Computing endomorphism rings of elliptic curves under the GRH

Gaetan Bisson

LORIA, 54506 Vandœuvre-lès-Nancy, France
TU/e, 5600 MB Eindhoven, The Netherlands

Abstract

We design a probabilistic algorithm for computing endomorphism rings of ordinary elliptic curves defined over finite fields that we prove has a subexponential runtime in the size of the base field, assuming solely the generalized Riemann hypothesis.

Additionally, we improve the asymptotic complexity of previously known, heuristic, subexponential methods by describing a faster isogeny-computing routine.

1 Introduction

Endomorphism rings of ordinary elliptic curves over finite fields are central objects in complex multiplication (CM) theory; as such, they appear in various computational number-theoretic contexts. For instance, the CM method for generating curves with a prescribed number of points relies on evaluating so-called Hilbert class polynomials, for which the state-of-the-art algorithm of [19] requires an endomorphism-ring-computing subroutine. They are also potentially relevant security parameters in certain cryptographic applications.

They were first studied by Kohel [13] who, assuming the generalized Riemann hypothesis (GRH), gave a deterministic method for computing them in time $O(q^{1/3+\epsilon})$ where q is the cardinality of the base field. Recently, a probabilistic algorithm with subexponential complexity in $\log q$ was obtained in [3] by relying on several additional assumptions; its runtime is

$$L(q)^{\sqrt{3}/2+o(1)} \qquad \text{where} \qquad L(x) = \exp\sqrt{\log x \log\log x}.$$

Here, we describe a variant of this method that computes endomorphism rings in proven probabilistic subexponential time, assuming only the GRH. The core idea remains to exploit complex multiplication theory to test orders using their ideal structure and evaluating corresponding isogenies.

However, our method differs from that of [3] in several aspects. First, we use a more direct, faster isogeny-computing routine which allows us to bring down the exponent in the overall complexity. In addition, to select which orders are to be tested, we give a generic lattice-ascending procedure; it is suited to work in general number fields which is a necessary step to generalize the whole procedure to higher-dimensional abelian varieties (see [2, Chapter 8] for details). Finally, to prove the complexity of the order-testing method, we adapt material from Seysen [17] and proofs due to Hafner and McCurley [10] to make use of a sharp bound derived from the GRH by Jao, Miller, and Venkatesan [12, Corollary 1.3].

All in all, on input an ordinary elliptic curve $\mathscr{E}$ defined over a finite field $\mathbb{F}_q$ our main algorithm returns the structure of its endomorphism ring $\operatorname{End}\mathscr{E}$ in proven (assuming the GRH) probabilistic time

$$L(q)^{1+o(1)} + L(q)^{1/\sqrt{2}+o(1)}$$

where the first term only accounts for factoring an integer less than $4q$ using the state-of-the-art proven method of Lenstra and Pomerance [15]; in other words, apart from that factorization, we adapted and proved under the GRH all parts of the heuristic method of [3] while improving its asymptotic complexity.

We stress that although its runtime is probabilistic and depends on the GRH, the output of our algorithm is unconditionally correct (it is a *Las Vegas* algorithm).

Section 2 fixes notations on endomorphism rings and orders. Section 3 then presents the order-testing method using relations. Section 4 gives the direct and fast isogeny-computing routine. Section 5 describes our lattice-ascending procedure and main algorithm. Section 6 proves that class groups are characterized by short relations. Section 7 finally shows how orders are determined by their class groups.

2 Background

Let $\mathscr{E}$ be an ordinary elliptic curve over a finite field $\mathbb{F}_q$. The Frobenius endomorphism π acts on geometric points of $\mathscr{E}$ by raising their coordinates to the q^{th} power; its characteristic polynomial $\chi_\pi(x)$ is of the form $x^2 - tx + q$ and computing the integer t is equivalent to finding the number of points on $\mathscr{E}$, namely $\chi_\pi(1)$, which can be done in deterministic polynomial time in $\log q$ as Schoof showed in [16].

Deuring proved in [8] that $\mathbb{Q}\otimes\operatorname{End}\mathscr{E} \simeq \mathbb{Q}(\pi)$. Since the number field $K = \mathbb{Q}[x]/(\chi_\pi(x))$ is isomorphic to $\mathbb{Q}(\pi)$, by computing the trace t we have already determined the endomorphism ring up to fractions. From now on, we make this isomorphism implicit by setting $\pi = x$.[1]

The number field K is called the CM field of $\mathscr{E}$; the implicit isomorphism maps $\operatorname{End}\mathscr{E}$ to an order in K so we have

$$\mathbb{Z}[\pi] \subseteq \operatorname{End}\mathscr{E} \subseteq \mathcal{O}_K$$

where $\mathcal{O}_K$ is the ring of integers of K. Conversely, Waterhouse proved in [21, Theorem 4.2] that all orders containing $\mathbb{Z}[\pi]$ arise as endomorphism rings. The index $[\mathcal{O}_K : \mathbb{Z}[\pi]]$ is essentially the square part of $\Delta = t^2 - 4q$ and measures how broad the search-range is: in the worst case, it can be exponential (in $\log q$).

The orders of K containing $\mathbb{Z}[\pi]$ form a finite lattice (in the set-theoretic sense) where $\mathcal{O}_K$ is the maximal order, $\mathbb{Z}[\pi]$ the minimal one, and $\operatorname{End}\mathscr{E}$ lies in between. Unfortunately it might have exponentially many orders so we need to devise a better way of finding $\operatorname{End}\mathscr{E}$ than testing each in turn; this is the purpose of the lattice-ascending algorithm of Section 5 which tests only polynomially many orders. For those orders $\mathcal{O}$, we *test* whether $\mathcal{O} \subseteq \operatorname{End}\mathscr{E}$ with the methodology of Section 3.

3 The CM approach

We now present the approach of [3] to testing whether $\mathcal{O} \subseteq \operatorname{End}\mathscr{E}$, in a somewhat more abstract flavor. For the theory of imaginary quadratic orders, we refer to [7].

It will be implicitly understood that we exclusively consider ideals of norm coprime to Δ, so that their images in $\mathbb{Z}[\pi]$ are unramified and invertible. Since every (invertible) ideal class

[1] The conjugate of x might equivalently be taken as π; this choice just needs to be fixed.

of each order containing $\mathbb{Z}[\pi]$ has a representative of this type, this has no effect on our use of class groups, which arises from the following result of CM theory (see [4, Chapter 8] for a concise proof).

Theorem 3.1. *For any ideal $\mathfrak{a}$ of $\operatorname{End}\mathcal{E}$, denote by $\phi_{\mathfrak{a}}$ the isogeny with kernel $\bigcap_{\alpha\in\mathfrak{a}} \ker\alpha$. The class group $\operatorname{cl}(\mathcal{O})$ acts faithfully and transitively on the set of isomorphism classes of elliptic curves with endomorphism ring $\mathcal{O}$ by $\mathfrak{a} : \mathcal{E} \mapsto \phi_{\mathfrak{a}}(\mathcal{E})$.*

Intuitively, it means that the structure of the class group dictates that of the isogeny graph; hence, by looking at the latter, we deduce things on the former and obtain information about the endomorphism ring. This action is effective, as Proposition 4.4 will show. In this setting, we formalize the notion of structure by the following concept.

Definition. *We define relations as multisets of ideals of $\mathbb{Z}[\pi]$. We say that a relation R holds in an order $\mathcal{O}$ (or that it is a relation of $\mathcal{O}$) if the product $\prod_{\mathfrak{a}\in R} \mathfrak{a}\mathcal{O}$ is trivial in $\operatorname{cl}(\mathcal{O})$; we say that it holds in the isogeny graph if the composition of the isogenies $\phi_{\mathfrak{a}\operatorname{End}\mathcal{E}}$ for $\mathfrak{a} \in R$ fixes $\mathcal{E}$.*

The theorem implies that a relation holds in $\operatorname{End}\mathcal{E}$ if and only if it holds in the isogeny graph, which gives a way to tell the endomorphism ring apart from other orders (we will see that $\phi_{\mathfrak{a}\operatorname{End}\mathcal{E}}$ can be computed without knowing $\operatorname{End}\mathcal{E}$).

To avoid testing all orders, we rely on this simple result from [7, Chapter 7]:

Lemma 3.2. *If a relation holds in some order, it holds in all orders containing it.*

Intuitively, as we ascend the lattice of orders, more and more relations hold, which also translates into class groups getting smaller. This is why we chose $\mathbb{Z}[\pi]$ to be the ring of our ideals: via the morphism $\mathfrak{a} \mapsto \mathfrak{a}\mathcal{O}$ we can map ideals of $\mathbb{Z}[\pi]$ to any order above in a way that induces surjective morphisms of class groups.

To search for the endomorphism ring $\operatorname{End}\mathcal{E}$ in the lattice, we will *test* whether orders $\mathcal{O}$ lie below it by selecting relations of them and checking whether they hold in the isogeny graph. Before we describe that procedure in detail, let us mention how to compute isogenies.

4 Computing the CM action

To make use of Theorem 3.1, we need to work with *isomorphism classes* of elliptic curves; for this, we rely on [7, Proposition 14.19] which states that two ordinary elliptic curves are isomorphic if and only if their cardinalities and j-invariants are the same. Computing the cardinality takes polynomial time, and since the j-invariant is a rational function in the coefficients of a Weierstrass equation, it does not take longer to evaluate it. From now on, it will be implicitly understood that we work with isomorphism classes via this representation.

To evaluate the action $\phi_{\mathfrak{a}}(\mathcal{E})$ of an ideal $\mathfrak{a}$, we combine classical tools as follows (each step is explained in detail on the following page). Note that by factoring the ideal $\mathfrak{a}$, one can reduce to the case where it is prime.

Algorithm 4.1.

```
 INPUT:  An elliptic curve $\mathcal{E}/\mathbb{F}_q$ with Frobenius polynomial $\chi_\pi$ and a prime ideal $\mathfrak{a}$.
OUTPUT:  The isogenous elliptic curve $\phi_{\mathfrak{a}}(\mathcal{E})$.
     1.  Find a basis $(P_i)$ of the $\ell$-torsion of $\mathcal{E}$ over $\mathbb{F}_{q^{\ell-1}}$ for $\ell = \operatorname{norm}(\mathfrak{a})$.
     2.  Write the matrix $M$ of the Frobenius endomorphism on $(P_i)$.
     3.  Compute the eigenspaces of $M \in \operatorname{Mat}_2(\mathbb{Z}/\ell\mathbb{Z})$.
     4.  Determine which is the kernel of the isogeny $\phi_{\mathfrak{a}}$.
     5.  Compute this isogeny.
```

Step 5 computes $\phi_{\mathfrak{a}}$ from its kernel, which Vélu's formulæ [20] do in $O(\ell)$ curve operations over $\mathbb{F}_{q^{\ell-1}}$. Step 4 uses the following idea from the SEA algorithm proved in [9, Stage 3]:

Proposition 4.2. *Let $\mathfrak{a}$ be an ideal of $\mathcal{O}$ of prime norm ℓ; write it as $\ell\mathcal{O} + u(\pi)\mathcal{O}$ where the polynomial u is an irreducible factor of $\chi_\pi \bmod \ell$. The kernel of the corresponding isogeny $\phi_{\mathfrak{a}}$ is an eigenspace of the Frobenius endomorphism, and we have $u = \operatorname{char}(\pi|_{\ker \phi_{\mathfrak{a}}})$.*

Since the map $\mathfrak{a} \mapsto \mathfrak{a}\mathcal{O}$ preserves the norm ℓ and polynomial u of ideals $\mathfrak{a}$ of $\mathbb{Z}[\pi]$, the order $\mathcal{O}$ need not be known to compute $\phi_{\mathfrak{a}\mathcal{O}}$, which we use for $\mathcal{O} = \operatorname{End}\mathcal{E}$.

Step 3 is classical and takes quasi-linear time in $\log \ell$; it outputs the $\mathbb{F}_q$-rational subgroups of $\mathcal{E}[\ell]$ isomorphic to $\mathbb{Z}/\ell\mathbb{Z}$. Step 2 decomposes $\pi(P_i)$ as $\sum_{j\in\{1,2\}} M_{ij}P_j$ for which a baby-step giant-step approach requires $O(\ell)$ operations in $\mathcal{E}/\mathbb{F}_{q^{\ell-1}}$.

Finally, Step 1 uses the fact that points of rational subgroups of order ℓ are defined over an extension of degree $\ell - 1$; it selects random ℓ^k-torsion points over this extension and lifts one along the other to obtain independent ℓ-torsion points. This procedure originates from [6, Theorem 1]; it does not only apply to elliptic curves, but we give a detailed algorithm specialized to this setting below.

Algorithm 4.3.

```
INPUT:   An elliptic curve $\mathcal{E}/\mathbb{F}_q$ with Frobenius polynomial $\chi_\pi$ and a prime $\ell$.
OUTPUT:  A basis of the $\ell$-torsion $\mathcal{E}[\ell]$ of $\mathcal{E}$ over $\mathbb{F}_{q^{\ell-1}}$.
  a.     Decompose $\#\mathcal{E}(\mathbb{F}_{q^{\ell-1}})$ as $m\ell^k$ where $\ell \nmid m$.
  b.     Let $P$ and $Q$ be $m$ times random points of $\mathcal{E}(\mathbb{F}_{q^{\ell-1}})$;
  c.     Compute the order $\ell^{k_P}$ of $P$ and $\ell^{k_Q}$ of $Q$ and assume $k_P \geq k_Q$.
  d.     Precompute the table $(i, i\,\ell^{k_P-1}P)$ for $i \in \mathbb{Z}/\ell\mathbb{Z}$.
  e.     For $j$ from $k_Q - 1$ down to $1$:
  f.         If $\ell^j Q = i\,\ell^{k_P-1}P$ for some $i$, set $Q \leftarrow Q - i\,\ell^{k_P-j-1}P$.
  g.         If $Q = 0_{\mathcal{E}}$ then go back to Step b.
  h.     Return $(\ell^{k_P-1}P, \ell^{k_Q-1}Q)$.
```

The cardinality of $\mathcal{E}(\mathbb{F}_{q^{\ell-1}})$ can be computed as $\operatorname{Res}_x(\chi_\pi(x), x^{\ell-1}-y)(1)$; since it is $O(q^\ell)$, extracting random points of it and multiplying them by m requires $O(\ell \log q)$ operations in $\mathbb{F}_{q^{\ell-1}}$. Similarly, both k_P and k_Q are bounded by $k = O(\ell \log q)$. The lookup in Step f is negligible if an efficient data structure such as a red-black tree is used to store the precomputed table of Step d. Finally, the probability of going back to Step b is $O(1/\ell)$ as proven in [6].

Using fast arithmetic, operations in $\mathbb{F}_{q^{\ell-1}}$ take at most $(\ell \log q)^{1+o(1)}$ time; hence:

Proposition 4.4. *Algorithm 4.1 returns the curve $\phi_{\mathfrak{a}\operatorname{End}\mathcal{E}}(\mathcal{E})$ isogenous to a prescribed curve $\mathcal{E}/\mathbb{F}_q$ in probabilistic time $O(\ell^{2+o(1)}\log^{2+o(1)} q)$, where $\ell = \operatorname{norm}(\mathfrak{a})$.*

5 Ascending the lattice of orders

Orders in an imaginary quadratic field K are of the form $\mathbb{Z} + f\mathcal{O}_K$ for some $f \in \mathbb{N}$ known as the conductor; consequently, inclusion of orders simply translates to divisibility of conductors. Those orders we are interested in contain $\mathbb{Z}[\pi]$, so their conductors divide the index $[\mathcal{O}_K : \mathbb{Z}[\pi]]$.

We will be ascending the lattice of orders one step at a time: each step consists in enumerating all orders lying directly above a prescribed order, that is, containing it with prime index ℓ. The possible values for ℓ are the prime factors of $[\mathcal{O}_K : \mathbb{Z}[\pi]]$ which can be listed by factoring (the square-part of) the discriminant Δ, for which the state-of-the-art proven method

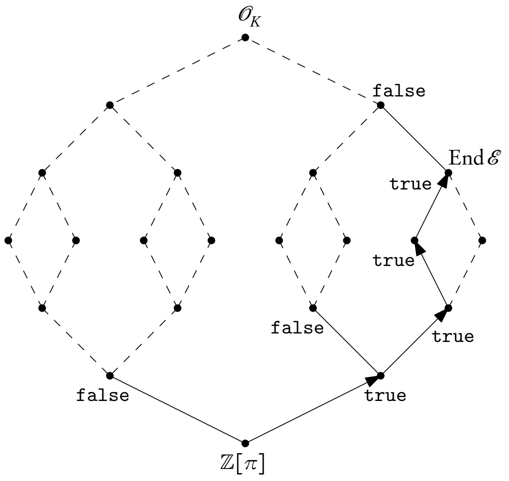

Figure 1: Locating End $\mathcal{E}$ by ascending a test-sequence of orders.

of Lenstra and Pomerance [15] uses $L(q)^{1+o(1)}$ operations. Enumerating orders above (resp. below) then simply amounts to dividing (resp. multiplying) the conductor by the possible ℓ's; naturally, since our orders are to contain $\mathbb{Z}[\pi]$, this is subject to the condition that the conductor remains a factor of the index $[\mathcal{O}_K : \mathbb{Z}[\pi]]$.

Our strategy to locate the endomorphism ring in this lattice by testing orders and ascending in corresponding directions works as follows: given some order $\mathcal{O}'$ contained in End $\mathcal{E}$ (starting with $\mathcal{O}' = \mathbb{Z}[\pi]$), find an order $\mathcal{O}$ directly above it that still lies below End $\mathcal{E}$; then replace $\mathcal{O}'$ by $\mathcal{O}$ and iterate the process. The ascension ends when no $\mathcal{O}$ directly above $\mathcal{O}'$ is contained in End $\mathcal{E}$; then, we must have End $\mathcal{E} \simeq \mathcal{O}'$. See Figure 1 where we start from the bottom and ascend towards orders $\mathcal{O}$ for which the statement $\mathcal{O} \subseteq$ End $\mathcal{E}$ holds.

We formalize this procedure into:

Algorithm 5.1.

```
INPUT:   An ordinary elliptic curve $\mathcal{E}$ over a finite field $\mathbb{F}_q$.
OUTPUT:  An order isomorphic to the endomorphism ring of $\mathcal{E}$.
1.  Compute the Frobenius polynomial $\chi_\pi(x)$ of $\mathcal{E}$.
2.  Factor the discriminant $\Delta$ and construct the order $\mathcal{O}' = \mathbb{Z}[\pi]$.
3.  For orders $\mathcal{O}$ directly above $\mathcal{O}'$:
4.      If $\mathcal{O} \subseteq \mathrm{End}\,\mathcal{E}$ set $\mathcal{O}' \leftarrow \mathcal{O}$ and go to Step 3.
5.  Return $\mathcal{O}'$.
```

Steps 1 and 2 only require polynomial time in $\log q$, except the factorization of the discriminant Δ, which uses $L(q)^{1+o(1)}$ operations.

Under the GRH, the following result will be established in Section 7.

Proposition 5.2 (GRH). *Let $\mathcal{O}$ be an order above $\mathbb{Z}[\pi]$. One can determine whether $\mathcal{O} \subseteq$ End $\mathcal{E}$ in probabilistic time $L(q)^{1/\sqrt{2}+o(1)}$ with failure probability $o(1/\log^2 q)$.*

The number of orders directly above $\mathbb{Z}[\pi]$ (to be tested in Step 4) is the number of prime factors of $[\mathcal{O}_K : \mathbb{Z}[\pi]]$ and it decreases as $\mathcal{O}'$ grows; the number of ascending steps (of times Step 3 is reached) is bounded by the sum of the exponents in the factorization of $[\mathcal{O}_K : \mathbb{Z}[\pi]]$

into prime powers. These two quantities are smaller than $\log_2 \Delta$ so the overall number of tests is at most quadratic in $\log q$. As a consequence, we have:

Theorem 5.3 (GRH). *The endomorphism ring of an ordinary elliptic curve defined over $\mathbb{F}_q$ can be computed, with failure probability $o(1)$, in probabilistic time $L(q)^{1+o(1)} + L(q)^{1/\sqrt{2}+o(1)}$ where the first term only accounts for the complexity of factoring the discriminant $\Delta = O(q)$.*

To unconditionally verify the output, we use the certification method of [3, Section 3.2], which is straightforwardly adapted to incorporate the isogeny-computing routine of Section 4 and the proof material of Section 6 and 7. Under the GRH, it then takes $L(q)^{1/\sqrt{2}+o(1)}$ operations to unconditionally check whether $\mathcal{O} = \operatorname{End} \mathcal{E}$ for a prescribed $\mathcal{O}$. We then obtain a so-called *Las Vegas* algorithm for which the above theorem holds without the failure probability statement.

The rest of this paper is devoted to the proof of Proposition 5.2.

6 Class groups from short relations

To test whether $\mathcal{O} \subseteq \operatorname{End} \mathcal{E}$ reliably, we *characterize* $\mathcal{O}$ by a set of relations R that hold in it but not collectively in any order not containing it. We will then test whether they hold in the isogeny graph, so we seek relations R for which the cost of computing the associated isogeny, roughly $\sum_{\mathfrak{a}\in R} \operatorname{norm}(\mathfrak{a})^2$, is small.

We start by bounding the norms of ideals to appear in our relations: form the set $\mathcal{B}$ of prime ideals $\mathfrak{p}$ of $\mathbb{Z}[\pi]$ with norm less than some integer N to be fixed later, and consider *smooth* ideals

$$\sigma(n) = \prod_{\mathfrak{p}\in\mathcal{B}} \mathfrak{p}^{n_\mathfrak{p}}$$

for vectors $n \in \mathbb{Z}^{\mathcal{B}}$. If $\sigma_{\mathcal{O}}(n)$ denotes the corresponding ideal class in $\operatorname{cl}(\mathcal{O})$, the kernel of the map $\sigma_{\mathcal{O}}$ is a lattice $\Lambda_{\mathcal{O}}$ in $\mathbb{Z}^{\mathcal{B}}$ consisting of all relations of $\mathcal{O}$ formed of ideals in $\mathcal{B}$: the coordinate $n_\mathfrak{p}$ is the multiplicity of the ideal $\mathfrak{p}$ in the relation. When $\sigma_{\mathcal{O}}$ is surjective, we have

$$\operatorname{cl}(\mathcal{O}) \simeq \mathbb{Z}^{\mathcal{B}}/\Lambda_{\mathcal{O}}.$$

Nothing of value is lost by only considering relations R of $\Lambda_{\mathcal{O}}$ since, assuming the GRH, Bach proved in [1] that $\sigma_{\mathcal{O}}$ is indeed surjective provided that $N \geq 12\log^2|\Delta|$.

The isogeny chain associated to a relation $n \in \Lambda_{\mathcal{O}}$ comprises at most $\|n\|_1 = \sum |n_\mathfrak{p}|$ isogenies of degree up to N so the complexity of evaluating it is crudely bounded by $\|n\|_1 N^{2+o(1)}$. This norm can be controlled by a result of Jao, Miller, and Venkatesan [12, Corollary 1.3] and more specifically its following specialization found in [5, Theorem 2.1].

Theorem 6.1. *Under the GRH, for all $\epsilon > 0$ there exists some $c > 1$ such that the following holds. Let $\mathcal{O}$ be an imaginary quadratic order, denote by D its discriminant, and let N and l be integers verifying*

$$N \geq \log^{2+\epsilon}|D| \qquad \text{and} \qquad l \geq c\frac{\log|D|}{\log\log|D|}.$$

If n is drawn uniformly at random from the set of vectors of $\mathbb{Z}^{\mathcal{B}}$ with norm l, the probability that the ideal $\sigma_{\mathcal{O}}(n)$ falls in any subset S of $\operatorname{cl}(\mathcal{O})$ is at least $\frac{1}{2}\frac{\#S}{\#\operatorname{cl}(\mathcal{O})}$.

Corollary 6.2 (GRH). *For $N = \log^{2+\epsilon}|D|$ the diameter of $\Lambda_{\mathcal{O}}$ is $o(\log^{4+\epsilon}|D|)$.*

Proof. To prove this, we construct a generating set for $\Lambda_{\mathscr{O}}$ formed by $O(\log^{2+\epsilon}|D|)$ relations of norm $o(\log^2|D|)$. Siegel showed in [18] that $\mathrm{cl}(\mathscr{O})$ is an abelian group of order $D^{1/2+o(1)}$ so there exist $O(\log|D|)$ ideal classes α_i such that $\mathbb{Z}^{\mathscr{B}}/\Lambda_{\mathscr{O}} \simeq \prod\langle\alpha_i\rangle$; we fix these and proceed to write a generating set for $\Lambda_{\mathscr{O}}$ consisting of:

- relations expressing that $\alpha_i^{\mathrm{ord}(\alpha_i)} = 1$;
- relations expressing the primes $\mathfrak{p} \in \mathscr{B}$ in terms of the α_i.

First define a map $\sigma_{\mathscr{O}}^{-1}$ by fixing a preimage of norm at most $c\log|D|/\log\log|D|$ for each ideal class; it exists by Theorem 6.1. Now use a double-and-add approach to ensure that norms remain small: for each i, express that $\alpha_i^{\mathrm{ord}(\alpha_i)} = 1$ by the relations

(i) $\sigma_{\mathscr{O}}^{-1}\left(\alpha_i^{2^j}\right) - 2\sigma_{\mathscr{O}}^{-1}\left(\alpha_i^{2^{j-1}}\right)$ for $j \in \{1, \dots, \lfloor\log_2 \mathrm{ord}(\alpha_i)\rfloor\}$;

(ii) $\sum_j b_j \sigma_{\mathscr{O}}^{-1}\left(\alpha_i^{2^j}\right)$ where b_j denotes the j^{th} least significant bit of $\mathrm{ord}(\alpha_i)$.

Now write each $\mathfrak{p} \in \mathscr{B}$ on the α_i by decomposing its class as a product $\prod \alpha_i^{n_i}$ where $n_i \in \{0, \dots, \mathrm{ord}(\alpha_i)\}$; noting $\delta_{\mathfrak{p}}$ the vector with coordinate one at $\mathfrak{p}$ and zero elsewhere, this gives the relations:

(iii) $\delta_{\mathfrak{p}} - \sum_i \sum_j c_{ij}\sigma_{\mathscr{O}}^{-1}\left(\alpha_i^{2^j}\right)$ where c_{ij} is the j^{th} least significant bit of n_i.

Preimages by $\sigma_{\mathscr{O}}$ have length $o(\log|D|)$ and there are at most $\sum\lfloor\log_2 \mathrm{ord}(\alpha_i)\rfloor = O(\log|D|)$ terms, therefore each such relation has length $o(\log|D|)^2$. □

To generate short relations, we simply plug this bound into the algorithm of Seysen [17] and rely on ingredients of Hafner and McCurley [10] for the proof. Note that Childs, Jao, and Soukharev [5] proposed a similar algorithm that finds one relation, while we seek several random relations to characterize the order $\mathscr{O}$.

Algorithm 6.3.

```
 INPUT:  An imaginary quadratic order $\mathscr{O}$ of discriminant $D$ and some $z > 0$.
OUTPUT:  A quasi-random relation $n \in \Lambda_{\mathscr{O}}$ with $\|n\|_1 = o(\log^{6+\epsilon}|D|)$.
     1.  Form the set $\mathscr{B}$ of primes $\mathfrak{p}$ of $\mathscr{O}$ with norm less than $N = L(q)^z$.
     2.  Draw uniformly at random a vector $x \in \mathbb{Z}^{\mathscr{B}}$ with coordinates
         $|x_{\mathfrak{p}}| < \log^{4+\epsilon}|D|$ if $\mathrm{norm}(\mathfrak{p}) < \log^{2+\epsilon}|D|$, else $x_{\mathfrak{p}} = 0$.
     3.  Compute the reduced ideal representative $\mathfrak{a}$ of $\sigma_{\mathscr{O}}(x)$.
     4.  If $\mathfrak{a}$ factors over $\mathscr{B}$ as $\prod \mathfrak{p}^{y_{\mathfrak{p}}}$ then return the vector $x - y$.
     5.  Otherwise, go back to Step 2.
```

Proposition 6.4 (GRH). *Let $\mathscr{O}$ be an order containing $\mathbb{Z}[\pi]$; its discriminant D is then at most $\Delta = O(q)$. The algorithm above requires $L(q)^{z+o(1)} + L(q)^{1/(4z)+o(1)}$ operations to find a relation of $\mathscr{O}$ whose associated isogeny can be computed in time $L(q)^{2z+o(1)}$.*

Proof. Step 4 consists in testing the smoothness of (the norm of) $\mathfrak{a}$; Lenstra, Pila, and Pomerance [14, Corollary 1.2] proved this requires $\exp\left(\log^{2/3+o(1)}N\right)\log^3 q$ operations, that is, $L(q)^{o(1)}$ since $N = L(q)^z$. The probability that this factorization is successful, in other words, that the norm of $\mathfrak{a}$ is N-smooth is $L(q)^{1/(4z)+o(1)}$ provided that it behaves as a random integer; this follows directly from combining the corollary above with [17, Proposition 4.4]; see also [10]. The relation involves $o(\log^{4+2+\epsilon} q)$ ideals of norm up to $L(q)^z$, whence the time bound for evaluating the associated isogeny by Proposition 4.4. □

The relations we generate discriminate between orders with distinct class groups:

Lemma 6.5 (GRH). *Take any two orders $\mathcal{O}$ and $\mathcal{O}'$; a relation of $\mathcal{O}$ generated by the algorithm above has a probability $[\Lambda_{\mathcal{O}} : \Lambda_{\mathcal{O}} \cap \Lambda_{\mathcal{O}'}]^{-1} + o(1)$ of holding in $\mathcal{O}'$.*

Proof. This follows directly from [10, Lemma 2] adapted to the context of our algorithm, which proves the quasi-randomness of the relations it generates. □

7 Orders from class groups

Now, to finally prove Proposition 5.2, let us establish the correctness and runtime of the following algorithm.

Algorithm 7.1.

```
INPUT:   An ordinary elliptic curve $\mathcal{E}/\mathbb{F}_q$ and an order $\mathcal{O} \supseteq \mathbb{Z}[\pi]$.
OUTPUT:  Whether $\mathcal{O} \subseteq \operatorname{End}\mathcal{E}$, with failure probability $o(1/\log^2 q)$.
    1. Generate a set of $3\log\log q$ relations of $\mathcal{O}$ with Algorithm 6.3.
    2. If one does not hold in the isogeny graph, return false.
    3. Check whether $\mathcal{O} \subseteq \operatorname{End}\mathcal{E}$ locally at 2 and 3; if not, return false.
    4. Return true.
```

By Proposition 6.4, Step 1 takes $L(q)^{z+o(1)} + L(q)^{1/(4z)+o(1)}$ time to find relations whose associated isogenies are then evaluated by Step 2 in $L(q)^{2z+o(1)}$ operations. To balance these quantities, we set $z = 1/2\sqrt{2}$ which yields an overall complexity of $L(q)^{1/\sqrt{2}+o(1)}$.

The correctness follows from Lemma 3.2 and Theorem 3.1, in that Steps 1 and 2 determine whether $\Lambda_{\mathcal{O}} \subseteq \Lambda_{\operatorname{End}\mathcal{E}}$; the failure probability is at most $(2+o(1))^{-3\log\log q} = o(1/\log^2 q)$, by Lemma 6.5 applied to $\mathcal{O}' = \operatorname{End}\mathcal{E}$. The proposition below proves that, combined with Step 3, this determines whether $\mathcal{O} \subseteq \operatorname{End}\mathcal{E}$.

Proposition 7.2. *Let $\mathcal{O}$ and $\mathcal{O}'$ be two orders in an imaginary quadratic field K. The lattice $\Lambda_{\mathcal{O}'}$ contains $\Lambda_{\mathcal{O}}$ if and only if the order $\mathcal{O}'$ contains $\mathcal{O}$ or:*

1. *$K = \mathbb{Q}(\sqrt{-4})$ and $\mathcal{O}'$ has conductor 2;*
2. *$K = \mathbb{Q}(\sqrt{-3})$ and $\mathcal{O}'$ has conductor 2 or 3;*
3. *The prime 2 splits in K and $\mathcal{O}'$ has index 2 in some order above $\mathcal{O}$ of odd conductor.*

Intuitively, this means that identifying orders by their class groups has a single blind spot locally at 2 and 3 where the two biggest orders cannot be distinguished; Step 3 is thus needed to ensure that the endomorphism ring is accurately determined even amongst those orders with identical class groups. This is a straightforward refinement of [3, Proposition 5], but we give a proof below for completeness.

Proof. Denote by $S_{\mathcal{O}}$ (resp. $S_{\mathcal{O}'}$) the set of primes ℓ that split into principal ideals in $\mathcal{O}$ (resp. $\mathcal{O}'$). Using relations formed of a single prime ideal, we see that $\Lambda_{\mathcal{O}} \subseteq \Lambda_{\mathcal{O}'}$ implies $S_{\mathcal{O}} \subseteq S_{\mathcal{O}'}$. Now, $S_{\mathcal{O}}$ (resp. $S_{\mathcal{O}'}$) is also the set of primes that split completely in the ring class field $L_{\mathcal{O}}$ of $\mathcal{O}$ (resp. $L_{\mathcal{O}'}$). By Chebotarev's density theorem, $S_{\mathcal{O}} \subseteq S_{\mathcal{O}'}$ thus implies $L_{\mathcal{O}'} \subseteq L_{\mathcal{O}}$, which means that the class field theory conductor $\mathfrak{f}(L_{\mathcal{O}'}/K)$ of $L_{\mathcal{O}'}$ divides $\mathfrak{f}(L_{\mathcal{O}}/K)$.

This conductor $\mathfrak{f}(L_{\mathcal{O}}/K)$ is related to $f_{\mathcal{O}}$ by (see [7, Exercises 9.20–9.23]):

$$\mathfrak{f}(L_{\mathcal{O}}/K) = \begin{cases} \mathcal{O}_K, & \text{when } K = \mathbb{Q}(\sqrt{-4}) \text{ and } f_{\mathcal{O}} = 2, \\ \mathcal{O}_K, & \text{when } K = \mathbb{Q}(\sqrt{-3}) \text{ and } f_{\mathcal{O}} = 2 \text{ or } 3, \\ u\,\mathcal{O}_K, & \text{when 2 splits in } K \text{ and } f_{\mathcal{O}} = 2u \text{ with } u \text{ odd}, \\ f_{\mathcal{O}}\mathcal{O}_K, & \text{otherwise.} \end{cases}$$

Naturally, the same stands for $\mathcal{O}'$. In the fourth case, the fact that $\mathfrak{f}(L_{\mathcal{O}}/K)$ divides $\mathfrak{f}(L_{\mathcal{O}'}/K)$ implies that $f_{\mathcal{O}'}$ divides $f_{\mathcal{O}}$, in other words $\mathcal{O} \subseteq \mathcal{O}'$; the three other cases correspond, in order, to the exceptions listed in the proposition. □

Finally, let us address Step 3. To check whether $\mathcal{O} \subseteq \operatorname{End}\mathscr{E}$ locally at some prime p, one uses a method of Kohel [13] known as "climbing the volcano", which can be done in the traditional "blind" way by following three p-isogeny paths from $\mathscr{E}$ and seeing which hits the "floor of rationality" first, or using the more advanced technique of [11] to directly determine the kernel of the ascending p-isogeny by pairing computations. Eventually, both methods return the valuation at p of the conductor of $\operatorname{End}\mathscr{E}$ by computing at most $O(\operatorname{val}_p[\mathcal{O}_K : \mathbb{Z}[\pi]])$ isogenies of degree p; since we use $p = 2, 3$, this takes polynomial time in $\log q$.

Acknowledgments

This work owes everything to the numerous fruitful discussions the author had with Karim Belabas, Andreas Enge, Pierrick Gaudry, Andrew V. Sutherland, and Emmanuel Thomé. This paper also benefited from the helpful remarks of Tanja Lange and of the anonymous referee.

References

[1] Eric Bach. Explicit bounds for primality testing and related problems. *Mathematics of Computation*, 55(191):355–380, 1990.

[2] Gaetan Bisson. Endomorphism rings in cryptography. PhD thesis, Eindhoven University of Technology and Institut National Polytechnique de Lorraine, 2011.

[3] Gaetan Bisson and Andrew V. Sutherland. Computing the endomorphism ring of an ordinary elliptic curve over a finite field. *Journal of Number Theory*, 131(5):815–831, 2011.

[4] John W. S. Cassels and Albrecht Fröhlich. *Algebraic Number Theory*. Academic Press, 1967.

[5] Andrew M. Childs, David Jao, and Vladimir Soukharev. Constructing elliptic curve isogenies in quantum subexponential time, 2010. Preprint available at `http://arxiv.org/abs/1012.4019`.

[6] Jean-Marc Couveignes. Linearizing torsion classes in the Picard group of algebraic curves over finite fields. *Journal of Algebra*, 301:2085–2118, 2009.

[7] David A. Cox. *Primes of the form $x^2 + ny^2$*. Pure and Applied Mathematics. John Wiley & Sons, 1989.

[8] Max Deuring. Die Typen der Multiplikatorenringe elliptischer Funktionenkörper. *Abhandlungen aus dem mathematischen Seminar der hamburgischen Universität*, 14:197–272, 1941.

[9] Steven D. Galbraith, Florian Hess, and Nigel P. Smart. Extending the GHS Weil descent attack. In *Advances in Cryptology—EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 29–44. Springer, 2002.

[10] James L. Hafner and Kevin S. McCurley. A rigorous subexponential algorithm for computing in class groups. *Journal of the American Mathematical Society*, 2(4):837–850, 1989.

[11] Sorina Ionica and Antoine Joux. Pairing the volcano. In *Algorithmic Number Theory Symposium—ANTS IX*, volume 6197 of *Lecture Notes in Computer Science*, pages 201–218. Springer, 2010.

[12] David Jao, Stephen D. Miller, and Ramarathnam Venkatesan. Expander graphs based on GRH with an application to elliptic curve cryptography. *Journal of Number Theory*, 129(6):1491–1504, 2009.

[13] David Kohel. Endomorphism rings of elliptic curves over finite fields. PhD thesis, University of California at Berkeley, 1996.

[14] Hendrik W. Lenstra, Jr., Jonathan Pila, and Carl Pomerance. A hyperelliptic smoothness test, I. *Philosophical Transactions of the Royal Society of London, A*, 345(1676):397–408, 1993.

[15] Hendrik W. Lenstra, Jr. and Carl Pomerance. A rigorous time bound for factoring integers. *Journal of the American Mathematical Society*, 5(3):483–516, 1992.

[16] René Schoof. Counting points on elliptic curves over finite fields. *Journal de Théorie des Nombres de Bordeaux*, 7:219–254, 1995.

[17] Martin Seysen. A probabilistic factorization algorithm with quadratic forms of negative discriminant. *Mathematics of Computation*, 48(178):757–780, 1987.

[18] Carl Ludwig Siegel. Über die Classenzahl quadratischer Zahlkörper. *Acta Arithmetica*, 1:83–86, 1935.

[19] Andrew V. Sutherland. Computing Hilbert class polynomials with the Chinese Remainder Theorem. Available at `http://arxiv.org/abs/0903.2785`, 2009.

[20] Jacques Vélu. Isogénies entre courbes elliptiques. *Comptes Rendus de l'Académie des Sciences de Paris*, 273, 1971.

[21] William C. Waterhouse. Abelian varieties over finite fields. *Annales Scientifiques de l'École Normale Supérieure*, 4(2):521–560, 1969.