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 [18] requires an endomorphism-ring-computing subroutine. They are also potentially relevant security parameters in certain cryptographic applications.

They were first studied by Kohel [12] 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 [2] 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; it "ascends" the lattice of orders in a generic manner, and "tests" orders using their class group structure. The lattice-ascending procedure is suited to work in general number fields, which is a necessary step for generalizing this algorithm to higher-dimensional abelian varieties; for now, only the method of Eisenträger and Lauter [7] and that of Wagner [20] apply to this setting but they are both of exponential nature. To prove the complexity of the order-testing method, we adapt material from Seysen [16] and proofs due to Hafner and McCurley [9] to make use of a sharp bound derived from the GRH by Jao, Miller, and Venkatesan [11, Corollary 1.3].

Additionally, we use a more direct, faster isogeny-computing routine than [2] which allows us to bring down the exponent in the complexity. Explicitly, on input an ordinary elliptic curve $\mathcal{E}$ defined over a finite field $\mathbb{F}_q$ our main algorithm outputs the structure of its endomorphism ring $\operatorname{End}\mathcal{E}$ in proven (under the GRH) probabilistic time

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

where the first term only accounts for the cost of factoring of a certain integer less than $4q$ using the state-of-the-art proven method of Lenstra and Pomerance [14]; in other words, apart from that factorization, we were able to adapt and prove under the GRH all parts of the heuristic subexponential method above while improving its asymptotic complexity.

Section 2 fixes notations on endomorphism rings and orders. Section 3 then presents the order-testing method using "relations". Section 4 gives the direct-but-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 $\mathcal{E}$ be an ordinary elliptic curve defined over a finite field $\mathbb{F}_q$. The Frobenius endomorphism π acts on geometric points of $\mathcal{E}$ 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 the curve, namely $\chi_\pi(1)$. Schoof showed in [15] how this can be done in deterministic polynomial time in the size of the base field, $\log q$.

Many endomorphisms stem from the Frobenius endomorphism, as Deuring proved in [6] that $\mathbb{Q} \otimes \operatorname{End}\mathcal{E} \simeq \mathbb{Q}(\pi)$. 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 $\mathcal{E}$; the implicit isomorphism maps $\operatorname{End}\mathcal{E}$ to an order in K so we have

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

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 the discriminant $\Delta = t^2 - 4q$; this 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 $\operatorname{End}\mathcal{E}$ lies in between. Unfortunately it might have exponentially many orders so we need to devise a better way of finding $\operatorname{End}\mathcal{E}$ 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 $\mathcal{O} \subseteq \operatorname{End}\mathcal{E}$ with the methodology of Section 3 which we develop in Sections 6 and 7.

---

[1]The conjugate of x might equivalently be taken as π; this choice just needs to be made once and for all.

3 The CM approach

We now present the approach of [2] to testing whether $\mathcal{O} \subseteq \operatorname{End} \mathcal{E}$ in a somewhat more abstract flavor. For the theory of imaginary quadratic orders, we refer to [5].

In this paper, it is 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 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.

Theorem 3.1. *When $\mathfrak{a}$ is an ideal of $\operatorname{End} \mathcal{E}$, denote by $\phi_{\mathfrak{a}}$ the isogeny with kernel $\bigcap_{\alpha \in \mathfrak{a}} \ker \alpha$. The ideal 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, the structure of the class group dictates that of the isogeny graph; hence, by looking at the latter, we might deduce things on the former and obtain information about the endomorphism ring. This action is effective, as embodied in Proposition 4.4. 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 of the lattice (we will see in the next section 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 [5, Chapter 7]:

Lemma 3.2. *If a relation holds in some order, it also does 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 [5, 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. In the following, it is 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:

Algorithm 4.1.
    INPUT: *An elliptic curve $\mathcal{E}/\mathbb{F}_q$ with Frobenius polynomial χ_π and an ideal $\mathfrak{a}$.*
    OUTPUT: *The isogenous elliptic curve $\phi_{\mathfrak{a}}(\mathcal{E})$.*

    1.   Find a basis (P_i) of the ℓ-torsion of $\mathcal{E}$ over $\mathbb{F}_{q^{\ell-1}}$ where $\ell = \operatorname{norm}(\mathfrak{a})$.
    2.   Write the matrix M of the Frobenius endomorphism on the basis (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 [19] do in $O(\ell)$ curve operations over $\mathbb{F}_{q^{\ell-1}}$. Step 4 relies on an idea from the SEA algorithm found in [8, 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 characteristic polynomial of the restriction to the kernel of $\phi_{\mathfrak{a}}$ of the Frobenius endomorphism is u.*

Since the map $\mathfrak{a} \mapsto \mathfrak{a}\mathcal{O}$ from ideals of $\mathbb{Z}[\pi]$ preserves their norm ℓ and polynomial u, there is no need to know $\mathcal{O}$ to compute $\phi_{\mathfrak{a}\mathcal{O}}$; this is particularly useful for $\mathcal{O} = \operatorname{End}\mathcal{E}$.

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 $\mathcal{E}/\mathbb{F}_{q^{\ell-1}}$. 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 $\mathbb{Z}/\ell\mathbb{Z}$.

Finally, Step 1 uses the fact that points of rational subgroups of order ℓ are necessarily defined over an extension of degree $\ell - 1$; it proceeds by selecting random ℓ^k-torsion points over this extension and lifting one along the other to obtain independent ℓ-torsion points. This idea originates from [4, Theorem 1] to which we refer for details.

Algorithm 4.3.
    INPUT: *An elliptic curve $\mathcal{E}/\mathbb{F}_q$ with Frobenius polynomial χ_π and a prime ℓ.*
    OUTPUT: *A basis of the ℓ-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 ℓ^{k_P} of P and ℓ^{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 [4].

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

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; inclusion of orders corresponds 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 of Lenstra and Pomerance [14] 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 $\operatorname{End} \mathcal{E}$ (we start with $\mathcal{O}' = \mathbb{Z}[\pi]$), find some order $\mathcal{O}$ directly above $\mathcal{O}'$ which lies below $\operatorname{End} \mathcal{E}$; then replace $\mathcal{O}'$ by $\mathcal{O}$ and iterate the process. The ascension ends when no $\mathcal{O}$ is contained in $\operatorname{End} \mathcal{E}$; then, we must have $\operatorname{End} \mathcal{E} \simeq \mathcal{O}'$. See Figure 1 where we start from the bottom and ascend towards orders $\mathcal{O}$ for which the statement $\mathcal{O} \subseteq \operatorname{End} \mathcal{E}$ 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 Δ and construct the order $\mathcal{O}' = \mathbb{Z}[\pi]$.*
  3. *For orders $\mathcal{O}$ directly above $\mathcal{O}'$:*
  4.   *If $\mathcal{O} \subseteq \operatorname{End} \mathcal{E}$ set $\mathcal{O}' \leftarrow \mathcal{O}$ and go to Step 3.*
  5. *Return $\mathcal{O}'$.*

Steps 1 and 2 are classical and only require polynomial time in $\log q$, except the factorization of Δ which takes $L(q)^{1+o(1)}$ time. Under the GRH, we will later prove:
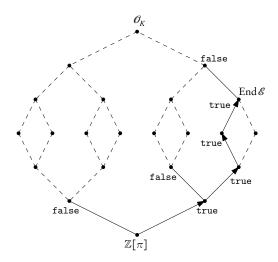
Figure 1: Locating $\text{End}\,\mathcal{E}$ by ascending a test-sequence of orders.

Proposition 5.2 (GRH). *Let $\mathcal{O}$ be an order above $\mathbb{Z}[\pi]$. One can determine whether $\mathcal{O} \subseteq \text{End}\,\mathcal{E}$ 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 $\Delta = O(q)$.*

The output may be unconditionally verified using the certification method of [2, Section 3.2]. This probabilistic procedure can be adapted to use the isogeny-computing routing of Section 4 and the proof material of Section 6; under the GRH, it then requires $L(q)^{1/\sqrt{2}+o(1)}$ operations. As a result, we obtain an algorithm for which the above theorem holds without the "failure probability" statement; this is sometimes called a *Las Vegas* algorithm.

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

6 Class groups from short relations

To test whether $\mathcal{O} \subseteq \text{End}\,\mathcal{E}$ reliably, we *characterize* $\mathcal{O}$ by a set of relations R that hold in it but not collectively in any order of the lattice not containing it. We will then test whether they hold in the isogeny graph, so we seek relations

R for which the (quasi-)quadratic cost of computing the associated isogeny $\sum_{\mathfrak{a} \in R} \operatorname{norm}(\mathfrak{a})^2$ 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 [11, Corollary 1.3] and more specifically its following specialization found in [3, Theorem 2.1].

Theorem 6.1. *Under the GRH, for all positive numbers ϵ there exists a constant $c > 1$ such that, for any imaginary quadratic order $\mathcal{O}$ of discriminant D and integers $N \geq \log^{2+\epsilon} |D|$ and*

$$l \geq c \frac{\log |D|}{\log \log |D|},$$

the probability, for random vectors $n \in \mathbb{Z}^{\mathcal{B}}$ of norm l, that the ideal class $\sigma_{\mathcal{O}}(n)$ 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 the lattice $\Lambda_{\mathcal{O}}$ is $o(\log^{4+\epsilon} |D|)$.*

Proof. To prove this, we construct a generating set for $\Lambda_{\mathcal{O}}$ formed by $O(\log^{2+\epsilon} |D|)$ relations of norm $o(\log^2 |D|)$. Siegel showed in [17] that $\operatorname{cl}(\mathcal{O})$ is an abelian group of order $D^{1/2+o(1)}$ so there exist $O(\log |D|)$ ideal classes α_i such that $\mathbb{Z}^{\mathcal{B}}/\Lambda_{\mathcal{O}} \simeq \prod \langle \alpha_i \rangle$; we fix these and proceed to write a generating set for $\Lambda_{\mathcal{O}}$ consisting of:

- relations expressing that $\alpha_i^{\operatorname{ord}(\alpha_i)} = 1$;

- relations expressing the primes $\mathfrak{p} \in \mathcal{B}$ in terms of the α_i.

First define a map $\sigma_{\mathcal{O}}^{-1}$ 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 $\alpha_i^{\operatorname{ord}(\alpha_i)} = 1$ by the relations

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

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

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

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

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

To generate short relations, we simply plug this bound into the algorithm of Seysen [16] and rely on ingredients of Hafner and McCurley [9] for the proof. Note that Childs, Jao, and Soukharev [3] proposed a similar algorithm for finding one relation, while we seek several random relations in order to characterize the order $\mathcal{O}$.

Algorithm 6.3.

    INPUT: *An imaginary quadratic order $\mathcal{O}$ of discriminant D.*

    OUTPUT: *A quasi-random relation $n \in \Lambda_{\mathcal{O}}$ with $\|n\|_1 = o(\log^{6+\epsilon}|D|)$.*

    1. Form the set $\mathcal{B}$ of primes $\mathfrak{p}$ of $\mathcal{O}$ with norm less than $N = L(q)^z$.

    2. Draw uniformly at random a vector $x \in \mathbb{Z}^{\mathcal{B}}$ with coordinates $|x_{\mathfrak{p}}| < \log^{4+\epsilon}|D|$ if $\operatorname{norm}(\mathfrak{p}) < \log^{2+\epsilon}|D|$, else $x_{\mathfrak{p}} = 0$.

    3. Compute the reduced ideal representative $\mathfrak{a}$ of $\sigma_{\mathcal{O}}(x)$.

    4. If $\mathfrak{a}$ factors over $\mathcal{B}$ 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 $\mathcal{O}$ 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 $\mathcal{O}$ 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 [13, 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 [16, Proposition 4.4]; see also [9]. 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. $\qquad\square$

Hopefully, 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 also holding in $\mathcal{O}'$.*

Proof. This follows directly from [9, Lemma 2] adapted to the context of our algorithm, which proves the quasi-randomness of the relations it generates. $\qquad\square$

7 Orders from class groups

Our proof of Proposition 5.2 now boils down to exhibiting the following.

Algorithm 7.1.
   INPUT: *An ordinary elliptic curve $\mathcal{E}/\mathbb{F}_q$ and an order $\mathcal{O} \supseteq \mathbb{Z}[\pi]$.*
   OUTPUT: *Whether $\mathcal{O} \subseteq \text{End}\,\mathcal{E}$.*

   1. *Compute a set of $3 \log\log q$ relations of $\mathcal{O}$.*
   2. *If one does not hold in the isogeny graph, return **false**.*
   3. *Check whether $\mathcal{O} \subseteq \text{End}\,\mathcal{E}$ locally at 2 and 3; if not, return **false**.*
   4. *Return **true**.*

By Proposition 6.4, Step 1 requires $L(q)^{z+o(1)} + L(q)^{1/(4z)+o(1)}$ operations to find relations whose associated isogenies are then evaluated by Step 2 in $L(q)^{2z+o(1)}$. To balance these quantities, we set $z = 1/2\sqrt{2}$ which gives 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 $\Lambda_\mathcal{O} \subseteq \Lambda_{\text{End}\,\mathcal{E}}$; the probability of failure is at most $(2 + o(1))^{-3\log\log q} = o(1/\log^2 q)$, by Lemma 6.5 applied to $\mathcal{O}' = \text{End}\,\mathcal{E}$. The proposition below argues that, combined with Step 3, this really determines whether $\mathcal{O} \subseteq \text{End}\,\mathcal{E}$.

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 required in our algorithm to ensure it exactly determines the endomorphism ring even amongst those orders with identical class groups. This statement is a straightforward refinement of [2, Proposition 5]; we nevertheless give the 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 that $f_\mathcal{O}$ of $\mathcal{O}$ as follows (see [5, Exercises 9.20–9.23]).

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

9

Naturally, the same stands for $\mathcal{O}'$. In the latter 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. $\square$

Finally, let us address Step 3. To check whether $\mathcal{O} \subseteq \operatorname{End}\mathcal{E}$ locally at some prime p, one uses a method of Kohel [12] known as "climbing the volcano", which can be done in the traditional "blind" way by following three p-isogeny paths from $\mathcal{E}$ and seeing which hits the "floor of rationality" first, or using the more advanced technique of [10] 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 $\operatorname{End}\mathcal{E}$ 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

References

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

[2] 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.

[3] 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`.

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

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

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

[7] Kirsten Eisenträger and Kristin Lauter. A CRT algorithm for constructing genus 2 curves over finite fields. In François Rodier and Serge Vladut, editors, *Arithmetic, Geometry and Coding Theory–AGCT 10*, 2004.

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

[9] 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.

[10] Sorina Ionica and Antoine Joux. Pairing the volcano. In *Algorithmic Number Theory Symposium—ANTS IX*, volume 6197 of *LNCS*, pages 201–218. Springer, 2010.

[11] 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.

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

[13] 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.

[14] 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.

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

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

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

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

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

[20] Markus Wagner. Über Korrespondenzen zwischen algebraischen Funktionenkörper. PhD thesis of the Technische Universität Berlin, 2009.

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