HAL
open science

Isogeny graphs, modular polynomials, and applications

Chloe Martindale

To cite this version:

Chloe Martindale. Isogeny graphs, modular polynomials, and applications. Number Theory [math.NT]. Université de Bordeaux; Universiteit Leiden (Leyde, Pays-Bas), 2018. English. NNT : 2018BORD0086 . tel-01992715

HAL Id: tel-01992715

https://theses.hal.science/tel-01992715v1

Submitted on 24 Jan 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE

présentée à

L'UNIVERSITÉ DE BORDEAUX

ÉCOLE DOCTORALE DE MATHÉMATIQUES ET INFORMATIQUE

par Chloe Martindale

POUR OBTENIR LE GRADE DE

DOCTEUR

SPECIALITÉ: Mathématiques Pures

Isogeny Graphs, Modular Polynomials, and Applications

Soutenue le : 14 juin 2018 à Leiden

Devant la commission d'examen formée de :

ENGE, Andreas	Professeur	Université de Bordeaux	Directeur
STRENG, Marco	Docteur	Universiteit Leiden	Directeur
KOHEL, David	Professeur	Université de Aix-Marseille	Rapporteur
JETCHEV, Dimitar	Docteur	École polytechnique fédérale de Lausanne	Rapporteur

Ce travail a été financé par Algant-Doc Erasmus Action et a été réalisé à l'Universiteit Leiden et à l'Université de Bordeaux.

Contents

Introduction **x**

1 The theory of canonical lifts and other preliminaries **1**
1.1 Principally polarised abelian varieties . . . 1
1.2 Lifting ordinary abelian varieties over $\mathbb{F}_q$ to ideals . . . 2
1.3 The Fixed Frobenius Lifting Theorem . . . 3
1.4 The theory of canonical lifts . . . 6
1.4.1 Serre-Tate lifts of ordinary abelian varieties . . . 7
1.4.2 Deligne lifts of ordinary abelian varieties . . . 8
1.4.3 Howe lifts of polarised ordinary abelian varieties . . . 9
1.4.4 Proof of the Fixed Frobenius Lifting Theorem . . . 11
1.5 Maximal real multiplication . . . 13
1.6 Hilbert modular forms . . . 14
1.7 A normalisation lemma for principally polarised ideals . . . 17

2 Hilbert modular polynomials **19**
2.1 Introduction and statement of the results . . . 19
2.2 Defining RM isomorphism invariants . . . 21
2.3 Algorithm to compute a set of Hilbert modular polynomials . . . 23
2.4 Computing the RM isomorphism invariants for a given genus 2 curve . . . 28
2.4.1 The algorithm . . . 31
2.5 Complexity and simplifications for genus 2 . . . 34

3 The structure of μ-isogeny graphs **38**
3.1 The Volcano Theorem . . . 38
3.2 Parametrising orders by their real conductors . . . 46
3.3 All μ-isogenies are ascending, descending or horizontal . . . 48
3.4 Principally polarised ideals are invertible . . . 50
3.5 The action of the Shimura class group . . . 53
3.6 Counting horizontal μ-isogenies . . . 55
3.7 A construction of ascending μ-isogenies . . . 58
3.8 Counting the degree of vertices in the μ-isogeny graph . . . 59
3.9 The order of the Shimura class group . . . 60
3.10 Example computation of a μ-isogeny graph . . . 66

4 Isogenies for point counting on genus two hyperelliptic curves with maximal real multiplication **51**
4.1 Introduction . . . 52
4.1.1 The state of the art . . . 53
4.1.2 Our contributions, and beyond . . . 53
4.1.3 Vanilla abelian varieties . . . 54
4.2 Genus one curves: elliptic curve point counting . . . 55
4.2.1 Schoof's algorithm . . . 56
4.2.2 Frobenius eigenvalues and subgroups . . . 57
4.2.3 Modular polynomials and isogenies . . . 57
4.2.4 Elkies, Atkin, and volcanic primes . . . 58

4.2.5 Computing the type of a prime . . . 59
4.2.6 Atkin's improvement . . . 59
4.2.7 Elkies' improvement . . . 60
4.3 The genus 2 setting . . . 61
4.3.1 The Jacobian . . . 61
4.3.2 Frobenius and endomorphisms of $J_{\mathcal{C}}$ . . . 62
4.3.3 Real multiplication . . . 62
4.3.4 From Schoof to Pila . . . 62
4.3.5 The Gaudry–Schost approach . . . 63
4.3.6 Point counting with efficiently computable RM . . . 64
4.3.7 Generalizing Elkies' and Atkin's improvements to genus 2 . . . 65
4.3.8 μ-isogenies . . . 66
4.4 Invariants . . . 66
4.4.1 Invariants for RM abelian surfaces . . . 67
4.4.2 Hilbert modular polynomials for RM abelian surfaces . . . 67
4.4.3 Invariants for curves and abelian surfaces . . . 68
4.4.4 Pulling back curve invariants to RM invariants . . . 70
4.5 Atkin theorems in genus 2 . . . 70
4.5.1 Roots of G_μ and the order of Frobenius . . . 70
4.5.2 The factorization of G_μ . . . 72
4.5.3 The characteristic polynomial of Frobenius . . . 73
4.5.4 Prime types for real multiplication by $\mathcal{O}_F$ . . . 74
4.5.5 The parity of the number of factors of G_μ . . . 75
4.6 The case $F = \mathbb{Q}(\sqrt{5})$: Gundlach–Müller invariants . . . 75
4.7 Experimental results . . . 77

A The notions of dual and polarisation in equivalent categories **80**

Bibliography **89**

Index **92**

Summary **92**

Samenvatting **97**

Résumé **103**

Acknowledgements **109**

Curriculum Vitae **110**

Introduction

Background

Algebraic curves have been studied in various forms for thousands of years, yet still today there are many unsolved problems relating to the subject. In the 19th century, Abel and Jacobi transformed this subject by associating to any algebraic curve its *Jacobian*, which is an additive group containing the curve itself, and which is in particular an example of an abelian variety.

Possibly the most studied algebraic curves are *elliptic curves*, the Jacobians of which are isomorphic to the elliptic curves themselves, meaning that a group law can be defined directly on the curve. One consequence of this nice property is that elliptic curves lend themselves in a natural way to modern cryptographic algorithms, as the rational points on an elliptic curve form a group.

When studying maps between abelian varieties, we will restrict to *isogenies*, which are surjective homomorphisms with finite kernel. In particular, they preserve the identity.

In curve-based cryptography, it is important to develop fast algorithms for computing isogenies, for computing endomorphism rings, and for counting points on curves defined over finite fields $\mathbb{F}_p$, where p is a large prime number. There are many elliptic curve algorithms that have been developed in recent years to this end, and due to their geometric nature, one may ask if these algorithms can be applied to more general algebraic curves (by studying their Jacobians).

Overview

In Chapter 1, we give a (polarisation-preserving) equivalence of categories between abelian varieties defined over a finite field with a given characteristic polynomial of Frobenius and ideals of an order in a number field. The main applications of the results of this thesis concern (Jacobians of) curves defined over finite fields, but in many cases it is much easier to prove theoretical results for ideals than for abelian varieties. We call the statement of this equivalence of categories the Fixed Frobenius Lifting Theorem (Theorem 1.3.11). This equivalence is well-known and often-used, but the precise statement does not, to our knowledge, appear in the literature. We prove it as a consequence of similar, but more general, results of Deligne and Howe, which use the theory of canonical lifts of Lubin, Serre, and Tate.

In Section 1.6, we also give an introduction to the theory of Hilbert modular forms, which we will need in Chapter 2.

In Chapter 2, we give a generalisation of the modular polynomial for elliptic curves. A modular polynomial makes use of the *j-invariant* of an elliptic curve. For a field k and an elliptic curve E/k of the form

$$y^2 = x^3 + Ax + B,$$

with $A, B \in k$, the j-invariant is defined by

$$j(E) = 1728\frac{4A^3}{4A^3 + 27B^2},$$

and determines the elliptic curve uniquely up to $\overline{k}$-isomorphism. For a prime ℓ, we will refer to an isogeny of degree ℓ as an *ℓ-isogeny*. There is an irreducible polynomial

$$\Phi_\ell(X, Y) \in \mathbb{Z}[X, Y]$$

called the *modular polynomial* such that given elliptic curves E and E' over a field k there exists an ℓ-isogeny $E \to E'$ over $\overline{k}$ if and only if

$$\Phi_\ell(j(E), j(E')) = 0.$$

We generalise the modular polynomial for elliptic curves to a tuple of modular polynomials for principally polarised ordinary abelian varieties with real multiplication by the maximal order of a given number field K_0. We can think of elliptic curves defined over $\mathbb{C}$ as points in the moduli space $\mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}$, and the modular polynomials for elliptic curves can be computed using this interpretation. There are two common generalisations of this moduli space: Siegel moduli space, which parametrises principally polarised abelian varieties, and Hilbert moduli space, which parametrises principally polarised abelian varieties with real multiplication by a fixed number field K_0. The generalisation of modular polynomials to abelian varieties using Siegel moduli space was studied by Dupont [Dup06]. However, even for abelian surfaces, most practical applications of these Siegel modular polynomials are computationally out of reach, as the smallest example is already 25.6MB. We give a generalisation of modular polynomials using Hilbert moduli space. Included in this generalisation is a generalisation of the j-invariant to *RM isomorphism invariants* for principally polarised abelian varieties with maximal real multiplication by a fixed number field K_0. Theoretically, it is relatively easy to show that such invariants exist, and computationally, we use the formulae given by Müller for $K_0 = \mathbb{Q}(\sqrt{5})$ [Mue83] and $K_0 = \mathbb{Q}(\sqrt{2})$ [Mue85]. The main contribution of this chapter is an algorithm to compute these Hilbert modular polynomials, and we have implemented this algorithm in MAGMA for $K_0 = \mathbb{Q}(\sqrt{5})$. The resulting Hilbert modular polynomials are much more manageable than the Siegel equivalent, although the algorithm is very slow so computing higher levels would still require some work.

In Chapter 3, we give a generalisation of Kohel's structure theorem for isogeny graphs. In his PhD thesis [Koh96], David Kohel studied the structure of *isogeny graphs* of elliptic curves. An ℓ-isogeny graph of elliptic curves is an undirected graph for which each vertex represents a j-invariant of an elliptic curve over a field k, and an edge between $j(E)$ and $j(E')$ represents a pair of ℓ-isogenies between E and E' that are dual to each other (up to isomorphism).

Kohel gave a structure theorem for ℓ prime and ordinary $E/\mathbb{F}_q$ (with special cases occurring at $j(E) = 0$ and 1728). Among other things, Kohel's structure theorem is a key component in efficiently computing the endomorphism ring of an ordinary elliptic curve over $\mathbb{F}_q$.

We generalise Kohel's theorem to a structure theorem for isogeny graphs of principally polarised ordinary abelian varieties over $\mathbb{F}_q$ with real multiplication by the maximal order of a fixed real number field K_0. The isogeny graphs we study in this thesis are graphs of isogenies depending on a parameter μ, which is a totally positive element of K_0 that generates a prime ideal in $\mathcal{O}_{K_0}$. In Section 3.1, we state the main theorem of this chapter, the Volcano Theorem, and the rest of the chapter is dedicated to the proof, except for Section 3.10, in which we give an example computation of an isogeny graph. This problem has also been studied by Ionica and Thomé [IT14], who give a structure theorem for Jacobians of curves of genus two with real multiplication by the maximal order of a fixed real quadratic number field of narrow class number 1, and in parallel to the work in this thesis, by Brooks, Jetchev, and Wesolowski [BJW17], who also prove Theorem 3.1.9, using different methods, with the added assumption that the CM-type is primitive. Brooks, Jetchev, and Wesolowski also studied the structure of isogeny graphs for which the isogenies depend on a parameter $\mathfrak{l}$, a prime ideal in $\mathcal{O}_{K_0}$ which is not necessarily generated by a totally positive element $\mu \in \mathcal{O}_{K_0}$.

Chapter 4 is a joint article [Bal+17] with Ballentine, Guillevic, Lorenzo-García, Massierer, Smith, and Top, in which we generalise the Atkin–Elkies–Schoof algorithm to count points on elliptic curves over finite fields. The Atkin–Elkies–Schoof algorithm makes use of factorisation patterns of modular polynomials to give a polynomial time algorithm for counting points on elliptic curves. We give a polynomial time algorithm to count points on genus 2 curves over a finite field with real mutiplication by the maximal order a fixed number field K_0 using the factorisation patterns of the Hilbert modular polynomials of Chapter 2.

Appendix A gives the technical category-theoretical details necessary for the proof of the equivalence of categories given in Chapter 1, Theorem 1.3.11.

Chapter 1

The theory of canonical lifts and other preliminaries

In much of this thesis we will study principally polarised ordinary abelian varieties over $\mathbb{F}_q$, where q is a prime or a power of a prime. In this chapter, we specialise results of Deligne and Howe that allow us to work with ideals and elements of CM-fields instead of with varieties over $\mathbb{F}_q$. The proofs of these results are based on the lifting theorems of Lubin, Serre and Tate. The main theorem of this chapter, the Fixed Frobenius Lifting Theorem (Theorem 1.3.11), is an equivalence between two categories, so we now proceed by defining these categories.

1.1 Principally polarised abelian varieties

We first summarise some preliminaries on abelian varieties. For details on this subject there are many good textbooks, for example Mumford's book [Mum08].

Definition 1.1.1. An *abelian variety* A over a field k is a complete group variety over k.

Remark 1.1.2. If A is an abelian variety defined over $\mathbb{C}$ then $\mathrm{A}(\mathbb{C})$ is complex analytically isomorphic to a complex torus.

Definition 1.1.3. An *isogeny* is a morphism of abelian varieties that is finite as a morphism of varieties and surjective. The *degree* of an isogeny is its degree as a morphism of varieties.

Definition 1.1.4. For an abelian variety A over a field k, we define the *Picard group* of A, written as $\mathrm{Pic}(\mathrm{A})$, to be the group of isomorphism classes of line bundles on A.

Proposition 1.1.5. For an abelian variety A over a field k and a line bundle $\mathcal{L}$ on A, the map defined by

$$\begin{array}{rccc} \phi_{\mathcal{L}}: & \mathrm{A}(k) & \longrightarrow & \mathrm{Pic}(\mathrm{A}) \\ & x & \mapsto & [T_x^*\mathcal{L} \otimes \mathcal{L}^{-1}], \end{array}$$

where T_x denotes translation by x and $[\cdot]$ denotes the isomorphism class of $\cdot$ in $\mathrm{Pic}(\mathrm{A})$, is a homomorphism.

Proof. This follows from the Theorem of the Square, see e.g. [GM07, Corollary 2.10]. □

Definition 1.1.6. For an abelian variety defined over an algebraically closed field k, we define $\mathrm{Pic}^0(\mathrm{A})$ to be the subgroup of $\mathrm{Pic}(\mathrm{A})$ consisting of classes of line bundles $\mathcal{L}$ such that the morphism $\phi_{\mathcal{L}}$ is identically 0.

Proposition 1.1.7. Given an abelian variety A over an algebraically closed field k, the group $\mathrm{Pic}^0(\mathrm{A})$ carries a canonical structure of an abelian variety over k.

Proof. See [Mum08, Chapter III, Corollary 5]. □

Definition 1.1.8. Given an abelian variety A over an algebraically closed field k, we define the *dual abelian variety* $\mathrm{A}^\vee$ of A to be $\mathrm{Pic}^0(\mathrm{A})$.

Remark 1.1.9. To define the dual abelian variety $\mathrm{A}^\vee$ of an abelian variety A over an arbitrary field k, we need some basic theory of schemes. We have omitted that here for simplicity; the interested reader can refer to [MFK94, Chapter 6].

Proposition 1.1.10. Given an abelian variety A over an algebraically closed field k, if $\mathcal{L}$ is an ample line bundle on A, then we associate to $\mathcal{L}$ an isogeny of abelian varieties $\phi_{\mathcal{L}} : \mathrm{A} \to \mathrm{A}^\vee$ which is given on points by

$$x \mapsto [T_x^* \mathcal{L} \otimes \mathcal{L}^{-1}].$$

Proof. See [GM07, Theorem 6.18]. □

Definition 1.1.11. For an abelian variety A over an arbitrary field k, we define a *polarisation* to be an isogeny (over k)

$$\xi : \mathrm{A} \longrightarrow \mathrm{A}^\vee$$

such that there exists an ample line bundle $\mathcal{L}$ of $\mathrm{A} \times \overline{k}$ for which $\xi = \phi_{\mathcal{L}}$, where $\phi_{\mathcal{L}}$ is the canonical isogeny of Proposition 1.1.10. We define a *principal polarisation* to be a polarisation that is an isomorphism.

1.2 Lifting ordinary abelian varieties over $\mathbb{F}_q$ to ideals

Definition 1.2.1. A *CM-field* K is a totally imaginary quadratic extension of a totally real number field K_0. We denote by $\bar{\cdot}$ the generator of the Galois group $\mathrm{Gal}(K/K_0)$, and we refer to this as *complex conjugation*. For a CM-field K of degree $2g$ over $\mathbb{Q}$, we define a *CM-type* of K to be a set of g embeddings

$$\{\phi : K \hookrightarrow \mathbb{C}\}$$

that are pairwise non-complex conjugate.

Definition 1.2.2. For q a prime power, write $\mathbf{Ord}_{\mathbb{F}_q}$ for the category of ordinary abelian varieties over $\mathbb{F}_q$. For a Weil q-number π, and a CM-field K such that $K = \mathbb{Q}(\pi)$, write $\chi(\pi)$ for the minimal polynomial of π over $\mathbb{Q}$. For $\mathrm{A} \in \mathbf{Ord}_{\mathbb{F}_q}$, write $\chi(\mathrm{Frob}_q(\mathrm{A}))$ for the characteristic polynomial of the q-power Frobenius endomorphism of A. We define $\mathbf{Ord}_\pi$ to be the full subcategory of $\mathbf{Ord}_{\mathbb{F}_q}$ with objects given by

$$\{\mathrm{A} \in \mathbf{Ord}_{\mathbb{F}_q} : \chi(\mathrm{Frob}_q(\mathrm{A})) = \chi(\pi)\}.$$

Note that for every $\mathrm{A} \in \mathbf{Ord}_\pi$ the complex conjugate $\overline{\pi}$ of π also defines an endomorphism on A as the multiplication-by-q map $[q]$ factors through $[\pi]$. In particular, every $\mathrm{A} \in \mathbf{Ord}_\pi$ in this category comes with a map

$$\begin{array}{rccc} \iota_{\mathrm{A}} : & \mathbb{Z}[\pi, \overline{\pi}] & \hookrightarrow & \mathrm{End}(\mathrm{A}) \\ & \pi & \mapsto & \mathrm{Frob}_q(\mathrm{A}) \\ & \overline{\pi} & \mapsto & \mathrm{Ver}_q(\mathrm{A}). \end{array}$$

Then for every $g \in \mathrm{Hom}_{\mathbf{Ord}_\pi}(\mathrm{A}, \mathrm{A}')$ and every $r \in \mathbb{Z}[\pi, \overline{\pi}]$, we have that $\iota_{\mathrm{A}}(r) \circ g = g \circ \iota_{\mathrm{A}'}(r)$. From now on, we omit ι from the notation.

Definition 1.2.3. Given a prime power q, a Weil q-number π, and a CM-field K such that $K = \mathbb{Q}(\pi)$, we define $\mathbf{Id}_\pi$ to be the category with objects given by the fractional $\mathbb{Z}[\pi, \overline{\pi}]$-ideals, where for any objects $\mathfrak{a}$ and $\mathfrak{b}$ of $\mathbf{Id}_\pi$, the morphisms in $\mathbf{Id}_\pi$ from $\mathfrak{a}$ to $\mathfrak{b}$ are given by

$$\mathrm{Hom}(\mathfrak{a}, \mathfrak{b}) = \{\alpha \in K : \alpha\mathfrak{a} \subseteq \mathfrak{b}\}.$$

Definition 1.2.4. We say that a Weil q-number π is *ordinary* if at least half of the roots in $\overline{\mathbb{Q}}_q$ of the minimal polynomial of π are q-adic units.

The Fixed Frobenius Lifting Theorem, Theorem 1.3.11, will state that if π is an ordinary Weil q-number then there is an equivalence of categories

$$\mathbf{Ord}_\pi \longrightarrow \mathbf{Id}_\pi,$$

and that this functor satisfies some useful properties.

1.3 The Fixed Frobenius Lifting Theorem

Definition 1.3.1. Write $\mathbf{Ord}_{\mathbb{F}_q}$ for the category of ordinary abelian varieties over $\mathbb{F}_q$. We define $\mathbf{POrd}_{\mathbb{F}_q}$ to be the category whose objects are pairs (A, ξ) where $\mathrm{A} \in \mathbf{Ord}_{\mathbb{F}_q}$ and $\xi : \mathrm{A} \to \mathrm{A}^\vee$ is a principal polarisation of A. We define a morphism $f : (\mathrm{A}, \xi) \to (\mathrm{A}', \xi')$ in $\mathbf{POrd}_{\mathbb{F}_q}$ to be an isomorphism of abelian varieties $f : \mathrm{A} \to \mathrm{A}'$ for which the following diagram commutes:

$$\begin{tikzcd} \mathrm{A} \arrow[r, "f"] \arrow[d, "\xi"'] & \mathrm{A}' \arrow[d, "\xi'"] \\ \mathrm{A}^\vee & (\mathrm{A}')^\vee. \arrow[l, "f^\vee"] \end{tikzcd}$$

Definition 1.3.2. Recall the notation χ from Definition 1.2.2. We define $\mathbf{POrd}_\pi$ to be the full subcategory of $\mathbf{POrd}_{\mathbb{F}_q}$ with objects given by

$$\{(\mathrm{A}, \xi) \in \mathbf{POrd}_{\mathbb{F}_q} : \chi(\mathrm{Frob}_q(\mathrm{A})) = \chi(\pi)\}.$$

Definition 1.3.3. Let π be a Weil q-number such that $K = \mathbb{Q}(\pi)$ is a CM-field of degree $2g$ over $\mathbb{Q}$. Let $(\mathrm{A}, \xi), (\mathrm{A}', \xi') \in \mathbf{POrd}_\pi$ be g-dimensional abelian varieties, and let K_0 be the maximal totally real subfield of K. Recall that $\mathrm{End}(\mathrm{A}) \subseteq \mathrm{End}(\mathrm{A}) \otimes \mathbb{Q} = K$. For $\mu \in \mathcal{O}_{K_0}$, if $\mu \in \mathrm{End}(\mathrm{A})$, we define a *$\mu$-isogeny*

$$f : (\mathrm{A}, \xi) \longrightarrow (\mathrm{A}', \xi')$$

to be a morphism $f : \mathrm{A} \to \mathrm{A}'$ in $\mathbf{Ord}_{\mathbb{F}_q}$ such that the diagram

$$\begin{tikzcd} \mathrm{A} \arrow[dr, "\xi"] & \mathrm{A} \arrow[l, "\mu"'] \arrow[r, "f"] & \mathrm{A}' \arrow[d, "\xi'"] \\ & \mathrm{A}^\vee & \mathrm{A}'^\vee \arrow[l, "f^\vee"] \end{tikzcd}$$

commutes.

Remark 1.3.4. Note that the morphisms in $\mathbf{POrd}_\pi$ are exactly the 1-isogenies.

We now define, in several steps, the notion of a polarisation on objects in $\mathbf{Id}_\pi$ that will be functorially compatible with the notion of polarisation on objects in $\mathbf{Ord}_\pi$. Fix a prime power q and a Weil q-number π such that $K = \mathbb{Q}(\pi)$ is a CM-field. We first show how to associate a CM-type of K to π, following Howe [How95, Notation 4.6].

Let k be an algebraic closure of $\mathbb{F}_q$ and write $\mathbb{Q}_q^{\mathrm{ur}} = W(k)$, where $W(k)$ denotes the ring of Witt vectors of k. Now fix one embedding $j : \mathbb{Q}_q^{\mathrm{ur}} \hookrightarrow \mathbb{C}$, and identify $\mathbb{Q}_q^{\mathrm{ur}}$ with its image under j so that $\mathbb{Q}_q^{\mathrm{ur}} \subseteq \mathbb{C}$. Now, write $\overline{\mathbb{Q}_q}$ and $\overline{\mathbb{Q}}$ for the algebraic closures of $\mathbb{Q}_q^{\mathrm{ur}}$ and $\mathbb{Q}$ inside $\mathbb{C}$ respectively. We then obtain the following diagram of inclusions (some of which depend on j):

$$\begin{tikzcd} & \mathbb{C} \arrow[d, dash] & \\ & \overline{\mathbb{Q}_q} \arrow[dl, dash] \arrow[dr, dash] & \\ \mathbb{Q}_q^{\mathrm{ur}} \arrow[d, dash] & & \overline{\mathbb{Q}} \arrow[dl, dash] \\ \mathbb{Q}_q \arrow[dr, dash] & & \\ & \mathbb{Q}, & \end{tikzcd}$$

so that in particular the q-adic valuation on $\mathbb{Q}_q^{\mathrm{ur}}$ extends uniquely to a q-adic valuation v_j on $\overline{\mathbb{Q}_q} \supseteq \overline{\mathbb{Q}}$.

Definition 1.3.5. (c.f. [How95, Notation 4.6])
For a rational prime power q, fix $j : \mathbb{Q}_q^{\mathrm{ur}} \hookrightarrow \mathbb{C}$ as above and define v_j to be the q-adic valuation on $\overline{\mathbb{Q}} \subseteq \mathbb{C}$ obtained from j. Then given a CM-field K and an algebraic integer π such that $K = \mathbb{Q}(\pi)$ and $\pi\overline{\pi} = q$, we define the *(π, j)–CM-type of K* to be

$$\Phi_{\pi,j} := \{\phi : K \hookrightarrow \mathbb{C} : v_j(\phi(\pi)) > 0\}.$$

Definition 1.3.6. With notation as in Definition 1.3.5, for any $x \in K$, we say that x is $\Phi_{\pi,j}$*–positive-imaginary* (respectively *non-positive-imaginary*) if, for every $\phi \in \Phi_{\pi,j}$, we have that $\phi(x)/i \in \mathbb{R}_{>0}$ (respectively $\mathbb{R}_{\leq 0}$).

Definition 1.3.7. For an object $\mathfrak{a} \in \mathbf{Id}_\pi$, we define the *dual* of $\mathfrak{a}$ to be the fractional $\mathbb{Z}[\pi, \overline{\pi}]$-ideal

$$\mathfrak{a}^\vee = \{\alpha \in K : \mathrm{tr}(\alpha\overline{\mathfrak{a}}) \subseteq \mathbb{Z}\}.$$

A *polarisation* of $\mathfrak{a}$ is a non-zero $\Phi_{\pi,j}$-positive-imaginary element $\beta \in \mathrm{Hom}_{\mathbf{Id}_\pi}(\mathfrak{a}, \mathfrak{a}^\vee)$. If in addition $\beta\mathfrak{a} = \mathfrak{a}^\vee$, then we say that β is *principal*. For a morphism $\alpha \in \mathrm{Hom}_{\mathbf{Id}_\pi}(\mathfrak{a}, \mathfrak{b})$, we define the *dual* of α to be

$$\alpha^\vee = \overline{\alpha} \in \mathrm{Hom}_{\mathbf{Id}_\pi}(\mathfrak{b}^\vee, \mathfrak{a}^\vee).$$

Remark 1.3.8. Suppose that $\mathfrak{a}, \mathfrak{b} \in \mathbf{Id}_\pi$ and β is a polarisation of $\mathfrak{a}$. Observe that for any totally real totally positive element μ of K, if $\mu\beta \in \mathrm{Hom}_{\mathbf{Id}_\pi}(\mathfrak{b}, \mathfrak{b}^\vee)$, then $\mu\beta$ is a polarisation of $\mathfrak{b}$.

Next, we define $\mathbf{PId}_\pi$ and the notion of μ-isogeny exactly as we defined $\mathbf{POrd}_\pi$.

Definition 1.3.9. Fix a prime power q, a Weil q-number π, and a CM-field K such that $K = \mathbb{Q}(\pi)$. We define the category $\mathbf{PId}_\pi$ to be the category with objects given by pairs $(\mathfrak{a}, \beta)$, where $\mathfrak{a} \in \mathbf{Id}_\pi$ and $\beta \in \mathrm{Hom}_{\mathbf{Id}_\pi}(\mathfrak{a}, \mathfrak{a}^\vee)$ is a principal polarisation of $\mathfrak{a}$. We define a morphism $(\mathfrak{a}, \beta) \to (\mathfrak{a}', \beta') \in \mathbf{PId}_\pi$ to be an isomorphism $\alpha \in \mathrm{Hom}_{\mathbf{Id}_\pi}(\mathfrak{a}, \mathfrak{a}')$ in $\mathbf{Id}_\pi$ such that

$$\beta = \overline{\alpha}\beta'\alpha.$$

Definition 1.3.10. For $(\mathfrak{a}, \beta), (\mathfrak{a}', \beta') \in \mathbf{PId}_\pi$ and $\mu \in \mathrm{End}(\mathfrak{a})$, a *$\mu$-isogeny*

$$\alpha : (\mathfrak{a}, \beta) \to (\mathfrak{a}', \beta')$$

is a morphism $\alpha \in \mathrm{Hom}_{\mathbf{Id}_\pi}(\mathfrak{a}, \mathfrak{a}')$ such that

$$\beta\mu = \overline{\alpha}\beta'\alpha.$$

Theorem 1.3.11 (Fixed Frobenius Lifting Theorem)**.** *Fix a prime power q, an ordinary Weil q-number π, and let K be the CM-field $\mathbb{Q}(\pi)$. Then there exists an equivalence of categories*

$$T_\pi : \mathbf{Ord}_\pi \longrightarrow \mathbf{Id}_\pi$$

that preserves the notions of dual and polarisation, and preserves the action of $\mathbb{Z}[\pi, \overline{\pi}]$. (See Remark 1.3.12 for formal definitions.)

Note that as T_π preserves the notion of polarisation, it is automatic that it preserves the notion of μ-isogeny.

Remark 1.3.12. Suppose that C and D are categories, each equipped with an involution called dual and denoted as

$$\vee : C \to C \quad \text{and} \quad \vee : D \to D.$$

We say that a functor $F : C \to D$ *preserves the notion of dual* if it comes with a natural isomorphism $f : F \circ \vee \tilde{\to} \vee \circ F$.

Suppose that for all objects A in C (resp. D) we have a subset $P_{\mathrm{A}} \subseteq \mathrm{Hom}(\mathrm{A}, \mathrm{A}^\vee)$ of 'polarisations' such that for every isomorphism $m : \mathrm{B} \to \mathrm{A}$ in C (resp. D), the map

$$\begin{array}{ccc} \mathrm{Hom}(\mathrm{A}, \mathrm{A}^\vee) & \longrightarrow & \mathrm{Hom}(\mathrm{B}, \mathrm{B}^\vee) \\ \varphi & \mapsto & m^\vee \varphi m \end{array}$$

induces a bijection between P_{A} and P_{B}. Given a functor $F = (F, f) : C \to D$ preserving the notion of duals, we say that F *preserves the notion of polarisation* if for all objects $\mathrm{A} \in C$ the map

$$\begin{array}{ccc} \mathrm{Hom}(\mathrm{A}, \mathrm{A}^\vee) & \longrightarrow & \mathrm{Hom}(F(\mathrm{A}), F(\mathrm{A})^\vee) \\ \xi & \mapsto & f_{\mathrm{A}} \circ F(\xi) \end{array}$$

induces a bijection between P_{A} and $P_{F(\mathrm{A})}$.

We show in Appendix A that if functors $F : C \to D$ and $G : D \to C$ define an adjoint equivalence of categories and F preserves the notions of dual and polarisation, then G also preserves the notions of dual and polarisation. By [Lan78, Theorem IV.4.1], every equivalence of categories is one direction of an adjoint equivalence of categories.

Let R be a commutative ring and suppose that C and D are R-linear categories (i.e. the sets $\mathrm{Hom}(\mathrm{A}, \mathrm{B})$ are R-modules and composition of morphisms is R-bilinear). We say that F *preserves the action of* R if

$$F : \mathrm{Hom}_C(\mathrm{A}, \mathrm{B}) \longrightarrow \mathrm{Hom}_D(F(\mathrm{A}), F(\mathrm{B}))$$

is R-linear.

The remainder of this chapter is dedicated to defining the functor T_π, and to showing how Theorem 1.3.11 follows from the work of Deligne and Howe in [Del69] and [How95] via the lifting theorems of Serre, Tate and Lubin.

1.4 The theory of canonical lifts

In order to write down the functor of Theorem 1.3.11, we require the notion of a 'Serre-Tate lift' of both an ordinary abelian variety A over a field k of positive characteristic and of a morphism of ordinary abelian varieties over k. Categorically lifting ordinary abelian varieties over $k = \overline{\mathbb{F}_p}$ to the ring $W_n(k)$ of Witt vectors of length n was first studied by Lubin, Serre, and Tate in a seminar, skeleton notes of which can be found at [LST64]. A simpler proof of their main lifting theorem was later found by Drinfeld and written down by Katz in [Kat81, Chapter 1]. The machinery required to use this theorem to lift ordinary abelian varieties over $\mathbb{F}_q$ to abelian schemes over $\mathbb{Q}_q^{\mathrm{ur}}$ was written down by Messing in [Mes72]; the version of the lifting theorems that we state here are as stated by Messing.

1.4.1 Serre-Tate lifts of ordinary abelian varieties

In this section we show how to lift ordinary abelian varieties over a finite field k to abelian schemes over the Witt vectors $W(k)$ of k; we first recall the definition of an abelian scheme (c.f. [MFK94, Definition 6.1]).

Definition 1.4.1. For a noetherian scheme S, an *abelian scheme* over S is defined to be a proper smooth group S–scheme of which all fibres are geometrically connected.

Proposition 1.4.2. It is equivalent to define an abelian scheme to be a proper smooth group scheme over S of which all fibres are abelian varieties. In particular, when k is a field, we have that A is an abelian $\mathrm{Spec}(k)$-scheme if and only if A is an abelian variety.

Proof. Suppose that A is an abelian scheme over a noetherian scheme S. Then every fibre of A is a proper smooth geometrically connected group scheme over a field. A fibre being proper implies in particular that it is of finite type and separated, and a fibre being smooth and geometrically connected implies that it is geometrically irreducible (see [Stack-Exchange]). Therefore every fibre of A is a finite type, separated, geometrically irreducible group scheme over a field, hence a variety. Every fibre is a group object by definition, so a group variety. The reverse direction is clear. □

Fix a perfect field k of characteristic $p > 0$, and write $W(k)$ for the ring of Witt vectors of k.

Theorem 1.4.3. *Let* A *be an ordinary abelian variety defined over* k*. Then, up to unique isomorphism, there is a projective abelian scheme* $\mathrm{B} \to W(k)$ *such that* $\mathrm{B} \times_{W(k)} k = \mathrm{A}$ *and the map* $\mathrm{End}(\mathrm{B}) \to \mathrm{End}(\mathrm{A})$ *is bijective.*

Proof. See [Mes72, p. V.3.3]. □

Definition 1.4.4. For an ordinary abelian variety A defined over k, we define the *Serre-Tate lift* of A to be the projective abelian $W(k)$-scheme satisfying the conditions of Theorem 1.4.3.

Theorem 1.4.5. *Let* A *and* A' *be ordinary abelian varieties over* k *and let* B *and* B' *be the Serre-Tate lifts of* A *and* A' *respectively. Then the map*

$$\phi : \mathrm{Hom}(\mathrm{B}, \mathrm{B}') \longrightarrow \mathrm{Hom}(\mathrm{A}, \mathrm{A}')$$

is bijective.

Proof. See [Mes72, p. V.3.4]. □

Definition 1.4.6. For A, A′, B, B′ and ϕ as in Theorem 1.4.5 and $f \in \mathrm{Hom}(\mathrm{A}, \mathrm{A}')$, we define the *Serre-Tate lift* of f to be $\phi^{-1} f \in \mathrm{Hom}(\mathrm{B}, \mathrm{B}')$.

1.4.2 Deligne lifts of ordinary abelian varieties

Deligne used the lifting theorems Theorem 1.4.3 and Theorem 1.4.5 to represent ordinary abelian varieties over finite fields as linear algebra objects over $\mathbb{Z}$, for which he defined the following category:

Definition 1.4.7. (c.f. [How95, Definition 4.1])
For a prime power q, we define the category $\mathbf{Del}_q$ to be the category whose objects are pairs (Λ, F), where the Λ are finitely generated free $\mathbb{Z}$-modules, and for a given Λ, the F are endomorphisms of Λ such that

1. the endomorphism $F \otimes \mathbb{Q}$ of $\Lambda \otimes \mathbb{Q}$ is semi-simple, and its eigenvalues in $\mathbb{C}$ have magnitude $q^{1/2}$,
2. at least half of the roots of the characteristic polynomial of F in $\overline{\mathbb{Q}}_q$, counting multiplicities, are p-adic units, and
3. there is an endomorphism V of Λ such that $F \circ V = q$.

The morphisms

$$(\Lambda, F) \longrightarrow (\Lambda', F')$$

of $\mathbf{Del}_q$ are homomorphisms $\varphi : \Lambda \longrightarrow \Lambda'$ of $\mathbb{Z}$-modules such that $\varphi \circ F = F' \circ \varphi$.

Remark 1.4.8. In Theorem 1.4.9 and in the rest of this thesis, for an abelian variety A over a field k, and a field embedding $j : k \hookrightarrow k'$, we will write $\mathrm{A} \times_j k'$ or $\mathrm{A} \times k'$ for $\mathrm{A} \times_{\mathrm{Spec}(k)} \mathrm{Spec}(k')$.

Theorem 1.4.9 (Deligne's lifting theorem)**.** *For a prime power q, fix an embedding $j : \mathbb{Q}_q^{\mathrm{ur}} \hookrightarrow \mathbb{C}$. Let $\mathbf{Ord}_{\mathbb{F}_q}$ be the category of ordinary abelian varieties defined over $\mathbb{F}_q$, and for an object A in $\mathbf{Ord}_{\mathbb{F}_q}$, let $\mathrm{B}/\mathbb{Q}_q^{\mathrm{ur}}$ be the Serre-Tate lift of $\mathrm{A} \times \overline{\mathbb{F}_q}$. Define*

$$D(\mathrm{A}) = H_1(\mathrm{B} \times_j \mathbb{C}, \mathbb{Z}).$$

Let $\mathrm{Frob}_q(\mathrm{A})$ be the q-power Frobenius endomorphism on A, let $\mathrm{Frob}_q(\mathrm{B})$ be its Serre-Tate lift, and let $\overline{\mathrm{Frob}_q(\mathrm{B})}$ be the endomorphism induced by $\mathrm{Frob}_q(\mathrm{B}) \times_j \mathbb{C}$ on $D(\mathrm{A})$. Then the functor defined by

$$\begin{array}{ccc} \mathbf{Ord}_{\mathbb{F}_q} & \longrightarrow & \mathbf{Del}_q \\ \mathrm{A} & \mapsto & (D(\mathrm{A}), \overline{\mathrm{Frob}_q(\mathrm{B})}) \end{array}$$

is an equivalence of categories.

Proof. See [Del69, Théorème 7]. □

We will in fact only use a special case of Deligne's lifting theorem, stated in Corollary 1.4.12.

Definition 1.4.10. For a prime power q, a non-negative integer g, an algebraic integer π, and a CM-field K of degree $2g$ over $\mathbb{Q}$ such that $K = \mathbb{Q}(\pi)$, where $q = \pi\overline{\pi}$, we define the category $\mathbf{Mod}_\pi$ to be the category of $\mathbb{Z}[\pi, \overline{\pi}]$-modules that are free of rank $2g$ over $\mathbb{Z}$.

Remark 1.4.11. Observe that $\mathbf{Mod}_\pi$ is equivalent to $\mathbf{Id}_\pi$.

Consider $\mathbf{Mod}_\pi$ as a subcategory of $\mathbf{Del}_q$ by viewing a $\mathbb{Z}[\pi, \overline{\pi}]$-module M as a pair (M, F) where F is the action of π, and the Verschiebung V is the action of $\overline{\pi}$.

Corollary 1.4.12. For a prime power q, an ordinary Weil q-number π, and a CM-field K such that $K = \mathbb{Q}(\pi)$, define $\mathbf{Ord}_\pi$ as in Definition 1.2.2 and $\mathbf{Mod}_\pi$ as in Definition 1.4.10. The functor of Theorem 1.4.9 defines a functor

$$D_\pi : \quad \mathbf{Ord}_\pi \quad \longrightarrow \quad \mathbf{Mod}_\pi$$

that is an equivalence of categories.

Proof. Note that $\mathbf{Mod}_\pi$ is exactly the full subcategory of pairs (Λ, F) for which the characteristic polynomial of the Frobenius F is exactly the minimal polynomial of π over $\mathbb{Q}$. The result then follows from Theorem 1.4.9. □

Remark 1.4.13. The equivalence of categories T_π of Theorem 1.3.11, and the fact that T_π preserves the action of $\mathbb{Z}[\pi, \overline{\pi}]$ follow immediately from Corollary 1.4.12. For Theorem 1.3.11, it remains only to show that the equivalence of categories respects the notions of dual and polarisation.

1.4.3 Howe lifts of polarised ordinary abelian varieties

Howe ([How95]) gave a notion of polarisation on the objects of $\mathbf{Del}_q$ which is compatible with the notion of polarisation in $\mathbf{Ord}_{\mathbb{F}_q}$ under the functor given in Theorem 1.4.9. We give in Theorem 1.4.21 the special case of Howe's lifting theorem that we need in order to prove the Fixed Frobenius Lifting Theorem. We first define polarisations of objects in $\mathbf{Mod}_\pi$, following Howe.

Definition 1.4.14. For a prime power q, a Weil q-number π that generates a CM-field $\mathbb{Q}(\pi)$, let $\Lambda \in \mathbf{Mod}_\pi$. We define the *dual* of Λ to be

$$\Lambda^\vee = \mathrm{Hom}_{\mathbb{Z}}(\Lambda, \mathbb{Z})$$

viewed as a $\mathbb{Z}[\pi, \overline{\pi}]$-module via

$$\begin{array}{ccc} \mathbb{Z}[\pi,\overline{\pi}] \times \Lambda^\vee & \longrightarrow & \Lambda^\vee \\ (r, f) & \mapsto & (\lambda \mapsto f(\overline{r}\lambda)). \end{array}$$

Then in particular, $\Lambda^\vee \in \mathbf{Mod}_\pi$.

Definition 1.4.15. Let R be a commutative ring with an involution

$$\begin{array}{ccc} R & \longrightarrow & R \\ r & \mapsto & \overline{r}, \end{array}$$

let M be an R-module, let N be an abelian group, and let f be a $\mathbb{Z}$-bilinear form

$$f : M \times M \longrightarrow N.$$

We define f to be *R–semi-balanced* if for every $r \in R$ and $\ell, m \in M$, we have that

$$f(r\ell, m) = f(\ell, \overline{r}m).$$

If furthermore N is an R-module and for every $r \in R$ and $\ell, m \in M$ we have that

$$f(r\ell, m) = rf(\ell, m) = f(\ell, \overline{r}m),$$

we say that f is *R–sesquilinear*.

Definition 1.4.16. (c.f. [How95, p. 2370])
For a prime power q, a Weil q-number π that generates a CM-field $\mathbb{Q}(\pi)$, an element $\Lambda \in \mathbf{Mod}_\pi$, and $\zeta \in \mathrm{Hom}_{\mathbf{Mod}_\pi}(\Lambda, \Lambda^\vee)$ we define the *$\mathbb{Z}$-bilinear form associated to* ζ to be

$$\begin{array}{cccc} b: & \Lambda \times \Lambda & \longrightarrow & \mathbb{Z} \\ & (s, t) & \mapsto & \zeta(s)(t). \end{array}$$

One can check that this is a non-degenerate $\mathbb{Z}[\pi, \overline{\pi}]$–semi-balanced form.

Proposition 1.4.17. For an order $\mathcal{O}$ in a number field K with an involution $\bar{\cdot}$ such that $\overline{\mathcal{O}} = \mathcal{O}$, given a non-degenerate $\mathcal{O}$–semi-balanced form $b : \Lambda \times \Lambda \to \mathbb{Z}$, there exists a unique non-degenerate K-sesquilinear form $S : (\Lambda \otimes \mathbb{Q}) \times (\Lambda \otimes \mathbb{Q}) \to K$ such that $b \otimes \mathbb{Q} = \mathrm{tr}_{K/\mathbb{Q}} \circ S$.

Proof. See [Knu91, Theorem I.7.4.1, p.44]. □

Definition 1.4.18. For $\Lambda \in \mathbf{Mod}_\pi$ and $\zeta \in \mathrm{Hom}_{\mathbf{Mod}_\pi}(\Lambda, \Lambda^\vee)$, let $b : \Lambda \times \Lambda \to \mathbb{Z}$ be the $\mathbb{Z}$-bilinear form associated to ζ. We define the *K-sesquilinear form associated to* ζ to be the unique non-degenerate K-sesquilinear form of Proposition 1.4.17.

Remark 1.4.19. For every $\Lambda \in \mathbf{Mod}_\pi$, given a non-degenerate $\mathbb{Z}[\pi, \overline{\pi}]$-semi-balanced form $b : \Lambda \times \Lambda \to \mathbb{Z}$, there is a unique morphism $\zeta \in \mathrm{Hom}(\Lambda, \Lambda^\vee)$ for which the associated $\mathbb{Z}$-bilinear form is b given by $\zeta(s)(t) = b(s, t)$.

Definition 1.4.20. (c.f. [How95, Definition 4.8])
Fix a Weil q-number π and a CM-field K such that $K = \mathbb{Q}(\pi)$, and fix an embedding $j : \mathbb{Q}_q^{\mathrm{ur}} \hookrightarrow \mathbb{C}$. Recall the definition of the CM-type $\Phi_{\pi,j}$ of K from Definition 1.3.5, and recall the definition of $\Phi_{\pi,j}$ - - non - positive - imaginary from Definition 1.3.6. For $\Lambda \in \mathbf{Mod}_\pi$, we define a *$j$-polarisation of* Λ to be a morphism

$$\zeta : \Lambda \longrightarrow \Lambda^\vee$$

such that the sesquilinear form S associated to ζ is skew-Hermitian (i.e. for every $u, v \in \Lambda \otimes \mathbb{Q}$ we have $S(u, v) = -\overline{S(v, u)}$) and such that for every $\lambda \in \Lambda$ we have that $S(\lambda, \lambda)$ is $\Phi_{\pi,j}$-non-positive-imaginary.

The following theorem, a special case of Howe's lifting theorem in [How95, Proposition 4.9], shows that this definition of polarisation is what we should use if we wish to study ordinary abelian varieties over finite fields.

Theorem 1.4.21. *For an abelian variety* $\mathrm{A} \in \mathbf{Ord}_\pi$ *with dual abelian variety* $\mathrm{A}^\vee \in \mathbf{Ord}_\pi$ *and an isogeny* $\xi : \mathrm{A} \to \mathrm{A}^\vee$ *in* $\mathbf{Ord}_\pi$*, let* Λ*,* $\Lambda^\vee \in \mathbf{Mod}_\pi$ *and* $\alpha \in \mathrm{Hom}(\Lambda, \Lambda^\vee)$ *be the images under the functor of Corollary 1.4.12 of* A*,* $\mathrm{A}^\vee$ *and* ξ *respectively. Let* $j : \mathbb{Q}_q^{\mathrm{ur}} \hookrightarrow \mathbb{C}$ *be the embedding on which the functor of Corollary 1.4.12 depends. Then* α *is a* j*-polarisation of* Λ *if and only if* ξ *is a polarisation of* A*.*

Proof. See [How95, Proposition 4.9]. □

Remark 1.4.22. The first step of Howe's proof is a reference to the well-known result that one may lift polarisations of ordinary abelian varieties over finite fields to polarisations of abelian varieties over $\mathbb{C}$, but the reference [Del69, Theorem 1] cited by Howe in [How95, Proof of Proposition 4.9] does not give a proof. In private correspondence, Howe was kind enough to provide the following argument for this step.

Recall that we fixed an embedding $j : \mathbb{Q}_q^{\mathrm{ur}} \hookrightarrow \mathbb{C}$. Let B be the Serre-Tate lift of $\mathrm{A} \in \mathbf{Ord}_\pi$, and write $\mathrm{B}_\mathbb{C}$ for $\mathrm{B} \times_j \mathbb{C}$.

Let $K = \mathrm{End}(\mathrm{B}_\mathbb{C}) \otimes \mathbb{Q}$ and fix an ample divisor $\mathcal{L}$ on $\mathrm{B}_\mathbb{C}$, so that we have a Rosati involution on K. Then the Neron-Severi group of $\mathrm{B}_\mathbb{C}$ (that is, $\mathrm{Pic}(\mathrm{B}_\mathbb{C})/\mathrm{Pic}^0(\mathrm{B}_\mathbb{C})$), when tensored with $\mathbb{Q}$, can be identified via $\mathcal{M} \to \phi_\mathcal{L}^{-1} \circ \phi_\mathcal{M}$ with the maximal additive subgroup K_0 of K fixed by the Rosati involution. (See Mumford [Mum08, Application III, page 208].) So to every line bundle $\mathcal{M}$, we can associate a real (i.e. fixed by Rosati) element of K. And the ample line bundles $\mathcal{M}$ are precisely the ones for which $\phi_\mathcal{L}^{-1} \circ \phi_\mathcal{M} \in K_0$ is totally positive. (See the last paragraph of Section 21 of Mumford [Mum08].)

Now, there is an ample line bundle $\mathcal{L}$ on $\mathrm{B}_\mathbb{C}$ whose reduction $\overline{\mathcal{L}}$ is ample (see [Gro61, Corollaire 4.5.14]). This gives us one polarisation λ of B that descends to a polarisation $\overline{\lambda}$ of A. But an isogeny $f : \mathrm{B}_\mathbb{C} \to \mathrm{B}_\mathbb{C}^\vee$ is a polarisation if and only if there exists an ample line bundle $\mathcal{M}$ on $\mathrm{B}_\mathbb{C}$ such that $f = \phi_\mathcal{M}$, which is if and only if $f^{-1}\lambda$ is a totally positive real element of $\mathrm{End}(\mathrm{B}) \otimes \mathbb{Q}$ (by the previous paragraph), and this condition holds for f if and only if it holds for the reduction $\overline{f}$ of f.

1.4.4 Proof of the Fixed Frobenius Lifting Theorem

The Fixed Frobenius Lifting Theorem, Theorem 1.3.11, is a consequence of Howe's lifting theorem Theorem 1.4.21. We only need to show that there is a functor defining an equivalence of categories between $\mathbf{Mod}_\pi$ of $\mathbb{Z}[\pi, \overline{\pi}]$-modules and the category $\mathbf{Id}_\pi$ of fractional $\mathbb{Z}[\pi, \overline{\pi}]$-ideals that preserves the notions of dual and polarisation.

Proof of Theorem 1.3.11. By Theorem 1.4.21, the equivalence of categories $\mathbf{Ord}_\pi \to \mathbf{Mod}_\pi$ given in Corollary 1.4.12 preserves the notions of dual and polarisation. By construction this equivalence also preserves the action of $\mathbb{Z}[\pi, \overline{\pi}]$. We show that the forgetful functor

$$O_\pi : \mathbf{Id}_\pi \to \mathbf{Mod}_\pi$$

1. preserves the action of $\mathbb{Z}[\pi, \overline{\pi}]$.
2. preserves the notion of duals.
3. preserves the notion of polarisation.

The preservation of the action of $\mathbb{Z}[\pi, \overline{\pi}]$ is immediate as O_π maps the morphism in $\mathbf{Id}_\pi$ defined by π to the morphism in $\mathbf{Mod}_\pi$ defined by π, and similarly for $\overline{\pi}$. Observe also that O_π is an equivalence of categories, and as stated in Remark 1.3.12, if O_π preserves the notions of dual and polarisation then so does the reverse functor. For (2), given $\mathfrak{a} \in \mathbf{Id}_\pi$, we claim that

$$f_\mathfrak{a} : \alpha \mapsto (\beta \mapsto \mathrm{tr}_{K/\mathbb{Q}}(\overline{\alpha}\beta)) \tag{1.1}$$

defines a natural isomorphism from

$$O_\pi(\mathfrak{a}^\vee) = O_\pi(\{\alpha \in K : \mathrm{tr}(\alpha\overline{\mathfrak{a}}) \subseteq \mathbb{Z}\})$$

to

$$O_\pi(\mathfrak{a})^\vee = \mathrm{Hom}(O_\pi(\mathfrak{a}), \mathbb{Z})).$$

As

$$\begin{array}{rccc} T: & K \times K & \longrightarrow & \mathbb{Q} \\ & \alpha, \beta & \mapsto & \mathrm{tr}(\overline{\alpha}\beta) \end{array}$$

is a non-degenerate bilinear form, it induces an isomorphism of $\mathbb{Q}$-vector spaces

$$\begin{array}{ccc} K & \longrightarrow & \mathrm{Hom}(K, \mathbb{Q}) \\ \alpha & \mapsto & T(\alpha, -), \end{array}$$

which when restricted to the subset $\mathfrak{a} \subseteq K$, gives us exactly the map of (1.1).

For (3), we have to show that the map

$$\begin{array}{ccc} \mathrm{Hom}_{\mathbf{Id}_\pi}(\mathfrak{a}, \mathfrak{a}^\vee) & \longrightarrow & \mathrm{Hom}_{\mathbf{Mod}_\pi}(O_\pi(\mathfrak{a}), O_\pi(\mathfrak{a})^\vee) \\ \beta & \mapsto & f_\mathfrak{a} \circ O_\pi(\beta) = (x \mapsto \mathrm{tr}_{K/\mathbb{Q}}(\overline{\beta x} -)) \end{array}$$

induces a bijection $P_\mathfrak{a} \leftrightarrow P_{O_\pi(\mathfrak{a})}$ of the set of polarisations of $\mathfrak{a}$ and the set of polarisations of $O_\pi(\mathfrak{a})$. The morphism

$$\begin{array}{rccc} \zeta: & O_\pi(\mathfrak{a}) & \longrightarrow & O_\pi(\mathfrak{a})^\vee \\ & x & \mapsto & \mathrm{tr}_{K/\mathbb{Q}}(\overline{\beta x} -) \end{array}$$

is in $P_{O_\pi(\mathfrak{a})}$ if and only if the sesquilinear form associated to ζ, given by

$$\begin{array}{rccc} S: & (O_\pi(\mathfrak{a}) \otimes \mathbb{Q}) \times (O_\pi(\mathfrak{a}) \otimes \mathbb{Q}) & \longrightarrow & K \\ & (s, t) & \mapsto & \overline{\beta} st, \end{array}$$

is skew-Hermitian and, for every $\lambda \in O_\pi(\mathfrak{a})$, we have that $S(\lambda, \lambda)$ is $\Phi_{\pi,j}$–non-positive-imaginary. But S is skew-Hermitian if and only if β is totally imaginary, and $S(\lambda, \lambda)$ is $\Phi_{\pi,j}$–non-positive-imaginary for every $\lambda \in O_\pi(\mathfrak{a})$ if and only if β is $\Phi_{\pi,j}$–positive-imaginary. Therefore $\beta \in \mathrm{Hom}(\mathfrak{a}, \mathfrak{a}^\vee)$ is in $P_\mathfrak{a}$ if and only if $f_\mathfrak{a} \circ O_\pi(\beta) \in P_{O_\pi(\mathfrak{a})}$, hence (3) holds. □

1.5 Maximal real multiplication

In much of this thesis, we will study principally polarised abelian varieties of dimension g defined over $\mathbb{C}$ that have *maximal real multiplication*, that is, the real part of the endomorphism ring is a maximal order in a totally real number field of degree g over $\mathbb{Q}$. We now give some preliminaries.

Definition 1.5.1. Fix an ordinary Weil q-number π and a CM-field $K = \mathbb{Q}(\pi)$, and denote by K_0 the maximal totally real subfield of K. Let $\mathbf{C}_\pi$ denote one of $(\mathbf{P})\mathbf{Mod}_\pi$, $(\mathbf{P})\mathbf{Ord}_\pi$, or $(\mathbf{P})\mathbf{Id}_\pi$. For each choice of $\mathbf{C}_\pi$, every object $\mathrm{A} \in \mathbf{C}_\pi$ comes together with an embedding $\mathbb{Z}[\pi, \overline{\pi}] \hookrightarrow \mathrm{End}(\mathrm{A})$, so we identify $K = \mathbb{Z}[\pi, \overline{\pi}] \otimes \mathbb{Q}$ with a subring of $\mathrm{End}(\mathrm{A}) \otimes \mathbb{Q}$. We define $\mathbf{C}_{\pi,K_0}$ to be the full category of $\mathbf{C}_\pi$ consisting of those $\mathrm{A} \in \mathbf{C}_\pi$ such that $\mathcal{O}_{K_0} \subseteq \mathrm{End}(\mathrm{A})$.

Definition 1.5.2. Let K_0 be a totally real number field of degree g over $\mathbb{Q}$ with ring of integers $\mathcal{O}_{K_0}$, and define $\mathbf{Ord}_{\mathbb{C},g}$ to be the category of abelian varieties over $\mathbb{C}$ of dimension g. We define the objects of the category $\mathbf{Ord}_{\mathbb{C},K_0}$ to be pairs (A, ι), where $\mathrm{A} \in \mathbf{Ord}_{\mathbb{C},g}$ and $\iota : \mathcal{O}_{K_0} \hookrightarrow \mathrm{End}(\mathrm{A})$ is an embedding. A morphism in $\mathbf{Ord}_{\mathbb{C},K_0}$ between two objects (A, ι) and (A', ι') is given by a morphism $f : \mathrm{A} \to \mathrm{A}'$ in $\mathbf{Ord}_{\mathbb{C},g}$ such that the diagram

$$\begin{tikzcd} \mathrm{End}(\mathrm{A}) \otimes \mathbb{Q} \arrow[r, "g \mapsto f \circ g \circ f^{-1}"] & \mathrm{End}(\mathrm{A}') \otimes \mathbb{Q} \\ K_0 \arrow[u, "\iota"] \arrow[ur, "\iota'"'] & \end{tikzcd}$$

commutes. We define the objects of the category $\mathbf{POrd}_{\mathbb{C},K_0}$ to be triples (A, ξ, ι), where $(\mathrm{A}, \iota) \in \mathbf{Ord}_{\mathbb{C},K_0}$ and $\xi : \mathrm{A} \to \mathrm{A}^\vee$ is a principal polarisation of A, and the image of ι is stable under the Rosati involution. A morphism in $\mathbf{POrd}_{\mathbb{C},K_0}$ between two objects (A, ξ, ι) and $(\mathrm{A}', \xi', \iota') \in \mathbf{POrd}_{\mathbb{C},K_0}$ is an isomorphism

$$f : (\mathrm{A}, \iota) \longrightarrow (\mathrm{A}', \iota')$$

in $\mathbf{Ord}_{\mathbb{C},K_0}$ that makes the diagram

$$\begin{tikzcd} \mathrm{A} \arrow[r, "f"] \arrow[d, "\xi"'] & \mathrm{A}' \arrow[d, "\xi'"] \\ \mathrm{A}^\vee & \mathrm{A}'^\vee \arrow[l, "f^\vee"] \end{tikzcd}$$

commute.

Definition 1.5.3. Let K_0 be a totally real number field with ring of integers $\mathcal{O}_{K_0}$. For

$$(\mathrm{A}, \xi, \iota), (\mathrm{A}', \xi', \iota') \in \mathbf{POrd}_{\mathbb{C},K_0}$$

and $\mu \in \mathcal{O}_{K_0}$, we define a *$\mu$-isogeny* $f : (\mathrm{A}, \xi, \iota) \to (\mathrm{A}', \xi', \iota')$ to be a morphism $f : (\mathrm{A}, \iota) \to (\mathrm{A}', \iota')$ in $\mathbf{Ord}_{\mathbb{C},K_0}$ such that the diagram

$$\begin{tikzcd} \mathrm{A} \arrow[dr, "\xi"] & \mathrm{A} \arrow[l, "\iota(\mu)"'] \arrow[r, "f"] & \mathrm{A}' \arrow[d, "\xi'"] \\ & \mathrm{A}^\vee & \mathrm{A}'^\vee \arrow[l, "f^\vee"] \end{tikzcd}$$

commutes.

Definition 1.5.4. Let q be a prime power, let π be an ordinary Weil q-number, and let $K = \mathbb{Q}(\pi)$ be a CM-field with maximal totally real subfield K_0. Let $g = [K_0 : \mathbb{Q}]$, and define

$$\mathbf{Ord}_{\mathbb{C},\pi}$$

to be the full subcategory of $\mathbf{Ord}_{\mathbb{C},g}$ with objects $(\mathrm{A}, e : \mathbb{Z}[\pi, \overline{\pi}] \hookrightarrow \mathrm{End}(\mathrm{A}))$, where e has CM-type $\Phi_{\pi,j}$. We define

$$\mathbf{Ord}_{\mathbb{C},\pi,K_0}$$

to be the full subcategory of $\mathbf{Ord}_{\mathbb{C},\pi}$ such for every object (A, e), the embedding $e : \mathbb{Z}[\pi, \overline{\pi}] \hookrightarrow \mathrm{End}(\mathrm{A})$ extends to an embedding $f : \mathcal{O}_{K_0}[\pi, \overline{\pi}] \hookrightarrow \mathrm{End}(\mathrm{A})$. (Note that in fact $\mathcal{O}_{K_0}[\pi, \overline{\pi}] = \mathcal{O}_{K_0}[\pi + \overline{\pi}]$ as $\pi + \overline{\pi} \in \mathcal{O}_{K_0}$.)

Observe that for $(\mathrm{A}, e) \in \mathbf{Ord}_{\mathbb{C},\pi,K_0}$, we have that $(\mathrm{A}, f|_{\mathcal{O}_{K_0}}) \in \mathbf{Ord}_{\mathbb{C},K_0}$.

Theorem 1.5.5. *Let q be a prime power, let π be an ordinary Weil q-number, and let $K = \mathbb{Q}(\pi)$ be a CM-field with maximal totally real subfield K_0. Let $g = [K_0 : \mathbb{Q}]$. There is an equivalence of categories*

$$\begin{array}{rccc} F_\pi : & \mathbf{Id}_{\pi,K_0} & \longrightarrow & \mathbf{Ord}_{\mathbb{C},\pi,K_0} \\ & \mathfrak{a} & \mapsto & \mathbb{C}^g/\Phi_{\pi,j}(\mathfrak{a}) \end{array}$$

that preserves the action of $\mathcal{O}_{K_0}[\pi, \overline{\pi}]$ and the notions of dual and polarisation.

Proof. This is Theorems 4.1 and 4.2 (1) of Lang [Lan83].

□

1.6 Hilbert modular forms

Definition 1.6.1. Let K_0 be a totally real number field of degree g over $\mathbb{Q}$ with ring of integers $\mathcal{O}_{K_0}$. Let $\mathcal{N}$ be an invertible $\mathcal{O}_{K_0}$-ideal. Then the matrix group $\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{N})$ is defined as

$$\left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(K_0) : a, d \in \mathcal{O}_{K_0},\, b \in \mathcal{N},\, c \in \mathcal{N}^{-1} \right\}.$$

Let $\mathbb{H}$ be the complex upper half plane. We want to view objects in $\mathbf{POrd}_{\mathbb{C},K_0}$ as elements of $\mathbb{H}^g$, where g is the degree of K_0 over $\mathbb{Q}$. We will be interested in the action of matrix groups with entries in K_0 on elements of $\mathbb{H}^g$, hence it is much more convenient to work with $K_0 \otimes \mathbb{C}$ instead of $\mathbb{C}^g$. To this end, we fix once for all a $\mathbb{C}$-algebra isomorphism

$$\mathbb{C}^g \longrightarrow K_0 \otimes \mathbb{C} \tag{1.2}$$

and we define $K_0 \otimes \mathbb{H}$ to be the image of $\mathbb{H}^g$ under this isomorphism. Observe that $K_0 \otimes \mathbb{H}$ does not depend on the choice of isomorphism. Let the group of 2×2 matrices with entries in K_0 that have totally positive determinant be denoted by $\mathrm{GL}_2(K_0)^+$. The group $\mathrm{GL}_2(K_0)^+$ acts on $K_0 \otimes \mathbb{H}$ as

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot \tau \mapsto (a\tau + b)(c\tau + d)^{-1}.$$

Lemma 1.6.2. Let K_0 be a totally real number field of degree g over $\mathbb{Q}$ with ring of integers $\mathcal{O}_{K_0}$, and write $\mathcal{O}_{K_0}^\vee$ for the trace dual of $\mathcal{O}_{K_0}$. Then there is a bijection

$$\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)\backslash(K_0 \otimes \mathbb{H}) \longrightarrow \{(\mathrm{A}, \xi, \iota) \in \mathbf{POrd}_{\mathbb{C},K_0}\}_{/\cong}$$

where the image of $\tau \in \mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)\backslash(K_0 \otimes \mathbb{H})$ is $\mathrm{A} = (K_0 \otimes \mathbb{C})/(\tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee)$ with the natural embedding ι and the polarisation induced by the Riemann form $E : (K_0 \otimes \mathbb{C}) \times (K_0 \otimes \mathbb{C}) \longrightarrow \mathbb{R}$ given by

$$E(\tau u_1 + u_2, \tau v_1 + v_2) = \mathrm{tr}_{(K_0\otimes\mathbb{R})/\mathbb{R}}(u_1 v_2 - u_2 v_1)$$

for $u_1, u_2, v_1, v_2 \in K_0 \otimes \mathbb{R}$.

Proof. See [Gee88, Chapter IX, Section 1]. □

Definition 1.6.3. Let κ be an integer, and let τ be in $K_0 \otimes \mathbb{H}$. Then the *weight function* w_κ is defined by

$$\begin{array}{rccc} w_\kappa : & \mathrm{GL}_2(K_0)^+ \times (K_0 \otimes \mathbb{H}) & \longrightarrow & \mathbb{C} \\ & (M, \tau) & \mapsto & \left(\mathrm{N}_{K_0/\mathbb{Q}}(\det(M))^{-1/2}\mathrm{N}_{(K_0\otimes\mathbb{C})/\mathbb{C}}(c\tau + d)\right)^\kappa, \end{array}$$

where we choose the positive square root.

Definition 1.6.4. Let $\mathrm{GL}_2(K_0)^+$ and $K_0 \otimes \mathbb{H}$ be as in Definition 1.6.3. Let M be any matrix in $\mathrm{GL}_2(K_0)^+$, and let $f : K_0 \otimes \mathbb{H} \to \mathbb{C}$ be a holomorphic map. Then we define $f|_{[M]_\kappa}$ by

$$\begin{array}{rccc} f|_{[M]_\kappa} : & K_0 \otimes \mathbb{H} & \to & \mathbb{C} \\ & \tau & \mapsto & w_\kappa(M, \tau)^{-1} f(M\tau). \end{array}$$

It is straightforward to check that for $M, N \in \mathrm{GL}_2(K_0)^+$, we have

$$(f|_{[M]_\kappa})|_{[N]_\kappa} = f|_{[MN]_\kappa}.$$

Definition 1.6.5. Let $\mathrm{GL}_2(K_0)^+$ and $K_0 \otimes \mathbb{H}$ be as above, and assume that $g > 1$. Let Γ be a congruence subgroup of $\mathrm{GL}_2(K_0)^+$. We say that $f : K_0 \otimes \mathbb{H} \to \mathbb{C}$ is a *Hilbert modular form* of weight κ for Γ if and only if it is holomorphic and for all $M \in \Gamma$ and $\tau \in K_0 \otimes \mathbb{H}$, we have

$$f|_{[M]_\kappa}(\tau) = f(\tau).$$

From this point on, if f is a Hilbert modular form of weight κ, then for $M \in \mathrm{GL}_2(K_0)^+$ we will write $f|_M$ for $f|_{[M]_\kappa}$.

Remark 1.6.6. For $g = 1$, we also have to impose holomorphicity at the cusps.

Definition 1.6.7. With notation as in Definition 1.6.5, if $\varphi = f/g$ is the quotient of Hilbert modular forms for Γ of equal weight, then we say that φ is a *Hilbert modular function* for Γ.

Definition 1.6.8. Suppose that $g = 2$. Then for $f \in \mathcal{M}_{K_0,k}$, if for every $(\tau_1, \tau_2) \in K_0 \otimes \mathbb{H} = \mathbb{H}^2$ we have

$$f(\tau_1, \tau_2) = f(\tau_2, \tau_1),$$

we say that f is *symmetric*.

Definition 1.6.9. Let $\mathcal{O}_{K_0}^\vee$ be the trace dual of $\mathcal{O}_{K_0}$. We define $\mathcal{M}_{K_0,\kappa}$ to be the $\mathbb{C}$-vector space of Hilbert modular forms for $\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$ of weight κ, and we define

$$\mathcal{M}_{K_0} = \bigoplus_\kappa \mathcal{M}_{K_0,\kappa}$$

to be the graded $\mathbb{C}$-algebra of all Hilbert modular forms for $\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$. For $f \in \mathcal{M}_{K_0}$, let $\mathrm{coeffs}(f)$ be the set of coefficients of the q-expansion of f around the cusp at infinity. For a ring R, we define

$$\mathcal{M}_{K_0,\kappa}(R) = \{f \in \mathcal{M}_{K_0,\kappa} : \ \mathrm{coeffs}(f) \subseteq R\},$$

and

$$\mathcal{M}_{K_0}(R) = \{f \in \mathcal{M}_{K_0} : \ \mathrm{coeffs}(f) \subseteq R\}.$$

Theorem 1.6.10. *(Baily-Borel Theorem)*
Let $\mathcal{M}_{K_0}$ be the graded ring of Hilbert modular forms for $\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$. Then the normal complex analytic space of $\mathrm{Proj}(\mathcal{M}_{K_0})$ is a compactification of

$$V = \mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)\backslash(K_0 \otimes \mathbb{H}).$$

Proof. See [Gee88, p. II.7.1]. □

Definition 1.6.11. We define the *Hilbert modular variety* $\overline{V}$ to be the normal complex analytic space of $\mathrm{Proj}(\mathcal{M}_{K_0})$. We will also refer to this as the *Baily-Borel compactification of* V.

Proposition 1.6.12. (Rapoport)
$\mathcal{M}_{K_0,\kappa}(\mathbb{Z})$ is a finitely generated $\mathbb{Z}$-module.

Proof. See [Rap78, Proposition 6.6]. □

Lemma 1.6.13. (Rapoport)

$$\mathcal{M}_{K_0}(\mathbb{Z}) \otimes_{\mathbb{Z}} \mathbb{C} = \mathcal{M}_{K_0}.$$

Proof. See the proof of [Rap78, Lemma 6.12]. □

Proposition 1.6.14. Let K_0 be a quadratic number field of discriminant 5, 8, 13 or 17. Then $\mathcal{M}_{K_0}(\mathbb{Q})$ is a finitely generated $\mathbb{Q}$-algebra, and the q-expansions of a choice of generators are known.

Proof. For discriminant 5 see [Mue85] or [May07], for discriminant 8 see [Mue83], and for discriminants 13 and 17 see [May07]. □

Remark 1.6.15. In everything that follows, we will assume that $\mathcal{M}_{K_0}(\mathbb{Q})$ is a finitely generated $\mathbb{Q}$-algebra.

1.7 A normalisation lemma for principally polarised ideals

Let q be a prime power, let π be an ordinary Weil q-number, and let $K = \mathbb{Q}(\pi)$ be a CM-field with maximal totally real subfield K_0. As we have seen in Theorem 1.3.11, we can study principally polarised ordinary abelian varieties over finite fields by studying principally polarised ideals. Recall from Definition 1.5.1 that we defined $\mathbf{PId}_{\pi,K_0}$ to be the category of principally polarised $\mathcal{O}_{K_0}[\pi,\overline{\pi}]$-ideals. This corresponds to studying principally polarised ordinary abelian varieties over $\mathbb{F}_q$ with Frobenius π and with real multiplication by $\mathcal{O}_{K_0}$, which are a main topic of interest throughout this thesis. In this section we prove a very useful property of objects $(\mathfrak{a},\beta) \in \mathbf{PId}_{\pi,K_0}$ that we will use throughout this thesis:

Lemma 1.7.1. Suppose that $(\mathfrak{a},\beta) \in \mathbf{PId}_{\pi,K_0}$ is a principally polarised fractional $\mathcal{O}_{K_0}[\pi,\overline{\pi}]$-ideal. Then there exists $\tau \in K - K_0$ such that

$$(\mathfrak{a},\beta) \cong_{\mathbf{PId}_{\pi,K_0}} (\tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee, (\overline{\tau}-\tau)^{-1}).$$

Proof. By assumption we have that $\mathcal{O}_{K_0} \subset \mathrm{End}(\mathfrak{a}) \subseteq \mathcal{O}_K$, where $K = \mathbb{Q}(\pi)$ is a totally imaginary quadratic extension of K_0. In particular, as $\mathcal{O}_{K_0}$ is a Dedekind domain, by Cohen [Coh93, Theorem 1.2.19] there exist $x, y' \in K$ and a fractional $\mathcal{O}_{K_0}$-ideal $\mathfrak{b}$ such that

$$\mathfrak{a} = x\mathcal{O}_{K_0} + y'\mathfrak{b}.$$

Now, from the polarisation β of $\mathfrak{a}$, we have a non-degenerate alternating $\mathbb{Z}$-bilinear form defined by

$$\begin{array}{rccc} E: & \mathfrak{a}\times\mathfrak{a} & \longrightarrow & \mathbb{Z} \\ & (u,v) & \mapsto & \mathrm{tr}_{K/\mathbb{Q}}(\beta\overline{u}v), \end{array}$$

which factors via the non-degenerate alternating $\mathcal{O}_{K_0}$-bilinear form

$$\begin{array}{rccc} S: & \mathfrak{a} \times \mathfrak{a} & \longrightarrow & \mathcal{O}_{K_0}^{\vee} \\ & (u,v) & \mapsto & \mathrm{tr}_{K/K_0}(\beta \overline{u} v) \end{array}$$

by definition of the trace dual $\mathcal{O}_{K_0}^{\vee}$. The matrix of $S \otimes \mathbb{Q}$ with respect to the K_0-basis $\langle x, y' \rangle$ is then given by

$$\begin{pmatrix} 0 & c \\ -c & 0 \end{pmatrix},$$

where $c = \mathrm{tr}_{K/K_0}(\beta \overline{x} y')$. Choose $\mathbb{Z}$-bases $(\omega_1, \ldots, \omega_g)$ and $(b_1, \ldots, b_g)$ for $\mathcal{O}_{K_0}$ and $\mathfrak{b}$ respectively. We compute the matrix of E with respect to the $\mathbb{Z}$-basis

$$\langle x\omega_1, \ldots, x\omega_g, y'b_1, \ldots, y'b_g \rangle,$$

to be

$$\begin{pmatrix} 0 & M \\ -M & 0 \end{pmatrix},$$

where

$$M = (\mathrm{tr}_{K_0/\mathbb{Q}}(c\omega_i b_j))_{i,j=1,\ldots,g}.$$

In turn, we get that M is the matrix of the $\mathbb{Z}$-bilinear form

$$\begin{array}{rccc} F: & \mathcal{O}_{K_0} \times \mathfrak{b} & \longrightarrow & \mathbb{Z} \\ & (u,v) & \mapsto & \mathrm{tr}_{K_0/\mathbb{Q}}(cuv) \end{array}$$

with respect to the $\mathbb{Z}$-bases $(\omega_1, \ldots, \omega_g)$ and $(b_1, \ldots, b_g)$. In particular, as E (and hence F) is non-degenerate and the matrix of E has determinant ± 1, we get that

$$c\mathfrak{b} = \mathcal{O}_{K_0}^{\vee}.$$

Hence, we have that

$$\mathfrak{a} = x\mathcal{O}_{K_0} + y'c^{-1}\mathcal{O}_{K_0}^{\vee}.$$

Then, setting $y = y'c^{-1}$, multiplication by y^{-1} defines an isomorphism in $\mathbf{PId}_{\pi,K_0}$ from $(\mathfrak{a}, \beta)$ to

$$(xy^{-1}\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^{\vee}, y\overline{y}\beta).$$

Now repeat the same argument with $\mathfrak{a}' = \tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^{\vee}$, where $\tau = xy^{-1}$ and $\beta' = y\overline{y}\beta$. Then choose $\mathfrak{b} = \mathcal{O}_{K_0}^{\vee}$ so that

$$c\mathcal{O}_{K_0}^{\vee} = c\mathfrak{b} = \mathcal{O}_{K_0}^{\vee},$$

hence

$$\mathcal{O}_{K_0}^{\times} \ni c = \mathrm{tr}_{K/K_0}(\beta'\overline{\tau}) = \beta'\overline{\tau} + \overline{\beta'}\tau = \beta'(\overline{\tau} - \tau).$$

So we can replace τ by $c^{-1}\tau$, giving

$$\beta' = (\overline{\tau} - \tau)^{-1}.$$

□

Remark 1.7.2. Note that, as $(\overline{\tau} - \tau)^{-1}$ is a polarisation, it is by definition $\Phi_{\pi,j}$-positive-imaginary. That is, for every $\phi \in \Phi_{\pi,j}$ we have that $\phi(\tau) \in \mathbb{H}$, hence with respect to $\Phi_{\pi,j}$ we have that $\tau \in K_0 \otimes \mathbb{H}$.

Chapter 2

Hilbert modular polynomials

2.1 Introduction and statement of the results

The modular polynomial for elliptic curves of prime level p is an irreducible polynomial $\Phi_p(X,Y) \in \mathbb{Z}[X,Y]$ which, for every pair of p-isogenous elliptic curves E and E', satisfies

$$\Phi_p(j(E), j(E')) = 0,$$

where $j(E)$ is the j-invariant of the elliptic curve E. Examples of these modular polynomials can be found for example on Sutherland's website [Sut18]. One of the reasons that modular polynomials interest us is that given the j-invariant of an elliptic curve E over a field k, we can find the j-invariants of all those elliptic curves that are p-isogenous to it by computing the roots of $\Phi_p(j(E), Y) \in k[Y]$. In this chapter, we describe an analogue of the modular polynomial for principally polarised abelian varieties of dimension g with real multiplication, which we call a *set of Hilbert modular polynomials*. This is a Hilbert modular function analogue of Dupont's work with Siegel modular functions in [Dup06]. The advantage of working in the Hilbert setting is that the coefficients and degrees of the polynomials are much more manageable than in the Siegel setting, making it possible to compute modular polynomials for higher prime levels than previously. Furthermore, Algorithm 2.4.8, which is implemented in MAGMA, computes these polynomials. This chapter gives a proof that the output of the algorithm is correct.

The modular polynomial for elliptic curves of level p parametrises p-isogenies of elliptic curves (for p prime) and is defined using the j-invariant. To generalise the modular polynomial to a Hilbert modular setting, we first fix a totally real number field K_0 of degree g over $\mathbb{Q}$, and we write $\mathcal{O}_{K_0}$ for its maximal order. We then need to replace j by an 'isomorphism invariant' for objects $(\mathrm{A}, \xi, \iota) \in \mathbf{POrd}_{\mathbb{C},K_0}$, the category of principally polarised complex abelian g-folds (A, ξ) with an appropriate embedding $\iota : \mathcal{O}_{K_0} \hookrightarrow \mathrm{End}(\mathrm{A})$ (see Definition 1.5.2 for the formal definition). Let $\overline{V}$ be the Hilbert modular variety for $\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$, as in Definition 1.6.11, where $\mathcal{O}_{K_0}^\vee$ is the trace dual of $\mathcal{O}_{K_0}$. Recall from Definition 1.6.9 that $\mathcal{M}_{K_0}(\mathbb{Z})$ denotes the ring of Hilbert modular forms with coefficients in $\mathbb{Z}$, and we write $\mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z}))$ for the field of quotients of modular forms in $\mathcal{M}_{K_0}(\mathbb{Z})$ of equal weight. We will see in Section 2.2 that for some $d \in \mathbb{Z}$, there exist d Hilbert modular functions

$$J_1, \ldots, J_d \in \mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z})),$$

such that the function field of $\overline{V}$ is $\mathbb{C}(J_1, \ldots, J_d)$, and for such $J_1, \ldots, J_d$, there exists a Zariski-open affine subvariety U of $\overline{V}$ such that the rational map

$$(J_1, \ldots, J_d) : U \dashrightarrow \mathbb{A}^d_{\mathbb{C}}$$

is an injective morphism.

Definition 2.1.1. A d-tuple of Hilbert modular functions $(J_1, \ldots, J_d) \in \mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z}))^{\times d}$ such that

$$\mathbb{C}(\overline{V}) = \mathbb{C}(J_1, \ldots, J_d)$$

is a choice of *RM isomorphism invariants* for K_0.

Remark 2.1.2. Fixing U as above, if $(\mathrm{A}, \xi, \iota) \in \mathbf{POrd}_{\mathbb{C},K_0}$ corresponds as in Lemma 1.6.2 to a point in U, then the d-tuple

$$(J_1, \ldots, J_d)(\mathrm{A}, \xi, \iota)$$

determines (A, ξ, ι) up to isomorphism. That is, on U, RM isomorphism invariants are isomorphism invariants in the intuitive sense.

Definition 2.1.3. For a totally positive prime element μ of $\mathcal{O}_{K_0}$, and for $\tau, \tau' \in K_0 \otimes \mathbb{H}$, we say that *there exists a μ-isogeny*

$$\tau \to \tau'$$

if there exists a μ-isogeny

$$(\mathrm{A}, \xi, \iota) \longrightarrow (\mathrm{A}', \xi', \iota',)$$

where the isomorphism classes of (A, ξ, ι) and $(\mathrm{A}', \xi', \iota') \in \mathbf{POrd}_{\mathbb{C},K_0}$ correspond as in Lemma 1.6.2 to the equivalence classes of τ and τ' in V respectively. (Recall from Theorem 1.5.5 and Lemma 1.7.1 that τ and τ' satisfy

$$H_1(\mathrm{A}(\mathbb{C}), \mathbb{Z}) = \tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^{\vee} \quad \text{and} \quad H_1(\mathrm{A}'(\mathbb{C}), \mathbb{Z}) = \tau'\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^{\vee}.)$$

Our higher dimensional analogue of the modular polynomial for elliptic curves will parametrise μ-isogenies of objects in $\mathbf{POrd}_{\mathbb{C},K_0}$, and will be defined using the isomorphism invariants of Definition 2.1.1. The first main theorem of this chapter, given below, gives this higher dimensional analogue of the modular polynomial.

Theorem 2.1.4. *For a totally real number field K_0 of degree g over $\mathbb{Q}$, and a totally positive prime element μ of $\mathcal{O}_{K_0}$, let $\overline{V}$ be the Hilbert modular variety for K_0 (as defined in Definition 1.6.11), and fix a choice of RM isomorphism invariant $(J_1, \ldots, J_d)$ for K_0 (as defined in Definition 2.1.1). Then Algorithm 2.4.8 below outputs a polynomial*

$$G_\mu(X_1, \ldots, X_d, Y) \in \mathbb{Z}[X_1, \ldots, X_d, Y]$$

that has degree $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1$ in Y and such that $\Delta G_\mu(J_1, \ldots, J_d, Y)$ is not constant zero on V, and outputs polynomials

$$H_{\mu,i}(X_1, \ldots, X_d, Y, Z_i) \in \mathbb{Z}[X_1, \ldots, X_d, Y, Z_i]$$

that are linear in Z_i, where $i = 2, \ldots d$. Furthermore, for any choice of Zariski-open subvariety U of $\overline{V}$ such that the map

$$(J_1, \ldots, J_d) : U \to \mathbb{A}^d_{\mathbb{C}}$$

is injective, for all but finitely many

$$[\tau], [\tau'] \in (U \cap V) - \{x \in (U \cap V) : \Delta G_\mu(J_1(x), \ldots, J_d(x), Y) = 0\},$$

there exists a μ-isogeny

$$\tau \to \tau'$$

if and only if

$$G(J_1(\tau), \ldots, J_d(\tau), J_1(\tau')) = 0,$$

and for $i = 2, \ldots, d$,

$$H_{\mu,i}(J_1(\tau), \ldots, J_d(\tau), J_1(\tau'), J_i(\tau')) = 0.$$

Definition 2.1.5. For a totally positive prime element $\mu \in K_0$, we define a *Hilbert modular polynomials of level μ* to be a set of polynomials

$$\left\{ \begin{array}{c} G_\mu(X_1, \ldots, X_d, Y) \in \mathbb{Z}[X_1, \ldots, X_d, Y], \\ H_{\mu,i}(X_1, \ldots, X_d, Y, Z_i) \in \mathbb{Z}[X_1, \ldots, X_d, Y, Z_i] \end{array} \right\}_{i=2,\ldots,d}$$

such that $G_\mu(X_1, \ldots, X_d, Y)$ and $H_{\mu,i}(X_1, \ldots, X_d, Y, Z_i)$ satisfy the conclusions of Theorem 2.1.4.

Remark 2.1.6. Even though Theorem 2.1.4 is over $\mathbb{C}$, in practise we can use it also over finite fields (see Section 2.5).

2.2 Defining RM isomorphism invariants

As before, let K_0 be a totally real number field of degree g over $\mathbb{Q}$, and let $\overline{V}$ be the Hilbert modular variety for $\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$, as defined in Definition 1.6.11. The aim of this section is to prove Proposition 2.2.1.

For completeness, we recall here the definition of RM isomorphism invariants from the previous section.

Definition 2.1.1. *A d-tuple of Hilbert modular functions*

$$(J_1, \ldots, J_d) \in \mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z}))^{\times d}$$

such that

$$\mathbb{C}(\overline{V}) = \mathbb{C}(J_1, \ldots, J_d)$$

is a choice of RM isomorphism invariants *for K_0.*

Proposition 2.2.1. Write $\mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z}))$ for the $\mathbb{Q}$-algebra of quotients of Hilbert modular forms in $\mathcal{M}_{K_0}(\mathbb{Z})$ of equal weight. There exists $d \in \mathbb{Z}$ and a choice

$$J_1, \ldots, J_d \in \mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z}))$$

of RM isomorphism invariant for K_0. Furthermore, for such $J_1, \ldots, J_d$, there exists a Zariski-open affine subvariety U of $\overline{V}$ such that the map

$$(J_1, \ldots, J_d) : U \longrightarrow \mathbb{A}^d_{\mathbb{C}}$$

is a well-defined injective morphism.

Proof. Write $\mathbb{C}(\mathcal{M}_{K_0})$ for the field of quotients of elements of $\mathcal{M}_{K_0}$ of equal weight. By definition of $\overline{V}$ (see Definition 1.6.11), we have that $\mathbb{C}(\overline{V}) = \mathbb{C}(\mathcal{M}_{K_0})$, and by Lemma 1.6.13, we know that

$$\mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z})) \otimes_{\mathbb{Q}} \mathbb{C} = \mathbb{C}(\mathcal{M}_{K_0}).$$

So let $J_1, \ldots, J_d$ be generators of the $\mathbb{Q}$-algebra $\mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z}))$, so that

$$\mathbb{C}(J_1, \ldots, J_d) = \mathbb{C}(\overline{V}),$$

and write W for the image of $(J_1, \ldots, J_d)$ in $\mathbb{A}^d_{\mathbb{C}}$. Then by [Har77, Corollary I.4.5], there are non-empty Zariski-open subsets $U \subseteq \overline{V}$ and $U' \subseteq W$ such that U is isomorphic to U'. □

Example 2.2.2. If $g = 1$, so that $K_0 = \mathbb{Q}$, then we have that

$$\mathrm{SL}_2(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)\backslash K_0 \otimes \mathbb{H} = \mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}.$$

The j-invariant for elliptic curves defines an isomorphism

$$j : \mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H} \longrightarrow \mathbb{A}^1_{\mathbb{C}}.$$

Hence setting

$$V = \mathrm{SL}_2(\mathbb{Z})\backslash\mathbb{H}, \quad \overline{V} = \mathbb{P}^1_{\mathbb{C}}, \quad U = V, \text{ and } \quad J_1 = j$$

gives us $\mathbb{C}(\overline{V}) = \mathbb{C}(J_1)$ and an injective morphism $J_1 : U \to \mathbb{A}^1_{\mathbb{C}}$.

2.3 Algorithm to compute a set of Hilbert modular polynomials

As before, in what follows, K_0 is a totally real number field of degree g over $\mathbb{Q}$ with ring of integers $\mathcal{O}_{K_0}$. From this point on, we fix RM isomorphism invariants $(J_1, \ldots, J_d) \in \mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z}))^{\times d}$, and a non-empty Zariski-open subvariety U of the Hilbert modular variety $\overline{V}$ such that

$$(J_1, \ldots, J_d) : U \longrightarrow \mathbb{A}^d_{\mathbb{C}}$$

defines an injective morphism.

For $i = 1, \ldots, d$, we choose f_i and g_i to be elements of $\mathcal{M}_{K_0}(\mathbb{Z})$ of weight k_i such that

$$J_i = f_i/g_i. \tag{2.1}$$

Definition 2.3.1. Let $\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$ be as in Definition 1.6.1 and let μ be a totally positive prime element of $\mathcal{O}_{K_0}$. Define

$$\Gamma^0(\mu) := \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee) : b \in \mu\mathcal{O}_{K_0}^\vee \right\}.$$

For any $x \in K_0$ define

$$\underline{x} := \begin{pmatrix} x & 0 \\ 0 & 1 \end{pmatrix}.$$

Given a Hilbert modular form $f \in \mathcal{M}_{K_0}(\mathbb{Z})$, for every $N \in \mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$, the function $f|_{\underline{\mu^{-1}}N}$ depends only on the class of N in $\Gamma^0(\mu)\backslash\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$.

Definition 2.3.2. Denote by $\mathcal{C}$ a choice of coset representatives for the quotient of groups

$$\Gamma^0(\mu)\backslash\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee).$$

We then further define

$$\Phi_\mu(Y) := \prod_{M \in \mathcal{C}} \left(g_1|_{\underline{\mu^{-1}}M} Y - f_1|_{\underline{\mu^{-1}}M} \right)$$

and for each $i = 2, \ldots, d$,

$$\Psi_{\mu,i}(Y, Z_i) := \sum_{M \in \mathcal{C}} \left\{ \left(g_i|_{\underline{\mu^{-1}}M} Z_i - f_i|_{\underline{\mu^{-1}}M} \right) \prod_{\substack{M' \in \mathcal{C} \\ M' \neq M}} \left(g_1|_{\underline{\mu^{-1}}M'} Y - f_1|_{\underline{\mu^{-1}}M'} \right) \right\}.$$

Note that the definitions of $\Phi_\mu(Y)$ and $\Psi_{\mu,i}(Y, Z_i)$ do not depend on the choice of coset representatives for $\Gamma^0(\mu)\backslash\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$.

Remark 2.3.3. We have that

$$\Phi_\mu(Y) \in \mathcal{M}_{K_0}(\mathbb{Z})[Y] \quad \text{and} \quad \Psi_{\mu,i}(Y, Z_i) \in \mathcal{M}_{K_0}(\mathbb{Z})[Y, Z_i].$$

Proof. Recall that for $M \in \mathcal{C}$ and $N \in \mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$, for every $f \in \mathcal{M}_{K_0}$, we have that

$$(f|_{\underline{\mu^{-1}}M})|_N(\tau) = f|_{\underline{\mu^{-1}}MN}(\tau).$$

In particular, acting by $|_N$ on the coefficients of $\Phi_\mu(Y)$ (and $\Psi_{\mu,i}(Y, Z_i)$) just permutes the factors (or terms) of the defining product (or sum), leaving $\Phi_\mu(Y)$ (and $\Psi_{\mu,i}(Y, Z_i)$) unchanged, hence the coefficients are modular forms for $\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$. □

As Φ_μ is a univariate polynomial with coefficients that are modular forms for $\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$ of equal weight, the discriminant $\Delta\Phi_\mu$ is also a modular form for $\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$. In particular, whether or not $(\Delta\Phi_\mu)(\tau) = 0$ depends only on the class of τ in V.

Proposition 2.3.4. Fix notation as in Definition 2.3.2 and recall from Definition 2.1.3 the definition of a μ-isogeny $\tau \to \tau'$ for $\tau, \tau' \in K_0 \otimes \mathbb{H}$. For any $\tau, \tau' \in K_0 \otimes \mathbb{H}$ such that the classes $[\tau]$ and $[\tau']$ of τ and τ' in $\overline{V}$ are in

$$(U \cap V) - \{x \in (U \cap V) : (\Delta\Phi_\mu)(x) = 0\},$$

there exists a μ-isogeny $\tau \to \tau'$ if and only if for every $i = 2, \ldots, d$, evaluating $\Phi_\mu(Y)$ and $\Psi_{\mu,i}(Y, Z_i)$ at $(Y, Z_2, \ldots, Z_d) = (J_1([\tau']), \ldots, J_d([\tau']))$, and then evaluating the resulting modular forms at τ, gives

$$\left(\Phi_\mu(J_1([\tau']))\right)(\tau) = 0 \quad \text{and} \quad \left(\Psi_{\mu,i}(J_1([\tau']), J_i([\tau']))\right)(\tau) = 0.$$

Lemma 2.3.5. If μ is a totally positive prime element of $\mathcal{O}_{K_0}$ then the set $\Gamma^0(\mu)\backslash\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$ has $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1$ elements.

Proof. Define

$$k := \max\{n \in \mathbb{Z} : (\mathcal{O}_{K_0}^\vee)^{-1} \subseteq \mu^n \mathcal{O}_{K_0}\}.$$

There is a bijection of sets

$$\begin{array}{ccc} \Gamma^0(\mu)\backslash \mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee) & \longleftrightarrow & (\underline{\mu^k}\Gamma^0(\mu)\underline{\mu^{-k}})\backslash(\underline{\mu^k}\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)\underline{\mu^{-k}}) \\ M & \mapsto & \underline{\mu^k} M \underline{\mu^{-k}}. \end{array}$$

We claim that

$$\left\{ \begin{pmatrix} 1 & a \\ 0 & 1 \end{pmatrix} : a \in \mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0} \right\} \cup \left\{ \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \right\}$$

is in bijection with $(\underline{\mu^k}\Gamma^0(\mu)\underline{\mu^{-k}})\backslash(\underline{\mu^k}\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)\underline{\mu^{-k}})$. Let

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \underline{\mu^k}\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)\underline{\mu^{-k}}.$$

Then $a, d \in \mathcal{O}_{K_0}$, $b \in \mu^k \mathcal{O}_{K_0}^\vee \subseteq (\mathcal{O}_{K_0})_{(\mu)}$ and $c \in \mu^{-k}(\mathcal{O}_{K_0}^\vee)^{-1} \subseteq (\mathcal{O}_{K_0})_{(\mu)}$, so that in particular, reduction by μ defines a group homomorphism

$$r : \underline{\mu^k}\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)\underline{\mu^{-k}} \to \mathrm{SL}_2(\mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0}).$$

Now $\mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0}$ is a field as $\mu\mathcal{O}_{K_0}$ is prime, and $\mathrm{SL}_2(\mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0})$ acts on $\mathbb{P}^1(\mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0})$ as

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot (x : y) \mapsto (ax + by : cx + dy).$$

The stabilizer of $(0 : 1)$ is

$$\left\{ \begin{pmatrix} a & 0 \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0}) \right\},$$

the pull-back of which under r is $\underline{\mu^k}\Gamma^0(\mu)\underline{\mu^{-k}}$, so the bijection follows from the orbit-stabilizer theorem. □

We will prove Proposition 2.3.4 by using the above lemma and a representation of μ-isogenies up to isomorphism.

Definition 2.3.6. We say μ-isogenies $f : (\mathrm{A}, \xi_\mathrm{A}, \iota_\mathrm{A}) \to (\mathrm{B}, \xi_\mathrm{B}, \iota_\mathrm{B})$ and $g : (\mathrm{A}, \xi_\mathrm{A}, \iota_\mathrm{A}) \to (\mathrm{B}', \xi_{\mathrm{B}'}, \iota_{\mathrm{B}'})$ are *isomorphic* if there exists a 1-isogeny $\varphi : (\mathrm{B}, \xi_\mathrm{B}, \iota_\mathrm{B}) \to (\mathrm{B}', \xi_{\mathrm{B}'}, \iota_{\mathrm{B}'})$ such that the diagram

$$\begin{tikzcd} (\mathrm{A}, \xi_\mathrm{A}, \iota_\mathrm{A}) \arrow[r, "f"] \arrow[dr, "g"'] & (\mathrm{B}, \xi_\mathrm{B}, \iota_\mathrm{B}) \arrow[d, "\varphi"] \\ & (\mathrm{B}', \xi_{\mathrm{B}'}, \iota_{\mathrm{B}'}) \end{tikzcd}$$

commutes.

Definition 2.3.7. For every $\begin{pmatrix} a & b \\ c & d \end{pmatrix} = M \in \mathrm{GL}_2(K_0)^+$ and for every $\tau \in K_0 \otimes \mathbb{H}$, we define $\varphi_{M,\tau}$ to be the element of $\mathrm{Hom}_{\mathbf{Ord}_{\mathbb{C},K_0}}(\tau, M\tau) \otimes \mathbb{Q}$ that is multiplication by $(c\tau + d)^{-1}$ on $K_0 \otimes \mathbb{C}$.

Note that

$$\varphi_{B,A\tau} \circ \varphi_{A,\tau} = \varphi_{BA,\tau} \tag{2.2}$$

and

$$\varphi_{M,\tau}^{-1} = \varphi_{M^{-1},M\tau}. \tag{2.3}$$

Lemma 2.3.8. We have that $\varphi_{M,\tau}$ is an isomorphism in $\mathbf{POrd}_{\mathbb{C},K_0}$ if and only if $M \in \mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$.

Proof. Write $M = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ and for any $\tau' \in K_0 \otimes \mathbb{H}$ let $E_{\tau'}$ be the Riemann form

$$E_{\tau'}(u_1\tau + u_2, v_1\tau' + v_2) = \mathrm{tr}_{K_0/\mathbb{Q}}(u_1v_2 - u_2v_1).$$

We get commutative diagram of unpolarised abelian varieties, where the dashed arrows are automorphisms of $K_0 \otimes \mathbb{C}$ that may or may not induce actual maps of abelian varieties:

$$\begin{tikzcd}
(K_0 \otimes \mathbb{C})/(\tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee) \arrow[r, dashed, "\varphi_{M,\tau}:=(c\tau+d)^{-1}"] \arrow[dr, dashed, "f:=\mathrm{id}_{(K_0\otimes\mathbb{C})}"] & (K_0 \otimes \mathbb{C})/(M\tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee) \arrow[d, "c\tau+d"] \\
& (K_0 \otimes \mathbb{C})/((a\tau+b)\mathcal{O}_{K_0} + (c\tau+d)\mathcal{O}_{K_0}^\vee).
\end{tikzcd}$$

Now f, and hence ϕ defines an isomorphism on lattices if and only if $M \in \mathrm{GL}(\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee)$. Suppose now that $M \in \mathrm{GL}(\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee)$. It remains to show that $\det(M) = 1$ if and only if ϕ is an isomorphism in $\mathbf{POrd}_{\mathbb{C},K_0}$, that is, if

$$E_\tau(\alpha, \beta) = E_{M\tau}(\phi(\alpha), \phi(\beta)).$$

Write $E_\tau = \mathrm{tr}_{K_0/\mathbb{Q}} \circ S_\tau$ and $E_{M\tau} = \mathrm{tr}_{K_0/\mathbb{Q}} \circ S_{M\tau}$. The matrices of S_τ and $\phi^* S_{M\tau}$ with respect to the $(K_0 \otimes \mathbb{R})$-basis $\{\tau, 1\}$ of $K_0 \otimes \mathbb{C}$ are

$$\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$$

and

$$M \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} M^t$$

respectively, so $S_\tau = \phi^* S_{M\tau}$ if and only if $\det(M) = 1$ and the result follows. □

Lemma 2.3.9. Fix a totally positive prime element $\mu \in K_0$. Then for any $\tau \in K_0 \otimes \mathbb{H}$, there is a map

$$\begin{array}{rccc} i: & \Gamma^0(\mu)\backslash\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee) & \longrightarrow & \left\{ \mu\text{-isogenies from } \tau \right\}_{/\cong} \\ & M & \mapsto & \varphi_{\underline{\mu^{-1}},M\tau} \circ \varphi_{M,\tau}, \end{array}$$

and i defines a bijection of sets.

Proof. Observe that $\mathrm{id}_{K_0\otimes\mathbb{C}}$ defines a μ-isogeny

$$\varphi_{\underline{\mu^{-1}},\tau} : \left((K_0 \otimes \mathbb{C})/(\tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee), \xi, \iota\right) \longrightarrow \left((K_0 \otimes \mathbb{C})/(\mu^{-1}\tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee), \mu\xi, \iota\right),$$

where $\xi = (\tau - \overline{\tau})^{-1}$, which in other words is a μ-isogeny $\tau \to \mu^{-1}\tau$. Replacing τ by $M\tau$ for

$$M \in \Gamma^0(\mu)\backslash\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$$

it is easy to see that i is well-defined on $\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$.

We claim further that i is a well-defined injection of sets. Let $M, N \in \mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$ and suppose that $\varphi_{\underline{\mu^{-1}},M\tau} \circ \varphi_{M,\tau}$ and $\varphi_{\underline{\mu^{-1}},N\tau} \circ \varphi_{N,\tau}$ are isomorphic as μ-isogenies. That is, there exists an isomorphism $\psi : \underline{\mu^{-1}}M\tau \to \underline{\mu^{-1}}N\tau$ such that

$$\psi \circ \varphi_{\underline{\mu^{-1}},M\tau} \circ \varphi_{M,\tau} = \varphi_{\underline{\mu^{-1}},N\tau} \circ \varphi_{N,\tau}, \tag{2.4}$$

hence by (2.2) and (2.3)

$$\psi = \varphi_{\underline{\mu^{-1}}NM^{-1}\underline{\mu},\underline{\mu^{-1}}M\tau}. \tag{2.5}$$

By Lemma 2.3.8, as ψ is an isomorphism, we have that $\underline{\mu^{-1}}NM^{-1}\underline{\mu} \in \mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$. Define $X = NM^{-1}$ and $T = \underline{\mu^{-1}}NM^{-1}\underline{\mu}$. As T and $X \in \mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$, we get further that $X \in \Gamma^0(\mu)$. Conversely, suppose that $NM^{-1} \in \Gamma^0(\mu)$. Then $\underline{\mu^{-1}}NM^{-1}\underline{\mu} \in \mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$, so ψ defined by (2.5) is an isomorphism. Hence i is a well-defined injection of sets.

To show that i is in fact a bijection we proceed by counting. By Lemma 2.3.5 the set $\mathcal{C}$ has $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1$ elements, so we just need to show that there are at most $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1$ non-isomorphic μ-isogenies from any given $\tau \in K_0 \otimes \mathbb{H}$. If $f : (\mathrm{A}, \xi_\mathrm{A}, \iota_\mathrm{A}) \to (\mathrm{B}, \xi_\mathrm{B}, \iota_\mathrm{B})$ is a μ-isogeny, then

$$\ker(f) \subseteq \ker(\mu) \cong (\mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0})^{\times 2}.$$

Also, as for every $\alpha \in \mathcal{O}_{K_0}$ the following diagram commutes:

$$\begin{tikzcd}
\ker(f) \arrow[r] \arrow[d, dashed] & \mathrm{A} \arrow[r, "f"] \arrow[d, "\iota_{\mathrm{A}}(\alpha)"'] & \mathrm{B} \arrow[d, "\iota_{\mathrm{B}}(\alpha)"'] \\
\ker(f) \arrow[r] & \mathrm{A} \arrow[r, "f"] & \mathrm{B},
\end{tikzcd}$$

the kernel of f is an $\mathcal{O}_{K_0}$-module, and hence an $\mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0}$ sub-vector space of $(\mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0})^{\times 2}$. Then, as $\deg(f) = \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu)$, there are at most $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1$ distinct kernels of μ-isogenies from any given τ (or equivalently any given $(\mathrm{A}, \xi, \iota) \in \mathbf{POrd}_{\mathbb{C},K_0}$). Therefore it remains to show that there do not exist non-isomorphic μ-isogenies $f : (\mathrm{A}, \xi_{\mathrm{A}}, \iota_{\mathrm{A}}) \to (\mathrm{B}, \xi_{\mathrm{B}}, \iota_{\mathrm{B}})$ and $f' : (\mathrm{A}, \xi_{\mathrm{A}}, \iota_{\mathrm{A}}) \to (\mathrm{B}', \xi_{\mathrm{B}'}, \iota_{\mathrm{B}'})$ with the same kernel. By the universal property of quotient maps there exists an isomorphism α (of unpolarised abelian varieties) such that the following diagram commutes:

$$\begin{tikzcd}
\mathrm{A} \arrow[r, "f"] \arrow[dr, "f'"'] & \mathrm{B} \arrow[d, "\alpha"] \\
& \mathrm{B}'.
\end{tikzcd}$$

We claim that α is a 1-isogeny. Consider the following diagram:

$$\begin{tikzcd}
\mathrm{A} \arrow[dr, "\xi_{\mathrm{A}}"'] & \mathrm{A} \arrow[l, "\iota_{\mathrm{A}}(\mu)"'] \arrow[r, "f"] & \mathrm{B} \arrow[r, "\alpha"] \arrow[d, "\xi_{\mathrm{B}}"] \arrow[dl, phantom, "(1)"] & \mathrm{B}' \arrow[d, "\xi_{\mathrm{B}'}"] \arrow[dl, phantom, "(2)"] \\
& \mathrm{A}^\vee & \mathrm{B}^\vee \arrow[l, "f^\vee"] & \mathrm{B}'^\vee. \arrow[l, "\alpha^\vee"]
\end{tikzcd}$$

Diagram (1) commutes as f is a μ-isogeny and the diagram formed by the outside arrows commutes as f' is a μ-isogeny, hence diagram (2) commutes. Similarly, consider the following diagram:

$$\begin{tikzcd}
\mathrm{End}(\mathrm{A}) \otimes \mathbb{Q} \arrow[r, "\beta \mapsto f \circ \beta \circ f^{-1}"] \arrow[dr, phantom, "(1)"] & \mathrm{End}(\mathrm{B}) \otimes \mathbb{Q} \arrow[r, "\beta \mapsto \alpha \circ \beta \circ \alpha^{-1}"] \arrow[dr, phantom, "(2)"] & \mathrm{End}(\mathrm{B}') \otimes \mathbb{Q} \\
& K_0 \arrow[ul, "\iota_{\mathrm{A}}"] \arrow[u, "\iota_{\mathrm{B}}"] \arrow[ur, "\iota_{\mathrm{B}'}"'] &
\end{tikzcd}$$

Diagram (1) commutes as f is a μ-isogeny and the diagram formed by the outside arrows commutes as f' is a μ-isogeny and

$$f' \circ \beta(f')^{-1} = (\alpha \circ f) \circ \beta \circ (\alpha \circ f)^{-1} = \alpha \circ (f \circ \beta \circ f^{-1}) \circ \alpha^{-1}.$$

Hence (2) commutes, so α is a 1-isogeny and f and f' are isomorphic as μ-isogenies. □

Proof of Proposition 2.3.4. Suppose first that there exists a μ-isogeny $\tau \to \tau'$. Then by Lemma 2.3.9, there exists $N \in \mathcal{C} = \Gamma^0(\mu)\backslash \mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)$ such that this μ-isogeny is isomorphic to a μ-isogeny $\tau \to \mu^{-1}N\tau$, so we can identify τ' with $\mu^{-1}N\tau$. Plugging this into the definitions of $\Phi_\mu(Y)$ and $\Psi_{\mu,i}(Y, Z_i)$, we get

$$\Phi_\mu(J_1(\mu^{-1}N\tau)) = 0$$

and

$$\Psi_{\mu,i}(J_1(\mu^{-1}N\tau), J_i(\mu^{-1}N\tau)) = 0.$$

Suppose now that $(Y_0, Z_{2,0}, \ldots, Z_{d,0})$ is a common root of $\Phi_\mu(Y)$ and $\Psi_{\mu,i}(Y, Z_i)$. One can see directly from the definition of Φ_μ and $\Psi_{\mu,i}$ that under the discriminant condition, the set of common roots of (2.3.2) is exactly the set

$$\{(J_1(\mu^{-1}M\tau), \ldots, J_d(\mu^{-1}M\tau)) : M \in \mathcal{C}\}.$$

Therefore, there exists $N \in \mathcal{C}$ such that

$$(Y_0, Z_{2,0}, \ldots, Z_{d,0}) = (J_1(\mu^{-1}N\tau), \ldots, J_d(\mu^{-1}N\tau)),$$

and by Lemma 2.3.9 there exists a μ-isogeny

$$\tau \to \mu^{-1}N\tau.$$

□

2.4 Computing the RM isomorphism invariants for a given genus 2 curve

In Definition 2.1.1, we defined RM isomorphism invariants for elements of $\mathbf{POrd}_{\mathbb{C},K_0}$. Restrict now to the dimension 2 case. It is however not immediately clear how to compute these given the equation of a genus 2 curve. We have a computational advantage in genus 2, which is that there already exist Igusa-Clebsch invariants to determine a curve up to isomorphism.

Definition 2.4.1. For a curve C of genus 2 over a field k with $\mathrm{char}(k) \neq 2$, there exists a hyperelliptic model $y^2 = f(x)$ of C, where f is a separable polynomial of degree 6. Fix such a model, denote by c the leading coefficient of f, fix an ordering $x_1, \ldots, x_6$ of the roots of f in its splitting field, and denote by (ij) the difference $x_i - x_j$. For $\mathrm{char}(k) \neq 2, 3, 5$, we define the *Igusa-Clebsch invariants* of C to be

$$\begin{aligned} I_2 &= c^2 \sum (12)^2(34)^2(56)^2, \\ I_4 &= c^4 \sum (12)^2(23)^2(31)^2(45)^2(56)^2(64)^2, \\ I_6 &= c^6 \sum (12)^2(23)^2(31)^2(45)^2(56)^2(64)^2(14)^2(25)^2(36)^2, \\ I_{10} &= c^{10} \prod (12)^2, \end{aligned}$$

where each sum and product runs over the distinct expressions obtained by applying a permutation to the index set $\{1, \ldots, 6\}$.

These invariants are integral whenever f is integral. The Igusa-Clebsch invariants are 'invariants for the Siegel moduli space'. Before making this more precise, we recall some facts about the Siegel moduli space.

Definition 2.4.2. We define

$$\mathrm{Sym}_2(\mathbb{C}) = \left\{ \begin{pmatrix} \tau_1 & \tau_2 \\ \tau_2 & \tau_3 \end{pmatrix} \in \mathrm{Mat}_{2\times 2}(\mathbb{C}) \right\},$$

and for $\tau \in \mathrm{Sym}_2(\mathbb{C})$, we write $\mathrm{Im}(\tau) > 0$ for '$\mathrm{Im}(\tau)$ is positive definite'.

Definition 2.4.3. The *Siegel upper half space* is defined to be

$$\mathbb{H}_2 = \left\{ \tau = \begin{pmatrix} \tau_1 & \tau_2 \\ \tau_2 & \tau_4 \end{pmatrix} \in \mathrm{Sym}_2(\mathbb{C}) : \mathrm{Im}(\tau) > 0 \right\},$$

and the symplectic group

$$\mathrm{Sp}_2(\mathbb{Z}) = \left\{ \gamma \in \mathrm{GL}_4(\mathbb{Z}) : \gamma \begin{pmatrix} 0 & I_2 \\ -I_2 & 0 \end{pmatrix} \gamma^{\mathrm{tr}} = \begin{pmatrix} 0 & I_2 \\ -I_2 & 0 \end{pmatrix} \right\}$$

acts on $\mathbb{H}_2$ via

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} \cdot \tau = (A\tau + B)(C\tau + D)^{-1}.$$

The field of rational functions of the coarse moduli space for hyperelliptic curves of genus 2 can be generated by three Siegel modular functions, as shown by Igusa in [Igu60]. Following the notation in the Echidna database [Echidna], we choose as generators three Siegel modular functions

$$i_1, i_2, i_3 : \mathrm{Sp}_2(\mathbb{Z}) \backslash \mathbb{H}_2 \longrightarrow \mathbb{C}$$

such that, if C is a curve of genus 2, and $[\tau] \in \mathrm{Sp}_2(\mathbb{Z}) \backslash \mathbb{H}_2$ is the point in the moduli space corresponding to C, then

$$i_1(\tau) = (I_4 I_6 / I_{10})(C), \tag{2.6}$$

$$i_2(\tau) = (I_2^3 I_4 / I_{10})(C), \tag{2.7}$$

$$i_3(\tau) = (I_2^2 I_6 / I_{10})(C). \tag{2.8}$$

Now, for a totally real quadratic number field K_0, the forgetful functor

$$\begin{array}{ccc} \mathbf{POrd}_{\mathbb{C},K_0} & \longrightarrow & \mathbf{POrd}_{\mathbb{C},2} \\ (\mathrm{A},\xi,\iota) & \mapsto & (\mathrm{A},\xi) \end{array}$$

induces a map

$$\phi : \mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^{\vee})\backslash K_0 \otimes \mathbb{H} \to \mathrm{Sp}_2(\mathbb{Z})\backslash \mathbb{H}_2,$$

which is generically 2-1. We will refer to this as the *modular map*. The image of this map is called the *Humbert surface for* K_0, and is denoted as $\mathcal{H}_{K_0}$. That is, the modular map ϕ induces a degree 2 map

$$\phi : \mathcal{M}_{K_0} \longrightarrow \mathcal{H}_{K_0}.$$

In particular, as there exist 2 algebraically independent Siegel modular functions f_1 and f_2 in

$$\mathbb{C}(\mathcal{H}_{K_0}) \subseteq \mathbb{C}(i_1, i_2, i_3),$$

we get 2 algebraically independent Hilbert modular functions

$$J_1 = \phi^* f_1 \quad \text{and} \quad J_2 = \phi^* f_2 \tag{2.9}$$

in $\mathbb{C}(\mathcal{M}_{K_0})$. Also, by construction, we get that J_1 and J_2 are *symmetric*, that is, that if σ is the generator of $\mathrm{Gal}(K_0/\mathbb{Q})$, then for all $\tau \in K_0 \otimes \mathbb{H}$, we have that

$$J_1(\sigma(\tau)) = J_1(\tau) \quad \text{and} \quad J_2(\sigma(\tau)) = J_2(\tau).$$

By Proposition 1.6.12 and Lemma 1.6.13, we have that $\mathbb{C}(\overline{V})$ is a finite separable field extension of $\mathbb{C}(J_1, J_2)$ and hence is generated by one element; choose such an element and denote it by J_3. Write $m(X) \in \mathbb{C}(J_1, J_2)[X]$ for the minimal polynomial of J_3; then $m(X)$ is the pullback along ϕ of a polynomial in $\mathbb{C}(i_1, i_2, i_3)[X]$.

The subtlety of how to choose the root of $m(X)$ in practice is addressed in Algorithm 2.5.4, Step 2.

Example 2.4.4. Gundlach [Gun63] and Müller [Mue85] computed formulae for a choice of isomorphism invariants J_1, J_2, and J_3 for $K_0 = \mathbb{Q}(\sqrt{5})$, and gave the functions from which J_1, J_2, and J_3^2 (here $m(X)$ is quadratic and without a linear term) are pulled back along ϕ:

$$J_1 = \phi^* \left(\frac{2^{-6}3^{-3}i_1^2 i_2^2 + 2^{-3}3^2 i_1 i_2^2 - 2^{-4}3^{-3} i_1 i_3^3 + 2^{-5}3^2 i_2 i_3^2}{i_1^2 i_2^2 + 2^2 3^5 i_1 i_2^2} \right), \tag{2.10}$$

$$J_2 = \phi^* \left(\frac{2^9 i_1^3 i_2^2 + 2^{11} 3^5 i_1^2 i_2^2}{i_1^2 i_2^2 + 2^2 i_1 i_3^3 - 2 \cdot 3^5 i_2 i_3^2} \right), \tag{2.11}$$

$$\begin{aligned} J_3^2 = & 5^5 - 2^{-1}5^3 J_1 J_2 + 2^{-4} J_2 + 2^{-1}3^2 5^2 J_2^2 J_1^3 - 2^{-3} J_1^2 J_2^2 - 2 \cdot 3^3 J_2^3 J_1^5 \\ & + 2^{-4} J_2^3 J_1^4 \end{aligned} \tag{2.12}$$

Remark 2.4.5. For each choice of K_0, we have to recalculate RM isomorphism invariants J_1, J_2, and J_3. In [LNY16, Theorem 2.2], Lauter, Naehrig, and Yang give a method to calculate a choice of Siegel modular functions f_1 and f_2 as in (2.9), but the minimal polynomial of J_3 over $\mathbb{Q}(J_1, J_2)$ is not known in general.

Recall from Lemma 1.6.13 that $\mathbb{C}(\overline{V}) = \mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z})) \otimes \mathbb{C}$, so that in particular a choice of $\mathbb{Q}$-algebra generators $J_1, \ldots, J_d$ for $\mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z})$ is also a choice of $\mathbb{C}$-algebra generators for $\mathbb{C}(\overline{V})$. In the cases for which a complete set of generators is known, namely K_0 of discriminant 5, 8, 13, and 17, we can choose RM isomorphism invariants $J_1, J_2, J_3 \in \mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z}))^{\times 3}$ for which J_1 and J_2 are symmetric Hilbert modular functions (as above) and $J_3^2 \in \mathbb{Q}(J_1, J_2)$. For simplicity, we restrict to this case in all that follows.

2.4.1 The algorithm

Given the coefficients of the q-expansions of the numerators and denominators of $J_1, \ldots, J_d$ up to a high enough precision (see the implementation at `www.martindale.info` for details on the precision), using Lemma 2.5.2 and the formulae for $\Phi_\mu(Y)$ and $\Psi_{\mu,i}(Y, Z_i)$ given in Definition 2.3.2 we can write out explicit formulae for the q-expansions of the coefficients (with respect to Y and Z_i) up to some precision of $\Phi_\mu(Y)$

and $\Psi_{\mu,i}(Y, Z_i)$. Fix $\mathbb{Q}$-algebra generators of $\mathcal{M}_{K_0}(\mathbb{Q})$ to be $\gamma_1, \ldots, \gamma_s \in \mathcal{M}_{K_0}(\mathbb{Z})$ of weights $\kappa_1, \ldots, \kappa_s$ respectively (recall from Remark 1.6.15 that we assumed s to be finite), and assume that we also know sufficiently many coefficients of the q-expansions of $\gamma_1, \ldots, \gamma_s$. Then for each coefficient $f \in \mathcal{M}_{K_0}(\mathbb{Z})$ of $\Phi_\mu(Y)$ or $\Psi_{\mu,i}(Y, Z_i)$ it is just linear algebra to determine integers $h_1, \ldots, h_s$ and rational numbers $b_{\underline{h}}$, where $\underline{h} = (h_1, \ldots, h_s)$, such that

$$f = \sum_{\left\{\underline{h} \in (\mathbb{Z}_{\geq 0})^s : \sum_{j=1}^{s} h_j \kappa_j = k\right\}} b_{\underline{h}} \prod_{j}^{s+1} \gamma_j^{h_j}, \tag{2.13}$$

where k is the weight of f. To deduce the Hilbert modular polynomials G_μ and $H_{\mu,i}$ from Φ_μ and $\Psi_{\mu,i}$, we first have to scale Φ_μ and $\Psi_{\mu,i}$ so that the coefficients are in $\mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z}))$. To do this, we construct a ring homomorphism

$$\mathcal{M}_{K_0}(\mathbb{Z}) \longrightarrow \mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z})).$$

To this end, we define

$$d = \gcd(\{\kappa : \mathcal{M}_{K_0,\kappa} \neq \emptyset\})$$

and choose w_1 and w_2 such that $\mathcal{M}_{K_0,w_1}(\mathbb{Z}), \mathcal{M}_{K_0,w_2}(\mathbb{Z}) \neq \emptyset$ and $d = w_1 - w_2$. Then choose

$$\varphi \in \mathcal{M}_{K_0,w_2}(\mathbb{Z}) \quad \text{and} \quad \psi \in \mathcal{M}_{K_0,w_1}(\mathbb{Z}), \tag{2.14}$$

and define

$$\varphi_i = \varphi^{\kappa_i/d} \quad \text{and} \quad \psi_i = \psi^{\kappa_i/d}.$$

This defines a map

$$\begin{array}{ccc} \mathcal{M}_{K_0,\kappa_i}(\mathbb{Z}) & \longrightarrow & \mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z})) \\ \gamma_i & \mapsto & \frac{\varphi_i}{\psi_i}\gamma_i \end{array}$$

which extends $\mathbb{Z}$-linearly to a map

$$\rho : \mathcal{M}_{K_0}(\mathbb{Z}) \longrightarrow \mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z})), \tag{2.15}$$

which is in fact a ring homomorphism. In Algorithm 2.4.8, we will assume that the representations of $\rho(\gamma_1), \ldots, \rho(\gamma_s)$ as rational functions in $J_1, \ldots, J_d$ are known.

Example 2.4.6. Müller [Mue85] defined four elements $(\gamma_1, \gamma_2, \gamma_3, \gamma_4) = (g_2, s_5, g_6, s_{15})$ of $\mathcal{M}_{\mathbb{Q}(\sqrt{5})}(\mathbb{Z})$ of weights 2, 5, 6, and 15 respectively that generate $\mathcal{M}_{\mathbb{Q}(\sqrt{5})}(\mathbb{Q})$ as a $\mathbb{Q}$-algebra and defined modular functions

$$(J_1, J_2, J_3) = \left(\frac{g_2^5}{s_5^2}, \frac{s_6}{g_2^3}, \frac{s_5^3}{s_{15}}\right), \tag{2.16}$$

such that $\mathbb{Q}(\mathcal{M}_{K_0}(\mathbb{Z})) = \mathbb{Q}(J_1, J_2, J_3)$. In this case, we get that $d = 1$, we choose $w_1 = 5$ and $w_2 = 4$, and we choose $\varphi = g_2^2$ and $\psi = s_5$. Then

$$\begin{aligned}
\gamma_1 = g_2 &\mapsto \frac{g_2^5}{s_5^2} = J_1 \\
\gamma_2 = s_5 &\mapsto \frac{g_2^{10}}{s_5^4} = \left(\frac{g_2^5}{s_5^2}\right)^2 = J_1^2 \\
\gamma_3 = s_6 &\mapsto \frac{g_2^{12} s_6}{s_5^6} = \left(\frac{g_2^5}{s_5^2}\right)^3 \frac{s_6}{g_2^3} = J_1^3 J_2 \\
\gamma_4 = s_{15} &\mapsto \frac{g_2^{30} s_{15}}{s_5^{15}} = \left(\frac{g_2^5}{s_5^2}\right)^6 \frac{s_{15}}{s_5^3} = J_1^6 J_3^{-1}.
\end{aligned}$$

The choice given in Equation (2.16) is the choice in the implementation of Algorithm 2.4.8 that can be found at `www.martindale.info`.

The following algorithm computes a set of Hilbert modular polynomials in the sense of Definition 2.1.5.

Lemma 2.4.7. Let k_i be the weight of ψ_i (the denominator of J_i). Let $\Phi_\mu(Y)$ and $\Psi_{\mu,i}(Y, Z_i)$ be as in Definition 2.3.2. There exist modular forms $y_0, \ldots, y_{|\mathcal{C}|} \in \mathcal{M}_{K_0}$ of weight $|\mathcal{C}|k_1$, and for $i = 2, \ldots, d$, there exist modular forms $z_{i,0}, z'_{i,0}, \ldots, z_{i,|\mathcal{C}|-1}, z'_{i,|\mathcal{C}|-1} \in \mathcal{M}_{K_0}$ of weight $(|\mathcal{C}| - 1)k_1 + k_i$ such that

$$\Phi_\mu(Y) = \sum_{n=0}^{|\mathcal{C}|} y_n Y^n$$

and

$$\Psi_{\mu,i}(Y, Z_i) = \sum_{n=0}^{|\mathcal{C}|-1} (z_{i,n} Z_i - z'_{i,n}) Y^n.$$

Proof. This follows from the explicit formulae in Definition 2.3.2. □

Algorithm 2.4.8.
INPUT: A totally real number field K_0 of degree g over $\mathbb{Q}$, the q-expansions of generators $\gamma_1, \ldots, \gamma_s$ of the $\mathbb{Q}$-algebra $\mathcal{M}_{K_0}(\mathbb{Q})$ (up to a certain precision), the images of $\gamma_1, \ldots, \gamma_s$ under ρ as rational functions of $J_1, \ldots, J_d$, and a totally positive element $\mu \in \mathcal{O}_{K_0}$ that generates a prime ideal.
OUTPUT: Polynomials

$$G_\mu(X_1, \ldots, X_d, Y) \in \mathbb{Z}[X_1, \ldots, X_d, Y]$$
$$H_{\mu,i}(X_1, \ldots, X_d, Y, Z_i) \in \mathbb{Z}[X_1, \ldots, X_d, Y, Z_i],$$

for $i = 2, \ldots, d$, satisfying the conclusions of Theorem 2.1.4.

1. Compute the q-expansions of the coefficients of Φ_μ and $\Psi_{\mu,i}$ up to precision P. For more details in genus 2, see Remark 2.5.3. For details on how to compute the required precision, see the MAGMA code, which can be found at `www.martindale.info`.

2. As in (2.13), write each coefficient of Φ_μ and $\Psi_{\mu,i}$ as elements of $\mathbb{Z}[\gamma_1, \ldots, \gamma_s]$ using linear algebra on the q-expansions (here it is necessary to have chosen the precision of the q-expansions to be sufficiently large).

3. For each i, the input contains an expression

   $$\tilde{\rho}(\gamma_i) \in \mathbb{Q}(X_1, \ldots, X_d)$$

   such that

   $$\tilde{\rho}(\gamma_i)(J_1, \ldots, J_d) = \rho(\gamma_i).$$

   Define

   $$G_\mu(X_1, \ldots, X_d, Y) \in \mathbb{Z}[X_1, \ldots, X_d, Y]$$

   to be the numerator of $\tilde{\rho}(\Phi_\mu(Y))$ and

   $$H_{\mu,i}(X_1, \ldots, X_d, Y, Z_i) \in \mathbb{Z}[X_1, \ldots, X_d, Y, Z_i]$$

   to be the numerator of $\tilde{\rho}(\Psi_{\mu,i}(Y, Z_i))$.

We have implemented a more optimised version of this in MAGMA for $K_0 = \mathbb{Q}(\sqrt{5})$ and $K_0 = \mathbb{Q}(\sqrt{2})$, see Section 2.5. That the output of Algorithm 2.4.8 is correct was in the statement of Theorem 2.1.4, which we now prove:

Proof of Theorem 2.1.4. Define $D_1 \in \mathcal{M}_{K_0}(\mathbb{Z})[Y]$ to be the denominator of $\rho(\Phi_\mu(Y))$ and

$$D_i \in \mathcal{M}_{K_0}(\mathbb{Z})[Y, Z_i]$$

to be the denominator of $\rho(\Psi_{\mu,i}(Y, Z_i))$. Let

$$S = \{[\tau] \in U \cap V : D_1(J_1(\tau)) = 0\} \cup \{[\tau] :\in U \cap V : D_i(J_1(\tau), J_i(\tau)) = 0\}.$$

Then S is a finite set, as D_1 and D_i have finitely many roots, and for any value $r \in \mathbb{C}$ and any $1 \leq i \leq d$, there are finitely many $[\tau]$ such that $J_i(\tau) = r$ as J_i extends to a holomorphic function on the compact set $\overline{V}$.

It is immediate from Proposition 2.3.4 that the roots of $(\Phi_\mu(Y))(\tau)$ are given by the first isomorphism invariant $J_1(\tau')$ of all the $\tau' \in K_0 \otimes \mathbb{H}$ that are μ-isogeneous to τ, up to isomorphism. If all the $J_1(\tau')$ are distinct then it also follows from Proposition 2.3.4 that the unique root of $(\Psi_{\mu,i}(J_1(\tau'), Z_i))(\tau)$ is $J_i(\tau')$. If they are not distinct then $(\Delta\Phi_\mu)(\tau) = 0$, so as $[\tau] \notin S$, we have that $\Delta G_\mu(J_1(\tau), \ldots, J_d(\tau), Y) = 0$. Hence, for every

$$[\tau], [\tau'] \in (U \cap V) - S \cup \{x \in (U \cap V) : \Delta G_\mu(J_1(x), \ldots, J_d(x), Y) = 0\},$$

there exists a μ-isogeny $\tau \to \tau'$ if and only if $(\Phi_\mu(J_1(\tau'))(\tau) = 0$ and for $i = 2, \ldots, d$, we have that $(\Psi_{\mu,i}(J_1(\tau'), J_i(\tau'))(\tau) = 0$. But for every

$$[\tau], [\tau'] \in (U \cap V) - S \cup \{x \in (U \cap V) : \Delta G_\mu(J_1(x), \ldots, J_d(x), Y) = 0\},$$

we have that $(\Phi_\mu(J_1(\tau'))(\tau) = 0$ if and only if

$$G_\mu(J_1(\tau), \ldots, J_d(\tau), J_1(\tau') = 0$$

and, for $i = 2, \ldots, d$, we have that $(\Psi_{\mu,i}(J_1(\tau'), J_i(\tau'))(\tau) = 0$ if and only if

$$H_{\mu,i}(J_1(\tau), \ldots, J_d(\tau), J_1(\tau'), J_i(\tau')) = 0,$$

so the theorem follows. □

2.5 Complexity and simplifications for genus 2

We only implemented an algorithm to compute the set of Hilbert modular polynomials in genus 2, and only for small quadratic fields K_0, due to the fact that we do not know explicit q-expansions for the RM invariants $J_1, \ldots, J_d$ in any other larger genus. Hence, we restrict now to the genus 2 case, and for simplicity, we set $d = 3$.

Lemma 2.5.2 gives one simplification of the formulae for genus 2: in this case K_0 is quadratic, so that $\mathcal{O}_{K_0}$ and $\mathcal{O}_{K_0}^\vee$ are isomorphic as $\mathcal{O}_{K_0}$-modules. This means that we may define the Hilbert modular variety as a compactification of $\mathrm{SL}_2(\mathcal{O}_{K_0})\backslash(K_0 \otimes \mathbb{H})$ instead of $\mathrm{SL}(\mathcal{O}_{K_0} \oplus \mathcal{O}_{K_0}^\vee)\backslash(K_0 \otimes \mathbb{H})$. When we do this, in Lemma 2.3.9, we must replace the matrix group $\Gamma^0(\mu)$ with the matrix group $\Gamma^0(\mu)'$, which we now define.

Definition 2.5.1. For a totally real number field K_0 of degree 2 over $\mathbb{Q}$, with ring of integers $\mathcal{O}_{K_0}$, and a totally positive element $\mu \in K_0$, we define

$$\Gamma^0(\mu)' = \left\{ \begin{pmatrix} a & \mu b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathcal{O}_{K_0}) : a, b, c, d \in \mathcal{O}_{K_0} \right\}.$$

Lemma 2.5.2. For a totally real number field K_0 of degree 2 over $\mathbb{Q}$ with ring of integers $\mathcal{O}_{K_0}$, and a totally positive element $\mu \in \mathcal{O}_{K_0}$ that generates a prime ideal, the set

$$\mathcal{C} = \left\{ \begin{pmatrix} 1 & \omega \\ 0 & 1 \end{pmatrix} : \omega \in \mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0} \right\} \bigcup \left\{ \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \right\}$$

is a choice of coset representatives for the quotient of groups $\Gamma^0(\mu)'\backslash\mathrm{SL}_2(\mathcal{O}_{K_0})$.

Proof. The matrix group $\mathrm{SL}_2(\mathcal{O}_{K_0})$ acts on $\mathbb{P}^1(\mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0})$ by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot (x : y) = (ax + by : cx + dy).$$

Then in particular, the stabilizer of $(0 : 1)$ is given by $\Gamma^0(\mu)'$, and hence by the orbit-stabilizer theorem, there exists a natural bijection from $\mathcal{C}$ to $\Gamma^0(\mu)'\backslash\mathrm{SL}_2(\mathcal{O}_{K_0})$. □

Remark 2.5.3. Using the representation of $\Gamma^0(\mu)'\backslash\mathrm{SL}_2(\mathcal{O}_{K_0})$ given in Lemma 2.5.2, we can write out explicit q-expansions of the coefficients of Φ_μ and $\Psi_{\mu,i}$ via the following. Let f be a modular form for $\mathrm{SL}_2(\mathcal{O}_{K_0})$ of weight k with q-expansion

$$f(\tau) = \sum_{t \in (\mathcal{O}_{K_0}^\vee)^+} \alpha(t) e^{2\pi i \mathrm{tr}(t\tau)},$$

and let $\ell = \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu)$.

1. For $\omega \in \mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0}$ and $M = \begin{pmatrix} 1 & \omega \\ 0 & 1 \end{pmatrix}$, we have that

$$f|_{\underline{\mu^{-1}}M\tau} = \ell^{-k/2} \sum_{t \in (\mathcal{O}_{K_0}^\vee)^+} \zeta_\ell^{\mathrm{tr}(\ell\mu^{-1}t\omega)} \alpha(t) e^{2\pi i \mathrm{tr}(\mu^{-1}t\tau)},$$

   where $(\mathcal{O}_{K_0}^\vee)^+$ denotes the totally positive elements of $\mathcal{O}_{K_0}^\vee$.

2. For $M = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$, we have that

$$f|_{\underline{\mu^{-1}}M\tau} = \ell^{k/2} \sum_{t \in (\mathcal{O}_{K_0}^\vee)^+} \alpha(t) e^{2\pi i \mathrm{tr}(\mu t\tau)},$$

   where $(\mathcal{O}_{K_0}^\vee)^+$ denotes the totally positive elements of $\mathcal{O}_{K_0}^\vee$.

Algorithm 2.4.8 is extremely slow and uses a lot of memory, and so we give here some practical improvements on the computation time and memory usage. First of all, we do not compute the third modular polynomial $H_{\mu,3}(X_1, X_2, X_3, Y, Z_3)$; Algorithm 2.5.4 shows that, given $(\mathrm{A}, \xi, \iota) \in \mathbf{POrd}_{\mathbb{C},K_0}$, we can compute every abelian surface μ-isogenous to it without using $H_{\mu,3}$.

Algorithm 2.5.4.
INPUT: The first 2 Hilbert modular polynomials $G_\mu(X_1, X_2, X_3, Y)$ and $H_{\mu,2}(X_1, X_2, X_3, Y, Z_2)$, as defined in Definition 2.1.5, the RM isomorphism invariants $(j_1, j_2, j_3) \in \mathbb{C}^3$ of some $(\mathrm{A}, \xi, \iota) \in \mathbf{POrd}_{\mathbb{C},K_0}$, as defined in Definition 2.1.1, and the minimal polynomial $m(X) \in \mathbb{Q}(J_1, J_2)[X]$ of J_3, as in Section 2.4.
OUTPUT: The RM isomorphism invariants of each $(\mathrm{A}', \xi', \iota') \in \mathbf{POrd}_{\mathbb{C},K_0}$ that is μ-isogenous to (A, ξ, ι), or failure.

1. Set L to be the list of the $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1$ roots of $G_\mu(j_1, j_2, j_3, Y)$. If the roots are not distinct, output failure.
2. For every $j_1' \in L$:
   (a) set j_2' to be the unique element of $\mathbb{C}$ for which $H_{\mu,2}(j_1, j_2, j_3, j_1', j_2') = 0$,
   (b) set L_0 to be the list of the roots of $m(X)$ evaluated at $(J_1, J_2) = (j_1', j_2')$.
   (c) for every $l \in L_0$, check if $G_\mu(j_1', j_2', l, j_1) = 0$. If true for exactly one l, set $j_3' = l$. Else, output failure.
   (d) add (j_1', j_2', j_3') to list L'.
3. Return L'.

The second major improvement is to do computations in finite fields in place of in $\mathbb{Q}$ and $\mathbb{Q}(\zeta_{\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu)})$ and then use the Chinese Remainder Theorem.

One advantage of working over a finite field in place of $\mathbb{Q}$ is that while the algorithm is running over $\mathbb{Q}$, the coefficients of the q-expansions blow up, using up memory space and slowing down computations, so that Algorithm 2.5.5 is significantly faster than Algorithm 2.4.8.

Algorithm 2.5.5.
INPUT:

1. A totally real number field K_0 of degree 2 over $\mathbb{Q}$.
2. The q-expansions of generators $\gamma_1, \ldots, \gamma_s \in \mathcal{M}_{K_0}(\mathbb{Z})$ of the $\mathbb{Q}$-algebra $\mathcal{M}_{K_0}(\mathbb{Q})$.
3. The images of $\gamma_1, \ldots, \gamma_s$ under ρ as rational functions of J_1, J_2, J_3, where ρ is as defined in (2.15).
4. A totally positive element $\mu \in K_0$ that generates a prime ideal.
5. An upper bound B on the absolute values and a common denominator D of the rational coefficients of the coefficients of $\Phi_\mu(Y)$ and $\Psi_{\mu,2}(Y, Z_2)$ when represented as formal polynomials $\gamma_1, \ldots, \gamma_s$.

6. A prime p_0 such that for every prime $p \geq p_0$, the q-expansion coefficients in Step 1 of Algorithm 2.4.8 have denominator coprime to p, and when replacing $\mathbb{Q}$ and $\mathbb{Q}(\zeta_\ell)$ by $\mathbb{F}_p$ and $\mathbb{F}_p(\zeta_\ell)$, the system of linear equations in Step 2 of Algorithm 2.4.8 still has a unique solution.

OUTPUT: The first 2 polynomials

$$G_\mu(X_1, X_2, X_3, Y) \in \mathbb{Z}[X_1, X_2, X_3, Y], \text{ and}$$
$$H_{\mu,2}(X_1, X_2, X_3, Y, Z_2) \in \mathbb{Z}[X_1, X_2, X_3, Y, Z_2]$$

of Definition 2.1.5.

1. Create a list L of primes in the following way:
   (a) Set $i = 0$.
   (b) Set $b = p_i$.
   (c) Set $p_{i+1} = \min\{n \in \mathbb{Z}_{>b} : n \text{ prime}, n \equiv 1 \bmod \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu)\}$. (This condition is to speed up the computations as the $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu)^{\text{th}}$ roots of unity are then in $\mathbb{F}_p$.)
   (d) Reduce the coefficients of the q-expansions of $\gamma_1, \ldots, \gamma_s$ mod p_{i+1} to get
   $$\overline{\gamma_1}, \ldots, \overline{\gamma_s} \in \mathcal{M}_{K_0}(\mathbb{Z})/p_{i+1}\mathcal{M}_{K_0}(\mathbb{Z}).$$
   If $\overline{\gamma_1}, \ldots, \overline{\gamma_s}$ generate $\mathcal{M}_{K_0}(\mathbb{Z})/p_{i+1}\mathcal{M}_{K_0}(\mathbb{Z})$ as a $\mathbb{F}_{p_{i+1}}$-algebra, go to step (e). Else, set $b = p_{i+1}$ and go to step (c).
   (e) If $\prod_{j=1}^{i+1} p_j < 2BD$ then set $i = i + 1$ and go to (b). Else return
   $$L = \{p_1, \ldots, p_{i+1}\}.$$
2. Write the coefficients mod p of $\Phi_\mu(Y)$ and $\Psi_{\mu,2}(Y)$ as formal polynomials in $\gamma_1, \ldots, \gamma_s$ for every $p \in L$ by following Step 1 and 2 of Algorithm 2.4.8, with $\mathbb{Q}$ (and $\mathbb{Q}(\zeta_{\mathrm{Norm}_{K_0/\mathbb{Q}}})$) replaced by $\mathbb{F}_p$. (This can be done in parallel.)
3. Use the Chinese Remainder Theorem to compute the coefficients of $D\Phi_\mu(Y)$ and $D\Psi_{\mu,2}(Y)$ as formal polynomials in $\gamma_1, \ldots, \gamma_s$ with integer coefficients.
4. Compute G_μ and $H_{\mu,2}$ following Step 3 of Algorithm 2.4.8.

Remark 2.5.6. Heuristically, we expect that for large primes p and most (A, ξ) and $(\mathrm{A}', \xi') \in \mathbf{POrd}_{\mathbb{F}_p, K_0}$, there exists a μ-isogeny $(\mathrm{A}, \xi) \to (\mathrm{A}', \xi')$ if and only if

$$G_\mu(J_1(\mathrm{A}), J_2(\mathrm{A}), J_3(\mathrm{A}), J_1(\mathrm{A}')) \equiv H_{\mu,1}(J_1(\mathrm{A}), J_2(\mathrm{A}), J_3(\mathrm{A}), J_1(\mathrm{A}'), J_2(\mathrm{A}')) \equiv 0 \bmod p$$

and $J_3(\mathrm{A}')$ is the same as the output of Step 2 of Algorithm 2.5.4 (with $\mathbb{C}$ replaced by $\mathbb{F}_p$) with

$$(j_1, j_2, j_3, j_1', j_2') = (J_1(\mathrm{A}), J_2(\mathrm{A}), J_3(\mathrm{A}), J_1(\mathrm{A}'), J_2(\mathrm{A}')).$$

The disadvantage of Algorithm 2.5.5 is that we have to guess the input values B, D, and p_0. However, the speed up is quite significant: for $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) = 11$, Algorithm 2.4.8 took 1 week and Algorithm 2.5.5 took 90 minutes (on the same machine). Also, we can heuristically check the output by looking at the behaviour of the polynomials, for example by attempting to run Algorithm 2.5.4. Even with these improvements, there is still a long way to go before this algorithm is practical for larger values of $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu)$; Table 2.1 gives the timings for the computations that we have done so far.

$\mathrm{Disc}(K_0)$	8	5	5	5	5	5
$\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu)$	2	4	5	9	11	19
Time	2 secs	63 secs	90 secs	$\sim$ 4 mins	$\sim$ 90 mins	$\sim$ 3 days

Table 2.1: Timings for computation of Hilbert modular polynomials G_μ and $H_{\mu,2}$

Chapter 3

The structure of μ-isogeny graphs

The main theorem of this chapter, the Volcano Theorem, Theorem 3.1.9, gives the complete structure of the graph of μ-isogenies of principally polarised dimension g abelian varieties defined over a finite field with real multiplication by a given maximal order. We defined μ-isogenies in Definition 1.3.2; the definition is recalled below. This is a generalisation of David Kohel's structure theorem for $g = 1$ in [Koh96], and Ionica and Thomé's work on genus 2 curves with maximal real multiplication by a given maximal order with narrow class number 1 in [IT14]. In parallel to the work in this thesis, Brooks, Jetchev, and Wesolowski recently obtained some overlapping results, proven using different methods, in [BJW17].

3.1 The Volcano Theorem

Let q be a prime power, let π be a Weil q-number, and let K be a CM-field of degree $2g$ over $\mathbb{Q}$ such that $K = \mathbb{Q}(\pi)$. Recall from Definition 1.3.2 that $\mathbf{POrd}_{\pi,K_0}$ denotes the category of principally polarised ordinary abelian varieties (A, ξ) over $\mathbb{F}_q$ such that the characteristic polynomial of the q-power Frobenius equals the minimal polynomial of π, together with an embedding $\mathcal{O}_{K_0} \hookrightarrow \mathrm{End}(\mathrm{A})$ that extends the embedding

$$\begin{array}{ccc} \mathbb{Z}[\pi + \overline{\pi}] & \longrightarrow & \mathrm{End}(\mathrm{A}) \\ \pi + \overline{\pi} & \mapsto & \mathrm{Frob}(\mathrm{A}) + \mathrm{Ver}(\mathrm{A}), \end{array}$$

where Frob and Ver denote the q-power Frobenius and Verschiebung morphisms respectively. Note that the only restrictions we are making for a principally polarised abelian variety to be in $\mathbf{POrd}_{\pi,K_0}$ are that there exists an embedding $\mathcal{O}_{K_0} \hookrightarrow \mathrm{End}(\mathrm{A})$ and that A is ordinary and geometrically simple. Indeed, given ordinary $\mathrm{A}/\mathbb{F}_q$ simple over $\overline{\mathbb{F}_q}$, let χ be the characteristic polynomial of Frobenius and define $K = \mathbb{Q}(\pi) = \mathbb{Q}[x]/(\chi(x))$. If A is simple and ordinary, then $\mathrm{End}(\mathrm{A}_{\overline{\mathbb{F}_q}}) \otimes \mathbb{Q} = K$. Recall also that the only morphisms in $\mathbf{POrd}_{\pi,K_0}$ are isomorphisms.

Fix a totally positive element $\mu \in \mathcal{O}_{K_0}$ such that $\mu\mathcal{O}_{K_0}$ is a prime ideal. Recall from Definition 1.3.3 that for $(\mathrm{A}, \xi), (\mathrm{A}', \xi') \in \mathbf{POrd}_{\pi,K_0}$ with the map $\iota : \mathcal{O}_{K_0} \hookrightarrow \mathrm{End}(\mathrm{A})$ induced by $\pi \mapsto \mathrm{Frob}_q$, we define a μ-isogeny

$$f : (\mathrm{A}, \xi) \longrightarrow (\mathrm{A}', \xi')$$

to be a morphism $\mathrm{A} \to \mathrm{A}'$ of abelian varieties such that the diagram

$$\begin{tikzcd} \mathrm{A} \arrow[dr, "\xi"'] & \mathrm{A} \arrow[l, "\iota(\mu)"'] \arrow[r, "f"] & \mathrm{A}' \arrow[d, "\xi'"] \\ & \mathrm{A}^\vee & (\mathrm{A}')^\vee \arrow[l, "f^\vee"] \end{tikzcd}$$

commutes. Recall that we denote $\iota(\mu)$ also by μ.

For principally polarised abelian varieties (A, ξ) and (A', ξ') and a morphism $f : \mathrm{A} \to \mathrm{A}'$ we write $f^\dagger = \xi^{-1} f^\vee \xi' : \mathrm{A}' \to \mathrm{A}$. Note that f is a μ-isogeny if and only if $f^\dagger f = \iota(\mu)$. Note that if $(\mathrm{A}, \xi) = (\mathrm{A}', \xi')$, then $(\cdot)^\dagger$ is the Rosati involution. We will also call $f^\dagger$ the dual of f.

Definition 3.1.1. Assume that the only roots of unity in $\mathcal{O}_K$ are ± 1. The *μ-isogeny graph for the Weil q-number π* is the weighted undirected graph for which:

1. The vertices are the isomorphism classes of objects in $\mathbf{POrd}_{\pi,K_0}$,
2. There is an edge between vertices x and x' if and only if there exists a μ-isogeny from x to x'.
3. If a μ-isogeny $f : (\mathrm{A}, \xi) \to (\mathrm{A}, \xi)$ satisfies $f^\dagger = \pm f$, then the edge corresponding to this isogeny has weight $\frac{1}{2}$. All other edges have weight 1.

Remark 3.1.2. In fact, given that the only roots of unity in $\mathcal{O}_K$ are ± 1, if there exists a μ-isogeny f between (A, ξ) and (A', ξ'), the edge between $x = [(\mathrm{A}, \xi)]$ and $x' = [(\mathrm{A}', \xi')]$ represents the μ-isogenies

$$\{f, -f, f^\dagger, -f^\dagger\}.$$

That is, the weight of an edge between x and x' in a μ-isogeny graph denotes the total number of isomorphic μ-isogenies between (A, ξ) and (A', ξ') divided by 4.

Definition 3.1.3. We define the graphs I, R_1, R_2, and for $n \in \mathbb{Z}_{\geq 1}$, the graph C_n in the following way:

- The graph I is a single vertex with no edges.

- The graph R_1 is a single vertex with one edge of weight $\frac{1}{2}$.

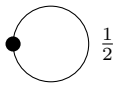

- The graph R_2 is a pair of vertices joined by a single edge of weight 1.

- For $n \in \mathbb{Z}_{\geq 1}$, the graph C_n is a cycle of length n where every edge has weight 1.

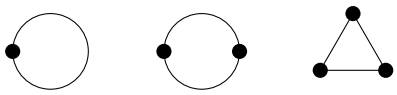

C_1 C_2 C_3

Definition 3.1.4. For $v \in \mathbb{Z}_{\geq 1}$, d, $n \in \mathbb{Z}_{\geq 0}$, and $\Gamma \in \{I, R_1, R_2, C_n\}$, a (Γ, v, d)-*volcano* is a weighted undirected connected graph whose vertices are partitioned into levels $V_0, \ldots, V_d$ such that the following hold:

1. The subgraph on level V_0 is Γ.
2. For all $i > 0$, each vertex in V_i has exactly one neighbour in level V_{i-1}.
3. There exists no edge between vertices in V_i for $i > 0$.
4. For all non-negative integers $i < d$, each vertex in V_i has degree v, where the degree is the weighted sum of the edges counted with intersection multiplicity.

Example 3.1.5. Here is a $(C_5, 3, 1)$-volcano and an $(R_1, 3, 2)$-volcano.

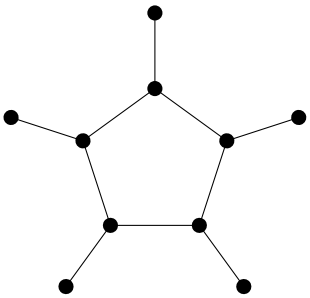

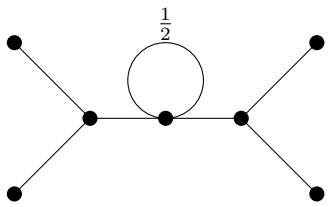

Definition 3.1.6. Let G be a (Γ, v, d)-volcano, and let E be an edge in G between w and w', where $w \in V_i$ and $w' \in V_j$. If $i > j$ we say that E *ascends* from w to w' and *descends* from w' to w. If $i = j$ we say that E is *horizontal*.

Definition 3.1.7. Let $\mathcal{O}$ be an order in $\mathcal{O}_K$ which contains $\mathcal{O}_{K_0}$. The *Shimura class group* of $\mathcal{O}$ is defined to be

$$\mathrm{SCl}(\mathcal{O}) = \frac{\{(\mathfrak{c}, \lambda) : \mathfrak{c} \text{ an invertible fractional } \mathcal{O}\text{-ideal}, \lambda \in K_0^+, \mathfrak{c}\bar{\mathfrak{c}} = \lambda\mathcal{O}\}}{\{(v\mathcal{O}, v\overline{v})) : v \in K^\times\}},$$

where K_0^+ denotes the subgroup of totally positive elements of $K_0^\times$.

Definition 3.1.8. Let μ be a totally positive element of K_0 that generates a prime ideal $\mu\mathcal{O}_{K_0}$, and let $\mathcal{O}$ be an order in K containing $\mathcal{O}_{K_0}[\pi, \overline{\pi}]$ that is locally maximal at μ. If $\mu\mathcal{O}$ factors in $\mathcal{O}$ as $\mu\mathcal{O} = \mathfrak{m}\overline{\mathfrak{m}}$, define $n(\mu, \mathcal{O})$ to be the order of the class of $(\mathfrak{m}, \mu)$ in $\mathrm{SCl}(\mathcal{O})$. For such $\mathcal{O}$ and μ, we define the graph $\Gamma_{\mathcal{O},\mu}$ by

$$\Gamma_{\mathcal{O},\mu} = \begin{cases} I & \text{if } \mu\mathcal{O}_{K_0} \text{ is inert in } K/K_0, \\ C_{n(\mathcal{O},\mu)} & \text{if } \mu\mathcal{O}_{K_0} \text{ is split in } K/K_0, \\ R_{n(\mathcal{O},\mu)} & \text{if } \mu\mathcal{O}_{K_0} \text{ is ramified in } K/K_0. \end{cases}$$

The purpose of this section will be to prove the Volcano Theorem, below, which is our analogue to the results for elliptic curves first given by David Kohel in [Koh96].

Theorem 3.1.9 (Volcano Theorem)**.** *Let K be a CM-field of degree $2g$, generated over $\mathbb{Q}$ by an ordinary Weil q-number π, and with maximal totally real subfield K_0. Suppose further that the only roots of unity in $\mathcal{O}_K$ are $\{\pm 1\}$. Let μ be a totally positive element of $\mathcal{O}_{K_0}$ such that $\mu\mathcal{O}_{K_0}$ is a prime ideal. Define*

$$v = \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1$$

and

$$d = \max\{k \in \mathbb{Z} : \mathcal{O}_{K_0}[\pi, \overline{\pi}] \subseteq \mathcal{O}_{K_0} + \mu^k\mathcal{O}_K\}.$$

For every connected component C of the μ-isogeny graph for the Weil q-number π, there exists an order $\mathcal{O}$ in K containing $\mathcal{O}_{K_0}[\pi, \overline{\pi}]$ that is locally maximal at μ such that C is a $(\Gamma_{\mathcal{O},\mu}, v, d)$-volcano.

Our first goal will be to understand how μ-isogenous abelian varieties can differ. As isogenies preserve the endomorphism algebra, looking at the endomorphism rings of μ-isogenous abelian varieties is a natural place to begin - in fact the endomorphism ring of any abelian variety in our μ-isogeny graph is an order in $\mathcal{O}_K$ that contains $\mathcal{O}_{K_0}[\pi, \overline{\pi}]$, by assumption. Furthermore, we will see in Proposition 3.3.1 that every vertex v in V_i satisfies $\mathrm{End}(v) = \mathcal{O}_{K_0} + \mu^i\mathcal{O}$. The following proposition, which we will prove in Section 3.2, gives a classification of the orders appearing as endomorphism rings of principally polarised ordinary abelian varieties with real multiplication by $\mathcal{O}_{K_0}$:

Proposition 3.2.1. *There is a bijection of sets*

$$\begin{array}{ccc} \left\{ \begin{array}{c} \textit{Orders } \mathcal{O} \textit{ in } \mathcal{O}_K \\ \textit{s.t. } \mathcal{O}_{K_0} \subseteq \mathcal{O} \end{array} \right\} & \leftrightarrow & \{\textit{Ideals of } \mathcal{O}_{K_0}\} \\ \mathcal{O} & \mapsto & (\mathcal{O} : \mathcal{O}_K) \cap \mathcal{O}_{K_0} \\ \mathcal{O}_{K_0} + \mathfrak{f}\mathcal{O}_K & \leftarrow\!\shortmid & \mathfrak{f}. \end{array}$$

Before giving the proof of the Volcano theorem we give some useful definitions regarding conductors and μ-isogenies.

Definition 3.1.10. Let $\mathcal{O}$ be an order in K containing $\mathcal{O}_{K_0}$. We define the *conductor* of $\mathcal{O}$ to be $(\mathcal{O} : \mathcal{O}_K)$, and we define the *real conductor* of $\mathcal{O}$ to be

$$\mathfrak{f}_{\mathcal{O}} = (\mathcal{O} : \mathcal{O}_K) \cap \mathcal{O}_{K_0} = \{x \in \mathcal{O}_{K_0} : x\mathcal{O}_K \subseteq \mathcal{O}\}.$$

For a prime element $\mu \in \mathcal{O}_{K_0}$, we define the *real conductor of $\mathcal{O}$ locally at μ* to be $\mu^k\mathcal{O}_{K_0}$, where

$$k = \mathrm{ord}_{\mu\mathcal{O}_{K_0}}(\mathfrak{f}_{\mathcal{O}}) := \max_{n \in \mathbb{Z}}\{\mathfrak{f}_{\mathcal{O}} \subseteq \mu^n\mathcal{O}_{K_0}\}.$$

We define the *non-μ-part of the real conductor* to be $\mu^{-k}\mathfrak{f}_{\mathcal{O}}$, which is an $\mathcal{O}_{K_0}$-ideal coprime to $\mu\mathcal{O}_{K_0}$.

Definition 3.1.11. Suppose that we have a μ-isogeny ϕ between objects of $\mathbf{POrd}_{\mathbb{C},K_0}$ given by

$$\phi : (\mathrm{A}, \xi) \longrightarrow (\mathrm{A}', \xi').$$

Write $\mathcal{O} = \mathrm{End}(\mathrm{A})$ and $\mathcal{O}' = \mathrm{End}(\mathrm{A}')$. If

(a) $\mu\mathfrak{f}_{\mathcal{O}} = \mathfrak{f}_{\mathcal{O}'}$, then we say that ϕ is *ascending*,

(b) $\mu\mathfrak{f}_{\mathcal{O}'} = \mathfrak{f}_{\mathcal{O}}$, then we say that ϕ is *descending*, and

(c) $\mathfrak{f}_{\mathcal{O}} = \mathfrak{f}_{\mathcal{O}'}$, then we say that ϕ is *horizontal*.

Proposition 3.3.1. *All μ-isogenies from $(\mathrm{A}, \xi, \iota) \in \mathbf{POrd}_{\mathbb{C},K_0}$ such that $\mathrm{End}(\mathrm{A}) \otimes \mathbb{Q} = K$ are ascending, descending, or horizontal.*

We will prove Proposition 3.3.1 in Section 3.3.

Remark 3.1.12. Recall from Theorem 1.5.5 that there is a faithful functor

$$\mathbf{Id}_{\pi,K_0} \longrightarrow \mathbf{Ord}_{\mathbb{C},K_0}$$

that preserves the notions of dual, polarisation, and μ-isogeny. Recall also from Theorem 1.3.11 that there is an equivalence of categories

$$\mathbf{Ord}_{\pi} \longrightarrow \mathbf{Id}_{\pi}$$

that preserves the notions of dual, polarisation, and the action of π and $\overline{\pi}$, hence this induces an equivalence of categories

$$\mathbf{Ord}_{\pi,K_0} \longrightarrow \mathbf{Id}_{\pi,K_0}.$$

In particular, Proposition 3.3.1 implies that all μ-isogenies between elements of $\mathbf{POrd}_{\pi,K_0}$ or $\mathbf{PId}_{\pi,K_0}$ are ascending, descending or horizontal.

In particular, in a μ-isogeny graph G, the non-μ-part of the real conductor is the same for all vertices of any given connected component C of G.

Definition 3.1.13. Given a connected component C of a μ-isogeny graph, we choose a vertex $(\mathrm{A}, \xi) \in C$, and we define the *real conductor of C* to be the non-μ-part of the real conductor of $\mathrm{End}(\mathrm{A})$. We denote this by $\mathfrak{f}_C$.

Below is the proof of Theorem 3.1.9, the Volcano Theorem. The strategy of the proof is as follows: let C be a connected component of the μ-isogeny graph for Weil q-number π, let $\mathfrak{f}_C$ be the conductor of C, and let $\mathcal{O}_C$ be the order of real conductor $\mathfrak{f}_C$ as in Proposition 3.2.1. We first prove that the full subgraph $C(\mathfrak{f}_C)$ of C that contains the vertices with endomorphism ring $\mathcal{O}_C$ is of the form $\Gamma_{\mathcal{O}_C,\mu}$. We then show that from every vertex v with $\mathrm{End}(v) = \mathcal{O}_{K_0} + \mu^i\mathcal{O}_C$ for $i < d$ there are exactly $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1$ edges, and that every edge from a vertex in $C - C(\mathfrak{f}_C)$ is either ascending or descending. Finally, we show that there is a unique ascending edge.

Proof of Theorem 3.1.9. For a connected component C of a μ-isogeny graph G, let G_C be the union of the connected components of G with real conductor $\mathfrak{f}_C$. (Note that $C \subseteq G_C \subseteq G$.) We first partition G_C by endomorphism ring and look at the action of the Shimura class group on these subsets. To this end, write $\mu^d\mathcal{O}_{K_0}$ for the real conductor of $\mathcal{O}_{K_0}[\pi, \overline{\pi}]$ locally at μ (this is equivalent to the formula given for d in the statement of the Volcano Theorem). By Proposition 3.3.1, we can partition the set of vertices of G_C as

$$\bigsqcup_{i=0}^{d} V(\mu^i\mathfrak{f}_C),$$

where for any ideal I in $\mathcal{O}_{K_0}$, we define

$$V(I) = \{(\mathrm{A}, \xi) \in \mathbf{POrd}_{\pi,K_0} : (\mathrm{End}(\mathrm{A}) : \mathcal{O}_K) \cap \mathcal{O}_{K_0} = I\}_{/\cong}. \tag{3.1}$$

To look at the action of the Shimura class group on these subsets, recall that by Theorem 1.3.11 there is a dual, polarisation, and action-of-$\mathcal{O}_{K_0}[\pi, \overline{\pi}]$ preserving equivalence of categories $\mathbf{Id}_{\pi,K_0} \leftrightarrow \mathbf{Ord}_{\pi,K_0}$, where $\mathbf{(P)Id}_{\pi,K_0}$ was defined to be the category of (principally polarised) fractional $\mathcal{O}_{K_0}[\pi, \overline{\pi}]$-ideals. In particular $V(I)$ can also be viewed as

$$\{(\mathfrak{a}, \beta) \in \mathbf{PId}_{\pi,K_0} : (\mathrm{End}(\mathfrak{a}) : \mathcal{O}_K) \cap \mathcal{O}_{K_0} = I\}_{/\cong}. \tag{3.2}$$

The following proposition gives us the action of the Shimura class group on $V(\mu^i\mathfrak{f}_C)$:

Proposition 3.5.1. *For an order $\mathcal{O}$ in $\mathcal{O}_K$ containing $\mathcal{O}_{K_0}[\pi,\overline{\pi}]$ of real conductor $\mathfrak{f}_{\mathcal{O}}$, if the set*

$$V(\mathfrak{f}_{\mathcal{O}}) = \{(\mathfrak{a},\beta) \in \mathbf{PId}_{\pi,K_0} : (\mathrm{End}(\mathfrak{a}) : \mathcal{O}_K) \cap \mathcal{O}_{K_0} = \mathfrak{f}_{\mathcal{O}}\}_{/\cong}$$

is non-empty, then the Shimura class group $\mathrm{SCl}(\mathcal{O})$ of $\mathcal{O}$ acts freely and transitively on $V(\mathfrak{f}_{\mathcal{O}})$ via

$$\begin{array}{ccc} \mathrm{SCl}(\mathcal{O}) \times V(\mathfrak{f}_{\mathcal{O}}) & \longrightarrow & V(\mathfrak{f}_{\mathcal{O}}) \\ ([(\mathfrak{c},\lambda)],[(\mathfrak{a},\beta)]) & \mapsto & [(\mathfrak{c}^{-1}\mathfrak{a},\lambda\beta)]. \end{array}$$

We will prove this in Section 3.5. We will now compute the structure of the full subgraph $G(\mathfrak{f}_C)$ of G_C with vertices $V(\mathfrak{f}_C)$. If $V(\mathfrak{f}_C)$ is non-empty, Proposition 3.5.1 tells us how many vertices are in $V(\mathfrak{f}_C)$, but to deduce the structure of $G(\mathfrak{f}_C)$ we have to also study the (necessarily horizontal) μ-isogenies between the vertices $V(\mathfrak{f}_C)$ of $G(\mathfrak{f}_C)$. Two μ-isogenies $f_1 : (\mathrm{A},\xi) \longrightarrow (\mathrm{A}_1',\xi_1')$ and $f_2 : (\mathrm{A},\xi) \longrightarrow (\mathrm{A}_2',\xi_2')$ are defined to be *isomorphic* if there exists a 1-isogeny $\phi' : (\mathrm{A}_1',\xi_1') \longrightarrow (\mathrm{A}_2',\xi_2')$ such that the diagram

$$\begin{tikzcd} (\mathrm{A},\xi) \arrow[r, "f_1"] \arrow[dr, "f_2"'] & (\mathrm{A}_1',\xi_1') \arrow[d, "\phi'"] \\ & (\mathrm{A}_2',\xi_2') \end{tikzcd}$$

commutes. We deduce the structure of $G(\mathfrak{f}_C)$ from the following proposition:

Proposition 3.6.1. *Given $(\mathrm{A},\xi) \in \mathbf{POrd}_{\pi,K_0}$, let $\mathcal{O} = \mathrm{End}(\mathrm{A})$, let $\mathfrak{f}_{\mathcal{O}} = (\mathcal{O} : \mathcal{O}_K) \cap \mathcal{O}_{K_0}$ be the real conductor of $\mathcal{O}$, and let $\mu \in \mathcal{O}_{K_0}$ be a totally positive prime element $\mu \in \mathcal{O}_{K_0}$. Suppose that there exists $(\mathrm{A},\xi) \in \mathbf{POrd}_{\pi,K_0}$ with $\mathrm{End}(\mathrm{A}) = \mathcal{O}$. Then there is a bijection of sets*

$$\left\{ \begin{array}{c} \textit{horizontal } \mu\textit{-isogenies} \\ \textit{from } (\mathrm{A},\xi) \end{array} \right\}_{/\cong} \longleftrightarrow \{\mathfrak{m} \textit{ an } \mathcal{O}\textit{-ideal} : \mathfrak{m}\overline{\mathfrak{m}} = \mu\mathcal{O}\}$$

such that:

1. *The codomain of the μ-isogeny from $[(\mathrm{A},\xi)]$ corresponding to $[(\mathfrak{m},\mu)]$ is given by $[(\mathfrak{m},\mu)] \cdot [(\mathrm{A},\xi)]$, where $\cdot$ is the action of Proposition 3.5.1.*
2. *The dual $f^\dagger$ of the μ-isogeny $f : [(\mathrm{A},\xi)] \to [(\mathfrak{m},\mu)] \cdot [(\mathrm{A},\xi)]$ corresponding to $\mathfrak{m}$ is isomorphic to the μ-isogeny from $[(\mathfrak{m},\mu)] \cdot [(\mathrm{A},\xi)]$ corresponding to $\overline{\mathfrak{m}}$.*

For the proof, see Section 3.6. This immediately gives us the following:

Corollary 3.1.14. For $(\mathrm{A},\xi) \in \mathbf{POrd}_{\pi,K_0}$ with $\mathrm{End}(\mathrm{A}) = \mathcal{O}$, if $\mathfrak{f}_{\mathcal{O}} \not\subseteq \mu\mathcal{O}_{K_0}$, then up to isomorphism, there are exactly m horizontal μ-isogenies from (A,ξ), where

$$m = \begin{cases} 0 & \text{if } \mu\mathcal{O}_{K_0} \text{ is inert in } K/K_0. \\ 1 & \text{if } \mu\mathcal{O}_{K_0} \text{ is ramified in } K/K_0. \\ 2 & \text{if } \mu\mathcal{O}_{K_0} \text{ splits in } K/K_0. \end{cases}$$

If $\mathfrak{f}_{\mathcal{O}} \subseteq \mu\mathcal{O}_{K_0}$, then there are no horizontal μ-isogenies from (A,ξ).

Now if $V(\mathfrak{f}_C)$ is non-empty, Proposition 3.5.1 and Proposition 3.6.1 tell us that

(a) if $\mu\mathcal{O}_{K_0}$ is inert in K/K_0 then there are no edges in $G(\mathfrak{f}_C)$,

(b) if $\mu\mathcal{O}_{K_0}$ is ramified in K/K_0 and the element $[(\mathfrak{m},\mu)] \in \mathrm{SCl}(\mathcal{O}_C)$ is trivial, then $G(\mathfrak{f}_C)$ is the disjoint union of loops of weight $\frac{1}{2}$,

(c) if $\mu\mathcal{O}_{K_0}$ is ramified in K/K_0 and the element $[(\mathfrak{m},\mu)] \in \mathrm{SCl}(\mathcal{O}_C)$ is non-trivial, then $G(\mathfrak{f}_C)$ is the disjoint union of pairs of vertices joined by a single edge, and

(d) if $\mu\mathcal{O}_{K_0}$ splits in K/K_0 as $\mathfrak{m}\overline{\mathfrak{m}}$, then $G(\mathfrak{f}_C)$ is the disjoint union of cycles of length n, where n is the order of $[(\mathfrak{m},\mu)]$ in $\mathrm{SCl}(\mathcal{O}_C)$.

That is, every non-empty connected component of $G(\mathfrak{f}_C)$ has exactly the form $\Gamma_{\mathcal{O}_C,\mu}$. Hence, if $d = 0$, then we are done, so assume now that $d > 0$. (Recall that d is the exponent of the real conductor $\mu^d\mathcal{O}_{K_0}$ of $\mathcal{O}_{K_0}[\pi, \overline{\pi}]$ locally at μ.) We first need to show that every non-empty connected component C' of G_C contains vertices in $V(\mathfrak{f}_C)$, which is an immediate corollary of the following proposition:

Proposition 3.7.1. *For $i \in \mathbb{Z}_{>0}$, from every vertex in $V(\mu^i\mathfrak{f}_C)$ there is an ascending μ-isogeny.*

For the proof, see Section 3.7. Defining

$$v = \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1$$

and

$$v' = \begin{cases} v & \text{if } \Gamma_{\mathcal{O}_C,\mu} = I, \\ v-1 & \text{if } \Gamma_{\mathcal{O}_C,\mu} = R_n, \\ v-2 & \text{if } \Gamma_{\mathcal{O}_C,\mu} = C_n, \end{cases} \tag{3.3}$$

where n is the order of $[(\mathfrak{m}, \mu)]$ in $\mathrm{SCl}(\mathcal{O}_C)$, it now remains to consider the non-maximal vertices. To this end we have the following proposition:

Proposition 3.8.1. *For $0 \leq i < d$, every vertex in $V(\mu^i\mathfrak{f}_C)$ has degree*

$$\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1.$$

We will prove this in Section 3.8. Proposition 3.8.1 together with Corollary 3.1.14 proves that

(i) there are exactly v' descending edges from each vertex in $V(\mathfrak{f}_C)$.

(ii) for $0 < i < d$, there are v edges from every vertex in $V(\mu^i\mathfrak{f}_C)$ and they are all either ascending or descending.

It remains only to show that if $i > 0$ then there is a unique ascending edge from each vertex in $V(\mu^i\mathfrak{f}_C)$, which will be proven in Section 3.9, as part of the following proposition:

Proposition 3.9.1. *Let all notation be as above. If $d > 0$, we have*

$$\#V(\mu\mathfrak{f}_C) = v'\#V(\mathfrak{f}_C)$$

and for $1 \leq i < d$,

$$\#V(\mu^{i+1}\mathfrak{f}_C) = (v-1)\#V(\mu^i\mathfrak{f}_C).$$

Also, for every $0 < i \leq d$, there is a unique ascending edge from every vertex in $V(\mu^i\mathfrak{f}_C)$.

This finishes the proof of the Volcano Theorem. □

The rest of this chapter is dedicated to proving the 'black-box' propositions from the above proof of the Volcano Theorem. For these propositions we will use the Fixed Frobenius Lifting Theorem (Theorem 1.3.11) to work instead in the category $\mathbf{PId}_{\pi,K_0}$ of principally polarised fractional $\mathcal{O}_{K_0}[\pi, \overline{\pi}]$-ideals that was defined in Definition 1.5.1. Recall that the Fixed Frobenius Lifting Theorem gave us an equivalence of categories

$$\mathbf{Ord}_{\pi,K_0} \leftrightarrow \mathbf{Id}_{\pi,K_0}$$

that preserves the notions of duals, of polarisation and the action of π.

Remark 3.1.15. Recall that in Chapter 1 we also defined the category $\mathbf{Ord}_{\mathbb{C},K_0}$ of complex abelian varieties with maximal real multiplication, and that in Theorem 1.5.5 we gave a faithful functor

$$\mathbf{Ord}_{\pi,K_0} \longrightarrow \mathbf{Ord}_{\mathbb{C},K_0}$$

that preserves the notions of dual, polarisation, and the action of $\mathcal{O}_{K_0}$. In particular, for each CM-field K the Volcano theorem can also be applied to complex abelian varieties in the set

$$\{(\mathrm{A}, \iota) : \mathrm{A} \in \mathbf{Ord}_{\mathbb{C}},\ \iota : K \tilde{\to} \mathrm{End}(\mathrm{A}) \otimes \mathbb{Q},\ \iota(\mathcal{O}_{K_0}[\pi, \overline{\pi}]) \subseteq \mathrm{End}(\mathrm{A})\}.$$

3.2 Parametrising orders by their real conductors

Let K be a CM-field with maximal totally real subfield K_0, and write $\mathcal{O}_K$ and $\mathcal{O}_{K_0}$ for the rings of integers of K and K_0 respectively. In this section we prove that orders $\mathcal{O}$ in K containing $\mathcal{O}_{K_0}$ are determined completely by their real conductors $(\mathcal{O} : \mathcal{O}_K) \cap \mathcal{O}_{K_0}$ (c.f. Definition 3.1.10).

Proposition 3.2.1. There is a bijection of sets

$$\begin{array}{ccc} \left\{ \begin{array}{c} \text{Orders } \mathcal{O} \text{ in } \mathcal{O}_K \\ \text{s.t. } \mathcal{O}_{K_0} \subseteq \mathcal{O} \end{array} \right\} & \leftrightarrow & \{\text{Ideals of } \mathcal{O}_{K_0}\} \\ \mathcal{O} & \mapsto & (\mathcal{O} : \mathcal{O}_K) \cap \mathcal{O}_{K_0} \\ \mathcal{O}_{K_0} + \mathfrak{f}\mathcal{O}_K & \leftarrow\!\shortmid & \mathfrak{f}. \end{array}$$

Proof. Define $\mathfrak{f}(\mathcal{O}) = (\mathcal{O} : \mathcal{O}_K) \cap \mathcal{O}_{K_0}$ and $\mathcal{O}_{\mathfrak{f}} = \mathcal{O}_{K_0} + \mathfrak{f}\mathcal{O}_K$. It suffices to show that for every order $\mathcal{O}$ in $\mathcal{O}_K$ containing $\mathcal{O}_{K_0}$, we have

$$\mathcal{O} = \mathcal{O}_{\mathfrak{f}(\mathcal{O})} \tag{3.4}$$

and that for every $\mathcal{O}_{K_0}$-ideal $\mathfrak{f}$, we have

$$\mathfrak{f} = \mathfrak{f}(\mathcal{O}_{\mathfrak{f}}). \tag{3.5}$$

We start by proving (3.4). As

$$((\mathcal{O} : \mathcal{O}_K) \cap \mathcal{O}_{K_0})\mathcal{O}_K \subseteq (\mathcal{O} : \mathcal{O}_K)\mathcal{O}_K \subseteq \mathcal{O},$$

it is clear that

$$(\mathcal{O}_{\mathfrak{f}(\mathcal{O})} =)\mathcal{O}_{K_0} + ((\mathcal{O} : \mathcal{O}_K) \cap \mathcal{O}_{K_0})\mathcal{O}_K \subseteq \mathcal{O}.$$

To prove equality, first note that K/K_0 is a 1-dimensional K_0-vector space. Choosing a basis gives a K_0-linear isomorphism

$$q : K/K_0 \xrightarrow{\sim} K_0.$$

Then

$$\mathcal{O}/\mathcal{O}_{\mathfrak{f}(\mathcal{O})} \cong (\mathcal{O}/\mathcal{O}_{K_0})/(\mathcal{O}_{\mathfrak{f}(\mathcal{O})}/\mathcal{O}_{K_0}) \cong q(\mathcal{O})/q(\mathcal{O}_{\mathfrak{f}(\mathcal{O})}),$$

hence to prove (3.4), it suffices to prove that

$$q(\mathcal{O}) = q(\mathcal{O}_{\mathfrak{f}(\mathcal{O})}). \tag{3.6}$$

To this end, we claim that

(a) for all orders $\mathcal{O}$ of K containing $\mathcal{O}_{K_0}$, we have that $q(\mathcal{O}) = \mathfrak{f}(\mathcal{O})q(\mathcal{O}_K)$.

(b) for all ideals $\mathfrak{f}$ of $\mathcal{O}_{K_0}$, we have that $q(\mathcal{O}_{\mathfrak{f}}) = \mathfrak{f}q(\mathcal{O}_K)$.

Observe that if both (a) and (b) hold, then

$$q(\mathcal{O}) = \mathfrak{f}(\mathcal{O})q(\mathcal{O}_K) = q(\mathcal{O}_{\mathfrak{f}(\mathcal{O})}),$$

so that (3.6) holds and hence so does (3.4). We first prove (a): note that $q(\mathcal{O})$ and $q(\mathcal{O}_K)$ are non-zero finitely generated $\mathcal{O}_{K_0}$-submodules of K_0, and $\mathcal{O}_{K_0}$ is a Dedekind domain, hence $q(\mathcal{O})$ and $q(\mathcal{O}_K)$ are non-zero invertible fractional ideals of $\mathcal{O}_{K_0}$, so that (a) holds if and only if

$$(q(\mathcal{O}) : q(\mathcal{O}_K)) = \mathfrak{f}(\mathcal{O}).$$

Now $\alpha \in K$ satisfies $\alpha \in (q(\mathcal{O}) : q(\mathcal{O}_K))$ if and only if $\alpha \in K_0$ and for all $x \in \mathcal{O}_K$ we have that

$$q(\alpha x) \in q(\mathcal{O}). \tag{3.7}$$

We claim that (3.7) holds if and only if for all $x \in \mathcal{O}_K$ we have that

$$\alpha x \in \mathcal{O} + \mathcal{O}_{K_0} = \mathcal{O}.$$

The 'if' statement is clear, so assume that for all $x \in \mathcal{O}_K$ we have that $q(\alpha x) \in q(\mathcal{O})$. Then for each $\alpha x \in \mathcal{O} + K_0$ there exists $y \in \mathcal{O}$ and $z \in K_0$ such that $\alpha x = y + z$. Also $\alpha \in (q(\mathcal{O}) : q(\mathcal{O}_K))$ and

$q(\mathcal{O}) \subseteq q(\mathcal{O}_K)$, so $\alpha \in (q(\mathcal{O}_K) : q(\mathcal{O}_K)) = \mathcal{O}_{K_0}$, hence $z = \alpha x - y$ is an algebraic integer, so is in $\mathcal{O}_{K_0}$. This proves the 'only if'. Therefore

$$(q(\mathcal{O}) : q(\mathcal{O}_K)) = \mathcal{O}_{K_0} \cap (\mathcal{O} : \mathcal{O}_K) = \mathfrak{f}(\mathcal{O}).$$

So (a) holds. We now prove (b):

$$\begin{aligned} q(\mathcal{O}_\mathfrak{f}) &= \{q(a+b) : a \in \mathcal{O}_{K_0},\, b \in \mathfrak{f}\mathcal{O}_K\} \\ &= \{q(a) + q(b) : a \in \mathcal{O}_{K_0},\, b \in \mathfrak{f}\mathcal{O}_K\} \\ &= \{q(b) : b \in \mathfrak{f}\mathcal{O}_K\} \\ &= q(\mathfrak{f}\mathcal{O}_K) \\ &= \mathfrak{f}q(\mathcal{O}_K). \end{aligned}$$

So (b), and hence (3.4), holds. It remains to prove (3.5), which is now almost automatic:

$$\mathfrak{f}(\mathcal{O}_\mathfrak{f})q(\mathcal{O}_K) \overset{(a)}{=} q(\mathcal{O}_\mathfrak{f}) \overset{(b)}{=} \mathfrak{f}q(\mathcal{O}_K)$$

and $q(\mathcal{O}_K)$ is an invertible fractional $\mathcal{O}_{K_0}$-ideal, hence

$$\mathfrak{f}(\mathcal{O}_\mathfrak{f}) = \mathfrak{f}.$$

□

This proposition has an easy corollary:

Corollary 3.2.2. Every order $\mathcal{O}$ in K containing $\mathcal{O}_{K_0}$ is stable under complex conjugation.

Proof. By Proposition 3.2.1, we have that

$$\begin{aligned} \mathcal{O} &= \mathcal{O}_{K_0} + ((\mathcal{O} : \mathcal{O}_K) \cap \mathcal{O}_{K_0})\mathcal{O}_K \\ &= \mathcal{O}_{K_0} + (\overline{(\mathcal{O} : \mathcal{O}_K) \cap \mathcal{O}_{K_0}})\mathcal{O}_K \\ &= \mathcal{O}_{K_0} + ((\overline{\mathcal{O}} : \mathcal{O}_K) \cap \mathcal{O}_{K_0})\mathcal{O}_K \\ &= \overline{\mathcal{O}}. \end{aligned}$$

□

Remark 3.2.3. In fact, it is necessary that the endomorphism ring of a principally polarised abelian variety is stable under complex conjugation, as the Rosati involution is just complex conjugation.

3.3 All μ-isogenies are ascending, descending or horizontal

Recall that for an order $\mathcal{O}$ in $\mathcal{O}_K$ containing $\mathcal{O}_{K_0}$, the *real conductor* was defined in Definition 3.1.10 to be

$$\mathfrak{f}_\mathcal{O} = (\mathcal{O} : \mathcal{O}_K) \cap \mathcal{O}_{K_0},$$

and for a totally positive prime element μ of $\mathcal{O}_{K_0}$. Suppose that we have a μ-isogeny ϕ between objects of $\mathbf{POrd}_{\mathbb{C},K_0}$ given by

$$\phi : (\mathrm{A}, \xi, \iota) \longrightarrow (\mathrm{A}', \xi', \iota'),$$

and that $\mathrm{End}(\mathrm{A}) \otimes \mathbb{Q}$ is a CM-field. Recall from Definition 3.1.11 that, writing $\mathcal{O} = \mathrm{End}(\mathrm{A})$ and $\mathcal{O}' = \mathrm{End}(\mathrm{A}')$, if

(a) $\mu\mathfrak{f}_{\mathcal{O}'} = \mathfrak{f}_\mathcal{O}$, then we say that ϕ is *ascending*,

(b) $\mu\mathfrak{f}_\mathcal{O} = \mathfrak{f}_{\mathcal{O}'}$, then we say that ϕ is *descending*, and

(c) $\mathfrak{f}_\mathcal{O} = \mathfrak{f}_{\mathcal{O}'}$, then we say that ϕ is *horizontal*.

In this section we prove the following:

Proposition 3.3.1. All μ-isogenies from $(\mathrm{A}, \xi, \iota) \in \mathbf{POrd}_{\mathbb{C},K_0}$ such that $\mathrm{End}(\mathrm{A}) \otimes \mathbb{Q} = K$ are ascending, descending, or horizontal.

Remark 3.3.2. Observe that by Theorem 1.5.5, it follows from Proposition 3.3.1 that all μ-isogenies between objects of $\mathbf{PId}_{\pi,K_0}$ are ascending, descending, or horizontal.

We first prove a useful lemma:

Lemma 3.3.3. The dual of a μ-isogeny between objects of $\mathbf{POrd}_{\mathbb{C},K_0}$ with complex multiplication or objects of $\mathbf{POrd}_{\pi,K_0}$ is also a μ-isogeny.

Proof. We prove this for objects of $\mathbf{POrd}_{\mathbb{C},K_0}$; it then follows for objects of $\mathbf{POrd}_{\pi,K_0}$ by Theorem 1.5.5 and Theorem 1.3.11. Let (A, ξ, ι) and $(\mathrm{A}', \xi', \iota') \in \mathbf{POrd}_{\mathbb{C},K_0}$ and suppose that

$$f : (\mathrm{A}, \xi, \iota) \longrightarrow (\mathrm{A}', \xi', \iota')$$

is a μ-isogeny. For $\iota : \mathcal{O}_{K_0} \hookrightarrow \mathrm{End}(\mathrm{A})$, define

$$\begin{array}{rccc} \iota^\vee : & \mathcal{O}_{K_0} & \longrightarrow & \mathrm{End}(\mathrm{A}^\vee) \\ & \alpha & \mapsto & \iota(\alpha)^\vee. \end{array}$$

As the multiplication-by-μ map commutes with isogenies that preserve the real multiplication, both

$$\begin{tikzcd} (\mathrm{A}')^\vee \arrow[dr, "(\xi')^{-1}"'] & (\mathrm{A}')^\vee \arrow[l, "\mu"'] \arrow[r, "f^\vee"] & \mathrm{A}^\vee \arrow[d, "\xi^{-1}"] \\ & \mathrm{A}' & \mathrm{A}, \arrow[l, "f"] \end{tikzcd}$$

and

$$\begin{tikzcd} \mathrm{End}((\mathrm{A}')^\vee) \otimes \mathbb{Q} \arrow[rr, "g \mapsto f^\vee g (f^\vee)^{-1}"] & & \mathrm{End}(\mathrm{A}^\vee) \otimes \mathbb{Q} \\ K_0 \arrow[u, "\iota'^\vee"] \arrow[urr, "\iota^\vee"] & & \end{tikzcd}$$

commute, so

$$f^\vee : ((\mathrm{A}')^\vee, (\xi')^{-1}, \iota'^\vee) \longrightarrow (\mathrm{A}^\vee, \xi^{-1}, \iota^\vee)$$

is a μ-isogeny. $\square$

Proof of Proposition 3.3.1. By assumption, the endomorphism ring $\mathcal{O}$ of A is an order in K. Let $f : (\mathrm{A}, \xi, \iota) \to (\mathrm{A}', \xi', \iota')$ be a μ-isogeny. We identify $\mathcal{O}' = \mathrm{End}(\mathrm{A}')$ with a subring of K via

$$\begin{array}{rccc} f^* : & \mathcal{O}' & \longrightarrow & \mathrm{End}(\mathrm{A}) \otimes \mathbb{Q} = K \\ & \alpha & \mapsto & f^{-1}\alpha f, \end{array}$$

where f^{-1} is the inverse of $f \in \mathrm{Hom}(\mathrm{A}, \mathrm{A}') \otimes \mathbb{Q}$. It suffices to show that if (A, ξ, ι) and $(\mathrm{A}', \xi', \iota')$ in $\mathbf{POrd}_{\mathbb{C},K_0}$ are μ-isogenous, with $\mathrm{End}(\mathrm{A}) = \mathcal{O}$ and $\mathrm{End}(\mathrm{A}') = \mathcal{O}'$, then

$$\mathfrak{f}_{\mathcal{O}} = \mu \mathfrak{f}_{\mathcal{O}'},\ \mathfrak{f}_{\mathcal{O}'} = \mu \mathfrak{f}_{\mathcal{O}},\ \text{ or } \mathfrak{f}_{\mathcal{O}} = \mathfrak{f}_{\mathcal{O}'},$$

where $\mathfrak{f}_{\mathcal{O}}$ and $\mathfrak{f}_{\mathcal{O}'}$ are the real conductors of $\mathcal{O}$ and $\mathcal{O}'$ respectively. So let

$$f : (\mathrm{A}, \xi, \iota) \longrightarrow (\mathrm{A}', \xi', \iota')$$

be a μ-isogeny; then the diagram

$$\begin{tikzcd} \mathrm{A} & \mathrm{A} \arrow[l, "\mu"'] \arrow[r, "f"] & \mathrm{A}' \arrow[d, "\xi'"] \\ & \mathrm{A}^\vee \arrow[ul, "\xi^{-1}"] & (\mathrm{A}')^\vee \arrow[l, "f^\vee"] \end{tikzcd} \tag{3.8}$$

commutes. In particular, we have a surjective morphism

$$f^\dagger := (\xi^{-1} \circ f^\vee \circ \xi') : \mathrm{A}' \longrightarrow \mathrm{A}$$

such that

$$f^\dagger \circ f = \xi^{-1} \circ f^\vee \circ \xi' \circ f = [\mu]_{\mathrm{A}}.$$

Then for every $\alpha \in \mathcal{O}'$, take $\varphi \in \mathrm{End}(\mathrm{A}')$ such that $\iota'(\varphi) = \alpha$, so that $f^\dagger \circ \varphi \circ f \in \mathrm{End}(\mathrm{A})$. Then $\mu\alpha = \iota(f^\dagger \circ \varphi \circ f) \in \mathcal{O}$, hence

$$\mu\mathcal{O}' \subseteq \mathcal{O}.$$

By Lemma 3.3.3, the dual of a μ-isogeny is also a μ-isogeny, so $\mathcal{O} \subseteq \mu^{-1}\mathcal{O}'$. In particular, this implies that

$$\mu\mathfrak{f}_{\mathcal{O}'} \subseteq \mathfrak{f}_{\mathcal{O}} \subseteq \mu^{-1}\mathfrak{f}_{\mathcal{O}'}.$$

□

3.4 Principally polarised ideals are invertible

We will use repeatedly for the rest of the chapter the following proposition:

Proposition 3.4.1. If $(\mathfrak{a}, \beta) \in \mathbf{PId}_{\pi,K_0}$, then $\mathfrak{a}$ is an invertible $\mathrm{End}(\mathfrak{a})$-ideal.

Before proving Proposition 3.4.1 we first prove a useful formula for $\mathrm{End}(\mathfrak{a})$. Recall from Lemma 1.7.1 that for $(\mathfrak{a}, \beta) \in \mathbf{PId}_{\pi,K_0}$ there exists $\tau \in K$ such that

$$(\mathfrak{a}, \beta) \cong (\tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee, (\tau - \overline{\tau})^{-1}),$$

where $\mathcal{O}_{K_0}^\vee$ is the trace dual of $\mathcal{O}_{K_0}$.

Lemma 3.4.2. Let $\mathfrak{a}$ be a fractional $\mathcal{O}_{K_0}[\pi, \overline{\pi}]$-ideal such that $\mathfrak{a} = \tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee$, where $\tau \in K$, and choose $A, B, C \in K_0$ not all zero such that

$$A\tau^2 + B\tau + C = 0. \tag{3.9}$$

Define the fractional $\mathcal{O}_{K_0}$-ideal $\mathfrak{d}$ by

$$\mathfrak{d} = A\mathcal{O}_{K_0}^\vee + B\mathcal{O}_{K_0} + C(\mathcal{O}_{K_0}^\vee)^{-1}. \tag{3.10}$$

Then

$$\mathrm{End}(\mathfrak{a}) = A\tau\mathfrak{d}^{-1} + \mathcal{O}_{K_0}.$$

Proof. For every $x \in \mathrm{End}(\mathfrak{a})$, as $\mathrm{End}(\mathfrak{a}) \subseteq K$, we know that there exist $a, b \in K_0$ such that

$$x = a\tau + b.$$

Then for every $a, b \in K_0$, we have that $a\tau + b \in \mathrm{End}(\mathfrak{a})$ if and only if

$$(a\tau + b)(\tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee) \subseteq \tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee.$$

That is, if and only if for every $\alpha \in \mathcal{O}_{K_0}$ and every $\beta \in \mathcal{O}_{K_0}^\vee$, we have that

$$\tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee \ni (a\tau + b)\tau\alpha = \tau^2\alpha a + \tau\alpha b = -A^{-1}(B\tau + C)\alpha a + \tau\alpha b \tag{3.11}$$

and

$$\tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee \ni (a\tau + b)\beta = \beta\tau a + \beta b. \tag{3.12}$$

Now, we have (3.11) for every $\alpha \in \mathcal{O}_{K_0}$ if and only if

$$\left(b - \frac{B}{A}a\right)\mathcal{O}_{K_0} \subseteq \mathcal{O}_{K_0} \tag{3.13}$$

and

$$a\frac{C}{A}\mathcal{O}_{K_0} \subseteq \mathcal{O}_{K_0}^\vee. \tag{3.14}$$

Similarly, we have (3.12) for every $\beta \in \mathcal{O}_{K_0}^\vee$ if and only if

$$a\mathcal{O}_{K_0}^\vee \subseteq \mathcal{O}_{K_0} \tag{3.15}$$

and

$$b \in \mathcal{O}_{K_0}. \tag{3.16}$$

Note that (3.16) is equivalent to $b\mathcal{O}_{K_0}^\vee \subseteq \mathcal{O}_{K_0}^\vee$ because $\mathcal{O}_{K_0}^\vee$ is an invertible $\mathcal{O}_{K_0}$-ideal (as $\mathcal{O}_{K_0}$ is a Dedekind domain). Now $a\tau + b \in \mathrm{End}(\mathfrak{a})$ if and only if (3.13)-(3.16) hold. Furthermore, (3.13) and (3.16) hold if and only if $b \in \mathcal{O}_{K_0}$ and

$$\frac{B}{A}a\mathcal{O}_{K_0} \subseteq \mathcal{O}_{K_0}, \tag{3.17}$$

and (3.14) holds if and only if

$$a\frac{C}{A}(\mathcal{O}_{K_0}^\vee)^{-1} \subseteq \mathcal{O}_{K_0}. \tag{3.18}$$

We now have that $a\tau + b \in \mathrm{End}(\mathfrak{a})$ if and only if (3.15), (3.16), (3.17), and (3.18) hold. But (3.15), (3.17), and (3.18) hold if and only if

$$a \in \left(\mathcal{O}_{K_0} : \mathcal{O}_{K_0}^\vee + \frac{B}{A}\mathcal{O}_{K_0} + \frac{C}{A}(\mathcal{O}_{K_0}^\vee)^{-1}\right) = A\mathfrak{d}^{-1}.$$

Hence

$$\mathrm{End}(\mathfrak{a}) = A\tau\mathfrak{d}^{-1} + \mathcal{O}_{K_0}.$$

□

Proof of Proposition 3.4.1. By Lemma 1.7.1, there exists $\alpha \in K^\times$ such that

$$\alpha\mathfrak{a} = \tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee. \tag{3.19}$$

Without loss of generality, set $\alpha = 1$. Let A, B, C, and $\mathfrak{d}$ be as in Lemma 3.4.2. We claim that

$$\mathfrak{a}\left(A\mathfrak{d}^{-1}\overline{\mathfrak{a}}(\mathcal{O}_{K_0}^\vee)^{-1}\right) = \mathrm{End}(\mathfrak{a}).$$

Note that as $\mathcal{O}_{K_0}$ is a Dedekind domain, all fractional $\mathcal{O}_{K_0}$-ideals are invertible. In particular, both $\mathfrak{d}$ and $\mathcal{O}_{K_0}^\vee$ are invertible $\mathcal{O}_{K_0}$-ideals. Note also that

$$\mathrm{tr}_{K/K_0}(\tau) = -B/A \quad \text{and} \quad \mathrm{N}_{K/K_0}(\tau) = -C/A. \tag{3.20}$$

Now

$$\begin{aligned}
&\mathfrak{a}(A\overline{\mathfrak{a}}\mathfrak{d}^{-1}(\mathcal{O}_{K_0}^\vee)^{-1}) \\
&= (\tau\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee)(\overline{\tau}\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee)A\mathfrak{d}^{-1}(\mathcal{O}_{K_0}^\vee)^{-1} && \text{by (3.19)} \\
&= (\tau\overline{\tau}(\mathcal{O}_{K_0}^\vee)^{-1} + \tau\mathcal{O}_{K_0} + \overline{\tau}\mathcal{O}_{K_0} + \mathcal{O}_{K_0}^\vee)A\mathfrak{d}^{-1} \\
&= (C(\mathcal{O}_{K_0}^\vee)^{-1} + B\mathcal{O}_{K_0} + A\tau\mathcal{O}_{K_0} + A\mathcal{O}_{K_0}^\vee)\mathfrak{d}^{-1} && \text{by (3.20)} \\
&= A\tau\mathfrak{d}^{-1} + \mathcal{O}_{K_0} \\
&= \mathrm{End}(\mathfrak{a}). && \text{by Lemma 3.4.2}
\end{aligned}$$

Hence $\mathfrak{a}$ is an invertible $\mathrm{End}(\mathfrak{a})$-ideal, with

$$\mathfrak{a}^{-1} = A\overline{\mathfrak{a}}\mathfrak{d}^{-1}(\mathcal{O}_{K_0}^\vee)^{-1}.$$

□

One nice corollary of Proposition 3.4.1 is the following formula:

Corollary 3.4.3. Let $(\mathfrak{a}, \beta) \in \mathbf{PId}_{\pi,K_0}$ with $\mathrm{End}(\mathfrak{a}) = \mathcal{O}$. Then

$$\beta\mathfrak{a}\overline{\mathfrak{a}} = \mathcal{O}^\vee. \tag{3.21}$$

Proof. Note first that by Definition 1.3.7, we have

$$\begin{aligned}\mathfrak{a}^\vee &= \{x \in K : \mathrm{tr}_{K/\mathbb{Q}}(\overline{x}\mathfrak{a}) \subseteq \mathbb{Z}\} \\ &= \{x \in K : \mathrm{tr}_{K/\mathbb{Q}}(x\overline{\mathfrak{a}}) \subseteq \mathbb{Z}\}.\end{aligned}$$

Also

$$\mathrm{tr}_{K/\mathbb{Q}}((\overline{\mathfrak{a}}^{-1}\mathcal{O}^\vee)\overline{\mathfrak{a}}) = \mathrm{tr}_{K/\mathbb{Q}}(\mathcal{O}^\vee) \subseteq \mathbb{Z},$$

hence

$$\overline{\mathfrak{a}}^{-1}\mathcal{O}^\vee \subseteq \mathfrak{a}^\vee = \beta\mathfrak{a},$$

hence as $\overline{\mathfrak{a}}$ is an invertible $\mathcal{O}$-ideal, we have that

$$\mathcal{O}^\vee \subseteq \beta\mathfrak{a}\overline{\mathfrak{a}}.$$

For the other inclusion, observe that by Definition 1.3.7

$$\mathcal{O}^\vee = \{x \in K : \mathrm{tr}_{K/\mathbb{Q}}(x\overline{\mathcal{O}}) \subseteq \mathbb{Z}\},$$

and that

$$\mathrm{tr}_{K/\mathbb{Q}}(\overline{\mathfrak{a}}\mathfrak{a}^\vee\overline{\mathcal{O}}) = \mathrm{tr}_{K/\mathbb{Q}}(\overline{\mathfrak{a}}\mathfrak{a}^\vee) \subseteq \mathbb{Z},$$

so that in particular $\overline{\mathfrak{a}}\mathfrak{a}^\vee \subseteq \mathcal{O}^\vee$. Hence

$$\beta\mathfrak{a}\overline{\mathfrak{a}} = \overline{\mathfrak{a}}\mathfrak{a}^\vee \subseteq \mathcal{O}^\vee.$$

□

3.5 The action of the Shimura class group

Let π be an ordinary Weil q-number that generates a CM-field $K = \mathbb{Q}(\pi)$ with maximal totally real subfield K_0. Let μ be a totally positive prime element of $\mathcal{O}_{K_0}$ and let G be the μ-isogeny graph for π. Let C be a connected component of G and let $\mathfrak{f}_C$ be the real conductor of C, as defined in Definition 3.1.13. Recall from Equation (3.1) that we defined

$$V(\mu^i\mathfrak{f}_C) = \{(\mathrm{A}, \xi) \in \mathbf{POrd}_{\pi,K_0} : (\mathrm{End}(\mathrm{A}) : \mathcal{O}_K) \cap \mathcal{O}_{K_0} = \mu^i\mathfrak{f}_C\}_{/\cong},$$

and that under the equivalence of categories of Theorem 1.3.11, we may also define

$$V(\mu^i\mathfrak{f}_C) = \{(\mathfrak{a}, \beta) \in \mathbf{PId}_{\pi,K_0} : (\mathrm{End}(\mathfrak{a}) : \mathcal{O}_K) \cap \mathcal{O}_{K_0} = \mu^i\mathfrak{f}_C\}_{/\cong}.$$

Proposition 3.5.1. For an order $\mathcal{O}$ in $\mathcal{O}_K$ containing $\mathcal{O}_{K_0}[\pi, \overline{\pi}]$ of real conductor $\mathfrak{f}_\mathcal{O}$, if the set

$$V(\mathfrak{f}_\mathcal{O}) = \{(\mathfrak{a}, \beta) \in \mathbf{PId}_{\pi,K_0} : (\mathrm{End}(\mathfrak{a}) : \mathcal{O}_K) \cap \mathcal{O}_{K_0} = \mathfrak{f}_\mathcal{O}\}_{/\cong}$$

is non-empty, then the Shimura class group $\mathrm{SCl}(\mathcal{O})$ of $\mathcal{O}$ acts freely and transitively on $V(\mathfrak{f}_\mathcal{O})$ via

$$\begin{array}{ccc}\mathrm{SCl}(\mathcal{O}) \times V(\mathfrak{f}_\mathcal{O}) & \longrightarrow & V(\mathfrak{f}_\mathcal{O}) \\ ([(\mathfrak{c}, \lambda)], [(\mathfrak{a}, \beta)]) & \mapsto & [(\mathfrak{c}^{-1}\mathfrak{a}, \lambda\beta)].\end{array}$$

Proof. Define

$$\mathrm{Frac}(\mathcal{O}) = \{(\mathfrak{c}, \lambda) : \mathfrak{c} \text{ a fractional } \mathcal{O}\text{-ideal}, \lambda \in K_0^+, \mathfrak{c}\overline{\mathfrak{c}} = \lambda\mathcal{O}\}$$

and

$$\mathrm{Prin}(\mathcal{O}) = \{(v\mathcal{O}, v\overline{v}) : v \in K^\times\}$$

so that $\mathrm{SCl}(\mathcal{O}) = \mathrm{Frac}(\mathcal{O})/\mathrm{Prin}(\mathcal{O})$. Now if $(\mathfrak{c}, \lambda) \in \mathrm{Frac}(\mathcal{O})$ and $[(\mathfrak{a}, \beta)] \in \mathrm{SCl}(\mathcal{O})$ then $\mathfrak{c}^{-1}\mathfrak{a}$ is an invertible $\mathcal{O}$-ideal and is in $\mathbf{Id}_\pi$. Also

$$\begin{aligned}(\mathfrak{c}^{-1}\mathfrak{a})^\vee &= (\overline{\mathfrak{c}^{-1}\mathfrak{a}})^{-1}\mathcal{O}^\vee && \text{by Corollary 3.4.3} \\ &= \overline{\mathfrak{c}}\mathfrak{a}^\vee && \text{by Corollary 3.4.3} \\ &= \lambda\beta\mathfrak{c}^{-1}\mathfrak{a} && \text{as } \mathfrak{c}\overline{\mathfrak{c}} = \lambda\mathcal{O} \text{ and } \beta\mathfrak{a} = \mathfrak{a}^\vee,\end{aligned}$$

and λ is totally positive, hence $\lambda\beta$ is a principal polarisation of $\mathfrak{c}^{-1}\mathfrak{a}$ by Remark 1.3.8. Therefore $\mathrm{Frac}(\mathcal{O})$ acts on $V(\mathfrak{f}_{\mathcal{O}})$ via

$$(\mathfrak{c}, \lambda) * [(\mathfrak{a}, \lambda)] = [(\mathfrak{c}^{-1}\mathfrak{a}, \lambda\beta)].$$

We now show that this in fact defines a free action

$$\mathrm{SCl}(\mathcal{O}) \times V(\mathfrak{f}_{\mathcal{O}}) \longrightarrow V(\mathfrak{f}_{\mathcal{O}}).$$

Take $(\mathfrak{c}, \lambda) \in \mathrm{Frac}(\mathcal{O})$. Then for $[(\mathfrak{a}, \beta)] \in V(\mathfrak{f}_{\mathcal{O}})$, we have that

$$[(\mathfrak{a}, \beta)] = [(\mathfrak{c}^{-1}\mathfrak{a}, \lambda\beta)]$$

if and only if there exists an isomorphism

$$(\mathfrak{a}, \beta) \longrightarrow (\mathfrak{c}^{-1}\mathfrak{a}, \lambda\beta)$$

in $\mathbf{PId}_{\pi,K_0}$. Recall from Definition 1.3.9 and Definition 1.5.1 that an isomorphism in $\mathbf{PId}_{\pi,K_0}$ is $\alpha \in K$ such that

$$\alpha\mathfrak{a} = \mathfrak{c}^{-1}\mathfrak{a} \quad \text{and} \quad \alpha\lambda\beta\overline{\alpha} = \beta.$$

This is equivalent to $\lambda = (\alpha\overline{\alpha})^{-1}$ and $\mathfrak{c} = \alpha^{-1}\mathcal{O}$. In particular, we have that $[(\mathfrak{a}, \beta)] = [(\mathfrak{c}^{-1}\mathfrak{a}, \lambda\beta)]$ if and only if

$$[(\mathfrak{c}, \lambda)] = [(\alpha^{-1}\mathcal{O}, (\alpha\overline{\alpha})^{-1})] = [(\mathcal{O}, 1)].$$

Hence the action is free and well-defined, so it remains to show that it is transitive. That is, it remains to show that if $[(\mathfrak{a}, \beta)]$ and $[(\mathfrak{a}', \beta')] \in V(\mathfrak{f}_{\mathcal{O}})$, then

$$(\mathfrak{a}(\mathfrak{a}')^{-1}, \beta^{-1}\beta') \in \mathrm{Frac}(\mathcal{O}). \tag{3.22}$$

First, note that as β and β' are polarisations, for every $\phi \in \Phi_{\pi,j}$, we have that

$$\phi(\beta)/i,\ \phi(\beta')/i \in \mathbb{R}_{>0}.$$

(Recall the definition of $\Phi_{\pi,j}$ from Definition 1.3.5.) In particular, for every $\phi \in \Phi_{\pi,j}$, we have that

$$\phi(\beta'\beta^{-1}) = \phi(\beta')\phi(\beta)^{-1} \in \mathbb{R}_{>0},$$

so $\beta'\beta^{-1}$ is totally positive. Finally, we have by Corollary 3.4.3 that $\mathfrak{a}\overline{\mathfrak{a}}\beta = \mathcal{O}^{\vee}$ and $\mathfrak{a}'\overline{\mathfrak{a}'}\beta' = \mathcal{O}^{\vee}$, so

$$\mathfrak{a}(\mathfrak{a}')^{-1}\overline{\mathfrak{a}(\mathfrak{a}')^{-1}} = \beta'\beta^{-1}\mathcal{O},$$

hence Equation (3.22) holds and the action is transitive. □

3.6 Counting horizontal μ-isogenies

The goal of this section is to prove Proposition 3.6.1, which was used in the proof of the Volcano Theorem, Theorem 3.1.9.

Proposition 3.6.1. Given $(\mathrm{A}, \xi) \in \mathbf{POrd}_{\pi,K_0}$, let $\mathcal{O} = \mathrm{End}(\mathrm{A})$, let $\mathfrak{f}_{\mathcal{O}} = (\mathcal{O} : \mathcal{O}_K) \cap \mathcal{O}_{K_0}$ be the real conductor of $\mathcal{O}$, and let $\mu \in \mathcal{O}_{K_0}$ be a totally positive prime element $\mu \in \mathcal{O}_{K_0}$. Suppose that there exists $(\mathrm{A}, \xi) \in \mathbf{POrd}_{\pi,K_0}$ with $\mathrm{End}(\mathrm{A}) = \mathcal{O}$. Then there is a bijection of sets

$$\left\{ \begin{array}{c} \text{horizontal } \mu\text{-isogenies} \\ \text{from } (\mathrm{A}, \xi) \end{array} \right\}_{/\cong} \longleftrightarrow \{\mathfrak{m} \text{ an } \mathcal{O}\text{-ideal} : \mathfrak{m}\overline{\mathfrak{m}} = \mu\mathcal{O}\}$$

such that:

1. The codomain of the μ-isogeny from $[(\mathrm{A}, \xi)]$ corresponding to $[(\mathfrak{m}, \mu)]$ is given by $[(\mathfrak{m}, \mu)] \cdot [(\mathrm{A}, \xi)]$, where $\cdot$ is the action of Proposition 3.5.1.

2. The dual $f^{\dagger}$ of the μ-isogeny $f : [(\mathrm{A}, \xi)] \to [(\mathfrak{m}, \mu)] \cdot [(\mathrm{A}, \xi)]$ corresponding to $\mathfrak{m}$ is isomorphic to the μ-isogeny from $[(\mathfrak{m}, \mu)] \cdot [(\mathrm{A}, \xi)]$ corresponding to $\overline{\mathfrak{m}}$.

We first prove two lemmas.

Lemma 3.6.2. Suppose that $(\mathfrak{a}, \beta)$ and $(\mathfrak{a}', \beta') \in \mathbf{PId}_{\pi,K_0}$ and that $\mathrm{End}(\mathfrak{a}) = \mathrm{End}(\mathfrak{a}') = \mathcal{O}$. If $\alpha \in (\mathfrak{a}' : \mathfrak{a})$ is a μ-isogeny

$$(\mathfrak{a}, \beta) \longrightarrow (\mathfrak{a}', \beta')$$

then

$$\alpha\mathfrak{a}(\mathfrak{a}')^{-1}\overline{\alpha\mathfrak{a}(\mathfrak{a}')^{-1}} = \mu\mathcal{O}.$$

Proof. By Corollary 3.4.3 we know that

$$\beta\mathfrak{a}\overline{\mathfrak{a}} = \beta'\mathfrak{a}'\overline{\mathfrak{a}'} = \mathcal{O}^\vee.$$

Now given $\alpha \in (\mathfrak{a}' : \mathfrak{a})$, by definition it is a μ-isogeny if and only if

$$\mu\beta = \alpha\overline{\alpha}\beta',$$

which implies that

$$\alpha\overline{\alpha}\beta\mathfrak{a}\overline{\mathfrak{a}} = \mu\beta\mathfrak{a}'\overline{\mathfrak{a}'}.$$

Also, by Proposition 3.4.1, we know that $\mathfrak{a}'$ and $\overline{\mathfrak{a}'}$ are invertible as $\mathcal{O}$-ideals, hence

$$\alpha\mathfrak{a}(\mathfrak{a}')^{-1}\overline{\alpha\mathfrak{a}(\mathfrak{a}')^{-1}} = \mu\mathcal{O}.$$

□

Lemma 3.6.3. Let $\mathcal{O}$ be an order in $\mathcal{O}_K$ containing $\mathcal{O}_{K_0}$ of real conductor $\mathfrak{f}_\mathcal{O}$ and suppose that $\mu\mathcal{O}_{K_0}|\mathfrak{f}_\mathcal{O}$. Then there is a unique prime ideal $\mathfrak{m}$ of $\mathcal{O}$ lying above $\mu\mathcal{O}_{K_0}$, given by

$$\mathfrak{m} = \mu\mathcal{O}_{K_0} + \mathfrak{f}_\mathcal{O}\mathcal{O}_K.$$

Proof. Recall that $\mathcal{O} = \mathcal{O}_{K_0} + \mathfrak{f}_\mathcal{O}\mathcal{O}_K$ so $\mathfrak{m}$ is clearly an ideal of $\mathcal{O}$. Furthermore, we have that

$$\mathcal{O}/\mathfrak{m} \cong \mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0}$$

which is a field (as $\mu\mathcal{O}_{K_0}$ is prime and $\mathcal{O}_{K_0}$ is a Dedekind domain), so $\mathfrak{m}$ is maximal. Suppose now that $\widetilde{\mathfrak{m}}$ is a prime ideal of $\mathcal{O}$ lying above $\mu\mathcal{O}_{K_0}$, so that

$$\mu\mathcal{O}_{K_0} + \mu\mathfrak{f}_\mathcal{O}\mathcal{O}_K = \mu\mathcal{O} \subseteq \widetilde{\mathfrak{m}}. \tag{3.23}$$

We can factor $\mu\mathfrak{f}_\mathcal{O}\mathcal{O}_K$ into $\mathcal{O}$-ideals as

$$\mu\mathfrak{f}_\mathcal{O}\mathcal{O}_K = \mathfrak{m}(\mathfrak{f}_\mathcal{O}\mathcal{O}_K),$$

which is contained in $\widetilde{\mathfrak{m}}$ by (3.23). Hence as $\widetilde{\mathfrak{m}}$ is a prime $\mathcal{O}$-ideal, either

$$\mathfrak{m} \subseteq \widetilde{\mathfrak{m}} \tag{3.24}$$

or

$$\mathfrak{f}_\mathcal{O}\mathcal{O}_K \subseteq \widetilde{\mathfrak{m}}, \tag{3.25}$$

and by (3.23), we have also that $\mu\mathcal{O}_{K_0} \subseteq \widetilde{\mathfrak{m}}$, so that (3.25) implies that

$$\mu\mathcal{O}_{K_0} + \mathfrak{f}_\mathcal{O}\mathcal{O}_K \subseteq \widetilde{\mathfrak{m}}. \tag{3.26}$$

Therefore, as $\mathfrak{m} = \mu\mathcal{O}_{K_0} + \mathfrak{f}_\mathcal{O}\mathcal{O}_K$ is a maximal ideal, we have by (3.24) or (3.26) that $\widetilde{\mathfrak{m}} = \mathfrak{m}$, so $\mathfrak{m}$ is unique. □

Proof of Proposition 3.6.1. We prove this in the equivalent category $\mathbf{Id}_{\pi,K_0}$ instead of in $\mathbf{Ord}_{\pi,K_0}$; that is, we count μ-isogenies from $(\mathfrak{a}, \beta) \in \mathbf{PId}_{\pi,K_0}$ such that $\mathrm{End}(\mathfrak{a}) = \mathcal{O}$. Suppose that there exists $(\mathfrak{a}', \beta') \in \mathbf{PId}_{\pi,K_0}$ with $\mathrm{End}(\mathfrak{a}') = \mathcal{O}$ and a μ-isogeny

$$\alpha : (\mathfrak{a}, \beta) \longrightarrow (\mathfrak{a}', \beta').$$

Then by Lemma 3.6.2, we have that

$$\alpha\mathfrak{a}(\mathfrak{a}')^{-1}\overline{\alpha\mathfrak{a}(\mathfrak{a}')^{-1}} = \mu\mathcal{O}. \tag{3.27}$$

Also $\alpha\mathfrak{a}(\mathfrak{a}')^{-1}$ is an $\mathcal{O}$-ideal as $\alpha\mathfrak{a} \subset \mathfrak{a}'$, so (3.27) implies that there is an $\mathcal{O}$-ideal $\mathfrak{m}$ such that $\mu\mathcal{O} = \mathfrak{m}\overline{\mathfrak{m}}$. If $\mu\mathcal{O}_{K_0}|\mathfrak{f}_{\mathcal{O}}$, we have by Lemma 3.6.3 that such an $\mathfrak{m}$ does not exist and hence there are no horizontal μ-isogenies in this case.

Suppose now that $\mathfrak{f}_{\mathcal{O}} = \mathfrak{f}_C$. Then $\mu\mathcal{O}$ is relatively prime to $\mathfrak{f}_{\mathcal{O}}\mathcal{O} = (\mathcal{O} : \mathcal{O}_K)$, and hence decomposes uniquely into prime ideals. As μ is a prime element of $\mathcal{O}_{K_0}$, the ideal $\mu\mathcal{O}_{K_0}$ is either inert, ramified, or split in K/K_0. If $\mu\mathcal{O}_{K_0}$ is inert, then there exists no $\mathcal{O}$-ideal $\mathfrak{m}$ such that $\mu\mathcal{O} = \mathfrak{m}\overline{\mathfrak{m}}$, so as before, there are no horizontal μ-isogenies from $(\mathfrak{a}, \beta)$ in this case.

It remains to consider the case in which $\mu\mathcal{O}_{K_0}$ is split or ramified in K/K_0, so suppose that $\mu\mathcal{O}$ decomposes as $\mu\mathcal{O} = \mathfrak{m}\overline{\mathfrak{m}}$. Then $\mu \in (\mathfrak{m}\mathfrak{a} : \mathfrak{a})$ and $\mathfrak{m}$ corresponds to the μ-isogeny

$$\mu : (\mathfrak{a}, \beta) \longrightarrow (\mathfrak{m}\mathfrak{a}, \mu^{-1}\beta),$$

and $\mu \in (\overline{\mathfrak{m}}\mathfrak{a} : \mathfrak{a})$ and $\overline{\mathfrak{m}}$ corresponds to the μ-isogeny

$$\mu : (\mathfrak{a}, \beta) \longrightarrow (\overline{\mathfrak{m}}\mathfrak{a}, \mu^{-1}\beta).$$

We claim that up to isomorphism these are the only horizontal μ-isogenies from $(\mathfrak{a}, \beta)$. Suppose that there is an object $(\mathfrak{a}', \beta') \in \mathbf{PId}_{\pi,K_0}$ with $\mathrm{End}(\mathfrak{a}') = \mathcal{O}$ for which some $\alpha \in (\mathfrak{a}' : \mathfrak{a})$ defines a μ-isogeny

$$(\mathfrak{a}, \beta) \longrightarrow (\mathfrak{a}', \beta').$$

Then $\beta' = (\alpha\overline{\alpha})^{-1}\mu\beta$, and by Lemma 3.6.2, we have that

$$\alpha\mathfrak{a}(\mathfrak{a}')^{-1}\overline{\alpha\mathfrak{a}(\mathfrak{a}')^{-1}} = \mu\mathcal{O},$$

so by unique factorisation, we have that

$$\alpha\mathfrak{a}(\mathfrak{a}')^{-1} = \mathfrak{m} \quad \text{or} \quad \alpha\mathfrak{a}(\mathfrak{a}')^{-1} = \overline{\mathfrak{m}},$$

that is,

$$\mathfrak{a}' = \mu^{-1}\alpha\overline{\mathfrak{m}}\mathfrak{a} \quad \text{or} \quad \mathfrak{a}' = \mu^{-1}\alpha\mathfrak{m}\mathfrak{a}.$$

It is then easy to see that if $\mathfrak{a}' = \mu^{-1}\alpha\mathfrak{m}\mathfrak{a}$ then $\mu^{-1}\alpha \in (\mathfrak{a}' : \mathfrak{m}\mathfrak{a})$ defines a 1-isogeny (i.e. isomorphism)

$$(\mathfrak{m}\mathfrak{a}, \mu^{-1}\beta) \longrightarrow (\mathfrak{a}', (\alpha\overline{\alpha})^{-1}\mu\beta)$$

corresponding to $\mu^{-1}\alpha\mathcal{O}$, in which case the μ-isogeny defined by $\mu \in (\mathfrak{m}\mathfrak{a} : \mathfrak{a})$ corresponding to $\mathfrak{m}$ and the μ-isogeny defined by $\alpha \in (\mathfrak{a}' : \mathfrak{a})$ corresponding to $\mu^{-1}\alpha\mathcal{O}$ make the diagram

$$\begin{tikzcd}
(\mathfrak{a}, \beta) \arrow[r, "\mu"] \arrow[d, "1"] & (\mathfrak{m}\mathfrak{a}, \mu^{-1}\beta) \arrow[d, "\mu^{-1}\alpha"] \\
(\mathfrak{a}, \beta) \arrow[r, "\alpha"] & (\mathfrak{a}', \beta')
\end{tikzcd}$$

commute and hence are isomorphic by definition. Similarly, if $\mathfrak{a}' = \mu^{-1}\alpha\overline{\mathfrak{m}}\mathfrak{a}$ then $\mu \in (\overline{\mathfrak{m}}\mathfrak{a} : \mathfrak{a})$ and $\alpha \in (\mathfrak{a}' : \mathfrak{a})$ are isomorphic as μ-isogenies.

We now show that $\mathfrak{m} = \overline{\mathfrak{m}}$ if and only if the μ-isogenies defined by $\mu \in (\mathfrak{m}\mathfrak{a} : \mathfrak{a})$ and $\mu \in (\overline{\mathfrak{m}}\mathfrak{a} : \mathfrak{a})$ corresponding to $\mathfrak{m}$ and $\overline{\mathfrak{m}}$ respectively are isomorphic. The 'only if' is clear, so we proceed by proving the 'if'. Suppose that there exists a 1-isogeny defined by $\alpha \in (\overline{\mathfrak{m}}\mathfrak{a} : \mathfrak{m}\mathfrak{a})$ such that the diagram

$$\begin{tikzcd}
(\mathfrak{a}, \beta) \arrow[r, "\mu"] \arrow[dr, "\mu"] & (\mathfrak{m}\mathfrak{a}, \mu^{-1}\beta) \arrow[d, "\alpha"] \\
& (\overline{\mathfrak{m}}\mathfrak{a}, \mu^{-1}\beta)
\end{tikzcd}$$

commutes. Then $\alpha = 1$ so $\mathfrak{m}\mathfrak{a} = \alpha\mathfrak{m}\mathfrak{a} = \overline{\mathfrak{m}}\mathfrak{a}$. Therefore as $\mathfrak{a}$ is an invertible $\mathcal{O}$-ideal, we get that $\overline{\mathfrak{m}} = \mathfrak{m}$.

It remains to show (2), that the dual $f^\dagger$ of the μ-isogeny

$$f = \mu : (\mathfrak{a}, \beta) \to (\mathfrak{m}\mathfrak{a}, \mu^{-1}\beta)$$

corresponding to $\mathfrak{m}$ is the μ-isogeny from

$$[(\mathfrak{m}\mathfrak{a}, \mu^{-1}\beta)]$$

corresponding to $\overline{\mathfrak{m}}$. By the definition of a μ-isogeny, we have that $f^\dagger f = \mu$, hence

$$f^\dagger = 1 : (\mathfrak{m}\mathfrak{a}, \mu^{-1}\beta) \to (\mathfrak{a}, \beta).$$

This is the composition of the μ-isogeny

$$g = \mu : (\mathfrak{m}\mathfrak{a}, \mu^{-1}\beta) \to (\overline{\mathfrak{m}}\mathfrak{m}\mathfrak{a}, \mu^{-2}\beta)$$

corresponding to $\overline{\mathfrak{m}}$ and the 1-isogeny

$$\mu^{-1} : (\mu\mathfrak{a}, \mu^{-2}\beta) \to (\mathfrak{a}, \beta).$$

Hence $f^\dagger$ is isomorphic to g so (2) holds. □

3.7 A construction of ascending μ-isogenies

The goal of this section is to prove Proposition 3.7.1, which was used in the proof of the Volcano Theorem, Theorem 3.1.9. Recall that C is a connected component of the μ-isogeny graph for the Weil q-number π, that $\mathfrak{f}_C$ is the real conductor $\mathfrak{f}_C$ (as defined in Definition 3.1.13), and that

$$V(\mu^i\mathfrak{f}_C) = \{(\mathrm{A}, \xi) \in \mathbf{POrd}_{\pi,K_0} : (\mathrm{End}(\mathrm{A}) : \mathcal{O}_K) \cap \mathcal{O}_{K_0} = \mu^i\mathfrak{f}_C\}_{/\cong}.$$

Proposition 3.7.1. For $i \in \mathbb{Z}_{>0}$, from every vertex in $V(\mu^i\mathfrak{f}_C)$ there is an ascending μ-isogeny.

Proof. We prove this in the category $\mathbf{PId}_{\pi,K_0}$. That is, we prove that for that $(\mathfrak{a}, \beta) \in \mathbf{PId}_{\pi,K_0}$ with $\mathrm{End}(\mathfrak{a}) = \mathcal{O}$, if $\mu\mathcal{O}_{K_0}$ divides $\mathfrak{f}_\mathcal{O}$, the real conductor of $\mathcal{O}$, then there is an ascending μ-isogeny from $(\mathfrak{a}, \beta)$. So suppose that $\mu\mathcal{O}_{K_0}|\mathfrak{f}_\mathcal{O}$ and write $\mathcal{O}' = \mathcal{O}_{K_0} + \mu^{-1}\mathfrak{f}_\mathcal{O}\mathcal{O}_K$ for the order in $\mathcal{O}_K$ of real conductor $\mu^{-1}\mathfrak{f}_\mathcal{O}$. We claim that

$$\mathrm{End}(\mathfrak{a}\mathcal{O}') = \mathcal{O}', \tag{3.28}$$

and that

$$(\mathfrak{a}\mathcal{O}', \mu\beta) \in \mathbf{PId}_{\pi,K_0}. \tag{3.29}$$

Note that (3.29) implies that $1 \in (\mathfrak{a}\mathcal{O}', \mathfrak{a})$ defines a μ-isogeny

$$(\mathfrak{a}, \beta) \longrightarrow (\mathfrak{a}\mathcal{O}', \mu\beta),$$

and (3.28) implies that it is ascending. For (3.28), observe that $\mathfrak{a}\mathcal{O}'$ is an invertible $\mathcal{O}'$-ideal with inverse $\mathfrak{a}^{-1}\mathcal{O}'$, and hence $\mathrm{End}(\mathfrak{a}\mathcal{O}') = \mathcal{O}'$. For (3.29), if we can can show that $(\mathfrak{a}\mathcal{O}')^\vee = \mu\beta\mathfrak{a}\mathcal{O}'$ then by Remark 1.3.8 we have that $\mu\beta$ is a principal polarisation of $\mathfrak{a}'\mathcal{O}'$ as μ is totally positive. We have that

$$\begin{aligned}
(\mathfrak{a}\mathcal{O}')^\vee &= \{x \in K : \mathrm{tr}_{K/\mathbb{Q}}(\overline{x}\mathfrak{a}\mathcal{O}') \subseteq \mathbb{Z}\} \\
&= \{x \in K : x\overline{\mathfrak{a}} \subseteq (\mathcal{O}')^\vee\} \\
&= ((\mathcal{O}')^\vee : \overline{\mathfrak{a}}) \\
&= \overline{\mathfrak{a}}^{-1}(\mathcal{O}')^\vee && \text{as } \overline{\mathfrak{a}} \text{ is invertible} \\
&= \beta\mathfrak{a}(\mathcal{O} : \mathcal{O}^\vee)(\mathcal{O}')^\vee && \text{by Corollary 3.4.3.}
\end{aligned}$$

In the last step we used also that $\mathcal{O}^\vee$ is an invertible $\mathcal{O}$-ideal; note that $\mathcal{O}^\vee = \beta\mathfrak{a}\overline{\mathfrak{a}}$ is the product of invertible $\mathcal{O}$-ideals and so is itself also an invertible $\mathcal{O}$-ideal. Hence it suffices to show that

$$(\mathcal{O} : \mathcal{O}^\vee)(\mathcal{O}')^\vee = \mu\mathcal{O}'. \tag{3.30}$$

Now as $\mu\mathcal{O}' \subseteq \mathcal{O}$, we have that

$$\mu\mathcal{O}'\mathcal{O}^\vee \subseteq \mu\mathcal{O}'(\mu\mathcal{O}')^\vee = \mu\mathcal{O}'\mu^{-1}(\mathcal{O}')^\vee = (\mathcal{O}')^\vee,$$

then multiplying by $(\mathcal{O} : \mathcal{O}^\vee)$ gives $\mu\mathcal{O}' \subseteq (\mathcal{O} : \mathcal{O}^\vee)(\mathcal{O}')^\vee$. Also, as $(\mathcal{O}')^\vee \subseteq \mathcal{O}^\vee$, this gives us that

$$\mu\mathcal{O}' \subseteq (\mathcal{O} : \mathcal{O}^\vee)(\mathcal{O}')^\vee \subseteq \mathcal{O}.$$

By Lemma 3.6.3 we know that $\mu\mathcal{O}'$ is the unique prime ideal of $\mathcal{O}$ lying above $\mu\mathcal{O}_{K_0}$, hence maximal, so either $\mu\mathcal{O}' = (\mathcal{O} : \mathcal{O}^\vee)(\mathcal{O}')^\vee$ or $(\mathcal{O} : \mathcal{O}^\vee)(\mathcal{O}')^\vee = \mathcal{O}$. But $\mathcal{O}$ is not an $\mathcal{O}'$-submodule of K, hence (3.30), and in turn (3.29), hold. □

3.8 Counting the degree of vertices in the μ-isogeny graph

Recall from (3.2) that for an ideal I of $\mathcal{O}_{K_0}$ we defined

$$V(I) = \{(\mathrm{A}, \xi) \in \mathbf{POrd}_{\pi,K_0} : (\mathrm{End}(\mathrm{A}) : \mathcal{O}_K) \cap \mathcal{O}_{K_0} = I\}/\cong$$

to be the set of vertices in the μ-isogeny graph for the Weil q-number π for which the endomorphism rings of the corresponding abelian varieties have real conductor I. In this section we prove the following proposition:

Proposition 3.8.1. For $0 \leq i < d$, every vertex in $V(\mu^i \mathfrak{f}_C)$ has degree

$$\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1.$$

Proof. Recall that in Theorem 1.5.5 we defined an equivalence of categories

$$F_\pi : \quad \mathbf{Id}_{\pi,K_0} \quad \longrightarrow \quad \mathbf{Ord}_{\mathbb{C},\pi,K_0}$$

that preserves the action of $\mathcal{O}_{K_0}[\pi, \overline{\pi}]$ and the notions of dual and polarisation. Let $(\mathfrak{a}, \beta) \in \mathbf{PId}_{\pi,K_0}$ and suppose that $[(\mathfrak{a}, \beta)] \in V(\mu^i \mathfrak{f}_C)$, where $0 \leq i < d$. Write $F_\pi(\mathfrak{a}, \beta) = (\mathrm{A}, \xi, e)$.

From Lemma 2.3.9 we have that there are $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu)+1$ non-isomorphic μ-isogenies from any object in $\mathbf{POrd}_{\mathbb{C},K_0}$. (See Definition 2.1.3 to recall how Lemma 2.3.9 relates to objects in $\mathbf{POrd}_{\mathbb{C},K_0}$). Hence, it suffices to show that every μ-isogeny from $(\mathrm{A}, \xi, e|_{\mathcal{O}_{K_0}})$ in $\mathbf{Ord}_{\mathbb{C},K_0}$ comes from a unique μ-isogeny in $\mathbf{Ord}_{\mathbb{C},\pi}$.

Given a μ-isogeny

$$f : (\mathrm{A}, \xi, e|_{\mathcal{O}_{K_0}}) \to (\mathrm{A}', \xi', \iota')$$

in $\mathbf{Ord}_{\mathbb{C},K_0}$, embed K into $\mathrm{End}(\mathrm{A}') \otimes \mathbb{Q}$ via $e' = f \circ e \circ f^{-1}$. Then

$$\mathrm{End}((\mathrm{A}, e)) = \mathrm{End}(\mathfrak{a}) = \mathcal{O}_{K_0} + \mu^i \mathfrak{f}_C \mathcal{O}_K \subseteq \mathcal{O}_K,$$

therefore by Proposition 3.3.1 we have that $(e')^{-1}(\mathrm{End}(\mathrm{A}')) = \mathcal{O}_{K_0} + \mu^j \mathfrak{f}_C \mathcal{O}_K$ where $j \in \{i-1, i, i+1\}$. In particular, as $i < d$ this implies that

$$f : (\mathrm{A}, \xi, e) \to (\mathrm{A}', \xi', e')$$

is in $\mathbf{Ord}_{\mathbb{C},\pi,K_0}$. Conversely, for every embedding e' such that $f : (\mathrm{A}, \xi, e) \to (\mathrm{A}', \xi', e')$ is a μ-isogeny in $\mathbf{Ord}_{\mathbb{C},\pi,K_0}$, by definition we have that $e' = f \circ e \circ f^{-1}$. □

3.9 The order of the Shimura class group

In this section we prove Proposition 3.9.1, which was used in the proof of the Volcano Theorem. We will use notation as in the proof of the volcano theorem: recall that for a connected component C of the μ-isogeny graph for the Weil q-number π, we defined $\mathfrak{f}_C$ to be the real conductor of C, and we defined $\mathcal{O}_C$ to be the order given by $\mathcal{O}_{K_0} + \mathfrak{f}_C \mathcal{O}_K$. Recall also that we defined $V(I)$ to be the set of vertices with endomorphism ring of real conductor I, we defined $v = \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1$, and we defined

$$v' = \begin{cases} v & \text{if } \Gamma_{\mathcal{O}_C,\mu} = I, \\ v-1 & \text{if } \Gamma_{\mathcal{O}_C,\mu} = R_n \text{ for some } n, \\ v-2 & \text{if } \Gamma_{\mathcal{O}_C,\mu} = C_n \text{ for some } n. \end{cases} \tag{3.31}$$

Proposition 3.9.1. Let all notation be as above. If $d > 0$, we have

$$\#V(\mu \mathfrak{f}_C) = v' \#V(\mathfrak{f}_C)$$

and for $1 \leq i < d$,

$$\#V(\mu^{i+1} \mathfrak{f}_C) = (v-1)\#V(\mu^i \mathfrak{f}_C).$$

Also, for every $0 < i \leq d$, there is a unique ascending edge from every vertex in $V(\mu^i \mathfrak{f}_C)$.

Lemma 3.9.2. Let K be a CM-field with maximal totally real subfield K_0 such that that the only roots of unity in $\mathcal{O}_K$ are ± 1, and let $\mathcal{O}' \subset \mathcal{O}$ be orders in $\mathcal{O}_K$ containing $\mathcal{O}_{K_0}$. Write $\mathfrak{f}_{\mathcal{O}'}$ and $\mathfrak{f}_{\mathcal{O}}$ for the real conductors of $\mathcal{O}'$ and $\mathcal{O}$ respectively, and suppose that that $\mathfrak{f}_{\mathcal{O}'} = \mu\mathfrak{f}_{\mathcal{O}}$. Then the map

$$\begin{array}{rccc} \rho: & \mathrm{SCl}(\mathcal{O}') & \longrightarrow & \mathrm{SCl}(\mathcal{O}) \\ & [(\mathfrak{a}', \lambda')] & \mapsto & [(\mathfrak{a}'\mathcal{O}, \lambda')]. \end{array}$$

is a surjective homomorphism. Furthermore, if $\mu\mathcal{O}_{K_0} \not| \mathfrak{f}_{\mathcal{O}}$, then

$$\#\ker(\rho) = \begin{cases} \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1 & \text{if } \mu\mathcal{O}_{K_0} \text{ is inert in } K/K_0 \\ \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) & \text{if } \mu\mathcal{O}_{K_0} \text{ is ramified in } K/K_0 \\ \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) - 1 & \text{if } \mu\mathcal{O}_{K_0} \text{ is split in } K/K_0, \end{cases}$$

and otherwise

$$\#\ker(\rho) = \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu).$$

To prove Lemma 3.9.2, we first prove some lemmas. The proofs of Lemma 3.9.3 and Lemma 3.9.6 are based on the proofs of two similar results in [BS17, Lemma 7] and [BS17, Lemma 8].

Lemma 3.9.3. Let K be a CM field with maximal totally real subfield K_0, let $\mathcal{O}$ be an order in K that contains K_0, and let $\mathfrak{f}_{\mathcal{O}}$ be the real conductor of $\mathcal{O}$. Suppose that $\mathfrak{a}$ is an invertible ideal of $\mathcal{O}$. Then $\mathfrak{a} + \mathfrak{f}_{\mathcal{O}}\mathcal{O}_K = \mathcal{O}$ if and only if $\mathfrak{a} + \mathfrak{f}_{\mathcal{O}}\mathcal{O} = \mathcal{O}$.

Proof. Recall from Proposition 3.2.1 that $\mathcal{O} = \mathcal{O}_{K_0} + \mathfrak{f}_{\mathcal{O}}\mathcal{O}_K$ so that in particular

$$\mathfrak{f}_{\mathcal{O}}\mathcal{O} \subseteq \mathfrak{f}_{\mathcal{O}}\mathcal{O}_K \subseteq \mathcal{O}$$

and hence

$$\mathfrak{a} + \mathfrak{f}_{\mathcal{O}}\mathcal{O} \subseteq \mathfrak{a} + \mathfrak{f}_{\mathcal{O}}\mathcal{O}_K \subseteq \mathfrak{a} + \mathcal{O} = \mathcal{O}.$$

So one implication is clear. It remains to prove the other implication, so assume that $\mathfrak{a} + \mathfrak{f}_{\mathcal{O}}\mathcal{O}_K = \mathcal{O}$. Observe that $\mathfrak{f}_{\mathcal{O}}\mathcal{O} = \mathfrak{f}_{\mathcal{O}}(\mathcal{O}_{K_0} + \mathfrak{f}_{\mathcal{O}}\mathcal{O}_K) = \mathfrak{f}_{\mathcal{O}} + \mathfrak{f}_{\mathcal{O}}^2\mathcal{O}_K$, giving

$$\mathfrak{f}_{\mathcal{O}}^2\mathcal{O}_K \subseteq \mathfrak{f}_{\mathcal{O}}\mathcal{O}, \tag{3.32}$$

and that

$$\mathfrak{a} + \mathfrak{f}_{\mathcal{O}}\mathcal{O}_K = \mathcal{O} = \mathcal{O}^2 = (\mathfrak{a} + \mathfrak{f}_{\mathcal{O}}\mathcal{O}_K)^2 \subseteq \mathfrak{a} + \mathfrak{f}_{\mathcal{O}}^2\mathcal{O}_K. \tag{3.33}$$

Together, (3.32) and (3.33) give us

$$\mathfrak{a} + \mathfrak{f}_{\mathcal{O}}\mathcal{O}_K \subseteq \mathfrak{a} + \mathfrak{f}_{\mathcal{O}}\mathcal{O},$$

hence

$$\mathfrak{a} + \mathfrak{f}_{\mathcal{O}}\mathcal{O} = \mathfrak{a} + \mathfrak{f}_{\mathcal{O}}\mathcal{O}_K = \mathcal{O}.$$

□

We will make implicit use of Lemma 3.9.3 from this point on.

Definition 3.9.4. For R a ring and $\mathfrak{f}$ an ideal in R, a fractional R-ideal $\mathfrak{a}$ is defined to be *coprime to* $\mathfrak{f}$ if for every prime ideal $\mathfrak{p}$ of R that divides $\mathfrak{f}$, the localisation $\mathfrak{a} \otimes_R R_{\mathfrak{p}}$ at $\mathfrak{p}$ of $\mathfrak{a}$ is $R_{\mathfrak{p}}$.

Lemma 3.9.5. Let K, $\mathcal{O}$, $\mathcal{O}'$, $\mathfrak{f}_{\mathcal{O}}$, and $\mathfrak{f}_{\mathcal{O}}$ be as in Lemma 3.9.2. For $R = \mathcal{O}$ or $\mathcal{O}'$, define $G_{\mathcal{O}}$ to be the group of invertible fractional R-ideals that are coprime to $\mathfrak{f} = \mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K$. Then there is an isomorphism of groups

$$\begin{array}{rccc} r: & G_{\mathcal{O}'} & \longrightarrow & G_{\mathcal{O}} \\ & \mathfrak{a}' & \mapsto & \mathfrak{a}'\mathcal{O}. \end{array}$$

Proof. We first check that r is well-defined. Suppose that $\mathfrak{a}'$ is an invertible fractional $\mathcal{O}'$-ideal coprime to $\mathfrak{f}$. Then for every prime ideal $\mathfrak{p}'$ of $\mathcal{O}'$ dividing $\mathfrak{f}$, we have that $\mathfrak{a}' \otimes_{\mathcal{O}'} \mathcal{O}'_{\mathfrak{p}'} = \mathcal{O}'_{\mathfrak{p}'}$. Let $\mathfrak{p}$ be a prime of $\mathcal{O}$ lying above $\mathfrak{p}'$. Then

$$\mathfrak{a}'\mathcal{O} \otimes_{\mathcal{O}} \mathcal{O}_{\mathfrak{p}} = \mathfrak{a}' \otimes_{\mathcal{O}'} \mathcal{O}_{\mathfrak{p}} = \mathfrak{a}' \otimes_{\mathcal{O}'} \mathcal{O}'_{\mathfrak{p}'} \otimes_{\mathcal{O}'_{\mathfrak{p}'}} \mathcal{O}_{\mathfrak{p}} = \mathcal{O}'_{\mathfrak{p}'} \otimes_{\mathcal{O}'_{\mathfrak{p}'}} \mathcal{O}_{\mathfrak{p}} = \mathcal{O}_{\mathfrak{p}}.$$

Hence $\mathfrak{a}'\mathcal{O}$ is coprime to $\mathfrak{f}$, so r is well-defined.

By [Ste08, Theorem 5.3], there is an isomorphism of groups

$$\begin{array}{rccc} \psi_R: & G_R & \longrightarrow & \bigoplus_{\text{prime } \mathfrak{p}\subseteq R} G_{R_\mathfrak{p}} \\ & \mathfrak{a} & \mapsto & (\mathfrak{a}\otimes\mathcal{O}_\mathfrak{p})_\mathfrak{p}. \end{array}$$

Furthermore, we claim that

$$\bigoplus_{\text{prime } \mathfrak{p}'\subseteq\mathcal{O}'} G_{\mathcal{O}'_{\mathfrak{p}'}} = \bigoplus_{\text{prime } \mathfrak{p}\subseteq\mathcal{O}} G_{\mathcal{O}_\mathfrak{p}}. \tag{3.34}$$

To see this, given a prime ideal $\mathfrak{p}\subseteq\mathcal{O}$, let $\mathfrak{p}'=\mathfrak{p}\cap\mathcal{O}'$. Then if $\mathfrak{p}\not|\mathfrak{f}$, we get an isomorphism $\mathcal{O}'_{\mathfrak{p}'}\cong\mathcal{O}_\mathfrak{p}$. This gives all primes of $\mathcal{O}'$ that are coprime to $\mathfrak{f}$. Also $G_{\mathcal{O}'_{\mathfrak{p}'}}$ is trivial if $\mathfrak{p}'|\mathfrak{f}$ and $G_{\mathcal{O}_\mathfrak{p}}$ is trivial if $\mathfrak{p}|\mathfrak{f}$, so (3.34) holds. Hence $\psi_\mathcal{O}^{-1}\circ\psi_{\mathcal{O}'}: G_{\mathcal{O}'}\to G_\mathcal{O}$ defines an isomorphism from $G_{\mathcal{O}'}$ to $G_\mathcal{O}$.

It remains to show that $r=\psi_\mathcal{O}^{-1}\circ\psi_{\mathcal{O}'}$. Let $\mathfrak{a}'\in G_{\mathcal{O}'}$. Then

$$\psi_{\mathcal{O}'}(\mathfrak{a}')=(\mathfrak{a}'\otimes\mathcal{O}'_{\mathfrak{p}'})_{\mathfrak{p}'}=(\mathfrak{a}'\otimes\mathcal{O}_\mathfrak{p})_\mathfrak{p}\in\bigoplus_{\text{prime } \mathfrak{p}\subseteq\mathcal{O}} G_{\mathcal{O}_\mathfrak{p}},$$

and

$$\psi_\mathcal{O}\circ r(\mathfrak{a}')=\psi_\mathcal{O}(\mathfrak{a}'\mathcal{O})=(\mathfrak{a}'\mathcal{O}\otimes\mathcal{O}_\mathfrak{p})_\mathfrak{p}=(\mathfrak{a}'\otimes\mathcal{O}_\mathfrak{p})_\mathfrak{p},$$

so $r=\psi_\mathcal{O}^{-1}\circ\psi_{\mathcal{O}'}$ and the lemma now follows. □

Lemma 3.9.6. Let K, $\mathcal{O}$, $\mathcal{O}'$, $\mathfrak{f}_{\mathcal{O}'}$, and ρ be as in Lemma 3.9.2. Then ρ is a surjective homomorphism and

$$\ker(\rho)\cong\frac{(\mathcal{O}/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K)^\times}{(\mathcal{O}'/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K)^\times}.$$

Proof. For any order R in K, define

$$I_R=\left\{(\mathfrak{a},\alpha):\ \begin{array}{l}\mathfrak{a}\text{ an invertible fractional } R\text{-ideal,}\\ \mathfrak{a}\overline{\mathfrak{a}}=\alpha R,\ \alpha\in K_0,\ \alpha>>0\end{array}\right\}$$

and

$$P_R=\{(xR,x\overline{x}):x\in K^*\}$$

so that

$$\mathrm{SCl}(R)=I_R/P_R.$$

Furthermore, for any ideal $\mathfrak{f}$ in R, define

$$I_R(\mathfrak{f})=\left\{(\mathfrak{a},\alpha):\ \begin{array}{c}\mathfrak{a}\text{ an invertible fractional } R\text{-ideal coprime to } \mathfrak{f},\\ \mathfrak{a}\overline{\mathfrak{a}}=\alpha R,\ \alpha\in K_0,\ \alpha>>0\end{array}\right\},$$

and define

$$P_R(\mathfrak{f})=I_R(\mathfrak{f})\cap P_R,$$

so that

$$I_R(\mathfrak{f})/P_R(\mathfrak{f})\hookrightarrow\mathrm{SCl}(R).$$

In fact, by [Ste08, Proposition 4.4] and the Chinese Remainder Theorem, we get an isomorphism

$$I_R(\mathfrak{f})/P_R(\mathfrak{f})\cong\mathrm{SCl}(R).$$

In particular, it suffices to show that

$$\begin{array}{rccc} \widetilde{\rho}: & I_{\mathcal{O}'}(\mathfrak{f})/P_{\mathcal{O}'}(\mathfrak{f}) & \longrightarrow & I_\mathcal{O}(\mathfrak{f})/P_\mathcal{O}(\mathfrak{f}) \\ & [(\mathfrak{a}',\alpha')] & \mapsto & [(\mathfrak{a}'\mathcal{O},\alpha')] \end{array}$$

with $\mathfrak{f}=\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K$ is a surjective homomorphism with kernel isomorphic to

$$\frac{(\mathcal{O}/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K)^\times}{(\mathcal{O}'/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K)^\times}.$$

We claim first that

$$\begin{array}{ccc} I_{\mathcal{O}'}(\mathfrak{f}) & \longrightarrow & I_\mathcal{O}(\mathfrak{f}) \\ (\mathfrak{a}',\alpha') & \mapsto & (\mathfrak{a}'\mathcal{O},\alpha') \end{array} \tag{3.35}$$

defines a bijection. This is well-defined and injective by Lemma 3.9.5, so we only prove surjectivity. Suppose that $(\mathfrak{a}, \alpha) \in I_{\mathcal{O}}(\mathfrak{f})$, and let $\mathfrak{a}' = r^{-1}(\mathfrak{a})$, where r is the isomorphism of Lemma 3.9.5. As r is an isomorphism, we only show that $\alpha\mathcal{O}'$ is coprime to $\mathfrak{f}$. It suffices to show that $\alpha\mathcal{O}_{K_0}$ and $\mathfrak{f}_{\mathcal{O}'}$ are coprime as $\mathcal{O}_{K_0}$-ideals: as $\mathcal{O}_{K_0} \subseteq \mathcal{O}'$ in this case

$$\mathcal{O}' = \mathcal{O}'\mathcal{O}_{K_0} = \mathcal{O}'(\alpha\mathcal{O}_{K_0} + \mathfrak{f}_{\mathcal{O}'}) \subseteq \alpha\mathcal{O}' + \mathfrak{f},$$

so $\alpha\mathcal{O}'$ is coprime to $\mathfrak{f}$. Note that $\alpha\mathcal{O}$ is coprime to $\mathfrak{f}$, so $(\alpha \bmod \mathfrak{f}) \in (\mathcal{O}/\mathfrak{f})^\times$ defines an automorphism on $\mathcal{O}/\mathfrak{f}$. Also, there is an injective ring homomorphism

$$\mathcal{O}_{K_0}/\mathfrak{f}_{\mathcal{O}'} \hookrightarrow \mathcal{O}/\mathfrak{f}$$

that sends

$$\alpha \bmod \mathfrak{f}_{\mathcal{O}'} \mapsto \alpha \bmod \mathfrak{f}.$$

Hence multiplication by $(\alpha \bmod \mathfrak{f}_{\mathcal{O}'})$ defines an injective endomorphism on $\mathcal{O}_{K_0}/\mathfrak{f}_{\mathcal{O}'}$, which is a finite ring, hence multiplication by $(\alpha \bmod \mathfrak{f}_{\mathcal{O}'})$ is an automorphism. In particular, this gives that

$$\alpha \bmod \mathfrak{f}_{\mathcal{O}'} \in (\mathcal{O}_{K_0}/\mathfrak{f}_{\mathcal{O}'})^\times$$

and hence $\alpha\mathcal{O}_{K_0}$ is coprime to $\mathfrak{f}_{\mathcal{O}'}$. We have now proved the surjectivity of (3.35), hence (3.35) is a bijection.

We have proven that $\widetilde{\rho}$ is surjective and has kernel $P_{\mathcal{O}}(\mathfrak{f})/P_{\mathcal{O}'}(\mathfrak{f})$. Suppose that $(x\mathcal{O}, x\overline{x}) \in P_{\mathcal{O}}(\mathfrak{f})$. Then there exist invertible $\mathcal{O}$-ideals $\mathfrak{b}$ and $\mathfrak{c}$, coprime to $\mathfrak{f}$, such that $x\mathcal{O} = \mathfrak{b}/\mathfrak{c}$. Furthermore, without loss of generality we may assume that $\mathfrak{b}$ and $\mathfrak{c}$ are principal: let $r \in \mathbb{Z}_{>0}$ be minimal such that $\mathfrak{c}^r$ is principal, then $x\mathcal{O} = (\mathfrak{b}\mathfrak{c}^{r-1})/\mathfrak{c}^r$, and $\mathfrak{b}\mathfrak{c}^{r-1}$ and $\mathfrak{c}^r$ are coprime to $\mathfrak{f}$ and principal. For $(x\mathcal{O}, x\overline{x}) \in P_{\mathcal{O}}(\mathfrak{f})$, choose α and $\beta \in \mathcal{O}$ such that $x\mathcal{O} = (\alpha\mathcal{O})/(\beta\mathcal{O})$ and $x\overline{x} = \alpha\beta^{-1}\overline{\alpha\beta^{-1}}$. We claim that

$$\begin{array}{rccc} i: & P_{\mathcal{O}}(\mathfrak{f}) & \longrightarrow & (\mathcal{O}/\mathfrak{f})^\times/(\mathcal{O}'/\mathfrak{f})^\times \\ & (x\mathcal{O}, x\overline{x}) & \mapsto & \alpha\beta^{-1} \end{array}$$

is a well-defined surjective morphism with kernel $P_{\mathcal{O}'}(\mathfrak{f})$. Well-defined is clear as $\alpha\beta^{-1}$ is uniquely defined up to roots of unity in $\mathcal{O}$, and $\mathcal{O}$ and $\mathcal{O}'$ have the same roots of unity by assumption. Surjectivity is also clear: for every $x + \mathfrak{f} \in (\mathcal{O}/\mathfrak{f})^\times$, we have an $\mathcal{O}$-ideal $x\mathcal{O}$ that is coprime to $\mathfrak{f}$. The kernel of i is given by

$$\{(x\mathcal{O}, x\overline{x}) \in P_{\mathcal{O}}(\mathfrak{f}) : \alpha\beta^{-1} + \mathfrak{f} \in (\mathcal{O}'/\mathfrak{f})^\times\},$$

hence

$$\begin{array}{ccc} \ker(i) & \longrightarrow & P_{\mathcal{O}'}(\mathfrak{f}) \\ (x\mathcal{O}, x\overline{x}) & \mapsto & (x\mathcal{O}', x\overline{x}) \end{array}$$

defines a bijection. This proves the lemma. □

Proof of Lemma 3.9.2. We have from Lemma 3.9.6 that ρ is a well-defined surjective homomorphism and that

$$\ker(\rho) \cong \frac{(\mathcal{O}/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K)^\times}{(\mathcal{O}'/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K)^\times}. \tag{3.36}$$

To count $\#\ker(\rho)$, we first show that

$$\ker(\rho) \cong \frac{(\mathcal{O}/\mu\mathcal{O})^\times}{(\mathcal{O}'/\mu\mathcal{O})^\times}.$$

We have that

$$\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K = \mu\mathfrak{f}_{\mathcal{O}}\mathcal{O}_K \subseteq \mu(\mathcal{O}_{K_0} + \mathfrak{f}_{\mathcal{O}}\mathcal{O}_K) = \mu\mathcal{O},$$

so that in particular there is a natural map

$$\mathcal{O}/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K \twoheadrightarrow \mathcal{O}/\mu\mathcal{O}$$

and an induced morphism of unit groups

$$(\mathcal{O}/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K)^\times \twoheadrightarrow (\mathcal{O}/\mu\mathcal{O})^\times.$$

Define

$$\varphi : (\mathcal{O}/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K)^\times \longrightarrow \frac{(\mathcal{O}/\mu\mathcal{O})^\times}{(\mathcal{O}'/\mu\mathcal{O})^\times}$$

to be the composition of this with the natural quotient morphism. We claim that

$$\ker(\varphi) = (\mathcal{O}'/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K)^\times. \tag{3.37}$$

Clearly $(\mathcal{O}'/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K)^\times \subseteq \ker(\varphi)$. To show that $\ker(\varphi) \subseteq (\mathcal{O}'/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K)^\times$, suppose that $x + \mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K \in \ker(\varphi)$. Then there exists $y \in \mathcal{O}'$ such that $x - y \in \mu\mathcal{O} \subseteq \mathcal{O}'$, so $x \in \mathcal{O}'$. Hence (3.37) holds, so that by the isomorphism theorem we have a group isomorphism

$$\frac{(\mathcal{O}/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K)^\times}{(\mathcal{O}'/\mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K)^\times} \cong \frac{(\mathcal{O}/\mu\mathcal{O})^\times}{(\mathcal{O}'/\mu\mathcal{O})^\times}. \tag{3.38}$$

Then by (3.36),

$$\#\ker(\rho) = \frac{\#(\mathcal{O}/\mu\mathcal{O})^\times}{\#(\mathcal{O}'/\mu\mathcal{O})^\times}. \tag{3.39}$$

We first count the denominator. By Proposition 3.2.1, we have that

$$\mathcal{O}' = \mathcal{O}_{K_0} + \mathfrak{f}_{\mathcal{O}'}\mathcal{O}_K = \mathcal{O}_{K_0} + \mu\mathfrak{f}_{\mathcal{O}}\mathcal{O}_K = \mathcal{O}_{K_0} + \mu\mathcal{O},$$

hence

$$\mathcal{O}'/\mu\mathcal{O} = (\mathcal{O}_{K_0} + \mu\mathcal{O})/\mu\mathcal{O} \cong \mathcal{O}_{K_0}/(\mu\mathcal{O} \cap \mathcal{O}_{K_0}) = \mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0}.$$

Write

$$\ell = \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu).$$

We assumed $\mu\mathcal{O}_{K_0}$ to be prime, so $\mathcal{O}'/\mu\mathcal{O} \cong \mathcal{O}_{K_0}/\mu\mathcal{O}_{K_0}$ is an integral domain with ℓ elements. Hence

$$\#(\mathcal{O}'/\mu\mathcal{O})^\times = \ell - 1.$$

We now count the numerator of (3.39). If either

(a) $\mu\mathcal{O}_{K_0} \nmid \mathfrak{f}\mathcal{O}$ and $\mu\mathcal{O}_{K_0}$ is ramified in K/K_0 so that $\mu\mathcal{O} = \mathfrak{m}^2$ is a square in $\mathcal{O}$, or

(b) $\mu\mathcal{O}_{K_0} \mid \mathfrak{f}_{\mathcal{O}}$, so that by Lemma 3.6.3, the $\mathcal{O}$-ideal $\mathfrak{m} = \mu\mathcal{O}_{K_0} + \mathfrak{f}_{\mathcal{O}}\mathcal{O}_K$ is the unique prime $\mathcal{O}$-ideal containing $\mu\mathcal{O}$,

then there is a unique maximal ideal $\mathfrak{m}/\mu\mathcal{O}$ in $\mathcal{O}/\mu\mathcal{O}$, and this is the set of non-units. Therefore in either case

$$\#(\mathcal{O}/\mu\mathcal{O})^\times = \#(\mathcal{O}/\mu\mathcal{O}) - \#(\mathfrak{m}/\mu\mathcal{O}) = \ell^2 - \ell = \ell(\ell - 1).$$

If $\mu\mathcal{O}_{K_0} \nmid \mathfrak{f}_{\mathcal{O}}$ and $\mu\mathcal{O}_{K_0}$ is inert in K/K_0, then $\mu\mathcal{O}$ is prime in $\mathcal{O}$ and hence $\mathcal{O}/\mu\mathcal{O}$ is an integral domain with $\mathrm{Norm}_{K/\mathbb{Q}}(\mu) = \ell^2$ elements, giving

$$\#(\mathcal{O}/\mu\mathcal{O})^\times = \ell^2 - 1.$$

Finally, if $\mu\mathcal{O}_{K_0} \nmid \mathfrak{f}_{\mathcal{O}}$ and $\mu\mathcal{O}_{K_0}$ splits in K/K_0, then there are 2 distinct prime ideals $\mathfrak{m}$ and $\overline{\mathfrak{m}}$ of $\mathcal{O}$ lying above $\mu\mathcal{O}$. This gives

$$\#(\mathcal{O}/\mu\mathcal{O})^\times = \#(\mathcal{O}/\mathfrak{m}\mathcal{O})^\times \#(\mathcal{O}/\overline{\mathfrak{m}}\mathcal{O})^\times = (\ell - 1)^2.$$

The result now follows from (3.39). □

Proof of Proposition 3.9.1. Recall that C is a connected component of the μ-isogeny graph for Weil q-number π, so in particular contains a vertex ν. If $\nu \notin V(\mathfrak{f}_C)$, there is an ascending μ-isogeny from ν by Proposition 3.7.1, so inductively we see that $V(\mathfrak{f}_C)$ is non-empty. We first show that

$$\#V(\mu\mathfrak{f}_C) = v'\#V(\mathfrak{f}_C),$$

where

$$v' = \begin{cases} \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1 & \text{if } \mu\mathcal{O}_{K_0} \text{ is inert in } K/K_0 \\ \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) & \text{if } \mu\mathcal{O}_{K_0} \text{ is ramified in } K/K_0 \\ \mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) - 1 & \text{if } \mu\mathcal{O}_{K_0} \text{ is split in } K/K_0. \end{cases}$$

By assumption $d > 0$, so by Proposition 3.8.1 there are $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) + 1$ edges from every vertex in $V(\mathfrak{f}_C)$, and by Proposition 3.6.1, we have that v' of these are non-horizontal, hence descending to $V(\mu\mathfrak{f}_C)$ by Proposition 3.3.1. Note that $\mathrm{Norm}_{K_0/\mathbb{Q}}(\mu) \geq 2$, so that $v' > 0$, and hence $V(\mu\mathfrak{f}_C)$ is non-empty. Also, by Proposition 3.5.1, for any K-order $\widetilde{\mathcal{O}}$ such that $\mathcal{O}_{K_0}[\pi, \overline{\pi}] \subseteq \widetilde{\mathcal{O}}$, if $V(\mathfrak{f}_{\mathcal{O}}) \neq \emptyset$ we have that $\#\mathrm{SCl}(\widetilde{\mathcal{O}}) = \#V(\mathfrak{f}_{\mathcal{O}})$, so by Lemma 3.9.2 we get that

$$\#V(\mu\mathfrak{f}_C) = \#\mathrm{SCl}(\mathcal{O}_{K_0} + \mu\mathfrak{f}_C\mathcal{O}_K) = v'\#\mathrm{SCl}(\mathcal{O}_{K_0} + \mathfrak{f}_C\mathcal{O}_K) = v'\#V(\mathfrak{f}_C).$$

Observe that there is a *unique* ascending μ-isogeny from every vertex in $V(\mu\mathfrak{f}_C)$: by Proposition 3.7.1 there is an ascending μ-isogeny from every vertex in $V(\mu\mathfrak{f}_C)$, which accounts for $\#V(\mu\mathfrak{f}_C)$ of the descending μ-isogenies from $V(\mathfrak{f}_C)$, but this is all of them as $\#V(\mu\mathfrak{f}_C) = v'\#V(\mathfrak{f}_C)$.

By induction, for every $1 < i < d$ we have that

$$\#V(\mu^{i+1}\mathfrak{f}_C) = (v-1)\#V(\mu^i\mathfrak{f}_C).$$

(The induction is the same argument as for $i = 1$ above, where we replace the horizontal edges between elements of $V(\mathfrak{f}_C)$ by the unique ascending edge from every element of $V(\mu^{i-1}\mathfrak{f}_C)$.) □

3.10 Example computation of a μ-isogeny graph

All the calculations for this example were done in Sage [Sage]. Let us consider the curve

$$\begin{aligned}\mathcal{C} : y^2 =& 902701461021360x^6 + 938022069033830x^5 + 2496384827106779x^4 \\ &+ 560788189813847x^3 + 2116308108498283x^2 \\ &+ 1865564692722366x + 2658210628678317\end{aligned}$$

defined over $\mathbb{F}_p$, with $p = 2681144777671301$, which is a prime. This curve was taken from the Echidna Database [Echidna, `https://www.i2m.univ-amu.fr/perso/david.kohel/dbs/cgi-bin/quartic_cm_field.py?D=5&A=37&B=281`] and has endomorphism ring isomorphic to the maximal order of the quartic CM-field

$$K := \mathbb{Q}[x]/(x^4 + 37x^2 + 281),$$

in which p splits completely and $p = \pi\overline{\pi}$, where $K = \mathbb{Q}(\pi)$, and π is the Frobenius morphism on the Jacobian of $\mathcal{C}$. The minimal polynomial of π is

$$\begin{aligned}\chi_\pi(x) =& x^4 - 605104x^3 - 5215893977257194x^2 - 1622371429548014920304x \\ &+ 7188537318834090069340399032601.\end{aligned}$$

The maximal totally real subfield of K is $K_0 = \mathbb{Q}(\sqrt{5})$, and we will now fix

$$\mu = (5 + \sqrt{5})/2.$$

Then μ is a totally positive algebraic integer in K_0 with norm 5, and $\mu\mathcal{O}_K$ splits into prime ideals of $\mathcal{O}_K$ as $\mathfrak{m}\overline{\mathfrak{m}}$, where

$$\begin{aligned}\mathfrak{m} =& 5\mathcal{O}_K + (-66584412017/97334935924869034947945714865000 0\pi^3 \\ &+ 17464102246896083/486674679624345174739728574325000\pi^2 \\ &+ 1793587767083954706901049 69/97334935924869034947945714865000 0\pi \\ &- 22924673687227109/181517493451825000)\mathcal{O}_K.\end{aligned}$$

We can easily check that the order of $[(\mathfrak{m}, \mu)]$ in $\mathrm{SCl}(\mathcal{O}_K)$ is 3, and that

$$d = \max\{k \in \mathbb{Z} : \mathcal{O}_{K_0}[\pi, \overline{\pi}] \subseteq \mathcal{O}_{K_0} + \mu^k\mathcal{O}_K\} = 4,$$

so that by Theorem 3.1.9, the connected component of the μ-isogeny graph in which $\mathcal{C}$ lies is a $(C_3, 6, 4)$-volcano, pictured below.

Colour	Blue	Purple	Green	Orange	Red
End(A)	$\mathcal{O}_K$	$\mathcal{O}_{K_0} + \mu\mathcal{O}_K$	$\mathcal{O}_{K_0} + \mu^2\mathcal{O}_K$	$\mathcal{O}_{K_0} + \mu^3\mathcal{O}_K$	$\mathcal{O}_{K_0} + \mu^4\mathcal{O}_K$

Table 3.1: Colour coding

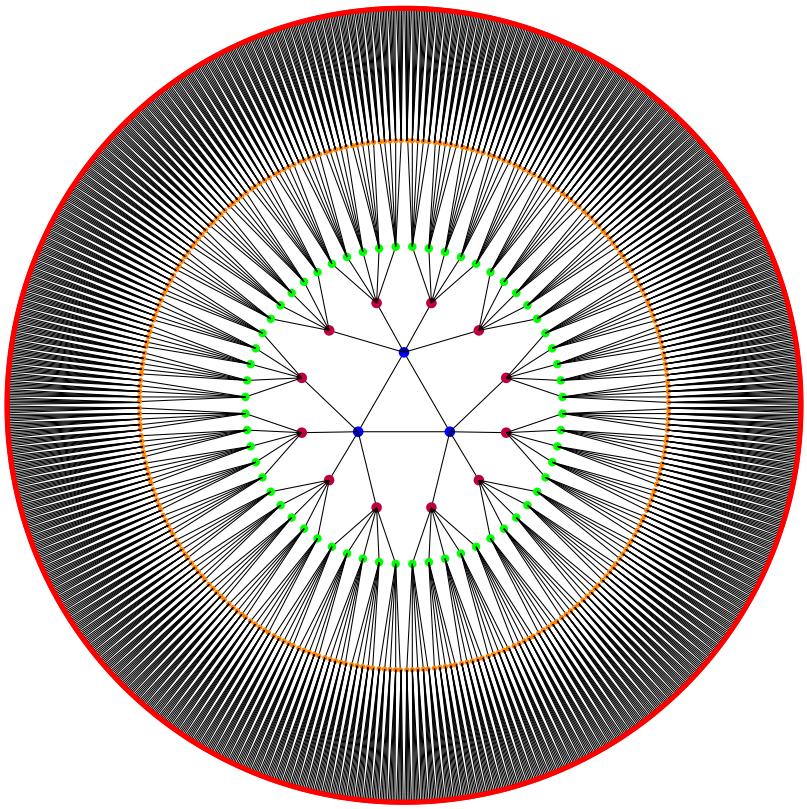

A $(C_3, 6, 4)$-volcano.

Chapter 4

Isogenies for point counting on genus two hyperelliptic curves with maximal real multiplication

This chapter is joint work with Ballentine, Guillevic, Lorenzo-García, Massierer, Smith, and Top, and has been published as [Bal+17].

This chapter reports on work carried out at the workshop *Algebraic Geometry for Coding Theory and Cryptography* at the Institute for Pure and Applied Mathematics (IPAM), University of California, Los Angeles, February 22–26, 2016. The authors thank IPAM for its generous support.

Please note that the numbering in this thesis is different from the published version; due to this being Chapter 4 of the thesis every number of the form $4.x$ appears in the published version as $3.x$.

Abstract

Schoof's classic algorithm allows point-counting for elliptic curves over finite fields in polynomial time. This algorithm was subsequently improved by Atkin, using factorizations of modular polynomials, and by Elkies, using a theory of explicit isogenies. Moving to Jacobians of genus-2 curves, the current state of the art for point counting is a generalization of Schoof's algorithm. While we are currently missing the tools we need to generalize Elkies' methods to genus 2, recently Martindale and Milio have computed analogues of modular polynomials for genus-2 curves whose Jacobians have real multiplication by maximal orders of small discriminant. In this article, we prove Atkin-style results for genus-2 Jacobians with real multiplication by maximal orders, with a view to using these new modular polynomials to improve the practicality of point-counting algorithms for these curves.

4.1 Introduction

Efficiently computing the number of points on the Jacobian of a genus 2 curve over a finite field is an important problem in experimental number theory and number-theoretic cryptography. When the characteristic of the finite field is small, Kedlaya's algorithm and its descendants provide an efficient solution (see [Ked01], [Har07], and [Har12]), while in extremely small characteristic we have extremely fast AGM-style algorithms (see for example [Mes01], [Mes02], and [Car04]). However, the running times of these algorithms are exponential in the size of the field characteristic; the hardest case, therefore (and also the most important case for contemporary cryptographic applications) is where the characteristic is large, or even where the field is a prime field.

So let q be a power of a large prime p, and let $\mathcal{C}$ be a genus-2 curve over $\mathbb{F}_q$. Our fundamental problem is to compute the number of $\mathbb{F}_q$-rational points on the Jacobian $J_{\mathcal{C}}$ of $\mathcal{C}$.

4.1.1 The state of the art

In theory, the problem is solved: we can compute $\#J_{\mathcal{C}}(\mathbb{F}_q)$ in polynomial time (that is, polynomial in $\log q$) using Pila's algorithm [Pil90], which is the immediate generalization of Schoof's elliptic-curve point-counting algorithm [Sch85] to higher-dimensional abelian varieties. But the exponent in Pila's polynomial time is extremely large; so, despite its theoretical importance, this algorithm is completely impractical (see §4.3.4). Indeed, to our knowledge it has never been implemented.

Gaudry and Schost have developed and successfully implemented a much more practical variant of Pila's algorithm for the case $q = p$ that runs in time $\widetilde{O}(\log^8 p)$; not just polynomial time, but on the edge of practicality [GS12]. Still, their algorithm requires an extremely intensive calculation for cryptographic-sized Jacobians: Gaudry and Schost estimated a running time of around one core-month (in 2008) to compute $\#J_{\mathcal{C}}(\mathbb{F}_p)$ when p has around 128 bits [GS12].

The situation improves dramatically if $J_{\mathcal{C}}$ is equipped with an efficiently computable *real multiplication* endomorphism. For such Jacobians, Gaudry, Kohel, and Smith [GKS11] give an algorithm to compute $\#J_{\mathcal{C}}(\mathbb{F}_q)$ in time $\widetilde{O}(\log^5 q)$. This allowed the computation of $\#J_{\mathcal{C}}(\mathbb{F}_p)$ for one curve $\mathcal{C}$ drawn from the genus-2 family in [TTV91] with $p = 2^{512} + 1273$ in about 80 core-days (in 2011); this remains, to date, the record for genus-2 point counting over prime fields. For 128-bit fields, the cost is reduced to 3 core hours (in 2011).

All of these algorithms are generalizations of Schoof's algorithm, which computes the Frobenius trace (and hence the order $\#E(\mathbb{F}_q)$) of an elliptic curve $E/\mathbb{F}_q$ modulo ℓ for a series of small primes ℓ by considering the action of Frobenius on the ℓ-torsion. But Schoof's algorithm is not the state of the art for elliptic-curve point counting: it has evolved into the much faster Schoof–Elkies–Atkin (SEA) algorithm, surveyed in [Sch95]. Atkin's improvements involve factoring the ℓ-th modular polynomial (evaluated at the j-invariant of the target curve) to deduce information on the Galois structure of the ℓ-torsion, which then restricts the possible values of the trace modulo ℓ (see §4.2.6). Elkies' improvements involve computing the kernel of a rational ℓ-isogeny, which takes the place of the full ℓ-torsion; deducing the existence of the isogeny, and computing its kernel, requires finding a root of the ℓ-th modular polynomial evaluated at the j-invariant of the target curve (see §4.2.7).

4.1.2 Our contributions, and beyond

Our ultimate goal is to generalize Atkin's and Elkies' improvements to genus 2. In this article, we concentrate on generalizing Atkin's methods to genus-2 Jacobians with known real multiplication. This project is prompted by the recent appearance of two new algorithms for computing modular ideals, the genus-2 analogue of modular polynomials: Milio [Mil15a] has computed modular ideals for general genus-2 Jacobians, while Milio [Mil15b, §5] and Martindale [Mar18] have independently computed modular ideals for genus-2 Jacobians with RM by orders of small discriminants.

To extend Elkies' methods to genus 2 we would need an analogue of Elkies' algorithm [Sch95, §§7-8], which computes defining equations for the kernel of an isogeny of elliptic curves (and the isogeny itself) corresponding to a root of the evaluated modular polynomial. We do not know of any such algorithm in genus 2. Couveignes and Ezome have recently developed an algorithm to compute explicit (ℓ, ℓ)-isogenies of genus-2 Jacobians [CE15], presuming that the kernel has already been constructed somehow—but kernel construction is precisely the missing step that we need.[1]

In contrast, Atkin's improvements for elliptic-curve Schoof require nothing beyond the modular polynomial itself; so we can hope to achieve something immediately in genus 2 by generalizing Atkin's results on factorizations of modular polynomials to the decomposition of genus-2 modular ideals. This is precisely what we do in this article.

We focus on the RM case for three reasons. First, the construction of explicit modular ideals is furthest advanced in this case: Milio has constructed modular ideals for primes in $\mathbb{Q}(\sqrt{5})$ of norm up to 31, while for general Jacobians the current limit is 3. It is therefore already possible to compute nontrivial and interesting examples in the RM case. Second, point counting is currently much more efficient for Jacobians with efficiently computable RM; we hope that, at some point, our methods can help tip RM point counting from "feasible" into "routine". Third, from a purely theoretical point of view, the RM case is more similar to the elliptic curve case in the sense that real multiplication allows us, in favorable circumstances, to split ℓ-torsion subgroups of the Jacobian into groups of the same size as encountered for elliptic curves.

After recalling the SEA algorithm for elliptic curves in §4.2, we describe the current state of genus 2 point counting, and set out our program for a generalized SEA algorithm in §4.3. We describe the modular invariants we need for this in §4.4, and the modular ideals that relate them in §4.4.2. We can then state and prove our main theoretical results, which are generalizations of Atkin's theorems for these modular ideals, in §4.5. In §4.6 we provide some concrete details on the special case of RM by $\mathbb{Q}(\sqrt{5})$, before concluding with some experimental results in §4.7.

4.1.3 Vanilla abelian varieties

We can substantially simplify the task ahead by restricting our attention to a class of elliptic curves and Jacobians (more generally, abelian varieties) with sufficiently general CM endomorphism rings. The following definition makes this precise.

Definition 4.1. We say that a g-dimensional abelian variety $\mathcal{A}/\mathbb{F}_q$ is *vanilla*[2] if its endomorphism algebra $\mathrm{End}_{\overline{\mathbb{F}}_q}(\mathcal{A}) \otimes \mathbb{Q}$ (over the algebraic closure) is a CM field of degree $2g$ that does *not* contain any roots of unity other than ± 1.

If an elliptic curve $\mathcal{E}/\mathbb{F}_q$ is vanilla, then $\mathcal{E}$ is nonsupersingular and $j(\mathcal{E})$ is neither 0 nor 1728: these are the conditions Schoof applies systematically in [Sch95]. We note that in particular, vanilla abelian varieties are absolutely simple.

To fix notation, we recall that if $\mathcal{A}$ is an abelian variety, then a *principal polarization* is an isomorphism $\xi\colon \mathcal{A} \to \mathcal{A}^\vee$ associated with an ample divisor class on $\mathcal{A}$, where $\mathcal{A}^\vee = \mathrm{Pic}^0(\mathcal{A})$ is the dual abelian variety (see for example [Mil86, §13]). We will be working with elliptic curves and Jacobians of genus-2 curves; these all have a canonical principal polarization. Each endomorphism ϕ of $\mathcal{A}$ has a corresponding dual

[1] We would also like mention Bisson, Cosset, and Robert's `AVIsogenies` software package [BCR], which provides some functionality in this direction. However, their methods apply to abelian surfaces with a lot of rational 2- and 4-torsion, and applying them to general genus-2 Jacobians (with or without known RM) generally requires a substantial extension of the base field to make that torsion rational. This is counterproductive in the context of point counting.

[2] Vanilla is the most common and least complicated flavour of abelian varieties over finite fields. Heuristically, over large finite fields, randomly sampled abelian varieties are vanilla with overwhelming probability. Indeed, being vanilla is invariant in isogeny classes, and Howe and Zhu have shown in [HZ02, Theorem 2] that the fraction of isogeny classes of g-dimensional abelian varieties over $\mathbb{F}_q$ that are ordinary and absolutely simple tends to 1 as $q \to \infty$. All absolutely simple ordinary abelian varieties are vanilla, except those whose endomorphism algebras contain roots of unity; but the number of such isogeny classes for fixed g is asymptotically negligible.

endomorphism $\phi^\vee$ of $\mathcal{A}^\vee$. If $(\mathcal{A}, \xi)$ is a principally polarized abelian variety, then ξ induces a *Rosati involution* on $\mathrm{End}(\mathcal{A})$, defined by

$$\phi \longmapsto \phi^\dagger := \xi^{-1} \circ \phi^\vee \circ \xi \quad \text{for } \phi \in \mathrm{End}(\mathcal{A}) \ .$$

In the world of elliptic curves, the Rosati involution is the familiar dual. For vanilla abelian varieties, the Rosati involution acts as complex conjugation on the endomorphism ring.

Fix a real quadratic field $F = \mathbb{Q}(\sqrt{\Delta})$, with fundamental discriminant $\Delta > 0$ and ring of integers $\mathcal{O}_F$. We write $\alpha \mapsto \bar{\alpha}$ for the involution of F over $\mathbb{Q}$; we emphasize that in this article, $\bar{\cdot}$ does *not* denote complex conjugation.

From a theoretical point of view, when talking about real multiplication, our fundamental data are triples $(\mathcal{A}, \xi, \iota)$ where $\mathcal{A}$ is an abelian surface, $\xi \colon \mathcal{A} \to \mathcal{A}^\vee$ is a principal polarization, and $\iota \colon \mathcal{O}_F \hookrightarrow \mathrm{End}(\mathcal{A})$ is an embedding stable under the Rosati involution (that is, $\iota(\mu)^\dagger = \iota(\mu)$ for all μ in $\mathcal{O}_F$; we can then think of the Rosati involution as complex conjugation on the endomorphism ring). While this notation $(\mathcal{A}, \xi, \iota)$ may seem quite heavy at first glance, we remind the reader that generally there are only two choices of embedding ι (corresponding to the two square roots of Δ), and we are only really interested in the case where $\mathcal{A}$ is a Jacobian, in which case the polarization ξ is canonically determined.

4.2 Genus one curves: elliptic curve point counting

We begin by briefly recalling the SEA algorithm for elliptic curve point counting in large characteristic. First we describe Schoof's original algorithm [Sch95], before outlining the improvements of Elkies and Atkin. This will provide a point of reference for comparisons with genus-2 algorithms.

Let $\mathcal{E}$ be an elliptic curve over a finite field $\mathbb{F}_q$ of large characteristic (or at least, with $\mathrm{char}(\mathbb{F}_q) \gg \log q$). We may suppose that $\mathcal{E}$ is defined by a (short) Weierstrass equation $\mathcal{E} : y^2 = x^3 + ax + b$, with a and b in $\mathbb{F}_q$.

Like all modern point-counting algorithms, the Schoof and SEA algorithms compute the characteristic polynomial

$$\chi_\pi(X) = X^2 - tX + q$$

of the Frobenius endomorphism π of $\mathcal{E}$. We call t the *trace* of Frobenius. Since the $\mathbb{F}_q$-rational points on $\mathcal{E}$ are precisely the fixed points of π, we have

$$\#\mathcal{E}(\mathbb{F}_q) = \chi_\pi(1) = q + 1 - t \ ;$$

so determining $\#\mathcal{E}(\mathbb{F}_q)$ is equivalent to determining t. Hasse's theorem tells us that

$$|t| \leq 2\sqrt{q} \ . \tag{4.1}$$

4.2.1 Schoof's algorithm

Schoof's basic strategy is to choose a set $\mathcal{L}$ of primes $\ell \neq p$ such that $\prod_{\ell \in \mathcal{L}} \ell > 4\sqrt{q}$. We then compute $t_\ell := t \bmod \ell$ for each of the primes ℓ in $\mathcal{L}$, and then recover the value of t from $\{(t_\ell, \ell) : \ell \in \mathcal{L}\}$ using the Chinese Remainder Theorem. The condition $\prod_{\ell \in \mathcal{L}} \ell > 4\sqrt{q}$ ensures that t is completely determined by the collection of t_ℓ (by Hasse's theorem, Equation (4.1)).

For Schoof's original algorithm, the natural choice is to let $\mathcal{L}$ be the set of the first $O(\log q)$ primes, stopping when the condition $\prod_{\ell \in \mathcal{L}} \ell > 4\sqrt{q}$ is satisfied. When applying Elkies' and Atkin's modifications, we will need to be more subtle with our choice of $\mathcal{L}$. It is also possible to replace primes with small prime powers; we will not explore this option here.

Now, let ℓ be one of our primes in $\mathcal{L}$; our aim is to compute t_ℓ. We know that $\pi^2(P) - [t]\pi(P) + [q]P = 0$ for all P in $\mathcal{E}$, and hence

$$\pi^2(P) - [t_\ell]\pi(P) + [q \bmod \ell]P = 0 \quad \text{for all } P \in \mathcal{E}[\ell] \ .$$

We can therefore compute t_ℓ as follows:

1. Construct a point P of order ℓ.
2. Compute $Q = \pi(P)$ and $R = \pi^2(P) + [q \bmod \ell]P$.

3. Search for $0 \le t_\ell < \ell$ such that $[t_\ell]Q = R$, using Shanks' baby-step giant-step algorithm in the cyclic subgroup of the ℓ-torsion generated by Q.

To construct such a P, we begin by computing the ℓ-th division polynomial Ψ_ℓ in $\mathbb{F}_q[X]$, which is the polynomial whose roots in $\overline{\mathbb{F}}_q$ are precisely the x-coordinates of the nontrivial points in $\mathcal{E}[\ell]$. When ℓ is odd and prime to q, we have $\deg \Psi_\ell = (\ell^2 - 1)/2$. We then define the ring

$$A = \mathbb{F}_q[X, Y]/(\Psi_\ell(X), Y^2 - X^3 - aX - b),$$

and take $P = (X, Y)$ in $\mathcal{E}(A)$.

In order to work efficiently with $Q = \pi(P) = (X^q, Y^q)$ in the search for t_ℓ, we need to compute a compact form for Q. This means computing reduced representatives for X^q and Y^q in the ring A—that is, reducing X^q modulo $\Psi_\ell(X)$ and Y^q modulo $(\Psi_\ell(X), Y^2 - X^3 - aX - b)$—which costs $O(\log q)$ $\mathbb{F}_q$-operations.

Having computed t_ℓ for each ℓ in $\mathcal{L}$, we recover t (and hence χ_π) using the Chinese Remainder Theorem; this then yields $\#\mathcal{E}(\mathbb{F}_q) = q + 1 - t$. In cryptographic contexts, we are generally interested in curves of (almost) prime order. One particularly convenient feature of Schoof's algorithm is that it allows us to detect small prime factors of $\#\mathcal{E}(\mathbb{F}_q)$ early: we can determine if any ℓ in $\mathcal{L}$ divides $\#\mathcal{E}(\mathbb{F}_q)$ by checking whether $t_\ell \equiv q + 1 \pmod{\ell}$. If we find such a factor, then we can immediately abort the calculation of t and move on to another candidate curve.

The cost to compute χ_ℓ is $\widetilde{O}(\ell^2 + (\log q)\ell^2 + \sqrt{\ell}\ell^2)$ $\mathcal{F}_q$-operations. We can take $\mathcal{L}$ to be a set of $O(\log q)$ primes, the largest of which is in $O(\log q)$; the total cost is therefore $\widetilde{O}(\log^4 q)$ $\mathbb{F}_q$-operations.

4.2.2 Frobenius eigenvalues and subgroups

Fix a basis of $\mathcal{E}[\ell]$, and thus an isomorphism $\mathcal{E}[\ell] \cong {\mathbb{F}_\ell}^2$. Now π acts on $\mathcal{E}[\ell]$ as an element of $\mathrm{GL}_2(\mathbb{F}_\ell)$. The local characteristic polynomial χ_ℓ is just the characteristic polynomial of this matrix.

Likewise, π permutes the ℓ-subgroups of $\mathcal{E}[\ell]$; that is, the one-dimensional subspaces of $\mathcal{E}[\ell] \cong {\mathbb{F}_\ell}^2$. These are the points of $\mathbb{P}(\mathcal{E}[\ell]) \cong \mathbb{P}^1(\mathbb{F}_\ell)$, and we can consider the image of π in $\mathrm{PGL}_2(\mathbb{F}_\ell) \cong \mathrm{Aut}(\mathbb{P}(\mathcal{E}[\ell]))$. The order of π as an element of $\mathrm{PGL}_2(\mathbb{F}_\ell)$ is clearly independent of the choice of basis.

Proposition 4.2. *Let $\mathcal{E}/\mathbb{F}_q$ be an elliptic curve with Frobenius endomorphism π, and let $\ell \neq p = \mathrm{char}(\mathbb{F}_q)$ be an odd prime. If e is the order of the image of π in $\mathrm{PGL}_2(\mathbb{F}_\ell)$, then the trace t of π satisfies*

$$t^2 = \eta_e q \quad \text{in } \mathbb{F}_\ell \ ,$$

where $\eta_e = \begin{cases} \zeta + \zeta^{-1} + 2 \text{ with } \zeta \in \mathbb{F}_{\ell^2}^\times \text{ of order } e & \text{if } \gcd(\ell, e) = 1 \ , \\ 4 & \text{otherwise} \ . \end{cases}$

Proof. We follow the proof of [Sch95, Proposition 6.2] (correcting the minor error that leads in the case e even to an $e/2$-th rather than e-th root of unity appearing in the last part of the statement). Let $\lambda_1, \lambda_2 \in \mathbb{F}_{\ell^2}$ be the eigenvalues of the image of π in $\mathrm{Aut}(\mathcal{E}[\ell]) \cong \mathrm{GL}_2(\mathbb{F}_\ell)$; then

$$\lambda_1 + \lambda_2 = t \quad \text{and} \quad \lambda_1\lambda_2 = q \quad \text{in} \quad \mathbb{F}_\ell \ .$$

In case $\lambda_1 = \lambda_2$ we have $e \mid \ell$ and the assertion follows. In case $\lambda_1 \neq \lambda_2$ the given e is the minimal integer > 0 with $\lambda_1^e = \lambda_2^e$. In particular $\gcd(e, \ell) = 1$ and $\lambda_2 = \lambda_1\zeta$ for some primitive e=th root of unity ζ (in $\mathbb{F}_{\ell^2}$; in fact $e \mid \ell - 1$ in case the eigenvalues are in $\mathbb{F}_\ell$ and $e \mid \ell + 1$ otherwise). Hence $q = \lambda_1\lambda_2 = \lambda_1^2\zeta$ which implies

$$t^2 = (\lambda_1 + \lambda_2)^2 = \lambda_1^2(1 + \zeta)^2 = q\zeta^{-1}(\zeta^2 + 2\zeta + 1) = (\zeta + \zeta^{-1} + 2)q.$$

□

4.2.3 Modular polynomials and isogenies

The order-ℓ subgroups of $\mathcal{E}[\ell]$ are precisely the kernels of ℓ-isogenies from $\mathcal{E}$ to other elliptic curves, and the set of all such ℓ-isogenies (up to isomorphism) corresponds to the set of roots of $\Phi_\ell(j(\mathcal{E}), x)$ in $\overline{\mathbb{F}}_q$. The classical modular polynomial $\Phi_\ell(X, Y)$, of degree $\ell + 1$ (in X and Y) over $\mathbb{Z}$, is defined by the property that $\Phi_\ell(j(\mathcal{E}_1), j(\mathcal{E}_2)) = 0$ precisely when there exists an ℓ-isogeny $\mathcal{E}_1 \to \mathcal{E}_2$. For ℓ in $O(\log q)$, one can compute $\Phi_\ell(j(\mathcal{E}), x)$ in $\widetilde{O}(\ell^3)$ $\mathcal{F}_q$-operations using Sutherland's algorithm [Sut13]. Alternatively, we can

use precomputed databases of modular polynomials over $\mathbb{Z}$, reducing them modulo p and specializing them at $j(\mathcal{E})$.

The Galois orbits of the roots of $\Phi_\ell(j(\mathcal{E}), x)$ correspond to orbits of ℓ-isogeny kernels under π, and to orbits of points of $\mathbb{P}^1(\mathbb{F}_\ell)$ under the image of π in $\mathrm{PGL}_2(\mathbb{F}_\ell)$. If $j(\mathcal{E}_1)$ and $j(\mathcal{E}_2)$ are both in $\mathbb{F}_{q^k}$, then the isogeny is defined over $\mathbb{F}_{q^k}$ (up to a possible twist); in particular, its kernel is defined over $\mathbb{F}_{q^k}$. More precisely, we have the following key lemma:

Lemma 4.3 (Proposition 6.1 of [Sch95])**.** *Let $\mathcal{E}/\mathbb{F}_q$ be a vanilla elliptic curve with Frobenius endomorphism π.*

1. *The polynomial $\Phi_\ell(j(\mathcal{E}), x)$ has a root in $\mathbb{F}_{q^e}$ if and only if the kernel of the corresponding ℓ-isogeny is a one-dimensional eigenspace of π^e in $\mathcal{E}[\ell]$.*

2. *The polynomial $\Phi_\ell(j(\mathcal{E}), x)$ splits completely over $\mathbb{F}_{q^d}$ if and only if π^d acts as a scalar matrix on $\mathcal{E}[\ell]$; that is, if and only if d is a multiple of the order e of the image of π in $\mathrm{PGL}_2(\mathbb{F}_\ell)$. In particular, the minimal such d is e.*

4.2.4 Elkies, Atkin, and volcanic primes

The primes $\ell \neq p$ are divided into 3 classes, or types, with respect to a given $\mathcal{E}/\mathbb{F}_q$: *Elkies*, *Atkin*, and *volcanic*. The type of ℓ simultaneously reflects the factorization of $\Phi_\ell(j(\mathcal{E}), x)$ and the Galois structure of the ℓ-subgroups of $\mathcal{E}[\ell]$. Here we recall a number of facts about these classes, all of which are proven in [Sch95, §6]; see also [Was08, §12.4].

A prime ℓ is **Elkies** if the ideal (ℓ) is split in $\mathbb{Z}[\pi]$; or, equivalently, if $t^2 - 4q$ is a nonzero square modulo ℓ. Each of the two prime ideals over (ℓ) defines the kernel of an ℓ-isogeny, $\phi_i \colon \mathcal{E} \to \mathcal{E}_i$ for $i = 1, 2$, say. This means that $j(\mathcal{E}_1)$ and $j(\mathcal{E}_2)$ must be roots in $\mathbb{F}_q$ of $\Phi_\ell(j(\mathcal{E}), x)$. Lemma 4.3 then implies that

$$\Phi_\ell(j(\mathcal{E}), x) = (x - j(\mathcal{E}_1))(x - j(\mathcal{E}_2)) \prod_{i=1}^{(\ell-1)/e} f_i(x)$$

where each of the f_i are irreducible of degree e, and $e > 1$ is the order of the image of π in $\mathrm{PGL}_2(\mathbb{F}_\ell)$, which must divide $\ell - 1$ in this case.

A prime ℓ is **Atkin** if the ideal (ℓ) is inert in $\mathbb{Z}[\pi]$; or, equivalently, if $t^2 - 4q$ is *not* a square modulo ℓ. There are *no* $\mathbb{F}_q$-rational ℓ-isogenies from $\mathcal{E}$, and no $\mathbb{F}_q$-rational ℓ-subgroups of $\mathcal{E}[\ell]$. Looking at the modular polynomial, Lemma 4.3 implies

$$\Phi_\ell(j(\mathcal{E}), x) = \prod_{i=1}^{(\ell+1)/e} f_i(x) \ ,$$

where each of the f_i is an irreducible polynomial of degree e, and $e > 1$ is the order of the image of π in $\mathrm{PGL}_2(\mathbb{F}_\ell)$, which must divide $\ell + 1$ in this case.

Finally, a prime ℓ is **volcanic** if the ideal (ℓ) is ramified in $\mathbb{Z}[\pi]$; or, equivalently, if ℓ divides $t^2 - 4q$. Applying Lemma 4.3, either

$$\Phi_\ell(j(\mathcal{E}), x) = \prod_{i=1}^{\ell+1} (x - j_i)$$

with all of the j_i in $\mathbb{F}_q$ (so there are $\ell + 1$ rational ℓ-isogenies, and $\ell + 1$ rational ℓ-subgroups of $\mathcal{E}[\ell]$); or

$$\Phi_\ell(j(\mathcal{E}), x) = (Y - j_1) \cdot f(x) \ ,$$

with f irreducible of degree ℓ (so there is a single rational ℓ-isogeny, and one rational ℓ-subgroup of $\mathcal{E}[\ell]$). In either situation, $\pi|_{\mathcal{E}[\ell]}$ acts on $\mathcal{E}[\ell]$ with eigenvalues $\lambda_1 = \lambda_2$, so its image in $\mathrm{PGL}_2(\mathbb{F}_\ell)$ therefore has order $e \mid \ell$.

We note an interesting and useful fact in passing: if $\mathcal{E}/\mathbb{F}_q$ is vanilla, $\ell \neq p$ is an odd prime, and r is the number of irreducible factors of $\Phi_\ell(j(\mathcal{E}), x)$, then

$$(-1)^r = \left(\frac{q}{\ell}\right) \tag{4.2}$$

(cf. [Sch95, Proposition 6.3]; the proof generalizes easily from $q = p$ to general prime powers).

4.2.5 Computing the type of a prime

The type of a given prime ℓ for $\mathcal{E}$ (that is, being volcanic, Atkin, or Elkies) is defined in terms of the structure of $\mathbb{Z}[\pi]$ and the trace t. When we are point-counting, these are unknown quantities; but we can still determine the type of ℓ *without* knowing t or $\mathbb{Z}[\pi]$, by factoring $\Phi_\ell(j(\mathcal{E}), x)$ and comparing with the possible factorization types above. This, in turn, gives us useful information about t and $\mathbb{Z}[\pi]$. Determining the type of ℓ in this way costs $\widetilde{O}(\ell^2 + (\log q)\ell)$ $\mathcal{F}_q$-operations.

In fact, computing the type of ℓ for $\mathcal{E}$ is a good way of checking the correctness of a claimed modular polynomial. Suppose somebody has computed a polynomial $F(J_1, J_2)$, and claims it is equal to Φ_ℓ. The factorization patterns for modular polynomials corresponding to the prime types above are so special that there is very little hope of getting these patterns for $F(j(\mathcal{E}), x)$ for varying $\mathcal{E}$ and p unless F and Φ_ℓ define the same variety in the (J_1, J_2)-plane. We will use the genus-2 analogue of this observation in §4.7 to check the correctness of some of Martindale's modular polynomials.

4.2.6 Atkin's improvement

Atkin's contribution to the SEA algorithm was to exploit the factorization type of the modular polynomial to restrict the possible values of $t \pmod{\ell}$. While this does not improve the asymptotic complexity of Schoof's algorithm, it did allow significant practical progress before the advent of Elkies' improvements.

For example: if ℓ is volcanic, then by definition

$$t^2 = 4q \quad \text{in } \mathbb{F}_\ell \ , \tag{4.3}$$

which determines t_ℓ up to sign: $t \equiv \pm 2\sqrt{q} \pmod{\ell}$. Note that this is also a consequence of Proposition 4.2, which we will now apply to the other two prime types.

If ℓ is Elkies or Atkin for $\mathcal{E}$, then Proposition 4.2 tells us that

$$t^2 = (\zeta + \zeta^{-1} + 2)q \quad \text{in } \mathbb{F}_\ell \tag{4.4}$$

for some primitive e-th root of unity ζ in $\mathbb{F}_{\ell^2}$, where $e \mid \ell - 1$ if ℓ is Elkies and $e \mid \ell + 1$ if ℓ is Atkin. The number of possible values of t_ℓ^2 is therefore half the number of primitive e-th roots in these cases. Note that modular polynomials can only give us information about t_ℓ^2—that is, t_ℓ up to sign—since their solutions tell us about isogenies only up to quadratic twists, and twisting changes the sign of the trace.

Obviously, the smaller the degree e of the non-linear factors of $\Phi_\ell(j(\mathcal{E}), x)$, the fewer the values that t_ℓ can possibly take. For example, if $e = 2$ then $t_\ell = 0$; if $e = 3$, then $t_\ell = \pm\sqrt{q}$ in $\mathbb{F}_\ell$; and if $e = 4$, then $t_\ell = \pm\sqrt{2q}$ in $\mathbb{F}_\ell$.

The challenging part of Atkin's technique is making use of these extra modular congruences. Atkin's *match-and-sort* algorithm (see for example [Ler97, §11.2]) is a sort of sophisticated baby-step giant-step in $\mathcal{E}(\mathbb{F}_q)$ exploiting this modular information. Alternatively, we can use Joux and Lercier's *Chinese-and-match* algorithm [JL01].

4.2.7 Elkies' improvement

Elkies' contribution to the SEA algorithm was to note that when computing t_ℓ, we can replace $\mathcal{E}[\ell]$ with the kernel of a rational ℓ-isogeny, if it exists. Looking at the classification of primes, we see that there exists a rational ℓ-isogeny precisely when ℓ is volcanic or Elkies (whence the terminology). Of course, as we saw above, if ℓ is one of the rare volcanic primes then t_ℓ is already determined up to sign; it remains to see what can be done for Elkies primes.

Let ℓ be an Elkies prime for $\mathcal{E}$, and let ϕ_1 and ϕ_2 be ℓ-isogenies corresponding to the two roots of $\Phi_\ell(j(\mathcal{E}), x)$ in $\mathbb{F}_q$. First, we note that $\pi(P_i) = [\lambda_i]P_i$ for P_i in $\ker \phi_i$, and $\lambda_1 + \lambda_2 \equiv t \pmod{\ell}$. We only need to compute one of the λ_i, since then the other is determined by the relation $\lambda_1\lambda_2 = q$.

So let ϕ be one of the two ℓ-isogenies; we want to compute its eigenvalue λ. The nonzero elements (x, y) of $\ker \phi$ satisfy $f_\phi(x) = 0$, where f_ϕ is a polynomial of degree $(\ell - 1)/2$ (if ℓ is odd; if $\ell = 2$, then $\deg f_\phi = 1$). To compute λ, we define the ring $A = \mathbb{F}_q[X, Y]/(f_\phi(X), Y^2 - X^3 - aX - b)$, set $P = (X, Y)$ in $\mathcal{E}(A)$, then compute $Q = \pi(P)$ and solve for λ in $Q = [\lambda]P$; then $t_\ell \equiv \lambda + q/\lambda \pmod{\ell}$.

This approach is substantially faster than Schoof's algorithm for Elkies ℓ, because the degree of f_ϕ is only $(\ell - 1)/2$, whereas the degree of Ψ_ℓ is $(\ell^2 - 1)/2$; so each operation in $\mathcal{E}(A)$ costs much less than it would if we used Ψ_ℓ instead of f_ϕ. (In practice, it is also nice to be able to reduce the number of costly Frobenius computations, since we only need to compute $\pi(P)$ and not $\pi(\pi(P))$.)

The crucial step is computing f_ϕ given only $\mathcal{E}$ and the corresponding root j_i of $\Phi_\ell(j(\mathcal{E}), X)$. We can do this using Elkies' algorithm, which is explained in [Sch95, §§7–8]. The total cost of computing t_ℓ is then $\widetilde{O}(\log^3 q)$ $\mathbb{F}_q$-operations: that is, a whole factor of $\log q$ faster compared to Schoof's algorithm.

Ideally, then, we should choose $\mathcal{L}$ to only contain Elkies and volcanic primes: that is, non-Atkin primes. The usual naive heuristic on prime classes is to suppose that as $q \to \infty$, the number of Atkin and non-Atkin primes less than B for $\mathcal{E}/\mathbb{F}_q$ is approximately equal when $B \sim \log q$; under this heuristic, taking $\mathcal{L}$ to contain only non-Atkin primes, the SEA algorithm computes t in $\widetilde{O}(\log^4 q)$ $\mathbb{F}_q$-operations.

While the heuristic holds on the average, assuming the GRH, Galbraith and Satoh have shown that it can fail for some curves [Sat02, Appendix A]: there exist curves $\mathcal{E}/\mathbb{F}_q$ such that if we try to compute t_ℓ using ℓ in the smallest possible set $\mathcal{L}$ containing only non-Atkin primes, then $\mathcal{L}$ must contain primes in $\Omega(\log^2 q)$.

Remark 4.4. It is important to note that Elkies' technique applies only to primes ℓ where there exists a rational ℓ-isogeny: that is, only Elkies and volcanic primes. Atkin's technique for restricting the possible values of t_ℓ applies to *all* primes—not only Atkin primes.

4.3 The genus 2 setting

Let $\mathcal{C}$ be a genus-2 curve defined over $\mathcal{F}_q$ (again, for q odd). We suppose that $\mathcal{C}$ is defined by an equation of the form $y^2 = f(x)$, where f is squarefree of degree 5.[3] The curve $\mathcal{C}$ then has a unique point at infinity, which we denote ∞.

4.3.1 The Jacobian

We write $J_\mathcal{C}$ for the Jacobian of $\mathcal{C}$. Our main algorithmic handle on $J_\mathcal{C}$ is Mumford's model for hyperelliptic Jacobians, which represents the projective $J_\mathcal{C}$ as a disjoint union of three affine subsets. In this model, points of $J_\mathcal{C}$ correspond to pairs of polynomials $\langle a(x), b(x)\rangle$ where a is monic, $\deg b < \deg a \leq 2$, and $b^2 \equiv f \pmod{a}$ (we call $\langle a, b\rangle$ the *Mumford representation* of the Jacobian point). Mumford's coordinates on the affine subsets of $J_\mathcal{C}$ are the coefficients of the polynomials a and b (and in particular, a point $\langle a, b\rangle$ of $J_\mathcal{C}$ is defined over $\mathbb{F}_q$ if and only if a and b have coefficients in $\mathbb{F}_q$). The three affine subsets are

$$\begin{aligned} W_2 &:= \{\langle a, b\rangle \in J_\mathcal{C} \mid \deg(a) = 2\} && (\text{“general” elements}) , \\ W_1 &:= \{\langle a, b\rangle \in J_\mathcal{C} \mid \deg(a) = 1\} && (\text{“special” elements}) , \\ W_0 &:= \{0_{J_\mathcal{C}} = \langle 1, 0\rangle\} && (\text{the trivial element}) , \end{aligned}$$

and $J_\mathcal{C} = W_2 \sqcup W_1 \sqcup W_0$. The group law on $J_\mathcal{C}$ can be explicitly computed on Mumford representatives using Cantor's algorithm [Can87].

The point of $J_\mathcal{C}$ corresponding to a general divisor class $[(x_P, y_P) + (x_Q, y_Q) - 2\infty]$ on $\mathcal{C}$ is represented by $\langle a, b\rangle$ where $a(x) = (x - x_P)(x - x_Q)$ and b is the linear polynomial such that $b(x_P) = y_P$ and $b(x_Q) = y_Q$. Special classes $[(x_P, y_P) - \infty]$ are represented by $\langle a, b\rangle = \langle x - x_P, y_P\rangle$, while $0_{J_\mathcal{C}} = [0]$ is represented by $\langle a, b\rangle = \langle 1, 0\rangle$.

4.3.2 Frobenius and endomorphisms of $J_\mathcal{C}$

The characteristic polynomial χ_π of the Frobenius endomorphism π has the form

$$\chi_\pi(X) = X^4 - tX^3 + (2q + s)X^2 - tqX + q^2 ,$$

where s and t are integers satisfying the inequalities (cf. [Rüc90])

$$|s| < 4q , \qquad |t| \leq 4\sqrt{q} , \qquad t^2 > 4s , \qquad s + 4q > 2|t|\sqrt{q} .$$

We have

$$\#J_\mathcal{C}(\mathcal{F}_q) = \chi_\pi(1) = 1 - t + 2q + s - tq + q^2 ,$$

as well as $\#\mathcal{C}(\mathcal{F}_q) = 1 - t + q$ and $\#\mathcal{C}(\mathcal{F}_{q^2}) = 1 - t^2 + 4q + 2s + q^2$. In genus 2, therefore, the point counting problem is to determine the integers s and t.

[3] For full generality, we should also allow $\deg f = 6$; the curve $\mathcal{C}$ then has two points at infinity. This substantially complicates the formulæ without significantly modifying the algorithms or their asymptotic complexity, so we will not treat this case here.

4.3.3 Real multiplication

We are interested in Jacobians $J_{\mathcal{C}}$ with real multiplication by a fixed order $\mathcal{O}$ in a quadratic real field $F := \mathbb{Q}(\sqrt{\Delta})$; that is, such that there is an embedding $\iota\colon \mathcal{O} \to \mathrm{End}(J_{\mathcal{C}})$. In this article, we will further restrict to the case where $\mathcal{O}$ is the maximal order $\mathcal{O}_F$ of F; note that if $\mathcal{O}$ is an order in F that is not locally maximal at a prime ℓ, then there exist no isogenies of degree ℓ that preserve the polarization (see Definition 4.5). These Jacobians can be constructed either from points in their moduli spaces (as in §4.4), or from a few known explicit families (as in §4.7).

The fixed field $\mathbb{Q}(\pi + \pi^\dagger)$ of the Rosati involution on $\mathbb{Q}(\pi)$ is a real quadratic field, and $\mathbb{Z}[\pi + \pi^\dagger]$ is a suborder of $\mathcal{O}_F$. The characteristic polynomial of $\pi + \pi^\dagger$ is

$$\chi_{\pi+\pi^\dagger}(X) = (X^2 - tX + s)^2 \ ,$$

so determining $\chi_{\pi+\pi^\dagger}$ also solves the point counting problem for $J_{\mathcal{C}}$.

Later, we will be particularly interested in $\mathcal{C}$ such that $J_{\mathcal{C}}$ has real multiplication by an order of small discriminant. While such curves are special, from a cryptographic perspective they are not "too special". From an arithmetic point of view, all curves (with ordinary simple Jacobians) over $\mathbb{F}_q$ have real multiplication. Here, we simply require that real multiplication to have small discriminant; the discriminant of the entire endomorphism ring of $J_{\mathcal{C}}$ can still be just as large as for a general choice of curve over the same field. From a geometric point of point view, the moduli of these $\mathcal{C}$ live on two-dimensional Humbert surfaces inside the three-dimensional moduli space of genus-2 curves. In concrete terms, this means that when selecting random curves over a fixed $\mathbb{F}_q$, only $\sim 1/q$ of them have real multiplication by a fixed order; but if we restrict our choice to those curves then there are still $O(q^2)$ of them to choose from.

4.3.4 From Schoof to Pila

The Schoof–Pila algorithm deals with higher dimensions [Sch85; Pil90]. Its input is a set of defining equations for a projective model of the abelian variety, and its group law. Jacobians of genus-2 curves are abelian varieties, and we can apply Pila's algorithm to them using the defining equations computed by Flynn [Fly90] or Grant [Gra90]. However, the complexity of Pila's algorithm is $O((\log q)^\Delta)$, where Δ (and the big-O constant) depends on the number of variables (i.e., the dimension of the ambient projective space) and the degree and number of the defining equations. Pila derives an upper bound for Δ in [Pil90, §4], but when we evaluate this bound in the parameters of Flynn's model for $J_{\mathcal{C}}$ (72 quadratic forms in 16 variables) we get a 30-bit Δ; Grant's model (13 quadratic and cubic forms in 9 variables) yields a 23-bit Δ.[4] While these are only upper bounds, we are clearly in the realm of the impractical here.

4.3.5 The Gaudry–Schost approach

Pila's algorithm requires a concrete (and necessarily complicated) nonsingular projective model for $J_{\mathcal{C}}$. The Gaudry–Schost algorithm applies essentially the same ideas to Mumford's affine models for subsets of $J_{\mathcal{C}}$.

Our first problem is to find an analogue for $J_{\mathcal{C}}$ of the elliptic division polynomials Ψ_ℓ. Ultimately, we want an ideal $I_\ell = (F_0, \ldots, F_r) \subset \mathbb{F}_q[A_1, A_0, B_1, B_0]$ such that $\langle a, b\rangle = \langle x^2 + a_1x + a_0, b_1x + b_0\rangle$ is in $J_{\mathcal{C}}[\ell]$ if and only if (a_1, a_0, b_1, b_0) is in the variety of I_ℓ: that is,

$$[\ell]\langle x^2 + a_1x + a_0, b_1x + b_0\rangle = 0 \Longleftrightarrow F(a_1, a_0, b_1, b_0) = 0 \text{ for all } F \in I_\ell \ .$$

Then, the image of $\langle x^2 + A_1x + A_0, B_1x + B_0\rangle$ in $J_{\mathcal{C}}(\mathbb{F}_q[A_1, A_0, B_1, B_0]/I_\ell)$ is an element of order ℓ that we can use for a Schoof-style computation of $\chi(T) \pmod{\ell}$.

The simplest approach here would be to take a general Mumford representative

$$\langle x^2 + A_1x + A_0, B_1x + B_0\rangle,$$

compute $L = [\ell]\langle x^2 + A_1x + A_0, B_1x + B_0\rangle$, and then equate coefficients in $L = 0_{J_{\mathcal{C}}}$ to derive the relations in I_ℓ. But we cannot do this, because L is in $W_2(\mathbb{F}_q(A_1, A_0, B_1, B_0))$ (that is, its a-polynomial has degree 2, and its b-polynomial degree 1), while $0_{J_{\mathcal{C}}} = \langle 1, 0\rangle$ is in W_0: these elements are not in the same affine subvariety, and cannot be directly compared or equated in this form.

[4] With polynomial time estimates like these, who needs enemies?

Gaudry and Harley [GH00] neatly stepped around this problem by observing that any element of $J_{\mathcal{C}}$ can be written as the difference of two elements of W_1 (which may be defined over a quadratic extension). They therefore start with $D = [(x_P, y_P) + (x_Q, y_Q) - 2\infty] = [(x_P, y_P) - (x_Q, -y_Q)]$ in $J_{\mathcal{C}}$, and find polynomial relations on x_P, y_P, x_Q, and y_Q such that $[\ell]D = 0$ by computing $[\ell]\langle x - x_P, y_P\rangle$ and $[\ell]\langle x - x_Q, -y_Q\rangle$, and equating coefficients in $[\ell]\langle x - x_P, y_P\rangle = [\ell]\langle x - x_Q, -y_Q\rangle$. There is a quadratic level of redundancy in these relations, which is a direct result of the redundancy in the initial representation of D: the involution $(x_P, y_P) \leftrightarrow (x_Q, y_Q)$ fixes D.

Gaudry and Schost remove this redundancy by resymmetrizing the relations with respect to this involution, re-expressing them in terms of $A_1 = -(x_P + x_Q)$, $A_0 = x_P x_Q$, $B_1 = (y_P - y_Q)/(x_P - x_Q)$, and $B_0 = (x_P y_Q - x_Q y_P)/(x_P - x_Q)$, and computing a triangular basis for the resulting *division ideal* I_ℓ. Their algorithm yields a triangular basis for I_ℓ, which facilitates fast reduction modulo I_ℓ.

Once we have I_ℓ, we can compute $t \pmod{\ell}$ and $s \pmod{\ell}$ as follows:

1. Construct the symbolic ℓ-torsion point
$$P := \langle x^2 + A_1 x + A_0, B_1 x + B_0\rangle \in J_{\mathcal{C}}(\mathbb{F}_q[A_1, A_0, B_1, B_0]/I_\ell) \ ;$$
2. Compute the points
$$\begin{aligned} Q_s &:= \pi^2(P) \ , \\ Q_t &:= \pi(\pi^2(P) + [q \bmod \ell]\pi(P)) \ , \\ R &:= \pi^4(P) + [2q \bmod \ell]\pi^2(P) + [q^2 \bmod \ell]P \end{aligned}$$
using Cantor arithmetic, with reduction of coefficients modulo I_ℓ;
3. Search for $0 \le s_\ell, t_\ell < \ell$ such that
$$[t_\ell]Q_t - [s_\ell]Q_s = R$$
(using, say, a two-dimensional baby-step giant-step algorithm).

The result is an algorithm that runs in time $\widetilde{O}(\log^8 q)$. Of course, once t has been determined, we can simplify Steps (2) and (3) above to find s_ℓ more quickly for the remaining ℓ, but this does not change the asymptotic complexity. In practice, the algorithm has been used to construct cryptographically secure curves: Gaudry and Schost computed a generic genus-2 curve over $\mathbb{F}_{2^{127}-1}$ such that both the Jacobian and its quadratic twist have prime order [GS12]. Instances of the discrete logarithm problem in this Jacobian offer a claimed security level of roughly 128 bits, which is the current minimum for serious cryptosystems. This computation also represents the current record for point counting for general genus-2 curves.

The Gaudry–Schost computation illustrates not only the state-of-the-art of genus-2 point counting, but also the practical challenge involved in producing cryptographically strong genus-2 Jacobians. The Schoof-like point counting algorithm was only applied using the prime powers 2^{17}, 3^9, 5^4, and 7^2, and the primes 11 through 31. Combining the information given by these prime powers completely determines t, but not s; but it still gives us enough modular information about s to be able to recover its precise value using Pollard's kangaroo algorithm in a reasonable time (≈ 2 hours, in this case). The kangaroo algorithm is exponential, and would not be practical for computing this Jacobian order alone without the congruence data generated by the Schoof-like computations. Gaudry and Schost estimated the average cost of these calculations as one core-month (in 2008) per curve.

4.3.6 Point counting with efficiently computable RM

In [GKS11], Gaudry, Kohel, and Smith described a number of improvements to the Gaudry–Schost algorithm that apply when $J_{\mathcal{C}}$ is equipped with an explicit and efficiently computable endomorphism ϕ generating a real quadratic subring of $\mathrm{End}(J_{\mathcal{C}})$. When we say that ϕ is *explicit* we mean that we can compute the images under ϕ of divisor classes on $J_{\mathcal{C}}$, including symbolic Mumford representatives for generic divisor classes. When we say that ϕ is *efficiently computable*, we mean that these images can be computed for a cost comparable with a few group operations: that is, from an algorithmic point of view, we may view evaluation of ϕ as an elementary group operation like adding or doubling.

Suppose that $\mathbb{Z}[\pi + \pi^\dagger]$ is contained in $\mathbb{Z}[\phi]$ (this is reasonable, since in the examples we know, $\mathbb{Z}[\phi]$ is a maximal order), and let Δ be the discriminant of $\mathbb{Z}[\phi]$. Then $\pi + \pi^\dagger = m\phi + n$ for some m and n, which

completely determine s and t: if the characteristic polynomial of ϕ is $(X^2 - t_\phi X + s_\phi)^2$, then $t = 2m + nt_\phi$ and $s = (t^2 - s_\phi^2 \Delta)/4$. It follows that m and n are both in $O(\sqrt{q})$.

We can compute m and n using a technique similar to Gaudry–Schost. Multiplying the relation $\pi + \pi^\dagger = m\phi + n$ through by π, we have $\pi^2 - (m\phi + n)\pi + q = 0$. Imitating Schoof's algorithm, we can compute $m_\ell := m \pmod{\ell}$ and $n_\ell := n \pmod{\ell}$ by taking a generic element D of $J_{\mathcal{C}}[\ell]$ (as in Gaudry–Schost), computing $(\pi^2 + q)(D)$, $\pi(D)$, and $\phi\pi(D)$ (using two applications of π), and then solving for m_ℓ and n_ℓ.

We can do even better by exploiting split primes in $\mathbb{Z}[\phi]$. If $\ell = \mathfrak{l}_1\mathfrak{l}_2$ is split, then the ℓ-torsion decomposes as $J_{\mathcal{C}}[\mathfrak{l}_1] \oplus J_{\mathcal{C}}[\mathfrak{l}_2]$, and once we have found a short generator (or generators) for $\mathfrak{l}_i$ we can take D to be an element of $J_{\mathcal{C}}[\mathfrak{l}_i]$ instead of $J_{\mathcal{C}}[\ell]$. Such generators can be found with coefficients in $O(\sqrt{\ell})$; the result is that we work modulo a much smaller ideal, of degree $O(\ell^2)$ rather than $O(\ell^4)$.

But going further, $\pi + \pi^\dagger$ acts as a scalar on $J_{\mathcal{C}}[\mathfrak{l}_i]$, and so we can compute its eigenvalue to determine m_ℓ and n_ℓ. The total cost of computing m_ℓ and n_ℓ, and hence t_ℓ and s_ℓ, is then $\widetilde{O}(\log^5 q)$ [GKS11, Theorem 1], a substantial improvement on Gaudry–Schost's $\widetilde{O}(\log^8 q)$.

The computation resembles what we would do for an Elkies prime in the elliptic case, except that there is no need for modular polynomials to compute the prime type, or for an analogue of Elkies' algorithm: we know in advance which primes split in $\mathbb{Z}[\phi]$, and we can compute the kernel using the decomposition. But if we did have an analogue of Elkies' algorithm, then we could further reduce the complexity by further decomposing some of the $J_{\mathcal{C}}[\mathfrak{l}_i]$ into cyclic factors, and thus working modulo ideals of degree $O(\ell)$. If we have an analogue of Atkin's algorithm, then we can restrict the possible values of m_ℓ and n_ℓ; this would not change the asymptotic complexity of the algorithm, but it could have a significant practical impact.

4.3.7 Generalizing Elkies' and Atkin's improvements to genus 2

Ultimately, we would like to generalize the SEA algorithm to genus 2. The first requirement is a genus-2 analogue of elliptic modular polynomials; so assume for the moment that we have a modular ideal relating suitable invariants of genus-2 curves.

To generalize Elkies' improvements to genus 2, we need an analogue of Elkies' algorithm: that is, an algorithm which, given two general moduli points corresponding to isogenous Jacobians, constructs defining polynomials for (the kernel of) the isogeny. The most convenient such presentation would be as an ideal cutting out the intersection of the kernel with W_2, since then the Gaudry–Schost approach could be adapted without too much difficulty (at least in theory). Unfortunately, at present, no such algorithm is known.

In contrast, Atkin's techniques for elliptic curves require only the factorization of (specializations of) elliptic modular polynomials; we deduce possible congruences on the trace from the degrees of the factors. It is clear how we should generalize Atkin's techniques to genus 2: we should deduce possible congruences on s and t from the degrees of primary components of specialized modular ideals.

The following sections make this concrete. In §4.4, we define the appropriate analogues of the elliptic j-invariant for genus-2 curves with real multiplication. We can then define real-multiplication analogues of the elliptic modular polynomials in §4.4.2, before investigating their factorization in §4.5.

4.3.8 μ-isogenies

Before defining any generalized invariants or modular polynomials, we must define an appropriate class of isogenies in genus 2: that is, isogenies that are compatible with the real multiplication structure. (This is not an issue for elliptic curves, because the elliptic analogue of the real endomorphism subring is just $\mathbb{Z}$—and everything is compatible with integer multiplications.)

Definition 4.5. Let $(\mathcal{A}, \xi, \iota)$ and $(\mathcal{A}', \xi', \iota')$ be triples encoding principally polarized abelian surfaces with real multiplication by $\mathcal{O}_F$. Here $\xi\colon \mathcal{A} \to \mathcal{A}^\vee$ and $\xi'\colon \mathcal{A}' \to (\mathcal{A}')^\vee$ are principal polarizations, and $\iota\colon \mathcal{O}_F \hookrightarrow \mathrm{End}(\mathcal{A})$ and $\iota'\colon \mathcal{O}_F \hookrightarrow \mathrm{End}(\mathcal{A}')$ are embeddings that are stable under the Rosati involution. If μ is a totally positive element of F, then a *μ-isogeny* $(\mathcal{A}, \xi, \iota) \to (\mathcal{A}', \xi', \iota')$ is an isogeny $f\colon \mathcal{A} \to \mathcal{A}'$ such

that the diagrams

$$\begin{tikzcd} \mathcal{A} \arrow[r, "\iota(\mu)"] \arrow[d, "f"'] & \mathcal{A} \arrow[r, "\xi"] & \mathcal{A}^\vee \\ \mathcal{A}' \arrow[rr, "\xi'"'] & & (\mathcal{A}')^\vee \arrow[u, "f^\vee"'] \end{tikzcd} \qquad \text{and} \qquad \begin{tikzcd} F \arrow[r, "\iota"] \arrow[dr, "\iota'"'] & \mathrm{End}(\mathcal{A}) \otimes \mathbb{Q} \arrow[d, "\phi"] \\ & \mathrm{End}(\mathcal{A}') \otimes \mathbb{Q} \end{tikzcd}$$

commute, where ϕ is the map induced by f on endomorphism algebras.

If $f\colon (\mathcal{A}, \xi, \iota) \to (\mathcal{A}', \xi', \iota')$ is a μ-isogeny, then the polarization ξ' pulls back via f to $\xi \circ \iota(\mu)$. For comparison, an elliptic ℓ-isogeny is an $f\colon \mathcal{E} \to \mathcal{E}'$ such that the canonical polarization on $\mathcal{E}'$ pulls back via f to ℓ times the polarization on $\mathcal{E}$ (in more concrete terms: the identity point $0_{\mathcal{E}'}$ on $\mathcal{E}$ pulls back via f to a divisor on $\mathcal{E}$ equivalent to $\ell \cdot 0_{\mathcal{E}}$).

4.4 Invariants

Elliptic modular polynomials relate isogenous elliptic curves in terms of their j-invariants; their genus-2 analogues must relate invariants of genus-2 Jacobians. This section describes and relates the various invariants that we will need. Since we are dealing with classical constructions in this section, we work over a field $k \subseteq \mathbb{C}$. However, the resulting algebraic expressions carry over to the case where $k = \mathbb{F}_q$ (at least for large enough p). All of the results in this section are well-known, and are shown here for completeness and easy reference; we refer the reader to [Lan82], [LNY16], [LY11], and [Mar18] for further detail.

4.4.1 Invariants for RM abelian surfaces

Let F be a real quadratic field with ring of integers $\mathcal{O}_F$. We need RM analogues of the elliptic j-invariant and elliptic modular polynomials for μ-isogenies of abelian surfaces with RM by $\mathcal{O}_F$. Our first step is to define appropriate replacements for the j-invariant that classify our triples (A, ξ, ι) up to isomorphism. Instead of a single j-invariant, we will have a triple (J_1, J_2, J_3) of *RM invariants*, which are functions on the corresponding Hilbert modular surface.

The invariants (J_1, J_2, J_3) are constructed as follows. For a field k, we consider the coarse moduli space $\mathcal{H}_F(k)$ of triples $(\mathcal{A}, \xi, \iota)$ (where as before, $\mathcal{A}/k$ is an abelian variety with a principal polarization $\xi\colon \mathcal{A} \to \mathcal{A}^\vee$ and an embedding $\iota\colon \mathcal{O}_F \hookrightarrow \mathrm{End}_k(\mathcal{A})$ stable under the Rosati involution). Then $\mathcal{H}_F(k)$ is coarsely represented by the Hilbert modular space $\mathrm{SL}_2(\mathcal{O}_F \oplus \mathcal{O}_F) \setminus (F \otimes \mathbb{H})$ (see [Gee88]), where $F \otimes \mathbb{H} := \{\tau \in F \otimes \mathbb{C} : \Im(\tau) > 0\}$ and for any fractional ideal $\mathfrak{f}$ of F,

$$\mathrm{SL}_2(\mathcal{O}_F \oplus \mathfrak{f}) := \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(F) : a, d \in \mathcal{O}_F,\ b \in \mathfrak{f},\ c \in \mathfrak{f}^{-1} \right\}$$

acts on $F \otimes \mathbb{H}$ by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot \tau = \frac{a\tau + b}{c\tau + d} \ .$$

Proposition 4.6. *Let V be the Baily–Borel compactification of $SL_2(\mathcal{O}_F) \setminus (F \otimes \mathbb{H})$, and $\mathbb{C}(V)$ the function field of V. There exist rational functions J_1, J_2, and J_3 on V such that*

$$\mathbb{C}(V) = \mathbb{C}(J_1, J_2, J_3) \ .$$

Proof. The transcendence degree of $\mathbb{C}(V)$ over $\mathbb{C}$ is 2, so there exist 2 algebraically independent functions J_1, J_2 in $\mathbb{C}(V)$. Furthermore, $\mathbb{C}(V)$ is a finite separable field extension of $\mathbb{C}(J_1, J_2)$, so it is generated by at most one further element, J_3. □

Definition 4.7. Fixing a choice of rational functions J_1, J_2, and J_3 as in Proposition 4.6, we call (J_1, J_2, J_3) the *RM invariants* for F.

4.4.2 Hilbert modular polynomials for RM abelian surfaces

We are now ready to define modular polynomials for abelian surfaces with RM structure. For elliptic curves we have a single j-invariant, and we can relate ℓ-isogenous j-invariants using a single bivariate polynomial $\Phi_\ell(X,Y)$. For our abelian surfaces, we have a tuple of three invariants (J_1,J_2,J_3), and to relate μ-isogenous tuples of invariants we need a *modular ideal* of polynomials in $\mathbb{Q}[X_1,X_2,X_3,Y_1,Y_2,Y_3]$, such that when we specialize the first three variables in the (J_1,J_2,J_3) corresponding to the isomorphism class of some triple (A,ξ,ι), the result is an ideal cutting out the moduli points (J_1',J_2',J_3') for triples (A',ξ',ι') that are μ-isogenous to (A,ξ,ι).

The *Hilbert modular polynomials* below represent a particularly convenient basis for this ideal. We refer the reader to [Mar18, Chapter 2] for theoretical details and proofs, as well as algorithms for computing the polynomials. Alternatively, Milio's algorithm can be used to compute Hilbert modular polynomials $\Phi_\ell(X,\mathfrak{J}_1,\mathfrak{J}_2)$ and $\Psi_\ell(X,\mathfrak{J}_1,\mathfrak{J}_2)$, in time $O(d_T d_{\mathfrak{J}_2})\tilde{O}(\ell N)+4(\ell+1)\tilde{O}(d_T d_{\mathfrak{J}_2} N)\subseteq \tilde{O}(d_T d_{\mathfrak{J}_2}\ell N)$ [Mil15b, Theorem 5.4.4], where N is the precision and $d_T, d_{\mathfrak{J}_2}$ are degrees involved in the computation, see [Mil15b, §5.4].

Definition 4.8. The *Hilbert modular polynomials*

$$\begin{aligned}
&G_\mu(X_1,X_2,X_3,Y_1)\ ,\\
&H_{\mu,2}(X_1,X_2,X_3,Y_1,Y_2)=H^{(1)}_{\mu,2}(X_1,X_2,X_3,Y_1)Y_2+H^{(0)}_{\mu,2}(X_1,X_2,X_3,Y_1)\ ,\\
&H_{\mu,3}(X_1,X_2,X_3,Y_1,Y_3)=H^{(1)}_{\mu,3}(X_1,X_2,X_3,Y_1)Y_3+H^{(0)}_{\mu,3}(X_1,X_2,X_3,Y_1)
\end{aligned}$$

in $\mathbb{Q}[X_1,X_2,X_3,Y_1,Y_2,Y_3]$ are defined such that for all triples $(\mathcal{A},\xi,\iota)$ and $(\mathcal{A}',\xi',\iota')$ representing points τ and τ' in a certain Zariski-open subset[5] of the Baily–Borel compactification of $\mathrm{SL}_2(\mathcal{O}_F\oplus\mathfrak{f})\setminus(F\otimes\mathbb{H})$, there exists a μ-isogeny $f\colon(\mathcal{A},\xi,\iota)\to(\mathcal{A}',\xi',\iota')$ if and only if

$$\begin{aligned}
G_\mu(J_1(\tau),J_2(\tau),J_3(\tau),J_1(\tau'))&=0\ ,\\
H_{\mu,2}(J_1(\tau),J_2(\tau),J_3(\tau),J_1(\tau'),J_2(\tau'))&=0\ ,\\
H_{\mu,3}(J_1(\tau),J_2(\tau),J_3(\tau),J_1(\tau'),J_3(\tau'))&=0\ .
\end{aligned}$$

The special form of G_μ, $H_{2,\mu}$, and $H_{3,\mu}$ are very convenient for computations. If (J_1,J_2,J_3) is a fixed moduli point, then each root α of $G(J_1,J_2,J_3,x)$ corresponds to a unique μ-isogenous moduli point

$$(J_1',J_2',J_3')=\left(\alpha,-\frac{H^{(0)}_{\mu,2}(J_1,J_2,J_3,\alpha)}{H^{(1)}_{\mu,2}(J_1,J_2,J_3,\alpha)},-\frac{H^{(0)}_{\mu,3}(J_1,J_2,J_3,\alpha)}{H^{(1)}_{\mu,3}(J_1,J_2,J_3,\alpha)}\right)\ .$$

We observe that the action of Galois on the set of μ-isogenies from an RM abelian variety representing (J_1,J_2,J_3) is completely described by the action of Galois on the roots of $G_\mu(J_1,J_2,J_3,x)$; in particular, over $\mathbb{F}_q$, rational cycles of μ-isogenies under Frobenius correspond to irreducible factors of $G_\mu(J_1,J_2,J_3,x)$. From the point of view of Atkin generalizations, therefore, we only really need G_μ to replace Φ_ℓ.

4.4.3 Invariants for curves and abelian surfaces

We need to relate the RM invariants (J_1,J_2,J_3) to the invariants for plain old principally polarized abelian surfaces, and in particular Jacobians of genus 2 curves without any special RM structure. The moduli space $\mathcal{A}_2$ of principally polarized abelian surfaces is coarsely represented by the Siegel modular space $\mathrm{Sp}_2(\mathbb{Z})\backslash\mathbb{H}_2$, where

$$\mathbb{H}_2:=\left\{\tau=\begin{pmatrix}\tau_1&\tau_2\\\tau_2&\tau_3\end{pmatrix}\in\mathrm{Sym}_2(\mathbb{C}):\Im(\tau)>0\right\}\ ,$$

and the symplectic group

$$\mathrm{Sp}_2(\mathbb{Z})=\left\{g\in\mathrm{GL}_4(\mathbb{Z}):g\begin{pmatrix}0&I_2\\-I_2&0\end{pmatrix}g^t=\begin{pmatrix}0&I_2\\-I_2&0\end{pmatrix}\right\}$$

[5] See [Mar18, Chapter 2, Section 2] for details on this subset. For point counting over large finite fields, it is enough to note that since the subset is Zariski open, randomly sampled Jacobians with real multiplication by $\mathcal{O}_F$ have their RM invariants in this subset with overwhelming probability.

acts on $\mathbb{H}_2$ via

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot \tau = \frac{a\tau + b}{c\tau + d} \ .$$

Every rational function on $\mathrm{Sp}_2(\mathbb{Z})\backslash\mathbb{H}_2$ is a quotient of elements of the graded ring of holomorphic Siegel modular forms for $\mathrm{Sp}_2(\mathbb{Z})$. Igusa proved in [Igu60] that this ring is generated by ψ_4, ψ_6, χ_{10}, and χ_{12}, where

$$\psi_k(\tau) = \sum_{\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in P\backslash \mathrm{Sp}_2(\mathbb{Z})} \det(c\tau + d)^{-k}$$

is the normalized Eisenstein series of weight k for even integers $k \geq 4$ (here P is the standard Siegel parabolic subgroup of $\mathrm{Sp}_2(\mathbb{Z})$), and

$$\begin{aligned} \chi_{10} &= -2^{-12} \cdot 3^{-5} \cdot 5^{-2} \cdot 7^{-1} \cdot 43867(\psi_4\psi_6 - \psi_{10}) \ , \\ \chi_{12} &= 2^{-13} \cdot 3^{-7} \cdot 5^{-3} \cdot 7^{-2} \cdot 337^{-1} \cdot 131 \cdot 593(3^2 \cdot 7^2\psi_4^3 + 2 \cdot 5^3\psi_6^2 - 691\psi_{12}) \end{aligned}$$

are Siegel modular cusp forms of weight 10 and 12 respectively.

Curves of genus 2 are typically classified up to isomorphism by their Igusa invariants (j_1, j_2, j_3), or by their Igusa–Clebsch invariants (A, B, C, D). Since the map $\mathcal{C} \mapsto J_{\mathcal{C}}$ is an open immersion of the (coarse) moduli space of genus-2 curves $\mathcal{M}_2$ into $\mathcal{A}_2$, the Igusa invariants j_i can be written as rational functions of ψ_4, ψ_6, χ_{10} and χ_{12} as follows [Igu67]:

$$\begin{aligned} j_1(\tau) &= 2 \cdot 3^5 \cdot \chi_{12}^5\chi_{10}^{-6} \ , \\ j_2(\tau) &= 2^{-3} \cdot 3^3 \cdot \psi_4\chi_{12}^3\chi_{10}^{-4} \ , \\ j_3(\tau) &= 2^{-5} \cdot 3 \cdot \left(\psi_6\chi_{12}^2\chi_{10}^{-3} + 2^2 \cdot 3 \cdot \psi_4\chi_{12}^3\chi_{10}^{-4}\right) \ . \end{aligned}$$

Here $j_i(\tau) = j_i(\mathcal{C})$ if there is a genus 2 curve $\mathcal{C}/\mathbb{C}$ such that $J_{\mathcal{C}}$ is isomorphic to the abelian surface $\mathbb{C}^2/(\mathbb{Z}^2\tau + \mathbb{Z}^2)$. If there is no such $\mathcal{C}$, which happens exactly when $\chi_{10}(\tau) = 0$, then $j_i(\tau)$ is not well-defined. The Igusa–Clebsch invariants are related to the Siegel modular forms by

$$(\psi_4,\ \psi_6,\ \chi_{10},\ \chi_{12}) = \left(2^{-2}B,\ 2^{-3}(AB - 3C),\ -2^{-14}C,\ 2^{-17}3^{-1}AD\right) \ . \tag{4.5}$$

4.4.4 Pulling back curve invariants to RM invariants

The natural maps $\mathbb{H}^2 \to \mathbb{H}_2$, $\mathrm{SL}_2(F) \to \mathrm{Sp}_2(\mathbb{Q})$, and $(\mathcal{O}_F/2\mathcal{O}_F)^2 \to (\mathbb{Z}/2\mathbb{Z})^4$ induce an embedding

$$\phi\colon \mathcal{H}_F(k) \hookrightarrow \mathcal{A}_2(k) \ ,$$

which we can use to pull back Igusa invariants to RM invariants, thus expressing the j_i in terms of the J_i. We will see detailed formulæ for this pullback for $F = \mathbb{Q}(\sqrt{5})$ in Proposition 4.17.

This pullback from curves and their invariants to RM invariants is essential for our computations: after all, in point counting one usually starts from a curve. In our applications, we are given the equation of a curve $\mathcal{C}/\mathbb{F}_q$ drawn from a family of curves with known RM by $\mathcal{O}_F$. Having computed the Igusa or Igusa–Clebsch invariants of $\mathcal{C}$, we can pull them back to RM invariants (J_1, J_2, J_3). This pullback is possible, because $\mathcal{C}$ was chosen from an appropriate family, but choosing a preimage (J_1, J_2, J_3) implicitly involves choosing one of the two embeddings of $\mathcal{O}_F$ into $\mathrm{End}(J_{\mathcal{C}})$. This choice cannot always be made over the ground field: a point in $\mathcal{A}_2(k)$ may not pull back to a pair of points in $\mathcal{H}_F(k)$, but rather a conjugate pair of points over a quadratic extension of k. Proposition 4.18 makes this subtlety explicit in the case $F = \mathbb{Q}(\sqrt{5})$.

4.5 Atkin theorems in genus 2

We are now ready to state some Atkin-style results for μ-isogenies in genus 2.

Let $(\mathcal{A}, \xi, \iota)$ be a triple describing a vanilla abelian surface over $\mathbb{F}_q$ with real multiplication by $\mathcal{O}_F$, and let μ be a totally positive element of $\mathcal{O}_F$ of norm ℓ. Then $\iota(\mu)$ is an endomorphism of degree ℓ^2, and we have a subgroup[6]

$$\mathcal{A}[\mu] := \ker(\iota(\mu)) \subset \mathcal{A}[\ell] \ .$$

[6] We emphasize that the subgroup $\mathcal{A}[\mu]$ depends on ι, but we have chosen to write $\mathcal{A}[\mu]$ instead of the more cumbersome $\mathcal{A}[\iota(\mu)]$.

If $(\bar{\mu}) \neq (\mu)$ (that is, $(\ell) \neq (\mu^2)$), then we have a decomposition $\mathcal{A}[\ell] = \mathcal{A}[\mu] \oplus \mathcal{A}[\bar{\mu}]$. The one-dimensional subspaces of $\mathcal{A}[\mu]$ are the kernels of μ-isogenies.

In §4.2 we used the elliptic modular polynomial Φ_ℓ to study the structure of $\mathcal{E}[\ell]$. Here, we will use the Hilbert modular polynomial G_μ to study the structure of $\mathcal{A}[\mu]$. The propositions of this section are generalizations for curves of genus 2 to Schoof's Propositions 6.1, 6.2 and 6.3 for elliptic curves in [Sch95].

4.5.1 Roots of G_μ and the order of Frobenius

Our first result relates the order of Frobenius acting on $\mathbb{P}(\mathcal{A}[\mu])$ to the extensions of $\mathbb{F}_q$ generated by roots of specialized Hilbert modular polynomials.

Proposition 4.9. [7] *Let $\mathcal{A}/\mathbb{F}_q$ be a vanilla abelian surface with RM by $\mathcal{O}_F$ and RM invariants (J_1, J_2, J_3) in $\mathbb{F}_q^3$, and with Frobenius endomorphism π. Let μ be a totally positive element of $\mathcal{O}_F$ of prime norm $\ell = \mu\bar{\mu}$.*

1. *The polynomial $G_\mu(J_1, J_2, J_3, x)$ has a zero $\tilde{J}_1$ in $\mathbb{F}_{q^e}$ if and only if the kernel of the corresponding μ-isogeny $\mathcal{A} \to \tilde{\mathcal{A}}$ is a 1-dimensional eigenspace of π^e in $\mathcal{A}[\mu]$.*

2. *The polynomial $G_\mu(J_1, J_2, J_3, x)$ splits completely in $\mathbb{F}_{q^e}[x]$ if and only if π^e acts as a scalar matrix on $\mathcal{A}[\mu]$.*

Proof. The proof follows that of [Sch95, Proposition 6.1] (stated as Lemma 4.3 here).

For (1): Let $f\colon \mathcal{A} \to \tilde{\mathcal{A}}$ be a μ-isogeny with kernel S, and let $(\tilde{J}_1, \tilde{J}_2, \tilde{J}_3)$ be the RM invariants of $\tilde{\mathcal{A}}$. If S is an eigenspace of π^e, then the quotient $\mathcal{A} \to \mathcal{A}/S$ is defined over $\mathbb{F}_{q^e}$. The Igusa invariants of $\mathcal{A}/S$ are therefore all in $\mathbb{F}_{q^e}$, and since $\mathcal{A}/S$ is isomorphic to $\tilde{\mathcal{A}}$ as a principally polarized abelian surface, the Igusa invariants of $\tilde{\mathcal{A}}$ are all in $\mathbb{F}_{q^e}$. To conclude that $\tilde{J}_1$ is in $\mathbb{F}_{q^e}$, we need to show that the injection $\tilde{\iota}\colon \mathcal{O}_F \hookrightarrow \mathrm{End}(\tilde{\mathcal{A}})$ is defined over $\mathbb{F}_{q^e}$; but this follows from the commutativity of the second diagram in Definition 4.5.

Conversely: suppose $G_\mu(J_1, J_2, J_3, \tilde{J}_1) = 0$ for some $\tilde{J}_1$ in $\mathbb{F}_{q^e}$. Then the fact that each of the $H_{\mu,i}$ is a linear polynomial in Y_i with coefficients in $\mathbb{F}_q[J_1, J_2, J_3, \tilde{J}_1] = \mathbb{F}_{q^e}$ shows that there exist $\tilde{J}_2$ and $\tilde{J}_3$ in $\mathbb{F}_{q^e}$ such that $(\tilde{J}_1, \tilde{J}_2, \tilde{J}_3)$ are the RM invariants of a triple $(\tilde{\mathcal{A}}, \tilde{\xi}, \tilde{\iota})$ that is μ-isogenous to $(\mathcal{A}, \xi, \iota)$. This means that there is an $\overline{\mathbb{F}}_q$-isomorphism $(\tilde{\mathcal{A}}, \tilde{\xi}, \tilde{\iota}) \to (\mathcal{A}', \xi', \iota')$ where $(\mathcal{A}', \xi', \iota')$ is defined over $\mathbb{F}_{q^e}$. Let $f\colon \mathcal{A} \to \mathcal{A}'$ be the composite μ-isogeny. Its kernel S is a one-dimensional subspace of $\mathcal{A}[\ell]$. It remains to show that S is an eigenspace of π^e; this is the case if and only if f is defined over $\mathbb{F}_{q^e}$. The $\mathbb{Z}$-module $\mathrm{Hom}_{\overline{\mathbb{F}}_q}(\mathcal{A}, \mathcal{A}')$ is free of rank 4 (because $\mathcal{A}$ is vanilla); and its submodule $\mathrm{Hom}_{\mathbb{F}_{q^e}}(\mathcal{A}, \mathcal{A}')$ of $\mathbb{F}_{q^e}$-isogenies is either 0 or equal to $\mathrm{Hom}_{\overline{\mathbb{F}}_q}(\mathcal{A}, \mathcal{A}')$. Hence, f is defined over $\mathbb{F}_{q^e}$ if $\mathrm{Hom}_{\mathbb{F}_{q^e}}(\mathcal{A}, \mathcal{A}') \neq 0$; and $\mathrm{Hom}_{\mathbb{F}_{q^e}}(\mathcal{A}, \mathcal{A}') \neq 0$ if and only if the Frobenius endomorphisms of $\mathcal{A}/\mathbb{F}_{q^e}$ and $\mathcal{A}'$ have the same characteristic polynomial.

Since $\mathcal{A}$ is vanilla, and $\mathcal{A}'$ is $\overline{\mathbb{F}}_q$-isogenous to $\mathcal{A}$, we have $\mathrm{End}_{\overline{\mathbb{F}}_q}(\mathcal{A}') \otimes \mathbb{Q} \cong \mathrm{End}_{\overline{\mathbb{F}}_q}(\mathcal{A}) \otimes \mathbb{Q} \cong K$ for some quartic CM-field K. So let ψ and ψ' be the images in K of the Frobenius endomorphisms of $\mathcal{A}/\mathbb{F}_{q^e}$ and $\mathcal{A}'$, respectively (note that $\psi = \pi^e$). Now up to complex conjugation, we have $\psi^s = (\psi')^s$ in K for some $s > 0$. If $\psi = \psi'$, then $\mathcal{A}$ and $\mathcal{A}'$ are $\mathbb{F}_{q^e}$-isogenous, and we are done. If $\psi = -\psi'$, then we replace $(\mathcal{A}', \xi', \iota')$ by its quadratic twist; and then $\mathcal{A}$ and $\mathcal{A}'$ are $\mathbb{F}_{q^e}$-isogenous. Otherwise, if $\psi \neq \pm\psi'$, then ψ/ψ' must be a root of unity of order at least 3 in K, which is impossible because $\mathcal{A}$ is vanilla. Hence $\psi = \psi'$, so ψ and ψ' have the same characteristic polynomial, and therefore f is defined over $\mathbb{F}_{q^e}$.

For (2): If all of the zeroes of $G_\mu(J_1, J_2, J_3, x)$ are contained in $\mathbb{F}_{q^e}$, then all of the 1-dimensional subspaces of $\mathcal{A}[\mu]$ are eigenspaces of π^e by Part (1). This implies that π^e acts as a scalar matrix on $\mathcal{A}[\mu]$. □

Remark 4.10. As an example of what can go wrong if the vanilla condition is dropped, consider the curve

$$\mathcal{C} : y^2 = x^5 + 1 \ .$$

The Jacobian $J_\mathcal{C}$ of this curve has complex multiplication by $\mathbb{Q}(\zeta_5)$, so it is not vanilla. While $J_\mathcal{C}$ has real multiplication by the maximal order of $\mathbb{Q}(\sqrt{5})$, the Siegel modular form ψ_4 is zero for this curve. Proposition 4.18 below gives explicit formulæ for J_1, J_2, and J_3^2 for Jacobians with maximal real multiplication by $\mathbb{Q}(\sqrt{5})$; and when we look at those formulæ, we see that J_1 is not well-defined when $\psi_4 = 0$.

[7] This is conditional under the heuristics of Remark 2.5.6

4.5.2 The factorization of G_μ

The Frobenius endomorphism π of $\mathcal{A}$ commutes with $\iota(\mu)$ (since $\mathcal{A}$ is vanilla), so it restricts to an endomorphism of $\mathcal{A}[\mu]$.

Lemma 4.11. *Let $\mathcal{A}/\mathbb{F}_q$ be a vanilla abelian surface with Frobenius endomorphism π, and let ℓ be an odd prime.*

1. *If ℓ splits in $\mathbb{Z}[\pi+\pi^\dagger]$ (or equivalently, if t^2-4s is a square in $\mathbb{F}_\ell$), then*

$$\chi_\pi(T) \equiv (T^2-uT+q)(T^2-u'T+q) \pmod{\ell}$$

for some u and u' in $\mathbb{Z}/\ell\mathbb{Z}$.

2. *If ℓ is ramified in $\mathbb{Z}[\pi+\pi^\dagger]$ (or equivalently, if ℓ divides t^2-4s), then*

$$\chi_\pi(T) \equiv (T^2-uT+q)^2 \pmod{\ell}$$

where $u = t/2$ in $\mathbb{Z}/\ell\mathbb{Z}$.

3. *If ℓ is inert in $\mathbb{Z}[\pi+\pi^\dagger]$ (or equivalently, if t^2-4s is a square in $\mathbb{F}_\ell$), then*

$$\chi_\pi(T) \not\equiv (T^2-uT+q)(T^2-u'T+q) \pmod{\ell}$$

for any $u, u' \in \mathbb{Z}/\ell\mathbb{Z}$.

Proof. This is a direct consequence of [Lan86, Chapter 1, Proposition 25]. □

Lemma 4.12. *Let $(\mathcal{A},\xi,\iota)$ be a triple describing a vanilla abelian surface over $\mathbb{F}_q$ with real multiplication by $\mathcal{O}_F$, and let μ be a totally positive element of $\mathcal{O}_F$ of prime norm $\mu\overline{\mu}=\ell$. The restriction of the Frobenius endomorphism π to $\mathcal{A}[\mu]$ has characteristic polynomial*

$$\chi_{\pi,\mu}(T) \equiv T^2-uT+q \pmod{\ell} \quad \textit{for some } u \in \mathbb{Z}/\ell\mathbb{Z}\ .$$

Proof. By definition, $\ell = \mu\bar{\mu}$ splits in $\mathcal{O}_F$, so it either splits or ramifies in the suborder

$$\mathbb{Z}[\pi+\pi^\dagger] \subseteq \mathcal{O}_F;$$

we are therefore in Case (1) or (2) of Lemma 4.11. In particular, both π and $\pi^\dagger$ restrict to endomorphisms of $\mathcal{A}[\mu]$, and they have the same eigenvalues λ and q/λ; so the characteristic polynomial of π is

$$T^2-(\lambda+q/\lambda)T+q.$$

The result follows with $u = \lambda + q/\lambda$. □

Proposition 4.13 uses the factorization of the modular polynomial G_μ, specialized at the RM invariants of $\mathcal{A}$, to derive information $\chi_{\pi,\mu}(T) \pmod{\ell}$.

Proposition 4.13. [8] *Let $(\mathcal{A},\xi,\iota)$ be a triple describing a vanilla abelian surface over $\mathbb{F}_q$ with real multiplication by $\mathcal{O}_F$ and with RM invariants (J_1,J_2,J_3), and let μ be a totally positive element of $\mathcal{O}_F$ of prime norm $\mu\overline{\mu}=\ell$. Let π be the Frobenius endomorphism of $\mathcal{A}$, with $\chi_{\pi,\mu}(T) = T^2-uT+q$ the characteristic polynomial of the restriction of π to $\mathcal{A}[\mu]$, and let e be the order of π in $\mathrm{Aut}(\mathbb{P}(\mathcal{A}[\mu])) \cong \mathrm{PGL}_2(\mathbb{F}_\ell)$.*

The polynomial $G_\mu(J_1,J_2,J_3,x)$ has degree $\ell+1$ in $\mathbb{F}_q[x]$, and its factorization type is as follows:

1. *If u^2-4q is not a square in $\mathbb{F}_\ell$, then $e>1$ and the factorization type is*

$$(e,\ldots,e) \quad \textit{where} \quad e \mid \ell+1\ .$$

2. *If u^2-4q is a nonzero square in $\mathbb{F}_\ell$, then the factorization type is*

$$(1,1,e,\ldots,e) \quad \textit{where} \quad e \mid \ell-1\ .$$

[8] This is conditional under Remark 2.5.6 and should include the assumption that $G_\mu(J_1,J_2,J_3,x)$ is separable.

3. If $u^2 - 4q = 0$ in $\mathbb{F}_\ell$, then the factorization type is

$$(1, e) \quad \textit{where} \quad e = \ell \ .$$

Proof. By Lemma 4.11, the endomorphism π acts on $\mathcal{A}[\mu]$ as a 2×2 matrix in $\mathrm{GL}_2(\mathbb{F}_\ell)$ with characteristic polynomial $T^2 - uT + q = 0$. If the matrix has two conjugate eigenvalues λ_1, λ_2 in $\mathbb{F}_{\ell^2}$, then we are in Case (1): there are no 1-dimensional eigenspaces of π in $\mathcal{A}[\mu]$, and all irreducible factors of $G_\mu(J_1, J_2, J_3, x)$ have degree e, where e is the smallest exponent such that λ_i^e is in $\mathbb{F}_\ell$.

If the matrix has two eigenvalues in $\mathbb{F}_\ell$ and is diagonalizable, then the discriminant $t^2 - 4s$ is a square modulo ℓ: we are in Case (2). This time $\mathcal{A}[\mu]$ is the direct product of two 1-dimensional eigenspaces, which account for two linear factors of $G_\mu(J_1, J_2, J_3, x)$. The remaining factors have degree e, where e is the smallest positive integer such that π^e acts as a scalar matrix.

If the matrix has a double eigenvalue and is not diagonalizable, then we are in Case (3): there is only one 1-dimensional eigenspace, and the matrix of π^ℓ is scalar. □

4.5.3 The characteristic polynomial of Frobenius

Now that we can compute the order of Frobenius, we want to use this to derive information on the characteristic polynomial. Proposition 4.14 generalizes Proposition 4.2 to genus 2.

Proposition 4.14. *Let $(\mathcal{A}/\mathbb{F}_q, \xi, \iota)$ be a triple describing a vanilla abelian surface with real multiplication by $\mathcal{O}_F$, and let μ be a totally positive element of prime norm $\ell = \mu\bar{\mu} \notin \{2, p\}$. Let π be the Frobenius endomorphism of $\mathcal{A}$, and $\chi_{\pi,\mu}(T) = T^2 - uT + q$ the characteristic polynomial of its restriction to $\mathcal{A}[\mu]$. If e is the order of the image of π in $\mathrm{Aut}(\mathbb{P}(\mathcal{A}[\mu])) \cong \mathrm{PGL}_2(\mathbb{F}_\ell)$, then*

$$u^2 = \eta_e q \quad \textit{in } \mathbb{F}_\ell \ ,$$

$$\textit{where } \eta_e = \begin{cases} \zeta + \zeta^{-1} + 2 \textit{ with } \zeta \in \mathbb{F}_{\ell^2}^\times \textit{ of order } e & \textit{if } \gcd(\ell, e) = 1 \ , \\ 4 & \textit{otherwise} \ . \end{cases}$$

Proof. The proof is identical to that of Proposition 4.2. □

Coming back to point counting: suppose we have a Jacobian $J_\mathcal{C}$ with real multiplication by $\mathcal{O}_F$; we want to compute the characteristic polynomial

$$\chi_\pi(T) = T^4 - tT^3 + (2q + s)T^2 - tqT + q^2 \ .$$

If we have a totally positive element μ in $\mathcal{O}_F$ such that $\mu\bar{\mu} = \ell$, then we know that $\chi_\pi(T) \pmod{\ell}$ splits into two quadratic factors:

$$\chi_\pi(T) \equiv \chi_{\pi,\mu}(T)\chi_{\pi,\bar{\mu}}(T) \equiv (T^2 - uT + q)(T^2 - u'T + q) \pmod{\ell} \ ,$$

so

$$t \equiv u + u' \pmod{\ell} \qquad \text{and} \qquad s \equiv uu' - 2q \pmod{\ell} \ . \tag{4.6}$$

Given precomputed Hilbert modular polynomials G_μ and $G_{\bar{\mu}}$, then, we can specialize them at the RM invariants of $J_\mathcal{C}$ and factor to determine the order of Frobenius on $J_\mathcal{C}[\mu]$ and on $J_\mathcal{C}[\bar{\mu}]$ using Proposition 4.13. We can then apply Proposition 4.14 and Equations (4.6) to restrict the possible values of s and t modulo ℓ.

The question of how best to exploit this extra modular information remains open. Atkin's match-and-sort and Joux and Lercier's Chinese-and-match algorithms for elliptic curves cannot be re-used directly here, because they were designed to solve the one-dimensional problem of determining the elliptic trace, while here we have the two-dimensional problem of determining (s, t).

4.5.4 Prime types for real multiplication by $\mathcal{O}_F$

The factorization patterns in Proposition 4.13 are the same as those we saw for specialized elliptic modular polynomials in §4.2.4. This leads us to define an analogous classification of prime types, for totally positive elements in $\mathcal{O}_F$ of prime norm.

Definition 4.15. Let μ be a totally positive element of $\mathcal{O}_F$ such that $\mu\bar{\mu} = (\ell)$ for some prime $\ell \neq 2, p$. We say that

- μ is $\mathcal{O}_F$**-Elkies** for a vanilla triple $(\mathcal{A}, \xi, \iota)$ with RM invariants (J_1, J_2, J_3) if the factorization type of $G_\mu(J_1, J_2, J_3, x)$ is $(1, 1, e, \ldots, e)$ with $e > 1$;
- μ is $\mathcal{O}_F$**-Atkin** for a vanilla triple $(\mathcal{A}, \xi, \iota)$ with RM invariants (J_1, J_2, J_3) if the factorization type of $G_\mu(J_1, J_2, J_3, x)$ is $(e, \ldots, e)$ with $e > 1$; and
- μ is $\mathcal{O}_F$**-volcanic** for a vanilla triple $(\mathcal{A}, \xi, \iota)$ with RM invariants (J_1, J_2, J_3) if the factorization type of $G_\mu(J_1, J_2, J_3, x)$ is $(1, e)$ or $(1, \ldots, 1)$.

If $K \cong \mathrm{End}_{\overline{\mathbb{F}}_q}(\mathcal{A}) \otimes \mathbb{Q}$ is Galois then the type of μ completely determines the type of $\bar{\mu}$ (and vice versa). For general K, however, this does not hold: the type of $\bar{\mu}$ is not determined by the type of μ.

4.5.5 The parity of the number of factors of G_μ

The following proposition is the genus-2 real multiplication analogue of Equation (4.2) (cf. [Sch95, Proposition 6.3]).

Proposition 4.16. *Let $(\mathcal{A}, \xi, \iota)$ be a triple describing a vanilla abelian surface over $\mathbb{F}_q$ with real multiplication by $\mathcal{O}_F$, and with RM invariants (J_1, J_2, J_3). Let μ be a totally positive element of $\mathcal{O}_F$ of prime norm $\mu\bar{\mu} = \ell$, let $\chi_{\pi,\mu}(T) = T^2 - uT + q$ be the characteristic polynomial of π restricted to $\mathcal{A}[\mu]$, and let r denote the number of irreducible factors in the factorization of $G_\mu(J_1, J_2, J_3, x)$. Then*

$$(-1)^r = \left(\frac{q}{\ell}\right) .$$

Proof. If ℓ divides $u^2 - 4q$ and π has order ℓ in Case (3) of Proposition 4.13, then the result is true. Suppose therefore that $u^2 - 4q \neq 0 \pmod{\ell}$, that is, we are in Cases (1) or (2) of Proposition 4.13, and let $\mathcal{T} \subseteq \mathrm{GL}_2(\mathbb{F}_\ell)$ be a maximal torus containing π. In other words, we take $\mathcal{T} = \{\mathrm{diag}(\alpha, \beta) : \alpha, \beta \in \mathbb{F}_\ell^\times\}$ split in Case (2), and $\mathcal{T}$ non-split (i.e., isomorphic to $\mathbb{F}_{\ell^2}^\times$) in Case (1). The image $\overline{\mathcal{T}}$ of $\mathcal{T}$ in $\mathrm{PGL}_2(\mathbb{F}_\ell)$ is cyclic of order $\ell + 1$ in Case (1) and $\ell - 1$ in Case (2). The determinant induces an isomorphism $\det\colon \overline{\mathcal{T}}/\overline{\mathcal{T}}^2 \to \mathbb{F}_\ell^\times/(\mathbb{F}_\ell^\times)^2$. The action of π is via $\det(\pi) = q$, and we obtain an isomorphism $\det\colon \overline{\mathcal{T}}/\langle\overline{\mathcal{T}}^2, \pi\rangle \to \mathbb{F}_\ell^\times/\langle(\mathbb{F}_\ell^\times)^2, q\rangle$. This shows that the index $[\overline{\mathcal{T}} : \pi]$ is odd if and only if q is not a square mod ℓ. Since the number r of irreducible factors of $G_\mu(J_1, J_2, J_3, x)$ over $\mathbb{F}_q$ is equal to $r = (l + 1)/e$ or $r = 2 + (l - 1)/e = [\overline{\mathcal{T}} : \pi]$, the proposition follows. □

4.6 The case $F = \mathbb{Q}(\sqrt{5})$: Gundlach–Müller invariants

All of the theory above can be made much more explicit in the case where $F = \mathbb{Q}(\sqrt{5})$, where the invariants J_1, J_2, and J_3 are known as Gundlach–Müller invariants [Gun63; Mue85]. Our computational results are based on this case, so we will work out the details here, following the treatment in [LY11].

Fixing a square root of 5 in $\mathbb{C}$, we set $\epsilon := (1 + \sqrt{5})/2$ and $\bar{\epsilon} := (1 - \sqrt{5})/2$; each is the image of the fundamental unit of $\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$ under one of its two embeddings into $\mathbb{C}$. Let

$$q_1 := e\left(\frac{\epsilon z_1 - \bar{\epsilon} z_2}{\sqrt{5}}\right) \quad \text{and} \quad q_2 := e\left(\frac{z_2 - z_1}{\sqrt{5}}\right) \quad \text{for} \quad z = (z_1, z_2) \in \mathbb{H}^2 .$$

The Eisenstein series of even weight $k \geq 2$ are defined by

$$g_k(z) = 1 + \sum_{t = a + b\bar{\epsilon} \in \mathcal{O}_F^+} b_k(t) q_1^a q_2^b ,$$

where the coefficients $b_k(t)$ are defined by

$$b_k(t) = \kappa_k \sum_{(\mu) \supseteq (t)} \mathrm{N}(\mu)^{k-1} \quad \text{with} \quad \kappa_k = \frac{(2\pi)^{2k}\sqrt{5}}{(k-1)!^2 5^k \zeta_F(k)} \in \mathbb{Q}$$

(here $\mathrm{N}(\mu)$ is the norm $\#\mathcal{O}_F/(\mu)$). The Hilbert modular forms s_6, s_{10}, s_{12}, and s_{15} of respective weight 6, 10, 12, and 15 for $\mathcal{H}_{\mathbb{Q}(\sqrt{5})}$ are defined by

$$\begin{aligned}
s_6 &:= -\frac{67}{2^5 \cdot 3^3 \cdot 5^2}(g_6 - g_2^3) \ , \\
s_{10} &:= \frac{1}{2^{10} \cdot 3^5 \cdot 5^5 \cdot 7}\left(191 \cdot 2161 g_{10} - 5 \cdot 67 \cdot 2293 g_2^2 g_6 + 2^2 \cdot 3 \cdot 7 \cdot 4231 g_2^5\right) \ , \\
s_{12} &:= \frac{1}{2^2}\left(s_6^2 - g_2 s_{10}\right) \ , \\
s_5^2 &:= s_{10} \ , \\
s_{15}^2 &:= 5^5 s_{10}^3 - \frac{5^3 g_2^2 s_6 s_{10}^2}{2} + \frac{g_2^5 s_{10}^2}{2^4} + \frac{3^2 \cdot 5^2 g_2 s_6^3 s_{10}}{2} - \frac{g_2^4 s_6^2 s_{10}}{2^3} - 2 \cdot 3^3 s_6^5 + \frac{g_2^3 s_6^4}{2^4} \ .
\end{aligned}$$

Finally, the Gundlach–Müller invariants for $\mathbb{Q}(\sqrt{5})$ are

$$J_1 := s_6/g_2^3 \ , \quad J_2 := g_2^5/s_5^2 \ , \quad \text{and} \quad J_3 := s_5^3/s_{15} \ .$$

The Hilbert modular polynomials for $\mathbb{Q}(\sqrt{5})$ are too large to reproduce here, but they can be downloaded from `martindale.info`.[9]

Proposition 4.17 ([LY11, Proposition 4.5] with correction to $\phi^*(j_1)$)**.** *For $F = \mathbb{Q}(\sqrt{5})$, the Igusa invariants pull back to*

$$\begin{aligned}
\phi^*(j_1) &= 4J_2(3J_1^2 J_2 - 2)^5 \ , \\
\phi^*(j_2) &= \frac{1}{2} J_2 (3J_1^2 J_2 - 2)^3 \ , \\
\phi^*(j_3) &= 2^{-3} J_2 (2J_1^2 J_2 - 2)^2 (4J_1^2 J_2 + 2^5 \cdot 3^2 J_1 - 3) \ .
\end{aligned}$$

For our computations, we want to write J_1, J_2 and J_3 in terms of the Siegel modular forms ψ_4, ψ_6, χ_{10} and χ_{12}. (For a canonical way of writing J_1, J_2 and J_3 in terms of Igusa–Clebsch invariants, we refer to [Mar18, Example 2.5.4].)

Proposition 4.18 ([Mar18, Example 2.5.4][10])**.** *For $F = \mathbb{Q}(\sqrt{5})$, we have*

$$\begin{aligned}
J_2 &= \phi^*\left((\psi_4\psi_6/\chi_{10} - 3^5 2^{12})(-2 - 2(\psi_6^2 - 2^{12}3^6\chi_{12})/\psi_4^3)^{-1}\right) \ , \\
J_1 &= 3^2 2^5 J_2^{-1} + \phi^*\left(2^{-6}3^{-3}(1 - (\psi_6^2 - 2^{12}3^6\chi_{12})/\psi_4^3)\right) \ , \\
J_3^2 &= 5^5 - 2^{-1}5^3 J_1 J_2 + 2^{-4} J_2 + 2^{-1}3^2 5^2 J_2^2 J_1^3 - 2^{-3} J_1^2 J_2^2 - 2 \cdot 3^3 J_2^3 J_1^5 + 2^{-4} J_2^3 J_1^4 \ .
\end{aligned}$$

The choice of square root for J_3 corresponds to the choice of embedding ι.

Proposition 4.18 can be used to find RM invariants for curves drawn from families with known real multiplication, before factoring specialized Hilbert modular polynomials in those RM invariants to derive information on Frobenius. However, it also crystallizes the rationality question alluded to at the end of §4.4.4: as we see, a set of values of the Hilbert modular forms over $\mathbb{F}_q$ (or, equivalently, a tuple of Igusa or Igusa–Clebsch invariants over $\mathbb{F}_q$) only determine J_1, J_2, and J_3^2 over $\mathbb{F}_q$.

To get J_3, we need to choose a square root of J_3^2; but J_3^2 is not guaranteed to be a square in $\mathbb{F}_q$. If J_3^2 is not a square in $\mathbb{F}_q$, then we cannot apply Propositions 4.9 or 4.13—not even if J_3 does not appear unsquared in the specialized polynomial G_μ.

4.7 Experimental results

In order to validate the factorization patterns of Proposition 4.13, we ran a series of experiments for $F = \mathbb{Q}(\sqrt{5})$, using the family of curves [TTV91]

$$\mathcal{C}_a : y^2 = x^5 - 5x^3 + 5x + a$$

[9] The polynomials $H_{\mu,3}$ do not appear there, but only G_μ is required to apply our results in §4.5.

[10] The number of this example has changed to Example 2.4.4

whose Jacobians all have real multiplication by $\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$. This family was used in the point-counting records of [GKS11]. The Igusa–Clebsch invariants of $\mathcal{C}_a$ are

$$(A, B, C, D) = \left(2^5 \cdot 5^2 \cdot 7, 2^{10} \cdot 5^4, -2^{13} \cdot 5^5 \cdot (9a^2 - 236), 2^{20} \cdot 5^5 \cdot (a^2 - 4)^2\right) .$$

Our experiments treated

1. the ramified prime $\ell = 5$, with $\mu = (5 + \sqrt{5})/2$, and the modular polynomial G_μ from `martindale.info`;
2. the split prime $\ell = 11$, with $\mu = (7+\sqrt{5})/2$, and the modular polynomial G_μ from `martindale.info`.

We collected statistics on the factorization patterns for 10000 tests. For each test, we chose a random prime q of ten decimal digits, and we chose a randomly from $\mathbb{F}_q$ subject to the requirement that $\mathcal{C}_a$ be nonsingular, which is $a^2 \neq 4$. We then applied the formulæ of Equation (4.5) and Proposition 4.18 to obtain the RM invariants J_2 and J_1 for the Jacobian of $\mathcal{C}_a$, as well as the squared invariant J_3^2.

In half the cases on average, J_3^2 had a square root in $\mathcal{F}_q$; in these cases we could obtain J_3, and proceed to factor $G_\mu(J_1, J_2, J_3, x)$. The average frequencies of the resulting factorization patterns appear in Tables 4.1 and 4.2 (here we take the averages over the roughly 5000 tests where J_3^2 has a root in $\mathcal{F}_q$; for the two roots J_3 and $-J_3$ in $\mathcal{F}_q$, we always obtained the same factorization pattern).

Factorization pattern, type of μ		Number found	Percentage
$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$-Elkies:	$(1, 1, e, \ldots, e)$ with $e > 1$	total 1835	total 36.8%
	$(1, 1, 4)$	1266	25.4%
	$(1, 1, 2, 2)$	569	11.4%
$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$-Atkin:	$(e, \ldots, e)$ with $e > 1$	total 2049	total 41.1%
	(6)	844	16.9%
	$(3, 3)$	794	15.9%
	$(2, 2, 2)$	411	8.2%
$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$-Volcanic:	$(1, e)$ or $(1, \ldots, 1)$	total 1105	total 22.1%
	$(1, 5)$	1058	21.2%
	$(1, 1, 1, 1, 1, 1)$	47	0.9%

Table 4.1: Factorization pattern frequencies for the modular polynomial $G_\mu(J_1, J_2, J_3, x)$ for $\mu = (5+\sqrt{5})/2$ of norm $\ell = 5$. The degree of $G_\mu(J_1, J_2, J_3, x)$ in x is 6. We only factored when J_3^2 was a square in $\mathcal{F}_q$, which happened in 4989 of the 10000 trials (49.9%).

According to Proposition 4.13, we would expect that $1/\ell$ of the time μ should be $\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$-volcanic, $(\ell-1)/2\ell$ of the time μ should be $\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$-Elkies, and $(\ell-1)/2\ell$ of the time μ should be $\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$-Atkin. The summary of our above results in Table 4.3 appears to confirm this. This gives us considerable confidence that the Hilbert modular polynomials computed in [Mar18, Chapter 2] are correct.

Finally, we ran the same tests on Milio's modular polynomial[11] $\Phi(\mathfrak{J}_1, \mathfrak{J}_2, X)$ for $\ell = 5$ and $\mu = (5 + \sqrt{5})/2$, where $\mathfrak{J}_1 = J_2$ and $\mathfrak{J}_2 = J_1 J_2$. We obtained exactly the same factorization patterns each time J_3 was in $\mathcal{F}_q$.

[11] Available from `https://members.loria.fr/EMilio/modular-polynomials/`

Factorization pattern, type of μ		Number found	Percentage
$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$-Elkies:	$(1,1,e,\ldots,e)$ with $e>1$	total 2262	total 44.7%
	$(1,1,10)$	994	19.7%
	$(1,1,5,5)$	1040	20.6%
	$(1,1,2,2,2,2,2)$	228	4.5%
$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$-Atkin:	$(e,\ldots,e)$ with $e>1$	total 2329	total 46.1%
	(12)	859	17.0%
	$(6,6)$	404	8.0%
	$(4,4,4)$	424	8.4%
	$(3,3,3,3)$	429	8.5%
	$(2,2,2,2,2,2)$	213	4.2%
$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$-volcanic:	$(1,e)$ or $(1,\ldots,1)$	total 466	total 9.2%
	$(1,11)$	461	9.1%
	$(1,1,1,1,1,1,1,1,1,1,1,1)$	5	0.1%

Table 4.2: Factorization pattern frequencies for the modular polynomial $G_\mu(J_1,J_2,J_3,x)$ for $\mu=(7+\sqrt{5})/2$ of norm $\ell=11$. The degree of $G_\mu(J_1,J_2,J_3,x)$ in x is 12. We only factored when J_3^2 was a square in $\mathbb{F}_q$, which happened in 5057 of the 10000 trials (50.6%).

		Prime type frequencies for μ		
		$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$-volcanic	$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$-Elkies	$\mathcal{O}_{\mathbb{Q}(\sqrt{5})}$-Atkin
$\mu=\frac{5-\sqrt{5}}{2}$	Theory	20.0%	40.0%	40.0%
	Experiments	22.1%	36.8%	41.1%
$\mu=\frac{7+\sqrt{5}}{2}$	Theory	9.1%	45.5%	45.5%
	Experiments	9.2%	44.7%	46.1%

Table 4.3: Experimental evidence supporting the correctness of Martindale's Hilbert modular polynomials.

Appendix A

The notions of dual and polarisation in equivalent categories

In all that follows C and D will denote categories with an involution called dual and denoted by $\vee$. Furthermore, functors $F : C \to D$ and $G : D \to C$ will denote an adjoint equivalence (cf. [Lan78, Definition on p. 93]) of categories via natural isomorphisms

$$\gamma = \{\gamma_{\mathrm{A}} : GF\mathrm{A} \xrightarrow{\sim} \mathrm{A} : \mathrm{A} \in C\} \tag{A.1}$$

and

$$\delta = \{\delta_{\mathrm{A}} : FG\mathrm{A} \xrightarrow{\sim} \mathrm{A} : \mathrm{A} \in D\}. \tag{A.2}$$

Recall from Mac Lane [Lan78, Theorem IV.1.1] that in this situation, for all $\mathrm{A} \in C$, we have that

$$F(\gamma_{\mathrm{A}}) = \delta_{F(\mathrm{A})}, \tag{A.3}$$

and for all $\mathrm{B} \in D$, we have that $G(\delta_{\mathrm{B}}) = \gamma_{G(\mathrm{B})}$.

The purpose of this Appendix is to prove that, as stated in Remark 1.3.12, if F preserves the notion of dual (resp. polarisation), then G preserves the notion of dual (resp. polarisation). For the convenience of the reader, we repeat the definitions from Remark 1.3.12 of the preservation of the notions of dual and polarisation.

Definition A.1. If F comes together with a natural isomorphism

$$f : F \circ \vee \xrightarrow{\sim} \vee \circ F,$$

we say that (F, f) *preserves the notion of dual.* For an object $\mathrm{A} \in C$, we write f_{A} for the isomorphism $F(\mathrm{A}^\vee) \xrightarrow{\sim} F(\mathrm{A})^\vee$.

Claim A.2. Suppose that (F, f) preserves the notion of dual. Then there is a natural isomorphism $g : G \circ \vee \xrightarrow{\sim} \vee \circ G$ given by

$$\{g_{\mathrm{B}} = \gamma_{(G\mathrm{B})^\vee} \circ G(f^{-1}_{G\mathrm{B}}) \circ G(\delta^\vee_{\mathrm{B}}) : \mathrm{B} \in D\}. \tag{A.4}$$

In particular, by definition (G, g) preserves the notion of dual.

Proof. The morphism g given in (A.4) is a natural isomorphism if for every $\varphi \in \mathrm{Hom}_D(\mathrm{A}, \mathrm{B})$, the diagram

$$\begin{tikzcd}
G(\mathrm{A}^\vee) \arrow[r, "G(\delta^\vee_{\mathrm{A}})"] & G((FG\mathrm{A})^\vee) \arrow[r, "G(f^{-1}_{G\mathrm{A}})"] & GF((G\mathrm{A})^\vee) \arrow[r, "\gamma_{(G\mathrm{A})^\vee}"] & (G\mathrm{A})^\vee \\
G(\mathrm{B}^\vee) \arrow[u, "G(\varphi^\vee)"] \arrow[r, "G(\delta^\vee_{\mathrm{B}})"'] \arrow[ur, phantom, "(1)\ G((FG\varphi)^\vee)"] & G((FG\mathrm{B})^\vee) \arrow[u] \arrow[r, "G(f^{-1}_{G\mathrm{B}})"'] \arrow[ur, phantom, "(2)\ GF((G\varphi)^\vee"] & GF((G\mathrm{B})^\vee) \arrow[u] \arrow[r, "\gamma_{(G\mathrm{B})^\vee}"'] \arrow[ur, phantom, "(3)\quad (G\varphi)^\vee"] & (G\mathrm{B})^\vee \arrow[u]
\end{tikzcd}$$

commutes. But (1) commutes as δ is a natural isomorphism, (2) commutes as f is a natural isomorphism, and (3) commutes as γ is a natural isomorphism. Hence

$$g = \{g_{\mathrm{B}} = \gamma_{(G\mathrm{B})^\vee)} \circ G(f^{-1}_{G\mathrm{B}}) \circ G(\delta^\vee_{\mathrm{B}}) : \mathrm{B} \in D\}$$

is a natural isomorphism

$$g : G \circ \vee \xrightarrow{\sim} \vee \circ G.$$

□

From now on, we suppose that for all objects in C (resp. D), we have a subset $P_{\mathrm{A}} \subseteq \mathrm{Hom}(\mathrm{A}, \mathrm{A}^\vee)$ of 'polarisations' such that for every isomorphism $m : \mathrm{B} \to \mathrm{A}$ in C (resp. D), the map

$$\begin{array}{ccc} \mathrm{Hom}(\mathrm{A}, \mathrm{A}^\vee) & \longrightarrow & \mathrm{Hom}(\mathrm{B}, \mathrm{B}^\vee) \\ \varphi & \mapsto & m^\vee \varphi m \end{array} \tag{A.5}$$

induces a bijection between P_{A} and P_{B}.

Definition A.3. Suppose that (F, f) preserves the notion of dual. We say that (F, f) *preserves the notion of polarisation* if for all objects $\mathrm{A} \in C$ the map

$$\begin{array}{ccc} \mathrm{Hom}(\mathrm{A}, \mathrm{A}^\vee) & \longrightarrow & \mathrm{Hom}(F(\mathrm{A}), F(\mathrm{A})^\vee) \\ \xi & \mapsto & f_{\mathrm{A}} \circ F(\xi) \end{array}$$

induces a bijection between P_{A} and $P_{F(\mathrm{A})}$.

Claim A.4. Suppose that (F, f) preserves the notions of dual and of polarisation. Then (G, g) with g as in (A.4) preserves the notion of polarisation.

Lemma A.5. Suppose that (F, f) preserves the notion of dual. Then (FG, h) preserves the notions of dual and of polarisation, where

$$h = \{h_B = f_{GB} \circ F(g_B) : B \in D\} \tag{A.6}$$

and g is as in (A.4).

Proof. We have that (G, g) preserves the notion of dual by Claim A.2, and it is easy to check that (FG, h) also preserves the notion of dual. Therefore it suffices to prove that the isomorphism

$$\begin{array}{ccc} \mathrm{Hom}_D(\mathrm{B}, \mathrm{B}^\vee) & \longrightarrow & \mathrm{Hom}_D(FG\mathrm{B}, (FG\mathrm{B})^\vee) \\ \xi & \mapsto & h_B \circ FG(\xi) \end{array}$$

induces a bijection between P_{B} and $P_{FG\mathrm{B}}$. By (A.5), the isomorphism
$\delta_{\mathrm{B}} : FG\mathrm{B} \xrightarrow{\sim} \mathrm{B}$ gives an isomorphism

$$\begin{array}{ccc} \mathrm{Hom}_D(\mathrm{B}, \mathrm{B}^\vee) & \longrightarrow & \mathrm{Hom}_D(FG\mathrm{B}, (FG\mathrm{B})^\vee) \\ \xi & \mapsto & \delta_{\mathrm{B}}^\vee \circ \xi \circ \delta_{\mathrm{B}}, \end{array}$$

which in turn induces a bijection between P_{B} and $P_{FG\mathrm{B}}$. Hence, it suffices to show that for every $\xi \in \mathrm{Hom}_D(\mathrm{B}, \mathrm{B}^\vee)$, we have that

$$h_B \circ FG(\xi) = \delta_{\mathrm{B}}^\vee \circ \xi \circ \delta_{\mathrm{B}}. \tag{A.7}$$

Note, by definition of h_B and g_B, that

$$\begin{aligned} h_{\mathrm{B}} \circ FG(\xi) &= f_{G\mathrm{B}} \circ F(g_B) \circ FG(\xi) \\ &= f_{GB} \circ F(\gamma_{(G\mathrm{B})^\vee}) \circ FG(f_{G\mathrm{B}}^{-1} \circ \delta_{\mathrm{B}}^\vee \circ \xi), \end{aligned}$$

and $F(\gamma_{(G\mathrm{B})^\vee}) = \delta_{F((G\mathrm{B})^\vee)}$ by (A.3), so

$$h_B \circ FG(\xi) = f_{G\mathrm{B}} \circ \delta_{F((G\mathrm{B})^\vee)} \circ FG(f_{G\mathrm{B}}^{-1} \circ \delta_B^\vee \circ \xi). \tag{A.8}$$

Furthermore, as $f_{G\mathrm{B}}^{-1} \circ \delta_{\mathrm{B}}^\vee \circ \xi \in \mathrm{Hom}_D(B, F((G\mathrm{B})^\vee)$ and δ is a natural isomorphism, the diagram

$$\begin{tikzcd} \mathrm{B} \arrow[r, "f_{G\mathrm{B}}^{-1} \circ \delta_{\mathrm{B}}^\vee \circ \xi"] & F((G\mathrm{B})^\vee) \\ FG\mathrm{B} \arrow[u, "\delta_B"] \arrow[r, "FG(f_{G\mathrm{B}}^{-1} \circ \delta_{\mathrm{B}}^\vee \circ \xi)"] & FGF((G\mathrm{B})^\vee) \arrow[u, "\delta_{F((G\mathrm{B})^\vee)}"'] \end{tikzcd}$$

commutes, so that

$$\delta_{F((G\mathrm{B})^\vee)} \circ FG(f_{G\mathrm{B}}^{-1} \circ \delta_{\mathrm{B}}^\vee \circ \xi) = f_{G\mathrm{B}}^{-1} \circ \delta_{\mathrm{B}}^\vee \circ \xi \circ \delta_{\mathrm{B}}.$$

Plugging this into (A.8) gives (A.7) and the result follows. □

Proof of Claim A.4. As (F, f) preserves the notion of dual, by Claim A.2 (G, g) also preserves the notion of dual. Therefore, it suffices to show that the isomorphism

$$\begin{array}{rccc} b_{\mathrm{B}}: & \mathrm{Hom}_D(\mathrm{B}, \mathrm{B}^\vee) & \longrightarrow & \mathrm{Hom}_C(G\mathrm{B}, (G\mathrm{B})^\vee) \\ & \xi & \mapsto & g_{\mathrm{B}} \circ G(\xi) \end{array}$$

induces a bijection between P_{B} and $P_{G\mathrm{B}}$. But by Claim A.5 (FG, h) preserves the notion of polarisation, where h is as in (A.6), hence the isomorphism

$$\begin{array}{rccc} c_{\mathrm{B}}: & \mathrm{Hom}_D(\mathrm{B}, \mathrm{B}^\vee) & \longrightarrow & \mathrm{Hom}_D(FG\mathrm{B}, (FG\mathrm{B})^\vee) \\ & \xi & \mapsto & f_{G\mathrm{B}} \circ F(g_{\mathrm{B}}) \circ FG(\xi) \end{array}$$

induces a bijection between P_{B} and $P_{FG\mathrm{B}}$, and as (F, f) preserves the notion of polarisation, the isomorphism

$$\begin{array}{rccc} d_{G\mathrm{B}}: & \mathrm{Hom}_C(G\mathrm{B}, (G\mathrm{B})^\vee) & \longrightarrow & \mathrm{Hom}_D(FG\mathrm{B}, (FG\mathrm{B})^\vee) \\ & \xi & \mapsto & f_{G\mathrm{B}} \circ F(\xi) \end{array}$$

induces a bijection between $P_{G\mathrm{B}}$ and $P_{FG\mathrm{B}}$. Therefore the sets P_{B} and $P_{G\mathrm{B}}$ are in bijection via the isomorphism $d_{G\mathrm{B}}^{-1} \circ c_{\mathrm{B}}$, which is indeed b_{B} as, for any $\xi \in \mathrm{Hom}_D(\mathrm{B}, \mathrm{B}^\vee)$, we have that

$$\begin{aligned} d_{G\mathrm{B}}(b_{\mathrm{B}}(\xi)) &= d_{G\mathrm{B}}(g_{\mathrm{B}} \circ G(\xi)) \\ &= f_{G\mathrm{B}} \circ F(g_{\mathrm{B}} \circ G(\xi)) \\ &= f_{G\mathrm{B}} \circ F(g_{\mathrm{B}}) \circ FG(\xi) \\ &= c_{\mathrm{B}}(\xi). \end{aligned}$$

□

Bibliography

[Bal+17] S. Ballentine, A. Guillevic, E. Lorenzo-García, M. Massierer, C. Martindale, B. Smith, and J. Top. "Isogenies for point counting on genus two hyperelliptic curves with maximal real multiplication". In: *Algebraic Geometry for Coding Theory and Cryptography.* Vol. 9. Association for Women in Mathematics Series. Springer Int. Pub., 2017, pp. 63–94. ISBN: 978-3-319-63931-4.

[BCR] G. Bisson, R. Cosset, and D. Robert. *AVIsogenies: a library for computing isogenies between abelian varieties.* `http://\discretionary{}{}{}avisogenies.gforge.inria.fr`.

[BJW17] E.H. Brooks, D. Jetchev, and B. Wesolowski. "Isogeny graphs of ordinary abelian varieties". In: *Research in Number Theory* 3 (2017). ISSN: 2363-9555.

[BS17] G. Bisson and M. Streng. "On polarised class groups of orders in quartic CM-fields". In: *Math. Res. Lett.* 24.2 (2017), pp. 247–270. ISSN: 1073-2780.

[Can87] D.G. Cantor. "Computing in the Jacobian of a hyperelliptic curve". In: *Math. Comp.* 48.177 (1987), pp. 95–101.

[Car04] R. Carls. "A generalized arithmetic geometric mean". PhD thesis. University of Groningen, The Netherlands, 2004. URL: `http://hdl.handle.net/11370/f47bd074-2c0d-4521-a92b-4e89af5c1840`.

[CE15] J.-M. Couveignes and T. Ezome. "Computing functions on Jacobians and their quotients". In: *LMS J. Comput. Math.* 18.1 (2015), pp. 555–577.

[Coh93] H. Cohen. *A Course in Computational Algebraic Number Theory.* Vol. 138. Graduate Texts in Mathematics. Spring-Verlag, 1993.

[Del69] P. Deligne. "Variétés abéliennes ordinaires sur un corps fini". In: *Invent. Math.* 8 (1969), pp. 238–243. ISSN: 0020-9910.

[Dup06] R. Dupont. "Moyenne Arithmético-géométrique, Suites de Borchant et Applications". PhD thesis. École Polytechnique, 2006. URL: `http://www.lix.polytechnique.fr/Labo/Regis.Dupont/these_soutenance.pdf`.

[Echidna] D. Kohel. *The Echidna Database.* URL: `https://www.i2m.univ-amu.fr/perso/david.kohel/dbs/index.html`.

[Fly90] E.V. Flynn. "The Jacobian and Formal Group of a Curve of Genus 2 over an Arbitrary Ground Field". In: *Math. Proc. Cambridge Philos. Soc.* 107.3 (1990), pp. 425–441.

[Gee88] G. van der Geer. *Hilbert modular surfaces.* Vol. 16. Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]. Springer-Verlag, Berlin, 1988, pp. x+291. ISBN: 3-540-17601-2.

[GH00] P. Gaudry and R. Harley. "Counting Points on Hyperelliptic Curves over Finite Fields". In: *Algorithmic Number Theory, 4th International Symposium, ANTS-IV (Leiden, The Netherlands).* Ed. by W. Bosma. Vol. 1838. Lecture Notes in Computer Science. Berlin: Springer, 2000, pp. 313–332.

[GKS11] P. Gaudry, D. Kohel, and B. Smith. "Counting Points on Genus 2 Curves with Real Multiplication". In: *Advances in Cryptology—ASIACRYPT 2011 (Seoul, South Korea).* Ed. by Dong Hoon Lee and Xiaoyun Wang. Vol. 7073. Lecture Notes in Computer Science. Heidelberg: Springer, 2011, pp. 504–519.

[GM07] G. van der Geer and B. Moonen. *Abelian varieties*. Book in preparation. 2007.

[Gra90] D. Grant. "Formal groups in genus two". In: *J. Reine Angew. Math.* 411 (1990), pp. 96–121.

[Gro61] A. Grothendieck. "Éléments de géométrie algébrique. II. Étude globale élémentaire de quelques classes de morphismes". In: *Inst. Hautes Études Sci. Publ. Math.* 8 (1961), p. 222. ISSN: 0073-8301. URL: http://www.numdam.org/item?id=PMIHES_1961__8__222_0.

[GS12] P. Gaudry and E. Schost. "Genus 2 point counting over prime fields". In: *J. Symbolic Comput.* 47.4 (2012), pp. 368–400. DOI: 10.1016/j.jsc.2011.09.003.

[Gun63] K.-B. Gundlach. "Die Bestimmung der Funktionen zur Hilbertschen Modulgruppe des Zahlkörpers $Q(\sqrt{5})$". In: *Math. Ann.* 152 (1963), pp. 226–256. ISSN: 0025-5831.

[Har07] D. Harvey. "Kedlaya's algorithm in larger characteristic". In: *Int. Math. Res. Not. IMRN* 22 (2007), Art. ID rnm095, 29. ISSN: 1073-7928.

[Har12] M.C. Harrison. "An extension of Kedlaya's algorithm for hyperelliptic curves". In: *J. Symbolic Comput.* 47.1 (2012), pp. 89–101. ISSN: 0747-7171.

[Har77] R. Hartshorne. *Algebraic geometry*. Graduate Texts in Mathematics, No. 52. Springer-Verlag, New York-Heidelberg, 1977, pp. xvi+496. ISBN: 0-387-90244-9.

[How95] E.W. Howe. "Principally polarized ordinary abelian varieties over finite fields". In: *Trans. Amer. Math. Soc.* 347.7 (1995), pp. 2361–2401. ISSN: 0002-9947.

[HZ02] E.W. Howe and H.J. Zhu. "On the existence of absolutely simple abelian varieties of a given dimension over an arbitrary field". In: *J. Number Theory* 92.1 (2002), pp. 139–163. ISSN: 0022-314X.

[Igu60] J. Igusa. "Arithmetic variety of moduli for genus two". In: *Ann. of Math. (2)* 72 (1960), pp. 612–649. ISSN: 0003-486X.

[Igu67] J. Igusa. "Modular Forms and Projective Invariants". In: *Amer. J. Math.* 89.3 (1967), pp. 817–855.

[IT14] S. Ionica and E. Thomé. *Isogeny graphs of genus 2 curves with Maximal Real Multiplication*. https://eprint.iacr.org/2014/230.pdf. 2014.

[JL01] A. Joux and R. Lercier. ""Chinese & Match", an alternative to Atkin's "Match and Sort" method used in the SEA algorithm". In: *Math. Comp.* 70.234 (2001), pp. 827–836.

[Kat81] N. Katz. "Serre-Tate local moduli". In: *Algebraic surfaces (Orsay, 1976–78)*. Vol. 868. Lecture Notes in Math. Springer, Berlin-New York, 1981, pp. 138–202.

[Ked01] K.S. Kedlaya. "Counting points on hyperelliptic curves using Monsky–Washnitzer cohomology". In: *J. Ramanujan Math. Soc.* 16.4 (2001), pp. 323–338. ISSN: 0970-1249.

[Knu91] M.-A. Knus. *Quadratic and Hermitian forms over rings*. Vol. 294. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. With a foreword by I. Bertuccioni. Springer-Verlag, Berlin, 1991, pp. xii+524. ISBN: 3-540-52117-8.

[Koh96] D. Kohel. "Endomorphism rings of elliptic curves over finite fields". PhD thesis. University of California, Berkeley, 1996, p. 117. ISBN: 978-0591-32123-4.

[Lan78] S. Mac Lane. *Categories for the Working Mathematician*. Vol. 5. Graduate Texts in Mathematics. Springer New York, 1978, p. 317.

[Lan82] S. Lang. *Introduction to Algebraic and Abelian Functions*. Vol. 89. Graduate Texts in Mathematics. New York: Springer-Verlag, 1982.

[Lan83] Serge Lang. *Complex multiplication*. Vol. 255. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer-Verlag, New York, 1983, pp. viii+184. ISBN: 0-387-90786-6.

[Lan86] S. Lang. *Algebraic Number Theory*. Vol. 16. Graduate Texts in Mathematics. New York: Springer-Verlag, 1986.

[Ler97] R. Lercier. "Algorithmique des courbes elliptiques dans les corps finis". PhD thesis. École Polytechnique, Palaiseau, France, 1997. URL: https://tel.archives-ouvertes.fr/tel-01101949.

[LNY16] K. Lauter, M. Naehrig, and T. Yang. "Hilbert theta series and invariants of genus 2 curves". In: *J. Number Theory* 161 (2016), pp. 146–174. ISSN: 0022-314X.

[LST64] J. Lubin, J.P. Serre, and J. Tate. "Elliptic Curves and Formal Groups". In: (1964). Unpublished skeleton seminar notes.

[LY11] K. Lauter and T. Yang. "Computing genus 2 curves from invariants on the Hilbert moduli space". In: *J. Number Theory* 131.5 (2011), pp. 936–958. ISSN: 0022-314X.

[Mar18] C. Martindale. "Isogeny Graphs, Modular Polynomials, and Applications". PhD thesis. Universiteit Leiden, 2018.

[May07] S. Mayer. "Hilbert Modular Forms for the Fields $Q(\sqrt{5})$, $Q(\sqrt{13})$ and $Q(\sqrt{17})$". PhD thesis. Rheinisch-Westfälischen Technischen Hochschule Aachen, 2007. URL: http://www.matha.rwth-aachen.de/~mayer/homepage/dissertation-S-Mayer-revised-edition.pdf.

[Mes01] J.-F. Mestre. *Lettre à Gaudry et Harley.* https://webusers.imj-prg.fr/~jean-francois.mestre/lettreGaudryHarley.ps. 2001.

[Mes02] J.-F. Mestre. *Algorithme pour compter des points de courbes en petite caractéristique et petit genre.* https://webusers.imj-prg.fr/~jean-francois.mestre/rennescrypto.ps. Notes from a talk given at the Rennes cryptography seminar. 2002.

[Mes72] W. Messing. *The crystals associated to Barsotti-Tate groups: with applications to abelian schemes.* Lecture Notes in Mathematics, Vol. 264. Springer-Verlag, Berlin-New York, 1972, pp. iii+190.

[MFK94] D. Mumford, J. Fogarty, and F. Kirwan. *Geometric invariant theory.* Third. Vol. 34. Ergebnisse der Mathematik und ihrer Grenzgebiete (2) [Results in Mathematics and Related Areas (2)]. Springer-Verlag, Berlin, 1994, pp. xiv+292. ISBN: 3-540-56963-4.

[Mil15a] E. Milio. "A quasi-linear time algorithm for computing modular polynomials in dimension 2". In: *LMS J. Comput. Math.* 18.1 (2015), pp. 603–632.

[Mil15b] E. Milio. "Computing modular polynomials in dimension 2". PhD thesis. Université de Bordeaux, Dec. 2015. URL: https://tel.archives-ouvertes.fr/tel-01240690.

[Mil86] J.S. Milne. "Abelian varieties". In: *Arithmetic Geometry (Storrs, Connecticut, 1984).* New York: Springer, 1986, pp. 103–150. DOI: 10.1007/978-1-4613-8655-1.

[Mue83] R. Mueller. "Hilbertsche Modulformen und Modulfunktionen zu $\mathbf{Q}(\sqrt{8})$". In: *Math. Ann.* 266.1 (1983), pp. 83–103. ISSN: 0025-5831.

[Mue85] R. Mueller. "Hilbertsche Modulformen und Modulfunktionen zu $\mathbf{Q}(\sqrt{5})$". In: *Arch. Math. (Basel)* 45.3 (1985), pp. 239–251. ISSN: 0003-889X.

[Mum08] D. Mumford. *Abelian varieties.* Vol. 5. Tata Institute of Fundamental Research Studies in Mathematics. With appendices by C. P. Ramanujam and Yuri Manin, Corrected reprint of the second (1974) edition. Published for the Tata Institute of Fundamental Research, Bombay; by Hindustan Book Agency, New Delhi, 2008, pp. xii+263. ISBN: 978-81-85931-86-9; 81-85931-86-0.

[Pil90] J. Pila. "Frobenius maps of abelian varieties and finding roots of unity in finite fields". In: *Math. Comp.* 55.192 (1990), pp. 745–763.

[Rap78] M. Rapoport. "Compactifications de l'espace de modules de Hilbert-Blumenthal". In: *Compositio Math.* 36.3 (1978), pp. 255–335. ISSN: 0010-437X.

[Rüc90] H.-G. Rück. "Abelian surfaces and Jacobian varieties over finite fields". In: *Compositio Math.* 76.3 (1990), pp. 351–366. ISSN: 0010-437X.

[Sage] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 8.2).* http://www.sagemath.org. 2018.

[Sat02] T. Satoh. "On p-adic Point Counting Algorithms for Elliptic Curves over Finite Fields". In: *Algorithmic Number Theory (Sydney, 2002).* Ed. by C. Fieker and D. R. Kohel. Vol. 2369. Lecture Notes in Computer Science. Berlin: Springer, 2002, pp. 43–66.

[Sch85] R. Schoof. "Elliptic curves over finite fields and the computation of square roots mod p". In: *Math. Comp.* 44.170 (1985), pp. 483–494.

[Sch95] R. Schoof. "Counting points on elliptic curves over finite fields". In: *J. Théor. Nombres Bordeaux* 7.1 (1995), pp. 219–254.

[Stack-Exchange] Matt E (`https://math.stackexchange.com/users/221/matt-e`). *Why is a smooth connected scheme irreducible?* Mathematics Stack Exchange. eprint: `https://math.stackexchange.com/q/20508`.

[Ste08] P. Stevenhagen. "The arithmetic of number rings". In: *Algorithmic number theory: lattices, number fields, curves and cryptography.* Vol. 44. Math. Sci. Res. Inst. Publ. Cambridge Univ. Press, Cambridge, 2008, pp. 209–266.

[Sut13] A.V. Sutherland. "On the evaluation of modular polynomials". In: *ANTS X—Proceedings of the Tenth Algorithmic Number Theory Symposium (San Diego, California).* Vol. 1. Open Book Series. Berkeley, CA: Mathematical Sciences Publishers, 2013, pp. 531–555.

[Sut18] A. Sutherland. *Modular Polynomials.* Online database. 2018. URL: `https://math.mit.edu/~drew/~ClassicalModPolys.html`.

[TTV91] W. Tautz, J. Top, and A. Verberkmoes. "Explicit hyperelliptic curves with real multiplication and permutation polynomials". In: *Canad. J. Math.* 43.5 (1991), pp. 1055–1064. ISSN: 0008-414X.

[Was08] L.C. Washington. *Elliptic Curves: Number Theory and Cryptography.* Second. Vol. 50. Discrete Mathematics and its Applications. Boca Raton, FL: Chapman & Hall/CRC, 2008.

Index

(Γ, v, d)-volcano, 29
(π, j)–CM-type of K, 3
$\mathbf{Mod}_{\pi,K_0}$, 9
$\mathbf{Mod}_{\pi}$, 6
$\Phi_{\pi,j}$–positive-imaginary, 4
$\mathbf{Del}_q$, 6
ℓ-isogeny, vi
$\mathbf{Id}_{\pi,K_0}$, 9
$\mathbf{Id}_{\pi}$, 2
μ-isogeny, 3, 4, 10, 85
μ-isogeny $\tau \to \tau'$, 15
μ-isogeny graph, 28
$\mathbf{Ord}_{\mathbb{C},K_0}$, 9
$\mathbf{Ord}_{\mathbb{C},\pi,K_0}$, 10
$\mathbf{Ord}_{\mathbb{C},\pi}$, 10
$\mathbf{Ord}_{\mathbb{C},g}$, 9
$\mathbf{Ord}_{\mathbb{F}_q}$, 2
$\mathbf{Ord}_{\pi,K_0}$, 9
$\mathbf{Ord}_{\pi}$, 2
$\mathbf{PId}_{\pi,K_0}$, 9
$\mathbf{PId}_{\pi}$, 4
$\mathbf{POrd}_{\mathbb{C},K_0}$, 10
$\mathbf{POrd}_{\mathbb{F}_q}$, 3
$\mathbf{POrd}_{\pi,K_0}$, 9
$\mathbf{POrd}_{\pi}$, 3
j-invariant, vi

abelian scheme, 5
abelian variety, 1
ascending edge, 30
ascending isogeny, 31
associated $\mathbb{Z}$-bilinear form, 7
associated sesquilinear form, 7
Atkin prime, 80

Baily-Borel compactification, 12

Chinese-and-match algorithm, 81
CM-field, 2
CM-type, 2
complex conjugation, 2
conductor, 30
coprime fractional ideal, 45

degree, 1
descending edge, 30
descending isogeny, 31
division ideal, 84
dual abelian variety, 1
dual ideal, 4
dual module, 7
dual morphism, 4

efficiently computable endomorphism, 84
Elkies prime, 80
elliptic curve, vi
explicit endomorphism, 84

Hilbert modular form, 11
Hilbert modular function, 12
Hilbert modular polynomials, 14, 87
Hilbert modular variety, 12
horizontal edge, 30
horizontal isogeny, 31
Humbert surface, 22

Igusa-Clebsch invariants, 21
isogeny, vi, 1
isogeny graphs, vii
isomorphic μ-isogenies, 18, 32

Jacobian, vi

match-and-sort algorithm, 81
maximal real multiplication, 9
modular ideal, 87
modular map, 22
modular poynomial, vi
Mumford representation, 82

non-μ-part of the real conductor, 30

ordinary Weil q-number, 2

Picard group, 1
polarisation, 2, 4, 8
preserves the action of R, 5
preserves the notion of dual, 4
preserves the notion of polarisation, 4
principal polarisation, 2, 4, 77

real conductor, 30
real conductor locally at μ, 30
real conductor of a connected component, 31
RM invariants, 86
RM isomorphism invariants, vii, 14, 16
Rosati involution, 78

semi-balanced, 7

Serre-Tate lift, 5, 6
sesquilinear, 7
set of Hilbert modular polynomials, 15
Shimura class group, 30
Siegel upper half space, 21
symmetric Hilbert modular forms, 12, 22

trace of Frobenius, 78

vanilla, 77
volcanic prime, 80

weight function, 11

Summary

This thesis is primarily concerned with topics in and around the study of isogenies of abelian varieties. The precise definitions of both abelian varieties and isogenies are unfortunately beyond the scope of this summary, but we aim to give the reader an intuitive notion of both.

The most common example of an abelian variety that occurs in number theory is that of an *elliptic curve*. Let us consider the equation

$$E : y^2 = x^3 - x + 1.$$

This equation has a solution $(x, y) = (1, 1)$, because $1^2 = 1^3 - 1 + 1$. Equation E is an example of an elliptic curve, and if we plot it, then it looks like this:

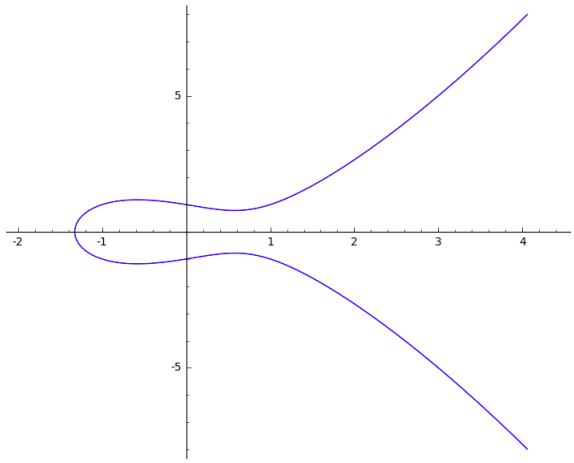

We can also see from the picture that we have a solution $(x, y) = (1, 1)$ to the equation, because the point with coordinates $(1, 1)$ lies on the curve. We can also spot other points that lie on the curve with integral (whole number) coordinates, such as $(x, y) = (0, -1)$, which also then give a solution to equation E.

Having found two points on the curve with integral coordinates, we can find more: in this example, drawing a straight line between the points $(0, -1)$ and $(1, 1)$ yields the following picture:

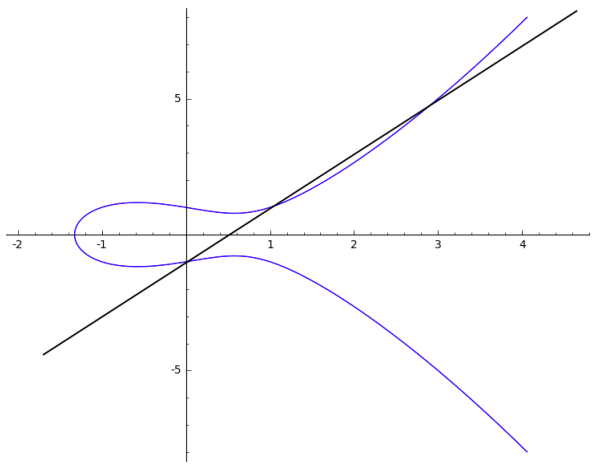

The straight line then intersects the curve in a third point $(x, y) = (3, 5)$, giving us a third solution to our equation. As it happens, this third point still has integral coordinates, although the same construction

starting from different points with integer coordinates could have yielded fractions (rational numbers). For example, the straight line passing through $(3, -5)$ and $(0, -1)$ intersects the curve in a third point $(-\frac{11}{9}, \frac{17}{27})$. However, this construction will never yield an irrational number like π or e!

In fact, in this example, infinitely many rational solutions (i.e. x and y can be written as fractions) to our equation can be found in this way - that is, by drawing a straight line between two points that we already know (or their reflections in the x-axis) and looking for a third point of intersection with the curve. Even better, in this example, it is possible to find *all* the rational solutions to our equation in this way as long as we use well-chosen points at each stage.

This is however, quite a 'special' example in this regard - an abelian variety is a geometric object that can be defined by polynomial equations like the one above, for which the rational solutions are related to one another in a prescribed way, for example by drawing straight lines and looking for intersection points as above. However, it's not always easy to find enough starting points to find all the solutions in this way, sometimes there are only finitely many solutions, and sometimes you don't even know if there should be a finite number or an infinite number of solutions. Also, for equations with higher degree than the example above (i.e. higher powers of x and y), or more variables (or more difficult in other ways), the relations between the points become more complicated.

Another way of finding solutions to such equations is via *isogenies*. An isogeny is a map between abelian varieties that 'preserves the geometric structure'. We explain by example what we mean by this: consider the map

$$(x, y) \mapsto (x - 11, y),$$

which sends the above equation $E : y^2 = x^3 - x + 1$ to

$$y^2 = x^3 - 33x^2 + 362x - 1319,$$

which looks like this:

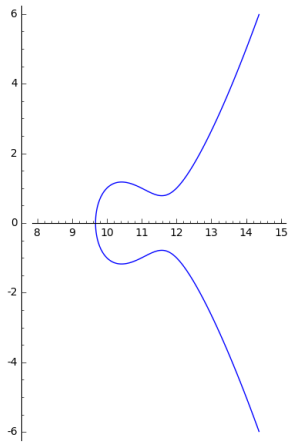

Above, we found three solutions $(0, -1)$, $(1, 1)$, and $(3, 5)$ to our equation E that lie on the same straight line. Note that if (x, y) is a solution to equation E, then $(x + 11, y)$ will be a solution to the new equation. So, we can use our map to compute three corresponding solutions to the new equation:

$$(0, -1) \mapsto (11, -1)$$
$$(1, 1) \mapsto (12, 1)$$
$$(3, 5) \mapsto (14, 5).$$

If you plot them, they still all lie on a straight line:

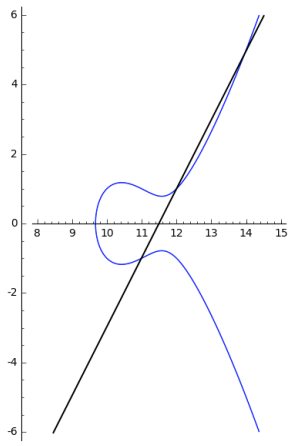

So the geometric relationship between these solutions is somehow preserved. This map is called an isogeny because this happens.

So, an isogeny from an abelian variety A to an abelian variety B not only maps points on A to points on B but preserves the relations between those points. Most importantly, if you are given an abelian variety A and asked to find points on A or relations between those points, it may be easier to look for an abelian variety B on which you can easily spot points, and an isogeny from B to A.

The take home message is: we are interested in knowing, given two abelian varieties, whether there exists an isogeny from one to the other. Normally, we are checking a bit more: whether or not there exists an isogeny of a certain type (for mathematicians: in the case of elliptic curves, this type is the degree). In Chapter 2, we give an algorithm to do this, which we have implemented for some 'small' abelian varieties. (Here 'small' means abelian varieties coming from genus 2 curves, which will be explained shortly.)

Another approach to help understand, given two abelian varieties, whether or not there exists an isogeny between them is to make a diagram of the information, called an isogeny graph. An isogeny graph is a diagram with: nodes labelled as abelian varieties, and an arrow between two nodes if there is an isogeny (of a certain prescribed type) from one node to the other.

For example, represent the equation $E : y^2 = x^3 - x + 1$ as a white node, and the equation $E' : y^2 = x^3 - 33x^2 + 362x - 1319$ as a black node. We saw already that there exists an isogeny from E to E' given by $(x, y) \mapsto (x - 11, y)$. Also, there exists an isogeny from E' to E given by $(x, y) \mapsto (x + 11, y)$ so part of our diagram would look like this:

We could also draw one undirected line instead of the two arrows, giving the following diagram:

In Chapter 3, we show that for our type of isogenies an isogeny graph [1] of abelian varieties consists of *volcano graphs*, an example of which is below:

[1] In this thesis, we also equate some nodes (for mathematicians: we identify isomorphic nodes), and E and E' would actually be represented by the same node. However, there do exist many graphs with lots of nodes even after equating some of the nodes.

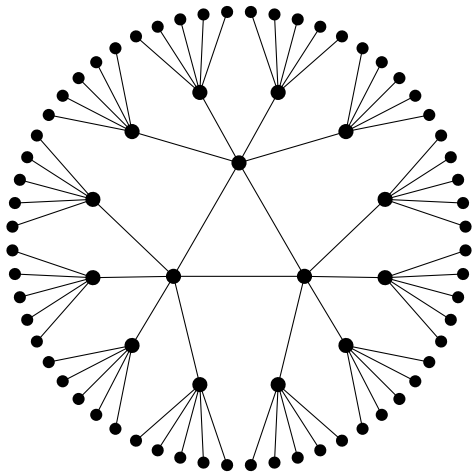

This type of graph has become known as a volcano as it resembles the bird's eye view of a volcano: the triangle in the centre is the 'rim' of the volcano (this could be replaced by any n-gon), and the lines going away from the rim is the 'lava' flowing down to the ground.

In Chapter 4 we study solutions of equations defining *curves of genus 2*. Except for some special cases, a genus 2 curve is given by an equation $y^2 = f(x)$, where $f(x)$ is a degree 5 or 6 polynomial (this means that the highest power of x that appears in $f(x)$ is 5 or 6). For example the equation

$$y^2 = x^5 + 1$$

represents a genus 2 curve. It looks like this:

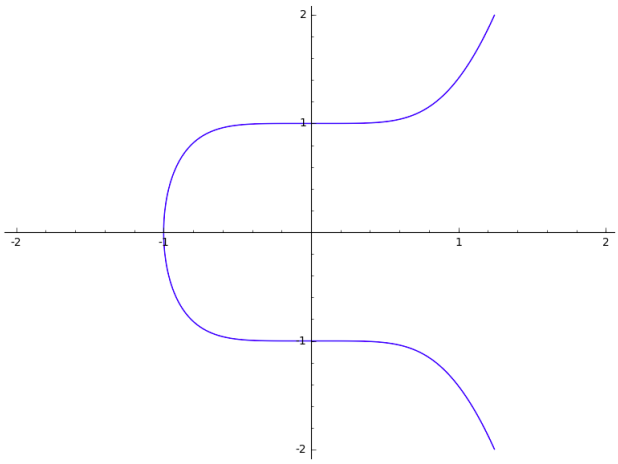

We can associate an abelian variety to any genus 2 curve; the study of genus 2 curves lies within the study of abelian varieties (in some sense). Also, instead of only looking for solutions of the equation given by whole numbers or fractions, we choose a prime number, say 101, and try to find integer coordinates (x, y) on the curve such that x^5+1-y^2 is divisible by 101, e.g. $x = 6$ and $y = 0$. One can count the number of choices for (x, y) with $0 \leq x, y < 101$ that yield x^5+1-y^2 divisible by 101 just by listing every possibility for x and y and checking whether you get a solution (in this case there are 97 solutions). However, if the prime is not 101, but 115792089237316195423570985008687907853269984665640564039457584007913129640233, then just counting all the solutions in this way cannot be done by modern computers, and many cryptographic protocols are based on the difficulty of this kind of problem. However, sometimes it is possible to count more efficiently by using the abelian variety structure. In Chapter 4, we give an efficient algorithm to count all the solutions (for a given large prime) for equations defining certain genus 2 curves. Chapter 4 is joint work with Sean Ballentine, Aurore Guillevic, Elisa Lorenzo-Garcia, Maike Massierer, Ben Smith, and Jaap Top.

Samenvatting

Dit proefschrift gaat hoofdzakelijk over isogenieën en abelse variëteiten. De precieze definities van de twee begrippen zijn helaas te geavanceerd voor deze samenvatting, maar wij proberen om wat intuïtie te geven voor beide concepten.

Het meest voorkomende voorbeeld van een abelse variëteit die zich in getaltheorie voordoet is een *elliptische kromme*. Beschouw de vergelijking

$$E : y^2 = x^3 - x + 1.$$

Deze vergelijking heeft een oplossing $(x, y) = (1, 1)$, want $1^2 = 1^3 - 1 + 1$. De vergelijking E is een voorbeeld van een elliptische kromme, en als we deze plotten krijgen we:

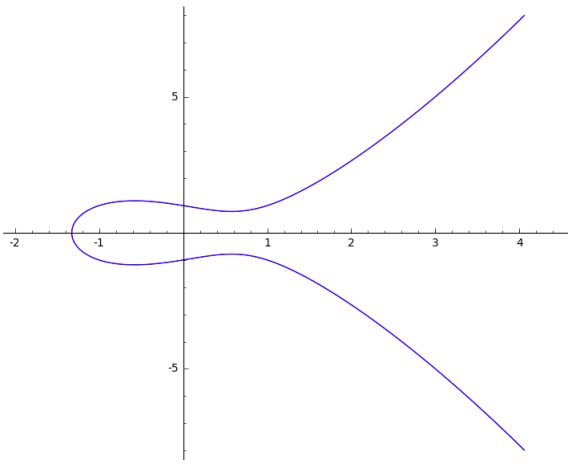

Wij zien ook direct uit de grafiek dat er een oplossing $(x, y) = (1, 1)$ is, omdat het punt met coördinaten $(1, 1)$ op de kromme ligt. Wij kunnen ook andere punten met geheeltallige coördinaten die op de kromme liggen nu zien, zoals $(x, y) = (0, -1)$, en die geven nog meer oplossingen van de vergelijking E.

Nu dat wij twee punten met geheeltallige coördinaten op de kromme hebben gevonden, kunnen wij er meer vinden: in dit voorbeeld tekenen wij een rechte lijn tussen de punten $(0, -1)$ en $(1, 1)$. Dit geeft de volgende grafiek:

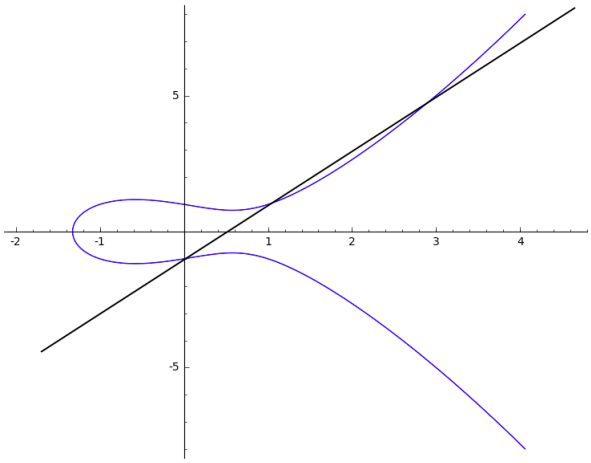

De rechte lijn snijdt de kromme in een derde punt $(x, y) = (3, 5)$, dus wij krijgen een derde oplossing van onze vergelijking. Het derde punt heeft weer geheeltallige coördinaten, maar dezelfde constructie kan ook breuken geven. Bijvoorbeeld, de rechte lijn die door $(3, -5)$ en $(0, -1)$ gaat, heeft een derde snijpunt met de kromme op $(-\frac{11}{9}, \frac{17}{27})$. Aan de andere kant kan deze constructie nooit een irrationaal getal zoals π of e geven!

In dit voorbeeld kunnen wij op deze manier, dat is door het tekenen van een rechte lijn tussen twee punten die wij al gevonden hebben (of hun reflecties in de x-as) en het zoeken naar een derde snijpunt, een oneindig aantal rationale oplossingen vinden, dat wil zeggen: x en y kunnen als breuk geschreven worden.

Beter zelfs: in dit voorbeeld is het mogelijk om *alle* rationale oplossingen van onze vergelijking op deze manier te vinden, zolang wij op elk moment de beste punten kiezen om te gebruiken.

Aan de andere kant, dit voorbeeld is best wel 'speciaal' – een abelse variëteit is een meetkundig object dat door polynomen gedefinieerd kan worden (de vergelijkingen hierboven zijn polynomen) waarvoor bovendien de rationale oplossingen een voorgeschreven relatie hebben, bijvoorbeeld door het tekenen van rechte lijnen die de kromme in drie rationale punten snijden. Vaak is het niet mogelijk om de beginpunten te vinden, of er is maar een eindige hoeveelheid rationale oplossingen, en soms weten wij niet of er een eindige of oneindige hoeveelheid rationale oplossingen is. Daarnaast worden voor vergelijkingen met een hogere graad dan ons voorbeeld (i.e. hogere machten van x en y), of meer variabelen, de relaties tussen de punten gecompliceerder.

Een andere manier van oplossingen vinden is via *isogenieën*. Een isogenie is een afbeelding tussen abelse variëteiten die 'de meetkundige structuur bewaart' (voor wiskundigen: de groepsstructuur). Wij leggen met een voorbeeld uit wat dit betekent: beschouw de afbeelding

$$(x, y) \mapsto (x - 11, y),$$

die de vergelijking $E : y^2 = x^3 - x + 1$ van hierboven naar

$$y^2 = x^3 - 33x^2 + 362x - 1319$$

stuurt. De grafiek van deze nieuwe vergelijking ziet er uit als:

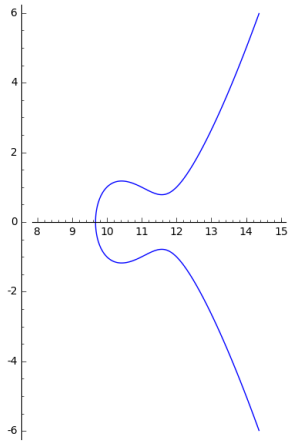

Wij vonden hierboven drie oplossingen $(0, -1)$, $(1, 1)$, en $(3, 5)$ van onze vergelijking E die op dezelfde rechte lijn liggen. Merk op dat als (x, y) een oplossing van de vergelijking E is, dan $(x+11, y)$ een oplossing van de nieuwe vergelijking wordt. Dus kunnen wij onze afbeelding gebruiken om drie overeenkomende oplossingen van de nieuwe vergelijking te berekenen:

$$\begin{aligned}(0, -1) &\mapsto (11, -1)\\ (1, 1) &\mapsto (12, 1)\\ (3, 5) &\mapsto (14, 5).\end{aligned}$$

Laten wij hen plotten en zien we dat zij nog steeds op een rechte lijn liggen:

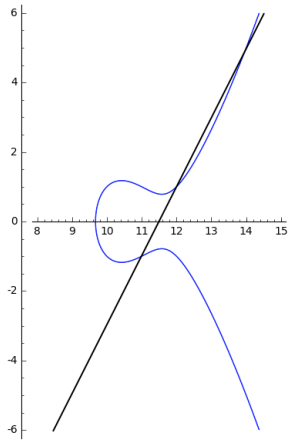

Dus de meetkundige relatie tussen deze oplossingen is in zekere zin behouden. De afbeelding is een *isogenie* omdat dit gebeurt.

Wij hebben gezien dat een isogenie van een abelse variëteit A naar een abelse variëteit B niet alleen maar punten van A naar punten van B stuurt, maar ook de relatie tussen de punten behoudt. Het is belangrijk om in te zien dat als een abelse variëteit A gegeven is en u bent gevraagd om punten van A te vinden, of relaties tussen de punten, het dan makkelijker kan zijn om naar een abelse variëteit B te zoeken waar het makkelijk is om punten en relaties te vinden, en een isogenie van B naar A te geven.

De boodschap is: wanneer twee abelse variëteiten zijn gegeven dan willen wij weten of er een isogenie bestaat van de ene naar de andere. Normaalchecken wij nog een beetje meer: of er een isogenie bestaat van een specifiek type (voor wiskundigen: in het geval van elliptische krommen is het type de graad). In hoofdstuk 2 geven wij een algoritme om dit te doen, dat wij voor een paar 'kleine' abelse variëteiten ook hebben geimplementeerd. (Hier betekent 'klein' abelse variëteiten die van geslacht twee krommen vandaan komen, dit zullen wij straks uitleggen.)

Nog een manier die kan helpen om te zien of er een isogenie tussen twee abelse variëteiten bestaat, is door een diagram te maken van de informatie; dit heet een isogenieëngraaf. Een isogenieëngraaf is een diagram dat bestaat uit knopen gemarkeerd als abelse variëteiten, met steeds een pijl van een knoop naar een andere als er een isogenie (van een gegeven type) bestaat van de ene abelse variëteit naar de andere.

Bijvoorbeeld, neem voor de vergelijking $E : y^2 = x^3 - x + 1$ een witte knoop, en voor de vergelijking $E' : y^2 = x^3 - 33x^2 + 362x - 1319$ een zwarte knoop. Wij hebben al gezien dat er een isogenie van E naar E' is gegeven door $(x, y) \mapsto (x - 11, y)$. Er bestaat ook een isogenie van E' naar E gegeven door $(x, y) \mapsto (x + 11, y)$ dus een deel van ons diagram ziet er als volgt uit:

Wij zouden ook een ongerichte lijn kunnen tekenen in plaats van de twee pijlen, zodat het diagram wordt:

In hoofdstuk 3 bewijzen wij dat voor onze type isogenieën bestaat de isogenieëngraaf [2] uit *vulkanen*. Een vulkaan ziet er bijvoorbeeld als volgt uit:

[2] In dit proefschrift stellen we sommige knopen gelijk (voor de wiskundigen: wij stellen isomorfe knopen gelijk), en E en E' worden eigenlijk door dezelfde knoop gerepresenteerd. Maar er bestaan wel grafen met heel veel knopen, zelfs na dit gelijkstellen.

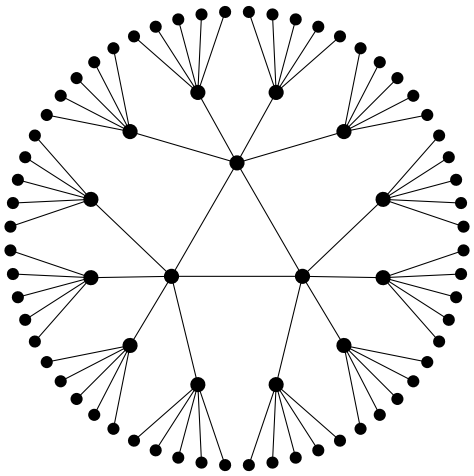

Dit type graaf wordt vulkaan genoemd omdat het eruitziet als het bovenaanzicht van een vulkaan: de driehoek in het centrum is de 'rand' van de vulkaan (deze kan ook door iedere n-hoek vervangen worden), en de lijnen die weg gaan van de rand stellen de lava voor dat naar beneden stroomt.

In hoofdstuk 4 bestuderen wij oplossingen van vergelijkingen die *geslacht 2 krommen* definiëren. Buiten wat speciale gevallen wordt een geslacht 2 kromme door een vergelijking $y^2 = f(x)$ gegeven, waar $f(x)$ een polynoom is van graad 5 of 6 (dit betekent dat de hoogste macht van x die in $f(x)$ voorkomt 5 of 6 is). Bijvoorbeeld de vergelijking

$$y^2 = x^5 + 1$$

is een geslacht 2 kromme. Deze ziet er uit als:

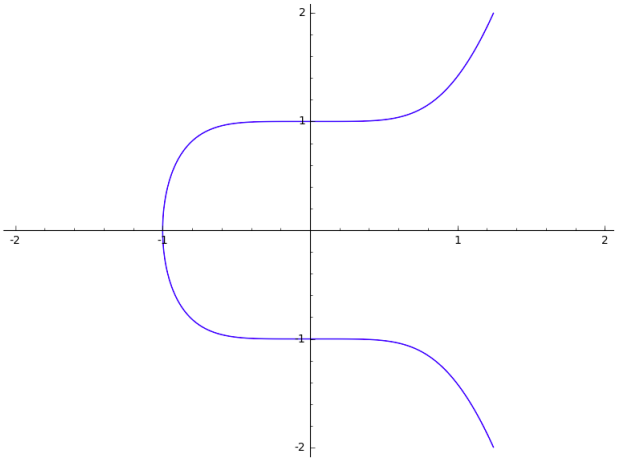

Wij kunnen met elke geslacht 2 kromme een abelse variëteit associëren; de studie naar geslacht 2 krommen is een substudie van die naar abelse variëteiten (in zekere zin). In plaats van alleen te zoeken naar oplossingen die zijn gegeven door gehele getallen of breuken, kiezen wij daarnaast een priemgetal, bijvoorbeeld 101, en proberen wij geheeltallige coördinaten (x, y) te vinden zodat $x^5 + 1 - y^2$ gedeeld kan worden door 101, bijvoorbeeld $x = 6$ en $y = 0$.

Wij kunnen het aantal mogelijkheden tellen voor (x, y) met $0 \leq x, y < 101$ zodat $x^5 + 1 - y^2$ gedeeld kan worden door 101 door elke optie voor x en y op te sommen en te checken of het een oplossing geeft (in dit geval er zijn 97 oplossingen). Aan de andere kant, als de priem niet 101 is, maar 115792089237316195423570985008687907853269984665640564039457584007913129640233, dan kan het op deze manier tellen van alle oplossingen niet door moderne computers gedaan worden. Cryptographische protocollen zijn op de moelijkheid van dit soort problemen gebaseerd. Maar soms is het efficiënter om te tellen door gebruik te maken van de structuur van de abelse variëteit. In hoofdstuk 4 geven wij een efficiënt algoritme om alle oplossingen (voor een gegeven grote priem) te tellen voor vergelijkingen van bepaalde geslacht 2 krommen. Hoofdstuk 4 is een samenwerking met Sean Ballentine, Aurore Guillevic, Elisa Lorenzo-Garcia, Maike Massierer, Ben Smith, en Jaap Top.

Résumé

Dans cette thèse nous étudions les isogénies entre variétés abéliennes. Les définitions précises de variété abélienne et d'isogénie dépassent malheureusement le cadre de ce résumé, mais nous essayons d'en donner une idée intuitive.

L'exemple le plus commun d'une variété abélienne dans la théorie des nombres est celui de *courbe elliptique*. Considérons l'équation

$$E : y^2 = x^3 - x + 1.$$

Cette équation admet $(x, y) = (1, 1)$ pour solution car $1^2 = 1^3 - 1 + 1$. L'équation E est un exemple d'une courbe elliptique, et si nous la dessinons, elle ressemble à ceci:

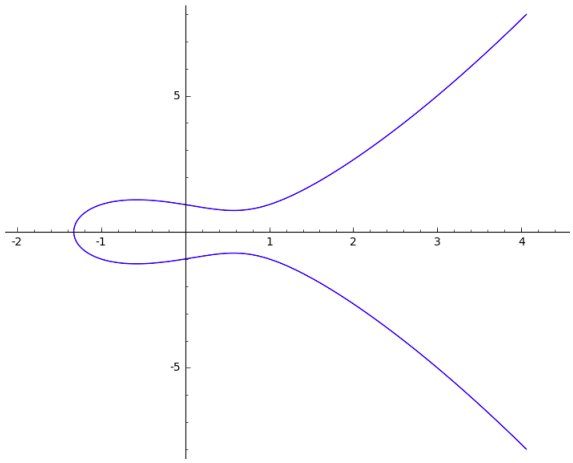

Il est évident sur l'image que $(x, y) = (1, 1)$ est une solution de l'équation, parce que le point de coordonnées $(1, 1)$ se trouve sur la courbe. Nous pouvons aussi trouver d'autres points sur la courbe avec des coordonnées entières (nombres entiers), telles que $(x, y) = (0, -1)$, qui donnent solutions de l'équation E.

Ayant trouvé deux points sur la courbe avec des coordonnées entières, nous pouvons les utiliser pour en trouver d'autres: dans notre exemple, traçons la ligne droite entre les points $(0, -1)$ et $(1, 1)$; nous obtenons l'image suivante:

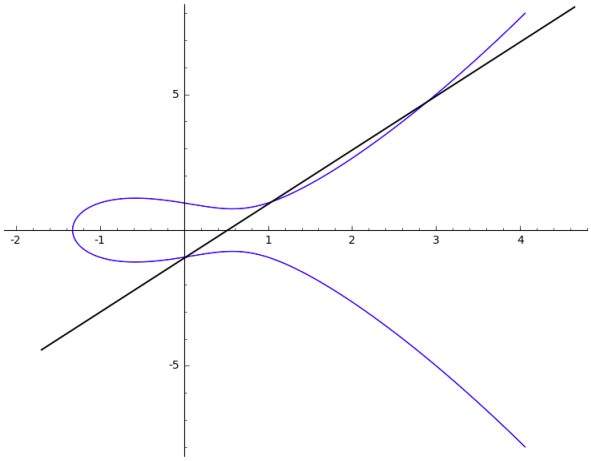

La droite coupe la courbe en un troisième point $(x, y) = (3, 5)$, ce qui nous fournit une troisième solution à notre équation. Dans cet exemple, ce troisième point a lui aussi des coordonnées entières, bien que la même construction à partir de points différents aurait pu donner des fractions (nombres rationnels). Par exemple, la ligne droite passant par $(3, -5)$ et $(0, -1)$ coupe la courbe en un troisième point $(-\frac{11}{9}, \frac{17}{27})$. Cependant, cette construction ne donnera jamais un nombre irrationnel comme π ou e!

En fait, dans cet exemple, il y a une infinité de solutions rationnelles (c'est-à-dire que x et y sont des fractions) à notre équation qui peuvent être trouvées de cette façon - en dessinant une ligne droite entre deux points connu (ou leurs réflexions par rapport à l'axe x) et en cherchant le troisième point d'intersection avec la courbe. Mieux encore, dans cet exemple, il est possible de trouver *toutes* les solutions rationnelles à notre équation de cette manière pour peu que l'on utilise des points bien choisis à chaque étape.

Cet exemple est, cependant, spécial à cet égard - une variété abélienne est un objet géométrique qui peut être défini par des équations polynômiales comme celle ci-dessus, pour lesquelles les solutions rationnelles ont des relations prescrites, par exemple en dessinant des lignes droites et en recherchant les points d'intersection comme ci-dessus. En général, ce n'est pas toujours possible à trouver des points de départ pour cette construction, ou parfois il y a seulement un nombre fini de solutions, et parfois on ne sait pas s'il y a un nombre fini ou infini de solutions. De même, pour des équations ayant un degré supérieur à l'exemple ci-dessus (c'est-à-dire des puissances supérieures de x et y) ou plus de variables (ou qui sont plus compliquées d'une autre manière), les relations deviennent plus compliquées.

Une autre façon de trouver des solutions à de telles équations est d'utiliser des *isogénies*. Une isogénie est une application entre variétés abéliennes qui «préserve la structure géométrique». Expliquons ceci sur un exemple: considérer l'application

$$(x, y) \mapsto (x - 11, y),$$

qui envoie les solutions de l'équation $E : y^2 = x^3 - x + 1$ vers des solutions de

$$y^2 = x^3 - 33x^2 + 362x - 1319,$$

qui ressemble à:

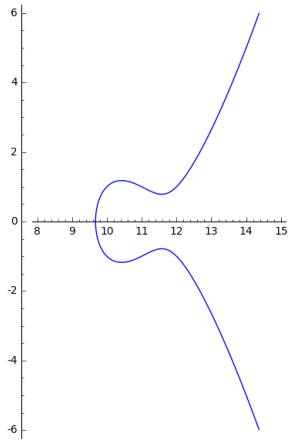

Nous avons trouvé ci-avant trois solutions $(0, -1)$, $(1, 1)$ et $(3, 5)$ à notre équation E et celles-si se trouvent sur la même droite. Notons que si (x, y) est une solution de l'équation E, alors $(x + 11, y)$ sera bien une solution de la nouvelle équation. Ainsi, on peut utiliser l'application pour exprimer les trois solutions correspondantes de la nouvelle équation:

$$\begin{aligned}(0, -1) &\mapsto (11, -1)\\ (1, 1) &\mapsto (12, 1)\\ (3, 5) &\mapsto (14, 5).\end{aligned}$$

Si l'on ajoute ces trois solutions sur le dessin ci-dessus, on constate que les trois points sont alignés:

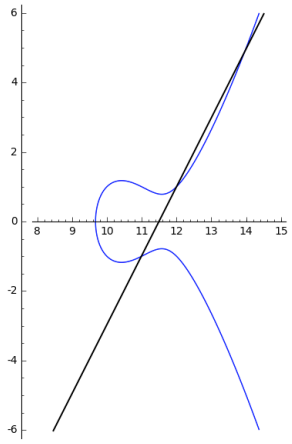

Alors, parce que l'application est une *isogénie*, la relation géométrique (i.e. le fait d'être alignées) entre ces solutions est préservée.

À l'aide d'une isogénie d'une variété abélienne A vers une variété abélienne B, et étant donnés des points sur A, on peut trouver des points sur B. L'isogénie préserve de plus les relations entres ces points. Plus important encore, si l'on se donne une variété abélienne A et que l'on veut trouver des points sur A ou des relations entre ces points, il peut être plus facile de chercher une variété abélienne B sur laquelle on peut plus facilement trouver des points et une isogénie de B à A.

Le message principal est donc: il est intéressant de savoir, étant donné deux variétés abéliennes, s'il existe une isogénie entre elles. En fait, on pose une question un peu plus précise: existe-t-il une isogénie d'un type spécifique entre deux variétés abéliennes données? (Pour les mathématiciennes, dans le cas des courbes elliptiques, le type d'une isogénie est son degré.) Dans le chapitre 2, nous donnons un algorithme qui répond à cette question. Nous avons de plus implémenté celui-ci pour certaines «petites» variétés abéliennes. (Ici «petit» signifie des variétés abéliennes qui viennent des courbes de genre 2, qui nous définissons plus loin.)

Une autre approche pour aider à voir, étant données deux variétés abélienne, s'il existe ou pas une isogénie entre elles, est de faire un diagramme de la situation, appelé graphe d'isogénie. Un graphe d'isogénie est un diagramme dont les sommets sont des variétés abéliennes, et dans lequel deux sommets sont reliés par une arête s'il y a une isogénie (d'un type donné) entre eux.

Par exemple, représentons l'équation $E : y^2 = x^3 - x + 1$ par un sommet blanc, et l'équation $E' : y^2 = x^3 - 33x^2 + 362x - 1319$ par un sommet noir. Nous avons déjà vu qu'il existe une isogénie de E vers E' donnée par $(x, y) \mapsto (x - 11, y)$. De plus, il existe une isogénie de E' vers E donnée par $(x, y) \mapsto (x + 11, y)$. Une partie de notre graphe ressemblerait alors à:

Nous aurions pu aussi dessiner une arête non orientée au lieu des deux arêtes orientées, ce qui donnerait le graphe suivant:

Au chapitre 3, nous prouvons que pour notre type d'isogénie, le graphe d'isogénie [3] se compose de volcans, dont un exemple est ci-dessous:

[3] Dans cette thèse, nous aussi assimilons des sommets (pour les mathéciennes: nous identifions des sommets isomorphiques), et en fait E et E' seraient représenté par la même sommet. Cependant, c'éxiste des graphes avec beaucoup de sommets même après en idenfier un part des sommets.

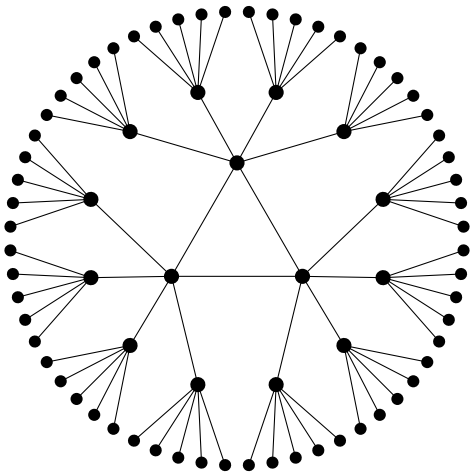

Ce type de graphe s'appelle un volcan car il ressemble à une vue aérienne d'un volcan: le triangle au centre correspond au «cratère» du volcan (on pourrait le remplacer par n'importe quel n-gone) et les lignes qui s'éloignent du cratère sont la «lave» qui coule vers le sol.

Au chapitre 4, nous étudions des solutions d'équations définissant des *courbes de genre 2*. En dehors de quelques cas particuliers, une courbe de genre 2 est donnée par une équation $y^2 = f(x)$, où $f(x)$ est un polynôme de degré 5 ou 6 (c'est-à-dire que la plus grande puissance de x apparaissant dans $f(x)$ est x^5 ou x^6). Par exemple l'équation

$$y^2 = x^5 + 1$$

représente une courbe de genre 2. Elle ressemble à ceci:

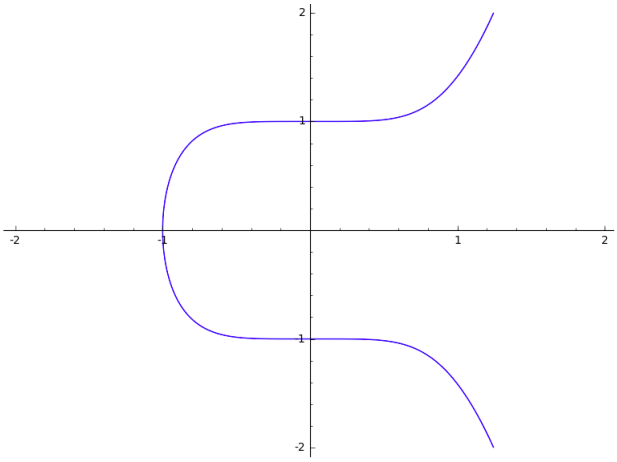

On peut associer une variété abélienne à n'importe quelle courbe de genre 2; l'étude des courbes de genre 2 se ramène donc à l'étude de certaines variétés abéliennes (dans un certain sens). Au lieu de chercher seulement des solutions de l'équation donnée par des nombres entiers ou des fractions, on choisit un nombre premier, disons 101, et on essaie de trouver des points à coordonnées entières (x, y) tels que $x^5 + 1 - y^2$ est divisible par 101, par exemple $x = 6$ et $y = 0$. On peut compter le nombre des choix de (x, y) avec $0 \leq x, y < 101$ qui donnent $x^5 + 1 - y^2$ divisible par 101, simplement en énumérant toutes les possibilités pour x et y et en vérifiant si on a une solution (dans ce cas, il y a 97 solutions). Toutefois, si le premier n'est pas 101, mais 1157920892373161954235709850086879078532699846656405640394575840079131296402 33, alors le comptage de toutes les solutions ne peut pas être effectué de cette manière par les ordinateurs modernes, et de nombreux protocoles cryptographiques sont basés sur la difficulté de ce type de problème. Cependant, parfois c'est possible à compter le nombre de solutions si on utilise la structure de variété abélienne. Au chapitre 4, nous donnons un algorithme efficace pour compter toutes les solutions (pour un grand nombre premier donné) pour les équations définissant des certaines courbes de genre 2. Le chapitre 4 est un travail en commun avec Sean Ballentine, Aurore Guillevic, Elisa Lorenzo-Garcia, Maike Massierer, Ben Smith et Jaap Top.

Acknowledgements

First of all, I would like to thank my supervisor, Marco Streng, for his patient guidance, for introducing me to the world of complex multiplication, and for suggesting research problems that not only kept my focus and interest for the duration of my PhD, but that led naturally to many new and interesting topics beyond that. I would also like to thank my promotors Peter Stevenhagen and Andreas Enge for their ongoing support and guidance.

I would like to thank Bas, Lenny, and David for helping me to understand canonical lifts of abelian varieties.

I would like to thank my new husband, former colleague, and former office mate for supporting me every step of the way. I would like to thank Mima, who welcomed me into the Leiden family on my first day, and who radiates joy to all around her.

I would like to thank my family: mum, dad, Jacob, Noel, Fiona, Isaac, Thea, Simon, and Hannah, who have always supported me in everything I do.

I would like to thank Richard Griffon for his invaluable help with the 'résumé'.

I would of course also like to thank my wonderful colleagues in Leiden and Bordeaux, without all of whom the PhD experience would have been very different and a lot less 'gezellig'. A special mention has to go to my first office mates Yan and Qijun, and to Julian and Steven, for helping me survive through 'Nederlandse donderdagen' and for the countless times we saved the world together.

Curriculum Vitae

Chloe Martindale was born in Huntingdon in the United Kingdom on 13th September 1990.

In 2004 she started at Chetham's School of Music, where she later obtained her A levels. Following Chetham's, she went on to read mathematics at the University of Oxford, matriculating in 2009. She obtained a first class honours bachelor degree from the University of Oxford in 2013.

She then remained at Oxford a further year in the masters programme for mathematics, from which she graduated with a first class masters degree in 2014. She wrote her masters' thesis on *Elliptic curves and Jacobians of genus two curves* under the supervision of Prof. dr. Victor Flynn.

In September 2014 she started her PhD studies with an Erasmus Mundus ALGANT-doc fellowship under the supervision of Marco Streng, Andreas Enge, and Peter Stevenhagen.

She has been working as a Postdoc at the Technische Universiteit Eindhoven since March 2017.