The Pennsylvania State University
The Graduate School

ABELIAN VARIETIES AND DECIDABILITY IN NUMBER THEORY

A Dissertation in
Mathematics
by
Caleb Springer

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

August 2021

The dissertation of Caleb Springer was reviewed and approved by the following:

Anne Kirsten Eisenträger
Professor of Mathematics
Francis R. Pentz and Helen M. Pentz Professor of Science
Dissertation Advisor
Chair of Committee

Mihran Papikian
Professor of Mathematics

Linda Westrick
Assistant Professor of Mathematics

Martin Furer
Professor of Computer Science and Engineering

Carina Curto
Professor of Mathematics
Chair of Graduate Program

Abstract

This dissertation consists of two parts, both of which are focused upon problems and techniques from algebraic number theory and arithmetic geometry. In the first part, we consider abelian varieties defined over finite fields, which are a key object in cryptography in addition to being inherently intriguing in their own right. First, generalizing a theorem of Lenstra for elliptic curves, we present an explicit description of the group of rational points $A(\mathbb{F}_q)$ of a simple abelian variety A over a finite field $\mathbb{F}_q$ as a module over the endomorphism ring $\mathrm{End}_{\mathbb{F}_q}(A)$, under some technical conditions. Next, we present an algorithm for computing $\mathrm{End}_{\mathbb{F}_q}(A)$ in the case of ordinary abelian varieties of dimension 2, again under certain conditions, building on the work of Bisson and Sutherland. We prove the algorithm has subexponential running time by exploiting ideal class groups and class field theory.

In the second part, we turn our attention to questions of decidability and definability for algebraic extensions of $\mathbb{Q}$, in the vein of Hilbert's Tenth Problem and its generalizations. First, we show that a key technique for proving undecidability results fails for "most" subfields $L \subseteq \overline{\mathbb{Q}}$. More specifically, we view the set of subfields of $\overline{\mathbb{Q}}$ as a topological space, and prove there is a meager subset containing all fields $L \subseteq \overline{\mathbb{Q}}$ for which the ring of integers $\mathcal{O}_L$ is existentially or universally definable in L. Finally, we present explicit families of infinite algebraic extensions of $\mathbb{Q}$ whose first-order theory is undecidable. This is achieved by leveraging the unit groups of totally imaginary quadratic extensions of totally real fields, building on the work of Martínez-Ranero, Utreras and Videla.

Table of Contents

List of Symbols

$\mathbb{C}$	The field complex numbers
$\mathbb{R}$	The field of real numbers
$\mathbb{Q}$	The field of rational numbers
$\overline{K}$	The algebraic closure of a field K
$\mathbb{Z}$	The ring of (rational) integers
$\mathbb{N}$	The set of natural numbers (positive integers)
$\mathcal{O}_K$	The ring of integers of an algebraic extension K of $\mathbb{Q}$
$\mathbb{F}_q$	The finite field containing q elements
$\mathrm{char}(k)$	The characteristic of the field k
$\mathrm{End}_k(A)$	The (k-rational) endomorphism ring of an abelian variety A defined over k

Acknowledgments

First, I would like to thank my advisor Kirsten Eisenträger for her constant support and wise direction throughout my time as a graduate student. As she guided me through researching fascinating problems, she also acted as a role model for clear communication, diligence, and professionalism. I also want to thank my collaborators Russell Miller and Linda Westrick, who have expanded my horizons in mathematics. Additionally, I thank Jean-François Biasse, Gaetan Bisson, Arno Fehm, David Kohel, Stefano Marseglia, Chloe Martindale, Hector Pasten, Alexandra Shlapentokh, Andrew Sutherland, Tom Tucker, Xavier Vidaux, Benjamin Wesolowski, and Yuri Zarhin for their helpful comments, conversations, and suggestions throughout the time in which this research was completed.

Finally, I want to thank the people who have personally shaped and supported me: my wife, Kelsey; my parents, Jim and Carol; my sisters, Steph and Christine; my grandparents, Helen, Jim, Rae and Bill; and all of my friends and colleagues throughout graduate school. I especially thank my friends Sonny Arora, Travis Morrison and Krzysztof Pawelec for helping me enter and navigate the world of number theory.

The author was partially supported by National Science Foundation grants DMS-1056703, CNS-1617802, and CNS-2001470. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Material from Chapter 3 was first published in *The European Journal of Mathematics*, © 2021 Springer Nature. [96] Reproduced with permission from Springer Nature. `https://www.springer.com/gp/rights-permissions/obtaining-permissions/882`

Material from Chapter 4 was first published in *Journal of Number Theory*, Volume 202, September 2019, © 2019 Elsevier Inc. [94]. Reproduced with permission. `https://www.elsevier.com/about/policies/copyright#Author-rights`

Material from Chapter 6 was produced in collaboration with Kirsten Eisenträger, Linda Westrick and Russell Miller [31]. Following the tradition formalized by the American Mathematical Society, all authors are listed alphabetically. `http://www.ams.org/profession/leaders/CultureStatement04.pdf`

Material in Chapter 7 was first published in the *Proceedings of the Amer. Math. Soc.*, Volume 148, Number 11, November 2020, published by the American Mathematical Society © 2020 American Mathematical Society. [95] Reproduced with permission. `http://www.ams.org/publications/authors/ctp`.

Dedication

To my wife, Kelsey.

Chapter 1
Introduction

This dissertation considers problems from algebraic number theory and arithmetic geometry. These problems are of theoretical interest and have applications to cryptography. In the first part of the dissertation, we study abelian varieties defined over finite fields. In the second part, we solve problems in the intersection of logic and number theory related to Hilbert's Tenth Problem and its generalizations.

1.1 Abelian varieties defined over finite fields

Elliptic curves have proven to be ubiquitous within number theory for the sake of both theoretical and applied interests. For example, elliptic curves defined over finite fields have played a crucial role within cryptography, including both the currently-used Elliptic Curve Cryptography (ECC) and proposed post-quantum cyptosystems. The usefulness and beauty of elliptic curves comes from the fact that elliptic curves combine an abelian group and a smooth projective variety in one package, mixing together the structures found in algebra and geometry. This is the motivating example we will address in a more general context.

In the first part of this dissertation, we will consider abelian varieties over finite fields, which are the higher-dimensional analogues of elliptic curves. Moving from elliptic curves to the more general setting of abelian varieties of arbitrary dimension provides the opportunity for both theoretical and practical advances. In the context of cryptography, which enlists Jacobian varieties defined over finite fields, moving from elliptic curves to higher dimensional varieties allows one to work over a dramatically smaller field without lowering the level of security, as seen in genus 2-based cryptography [10] and possible improvements to supersingular isogeny-based cryptography [21].

1.1.1 The structure of the group of rational points

Given an abelian variety A defined over a field k, there are two important and interrelated objects of interest: (1) the *group of k-rational points* $A(k)$, and (2) the *k-rational endomorphism ring* $\mathrm{End}_k(A)$. The group $A(k)$ has the structure of a module over the ring $\mathrm{End}_k(A)$, and is the underlying group used for ECC (Elliptic Curve Cryptography) and its higher-dimensional generalizations. The work that follows will seek to understand and compute these objects. First, generalizing a result of Lenstra for elliptic curves, we describe the group of rational points under certain technical conditions.

Theorem 3.1.4. *For $g \geq 1$, let A be a simple abelian variety over $\mathbb{F}_q$ of dimension g with Frobenius endomorphism π. Write $K = \mathbb{Q}(\pi)$ and $R = \mathrm{End}_{\mathbb{F}_q}(A)$, and let Z be the center of R.*

(a) If $[K : \mathbb{Q}] = 2g$ and R is a Gorenstein ring, then

$$A(\mathbb{F}_{q^n}) \cong R/R(\pi^n - 1).$$

(b) If $(\pi^n - 1)Z$ is the product of invertible prime ideals in Z, then there is an isomorphism of Z-modules

$$A(\mathbb{F}_{q^n}) \cong (Z/Z(\pi^n - 1))^d,$$

where $d = 2g/[K : \mathbb{Q}]$. Moreover, this Z-module has exactly one left R-module structure, up to isomorphism. This R-module structure comes from the isomorphism of rings $R/R(\pi^n - 1) \cong \mathrm{M}_d(Z/Z(\pi^n - 1))$, and there is an isomorphism of R-modules

$$A(\mathbb{F}_{q^n})^d \cong R/R(\pi^n - 1).$$

As a corollary, we show that every simple ordinary isogeny class of abelian varieties over $\mathbb{F}_q$ contains an abelian variety A such that $A(\mathbb{F}_q)$ is a cyclic group; see Corollary 3.1.3.

1.1.2 Endomorphism ring computation

There has recently been great attention on computing endomorphism rings for abelian varieties defined over finite fields. For example, when constructing ordinary abelian varieties with a prescribed number of points for use in cryptography, the computation of endomorphism rings is a serious bottleneck [1, 30, 58]. If A is a simple ordinary abelian

variety over a finite field $\mathbb{F}_q$ with Weil q-integer π, then $\mathrm{End}(A)$ is isomorphic to an order in the field $\mathbb{Q}(\pi)$ and *computing the endomorphism ring* means determining this order.

Theorem 4.1.1. *There is a subexponential algorithm which, given an ordinary abelian variety A of dimension 2 over a finite field $\mathbb{F}_q$ satisfying certain technical conditions, computes the endomorphism ring* $\mathrm{End}\, A$.

The technical conditions are clearly and explicitly described in Chapter 4. The key idea for the algorithm, generalizing an algorithm of Bisson and Sutherland for elliptic curves [8], is the action of the class group $\mathrm{Cl}(\mathcal{O})$ on the set of all abelian varieties with endomorphism ring $\mathcal{O}$ via isogenies. Thus, computing isogenies probes the class group $\mathrm{Cl}(\mathcal{O})$. Using class field theory, we show that knowing $\mathrm{Cl}(\mathcal{O})$ is enough to determine the order $\mathcal{O}$ itself in the general case, thereby reducing the problem to simply computing certain isogenies.

1.2 Decidability and definability in number theory

In the second part of this dissertation, we turn our attention to decidability. Hilbert's Tenth Problem, in its original form, asks for an algorithm which takes as input a multivariable polynomial equation $f(x_1, \ldots, x_n) = 0$ with integer coefficients and outputs YES or NO depending on whether or not the equation has an integer solution $(a_1, \ldots, a_n) \in \mathbb{Z}^n$. Matiyasevich [66], building on earlier work by Davis, Putnam, and Robinson [23], proved that no such algorithm exists, i.e., Hilbert's Tenth Problem is *undecidable*. Since then, analogues of this problem have been studied by asking the same question for polynomial equations with coefficients and solutions in other recursive commutative rings R, under the name *Hilbert's Tenth Problem over R*. One of the most important unsolved questions in this area is Hilbert's Tenth Problem over the field of rational numbers $\mathbb{Q}$, and more generally over number fields.

1.2.1 A topological approach to undefinability

Definability is one of the key techniques used to prove undecidability results in general. If $\mathbb{Z}$ were existentially definable in $\mathbb{Q}$, then it would follow immediately that Hilbert's Tenth Problem over $\mathbb{Q}$ is undecidable. However, it is conjectured that $\mathbb{Z}$ is in fact not existentially definable in $\mathbb{Q}$, i.e., definability is an unusable technique in this context. Unfortunately, proving unconditionally that $\mathbb{Z}$ is not existentially definable in $\mathbb{Q}$ currently

appears to be out of reach. In fact, it is generally very difficult to find examples of fields $L \subseteq \overline{\mathbb{Q}}$ where existential undefinability is currently provable.

However, we can instead consider the set of all subfields of $\overline{\mathbb{Q}}$ simultaneously, and prove that for "most" subfields $L \subseteq \overline{\mathbb{Q}}$, the ring of integers $\mathcal{O}_L$ is existentially undefinable. To be precise, we view the set $\mathrm{Sub}(\overline{\mathbb{Q}})$ of subfields of $\overline{\mathbb{Q}}$ to be a topological space which is computably homeomorphic to Cantor space. With this topology, the notion of a meager set provides a way to describe which sets are very small, and we prove the following theorem.

Theorem 6.1.1. *The set of algebraic extensions K of $\mathbb{Q}$ for which $\mathcal{O}_K$ is existentially or universally definable is a meager subset of* $\mathrm{Sub}(\overline{\mathbb{Q}})$.

1.2.2 Undecidability of some infinite extensions of $\mathbb{Q}$

While Hilbert's Tenth Problem over $\mathbb{Q}$ currently remains unknown, the decidability of the entire first-order theory of $\mathbb{Q}$, or more generally an algebraic extension of $\mathbb{Q}$, is another interesting question in its own right. Julia Robinson proved that $\mathbb{Q}$ has undecidable first-order theory [80], and extended her result to every number field [81]. Rumely generalized this result further and proved that every global field is undecidable [83]. However, the problem of decidability remains open for infinite algebraic extensions of $\mathbb{Q}$ in general. In Chapter 7, we outline the current state of knowledge with a demonstration of examples, including both decidable and undecidable infinite extensions of $\mathbb{Q}$.

Before stating the following theorem, we recall some standard notation and terminology. Given a number field F, let $F^{(d)}$ be the compositum of all extensions of F of degree at most d, and let $F_{ab}^{(d)}$ be the maximal abelian subextension of $F \subseteq F^{(d)}$. Given an algebraic number $\alpha \in \overline{\mathbb{Q}}$ whose minimal polynomial over $\mathbb{Q}$ is $f(x)$, we say that α is *totally real* if every root of $f(x)$ in $\mathbb{C}$ is a real number. Similarly, α is *totally imaginary* if every root is non-real. We say that a field $L \subseteq \overline{\mathbb{Q}}$ is totally real if all of its elements are totally real, and totally imaginary if at least one element is totally imaginary.

Following recent work of Martínez-Ranero, Utreras and Videla, we prove the first-order undecidability for the following class of totally imaginary fields. The key idea involves exploiting the fact that the unit group $\mathcal{O}_K^\times$ is a finite-index subgroup of $\mathcal{O}_L^\times$, which allows us to reduce the problem to the totally real subfield and apply a modification of a method originally developed by Julia Robinson for totally real fields.

Theorem 7.3.6. *Let K be an infinite totally real extension of $\mathbb{Q}$ which is contained in $F_{ab}^{(d)}$ for some $d \geq 2$ and some number field F. Assume K contains all roots of a parametrized family of polynomials*

$$\{f_a(x) = x^n + p_{n-1}(a)x^{n-1} + \dots p_1(a)x + p_0(a) \mid a \in \mathbb{Z}_{>N_0}\}$$

where each $p_i(t) \in \mathbb{Z}[t]$ is a polynomial, $p_0(t) = \pm 1$ is constant and $p_j(t)$ is nonconstant for some $1 \leq j \leq n-1$. If L is any totally imaginary quadratic extension of K, then the first-order theory of L is undecidable.

Part I

Abelian Varieties over Finite Fields

Chapter 2
Background for Abelian Varieties

2.1 Motivation: Elliptic curves

Elliptic curves are one of the most important objects of modern number theory, and the most natural starting point for this segment of the dissertation because all of the results that follow are generalizations of theorems that were first discovered for elliptic curves. Since the work that follows can be technical at points, we start with a basic introduction to elliptic curves to establish a strong motivating example to build upon. The subsequent sections will require additional background knowledge in algebraic geometry (e.g., [62]), although the first section is aimed at more a general audience. For a complete introduction to elliptic curves, we refer to Silverman's textbook [93].

Definition 2.1.1. *An* elliptic curve E *over a field* k *is a smooth curve, i.e. a smooth projective variety of dimension 1, with a specified base point* P*.*

When the characteristic of the field $\mathrm{char}(k) \neq 2, 3$, it is possible to reduce to the case of *Weierstrass equations*, taking the concrete form

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

where $x^3 + ax + b$ is a polynomial with 3 distinct roots (in the algebraic closure $\overline{k}$). The latter condition is required simply to ensure that the curve is smooth. To be precise, the elliptic curve E is the *projectivization* of the equation above, which ultimately means there is another point, typically denoted ∞, on the elliptic curve in addition to all of the points (x, y) satisfying the given (affine) equation $y^2 = x^3 + ax + b$. See [93, III.1] for more details.

One of the main causes for interest in elliptic curves is their abundance of structure. Indeed, the fact that elliptic curves are projective varieties leads to a full array of

geometric methods and tools. However, the set of points on an elliptic curve is also importantly an *abelian group* in a natural sense. In the Weierstrass equation described above, the point ∞ is the group identity, and a full account of the group law is found in [93, III.2]. We pause to note that elliptic curves (and abelian varieties) also have a rich analytic theory, although this dissertation will focus on algebraic and geometric aspects of the theory.

To complete the introduction of elliptic curves, we also need to understand the relevant maps between elliptic curves, known as *isogenies* [93, III.4]. This requires the geometric notion of a *morphism*, which can be found in [93, I.3].

Definition 2.1.2. *Given two elliptic curves E_1, E_2 over a field k, an* isogeny *is a nonconstant morphism*

$$\phi : E_1 \to E_2$$

which maps the base point of E_1 to the base point of E_2. We say E_1 and E_2 are isogenous *if there exists an isogeny between them.*

In particular, an isogeny is both a morphism and homomorphism [93, Theorem III.4.8], combining geometric and algebraic structure. Moreover, isogenies are surjective maps with finite kernels [93, Corollary III.4.9]. Of particular interest are the maps from an elliptic curve to itself, which form the following ring.

Definition 2.1.3. *Given an elliptic curve E over a field k, the set of all isogenies from E to itself, together with the constant 0 map, form the* endomorphism ring *of E, denoted by* $\mathrm{End}(E)$. *We write* $\mathrm{End}_k(E)$ *if it is necessary to emphasize that the maps are defined over the field k.*

These rings can be concisely classified, as follows. An analogous object and classification will follow for abelian varieties over finite fields below in Theorem 2.3.2.

Theorem 2.1.4 (Corollary III.9.4, [93])**.** *Given an elliptic curve E over a field k, the endomorphism ring* $\mathrm{End}(E)$ *is either $\mathbb{Z}$, an order in a quadratic imaginary field, or an order in a quaternion algebra.*

In fact, if k is a finite field, then only the latter two options are possible [93, Theorem V.3.1], and if $\mathrm{char}(k) = 0$ then only the former two are possible [93, Corollary III.9.4].

2.2 Some Definitions

Now we turn our attention to the higher-dimensional analogues of elliptic curves which are the content of the first half of this dissertation. Like elliptic curves, we will find that abelian varieties have an abundance of structure leading to a wonderful combination of flavors. We refer to [70, 71] for a complete introduction to abelian varieties and Jacobians.

Definition 2.2.1. *An* abelian variety *A over a field k is a complete group variety.*

Importantly, abelian varieties are smooth and projective [70, Theorem V.7.1], and their group law is always abelian [70, Corollary V.2.4], as their name suggests. Although it may appear that we leave curves behind when moving to varieties of higher dimension, we find that one of the nicest sources of abelian varieties of dimension g are *Jacobians* of curves of genus g; see [71] for details.

Another source of abelian varieties of higher dimension comes from forming products of abelian varieties of smaller dimension. For example, one could consider an n-dimensional abelian variety of the form $E_1 \times E_2 \times \cdots \times E_n$ where each E_i is an elliptic curve. However, this dissertation turns in the opposite direction to rather consider *simple* abelian varieties, defined as follows.

Definition 2.2.2. *An abelian variety A is* simple *(or* elementary*) if A has no nontrivial abelian subvarieties.*

2.2.1 Maps between abelian varieties

Next, we must define the relevant maps between these objects. Note that, unlike the case of elliptic curves, the relevant maps are not always guaranteed to be isogenies. However, in the case of simple abelian varieties, every nonzero homomorphism is an isogeny by definition.

Definition 2.2.3. *Given two abelian varieties A_1, A_2 over a field k, a* homomorphism of abelian varieties *is a morphism*

$$\phi : A_1 \to A_2$$

which maps the identity element of A_1 to the identity element of A_2. Moreover, we say that ϕ is an isogeny *if the morphism is surjective and* $\ker(\phi)$ *is a finite group scheme. We say A_1 and A_2 are* isogenous *if there exists an isogeny between them.*

Definition 2.2.4. *Given an isogeny $\phi : A_1 \to A_2$, there is a corresponding map of function fields $\phi^* : K(A_2) \to K(A_1)$. The* degree *of the isogeny ϕ is defined to be the degree of the field extension $[K(A_1) : \phi^* K(A_2)]$. Moreover, ϕ is said to be* separable *if its kernel is a finite étale subgroup scheme.*

For background material on étale group schemes, we refer to an introductory article of Tate [101]. For our purposes, we will not need the full formal definition. Rather, we only need to know that the kernel of a separable isogeny $\phi : A_1 \to A_2$ over a field k can be recognized (via an equivalence of categories) as a subgroup of $A_1(\overline{k})$ whose cardinality is $\deg(\phi)$; see [101, §3.6].

Of primary importance are endomorphisms, i.e. maps from an abelian variety to itself. Endomorphism rings will be one of the key objects in all results that follow.

Definition 2.2.5. *Given an abelian variety A defined over a field k, the set of all homomorphisms from A to itself form the* endomorphism ring *of A, denoted* $\mathrm{End}(A)$. *We write* $\mathrm{End}_k(A)$ *to emphasize that all homomorphisms are defined over k if necessary.*

One important tool for studying an abelian variety A concerns the *dual abelian variety* $A^\vee$. To save space, we refer readers to [70, §V.9] for the definition and properties of the dual variety. Since we will not need to work directly with polarizations in what follows, readers can take polarizations to intuitively be some extra structure, like a kind of "orientation", imposed on the abelian variety. We follow [70, §V.13] exactly.

Definition 2.2.6. *Given an abelian variety A, a* polarization *on A is an isogeny $\lambda : A \to A^\vee$ such that $\lambda_{\overline{k}} = \varphi_{\mathcal{L}}$ for some ample invertible sheaf $\mathcal{L}$ on $A_{\overline{k}}$. If λ is an isomorphism, then we say ϕ is a* principal polarization.

We say A is *principally polarized* if it admits a principal polarization. A pair (A, λ) containing an abelian variety A and a fixed principal polarization λ is called a *principally polarized* abelian variety. Elliptic curves, and more generally Jacobian varieties, have a canonical principal polarization [71, Summary VII.6.11]. Thus theorems for elliptic curves are sometimes only generalizable to the special cases of principally polarized abelian varieties. This will come into play when we consider algorithmic aspects of abelian varieties defined over finite fields and the computation of endomorphism rings in Chapter 4.

2.2.2 Torsion and Tate modules

Recall that abelian varieties have an abelian group structure, i.e., a $\mathbb{Z}$-module structure. Thus, we can study the torsion elements. The following definitions and results are well-known and are clearly presented in [111], for example.

Definition 2.2.7. *Given an abelian variety A over a field k and $n \in \mathbb{Z}$, the* n-torsion *of A is*

$$A[n] = \{P \in A(\overline{k}) : [n]P = 0\}.$$

Following [111, §2], we recall how to stitch together the torsion subgroups to obtain the Tate module.

Definition 2.2.8. *Let A be an abelian variety over a field k and let $\ell \neq \mathrm{char}(k)$ be a prime. The (ℓ-adic)* Tate module *of A is*

$$T_\ell A = \lim_{\leftarrow} A[\ell^n]$$

where the inverse system is given by the multiplication maps $\ell^m : A[\ell^n] \to A[\ell^{n-m}]$.

Since $A[n] \cong (\mathbb{Z}/n\mathbb{Z})^{2\dim A}$ whenever $\gcd(n, \mathrm{char}(k)) = 1$, we have $T_\ell A \cong \mathbb{Z}_\ell^{2\dim A}$ where $\mathbb{Z}_\ell$ denotes the ℓ-adic integers.

2.3 Abelian varieties over finite fields

In this section, we restrict to the special case of abelian varieties defined over finite fields. We refer to [99, 110, 111] for introductions to the material that follows. Our first goals are to understand the classification of abelian varieties over finite fields up to isogeny, and to describe their endomorphism rings.

Recall that the Galois group $\mathrm{Gal}(\overline{\mathbb{F}_q}/\mathbb{F}_q)$ is topologically generated by a notable automorphism σ_q, the (q-)*Frobenius automorphism*, which is defined via $x \mapsto x^q$. Similarly, given an abelian variety $A/\mathbb{F}_q$, there is a Frobenius endomorphism $\pi : A \to A$ which is induced by the q-th power map and acts the same as σ_q on $A(\overline{\mathbb{F}_q})$. Further, we can view π as a linear map acting on the vector space $T_\ell A \otimes_{\mathbb{Z}_\ell} \mathbb{Q}_\ell$ for any prime $\ell \nmid q$, and define f_A to be the associated characteristic polynomial of degree $2\dim A$.

Theorem 2.3.1 (Theorem 7, [111])**.** *Let A and B be Abelian varieties over a finite field $\mathbb{F}_q$. The following are equivalent:*

(a) A and B are isogenous

(b) $f_A = f_B$.

(c) $\#A(\mathbb{F}_{q^n}) = \#B(\mathbb{F}_{q^n})$ for all $n \geq 1$.

Moreover, we can use the Frobenius endomorphism and the classification given above to precisely describe the endomorphism rings of abelian varieties defined over finite fields.

Theorem 2.3.2 (Theorem 8, [111])**.** *Let A be a simple abelian variety over $\mathbb{F}_q$ with Frobenius endomorphism π.*

(a) $f_A = m_A^d$ for some monic irreducible polynomial $m_A(t) \in \mathbb{Z}[t]$.

(b) All roots of f_A have absolute value $q^{1/2}$.

(c) $D = \mathrm{End}(A) \otimes \mathbb{Q}$ is a division algebra whose center is the field $\mathbb{Q}(\pi)$.

(d) $[D : \mathbb{Q}] = d^2[\mathbb{Q}(\pi) : \mathbb{Q}]$ and $2\dim(A) = d[\mathbb{Q}(\pi) : \mathbb{Q}]$.

It therefore follows that the isogeny class of a simple abelian variety A over $\mathbb{F}_q$ is uniquely determined by the irreducible polynomial m_A. In particular, it is uniquely determined by a root of f_A, up to Galois conjugacy. This leads to the following definition, in light of the theorem above.

Definition 2.3.3. *A* Weil q-integer *is an algebraic integer whose conjugates all have absolute value $q^{1/2}$.*

With this terminology, we have the following correspondence.

Theorem 2.3.4 (Honda-Tate Theorem, [100])**.** *There is a bijective correspondence between*

$$\{\textit{isogeny classes of simple abelian varieties over } \mathbb{F}_q\}$$

and

$$\{\textit{conjugacy classes of Weil } q\textit{-integers}\}.$$

Finally, we recall the definition of an ordinary abelian variety. Such varieties will be a key example in Chapter 3 and the context of Chapter 4.

Definition 2.3.5. *An abelian variety A over $\mathbb{F}_q$ with $p = \mathrm{char}(\mathbb{F}_q)$ is called* ordinary *if A contains the maximal number $p^{2\dim A}$ points of order dividing p.*

Waterhouse showed the following for ordinary abelian varieties, leading us to drop the field subscript on $\mathrm{End}(A)$ throughout Chapter 4.

Theorem 2.3.6 (Theorem 7.2, [110])**.** *Let A be a simple ordinary abelian variety over $\mathbb{F}_q$. Then $\mathrm{End}_{\mathbb{F}_q}(A)$ is commutative and unchanged by base field extension.*

Chapter 3
The Group of Rational Points

3.1 Main result and comparison to prior work

Given an abelian variety A over a finite field $\mathbb{F}_q$, one may view the group of rational points $A(\mathbb{F}_q)$ as a module over the ring $\mathrm{End}_{\mathbb{F}_q}(A)$ of endomorphisms defined over $\mathbb{F}_q$. Lenstra completely described this module structure for elliptic curves over finite fields in the following theorem. In addition to being useful and interesting in its own right, this theorem also determines *a fortiori* the underlying abelian group structure of $A(\mathbb{F}_q)$ purely in terms of the endomorphism ring. The latter perspective has been leveraged for the sake of computational number theory and cryptography; see, for example, the work of Galbraith [37, Lemma 1], Ionica and Joux [45, §2.3], and Kohel [56, Chapter 4]. The goal of this chapter is to generalize Lenstra's theorem beyond elliptic curves to abelian varieties of arbitrary dimension.

Theorem 3.1.1 ([61], Theorem 1)**.** *Let E be an elliptic curve over $\mathbb{F}_q$. Write $R = \mathrm{End}_{\mathbb{F}_q}(E)$ and let $\pi \in R$ be the Frobenius endomorphism of E.*

(a) Suppose that $\pi \notin \mathbb{Z}$. Then R has rank 2 over $\mathbb{Z}$ and there is an isomorphism of R-modules

$$E(\mathbb{F}_{q^n}) \cong R/(\pi^n - 1)R.$$

(b) Suppose that $\pi \in \mathbb{Z}$. Then R has rank 4 over $\mathbb{Z}$, we have

$$E(\mathbb{F}_{q^n}) \cong \mathbb{Z}/\mathbb{Z}(\pi^n - 1) \oplus \mathbb{Z}/\mathbb{Z}(\pi^n - 1)$$

as abelian groups. Further, this group has up to isomorphism exactly one left

R-module structure, and one has an isomorphism of R-modules

$$E(\mathbb{F}_{q^n}) \oplus E(\mathbb{F}_{q^n}) \cong R/R(\pi^n - 1).$$

Notice that E is supersingular in the second case, but not conversely. To prove the theorem, Lenstra notes that $E(\mathbb{F}_{q^n}) = E[\pi^n - 1]$, and $\pi^n - 1$ is a separable isogeny. For part (b), the abelian group structure is simply the well-known structure of the n-torsion of an elliptic curve for $n \in \mathbb{Z}$. The additional statements in part (b) follow from Morita equivalence and an isomorphism of rings, for integers n coprime to q, between R/Rn and the ring $\mathrm{M}_2(\mathbb{Z}/n\mathbb{Z})$ of 2×2 matrices with coefficients in $\mathbb{Z}/n\mathbb{Z}$.

For part (a) of the theorem, Lenstra uses the following proposition; see [61, Proposition 2.1].

Proposition 3.1.2. *Let E be an elliptic curve over $\mathbb{F}_q$, and let $R = \mathrm{End}_{\mathbb{F}_q} E$. If $[R : \mathbb{Z}] = 2$, then for every separable element $s \in R$ there is an isomorphism $E[s] \cong R/Rs$ of R-modules.*

Lenstra showed in his original paper that the preceding proposition does not immediately generalize to all "nice" abelian varieties of higher dimension, i.e. principally polarized ordinary abelian varieties; see [61, Proposition 6.4]. Although this means that a certain natural generalization is not correct, the examples that Lenstra produces must have very particular endomorphism rings. By inspecting Lenstra's theorem through two perspectives and imposing restrictions on the endomorphism ring, we can recover a natural generalization to certain abelian varieties of higher dimension.

3.1.1 First perspective: Gorenstein rings

First, consider part (a) of Lenstra's theorem, or more generally, Proposition 3.1.2. In this case, the endomorphism ring of the elliptic curve is commutative, specifically an order in an imaginary quadratic number field. In general, a simple abelian variety A of dimension g over $\mathbb{F}_q$ with Frobenius endomorphism π has commutative endomorphism ring exactly when $[\mathbb{Q}(\pi) : \mathbb{Q}] = 2g$, and in this case, $\mathrm{End}_{\mathbb{F}_q}(A)$ is an order in the field $\mathbb{Q}(\pi)$ [111, Theorem 8]. In fact, if π is an ordinary Weil q-integer, then the rings which arise as the endomorphism rings of abelian varieties in the corresponding isogeny class over $\mathbb{F}_q$ are precisely the orders of $\mathbb{Q}(\pi)$ which contain the minimal order $\mathbb{Z}[\pi, \overline{\pi}]$ [110, Theorem 7.4]. Since every order in a quadratic number field is Gorenstein, restricting to the

Gorenstein case for abelian varieties of arbitrary dimension provides us with our first natural generalization.

Proposition 3.2.1. *Let A be a simple abelian variety over $\mathbb{F}_q$ of dimension g with Frobenius endomorphism π. If $[\mathbb{Q}(\pi) : \mathbb{Q}] = 2g$ and $R = \mathrm{End}_{\mathbb{F}_q}(A)$ is a Gorenstein ring, then there is an isomorphism of R-modules*

$$A[s] \cong R/Rs$$

for every separable $s \in R$.

This proposition will be proved in Section 3.2 by using properties of finite local Gorenstein rings. To see examples where the proposition applies, note that $\mathrm{End}_{\mathbb{F}_q}(A)$ is guaranteed to be Gorenstein if A has *maximal real multiplication*, i.e. if $\mathrm{End}_{\mathbb{F}_q}(A)$ contains the ring of integers of the maximal totally real subfield of $\mathbb{Q}(\pi)$; see [13, Lemma 4.4]. Many recent results in the algorithmic study of abelian varieties over finite fields have productively focused on the case of maximal real multiplication, including results on point counting [2,39], isogeny graphs [13,46,64], and endomorphism ring computation, as in Chapter 4. At the other extreme, Centeleghe and Stix have shown that the minimal order $\mathbb{Z}[\pi, \overline{\pi}]$ is also always Gorenstein, where π is a Weil integer [18, Theorem 11]. In fact, we can use this fact with the result above to prove that every simple ordinary isogeny class over $\mathbb{F}_q$ contains an abelian variety whose group of rational points is cyclic, generalizing a result of Galbraith [37, Lemma 1].

Corollary 3.1.3. *If A is a simple ordinary abelian variety over $\mathbb{F}_q$ with Frobenius π and endomorphism ring $\mathrm{End}_{\mathbb{F}_q} = \mathbb{Z}[\pi, \overline{\pi}]$, then $A(\mathbb{F}_q)$ is a cyclic group.*

Proof. It is convenient to rewrite $\mathbb{Z}[\pi, \overline{\pi}] \cong \mathbb{Z}[F, V]/(FV - q, f_A(F))$ where $f_A(F)$ is the characteristic polynomial of π. Here, F and V are merely polynomial variables allowing us to recognize $\mathbb{Z}[\pi, \overline{\pi}]$ as a quotient of a polynomial ring. We identify F with π and V with $\overline{\pi}$, as in [18]. With this notation, Proposition 3.2.1 implies that

$$A(\mathbb{F}_q) \cong \mathbb{Z}[\pi, \overline{\pi}]/(\pi - 1) \cong \mathbb{Z}[F, V]/(FV - q, f_A(F), F - 1) \cong \mathbb{Z}/(f_A(1)),$$

which shows $A(\mathbb{F}_q)$ is a cyclic group. □

3.1.2 Second perspective: Modules over the center

Now consider part (b) of Lenstra's theorem, where E is a supersingular elliptic curve over $\mathbb{F}_q$ with all endomorphisms defined. Before describing the group of rational points $E(\mathbb{F}_{q^n})$ as a module over the endomorphism ring $\mathrm{End}_{\mathbb{F}_q}(E)$, Lenstra first identifies $E(\mathbb{F}_{q^n})$ as an abelian group, i.e. a module over $\mathbb{Z}$. Importantly, $\mathbb{Z}$ is the center of the endomorphism ring in this case.

Following this point of view, given a simple abelian variety A over $\mathbb{F}_q$ with Frobenius endomorphism π, we will first consider the structure of $A(\mathbb{F}_{q^n})$ as a module of the center of $\mathrm{End}_{\mathbb{F}_q}(A)$. Recall that the center of the endomorphism algebra $\mathrm{End}_{\mathbb{F}_q}(A) \otimes \mathbb{Q}$ is the field $\mathbb{Q}(\pi)$ [111, Theorem 8]. More generally, we can study $A[s]$ as a module over the center of the endomorphism ring $\mathrm{End}_{\mathbb{F}_q}(A)$ for any separable endomorphism s in the center, which leads us to the following result.

Proposition 3.3.1. *Let A be a simple abelian variety over $\mathbb{F}_q$ of dimension g, and let Z be the center of $R = \mathrm{End}_{\mathbb{F}_q}(A)$. If s is a separable element of Z for which sZ is the product of invertible prime ideals in Z, then there is an isomorphism of Z-modules*

$$A[s] \cong (Z/Zs)^d$$

where $d = 2g/[\mathbb{Q}(\pi) : \mathbb{Q}]$. Moreover, this Z-module has exactly one R-module structure, up to isomorphism. The unique R-module structure comes from the isomorphism of rings $R/Rs \cong \mathrm{M}_d(Z/Zs)$, and there is an isomorphism

$$A[s]^d \cong R/Rs$$

as R-modules.

This proposition will be proved in Section 3.3 through the study of kernel ideals. The latter parts of this proposition will follow from Morita equivalence, similarly to Theorem 3.1.1.(b). Notice that we must require that sZ is the product of invertible prime ideals, which is automatically true when Z is a maximal order. For example, let A be an abelian surface defined over $\mathbb{F}_p$ in the isogeny class corresponding to the Weil polynomial $(t^2 - p)^2$ for a prime $p \not\equiv 1 \bmod 4$. This Weil polynomial corresponds to the Weil restriction of a supersingular elliptic curve over $\mathbb{F}_{p^2}$, and A is simple over $\mathbb{F}_p$. The endomorphism ring $\mathrm{End}_{\mathbb{F}_p}(A)$ is a noncommutative ring whose center is $\mathbb{Z}[\sqrt{p}]$, which is a maximal order by construction because $p \not\equiv 1 \bmod 4$. Hence the proposition automatically applies in this

case for any separable $s \in \mathbb{Z}[\sqrt{p}]$.

3.1.3 Main Result

Combining the perspectives outlined above, we have the following main result.

Theorem 3.1.4. *For $g \geq 1$, let A be a simple abelian variety over $\mathbb{F}_q$ of dimension g with Frobenius endomorphism π. Write $K = \mathbb{Q}(\pi)$ and $R = \mathrm{End}_{\mathbb{F}_q}(A)$, and let Z be the center of R.*

(a) If $[K : \mathbb{Q}] = 2g$ and R is a Gorenstein ring, then

$$A(\mathbb{F}_{q^n}) \cong R/R(\pi^n - 1).$$

(b) If $(\pi^n - 1)Z$ is the product of invertible prime ideals in Z, then there is an isomorphism of Z-modules

$$A(\mathbb{F}_{q^n}) \cong (Z/Z(\pi^n - 1))^d,$$

where $d = 2g/[K : \mathbb{Q}]$. Moreover, this Z-module has exactly one left R-module structure, up to isomorphism. This R-module structure comes from the isomorphism of rings $R/R(\pi^n - 1) \cong \mathrm{M}_d(Z/Z(\pi^n - 1))$, and there is an isomorphism of R-modules

$$A(\mathbb{F}_{q^n})^d \cong R/R(\pi^n - 1).$$

Notice that parts (a) and (b) of the theorem provide the same answer in the case when all hypotheses are simultaneously satisfied, e.g. when A is a simple ordinary abelian variety with maximal endomorphism ring. The theorem follows immediately from the propositions above, given that $A(\mathbb{F}_{q^n}) = A[\pi^n - 1]$ and $\pi^n - 1$ is a separable isogeny, as in the elliptic curve case. Propositions 3.2.1 and 3.3.1 will be proved in Sections 3.2 and 3.3, respectively, which completes the proof of our main theorem. Finally, in Section 3.4, we stitch together all of the isomorphisms described above to understand the structure of $A(\overline{\mathbb{F}}_q)$ as a module of the endomorphism ring $\mathrm{End}_{\mathbb{F}_q}(A)$.

3.2 Gorenstein rings

The goal of this section is to prove the following generalization of Proposition 3.1.2, as outlined in the introduction.

Proposition 3.2.1. *Let A be a simple abelian variety over $\mathbb{F}_q$ of dimension g with Frobenius endomorphism π. If $[\mathbb{Q}(\pi) : \mathbb{Q}] = 2g$ and $R = \mathrm{End}_{\mathbb{F}_q}(A)$ is a Gorenstein ring, then there is an isomorphism of R-modules*

$$A[s] \cong R/Rs$$

for every separable $s \in R$.

In order to prove this proposition, we will follow a strategy that is largely similar to the proof of Theorem 3.1.1.(a) in Lenstra's original paper. Our approach differs from Lenstra by working directly with finite local Gorenstein rings, rather than using duality. Background for Gorenstein rings can be found in Matsumura's book [67, Chapter 18].

Lemma 3.2.2. *Let R be a Gorenstein domain and s a nonzero element of R. If the quotient $S = R/Rs$ is finite, then every faithful S-module M contains a submodule that is free of rank 1 over S.*

Proof. Notice that S is Gorenstein because R is Gorenstein; see [67, Exercise 18.1]. Additionally, the fact that S is finite implies that it is an Artinian ring. In particular, it is canonically isomorphic to a finite product of its localizations $S = S_1 \times \cdots \times S_r$. Thus every S-module M has the form $M \cong M_1 \times \cdots \times M_r$ where M_i is an S_i-module for each $1 \leq i \leq r$. This lemma therefore reduces to the following lemma. □

Lemma 3.2.3. *Let $(T, \mathfrak{m})$ be a finite local Artinian ring that is Gorenstein.*

(a) Every nonzero ideal $J \subseteq T$ contained in $\mathfrak{m}$ contains a nonzero element that is killed by all elements of $\mathfrak{m}$.

(b) Every faithful T-module N contains a submodule that is free of rank 1 over T.

Proof. To prove part (a), list the elements of the maximal ideal $\mathfrak{m} = \{a_1, \ldots, a_d\}$. Define $J_0 = J$, and for each $1 \leq i \leq d$, let J_i be the set of elements of J which are annihilated by $\{a_1, \ldots, a_i\}$. In other words, for each $1 \leq i \leq d$, the ideal J_i is the kernel of the map $f_i : J_{i-1} \to J_{i-1}$ defined by $x \mapsto a_i x$. All elements of $\mathfrak{m}$ are nilpotent, and therefore the kernel J_i of the map f_i is nontrivial precisely when $J_{i-1} \neq 0$. Since $J_0 \neq 0$ by hypothesis, it is clear by induction that $J_i \neq 0$ for all $1 \leq i \leq d$. In particular, there are nonzero elements in $J_d \subseteq J$ which are annihilated by every element of $\mathfrak{m}$.

For part (b), let $k = T/\mathfrak{m}$ be the residue field of T. Because T is a zero-dimensional Gorenstein ring, the k-vector space $\mathrm{Ext}^0_T(k, T) = \mathrm{Hom}_T(k, T)$ is one-dimensional; see [67,

Theorem 18.1]. Thus the annihilator of $\mathfrak{m}$ in T is a principal ideal $I = tT$ where $t = \phi(1)$ for some nonzero $\phi : k \to T$. Because N is a faithful module, there is some $n \in N$ such that $tn \neq 0$. Let $\mathrm{Ann}(n)$ be the annihilator of n, which is an ideal contained in $\mathfrak{m}$.

If $\mathrm{Ann}(n) = 0$, then the submodule $Tn \subseteq N$ is free of rank 1 and we are done. If $\mathrm{Ann}(n) \neq 0$, then part (a) implies that $\mathrm{Ann}(n)$ contains a nonzero element x which is killed by all elements of $\mathfrak{m}$. Since I is the annihilator of $\mathfrak{m}$, this means that $x \in \mathrm{Ann}(n)$ is also a nonzero element of I. However, I is a principal ideal that can be viewed as a module over the field $k = T/\mathfrak{m}$, hence every nonzero element of I is a generator. In particular, $xn \neq 0$ because $t \in I = xT$ and $tn \neq 0$. This contradiction completes the proof. □

We are now ready to prove the key proposition.

Proof of Proposition 3.2.1. Put $S = R/Rs$ and $M = A[s]$ for ease of notation. Notice that M is a *faithful* S-module: Any $r \in R$ such that $rM = rA[s] = 0$ factors as $r = ts$ for some $t \in R$, i.e. $r \in Rs$. Indeed, this follows immediately from the universal property of quotients; see [51, Remark 7.(c)].

Therefore, Lemma 3.2.2 implies that M contains a free S-submodule of rank 1. Now, we can count the cardinalities of these sets:

$$\#M = \deg s = N_{K/\mathbb{Q}}s = \#R/Rs = \#S.$$

The first equality comes from the separability of s, and the second equality above is a well-known theorem [70, Proposition 12.12]. Therefore, $M \cong S$ as an S-module because their cardinalities are the same. This proves Proposition 3.2.1. □

3.3 Using kernel ideals

In this section, A is a simple abelian variety over $\mathbb{F}_q$ with Frobenius endomorphism π. Then the endomorphism algebra $D = \mathrm{End}_{\mathbb{F}_q}(A) \otimes \mathbb{Q}$ is a division algebra with center $K = \mathbb{Q}(\pi)$ [111, Theorem 8]. Write $R = \mathrm{End}_{\mathbb{F}_q}(A)$, and let Z be the center of the endomorphism ring. Our goal in this section is to prove Proposition 3.3.1, which we repeat below for convenience.

Proposition 3.3.1. *If s is a separable element of Z for which sZ is the product of invertible prime ideals in Z, then there is an isomorphism of Z-modules*

$$A[s] \cong (Z/Zs)^d$$

where $d = 2g/[\mathbb{Q}(\pi) : \mathbb{Q}]$. Moreover, this Z-module has exactly one R-module structure, up to isomorphism. This R-module structure comes from the isomorphism of rings $R/Rs \cong \mathrm{M}_d(Z/Zs)$, and there is an isomorphism

$$A[s]^d \cong R/Rs$$

as R-modules.

To prove this proposition, we will inspect the isogenies associated to (left) ideals, inspired by Waterhouse [110]; see also [51, §2] for additional background. In the construction of Waterhouse, a nonzero ideal $I \subseteq R$ is associated to an isogeny whose kernel is $A[I] = \cap_{\alpha \in I} A[\alpha]$, where $A[\alpha]$ is the kernel of the endomorphism α. In other words, if I is generated by the elements $\alpha_1, \ldots, \alpha_m$, then the abelian variety $A/A[I]$ is isomorphic to the image of the map $(\alpha_1, \ldots, \alpha_m) : A \to A^m$.

Similarly, we can also associate a finite subgroup scheme H of A to a left ideal $I(H) \subseteq R$, given by

$$I(H) = \{\alpha \in R : H \subseteq A[\alpha]\}.$$

Given a nonzero ideal $I \subseteq R$, we always have $I \subseteq I(A[I])$. If equality holds, then I is called a *kernel ideal*. Every nonzero ideal I is contained in a kernel ideal J such that $A[I] = A[J]$.

For our purposes, we will be concerned with isogenies that are associated to ideals contained in the center $I_0 \subseteq Z$. For convenience, we will write $A[I_0]$ in place of $A[I_0 R]$. The goal of this section is to describe $A[s]$ in terms of $A[\mathfrak{p}_j^{e_j}]$ where $sZ = \mathfrak{p}_1^{e_1} \ldots \mathfrak{p}_r^{e_r}$ is the factorization of s into invertible prime ideals in Z, which will allow us to prove Proposition 3.3.1.

3.3.1 Basics of invertible ideals

First, we recall some basic key properties about invertible ideals in algebraic number theory. Within this section, let L denote a number field and let $\mathcal{O} \subseteq L$ be an order. The *conductor ideal* of $\mathcal{O}$ is defined to be $\mathfrak{f}_{\mathcal{O}} = \{a \in L : a\mathcal{O}_L \subseteq \mathcal{O}\}$. The following lemmas show the connection between the conductor ideal and the invertibility of ideals.

Lemma 3.3.2. *If $\mathfrak{p} \subseteq \mathcal{O}$ is a nonzero prime ideal, then the following are equivalent:*

1. *$\mathfrak{p}$ is invertible, i.e. $\mathfrak{p}I = a\mathcal{O}$ for some ideal $I \subseteq \mathcal{O}$ and some $a \in \mathcal{O}$;*

2. *$\mathfrak{p}$ is regular, i.e. the localization $\mathcal{O}_{\mathfrak{p}}$ is integrally closed;*

3. $\mathfrak{p}$ *is coprime to the conductor ideal* $\mathfrak{f}_{\mathcal{O}}$*, i.e.* $\mathfrak{p} + \mathfrak{f}_{\mathcal{O}} = \mathcal{O}$*.*

Moreover, when these equivalent conditions hold, the localization $\mathcal{O}_{\mathfrak{p}}$ *is a discrete valuation ring.*

Proof. The prime ideal $\mathfrak{p}$ is invertible if and only if it is regular by [72, Exercise I.12.5], which is true if and only if $\mathfrak{p} \not\supseteq \mathfrak{f}_{\mathcal{O}}$ [72, Proposition 12.10]. To obtain the last equivalent condition, observe that $\mathcal{O}$ is a one-dimensional Noetherian integral domain [72, Proposition I.12.2], so any nonzero prime ideal of $\mathcal{O}$ is maximal. In particular, $\mathfrak{p} \not\supseteq \mathfrak{f}_{\mathcal{O}}$ is equivalent to $\mathfrak{p} + \mathfrak{f}_{\mathcal{O}} = \mathcal{O}$.

Finally, if $\mathfrak{p}$ is regular, then the localization $\mathcal{O}_{\mathfrak{p}}$ is equal to the localization of the ring of integers $\mathcal{O}_L$ at the prime ideal $\hat{\mathfrak{p}} = \mathfrak{p}\mathcal{O}_L$ [72, Proposition 12.10], and the latter localization $\mathcal{O}_{L,\hat{\mathfrak{p}}}$ is known to be a discrete valuation ring [72, Proposition I.11.5]. □

While the preceding lemma focuses on prime ideals, the following result shows the connection between invertibility and the conductor ideal in general. In particular, we see that Proposition 3.3.1 can be rephrased to require that sZ is coprime to the conductor ideal $\mathfrak{f}_Z$ of Z instead of requiring that sZ is the product of invertible ideals.

Lemma 3.3.3 (Proposition 3.2, [63])**.** *If* $\mathfrak{a} \subseteq \mathcal{O}$ *is any ideal coprime to the conductor* $\mathfrak{f}_{\mathcal{O}}$*, then* $\mathfrak{a}$ *is invertible and is uniquely factored into (invertible) prime ideals.*

Recall that the *Picard group* $\mathrm{Cl}(\mathcal{O})$ is defined to be the quotient of the set of invertible fractional ideals of $\mathcal{O}$ by the set of principal fractional ideals. We refer readers to [72, §I.12] and [63] for additional background.

Lemma 3.3.4. *Every class of ideals in* $\mathrm{Cl}(\mathcal{O})$ *contains infinitely many prime ideals.*

Proof. The extension and contraction of ideals provides a natural bijection between the set of invertible prime ideals of $\mathcal{O}$ and the set of prime ideals of $\mathcal{O}_L$ which are coprime to the conductor ideal $\mathfrak{f}_{\mathcal{O}}$ [63, Lemma 3.3]. Using this bijection, there is a natural isomorphism of groups that allows us to interpret the Picard group $\mathrm{Cl}(\mathcal{O})$ in terms of fractional ideals of $\mathcal{O}_L$ which are coprime to the ideal $\mathfrak{f}_{\mathcal{O}}$ [63, Theorem 3.11]. This reduces the claim to a question concerning ideals in $\mathcal{O}_L$, and a generalization of the Dirichlet density theorem immediately shows that there are infinitely many suitable prime ideals [72, Theorem VII.13.2]. □

3.3.2 Isogenies associated to invertible ideals

Now we focus our attention on the invertible ideals of the center Z of the endomorphism ring R, and investigate the corresponding isogenies.

Lemma 3.3.5. *If $I_0 \subseteq Z$ is an invertible ideal, then I_0R is an invertible two-sided ideal of R. In particular, I_0R is a kernel ideal.*

Proof. Clearly I_0R is naturally a right ideal, and RI_0 is naturally a left ideal, and these two sets are equal as $I_0 \subseteq Z$ is in the center. Thus, I_0R is a two-sided ideal.

Because I_0 is invertible, there is a fractional ideal J_0 of Z such that $I_0J_0 = Z$. Since Z is the center of R, it also follows that

$$(I_0R)(J_0R) = (J_0R)(I_0R) = R.$$

Moreover, if J is any fractional two-sided ideal of R such that $J \cdot (I_0R) = (I_0R) \cdot J = R$, then $J_0R = (J_0R)(I_0R)J = J$. This proves that J_0R is the unique two-sided fractional ideal of R with this property, which we denote $(I_0R)^{-1}$. It follows immediately from uniqueness that $((I_0R)^{-1})^{-1} = I_0R$.

Now for any ideal I of R, define $(R : I) = \{x \in D : xI \subseteq R\}$. Then we have

$$(R : I_0R) = \{x \in D : xI_0 \subseteq R\} = \{x \in D : I_0x \subseteq R\}$$

because $xI_0R \subseteq R$ if and only if $xI_0 \subseteq R$, and $xI_0 = I_0x$ for all $x \in D$ because I_0 is contained in the center Z. In particular, $(R : I_0R)$ is a two-sided fractional ideal and it is easy to verify that $(R : I_0R) = (I_0R)^{-1}$. Indeed, the containments

$$R \supseteq (R : I_0R) \cdot I_0R \supseteq (I_0R)^{-1} \cdot (I_0R) = R$$

show that $(R : I_0R) \cdot I_0R = R$, and similarly $I_0R \cdot (R : I_0R) = R$. Therefore, we have

$$(R : (R : I_0R)) = ((I_0R)^{-1})^{-1} = I_0R.$$

By [51, Remark 7.(d)], we know that

$$I(A[I_0R]) \subseteq \bigcap_{Rf \supseteq I_0} Rf$$

where the intersection is taken over all elements $f \in D$.

A routine verification shows that

$$\begin{aligned}(R : (R : I_0R)) &= \{x \in D : x \cdot (R : I_0R) \subseteq R\} \\ &= \{x \in D : \forall y \in D, \text{ if } I_0y \subseteq R, \text{ then } xy \in R\} \\ &= \{x \in D : \forall y \in D \setminus \{0\}, \text{ if } I_0 \subseteq Ry^{-1}, \text{ then } x \in Ry^{-1}\} \\ &= \bigcap_{Ry^{-1} \supseteq I_0R} \{x \in D : x \in Ry^{-1}\} \\ &= \bigcap_{Ry^{-1} \supseteq I_0R} Ry^{-1} \\ &= \bigcap_{Rf \supseteq I_0R} Rf\end{aligned}$$

where the final equality comes from simply reindexing the intersection with $f = y^{-1}$.

Combining all of the containments above, we see that

$$I_0R \subseteq I(A[I_0R]) \subseteq \bigcap_{Rf \supseteq I} Rf = (R : (R : I_0R)) = I_0R$$

which shows that I_0R is a kernel ideal by definition. □

The lemma above is useful because it shows that the prime ideals appearing in Proposition 3.3.1 are actually kernel ideals, which gives us the following important information. We will write $|H|$ for the rank of a finite subgroup scheme H of A, or equivalently, the degree of the isogeny $\pi_H : A \to A/H$.

Proposition 3.3.6. *If $I_0 \subseteq Z$ is an invertible ideal, then*

$$\mathrm{End}_{\mathbb{F}_q}(A/A[I_0]) = \mathrm{End}_{\mathbb{F}_q}(A) = R.$$

Moreover,

$$|A[I_0]| = N_{K/\mathbb{Q}}(I_0)^{2g/[K:\mathbb{Q}]}.$$

Proof. For convenience, write $B = A/A[I_0]$. Because I_0R is a kernel ideal by Lemma 3.3.5, the endomorphism ring $\mathrm{End}_{\mathbb{F}_q}(B)$ is equal to the right order of I_0R [110, Proposition 3.9], which we denote by

$$\mathcal{O}_r(I_0R) = \{x \in D : (I_0R) \cdot x \subseteq I_0R\}$$

Since I_0R is a two-sided ideal, clearly $R \subseteq \mathcal{O}_r(I_0R)$. Conversely, let $x \in \mathcal{O}_r(I_0R)$. Then

$$Rx = (I_0R)^{-1}(I_0R)x \subseteq (I_0R)^{-1}I_0R = R$$

because I_0R is an invertible ideal. Therefore, $x \in R$ and $\mathrm{End}_{\mathbb{F}_q}(B) = \mathcal{O}_r(I_0R) = R$.

To prove the second claim, first assume that $I_0 = \alpha Z$ is a principal ideal. Then $A[I_0] = A[\alpha]$ and $|A[I_0]| = \deg(\alpha)$, so the claim is known [70, Proposition 12.12].

Now suppose I_0 is not principal. Because I_0 is an invertible ideal of Z, we can pick an ideal $J_0 \subseteq Z$ such that $I_0J_0 = \lambda Z$ and $N_{K/\mathbb{Q}}(J_0)$ is coprime to $|A[I_0]|$. Indeed, there are only finitely many prime factors of $|A[I_0]|$, while there are infinitely many prime ideals in the equivalence class $[I_0]^{-1} \in \mathrm{Cl}(Z)$ by Lemma 3.3.4. Multiplication of ideals corresponds to composition of isogenies [110, Proposition 3.12], and therefore

$$\begin{aligned} |A[I_0]| \cdot |B[J_0]| &= |A[I_0J_0]| \\ &= |A[\lambda]| \\ &= N_{K/\mathbb{Q}}(\lambda)^{2g/[K:\mathbb{Q}]} \\ &= N_{K/\mathbb{Q}}(I_0)^{2g/[K:\mathbb{Q}]} N_{K/\mathbb{Q}}(J_0)^{2g/[K:\mathbb{Q}]} \end{aligned}$$

Now the fact that the rank of $A[I_0]$ is coprime to $N_{K/\mathbb{Q}}(J_0)$ means that $|A[I_0]|$ divides $N_{K/\mathbb{Q}}(I_0)^{2g/[K:\mathbb{Q}]}$. But the same must be true for J_0, so $|B[J_0]|$ divides $N_{K/\mathbb{Q}}(J_0)^{2g/[K:\mathbb{Q}]}$ as well. Therefore, equality must hold, as claimed. □

Because we are ultimately only concerned with separable isogenies, we will restrict our attention to this case now. Recall that the kernel of a separable isogeny $\phi : A \to A'$ can be identified with a finite subgroup of $A(\overline{\mathbb{F}}_q)$ of cardinality $\deg \phi$.

Lemma 3.3.7. *If $r \geq 1$, and $\mathfrak{p} \subseteq Z$ is an invertible prime ideal which corresponds to a separable isogeny, then*

$$A[\mathfrak{p}^r] \cong (Z/\mathfrak{p}^r)^{2g/[K:\mathbb{Q}]}$$

is an isomorphism of Z-modules.

Proof. First, $A[\mathfrak{p}]$ is a $Z/\mathfrak{p}$-module. But $Z/\mathfrak{p}$ is a field, so $A[\mathfrak{p}]$ is a vector space, and therefore $A[\mathfrak{p}] \cong (Z/\mathfrak{p})^m$ for some m. We have $m = 2g/[K:\mathbb{Q}]$ by counting the cardinality of each side with Proposition 3.3.6.

Now we proceed by induction. Given $r \geq 2$, we know that $A[\mathfrak{p}^r]$ is a finitely generated module over $Z/\mathfrak{p}^r \cong Z_\mathfrak{p}/\mathfrak{p}^r Z_\mathfrak{p}$. Because $Z_\mathfrak{p}$ is a discrete valuation ring by Lemma 3.3.2,

we can apply the structure theorem for finitely generated modules [27, Theorem 12.1.6] to deduce that $A[\mathfrak{p}^r]$ is the direct sum of modules of the form $Z_\mathfrak{p}/\mathfrak{p}^i Z_\mathfrak{p} \cong Z/\mathfrak{p}^i$ for $1 \leq i \leq r$.

Further, $A[\mathfrak{p}^r]$ contains $A[\mathfrak{p}^{r-1}]$, which is of the form $(Z/\mathfrak{p}^{r-1})^{2g/[K:\mathbb{Q}]}$ by assumption. Thus, writing $A[\mathfrak{p}^r] \cong Z/\mathfrak{p}^{r_1} \times \cdots \times Z/\mathfrak{p}^{r_s}$ implies that $s = 2g/[K:\mathbb{Q}]$. By counting the cardinality, we must have $r_j = r$ for all $1 \leq j \leq s$. □

3.3.3 Proof of main result

Now we are ready to prove the main result of this section.

Proof of Proposition 3.3.1. We factor $(s) = \mathfrak{p}_1^{e_1} \dots \mathfrak{p}_r^{e_r}$. Notice that for any nonzero $I, J \subseteq R$, we have $A[I] \cap A[J] = A[I+J]$ by definition because $I + J$ is generated by $I \cup J$. Thus, coprime ideals correspond to subgroups with trivial intersection, and we conclude that we have an isomorphism of Z-modules:

$$A[s] \cong A[\mathfrak{p}_1^{e_1}] \times \cdots \times A[\mathfrak{p}_r^{e_r}].$$

For each $1 \leq i \leq r$, we see that $A[\mathfrak{p}_i^{e_i}] \cong (Z/\mathfrak{p}_i^{e_i})^{2g/[K:\mathbb{Q}]}$ by the proposition above. By the Chinese Remainder Theorem, we conclude that

$$A[s] \cong (Z/Zs)^{2g/[K:\mathbb{Q}]}$$

as desired.

Now write $d = 2g/[K:\mathbb{Q}]$ for convenience. To prove the second claim, we notice that the endomorphism ring of the Z-module $A[s] \cong (Z/Zs)^d$ is the ring of $d \times d$ matrices over Z/Zs, which we write as $\mathrm{End}_Z(A[s]) = \mathrm{M}_d(Z/Zs)$. As in the proof of Proposition 3.2.1, we see that $A[s]$ is a faithful R/Rs-module, so the map $R/Rs \to \mathrm{End}_Z(A[s])$ induced by the natural R-module structure on $A[s]$ is injective. Moreover, s defines a linear map on the lattice $R \subseteq D$, so we have

$$\#(R/Rs) = N_{D/\mathbb{Q}}(s) = N_{K/\mathbb{Q}}(N_{D/K}(s)) = N_{K/\mathbb{Q}}(s)^{[D:K]},$$

where $N_{D/\mathbb{Q}}(s)$ and $N_{D/K}(s)$ denote the determinants of $s : D \to D$ as a linear map over $\mathbb{Q}$ and K, respectively. On the other hand, it is clear that

$$\#\mathrm{M}_d(Z/Zs) = N_{K/\mathbb{Q}}(s)^{d^2} = N_{K/\mathbb{Q}}(s)^{[D:K]}$$

because $d^2 = [D:K]$; see [111, Theorem 8]. Therefore, R/Rs and $\mathrm{M}_d(Z/Zs)$ have the

same cardinality, so the injective ring map $R \to \mathrm{M}_d(Z/Zs)$ is an isomorphism.

Therefore, to prove that $A[s]$ has exactly one R-module structure, it suffices to show that $(Z/Zs)^d$ has exactly one $\mathrm{M}_d(Z/Zs)$-module structure. Morita equivalence states that every $\mathrm{M}_d(Z/Zs)$-module M' is isomorphic to M^d for some Z/Zs-module M, where M^d is given the natural left $\mathrm{M}_d(Z/Zs)$-module structure defined by applying matrices to column vectors; see [47, Proposition 1.4]. Thus we simply need to know that if a Z-module M satisfies $M^d \cong (Z/Zs)^d$, then $M \cong Z/Zs$. But, as above, s is the product of invertible primes, so M must be of the desired form.

Finally, we notice that $\mathrm{M}_d(Z/Zs)$ is isomorphic to $((Z/Zs)^d)^d$ as a module over itself, which proves the final claim. □

3.4 Considering the algebraic closure

Now that we have considered the module structure of the group of rational points of a simple abelian variety over a finite field $\mathbb{F}_q$, we turn our attention towards the algebraic closure $\overline{\mathbb{F}}_q$. Because $\overline{\mathbb{F}}_q$ is the union of all its finite subfields, we can stitch together the isomorphisms from Propositions 3.2.1 and 3.3.1 to recover the following theorem.

As before, given a simple abelian variety A of dimension g over $\mathbb{F}_q$, we write $R = \mathrm{End}_{\mathbb{F}_q}(A)$ and define Z to be the center of R. Let $[Z : \mathbb{Z}]$ denote the rank of Z as a $\mathbb{Z}$-module. Write $S \subseteq Z$ for the set of separable isogenies in Z, and R_S (resp. Z_S) for the left R-submodule (resp. Z-submodule) of the endomorphism algebra $R \otimes \mathbb{Q}$ generated by the set $\{s^{-1} : s \in S\}$. Equivalently, these can be recognized as localizations by the set S.

Theorem 3.4.1. *For $g \geq 1$, let A be a simple abelian variety over $\mathbb{F}_q$ of dimension g. Let $R = \mathrm{End}_{\overline{\mathbb{F}}_q}(A)$, and let Z be the center of R.*

(a) If $[\mathbb{Q}(\pi) : \mathbb{Q}] = 2g$ and R is a Gorenstein ring, then

$$A(\overline{\mathbb{F}}_q) \cong R_S/R.$$

is an isomorphism of R-modules.

(b) If Z is a maximal order, then

$$A(\overline{\mathbb{F}}_q) \cong (Z_S/Z)^d.$$

is an isomorphism of Z-modules where $d = 2g/[Z : \mathbb{Z}]$. Moreover, this Z-module has

exactly one left R-module structure, up to isomorphism, and there is an isomorphism

$$A(\overline{\mathbb{F}}_q)^d \cong R_S/R$$

as R-modules.

Proof. Notice that, in any case, we have

$$A(\overline{\mathbb{F}}_q) = \bigcup_{s \in S} A[s] = \bigcup_{n \geq 1} A[\pi^n - 1] = \bigcup_{n \geq 1} A(\mathbb{F}_{q^n})$$

where π denotes the Frobenius endomorphism of A over $\mathbb{F}_q$. Indeed, it is clear that each term contains the next, and the final term equals the first. This allows us to deduce the theorem after describing only $A[s]$ for $s \in S$.

For part (a), the hypotheses allow us to apply Proposition 3.2.1 to obtain isomorphisms $A[s] \cong R/Rs \cong s^{-1}R/R$ for every separable $s \in R$. In other words, for each $s \in S$, the set W_s of isomorphisms between $A[s]$ and $s^{-1}R/R$ is nonempty. Moreover, if s and t are two separable endomorphisms such that s divides t, then the isomorphism $A[t] \xrightarrow{\sim} t^{-1}R/R$ maps the submodule $A[s]$ isomorphically to $s^{-1}R/R$. Thus the set $\{W_s\}_{s \in S}$ form a projective system of nonempty finite sets, and the projective limit of this system is nonempty [12, Théorème 1, §7.4]. In particular, there exists a simultaneous choice of isomorphisms $A[s] \to s^{-1}R/R$ for all $s \in S$ that commutes with the natural inclusions of sets, and the result follows by taking the union over all $s \in S$.

Part (b) follows similarly. Indeed, for each $s \in S$, Proposition 3.3.1 provides an isomorphism $A[s] \cong (Z/Zs)^d \cong (s^{-1}Z/Z)^d$. By the same projective limit argument given for part (a), we obtain the desired isomorphism $A(\overline{\mathbb{F}}_q) \cong (Z_S/Z)^d$. Similarly, we obtain the isomorphism $A(\overline{\mathbb{F}}_q)^d \cong R_S/R$.

Finally, any two R-module structures on $(Z_S/Z)^d$ give rise to two R-module structures on $(s^{-1}Z/Z)^d$ for each $s \in S$. Since this structure is known to be unique by Proposition 3.3.1, we obtain compatible isomorphisms for all $s \in S$, and yet again obtain the desired isomorphism through the projective limit construction. □

Chapter 4
Computing Endomorphism Rings

4.1 Main result and comparison to prior work

4.1.1 Background

Computing the endomorphism ring of an abelian variety is a fundamental problem of computational number theory with applications in cryptography. Consider the foundational case when E is an ordinary elliptic curve over $\mathbb{F}_q$. In this case, $\operatorname{End} E$ is isomorphic to an order in the quadratic imaginary field $K = \mathbb{Q}(\pi)$ where π is the Frobenius of E. Further, the orders in K which contain π are precisely the orders which arise as the endomorphism ring of an elliptic curve E' over $\mathbb{F}_q$ which is isogenous to E [110, Theorem 4.2]. Because K is a quadratic imaginary field, the orders of K are uniquely identified by their index in $\mathcal{O}_K$. Thus, computing $\operatorname{End} E$ is equivalent to computing the index $[\mathcal{O}_K : \operatorname{End} E]$, which is a divisor of $[\mathcal{O}_K : \mathbb{Z}[\pi]]$.

The first algorithm for computing $\operatorname{End} E$ for an ordinary elliptic curve E was given by Kohel in his thesis [56]. Kohel's algorithm is deterministic and has exponential expected runtime $O(q^{\epsilon+1/3})$, assuming GRH. The key fact needed for this algorithm was Kohel's discovery that certain ℓ-isogeny graphs (i.e., graphs whose vertices can be identified with isomorphism classes of elliptic curves and whose edges can be identified with isogenies of prime degree ℓ) have a special "volcano" structure. By navigating the volcano graphs, one may deduce the prime factors of $[\mathcal{O}_K : \operatorname{End} E]$, and thus determine $\operatorname{End} E$ as desired, although this method is slow when working with large prime factors.

To make a faster algorithm for computing $\operatorname{End} E$, Bisson and Sutherland exploited the following fact, which was proven by Waterhouse [110, Theorem 4.5]. Given an order $\mathcal{O} \subseteq K$, the class group $\operatorname{Cl}(\mathcal{O})$ acts faithfully on the set of isomorphism classes of elliptic curves isogenous to E with endomorphism ring isomorphic to $\mathcal{O}$, where an

ideal of norm ℓ acts via an isogeny of degree ℓ. Thus, one may deduce information about the class group of $\operatorname{End} E$ by computing various isogenies. In particular, one may compare $\mathrm{Cl}(\operatorname{End} E)$ to $\mathrm{Cl}(\mathcal{O})$ for known "testing" orders $\mathcal{O} \subseteq K$ by simply computing class group relations. Bisson and Sutherland prove that determining which relations hold in $\mathrm{Cl}(\operatorname{End} E)$ is sufficient for determining the large prime factors of $[\mathcal{O}_K : \operatorname{End} E]$, which leads to a probabilistic algorithm to determine $[\mathcal{O}_K : \operatorname{End} E]$ in subexponential time [8].

Bisson extended the algorithm above to absolutely simple, principally polarized, ordinary abelian varieties of dimension 2 in [6]. The endomorphism ring of such an abelian variety is an order in a quartic CM field. Because the isogeny graph structure of ordinary abelian varieties of dimension 2 was essentially unknown at the time, the correctness of Bisson's algorithm required multiple heuristic assumptions about the relevant isogeny graphs and polarized class groups. Additionally, these orders are no longer uniquely identified by their index, as in the elliptic curve case. Instead, Bisson identifies an order $\mathcal{O}$ by a "lattice of relations" which hold in $\mathcal{O}$.

This chapter gives a different generalization of Bisson and Sutherland's elliptic curve algorithm. Our algorithm identifies orders by ideals in the maximal totally real subfield of the endomorphism algebra, and its correctness only relies on the assumptions required for computing the class group of an order in a number field in subexponential time, as in [4]. As a trade-off, we must explicitly restrict the class of abelian varieties that we consider.

4.1.2 Main result

Let A be an absolutely simple, principally polarized, ordinary abelian variety of dimension 2 defined over a finite field $\mathbb{F}_q$ with Frobenius π. Then $\operatorname{End} A$ is an order containing $\mathbb{Z}[\pi, \overline{\pi}]$ in the quartic CM field $K = \mathbb{Q}(\pi)$. We will assume that A has maximal RM, or maximal *real multiplication*, which means that $\operatorname{End} A \supseteq \mathcal{O}_F$ where F is the maximal totally real subfield of K. According to Brooks, Jetchev and Wesolowski [13, Theorem 2.1], the orders $\mathcal{O} \subseteq K$ which contain $\mathcal{O}_F$ are in bijective correspondence with ideals of $\mathcal{O}_F$, where $\mathcal{O}$ is associated to $\mathfrak{f} \cap \mathcal{O}_F$ and $\mathfrak{f}$ is the conductor ideal of $\mathcal{O}$. This allows our algorithm to have a simple output, namely the ideal of $\mathcal{O}_F$ which uniquely identifies $\operatorname{End} A$.

The foundation of our algorithm is the free action of the ideal class group $\mathrm{Cl}(\mathcal{O})$ on the set of isomorphism classes of abelian varieties isogenous to A with endomorphism ring $\mathcal{O} \subseteq K$. Just like in the elliptic curve case, this means that a product of ideals is trivial in $\mathrm{Cl}(\operatorname{End} A)$ if and only if the corresponding isogeny maps A to an isomorphic

abelian variety. Using the same tactic as Bisson and Sutherland, we compare $\mathrm{Cl}(\mathrm{End}\, A)$ to $\mathrm{Cl}(\mathcal{O})$ for known "testing" orders $\mathcal{O} \subseteq K$ by generating class group relations. When $\mathrm{End}\, A$ contains $\mathcal{O}_F$, we show through class field theory that class group relations are sufficient for determining the large prime factors of the ideal $\mathfrak{f}^+$ which identifies the endomorphism ring $\mathrm{End}\, A$.

However, unlike the elliptic curve case, it is not always possible to compute the action of $\mathrm{Cl}(\mathcal{O})$ on principally polarized abelian varieties of dimension $g > 1$. In fact, the isogenies corresponding to elements of $\mathrm{Cl}(\mathcal{O})$ do not even preserve principal polarizability in general. Bisson avoided this obstruction for his abelian surface algorithm by instead working with a convenient subgroup of the *polarized class group* $\mathfrak{C}(\mathcal{O})$, which has a nicer action that respects polarizations. Unfortunately, the structure of $\mathfrak{C}(\mathcal{O})$ cannot be analyzed through machinery like class field theory in the same way that we can analyze $\mathrm{Cl}(\mathcal{O})$, so it is difficult to demonstrate the existence of sufficiently many relations in $\mathfrak{C}(\mathcal{O})$.

For our algorithm, we ensure that it is possible to compute the action of $\mathrm{Cl}(\mathcal{O})$ by restricting our attention to the case when F has narrow class number 1. In this case, there is a surjective map from $\mathfrak{C}(\mathcal{O})$ to $\mathrm{Cl}(\mathcal{O})$ and the action of $\mathrm{Cl}(\mathcal{O})$ preserves principal polarizability. We also assume that the index $[\mathcal{O}_F : \mathbb{Z}[\pi + \overline{\pi}]]$ is not even, and that $\mathcal{O}_K^* = \mathcal{O}_F^*$ so that all isogenies are computable via Cosset, Robert, Dudeanu, et al. [20,26], and certain isogeny graphs are volcanoes [13]. This latter fact allows us to determine the small prime factors of $\mathfrak{f}^+(A)$, much like in the elliptic curve case.

In summary, we focus our attention on absolutely simple, principally polarized, ordinary abelian surfaces A for which the following is true:

1. A has maximal real multiplication by F.
2. $\mathcal{O}_K^* = \mathcal{O}_F^*$.
3. F has narrow class number 1.
4. The conductor gap $[\mathcal{O}_F : \mathbb{Z}[\pi + \overline{\pi}]]$ is not divisible by 2.

Theorem 4.1.1. *There is a subexponential algorithm which, given an abelian variety A of dimension 2 satisfying the conditions above, outputs the ideal of $\mathcal{O}_F$ uniquely identifying* $\mathrm{End}\, A$.

Note that the restrictions imposed by Theorem 4.1.1 make this algorithm considerably less general than the algorithm of Bisson, which accepts all absolutely simple, principally

polarized, ordinary abelian surfaces, with only the exclusion of a certain zero-density set of worst-case varieties. However, our extra restrictions provide benefits as a trade-off. First, the correctness of our algorithm relies solely on the heuristic assumptions required for Biasse and Fieker's subexponential algorithm for solving the principal ideal problem [3,4], thereby avoiding the nonstandard heuristic assumptions needed for Bisson's algorithm. Additionally, the zero-density set excluded by Bisson's algorithm is not explicitly known, while the conditions for Theorem 4.1.1 are explicit and verifiable.

With the notation,

$$L[a,c](n) = \exp\left((c+o(1))\log(n)^a(\log\log n)^{1-a}\right)$$

we heuristically bound the asymptotic runtime of the algorithm by

$$L[1/2, 2c](q) + L[1/2, 2\sqrt{d+1}](q).$$

where c and d are constants corresponding to the difficulty of the principal ideal problem [3, 4] and isogeny computation [20, 26], respectively. The algorithm can also be modified to produce a short certificate which allows a third party to verify the correctness of the output. This takes subexponential time using the same heuristic assumptions as before.

4.2 Background

4.2.1 Notation

Fix an absolutely simple, principally polarized, ordinary abelian variety A of dimension g over $\mathbb{F}_q$ with Frobenius endomorphism π. Using Pila's algorithm [74], the characteristic polynomial f_π of π can be computed in time polynomial in $\log q$. Alternatively, there are faster methods for the case of maximal RM; see [2, 39, 64]. The polynomial $f_\pi(x) \in \mathbb{Z}[x]$ encodes many features of A, such as the isogeny class, the number of $\mathbb{F}_q$-rational points, and the fact that A is ordinary [99]. In the remainder of the chapter, we will assume that f_π is known and an embedding $\operatorname{End} A \hookrightarrow \mathbb{Q}(\pi) \cong \mathbb{Q}[x]/(f_\pi)$ is given for simplicity. The field $\mathbb{Q}(\pi)$ is a CM field of degree $2g$ which we denote by K. Let F denote its maximal totally real subfield. We will write $\mathcal{O}$ for an arbitrary order in K which is possibly non-maximal.

4.2.2 Identifying orders

Waterhouse [110, Theorem 7.4] showed that orders of K arising as endomorphism rings of abelian varieties isogenous to A are precisely the orders $\mathcal{O}$ satisfying $\mathbb{Z}[\pi, \overline{\pi}] \subseteq \mathcal{O} \subseteq \mathcal{O}_K$. Our goal is to determine which order is $\operatorname{End} A$, although it is not clear how to uniquely identify such orders in a simple way. In the elliptic curve case, K is a quadratic imaginary field and for each positive integer n there is precisely one order $\mathcal{O}$ of index n inside $\mathcal{O}_K$, hence the orders of K can be uniquely identified by positive integers. To identify orders in CM fields in general, we will restrict our attention to abelian varieties with *maximal RM*, i.e. $\operatorname{End} A$ is an order containing $\mathcal{O}_F[\pi]$.

Recall that, given $\mathcal{O} \subseteq \mathcal{O}_K$, the *conductor ideal* of $\mathcal{O}$ is defined as the ideal

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

Notice that $\mathfrak{f}$ is the largest subset of K which is an ideal in both $\mathcal{O}_K$ and $\mathcal{O}$. When K is a quadratic imaginary field, the conductor ideal $\mathfrak{f}$ of an order $\mathcal{O}$ is the ideal of K generated by the integer $[\mathcal{O}_K : \mathcal{O}]$. Hence the following theorem is a generalization of the fact that orders in a quadratic imaginary field are uniquely determined by their index.

Theorem 4.2.1 (Theorem 2.1, [13])**.** *Orders in K containing $\mathcal{O}_F$ are in bijective correspondence with ideals of $\mathcal{O}_F$, as follows. To each order $\mathcal{O} \subseteq K$ containing $\mathcal{O}_F$, associate the ideal $\mathfrak{f}^+ = \mathfrak{f} \cap \mathcal{O}_F$ where $\mathfrak{f} \subseteq \mathcal{O}_K$ is the conductor ideal of $\mathcal{O}$. Similarly, to any ideal $\mathfrak{f}^+ \subseteq \mathcal{O}_F$, associate the order $\mathcal{O}_F + \mathfrak{f}^+\mathcal{O}_K$.*

The following algorithm is a modification of the algorithm of Klüners and Pauli [53, §6], and it allows one to compute the identifying $\mathcal{O}_F$-ideal of any order $\mathcal{O} \subseteq \mathcal{O}_K$ containing $\mathcal{O}_F$, thereby making the correspondence of Theorem 4.2.1 computable.

Algorithm 1: COMPUTE THE IDENTIFYING IDEAL

Input : An order $\mathcal{O}$ in a CM field K of degree $2g$ with maximal totally real subfield F satisfying $\mathcal{O} \cap F = \mathcal{O}_F$, and $\mathbb{Z}$-bases $\{\alpha_1, \ldots, \alpha_g\}$, $\{\tau_1, \ldots, \tau_{2g}\}$ and $\{\omega_1, \ldots, \omega_{2g}\}$ for $\mathcal{O}_F$, $\mathcal{O}$ and $\mathcal{O}_K$, respectively.

Output : Ideal $\mathfrak{f}^+ \subseteq \mathcal{O}_F$ identifying the order $\mathcal{O} \subseteq \mathcal{O}_K$ containing $\mathcal{O}_F$.

1 Define $b_{i,j,k} \in \mathbb{Q}$ by $\alpha_i \omega_j = \sum_{k=1}^{2g} b_{i,j,k} \tau_k$

2 Define the matrix

$$M := \begin{pmatrix} M_1 \\ \vdots \\ M_{2g} \end{pmatrix}$$

where

$$M_j := \begin{pmatrix} b_{1,j,1} & \ldots & b_{g,j,1} \\ \vdots & \ddots & \vdots \\ b_{1,j,2g} & \ldots & b_{g,j,2g} \end{pmatrix}$$

3 Let d be the greatest common divisor of all $d' \in \mathbb{Z}$ with $d'M \in \mathbb{Z}^{(2g)^2 \times g}$;

4 Let $H \in \mathbb{Z}^{g \times g}$ be the row Hermite normal form of dM;

5 $\mathfrak{f}^+ = \mathbb{Z}\beta_1 + \cdots + \mathbb{Z}\beta_g$ where $(\beta_1, \ldots, \beta_g) = (\alpha_1, \ldots, \alpha_g) dH^{-1}$.

Proposition 4.2.2. *Algorithm 1 is correct.*

Proof. Let $\beta = \sum_{i=1}^{g} a_i \alpha_i$ be an element of F. By definition, β is in $\mathfrak{f}^+$ if and only if $\beta \mathcal{O}_K \subseteq \mathcal{O}$, i.e. if and only if

$$\beta \omega_j = \sum_{i=1}^{g} a_i \alpha_i \omega_j = \sum_{i=1}^{g} a_i \sum_{k=1}^{2g} b_{i,j,k} \tau_k = \sum_{k=1}^{2g} \left(\sum_{i=1}^{g} a_i b_{i,j,k} \right) \tau_k \in \mathcal{O}$$

for every $1 \leq j \leq 2g$. Equivalently, $\sum_{i=1}^{g} a_i b_{i,j,k} \in \mathbb{Z}$ for all $1 \leq j \leq 2g$ because τ_k is a $\mathbb{Z}$-basis for the order $\mathcal{O}$.

Define M_j as the "multiplication by ω_j" matrix:

$$M_j := \begin{pmatrix} b_{1,j,1} & \ldots & b_{g,j,1} \\ \vdots & \ddots & \vdots \\ b_{1,j,2g} & \ldots & b_{g,j,2g} \end{pmatrix}$$

and put

$$M := \begin{pmatrix} M_1 \\ \vdots \\ M_{2g} \end{pmatrix}$$

Then, $\beta \in \mathfrak{f}^+$ if and only if $M\vec{a} \in \mathbb{Z}^{2g \times 2g}$ where $\vec{a} = (a_1, \ldots, a_g)^{tr}$. Note that M has maximal rank g because each M_j is injective (i.e. multiplication by ω_j has no kernel).

Write d for the greatest common divisor of the integers d' such that $d'M$ has integer coefficients. Let H be the row Hermite Normal Form of dM. View H as a $g \times g$ matrix by removing all of the zero rows. Because M has maximal rank g, H is invertible. Then

$$\beta \in \mathfrak{f}^+ \iff M\vec{a} \in \mathbb{Z}^{(2g)^2} \iff dM\vec{a} \in d\mathbb{Z}^{(2g)^2} \iff H\vec{a} \in d\mathbb{Z}^g \iff \vec{a} \in dH^{-1}\mathbb{Z}^g$$

Because $\mathfrak{f}^+$ is an integral ideal, each a_i must be an integer. We deduce that dH^{-1} is an integer matrix whose columns are a basis for $\mathfrak{f}^+$, written in the basis $\{\alpha_1, \ldots, \alpha_g\}$. Thus we can write a basis of $\mathfrak{f}^+$ as

$$(\beta_1, \ldots, \beta_g) = (\alpha_1, \ldots, \alpha_g)dH^{-1}.$$

□

Before analyzing the running time of Algorithm 1, we must choose integral bases which have multiplication tables of small size. Given a basis $b_1, \ldots, b_n$ for an order of discriminant Δ in a number field of degree n, the *multiplication table* $(x_{i,j,k})$ is defined by $b_i b_j = \sum_{k=1}^n x_{i,j,k} b_k$. There is a polynomial time basis reduction algorithm that provides a basis whose multiplication table has size $O(n^4(2 + \log|\Delta|))$; see [59, §2.10] or [14, §5]. We apply this to $\mathcal{O}_F$ and $\mathcal{O}_K$ to get reduced bases $\{\alpha_1, \ldots, \alpha_g\}$ and $\{\omega_1, \ldots, \omega_{2g}\}$, respectively. Any $\mathcal{O} \subset \mathcal{O}_K$ can be represented by a basis $\{a_1\omega_1, \ldots, a_{2g}\omega_{2g}\}$ where $|a_i| \leq [\mathcal{O}_K : \mathcal{O}]$.

Proposition 4.2.3. *Assuming that the bases for $\mathcal{O}_K$, $\mathcal{O}$ and $\mathcal{O}_F$ are chosen as outlined above, Algorithm 1 has running time polynomial in g and* $\log|\operatorname{disc}(\mathcal{O})|$.

Proof. Consider the integers $c_{i,j,k}$ defined by

$$\alpha_i\omega_j = \sum_{k=1}^{2g} c_{i,j,k}\omega_k. \tag{4.1}$$

We see that $c_{i,j,k} = a_k b_{i,j,k}$ by the choice of bases. Write $b_{i,j,k}$ as a rational number in

reduced form:

$$b_{i,j,k} = \frac{c_{i,j,k}}{a_k} = \frac{c_{i,j,k}/\gcd(c_{i,j,k}, a_k)}{a_k/\gcd(c_{i,j,k}, a_k)}.$$

Hence the integer d used in Step 3 is

$$d = \operatorname{lcm}\left\{\frac{a_k}{\gcd(a_k, c_{i,j,k})}\right\}.$$

Algorithm 1 essentially consists of computing the integer d in Step 3, the Hermite Normal Form H of the matrix dM in Step 4, and the inverse H^{-1} in Step 5. Each of these steps are polynomial time in the dimension and the size of the coefficients. We will modify [14, Proposition 5.3] to bound the size of $c_{i,j,k}$ by a polynomial in g, $\log|\operatorname{disc}(\mathcal{O}_K)|$ and $\log|\operatorname{disc}(\mathcal{O}_F)|$. Because $[\mathcal{O}_K : \mathcal{O}]$, $\log|\operatorname{disc}(\mathcal{O}_K)|$ and $\log|\operatorname{disc}(\mathcal{O}_F)|$ are all bounded by $\log|\operatorname{disc}(\mathcal{O})|$, this is sufficient for bounding the running time of the algorithm.

Recall that there are $2g$ conjugate pairs of complex embeddings of K, which we denote $\sigma_1, \ldots, \sigma_g, \overline{\sigma}_1, \ldots, \overline{\sigma}_g$ This provides an embedding $K \to \mathbb{R}^{2g}$ via

$$\alpha \mapsto \underline{\alpha} := (\operatorname{Re}(\sigma_1(\alpha)), \operatorname{Im}(\sigma_1(\alpha)), \ldots, \operatorname{Re}(\sigma_g(\alpha)), \operatorname{Im}(\sigma_g(\alpha))).$$

Embedding the equation (4.1) gives the linear equation

$$\underline{\alpha_i \omega_j} = \sum_{k=1}^{2g} c_{i,j,k} \underline{\omega_k}.$$

The matrix $(\underline{\omega_k})_k$ has determinant $2^{-g}|\operatorname{disc}(\mathcal{O}_K)|^{1/2}$ by definition. Hence by Cramer's rule

$$|c_{i,j,k}| \leq \frac{\|\underline{\alpha_i \omega_j}\| \cdot \prod_{r \neq k} \|\underline{\omega_r}\|}{2^{-g}|\operatorname{disc}(\mathcal{O}_K)|^{1/2}}$$

Because the bases for $\mathcal{O}_K$ and $\mathcal{O}_F$ are reduced, the norms of the vectors appearing in the numerator are bounded. Specifically, $\log\|\underline{\omega_k}\|$ and $\log\|\underline{\alpha_i}\|$ are bounded by a polynomial in g, $\log|\operatorname{disc}(\mathcal{O}_K)|$ and $\log|\operatorname{disc}(\mathcal{O}_F)|$; see [14, Proposition 5.2]. We conclude that the sizes of the integers $c_{i,j,k}$ are bounded as desired. □

For the purposes of this chapter, we want bound the running time in terms of the size of the finite field $\mathbb{F}_q$. This is achieved in the following corollary. Even though we will consider g to be a constant in the remainder of the chapter, we note that the algorithm is also polynomial in g.

Corollary 4.2.4. *Maintain the notation of the previous proposition, but additionally assume $K = \mathbb{Q}(\pi)$ where π is the Frobenius of a simple ordinary abelian variety A of dimension g over $\mathbb{F}_q$. If $\mathcal{O}$ contains π, then Algorithm 1 computes the ideal $\mathfrak{f}^+$ identifying $\mathcal{O}$ in time polynomial in g and* $\log q$.

Proof. Since the algorithm is proven to be polynomial in g and $\log|\operatorname{disc}(\mathcal{O})|$, we simply need to notice that $\mathcal{O} \supseteq \mathbb{Z}[\pi]$ implies

$$|\operatorname{disc}(\mathcal{O})| \leq |\operatorname{disc}(\mathbb{Z}[\pi])| \leq (2\sqrt{q})^{2g(2g-1)}.$$

The second part of the inequality follows immediately from the definition of discriminant because each root of the characteristic polynomial of π has absolute value $\sqrt{q}$. Upon taking logarithms, we obtain our result. □

For convenient notation, denote the ideal identifying an order $\mathcal{O} \supseteq \mathcal{O}_F$ by $\mathfrak{f}^+(\mathcal{O})$, and denote the order identified by $\mathfrak{f}^+$ as $\mathcal{O}(\mathfrak{f}^+)$. Denote by $\mathfrak{f}^+(A)$ the ideal identifying the order isomorphic to $\operatorname{End} A$ under the fixed embedding $\operatorname{End} A \hookrightarrow K$. Now computing $\operatorname{End} A$ is equivalent to computing $\mathfrak{f}^+(A)$, which is a divisor of $\mathfrak{f}^+(\mathcal{O}_F[\pi])$. Hence we consider the factors of $\mathfrak{f}^+(\mathcal{O}_F[\pi])$ and determine which prime powers divide $\mathfrak{f}^+(A)$. We will present different methods for dealing with small and large prime factors.

4.2.3 Class group action

Following Brooks, Jetchev and Wesolowski [13], we will focus on $\mathfrak{l}$-isogenies, defined below. Assume that A has maximal RM and write $\mathfrak{v}$ for the ideal $\mathfrak{f}^+(\mathcal{O}_F[\pi])$. Given $\alpha \in \operatorname{End} A$, we write $A[\alpha]$ for the kernel of α.

Definition 4.2.5 ($\mathfrak{l}$-isogeny)**.** *Let $\mathfrak{l} \subseteq \mathcal{O}_F$ be a prime ideal coprime to $\mathfrak{f}^+(A)$. An $\mathfrak{l}$-isogeny from A is an isogeny whose kernel is a proper subgroup of $A[\mathfrak{l}] = \cap_{\alpha \in \mathfrak{l}} A[\alpha]$ which is stable under the action of $\mathcal{O}_F$.*

The following lemma was stated without proof in the definition of $\mathfrak{l}$-isogeny in [13]. We record a proof here for completeness.[1]

Lemma 4.2.6. *Using the notation above, the degree of an $\mathfrak{l}$-isogeny is the norm $N_{F/\mathbb{Q}}(\mathfrak{l})$.*

Proof. Recall that $\#A[\mathfrak{l}] = N_{K/\mathbb{Q}}(\mathfrak{l}) = N_{F/\mathbb{Q}}(\mathfrak{l})^2$. A proof of this equality when $\operatorname{End} A$ is maximal is given in [110, Theorem 3.15] and the proof generalizes immediately. The

[1] The author thanks Benjamin Wesolowski for explaining this proof in private communication.

action of $\mathcal{O}_F$ on $A[\mathfrak{l}]$ induces an action of the field $\mathcal{O}_F/\mathfrak{l}$. In particular, $A[\mathfrak{l}]$ is an $\mathcal{O}_F/\mathfrak{l}$-vector space of dimension 2. Therefore proper $\mathcal{O}_F$-stable subgroups are precisely the 1-dimensional $\mathcal{O}_F/\mathfrak{l}$-subspaces of $A[\mathfrak{l}]$. Thus, the size of the kernel of an $\mathfrak{l}$-isogeny is $\#(\mathcal{O}_F/\mathfrak{l}) = N(\mathfrak{l})$. For separable isogenies, this is the same as the degree of the isogeny. □

Notice that the definition of an $\mathfrak{l}$-isogeny generalizes the definition of an ℓ-isogeny for elliptic curves. This definition is useful because of the following theorems about $\mathfrak{l}$-isogeny graphs, which mirror the foundational theorems used in Bisson and Sutherland's original algorithm for elliptic curves [8, §2.1]. Write $\mathcal{O}(\mathfrak{f}^+)$ for the order of K defined by the ideal $\mathfrak{f}^+ \subseteq \mathcal{O}_F$ and $\mathrm{Cl}(\mathfrak{f}^+)$ for the class group of $\mathcal{O}(\mathfrak{f}^+)$. Denote by $\mathrm{Ab}_{\pi,\mathfrak{f}^+}$ the set of abelian varieties defined over $\mathbb{F}_q$ in the isogeny class defined by π whose endomorphism ring has identifying ideal $\mathfrak{f}^+$. Given an ideal $\mathfrak{l} \subseteq \mathcal{O}_F$, define the symbol $(K/\mathfrak{l})$ to be $1, 0$, or -1 when $\mathfrak{l}$ is split, ramified, or inert in K, respectively.

Theorem 4.2.7. *There is a faithful action of* $\mathrm{Cl}(\mathfrak{f}^+)$ *on* $\mathrm{Ab}_{\pi,\mathfrak{f}^+}$*, where an ideal lying over a prime* $\mathfrak{l} \subseteq \mathcal{O}_F$ *acts by an* $\mathfrak{l}$*-isogeny.*

Proof. The action of the class group is a classical result [87, 110]. The fact that these isogenies are $\mathfrak{l}$-isogenies is clear from the definition, as pointed out in [13, Theorem 4.3]. □

Theorem 4.2.8. *Let* $\mathfrak{l}$ *be a prime of* $\mathcal{O}_F$ *which does not divide* $\mathfrak{v} = \mathfrak{f}^+(\mathcal{O}_F[\pi])$*, and let* $A \in \mathrm{Ab}_{\pi,\mathfrak{f}}$*. Then there are exactly* $1 + (K/\mathfrak{l})$ $\mathfrak{l}$*-isogenies starting from* A *which lead to varieties with endomorphism ring isomorphic to* $\mathcal{O}(\mathfrak{f})$*, and these are the only varieties defined over* $\mathbb{F}_q$ *which are* $\mathfrak{l}$*-isogenous to* A*.*

Proof. This follows immediately from Theorem 4.3 and Proposition 4.10 in [13]. □

4.2.4 Computing isogenies

Given an ideal $\mathfrak{a} \subseteq \mathrm{End}\, A$, we must determine the isogeny $\phi_{\mathfrak{a}}$ which corresponds to the action of $\mathfrak{a}$. One option would be to simply compute all possible isogenies, in the style of [8]. However, we do not have the benefit of an easily calculable modular polynomial when the dimension of A is $g > 1$. Instead, we compute the target of an isogeny from its kernel. We can determine the $\ker \phi_{\mathfrak{a}} = \cap_{a \in \mathfrak{a}} A[a]$ in the same way as Bisson [6]. If $\mathfrak{L} \subseteq \mathcal{O}$ is a prime ideal lying over ℓ which does not divide $[\mathcal{O}_K : \mathbb{Z}[\pi]]$, then we can write $\mathfrak{L} = \ell\mathcal{O} + r(\pi)\mathcal{O}$ where $r(x)$ is a factor modulo ℓ of the characteristic polynomial f_π of Frobenius. Then the kernel of $\phi_{\mathfrak{L}}$ is simply the kernel of the isogeny $r(\pi)$ restricted to $A[\ell]$.

After obtaining the subgroup corresponding to the given ideal, we must compute the target variety. This difficulty will be one of the main bottlenecks of the algorithm. We will only consider computing isogenies for the case $g = 2$. Consider a prime $\mathfrak{l} \subseteq \mathcal{O}_F$ lying over ℓ, and let ϕ be an $\mathfrak{l}$-isogeny whose kernel is known. There are two cases, depending on the norm of the ideal $\mathfrak{l}$. If $N\mathfrak{l} = \ell^2$, then ϕ is commonly called an (ℓ, ℓ)-*isogeny*, and the target variety can be determined by the algorithm in [20]. If $N\mathfrak{l} = \ell$, then ϕ is known as a *cyclic isogeny* and the target variety can be determined by the algorithm in [25, Chapter 4] and [26]. In both cases, the algorithm is polynomial in ℓ and $\log q$. By an algorithm of Mestre [68], we assume that the output of these algorithms is the Jacobian of a curve defined over the minimal field; see also [20]. By combining these algorithms and analyzing the running times of each algorithm, we find the following theorem.

Theorem 4.2.9. *Let A be an ordinary, principally polarized abelian variety defined over $\mathbb{F}_q$ of dimension $g = 2$ with maximal RM. Write $K = \mathbb{Q}(\pi)$, where π is Frobenius, and let F be the maximal totally real subfield of K. Let $\mathfrak{l} \subseteq \mathcal{O}_F$ be a prime lying over ℓ and assume the following:*

1. *$\mathfrak{l} = \beta\mathcal{O}_F$ for a totally positive element β.*
2. *The index $[\mathcal{O}_F : \mathbb{Z}[\pi + \overline{\pi}]]$ is not divisible by 2ℓ.*

¥ If ℓ is bounded by $L[1/2, c_0](q)$ for some constant $c_0 > 0$, then an $\mathfrak{l}$-isogeny with a given kernel can be found in time $L[1/2, dc_0](q)$ where $d > 0$ is a constant.

The extra assumptions in Theorem 4.2.9 are required only for isogenies of degree ℓ, which are called cyclic isogenies. These isogenies are the hardest case to handle when computing $\mathfrak{l}$-isogenies. If the extra assumptions are dropped, then it is possible that the target variety does not admit any principal polarization, which presents a major problem to computing the isogeny. Bisson avoided this problem by using CM-types and reflex fields to generate relations which are easily computable [6, §4] . However, it is only known that these easily computed relations are sufficient for determining the endomorphism ring under certain heuristic assumptions [6, Theorem 7.1].

4.2.5 Navigating isogeny graphs and identifying abelian varieties

We need a way to identify abelian varieties as we navigate the various $\mathfrak{l}$-isogeny graphs with the class group action. To do this, we follow the ideas of [49, §4.2]. Rosenhain

invariants[2] can be used to identify isomorphism classes of principally polarized simple ordinary abelian varieties of dimension $g = 2$ over $\mathbb{F}_q$, i.e. pairs (A, λ) where A is an abelian variety over $\mathbb{F}_q$ with a fixed principal polarization λ. Meanwhile, $\mathrm{Cl}(\mathcal{O})$ acts on unpolarized abelian varieties, i.e. abelian varieties A with no fixed polarization. We say that A is *principally polarizable* if a principal polarization λ exists, but is not fixed. By definition, a single isomorphism class of principally polarizable abelian varieties can be partitioned into isomorphism classes of principally polarized abelian varieties according to the different choices for principal polarization. Thus the isomorphism class of A as an unpolarized principally polarizable abelian variety is uniquely represented by the list of invariants which correspond to the different polarizations of A, up to isomorphism.

To investigate how the class group action relates to the different choices of polarizations, it is useful to remember the *polarized class group*[3], as discussed in [6, 13, 49]. For a given order $\mathcal{O} \subseteq K$, the polarized class group is defined as follows.

$$\mathfrak{C}(\mathcal{O}) = \{(\mathfrak{a}, \alpha) \mid \mathfrak{a} \subseteq \mathcal{O}, \mathfrak{a}\overline{\mathfrak{a}} = \alpha\mathcal{O}, \alpha \in F \text{ totally positive}\}/\sim .$$

Multiplication is performed component-wise, and $(\mathfrak{a}, \alpha) \sim (\mathfrak{b}, \beta)$ if there is an element $u \in K^\times$ such that $\mathfrak{a} = u\mathfrak{b}$ and $\beta = u\overline{u}\alpha$. Given an order $\mathcal{O} \subseteq K$ containing $\mathcal{O}_F$, the structure of $\mathfrak{C}(\mathcal{O})$ can be seen through the following exact sequence

$$1 \to (\mathcal{O}_F^\times)^+/N_{K/F}(\mathcal{O}^\times) \xrightarrow{u\mapsto(\mathcal{O},u)} \mathfrak{C}(\mathcal{O}) \xrightarrow{(\mathfrak{a},\alpha)\mapsto\mathfrak{a}} \mathrm{Cl}(\mathcal{O}) \xrightarrow{\mathfrak{a}\mapsto N_{K/F}(\mathfrak{a})} \mathrm{Cl}^+(\mathcal{O}_F), \tag{4.2}$$

where $(\mathcal{O}_F^\times)^+$ is the set of totally positive units in $\mathcal{O}_F$, $N_{K/F}$ is the relative norm from K to F, and $\mathrm{Cl}^+(\mathcal{O}_F)$ is the narrow class group of F. While the usual class group $\mathrm{Cl}(\mathcal{O})$ acts on isomorphism classes of abelian varieties, the polarized class group $\mathfrak{C}(\mathcal{O})$ acts on isomorphism classes of principally polarized abelian varieties. The image of $\mathfrak{C}(\mathcal{O})$ inside of $\mathrm{Cl}(\mathcal{O})$ is the subgroup of ideals which act freely on the set of principally polarizable abelian varieties. By assuming that $\mathrm{Cl}^+(\mathcal{O}_F)$ is trivial, we ensure that $\mathfrak{C}(\mathcal{O}) \to \mathrm{Cl}(\mathcal{O})$ is surjective, hence all isogenies arising from the action of $\mathrm{Cl}(\mathcal{O})$ preserve principal polarizability, which is necessary for the sake of computing isogenies. The exact sequence also shows how to count the number of isomorphism classes of principal polarizations.

Proposition 4.2.10. *If A is a principally polarizable ordinary simple abelian variety*

[2]There are other alternative invariants that can be used in a similar way, such as Gündlach or Igusa invariants. For our purposes, the choice of invariants used is not important.

[3]Also known as the *Shimura class group*.

over $\mathbb{F}_q$ *with maximal RM and* $\operatorname{End} A = \mathcal{O}$*, then the group*

$$\mathcal{U}(\mathcal{O}) := (\mathcal{O}_F^\times)^+ / N_{K/F}(\mathcal{O}^\times)$$

acts freely and transitively on the set of isomorphism classes of principal polarizations of A*. When the dimension of* A *is* $g = 2$*, then* $\mathcal{U}(\mathcal{O})$ *is an* $\mathbb{F}_2$*-vector space of dimension* $0 \leq d \leq 1$ *and* d *only depends on* F *and* K*.*

Proof. This Proposition summarizes Proposition 5.4 and Lemma 5.5 of [13]. □

This proposition implies that one may check whether an ordinary principally polarized abelian surface A is fixed under the action of an ideal $\mathfrak{a}$ by simply computing the list of one or two invariants for the target variety with different polarizations, and checking if the invariants of A are on this list.

There are two practical improvements to mention. First, if $\mathfrak{L} \subseteq \mathcal{O}$ corresponds to an (ℓ, ℓ)-isogeny, then one can compute $\phi_{\mathfrak{L}}$ as polarization-preserving by working with the element $(\mathfrak{L}, \ell) \in \mathfrak{C}(\mathcal{O})$, as in [6]. Thus, it is always sufficient to only check one invariant for (ℓ, ℓ)-isogenies. Second, if one is computing a chain of isogenies $\phi_{\mathfrak{L}_1}, \ldots, \phi_{\mathfrak{L}_k}$ corresponding to prime ideals $\mathfrak{L}_i$ and $\phi_{\mathfrak{L}_k}$ is a cyclic isogeny, then we can simply make an arbitrary choice of polarization for the target varieties at each step $\phi_{\mathfrak{L}_1}, \ldots, \phi_{\mathfrak{L}_{k-1}}$, and compute both polarizations for $\phi_{\mathfrak{L}_k}$, if necessary. Hence, the issue of polarizations will not change the complexity of any algorithms, and we will let this detail be implicit in the remainder of our exposition.

4.2.6 Small primes

We only use the class group action to determine the power of $\mathfrak{p}$ dividing $\mathfrak{f}^+(\operatorname{End} A)$ when $\mathfrak{p}$ is a prime ideal dividing $\mathfrak{v} = \mathfrak{f}^+(\mathcal{O}_F[\pi])$ which has large norm. For small primes $\mathfrak{p}$ such that $\mathfrak{p}^k \mid \mathfrak{v}$, one could follow Bisson [6] and use the method of Eisenträger and Lauter [29] to test whether $\operatorname{End} A$ contains $\mathcal{O}(\mathfrak{p}^k)$. However, this method does not immediately produce an isogenous abelian variety A' such that $\mathfrak{f}^+(\operatorname{End} A)$ is not divisible by the small prime factors of $\mathfrak{v}$, which is a feature in Bisson and Sutherland's original elliptic curve algorithm [8]. For this, we need the following additional result about isogeny graphs. More background may be found in [13].

Theorem 4.2.11 ([13, Theorem 4.3])**.** *Let* V *be any connected component of the* $\mathfrak{l}$*-isogeny graph for the isogeny class of an ordinary, absolutely simple abelian variety* A *with Frobenius* π *and maximal RM by* F*.*

1. *The graph V consists of levels $\{V_i\}_{i\geq 0}$ such that each level V_i shares a common endomorphism ring $\mathcal{O}_i$ and the valuation at $\mathfrak{l}$ of $\mathfrak{f}^+(\mathcal{O}_i)$ is i. The valuation of $\mathfrak{f}^+(\mathcal{O}_i)$ at other primes is the same for all i, and the number of levels is equal to the valuation at $\mathfrak{l}$ of $\mathfrak{f}^+(\mathcal{O}_F[\pi])$.*

2. *The graph V is an $N(\mathfrak{l})$-volcano if and only if $\mathcal{O}_0^\times \subseteq F$ and $\mathfrak{l}$ is principal in $\mathcal{O}_0 \cap F$.*

This theorem implies that finding the power of each $\mathfrak{l} \mid \mathfrak{v}$ which divides $\mathfrak{f}^+(A)$ is equivalent to finding the level of A in the $\mathfrak{l}$-isogeny graph. When navigating the graph, repeatedly moving from level V_{i+1} to level V_i until reaching V_0 is known as *ascending* the graph. We know that the $\mathfrak{l}$-isogeny graph is a volcano in our applications with no additional assumptions because we will assume in our algorithms that $\mathrm{Cl}^+(F) = 1$ and $\mathcal{O}_F^\times = \mathcal{O}_K^\times$ for the sake of Theorem 4.2.9 and Lemma 4.3.11. Hence the algorithms presented for volcano navigation in [98] immediately generalize and provide a way to find the level of a variety and ascend the graph. We call this method *isogeny climbing*, following the terminology given in the elliptic curve case. Isogeny climbing allows one to determine the power of $\mathfrak{l}$ dividing $\mathfrak{f}^+(A)$, and also allow one to find a new abelian variety A' such that $\mathfrak{f}^+(A')$ is the same as $\mathfrak{f}^+(A)$ with a given small prime factor removed. Notice that these methods are only efficient for small primes because we have to compute a large number of isogenies.

4.3 Class group relations

Let A be an absolutely simple, ordinary, principally polarized abelian variety with maximal RM, of dimension g over $\mathbb{F}_q$ with Frobenius π. We do not restrict the dimension g or the class group of F for this section because no simplicity is gained by focusing on special cases. Moreover, we recover many of the statements of [8] when setting $g = 1$. As before, write $K = \mathbb{Q}(\pi)$ and let F be the maximal totally real subfield. We begin by recalling a well-known fact about class groups.

Lemma 4.3.1. *If $\mathcal{O}(\mathfrak{f}_1^+) \subseteq \mathcal{O}(\mathfrak{f}_2^+)$, then there is a surjective map $\mathrm{Cl}(\mathcal{O}(\mathfrak{f}_1^+)) \to \mathrm{Cl}(\mathcal{O}(\mathfrak{f}_2^+))$ induced by mapping $\mathfrak{a} \mapsto \mathfrak{a}\mathcal{O}(\mathfrak{f}_2^+)$.*

Proof. For any order $\mathcal{O} \subseteq K$ of conductor $\mathfrak{f}$, let $I_{K,\mathcal{O}}(\mathfrak{f})$ be the group generated by ideals of $\mathcal{O}$ coprime to $\mathfrak{f}$ and let $I_K(\mathfrak{f})$ be the group generated by ideals of $\mathcal{O}_K$. There is an isomorphism $I_{K,\mathcal{O}}(\mathfrak{f}) \to I_K(\mathfrak{f})$ given by the map $\mathfrak{a} \mapsto \mathfrak{a}\mathcal{O}_K$ where the inverse is given by $\mathfrak{b} \mapsto \mathfrak{b} \cap \mathcal{O}$ [63, Proposition 3.4]. This map induces the surjective map of class groups. □

This map allows us to define the concept of a relation analogously to Bisson [5, 6]. Note that we do not follow Bisson and Sutherland's definition of relation in [8] because we do not have the benefit of an easily computable modular polynomial.

Definition 4.3.2. *A **relation** R is a tuple of ideals $(\mathfrak{a}_1, \ldots, \mathfrak{a}_k)$ of ideals of $\mathcal{O}_F[\pi]$. A relation **holds** in an order $\mathcal{O}$ if the product is trivial under the map from the lemma. Similarly, a relation **holds** for A if the corresponding composition of isogenies fixes A.*

Combining the lemma and the definition of relation, we obtain the following key corollary which is analogous to [8, Corollary 4]. This allows us to create an algorithm where testing class group relations determines the prime powers $\mathfrak{p}^k$ which divide $\mathfrak{f}^+(A)$.

Corollary 4.3.3. *Let $\mathfrak{v} \subseteq \mathcal{O}_F$ be an ideal which is divisible by a prime power $\mathfrak{p}^k \subseteq \mathcal{O}_F$. Write*

$$\mathfrak{f}_1^+ = \mathfrak{p}^{k-1-v_{\mathfrak{p}}(\mathfrak{v})}\mathfrak{v};$$

$$\mathfrak{f}_2^+ = \mathfrak{p}^k.$$

Assume there is some relation R which holds in $\mathcal{O}(\mathfrak{f}_1^+)$ but not in $\mathcal{O}(\mathfrak{f}_2^+)$. Given any $\mathfrak{f}^+ \subseteq \mathcal{O}_F$ which divides $\mathfrak{v}$, $\mathfrak{p}^k \mid \mathfrak{f}^+$ if and only if the relation R does not hold in $\mathcal{O}(\mathfrak{f}^+)$.

Proof. If $\mathfrak{p}^k \mid \mathfrak{f}^+$, then $\mathcal{O}(\mathfrak{f}^+) \subseteq \mathcal{O}(\mathfrak{f}_2^+)$, so the lemma implies the relation R does not hold in $\mathcal{O}(\mathfrak{f}^+)$. Conversely, suppose R does not hold in $\mathcal{O}(\mathfrak{f}^+)$. The lemma implies $\mathcal{O}(\mathfrak{f}_1^+) \not\subseteq \mathcal{O}(\mathfrak{f}^+)$. But $\mathfrak{f}_1$ is the same as $\mathfrak{v}$, except that the power of $\mathfrak{p}$ is decremented. Because $\mathfrak{f}^+$ divides $\mathfrak{v}$, we deduce that $\mathfrak{p}^k$ divides $\mathfrak{f}^+$. □

4.3.1 Review of class field theory

In order to prove that an algorithm based on Corollary 4.3.3 is unconditionally correct, we need to find infinitely many relations R that satisfy the necessary conditions. This is the reason why we use the traditional class group $\mathrm{Cl}(\mathcal{O})$ instead of the polarized class group $\mathfrak{C}(\mathcal{O})$ used by Bisson. Specifically, we can use ring class fields and class field theory to understand the structure of $\mathrm{Cl}(\mathcal{O})$ because prime ideals of $\mathcal{O}$ are principal if and only if they split completely in the ring class field of K; see Corollary 4.3.7 below.

Throughout this section, K remains a fixed CM-field, as before. Let us review the definitions and notation of class field theory so that we can recall the correspondence between finite abelian extensions of K and certain subgroups of fractional ideals of K. We refer the reader to [48, 63, 72] for additional background on class field theory.

For our purposes, a *modulus* $\mathfrak{m}$ for K is an ideal of $\mathcal{O}_K$. In general, a modulus can include real infinite primes which correspond to real embeddings, but CM-fields do not have any real embeddings, hence we will ignore the infinite primes in this exposition for simplicity. Write I_K for the group of all fractional ideals in K. Given a fixed modulus $\mathfrak{m}$, we denote by $I_K(\mathfrak{m})$ the subgroup of I_K generated by ideals of $\mathcal{O}_K$ coprime to $\mathfrak{m}$. Similarly, let P_K be the group of all principal ideals in K. Write $P_{K,\mathcal{O}}(\mathfrak{m})$ for the group generated by

$$\{\alpha\mathcal{O}_K : \alpha \in \mathcal{O}, \quad \alpha\mathcal{O} + \mathfrak{m} = \mathcal{O}\}$$

and $P_{K,1}(\mathfrak{m})$ for the group generated by

$$\{\alpha\mathcal{O}_K : \alpha \in \mathcal{O}_K, \ \alpha \equiv 1 \bmod \mathfrak{m}\mathcal{O}_K\}.$$

Note that the definition of $P_{K,1}(\mathfrak{m})$ must be modified if $\mathfrak{m}$ is divisible by infinite primes, but this is the simplest definition in our case because we will always assume $\mathfrak{m}$ is a product of finite primes; see [63, Lemma 3.5].

A subgroup $H \subset I_K$ is a *congruence subgroup (defined modulo $\mathfrak{m}$)* if

$$P_{K,1}(\mathfrak{m}) \subseteq H \subseteq I_K(\mathfrak{m}).$$

If $\mathfrak{m}$ divides $\mathfrak{m}'$, then $H \cap I_K(\mathfrak{m}')$ is also a congruence subgroup, but it is defined modulo $\mathfrak{m}'$ rather than modulo $\mathfrak{m}$. In this case, we call $H \cap I_K(\mathfrak{m}')$ a *restricted* congruence subgroup. This leads us to introduce the following equivalence relation. If H_1 and H_2 are congruence subgroups modulo $\mathfrak{m}_1$ and $\mathfrak{m}_2$, respectively, then we say H_1 and H_2 are *equivalent* if they have a common restriction, i.e. if there is a modulus $\mathfrak{m}_3$ such that $H_1 \cap I_K(\mathfrak{m}_3) = H_2 \cap I_K(\mathfrak{m}_3)$. In this case, there is an isomorphism $I_K(\mathfrak{m}_1)/H_1 \cong I_K(\mathfrak{m}_2)/H_2$ [48, Lemma V.6.1]. For each equivalence class $\mathbf{H}$ of congruence subgroups, there is a unique modulus $\mathfrak{f}$ and a congruence subgroup $H_\mathfrak{f}$ defined modulo $\mathfrak{f}$ such that $H_\mathfrak{f} \in \mathbf{H}$, and $\mathfrak{f}$ divides the defining modulus of every congruence subgroup in $\mathbf{H}$ [48, Lemma V.6.2]. Such an $\mathfrak{f}$ is called the *conductor* of $\mathbf{H}$.

Let L be a finite abelian extension of K, and let I_K^S be the subgroup of I_K generated by prime ideals which do not ramify in L. There is an *Artin map*

$$\Phi_{L/K} : I_K^S \to \mathrm{Gal}(L/K)$$

where a prime $\mathfrak{p}$ is sent to the *Frobenius automorphism* $\sigma_\mathfrak{p} \in \mathrm{Gal}(L/K)$. Specifically, $\sigma_\mathfrak{p}$ is the unique automorphism such that $\sigma_\mathfrak{p}(x) \equiv x^{N(\mathfrak{p})} \bmod \mathfrak{P}$ for all $x \in \mathcal{O}_L$ for any

ideal $\mathfrak{P} \subseteq \mathcal{O}_L$ lying over $\mathfrak{p}$. The map is extended to I_K^S multiplicatively. Write $\Phi_{L/K,\mathfrak{m}}$ for the restriction of $\Phi_{L/K}$ to the subgroup $I_K(\mathfrak{m})$ where $\mathfrak{m}$ is a modulus divisible by all primes of K which ramify in L. We say that the *reciprocity law holds* for $(L, K, \mathfrak{m})$ if $\ker \Phi_{L/K} \supseteq P_{K,1}(\mathfrak{m})$. In this case, $\ker \Phi_{L/K,\mathfrak{m}}$ is a congruence subgroup defined modulo $\mathfrak{m}$.

Definition 4.3.4. *Let L be an abelian extension of K. Write $\mathbf{H}(L/K)$ for the equivalence class of all congruence subgroups $H_{\mathfrak{m}}(L/K) = \ker \Phi_{L/K,\mathfrak{m}}$ where $\mathfrak{m}$ is modulus such that the reciprocity law holds for $(L, K, \mathfrak{m})$. The* class field theory conductor $\mathfrak{f}(L/K)$ *of the extension L/K is the conductor of the equivalence class $\mathbf{H}(L/K)$, i.e. the greatest common divisor of all moduli defining congruence subgroups in $\mathbf{H}(L/K)$.*

Theorem 4.3.5 (The Classification Theorem, [48, Theorem V.9.9])**.** *The correspondence $L \mapsto \mathbf{H}(L/K)$ is a one-to-one, inclusion-reversing correspondence between finite abelian extensions L of K and equivalence classes of congruence groups of K.*

Lv and Deng showed the following consequences of the classification theorem.

Theorem 4.3.6 ([63, Theorem 4.2])**.** *Let $\mathcal{O} \subseteq K$ be an order with conductor $\mathfrak{f}$. Then there exists a unique abelian extension L of K such that all primes of K ramified in L divide $\mathfrak{f}$, and the Artin map*

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

satisfies $\ker \Phi_{L/K,\mathfrak{f}} = P_{K,\mathcal{O}}(\mathfrak{f})$*, providing an isomorphism*

$$\mathrm{Cl}(\mathcal{O}) \cong \mathrm{Gal}(L/K)$$

The field L is called the ring class field of $\mathcal{O}$. In the case where $\mathcal{O} = \mathcal{O}_K$, the ring class field coincides with the Hilbert class field of K, which is the maximal abelian unramified extension of K.

By observing the basic properties of the Artin map, this theorem provides the following corollary.

Corollary 4.3.7 ([63, Corollary 4.4])**.** *Let $\mathcal{O} \subseteq K$ be an order. If $\mathfrak{p} \subseteq \mathcal{O}$ is a prime ideal coprime to $\mathfrak{f}$, then $\mathfrak{p}$ is principal if and only if $\mathfrak{p}$ splits completely in the ring class field of $\mathcal{O}$.*

We conclude this section by making one final observation about class field theory conductors.

Lemma 4.3.8. *If L_1 and L_2 are finite abelian extensions of K such that $L_2 \subseteq L_1$, then the class field theory conductor $\mathfrak{f}(L_2/K)$ divides $\mathfrak{f}(L_1/K)$.*

Proof. By definition, $\mathfrak{f}(L_2/K)$ divides every modulus for which a congruence subgroup in $\mathbf{H}(L_2/K)$ is defined, so we simply need to show that there is a congruence subgroup in $\mathbf{H}(L_2/K)$ defined modulo $\mathfrak{f}(L_1/K)$. This is equivalent to showing that $\ker \Phi_{L_2/K}$ contains $P_{K,1}(\mathfrak{f}(L_1/K))$. This is easy to see by the inclusion-reversing correspondence of Theorem 4.3.5 and the definition of $\mathfrak{f}(L_1/K)$, which imply

$$\ker \Phi_{L_2/K} \supseteq \ker \Phi_{L_1/K} \supseteq P_{K,1}(\mathfrak{f}(L_1/K)).$$

□

4.3.2 Existence of relations

Using the ring class fields defined in the preceding section, we will now find infinitely many relations sufficient for Corollary 4.3.3. To begin, we prove the following technical lemma.

Lemma 4.3.9. *Let $\mathcal{O}(\mathfrak{f}^+)$ be the order of K which contains $\mathcal{O}_F$ and corresponds to the ideal $\mathfrak{f}^+ \subseteq \mathcal{O}_F$. The ring class field L of K of $\mathcal{O}(\mathfrak{f}^+)$ is a Galois extension of F.*

Proof. The proof is analogous to the proof of Lemma 9.3 in [22]. Because K/F is an imaginary quadratic extension, its Galois group is generated by complex conjugation, which we denote by τ. Hence showing that L/F is Galois is equivalent to showing that $\tau(L) = L$. Theorem 4.3.5 states that there is an inclusion-reversing one-to-one correspondence between equivalence classes of congruence subgroups and abelian extensions of K. Thus, we simply need $\ker(\Phi_{\tau(L)/K,\mathfrak{f}}) = \ker(\Phi_{L/K,\mathfrak{f}})$ where $\mathfrak{f} = \mathfrak{f}^+\mathcal{O}_K$. Notice that $\tau(\mathfrak{f}) = \mathfrak{f}$ because $\mathfrak{f}^+$ is an ideal in a totally real field, and $\tau(\mathcal{O}_K) = \mathcal{O}_K$.

Theorem 4.3.6 tells us that

$$\ker(\Phi_{\tau(L)/K,\mathfrak{f}}) = P_{K,\mathcal{O}}(\mathfrak{f}).$$

It is easy to see that $\tau(P_{K,\mathcal{O}}(\mathfrak{f})) = P_{K,\mathcal{O}}(\mathfrak{f})$ because an ideal $\mathfrak{a}$ is coprime to $\mathfrak{f}$ if and only if $\tau(\mathfrak{a})$ is coprime to $\tau(\mathfrak{f}) = \mathfrak{f}$. We have $\ker(\Phi_{\tau(L)/K,\mathfrak{f}}) = \ker(\Phi_{\tau(L)/\tau(K),\mathfrak{f}}) = \tau(\ker(\Phi_{L/K,\mathfrak{f}}))$

by definition because $\tau(K) = K$. Therefore

$$\ker(\Phi_{\tau(L)/K,\mathfrak{f}}) = \tau(\ker(\Phi_{L/K,\mathfrak{f}}))) = \tau(P_{K,\mathcal{O}}(\mathfrak{f})) = P_{K,\mathcal{O}}(\mathfrak{f}) = \ker(\Phi_{L/K,\mathfrak{f}}),$$

proving that $\tau(L)$ and L corresponding to the same congruence subgroup, as desired. □

Now we can prove the existence of the needed relations, using the same idea as Bisson and Sutherland.

Proposition 4.3.10. *Assume that $\mathcal{O}_K^* = \mathcal{O}_F^*$. Let $\mathfrak{v} \subseteq \mathcal{O}_F$ be an ideal which is divisible by $\mathfrak{p}^k \subseteq \mathcal{O}_F$. If $N(\mathfrak{p}) \geq 3$, then there are infinitely many relations R satisfying the assumption of Corollary 4.3.3 above. Specifically, write $\mathfrak{f}_1^+ = \mathfrak{p}^{k-1-v_\mathfrak{p}(\mathfrak{v})}\mathfrak{v}$ and $\mathfrak{f}_2^+ = \mathfrak{p}^k$. Then there are infinitely many primes $\mathfrak{l} \subseteq \mathcal{O}_F$ which split in $\mathcal{O}_K$ such that R holds in $\mathcal{O}(\mathfrak{f}_1^+)$ but not $\mathcal{O}(\mathfrak{f}_2^+)$, where R is the one-element relation consisting of any prime lying over $\mathfrak{l}$.*

Proof. Let S_1 and S_2 be the set of primes of $\mathcal{O}_F$ which split into principal ideals in $\mathcal{O}(\mathfrak{f}_1^+)$ and $\mathcal{O}(\mathfrak{f}_2^+)$, respectively. Our goal is to show that $S_1 \setminus S_2$ is infinite. This proves the claim because every prime of $\mathcal{O}_F[\pi]$ lying over a prime in $S_1 \setminus S_2$ is a relation which holds in $\mathcal{O}(\mathfrak{f}_1^+)$ but not in $\mathcal{O}(\mathfrak{f}_2^+)$, as desired.

Write L_1 and L_2 for the ring class fields of $\mathcal{O}(\mathfrak{f}_1^+)$ and $\mathcal{O}(\mathfrak{f}_2^+)$, respectively. An ideal of $\mathcal{O}(\mathfrak{f}_i^+)$ is principal if and only if it splits completely in L_i by Corollary 4.3.7 because L_i is a ring class field. Hence the sets S_i are the sets of primes in $\mathcal{O}_F$ which split completely in L_i, respectively. The Chebotarev Density Theorem implies that $L_2 \subseteq L_1$ if and only if $S_1 \setminus S_2$ is finite [72, Proposition VII.13.9]. But $\mathfrak{p}^k$ divides the conductor of $\mathcal{O}(\mathfrak{f}_2)$ and does not divide the conductor of $\mathcal{O}(\mathfrak{f}_1^+)$ by construction. Because $N(\mathfrak{p}) \geq 3$, the lemma below implies that $\mathfrak{p}^k$ divides the class field theory conductor $\mathfrak{f}(L_2/K)$ but does not divide $\mathfrak{f}(L_1/K)$. By Lemma 4.3.8 we have $L_2 \not\subseteq L_1$, which completes the proof. □

The following lemma and its proof are a generalization of [22, Exercises 9.20-22].

Lemma 4.3.11. *Let K be a CM field with totally real subfield F, and let $\mathcal{O} \subseteq K$ be an order containing $\mathcal{O}_F$. Suppose that $\mathcal{O}_K^* = \mathcal{O}_F^*$. If L is the ring class field of $\mathcal{O}$, then the conductor ideal $\mathfrak{f}$ of $\mathcal{O}$ and the conductor ideal $\mathfrak{f}(L/K)$ of the ring class field L may only differ by primes of K of norm 2.*

Proof. By definition, the class field theory conductor $\mathfrak{f}(L/K)$ is a divisor of the conductor ideal $\mathfrak{f}$ because $\mathfrak{f}$ is the modulus used to define the extension L/K. Thus, if $\mathfrak{f}(L/K) \neq \mathfrak{f}$,

then we may write $\mathfrak{f} = \mathcal{P}\mathfrak{m}$ where $\mathcal{P}$ is a prime ideal of K and $\mathfrak{f}(L/K)$ divides $\mathfrak{m}$. Assume that this is the case for some prime $\mathcal{P}$ with $N(\mathcal{P}) \neq 2$. We will find a contradiction.

According to Theorem 4.3.5, there is an equivalence class $\mathbf{H}(L/K)$ of congruence subgroups defined with various moduli corresponding to the extension L/K. Because $\mathfrak{f}(L/K)$ divides $\mathfrak{m}$, there is a congruence subgroup $\ker \Phi_{L/K,\mathfrak{m}}$ defined modulo $\mathfrak{m}$. By construction, $P_{K,\mathcal{O}}(\mathfrak{f})$ is the restriction of $\ker \Phi_{L/K,\mathfrak{m}}$ to $I_K(\mathfrak{f})$, i.e. $\ker(\Phi_{L/K,\mathfrak{m}}) \cap I_K(\mathfrak{f}) = P_{K,\mathcal{O}}(\mathfrak{f})$. This proves that

$$P_{K,1}(\mathfrak{m}) \cap I_K(\mathfrak{f}) \subseteq P_{K,\mathcal{O}}(\mathfrak{f}). \tag{4.3}$$

because $\ker(\Phi_{L/K,\mathfrak{m}}) \supseteq P_{K,1}(\mathfrak{m})$.

By inspecting definitions, the sequence below is exact.

$$\mathcal{O}_K^* \to (\mathcal{O}_K/\mathfrak{f}\mathcal{O}_K)^* \xrightarrow{\phi} P_K \cap I_K(\mathfrak{f})/P_{K,1}(\mathfrak{f}) \to 1.$$

Define $\pi : (\mathcal{O}_K/\mathfrak{f})^* \to (\mathcal{O}_K/\mathfrak{m})^*$ and $\beta : (\mathcal{O}_F/\mathfrak{f}^+)^* \to (\mathcal{O}_K/\mathfrak{f}\mathcal{O}_K)^*$ as the natural maps induced by quotients. Notice that π is surjective and β is injective. One may show that

$$\mathcal{O}_K^* \cdot \ker \pi = \phi^{-1}(I_K(\mathfrak{f}) \cap P_{K,1}(\mathfrak{m}))$$

and similarly

$$\mathcal{O}_K^* \cdot \mathrm{Im}\beta = \phi^{-1}(P_{K,\mathcal{O}}(\mathfrak{f})),$$

which proves that $\ker \pi \subseteq (\mathcal{O}_K)^* \cdot \mathrm{Im}\beta$ by containment (4.3) above. Since $\mathcal{O}_K^* = \mathcal{O}_F^*$ and $\mathrm{Im}\beta$ is closed under the action of $\mathcal{O}_F^*$, this shows

$$\ker \pi \subseteq \mathrm{Im}\beta.$$

Recall that if $\mathfrak{a} \subseteq \mathcal{O}_K$ is an ideal, then

$$|(\mathcal{O}_K/\mathfrak{a})^*| = N(\mathfrak{a}) \prod_{\substack{\mathfrak{q}|\mathfrak{a} \\ \text{prime}}} \left(1 - \frac{1}{N(\mathfrak{q})}\right)$$

This implies that

$$|\ker(\pi)| = \frac{|(\mathcal{O}_K/\mathfrak{f})^*|}{|(\mathcal{O}_K/\mathfrak{m})^*|} = \begin{cases} N(\mathcal{P}) & \text{if } \mathcal{P} \mid \mathfrak{m}; \\ N(\mathcal{P}) - 1 & \text{if } \mathcal{P} \nmid \mathfrak{m}. \end{cases}$$

We will now conclude the proof by showing that $|\ker \pi| = 1$, which is only possible if

$N(\mathcal{P}) = 2$ and $\mathcal{P} \nmid \mathfrak{m}$, according to the formula for $|\ker \pi|$. To prove this, we consider two cases. Writing $\mathfrak{p} = \mathcal{P} \cap \mathcal{O}_F$, either $\mathfrak{p}\mathcal{O}_K = \mathcal{P}\overline{\mathcal{P}}$ or $\mathfrak{p}\mathcal{O}_K = \mathcal{P}$. In the former case, one observes that $\pi \circ \beta$ is injective, hence $|\ker \pi \cap \mathrm{Im}\beta| = |\ker \pi| = 1$. Clearly this is only possible if $N(\mathcal{P}) = 2$ and $\mathcal{P} \nmid \mathfrak{m}$.

Now consider the latter case, and suppose $\mathfrak{p}\mathcal{O}_K = \mathcal{P}$. Because $\mathfrak{p}$ is an inert prime of $\mathcal{O}_F$ and $\mathcal{P}$ divides $\mathfrak{f}$, it follows that $\mathfrak{p}$ divides $\mathfrak{f}^+$. Write $\mathfrak{f}^+ = \mathfrak{p}\mathfrak{m}_0$ for some $\mathfrak{m}_0 \subseteq \mathcal{O}_F$. Consider the (not necessarily exact) diagram

$$\begin{array}{ccccc} (\mathcal{O}_F/\mathfrak{f}^+)^* & \xrightarrow{\beta} & (\mathcal{O}_K/\mathfrak{f}\mathcal{O}_K)^* & \xrightarrow{\pi} & (\mathcal{O}_K/\mathfrak{m})^* \\ \downarrow{\scriptstyle \theta} & & & & \\ (\mathcal{O}_F/\mathfrak{m}_0)^* & & & & \end{array}$$

One shows $\ker \pi = \ker \pi \cap \mathrm{Im}\beta \cong \ker \theta$. However, we can compute $\ker(\theta)$ in the same way that we computed $\ker(\pi)$ to find that

$$|\ker(\theta)| = \frac{|(\mathcal{O}_F/\mathfrak{f}^+)^*|}{|(\mathcal{O}_F/\mathfrak{m}_0)^*|} = \begin{cases} N(\mathfrak{p}) & \text{if } \mathfrak{p} \mid \mathfrak{m}_0; \\ N(\mathfrak{p}) - 1 & \text{if } \mathfrak{p} \nmid \mathfrak{m}_0. \end{cases}$$

Because P is inert, $N(\mathcal{P}) = N(\mathfrak{p}\mathcal{O}_K) > N(\mathfrak{p})$. Hence $|\ker \pi| = |\ker \theta|$ is impossible, which gives the final contradiction. □

4.4 Algorithms

Now that we have proven that there are sufficiently many class group relations for determining the large prime factors of $\mathfrak{f}^+(A)$, we are able to present generalizations of the elliptic curve algorithms in [8]. Even though the results of the previous section apply whenever A is an ordinary simple abelian variety of arbitrary dimension with maximal RM, there are no known results for computing an arbitrary isogeny in such generality. We restrict our attention to a manageable case with the following requirements so that we can compute all isogenies arising from the action of the ideal class group.

Requirements (R)

For the remainder of the chapter, we will focus on the case where the ordinary, absolutely simple, principally polarized abelian variety A has dimension $g = 2$. We also assume the following are all true, which summarizes the hypotheses found in Theorems 4.2.1, 4.2.9

and 4.2.11, and Lemma 4.3.11. In particular, $\operatorname{End} A$ is uniquely identified by an ideal $\mathfrak{f}^+(A) \subseteq \mathcal{O}_F$ and we can compute the isogenies corresponding to the action of $\mathrm{Cl}(\operatorname{End} A)$.

1. A has maximal real multiplication by F.
2. $\mathcal{O}_K^* = \mathcal{O}_F^*$.
3. F has narrow class number 1.
4. The conductor gap $[\mathcal{O}_F : \mathbb{Z}[\pi + \overline{\pi}]]$ is not divisible by 2.

Notice that that the second assumption is very mild. It is shown in [97, Lemma II.3.3] that $\mathcal{O}_K^* = \mathcal{O}_F^*$ is true for every primitive quartic CM field except the cyclotomic field $K = \mathbb{Q}(\zeta_5)$.

Heuristic assumptions (H)

We collect the following heuristic assumptions that we require, and we denote all statements requiring these assumptions will be decorated by the letter (**H**). The first three assumptions allow us to compute a class group via [4], while the final two are moderate heuristic assumptions, similar to those in [8], which are only needed to bound the expected run time of the algorithms. In particular, the latter three assumptions are not needed in order to prove the algorithm is correct. All running times will be analyzed with the subexponential function

$$L[a, c](n) = \exp\left((c + o(1)) \log(n)^a (\log\log n)^{1-a}\right)$$

1. **GRH** is true.
2. **Smoothness assumption.** The probability $P(\iota, \mu)$ that an ideal of $\mathcal{O}$ of norm bounded by e^ι is a power-product of prime ideals of norm bounded by e^μ satisfies $P(\iota, \mu) \geq \exp(-u \log u(1 + o(1)))$ for $u = \log(\iota)/\log(\mu)$.
3. **Spanning relations** If $\{\mathfrak{p}_1, \ldots, \mathfrak{p}_N\}$ is a factor base of ideals generating $\mathrm{Cl}(\mathcal{O})$ where $N = L[a, c_1](\Delta)$ for some $0 < a < 1$ and $c_1 > 0$, then it suffices to collect N' relations to generate all possible relations if $N'/N = L[b, c_2](\Delta)$ for $0 < b < a$ and $c_2 > 0$.

4. **Random Norms:** Assume that the norms of the ideals generated in the reduction step of FindRelation have approximately the distribution of random integers in $[1, n]$. This allows us to analyze the probability that a norm is B-smooth.

5. **Random relations.** If $\mathcal{O}_1$ and $\mathcal{O}_2$ are as in Corollary 4.3.3 with sufficiently large discriminants, then the relation for $\mathcal{O}_1$ generated in Find Relation algorithm does not hold in $\mathcal{O}_2$ with probability bounded above 0.

6. **Integer factorization and smoothness checking.** The number field sieve with expected running time $L[1/3, c_f](n)$ by [15] and ECM finds a prime factor p of integer n in expected time $L[1/2, 2](p) \log^2 n$ by [60].

4.4.1 Finding relations

Following [5, 6, 8], we will use the reduction of random ideals to produce relations which hold in a given order $\mathcal{O} \subseteq K$ of discriminant $\Delta = \text{disc}(\mathcal{O})$. This is the same idea that is used in ideal class group computation algorithms, such as [4]. For background on ideal reduction, see [102, Chapter 5]. Given an ideal $\mathfrak{a} \subseteq \mathcal{O} \subseteq K$, reduction outputs an ideal $\mathfrak{b}$ which is equivalent to $\mathfrak{a}$ in $\text{Cl}(\mathcal{O})$ such that $N\mathfrak{b} \leq \Delta^2$. By taking $N\mathfrak{a} > \Delta^2$, we ensure that $\mathfrak{a}\mathfrak{b}^{-1}$ is a nontrivial relation which holds in $\mathcal{O}$. The reduction algorithm is polynomial in $\log |\Delta|$ for fields of fixed degree.

To test whether this given relation also holds in a second order, we need to solve the principal ideal problem. As shown in [3, §4.3] and [4], we can solve the principal ideal problem in subexponential time by using the heuristic assumptions **(H)**. In practice, one should compute the class group once and for all at the beginning of the algorithm, then simply do reductions as necessary when checking relations. To give some flexibility in applications, we use a parameter $\mu > 0$ which can be chosen arbitrarily.

Notice that a prime ideal $\mathfrak{L}$ is inverse to its complex conjugate $\overline{\mathfrak{L}}$ because the totally real subfield F has class number 1. Hence we can ensure that at most one of $\mathfrak{L}$ and $\overline{\mathfrak{L}}$ appears in the relation R. In particular, if $x_{\mathfrak{L}} - y_{\mathfrak{L}} < 0$ for any prime ideal $\mathfrak{L}$, then replace the prime power $\mathfrak{L}^{x_{\mathfrak{L}} - y_{\mathfrak{L}}}$ in the relation R with $\overline{\mathfrak{L}}^{y_{\mathfrak{L}} - x_{\mathfrak{L}}}$, which is an equivalent ideal in the class group. Implicitly, we throw out every relation which includes an undesirable prime, i.e. a prime ideal of $\mathcal{O}_F$ that divides the ideal $\mathfrak{f}^+(\mathcal{O}_F[\pi])$, or the index $[\mathcal{O}_F : \mathbb{Z}[\pi + \overline{\pi}]]]$, or the index $[\mathcal{O}_K : \mathbb{Z}[\pi]]$. There are only $O(\log q)$ such primes, so this does not change the complexity of the algorithm.

Algorithm 2: FindRelation

Input : Orders $\mathcal{O}_1$ and $\mathcal{O}_2$
Output : Relation R which holds in $\mathcal{O}_1$ but not $\mathcal{O}_2$

1 Set $B = L[1/2, \mu](n)$, $D_1 = \text{disc}(\mathcal{O}_1)$ and $n = |D_1|^2$.
2 Pick $x_{\mathfrak{L}}$ such that $x_{\mathfrak{L}} \leq B/N(\mathfrak{L})$ for prime ideals $\mathfrak{L}$ of norm bounded by B such that at most k_0 are nonzero and $\prod N(\mathfrak{L}) > n$.
3 Compute the reduced ideal $\mathfrak{b} = \prod_{\mathfrak{L}} \mathfrak{L}^{y_{\mathfrak{L}}}$ of $\mathfrak{a} = \prod \mathfrak{L}^{x_{\mathfrak{L}}}$.
4 **if** *$N(\mathfrak{b})$ is a B-smooth integer and the number of nonzero $y_{\mathfrak{L}}$ is at most $8 \log(|D_1|)^{1/2}$* **then**
5 Let R be the relation $(\mathfrak{L}^{x_{\mathfrak{L}} - y_{\mathfrak{L}}})_{N\mathfrak{L} < B}$.
6 **if** *R does not hold in $\mathcal{O}_2$* **then**
7 **return** R.
8 **end**
9 **end**
10 Go to Step 2.

Proposition 4.4.1 (H). *Given orders $\mathcal{O}_1$ and $\mathcal{O}_2$ of discriminants D_1 and D_2, the algorithm FindRelation has expected running time*

$$L[1/2, 1/\mu\sqrt{2}](|D_1|) + \log|D_1|^{1+\epsilon} L[1/2, c](|D_2|)$$

where $\mu > 0$ is a parameter that can be chosen arbitrarily, ϵ is arbitrarily small, and c is the constant from the running time bound of principal ideal testing. The output relation $R = (\mathfrak{L}_1^{e_1}, \ldots, \mathfrak{L}_k^{e_k})$ has exponents e_i bounded by $L\left[1/2, \mu\sqrt{2}\right](|D_1|)$ and k bounded by $\frac{16\sqrt{2}}{\mu}(\log|D_1|)^{1/2}$. The prime numbers ℓ_i lying under the $\mathfrak{L}_i$ are bounded by $L\left[1/2, \mu\sqrt{2}\right](|D_1|)$.

Proof. Recall that $B = L[1/2, \mu](n) = L[1/2, \mu\sqrt{2}](|D_1|)$ is the smoothness bound on the norms of the primes where $n = |D_1|^2$. Notice $\log n / \log B = \frac{1}{\mu}(\log n)^{1/2}(\log\log n)^{-1/2}$. With the bound $1/\rho(z)^{-1} = z^{z+o(1)}$ for the Dickman function $\rho(z)$ and the same argument as [8, Proposition 6], we expect the number of attempts required to find an ideal $\mathfrak{b}$ with B-smooth norm to be be asymptotically bounded by $\rho(\log n/\log B)^{-1} = L[1/2, 1/\mu\sqrt{2}](|D_1|)$, and a B-smooth integer in $[1, n]$ is expected to have $(2+o(1))\log n/\log B$ distinct prime factors. Thus we expect the number of prime ideals $\mathfrak{L}$ appearing in step 4 to be at most

$$k_0 + 8\log n/\log B \leq k_0 + \frac{8\sqrt{2}}{\mu}(\log|D_1|)^{1/2}(\log\log|D_1|)^{-1/2}$$

By heuristic assumption **(H)**, elliptic curve factorization [60] identifies a B-smooth integer in time $L[1/2, 2](B) = L[1/4, \sqrt{2\mu}](n)$ with high probability. Therefore, $L[1/2, 1/\mu\sqrt{2}](|D_1|)$ is a bound on the amount of time spent finding a relation which holds in $\mathcal{O}_1$ with B-smooth norm.

Step 6 is solving the Principal Ideal Problem in $\mathcal{O}_2$, which can be done by following Biasse's algorithm [3, Algorithm 7]. First, one computes the class group in time $L[1/2, c](|D_2|)$ by [4, Theorem 6.1]. This involves finding a set of relations which span all relations on a set of primes $\{\mathcal{P}_1, \ldots, \mathcal{P}_M\}$ which generate $\mathrm{Cl}(\mathcal{O}_2)$. Then, each ideal $\mathfrak{L}$ appearing in the relation R is reduced, if necessary, to an equivalent product over the generating set, i.e. $\mathfrak{L} = (\alpha)\mathcal{P}_1^{e_1} \ldots \mathcal{P}_M^{e_M}$ for some $\alpha \in K$ and $e_i \geq 0$. Each reduction takes time $\log(N(\mathfrak{L}))^{1+o(1)} L[1/2, c_1](|D_2|)$ for some constant $c_1 > 0$ by [3, Proposition 3.1]. Recall that the ideals are bounded by

$$\log N(\mathfrak{L}) < \log B \leq \log L[1/2, \mu](|D_1|^2) \ll (\log |D_1|)^{1/2} (\log \log |D_1|)^{1/2}$$

By combining this bound with the bound on the number of primes appearing in R, the cost of solving the principal ideal problem is bounded, for some constant $c > 0$, by $\log |D_1|^{1+\epsilon} L[1/2, c](|D_2|)$. We succeed in finding a relation which does not hold in $\mathcal{O}_2$ within $O(1)$ tries because **(H)** assumes that the probability of success is bounded above 0. Thus, we obtain the final bound on the running time. The bounds on k, ℓ_i and e_i follow directly from the construction of the relation in the algorithm. □

4.4.2 Computing from above

Now we present our algorithm for computing $\mathfrak{f}^+(A)$. As noted in Section 4.2.6, we can take care of all "small" prime factors, i.e. primes $\mathfrak{p}$ with $N(\mathfrak{p}) \leq C$ for some $C \geq 2$, by "isogeny climbing". The bound C can be chosen arbitrarily In the presentation of this and all remaining algorithms, we assume that $\mathfrak{f}^+(A)$ is not divisible by the square of any large primes, although the algorithms can be easily modified to handle this possibility. For example, we can modify Algorithm 3 by finding relations in Step 10 corresponding to $\mathfrak{p}^k$ for every $k \geq 1$ such that $\mathfrak{p}^k \mid \mathfrak{v}$. Instead, we simplify the presentation by only checking $k = 1$. The correctness of this algorithm immediately follows immediately from Corollary 4.3.3. There are infinitely many class group relations R as required in Step 10 according to Proposition 4.3.10.

Algorithm 3: Computing $\mathfrak{f}^+(A)$ from above

Input : Abelian variety A satisfying requirements **(R)**
Output : The ideal $\mathfrak{f}^+ \subseteq \mathcal{O}_F$ identifying $\operatorname{End} A$

```
Determine the ideal 𝔳 ⊆ O_F defining the order O_F[π] ⊆ O_K.
Fix bound C ≥ 3 and initiate 𝔲 = 1.
for every 𝔭 ⊆ O_F with N𝔭 < C which divides 𝔳 do
    Use isogeny climbing to determine v_𝔭 = v_𝔭(f⁺(A)).
    Replace 𝔲 with the product 𝔲𝔭^{v_𝔭}.
    Update A to be the abelian variety found by isogeny climbing
    Remove powers of 𝔭 from 𝔳.
end
for every 𝔭 ⊆ O_F with N𝔭 ≥ C which divides 𝔳 do
    Find relation R which holds in O(𝔳/𝔭) but not O(𝔭).
    if R does not hold in O(f⁺(A)) then
        Update 𝔲 to be the product 𝔲𝔭
    end
end
return f⁺(A) = 𝔲.
```

Proposition 4.4.2 (**H**). *Algorithm 3 has expected running time*

$$L[1/2, 2c](q) + L[1/2, 2\sqrt{d+1}](q).$$

where c is the constant from principal ideal testing, d is the constant from Theorem 4.2.9.

Proof. Computing $\mathfrak{v}$ takes polynomial time by Corollary 4.2.4, and factoring $\mathfrak{v}$ reduces in polynomial time to factoring its norm $N(\mathfrak{v})$. This is done in time $L[1/3, c_f](q)$ by **(H)**.

The number of iterations in the algorithm is the number of primes dividing $\mathfrak{v}$, so there are only $O(\log q)$ iterations needed. Recall that $\log|\operatorname{disc}(\mathbb{Z}[\pi, \overline{\pi}])| < 4\log q + 12\log 2$ by [6, Lemma 6.1]. Since we only consider orders containing $\mathcal{O}_F[\pi] \supseteq \mathbb{Z}[\pi, \overline{\pi}]$, this means that FindRelation takes time

$$L[1/2, 1/\mu\sqrt{2}](|D_1|) + \log|D_1|^{1+\epsilon} L[1/2, c](|D_2|) = L[1/2, \sqrt{2}/\mu](q) + L[1/2, 2c](q)$$

for each choice of $\mathcal{O}_1$ and $\mathcal{O}_2$, by the proposition above.

Computing whether a relation $R = (\mathfrak{L}_1^{e_1}, \ldots, \mathfrak{L}_k^{e_k})$ holds for A requires computing $\sum_{i=1}^k e_i = L[1/2, 2\mu\sqrt{2}](q)$ many isogenies which each take time $L[1/2, 2\mu d\sqrt{2}](q)$ to

compute by Theorem 4.2.9. Thus the expected running time is

$$L[1/2, \sqrt{2}/\mu](q) + L[1/2, 2c](q) + L[1/2, 2\mu(d+1)\sqrt{2}](q).$$

Solving $\sqrt{2}/\mu = 2\mu(d+1)\sqrt{2}$ presents $\mu = 1/\sqrt{2(d+1)}$ as an optimal choice for μ. Inserting this into the bound above, we obtain the desired bound:

$$L[1/2, 2c](q) + L[1/2, 2\sqrt{d+1}](q).$$

□

4.4.3 Certifying and verifying

Algorithm 4: CERTIFY

Input : CM field K with maximal totally real subfield F, ideals $\mathfrak{u}, \mathfrak{v} \subseteq \mathcal{O}_F$
Output : Certificate C

for *For every prime* $\mathfrak{p}$ *with* $v_\mathfrak{p}(\mathfrak{v}) - v_\mathfrak{p}(\mathfrak{u}) > 0$ **do**
  Find a Relation $R_\mathfrak{p}$ which holds in $\mathcal{O}(\mathfrak{u})$ but not $\mathcal{O}(\mathfrak{p})$
end
for *For every prime* $\mathfrak{p}$ *of* $\mathfrak{u}$ **do**
  Find a Relation $R_\mathfrak{p}$ which holds in $\mathcal{O}(\mathfrak{u}\mathfrak{p}^{-1})$ but not $\mathcal{O}(\mathfrak{p})$.
end
return the certificate $C = (\mathfrak{u}, \mathfrak{v}, K, \{R_\mathfrak{p}\}_{\mathfrak{p}|\mathfrak{v}})$.

The relations generated throughout Algorithm 3 work for any abelian variety in the given isogeny class which satisfies the requirements **(R)**. Collecting these relations without fixing a specific abelian variety A gives the CERTIFY and VERIFY algorithms below. Indeed, CERTIFY does not require an abelian variety as input; it simply gives a certificate that allows anyone to check the claim that $\mathfrak{f}^+(A) = \mathfrak{u}$ via the subsequent VERIFY algorithm, up to small prime factors. In the presentation below, we ignore the small prime factors, and assume they are taken care of by isogeny climbing, as before. The two loops in CERTIFY and VERIFY correspond to checking that $\mathfrak{u}$ divides $\mathfrak{f}^+(A)$ and $\mathfrak{u}$ is not a proper divisor of $\mathfrak{f}^+(A)$, respectively. Again, the correctness follows immediately from Corollary 4.3.3.

Corollary 4.4.3 (H). *The expected running time of* CERTIFY *is within a factor of*

$O(\log q)$ *of the expected running time of* FindRelation. *If* D_1 *is discriminant of the order corresponding to the ideal* $\mathfrak{u}$*, then the certificate has size* $O(\log |D_1| \log\log |D_1|)$.

Proof. Using the bound $|\operatorname{disc}(\mathcal{O}_F[\pi])| \leq |\operatorname{disc}(\mathbb{Z}[\pi, \overline{\pi}])| \leq 4^6 q^4$ from [6, Lemma 6.1], we find that the number of prime factors of $\mathfrak{v}$ is $O(\log q)$. Thus, FindRelation is called $O(\log q)$ many times in Certify.

By the bounds on the ℓ_i, k and e_i, we see that the size of the certificate is $O(\log |D_1| \log\log |D_1|)$. Here, we use the fact that $N(\mathfrak{L}_i) \leq \ell^4$ and the size of an ideal $\mathfrak{a}$ in a number field of fixed degree is $O(\log N(\mathfrak{a}))$ [102, Corollary 3.4.13]. □

Algorithm 5: Verify

Input : Abelian variety A satisfying requirements **(R)**, certificate $C = (\mathfrak{u}, \mathfrak{v}, K, \{R_\mathfrak{p}\}_{\mathfrak{p}|\mathfrak{v}})$

Output : `true` or `false`

1 **for** *For every prime* $\mathfrak{p}$ *with* $v_\mathfrak{p}(\mathfrak{v}) - v_\mathfrak{p}(\mathfrak{u}) > 0$ **do**
2 | Check that $R_\mathfrak{p}$ holds in $\operatorname{End} A$ by computing isogenies
3 | Immediately **return** `false` if failure occurs.
4 **end**
5 **for** *For every prime* $\mathfrak{p}$ *dividing* $\mathfrak{u}$ **do**
6 | Check that $R_\mathfrak{p}$ does not hold in $\operatorname{End} A$ by computing isogenies.
7 | Immediately **return** `false` if a relation holds.
8 **end**
9 **return** `true`

Proposition 4.4.4 (H). *Given a certificate* $C = (\mathfrak{u}, \mathfrak{v}, K, \{R_\mathfrak{p}\}_{\mathfrak{p}|\mathfrak{v}})$ *produced by* Certify *with parameter* $\mu > 0$ *and* $A/\mathbb{F}_q$*, the expected run time of* Verify *is*

$$L[1/2, \mu(d+1)\sqrt{2}](|D_1|)$$

where d *is the constant from Theorem 4.2.9 and* D_1 *is the discriminant of the order identified by* $\mathfrak{u}$.

Proof. There are at most $O(\log q)$ relations in a certificate. By Proposition 4.4.1, each relation has at most $O(\log^{1/2} |D_1|)$ distinct primes with exponents e_i bounded by $L[1/2, \mu\sqrt{2}](|D_1|)$. The primes ℓ_i are bounded by $L[1/2, \mu\sqrt{2}](|D_1|)$, which produces the claim by Theorem 4.2.9. □

Remark:

In [8, Algorithm 2], Bisson and Sutherland define an additional algorithm to compute the endomorphism ring of an ordinary elliptic curve by repeatedly calling their versions of the CERTIFY and VERIFY algorithms. This algorithm performs well for large orders because their version of FINDRELATION uses binary quadratic forms to quickly check class group relations in various orders. In our case, FINDRELATION is much more costly because no analogue of binary quadratic forms exists for performing computations in the class group of orders in general CM fields. As a result, the immediate generalization of [8, Algorithm 2] which uses the CERTIFY and VERIFY algorithms above correctly computes $\mathfrak{f}^+(\operatorname{End} A)$, but does not gain an performance improvement by focusing on abelian varieties with nearly-maximal endomorphism ring.

4.5 Computational Example

We now give an illustrative example to demonstrate how Algorithm 3 computes an endomorphism ring by using class group relations. All computations where performed using Magma [11] and the AVIsogenies library [7] on a 2.3 GHz Intel Core i5 processor. Although the programs were not optimized for maximum performance, running times are given below to display how isogeny computation is the major bottleneck of the algorithm.

4.5.1 Example

Let $q = 82307$ and let A be the Jacobian of the hyperelliptic curve C defined over $\mathbb{F}_q$ by

$$y^2 = x^5 - 3x^4 + 5x^3 - x^2 - 2x + 1.$$

The curve C is the specialization to $s = t = 1$ of the 2-parameter family given in [39, §4.3]. By [39, Proposition 3], A has maximal real multiplication by $F = \mathbb{Q}(\sqrt{5})$, which has narrow class number 1. The characteristic polynomial of the Frobenius endomorphism π is

$$f_\pi(t) = t^4 + 658t^3 + 263610t^2 + 658qt + q^2.$$

By [44, Theorem 6], we deduce that A is absolutely simple.

We begin by computing the ideal $\mathfrak{v} \subseteq \mathcal{O}_F$ which identifies $\mathcal{O}_F[\pi]$. Using Algorithm 1, it takes 0.02 seconds to compute that $\mathfrak{v} = \mathfrak{p}_{11}\mathfrak{p}_{131}$ where $\mathfrak{p}_{11}$ and $\mathfrak{p}_{131}$ are prime ideals of $\mathcal{O}_F$ of norm 11 and 131, respectively. Therefore, $\operatorname{End} A$ is one of the following four

orders, using the notation of Section 4.2.2:

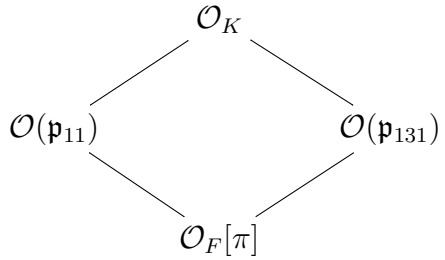

Next, we need suitable class group relations, as described in Corollary 4.3.3. Consider the prime number 7, which is inert in $\mathcal{O}_F$ and splits in $\mathcal{O}_K$. We find

$$f_\pi(t) \equiv (t^2 + t + 6)(t^2 + 6t + 6) \mod 7,$$

hence the prime ideals of $\mathcal{O}_F[\pi]$ lying over 7 are $\mathfrak{L}_1 = (7, \pi^2{+}\pi{+}6)$ and $\mathfrak{L}_1' = (7, \pi^2{+}6\pi{+}6)$. It takes 0.22 and 0.16 seconds, respectively, to find that the ideal $\mathfrak{L}_1\mathcal{O}(\mathfrak{p}_{11})$ has order 55 in $\mathrm{Cl}(\mathcal{O}(\mathfrak{p}_{11}))$, and the ideal $\mathfrak{L}_1\mathcal{O}(\mathfrak{p}_{131})$ has order 60 in $\mathrm{Cl}(\mathcal{O}(\mathfrak{p}_{131}))$. Therefore the relation $R_1 = (\mathfrak{L}_1^{55})$ holds in $\mathcal{O}(\mathfrak{p}_{11})$ but not $\mathcal{O}(\mathfrak{p}_{131})$. Thus $\mathfrak{p}_{11}$ divides the identifying ideal $\mathfrak{f}^+(\mathrm{End}\, A)$ if and only if the relation R_1 does not hold for A. Similarly, $R_2 = (\mathfrak{L}_1^{60})$ is a relation which holds in $\mathcal{O}(\mathfrak{p}_{131})$ but not $\mathcal{O}(\mathfrak{p}_{11})$, hence $\mathfrak{p}_{131}$ divides $\mathfrak{f}^+(\mathrm{End}\, A)$ if and only if R_2 does not hold for A.

Finally, we need to compute the chains of isogenies corresponding to the two relations R_1 and R_2. The kernel of the isogeny corresponding to the action of the ideal $\mathfrak{L}_1 = (7, \pi^2 + \pi + 6)$ is the rational symplectic subgroup $G \subseteq A[7]$ whose generators are annihilated by $\pi^2 + \pi + 6$. Using AVIsogenies, we enumerate all possible rational symplectic subgroups of $A[7]$, find the desired G in this list, and compute the corresponding $(7, 7)$-isogeny. Continuing in this way, it takes 154 seconds to compute all 60 isogenies, and we find that neither R_1 nor R_2 holds for A. By Corollary 4.3.3, this implies that $\mathfrak{f}^+(\mathrm{End}\, A) = \mathfrak{v} = \mathfrak{p}_{11}\mathfrak{p}_{131}$, i.e. $\mathrm{End}\, A \cong \mathcal{O}_F[\pi]$.

This example demonstrates how our algorithm is more efficient than the algorithm of Eisenträger and Lauter [29] in cases when the index $[\mathcal{O}_K : \mathcal{O}_F[\pi]]$ is divisible by large primes. Indeed, using Eisenträger and Lauter's algorithm on the example above requires the full 131-torsion of A, which is defined over an extension of degree 17030. As a result, this method is very costly compared to the computation of the $(7, 7)$-isogenies performed above. The same benefit is seen in examples of Bisson's algorithm [6, §8],

since it also exploits class group relations. Our algorithm differs from Bisson's algorithm by considerably restricting the class of abelian varieties considered, which allows us to avoid certain unconventional heuristic assumptions.

Part II

Decidability and Definability

Chapter 5
Background for Decidability

This chapter serves as an introduction for the remaining chapters by recalling some of the basic background for problems of decidability in number theory. We refer to [54, 75, 89] as helpful introductions to the topics at hand.

5.1 Formulas in the language of rings

Recall that Hilbert's Tenth Problem, in its original form, asks for an algorithm which takes as input a multivariable polynomial equation $f(x_1, \ldots, x_n) = 0$ with integer coefficients and outputs YES or NO depending on whether or not the equation has an integer solution $(a_1, \ldots, a_n) \in \mathbb{Z}^n$. More generally, we can ask the same question with $\mathbb{Z}$ replaced by some other ring R, known as *Hilbert's Tenth Problem over* R. When pursuing questions in this vein, it is useful to rephrase Hilbert's Tenth Problem in terms of formulas in the language of rings. The goal of this section is to introduce this terminology for future use.

Definition 5.1.1. *For an integral domain R, a* first-order *formula (in the language of rings) is a formula built using the ring operations $+$ and $\cdot$, the identity elements* 0 *and* 1*, along with the symbols $\neg$ (NOT), $\forall$ (FOR ALL), $\exists$ (THERE EXISTS), $\vee$ (OR), and $\wedge$ (AND). A formula containing additional coefficients from R has* parameters *from R. Moreover:*

(a) A formula without the symbol $\forall$ is an existential *formula;*

(b) A formula without the symbol $\exists$ is a universal *formula;*

(c) A formula without the symbol $\neg$ is a positive *formula.*

Example 5.1.2. *The formula $\phi_1(X)$ defined by*

$$\exists Y_1 \exists Y_2 \exists Y_3 \exists Y_4 \; X = Y_1^2 + Y_2^2 + Y_3^2 + Y_4^2$$

is a positive existential formula. By Lagrange's four squares theorem, this formula is satisfiable over the ring $\mathbb{Z}$ by precisely the non-negative integers.

In the end, we are concerned with which formulas are or are not satisfiable over a given ring R.

Definition 5.1.3. *We say the first-order theory of R is* decidable *if there is an algorithm which takes as input a first-order formula and outputs* YES *or* NO*, according to whether or not the formula is satisfiable over R. The decidability of other theories (e.g., the positive existential theory of R) is defined analogously.*

5.2 Equivalent problems

There is a well-established connection between Hilbert's Tenth Problem over an integral domain R and the theories defined above. Intuitively, the lemma below is summarized as "one equals finitely many", i.e. the problem of deciding whether a single polynomial equation is solvable is equivalent to the same problem for a system of equations.

Lemma 5.2.1 (§1.1, [89])**.** *Let R be an integral domain whose field of fractions is not algebraically closed. Any formula using $\vee$ and $\wedge$ can be rewritten as a formula which does not use these symbols.*

In particular, if R is a recursive integral domain whose field of fractions is not algebraically closed, then the decidability of Hilbert's Tenth Problem is equivalent to the decidability of the positive existential theory of R. Moreover, in the context which is relevant for the following chapter, we find that the restriction to *positive* existential formulas is unnecessary.

Theorem 5.2.2 (Theorem 4.2, [88])**.** *Let $L \subseteq \overline{\mathbb{Q}}$ be a field. Any formula over $\mathcal{O}_L$ (resp. L) using the symbol $\neg$ can be rewritten as an equivalent positive formula.*

5.3 Definability

One of the key techniques for showing undecidability is known as definability. In particular, we are concerned with the following.

Definition 5.3.1. *Let $R \subseteq S$ be rings. We say that S is existentially, universally, or first-order* definable *in S if there is a formula $\phi(x)$ of the appropriate form such that for all $x \in S$, $\phi(x)$ holds if and only if $x \in R$.*

With this definition the following is immediate:

Lemma 5.3.2. *Let $R \subseteq T$ be rings, and suppose that R is first-order definable in S. If the first-order theory of R is undecidable, then the same is true for S. The same result holds when "first-order" is replaced by "existential" in all cases.*

For example, suppose $L \subseteq \overline{\mathbb{Q}}$ is a field. If there is a formula $\phi(X)$ such that for any $x \in L$, $\phi(x)$ holds if and only if $x \in \mathcal{O}_L$, then we can reduce problems of decidability for L to the same problem for its ring of integers. Using this idea in the context of first-order decidability is a key component of Chapter 7. On the other hand, proving that this strategy rarely works for proving the existential undecidability of fields is the content of Chapter 6, where we will explicitly show (in a precise topological sense) that for most fields $L \subseteq \overline{\mathbb{Q}}$, the ring of integers $\mathcal{O}_L$ is neither existentially nor universally definable in L.

Chapter 6
Undefinability and Topology

6.1 Main result and comparison to prior work

As outlined in the previous chapter, we are primarily motivated by Hilbert's Tenth Problem and its generalizations, with an eye towards the important technique of definability. Recall that if $\mathbb{Z}$ is existentially definable in $\mathbb{Q}$, then a reduction argument shows that Hilbert's Tenth Problem for $\mathbb{Q}$ must be undecidable. However, if Mazur's Conjecture holds, then $\mathbb{Z}$ is actually not existentially definable in $\mathbb{Q}$. Proving this unconditionally currently appears to be out of reach. In fact, it seems generally very difficult to prove undefinability results for individual fields. One example of success is the field of all totally real algebraic numbers $\mathbb{Q}^{\text{tr}}$. Fried, Haran and Völklein showed that its first-order theory is decidable [34], while J. Robinson showed that the first-order theory of the ring of all totally real integers $\mathbb{Z}^{\text{tr}}$ is undecidable [82]. This difference in decidability implies that $\mathbb{Z}^{\text{tr}}$ cannot be first-order definable in the field $\mathbb{Q}^{\text{tr}}$. Another example is the ring $\overline{\mathbb{Z}}$ of all algebraic integers inside $\overline{\mathbb{Q}}$, which is undefinable by the strong minimality of $\overline{\mathbb{Q}}$. In both examples, the facts used for proving undefinability are not remotely close to necessary conditions for undefinability. Instead, they simply reflect the available pathways for unconditionally proving undefinability in a limited number of cases.

While it is still an open question whether $\mathbb{Z}$ is existentially definable in $\mathbb{Q}$, it is possible to give a first-order definition of $\mathbb{Z}$ in $\mathbb{Q}$, i.e. a definition that uses both existential and universal quantifiers. This was first done by J. Robinson [80], who generalized this result to define the ring of integers $\mathcal{O}_K$ inside any number field K [81]. Later, Rumely [83] was able to make the definition of the ring of integers uniform across number fields. Robinson's definition was improved by Poonen [76] who gave a $\forall\exists$-definition that in every number field K defines its ring of integers. Following this, Koenigsmann [55] proved that it is possible to give a universal definition of $\mathbb{Z}$ in $\mathbb{Q}$, i.e. a definition that only involves

universal ($\forall$) quantifiers, and Park extended his result to show that $\mathcal{O}_K$ is universally definable in K for every number field K [73]. This raises the question of whether we can expect universal and first-order definability to continue to hold for many infinite algebraic extensions of $\mathbb{Q}$.

Currently, first-order definability results are only known for certain classes of infinite extensions of the rationals. These are usually proved in order to prove the first-order undecidability of certain infinite extensions via reductions. For example, Videla proved the definability of the ring of integers over certain infinite algebraic pro-p extensions of $\mathbb{Q}$ [106], while Fukuzaki was able to define the ring of integers in infinite extensions in which every finite subextension has odd degree and that satisfy certain ramification conditions [36]. These results were further generalized by Shlapentokh in [91], to which we refer readers for more extensive background on known results for the first-order definability and decidability of infinite algebraic extensions of $\mathbb{Q}$. In Shlapentokh's framework, all known examples of algebraic extensions of $\mathbb{Q}$ with first-order definable rings of integers can be viewed as relatively small extensions which are somehow "close" to $\mathbb{Q}$. On the other hand, although first-order definability seems less likely for extensions which are similarly "far from" $\mathbb{Q}$, very few such examples are known, as mentioned above.

In this chapter, we take the perspective of considering all algebraic extensions of $\mathbb{Q}$ simultaneously. From this vantage point, we prove that $\mathcal{O}_K$ is both existentially and universally undefinable in K for "most" algebraic extensions K of $\mathbb{Q}$. To make this notion precise, we view the set $\mathrm{Sub}(\overline{\mathbb{Q}})$ of subfields of $\overline{\mathbb{Q}}$ as a topological space by considering it as a subset of $2^{\overline{\mathbb{Q}}}$, from which it inherits the product topology. In this topology, every nonempty open set is non-meager. The precise version of our theorem then can be written as follows:

Theorem 6.1.1. *The set of algebraic extensions K of $\mathbb{Q}$ for which $\mathcal{O}_K$ is existentially or universally definable is a meager subset of* $\mathrm{Sub}(\overline{\mathbb{Q}})$.

In particular, there are uncountably many algebraic extensions K of $\mathbb{Q}$ for which the ring of integers $\mathcal{O}_K$ is neither existentially nor universally definable in K. While our initial interest in this area came from questions about the definability of rings of integers in fields, our proofs never use the fact that the definable set in question is a ring, and the most general version of our theorem, stated below as Theorem 6.5.9, concerns definability for arbitrary subsets of algebraic fields. After seeing one of the authors speak on these results, Philip Dittmann and Arno Fehm generalized Theorem 6.1.1 in a different way, in [24], improving "existentially or universally definable" to "definable," using the fact

that $\mathcal{O}_K$ forms a ring. Their proof uses techniques from model theory, entirely different from those employed here.

Our original motivation, following [69], was to obtain Theorem 6.1.1 for the quotient space $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$ which only considers fields up to isomorphism. This goal is achieved in Corollary 6.5.15 as a consequence of Theorem 6.1.1. However, the topology on $\mathrm{Sub}(\overline{\mathbb{Q}})$ is itself quite natural, and coincides (via the Galois correspondence) with the Vietoris topology on the space of closed subgroups of $\mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$. We thank Florian Pop for pointing out this connection to us, and for alerting us that the same topology appears in [77], where it is called the strict topology. The topology has also been used by other authors: for examples, see [32, 33, 38, 43].

To prove this theorem, we study the existential definability of infinite sets $Z \subseteq \mathbb{Q}$ whose complement is not thin, in the sense of Serre. The necessary background of algebraic number theory, arithmetic geometry and thin sets is recalled in Section 6.2. In order to prove the main theorem, we introduce a notion of rank in Section 6.3 that formalizes which existential formulas are the "simplest." In particular, rank is a well-ordering of existential formulas, so if Z is existentially definable in $\mathbb{Q}$ over some field $L \subseteq \overline{\mathbb{Q}}$, then there is a formula of minimal rank which does the job. By studying such minimal-rank formulas in Section 6.4, we deduce a convenient normal form for existential definitions; see Theorem 6.4.8. Finally, we introduce the topological spaces of $\mathrm{Sub}(\overline{\mathbb{Q}})$ and $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$ in Section 6.5, and use the normal form to deduce the main result via Hilbert's Irreducibility Theorem. In fact, the proof also leads to an algorithm which, given a basic open subset $U \subseteq \mathrm{Sub}(\overline{\mathbb{Q}})$, produces a computable field $L \in U$ in which the ring of integers $\mathcal{O}_L$ is neither existentially or universally definable; see Theorem 6.5.12.

6.2 Background from number theory

In this section, we will recall some of the basic facts that we will require for fields, thin sets, and affine varieties. Readers can find additional background in the books of Lang [57], Serre [85] and Liu [62], respectively.

6.2.1 Field extensions and the irreducibility of polynomials

In the material that follows, we will be presented with the following question: Given number fields $F \subseteq K$, which field extensions of F contain elements of the complement $K \setminus F$? This question is intimately related to the irreducibility of polynomials. First, we

recall a basic result on the irreducibility of multivariable polynomials.

Lemma 6.2.1. *If K/F is an extension of fields within a larger field L, and $z \in L$ is algebraic over F with $F(z) \cap K \neq F$, then the minimal polynomial $h(Z)$ of f over F must be reducible over K.*

Proof. By hypothesis $1 < [F(z) \cap K : F]$, so

$$[F(z) : F(z) \cap K] < [F(z) : F(z) \cap K] \cdot [F(z) \cap K : F] = [F(z) : F].$$

From this it follows that $h(Z)$ must factor over $F(z) \cap K$, so it certainly also factors over the larger field K. □

The next proposition forms a kind of converse to Lemma 6.2.1 when K/F is a finite Galois extension. Given an algebraic function field $E = \mathrm{Frac}(F[Y_0, Y_1, \ldots, Y_m]/(f))$ where $f \in F[Y_0, Y_1, \ldots, Y_m]$ is an irreducible polynomial, the *constant field* of E is the set of elements which are algebraic over F.

Proposition 6.2.2. *Let F be a number field, and K a finite Galois extension of F. If $m \geq 0$ and $f \in F[Y_0, Y_1, \ldots, Y_m]$ is an irreducible polynomial that becomes reducible in $K[Y_0, Y_1, \ldots, Y_m]$, then the constant field of $E = \mathrm{Frac}(F[Y_0, Y_1, \ldots, Y_m]/(f))$ is larger than F. In particular, there is an element $z \in E \setminus F$ such that there is an F-linear field embedding of $F(z)$ into K with the image of z lying in $K \setminus F$.*

Proof. Assume without loss of generality that Y_m appears nontrivially in f, and write $L = F(Y_0, Y_1, \ldots, Y_{m-1})$. We will view $E = L(\theta)$ for an element θ in the algebraic closure $\overline{L}$ with minimal polynomial f. Similarly consider K to be an extension of F inside $\overline{L}$.

Suppose that E contains no elements of $K \setminus F$. Then $E \cap K = L \cap K = F$, and a basic theorem of Galois theory [57, Theorem 1.12] implies the following because K is a Galois extension of F:

$$[EK : E] = [K : E \cap K] = [K : F] = [K : L \cap K] = [LK : L].$$

Using the diamond written below, we deduce that $[E : L] = [EK : LK]$. Importantly, these field extension degrees are also the degrees of the minimal polynomial of θ over L

and LK, respectively.

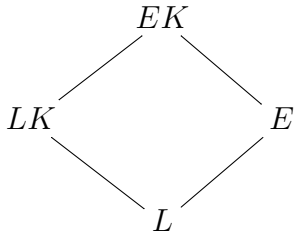

This shows that f remains irreducible over the field $L = K(Y_0, Y_1, \ldots, Y_{m-1})$ as a polynomial in Y_m. We claim that f is actually irreducible as an element of the ring $K[Y_0, Y_1, \ldots, Y_m]$, which contradicts the hypothesis. To prove this, it only remains to show that the coefficients of f lying in $K[Y_0, \ldots, Y_{m-1}]$ have no common factor; see [57, IV.2.3]. Clearly, as a polynomial in Y_m, the coefficients of f lying in $F[Y_0, \ldots, Y_m]$ have no common factor over F because f is irreducible over F. In fact, this implies that the coefficients also have no common factor over any algebraic extension of F by the following lemma, which completes the proof. □

Lemma 6.2.3. *Let F be a field and let F' be a separable extension. If $f_0, f_1, \ldots, f_k$ are a collection of polynomials in $F[Y_0, \ldots, Y_m]$ with no common factor, then $f_0, \ldots, f_k$ also have no common factor over the extension F'.*

Proof. By writing $f_0, \ldots, f_k$ in terms of their irreducible factors, we can reduce without loss of generality to the case of two irreducible polynomials $f_0, f_1 \in F[Y_0, \ldots, Y_m]$. Indeed, for every irreducible factor p of f_0, there is a polynomial f_j for $1 \leq j \leq m$ which is not divisible by p, and it suffices to show that the irreducible factors of f_j remain relatively prime to p over the larger field F'.

Notice that irreducible polynomials f_0 and f_1 are relatively prime over F if and only if $f_0 f_1$ generates a radical ideal in $F[Y_0, \ldots, Y_m]$, i.e. if and only if $F[Y_0, \ldots, Y_m]/(f_0 f_1)$ is a reduced ring. The latter condition is stable under separable field extensions, i.e. $F'[Y_0, \ldots, Y_m]/(f_0 f_1)$ is also reduced; see [62, Proposition 3.2.7.(b)]. Therefore f_0 and f_1 have no common factors over F'. □

6.2.2 Dimensions of rings and affine varieties

We will require a usable notion of dimension, which can equivalently be viewed as a geometric or algebraic phenomenon. In particular, there are related notions of the dimension of a commutative ring A, and the dimension of the associated topological

space Spec A consisting of all prime ideals of A with the Zariski topology. In this section, we will review some basic facts of commutative algebra and algebraic geometry, limiting the discussion to only what is necessary for our purposes.

First, let us recall this topology and some basic notation. Given a commutative ring A, the set Spec A is endowed with the *Zariski topology* by defining the following as basic closed and open sets, respectively. For any ideal $I \subseteq A$, we define $V(I)$ to be the subset of Spec A consisting of all prime ideals that contain I, and $D(f) = \text{Spec } A \setminus V(f)$. Notice that it is natural via the isomorphism theorems for rings to identify $V(I)$ with Spec A/I. With this notation, the closed subsets of Spec A in the Zariski topology are precisely the sets of the form $V(I)$ where $I \subseteq A$ is an ideal, and sets of the form $D(f)$ for $f \in A$ form a base for the open subsets of Spec A. In fact, Spec A is an *affine scheme*, meaning that it has even more structure than just a topology, although we will not require this full structure; see [62, Chapter 2] for more background.

In this chapter, we consider the ring $A = F[Y_0, \ldots, Y_m]$ and its quotients, where F is a subfield of $\overline{\mathbb{Q}}$. An *affine variety over* F is an object of the form $V(I) = \text{Spec } F[Y_0, \ldots, Y_m]/I$ for some $m \geq 0$ and some ideal $I \subseteq F[Y_0, Y_1, \ldots, Y_m]$. Furthermore, if the quotient $F[Y_0, \ldots, Y_m]/I$ is an integral domain, then the corresponding affine variety is called *integral*. We will write $V(I) = V(f_1, \ldots, f_k)$ when the ideal $I \subseteq F[X, Y_1, \ldots, Y_m]$ is generated by $\{f_1, \ldots, f_k\}$. If there is ambiguity about the base field, then we will write V_F instead of V for clarity.

Given an affine variety $V = \text{Spec } F[Y_0, \ldots, Y_m]/I$, the *rational points* of V (over F) are the tuples $(y_0, \ldots, y_m) \in F^m$ such that $f(y_0, \ldots, y_m) = 0$ for all $f \in I$. The set of rational points can be identified with the set of all F-algebra homomorphisms $\varphi : F[Y_0, \ldots, Y_m]/I \to F$. We refer the reader to [62, Section 2.3.2] for more details. As we are frequently working over non-algebraically closed fields, it is possible for nontrivial affine varieties to have no rational points, such as the affine variety Spec $\mathbb{Q}[Y_0, \ldots, Y_m]/(Y_0^2 + \cdots + Y_m^2 + 1)$ for any $m \geq 0$. We can view the varieties as geometric objects which help us find and describe the rational points.

The *Krull dimension* of a ring A, written $\dim A$, is the supremal length r of a chain of prime ideals $\mathfrak{p}_0 \subsetneq \cdots \subsetneq \mathfrak{p}_r$ in A. Similarly, given a topological space X, we define $\dim X$ to be the supremal length r of a chain of irreducible closed subsets $Z_0 \subsetneq \cdots \subsetneq Z_r$ in X. The following proposition equates these two notions of dimension. Recall that the nilradical of a commutative ring is the set of all nilpotent elements, or equivalently the intersection of all prime ideals.

Proposition 6.2.4 (Proposition 2.5.8, [62])**.** *Let A be a (commutative) ring and let N*

be the nilradical of A. Then $\dim \operatorname{Spec} A = \dim A = \dim A/N$.

In our applications, we need to understand the dimension of subsets of affine varieties. Recall that if X is any topological space and Y is any subset of X endowed with the subset topology, then $\dim Y \leq \dim X$ [62, Proposition 2.5.5]. In the context of affine varieties and open subsets, this inequality is often an equality due to the fact that open subsets in the Zariski topology are "large". This idea is formulated precisely in the following proposition. Given a field extension L/F, we write $\operatorname{trdeg}_F L$ for the *transcendence degree* of L over F. If $X = \operatorname{Spec} A$ is an integral affine variety, we call $\operatorname{Frac}(A)$ the *function field* of X.

Proposition 6.2.5 (Proposition 2.5.19, [62])**.** *If $X = \operatorname{Spec} A$ is an integral affine variety over a field F, then*

$$\dim U = \dim X = \operatorname{trdeg}_F \operatorname{Frac}(A)$$

for any nonempty open subset $U \subseteq X$.

Similarly, it is helpful to know when a subset of a topological space X has strictly smaller dimension than X. In contrast to the result immediately above, this often happens for proper closed subsets of an affine variety.

Proposition 6.2.6 (Corollary 2.5.26, [62])**.** *Let $X = \operatorname{Spec} A$ be an integral affine variety. If $f \in A$ is nonzero, then every irreducible component of $V(f)$ has dimension $\dim X - 1$. In particular, every proper closed subset of X has strictly smaller dimension than X.*

So far in this section, the definition of dimension depends on the base field $F \subseteq \overline{\mathbb{Q}}$, *a priori*. However, the result below clarifies that dimension stays the same under base extension. This allows us to ignore the field of definition to some extent, especially when defining the rank of a formula below, although the notion of integrality truly does depend on the base field, so care is still required when applying the previous two propositions.

Proposition 6.2.7 (Proposition 3.2.7, [62])**.** *Let $F \subseteq L \subseteq \overline{\mathbb{Q}}$ be fields. Given an affine variety $V_F(f_1, \ldots, f_k) = \operatorname{Spec} F[Y_0, \ldots, Y_m]/(f_1, \ldots, f_k)$, the affine variety*

$$V_L(f_1, \ldots, f_k) = \operatorname{Spec} L[Y_0, \ldots, Y_m]/(f_1, \ldots, f_k)$$

is the base extension *of the variety $V_F(f_1, \ldots, f_k)$ to L, and these affine varieties have the same dimension.*

To apply this proposition to open sets, we remark that open sets can be equivalently viewed as affine varieties themselves, albeit in a different ambient space with an extra variable.

Corollary 6.2.8. *Let $F \subseteq \overline{\mathbb{Q}}$ be a field. For polynomials $g, f_1, \ldots, f_k \in F[Y_0, \ldots, Y_m]$, define $A = F[Y_0, \ldots, Y_m]/(f_1, \ldots, f_k)$ and let A_g be the localization of A be the element g. Then there are isomorphisms of ringed topological spaces*

$$V_F(f_1, \ldots, f_k) \cap D(g) \cong \operatorname{Spec}(A_g) \cong V_F(f_1, \ldots, f_k, Y_{m+1}g - 1).$$

In particular, $\dim V_F(f_1, \ldots, f_k) \cap D(g) = \dim V_L(f_1, \ldots, f_k) \cap D(g)$ *for any algebraic extension of fields $L \supseteq K$.*

Proof. The first isomorphism is [62, Lemma 2.3.7]. The second isomorphism actually follows from a well-known isomorphism of underlying rings

$$A_g \cong F[Y_0, \ldots, Y_{m+1}]/(f_1, \ldots, f_k, Y_{m+1}g - 1);$$

see [79, Lemma §6.2]. Therefore, the statement on dimension follows immediately from Proposition 6.2.7. □

6.2.3 Thin sets

Hilbert's Irreducibility Theorem can take many different forms, but we put a simple version here that suffices for the purposes of this article. For brevity, we present *thin sets* as a black box, and refer the reader to [85, Prop. 3.3.5] for more details. Essentially, a thin subset $T \subseteq K$ of a number field is small, in the view of arithmetic geometry. For example, any set of points that is contained in a closed subvariety of affine n-space K^n, and which is different from the entire space, is thin with respect to K. All necessary details can be deduced from the results we recall below .

Theorem 6.2.9 (Hilbert's Irreducibility Theorem)**.** *Let $f(Y_0, Y_1, \ldots, Y_m)$ be a polynomial with coefficients in a number field K which is irreducible as an $(m+1)$-variable polynomial. There exists a thin set $T \subseteq K^m$ such that if $(y_1, \ldots, y_m) \in K^m \setminus T$, then $f(Y_0, y_1, \ldots, y_m)$ is an irreducible single-variable polynomial of degree* $\deg_{Y_0}(f)$.

In order for the theorem above to be non-trivial, we need to know that K^m is not a thin subset of itself, and this is indeed true for all number fields [85, Prop 3.4.1].

Moreover, the propositions below show that thin sets cannot contain arithmetically important subsets, which will allow us to use Hilbert's Irreducibility Theorem in the cases we care about.

Proposition 6.2.10 (Proposition 3.2.1, [85])**.** *If L/K is a finite extension of fields and $T \subseteq L^m$ is thin with respect to L, then $T \cap K^m$ is thin with respect to K.*

Proposition 6.2.11. *If K is a number field, then no thin subset of K contains either $\mathbb{Z}$ or $\mathbb{Q} \setminus \mathbb{Z}$.*

Proof. Thin sets of $\mathbb{Q}$ cannot contain $\mathbb{Z}$ or $\mathbb{Q} \setminus \mathbb{Z}$ by [85, Theorem 3.4.4] and [85, Prop. 3.4.2], respectively. Thus, the result for arbitrary number fields follows from Proposition 6.2.10. □

Moreover, we can understand thin sets in products. This lemma will be used to show that if a set $Z \subseteq \mathbb{Q}$ is not thin, then the product $Z \times \mathbb{Q}^n$ cannot be thin, either.

Lemma 6.2.12. *If $n \geq 0$ and $S \subseteq \mathbb{Q}$ is a set such that $S \times \mathbb{Q}^n \subseteq \mathbb{Q}^{n+1}$ is thin, then $S \subseteq \mathbb{Q}$ is thin.*

Proof. There is a line $\mathcal{L} \subseteq \mathbb{Q}^{n+1}$ such that $\mathcal{L} \cap (S \times \mathbb{Q}^n)$ is thin in $\mathcal{L}$ and the projection of $\mathcal{L}$ to the first coordinate is all of $\mathbb{Q}$ [85, Proposition 3.2.3]. As $\mathcal{L}$ is a line, this projection is an isomorphism and $\mathcal{L} \cap (S \times \mathbb{Q}^n)$ maps onto to the set S. Therefore, S is thin in $\mathbb{Q}$. □

Finally, we prove a proposition that lets us stitch this material together. This is ultimately the result that is required in the proof of our main theorem.

Proposition 6.2.13. *Let K be a number field and let $f(X, Y_1, \ldots, Y_m), g(X, Y_1, \ldots, Y_m) \in K[X, Y_1, \ldots, Y_m]$ be relatively prime irreducible polynomials. Then there is a thin set $T \subseteq K^m$ such that $f(x, y_1, \ldots, y_{m-1}, Y)$ and $g(x, y_1, \ldots, y_{m-1}, Y)$ are relatively prime irreducible single-variable polynomials for every $(x, y_1, \ldots, y_{m-1}) \in K^m \setminus T$, of degrees $\deg_{Y_m}(f)$ and $\deg_{Y_m}(g)$, respectively.*

Proof. Take T_0 to be the union of the two thin sets given by applying Hilbert's Irreducibility Theorem to f and g separately. By construction, $f(x, y_1, \ldots, y_{m-1}, Y)$ and $g(x, y_1, \ldots, y_{m-1}, Y)$ are irreducible polynomials in Y for every $(x, y_1, \ldots, y_{m-1}) \in K^m \setminus T_0$, and it only remains to check the claim of relative primality.

If $\deg_{Y_m}(f) \neq \deg_{Y_m}(g)$, then this claim is trivial. Therefore, write $d = \deg_{Y_m}(f) = \deg_{Y_m}(g)$, and consider $(x, \ldots, y_{m-1}) \in K^m \setminus T$. Since the polynomials $f(x, y_1, \ldots, y_{m-1}, Y)$

and $g(x, y_1, \ldots, y_{m-1}, Y)$ are irreducible, the failure of relative primality implies that they are unit multiples of each other, i.e., $f(x, y_1, \ldots, y_{m-1}, Y) = zg(x, y_1, \ldots, y_{m-1}, Y)$ for some nonzero $z \in K$. In particular, if we write

$$f(X, Y_1, \ldots, Y_m) = \sum_{i=0}^{d} f_i(X, Y_1, \ldots, Y_{m-1}) Y_m^i,$$

$$g(X, Y_1, \ldots, Y_m) = \sum_{i=0}^{d} g_i(X, Y_1, \ldots, Y_{m-1}) Y_m^i,$$

where $f_i, g_i \in K[X, Y_1, \ldots, Y_{m-1}]$ are polynomials, then this condition is the same as

$$f_i(x, y_1, \ldots, y_{m-1}) = zg_i(x, y_1, \ldots, y_{m-1})$$

for all $0 \leq i \leq d$. Multiplying these conditions together, we get the equations

$$f_i g_j = z g_i g_j = g_i f_j$$

for $0 \leq i, j \leq d$. We will show that this system of equations holds only inside a thin set, which completes the proof.

We claim that the polynomial

$$f_i(X, Y_1, \ldots, Y_{m-1}) g_j(X, Y_1, \ldots, Y_{m-1}) - g_i(X, Y_1, \ldots, Y_{m-1}) f_j(X, Y_1, \ldots, Y_{m-1})$$

is nonzero for some choice of i and j. Indeed, if this were not the case, then we would find that

$$\begin{aligned} f_i(X, Y_1, \ldots, Y_{m-1}) g(X, Y_1, \ldots, Y_m) &= \sum_{j=0}^{d} f_i(X, Y_1, \ldots, Y_{m-1}) g_j(X, Y_1, \ldots, Y_{m-1}) Y_m^j \\ &= \sum_{j=0}^{d} g_i(X, Y_1, \ldots, Y_{m-1}) f_j(X, Y_1, \ldots, Y_{m-1}) Y_m^j \\ &= g_i(X, Y_1, \ldots, Y_{m-1}) f(X, Y_1, \ldots, Y_m) \end{aligned}$$

for all i. As g and f are irreducible and the only polynomials on the left and right sides of the equation containing the variable Y_m, we conclude that they are unit multiples of each other, which contradicts the hypothesis of relative primality.

Therefore, let T_1 be the set of all K-rational points on the affine variety

$$V_K(\{f_i g_j - g_i f_j : 0 \leq i < j \leq \deg_{Y_m}(f)\}).$$

Since one of the polynomials in the defining set is nonzero, the affine variety is a proper closed variety, which implies that T_1 is a thin set by definition. By construction, the set $T = T_0 \cup T_1$ is the desired thin set. □

6.3 Rank of a Formula

The goal of this section is to define a notion of rank for existential formulas in the language of fields, using degrees of polynomials and dimensions of varieties, as well as the number of ∃-quantifiers used. Certain formulas will have the same rank, just as certain polynomials have the same degree. Crucially, the ranks are well-ordered.

6.3.1 A useful well-ordering

Definition 6.3.1. *Let $(\mathcal{L}, <)$ be a linear order. For a finite tuple $(a_0, \ldots, a_n) \in \mathcal{L}^{<\omega}$, write $\vec{a}^*$ for the tuple of the same $(n+1)$ elements (including repetitions) arranged in $<$-descending order: $\vec{a}^* = (a_{\alpha(0)}, \ldots, a_{\alpha(n)})$ where α is a permutation and $a_{\alpha(i+1)} \leq a_{\alpha(i)}$ for all $i < n$. Write $\vec{a} =^* \vec{b}$ just if $\vec{a}^* = \vec{b}^*$.*

Then the $*$-order $(\mathcal{L}^*, <^*)$ is the lexicographic order $<^*$ (defined using $<$ on individual coordinates) on the set $\mathcal{L}^*$ of $=^*$-equivalence classes in $\mathcal{L}^{<\omega}$. To be clear: if $\vec{a}^*$ is a proper initial segment of $\vec{b}^*$, then $\vec{a}^* <^* \vec{b}^*$.

Equivalently, one can view the elements of $\mathcal{L}^*$ as finite multisets of elements of $\mathcal{L}$, with the elements of each multiset listed in $<$-nonincreasing order.

Lemma 6.3.2. *If $(\mathcal{L}, <)$ is a well order, then so is $(\mathcal{L}^*, <^*)$.*

Proof. Clearly $<^*$ is a linear order. If it were not a well order, there would be a least $a \in \mathcal{L}$ such that some infinite $<^*$-descending sequence begins with an $\vec{a}^*$ whose greatest element is a. Choose such an $\vec{a}^* = (a^k, a_1, \ldots, a_n)$, in nonincreasing order with $a_1 < a$ after a appears k times, with k as small as possible (and allowing $n = 0$). Then the infinite descending sequence beginning with this $\vec{a}^*$ can only have finitely many terms that begin with a^k, for if there were infinitely many, then by "chopping off" the a^k from each term, we would get an infinite sequence contradicting the choice of a. But then,

immediately after the last term beginning with a^k comes a term beginning with a^j for $j < k$, and this term also begins an infinite descending sequence in $\mathcal{L}^*$, contradicting either the minimality of k (if $j > 0$) or the minimality of a (if $j = 0$). □

6.3.2 Definition of rank

We present an explicit way to put a well-ordering on the set of existential formulas with parameters in any given field. This is done by associating a *rank* to every existential formula.

Any existential formula $\alpha(X)$ can be written in disjunctive normal form

$$\alpha(X) = \exists \vec{Y}(\alpha_1 \vee \alpha_2 \vee \cdots \vee \alpha_n),$$

where each $\alpha_i(X, \vec{Y})$ is a conjunction of equations and inequations. Bringing the existential quantifiers inside the disjunctions and discarding any unused quantifiers, any existential formula can be rewritten as

$$((\exists Y_1 \cdots \exists Y_{m_1})\alpha_1) \vee \cdots \vee ((\exists Y_1 \cdots \exists Y_{m_n})\alpha_n),$$

where all variables $Y_1, \ldots, Y_{m_i}$ appear in α_i. One can also easily rearrange any $\alpha_i(X, \vec{Y})$ into a conjunction of the form

$$f_1(X, \vec{Y}) = \cdots = f_k(X, \vec{Y}) = 0 \ \& \ g(X, \vec{Y}) \neq 0.$$

Only one inequation $g \neq 0$ is needed, as several $g_i(X, \vec{Y})$ could be multiplied together. It is allowed for g to be the constant 1. We call an existential formula *rankable* if it is given in the above format. It is trivial to rearrange any existential formula into rankable format, so in this chapter any existential formula which appears is assumed to be rankable.

Before defining rank, we present a way to order tuples of polynomials. Notice that this notion depends on a specific order for the variables.

Definition 6.3.3. *For the variables $X, Y_1, \ldots, Y_m$, the* multidegree *of a monomial $X^c Y_1^{d_1} \cdots Y_m^{d_m}$ is $(c, d_1, \ldots, d_m)$, and these $(m+1)$-tuples are ordered by the reverse lexicographic order. The* multidegree $\mathrm{mdeg}(f)$ of a polynomial f *is the maximum of the multidegrees of each monomial appearing (with nonzero coefficient) in it.*

Observe that the linear order defined above on multidegrees is a well-ordering.

Definition 6.3.4. *A* basic rankable formula *is an existential formula of the form*

$$\exists Y_1 \cdots \exists Y_m \ [f_1(X, Y_1, \ldots, Y_m) = \cdots = f_k(X, Y_1, \ldots, Y_m) = 0 \ \& \ g(X, \vec{Y}) \neq 0],$$

and the rank *of such a formula is the triple*

$$\mathrm{rk}(\beta) = (m, e, (\mathrm{mdeg}(f_1), \ldots, \mathrm{mdeg}(f_k))^*),$$

where the second component is the dimension e of $V_{\overline{\mathbb{Q}}}(\vec{f}) \cap D(g)$, as defined in Section 6.2.2, and the third component uses the $=^*$-classes of tuples of multidegrees, as in Definition 6.3.1.

In this definition, we see that $V_{\overline{\mathbb{Q}}}(\vec{f}) \cap D(g)$ is a subset of an ambient space of dimension $m+1$. Therefore, the first coordinate of the definition of rank can be equivalently viewed as a measure of the dimension of this ambient space. Additionally, by Corollary 6.2.8, the base field does not matter in the definition of the dimension e, so we will usually drop the $\overline{\mathbb{Q}}$ from this notation.

We define an order $\prec$ on ranks of basic rankable formulas in forwards lexicographic order, meaning that

$$(m, e, (d_1, \ldots, d_k)^*) \prec (m', e', (d'_1, \ldots, d'_{k'})^*)$$

if and only if one of the following holds:

- $m < m'$, i.e., the first formulas uses fewer $\exists$-quantifiers; or
- $m = m'$ and $e < e'$, so the first formula defines an open variety of lesser dimension than the second; or
- $m = m'$ and $e = e'$ and $(d_1, \ldots, d_k)^* <^* (d'_1, \ldots, d'_{k'})^*$, so the first formula uses polynomials of lower multidegree.

The least possible rank of a (satisfiable) basic rankable formula is $(0, 0, (1)^*)$, which is the rank of the quantifier-free formula $X = x$ for any specific value x: here $m = 0$, $k = 1$ and the variety, which has a single component whose dimension is 0, is defined by $f_1 = X - x = 0$ whose multidegree (in the single variable X, since $m = 0$) is simply 1. (The variety defined by $0 = 0$ has dimension 1, so the formula $0 = 0$ has higher rank.)

Let $\mathcal{R}$ denote the set of all possible ranks of basic rankable formulas. Then $(\mathcal{R}, \prec)$ is a well-ordering. (The third component of $\prec$ is well-ordered by Lemma 6.3.2.) Let $(R^*, \prec^*)$ be the result of applying Definition 6.3.1 to $(\mathcal{R}, \prec)$.

Observe that an existential formula is rankable if and only if it is the finite disjunction of basic rankable formulas.

Definition 6.3.5. *If $\alpha = \vee_{i=1}^{r} \beta_i$ is a rankable formula, the* rank *of α is defined to be*

$$\mathrm{rk}(\alpha) = (\mathrm{rk}(\beta_1), \ldots, \mathrm{rk}(\beta_n))^* \in \mathcal{R}^*$$

The rankable formulas can then be compared using the ordering $\prec^*$. By Lemma 6.3.2, $(\mathcal{R}^*, \prec^*)$ is a well-order.

6.4 Minimal formulas and hypersurfaces

The well-ordering of ranks means that every nonempty set of existential formulas has an element of minimal rank. For example, if there exists an existential formula that defines $\mathcal{O}_L$ in L, then there is an existential formula α that accomplishes this which has minimal rank among all such formulas. Such a formula can be considered a minimal successful formula. This motivates the following general definition.

Definition 6.4.1. *For a field $L \subseteq \overline{\mathbb{Q}}$ and an existential formula $\alpha(X)$ with coefficients from L, we say α is* L-minimal *if α has minimal rank among all existential formulas α' for which*

$$\forall x(\alpha(x) \iff \alpha'(x))$$

holds over L.

In order for the above to make sense, α' ranges only over those existential formulas which have coefficients from L. We will show that any L-minimal formula must take the form of a disjunction of formulas with two very simple formats: quantifier-free formulas, and formulas with only one equation and one inequation.

We will start by considering a general rankable formula, then minimize it as much as possible. First, we want to minimize the number of quantifiers, which is the first component of rank. Clearly, we can eliminate the quantifier for any variable that does not appear in any polynomial of the formula. The following simple lemma allows us also to remove any variables that appear in the inequation, but none of the equations.

Lemma 6.4.2. *Let $1 \leq e < m$ and let $\delta(X)$ be the basic rankable existential formula*

$$\exists Y_1 \cdots \exists Y_m \; [f_1(X, Y_1, \ldots, Y_e) = \cdots = f_k(X, Y_1, \ldots, Y_e) = 0 \neq g(X, Y_1, \ldots, Y_m)],$$

where $f_i \in F[X, Y_1, \ldots, Y_e]$ and $g \in F[X, Y_1, \ldots, Y_m]$ for some field F.

Then there are polynomials $g_1, \ldots, g_r \in F[X, Y_1, \ldots, Y_e]$ such that $\delta(X)$ is equivalent over F to the disjunction of formulas

$$\vee_{i=1}^{r} \exists Y_1 \cdots \exists Y_e \ [f_1(X, \vec{Y}) = \cdots = f_k(X, \vec{Y}) = 0 \neq g_i(X, \vec{Y})].$$

Proof. Write out $g = \sum_{i=0}^{d_m} g_i(X, Y_1, \ldots, Y_{m-1}) Y_m^i$ as a polynomial in Y_m. Notice that if $(x, y_1, \ldots, y_{m-1}) \in \overline{\mathbb{Q}}^m$ is any tuple, then there is a $y_m \in \mathbb{Q}$ such that $g(x, y_1, \ldots, y_m) \neq 0$ if and only if $g_i(x, y_1, \ldots, y_{m-1}) \neq 0$ for some $0 \leq i \leq d_m$. Therefore, we can remove the quantifier for Y_m and instead use a disjunction where g is replaced by g_j for $0 \leq j \leq d_m$ in each formula. By induction, this completes the proof. □

To continue minimizing the number of quantifiers, we can take a more geometric perspective. A basic rankable formula $\beta(X)$ with m quantifiers

$$\exists Y_1 \cdots \exists Y_m \ [f_1(X, \vec{Y}) = \cdots = f_k(X, \vec{Y}) = 0 \neq g(X, \vec{Y})]$$

corresponds to the projection to the X-coordinate of the points on the variety $D(g) \cap V(f_1, \ldots, f_k)$. Minimizing the number of quantifiers m is equivalent to minimizing the dimension $m + 1$ of the ambient space where the variety lives. If k is large, then we expect the dimension e of the variety to be much smaller than $m+1$, and we can consider this "wasteful," as it uses more variables than necessary. The following proposition uses a basic result of algebraic geometry to show that, in a special case with integral affine varieties, we only need $m = e$ quantifiers and a single equation to describe all but a lower-dimensional closed subset. To complete the section, we will the show that this is enough to deduce the result in general.

Proposition 6.4.3. *Let $F \subseteq \overline{\mathbb{Q}}$ be a field and $\mathfrak{p} = (f_1, \ldots, f_k) \subseteq F[X, Y_1, \ldots, Y_m]$ a prime ideal. Define $\beta(X)$ to be the formula*

$$\beta(X) = \exists Y_1, \ldots, Y_m [f_1(X, Y_1, \ldots, Y_m) = \cdots = f_k(X, Y_1, \ldots, Y_m) = 0]$$

and set $e = \dim V_F(\mathfrak{p})$. If $\beta(X)$ is satisfied by infinitely many values of X in $\overline{\mathbb{Q}}$ and $e \leq m - 1$, then after possibly reordering indices, there are polynomials $h \in F[X, Y_1, \ldots, Y_e]$ and $s \in F[X, Y_1, \ldots, Y_{e-1}]$ with h irreducible and $s \notin \mathfrak{p}$ such that $\beta(X)$ is equivalent to $\gamma_1(X) \vee \gamma_2(X)$ over F, where

$$\gamma_1(X) = \exists Y_1 \cdots \exists Y_e \ [h(X, \ldots, Y_e) = 0 \neq s(X, \vec{Y})],$$

$$\gamma_2(X) = \exists Y_1 \cdots \exists Y_m \ [s(X, \ldots, Y_m) = f_1(X, \ldots, Y_m) = \cdots = f_k(X, \ldots, Y_m) = 0].$$

Proof. Write $L = \mathrm{Frac}(F[X, Y_1, \ldots, Y_m]/\mathfrak{p})$. By Proposition 6.2.5, we know that e is equal to the transcendence degree of L over F. Since the images of $\{X, Y_1, \ldots, Y_m\}$ generate L over F, there is a transcendence basis consisting of a subset of these elements, and we can force $\bar{X}$ to be in this basis because $\bar{X}$ is not algebraic over F [57, Theorem VIII.1.1]. Indeed, if $\bar{X}$ were algebraic over F, then it would be the root of a single-variable polynomial over F, and therefore $\beta(X)$ would only be solvable over $\overline{\mathbb{Q}}$ by finitely many X, which is not the case by hypothesis.

Reorder the variables so that $\{\bar{X}, \bar{Y}_1, \ldots, \bar{Y}_{e-1}\}$ is a transcendence basis of L over F. Write $L_0 = F(X, Y_1, \ldots, Y_{e-1})$. Although a particular ordering of the variables is used when defining the multidegree component of rank in Definition 6.3.4, we will produce lower-rank formulas purely in terms of quantifiers and dimension, and therefore the multidegree will not matter here. As L is a finite separable extension of L_0, the primitive element theorem states that $L = L_0(\theta)$ for a single element θ. Write $h \in L_0[Y]$ for the minimal polynomial of θ. By clearing denominators if necessary, we can assume without loss of generality that $h \in F[X, Y_1, \ldots, Y_{e-1}, Y]$ is an irreducible multivariable polynomial. Therefore, writing $\mathfrak{p} = (f_1, \ldots, f_k)$, we have an isomorphism of fields:

$$L_0[Y_e, \ldots, Y_m]/(f_1, \ldots, f_k) \cong L \cong L_0[Y]/(h) \cong \mathrm{Frac}(F[X, Y_1, \ldots, Y_{e-1}, Y]/(h)).$$

Geometrically, this says that the integral affine variety $V_F(\mathfrak{p})$ is birational to the hypersurface $V_F(h)$. In fact, we can see that the two varieties contain isomorphic open sets, as follows.

Using the isomorphism of fields we can write $Y_j = \sum_{\ell=0}^{N_j} c_{j,\ell} Y^\ell$ for every $e_\mathfrak{p} \leq j \leq m$, and $Y = \sum_{\vec{a}} d_{\vec{a}} Y_{e_\mathfrak{p}}^{a_0} \ldots Y_m^{a_{m-e_\mathfrak{p}}}$, where $c_{j,\ell}$ and $d_{\vec{a}}$ are elements of L_0, and in particular not contained in $\mathfrak{p}$ because L_0 is a subfield of the function field of $V_F(\mathfrak{p})$. Let s be the products of all denominators appearing in these terms. Then these equations give an isomorphism of the open sets $V_F(\mathfrak{p}) \cap D(s)$ and $V_F(h) \cap D(s)$; see [62, Lemma 3.7]. Moreover, the X-coordinate of rational points is unchanged by the isomorphism because we included X in the transcendence basis. As $V_F(\mathfrak{p}) = (V_F(\mathfrak{p}) \cap D(s)) \cup V_F(\mathfrak{p} + (s))$, this proves the claim that the formula β is equivalent over F to the disjunction stated above. □

Next we show that minimal formulas all have a very convenient structure.

Proposition 6.4.4. *If $\alpha(X) = \vee_{i=1}^{r} \beta_i(X)$ is the disjunction of basic rankable formulas which is L-minimal for some field $L \subseteq \overline{\mathbb{Q}}$, then each $\beta_i(X)$ has one of the following forms:*

(i) The quantifier-free formula $X = z_0$ for a fixed $z_0 \in L$.

(ii) The "hypersurface formula", $\exists Y_1 \ldots \exists Y_e \ [f(X, Y_1, \ldots, Y_e) = 0 \neq g(X, Y_1, \ldots, Y_e)]$ for an irreducible $f \in L[X, Y_1, \ldots, Y_e]$ and a polynomial $g \in L[X, Y_1, \ldots, Y_e]$.

Proof. Let $\beta(X)$ be a fixed $\beta_i(X)$ which does not have the desired form. Write β in the form

$$\exists Y_1 \cdots \exists Y_m \ [f_1(X, \ldots, Y_m) = \cdots = f_k(X, \ldots, Y_m) = 0 \neq g(X, \vec{Y})]$$

and consider the ideal $I = (f_1, \ldots, f_k)$. Define $e = \dim V(I) \cap D(g)$. Without loss of generality, we can assume that each f_i is irreducible. Otherwise, if $f_1 = h_1 h_2$ is a nontrivial factorization, then we could write β as the disjunction of two formulas with f_1 replaced by h_1 or h_2, respectively, which have smaller multidegree.

Since f_1 is irreducible, $V_L(I)$ is a closed subset of the integral affine variety $V_L(f_1)$ which has dimension $\dim V_L(f_1) = m$ by Proposition 6.2.6. In fact, we see that either $V(f_1) = V(I)$, in which case we are done, or we have

$$e = \dim V(I) \cap D(g) \leq \dim V(I) < \dim V(f_1) = m.$$

By assumption, we are in the latter case, and we will produce a set of formulas with parameters in L which explicitly contradicts the minimality of α.

The ideal I has a primary decomposition $I = \mathfrak{q}_1 \cap \cdots \cap \mathfrak{q}_r$ where each $\mathfrak{q}_i$ is a primary ideal associated to a prime ideal $\mathfrak{p}_i$. Indeed, the rational points on $V(I) \cap D(g)$ are the same as the rational points on $\cup_{i=1}^{r} V(\mathfrak{p}_i) \cap D(g)$. Notice that the open set $V(\mathfrak{p}_i) \cap D(g)$ might be empty for some i, but whenever it is nonempty, $V(\mathfrak{p}_i) \cap D(g)$ has the same dimension as $V(\mathfrak{p}_i)$ by Proposition 6.2.5.

To summarize, we have shown that the formula $\beta(X)$ is equivalent to the disjunction $\vee_{i=1}^{r} \delta_{\mathfrak{p}_i}(X)$ where each $\delta_{\mathfrak{p}_i}(X)$ is defined as a formula

$$\delta_{\mathfrak{p}_i}(X) = \exists Y_1, \ldots, Y_m [p_1^i(X, \vec{Y}) = \cdots = p_{n(i)}^i(X, \vec{Y}) = 0 \neq g(X, \vec{Y})],$$

where $\mathfrak{p}_i = (p_1^i, \ldots, p_{n(i)}^i)$. For each i, we will replace $\delta_{\mathfrak{p}_i}(X)$ itself with an equivalent disjunction of basic rankable formulas, each of which has rank strictly smaller than β. By definition, this contradicts the minimality of α, and the proof will be done.

To this end, we analyze the primes $S = \{\mathfrak{p}_1, \ldots, \mathfrak{p}_r\}$ and divide them accordingly. Let S_{finite} be the set of primes $\mathfrak{p} \in S$ such that only finitely many elements of L satisfy $\delta_{\mathfrak{p}}(X)$ in F, and let S_∞ be all other primes of S. Partition $S_\infty = S_{\text{big}} \cup S_{\text{small}}$ where

$$S_{\text{big}} = \{\mathfrak{p} \in S_\infty \mid \dim V(\mathfrak{p}) = e\},$$
$$S_{\text{small}} = \{\mathfrak{p} \in S_\infty \mid \dim V(\mathfrak{p}) < e\}.$$

For any prime $\mathfrak{p} \in S_{\text{finite}}$, let $\{z_1, \ldots, z_n\}$ be the finite set of elements of L which satisfy $\delta_{\mathfrak{p}}(X)$ in L. We may therefore replace $\delta_{\mathfrak{p}}(X)$ with the disjunction of quantifier-free formulas $\vee_{i=1}^n (X - z_i)$. Each of these quantifier-free formulas consisting of a single-variable polynomial of degree 1 has the smallest rank possible for a nontrivial basic rankable formula and $\beta(X)$ has strictly larger rank.

For any $\mathfrak{p} \in S_{\text{small}}$, the formula $\delta_{\mathfrak{p}}(X)$ is already of smaller rank than β. Indeed, the ambient space is the same, and the dimension is strictly smaller by definition.

For any $\mathfrak{p} \in S_{\text{big}}$, letting $\mathfrak{p} = (p_1, \ldots, p_n)$, we apply Proposition 6.4.3 to see that $\exists \vec{Y}[p_1(X, \vec{Y}) = \cdots = p_n(X, \vec{Y}) = 0]$ is equivalent to the disjunction of two formulas

$$\exists Y_1 \cdots \exists Y_e \ [f(X, \ldots, Y_e) = 0 \neq s(X, \vec{Y})],$$
$$\exists Y_1 \cdots \exists Y_m \ [s(X, \ldots, Y_m) = p_1(X, \ldots, Y_m) = \cdots = p_n(X, \ldots, Y_m) = 0],$$

where $f \in L[X, Y_1, \ldots, Y_e]$ is irreducible and $s \in L[X, Y_1, \ldots, Y_{e-1}]$ is not contained in $\mathfrak{p}$. Thus, $\delta_{\mathfrak{p}}(X)$ is equivalent to the disjunction of the following two formulas

$$\exists Y_1 \cdots \exists Y_m \ [f(X, \ldots, Y_e) = 0 \neq g(X, Y_1, \ldots, Y_m) s(X, Y_1, \ldots, Y_{e-1})], \tag{6.1}$$
$$\exists Y_1 \cdots \exists Y_m \ [s(X, \ldots, Y_{e-1}) = p_1(X, \ldots, Y_m) = \cdots = p_n(X, \ldots, Y_m) = 0 \neq g(X, \vec{Y})]. \tag{6.2}$$

By Lemma 6.4.2, we can replace the formula (6.1) with a disjunction of basic rankable formulas, each of which uses only e quantifiers. Since $e < m$, all these formulas have strictly smaller rank than β.

On the other hand, formula (6.2) has m quantifiers just like β, but we claim the associated variety has smaller dimension. Indeed, we see that

$$\dim V(\mathfrak{p} + (s)) \cap D(g) \leq \dim V(\mathfrak{p} + (s)) < \dim V(\mathfrak{p}) = \dim V(\mathfrak{p}) \cap D(g) = e,$$

where the strict inequality follows by Proposition 6.2.6. Therefore this formula also has

strictly smaller rank than β. This completes the proof. $\square$

We can say more about the hypersurface formula appearing in the previous result. First, we present a simple result on elements of the function field of an irreducible hypersurface.

Lemma 6.4.5. *Let $F \subseteq \overline{\mathbb{Q}}$ be a field and $f \in F[X, Y_1, \ldots, Y_m]$ an irreducible polynomial whose degree in Y_m is positive. If $\bar{p}/\bar{q} \in \mathrm{Frac}(F[X, Y_1, \ldots, Y_m]/(f))$, then there are lifts of p and q to $F[X, Y_1, \ldots, Y_m]$ such that $\deg_{Y_m}(q) < \deg_{Y_m}(f)$.*

Proof. Write $f = \sum_{i=0}^{d} b_i Y_m^i$ where $b_i \in F[X, Y_1, \ldots, Y_{m-1}]$. Choose arbitrary lifts $p_0, q_0 \in F[Y_0, \ldots, Y_m]$ of $\bar{p}$ and $\bar{q}$. If $\deg_{Y_m} q_0 < \deg_{Y_m} f$, then we are already done. Otherwise, define $p_1 = b_d p_0$ and $q_1 = b_d q_0$, which define the same fraction in the function field because $b_d, q_0 \notin (f)$. Then the leading coefficient of q_1 is divisible by b_d, so we write it as $h_1 b_d$ for $h_1 \in F[Y_0, \ldots, Y_{m-1}]$. Define $q_2 = q_1 - h_1 Y_m^{(\deg_{Y_m} q_1)-d} f_1$, and notice that $\deg_{Y_m} q_2 < \deg_{Y_m} q_1$. Continuing in this way, the claim follows. $\square$

Proposition 6.4.6. *Suppose $L \subseteq \overline{\mathbb{Q}}$ is a field and $\beta(X)$ is a formula with coefficients from L of the following form*

$$\beta(X) = \exists Y_1 \ldots \exists Y_e [f(X, Y_1, \ldots, Y_e) = 0 \neq g(X, Y_1, \ldots, Y_e)]$$

Suppose $\beta(X)$ is L-minimal. Then f is absolutely irreducible.

Proof. First, it is clear that f is irreducible in L; if it were reducible then β could be equivalently expressed as the disjunction of two hypersurface formulas of strictly smaller rank.

Suppose for contradiction that f is not absolutely irreducible. We will use this fact to define $\{x \in L : \beta(x) \text{ holds in } L\}$ by a smaller rank formula using coefficients from L.

Let $F \subseteq L$ be a number field containing all the coefficients which appear anywhere in β. Let K be a finite Galois extension of F containing the coefficients of the absolutely irreducible factors of f over $\overline{\mathbb{Q}}$, and let $F' = K \cap L$. Then $F \subseteq F' \subseteq K$, and K is Galois over F' because it was Galois over F. We remark that F' is a subfield of L, and therefore f is irreducible over F'.

For each of the finitely many number fields E with $F' \subset E \subseteq K$, let $p_E \in F'[Z]$ be a minimal polynomial for a primitive generator of E over F'. Since K is Galois over F', none of these finitely many p_E have a root in L. Let $h = \prod_{E : F' \subset E \subseteq K} p_E$.

We claim that L has a lower-ranked formula ϕ with coefficients from F' and with the property that for all $x \in L$, $\phi(x)$ holds over L if and only if $\beta(x)$ does.

Let M be the function field of f over F'. By Proposition 6.2.2, M therefore contains some element $z_0 \in K \setminus F'$. Moreover, $F'(z_0)$ is a subfield of K which strictly contains F'. So $F'(z_0)$ contains a root z of h.

As an element of M, the root z will be of the form $\frac{p(X,\vec{Y})+(f)}{q(X,\vec{Y})+(f)}$, with $p, q \in F'[X, \vec{Y}]$. We may view p and q as polynomials $p, q \in F'[X, Y_1, \ldots, Y_e]$, modulo the ideal (f). These polynomials will satisfy

$$h\left(\frac{p(x,\vec{y})}{q(x,\vec{y})}\right) = 0$$

whenever $(x, \vec{y})$ is a solution to $f = 0$ and $q(x, \vec{y}) \neq 0$. Therefore, every solution $(x, \vec{y}) \in L^{m+1}$ to $f = 0$ has $q(x, \vec{y}) = 0$.

By Lemma 6.4.5, we may choose our specific $q \in F'[X, \vec{Y}]$ so that $\deg_{Y_e}(q) < \deg_{Y_e}(f)$. Notice that $q \notin (f)$ because $q + (f)$ is the denominator of an element of the function field, hence nonzero. Below we will consider q as a polynomial of degree d in Y_e, writing $q = \sum_{i \leq d} c_i Y_e^i$ with all $c_i \in F[X, Y_1, \ldots, Y_{e-1}]$. Without loss of generality, the leading nonzero coefficient c_d does not lie in (f). If it happens that Y_e does not appear in q, then $d = 0$ and $c_0 = q$.

But now we can use these facts to give a lower-ranked disjunction $\phi(X) = \gamma_0(X) \vee \gamma_1(X)$ which is equivalent to $\beta(X)$ in L. Since Y_e has lower degree in q than in f, the trick is to use the Euclidean algorithm here, using the leading term in the expansion $f = \sum_{i=0}^{d_1} Y_e^i \cdot b_i(X, Y_1, \ldots, Y_{e-1})$ and writing

$$r(X, \vec{Y}) = c_d(X, \ldots, Y_{e-1}) \cdot f(X, \vec{Y}) - b_{d_1}(X, Y_1, \ldots, Y_{e-1}) \cdot Y_e^{d_1 - d} \cdot q(X, \vec{Y})$$

as a remainder with $\deg_{Y_e}(r) < \deg_{Y_e}(q)$. Recall that the polynomial c_d is the coefficient of Y_e^d in q, hence does not involve Y_e. Observe also that all coefficients of r are in F'.

We claim that in this situation, a tuple $(x, \vec{y}) \in L^{m+1}$ is a point on $V(f) \cap D(g)$ if and only if one of the following conditions holds:

$$q(x, \vec{y}) = r(x, \vec{y}) = 0 \neq g(x, \vec{y}) \cdot c_d(x, y_1, \ldots, y_{e-1}) \tag{6.3}$$

or

$$f(x, \vec{y}) = c_d(x, y_1, \ldots, y_{e-1}) = 0 \neq g(x, \vec{y}). \tag{6.4}$$

To see the claim, first let $(x, \vec{y})$ be a point on $V(f) \cap D(g)$. As shown above, we must have $q(x, \vec{y}) = 0$. But the Euclidean equation shows that $r(x, \vec{y}) = 0$ as well, so the tuple satisfies one of the conditions, according to whether $c_d(x, y_1, \ldots, y_{m-1}) = 0$ or not. The converse of the claim follows by applying the Euclidean equation to the first condition,

and the latter condition directly defines a subset of $V(f) \cap D(g)$.

The formulas $\gamma_0(X)$ and $\gamma_1(X)$ that we promised above are simply the conditions in (6.3) and (6.4), each prefixed by $\exists Y_1 \cdots \exists Y_e$. Clearly these formulas have the same number of quantifiers as β. The first formula corresponds to a subset $V(r, q) \cap D(gc_d)$ of $V(f) \cap D(g)$ because $r + b_{d_1} Y_e^{d_1 - d} q = c_d f$. Hence the dimension of the subset cannot exceed the dimension of $V(f) \cap D(g)$. However, r and q were constructed to have lower multidegree than f, so γ_0 has strictly smaller rank than β.

On the other hand, the affine variety over F' defined by the latter formula is a proper closed subset of $V(f)$, hence

$$\dim V(f, c_d)) \cap D(g) \leq \dim V(f, c_d)) < \dim V(f) = \dim V(f)) \cap D(g)$$

showing that γ_1 has strictly smaller rank than β. □

Putting these results together yields the following normal form theorem for existential formulas in algebraic extensions of $\mathbb{Q}$.

Definition 6.4.7. *An* absolutely irreducible hypersurface formula *is a formula of the form*

$$\exists Y_1 \ldots \exists Y_e \ [f(X, Y_1, \ldots, Y_e) = 0 \neq g(X, Y_1, \ldots, Y_e)]$$

for polynomials $f, g \in \overline{\mathbb{Q}}[X, Y_1, \ldots, Y_e]$, where f is absolutely irreducible and does not divide g.

Theorem 6.4.8 (Normal Form for Existential Definitions)**.** *For any field $L \subseteq \overline{\mathbb{Q}}$, if $A \subseteq L$ is existentially definable in L, then A is definable in L by a formula of the form*

$$\alpha(X) = \vee_{i=1}^{r} \beta_i(X),$$

where each $\beta_i(X)$ has one of the following forms:

(i) The quantifier-free formula $X = z_0$ for a fixed $z_0 \in L$.

(ii) An absolutely irreducible hypersurface formula with coefficients from L which is satisfied by infinitely many $x \in L$.

Proof. Apply Propositions 6.4.4 and 6.4.6, plus the following two observations. If f divides g in any of the hypersurface formulas, then that formula is unsatisfiable. If a hypersurface formula is satisfied by at most finitely many $x \in L$ (including if it is unsatisfiable), then it could be replaced by a (possibly empty) disjunction of formulas of the form $X = z_0$, lowering the rank. □

6.5 The meagerness of definability

Recall that by identifying a subset of $\overline{\mathbb{Q}}$ with its characteristic function, we can consider the set $\mathrm{Sub}(\overline{\mathbb{Q}}) = \{L \subseteq \overline{\mathbb{Q}} : L \text{ is a field}\}$ as a subset of $2^{\overline{\mathbb{Q}}}$, from which it inherits the product topology. A basis for the topology is given by the sets

$$U_{\vec{a},\vec{b}} = \{L \in \mathrm{Sub}(\overline{\mathbb{Q}}) : a_1, \dots a_n \in L \text{ and } b_1, \dots, b_k \notin L\}$$

for any finite sequences of elements $\vec{a}, \vec{b}$ from $\overline{\mathbb{Q}}$. If $\vec{b}$ is empty, we write simply $U_{\vec{a}}$.

Recall that *Cantor space*, denoted 2^ω, is the set of infinite binary sequences with the product topology.

Proposition 6.5.1. *The space* $\mathrm{Sub}(\overline{\mathbb{Q}})$ *is homeomorphic to Cantor space.*

Proof. Since $\mathrm{Sub}(\overline{\mathbb{Q}})$ is a closed subset of the Cantor-homeomorphic space $2^{\overline{\mathbb{Q}}}$, it suffices to show that $\mathrm{Sub}(\overline{\mathbb{Q}})$ has no isolated points. But it is clear that whenever $U_{\vec{a},\vec{b}}$ is non-empty, there is $c \in \overline{\mathbb{Q}}$ such that both $U_{(\vec{a},c),\vec{b}}$ and $U_{\vec{a},(\vec{b},c)}$ are nonempty. □

The upshot of Proposition 6.5.1 is a structure on the set $\mathrm{Sub}(\overline{\mathbb{Q}})$ which allows us to describe when a set is "large" or "small" in terms of topology. In particular, we enlist the notions of meager sets and the property of Baire.

Definition 6.5.2. *A subset of a topological space is called* nowhere dense *if its closure has empty interior, and* meager *if it is the countable union of nowhere dense sets. A topological space is* Baire[1] *if every non-empty open subset is non-meager.*

Cantor space 2^ω is Baire, and by Proposition 6.5.1 the same is true for $\mathrm{Sub}(\overline{\mathbb{Q}})$, which allows us to consider meager sets to be small.

Definition 6.5.3. *For any $Z \subseteq \mathbb{Q}$, and formula $\beta(X)$ with coefficients $\vec{a}$ from $\overline{\mathbb{Q}}$, we define $S_\beta(Z)$ to be the set of algebraic fields in which β defines a subset of Z in $\mathbb{Q}$:*

$$S_\beta(Z) = \{L \in U_{\vec{a}} : \{x \in \mathbb{Q} : \beta(x) \text{ holds over } L\} \subseteq Z\}.$$

Proposition 6.5.4. *Let Z be a subset of $\mathbb{Q}$ such that $\mathbb{Q} \setminus Z$ is not thin in $\mathbb{Q}$. Then for every absolutely irreducible hypersurface formula*

$$\beta(X) = \exists \vec{Y}[f(X, \vec{Y}) = 0 \neq g(X, \vec{Y})]$$

[1]Some authors use the terminology *Baire space* to refer to topological spaces with this property. However, we reserve the name Baire space for the particular topological space ω^ω, which is discussed in related papers, such as [69], although we will not use it in this chapter.

with coefficients $\vec{a}$ from $\overline{\mathbb{Q}}$, the set $S_\beta(Z)$ is nowhere dense in $U_{\vec{a}}$.

Proof. Let $\vec{b}$ and $\vec{c}$ be any sequences of elements of $\overline{\mathbb{Q}}$ such that $U_{(\vec{a},\vec{c}),\vec{b}} \neq \emptyset$. Let $F = \mathbb{Q}(\vec{a},\vec{c})$ and let $K = F(\vec{b})$. By the application of Hilbert's Irreducibility Theorem in Proposition 6.2.13, there is a thin set $T \subseteq K^e$ such that for any $(x, y_1, \ldots, y_{e-1}) \in K^e \setminus T$, the polynomial $f(x, y_1, \ldots, y_{e-1}, Y_e)$ is irreducible of degree $\deg_{Y_e}(f)$, and $g(x, y_1, \ldots, y_{e-1}, Y_e)$ is not divisible by f. Because K is a number field, $T_{\mathbb{Q}} = T \cap \mathbb{Q}^e$ is also a thin set in $\mathbb{Q}^e$ by Proposition 6.2.10. Further, since $\mathbb{Q} \setminus Z$ is not thin in $\mathbb{Q}$, the thin set $T_{\mathbb{Q}}$ does not contain all of $(\mathbb{Q} \setminus Z) \times \mathbb{Q}^{e-1}$ by Lemma 6.2.12. For any such tuple $(x, y_1, \ldots, y_{e-1}) \in (\mathbb{Q} \setminus Z) \times \mathbb{Q}^{e-1}$ outside this thin set, the irreducibility of $f(x, y_1, \ldots, y_{e-1}, Y)$ over K implies that adjoining to F any root y of $f(x, y_1, \ldots, y_{e-1}, Y)$ will not generate any element of K: we will have $F(y) \cap K = F$, by Lemma 6.2.1. Thus $U_{(\vec{a},\vec{c},y),\vec{b}}$ is nonempty. Additionally, the divisibility condition implies that $g(x, y_1, \ldots, y_{m-1}, y) \neq 0$. Therefore, $U_{(\vec{a},\vec{c},y),b} \cap S_\beta(Z) = \emptyset$. □

Theorem 6.5.5. *The set of all fields $L \in \mathrm{Sub}(\overline{\mathbb{Q}})$ such that $\mathcal{O}_L$ is either existentially or universally definable in L is meager.*

Proof. By Proposition 6.2.11, neither $\mathbb{Z}$ nor $\mathbb{Q} \setminus \mathbb{Z}$ is thin in $\mathbb{Q}$. Therefore, Proposition 6.5.4 shows that for all absolutely irreducible hypersurface formulas β with coefficients from $\overline{\mathbb{Q}}$, the sets $S_\beta(\mathbb{Z})$ and $S_\beta(\mathbb{Q} \setminus \mathbb{Z})$ are nowhere dense. Let

$$S = \bigcup_\beta (S_\beta(\mathbb{Z}) \cup S_\beta(\mathbb{Q} \setminus \mathbb{Z})). \tag{6.5}$$

This is a countable union of nowhere dense sets, and is thus meager. We claim that if $L \in \mathrm{Sub}(\overline{\mathbb{Q}}) \setminus S$, then neither $\mathcal{O}_L$ nor $L \setminus \mathcal{O}_L$ are existentially definable in L.

If $\mathcal{O}_L$ is existentially definable in L, then it is definable in L by a formula $\alpha = \vee_{i<r} \beta_i$ in normal form, according to Theorem 6.4.8. This formula also defines $\mathbb{Z} = \mathbb{Q} \cap \mathcal{O}_L$ inside $\mathbb{Q}$ over L. Because $\mathcal{O}_L$ is infinite and r is finite, some β_i must be a hypersurface formula. Since $L \notin S_{\beta_i}$, there is some $x \in \mathbb{Q} \setminus \mathbb{Z}$ for which $\beta_i(x)$, and therefore also $\alpha(x)$, holds over L, contradicting that α defines $\mathcal{O}_L$ in L.

The same argument applied to $L \setminus \mathcal{O}_L$ shows that this set cannot be existentially definable in L, and thus $\mathcal{O}_L$ is not universally definable in L. □

The proof has shown a slightly stronger consequence.

Porism 6.5.6. *Let $L \in \mathrm{Sub}(\overline{\mathbb{Q}})$ and suppose $A \subseteq L$ is infinite and either existentially or universally definable in L. If either $A \cap \mathbb{Q} \subseteq \mathbb{Z}$ or $A \cap \mathbb{Q} \subseteq \mathbb{Q} \setminus \mathbb{Z}$, then L must lie in the meager set S defined in (6.5) above.*

Corollary 6.5.7. *The set of fields $L \in \mathrm{Sub}(\overline{\mathbb{Q}})$ such that $\mathbb{Z}$ itself is either existentially or universally definable in L is meager.*

We can also use the same approach when considering the definability of subfields. The set $\mathbb{Q}$ is co-thin in itself, so Proposition 6.5.4 does not apply to it. Nevertheless, Porism 6.5.6 yields a further result about the definability of $\mathbb{Q}$ in algebraic field extensions of itself, and more generally about the definability of number fields.

Corollary 6.5.8. *If F is a number field, then the set of fields $L \in \mathrm{Sub}(\overline{\mathbb{Q}})$ containing F such that F has an existential definition in L is a meager set.*

Proof. By Park's generalization [73] of a theorem of Koenigsmann [55], there is a quantifier-free formula $\phi(X, Y_1, \ldots, Y_n)$, in the language of fields, such that $\exists \vec{Y} \phi(X, \vec{Y})$ defines the algebraic non-integers $F \setminus \mathcal{O}_F$ in the field F. In particular, it defines $\mathbb{Q} \setminus \mathbb{Z}$ in $\mathbb{Q}$ over F. Now if $\gamma(Y)$ is existential and defines F in L, then the following formula with free variable X,

$$\gamma(X) \ \& \ \exists \vec{Y} [\ \gamma(Y_1) \ \& \cdots \& \ \gamma(Y_n) \ \& \ \phi(X, \vec{Y}) \ ],$$

is an existential definition of $(\mathbb{Q} \setminus \mathbb{Z})$ in $\mathbb{Q}$ over L. By Porism 6.5.6, $L \in S$. □

Our initial interest in these questions involved definability of rings of integers in fields, but the proof of Proposition 6.5.4 never used the fact that $\mathbb{Z}$ is a ring. Therefore, we can apply the same proof to arbitrary subsets Z of $\mathbb{Q}$. Recall that a formula $\phi(X)$ *defines Z within* $\mathbb{Q}$ in a field extension L of $\mathbb{Q}$ if

$$(\forall x \in \mathbb{Q}) \ [x \in Z \iff \phi(x) \text{ holds in } L].$$

Then our theorem, in its full strength, is as follows.

Theorem 6.5.9. *If Z is any coinfinite subset of $\mathbb{Q}$ that is not thin in $\mathbb{Q}$ (in the sense of Subsection 6.2.3), then the set*

$$\{L \in \mathrm{Sub}(\overline{\mathbb{Q}}) : Z \text{ is } \forall\text{-definable within } \mathbb{Q} \text{ in } L\}$$

is a meager set. Dually, if Z is infinite and $(\mathbb{Q} \setminus Z)$ is not thin in $\mathbb{Q}$, then

$$\{L \in \mathrm{Sub}(\overline{\mathbb{Q}}) : Z \text{ is } \exists\text{-definable within } \mathbb{Q} \text{ in } L\}$$

is a meager set.

Proof. This follows directly from Theorem 6.4.8 and Proposition 6.5.4. □

6.5.1 Computable fields

Next we effectivize Theorem 6.5.5 to obtain many computable algebraic extensions of $\mathbb{Q}$ whose algebraic integers are not existentially or universally definable.

Our arguments below will require the decidability of absolute irreducibility. Recall some standard terminology: a computable field E has a *splitting algorithm* if the *splitting set* $S_E = \{f \in E[T] : f \text{ is reducible in } E[T]\}$ is decidable, and has a *root algorithm* if the *root set* $R_E = \{f \in E[T] : f \text{ has a root in } E\}$ is decidable. Notice that these are both stated for single-variable polynomials. The next lemma is a specific case of the fact that splitting algorithms can be extended to more variables.

Lemma 6.5.10. *Fix any computable presentation of $\overline{\mathbb{Q}}$. Then it is decidable which polynomials in $\overline{\mathbb{Q}}[X_1, X_2, \ldots]$ are absolutely irreducible.*

Proof. $\overline{\mathbb{Q}}$ has a splitting algorithm, of course: all polynomials in $\overline{\mathbb{Q}}[T]$ of degree > 1 are reducible. The lemma now follows from another theorem of Kronecker (found in [28, §§ 58-59]), stating that whenever a computable field F has a splitting algorithm and t is transcendental over F (within a larger computable field), the field $F(t)$ also has a splitting algorithm. The irreducible polynomials of $\overline{\mathbb{Q}}[X_1, X_2]$ are precisely the irreducible polynomials of $\overline{\mathbb{Q}}[X_1]$ along with the polynomials which are irreducible in $\overline{\mathbb{Q}}(X_1)[X_2]$ and have no common factor among the coefficients lying in $\overline{\mathbb{Q}}[X_1]$; see [57, Theorem IV.2.3]. Therefore, reducibility is clearly decidable using Kronecker's result. Thus we can decide reducibility in $\overline{\mathbb{Q}}[X_1, X_2]$, and one continues by induction on the number n of variables, noting that the resulting decision procedures are uniform in n. □

Therefore, there is a computable listing $\beta_1, \beta_2, \ldots$ of all absolutely irreducible hypersurface formulas. Furthermore, we have the following effective version of Proposition 6.5.4.

Proposition 6.5.11. *Let Z be a computable subset of $\mathbb{Q}$ such that $\mathbb{Q} \setminus Z$ is not thin in $\mathbb{Q}$. Then there is an algorithm which, given any absolutely irreducible hypersurface formula β with coefficients $\vec{a}$, and any $\vec{c}, \vec{b}$ such that $U_{(\vec{a},\vec{c}),\vec{b}} \neq \emptyset$, returns y such that $U_{(\vec{a},\vec{c},y),\vec{b}}$ is non-empty and has empty intersection with $S_\beta(Z)$.*

Proof. The proof of Proposition 6.5.4 shows that there is a tuple $(x, y_1, \ldots, y_{e-1}, y) \in (\mathbb{Q} \setminus Z) \times \mathbb{Q}^{e-1} \times \overline{\mathbb{Q}}$ which witnesses that $\beta(x)$ holds in any field extending $\mathbb{Q}(y)$ while keeping $U_{(\vec{a},\vec{c},y),\vec{b}}$ non-empty. So an algorithm can search all such $x, y_1, \ldots, y_{e-1}, y$ until it finds one. This works because Z is computable, and it is computable to check whether a

given tuple from $\overline{\mathbb{Q}}$ satisfies the polynomials appearing in β, and computable to check whether $U_{(\vec{a},\vec{c},y),\vec{b}}$ is empty. □

Theorem 6.5.12. *For every pair of $\overline{\mathbb{Q}}$-tuples $\vec{a}, \vec{b}$, there is a computable $L \in U_{\vec{a},\vec{b}}$ such that $\mathcal{O}_L$ is neither existentially or universally definable in L. Moreover, every computable presentation of L has a splitting algorithm.*

Proof. We define sequences $\vec{a} = \vec{a}_0, \vec{a}_1, \ldots$ and $\vec{b} = \vec{b}_0, \vec{b}_1, \ldots$ in stages as follows. Recall that $\beta_1, \beta_2, \ldots$ is a computable listing of all absolutely irreducible hypersurface formulas. Let $c_1, c_2, \ldots$ be a computable listing of all elements of $\overline{\mathbb{Q}}$.

At stages of the form $s = 3t + 1$, given $U_{\vec{a}_{s-1},\vec{b}_{s-1}}$ nonempty, use Proposition 6.5.11 to find a y such that $U_{(\vec{a}_{s-1},y),\vec{b}_{s-1}}$ is non-empty and disjoint from $S_{\beta_t}(Z)$. Let $\vec{a}_s = (\vec{a}_{s-1}, y)$ and $\vec{b}_s = \vec{b}_{s-1}$.

At stages of the form $s = 3t + 2$, use an analogous process to avoid $S_{\beta_t}(\mathbb{Q} \setminus Z)$.

At stages of the form $s = 3t + 3$, consider $U_{(\vec{a}_{s-1},c_t),\vec{b}_{s-1}}$ and if it is nonempty, set $\vec{a}_s = (\vec{a}_{s-1}, c_t), \vec{b}_s = \vec{b}_{s-1}$. Otherwise, set $\vec{a}_s = \vec{a}_{s-1}$ and $\vec{b}_s = (\vec{b}_{s-1}, c_t)$.

Let $E = \{a \in \overline{\mathbb{Q}} : a \text{ appears in some } \vec{a}_s\}$. Then E is computable because by stage $3(t + 1)$ it has been decided whether c_t is included. And E does not have any existential or universal definition of $\mathcal{O}_E$ because by construction E avoids the set S from Theorem 6.5.5.

The splitting algorithm for L follows from Rabin's Theorem (see [78]), since L is given as a decidable subfield of (our computable presentation of) $\overline{\mathbb{Q}}$. Finally, whenever $L \cong \widetilde{L}$ are computable algebraic fields, their splitting sets are Turing-equivalent, so all computable presentations of L have splitting algorithms. □

6.5.2 The topological space of algebraic extensions of $\mathbb{Q}$ up to isomorphism

The questions of definability we have considered have the same answer over isomorphic fields. Although $\mathrm{Sub}(\overline{\mathbb{Q}})$ contains at least one isomorphic copy of every possible algebraic extension of $\mathbb{Q}$, it contains exactly one copy of an algebraic extension L of $\mathbb{Q}$ if and only if L is Galois over $\mathbb{Q}$. A number field F of degree n is isomorphic to at most n fields in $\mathrm{Sub}(\overline{\mathbb{Q}})$, but there are some infinite non-Galois extensions of $\mathbb{Q}$ which are isomorphic to uncountably many elements in $\mathrm{Sub}(\overline{\mathbb{Q}})$. Therefore, given the isomorphism invariance of the property under consideration, one might wonder if the results of the previous section have been skewed by the fact that some isomorphism classes are more represented in $\mathrm{Sub}(\overline{\mathbb{Q}})$ than others.

Thus it is also of interest to consider the collection of algebraic extensions of $\mathbb{Q}$ up to isomorphism as a topological space, as was done in [69]. We denote this set by $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$. From the perspective of number theory, the set $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$ can be identified as a quotient of $\mathrm{Sub}(\overline{\mathbb{Q}})$ by the absolute Galois group $G = \mathrm{Gal}\,\overline{\mathbb{Q}}\mathbb{Q}$, which equates isomorphic fields. The topology on $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$ is the quotient topology which it inherits from $\mathrm{Sub}(\overline{\mathbb{Q}})$.

Alternatively, from the perspective of computability theory, one could begin with the space $\mathcal{ALG}_0^*$ of all possible presentations of algebraic extensions of $\mathbb{Q}$ in a certain language. This is done in [69] and the relevant language in this case is the language of rings enlarged to include additional predicates for the existence of roots of monic one-variable polynomials. Equating isomorphic fields and taking the quotient topology leads to the space $\mathcal{ALG}_0^*/\cong$, which coincides with $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$ despite various differences between $\mathcal{ALG}_0^*$ and $\mathrm{Sub}(\overline{\mathbb{Q}})$. For example, in $\mathcal{ALG}_0^*$, every isomorphism class is represented with uncountably many copies. For details about $\mathcal{ALG}_0^*$, we refer the reader to [69].

Returning now to $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$, observe that for any $U_{\vec{a},\vec{b}}$, the following set is the smallest G-invariant subset of $\mathrm{Sub}(\overline{\mathbb{Q}})$ containing $U_{\vec{a},\vec{b}}$. It is also clopen, as there are only finitely many images $\phi(\vec{a}), \phi(\vec{b})$.

$$GU_{\vec{a},\vec{b}} := \{\phi(L) : L \in U_{\vec{a},\vec{b}}, \phi \in G\} = \bigcup_{\phi \in G} U_{\phi(\vec{a}),\phi(\vec{b})}$$

It follows that the quotient map $q : \mathrm{Sub}(\overline{\mathbb{Q}}) \to \mathrm{Sub}(\overline{\mathbb{Q}})/\cong$ is open and the images of the sets $GU_{\vec{a},\vec{b}}$ form a clopen basis for $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$.

Proposition 6.5.13 (Theorem 3.3, [69])**.** $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$ *is homeomorphic to Cantor space.*

Proof. The follows because $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$ is compact, has a countable clopen basis, and has no isolated points. The last condition follows because any non-empty $GU_{\vec{a},\vec{b}}$ contains at least two non-isomorphic fields. □

Therefore, notions of meager and co-meager make sense in $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$. In order to transfer the all our results about $\mathrm{Sub}(\overline{\mathbb{Q}})$ to results about $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$, we only need to check the following.

Proposition 6.5.14. *Let* $S \subseteq \mathrm{Sub}(\overline{\mathbb{Q}})$ *be as defined in (6.5), and let* $q : \mathrm{Sub}(\overline{\mathbb{Q}}) \to \mathrm{Sub}(\overline{\mathbb{Q}})/\cong$ *be the quotient map. Then* $q(S)$ *is meager in* $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$.

Proof. Observe that $S = \cup_\beta GS_\beta$, where β ranges over the absolutely irreducible hypersurface formulas and

$$GS_\beta := \bigcup_{\phi \in G} S_{\phi(\beta)}.$$

Here $\phi(\beta)$ denotes the result of applying ϕ to all coefficients appearing in β. There are only finitely many possible outcomes, so GS_β is a finite union of nowhere dense sets, and thus is nowhere dense. Additionally, since each S_β is closed, so is GS_β. Also, by construction GS_β is G-invariant. Since $\mathrm{Sub}(\overline{\mathbb{Q}}) \setminus GS_\beta$ is dense open and q is an open map, its image $q(\mathrm{Sub}(\overline{\mathbb{Q}}) \setminus GS_\beta)$ is dense open. Therefore, by G-invariance of GS_β, $q(GS_\beta)$ is nowhere dense. Thus $q(S) = \cup_\beta q(GS_\beta)$ is meager. □

Therefore, we have the following analogues of the results of the previous section.

Corollary 6.5.15. *The following sets are meager in* $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$*:*

1. *The set of isomorphism types of fields L in which $\mathcal{O}_L$ is existentially or universally definable.*

2. *The set of isomorphism types of fields in which $\mathbb{Z}$ is existentially or universally definable.*

3. *The set of isomorphism types of fields L in which some number field $F \subset L$ is existentially definable.*

Proof. These sets are all contained in $q(S)$. □

It may seem equally natural to consider the Lebesgue measure on Cantor space and transfer it to $\mathrm{Sub}(\overline{\mathbb{Q}})/\cong$, using some computable homeomorphism such as that obtained in [69, Theorem 3.3]. This is attempted to some extent in [69], but the resulting measure is not canonical: it depends to a great extent on arbitrary choices that are made during the construction of the homeomorphism. Indeed, the notion of *Haar-compatible measure*, put forth in [69], has had to be abandoned, as the reality is more complicated than the analysis in that article recognized. We hope to investigate this situation, and measure-theoretic perspectives in general, more fully in the near future.

Chapter 7
Undecidability and Unit Groups

7.1 Main result and comparison to prior work

In this chapter, we are concerned with the question of which infinite algebraic extensions of $\mathbb{Q}$ have undecidable first-order theory. There are many results which demonstrate that some infinite algebraic extensions of $\mathbb{Q}$ have undecidable first-order theory, while others do not. Our framing of this question comes from the work of Shalpentokh [91], whose definability results play a crucial role in the proof of the main result of this chapter.

To start, Shlapentokh proved that every abelian extension of $\mathbb{Q}$ which is ramified at only finitely many primes is undecidable [91, Theorem 5.5], generalizing a result of Videla [108]. On the other hand, recall that an algebraic number is said to be totally real if its minimal polynomial has only real roots. Fried, Haran and Völklein proved that the field $\mathbb{Q}^{\mathrm{tr}}$ of all totally real numbers is decidable [35].

One way to prove that a field K is undecidable is to show that the ring of integers $\mathcal{O}_K$ is undecidable and definable in K. For example, this method proves that the field $K = \mathbb{Q}(\{\sqrt{p} : p \text{ prime}\})$ is undecidable. In this case, the definability of $\mathcal{O}_K$ in K was shown by Videla [107, §5.4], while the undecidability of $\mathcal{O}_K$ was proven by Julia Robinson [82] and follows from a general "blueprint" that she described. The blueprint, in a more general form due to C.W. Henson [103, §3.3], states that a ring of integers is undecidable if there is a definable family of subsets which contains finite sets of arbitrarily large cardinality. As an application of this blueprint, Julia Robinson showed how such families of sets could be constructed for rings of totally real numbers by using totally positive elements. Vidaux and Videla expanded on her ideas to prove the undecidability of a large class of rings of integers $\mathcal{O}_K$ in totally real fields K [104, 105]. When combined with definability results, such as those of Fukuzaki [36], Shlapentokh [91], and Videla [107], this proves the undecidability of many totally real fields.

Recently, Martínez-Ranero, Utreras and Videla [65] leveraged these methods, which were developed for totally real fields, to instead prove the undecidability of the totally imaginary field $\mathbb{Q}^{(2)}$, the compositum of all degree 2 extensions of $\mathbb{Q}$. The key to their strategy was noticing that $\mathbb{Q}^{(2)}$ is a totally imaginary quadratic extension of the totally real field $K = \mathbb{Q}(\{\sqrt{p} : p \text{ prime}\})$, and therefore $\mathcal{O}_K^\times$ is a finite-index subgroup of $\mathcal{O}_{\mathbb{Q}^{(2)}}^\times$. Using this fact, they produce a special "large" set W, which is definable in $\mathcal{O}_{\mathbb{Q}^{(2)}}$ and contains only totally real elements. The undecidability of $\mathbb{Q}^{(2)}$ then follows from the aforementioned methods for totally real rings of integers and the definability of $\mathcal{O}_{\mathbb{Q}^{(2)}}$ in $\mathbb{Q}^{(2)}$. The goal of this chapter is to generalize their strategy and produce more examples of totally imaginary infinite extensions of $\mathbb{Q}$ with undecidable first-order theory.

Now we recall the necessary notation and present the main result of the chapter. Given a number field F, let $F^{(d)}$ be the compositum of all extensions of F of degree at most d, and let $F_{ab}^{(d)}$ be the maximal abelian subextension of $F \subseteq F^{(d)}$. The following theorem will be proved as Theorem 7.3.6, and the undecidability of $\mathbb{Q}^{(2)}$ follows as a special case. In Section 7.4, we give additional explicit examples of totally real fields K and families of polynomials $\{f_a(x) \mid a \in \mathbb{Z}_{\geq N_0}\}$ which satisfy the conditions of the theorem.

Theorem 7.3.6. *Let K be an infinite totally real extension of $\mathbb{Q}$ which is contained in $F_{ab}^{(d)}$ for some $d \geq 2$ and some number field F. Assume K contains all roots of a parametrized family of polynomials*

$$\{f_a(x) = x^n + p_{n-1}(a)x^{n-1} + \ldots p_1(a)x + p_0(a) \mid a \in \mathbb{Z}_{>N_0}\}$$

where each $p_i(t) \in \mathbb{Z}[t]$ is a polynomial, $p_0(t) = \pm 1$ is constant and $p_j(t)$ is nonconstant for some $1 \leq j \leq n-1$. If L is any totally imaginary quadratic extension of K, then the first-order theory of L is undecidable.

In this theorem, the ring of integers $\mathcal{O}_K$ is undecidable by a result of Vidaux and Videla [105]. This fact is used to deduce the undecidability of $\mathcal{O}_L$, and therefore the undecidability of L because $\mathcal{O}_L$ is definable in L by a result of Shlapentokh [91]. The strategy, as in the case of $\mathbb{Q}^{(2)}$, is to exploit the unit group $\mathcal{O}_L^\times$ to define a sufficiently large subset W of $\mathcal{O}_L$ which contains only totally real elements. This is done by explicitly using a polynomial whose values are power-sums of the units defined by the family $\{f_a(x)\}$. Producing this special polynomial is the main ingredient required to extend the method for $\mathbb{Q}^{(2)}$ to this more general setting. It would be interesting to investigate whether the the undecidability of a totally real field K implies the undecidability of its

totally imaginary quadratic extensions in a more general context, without the need for the explicit parametrized family of polynomials or the field $F_{ab}^{(d)}$ containing K.

For completeness, we pause to note that the work of Martínez-Ranero, Utreras and Videla described above is not the first time that a problem of decidability for the ring of integers of an infinite algebraic extension has been reduced to a subextension of degree 2. In the earlier work of Shlapentokh [90], a similar reduction argument proves, for certain totally real fields K, the undecidability of Hilbert's Tenth Problem for the ring of integers $\mathcal{O}_L$ of any quadratic extension L of K. In contrast to the methods above which rely on unit groups and the so-called JR-number, Shlapentokh's theorem [90, Theorem 7.9] requires an elliptic curve E defined over a finite extension K' of K such that $E(K')$ is finitely generated with positive rank. There are examples of totally real fields K for which such an elliptic curve exists [90, Example 10.1]. More work is required to determine whether such elliptic curves exist for the fields under consideration in this paper.

7.2 Sufficient Conditions For Undecidability

Throughout this chapter, K will denote a totally real infinite extension of $\mathbb{Q}$, and L a totally imaginary quadratic extension of K. The goal of this section is to find a suitable sufficient condition for when L has undecidable first-order theory. We will begin by reviewing some definability results and recall the methods used to prove undecidability in the totally real case.

Clearly, if $\mathcal{O}_L$ is definable in L and the first-order theory of $\mathcal{O}_L$ is undecidable, then the first-order theory of L is also undecidable. There are many results on the definability of rings of integers in infinite algebraic extensions of $\mathbb{Q}$; see, for example, the work of Fukuzaki [36], Shlapentokh [91], and Videla [107]. The following result of Shlapentokh, presented in [91, Example 4.3], will suffice for our purposes in this chapter.

Theorem 7.2.1. *If L is the compositum of finite extensions of $\mathbb{Q}$ of degree less than some positive integer d, then $\mathcal{O}_L$ is first-order definable in L.*

Now our problem reduces to finding undecidable rings of integers. The following condition, first presented by Julia Robinson [82, Theorem 2] and later generalized by C.W. Henson [103, §3.3], gives a sufficient condition for when a ring of integers has undecidable first-order theory.

Lemma 7.2.2. *Let $\mathcal{O}$ be a ring of algebraic integers. If there is a family $\mathcal{F}$ of subsets of $\mathcal{O}$, parametrized by an $\mathcal{L}_{ring}$-formula, which contains finite sets of arbitrarily large*

cardinality, then $\mathcal{O}$ has undecidable first-order theory.

7.2.1 Totally real fields

For totally real rings of integers, Lemma 7.2.2 enables a concrete method, first used by Julia Robinson [82] and developed further by Vidaux and Videla [104, 105], to prove undecidability. Given a set X of totally real algebraic numbers, define

$$X_t = \{\alpha \in X : 0 \ll \alpha \ll t\}$$

where $0 \ll \alpha \ll t$ means that every conjugate of α lies in the interval $(0, t)$. The *JR-number* of X is

$$\mathrm{JR}(X) = \inf\{t \in \mathbb{R} : \#X_t = \infty\}.$$

When $X = \mathcal{O}_K$ for a totally real field K, the sets X_t are definable in X because every totally positive algebraic number is the sum of four squares by a theorem of Siegel [92], and therefore the following theorem follows from Lemma 7.2.2.

Theorem 7.2.3 ([82]). *If K is a totally real field and the JR-number $\mathrm{JR}(\mathcal{O}_K)$ is either a minimum or infinite, then the first-order theory of $\mathcal{O}_K$ is undecidable.*

Examples where this theorem applies include the ring of all totally real algebraic integers $\mathbb{Z}^{\mathrm{tr}}$, and the ring of integers of $\mathbb{Q}(\{\sqrt{p} : p \text{ prime}\})$. The JR-numbers are 4 and ∞, respectively, in these cases [82]. For many years, there were no known examples of rings of totally real integers whose JR-numbers were finite and either different from 4, or not a minimum. Recently, infinitely many such examples have been constructed by Castillo Fernandez, Vidaux and Videla [16, 17, 104], and by Gillibert and Ranieri [40]. For the purposes of this chapter, we will only be concerned with totally real fields whose JR-number is infinite.

More examples of totally real fields K with $\mathrm{JR}(\mathcal{O}_K) = \infty$ come from a connection to the Northcott property, discovered by Vidaux and Videla [105]. A set X is said to have the *Northcott property* if

$$\{\alpha \in X : \hat{h}(\alpha) < t\}$$

is a finite set for every positive real number t, where $\hat{h}(\alpha)$ denotes the logarithmic Weil height.

Proposition 7.2.4 ([105, Theorem 2]). *If a totally real field K has the Northcott property, then $\mathrm{JR}(\mathcal{O}_K) = \infty$.*

The following theorem of Bombieri and Zannier provides many examples of fields with the Northcott property; see also [19, 112] for more examples. By definition, it suffices to show that the totally real field K is contained in a (possibly imaginary) field which has the Northcott property.

Theorem 7.2.5 ([9, Theorem 1]). *If F is a number field, then $F_{ab}^{(d)}$ has the Northcott property.*

7.2.2 Moving to totally imaginary fields

Recall that the field $\mathbb{Q}^{(2)}$ is a totally imaginary quadratic extension of the totally real field $K = \mathbb{Q}(\{\sqrt{p} : p \text{ prime}\})$. The proof of the undecidability of $\mathbb{Q}^{(2)}$ given by Martínez-Ranero, Utreras and Videla [65] uses the fact that $\mathrm{JR}(\mathcal{O}_K) = \infty$ to deduce that $\mathcal{O}_{\mathbb{Q}^{(2)}}$ is also undecidable. Essentially, they find a special set which is definable in $\mathcal{O}_{\mathbb{Q}^{(2)}}$ and contains only totally real elements, which allows them to apply some of the methods created for the totally real case. Below, we make this explicit by generalizing their strategy to a set of lemmas which provide a sufficient condition for undecidability.

Lemma 7.2.6. *If there is a first-order definable subset $W \subseteq \mathcal{O}_L$ such that $\mathbb{N} \subseteq W \subseteq \mathcal{O}_K$, then the first-order theory of $\mathcal{O}_L$ is undecidable.*

Proof. This result follows quickly from Lemma 7.2.2. Using W, define a family $\mathcal{F}$ of subsets of $\mathcal{O}_L$ parametrized by the formula $\phi_W(x; a, b)$:

$$ax \neq 0 \wedge ax \neq b \wedge \exists x_1, \ldots, x_8 \in W[ax = x_1^2 + \cdots + x_4^2 \wedge (b - ax) = x_5^2 + \cdots + x_8^2]$$

If $a, b \in \mathbb{N}$, then $\phi_W(x; a, b)$ implies $0 \ll ax \ll b$, and $\phi_W(n; a, b)$ holds for every natural number $0 < n < \frac{b}{a}$ by Lagrange's four square theorem. Therefore, we have

$$\left\{n \in \mathbb{N} : 0 < n < \frac{b}{a}\right\} \subseteq \{x \in \mathcal{O}_L : \phi_W(x; a, b)\} \subseteq \left\{x \in \mathcal{O}_K : 0 \ll x \ll \frac{b}{a}\right\}. \tag{7.1}$$

for every $\frac{b}{a} \in \mathbb{Q}_{>0}$. Because $\mathrm{JR}(\mathcal{O}_K) = \infty$, the sets on the righthand side above are finite for all $\frac{b}{a} \in \mathbb{Q}_{>0}$ by definition. Hence, the family of sets $\mathcal{F}$ contains finite sets of arbitrarily large size as $\frac{b}{a} \to \infty$, and therefore Lemma 7.2.2 applies. □

The proof above demonstrates why we restrict to the case of $\mathrm{JR}(\mathcal{O}_K) = \infty$. If we instead have that $\mathrm{JR}(\mathcal{O}_K) < \infty$ is a minimum, then we could attempt to perform a similar trick. However, the sets on the right-hand side of (7.1) are only finite for

$\frac{b}{a} < \mathrm{JR}(\mathcal{O}_K) < \infty$ in this case. Thus we would need to replace the lower bound given by $\mathbb{N}$ on the left-hand side of (7.1) with a larger lower bound to obtain our finite sets of arbitrarily large cardinality. This corresponds to modifying the hypothesis of the lemma to require $S \subseteq W \subseteq \mathcal{O}_K$ for some subset S such that $\#\{x \in S : \phi_S(x; a, b)\} \to \infty$ as $\frac{b}{a}$ approaches $\mathrm{JR}(\mathcal{O}_K)$ from below. There are no obvious candidates for a convenient and useful choice of S, due in part to the difficulties involved in constructing examples of totally real fields K such that $\mathrm{JR}(\mathcal{O}_K)$ is finite. Hence we will only consider the case $\mathrm{JR}(\mathcal{O}_K) = \infty$.

The upshot of the previous lemma is that we simply need to produce the desired set W. In the case of $\mathbb{Q}^{(2)}$, Martínez-Ranero, Utreras and Videla use a discrete derivative trick and Hilbert's solution to Waring's problem to produce the set W; see [65, Lemma 7]. Their strategy also works in our more general setting, but we will instead use a theorem of Kamke which solves a more general conjecture of Waring. In the following section, we will use families of polynomials and the unit group $\mathcal{O}_L^\times$ to produce the polynomial $f(x)$ and definable set X_0 required by the lemma.

Lemma 7.2.7. *Let $f \in \mathbb{Z}[x]$ be a nonconstant polynomial, and let $X_0 \subseteq \mathcal{O}_K$ be a subset which is definable in $\mathcal{O}_L$. If $f(n) \in X_0$ for each sufficiently large natural number $n \geq N_0$, then there is a first-order definable subset W of $\mathcal{O}_L$ such that $\mathbb{N} \subseteq W \subseteq \mathcal{O}_K$.*

Proof. Define $X_1 = \{\pm x : x \in X_0\}$. By replacing $f(x)$ with $\pm f(x + k)$ for some $k \geq N_0$, we can assume that $f(n) \in X_1$ is a nonnegative integer for all integers $n \geq 0$ without loss of generality. Then Kamke's theorem [50] states that there is an integer $r \geq 1$ such that every $m \in \mathbb{N}$ can be written in the form

$$m = f(a_1) + \cdots + f(a_{s_1}) + s_2$$

where $s_1, s_2 \in \mathbb{N}$ satisfy $s_1 + s_2 \leq r$ and $a_1, \ldots, a_{s_1} \in \mathbb{N}$.

Thus, we may simply define

$$W = \bigcup_{s_1=0}^{r} \bigcup_{s_2=0}^{r-s_1} \{x_1 + \cdots + x_{s_1} + s_2 : x_1, \ldots, x_{s_1} \in X_1\}.$$

□

7.3 Using the group of units

To complete the proof of our main theorem, we will exploit the structure of the group of units $\mathcal{O}_L^\times$ to apply Lemma 7.2.7. We begin by recalling two basic facts, which generalize [65, Lemmas 5-6]. Throughout this section, we assume that $\mathrm{JR}(\mathcal{O}_K) = \infty$.

Lemma 7.3.1. *The group of roots of unity $\mu_L \subset \mathcal{O}_L^\times$ is finite.*

Proof. If $\omega \in \mu_L$ is a root of unity, then $2 + \omega + \omega^{-1} \in K$ satisfies $0 \ll 2 + \omega + \omega^{-1} \ll 4$. Because $\mathrm{JR}(\mathcal{O}_K) > 4$, there are only finitely many elements $\alpha \in K$ satisfying $0 \ll \alpha \ll 4$. □

Lemma 7.3.2. *Write $\#\mu_L = 2N$. If $u \in \mathcal{O}_L^\times$, then $u^{2N} \in \mathcal{O}_K^\times$.*

Proof. Use the following notation. Let $K = \cup_{i=0}^\infty K_i$ where $K_0 \subseteq K_1 \subseteq \dots$ is an infinite tower of totally real number fields such that L_0 is a totally imaginary quadratic extension of K_0, $L_n = K_n L_0$ and $L = KL_0$. The lemma then immediately follows from the fact that $[\mathcal{O}_{L_n}^\times : \mu_{L_n}\mathcal{O}_{K_n}^\times] \in \{1, 2\}$ for all $n \geq 0$ [109, Theorem 4.12]. □

The previous lemma implies $(\mathcal{O}_L^\times)^{2N} \subseteq \mathcal{O}_K$ is definable in $\mathcal{O}_L$, which will allow us to produce a subset $X_0 \subseteq \mathcal{O}_K$ which is definable in $\mathcal{O}_L$. Next, we will define a useful multivariable polynomial, then specialize it to a certain single-variable polynomial $f(x)$ which satisfies Lemma 7.2.7.

We will use the following notation. For each $k, n \geq 1$, write $q_k(x_1, \dots, x_n) = x_1^k + \cdots + x_n^k$ for the k-th power-sum polynomial, and let s_k be the k-th elementary symmetric polynomial. The *Newton-Girard formulae* provide the following relation for any $k, n \geq 1$. For ease of notation, we suppress the variables on the right-hand side.

$$q_k(x_1, \dots, x_n) = (-1)^{k-1} k s_k + \sum_{i=1}^{k-1} (-1)^{k+i-1} s_{k-i} q_i.$$

Lemma 7.3.3. *Given any integers $m, n \geq 1$, there is a polynomial*

$$Q_m(x_0, \dots, x_{n-1}) \in \mathbb{Z}[x_0, \dots, x_{n-1}]$$

such that for any $(a_0, \dots, a_{n-1}) \in \mathbb{Z}^n$,

$$Q_m(a_0, \dots, a_{n-1}) = \alpha_1^m + \cdots + \alpha_n^m$$

where $\alpha_1, \dots, \alpha_n$ are the roots of $f(x) = x^n + a_{n-1}x^{n-1} + \cdots + a_0 \in \mathbb{C}[x]$.

Proof. For any polynomial $x^n + a_{n-1}x^{n-1} + \cdots + a_0 \in \mathbb{C}[x]$, the roots $\alpha_1, \ldots, \alpha_n \in \mathbb{C}$ satisfy

$$\begin{aligned} s_1(\alpha_1, \ldots, \alpha_n) &= \alpha_1 + \cdots + \alpha_n = -a_{n-1} \\ &\vdots \\ s_n(\alpha_1, \ldots, \alpha_n) &= \alpha_1 \ldots \alpha_n = (-1)^n a_0. \end{aligned}$$

Therefore, by using the Newton-Girard formulae and induction, we can write $q_m(\alpha_1, \ldots, \alpha_n)$ as a polynomial in $a_{n-k} = (-1)^k s_k(\alpha_1, \ldots, \alpha_n)$ for $1 \leq k \leq n$, as claimed. □

We focus our attention on the values that Q_m takes on the coefficients of particular families of polynomials whose roots are totally real units. Importantly, we need the resulting single-variable polynomial to be nonconstant to apply Lemma 7.2.7.

Lemma 7.3.4. *Let $p_0(t), \ldots, p_{n-1}(t) \in \mathbb{Z}[t]$ be polynomials which parametrize a family of polynomials*

$$\{f_a(x) = x^n + p_{n-1}(a)x^{n-1} + \ldots p_1(a)x + p_0(a) : a \in \mathbb{Z}_{\geq N_0}\}$$

where $p_j(x)$ is nonconstant for some $0 \leq j \leq n-1$. For any $N \geq 1$, there is some $k \geq 1$ such that

$$Q_{kN}(p_0(x), \ldots, p_{n-1}(x))$$

is nonconstant.

Proof. Factor each polynomial $f_a(x) = \prod_{i=1}^{n}(x - \alpha_{i,a})$ over the algebraic closure. First consider the case of $N = 1$. By assumption, there is a smallest index $1 \leq j_0 \leq n$ such that $s_{j_0}(\alpha_{1,a}, \ldots, \alpha_{n,a}) = (-1)^{j_0} p_{n-j_0}(a)$ is nonconstant as a varies. For each $1 \leq k \leq n$, the Newton-Girard formulae

$$q_k = (-1)^{k-1} k s_k + \sum_{i=1}^{k-1} (-1)^{k+i-1} s_{k-i} q_i.$$

can be expanded recursively to write q_k in terms of $s_1, \ldots, s_k$. This implies that

$$Q_k(p_0(a), \ldots, p_{n-1}(a)) = q_k(\alpha_{1,a}, \ldots, \alpha_{n,a})$$

is constant for $1 \leq k < j_0 - 1$, and nonconstant for $k = j_0$.

Now let $N > 1$. We will reduce to the previous case by defining a new family of polynomials $\hat{f}_a(x) = \prod_{i=1}^{n}(x - \alpha_{i,a}^N)$. First, we show that the new family $\{\hat{f}_a(x) : a \in \mathbb{Z}_{\geq N_0}\}$ is also parametrized. For each $1 \leq j \leq n$, the $(n-j)$-th coefficient of $\hat{f}_a(x)$ is equal to $(-1)^j s_j(\alpha_1^N, \ldots, \alpha_n^N)$. Clearly the polynomial $s_j(x_1^N, \ldots, x_n^N)$ is invariant under permutation of variables, so there is a polynomial $g_j(t_1, \ldots, t_n) \in \mathbb{Z}[t_1, \ldots, t_n]$ such that

$$s_j(x_1^N, \ldots, x_n^N) = g_j(s_1(x_1, \ldots, x_n), \ldots, s_n(x_1, \ldots, x_n));$$

see [57, Theorem IV.6.1]. Therefore,

$$\hat{f}_a(x) = x^n + \hat{p}_{n-1}(a)x^{n-1} + \cdots + \hat{p}_0(a)$$

where $\hat{p}_{n-j}(x) = (-1)^j g_j(-p_{n-1}(x), \ldots, (-1)^n p_0(x))$ for each $1 \leq j \leq n$.

This shows that $\{\hat{f}_a(x) : a \in \mathbb{Z}_{\geq N_0}\}$ is a parametrized family of polynomials, so it only remains to see that some $\hat{p}_j(x)$ is nonconstant. If $\hat{p}_j(x)$ is constant for all $1 \leq j \leq n-1$, then the family $\{\hat{f}_a(x) : a \in \mathbb{Z}_{\geq N_0}\}$ contains a single polynomial. But this is clearly impossible by the definition of $\hat{f}_a(x)$ because the family $\{f_a(x) : a \in \mathbb{Z}_{\geq N_0}\}$ is infinite by assumption. Hence applying the base case to the family $\{\hat{f}_a(x) : a \in \mathbb{Z}_{\geq N_0}\}$ completes the proof because

$$Q_{kN}(p_0(a), \ldots, p_{n-1}(a)) = \alpha_{1,a}^{kN} + \cdots + \alpha_{n,a}^{kN} = Q_k(\hat{p}_0(a), \ldots, \hat{p}_{n-1}(a))$$

for all $k \geq 1$ by construction. □

We are now ready to prove the main theorem of this section on the undecidability of rings of integers in totally imaginary fields.

Theorem 7.3.5. *Let K be a totally real field with $\mathrm{JR}(\mathcal{O}_K) = \infty$, and let $p_0(t), \ldots, p_{n-1}(t) \in \mathbb{Z}[t]$ be polynomials which parametrize a family of polynomials*

$$\{f_a(x) = x^n + p_{n-1}(a)x^{n-1} + \ldots p_1(a)x + p_0(a)\}$$

where $p_0(a) = \pm 1$ is constant, and $p_{j_0}(t)$ is nonconstant for some $1 \leq j_0 \leq n-1$. Assume that K contains all roots of $f_a(x)$ for all natural numbers $a \geq N_0$. If L is a totally imaginary quadratic extension of K, then the first-order theory of $\mathcal{O}_L$ is undecidable.

Proof. By Lemma 7.2.7, it suffices to find a nonconstant polynomial f and a definable subset $X_0 \subseteq \mathcal{O}_K$ such that $f(n) \in X_0$ for all sufficiently large $n \in \mathbb{Z}_{\geq N_0}$.

Let $N = \#\mu_L$, as in Lemma 7.3.2. We choose our polynomial to be

$$f(x) = Q_{2Nk}(p_0(x), \ldots, p_{n-1}(x)),$$

where Q_{2Nk} is defined in Lemma 7.3.3 and k is chosen according to Lemma 7.3.4 so that $f(x)$ is nonconstant. By definition, $f(a) = Q_{2N}(p_0(a), \ldots, p_{n-1}(a))$ is the sum of $2N$-th powers of units of $\mathcal{O}_K$ for each $a \geq N_0$, so we may define

$$X_0 = \{\alpha_1^{2N} + \cdots + \alpha_n^{2N} \mid \alpha_i \in \mathcal{O}_L^\times\}.$$

This subset is definable in $\mathcal{O}_L$, contains $f(n)$ for $n \geq N_0$ by assumption, and $X_0 \subseteq \mathcal{O}_K$ by Lemma 7.3.2. □

By using the results discussed in the previous section, this theorem implies the our main theorem on the undecidability of totally imaginary fields.

Theorem 7.3.6. *Let K be an infinite totally real extension of $\mathbb{Q}$ which is contained in $F_{ab}^{(d)}$ for some $d \geq 2$ and some number field F. Assume K contains all roots of a parametrized family of polynomials*

$$\{f_a(x) = x^n + p_{n-1}(a)x^{n-1} + \ldots p_1(a)x + p_0(a) \mid a \in \mathbb{Z}_{>N_0}\}$$

where each $p_i(t) \in \mathbb{Z}[t]$ is a polynomial, $p_0(t) = \pm 1$ is constant and $p_j(t)$ is nonconstant for some $1 \leq j \leq n-1$. If L is any totally imaginary quadratic extension of K, then the first-order theory of L is undecidable.

Proof. Using Proposition 7.2.4 and Theorem 7.2.5, we see that $\mathrm{JR}(\mathcal{O}_K) = \infty$. By Theorem 7.2.1, $\mathcal{O}_L$ is definable in L. Thus the undecidability follows from Theorem 7.3.5. □

7.4 Examples

We will now give some concrete examples of families of polynomials $\{f_a(x) : a \in \mathbb{Z}_{\geq N_0}\}$ which satisfy Theorem 7.3.6. In each case, K can be taken to be the totally real field generated by all roots of the polynomials $\{f_a(x)\}$, or any extension thereof which is contained in $F_{ab}^{(d)}$ for some number field F and some integer $d \geq 1$. Then Theorem 7.3.6 implies that any totally imaginary quadratic extension L of K has undecidable first-order theory.

7.4.1 Polynomials generating cyclic extensions of $\mathbb{Q}$

1. **The quadratic case:** Choosing $f_a(x) = x^2 - 2ax - 1$ produces the family of polynomials considered by Martínez-Ranero, Utreras and Videla to prove the undecidability of $\mathbb{Q}^{(2)}$ [65, Lemma 7] . More generally, one may use the family of polynomials $x^2 - p(a)x - 1$, where $p(t)$ is any nonconstant polynomial. Using these polynomials also shows that the fields $\mathbb{Q}_{ab}^{(d)}$ are undecidable for all $d \geq 2$, since each field is a totally imaginary quadratic extension of a totally real field.

2. **The cubic case:** Shanks describes some "simplest cubic extensions" [86] which are totally real and generated by roots of polynomials of the form

$$x^3 - ax^2 - (a+3)x - 1$$

for $a \geq -1$. Similarly, Kishi [52] gives the family of polynomials

$$x^3 - n(n^2+n+3)(n^2+2)x^2 - (n^3+2n^2+3n+3)x - 1$$

for $n \in \mathbb{Z}$. Each polynomial in both families generates a totally real cyclic cubic extension of $\mathbb{Q}$.

3. **The quartic case:** For $t \geq 4$, the following polynomials, constructed by Gras [41, Proposition 6], generate cyclic quartic totally real extensions of $\mathbb{Q}$.

$$x^4 - tx^3 - 6x^2 + tx + 1.$$

4. **The quintic case:** The following quintic polynomials, found by E. Lehmer, generate cyclic quintic totally real extensions of $\mathbb{Q}$ for any $a \in \mathbb{Z}$; see the paper of Schoof and Washington [84, §3].

$$\begin{aligned} x^5 + a^2x^4 - (2a^3+6a^2+10a+10)x^3 \\ + (a^4+5a^3+11a^2+15a+5)x^2 + (a^3+4a^2+10a+10)x + 1. \end{aligned}$$

5. **The sextic case:** For $a \geq 7$, the following polynomials generate cyclic sextic totally real extensions of $\mathbb{Q}$, as proved by Gras, [42].

$$x^6 - 2(a-1)x^5 - 5(a+2)x^4 - 20x^3 + 5(a-1)x^2 + (2a+4)x + 1.$$

7.4.2 Polynomials generating non-abelian extensions of $\mathbb{Q}$

1. Let $F = \mathbb{Q}(\sqrt{d})$ for some fixed square-free $d \in \mathbb{N}$. For any $a, b \in \mathbb{Z}$, the element $(a + b\sqrt{d})^2 + 1$ is totally positive, so $a + b\sqrt{d} + \sqrt{(a + b\sqrt{d})^2 + 1}$ is a totally real unit whose minimal polynomial is

$$x^4 - 4ax^3 + (4(a^2 - b^2d) - 2)x^2 + 4ax^2 + 1$$

   We therefore get an infinite 2-parameter family of suitable polynomials. Although the roots of such polynomials do not generally generate abelian extensions of $\mathbb{Q}$, the roots lie in $F^{(2)} = F^{(2)}_{ab}$, so our theorem applies.

2. More generally, if θ is any totally real algebraic integer, then $u(\theta) = \theta + \sqrt{\theta^2 + 1}$ is a totally real unit which satisfies $x^2 - 2\theta x - 1$. Let α be a fixed totally real algebraic integer with conjugates $\{\alpha = \alpha_1, \ldots, \alpha_n\}$ and let F be a number field containing α, enlarged to be Galois without loss of generality. Let $h(t_1, t_2) \in \mathbb{Z}[t_1, t_2]$ be a polynomial satisfying $\deg_{t_1}(h) > 0$ and $\deg_{t_2}(h) = [\mathbb{Q}(\alpha) : \mathbb{Q}] - 1$. Define $\theta(a) = h(a, \alpha)$. Then $\theta(a)$ is totally real for all $a \in \mathbb{Z}$ and we can take

$$f_a(x) = \prod_{i=1}^{n} (x^2 - 2h(a, \alpha_i)x - 1)$$

   which has $u(\theta(a))$ as a root by design. The coefficients of $f_a(x)$ are polynomials in a which depend on α, and at least one is nonconstant because the degree restrictions on $h(t_1, t_2)$ ensures that $h(a, \alpha)$ outputs infinitely many values as a varies. Again, the roots of all the polynomials $f_a(x)$ lie in $F^{(2)} = F^{(2)}_{ab}$, so our theorem applies, although it will not generally be contained in an abelian extension of $\mathbb{Q}$.

Bibliography

[1] Arora, S., and Eisenträger, K. Constructing Picard curves with complex multiplication using the Chinese remainder theorem. In *Proceedings of the Thirteenth Algorithmic Number Theory Symposium* (2019), vol. 2 of *Open Book Ser.*, Math. Sci. Publ., Berkeley, CA, pp. 21–36.

[2] Ballentine, S., Guillevic, A., Lorenzo García, E., Martindale, C., Massierer, M., Smith, B., and Top, J. Isogenies for point counting on genus two hyperelliptic curves with maximal real multiplication. In *Algebraic geometry for coding theory and cryptography*, vol. 9 of *Assoc. Women Math. Ser.* Springer, Cham, 2017, pp. 63–94.

[3] Biasse, J. Subexponential time relations in the class group of large degree number fields. *Advances in Mathematics of Communications 8*, 4 (2014), 407–425.

[4] Biasse, J.-F., and Fieker, C. Subexponential class group and unit group computation in large degree number fields. *LMS Journal of Computation and Mathematics 17*, A (2014), 385Ð403.

[5] Bisson, G. Computing endomorphism rings of elliptic curves under the GRH. *Journal of Mathematical Cryptology 5* (01 2011).

[6] Bisson, G. Computing endomorphism rings of abelian varieties of dimension two. *Mathematics of Computation of the American Mathematical Society 84*, 294 (2015).

[7] Bisson, G., Cosset, R., and Robert, D. AVIsogenies, 2010. http://avisogenies.gforge.inria.fr/.

[8] Bisson, G., and Sutherland, A. V. Computing the endomorphism ring of an ordinary elliptic curve over a finite field. *Journal of Number Theory 131*, 5 (2011;2009;), 815–831.

[9] Bombieri, E., and Zannier, U. A note on heights in certain infinite extensions of $\mathbb{Q}$. *Atti della Accademia Nazionale dei Lincei. Classe di Scienze Fisiche, Matematiche e Naturali. Rendiconti Lincei. Matematica e Applicazioni 12*, 1 (3 2001), 5–14.

[10] Bos, J. W., Costello, C., Hisil, H., and Lauter, K. Fast cryptography in genus 2. *J. Cryptology 29*, 1 (2016), 28–60.

[11] Bosma, W., Cannon, J., and Playoust, C. The Magma algebra system. I. The user language. *J. Symbolic Comput. 24*, 3-4 (1997), 235–265. Computational algebra and number theory (London, 1993).

[12] Bourbaki, N. *Éléments de mathématique. Théorie des ensembles.* Hermann, Paris, 1970.

[13] Brooks, E. H., Jetchev, D., and Wesolowski, B. Isogeny graphs of ordinary abelian varieties. *Res. Number Theory 3* (2017), Art. 28, 38.

[14] Buchmann, J., and Kessler, V. Computing a reduced lattice basis from a generating system. *Preprint* (1993).

[15] Buhler, J. P., Lenstra Jr., H. W., and Pomerance, C. Factoring integers with the number field sieve. In *The development of the number field sieve*, A. K. Lenstra and H. W. Lenstra, Eds., vol. 1554.;1554;. Springer-Verlag, New York;Berlin;, 1993.

[16] Castillo Fernández, M. *On the Julia Robinson number of rings of totally real algebraic integers in some towers of Nested Square Roots.* PhD thesis, Universidad de Concepción, 2018. URL: http://repositorio.udec.cl/handle/11594/3003.

[17] Castillo Fernández, M., Vidaux, X., and Videla, C. R. Julia robinson numbers and arithmetical dynamic of quadratic polynomials. *arXiv e-prints* (Nov 2017).

[18] Centeleghe, T. G., and Stix, J. Categories of abelian varieties over finite fields, I: Abelian varieties over $\mathbb{F}_p$. *Algebra Number Theory 9*, 1 (2015), 225–265.

[19] Checcoli, S., and Widmer, M. On the northcott property and other properties related to polynomial mappings. *Mathematical Proceedings of the Cambridge Philosophical Society 155*, 1 (07 2013), 1–12.

[20] Cosset, R., and Robert, D. Computing (ℓ,ℓ)-isogenies in polynomial time on Jacobians of genus 2 curves. *Mathematics of Computation of the American Mathematical Society 84*, 294 (2015;2014;), 1953–1975.

[21] Costello, C. Computing supersingular isogenies on Kummer surfaces. In *Advances in cryptology—ASIACRYPT 2018. Part III*, vol. 11274 of *Lecture Notes in Comput. Sci.* Springer, Cham, 2018, pp. 428–456.

[22] Cox, D. A. *Primes of the form $x^2 + ny^2$: Fermat, class field theory, and complex multiplication*, second;2; ed. John Wiley & Sons, Inc, Hoboken, New Jersey, 2013;2014;2012;.

[23] Davis, M., Putnam, H., and Robinson, J. The decision problem for exponential diophantine equations. *Ann. of Math. (2) 74* (1961), 425–436.

[24] Dittmann, P., and Fehm, A. Non-definability of rings of integers in most algebraic fields. Available arXiv:2011.14367.

[25] Dudeanu, A. *Computational aspects of Jacobians of hyperelliptic curves*. PhD thesis, École Polytechnique Fédérale de Lausanne, 2016.

[26] Dudeanu, A., Jetchev, D., Robert, D., and Vuille, M. Cyclic isogenies for abelian varieties with real multiplication. arXiv:1710.05147v2.

[27] Dummit, D. S., and Foote, R. M. *Abstract algebra*, third ed. John Wiley & Sons, Inc., Hoboken, NJ, 2004.

[28] Edwards, H. M. *Galois theory*, vol. 101 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1984.

[29] Eisenträger, K., and Lauter, K. A CRT algorithm for constructing genus 2 curves over finite fields. In *Arithmetic, Geometry, and Coding Theory* (2009), F. Rodier and S. Vladut, Eds., vol. 21, Société Mathématique de France, pp. 161–167.

[30] Eisenträger, K., and Lauter, K. A CRT algorithm for constructing genus 2 curves over finite fields. In *Arithmetics, geometry, and coding theory (AGCT 2005)*, vol. 21 of *Sémin. Congr.* Soc. Math. France, Paris, 2010, pp. 161–176.

[31] Eisenträger, K., Miller, R., Springer, C., and Westrick, L. A topological approach to undefinability in algebraic extensions of $\mathbb{Q}$. *arXiv e-prints* (Oct. 2020), arXiv:2010.09551.

[32] Ershov, Y. L. Fields with continuous local elementary properties. II. *Algebra i Logika 34*, 3 (1995), 262–273, 363.

[33] Fisher, S., and Gartside, P. On the space of subgroups of a compact group. I. *Topology Appl. 156*, 5 (2009), 862–871.

[34] Fried, M. D., Haran, D., and Völklein, H. Real Hilbertianity and the field of totally real numbers. In *Arithmetic geometry (Tempe, AZ, 1993)*, vol. 174 of *Contemp. Math.* Amer. Math. Soc., Providence, RI, 1994, pp. 1–34.

[35] Fried, M. D., Haran, D., and Völklein, H. Real Hilbertianity and the field of totally real numbers. In *Arithmetic geometry (Tempe, AZ, 1993)*, vol. 174 of *Contemp. Math.* Amer. Math. Soc., Providence, RI, 1994, pp. 1–34.

[36] Fukuzaki, K. Definability of the ring of integers in some infinite algebraic extensions of the rationals. *Math. Log. Q. 58* (2012), 317–332.

[37] Galbraith, S. D. Constructing isogenies between elliptic curves over finite fields. *LMS J. Comput. Math. 2* (1999), 118–138.

[38] Gartside, P., and Smith, M. Counting the closed subgroups of profinite groups. *J. Group Theory 13*, 1 (2010), 41–61.

[39] Gaudry, P., Kohel, D., and Smith, B. Counting points on genus 2 curves with real multiplication. In *Advances in cryptology—ASIACRYPT 2011*, vol. 7073 of *Lecture Notes in Comput. Sci.* Springer, Heidelberg, 2011, pp. 504–519.

[40] Gillibert, P., and Ranieri, G. Julia Robinson's numbers. *arXiv e-prints* (Oct 2017).

[41] Gras, M.-N. Table numérique du nombre de classes et des unités des extensions cycliques réelles de degré 4 de $\mathbb{Q}$. *Publications Mathématiques de Besançon* (1977/78), 1–79.

[42] Gras, M.-N. Special units in real cyclic sextic fields. *Mathematics of Computation 48*, 177 (1987), 179–182.

[43] Haran, D., and Jarden, M. The absolute Galois group of a pseudo p-adically closed field. *J. Reine Angew. Math. 383* (1988), 147–206.

[44] Howe, E. W., and Zhu, H. J. On the existence of absolutely simple abelian varieties of a given dimension over an arbitrary field. *Journal of Number Theory 92*, 1 (2002), 139 – 163.

[45] Ionica, S., and Joux, A. Pairing the volcano. *Math. Comp. 82*, 281 (2013), 581–603.

[46] Ionica, S., and Thomé, E. Isogeny graphs with maximal real multiplication. *J. Number Theory 207* (2020), 385–422.

[47] Jacobson, N. *Basic algebra. II*, second ed. W. H. Freeman and Company, New York, 1989.

[48] Janusz, G. J. *Algebraic number fields*, 2nd ed., vol. 7;7.;. American Mathematical Society, Providence, R.I, 1996.

[49] Jetchev, D., and Wesolowski, B. Horizontal isogeny graphs of ordinary abelian varieties and the discrete logarithm problem. arXiv:1506.00522v2.

[50] Kamke, E. Verallgemeinerungen des Waring-Hilbertschen Satzes. *Math. Ann. 83*, 1-2 (1921), 85–112.

[51] Kani, E. Products of CM elliptic curves. *Collect. Math. 62*, 3 (2011), 297–339.

[52] Kishi, Y. A family of cyclic cubic polynomials whose roots are systems of fundamental units. *Journal of Number Theory 102*, 1 (2003), 90 – 106.

[53] Klüners, J., and Pauli, S. Computing residue class rings and picard groups of orders. *Journal of Algebra 292*, 1 (2005), 47–64.

[54] Koenigsmann, J. Undecidability in number theory. In *Model theory in algebra, analysis and arithmetic*, vol. 2111 of *Lecture Notes in Math.* Springer, Heidelberg, 2014, pp. 159–195.

[55] Koenigsmann, J. Defining $\mathbb{Z}$ in $\mathbb{Q}$. *Ann. of Math. (2) 183*, 1 (2016), 73–93.

[56] Kohel, D. R. *Endomorphism rings of elliptic curves over finite fields*. ProQuest LLC, Ann Arbor, MI, 1996. Thesis (Ph.D.)–University of California, Berkeley.

[57] Lang, S. *Algebra*, third ed., vol. 211 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.

[58] Lauter, K. E., and Robert, D. Improved CRT algorithm for class polynomials in genus 2. In *ANTS X—Proceedings of the Tenth Algorithmic Number Theory Symposium* (2013), vol. 1 of *Open Book Ser.*, Math. Sci. Publ., Berkeley, CA, pp. 437–461.

[59] Lenstra, H. Algorithms in algebraic number-theory. *Bulletin of the American Mathematical Society 26*, 2 (1992), 211–244.

[60] Lenstra, H. W. Factoring integers with elliptic curves. *Annals of Mathematics 126*, 3 (1987), 649–673.

[61] Lenstra, Jr., H. W. Complex multiplication structure of elliptic curves. *J. Number Theory 56*, 2 (1996), 227–241.

[62] Liu, Q. *Algebraic geometry and arithmetic curves*, vol. 6 of *Oxford Graduate Texts in Mathematics*. Oxford University Press, Oxford, 2002. Translated from the French by Reinie Erné, Oxford Science Publications.

[63] Lv, C., and Deng, Y. On orders in number fields: Picard groups, ring class fields and applications. *Science China Mathematics 58*, 8 (Aug 2015), 1627–1638.

[64] Martindale, C. *Isogeny Graphs, Modular Polynomials, and Applications*. ProQuest LLC, Ann Arbor, MI, 2018. Thesis (Ph.D.)–University of Leiden.

[65] Martínez-Ranero, C., Utreras, J., and Videla, C. R. Undecidability of $\mathbb{Q}^{(2)}$. *Proc. Amer. Math. Soc. 148*, 3 (2020), 961–964.

[66] Matiyasevich, J. V. The Diophantineness of enumerable sets. *Dokl. Akad. Nauk SSSR 191* (1970), 279–282.

[67] Matsumura, H. *Commutative ring theory*, vol. 8 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1986. Translated from the Japanese by M. Reid.

[68] Mestre, J.-F. Construction de courbes de genre 2 à partir de leurs modules. In *Effective methods in algebraic geometry (Castiglioncello, 1990)*, vol. 94 of *Progr. Math.* Birkhäuser Boston, Boston, MA, 1991, pp. 313–334.

[69] Miller, R. Isomorphism and classification for countable structures. *Computability 8*, 2 (2019), 99–117.

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

[71] Milne, J. S. Jacobian varieties. In *Arithmetic geometry (Storrs, Conn., 1984)*. Springer, New York, 1986, pp. 167–212.

[72] Neukirch, J. *Algebraic number theory*, vol. 322 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999. Translated from the 1992 German original and with a note by Norbert Schappacher, With a foreword by G. Harder.

[73] Park, J. A universal first-order formula defining the ring of integers in a number field. *Math. Res. Lett. 20*, 5 (2013), 961–980.

[74] Pila, J. Frobenius maps of abelian varieties and finding roots of unity in finite fields. *Mathematics of Computation 55*, 192 (1990), 745–763.

[75] Poonen, B. Undecidability in number theory. *Notices Amer. Math. Soc. 55*, 3 (2008), 344–350.

[76] Poonen, B. Characterizing integers among rational numbers with a universal-existential formula. *Amer. J. Math. 131*, 3 (2009), 675–682.

[77] Pop, F. Classically projective groups and pseudo classically closed fields. In *Valuation theory and its applications, Vol. II (Saskatoon, SK, 1999)*, vol. 33 of *Fields Inst. Commun.* Amer. Math. Soc., Providence, RI, 2003, pp. 251–283.

[78] Rabin, M. Computable algebra, general theory, and theory of computable fields. *Trans. AMS 95* (1960), 341–360.

[79] Reid, M. *Undergraduate commutative algebra*, vol. 29 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1995.

[80] Robinson, J. Definability and decision problems in arithmetic. *J. Symbolic Logic 14* (1949), 98–114.

[81] Robinson, J. The undecidability of algebraic rings and fields. *Proc. Amer. Math. Soc. 10* (1959), 950–957.

[82] Robinson, J. On the decision problem for algebraic rings. In *Studies in mathematical analysis and related topics*. Stanford Univ. Press, Stanford, Calif, 1962, pp. 297–304.

[83] Rumely, R. S. Undecidability and definability for the theory of global fields. *Transactions of the American Mathematical Society 262*, 1 (1980), 195–217.

[84] Schoof, R., and Washington, L. C. Quintic polynomials and real cyclotomic fields with large class number. *Mathematics of Computation 50*, 182 (1988), 543–556.

[85] Serre, J.-P. *Topics in Galois theory*, second ed., vol. 1 of *Research Notes in Mathematics*. A K Peters, Ltd., Wellesley, MA, 2008. With notes by Henri Darmon.

[86] Shanks, D. The simplest cubic fields. *Mathematics of Computation 28*, 128 (1974), 1137–1152.

[87] Shimura, G., and Taniyama, Y. *Complex multiplication of Abelian varieties and its applications to number theory*, vol. 6. Mathematical Society of Japan, Tokyo, 1961.

[88] Shlapentokh, A. Diophantine classes of holomorphy rings of global fields. *J. Algebra 169*, 1 (1994), 139–175.

[89] Shlapentokh, A. Hilbert's tenth problem over number fields, a survey. In *Hilbert's tenth problem: relations with arithmetic and algebraic geometry (Ghent, 1999)*, vol. 270 of *Contemp. Math.* Amer. Math. Soc., Providence, RI, 2000, pp. 107–137.

[90] Shlapentokh, A. Rings of algebraic numbers in infinite extensions of $\mathbb{Q}$ and elliptic curves retaining their rank. *Archive for Mathematical Logic 48*, 1 (2009), 77–114.

[91] Shlapentokh, A. First-order decidability and definability of integers in infinite algebraic extensions of the rational numbers. *Israel Journal of Mathematics 226*, 2 (2018), 579–633.

[92] Siegel, C. Darstellung total positiver Zahlen durch Quadrate. *Math. Z. 11*, 3-4 (1921), 246–275.

[93] Silverman, J. H. *The arithmetic of elliptic curves*, second ed., vol. 106 of *Graduate Texts in Mathematics*. Springer, Dordrecht, 2009.

[94] Springer, C. Computing the endomorphism ring of an ordinary abelian surface over a finite field. *J. Number Theory 202* (2019), 430–457.

[95] Springer, C. Undecidability, unit groups, and some totally imaginary infinite extensions of $\mathbb{Q}$. *Proc. Amer. Math. Soc. 148*, 11 (2020), 4705–4715.

[96] Springer, C. The structure of the group of rational points of an abelian variety over a finite field. *European Journal of Mathematics* (2021).

[97] Streng, M. *Complex multiplication of abelian surfaces*. PhD thesis, Universiteit Leiden, 2010.

[98] Sutherland, A. V. Computing Hilbert class polynomials with the Chinese remainder theorem. *Mathematics of Computation of the American Mathematical Society 80*, 273 (2011;2010;2009;), 501–538.

[99] Tate, J. Endomorphisms of abelian varieties over finite fields. *Inventiones Mathematicae 2*, 2 (1966), 134–144.

[100] Tate, J. Classes d'isogénie des variétés abéliennes sur un corps fini (d'après T. Honda). In *Séminaire Bourbaki. Vol. 1968/69: Exposés 347–363*, vol. 175 of *Lecture Notes in Math.* Springer, Berlin, 1971, pp. Exp. No. 352, 95–110.

[101] Tate, J. Finite flat group schemes. In *Modular forms and Fermat's last theorem (Boston, MA, 1995)*. Springer, New York, 1997, pp. 121–154.

[102] Thiel, C. *On the complexity of some problems in algorithmic algebraic number theory*. PhD thesis, University of Saarland, 1995.

[103] Van Den Dries, L. Elimination theory for the ring of algebraic integers. *Journal fur die Reine und Angewandte Mathematik 1988*, 388 (1988), 189–205.

[104] Vidaux, X., and Videla, C. R. Definability of the natural numbers in totally real towers of nested square roots. *Proceedings of the American Mathematical Society 143*, 10 (2015), 4463–4477.

[105] Vidaux, X., and Videla, C. R. A note on the northcott property and undecidability. *Bulletin of the London Mathematical Society 48*, 1 (2015), 58–62.

[106] Videla, C. R. Definability of the ring of integers in pro-p Galois extensions of number fields. *Israel J. Math. 118* (2000), 1–14.

[107] Videla, C. R. Definability of the ring of integers in pro-p galois extensions of number fields. *Israel Journal of Mathematics 118*, 1 (Dec 2000), 1–14.

[108] Videla, C. R. The undecidability of cyclotomic towers. *Proc. Amer. Math. Soc. 128*, 12 (2000), 3671–3674.

[109] Washington, L. C. *Introduction to cyclotomic fields*, 2nd ed., vol. 83. Springer, New York, 1997.

[110] Waterhouse, W. C. Abelian varieties over finite fields. *Ann. Sci. École Norm. Sup. (4) 2* (1969), 521–560.

[111] Waterhouse, W. C., and Milne, J. S. Abelian varieties over finite fields. In *1969 Number Theory Institute (Proc. Sympos. Pure Math., Vol. XX, State Univ. New York, Stony Brook, N.Y., 1969)* (1971), pp. 53–64.

[112] Widmer, M. On certain infinite extensions of the rationals with northcott property. *Monatshefte für Mathematik 162*, 3 (2011), 341–353.

Vita

Caleb Springer

The Pennsylvania State University
Department of Mathematics
015 McAllister Building
University Park, PA 16802, USA

cks5320@psu.edu
http://personal.psu.edu/cks5320/
Citizenship: USA
Languages: English

Education

August 2021	Ph.D. Mathematics, The Pennsylvania State University Advisor: Kirsten Eisenträger
May 2015	B.Sc. Mathematics, University of Michigan

Awards

Fall 2020	Pritchard Dissertation Fellowship, The Pennsylvania State University
May 2015	Outstanding Achievement in Mathematics, The University of Michigan

Papers

1. Doubly isogenous genus-2 curves with D_4-action (with Vishal Arul, Jeremy Booher, Steven Groen, Everett Howe, Wanlin Li, Vlad Matei, and Rachel Pries.)
Submitted for publication. 33 pages. arxiv:2102.11419

2. A topological approach to undefinability in algebraic extensions of $\mathbb{Q}$. (with Kirsten Eisenträger, Russell Miller and Linda Westrick.)
Submitted for publication. 22 pages. arxiv:2010.09551

3. The structure of the group of rational points of an abelian variety over a finite field.
European Journal of Mathematics. (23 March 2021). 13 pages. arxiv:2006.00637

4. Restrictions on Weil polynomials of Jacobians of hyperelliptic curves. (with Edgar Costa, Ravi Donepudi, Ravi Fernando, Valentijn Karemaker, and Mckenzie West.)
To appear in the Simons Symposia volume for the Simons Collaboration "Arithmetic Geometry, Number Theory, and Computation". 15 pages. arxiv:2002.02067

5. Undecidability, unit groups, and some totally imaginary infinite extensions of $\mathbb{Q}$.
Proceedings of the AMS. Volume 148, Number 11, November 2020, Pages 4705–4715
10 pages. arxiv:1910.01239

6. Computing the endomorphism ring of an ordinary abelian surface over a finite field.
Journal of Number Theory. Volume 202, September 2019, Pages 430-457.
27 pages. arXiv:1810.12270