

CONSTRUCTING ISOGENIES BETWEEN ELLIPTIC CURVES OVER FINITE FIELDS

STEVEN D. GALBRAITH

Abstract

Let E_1 and E_2 be ordinary elliptic curves over a finite field \mathbb{F}_p such that $\#E_1(\mathbb{F}_p) = \#E_2(\mathbb{F}_p)$. Tate's isogeny theorem states that there is an isogeny from E_1 to E_2 which is defined over \mathbb{F}_p . The goal of this paper is to describe a probabilistic algorithm for constructing such an isogeny.

The algorithm proposed in this paper has exponential complexity in the worst case. Nevertheless, it is efficient in certain situations (that is, when the class number of the endomorphism ring is small). The significance of these results to elliptic curve cryptography is discussed.

1. Introduction

This paper concerns some computational problems related to isogenies between elliptic curves over finite fields. The primary goal is to give an algorithmic solution to Tate's isogeny theorem [34], which states that two elliptic curves E_1 and E_2 over a finite field \mathbb{F}_q are isogenous over \mathbb{F}_q if and only if they have the same number of \mathbb{F}_q -points.

Schoof [28] proposed a polynomial-time algorithm for counting the number of points on an elliptic curve over a finite field. There has been a considerable amount of research building on Schoof's idea (for instance, by Atkin [1], [2], Elkies [10], [11], Couveignes [8], Couveignes and Morain [9], Lercier [19], and Lercier and Morain [20]). This means that there is an efficient solution to the problem of determining whether two elliptic curves over \mathbb{F}_q are isogenous; namely, compute the number of points on each curve and see if the total is the same.

For this paper, we concentrate on the case of ordinary elliptic curves that are defined over \mathbb{F}_p . We make some comments about the case of supersingular curves and non-prime finite fields later in this section. More specifically, we describe a probabilistic algorithm for solving the following problem.

Problem 1. Let E_1 and E_2 be ordinary elliptic curves over \mathbb{F}_p such that $\#E_1(\mathbb{F}_p) = \#E_2(\mathbb{F}_p)$. Construct an isogeny $\varphi : E_1 \rightarrow E_2$ which is defined over \mathbb{F}_p .

For some applications one might want the isogeny φ of minimal possible degree. The method discussed in this paper will not necessarily produce such an isogeny; however, it will produce one of relatively smooth degree.

The main result of this paper is the following theorem (in all complexity estimates a unit cost is assumed for all operations in \mathbb{F}_p).

The author thanks the EPSRC for their support.

Received 6 October 1998, revised 19 April 1999; published 27 September 1999.

1991 Mathematics Subject Classification 11G20, 11Y16, 14K02

© 1999, Steven D. Galbraith

Theorem 1. Let E_1 and E_2 be ordinary elliptic curves over \mathbb{F}_p such that $\#E_1(\mathbb{F}_p) = \#E_2(\mathbb{F}_p)$. Assuming the truth of the Riemann hypothesis for imaginary quadratic number fields, then the probabilistic algorithm proposed in Section 4 will construct an isogeny $\varphi : E_1 \rightarrow E_2$ over the field \mathbb{F}_p . In the worst case, the algorithm requires expected time $O(p^{3/2} \ln p)$ and expected space $O(p \ln p)$.

In certain special cases (for instance, when the endomorphism rings of the elliptic curves have small class number) the algorithm runs in polynomial time. We make some comments in Section 8 about why we do not expect a polynomial-time solution to this problem in general.

For some applications one is concerned with the group structure of elliptic curves over finite fields. As part of our investigation we also study the following two problems.

Problem 2. Let E_1 and E_2 be ordinary elliptic curves over \mathbb{F}_p such that $E_1(\mathbb{F}_p)$ and $E_2(\mathbb{F}_p)$ have the same number of points. Construct a non-trivial group homomorphism between them which may be evaluated in polynomial time.

Problem 3. Let E be an ordinary elliptic curve over \mathbb{F}_p . Construct an isogenous elliptic curve E' over \mathbb{F}_p such that the group $E'(\mathbb{F}_p)$ is cyclic, and construct an \mathbb{F}_p -isogeny $\varphi : E \rightarrow E'$.

Of course, for most elliptic curves over \mathbb{F}_p , Problem 2 can be solved in various elementary ways. However, we are particularly interested in the case where $\#E(\mathbb{F}_p)$ has a large prime divisor and where the group homomorphism has small kernel and cokernel. We discuss Problem 2 in Section 7.

The methods used in this paper would generalise to the case of elliptic curves over non-prime finite fields \mathbb{F}_{p^m} . When p is large then the algorithm is essentially unchanged. When p is small it would be necessary to utilise the methods of Couveignes [8] and Lercier [19] (see [20] for a nice comparison) for computing isogenies.

For the case of supersingular curves an algorithm to solve Problem 1 is implicit in the work of Mestre [21]. The main simplification in the supersingular case is that all isomorphism classes of elliptic curves in a given isogeny class can be obtained by composing isogenies of any fixed prime degree (e.g., the prime 2, since all supersingular curves modulo p have a rational 2-torsion point). Therefore, there is an algorithm to construct an isogeny between any two supersingular curves with the same number of points. Since the number of supersingular elliptic curves in characteristic p is bounded by $\frac{p+1}{12} + 1$ [32, Theorem 4.1(c)], one can show that the algorithm runs in time $O(p \ln p)$ and requires space $O(p)$. For more details about computations with supersingular curves see [14] and [21].

The paper is organised as follows: In Section 2 we discuss the relevance of the results to the study of elliptic curve cryptography. In Section 3 we describe some of the necessary background results (more background information is given in Appendix A). The proof of Theorem 1 is split across several sections. Section 4 describes the algorithm and explains why it terminates. Sections 5 and 6 are concerned with an analysis of the complexity of the algorithm.

2. Application to cryptography

Elliptic curves over finite fields are being studied intensively with an eye to their use in cryptography. Given an elliptic curve E/\mathbb{F}_q and a point $P = (x, y) \in E(\mathbb{F}_q)$, the elliptic

curve discrete logarithm problem is the following: given a point $Q = (x', y')$ lying in the subgroup generated by P , find an integer λ such that $Q = \lambda P$.

The elliptic curve discrete logarithm problem is known to be easy to solve in certain cases, specifically when the subgroup generated by P has smooth order or when there is a mapping from E into a small-degree extension of the base field \mathbb{F}_q (such a mapping can arise from the Weil or Tate pairings [22], [12], or from taking p -adic logarithms [30], [27], [33], [26]).

For the rest of this section we will assume that all elliptic curves in question do not belong to one of these special cases. For these remaining elliptic curves, the only known methods for solving the elliptic curve discrete logarithm problem involve reducing to prime order subgroups (the Silver/Pohlig-Hellman reduction) and then applying generic group methods such as the baby-step-giant-step method (see [7]) and Pollard's Rho and Lambda methods [25]. The complexity of these methods depends only on the size of the largest prime factor (say L) of the number of points on the elliptic curve. This gives the naive impression that the difficulty of the elliptic curve discrete logarithm problem depends only on the prime L .

The experts in elliptic curve cryptography do not expect that the difficulty of the elliptic curve discrete logarithm problem does depend only on the prime L , or even on only the pair of primes L and p . One of the aims of the present work is to give a more thorough analysis of this issue. Indeed, we study the following question.

Question 1. Let E_1 and E_2 be two elliptic curves over \mathbb{F}_p . Suppose that the largest prime divisor of $\#E_i(\mathbb{F}_p)$ is the same prime L for both curves. Is there a method to reduce the discrete logarithm problem on $E_1(\mathbb{F}_p)$ to the discrete logarithm problem on $E_2(\mathbb{F}_p)$?

It is perhaps not so obvious why this question is relevant for cryptography (since only one elliptic curve is ever considered at a time). The issue arises when constructing elliptic curves on which to base a cryptosystem; since there are many different ways to write down elliptic curves, it is important to know whether some of these representations are 'stronger' or 'weaker' than others.

Lenstra [17] has studied a similar problem for finite fields \mathbb{F}_{p^n} . The result is that there is a polynomial-time algorithm to convert between any two different bases for \mathbb{F}_{p^n} as a vector space over \mathbb{F}_p . This is an important result as it means that all representations used for finite fields in cryptosystems based on discrete logarithms are equally secure (or insecure). The problems considered in this paper seem to be harder than the problem of constructing isomorphisms between finite fields, and we do not expect a polynomial-time solution to the question stated above.

The statement of Question 1 first appears to be more general than the rest of this paper. However, for the elliptic curves used in cryptography it is desirable that the largest prime divisor L of the order of the group $E_i(\mathbb{F}_p)$ be of size close to p , and so certainly $L > 4\sqrt{p}$. It follows from the Hasse bound that there is only one possible value for $\#E_i(\mathbb{F}_p)$ and so the two elliptic curves are isogenous. An isogeny between the curves provides a method for reducing the discrete logarithm problem on one to the other.

Although the algorithm given in this paper is exponential, in most cases (specifically, when the conductors of the endomorphism rings of the elliptic curves in question are not divisible by a large prime) it is more efficient than the known general-purpose algorithms for computing discrete logarithms. Therefore our results give some support to the idea that, in general, the discrete logarithm problem is equally difficult for all elliptic curves over a finite field \mathbb{F}_p which have the same number of points. The main exception to this conclusion is when we have two elliptic curves E_1 and E_2 over \mathbb{F}_p such that $\text{End}(E_1) \subset \text{End}(E_2)$ and such

that the index is divisible by a large prime. In this case it is not clear what the relationship is between the difficulty of the discrete logarithm problems on these elliptic curves.

We should emphasise the following point: this paper does not imply that elliptic curves that have complex multiplication of low class number are either weaker or stronger than random curves.

3. Background

Let E be an elliptic curve over the field \mathbb{F}_p with $p \geq 5$. Since it is easy to compute isomorphisms (see the Appendix, Subsection A.2) we will assume that E is written in the short Weierstrass form

$$E : y^2 = x^3 + Ax + B.$$

We denote by π the Frobenius endomorphism

$$\pi : (x, y) \mapsto (x^p, y^p)$$

on E which satisfies

$$\pi^2 - t\pi + p = 0$$

in $\text{End}(E)$ where $t = p + 1 - \#E(\mathbb{F}_p)$.

An elliptic curve E over \mathbb{F}_p is *supersingular* if the multiplication by p map has trivial kernel. The curve is *ordinary* if the kernel is *non-trivial*. For elliptic curves over \mathbb{F}_p we know that E is supersingular if and only if $t = 0$.

If an isogeny φ is defined over \mathbb{F}_p then its kernel C is also defined over \mathbb{F}_p (that is, if $(x_0, y_0) \in C$ and if $\sigma \in \text{Gal}(\overline{\mathbb{F}_p}/\mathbb{F}_p)$ then $(\sigma(x_0), \sigma(y_0)) \in C$). Since we are assuming that all elliptic curves are given by a short Weierstrass equation it follows that $(x_0, y_0) \in C$ if and only if $(x_0, -y_0) \in C$, from which it follows that the kernel C of the isogeny is determined by the polynomial $\psi(x) = \prod_{(x_0, \pm y_0) \in C} (x - x_0) \in \mathbb{F}_p[x]$. A separable isogeny of degree d is called a d -isogeny and its kernel has size d .

Every isogeny $\varphi : E_1 \rightarrow E_2$ over \mathbb{F}_p can be described as a rational map of the following form

$$\varphi(x, y) = \left(\frac{\varphi_1(x, y)}{\psi(x)^2}, \frac{\varphi_2(x, y)}{\psi(x)^3} \right), \tag{1}$$

where φ_1 and φ_2 are polynomials over \mathbb{F}_p in the variables x and y of the original curve E , and $\psi(x)$ determines the kernel of the isogeny as described above. In this paper, when we say ‘construct an isogeny’, we mean explicitly computing the polynomials $\psi(x)$, $\varphi_1(x, y)$ and $\varphi_2(x, y)$.

Vélu [35] showed how ψ determines the isogeny, and gave formulae to obtain the polynomials φ_1 and φ_2 in equation (1). Elkies [10] developed methods to calculate $\psi(x)$, given only the j -invariants of the curves E_1 and E_2 . The j -invariants of isogenous elliptic curves may be found from the modular polynomials (equations for $X_0(N)$) as is well-known in the theory of counting points. Some details of these methods are given in Appendix A.

Suppose that E is an ordinary elliptic curve over \mathbb{F}_p with $p + 1 - t$ points, and let $d = t^2 - 4p < 0$. Then $\text{End}(E)$ is some order in the field $K = \mathbb{Q}(\sqrt{d})$ and indeed $\mathbb{Z}[\pi] = \mathbb{Z}[(d + \sqrt{d})/2] \subseteq \text{End}(E) \subseteq \mathcal{O}_K$ where \mathcal{O}_K is the ring of integers of K . The conductor $c = [\mathcal{O}_K : \mathbb{Z}[\pi]]$ is precisely the largest integer such that $d/c^2 \equiv 0$ or $1 \pmod{4}$. The discriminant of K is therefore $D = d/c^2$. There are a finite number of possibilities for $\text{End}(E)$, namely, all the rings $\mathcal{O} = \mathbb{Z} + c'\mathcal{O}_K$ where c' is a divisor of c .

The primes dividing the conductor c of $\mathbb{Z}[\pi]$ will require special attention. This is due to the following fact.

Proposition 1. *Let $\varphi : E_1 \rightarrow E_2$ be an isogeny such that $\text{End}(E_1) \subseteq \text{End}(E_2)$ (respectively, $\text{End}(E_2) \subseteq \text{End}(E_1)$). Suppose l is a prime which divides the index $[\text{End}(E_2) : \text{End}(E_1)]$ (respectively, $[\text{End}(E_1) : \text{End}(E_2)]$). Then the degree of φ is divisible by l .*

Proof. See [14, Proposition 21]. □

Indeed, by taking a sequence of l -isogenies (for those primes l dividing the conductor of $\text{End}(E)$) it is possible to find an elliptic curve E' isogenous to E such that $\text{End}(E') = \mathcal{O}_K$ (the maximal order). These ideas are central to Kohel’s algorithm [14] for finding the exact endomorphism ring of a given ordinary elliptic curve E/\mathbb{F}_p . Some of the details of this algorithm are given in Appendix A and in the description of our method. We will use the following notation from [14]. Let l be a prime number and suppose that $\varphi : E_1 \rightarrow E_2$ is an isogeny of degree l . Then one of the three following cases holds.

1. $\text{End}(E_1) \simeq \text{End}(E_2)$ in which case the isogeny φ is said to be ‘horizontal’ at l .
2. $\text{End}(E_1)$ contains $\text{End}(E_2)$ with index l , in which case the isogeny φ is said to be an isogeny ‘down’ at l .
3. $\text{End}(E_1)$ is contained in $\text{End}(E_2)$ with index l , in which case the isogeny φ is said to be an isogeny ‘up’ at l .

If $l \nmid [\mathcal{O}_K : \text{End}(E)]$ then it is not possible to go up, and so we say that E is ‘on the surface at l ’. Similarly, if $l \nmid [\text{End}(E) : \mathbb{Z}[\pi]]$ then it is not possible to go down, and we say that E is ‘on the floor at l ’.

For each possible endomorphism ring \mathcal{O} , the theory of complex multiplication (see [15] or [31]), combined with Deuring’s theory (see [15]) about lifting endomorphisms from characteristic p to characteristic zero, shows that the number of isomorphism classes of elliptic curves over \mathbb{F}_p which have endomorphism ring isomorphic to \mathcal{O} is equal to $h_{\mathcal{O}}$ (the number of elements of the group $\text{Pic}(\mathcal{O})$ of classes of projective \mathcal{O} -modules). Furthermore, the action of isogenies is the same as composition of ideal classes. Indeed, elliptic curves over \mathbb{F}_p can be lifted to elliptic curves over \mathbb{C} with the same endomorphism ring \mathcal{O} , and can thus be interpreted as \mathbb{C}/\mathfrak{a} where \mathfrak{a} is an projective \mathcal{O} -module. Given some other projective \mathcal{O} -module \mathfrak{b} one has the isogeny

$$\mathbb{C}/\mathfrak{a} \longrightarrow \mathbb{C}/\mathfrak{a}\mathfrak{b}^{-1}$$

with kernel $\mathfrak{a}\mathfrak{b}^{-1}/\mathfrak{a} \simeq \mathcal{O}/\mathfrak{b}$. These facts will be useful in our analysis of the algorithm.

Over \mathbb{C} the correspondence between the isomorphism classes of elliptic curves with endomorphism ring equal to \mathcal{O} , and the ideal classes in the group $\text{Pic}(\mathcal{O})$ is canonical. However, the reduction of these curves from \mathbb{C} to \mathbb{F}_p depends on a choice of (totally split) prime $\wp \mid p$ in the ring class field. If we fix such a prime \wp then one can talk about the ideal class associated with an isomorphism class of elliptic curves over \mathbb{F}_p with endomorphism ring equal to \mathcal{O} . If it were possible to efficiently compute this correspondence between $\text{Pic}(\mathcal{O})$ and a set of j -invariants in \mathbb{F}_p , then the problems discussed in this paper would be easy to solve. However, the only method known to the author for computing this correspondence would be to compute approximations to all $j(\mathfrak{a}) \in \mathbb{C}$ (as \mathfrak{a} runs over the classes in $\text{Pic}(\mathcal{O})$), to construct the class polynomial and therefore the ring class field $H_{\mathcal{O}}$, to find a prime \wp above p , and then to reduce these values modulo \wp . This approach would be more arduous than the methods used in this paper.

Quadratic imaginary fields are well-understood and the structure and size of the class group may be computed in subexponential time using binary quadratic forms (see [7]). The class group of K is generated by the primes p which split or ramify in K . Bach [4] has proved that, assuming the Riemann hypothesis for the zeta function of K , the class group is generated by the prime ideals of norm less than $6(\ln |D|)^2$.

Class numbers of orders are closely related to the class number h_K of the maximal order by the following formula (see [15, Theorem 8.7] or [31, Exercise 4.12]). Suppose \mathcal{O} is an order of conductor c in \mathcal{O}_K . Let \mathcal{O}_K^* and \mathcal{O}^* be the units (that is, multiplicative subgroup of invertible elements). Define $(\frac{K}{p})$ to be $-1, 0,$ or $+1$ depending on whether the prime p is inert, ramified or split in K respectively. Then the relation between the class numbers is

$$h_{\mathcal{O}} = h_K c [\mathcal{O}_K^* : \mathcal{O}^*] \prod_{p|c} \left(1 - \left(\frac{K}{p}\right) p^{-1}\right). \tag{2}$$

In particular we see that the class number grows as the conductor grows.

For the purposes of this paper the following result (Exercise 5.27 of [7]) will be useful.

Proposition 2. *Let K be an imaginary quadratic field of discriminant D . Then*

$$h_K \leq \frac{1}{\pi} \sqrt{|D|} \ln |D|. \tag{3}$$

We now mention the relationship between the endomorphism ring and the group structure. Suppose that $\text{End}(E) = \mathcal{O}$ is an order containing $\mathbb{Z}[\pi]$. Let m be the index $m = [\mathcal{O} : \mathbb{Z}[\pi]]$ (which is the conductor of $\mathbb{Z}[\pi]$ divided by the conductor of \mathcal{O}). Then there is some integer a such that $\mathcal{O} = \mathbb{Z}[(\pi - a)/m]$. Lenstra [18] has shown that the group structure of $E(\mathbb{F}_p)$ is isomorphic to $\mathcal{O}/(\pi - 1)$ and thus, as a group, $E(\mathbb{F}_p) \cong \mathbb{Z}/l\mathbb{Z} \times \mathbb{Z}/n\mathbb{Z}$ where $l = \gcd(a - 1, m)$ and $l \mid n$. The following result is therefore immediate.

Lemma 1. *If E/\mathbb{F}_p is an ordinary elliptic curve such that $\text{End}(E) = \mathbb{Z}[\pi]$ then $E(\mathbb{F}_p)$ is a cyclic group.*

This lemma provides a method for solving Problem 3 of the introduction; given an elliptic curve E/\mathbb{F}_p such that the group $E(\mathbb{F}_p)$ is not cyclic, we can use the methods of Kohel’s thesis [14] to construct an isogenous elliptic curve E'/\mathbb{F}_p such that $\text{End}(E') = \mathbb{Z}[\pi]$. It then follows that $E'(\mathbb{F}_p)$ is a cyclic group.

4. The method

In this section we present a method to explicitly construct an isogeny between two elliptic curves E_1 and E_2 which have the same number of points.

From Tate’s isogeny theorem [34] we know that there is an isogeny $\varphi : E_1 \rightarrow E_2$ over \mathbb{F}_p . In practice it will be more efficient to consider a chain

$$E_1 \simeq E'_0 \xrightarrow{\varphi_1} E'_1 \xrightarrow{\varphi_2} E'_2 \xrightarrow{\varphi_3} \dots \xrightarrow{\varphi_n} E'_n \simeq E_2$$

of isogenies such that each φ_k has prime degree l_k .

For the input to the method, assume that we have short Weierstrass equations for E_1 and E_2 and that we know the trace t of Frobenius. It is necessary to compute $d = t^2 - 4p$ and to calculate the conductor of $\mathbb{Z}[\pi]$ in \mathcal{O}_K where $K = \mathbb{Q}(\sqrt{d})$. The conductor of $\mathbb{Z}[\pi]$ is the largest positive integer c such that $D := d/c^2 \equiv 0, 1 \pmod{4}$ and so this stage involves factoring the integer d . Factoring can be performed in subexponential time, and we will not

include this cost in our analysis. Note, however, that this means we have already lost any chance of achieving polynomial time!

We must decide upon a set \mathcal{L} of primes; this is the set of all primes l for which we will need to take l -isogenies (note that we will also need all primes $l \mid c$ in order to apply Kohel’s algorithm). We take $\mathcal{L} := \{\text{primes } l : (\frac{D}{l}) \in \{0, 1\} \text{ and } l < 6(\ln |D|)^2\}$ (where D is the discriminant of the maximal order). Assuming the generalised Riemann hypothesis, the primes in this set \mathcal{L} generate the ideal class group of the field K , and so they are sufficient for the task.

Here is an outline of the method:

Stage 0 Compute modular polynomials $\Phi_l(x, y) \in \mathbb{F}_p[x, y]$ for each prime l dividing c and each $l \in \mathcal{L}$.

Stage 1 For each curve E_i , use Kohel’s algorithm [14] to find a chain of isogenies from E_i to an elliptic curve E'_i whose endomorphism ring is the maximal order. Set j_1 and j_2 to be the j -invariants of these new curves E'_1 and E'_2 .

Stage 2 Starting from j_1 and j_2 , construct trees, whose vertices are labelled as j -invariants and whose edges (labelled by l) correspond to isogenies of prime degree l .

Stage 3 Once the two trees become connected, find the path connecting j_1 to j_2 . Add this chain of l -isogenies to those already found in Stage 1 and use the theory of Elkies and Vélú to construct the explicit polynomial form of each l -isogeny.

We now give further details about the stages. **Stage 0** involves standard methods; see [11] for details. Of course, rather than computing all the equations in advance, one would just compute and store them as they are required.

Stage 1. We know that $\text{End}(E_i)$ is an order in \mathcal{O}_K with conductor dividing c . Note that, due to equation (2), it is more likely that $\text{End}(E_1)$ and $\text{End}(E_2)$ have conductor close to c .

Let E be each E_i in turn. For each prime l dividing the conductor c of $\mathbb{Z}[\pi]$ write l^b for the exact power of l in c .

We use the following method (due to Kohel [14]) to determine the exact power of l dividing $\text{End}(E)$. Consider the roots of $\Phi_l(j(E), y) \pmod{p}$. If there is only one root then we know from Propositions 22 and 23 of Kohel [14] (or see Theorem 4 of Appendix A) that E is on the floor at l and therefore that the conductor of $\text{End}(E)$ is divisible by exactly l^b . In the other case, it follows that $\Phi_l(j(E), y)$ splits completely modulo p . If $b = 1$ then we know that l does not divide the conductor of $\text{End}(E)$. If $b > 1$ then it is not yet possible to distinguish the possibilities. However, if we now let y_0 run over the roots of $\Phi_l(j(E), y)$ then we can determine if y_0 is the j -invariant of an isogeny on the floor by considering the roots of $\Phi(y_0, y) \pmod{p}$. If y_0 is on the floor this polynomial will have only one root. If $\Phi(y_0, y) \pmod{p}$ has more than one root, then we repeat the process using the roots y . If it always takes at least b' steps down from $j(E)$ before hitting the floor then it follows that the power of l dividing the conductor of $\text{End}(E)$ is $b - b'$.

Once we have determined the power of l dividing the conductor of $\text{End}(E)$ we must also determine the unique l -isogenies up to a curve on the surface at l . We then know the intermediate j -invariants in the chain of l isogenies up to the elliptic curve E' such that l does not divide the conductor of $\text{End}(E')$. We now let $E = E'$ and continue the process for the remainder of the primes $l \mid c$.

At the completion of Stage 1 we have the chains of isogenies

$$E_1 \longrightarrow E_{1,1} \longrightarrow \cdots \longrightarrow E_{1,n} = E'_1 \text{ and } E_2 \longrightarrow E_{2,1} \longrightarrow \cdots \longrightarrow E_{2,m} = E'_2 \quad (4)$$

where $\text{End}(E'_1)$ and $\text{End}(E'_2)$ are both the maximal order \mathcal{O}_K . Put $j_1 = j(E'_1)$ and $j_2 = j(E'_2)$. We now want to link j_1 and j_2 by a chain of isogenies.

Stage 2. We grow two trees of l -isogenies, one starting from j_1 and the other from j_2 . Initially, begin with the two trivial graphs, each with a single vertex (labelled, respectively, j_1 and j_2) and no edges. Repeat the following steps: For each tree in turn choose a random $l \in \mathcal{L}$ and, for every vertex j in the tree, compute the roots of $\Phi_l(j, y) \pmod{p}$. For each of these roots y (apart from when $l \mid c$, in which case one must only use the roots which correspond to elliptic curves which lie on the surface), if there is not already a vertex in either of the two trees corresponding to y , then add a new vertex y and add a labelled edge between j and y (with labelling l). This process should be repeated until there is an edge added which connects the two trees together. This concludes Stage 2. Note that this process could be made more efficient by keeping track of which ideal classes have already been utilised (since it is possible to compute efficiently in the class group), and thus choosing $l \in \mathcal{L}$ to maximise the number of new classes found.

Since the vertices in our trees do not have bounded degree, and since it is necessary to detect if a j -invariant already appears in the tree, some kind of easily searched list of vertices would be an appropriate data structure (rather than the usual binary tree structure).

Stage 3. One must traverse the tree to find the chain of isogenies between the elliptic curves with j -invariants j_1 and j_2 . These are then combined with the chains (4) of l -isogenies found in Stage 1. For each l -isogeny in the chain we use the formulae of Elkies and Vélú to give explicit values for the polynomials ψ, φ_1 and φ_2 of (1). It is advised to store these isogenies individually, rather than combining them into huge polynomials.

Finally, the image curve will probably be given in terms of an equation different from that of E_2 , so it is necessary to construct an isomorphism to E_2 (see the Appendix, A.2).

It is clear that, assuming the Riemann hypothesis for the quadratic imaginary field K , this algorithm does terminate, having produced an isogeny from E_1 to E_2 .

Note that it is possible to construct an isogeny from E_2 to E_1 by working along the chain backwards.

5. The depth of the tree of isogenies

The goal of Stage 2 is to produce two ‘very bushy’ trees of small-degree isogenies, starting from each of the two given j -invariants. In this section we will estimate the depth (maximum distance from the root to the leaves) of these trees, as it determines the length of the resulting chain of isogenies.

The two initial j -invariants correspond (non-canonically) with a pair of ideals α_1 and α_2 in the ideal class group of the maximal order \mathcal{O}_K . The process in Stage 2 takes a sequence of small primes l_i which split (or ramify) in \mathcal{O}_K as $(l) = \mathfrak{l}'$. In the following we write \mathfrak{l}^{-1} instead of \mathfrak{l}' , since \mathfrak{l}' is a representative of the ideal class corresponding to the inverse of \mathfrak{l} . The two trees themselves contain the j -invariants of the elliptic curves corresponding to ideal classes represented by the ideals

$$\alpha_1 \prod_{i=1}^n \mathfrak{l}_{2i-1}^{a_{2i-1}} \quad \text{and} \quad \alpha_2 \prod_{i=1}^n \mathfrak{l}_{2i}^{a_{2i}}$$

with all $a_i \in \{-1, 0, +1\}$. It is clear that, if non-ramified prime ideals are chosen, the size of the trees grows initially like 3^n . However, as n grows there will be an increased probability of non-trivial relations between the ideals, and therefore there will be fewer new classes

arising. Note that when ramified primes are chosen then the number of ideal classes is expected to double, and if primes are repeated then the number of classes grows by a factor which is between 1 and 2.

The condition that the trees share a common vertex is equivalent to the condition that, for some choice of the $a_i \in \{-1, 0, +1\}$, we have

$$a_1 a_2^{-1} \sim \prod_{i=1}^{2n} l_i^{a_i}.$$

The complexity analysis depends on the following heuristic assumption. We say that two probabilities P_1 and P_2 are ‘close’ if $\frac{1}{2}P_1 \leq P_2 \leq 2P_1$.

Assumption 1. Let n be a positive integer and let L be a sequence $l_1, l_2, \dots, l_n \in \mathcal{L}$. Consider the set S of all ideal classes represented by ideals of the form $\prod_{i=1}^n l_i^{a_i}$ over all choices for $a_i \in \{-1, 0, +1\}$. Then, for a large proportion of sequences L , the set S has the following property: for a random class \mathfrak{a} in the class group of K , the probability that there are ideals $\mathfrak{s}_1, \mathfrak{s}_2 \in S$ such that $\mathfrak{s}_1 \mathfrak{a} \sim \mathfrak{s}_2$ is close to the probability in the case where the set S was chosen (of the same size) uniformly at random from the ideal class group of K .

The reason the assumption does not apply to all sequences L is that there are some sequences (such as a sequence of l which generates a proper subgroup of the ideal class group) for which the statement is not true.

The set \mathcal{L} has been defined so that (assuming the Generalised Riemann Hypothesis) the set S will eventually contain a representative of every ideal class of K . The heuristic assumption above is essentially a claim about the ‘uniform’ way in which the ideal classes are generated.

Of course, the set S is easily distinguished from a set chosen uniformly, as it has the strong property that $\mathfrak{a} \in S \Leftrightarrow \mathfrak{a}^{-1} \in S$. Nevertheless, we claim that this fact does not significantly affect the probabilities we are considering.

The assumption seems plausible due to the fact that the prime ideals lying above distinct rational primes do not have any obvious dependence on each other. At the end of this section we give an example that supports the assumption.

Lemma 2. *Subject to the heuristic assumption 1, Stage 2 is expected to terminate after $O(\ln h_K)$ iterations.*

Proof. We consider the two trees separately. At iteration n of the algorithm suppose that we have found representatives $\mathfrak{a}_1, \dots, \mathfrak{a}_N$ of N distinct ideal classes. Choose a splitting prime l (that is, $(l) = \mathfrak{l}^{-1}$) which is distinct from the primes already chosen. Consider the $3N$ ideals $\mathfrak{a}_i l^\epsilon$ where $\epsilon = 0, \pm 1$.

Some of these ideals may be equivalent, but this can only arise from $\mathfrak{a}_i l^{\epsilon_1} \sim \mathfrak{a}_j l^{\epsilon_2}$ where $\epsilon_1 \neq \epsilon_2$ and $i \neq j$. Without loss of generality we may assume that the pair (ϵ_1, ϵ_2) is equal to one of $(0, -1), (0, +1), (-1, +1)$. In each case there are $N(N - 1)$ possible values for $i \neq j$ and, under assumption 1, we expect that approximately $N(N - 1)/h_K$ of these pairs of ideals will be equivalent.

Therefore, we expect the number of distinct ideal classes obtained at the end of the n th iteration to be approximately $3N - 3N(N - 1)/h_K = (3 - 3(N - 1)/h_K)N$. When $N \leq \sqrt{h_K}$ it follows that the size of the tree increases by a factor which is close to 3. If a ramified prime is chosen then a similar argument shows that the size of the tree increases by a factor close to 2. Repeated primes can also be handled by this argument, though they

are unlikely to arise within Stage 2 as the algorithm specifies that the primes l be chosen uniformly at random. In conclusion, we expect the tree to have size $\sqrt{\pi h_K/2}$ after $O(\ln h_K)$ iterations.

We now have two trees of size $\sqrt{\pi h_K/2}$ which, subject to our heuristic Assumption 1, behave like two sets of randomly chosen ideal classes. By the ‘birthday paradox’ we expect the trees to share a common vertex. In other words, the two trees are joined by an edge, and Stage 2 of the algorithm is complete. \square

We give a toy example to illustrate the situation.

Consider the discriminant $D = -399992579$ (which would arise from an elliptic curve modulo $p = 10^8 + 37$ having $t = 87$). The class number is $h = 11920$, and so $\lfloor \sqrt{h} \rfloor = 109$. Consider the sequence 199, 137, 59, 41, 89, 53, 43, 89 of splitting primes and construct trees as in Stage 2. The resulting trees of depth 4 have cardinality 81, and there are a total of 5056 ideal classes represented as a product of a class from one tree and a class from another (in other words, 42% of the ideal classes appear). Taking the further two primes 107 and 71 yields trees of depth 5 with 242 and 243 vertices respectively (this illustrates that the size of the trees really does grow like 3^n while the size is less than $\sqrt{h_K}$). The total number of ideal classes represented as products is 11553, which is 97% of the class group.

6. Complexity analysis

It remains to give an analysis of the complexity of the method. We assume a unit cost for all operations in the field \mathbb{F}_p .

The analysis given in this section uses the results of the previous section, which depend on a heuristic assumption (Assumption 1). These results lead to the precise bounds given in equations (5) and (6) below. Actually, the worst-case complexity stated in Theorem 1 can be obtained without relying on the heuristic Assumption 1 of Section 5 (by running Stage 2 for $O(h)$ iterations).

Let $h = h_K$ be the class number of the maximal order. Note that, by (2) the class number of the maximal order is usually much smaller than the class numbers of non-maximal orders. Nevertheless, in the worst case, formula (3) shows that h is $O(p^{1/2} \ln p)$. There are efficient (subexponential) methods to compute the class number, so we may assume that we know the exact value for h . Let c be the conductor of $\mathbb{Z}[\pi]$. We ignore the cost of factoring $t^2 - 4p$ to find the conductor. In the worst case c is $O(p^{1/2})$, though in practice it will usually be divisible by only very small primes.

Rather than give our complexity estimate in terms of p alone, we will express it in terms of h and c also. The reason for this is that we want to emphasise which parts of the method depend on which parameters. We will now consider the various stages in turn.

Stage 0. Firstly, we must obtain the equations $\Phi_l(X, Y)$ for all the split/ramified primes l less than our bound $6(\ln |D|)^2 < 6(\ln 4p)^2$. By the prime number theorem there are $O((\ln p)^2 / (\ln \ln p)^2)$ such primes. We also need $\Phi_l(X, Y)$ for those primes dividing the conductor c . The first set takes time $O((\ln p)^8 / \ln \ln p)$ and space $O((\ln p)^6 / \ln \ln p)$ (in terms of \mathbb{F}_p operations). The primes dividing c contribute $O(c^3)$ time and $O(c^2)$ space.

We now must deal with the core of the method. The principal task is calculating the roots of the degree l polynomial $\Phi_l(j(E), Y) \pmod{p}$. Finding roots of a degree l polynomial modulo p may be performed in probabilistic polynomial time $O(l^2 \ln p)$ operations in \mathbb{F}_p (see [7, Section 1.6]).

Stage 1. For each prime l dividing c we must find the roots of $\Phi_l(j, y)$. If $l^a \parallel c$ then we will, at worst, need to take $O(l^{a-1})$ (which is $O(c)$) different values of j before we know we have reached the surface at l . In other words, we have to perform $O(c^3(\ln p))$ field operations. The length of the chain of isogenies from the E_i to elliptic curves on the surface will be $O(\ln c)$.

Stage 2. We must generate the two trees using l -isogenies. We expect that both trees need to be of size approximately $O(\sqrt{h})$. By Lemma 2 this should require $O(\ln h)$ iterations of Stage 2.

For each prime l we must find the roots $\Phi_l(j, y) \equiv 0 \pmod{p}$ for every vertex j in the tree. The tree has size $O(\sqrt{h})$ and finding each root takes $O(l^2 \ln p) = O((\ln p)^5)$ time. Furthermore, once each root y is found we need to search through both trees to see if it has already appeared, and this requires time $O(\ln h)$ if the trees are represented in a suitable way.

Hence Stage 2 requires $O(\ln h \sqrt{h}((\ln p)^5 + (\ln h)))$ time. The trees require $O(\sqrt{h})$ storage space and the chain of isogenies itself has length $O(\ln h)$.

Stage 3. Finding the chain of j -invariants which connects the roots of the trees takes time $O(\sqrt{h})$ and the length of the chain is $O(\ln h)$. Once we combine with the chains found in Step 1 the total chain will have length $O(\ln h + \ln c)$.

For each l -isogeny in the chain we then compute the isogeny using the methods of Elkies and Vélú. The primes $l \in \mathcal{L}$ require time $O((\ln p)^6)$ for computing the isogeny. The primes $l \mid c$ require time $O(c^3)$.

Therefore, Stage 3 requires $O(\sqrt{h} + \ln h(\ln p)^6 + (\ln c)c^3)$ time and $O(\ln h(\ln p)^4 + (\ln c)c^2)$ space.

Putting it all together, the algorithm takes expected time

$$O((\ln p)^8 / \ln \ln p + c^3 + c^3(\ln p) + \ln h \sqrt{h}((\ln p)^5 + (\ln h)) + \sqrt{h} + \ln h(\ln p)^6 + (\ln c)c^3) \tag{5}$$

and requires expected space

$$O((\ln p)^6 / \ln \ln p + c^2 + \ln c + \sqrt{h} + \ln h(\ln p)^4 + (\ln c)c^2). \tag{6}$$

In the worst case, since c could be $O(p^{1/2})$ and, by (3), h could be $O(p^{1/2} \ln p)$, the terms involving c dominate and so the expected running time is $O(p^{3/2}(\ln p))$ and the method requires $O(p(\ln p))$ space.

We emphasise that the algorithm performs much more efficiently in most cases. For instance, in most examples the conductor will be $O(\ln p)$ -smooth, in which case all the terms featuring c become polynomial-time, and the algorithm has expected running time $O(p^{1/4}(\ln p)^{13/2})$.

In addition to this, in the case when the maximal order has small class number (say h of size a power of $(\ln p)$) then the algorithm becomes polynomial-time. In particular, those elliptic curves generated by the CM method (see [3], [23] and [16]) have this property.

In cryptography it is usually suggested that elliptic curves be chosen at random so that the endomorphism rings do not have small class number. Actually, the elliptic curves used in cryptography seem to have slightly smaller class number than expected (see [13]).

7. Polynomial-time group homomorphism

We now address Problem 2. Given two elliptic curves E_1 and E_2 we wish to construct a non-trivial isogeny between them which may be evaluated in polynomial time.

The method proposed for constructing a chain of isogenies between two given elliptic curves requires exponential time; however the chain of isogenies has polynomial-sized length. Furthermore, if c is $O(\ln p)$ -smooth then all isogenies in the chain can be computed in polynomial time.

Thus, for those elliptic curves whose endomorphism ring has smooth conductor, we have produced a solution to Problem 2: once the exponential ‘pre-computation’ has been performed, the group homomorphism requires polynomial storage space and can be computed in polynomial time. This is an improvement over the ‘naive’ exponential solution to Problem 2 (that is, taking discrete logarithms) since that method would require storing an exponentially sized lookup table.

8. When the conductor is not smooth

The formula (2) shows that a random curve usually has conductor close to the full index $[\mathcal{O}_K : \mathbb{Z}[\pi]]$. The conductor c is essentially the square part of $t^2 - 4p$ and so we expect it to be quite smooth in general. However, if the conductor is divisible by a large prime then our method is impractical.

If the conductors of both $\text{End}(E_1)$ and $\text{End}(E_2)$ are divisible by the large prime then it is advisable not to perform Stage 1 of the method. Stage 2 may be equally well performed in a non-maximal order. The drawback is that the class number h will be increased by the large prime. Note that one can detect whether $\text{End}(E)$ is maximal by using the methods of this paper to find more than h_K isogenous elliptic curves E' such that $\text{End}(E') = \text{End}(E)$.

If one of the conductors of $\text{End}(E_1)$ and $\text{End}(E_2)$ is divisible by the large prime and the other is not then, due to Proposition 1, there is no shortcut available. This fact, together with the fact that determining the conductor requires finding the square factors of the discriminant, leads us to the opinion that Problems 1, 2 and 3 cannot have polynomial-time solutions in the general case.

9. Onwards

One may ask whether the methods of this paper are the only way to proceed. Suppose there is a translation of discrete logarithm problems on elliptic curves, given as a ‘geometrically defined’ map. Must this necessarily be an isogeny of curves?

For instance, as David Kohel has pointed out, an isogeny is defined as a map of one-dimensional group schemes

$$E_1 \xrightarrow{\varphi} E_2$$

over the field \mathbb{F}_p (so that $\varphi : E_1(\mathbb{F}_{p^n}) \rightarrow E_2(\mathbb{F}_{p^n})$ for all n). For our problem it would be sufficient to have a homomorphism of zero-dimensional constant group schemes $E_1[L, \pi - 1] \rightarrow E_2[L, \pi - 1]$.

It is not known to the author whether there is any way to find such objects computationally without taking isogenies between the one-dimensional group schemes they arise from. This would be an interesting question for further study.

Acknowledgements. This paper owes a huge debt to the work described in David Kohel’s thesis. I would like to thank David for patiently explaining some of these ideas to me.

I would also like to thank Bryan Birch, Alfred Menezes, Takakazu Satoh and the anonymous referees for helpful comments on earlier drafts of the paper.

Appendix A. Background material

A.1. Elliptic curves

In this paper we assume that all elliptic curves E over \mathbb{F}_p are given in Weierstrass form

$$E : y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$$

where $a_1, a_2, a_3, a_4, a_6 \in \mathbb{F}_p$. If E is written as a more general plane affine or projective curve then there are well-known methods (see [6, Chapter 8], and [32, Section III.1]) to construct an isomorphism from E to the Weierstrass form above. We will not consider the case when E is written as a highly singular curve or as a subvariety of some higher-dimensional affine or projective space.

Indeed, since we are assuming $p \geq 5$ it is possible to reduce to the short Weierstrass equation by substituting $y - a_1/2x - a_3/2$ for y and then substituting $x - (a_1^2/4 + a_2)/3$ for x .

Therefore, we will always assume that elliptic curves are given as short Weierstrass models

$$E : y^2 = x^3 + Ax + B$$

with $A, B \in \mathbb{F}_p$.

We deal first with the case of isomorphic curves.

A.2. Isomorphisms

Let $E_1 : y^2 = x^3 + A_1x + B_1$ and $E_2 : (y')^2 = (x')^3 + A_2x' + B_2$ over \mathbb{F}_p be elliptic curves with the same number of points. The j -invariants $j_1 = j(E_1)$ and $j_2 = j(E_2)$ may be easily computed from the formula $j(E_i) = 6912A_i^3 / (4A_i^3 + 27B_i^2)$.

If two elliptic curves have the same j -invariant then [32, Propositions III.1.4(b) and III.3.1(b)] they are isomorphic over $\overline{\mathbb{F}_p}$. Indeed, since E_1 and E_2 are in short Weierstrass form, the isomorphism must be of the form $\psi : (x, y) \mapsto (u^2x, u^3y)$ where u satisfies at least one of the following identities:

$$\begin{aligned} u^2 &= (B_2A_1/B_1A_2) & \text{if } B_1A_2 \neq 0; \\ u^4 &= A_2/A_1 & \text{if } A_1 \neq 0; \\ u^6 &= B_2/B_1 & \text{if } B_1 \neq 0. \end{aligned} \tag{7}$$

Note that these roots, if they exist in \mathbb{F}_p , may be easily computed (using the standard probabilistic algorithm for calculating roots modulo p as described in [7, Section 1.6]).

In the case when $\#E_1(\mathbb{F}_p) = \#E_2(\mathbb{F}_p)$ and the elliptic curves are not supersingular, it can be shown by an elementary argument that the isomorphism is actually defined over \mathbb{F}_p (in other words, $u \in \mathbb{F}_p$). Hence, the following result holds.

Proposition 3. *Let E_1 and E_2 be elliptic curves over \mathbb{F}_p such that $\#E_1(\mathbb{F}_p) = \#E_2(\mathbb{F}_p)$. If E_1 and E_2 are isomorphic as elliptic curves (that is, $j(E_1) = j(E_2)$) then there is a (probabilistic) polynomial time algorithm to find the isomorphism from $E_1(\mathbb{F}_p)$ to $E_2(\mathbb{F}_p)$.*

If long Weierstrass models are being used (for instance, if one is working in characteristic 2) then isomorphisms are equally easy to handle.

A.3. Isogenies

For the background on isogenies we refer to [32, Section III.4]. We denote by π the Frobenius isogeny

$$\pi : (x, y) \mapsto (x^p, y^p)$$

on an elliptic curve E/\mathbb{F}_p . This isogeny satisfies the relation

$$\pi^2 - t\pi + p = 0$$

where the number of points on the elliptic curve is given by $\#E(\mathbb{F}_p) = p + 1 - t$.

An isogeny $\varphi : E_1 \rightarrow E_2$ over \mathbb{F}_p may be expressed as a rational map

$$\varphi(x, y) = \left(\varphi_1(x, y)/\psi(x, y)^2, \varphi_2(x, y)/\psi(x, y)^3 \right),$$

where φ_1, φ_2 and ψ are polynomials over \mathbb{F}_p in the variables x and y of the original curve E . Another way to say this is that the polynomials must satisfy the relation

$$\varphi_2^2 = \varphi_1^3 + A_2\varphi_1\psi^4 + B_2\psi^6.$$

The points $(x, y) \in E(\overline{\mathbb{F}_p})$ which are roots of the polynomial $\psi(x, y)$ are the points in the kernel of the isogeny.

We gather a few important facts about isogenies.

Proposition 4. *Let E_1 be an elliptic curve over the finite field \mathbb{F}_p .*

1. *Any isogeny $\varphi : E_1 \rightarrow E_2$ factors as*

$$E_1 \xrightarrow{\pi^m} E_1^{(p^m)} \xrightarrow{\varphi'} E_2$$

where π^m is the m th power Frobenius map (where p^m is the inseparable degree of the isogeny φ) and where φ' is a separable isogeny.

2. *The degree of a separable isogeny $\varphi : E_1 \rightarrow E_2$ is equal to the number of points in the subgroup $\ker(\varphi)$.*
3. *Given a finite subgroup C of E_1 there is a unique elliptic curve E_2 and a unique separable isogeny $\varphi : E_1 \rightarrow E_2$ such that $\ker(\varphi) = C$. If C and E_1 are defined over \mathbb{F}_p then so are φ and E_2 .*

Proof. See [32, II.2.12, III.4.10 and III.4.12]. □

Property 1 of the Proposition means that, for the purposes of this paper, we may restrict attention to separable isogenies. A separable isogeny of degree d will often be called a d -isogeny.

A.4. Endomorphism rings of elliptic curves

The *endomorphism ring* $\text{End}(E)$ of an elliptic curve E is the set of all isogenies $\varphi : E \rightarrow E$ which are defined over $\overline{\mathbb{F}_p}$. This is a ring where addition is inherited from addition on E and where multiplication is composition of isogenies.

The fundamental result (see [32, Theorem V.3.1]) is that, for elliptic curves over finite fields, $\text{End}(E)$ is either an order in an imaginary quadratic field, or an order in a quaternion algebra. The latter case occurs if and only if E is a supersingular elliptic curve. In the non-supersingular (usually called ‘ordinary’) case, the endomorphisms are all defined over the base field \mathbb{F}_p .

An *imaginary quadratic field* is $K = \mathbb{Q}(\sqrt{d})$ where d is a negative integer which is assumed to have no square factors. The *discriminant* of the field K is either $D = 4d$ if $d \not\equiv 1 \pmod{4}$ or $D = d$ if $d \equiv 1 \pmod{4}$. An *order* in K is a subring of K (which contains 1). The *maximal order* in K is the ring $\mathcal{O}_K = \mathbb{Z}[(D + \sqrt{D})/2]$. Every order in K is of the form $\mathcal{O} = \mathbb{Z} + c\mathcal{O}_K = \mathbb{Z}[c(D + \sqrt{D})/2]$. The integer $c = [\mathcal{O}_K : \mathcal{O}]$ is called the

conductor of the order \mathcal{O} . The discriminant of the order \mathcal{O} is c^2D which shows that orders in quadratic imaginary fields are uniquely determined by their discriminant.

A.5. *More on endomorphisms and isogenies*

We will need a deeper understanding of isogenies. The theory of complex multiplication will be our main tool.

The theory of complex multiplication is most easily stated in the context of elliptic curves over the complex numbers. Elliptic curves over \mathbb{C} are isomorphic to complex tori $\mathbb{C}/\langle 1, \tau \rangle$ where $\tau \in \mathbb{C}$ has strictly positive imaginary part.

Suppose E has complex multiplication. In this case, $\text{End}(E) = \{\alpha \in \mathbb{C} : \alpha, \alpha\tau \in \langle 1, \tau \rangle\}$ and so $\alpha = a + b\tau \in \text{End}(E)$. The condition $\alpha\tau \in \langle 1, \tau \rangle$ implies that τ satisfies some equation $A\tau^2 + B\tau + C = 0$ (that is, τ lies in a quadratic imaginary field K). We may choose A, B, C to be integers such that $(A, B, C) = 1$ (that is, there is no prime dividing all three). The field K is equal to $\mathbb{Q}(\sqrt{B^2 - 4AC})$ and $\text{End}(E)$ is an order in K . The following lemma is extremely useful.

Lemma 3. (See [15, Theorem 8.1].) *Let $E = \mathbb{C}/\langle 1, \tau \rangle$ be an elliptic curve where τ satisfies $A\tau^2 + B\tau + C$ with $A, B, C \in \mathbb{Z}$ and $(A, B, C) = 1$. Then $\text{End}(E)$ is the order with discriminant $B^2 - 4AC$.*

Furthermore, writing $\mathcal{O} = \text{End}(E)$, the lattice $\langle 1, \tau \rangle$ is a projective \mathcal{O} -module, and two elliptic curves $\mathbb{C}/\langle 1, \tau \rangle$ and $\mathbb{C}/\langle 1, \tau' \rangle$ are isomorphic if and only if these \mathcal{O} -modules are in the same class in the group $\text{Pic}(\mathcal{O})$ (which is the group of projective \mathcal{O} -modules modulo the principal \mathcal{O} -modules). Hence, we may represent the isomorphism classes of elliptic curves over \mathbb{C} with $\text{End}(E) = \mathcal{O}$ by the elliptic curves \mathbb{C}/\mathfrak{a} where \mathfrak{a} runs over the classes in $\text{Pic}(\mathcal{O})$. These elliptic curves are defined over the ring class field $H_{\mathcal{O}}$. It follows that the number of isomorphism classes of elliptic curves E having $\text{End}(E) = \mathcal{O}$ is equal to the class number $h_{\mathcal{O}}$ of the order \mathcal{O} (that is, the number of elements of the group $\text{Pic}(\mathcal{O})$).

Let \mathfrak{b} be an \mathcal{O} -ideal. Then the identity map from \mathbb{C} to itself induces the map

$$\mathbb{C}/\mathfrak{a} \rightarrow \mathbb{C}/\mathfrak{a}\mathfrak{b}^{-1}$$

which has kernel $\mathfrak{a}\mathfrak{b}^{-1}/\mathfrak{a} \simeq \mathcal{O}/\mathfrak{b}$. This is an isogeny of degree $N_{K/\mathbb{Q}}(\mathfrak{b})$, and all isogenies arise in this way.

We require the theorems of Deuring on reduction/lifting of the endomorphism ring (a good reference is [15, Theorem 13.12 and Theorem 13.14]). It is necessary to assume that the conductor of $\text{End}(E)$ is coprime to p . This condition is automatic for elliptic curves modulo p . When generalising to the case of small characteristic fields we would use techniques due to Couveignes [8] and Lercier [19], [20] to find an isogeny to an elliptic curve with no power of p dividing the conductor of its endomorphism ring. We state the theorems in the form which we will need, and refer to [15] for the details.

Theorem 2. (Deuring) *Let \bar{E}/\mathbb{F}_p be an elliptic curve with a non-trivial endomorphism $\bar{\varphi}$. Then there is an elliptic curve E over a number field L , an endomorphism φ of E , and a prime \wp of L above p such that E and φ reduce to \bar{E} and $\bar{\varphi}$ modulo \wp .*

Theorem 3. (Deuring) *Let E be an elliptic curve over a number field L such that $\text{End}(E)$ is an order \mathcal{O} in a quadratic imaginary field K . Let p be a rational prime such that E has good reduction at p and such that p is coprime to the conductor of the order \mathcal{O} . Let \wp be a prime of \mathbb{Q} above p . Write \bar{E} for the reduction of E modulo \wp . The curve \bar{E} is*

supersingular if and only if the prime p is inert or ramified in the quadratic extension K/\mathbb{Q} . If the prime p splits in K then the modulo \wp reduction map $E \rightarrow \bar{E}$ induces an isomorphism $\text{End}(E) \cong \text{End}(\bar{E})$.

The upshot of these theorems is that elliptic curves and their endomorphism rings lift from \mathbb{F}_p to number fields (in fact, to the ring class field), and that they reduce well modulo p too. We may therefore apply the complex multiplication theory of elliptic curves over \mathbb{C} to the situation of elliptic curves modulo p .

In the application we consider two isogenous elliptic curves E_1 and E_2 over \mathbb{F}_p such that $\text{End}(E_1)$ and $\text{End}(E_2)$ are orders lying between \mathcal{O}_K and $\mathbb{Z}[\pi]$. Therefore, their conductors both divide the conductor $c = [\mathcal{O}_K : \mathbb{Z}[\pi]]$. It will be necessary to construct isogenies between curves whose endomorphism rings have different conductor, so we study some of these issues here.

As mentioned in the main body of the paper (Proposition 1), the only way for the conductor to change by a prime l is if one takes an isogeny whose degree is a multiple of l . Indeed, by taking a sequence of l -isogenies (for those primes l dividing the conductor of $\text{End}(E)$) it is possible to find an elliptic curve E' such that $\text{End}(E') = \mathcal{O}_K$ (the maximal order). These ideas are central to Kohel’s method (see [14]) for finding the exact endomorphism ring of a given ordinary elliptic curve E/\mathbb{F}_p .

We can be even more complete in our analysis of isogenies. The following theorem (which is essentially Propositions 22 and 23 of [14]) tells us everything we need. We use the language of Kohel, so an l -isogeny ‘down’ is an isogeny $\varphi : E_1 \rightarrow E_2$ of degree l such that $[\text{End}(E_1) : \text{End}(E_2)] = l$ whilst an l -isogeny ‘up’ is one with $[\text{End}(E_2) : \text{End}(E_1)] = l$. In the case where the endomorphism rings are preserved we call the isogeny ‘horizontal’. We say that an elliptic curve E is ‘on the surface at l ’ if $l \nmid [\mathcal{O}_K : \text{End}(E)]$. We say that E is ‘on the floor at l ’ if $l \nmid [\text{End}(E) : \mathbb{Z}[\pi]]$. We might also say that E is ‘of level n at l ’ if $l^n \parallel [\mathcal{O} : \text{End}(E)]$.

Theorem 4. [14] *Let E be an elliptic curve over \mathbb{F}_p having endomorphism ring \mathcal{O} where $\mathbb{Z}[\pi] \subseteq \mathcal{O} \subseteq \mathcal{O}_K$. Let l be a prime number. The following list classifies the possibilities for the l -isogenies defined over \mathbb{F}_p .*

- If $l \nmid [\mathcal{O}_K : \mathcal{O}]$ then the number of l -isogenies to elliptic curves with endomorphism ring equal to \mathcal{O} is $1 + \left(\frac{D}{l}\right)$.
- If $l \mid [\mathcal{O}_K : \mathcal{O}]$ then there is one l -isogeny up to an elliptic curve.
- If $l \nmid [\mathcal{O} : \mathbb{Z}[\pi]]$ then there are no l -isogenies down.
- If $l \mid [\mathcal{O} : \mathbb{Z}[\pi]]$ and $l \mid [\mathcal{O}_K : \mathcal{O}]$ then the number of l -isogenies down is l .
- If $l \mid [\mathcal{O} : \mathbb{Z}[\pi]]$ and $l \nmid [\mathcal{O}_K : \mathcal{O}]$ then the number of l -isogenies down is $l - \left(\frac{D}{l}\right)$.

Note. In each case, when there are several different isogenies to elliptic curves of the same level then some of the image elliptic curves may actually be isomorphic. This behaviour can be explained (see [14]) but it will not significantly affect our discussion.

Proof. The Deuring lifting theorems allow us to lift the elliptic curve E from \mathbb{F}_p to \mathbb{C} in such a way that the endomorphism ring is preserved. Suppose that $E = \mathbb{C}/\langle 1, \tau \rangle$ where τ satisfies $A\tau^2 + B\tau + C$ (where $(A, B, C) = 1$). So $\text{disc}(\mathcal{O}) = D = B^2 - 4AC$. Then there are $l + 1$ possibilities for the kernel of an l -isogeny, and they are $\iota_0 = \langle \frac{1}{l}, \tau \rangle$ and $\iota_k = \langle 1, (\tau + k)/l \rangle$ where $k = 1, \dots, l$. The image curves in these cases are

$$E_0 = E/\iota_0 \simeq \mathbb{C}/\langle 1, l\tau \rangle$$

and

$$E_k = E/\mathfrak{l}_k \simeq \mathbb{C}/\langle l, k + \tau \rangle.$$

Our goal is to determine $\text{End}(E_k)$. We will do this using Lemma 3.

Firstly we consider E_0 . The number $\alpha = l\tau$ satisfies $A\alpha^2 + lB\alpha + l^2C = 0$. If $(A, lB, l^2C) = 1$ then $\text{End}(E_0)$ has discriminant l^2D , and so \mathfrak{l}_0 is an l -isogeny down. On the other hand, if $(A, lB, l^2C) \neq 1$ then it follows that $l \mid A$ (so put $A' = A/l$). If $l \nmid B$ then α is actually a root of $A'\alpha^2 + B\alpha + lC = 0$, from which we see that $\text{End}(E_0) = \mathcal{O}$ (that is, the isogeny is horizontal). If $l \mid B$ (so write $B' = B/l$) and $l \mid [\mathcal{O}_K : \mathcal{O}]$ then, since $D = B^2 - 4AC$, it follows that $l^2 \mid A$ (so put $A' = A/l^2$). This means that α satisfies $A'\alpha^2 + B'\alpha + C = 0$ (which has discriminant D/l^2), and hence we have an l -isogeny up.

Now consider the elliptic curves E_k for $k = 1, 2, \dots, l$. The number $\alpha = (k + \tau)/l$ is a root of

$$l^2A\alpha^2 + l(B - 2Ak)\alpha + (Ak^2 - Bk + C) = 0.$$

In the case $(l^2A, l(B - 2Ak), (Ak^2 - Bk + C)) = 1$ we see that $\text{End}(E_k)$ has discriminant l^2D , and we have an isogeny down.

The condition $(l^2A, l(B - 2Ak), (Ak^2 - Bk + C)) = 1$ fails if and only if $l \mid (Ak^2 - Bk + C)$. Note that there are several possibilities for the solubility of the equation $Ak^2 - Bk + C \equiv 0 \pmod{l}$.

If $l \mid A$ and $l \mid B$ then there is no solution. In this case $l \mid D$. In the case where $l \mid [\mathcal{O}_K : \mathcal{O}]$ we have already found a single l -isogeny up. If l does not divide this index then it follows that l ramifies in K , and we have already found a single isogeny to an elliptic curve with endomorphism ring equal to \mathcal{O} .

If $l \mid A$ but $l \nmid B$ then we had already found an l -isogeny previously. In this case there is also the value $k = C/B$ which will give a horizontal l -isogeny. Thus, in this case, we have two horizontal isogenies, and the prime l splits in K .

If $l \nmid A$ then the equation is a true quadratic. There is a repeated root if and only if $l \mid B^2 - 4AC$ (which again corresponds to the ramified case handled above), and so there is only one horizontal solution. Otherwise there are two distinct solutions (equivalently, the prime l splits in K) and we obtain two horizontal l -isogenies. The final case is when l is inert in K . In this case there will be no solutions to the quadratic (that is, all values for k will give an l -isogeny down, and so we get a total of l isogenies down).

Finally, we must contemplate the Deuring reduction step. We reduce the elliptic curves E_k from \mathbb{C} to some finite field \mathbb{F}_{p^m} . These elliptic curves will actually be defined over \mathbb{F}_p if and only if their endomorphism ring contains $\mathbb{Z}[\pi]$. This completes the proof of the classification. \square

A.6. Modular curves

The modular curves $X_0(N)$ are a geometric tool which is of great value when studying isogenies. The standard equation $\Phi_N(x, y) \in \mathbb{Z}[x, y]$ for $X_0(N)$ as a plane algebraic curve is given by the relation

$$\Phi_N(j(\tau), j(N\tau)) = 0 \tag{8}$$

between the classical modular functions $j(\tau)$ and $j(N\tau)$ on $X_0(N)$.

The most important fact for the current application is that, if E is a given elliptic curve, then the elliptic curves E' which are N -isogenous to E are precisely those curves (up to

isomorphism) whose j -invariant is a root of

$$\Phi_N(j(E), y) = 0.$$

The equation (8) is the classical equation for $X_0(N)$ and it is the most useful for theoretical purposes. For practical computation, its degree and coefficients are much too large. There are many other ways to get modular equations which are just as useful but which have smaller degree and smaller integer coefficients. See the paper of Elkies [11] for information about this.

For this article we will stick to using (8). We will need a crude complexity estimate for calculating (8). Elkies [11, Section 3] notes that $\Phi_l(X, Y)$ requires $O(l^3)$ arithmetic operations in \mathbb{Z} to compute. If we work over \mathbb{F}_p then we require $O(l^3)$ arithmetic operations in \mathbb{F}_p and $O(l^2)$ elements of \mathbb{F}_p to store the resulting equation.

A.7. Vélu's formulae

The following explicit formulae come directly from Vélu's paper [35]. We have simplified them a little (to the case where the characteristic of K is not 2 or 3, and where we use a simpler Weierstrass model) and we reproduce them for the convenience of the reader.

Let $E : y^2 = x^3 + a_2x^2 + a_4x + a_6$ be an elliptic curve over a field K (so $a_1 = a_3 = 0$). Suppose there is a cyclic subgroup C of order n given by a polynomial $\psi(x)$, by which we mean

$$C = \{\infty\} \cup \left\{ (\alpha, \beta = \pm\sqrt{\alpha^3 + a_2\alpha^2 + a_4\alpha + a_6}) \in E(\overline{K}) : \psi(\alpha) = 0 \right\}.$$

We list the roots of $\psi(x)$ in two sets: the set F_2 will be those $\alpha \in \overline{K}$ such that $\psi(\alpha) = 0$ and $\alpha^3 + a_2\alpha^2 + a_4\alpha + a_6 = 0$ (that is, we have the 2-torsion point $(\alpha, 0)$), and the set R will be the set consisting of the rest of the roots of $\psi(x)$. Hence $\#C = 1 + \#F_2 + 2\#R$ and $\deg(\psi) = \#F_2 + \#R$.

For each $\alpha \in F_2 \cup R$, let (α, β) be one of the (one or two) corresponding points on $E(\overline{K})$ and define

$$\begin{aligned} g_\alpha^x &= 3\alpha^2 + 2a_2\alpha + a_4; \\ g_\alpha^y &= -2\beta; \\ t_\alpha &= g_\alpha^x \text{ if } \alpha \in F_2, \text{ or } 2g_\alpha^x \text{ if } \alpha \in R; \\ u_\alpha &= (g_\alpha^y)^2. \end{aligned}$$

The isogeny with kernel C is given by

$$\begin{aligned} X &= x + \sum_{\alpha \in F_2 \cup R} \left(t_\alpha / (x - \alpha) + u_\alpha / (x - \alpha)^2 \right), \\ Y &= y - \sum_{\alpha \in F_2 \cup R} \left(u_\alpha 2y / (x - \alpha)^3 + t_\alpha (y - \beta) / (x - \alpha)^2 - g_\alpha^x g_\alpha^y / (x - \alpha)^2 \right). \end{aligned}$$

Let $t = \sum_{\alpha \in F_2 \cup R} t_\alpha$ and $w = \sum_{\alpha \in F_2 \cup R} (u_\alpha + \alpha t_\alpha)$. Then the image of the isogeny is the curve

$$Y^2 = X^3 + a_2X^2 + (a_4 - 5t)X + (a_6 - 4a_2t - 7w).$$

Evaluating Vélu's formulae for an isogeny of degree N takes $O(N)$ time and space.

A.8. *Work of Elkies*

Vélu's formulae make the isogeny described in Proposition 4(3) explicit. Given the kernel of the isogeny (in terms of $\psi(x)$) Vélu's formulae provide us with equations for the image elliptic curve and the polynomials $\varphi_1(x, y)$, $\varphi_2(x, y)$ of equation (1).

We need a method to calculate the polynomial $\psi(x)$ associated with an N -isogeny between two elliptic curves given only their j -invariants. Such a method has been developed by Elkies in the context of algorithms for counting points on elliptic curves over finite fields.

The basic idea is to use relations (which are found by working over \mathbb{C}) between certain classical modular forms and functions. From using only the two j -values one may obtain all the data necessary to construct the polynomial $\psi(x)$. There are several references for these formulae (e.g., [10], [2], [29], [24], [5]).

Finding the equation $\psi(x)$ for an N -isogeny using these methods takes $O(N^3)$ field operations and space $O(N^2)$ field elements (though the final result only requires space $O(N)$).

References

1. A. O. L. ATKIN, 'The number of points on an elliptic curve modulo a prime', Preprint, 1988. 118
2. A. O. L. ATKIN, 'The number of points on an elliptic curve modulo a prime 2', Preprint, 1992. 118, 136
3. A. O. L. ATKIN and F. MORAIN, 'Finding suitable curves for the elliptic curve method of factorization', *Math. Comp.* 60 (1993) 399–405. 128
4. E. BACH, *Analytic methods in the analysis and design of number-theoretic algorithms*, Berkeley Doctoral Thesis, ACM Distinguished Dissertations (MIT Press, 1984) 123
5. I. BLAKE, G. SEROUSSI and N. P. SMART, *Elliptic curves in cryptography*, LMS Lecture Notes 265 (Cambridge University Press, 1999). 136
6. J. W. S. CASSELS, *Lectures on elliptic curves*, LMS Student Texts 24 (Cambridge University Press, 1991) 130
7. H. COHEN, *A course in computational algebraic number theory*, Grad. Texts in Math. 138 (Springer, New York, 1993). 120, 123, 123, 127, 130
8. J.-M. COUVEIGNES, 'Computing l -isogenies using the p -torsion', *ANTS-II*, Lecture Notes in Comput. Sci. 1122 (ed. H. Cohen, Springer, New York, 1996) 59–65. 118, 119, 132
9. J.-M. COUVEIGNES and F. MORAIN, 'Schoof's algorithm and isogeny cycles', *ANTS-I*, Lecture Notes in Comput. Sci. 877 (ed. L. M. Adleman, Springer, New York, 1994) 43–58. 118
10. N. ELKIES, 'Explicit isogenies', Preprint, 1991. 118, 121, 136
11. N. ELKIES, 'Elliptic and modular curves over finite fields and related computational issues', *Computational perspectives on number theory: proceedings of a conference in honor of A. O. L. Atkin* (ed. D. A. Buell and J. T. Teitelbaum, AMS, 1997) 21–76. 118, 124, 135, 135
12. G. FREY and H.-G. RÜCK, 'A remark concerning m -divisibility and the discrete logarithm problem in the divisor class group of curves', *Math. Comp.* 62 (1994) 865–874. 120

13. S. D. GALBRAITH and J. MCKEE, 'The probability that the number of points on an elliptic curve over a finite field is prime', Preprint, 1999. 128
14. D. KOHEL, 'Endomorphism rings of elliptic curves over finite fields', Berkeley PhD thesis, 1996. 119, 122, 122, 122, 123, 124, 124, 124, 133, 133, 133, 133
15. S. LANG, *Elliptic functions*, 2nd edn, Grad. Texts in Math. 112 (Springer, New York, 1987). 122, 122, 123, 132, 132, 132
16. G.-J. LAY and H. G. ZIMMER, 'Constructing elliptic curves with given group order over large finite fields', *ANTS-I*, Lecture Notes in Comput. Sci. 877 (ed. L. M. Adleman, Springer, New York, 1994) 250–263. 128
17. H.W. LENSTRA JR., 'Finding isomorphisms between finite fields', *Math. Comp.* 56 (1991) 329–347. 120
18. H.W. LENSTRA JR., 'Complex multiplication structure of elliptic curves', *J. Number Theory* 56 (1996) 227–241. 123
19. R. LERCIER, 'Computing isogenies in \mathbb{F}_{2^n} ', *ANTS-II*, Lecture Notes in Comput. Sci. 1122 (ed. H. Cohen, Springer, New York, 1996) 197–212. 118, 119, 132
20. R. LERCIER and F. MORAIN, 'Algorithms for computing isogenies between elliptic curves', *Computational perspectives on number theory: proceedings of a conference in honor of A. O. L. Atkin* (ed. D. A. Buell and J. T. Teitelbaum, AMS, 1997) 77–96. 118, 119, 132
21. J.-F. MESTRE, 'La methode des graphes. Exemples et applications', *Class numbers and fundamental units of algebraic number fields*, Proc. Int. Conf., Katata, Japan (Nagoya University, 1986) 217–242. 119, 119
22. A. MENEZES, T. OKAMOTO and S. VANSTONE, 'Reducing elliptic curve discrete logarithms to logarithms in a finite field', *IEEE Trans. Inform. Theory* 39 (1993) 1639–1646. 120
23. F. MORAIN, 'Building cyclic elliptic curves modulo large primes', *EUROCRYPT '91*, Lecture Notes in Comput. Sci. 549 (Springer, New York, 1991) 328–336. 128
24. V. MUELLER, 'Ein algorithmus zur bestimmung der punktzahl elliptischer kurven über endlichen körpern der charakteristik grösser drei', PhD. Thesis, Universität des Saarlandes, 1995. 136
25. J. POLLARD, 'Monte Carlo methods for index computation (mod p)', *Math. Comp.* 32 (1978) 918–924. 120
26. H.-G. RÜCK, 'On the discrete logarithm problem in the divisor class group of curves', *Math. Comp.*, to appear. 120
27. T. SATOH and K. ARAKI, 'Fermat quotients and the polynomial time discrete log algorithm for anomalous elliptic curves', *Comment. Math. Univ. St. Paul*, 47 (1998) 81–92. 120
28. R. SCHOOF, 'Elliptic curves over finite fields and the computation of square roots mod p ', *Math. Comp.* 44 (1985) 483–494. 118
29. R. SCHOOF, 'Counting points on elliptic curves over finite fields', *J. Théor. Nombres Bordeaux* 7 (1995) 219–254. 136
30. I. SEMAEV, 'Evaluation of discrete logarithms in a group of p -torsion points of an elliptic curve in characteristic p ', *Math. Comp.* 67 (1998) 353–356. 120

31. G. SHIMURA, *Introduction to the arithmetic theory of automorphic functions* (Iwanami/Princeton, 1971). 122, 123
32. J. H. SILVERMAN, *The arithmetic of elliptic curves*, Grad. Texts in Math. 106 (Springer, New York, 1986). 119, 130, 130, 130, 131, 131
33. N. P. SMART, ‘The discrete logarithm problem on elliptic curves of trace one’, *J. Cryptology*, to appear. 120
34. J. TATE, ‘Endomorphisms of abelian varieties over finite fields’, *Invent. Math.* 2 (1966) 134–144. 118, 123
35. J. VÉLU, ‘Isogénies entre courbes elliptiques’, *C. R. Acad. Sci. Paris Sér. I*, 273 (1971) 238–241. 121, 135

Steven D. Galbraith s.galbraith@rhbnc.ac.uk

Mathematics Department
Royal Holloway University of London
Egham, Surrey TW20 0EX