and Hilbert Class Polynomials

While we were able to characterize all curves with a given endomorphism ring in Section 3, our treatment failed to answer one important question. How does one compute the equation of a curve with complex multiplication by a given order? As we shall see in Section 5.5, being able to find such a curve is a key step in one powerful approach to finding a curve with a given number of points over a particular finite field.

In this section, we will discuss the problem over the complex numbers. For the remainder of this section, all curves are assumed to be over $\mathbb{C}$ unless stated otherwise. From proposition 3.17, if $\mathrm{End}(E) \cong \mathcal{O}$, an order in the imaginary quadratic field K, then we know that E corresponds to some proper (fractional) ideal $\mathfrak{a}$ of $\mathcal{O}$. While we can use the Eisenstein polynomials $G_4(\mathfrak{a})$ and $G_6(\mathfrak{a})$ to find E, these are infinite series and are often difficult to compute. Even worse, they are defined analytically, so there is no obvious way to connect them to the finite field case.

While the results of this section may not give a completely explicit description of an elliptic curve with $\mathrm{End}(E) = \mathcal{O}$, we will be able to convert our analytic description into an algebraic one, which we will be able to relate to the finite field case. Along the way, we will discuss some remarkable connections with Class Field Theory - the so-called Main Theorems of Complex Multiplication. It will turn out that CM elliptic curves are the key to answering Hilbert's 12th problem, which asks which algebraic numbers are necessary to generate all abelian extensions of a number field K, in the case where K is an imaginary quadratic field.

Our treatment will assume several basic results from the theory of modular forms.

5.1. Introducing the j-Function.

While it might seem natural to study Eisenstein series, it will be more convenient to work with a single object, the j-invariant of the corresponding elliptic curves, which we know parameterizes elliptic curves up to isomorphism. We use $\mathbb{H}$ to denote the upper half-plane of $\mathbb{C}$. Note that given any lattice, we may always take generators in $\mathbb{H}$.

Definition 5.1. The j-function is the meromorphic function $j : \mathbb{H} \to \mathbb{C}$ defined by:

$$j(\tau) = j([1,\tau]) = j(E_{[1,\tau]}) = \frac{1728 g_2([1,\tau])^3}{g_2([1,\tau])^3 - 27 g_3([1,\tau])^2}$$

It is easy to see that for any $\alpha = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z})$,

$$j(\alpha\tau) = j\left(\left[1, \frac{a\tau+b}{c\tau+d}\right]\right) = j([c\tau+d, a\tau+b]) = j([1,\tau]) = j(\tau),$$

so j is $\mathrm{SL}_2(\mathbb{Z})$-invariant and is therefore a modular function for $\Gamma(1) = \mathrm{SL}_2(\mathbb{Z})$. Then, writing $q = e^{2\pi i \tau}$, a basic computation from the theory of modular forms shows that j has a pole of order 1 at infinity and has q-expansion

$$j(\tau) = \frac{1}{q} + \sum_{n=0}^{\infty} c_n q^n = \frac{1}{q} + 744 + 196884q + \cdots, \tag{5.1}$$

where $c_n \in \mathbb{Z}$ for all $n \geq 0$ and we cite the first two values from [1].

In fact, every modular function for $\Gamma(1) = \mathrm{SL}_2(\mathbb{Z})$ is a rational function in j. This follows quickly since we can multiply by a suitable polynomial in $j(\tau)$ to remove poles away from infinity and subtract a polynomial in $j(\tau)$ to get a (level 1, weight 1) modular form that vanishes at infinity (see the proof of part (3) of 5.4.)

Much more could be said about the connections between the j-function and the theory of modular forms. We restrain ourselves from addressing these topics and continue working towards the main theorems of complex multiplication.

5.2. **The Modular Equation $\Phi_n(X,Y)$ and Some Properties.** In order to motivate our construction of the modular equation $\Phi_n(X,Y)$, we first make some quick comments about endomorphisms of elliptic curves with complex multiplication.

Note first of all that if $E/\mathbb{C}$ is an elliptic curve, then the kernel of the multiplication-by-n map is $E[n] \cong (\mathbb{Z}/n\mathbb{Z})^2$. In particular, it *is not* cyclic. However, if $\phi \in \mathrm{End}(E) \smallsetminus \mathbb{Z}$ is a *primitive* endomorphism (i.e. it does not factor over $[n]$ for any $n \in \mathbb{Z}$ except $n = \pm 1$), then $E[\phi]$ *is* cyclic of order $\deg(\phi)$. Hence, $E/\mathbb{C}$ has complex multiplication if and only if there exists $\phi \in \mathrm{End}(E)$ with cyclic kernel. This suggests that it may be wise to study endomorphisms with kernel isomorphic to $\mathbb{Z}/n\mathbb{Z}$ for some fixed n.

In fact, we take this idea one step further. Following the general mathematical principle that the additional degree of freedom provided by introducing an extra variable/object/dimension can often make more general cases easier to understand, we will study cyclic n-isogenies (i.e. isogenies with kernel $\mathbb{Z}/n\mathbb{Z}$). The modular equation $\Phi_n(X,Y)$ will parameterize (the j-invariants of) elliptic curves related by a cyclic n-isogeny. Our treatment is inspired by the Section 11.C of [1] and the excellent exposition in Section II.6 of [8].

By the equivalence of categories from 3.11, we can see that there is a cyclic n-isogeny from E_{Λ_1} to E_{Λ_2} if and only if Λ_2 is (up to scaling) a cyclic sublattice of index n in Λ_1, i.e. a sublattice with $\Lambda_1/\Lambda_2 \cong \mathbb{Z}/n\mathbb{Z}$.

It is easy to see that if

$$\Lambda_1 = [\omega_1, \omega_2] \text{ and } \Lambda_2 = [a\omega_1 + b\omega_2, c\omega_1, d\omega_2]$$

for $a, b, c, d \in \mathbb{Z}$, then Λ_2 is a cyclic sublattice if and only if $\gcd(a, b, c, d) = 1$, and has index $\det \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ in Λ_1.

Of course, this description is somewhat redundant. If $\alpha = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ and $\alpha' = \begin{pmatrix} a' & b' \\ c' & d' \end{pmatrix}$ with $\alpha = \gamma\alpha'$ for some $\gamma \in \mathrm{SL}_2(\mathbb{Z})$ correspond to lattices Λ and Λ', then $\Lambda = \Lambda'$. Thus, we consider orbit representatives for the action of $\mathrm{SL}_2(\mathbb{Z})$ on the set of matrices

$$M_n := \left\{ \alpha = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{Mat}_2(\mathbb{Z}) : \gcd(a, b, c, d) = 1 \text{ and } \det(\alpha) = n \right\}. \tag{5.2}$$

For computational purposes, the following set is particularly convenient.

Proposition 5.2. *Let M_n be the set of matrices defined in* (5.2)*. Then, the set*

$$C_n := \left\{ \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \in \mathrm{Mat}_2(\mathbb{Z}) : \gcd(a, b, d) = 1, ad = 1, 0 \leq b < d \right\} \tag{5.3}$$

is a full set of orbit representatives for the action of $\mathrm{SL}_2(\mathbb{Z})$ *by left multiplication on M_n.*

The proof is an easy exercise.

Next, we describe these sublattices when $\Lambda = [1, \tau]$ for some τ.

Remark 5.3. It $\Lambda_1 = [1, \tau], \Lambda_2 = [a+b\tau, c+d\tau]$, and $\alpha = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$, then, rescaling by $c+d\tau$, we have that $\Lambda_2 \sim [1, \alpha\tau]$, where $\alpha\tau = \frac{a+b\tau}{c+d\tau}$ is the usual action of $\mathrm{GL}_2^+(\mathbb{Z})$ on $\mathbb{H}$.

With this in hand, we are ready to define the modular equation.

Proposition 5.4. *For a matrix $\alpha \in \mathrm{GL}_2^+(\mathbb{Z})$, let $j \circ \alpha : \mathbb{H} \to \mathbb{C}$ be defined by $j \circ \alpha(\tau) = j(\alpha\tau)$. With this notation, set*

$$\Phi_n(X) = \prod_{\alpha \in C_n} (X - j \circ \alpha) = \sum_k s_k X^k. \tag{5.4}$$

Then, $\Phi_n(X) \in \mathbb{Z}[j](X)$.

Remark 5.5. The polynomial $\Phi_n(X, Y) \in \mathbb{Z}[X, Y]$ such that $\Phi_n(X) = \Phi_n(X, j)$ is called the *modular equation of order n*. By construction $\Phi_n(j_1, j_2) = 0$ if and only if there is a cyclic isogeny of order n from the elliptic curve with j-invariant j_1 to the elliptic curve with j-invariant j_2.

Proof. We follow the proof from Section II.6 of [8]. We prove the proposition with the following three steps.

(1) For each k, s_k is a modular function for $\mathrm{SL}_2(\mathbb{Z})$.
(2) The q-expansion of s_k has coefficients in $\mathbb{Z}$.
(3) For each k, $s_k \in \mathbb{Z}[j]$.

Claims (1) and (2) are essentially statements about the invariance of the set $\{j \circ \alpha : \alpha \in C_m\}$ under two operations - the action of $\mathrm{SL}_2(\mathbb{Z})$ on $\mathbb{H}$ and the action of $\mathrm{Gal}(\mathbb{Q}(\zeta_m), \mathbb{Q})$ on the q-expansions.

For (1), note that for every $\alpha \in C_n$ and $\gamma \in \mathrm{SL}_2(\mathbb{Z})$, there is a unique $\gamma_\alpha \in \mathrm{SL}_2(\mathbb{Z})$ such that $\gamma_\alpha \alpha \gamma \in C_n$ and so γ induces a map $C_n \to C_n$, $\alpha \mapsto \gamma_\alpha \alpha \gamma$. Since the

s_k are symmetric polynomials in the $j \circ \alpha$, it suffices to check that this map is a bijection. Indeed, it is an injective map between finite sets, since if $\gamma_\alpha \alpha \gamma = \gamma_\beta \beta \gamma$, then $\alpha = (\gamma_\alpha^{-1}\gamma_\beta)\beta$. This implies $\mathrm{SL}_2(\mathbb{Z})$-invariance.

By setting $q_m = e^{2\pi i \tau/m}$ and looking at the q_m-expansions of the $j \circ \alpha$, we see that each s_k has a pole of finite index at ∞ and so s_k is a modular function for $\mathrm{SL}_2(\mathbb{Z})$.

For (2), let $\zeta_m = e^{2\pi i/m}$ and let $\sigma \in \mathrm{Gal}(\mathbb{Q}(\zeta_m) : \zeta_m)$ be the automorphism that takes ζ_m to ζ_m^t. Then, from (5.1), we see that for $\alpha = \begin{pmatrix} a & b \\ 0 & d \end{pmatrix}$, the q_m expansions of $j \circ \alpha$ and $\sigma(j \circ \alpha)$ are

$$j \circ \alpha(\tau) = \sum_{n=-1}^{\infty} c_n q_m^{a^2 n} \zeta_m^{abn} \tag{5.5}$$

and so

$$\sigma(j \circ \alpha(\tau)) = \sum_{n=-1}^{\infty} c_n q_m^{a^2 n} \zeta_m^{abtn} \tag{5.6}$$

From (5.5) and (5.6), it is clear that the q-expansions of

$$\sigma\left(j \circ \begin{pmatrix} a & b \\ 0 & d \end{pmatrix}\right) \text{ and } j \circ \begin{pmatrix} a & bt \pmod d \\ 0 & d \end{pmatrix},$$

are equal, so the coefficients of the q-expansion of s_k are σ-invariant and therefore lie in $\mathbb{Z}$.

For (3), note first that s_k is holomorphic on $\mathbb{H}$. If s_k has q-expansion

$$s_k(\tau) = \sum_{t=-N}^{\infty} d_t q^t$$

then $s_k - d_t j^t$ is a modular function with a pole of lower order at ∞. Continuing in this manner inductively, there is some polynomial f such that $s_k - f(j)$ is a holomorphic modular function for $\mathbb{H}$ that vanishes at infinity. Hence, $s_k = f(j)$. Since at each stage, all of the coefficients d_t are integers, $s_k \in \mathbb{Z}[j]$, as claimed. □

Proposition 5.6. *The modular equation $\Phi_n(X, Y)$ has the following properties:*

1. *$\Phi_n(X, Y)$ is irreducible.*
2. *If p is prime, then $\Phi_p(X, Y) \equiv (X^p - Y)(X - Y^p) \mod p\mathbb{Z}[X, Y]$. (This is often called the* Kronecker congruence*.)*
3. *If n is not a square, then $\Phi_n(X, X)$ has leading coefficient ± 1.*
4. *$\Phi_n(X, Y) = \Phi_n(Y, X)$.*

Proof. As in the proof of Proposition 5.4, (1) will follow from the action of $\mathrm{SL}_2(\mathbb{Z})$ on C_n, while (2) and (3) come from the q-expansions (together with some properties of $\mathbb{Q}(\zeta_m)$).

For (1), we consider the field extensions $\mathbb{C}(j)$ and $\mathbb{C}(j, j \circ \alpha)$. By Proposition 5.4 that for all $\alpha \in C_n$,

$$[\mathbb{C}(j, j \circ \alpha) : \mathbb{C}(j)] \leq |C_n|.$$

Now, if $C^\infty(\mathbb{H})$ is the field of meromorphic functions on $\mathbb{H}$, we know that for all $\gamma \in \mathrm{SL}_2(\mathbb{Z})$, $f \mapsto f \circ \gamma$ is an automorphism of $C^\infty(\mathbb{H})$ that fixes $\mathbb{C}(j)$. So, recalling the map $C_n \to C_n$, $\alpha \mapsto \gamma_\alpha \alpha \gamma$, from the proof of Proposition 5.4, we see that

$j \circ \alpha$ and $j \circ (\gamma_\alpha \alpha \gamma)$ are conjugate over $\mathbb{C}(j)$. So, it suffices to check that for every $\alpha \in C_n$, there exists $\gamma_1, \gamma_2 \in \mathrm{SL}_2(\mathbb{Z})$ such that $\begin{pmatrix} n & 0 \\ 0 & 1 \end{pmatrix} \gamma_1 = \gamma_2 \alpha$.

This is a straight-forward number theory computation. Given $a, b, d \in \mathbb{Z}$ with $(a, b, d) = 1$, we can choose g, h with $gb + hd = \gcd(b, d)$. Then, $\gcd\left(g, \frac{d}{\gcd(b,d)}\right) = 1$, possibly adjusting g by a multiple of $\frac{d}{\gcd(b,d)}$, we may assume that $\gcd(g, \gcd(b, d)) = 1$. Then, we can choose $u, y \in \mathbb{Z}$ such that $g|u$. Set $x = u/g$. Then, $xga + ygb + yhd = 1$ and

$$\begin{pmatrix} m & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} y & -x \\ ag & gb + hd \end{pmatrix} = \begin{pmatrix} my & -mx \\ ag & gb + hd \end{pmatrix} = \begin{pmatrix} dy & -ax - by \\ g & h \end{pmatrix} \begin{pmatrix} a & b \\ 0 & d \end{pmatrix}.$$

Hence, $[\mathbb{C}(j, j \circ \alpha) : \mathbb{C}(j)] \geq |C_n|$ and so $\Phi_n(X, j)$ is irreducible over $\mathbb{C}(j)$ and so $\Phi_n(X, Y)$ is irreducible.

For (2), looking at the q-expansions as in (5.5) shows that if $\alpha_0 = \begin{pmatrix} 1 & 0 \\ 0 & p \end{pmatrix}$,

$$\begin{aligned} j \circ \begin{pmatrix} 1 & i \\ 0 & p \end{pmatrix} &\equiv j \circ \sigma_0 \pmod{1 - \zeta_p^i}. \\ j \circ \begin{pmatrix} p & 0 \\ 0 & 1 \end{pmatrix} &\equiv j^p \pmod{p}, \\ j &\equiv (j \circ \alpha_0)^p \pmod{p}. \end{aligned}$$

Now, $(1 - \zeta_p)|p$, so in the q-expansions, we have

$$\begin{aligned} \Phi_p(X) &\equiv (X - j \circ \sigma_0)^p (X - j^p) \\ &\equiv (X^p - (j \circ \sigma_0)^p)(X - j^p) \\ &\equiv (X^p - j)(X - j^p) \pmod{1 - \zeta_p}. \end{aligned}$$

Since the coefficients of the q expansions of both $\Phi_p(X)$ and $(X^p - j)(X - j^p)$ are in $\mathbb{Z}$, this implies that the coefficients are divisible by p, whence

$$\Phi_p(X, Y) \equiv (X^p - Y)(X - Y^p) \mod p\mathbb{Z}[X, Y],$$

as claimed.

(3) This time, we look at the q_n-expansions. Since n is not a square, the "leading" coefficient of the q-expansion of $j - j \circ \begin{pmatrix} a & b \\ 0 & d \end{pmatrix}$ is either 1 or $-\zeta_n^{ab}$. Then the product is an integer of norm 1, whence the computation from part (3) of the proof of Proposition 5.4 shows that $\pm\Phi(X, Y)$ is monic.

(4) For this proof, we first note that if E_j denotes the curve with j-invariant j, then the roots of $\Phi_n(X, j)$ are exactly the j-invariants of curves related to E_j by a cyclic n-isogeny. Then, because the dual of a cyclic isogeny is a cyclic isogeny of the same degree, the roots of $\Phi_n(x_0, X)$ are the same as the roots of $\Phi_n(X, x_0)$. Hence, $\Phi_n(X, x_0) = a(x_0)\Phi_n(x_0, X)$ where $a(x_0)$ is some function of x_0 which must be a rational function so that the coefficients of X agree. But $\Phi_n(x_0, x_0) \neq 0$ so long as x_0 is the j-invariant of a curve without a cyclic n-isogeny, so $a(x_0) = 1$ away from a finite set, and hence, $a(x_0) = 1$ everywhere. □

Part (3) of Proposition 5.6 has the following important consequence.

Corollary 5.7. *Suppose that $E/\mathbb{C}$ is an elliptic curve with complex multiplication. Then, $j(E)$ is an algebraic integer.*

Proof. As we discussed previously, if E has complex multiplication, then E has some endomorphism with cyclic kernel. By Proposition 4.18 applied to the class of principal ideals, we can choose a principal prime $\alpha\mathcal{O}$ of $\mathcal{O} = \mathrm{End}_{\mathbb{C}}(E)$ isogeny with degree p, which corresponds to an endomorphism with kernel $\mathbb{Z}/p\mathbb{Z}$. Then, by Part (3) of Proposition 5.6, $j(E)$ is a root of $\Phi_p(X, X)$, a monic polynomial with integer coefficients, so $j(E)$ is an algebraic integer. □

Remark 5.8. The construction of Φ_n gives a very explicit proof that $j(E)$ is an algebraic integer. However, this approach does not generalize well to higher dimensional abelian varieties. Two alternate approaches that first prove that $j(E)$ is algebraic using Galois theory and then look at reduction modulo prime ideals are presented in [7].

Another unfortunate consequence of our construction of Φ_n is that it is not clear how Φ_n relates to elliptic curves in fields F of characteristic $p > 0$. In fact, so long as $p \nmid n$, Φ_n also parameterizes pairs of elliptic curves related by a cyclic n-isogeny over $\overline{F}$. Igusa proved this in [2] by developing a theory of modular forms in positive characteristic. Alternately, one can note that since the multiplication-by-n map is a rational function, the cyclic subgroups of index n are determined by algebraic conditions. Then, the Velu formulas give algebraic formulas for the j-invariants of curves related by cyclic n-isogenies. These formulas must give the same parameterization in all characteristics. Then, since the modular polynomials Φ_n have coefficients in $\mathbb{Z}$, the coefficients of the modular polynomial in characteristic p must be the reduction of Φ_n modulo $p\mathbb{Z}[X, Y]$.

5.3. **The Hilbert Class Polynomial.** So far, we have seen that if $\mathcal{O}$ is an order in the imaginary quadratic field K and $\mathfrak{a} \subset \mathcal{O}$ is a proper ideal, then, $j(E_{\mathfrak{a}})$ is a root of some $\Phi_p(X, X)$ and in particular is an algebraic integer. However, we know from the Kronecker congruence (part (2) of Proposition 5.6) and the fact that $\Phi_p(X, X) \in \mathbb{Z}[X]$ is monic that $\deg(\Phi_p(X, X)) = 2p$, so this will not typically be the minimal polynomial of $j(E_{\mathfrak{a}})$. To understand $j(E_{\mathfrak{a}})$ better, we make the following definition.

Definition 5.9. The *Hilbert Class Polynomial* of the order $\mathcal{O}$ in the imaginary quadratic field K is defined by

$$H_{\mathcal{O}}(X) = \prod_{[\mathfrak{a}] \in C(\mathcal{O})} (X - j(E_{\mathfrak{a}})).$$

In the remainder of this section we will show that $H_{\mathcal{O}}(X) \in \mathbb{Z}[X]$ and in fact that it is irreducible over K and therefore $\mathbb{Q}$. In the process, we will see that the $K(j(E_{\mathfrak{a}}))$ is the splitting field of $H_{\mathcal{O}}(X)$ over K and $\mathrm{Gal}(K(j(E_{\mathfrak{a}})), K) \cong C(\mathcal{O})$ by an isomorphism that respects the corresponding group actions on $\mathcal{ELL}_{\mathbb{C}}(\mathcal{O})$. Our treatment is mostly inspired by Lecture 22 of [9]. Most of these results can also be found in Section II of [8].

Proposition 5.10. *Let $H_{\mathcal{O}}(X)$ be the Hilbert class polynomial of the order $\mathcal{O}$ in the imaginary quadratic field K. Then, $H_{\mathcal{O}}(X) \in \mathbb{Z}[X]$.*

Proof. First, note that if $\sigma \in \mathrm{Gal}(\overline{\mathbb{Q}}, \mathbb{Q})$, and E is any elliptic curve, then $\sigma : E \to E^{\sigma}$ induces an isomorphism $\mathrm{End}_{\mathbb{C}}(E) \cong \mathrm{End}_{\mathbb{C}}(E^{\sigma})$ with inverse σ^{-1} by acting on the coefficients of the endomorphisms. Hence, σ preserves the set

$$\{j(E) : E \in \mathcal{ELL}_{\mathbb{C}}(\mathcal{O})\}$$

of roots of $H_{\mathcal{O}}(X)$. Thus, $H_{\mathcal{O}}(X) \in \mathbb{Q}[X]$. Moreover, for any proper $\mathcal{O}$ ideal $\mathfrak{a}$, $\mathcal{O}$ has an element of norm p if and only if $j(E_{\mathfrak{a}})$ is a root of $\Phi_p(X)$ and so $H_{\mathcal{O}}(X)$ divides $\Phi_p(X)$. Hence, by Gauss's lemma $H_{\mathcal{O}}(X) \in \mathbb{Z}[X]$. □

Remark 5.11. We saw in the proof of 5.10 that $\sigma \in \mathrm{Gal}(\overline{\mathbb{Q}}, \mathbb{Q})$. Using the invariant differential, one can define a canonical embedding $\mathrm{End}_{\mathbb{C}}(E) \hookrightarrow \mathbb{C}$. Then, as in Theorem II.2.1 of [8], let $[\alpha]_E \in \mathrm{End}_{\mathbb{C}}(E)$ correspond to $\alpha \in \mathcal{O}$ under this isomorphism. Under these isomorphisms, $[\alpha^\sigma]_{E^\sigma}$ and $([\alpha]_E)^\sigma$ have the same action on the invariant differential on E^σ, where the action of σ on $\mathcal{O} \subset \overline{\mathbb{Q}}$ is the usual action and σ acts on $\mathrm{End}_{\mathbb{C}}(E)$ by acting on the coefficients of the isogenies. Since an endomorphism over $\mathbb{C}$ is uniquely determined by its effect on the invariant differential, $[\alpha^\sigma]_{E^\sigma} = ([\alpha]_E)^\sigma$.

While the equality in Remark 5.11 might, at first glance, look like a bit of abstract nonsense, it actually carries a great deal of content. Recall that the isomorphism $\mathbb{C}/\Lambda \xrightarrow{\cong} E_\Lambda$ is defined via power series - the Weierstrass $\wp$-function and its derivative - and is therefore analytic. For instance, we shall see in Theorem 5.15 that when $\mathcal{O}$ is an order in the imaginary quadratic field K, $\mathrm{Gal}(\overline{\mathbb{Q}}, K)$ acts transitively on $\mathcal{ELL}_{\mathbb{C}}(\mathcal{O}) = \{E_{\mathfrak{a}} : \mathfrak{a} \in C(\mathcal{O})\}$. However, it fixes elements of $C(\mathcal{O})$. Hence, it is not possible to use Λ^σ to determine E^σ_Λ. A priori, multiplication by α^σ "should" only give an endomorphism of E_{Λ^σ}. With this in mind, the fact that the action of $\mathrm{Gal}(\overline{\mathbb{Q}}, \mathbb{Q})$ on endomorphisms is the same on the coefficients of the isogenies and on the canonical embedding in $\mathbb{C}$ is a minor miracle.

As we shall soon see, the equivalence between the action of σ on "multiplication by elements of $\mathcal{O}$" and $\mathrm{End}_{\mathbb{C}}(\mathcal{O})$, which would otherwise only be related analytically, has a very important consequence. Specifically, if L is the splitting field of $H_{\mathcal{O}}(X)$, then $\mathrm{Gal}(L/K)$ embeds canonically in the ideal class group $C(\mathcal{O})$.

But first, we shall see that the actions of the Galois group $\mathrm{Gal}(\overline{\mathbb{Q}}, K)$ and the ideal class group $C(\mathcal{O})$ on $\mathcal{ELL}_{\mathbb{C}}(\mathcal{O})$ commute. Our treatment seeks to explain (and mildly generalize) the proof of Theorem II.2.5 in [8], which also serves as a reference for some details that we omit.

Proposition 5.12. *Let $\mathcal{O}$ be an order in the imaginary quadratic field K and let $\mathfrak{a}$ and $\mathfrak{b}$ be proper fractional $\mathcal{O}$-ideals. Then, for any $\sigma \in \mathrm{Gal}(\overline{\mathbb{Q}}, \mathbb{Q})$,*

$$(\mathfrak{b}E_{\mathfrak{a}})^\sigma = \mathfrak{b}^\sigma E^\sigma_{\mathfrak{a}}.$$

Proof. To prove this equality, we again need to show that applying σ in two different contexts (which at first glance are only related analytically) yields the same result. Since $\mathfrak{b}$ determines the kernel of an isogeny, one might hope to apply the Velu formulas to describe both sides algebraically, but since the kernel must pass through the Weierstrass $\wp$-function, it quickly becomes clear that this approach is hopeless. Hence, we will need to use a different tool - Remark 5.11 - to relate the algebra and the analysis.

In order to apply Remark 5.11, we first need to describe $\mathfrak{b}$ in terms of individual elements of $\mathcal{O}$. With this in mind, consider a free resolution of $\mathfrak{b}$ as an $\mathcal{O}$-module.

$$\to \mathcal{O}^{n_2} \xrightarrow{B} \mathcal{O}^{n_1} \to \mathfrak{b} \to 0$$

In particular, B has coefficients in $\mathcal{O}$ and since $\sigma(\mathcal{O}) = \mathcal{O}$,

$$\to \mathcal{O}^{n_2} \xrightarrow{B^\sigma} \mathcal{O}^{n_1} \to \mathfrak{b}^\sigma \to 0$$

is a free resolution of the $\mathcal{O}$-module $\mathfrak{b}^\sigma$, where B^σ is given by the action of σ on the coefficients of B as a matrix over $\mathcal{O}$.

The second key idea is that we can learn about $\mathfrak{b}E_\mathfrak{a}$ by studying $\hom(\mathfrak{b}, E_\mathfrak{a})$. As motivation, consider the following lemma, which modifies II.2.5.1 of [8] so the proof holds for general orders.

Lemma 5.13. *Let $\mathcal{O}$ be an order in an imaginary quadratic field K and let $\mathfrak{a}$ be a $\mathcal{O}$-ideal and M a torsion-free $\mathcal{O}$-module. Then,*

$$\Phi : \mathfrak{b}^{-1}M \to \hom_\mathcal{O}(\mathfrak{b}, M), \alpha \mapsto (\phi_\alpha : \mathfrak{b} \to M, x \mapsto \alpha x)$$

is an isomorphism.

Proof. Since $\mathcal{O}$ is a domain, $\mathfrak{b}, M$ are torsion free and $\operatorname{Frac}\mathcal{O} = K$, $\phi : \mathfrak{b} \to M$ extends to a map $\mathfrak{b} \otimes K \to M \otimes K$. This is a linear map from a one-dimensional vector space, so ϕ is multiplication by some element of $M \otimes K$. Since $\phi(\mathfrak{b}) \subset M$, ϕ must be multiplication by some element of $\mathfrak{b}^{-1}M$. Any such a multiplication clearly defines a homomorphism, which completes the proof. □

There are two obvious ways to study $\hom(\mathfrak{b}, E_\mathfrak{a})$ - we can vary either the first or the second component.

Next, we vary the first component using our free resolution of $\mathfrak{b}$. Noting that $\hom(\mathcal{O}^n, E_\mathfrak{a}) \cong E_\mathfrak{a}^n$ as an $\mathcal{O}$-module, we have the following exact sequence of $\mathcal{O}$ modules:

$$0 \to \hom(\mathfrak{b}, E_\mathfrak{a}) \to E_\mathfrak{a}^{n_1} \xrightarrow{B^t} E_\mathfrak{a}^{n_2}$$

Here, B^t is both a matrix over $\mathcal{O}$ and also a morphism of abelian varieties. The key point is that by Remark 5.11 the actions of $\sigma \in \operatorname{Gal}(\overline{\mathbb{Q}}, \mathbb{Q})$ is the same on B^t as a matrix over $\mathcal{O}$ and as a a morphism of abelian varieties. Since $\hom(\mathfrak{b}, E_\mathfrak{a}) = \ker(B^t)$, this means that $\hom(\mathfrak{b}, E_\mathfrak{a})$ is an abelian variety and moreover

$$\hom(\mathfrak{b}, E_\mathfrak{a})^\sigma \cong \hom(\mathfrak{b}^\sigma, E_\mathfrak{a}^\sigma).$$

If instead we vary the second component using the exact sequence

$$0 \to \mathfrak{a} \to \mathbb{C} \to E_\mathfrak{a} \to 0$$

and apply Lemma 5.13 (twice), we get an exact sequence

$$0 \to \mathfrak{b}^{-1}\mathfrak{a} \to \mathbb{C} \to \hom(\mathfrak{b}, E_\mathfrak{a}).$$

Hence, $\mathfrak{b}E_\mathfrak{a} \cong \mathbb{C}/(\mathfrak{b}^{-1}\mathfrak{a})$ embeds naturally in $\hom(\mathfrak{b}, E_\mathfrak{a})$. In fact, we can say more. Taking the "hom product" of the free resolution of $\mathfrak{b}$ and the exact sequence for $E_\mathfrak{a}$ and applying the snake lemma, we can extend this exact sequence to

$$0 \to \mathfrak{b}^{-1}\mathfrak{a} \to \mathbb{C} \to \hom(\mathfrak{b}, E_\mathfrak{a}) \to \mathfrak{a}^{n_1}/(B^t\mathfrak{a}^{n_2}).$$

Viewing $\hom(\mathfrak{b}, E_\mathfrak{a})$ as an abelian subvariety of $E_\mathfrak{a}^{n_1}$, all of these maps are continuous in the complex topology. $\mathfrak{a}^{n_1}/(B^t\mathfrak{a}^{n_2})$ is discrete, so Image$(\mathbb{C})$ is a union of connected components and it is connected, since $\mathbb{C}$ is connected. Since also $(0, \ldots, 0) \in$ Image$(\mathbb{C})$, $\mathfrak{b}E_\mathfrak{a} \cong$ Image$(\mathbb{C})$ is the connected component of $\hom(\mathfrak{b}, E_\mathfrak{a})$ containing $(0, \ldots 0)$. Hence,

$$\begin{aligned}
(\mathfrak{b}E_\mathfrak{a})^\sigma &\cong (\text{connected component of } (0, \ldots, 0) \text{ in } \hom(\mathfrak{b}, E_\mathfrak{a}))^\sigma \\
&\cong \text{ connected component of } (0, \ldots, 0)^\sigma \text{ in } \hom(\mathfrak{b}, E_\mathfrak{a})^\sigma \\
&\cong \text{ connected component of } (0, \ldots, 0) \text{ in } \hom(\mathfrak{b}^\sigma, E_\mathfrak{a}^\sigma) \\
&\cong \mathfrak{b}^\sigma E_\mathfrak{a}^\sigma,
\end{aligned}$$

as desired. □

The following corollary (c.f. Theorem 22.7 of [9]) is (almost) immediate.

Corollary 5.14. *Let $\mathcal{O}$ be an order in the imaginary quadratic field K and let L be the splitting field of $\mathcal{O}$. For $[\mathfrak{a}] \in C(\mathcal{O})$ and $\sigma \in \mathrm{Gal}(L/K)$, let $[\mathfrak{a}_\sigma]$ denote the unique element of $C(\mathcal{O})$ such that $E_\mathfrak{a}^\sigma = \mathfrak{a}_\sigma E_\mathfrak{a}$. Then,*

1. *$\mathfrak{a}_\sigma$ is independent of the ideal class of $\mathfrak{a}$.*
2. *The map $F : \mathrm{Gal}(L/K) \to C(\mathcal{O}), \sigma \mapsto \mathfrak{a}_\sigma$ is an injective group homomorphism that commutes with the group actions on $\mathcal{ELL}_\mathbb{C}(\mathcal{O})$.*

Proof. If $[\mathfrak{a}], [\mathfrak{b}] \in C(\mathcal{O})$, then

$$\mathfrak{a}_\sigma E_\mathfrak{a} \cong E_\mathfrak{a}^\sigma \cong ((\mathfrak{b}\mathfrak{a}^{-1})E_\mathfrak{b})^\sigma \cong (\mathfrak{b}\mathfrak{a}^{-1})^\sigma E_\mathfrak{b}^\sigma \cong \mathfrak{b}\mathfrak{a}^{-1}\mathfrak{b}_\sigma E_\mathfrak{b} = \mathfrak{b}_\sigma E_\mathfrak{b},$$

where $(\mathfrak{b}\mathfrak{a}^{-1})^\sigma = \mathfrak{b}\mathfrak{a}^{-1}$ since σ fixes K and $((\mathfrak{b}\mathfrak{a}^{-1})E_\mathfrak{b})^\sigma \cong (\mathfrak{b}\mathfrak{a}^{-1})^\sigma E_\mathfrak{b}^\sigma$ by Proposition 5.12. Both (1) and (2) follow immediately since the action of $C(\mathcal{O})$ on $\mathcal{ELL}_\mathbb{C}(\mathcal{O})$ is simply transitive. □

In fact, using the results of Deuring from Section 3.4 we can say much more. Since $\mathrm{Gal}(L/K)$ injects into $C(\mathcal{O})$, $\mathrm{Gal}(L/K)$ is abelian and so the Artin map defines a morphism $C(\mathcal{O}) \to \mathrm{Gal}(L/K)$. Moreover, the map from Corollary 5.14 is an isomorphism, with inverse given by the Artin map $\left(\frac{L/K}{\cdot}\right)$. Formally, we have the following:

Theorem 5.15. *Let F be as in Proposition 5.14. Then, for all $[\mathfrak{a}] \in C(\mathcal{O})$ and all $\sigma \in \mathrm{Gal}(L/K)$,*

$$\left(\frac{L/K}{F(\sigma)}\right) = \sigma \text{ and } \left[F\left(\left(\frac{L/K}{\mathfrak{a}}\right)\right)\right] = [\mathfrak{a}].$$

Moreover, both isomorphisms F and $\left(\frac{K(j(\mathfrak{a}))/K}{\cdot}\right)$ commute with the group action on $\mathcal{ELL}_\mathbb{C}(\mathcal{O})$.

Proof. Our proof follows the proof of Theorem 22.8 of [9] and II.4.2 of [8]. We first prove that $\left[F\left(\left(\frac{L/K}{\mathfrak{a}}\right)\right)\right] = [\mathfrak{a}]$

Let S be the set of primes p such that not all of the following hold.

1. p is relatively prime to the conductor of f.
2. p is unramified in L
3. Every curve in $\mathcal{ELL}_\mathbb{C}(\mathcal{O})$ has good reduction modulo every prime $\mathfrak{P}$ of L lying over p. (Alternately, p is prime to the discriminant of every minimal model over L of these elliptic curves.)
4. p is relatively prime to the discriminant of $H_\mathcal{O}(X)$, i.e. $\prod_{[\mathfrak{a}]\neq[\mathfrak{b}]\in C(\mathcal{O})}(j(E_\mathfrak{a}) - j(E_\mathfrak{b}))^2$.

Since each condition is violated by only finitely many primes, S is finite. Hence, by Proposition 4.18, which says that every ideal class contains infinitely many ideals of prime norm and the proof of Theorem 2.15 which gives a correspondence between proper ideals of $\mathcal{O}$ and $\mathcal{O}_K$ prime to the conductor of $\mathcal{O}$, given any ideal class $[\mathfrak{a}] \in C(\mathcal{O})$, there is some prime ideal $\mathfrak{p}$ of $\mathcal{O}_K$ with norm p such that $\mathfrak{p} \cap \mathcal{O} \in [\mathfrak{a}]$.

Now, fix some elliptic curve $E_\mathfrak{b}$ with $\mathrm{End}_\mathbb{C}(E_\mathfrak{b}) = \mathcal{O}$ and let $\mathfrak{P}$ be a prime of L lying over $\mathfrak{p}$. Let $\widetilde{\cdot}$ denote reduction mod $\mathfrak{P}$. Then, $\mathfrak{p}$ induces an isogeny $\phi : E_\mathfrak{b} \to \mathfrak{p}E_\mathfrak{b}$ with $\deg(\phi) = p$. Then, Proposition 3.21 says that $\deg(\widetilde{\phi} : \widetilde{E_\mathfrak{b}} \to \widetilde{\mathfrak{p}E_\mathfrak{b}}) = p$.

Now, if $\mathfrak{a}' \in [\overline{\mathfrak{p}}] = [\mathfrak{p}]^{-1}$ has norm prime to $\mathfrak{p}$, then $\overline{\mathfrak{a}'}$ induces an isogeny $\psi : \mathfrak{p}E_{\mathfrak{b}} \to \mathfrak{a}'\mathfrak{p}E_{\mathfrak{b}} = E_{\mathfrak{b}}$. Then, the $\psi \circ \phi$ acts by multiplication by α on the invariant differential, where $\mathfrak{a}'\mathfrak{p} = \alpha\mathcal{O}_K$. Now, $\widetilde{\alpha} = \widetilde{\deg(\psi)\deg(\phi)} = \widetilde{0}$ and so $\widetilde{\psi} \circ \widetilde{\phi}$ is inseparable. $\widetilde{\psi}$ has degree equal to the norm of $\mathfrak{a}'$, which is prime to p, so $\widetilde{\psi}$ is separable, whence $\widetilde{\phi}$ is inseparable.

Now, since $\tilde{\phi}$ is inseparable, $\tilde{\phi}$ factors over π, the pth power Frobenius endomorphism as $\tilde{\phi} = \phi' \circ \pi$. Now, $\deg(\phi') = \deg(\tilde{\phi})/\deg(\pi) = 1$, so ϕ' defines an isomorphism between $\widetilde{E_{\mathfrak{b}}}^{\pi}$ and $\widetilde{\mathfrak{p}E_{\mathfrak{b}}}$. Moreover, by the definition of the Artin symbol, $\pi = \widetilde{\left(\frac{L/K}{\mathfrak{p}}\right)}$ and so

$$\widetilde{E_{\mathfrak{b}}^{\left(\frac{L/K}{\mathfrak{p}}\right)}} \cong \widetilde{\mathfrak{p}E_{\mathfrak{b}}}.$$

The assumption that p is prime to the discriminant of $H_{\mathcal{O}}(X)$ means that the reductions mod $\mathfrak{P}$ of non-isomorphic curves in $\mathcal{ELL}_L(\mathcal{O})$ are not isomorphic. Hence, $E_{\mathfrak{b}}^{\left(\frac{L/K}{\mathfrak{p}}\right)} = \mathfrak{p}E_{\mathfrak{b}}$ and so $F\left(\left(\frac{L/K}{\mathfrak{p}}\right)\right) = [\mathfrak{p}]$, as desired.

In particular, F is surjective. We have already seen that F is injective, so F is a group isomorphism. The remainder of the theorem follows immediately from our earlier results. □

We conclude this subsection with the following corollary, which answers our question about the minimal polynomial of $j(E_{\mathfrak{a}})$.

Corollary 5.16. *Let $\mathcal{O}$ be an order in an imaginary quadratic field K and let $\mathfrak{a}$ be any proper fractional $\mathcal{O}$-ideal. Then, the Hilbert Class polynomial $H_{\mathcal{O}}(X)$ is irreducible over K and has splitting field $K(j(E_{\mathfrak{a}}))$.*

Proof. Let L be the splitting field of $H_{\mathcal{O}}(X)$. Since the action of $C(\mathcal{O})$ on $\mathcal{ELL}_{\mathbb{C}}(\mathcal{O})$ is transitive, Theorem 5.15 implies that the action of $\mathrm{Gal}(L/K)$ on $\mathcal{ELL}_{\mathbb{C}}(\mathcal{O})$ is transitive, so, $H_{\mathcal{O}}(X)$ is irreducible over K.

Moreover, $[L : K] = |C(\mathcal{O})| = [K(j(E_{\mathfrak{a}})) : K]$ and $j(E_{\mathfrak{a}}) \in L$, so $L = K(j(E_{\mathfrak{a}}))$, as desired. □

5.4. **The Main Theorems of Complex Multiplication.** For completeness sake, we now give the two main theorems of complex multiplication, using the statements from [1]. These theorems show how complex multiplication elliptic curves can be used to generate ring class fields and ray class fields, respectively. We essentially proved the first main theorem in the previous section and omit the proof of the second main theorem.

Theorem 5.17 (The First Main Theorem of Complex Multiplication)**.** *Let $\mathcal{O}$ be an order in an imaginary quadratic field K and let $\mathfrak{a}$ be a proper fractional $\mathcal{O}$-ideal. Then, $j(E_{\mathfrak{a}})$ is an algebraic integer and $K(j(E_{\mathfrak{a}}))$ is the ring class field of the order $\mathcal{O}$.*

Proof. We saw in the previous section that $K(j(E_{\mathfrak{a}}))$ is a Galois extension of K with Galois group $C(\mathcal{O})$. By the uniqueness assertion of the Existence Theorem 4.10 of Class Field Theory, $K(j(E_{\mathfrak{a}}))$ is the ring class field of $\mathcal{O}$. We proved that $j(E_{\mathfrak{a}})$ is an algebraic integer in Corollary 5.7. □

We now briefly comment on the proof given in section 11.D of [1], which at first glance may seem quite different from our proof, but is actually quite similar. In

particular, both proofs first show that all of the $j(E_{\mathfrak{a}})$ lie in the ring class field L, which is critically important in the second half of the proof. In our proof, this first result was disguised as the artin map mapping injectively into the class group, but it is easy to see using the correspondence theorems of class field theory that these result are really the same. Also, where we used the Artin map to find the Frobenius elements, while Cox's proof uses a corresponding argument using the Kronecker congruence. Again these are two sides of the same coin. Finally, the possible congruence relations in Cox's argument correspond to the possibility that our ϕ was separable or inseparable and a similar argument is needed to show that the separable case always holds.

We now quickly define the Weber function as in 11.D of [1] (omitting a couple of special cases), which we will need to state the Second Main Theorem of Complex Multiplication.

Definition 5.18. Let Λ be a lattice with $g_2(\Lambda) \neq 0, g_3(\Lambda) \neq 0$. Then, the Weber function $\tau(z;\Lambda)$ is defined by

$$\tau(z;\Lambda) = \frac{g_2(\Lambda)g_3(\Lambda)}{g_2(\Lambda)^3 - 27g_3(\Lambda)^2}\wp(z;\Lambda)$$

The advantage of $\tau(z;\Lambda)$ over $\wp(z;\Lambda)$ is that it is invariant upon rescaling the lattice. We can now cite [1] and state:

Theorem 5.19. *Let K be an imaginary quadratic field of discriminant d_K, let $w_K = (d_K + \sqrt{d_K})/2$ and let N be a positive integer. Then,*

1. *$K(j(E_{\mathcal{O}_K}), \tau(1/N;\mathcal{O}_K))$ is the ray class field of the ideal $N\mathcal{O}_K$.*
2. *Let $\mathcal{O}$ be the order of conductor N in K. Then, $K(j(E_{\mathcal{O}}), \tau(w_K;\mathcal{O}))$ is the ray class for the modulus $N\mathcal{O}_K$.*

Remark 5.20. We omit the proof. It is worth noting that $\tau(1/N;\mathcal{O}_K)$ (and similarly $\tau(w_K;\mathcal{O})$) is an N-torsion point on the elliptic curve. Hence, this result (which gives generating sets for many of the finite abelian extensions of K) is analogous to the result that the abelian extensions of $\mathbb{Q}$ are given by adjoining torsion points of the multiplicative group $\mathbb{C}^\times$, i.e. roots of unity. Indeed, this analogy is at the heart of the theory of complex multiplication.

5.5. **The CM Method and Other Computations.** Having seen the applications of $H_{\mathcal{O}}(X)$ and $\Phi_n(X,X)$ to the theory of complex multiplication and class field theory, we now discuss some computational applications.

First, we discuss the CM method for computing curves with a specified number of points. Our goal is to find an elliptic curve E over the finite field $\mathbb{F}_q$ with some specified number of points - say $q+1-t$ points. The ability to find such curves has important applications for several cryptographic and primality-proving algorithms. Our treatment is influenced by Sutherland's exposition in [9] and [10].

We motivate the CM method by noting that if $\mathrm{End}(E) \cong \mathcal{O}$, the endomorphisms of $E/\mathbb{F}_q$ correspond to principal ideals. Then, E has an endomorphism of order q if and only if $\mathcal{O}$ has a principal ideal of norm q, i.e. if and only if we can write

$$q = \left(\frac{t + v\sqrt{D}}{2}\right)\left(\frac{t - v\sqrt{D}}{2}\right) = \frac{t^2 - v^2 D}{4}$$

for some integers t and v. Moreover, we claim that for fixed $D < -4$, that if t and v are prime to q, then they are unique up to sign.

It clearly suffices to prove uniqueness when $D = d_K$, the discriminant of $\mathcal{O}_K$, as taking the ideal of conductor f multiplies the discriminant by f^2. In this case, taking $p = \text{char}(q)$, if q factors as above, then either factor generates a principal ideal of $\mathcal{O}_K$ of norm q that is not a p-th power. Since the only units of $\mathcal{O}_K$ are ± 1, these factors are not equal. Hence, p must split completely in $\mathcal{O}_K$ as $p\mathcal{O}_K = \mathfrak{p}\overline{fkp}$. Every $\mathcal{O}_K$-ideal of norm $q = p^a$ that is not divisible by p must therefore factor as either $\mathfrak{p}^a$ or $\overline{\mathfrak{p}}^a$. Since the only units of $\mathcal{O}_K$ are ± 1 and these ideals are principal, their generator is unique up to sign, as desired.

Now, if $E/\mathbb{F}_q$, we always have a endomorphism of norm q, namely the qth power Frobenius endomorphism π. Hence, if $\text{End}(E) = \mathcal{O}$, we must have that $\pi = \frac{\pm t \pm v\sqrt{D}}{2}$ for our unique choice of t and v. Also, we know that π satisfies the polynomial $x^2 - \text{tr}(\pi)x + N(\pi) = 0$ and so by the quadratic formula, $\pi = \frac{\text{tr}(\pi) \pm \sqrt{\text{tr}(\pi)^2 - 4N(\pi)}}{2}$, so $\text{tr}(\pi) = \pm t$. Hence, the elliptic curve $E/\mathbb{F}_p$ has $q + 1 \mp t$ points and so the quadratic twist will have $q + 1 \pm t$ points. In particular, exactly one of $E/\mathbb{F}_q$ and its quadratic twist will have $q + 1 - t$ points. If q is large, testing the order of a random point on each curve will distinguish between the two possibilities with high probability. Hence, we have reduced the problem of finding an elliptic curve with a specified number of points to the problem of finding an elliptic curve with a specified CM endomorphism ring. This approach to finding curves is known as the CM method.

Fortunately, we know exactly how to find a curve E with $\text{End}(E) \cong \mathcal{O}$ - we use the Hilbert Class polynomial $H_{\mathcal{O}}(X)$! So long as q is prime to the conductor of $\mathcal{O}$ and $p = \text{char}(q)$ splits completely in K - as we have seen is always the case when applying the CM method - by Corollary 3.25 of the Deuring Lifting Theorem, the roots of $H_{\mathcal{O}}(X)$ correspond under reduction modulo some prime $\mathfrak{P}$ to elliptic curves over $\overline{\mathbb{F}_q}$ with endomorphism ring $\mathcal{O}$. In particular, roots of $H_{\mathcal{O}}(X)$ lying in $\mathbb{F}_q$ are the j-invariants of elliptic curves over $\mathbb{F}_p$ with endomorphism ring $\mathcal{O}$. Thus, if we can compute the Hilbert Class polynomial $H_{\mathcal{O}}(X)$, we can easily find curves with endomorphism ring $\mathcal{O}$ and hence find curves with a specified number of points.

To illustrate how powerful knowing the Hilbert Class polynomial can be, we now demonstrate an ad-hoc method to finding an elliptic curve (or rather it's j-invariant) with a particular endomorphism ring, namely the ring of integers $\mathcal{O}_K$ in the field $K = \mathbb{Q}(\sqrt{-7})$.

Example 5.21. Our goal is to find (the j-invariant of) an elliptic curve (over $\mathbb{C}$) with endomorphism ring $\mathcal{O}_{\mathbb{Q}(\sqrt{-7})}$. Since this ring has class number 1, setting $x = \frac{1+\sqrt{-7}}{2}$, this curve must correspond to the lattice $\Lambda = C[1, x]$, where we will choose the constant C for computational convenience later on. Our computation generalizes the similar computation of $j(E_{\mathcal{O}_{K(\sqrt{-2})}})$ from Section 10.C of [1].

We will compute $j(E_\Lambda)$ using the power series expansion of $\wp(z; \Lambda)$, noting that both j and the coefficients of this power series are functions of $g_2(\Lambda)$ and $g_3(\Lambda)$. From this point forward, we tread Λ as fixed (except for our potential rescaling) and omit it from the notation.

Since the following arguments apply to compute $j\left(E_{[1,x]}\right)$ for any algebraic integer $x(\neq \pm 2)$ of norm 2, we leave x as a variable until we reach the point where it makes the equations unpleasant to write.

$\wp(z)$ is an even Λ-periodic function, so as $x\Lambda \subset \Lambda$, $\wp(xz)$ is also an even Λ-periodic function. In particular, part (ii) of Proposition 3.13 says that $\wp(xz)$ is a

rational function of $\wp$. Since x is an algebraic integer of norm 2, $\wp(xz)$ has two poles, each of order 2 on a fundamental domain for Λ, while $\wp(z)$ has only one pole of order 2, so it is not hard to see that $\wp(xz) = \frac{A(\wp(z))}{B(\wp(z))}$, where A has degree 2 and B has degree 1. (In fact, we can see that $B(\wp(z)) = c(\wp(z) - \wp(C/x))$.) Hence, we may write:

$$\wp(xz) = a\wp(z) + b + \frac{1}{c\wp(z) + d} \tag{5.7}$$

for some complex numbers a, b, c, d. Now, note that $\wp(z)$ has the power series expansion:

$$\wp(z) = \frac{1}{z^2} + \frac{g_2}{20}z^2 + \frac{g_3}{28}z^4 + \frac{g_2^2}{1200}z^6 + \cdots$$

By choosing an appropriate scale-factor C in the definition of Λ, we may assume that $g_2 = 20g$ and $g_3 = 28g$, so that

$$\wp(z) = \frac{1}{z^2} + gz^2 + gz^4 + g^2 3z^6 + \cdots$$

Looking at the coefficients of z^{-2} and z^0 in (5.7), we see that $a = \frac{1}{x^2}$ and $b = 0$. Then, we have that

$$\begin{aligned}
&c\wp(z) + d \\
=&\frac{c}{z^2} + d + gcz^2 + gcz^4 + \frac{g^2}{3}cz^6 + \cdots \\
=&(\wp(xz) - a\wp(z))^{-1} \\
=&\left(\left(x^2 - \frac{1}{x^2}\right)gz^2 + \left(x^4 - \frac{1}{x^2}\right)gz^4 + \left(x^6 - \frac{1}{x^2}\right)\frac{g^2}{3}z^6 + \cdots\right)^{-1} \\
=&\frac{z^{-2}}{g\left(x^2 - \frac{1}{x^2}\right)} - \frac{g\left(x^4 - \frac{1}{x^2}\right)}{g^2\left(x^2 - \frac{1}{x^2}\right)^2} + \frac{-z^2}{g\left(x^2 - \frac{1}{x^2}\right)}\left(\frac{-\left(x^4 - \frac{1}{x^2}\right)^2 + \frac{g}{3}\left(x^6 - \frac{1}{x^2}\right)\left(x^2 - \frac{1}{x^2}\right)}{\left(x^2 - \frac{1}{x^2}\right)^2}\right) + \cdots
\end{aligned}$$

From the coefficient of z^{-2}, $c = \frac{1}{g\left(x^2 - \frac{1}{x^2}\right)}$. Looking at the coefficient of z^2, we have

$$g\left(x^2 - \frac{1}{x^2}\right)^2 = \left(x^4 - \frac{1}{x^2}\right)^2 + \frac{g}{3}\left(x^6 - \frac{1}{x^2}\right)\left(x^2 - \frac{1}{x^2}\right)$$

Rearranging terms,

$$g = \frac{3(x^6 - 1)^2}{(x^4 - 1)(x^8 + 3x^4 - 4)} = \frac{3(x^6 - 1)^2}{(x^4 - 1)^2(x^4 + 4)}$$

Now, plugging in $x = \frac{1+\sqrt{-7}}{2}$, we see that $g = \frac{7}{4}$. Hence,

$$j\left(\frac{1 + \sqrt{-7}}{2}\right) = \frac{1728g_2^3}{g_2^3 - 27g_3^2} = \frac{1728 \cdot 8000g}{8000g - 27 \cdot 28^2} = \frac{1728 \cdot 14000}{14000 - 21168} = -3375.$$

Hence, the elliptic curve with j-invariant -3375 has complex multiplication by $\mathcal{O}_{\mathbb{Q}(\sqrt{-7})}$ and since $|C(\mathcal{O}_{\mathbb{Q}(\sqrt{-7})})| = 1$, it is the unique such curve (up to isomorphism.)

By ad-hoc approaches, one can compute many more j-invariants, and the corresponding elliptic curves, but these methods become complicated quite quickly. [1]

spends 7 full pages computing $j(\sqrt{-14})$. While some of the machinery developed there can be applied to compute other j-invariants, the more general approach provided by the Hilbert Class polynomial is clearly needed.

As Cox discusses in Section 13 of [1], one can use the modular polynomials $\Phi_m(X, Y)$ to compute the Hilbert Class polynomials $H_{\mathcal{O}}(X)$. In fact, there are explicit formulas for the factorization of $\Phi_m(X, X)$ into Hilbert Class polynomials. Hence, choosing m appropriately, factoring $\Phi_m(X, X)$, and looking at the irreducibles that appear only once in the factorization, gives an algorithm to compute any $H_{\mathcal{O}}(X)$. It is also worth noting that this gives a deterministic algorithm to determine the class number of any order, given access to appropriate $\Phi_m(X, Y)$.

While there are deterministic algorithms to compute $\Phi_m(X, Y)$, by matching up coefficients of j-expansions, these algorithms are typically very slow, because the coefficients of these polynomials grow incredibly quickly. See Section 13 of [1] for more detail.

In the remaining section, we will discuss a new algorithm, based on "isogeny volcanoes" that has tremendously increased the range of orders for which it is possible to compute $H_{\mathcal{O}}(X)$.

6. Isogeny Volcanoes

In this section, we discuss a powerful tool for computing elliptic curves - isogeny volcanoes. The main idea is that instead of studying an elliptic curve in isolation, we can learn about an elliptic curve more efficiently by studying the isogenous elliptic curves. As we shall see, if we restrict ourselves to following ℓ-isogenies (for ℓ a prime not equal to the characteristic of the base field), then the resulting graph of elliptic curves has a very special structure, called an ℓ-volcano. According to Sutherland in [10], this was first discovered/explained by David Kohel in his PhD thesis [3].

The structure of the ℓ-isogeny graph allows for several special purpose algorithms that allow various computations - of endomorphism rings, Hilbert class polynomials, and modular polynomials, among other things - to be carried out for much larger parameters than was previously possible. As an application, we will discuss the isogeny volcano-based algorithm for computing the Hilbert Class polynomial.

6.1. The ℓ-Isogeny Graph.

In this section, we will study ℓ-isogenies of elliptic curves over a field k, for a prime $\ell \neq \mathrm{char}(k)$. Our exposition is influenced by both [10] and Lecture 23 of [9].

We start by citing the following definition of the ℓ-isogeny graph from [10].

Definition 6.1. Let k be a field. The ℓ-isogeny graph $G_\ell(k)$ has vertex set k and directed edges (j_1, j_2) present with multiplicity equal to the multiplicity of j_2 as a root of $\Phi_\ell(j_1, Y)$.

By Remark 5.5, we have that there edges (j_1, j_2) correspond exactly to isogenies from the elliptic curve with j-invariant j_1 to the elliptic curve with j-invariant j_2. We will see that this graph has a remarkable structure that is dicatated by the endomorphism rings of the curves.

Hence, in order discuss the ℓ-isogeny graph, we first investigate the relationship between the endomorphism rings of ℓ-isogenous curves.

Proposition 6.2. *Let K be an imaginary quadratic field, k be any field and let E/k be an elliptic curve with*

$$\mathrm{End}_{\overline{k}}(E) = \mathcal{O} = \mathbb{Z} + f\mathcal{O}_K.$$

Suppose that $\phi : E \to E'$ is an isogeny with $\deg(\phi) = \ell$ a prime. Then,

$$\mathrm{End}_{\overline{k}}(E') = \mathcal{O}' = \mathbb{Z} + f'\mathcal{O}_K,$$

where either $f' = f$, $f' = \ell f$, or $f' = f/\ell$. In particular, if $\mathcal{O}$ is an order in K, $\mathcal{O}'$ is also an order in K.

Proof. First, we prove that if $\mathcal{O}$ is an order in K, then $\mathcal{O}'$ is an order in K. Since the endomorphism rings are torsion-free, the additive group homomorphism

$$\Psi : \mathrm{End}(E) \to \mathrm{End}(E'), \quad \psi \mapsto \phi \circ \psi \circ \widehat{\phi}$$

extends to a map $\Psi' : \mathrm{End}(E) \otimes \mathbb{Q} \to \mathrm{End}(E') \otimes \mathbb{Q}$. Then, define $L : \mathrm{End}(E') \otimes \mathbb{Q} \to \mathrm{End}(E') \otimes \mathbb{Q}$ by composition with the multiplication-by-$(1/\ell)$ map (on either side since it commutes with isogenies). It is clear that $L \circ \Psi'$ defines an isomorphism of rings between $\mathrm{End}(E) \otimes \mathbb{Q}$ and $\mathrm{End}(E') \otimes \mathbb{Q}$ with inverse defined similarly, but changing the role of ϕ and its dual $\widehat{\phi}$. Hence, $\mathrm{End}(E) \otimes \mathbb{Q} \cong \mathrm{End}(E') \otimes \mathbb{Q}$, so they are orders in the same imaginary quadratic field.

Furthermore, the map $L \circ \Psi'$ and its inverse tell us that $\ell\mathcal{O} \subset \mathcal{O}'$ and $\ell\mathcal{O}' \subset \mathcal{O}$, whence $f' = f$, $f' = f\ell$, or $f'\ell = f$. □

It will be useful to have some terminology to refer to these three cases, so we state the following definition, paraphrased from 2.7 of [10] (see also Definition 23.2 of [9]).

Definition 6.3. Let $\phi : E \to E'$ be an ℓ-isogeny between elliptic curves with endomophism rings $\mathcal{O} = \mathbb{Z} + f\mathcal{O}_K$ and $\mathcal{O}' = \mathbb{Z} + f'\mathcal{O}_K$ orders in the imaginary quadratic field K. If

(1) $\mathcal{O} = \mathcal{O}'$, then ϕ is *horizontal*.
(2) $\mathcal{O} \supsetneq \mathcal{O}'$, then $[\mathcal{O} : \mathcal{O}'] = \ell$ and ϕ is *descending*.
(3) $\mathcal{O} \subsetneq \mathcal{O}'$, then $[\mathcal{O}' : \mathcal{O}] = \ell$ and ϕ is *ascending*.

If ϕ is either descending or ascending, then ϕ is *vertical*.

Now, we know that so long as $\ell \neq \mathrm{char}(k)$, there are exactly $\ell + 1$ isogenies of degree ℓ from any elliptic curve $E/\overline{k}$. These isogenies correspond exactly to the $\ell + 1$ cyclic subgroups of the ℓ-torsion $(\mathbb{Z}/\ell\mathbb{Z})^2$. Equivalently, they correspond to index ℓ (additive) subgroups of the lattice $\mathcal{O}$. It would be nice to know given a curve E how many of these isogenies are horizontal, how many are descending, and how many are ascending. Remarkably, under mild restrictions on k, the answer depends only on the prime ℓ and the order $\mathcal{O} = \mathrm{End}_{\overline{k}}(E)$.

Theorem 6.4. *Let $\overline{k}$ be an algebraically closed field and let $E/\overline{k}$ be an elliptic curve with $\mathrm{End}_{\overline{k}}(E) \cong \mathcal{O} = \mathbb{Z} + f\mathcal{O}_K$ be an order in the imaginary quadratic field K. Let $\ell \neq \mathrm{char}(\overline{k})$ be a prime. If ℓ divides f, there are 0 horizontal ℓ-isogenies, ℓ descending isogenies, and 1 ascending isogeny from E. If $\ell \nmid f$, then the number of horizontal ℓ-isogenies*

- 0 *if ℓ is inert in K.*
- 1 *if ℓ is ramified in K.*
- 2 *is ℓ splits completely in K.*

and the remaining ℓ-isogenies are (necessarily) descending.

Proof. Our proof for the horizontal isogenies expands on section 2.9 of [10], while the proof in the vertical case is this author's own (although it has likely been discovered before).

Since the j-invariants are (the reductions of) algebraic integers, we may assume that our curve and all ℓ-isogenous curves are defined over some finite extension of $\mathbb{F}_p$ or $\mathbb{Q}$. Then, applying the result of Deuring that reduction of isogenies is injective and preserves degrees (Proposition 3.21), it suffices to check the case where $\overline{k} = \mathbb{C}$.

Now, for any ℓ-isogeny $\phi : E \to E'$, then E and E' must correspond to proper fractional ideals $\mathfrak{a}$ of $\mathcal{O}$ and $\mathfrak{b}$ of $\mathcal{O}'$, where $\mathfrak{b} \subset \mathfrak{a}$ by Proposition 3.17.

If $\mathcal{O} = \mathcal{O}'$, i.e. the isogeny is horizontal, $\mathfrak{b}$ is a proper fractional $\mathcal{O}$-ideal. Hence, $\mathfrak{a}^{-1}\mathfrak{b}$ is a proper $\mathcal{O}$-ideal of norm ℓ if and only if the isogeny is horizontal.

If $\mathcal{O} \subsetneq \mathcal{O}'$, i.e. the isogeny is ascending, $\mathfrak{b}$ is a fractional $\mathcal{O}$-ideal, but is not proper, whence $\mathfrak{a}^{-1}\mathfrak{b}$ is a $\mathcal{O}$-ideal of norm ℓ that is not proper. Hence, $\mathfrak{a}^{-1}\mathfrak{b}$ is a non-proper $\mathcal{O}$-ideal of norm ℓ if and only if the isogeny is ascending.

In the remaining case, $\mathcal{O} \supset \mathcal{O}'$, we see that the isogeny is descending if and only if $\mathfrak{a}^{-1}\mathfrak{b}$ is not a $\mathcal{O}$-ideal.

So, it suffices to determine how many of the $\ell + 1$ (additive) subgroups of $\mathcal{O}$ of order ℓ are proper ideals, how many are non-proper ideals, and how many are not ideals at all.

If ℓ divides the conductor $f = \ell f'$ of $\mathcal{O}$, then we may write $\mathcal{O}$ as the lattice $[\ell f' \omega_K, 1]$, where $\omega_K = \frac{D+\sqrt{D}}{2}$, where D is the discriminant of $\mathcal{O}_K$. Now, we know by Proposition 5.2 that every integer matrix with determinant ℓ has the form MN, where M is in $\mathrm{SL}_2(\mathbb{Z})$ and N is in C_ℓ. Hence, the sublattices of index ℓ are exactly $[\ell f' \omega_K + a, \ell]$, where $0 \leq a < \ell$ and $[\ell^2 f' \omega_K, 1]$.

Considering the imaginary part, $\ell f' \omega_K \notin [\ell^2 f' \omega_K, 1]$, so $[\ell^2 f' \omega_K, 1]$ is not a $\mathcal{O}$-ideal. Similarly,

$$(\ell f' \omega_K + a)\ell f' \overline{\omega_K} = \ell[\ell (f')^2 |\omega_K|^2 + f' a(\overline{\omega_K} + \omega_K)] - a(\ell f' \omega_K + a) + a^2$$

is in $[\ell f' \omega_K + a, \ell]$ if and only if $a \equiv 0 \pmod{\ell}$. So, unless $a = 0$, this is not an ideal. Finally, it is easy to see that in fact, $[\ell f' \omega_K, \ell]$ is not a proper $\mathcal{O}$-ideal, but rather a proper $\mathbb{Z} + f'\mathcal{O}_K$ ideal.

Hence, if ℓ divides the conductor of $\mathcal{O} = \mathrm{End}(E)$, then there is one ascending ℓ-isogeny from E and the remaining ℓ of the ℓ-isogenies are descending.

If ℓ does not divide the conductor of $\mathcal{O} = \mathrm{End}(E)$, then as in the proof of Theorem 2.15, there is a bijection between (proper) $\mathcal{O}$-ideals of norm ℓ and $\mathcal{O}_K$-ideals of norm ℓ, whence the claim about the number of horizontal isogenies is clear. Since it is impossible to have an ascending isogeny in this case, the remaining isogenies must be descending.

□

In fact, we can say slightly more about the components of CM elliptic curves in the ℓ-isogeny graph (where vertices are elliptic curves and edges are ℓ-isogenies from one to another). First, while a-priori the graph is bidirected, we may view it as an undirected graph, since the dual of an ℓ-isogeny is also an ℓ-isogeny in the opposite direction. Moreover, it is clear that all cycles must use only elliptic cures with endomorphism ring $\mathcal{O}$ having conductor prime to ℓ. Considering the $C(\mathcal{O})$ action on $\mathcal{ELL}_{\overline{k}}(\mathcal{O})$, we see that the size of a cycle must be the order in $C(\mathcal{O})$ of the norm ℓ proper ideal of $\mathcal{O}$.

Since the proof of Lemma 6 (the vertical isogeny portion of our Theorem 6.4) in [10] is superficially different looking, we make a few brief comments. First, our decision to multiply by the ideal $\mathfrak{a}^{-1}$ corresponds to the change of curves that Sutherland uses to prove that if one curve with a given endomorphism ring has an ascending isogeny, then all such curves do. The remainder of his proof is a counting argument and induction, using the result from Theorem 2.20 on the relative sizes of the ideal class groups. Note that our proof of Theorem 6.4, together with the following discussion is a stronger argument, in that a simple counting argument provides an easy proof of Theorem 2.20, rather than requiring this Theorem as a step in the proof.

Remark 6.5. Next, we consider the case where k is not algebraically closed. If $\mathcal{O}$ is an order in the imaginary quadratic field K and k contains $\sqrt{d_K}$, where d_K is the discriminant of $\mathcal{O}_K$ (equivalently, k contains the discriminant of $\mathcal{O}$), then by Corollary 5.16, $H_{\mathcal{O}}(X)$ splits completely in k if and only if it has a root in k. This means that either elliptic curve with endomorphism ring $\mathcal{O}$ can be defined over k or none of the elliptic curves with endomorphism ring $\mathcal{O}$ can be defined over k.

In fact, we can say more. First recall that if $\mathcal{O}$ and $\mathcal{O}'$ are the orders of conductor f and ℓf respectively, then we have

$$C(\mathcal{O}) \cong I_K(f)/P_{K,\mathbb{Z}}(f), \qquad C(\mathcal{O}') \cong I_K(\ell f)/P_{K,\mathbb{Z}}(\ell f)$$

as in Theorem 2.15. Now, since every ideal class of $C(\mathcal{O})$ contains a representative that is relatively prime to ℓf, we can rewrite $C(\mathcal{O})$ as

$$C(\mathcal{O}) \cong I_K(\ell f)/(P_{K,\mathbb{Z}}(f) \cap I_K(\ell f)).$$

Now, it is clear that

$$(P_{K,\mathbb{Z}}(f) \cap I_K(\ell f)) \supset P_{K,\mathbb{Z}}(\ell f),$$

so by Galois Theory and the uniqueness assertion of the existence theorem of class field theory (Theorem 4.10), the splitting field of $H_{\mathcal{O}}(X)$ over K is contained in the splitting field of $H_{\mathcal{O}'}(X)$ over K.

Reducing modulo a prime if necessary, it follows that, so long as k contains $\sqrt{d_K}$, if one elliptic curve with endomorphism ring $\mathbb{Z} + \ell^t f\mathcal{O}_K$ is defined over k, then for $0 \leq s \leq t$, then every elliptic curve with endomorphism ring $\mathbb{Z} + \ell^s f\mathcal{O}_K$ can be defined over k.

According to 2.3 of [10], given two ℓ-isogenous elliptic curves, both defined over k, we may assume that the ℓ-isogeny is also defined over k by choosing appropriate twists.

Remark 6.6. Unfortunately, when k does not contain $\sqrt{d_K}$, the situation is mildly more complicated. It is possible that some, but not all elliptic curves over $\overline{k}$ with endomorphism ring $\mathcal{O} = \mathrm{End}_{\overline{k}}(E)$ are defined over k.

However, as claimed in Lemma 23.5 of [9], it is is still true that either 0 or $|C(\mathcal{O})|$ curves have $\mathrm{End}_k(E) = \mathcal{O}$.

If $\mathrm{End}_{\overline{k}}(E) = \mathcal{O}$, then every isogeny acts on the invariant differential by $a + b\sqrt{d_K}$ for some $a, b \in \mathbb{Z}$. Now, the action on the invariant differential is defined algebraically, so if $b \neq 0$ for some endomorphism defined over k, then $a + b\sqrt{d_K} \in k$, and so $\sqrt{d_K} \in k$. If this is not the case, then every endomorphism acts on the invariant differential by multiplication by a, and so $\mathrm{End}_k(E) = \mathbb{Z}$. This confirms Kohel's dictionary at the end of Section 3.2 of his thesis [3], which says that in

characteristic zero $K(j(E_{\mathfrak{a}}))$ is the field of definition for the endomorphism ring of $E_{\mathfrak{a}}$, while $\mathrm{End}_{\mathbb{Q}(j(E_{\mathfrak{a}}))}(E_{\mathfrak{a}}) = \mathbb{Z}$.

Fortunately, when working with elliptic curves over finite fields, we will always have a Frobenius endomorphism which does not lie in $\mathbb{Z}$, so we will never have to worry about this possible issue.

This fact is stated somewhat unclearly in the existing literature on the topic, since the field of definition of the endomorphism is not always specified when talking about the endomorphism ring of an elliptic curve over a field k that is not algebraically closed.

For completeness, we give an example where some, but not all elliptic curves over $\overline{k}$ with endomorphism ring $\mathcal{O} = \mathrm{End}_{\overline{k}}(E)$ are defined over k, noting that this is *not* the correct interpretation of Sutherland's statements in 2.8 of [10] and Lemma 23.5 of [9].

One additional source of confusion is Sutherland's claim in his "proof" of Lemma 23.5 that if $\mathrm{End}(E) = \mathcal{O}$, then $\mathbb{Q}(j(E))$ is the splitting field of $H_{\mathcal{O}}(X)$ over $\mathbb{Q}$. However, $j(E_{\mathfrak{a}})$ is real if and only if $\mathfrak{a} = \overline{\mathfrak{a}} = \mathfrak{a}^{-1}$ in $C(\mathcal{O})$ since $j(E_{\overline{\mathfrak{a}}}) = \overline{j(E_{\mathfrak{a}})}$. In particular, $\mathbb{Q}(j(E_{\mathcal{O}}))$ is totally real. Hence, if $\mathcal{O}$ is any order with class group *not* isomorphic to $(\mathbb{Z}/2\mathbb{Z})^n$ for some n, $\mathbb{Q}(j(E_{\mathcal{O}}))$ contains some, but not all of the elliptic curves over $\mathbb{C}$ with endomorphism ring $\mathcal{O} = \mathrm{End}_{\mathbb{C}}(E)$.

Ultimately, the problem that this example demonstrates is that if L is the Hilbert Class field of the order $\mathcal{O}$ in the imaginary quadratic field K, then

$$\mathrm{Gal}(L/\mathbb{Q}) \cong \mathrm{Gal}(L/K) \rtimes \mathrm{Gal}(K/\mathbb{Q}) \cong C(\mathcal{O}) \rtimes \mathbb{Z}/2\mathbb{Z},$$

where the action of $\mathbb{Z}/2\mathbb{Z}$ is to invert elements of $C(\mathcal{O})$ - equivalently, it acts by conjugation. Since $\mathbb{Z}/2\mathbb{Z}$ is not normal in $\mathrm{Gal}(L/\mathbb{Q})$ if $C(\mathcal{O}) \ncong (\mathbb{Z}/2\mathbb{Z})^n$, we cannot quotient by it. Hence, $\mathbb{Q}(j(E_{\mathcal{O}}))$, the fixed field of the conjugation action, is not Galois over $\mathbb{Q}$, which yields the previous example.

With this out of the way, we are classify the components of the ℓ-isogeny graph containing CM curves. First, we define a special type of graph, known as an ℓ-volcano, citing the definition from [10]. The properties should seem familiar from our recent discussions.

Definition 6.7. An ℓ-volcano is a connected undirected graph whose vertices are partitioned into one or more levels $V_0, \ldots, V_d$ such that the following hold:

(1) The subgraph of V_0 is a regular graph of degree at most 2.
(2) For $i > 0$, each vertex in V_i has exactly one neighbor in level V_{i-1} and this accounts for every edge not on the surface.
(3) For $i < d$, each vertex in V_i has degree $\ell + 1$.

V_0 is called the *surface* of the volcano, while V_d is called the *floor*.

The name "volcano" comes from the shape of the graph, where V_0 is a cycle, surrounded by a rapidly branching graph, which suggests an aerial view of volcanic crater surrounded by the mountainside. See Figure 6.1 for an example of a 3-volcano with $|V_0| = 5$ and $d = 2$.

It is clear that the conditions defining a volcano are related to the horizontal and vertical isogenies we discussed in Theorem 6.4. Take f prime to ℓ, and take W_i to be the set of curves (over k) with endomorphism ring $\mathbb{Z} + f\ell^i\mathcal{O}_K$. Then, $E \in W_0$ has at most two horizontal isogenies, with all others mapping to curves in W_1, corresponding to condition (1). For $i > 0$, $E \in W_i$ has one isogeny to a

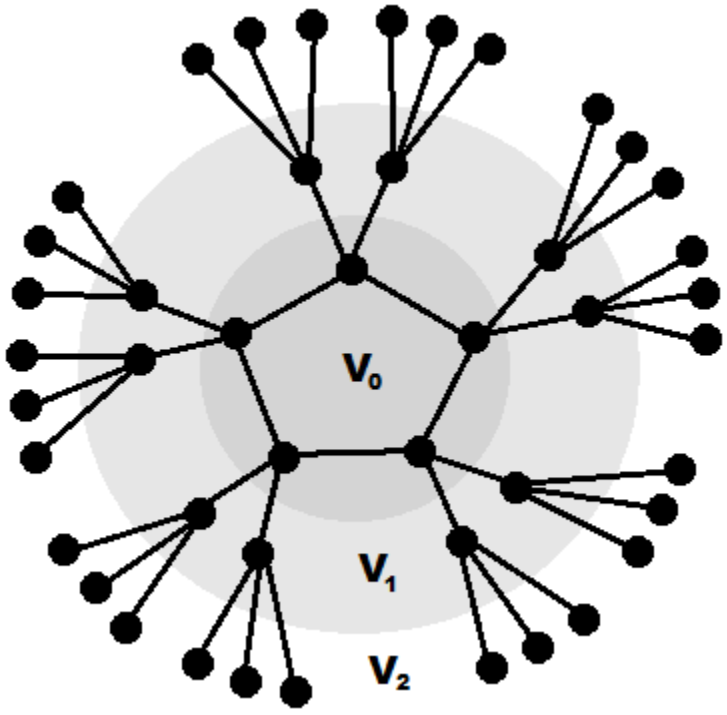

FIGURE 1. A 3-volcano with $|V_0| = 5$ and $d = 2$.

curve in W_{i-1} and ℓ isogenies to curves in W_{i+1} - unless W_{i+1} is empty, in which case we set $i = d$. If k is a finite field or number field, some such d must exist as the curves are defined over larger and larger extensions. Moreover, as we saw in 6.5, if i is minimal with W_{i+1} empty, then for all $j > i$, W_j is empty as well. This demonstrates condition (2) and (3).

Finally, set $\mathcal{O} = \mathbb{Z} + f\mathcal{O}_K$. If there is an ℓ-isogeny between curves in W_0, then it corresponds to some proper $\mathcal{O}$-ideal $\mathfrak{l}$ of norm ℓ, which together with its inverse in $C(\mathcal{O})$, determines all horizontal ℓ-isogenies between curves in W_0. Then, following the action of ℓ repeatedly, we see that the restriction of the ℓ-isogeny graph $G_\ell(k)$ to (the j-invariants of curves in) W_0 is a collection of disjoint cycles of order equal to the order of $[\mathfrak{l}]$ in $C(\mathcal{O})$.

Hence, recalling the definition of $\alpha_{K,p}$ from Definition 2.18, we have the following result, originally due to Kohel [3] and stated as in [10].

Theorem 6.8. *Let $\mathbb{F}_q$ be a finite field and let V be a component of $G_\ell(\mathbb{F}_q)$ containing the j-invariants of ordinary elliptic curves, but that does not contain 0 or 1728. Then, V is an ℓ-volcano of depth d for which the following hold:*

1. *The vertices in level V_i all have the same endomorphism ring $\mathcal{O}_i$.*
2. *The subgraph on V_0 has degree $1 + \alpha_{\mathrm{Frac}(\mathcal{O}_0),p}$.*
3. *If $\alpha_{\mathrm{Frac}(\mathcal{O}_0),p} \geq 0$, then $|V_0|$ is the order of $[\mathfrak{l}]$ in $C(\mathcal{O}_0)$ and otherwise $|V_0| = 1$.*
4. *$\ell \nmid [\mathcal{O}_K : \mathcal{O}_0]$ and $[\mathcal{O}_i : \mathcal{O}_{i+1}] = \ell$ for $0 \leq i < d$.*

(5) $d = \lfloor v_\ell((\mathrm{tr}(\pi)^2 - 4q))/2 \rfloor$ *where* π *is the Frobenius endomorphism of a curve with* j*-invariant in* V*.*

Proof. We have already proved everything except part 5. First, we note that for CM elliptic curves, the square of the trace of Frobenius is preserved by ℓ-isogenies. This follows from the fact proved when discussing the CM method in Section 5.5 that the integers t and v in the expression $4q = t^2 - v^2 D$, are unique up to sign for fixed q and D (if they exist). An ℓ-isogeny changes the discriminant D by multiplying it by either 1, ℓ^2, or ℓ^{-2}, whence v is multiplied by the inverse, and $t^2 = \mathrm{tr}(\pi)^2$ does not change.

Then, if we fix $t^2 = \mathrm{tr}(\pi)^2$, q, and the maximal factor of D that is relatively prime to ℓ, we are left with $\lfloor v_\ell((\mathrm{tr}(\pi)^2 - 4q))/2 \rfloor + 1$ solutions (up to the sign of t and v) for v and D, and so the floor of the volcano is at level $d = \lfloor v_\ell((\mathrm{tr}(\pi)^2 - 4q))/2 \rfloor$. □

As Sutherland notes in Remark 8 of [10], we exclude the cases of j-invariants 0 and 1728 because they correspond to the orders $\mathbb{Z}[i]$ and $\mathbb{Z}\left[\frac{1+\sqrt{-3}}{2}\right]$, which have additional units (besides ± 1) and are therefore slightly more complicated at level V_1, where there are triple (respectively double) edges going into each vertex from 0 (respectively 1728) at level V_1 and only single edges the other direction. Our counts on the size of V_1 are also off by a factor of approximately 3 (or 2) in these cases. Rather than discuss these special cases further, we note that the same algorithms that we discuss in the remainder of the paper apply in these cases as well.

6.2. **Computing With Isogeny Volcanoes.** In this section, we describe how isogeny volcanoes of elliptic curves can be used to facilitate computations involving elliptic curves. In particular, we focus on computing elliptic curves with a specified number of points (or equivalently a specified endomorphism ring via the CM method discussed in Section 5.5). We will find such curves via a more efficient algorithm for computing the Hilbert Class polynomials $H_{\mathcal{O}}(X)$. Then, computing its roots over our finite field gives us all of the curves with endomorphism ring $\mathcal{O}$, which gives us all of the desired curves. All of the algorithms we discuss, and several more, can be found in [10]. Sutherland also discusses the computational complexity in greater detail and presents more particular implementation details. We content ourselves to give a very high-level description of the algorithms and refer the reader interested in implementing these procedures or in their computational complexity to [10].

To start, we briefly describe a basic algorithms on ℓ-volcanoes in order to orient the reader with the structure of these graphs.

Proposition 6.9. *Given an* ℓ*-volcano* V*, the following algorithm finds a shortest path from* v *to the floor of* V*.*

(1) *Choose three distinct vertices adjacent to* v*. (If fewer than* $3 \leq \ell+1$ *adjacent vertices exist, return* v*.)*
(2) *Extend these adjacent vertices to paths with no backtracking (keeping each path the same length as the others) until one of them reaches a vertex with degree less than or equal to* 2*.*
(3) *Return the first such path.*

Proof. Note that a vertex is on the floor of V if and only if it has degree less than 3, since every vertex not on the floor has degree $\ell + 1 \geq 3$.

Now, if a edge in a path goes towards the floor (i.e. from $w \in V_i$ to $w' \in V_{i+1}$) and the path does not backtrack, every subsequent edge will also move towards the

floor, since the only edge from w' that does not go to V_{i+2} is the edge to w. Hence, if a path starts by moving from v towards the floor of V, when it terminates at the floor, it will be a shortest path from v to the floor.

Finally, since a vertex has at most two adjacent vertices that are not closer to the floor, given any three distinct adjacent vertices v, at least one will be descend towards the floor, which completes the proof of correctness. □

Remark 6.10. When we are working with an ℓ-isogeny volcano, it is worth noting that (given access to the modular polynomial $\Phi_\ell(X, Y)$,) we can efficiently compute neighboring vertices as follows.

Given a vertex $j(E)$, the neighboring vertices are the roots of the polynomial $\Phi_\ell(j(E), Y)$. According to the proof of Proposition 11 of [10], over a finite field $\mathbb{F}_q$, these can be computed efficiently enough that the computational bottle-neck is actually in substituting $j(E)$ into $\Phi_\ell(X, Y)$. Overall, Sutherland shows that if $M(n)$ is the time required to multiply two n-bit integers, then the previous algorithm can be computed in time $O(\ell^2 M(n) + M(\ell n)n)$.

According to the commentary immediately preceding Section 3.2 of [10], this algorithm is not state-of-the-art when the depth d is large - Ionica and Joux have recently developed a more efficient "pairing-based" approach to compute the distance from a vertex - but we hope that it has helped to acquaint the reader with the structure of volcanoes while giving some indication of their use in computations.

Now, as [10] points out, we can find the trace of Frobenius of an elliptic curve in polynomial time and so by (5) of Theorem 6.8, we can determine the depth of the volcano as well. Hence, this algorithm allows us to efficiently determine the power of ℓ dividing the discriminant of the endomorphism ring of our curve.

Running this procedure on each prime ℓ dividing $4q - (\operatorname{tr}\pi)^2$ gives an algorithm to compute the endomorphism ring of an elliptic curve E. Unfortunately, this algorithm is not very efficient if $4q - (\operatorname{tr}\pi)^2$ has large prime factors. [10] discusses how a more clever use of isogeny volcanoes can make this computation efficient.

6.2.1. *Computing Hilbert Class Polynomials.* At long last, we can describe how to use isogeny volcanoes to compute Hilbert Class polynomials $H_{\mathcal{O}}(X)$, which, as was discussed in Section 5.5 is a critical component of an efficient implementation of the CM method for finding elliptic curves with a specified number of points over a given finite field. Our discussion expands on 3.4 of [10]. We will assume throughout that the structure of the class group $C(\mathcal{O})$ is known.

We first recall that the roots of the Hilbert class polynomial modulo p are exactly the j-invariants of the elliptic curves $E/\overline{\mathbb{F}_p}$ with endomorphism ring $\operatorname{End}_{\mathbb{F}_p}(E) \cong \mathcal{O}$. Hence, if we can find the (j-invariants of) all of the elliptic curves over $\overline{\mathbb{F}_p}$ with endomorphism ring $\mathcal{O}$, then we can compute $H_{\mathcal{O}}(X) \pmod p$. If we repeat this computation for sufficiently many primes, then we can use the Chinese Remainder Theorem together with known bounds on the sizes of the coefficients of $H_{\mathcal{O}}(X)$ to compute $H_{\mathcal{O}}(X)$ over $\mathbb{Q}$.

Note that if $E/\mathbb{F}_q$ has endomorphism ring $\mathcal{O}$, where $\mathcal{O}$ is the order of discriminant D, then considering the Frobenius endomorphism $\pi_{\mathbb{F}_q}$,

$$q = N(\pi_{\mathbb{F}_q}) = \left(\frac{\operatorname{tr}(\pi_{\mathbb{F}_q}) + v\sqrt{D}}{2}\right)\left(\frac{\operatorname{tr}(\pi_{\mathbb{F}_q}) - v\sqrt{D}}{2}\right) = \frac{t^2 - v^2 D}{4},$$

so we need to find a large set of primes p together with some power $q = p^a$ satisfying $4q = t^2 - v^2 D$ for some integers t and v (that depend on q). So that we can work

in the finite field $\mathbb{F}_p$, which in many implementations has quicker arithmetic than more general finite fields, we prefer to use primes where $q = p^1$ We specialize to this case in what follows, although the algorithm also works for more general finite fields.

Now, given a prime p, and $q = p^a$ such that $4q = t^2 - v^2 D$, we wish to compute $H_{\mathcal{O}}(X) \pmod p$ by finding all of the j-invariants of elliptic curves with endomorphism ring $\mathcal{O}$.

First, we need to find at least one elliptic curve $E/\mathbb{F}_q$ with $\mathrm{End}_{\mathbb{F}_q}(E) = \mathcal{O}$. All we know about such curves is that they have trace of Frobenius $\pm t$. However, we also know that the endomorphism ring $\mathrm{End}_{\mathbb{F}_q}(E')$ of any other curve with trace of Frobenius $\pm t$ will also have $\mathrm{End}_{\mathbb{F}_q}(E) = \mathcal{O}'$ an order in the same imaginary quadratic field $\mathbb{Q}(\sqrt{t^2 - 4q}) = \mathbb{Q}(\sqrt{v^2 D}) = \mathbb{Q}(\sqrt{D})$. Hence, given a curve $E'/\mathbb{F}_q$ with trace of Frobenius $\pm t$, we can move to a curve $E/\mathbb{F}_q$ by moving along ℓ-isogeny volcanoes for the prime factors ℓ of $4q - t^2$. Using our depth finding algorithm, we can make sure that we always move in the right direction. (Note that the procedure we used to compute the endomorphisms will typically be efficient here, since we care most about the cases where D has few square divisors and where v is as small as possible.)

Finding an elliptic curve with trace of Frobenius $\pm t$ is simply a matter of choosing random curves and counting points via Schoof's algorithm. While this may seem somewhat circular - we are finding (many) curves with a specified number of points so that we can compute a polynomial that will allow us to compute more curves with a specified number of points - recall that the primes p that we are working with are much smaller than the primes of cryptographic size that we are interested in (and which may have hundreds of digits). Since the guess-and-check method for finding curves with a given number of points does not scale well, it could be impractical to carry out over a field with over 10^{100} elements, but still be efficient to run hundreds of times over fields of much smaller order.

Once we have found an elliptic curve $E/\mathbb{F}_q$ with $\mathrm{End}_{\mathbb{F}_q}(E) = \mathcal{O}$, we need to find all of the other elliptic curves with endomorphism ring $\mathcal{O}$. For this, we can again make use of the isogeny volcano structure.

We know that $C(\mathcal{O})$ acts simply transitively on $\mathcal{ELL}_{\mathbb{F}_q}(\mathcal{O})$. Moreover, the action of $[\mathfrak{a}] \in C(\mathcal{O})$ on $E \in \mathcal{ELL}_{\mathbb{F}_q}(\mathcal{O})$ is the same as following the appropriate horizontal $N(\mathfrak{a})$-isogeny. Then, if $C(\mathcal{O})$ is generated by $[\mathfrak{a}_1], \ldots, [\mathfrak{a}_d]$, where each $\mathfrak{a}_i$ has small prime norm, we can use our depth-finding algorithm to efficiently walk the surface of these $N(\mathfrak{a}_i)$-volcanoes until we have listed all $|C(\mathcal{O})|$ elements of $\mathcal{ELL}_{\mathbb{F}_q}(\mathcal{O})$.

Remark 6.11. In practice, as Sutherland notes in 3.4 of [10], when enumerating $\mathcal{ELL}_{\mathbb{F}_q}(\mathcal{O})$, it is important to choose a good presentation. For a naive implementation, the best choice of generators is typically given by a polycyclic presentation for $C(\mathcal{O})$ where the each generator is chosen to have the minimum possible norm among elements not in the span of the preceding generators. Such a norm will always be prime. This minimizes the primes ℓ for the modular polynomials $\Phi_\ell(X, Y)$ involved in the computations.

Note that this presentation only needs to be computed once in the computation of $H_{\mathcal{O}}(X)$ and can be used for all of the prime moduli, so it is worth the effort to find a good presentation.

Once we have found all j-invariants of curves in $\mathcal{ELL}_{\mathbb{F}_q}(\mathcal{O})$, it is a simple matter to compute that in $\mathbb{F}_q$ and therefore modulo p,

$$H_{\mathcal{O}}(X) \equiv \prod_{E \in \mathcal{ELL}_{\mathbb{F}_q}(\mathcal{O})} (X - j(E)) \pmod{p}.$$

Repeating the computation for many primes and applying the Chinese Remainder Theorem, we can compute $H_{\mathcal{O}}(X)$ over $\mathbb{Q}$.

Remark 6.12. Recall that in Section 5.5, the main obstruction to applying the CM method to compute elliptic curves was that we needed to be able to compute $\Phi_D(X, X)$ (and possibly also $\Phi_{4D}(X, X)$) in order to compute $H_{\mathcal{O}}(X)$. Using the isogeny volcano-based approach presented here, if

$$\begin{aligned} S = &\{p' = N(\mathfrak{a}_i) \text{ for some } \mathfrak{a}_i \text{ in our presentation of } C(\mathcal{O})\} \\ &\cup \{p' : (p')^2 | 4q - t^2 \text{ for some } q \text{ in our collection}\} \end{aligned}$$

then we only need to know Φ_p for $p \in S$. Typically, $\max(S)$ is much smaller than D, leading to huge computational savings.

When implementing this algorithm, a number of tricks, ranging from better methods for walking the surface of the volcanoes from our polycyclic presentation to using alternate versions of the modular polynomials can lead to substantial savings. It is also possible to use a similar algorithm to compute modular polynomials, which can then be used in turn to compute even larger modular polynomials or Hilber class polynomials. See sections 3.4 and 3.5 of [10] for further details.

To summarize our discussion, we state the algorithm that we have presented for computing the Hilbert class polynomial $H_{\mathcal{O}}(X)$, essentially quoting from 3.4 of [10].

Proposition 6.13. *Given an order $\mathcal{O}$ of discriminant D in the imaginary quadratic field K, the following steps compute $H_{\mathcal{O}}(X)$.*

1. *Find a large set of primes p satisfying $4p^a = t_p^2 - v_p^2 D$ for some integers t, v, a depending on p.*
2. *For each prime p, compute $H_{\mathcal{O}}(X)$ over $\mathbb{F}_{p^a}$ (equivalently mod p) as follows:*
   (a) *Search at random to find an elliptic curve $E/\mathbb{F}_{p^a}$ with $|E(\mathbb{F}_{p^a})| = p + 1 \pm t$.*
   (b) *Use the volcano depth-finding algorithm and appropriate isogenies to find a curve $E'/\mathbb{F}_{p^a}$ with $\mathrm{End}_{\mathbb{F}_{p^a}}(E') = \mathcal{O}$.*
   (c) *Use the $C(\mathcal{O})$-action, together with suitable presentation of $\mathcal{O}$ to enumerate $\mathcal{ELL}_{\mathbb{F}_{p^a}}(\mathcal{O})$ by walking the surface of various ℓ-isogeny volcanoes.*
   (d) *Compute $H_{\mathcal{O}}(X) \equiv \prod_{E \in \mathcal{ELL}_{\mathbb{F}_{p^a}}(\mathcal{O})}(X - j(E)) \pmod{p}$*
3. *Compute $H_{\mathcal{O}}(X)$ by the Chinese Remainder Theorem, using bounds on the coefficients.*

We now give an extended example of step 2 of the algorithm from Proposition 6.13, demonstrating the computation of $H_{\mathcal{O}_{\mathbb{Q}(\sqrt{-251})}}(X) \pmod{6311}$ starting from an elliptic curve over $\mathbb{F}_{6311}$ with trace of Frobenius ± 12.

Example 6.14. Let $K = \mathbb{Q}(\sqrt{-251})$ and $\mathcal{O} = \mathcal{O}_K$ be the order of discriminant 251. The prime 6311 satisfies $4 \cdot 6311 = 12^2 + 10^2 \cdot 251$. We will compute $H_{\mathcal{O}}(X) \pmod{6311}$.

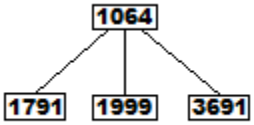

FIGURE 2. The connected component of 1999 in $G_2(\mathbb{F}_{6311})$.

For convenience, throughout this example, we let E_j denote the ($\overline{\mathbb{F}_{6311}}$-isomorphism class) of curves with j-invariant j.

To start, we a curve with trace of Frobenius ± 12. By random sampling, we find that we can take the curves E with $j(E) = 1999$. This completes step 2a.

Now, since $4 \cdot 6311 - 12^2 = 10^2 \cdot d_K$, we know that the endomorphism ring of E is an order $\mathcal{O}$ of conductor f dividing 10. Since $2^1|10$ exactly, the 2-volcanoes consisting orders of K will have depth $d = 1$, and similarly for the 5-volcanoes. Hence, every vertex that is not a leaf is on the surface, so we need to find a j-invariant that is not a leaf of a 2-isogeny-volcano or a 5-isogeny volcano.

First, we consider the 2-isogenies. In $\mathbb{F}_{6311}$,

$$\begin{aligned}\Phi_2(X, 1999) &\equiv X^3 + 2979X^2 + 1583X + 1969 \\ &\equiv (X - 1064)(X^2 - 2268X - 767)\end{aligned}$$

so 2 divides the conductor of $\mathrm{End}_{\mathbb{F}_{6311}}(E_{1999})$, but will not divide the conductor of $\mathrm{End}_{\mathbb{F}_{6311}}(E_{1064})$. As a double check

$$\Phi_2(X, 1999) \equiv (X - 1791)(X - 1999)(X - 3691)$$

so, indeed, 1999 is not a leaf and must therefore be on the surface of the 2-volcano.

For good measure, we can "map" the connected component of 1999, resulting in the 2-volcano in Figure 6.14.

Next, we consider the 5-isogenies. In $\mathbb{F}_{6311}$,

$$\begin{aligned}\Phi_5(X, 1064) &= X^6 - X^5 - 614X^4 - 1857X^3 + 2648X^2 + 1337X + 750 \\ &= (X - 4492)(X^5 - 1820X^4 + 3002X^3 + 2831X^2 + 2835X + 559),\end{aligned}$$

so $1064 \in V_1$, while it's neighbor 4492 will be on the surface. Hence, our original curve E_{1999} had $\mathrm{End}_{\mathbb{F}_{6311}}(E_{1999}) = \mathbb{Z} + 10\mathcal{O}_K$, while $\mathrm{End}_{\mathbb{F}_{6311}}(E_{1064}) = \mathbb{Z} + 5\mathcal{O}_K$ and finally, $\mathrm{End}_{\mathbb{F}_{4492}}(E_{4492}) = \mathcal{O}_K$. This completes step 2b.

Again, for good measure, we can "map" the connected component of 1064, resulting in the 5-volcano in Figure 6.14.

Now, we need to choose a presentation for $C(\mathcal{O})$. Since $|C(\mathcal{O})| = 7$ is prime, the group is cyclic and any element is a generator. Now, since $\mathcal{O}_K$ has an element $1 + \sqrt{-251}$ not divisible by 3, but with norm $N(1 + \sqrt{-251}) = 252$ divisible by 3, some element of $\mathcal{O}_K$ has a representative ideal of norm 3. Hence, we can apply the $C(\mathcal{O})$ action by following 3-isogenies. Even better, since $3 \nmid 10$, we know that we will always remain on the surface of the volcano, which will have depth 0.

Starting with our curve E_{4492}, we compute that

$$\begin{aligned}\Phi_3(X, 4492) &= X^4 - 1124X^3 - 1735X^2 + 2035X - 2964 \\ &= (X - 4357)(X - 1830)(X^2 - 1248X - 964).\end{aligned}$$

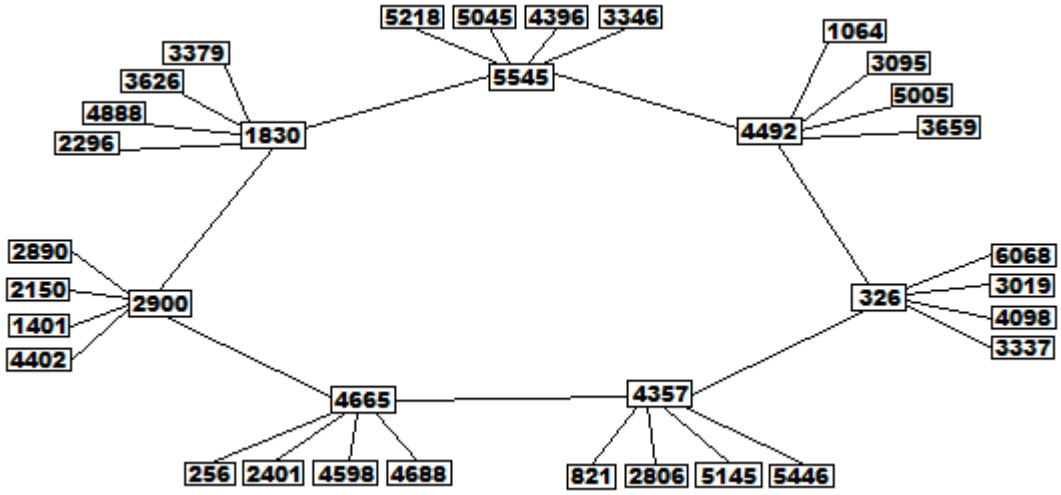

FIGURE 3. The connected component of 1064 in $G_5(\mathbb{F}_{6311})$.

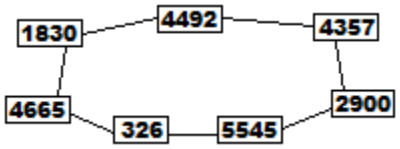

FIGURE 4. Mapping $\mathcal{O}_K$ for $K = \mathbb{Q}(\sqrt{-251})$ in $G_3(\mathbb{F}_{6311})$.

Choosing a root arbitrarily (this corresponds to choosing the generator or it's inverse), we find the elliptic curve E_{4357}. Repeating the process, we compute

$$\begin{aligned}\Phi_3(X, 4357) &= X^4 + 847X^3 - 1015X^2 - 2417X + 275 \\ &= (X - 4492)(X - 2900)(X^2 + 1928X - 373).\end{aligned}$$

We already visited the curve E_{4492}, so we next use the curve E_{2900}. Continuing in this manner, we see

$$\begin{aligned}\Phi_3(X, 2900) &= (X - 5545)(X - 4357)(X^2 + 703X + 161) \\ \Phi_3(X, 5545) &= (X - 2900)(X - 326)(X^2 - 1763X - 2896) \\ \Phi_3(X, 326) &= (X - 5545)(X - 4665)(X^2 + 240X - 981) \\ \Phi_3(X, 4665) &= (X - 1830)(X - 326)(X^2 + 1461X + 491) \\ \Phi_3(X, 1830) &= (X - 4665)(X - 4492)(X^2 - 600X - 835),\end{aligned}$$

and so

$$\mathcal{ELL}_{\mathbb{F}_{6311}}(\mathcal{O}_K) = \{E_{4492}, E_{4357}, E_{2900}, E_{5545}, E_{326}, E_{4665}, E_{1830}\},$$

which completes step 2c.

Mapping as before, we have the 3-volcano in Figure 6.14.

Finally, it is an easy matter to compute

$$\begin{aligned} H_{\mathcal{O}_K}(X) &= (X-4492)(X-4357)(X-2900)(X-5545)(X-326)(X-4665)(X-1830) \\ &= X^7 + 1129X^6 + 1073X^5 + 815X^4 + 5036X^3 + 2823X^2 + 5674X + 849, \end{aligned}$$

which completes part 2d of the algorithm.

Remark 6.15. We will not always be so lucky as in Example 6.14 when computing part 2c of the algorithm from 6.13. In general, the volcano will not have depth $d = 0$. For instance, as is clear from Figure 6.14, if we had chosen to generate $C(\mathcal{O})$ with an ideal of norm 5, at each stage of our walk, we would have had to check that our vertex remained on the surface using our height finding algorithm or some similar procedure.

Remark 6.16. It is worth noting that to compute $H_{\mathcal{O}}(X) \pmod{6311}$, we only needed to know Φ_2, Φ_3, and Φ_5. (In fact, we could have completed the computation without Φ_3!)

While Example 6.14 was "cooked up" in the sense that the order was chosen with prime class number (for simplicity of exposition) and E_{1999} was chosen to have endomorphism ring of conductor 10 to illustrate the process of finding the correct endomorphism ring, these computations are illustrative of the usual procedure for computing Hilbert Class polynomials. The savings in the size of the modular polynomials are tremendous!

For more examples of isogeny volcanoes, complexity analysis, various implementation details, and other computations involving isogeny volcanoes, see [10].

7. Conclusions

Throughout this paper, we have seen the power of complex multiplication. In characteristic zero, the elliptic curves with complex multiplication - i.e. with endomorphisms other than multiplication-by-an-integer - are quite rare. As we have seen, the endomorphism rings are always orders in imaginary quadratic fields, and the elliptic curves correspond to proper ideals in these orders. This suggested that the ideal class groups of the orders and these elliptic curves should be related. Indeed, this relationship is incredibly strong. As we showed, in Corollary 5.16 and Theorem 5.17 (the First Main Theorem of Complex Multiplication), if $\mathrm{End}_{\mathbb{C}}(E) = \mathcal{O}$ an order in the imaginary quadratic field K, then $K(j(E))$ is the Hilbert class field of the order $\mathcal{O}$, providing most of the modern progress towards Hilbert's 12th problem - to classify abelian extensions of number fields. Moreover, we saw that $\mathrm{Gal}(K(j(E))/K) \cong C(\mathcal{O})$ by a group isomorphism that preserves the group actions on $\mathcal{ELL}_{\mathbb{C}}(\mathcal{O})$.

Along the way, we developed a pair of very important concrete tools: the modular polynomials $\Phi_m(X, Y)$, which parameterize pairs of elliptic curves related by a cyclic m-isogeny and the Hilbert class polynomials, $H_{\mathcal{O}}(X)$, which have as their roots the j-invariants of all curves with endomorphism ring $\mathcal{O}$ (over an algebraically closed field). Remarkably, from the results of Deuring discussed in Section 3.4, we were able to conclude that these interpretations remain valid over *any* field, not merely fields of characteristic zero. This had important computational consequences. Since we saw that the elliptic curves with $q + 1 - t$ points over the finite field $\mathbb{F}_q$ are those curves whose endomorphism ring is an order in $K(\sqrt{t^2 - q})$, we

could use knowledge of the Hilbert class polynomials $H_{\mathcal{O}}(X)$ to find such curves. This is the "CM method" discussed in Section 5.5.

Until recently, if was difficult to compute Hilbert class polynomials, because traditional algorithms required knowledge of $\Phi_m(X, Y)$ for large m. These polynomials grow very quickly and were impractical either to store in memory or to recompute as needed. Starting with the observation in Kohel's PhD Thesis [3] that the graph of elliptic curves related by ℓ-isogenies has a very particular structure, efficient algorithms were developed to compute $H_{\mathcal{O}}(X) \pmod p$ for small primes p. Again applying the work of Deuring, this allows us to find $H_{\mathcal{O}}(X)$ - and therefore all curves with a given endomorphism ring. We discussed one such algorithm and the structure of this ℓ-isogeny graph in Section 6.

While we have discussed answers to many important questions, ranging from the purely theoretical to the computational, many questions still remain. The Main Theorems of Complex Multiplication have been generalized to certain abelian varieties, but in the vast majority of cases, Hilbert's 12th problem remains incredibly open. Given an arbitrary number field K, nobody knows how to classify it's abelian extensions.

On the computational side, there is also much work left to do. It seems natural to expect that a similar structure would exist for the analogue of ℓ-isogenies of hyperelliptic curves, or possibly even more general abelian varieties. Much work still needs to be done to determine what structure exists in this case and how it can be exploited computationally - say to construct hyperelliptic curves with a given number of points.

In any case, it is clear that Complex Multiplication and Isogeny Volcanoes are important topics in the modern world of number theory and arithmetic geometry. They will likely to remain major areas of research for the foreseeable future.

8. Acknowledgements

I would like to thank the following people/organizations for their support in writing this paper:

Kevin Sackel, for being a willing sounding board during my early struggles coming to grips with the material presented here.

Professor Tom Fisher, for several useful conversations that played an important role in helping me to organize this essay, from determining which parts were particularly important, to which are the most difficult to understand, to which could be left out entirely. I would also like to thank Dr. Fisher for setting this essay in the first place, as the wonderful topic made the writing a pleasure.

John Bootle, with whom I had several productive exchanges leading up to the Part III Seminars, and whose excellent talk helped me to realize how much there was to say about orders in Imaginary Quadratic Fields and how important it was to say these things. Our correspondence about possible topics for the part III seminar helped me to make sure that I covered all of the major topics and proved extremely helpful when outlining this essay.

Jack Lamplugh, for chairing the Part III Number Theory Seminar, and for a useful conversation about my talk that gave me the necessary impetus to start on the project of writing this essay.

Dr. Marjorie Batchelor, whose advice on scheduling one's time for writing a Part III essay was extremely helpful in making sure I alloted enough time despite the inevitable numerous delays.

My cousins, Nazzareno, Maria, and Giuseppe Mariucci, who were wonderfully accepting of the fact that I needed to work during my visit to their house in early April, despite my limited time with them.

My parents, George Triantafillou and Jean Farrington, who were happy to talk to me and provide encouragement when I took much-needed breaks from writing and studying and who have always supported me through my every endeavor.

The Churchill Scholarship Foundation of America, which made it possible for me to study Maths Part III at Cambridge in the first place.

The National Science Foundation, which is supporting my current and future studies through an NSF Graduate Fellowship.

References

[1] David A. Cox. *Primes of the Form x + ny: Fermat, Class Field Theory, and Complex Multiplication (Pure and Applied Mathematics: A Wiley Series of Texts, Monographs and Tracts)*. Wiley-Interscience, 1989.

[2] J. Igusa. Kroneckerian model of fields of elliptic modular functions. *Am. J. Math.*, 81:561–577, 1959.

[3] D. Koehl. *Endomorphism rings of elliptic curves over finite fields*. PhD thesis, University of Calfornia, Berkeley, 1996. Accessed on 15 April 2014 at `http://iml.univ-mrs.fr/~kohel/pub/index.html`.

[4] Serge Lang. *Elliptic Functions (Graduate Texts in Mathematics, Vol. 112)*. Springer, 1987.

[5] Jürgen Neukirch. *Algebraic Number Theory (Grundlehren der mathematischen Wissenschaften) (v. 322)*. Springer, 1999.

[6] Brian Osserman. Orders and their class groups.

[7] Joseph H. Silverman. *The Arithmetic of Elliptic Curves (Graduate Texts in Mathematics) (v. 106)*. Springer, 1994.

[8] Joseph H. Silverman. *Advanced Topics in the Arithmetic of Elliptic Curves (Graduate Texts in Mathematics)*. Springer, 1999.

[9] A. V. Sutherland. 18.783 elliptic curves lecture notes. `http://math.mit.edu/classes/18.783/lectures.html`, accessed January-April 2014, May 2013.

[10] A. V. Sutherland. Isogeny volcanoes. arXiv:1208.5370v3. *ArXiv e-prints*, May 2013.

Department of Mathematics, University of Cambridge, Cambridge, Cambridgeshire, UK CB3 0WB

E-mail address: `ngt24@cam.ac.uk`