Computing the endomorphism ring of an ordinary abelian surface over a finite field

Caleb Springer [1]

Department of Mathematics, The Pennsylvania State University, University Park, PA 16802, USA

A R T I C L E   I N F O

A B S T R A C T

We present a new algorithm for computing the endomorphism ring of an ordinary abelian surface over a finite field which is subexponential and generalizes an algorithm of Bisson and Sutherland for elliptic curves. The correctness of this algorithm only requires the heuristic assumptions required by the algorithm of Biasse and Fieker [2] which computes the class group of an order in a number field in subexponential time. Thus we avoid the multiple heuristic assumptions on isogeny graphs and polarized class groups which were previously required. The output of the algorithm is an ideal in the maximal totally real subfield of the endomorphism algebra, generalizing the elliptic curve case.

© 2019 Elsevier Inc. All rights reserved.

1. Introduction

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, $\mathrm{End}(E)$ 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 [33, 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 $\mathrm{End}(E)$ is equivalent to computing the index $[\mathcal{O}_K : \mathrm{End}(E)]$, which is a divisor of $[\mathcal{O}_K : \mathbb{Z}[\pi]]$.

The first algorithm for computing $\mathrm{End}(E)$ for an ordinary elliptic curve E was given by Kohel in his thesis [21]. 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 $[\mathcal{O}_K : \mathrm{End}(E)]$, and thus determine $\mathrm{End}(E)$ as desired, although this method is slow when working with large prime factors.

To make a faster algorithm for computing $\mathrm{End}(E)$, Bisson and Sutherland exploited the following fact, which was proven by Waterhouse [33, Theorem 4.5]. Given an order $\mathcal{O} \subseteq K$, the class group $\mathrm{Cl}(\mathcal{O})$ 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 $\mathrm{End}(E)$ by computing various isogenies. In particular, one may compare $\mathrm{Cl}(\mathrm{End}(E))$ 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 $\mathrm{Cl}(\mathrm{End}(E))$ is sufficient for determining the large prime factors of $[\mathcal{O}_K : \mathrm{End}(E)]$, which leads to a probabilistic algorithm to determine $[\mathcal{O}_K : \mathrm{End}(E)]$ in subexponential time [6].

Bisson extended the algorithm above to absolutely simple, principally polarized, ordinary abelian varieties of dimension 2 in [4]. 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 paper 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 [2]. As a trade-off, we must explicitly restrict the class of abelian varieties that we consider.

1.1. 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 $\text{End}(A)$ 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 $\text{End}(A) \supseteq \mathcal{O}_F$ where F is the maximal totally real subfield of K. According to Brooks, Jetchev and Wesolowski [8, 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 $\text{End}(A)$.

The foundation of our algorithm is the free action of the ideal class group $\text{Cl}(\mathcal{O})$ 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 $\text{Cl}(\text{End}(A))$ if and only if the corresponding isogeny maps A to an isomorphic abelian variety. Using the same tactic as Bisson and Sutherland, we compare $\text{Cl}(\text{End}(A))$ to $\text{Cl}(\mathcal{O})$ for known "testing" orders $\mathcal{O} \subseteq K$ by generating class group relations. When $\text{End}(A)$ 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 $\text{End}(A)$.

However, unlike the elliptic curve case, it is not always possible to compute the action of $\text{Cl}(\mathcal{O})$ on principally polarized abelian varieties of dimension $g > 1$. In fact, the isogenies corresponding to elements of $\text{Cl}(\mathcal{O})$ 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 $\text{Cl}(\mathcal{O})$, 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 $\text{Cl}(\mathcal{O})$ 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 $\text{Cl}(\mathcal{O})$ and the action of $\text{Cl}(\mathcal{O})$ 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. [11, 14], and certain isogeny graphs are volcanoes [8]. 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 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 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 [1,2], thereby avoiding the unconventional 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 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 [1,2] and isogeny computation [11,14], 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.

1.2. Organization

The paper is organized as follows. In Sections 2 and 3, we present background material for abelian varieties over finite fields with maximal RM and prove that class group relations are sufficient to distinguish orders in CM fields. In Section 4, we restrict to the case of abelian varieties of dimension $g = 2$ with the special properties outlined above, present the algorithm and analyze the runtime under certain heuristic assumptions. In Section 5, we present an illustrative example.

1.3. Acknowledgments

2. Background

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 [27], the characteristic polynomial f_π of π can be computed in time polynomial in $\log q$. 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 [31]. In the remainder of the paper, we will assume that f_π is known and an embedding $\mathrm{End}(A) \hookrightarrow \mathbb{Q}(\pi) \cong \mathbb{Q}[x]/(f_\pi)$ 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.

2.2. Identifying orders

Waterhouse [33, 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 $\mathrm{End}(A)$, 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. $\mathrm{End}(A)$ 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 2.1 (Theorem 2.1, [8]). *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 [20, §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 2.1 computable.

Proposition 2.2. *Algorithm 1 is correct.*

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

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

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

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

Define M_j as the "multiplication by ω_j" matrix:

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

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

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

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 [22, §2.10] or [9, §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 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 $\log|\mathrm{disc}(\mathcal{O})|$.*

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

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

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

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

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 [9, Proposition 5.3] to bound the size of $c_{i,j,k}$ by a polynomial in g, $\log|\mathrm{disc}(\mathcal{O}_K)|$ and $\log|\mathrm{disc}(\mathcal{O}_F)|$. Because $[\mathcal{O}_K : \mathcal{O}]$, $\log|\mathrm{disc}(\mathcal{O}_K)|$ and $\log|\mathrm{disc}(\mathcal{O}_F)|$ are all bounded by $\log|\mathrm{disc}(\mathcal{O})|$, 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

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

Embedding the equation (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 $2^{-g}|\mathrm{disc}(\mathcal{O}_K)|^{1/2}$ by definition. Hence by Cramer's rule

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

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, $\log|\mathrm{disc}(\mathcal{O}_K)|$ and $\log|\mathrm{disc}(\mathcal{O}_F)|$; see [9, Proposition 5.2]. We conclude that the sizes of the integers $c_{i,j,k}$ are bounded as desired. $\square$

For the purposes of this paper, 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 paper, we note that the algorithm is also polynomial in g.

Corollary 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 $\log|\mathrm{disc}(\mathcal{O})|$, we simply need to notice that $\mathcal{O} \supseteq \mathbb{Z}[\pi]$ implies

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

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. $\square$

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 $\mathrm{End}(A)$ under the fixed embedding $\mathrm{End}(A) \hookrightarrow K$. Now computing $\mathrm{End}(A)$ 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.

2.3. Class group action

Following Brooks, Jetchev and Wesolowski [8], 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 $\alpha \in \mathrm{End}(A)$, we write $A[\alpha]$ for the kernel of α.

Definition 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 [8]. We record a proof here for completeness.[2]

Lemma 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 $\mathrm{End}(A)$ is maximal is given in [33, Theorem 3.15] and the proof generalizes immediately. The 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. $\quad\square$

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 [6, §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 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 [28,33]. The fact that these isogenies are $\mathfrak{l}$-isogenies is clear from the definition, as pointed out in [8, Theorem 4.3]. $\quad\square$

Theorem 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.*

---

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

Proof. This follows immediately from Theorem 4.3 and Proposition 4.10 in [8]. □

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 [6]. 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 [4]. 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 [11]. If $N\mathfrak{l} = \ell$, then ϕ is known as a *cyclic isogeny* and the target variety can be determined by the algorithm in [13, Chapter 4] and [14]. In both cases, the algorithm is polynomial in ℓ and $\log q$. By an algorithm of Mestre [25], we assume that the output of these algorithms is the Jacobian of a curve defined over the minimal field; see also [11]. By combining these algorithms and analyzing the running times of each algorithm, we find the following theorem.

Theorem 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 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 [4, §4]. However, it is only known that these easily

computed relations are sufficient for determining the endomorphism ring under certain heuristic assumptions [4, Theorem 7.1].

2.5. Navigating isogeny graphs and identifying abelian varieties with invariants

We need a way to identify abelian varieties as we navigate the various ℓ-isogeny graphs with the class group action. To do this, we follow the ideas of [19, §4.2]. Rosenhain invariants[3] 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*,[4] as discussed in [4,8,19]. For a given order $\mathcal{O} \subseteq K$, the polarized class group is defined as follows

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

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 $\beta = u\bar{u}\alpha$. 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

$$1 \to (\mathcal{O}_F^\times)^+/N_{K/F}(\mathcal{O}^\times) \xrightarrow{u \mapsto (\mathcal{O}, u)} \mathfrak{C}(\mathcal{O}) \xrightarrow{(\mathfrak{a}, \alpha) \mapsto \mathfrak{a}} \mathrm{Cl}(\mathcal{O}) \xrightarrow{\mathfrak{a} \mapsto N_{K/F}(\mathfrak{a})} \mathrm{Cl}^+(\mathcal{O}_F), \quad (2)$$

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.

---

[3] 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.

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

Proposition 2.10. *If A is a principally polarizable ordinary simple abelian variety over $\mathbb{F}_q$ with maximal RM and $\mathrm{End}(A) = \mathcal{O}$, 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 [8]. □

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 [4]. 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.

2.6. Small primes

We only use the class group action to determine the power of $\mathfrak{p}$ dividing $\mathfrak{f}^+(\mathrm{End}(A))$ 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 [4] and use the method of Eisenträger and Lauter [15] to test whether $\mathrm{End}(A)$ contains $\mathcal{O}(\mathfrak{p}^k)$. However, this method does not immediately produce an isogenous abelian variety A' such that $\mathfrak{f}^+(\mathrm{End}(A))$ is not divisible by the small prime factors of $\mathfrak{v}$, which is a feature in Bisson and Sutherland's original elliptic curve algorithm [6]. For this, we need the following additional result about isogeny graphs. More background may be found in [8].

Theorem 2.11 *([8, 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 2.9 and Lemma 3.11. Hence the algorithms presented for volcano navigation in [30] 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.

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 [6] 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 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 $\mathfrak{b} \mapsto \mathfrak{b} \cap \mathcal{O}$ [24, 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 [3,4]. Note that we do not follow Bisson and Sutherland's definition of relation in [6] because we do not have the benefit of an easily computable modular polynomial.

Definition 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 [6, 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 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 $\mathfrak{f}^+$. $\square$

3.1. Review of class field theory

In order to prove that an algorithm based on Corollary 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 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 [18,24,26] 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

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

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 [24, 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 $I_K(\mathfrak{m}_1)/H_1 \cong I_K(\mathfrak{m}_2)/H_2$ [18, 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 $\mathbf{H}$ [18, 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 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 3.5 (*The Classification Theorem, [18, 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 3.6 *([24, 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 3.7 *([24, 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 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 3.5 and the definition of $\mathfrak{f}(L_1/K)$, which imply

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

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 3.3. To begin, we prove the following technical lemma.

Lemma 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 [12]. 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 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 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. $\quad\square$

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

Proposition 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 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 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 [26, 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 3.8 we have $L_2 \not\subseteq L_1$, which completes the proof. $\quad\square$

The following lemma and its proof are a generalization of [12, Exercises 9.20-22].

Lemma 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 3.5, there is a 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

$$P_{K,1}(\mathfrak{m}) \cap I_K(\mathfrak{f}) \subseteq P_{K,\mathcal{O}}(\mathfrak{f}), \tag{3}$$

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 (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

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

This implies that

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

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 $|\ker \pi \cap \operatorname{Im}\beta| = |\ker \pi| = 1$. 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

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

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

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

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. $\quad\square$

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 [6]. 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.

4.1. Requirements (**R**)

For the remainder of the paper, 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 2.1, 2.9 and

2.11, and Lemma 3.11. In particular, $\mathrm{End}(A)$ is uniquely identified by an ideal $\mathfrak{f}^+(A) \subseteq \mathcal{O}_F$ and we can compute the isogenies corresponding to the action of $\mathrm{Cl}(\mathrm{End}(A))$. Notice that the conditions below are all easily verifiable.

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 because $\mathcal{O}_K^* = \mathcal{O}_F^*$ is true for every primitive quartic CM field except the cyclotomic field $K = \mathbb{Q}(\zeta_5)$ [29, Lemma II.3.3].

4.2. 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 [2], while the final two are moderate heuristic assumptions, similar to those in [6], 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 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 [10] and ECM finds a prime factor p of integer n in expected time $L[1/2, 2](p) \log^2 n$ by [23].

4.3. Finding relations

Following [3,4,6], we will use the reduction of random ideals to produce relations which hold in a given order $\mathcal{O} \subseteq K$ of discriminant $\Delta = \mathrm{disc}(\mathcal{O})$. This is the same idea that is used in ideal class group computation algorithms, such as [2]. For background on ideal reduction, see [32, 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 $\mathrm{Cl}(\mathcal{O})$ 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 [1, §4.3] and [2], 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.

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 $B = L[1/2, \mu](n)$, $D_1 = \mathrm{disc}(\mathcal{O}_1)$ 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.

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.

Proposition 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 [6, Proposition 6], we expect the number of attempts required to find an ideal $\mathfrak{b}$ with B-smooth norm to 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 [23] 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 [1, Algorithm 7]. First, one computes the class group in time $L[1/2, c](|D_2|)$ by [2, 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 [1, 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. $\square$

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 $\mathrm{End}(A)$

1   Determine the ideal $\mathfrak{v} \subseteq \mathcal{O}_F$ defining the order $\mathcal{O}_F[\pi] \subseteq \mathcal{O}_K$.
2   Fix bound $C \geq 3$ and initiate $\mathfrak{u} = 1$.
3   **for** *every* $\mathfrak{p} \subseteq \mathcal{O}_F$ *with* $N\mathfrak{p} < C$ *which divides* $\mathfrak{v}$ **do**
4   Use isogeny climbing to determine $v_\mathfrak{p} = v_\mathfrak{p}(\mathfrak{f}^+(A))$.
5   Replace $\mathfrak{u}$ with the product $\mathfrak{u}\mathfrak{p}^{v_\mathfrak{p}}$.
6   Update A to be the abelian variety found by isogeny climbing
7   Remove powers of $\mathfrak{p}$ from $\mathfrak{v}$.
8   **end**
9   **for** *every* $\mathfrak{p} \subseteq \mathcal{O}_F$ *with* $N\mathfrak{p} \geq C$ *which divides* $\mathfrak{v}$ **do**
10   Find relation R which holds in $\mathcal{O}(\mathfrak{v}/\mathfrak{p})$ but not $\mathcal{O}(\mathfrak{p})$.
11   **if** *R does not hold in $\mathcal{O}(\mathfrak{f}^+(A))$* **then**
12   Update $\mathfrak{u}$ to be the product $\mathfrak{u}\mathfrak{p}$
13   **end**
14   **end**
15   **return** $\mathfrak{f}^+(A) = \mathfrak{u}$.

4.4. Computing from above

Now we present our algorithm for computing $\mathfrak{f}^+(A)$. As noted in Section 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 3.3. There are infinitely many class group relations R as required in Step 10 according to Proposition 3.10.

Proposition 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 2.9.

Proof. Computing $\mathfrak{v}$ takes polynomial time by Corollary 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 $\log |\mathrm{disc}(\mathbb{Z}[\pi, \overline{\pi}])| < 4\log q + 12\log 2$ by [4, 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 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:

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

4.5. 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
1 **for** *For every prime $\mathfrak{p}$ with $v_{\mathfrak{p}}(\mathfrak{v}) - v_{\mathfrak{p}}(\mathfrak{u}) > 0$* **do**
2   Find a Relation $R_{\mathfrak{p}}$ which holds in $\mathcal{O}(\mathfrak{u})$ but not $\mathcal{O}(\mathfrak{p})$
3 **end**
4 **for** *For every prime $\mathfrak{p}$ of $\mathfrak{u}$* **do**
5   Find a Relation $R_{\mathfrak{p}}$ which holds in $\mathcal{O}(\mathfrak{u}\mathfrak{p}^{-1})$ but not $\mathcal{O}(\mathfrak{p})$.
6 **end**
7 **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 3.3.

Corollary 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 $|\mathrm{disc}(\mathcal{O}_F[\pi])| \leq |\mathrm{disc}(\mathbb{Z}[\pi, \overline{\pi}])| \leq 4^6 q^4$ from [4, 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 $O(\log N(\mathfrak{a}))$ [32, Corollary 3.4.13]. $\quad \square$

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 $\text{End}(A)$ 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 $\text{End}(A)$ by computing isogenies.
7   Immediately **return false** if a relation holds.
8 end
9 return true

Proposition 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 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 4.1, 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 2.9. $\square$

4.6. Remark

In [6, 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 [6, Algorithm 2] which uses the CERTIFY and VERIFY algorithms above correctly computes $\mathfrak{f}^{+}(\text{End}(A))$, but does not gain an performance improvement by focusing on abelian varieties with nearly-maximal endomorphism ring.

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 were performed using Magma [7] and the AVIsogenies library [5] 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.

5.1. Example

Let $q = 82307$ and let A be the Jacobian of the hyperelliptic curve C defined over $\mathbb{F}_q$ by
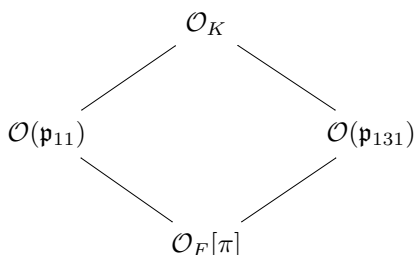
$$y^2 = x^5 - 3x^4 + 5x^3 - x^2 - 2x + 1.$$

The curve C is the specialization to $s = t = 1$ of the 2-parameter family given in [16, §4.3]. By [16, 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 [17, 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, $\mathrm{End}(A)$ is one of the following four orders, using the notation of Section 2.2:

$$
\begin{array}{ccc}
 & \mathcal{O}_K & \\
\diagup & & \diagdown \\
\mathcal{O}(\mathfrak{p}_{11}) & & \mathcal{O}(\mathfrak{p}_{131}) \\
\diagdown & & \diagup \\
 & \mathcal{O}_F[\pi] &
\end{array}
$$

Next, we need suitable class group relations, as described in Corollary 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 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 [15] 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 [4, §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.

References

[1] J. Biasse, Subexponential time relations in the class group of large degree number fields, Adv. Math. Commun. 8 (4) (2014) 407–425.

[2] J.-F. Biasse, C. Fieker, Subexponential class group and unit group computation in large degree number fields, LMS J. Comput. Math. 17 (A) (2014) 385–403.

[3] G. Bisson, Computing endomorphism rings of elliptic curves under the GRH, J. Math. Cryptol. 5 (2011).

[4] G. Bisson, Computing endomorphism rings of abelian varieties of dimension two, Math. Comput. 84 (2015) 294.

[5] G. Bisson, R. Cosset, D. Robert, AVIsogenies, http://avisogenies.gforge.inria.fr/, 2010.

[6] G. Bisson, A.V. Sutherland, Computing the endomorphism ring of an ordinary elliptic curve over a finite field, J. Number Theory 131 (5) (2011) 815–831, 2009.

[7] W. Bosma, J. Cannon, C. Playout, The magma algebra system. I. The user language, in: Computational Algebra and Number Theory, London, 1993, J. Symbolic Comput. 24 (3–4) (1997) 235–265.

[8] E.H. Brooks, D. Jetchev, B. Wesolowski, Isogeny graphs of ordinary abelian varieties, Res. Number Theory 3 (1) (2017) 1–38.

[9] J. Buchmann, V. Kessler, Computing a reduced lattice basis from a generating system, Preprint, 1993.

[10] J.P. Buhler, H.W. Lenstra Jr., C. Pomerance, Factoring integers with the number field sieve, in: A.K. Lenstra, H.W. Lenstra (Eds.), The Development of the Number Field Sieve, Vol. 1554, Springer-Verlag, New York, Berlin, 1993.

[11] R. Cosset, D. Robert, Computing (ℓ, ℓ)-isogenies in polynomial time on Jacobians of genus 2 curves, Math. Comp. 84 (294) (2015) 1953–1975.

[12] D.A. Cox, Primes of the Form $x^2 + ny^2$: Fermat, Class Field Theory, and Complex Multiplication, second ed., John Wiley & Sons, Inc, Hoboken, New Jersey, 2013.

[13] A. Dudeanu, Computational Aspects of Jacobians of Hyperelliptic Curves, PhD thesis, École Polytechnique Fédérale de Lausanne, 2016.

[14] A. Dudeanu, D. Jetchev, D. Robert, M. Vuille, Cyclic isogenies for abelian varieties with real multiplication, arXiv:1710.05147v2.

[15] K. Eisenträger, K. Lauter, A CRT algorithm for constructing genus 2 curves over finite fields, in: F. Rodier, S. Vladut (Eds.), Arithmetic, Geometry, and Coding Theory, in: Société Mathématique de France, vol. 21, 2009, pp. 161–167.

[16] P. Gaudry, D. Kohel, B. Smith, Counting points on genus 2 curves with real multiplication, in: D.H. Lee, X. Wang (Eds.), Advances in Cryptology – ASIACRYPT 2011, Berlin, Heidelberg, 2011, Springer, Berlin, Heidelberg, 2011, pp. 504–519.

[17] E.W. Howe, H.J. Zhu, On the existence of absolutely simple abelian varieties of a given dimension over an arbitrary field, J. Number Theory 92 (1) (2002) 139–163.

[18] G.J. Janusz, Algebraic Number Fields, 2nd ed., vol. 7, American Mathematical Society, Providence, R.I., 1996.
[19] D. Jetchev, B. Wesolowski, Horizontal isogeny graphs of ordinary abelian varieties and the discrete logarithm problem, arXiv:1506.00522v2.
[20] J. Klüners, S. Pauli, Computing residue class rings and Picard groups of orders, J. Algebra 292 (1) (2005) 47–64.
[21] D.R. Kohel, Endomorphism Rings of Elliptic Curves over Finite Fields, PhD thesis, University of California at Berkley, 1996, http://echidna.maths.usyd.edu.au/kohel/pub/thesis.pdf.
[22] H. Lenstra, Algorithms in algebraic number-theory, Bull. Amer. Math. Soc. 26 (2) (1992) 211–244.
[23] H.W. Lenstra, Factoring integers with elliptic curves, Ann. Math. 126 (3) (1987) 649–673.
[24] C. Lv, Y. Deng, On orders in number fields: Picard groups, ring class fields and applications, Sci. China Math. 58 (8) (Aug 2015) 1627–1638.
[25] J.-F. Mestre, Construction de courbes de genre 2 á partir de leurs modules, in: T. Mora, C. Traverso, S.O. Service (Eds.), Effective Methods in Algebraic Geometry, Vol. 94, Birkhäuser Boston, Boston, MA, 1991, pp. 313–334.
[26] J. Neukirch, Algebraic Number Theory, Vol. 322, Springer Berlin Heidelberg, Berlin, Heidelberg, 1999.
[27] J. Pila, Frobenius maps of abelian varieties and finding roots of unity in finite fields, Math. Comp. 55 (192) (1990) 745–763.
[28] G. Shimura, Y. Taniyama, Complex Multiplication of Abelian Varieties and Its Applications to Number Theory, Vol. 6, Mathematical Society of Japan, Tokyo, 1961.
[29] M. Streng, Complex Multiplication of Abelian Surfaces, PhD thesis, Universiteit Leiden, 2010.
[30] A.V. Sutherland, Computing Hilbert class polynomials with the Chinese remainder theorem, Math. Comp. 80 (273) (2011) 501–538.
[31] J. Tate, Endomorphisms of abelian varieties over finite fields, Invent. Math. 2 (2) (1966) 134–144.
[32] C. Thiel, On the Complexity of Some Problems in Algorithmic Algebraic Number Theory, PhD thesis, University of Saarland, 1995.
[33] W.C. Waterhouse, Abelian varieties over finite fields, Annales scientifiques de l'École Normale Supérieure 2 (4) (1969) 521–560.