Computing the trace of an endomorphism of a supersingular elliptic curve

Michael Thomas Wills

Thesis submitted to the faculty of the Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Master of Science
in
Mathematics

Gretchen Matthews, Chair
Travis Morrison, Co-chair
Daniel Orr
Mark Shimozono

May 5th, 2021
Blacksburg, Virginia

Computing the trace of an endomorphism of a supersingular elliptic curve

Michael Thomas Wills

Abstract

We provide an explicit algorithm for computing the trace of an endomorphism of an elliptic curve which is given by a chain of small-degree isogenies. We analyze its complexity, determining that if the length of the chain, the degree of the isogenies, and the log of the field size are all $O(n)$, the trace of the endomorphism can be computed in $O(n^6)$ bit operations. This makes explicit a theorem of Kohel which states that such a polynomial time algorithm exists. The given procedure is based on Schoof's point-counting algorithm.

Computing the trace of an endomorphism of a supersingular elliptic curve

Michael Thomas Wills

General Audience Abstract

The developing technology of quantum computers threatens to render current cryptographic systems (that is, systems for protecting stored or transmitted digital information from unauthorized third parties) ineffective. Among the systems proposed to ensure information security against attacks by quantum computers is a cryptographic scheme known as SIKE. In this thesis, we provide and analyze an algorithm that comprises one piece of a potential attack against SIKE by a classical computer. The given algorithm is also useful more generally in the field of arithmetic geometry.

Acknowledgements

This thesis would not have been possible without the beneficial influence of many people. First and foremost is Travis Morrison, who led me into this project, answered hundreds of my questions (often more than once), and has supplied encouragement, criticism, and a genuine interest in my well-being and mathematical future. Of equal importance is Gretchen Matthews, for shepherding me into the world of real math and persistently providing me with opportunities to become a better mathematician. Additionally, I thank Daniel Valvo (and though he doesn't know it, Ben Goodberry) for help with formatting.

Finally, I thank all of my non-mathematical family and friends who listened to me trying to explain what I was doing time and time again for their patience and support, even if they're still not entirely sure what this thesis is about.

This work was supported in part by the Commonwealth Cyber Initiative, an investment in the advancement of cyber R&D, innovation, and workforce development. For more information about CCI, visit www.cyberinitiative.org.

Contents

1 Introduction

With the advent of quantum computing, many of the current cryptographic systems used to protect digital information will become ineffective [Sho94]. As such, there is much current interest (including a competition sponsored by the National Institute of Standards and Technology [NIS17]) to find new cryptographic systems that are resistant to attacks by quantum computers. One proposed system is SIKE, or the Supersingular Isogeny Key Encapsulation protocol [Aza+20]. The security of this protocol is based on the believed hardness of finding paths in the ℓ-isogeny graph from one given supersingular elliptic curve to another for some small prime ℓ.

The problem of pathfinding in the ℓ-isogeny graph is heuristically equivalent to that of efficiently computing the endomorphism ring of a supersingular elliptic curve [GPS17, Section 2.2]. Such curves are defined by their endomorphism rings being orders in quaternion algebras. The analogous problem of determining the endomorphism ring of an ordinary elliptic curve (which has endomorphism ring isomorphic to an order in an imaginary quadratic field) *has* been solved in sub-exponential time under certain heuristics in [BS11], so this approach warrants further investigation.

An order $\mathcal{O}$ in a quaternion algebra is completely described by giving a basis for $\mathcal{O}$ as a rank four $\mathbb{Z}$-module along with a multiplication table for inner products on that basis; see [Voi21] for a detailed reference on quaternion algebras. In the case of endomorphism rings, the basis can be provided in the form $\{[1], \alpha, \beta, \gamma\}$, where [1] denotes the identity map and α, β, and γ are linearly independent non-integer endomorphisms. Inner products on the endomorphism ring are computed as traces of those endomorphisms:

$$\langle \alpha, \beta \rangle = \frac{1}{2} \operatorname{tr} \alpha\hat{\beta},$$

where $\hat{\beta}$ is the dual endomorphism to β.

Unfortunately, the average size of the smallest endomorphism of an elliptic curve scales up with the size of the field of definition; the number of elliptic curves defined over the finite field $\mathbb{F}_{p^2}$ having an endomorphism of degree less than M is roughly $O(M^{3/2})$, about half of which are supersingular [LB20]. For p of cryptographic size (that is, $p \approx 2^{256}$), it is only reasonable to explicitly store and manipulate endomorphisms of degree polylogarithmic in p, but by the above we cannot generally expect to find such small endomorphisms. We are thus moved to represent endomorphisms in a different way: as chains of small-degree isogenies whose composition gives a large-degree endomorphism.

The first question here is how exactly one finds a basis for the endomorphism ring. More directly, how can we find just one chain of small-degree isogenies whose composition is a non-integer endomorphism of an arbitrary supersingular elliptic curve? This is a question of cycle-finding in isogeny graphs, and is currently understood to be the asymptotically difficult part of the problem. It will not be treated in this thesis.

The second question is this: given an endomorphism represented as a chain of small-degree isogenies, how does one compute inner products with it? That is, how does one compute the trace? In [Koh96, Theorem 81], Kohel asserts the existence of a fast algorithm for computing such a trace using a modified version of Schoof's algorithm. Later, in [Ban+19, Appendix], Banks, Camacho-Navarro, Eisenträger, Morrison, and Park sketch a

more detailed outline of Kohel's proposed algorithm and give an initial complexity analysis, showing that the algorithm will run in time roughly $O(\log^7 p)$ for common sizes of parameter. In this thesis, we present the algorithm in full detail, perform a more careful complexity analysis giving a runtime of roughly $O(\log^6 p)$ for the use cases described in [Ban+19], and offer a small improvement that improves the asymptotic complexity to about $O(\log^5 p)$ in specific parameter regimes.

The structure of this thesis is as follows. In Section 2, we provide a brief background on the theory of elliptic curves and isogenies between them, and in Section 3 we discuss endomorphisms of elliptic curves in more detail. After that, we formalize and give algorithms for working with isogenies restricted to subgroups of elliptic curves in Section 4, and then analyze the complexity of these operations as well as some common polynomial operations in Section 5. In Section 6, we say a few words specifically about supersingular elliptic curves. In Section 7, we discuss Schoof's original algorithm for point-counting on elliptic curves. Finally, in Section 8 we present a modified version of Schoof's algorithm for computing traces of endomorphisms given by chains of isogenies, analyze its complexity, and discuss an attempted improvement for the supersingular case.

2 Preliminary definitions

Much of the following exposition is paraphrased from [Sil86, Chapters I-III] and [Sut19, Lectures 2, 5, and 6]. The reader who has not seen this material before or wants a deeper understanding is recommended to consult these sources.

2.1 Elliptic curves as varieties

Let F be a perfect field. An **elliptic curve** defined over F is a smooth projective plane curve of genus 1 defined over F with a marked F-rational point $\mathcal{O}$. As a first result, it can be shown that every elliptic curve over F is defined by an equation of the form

$$E : Y^2Z + a_1XYZ + a_3YZ^2 = X^3 + a_2X^2Z + a_4XZ^2 + a_6Z^3,$$

where $a_1, a_2, a_3, a_4, a_6 \in F$ and the marked point $\mathcal{O}$ is the unique point at infinity given by $\mathcal{O} = [0 : 1 : 0]$; see [Sil86, Proposition III.3.1]. An equation of this form is called a **Weierstrass equation**, and we write the affine component of E as

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

Further, if we assume that $\operatorname{char} F \neq 2, 3$, then every elliptic curve over F is defined by an equation of the form

$$E : y^2 = x^3 + Ax + B. \tag{1}$$

for some $A, B \in F$, where again $\mathcal{O}$ is the unique point at infinity [Sil86, Section III.1]. Equations following this pattern are called **short Weierstrass equations**, and elliptic curves expressed as such are referred to as being in **short Weierstrass form**. For a curve C given by a short Weierstrass equation, one can show that C is smooth if and only if $\Delta = -16(4A^3 + 27B^2)$ is non-zero, where Δ is called the **discriminant** of C [Sil86, Section

III.1]. Thus, just as every elliptic curve can be described by a short Weierstrass equation with non-zero discriminant, so too does every short Weierstrass equation with non-zero discriminant give rise to an elliptic curve. However, this correspondence is not in general bijective, as multiple short Weierstrass equations may describe the same elliptic curve.

Let E be an elliptic curve defined over F, and let F be a subfield of K. We denote by $E(K)$ the set of all **K-rational points of** E, which are the projective points $[X : Y : Z]$ of E with some representative (X_0, Y_0, Z_0) such that $X_0, Y_0, Z_0 \in K$. Note that if we take $F = \mathbb{F}_q$ to be a finite field of size q, then because any specific element in an algebraic extension of $\mathbb{F}_q$ is contained in a finite subfield $\mathbb{F}_{q^r}$, we must have that all points of $E(\overline{\mathbb{F}_q})$ are in fact contained in $E(K)$ for some finite field K, where $\overline{\mathbb{F}_q}$ is an algebraic closure of $\mathbb{F}_q$.

The points of E other than $\mathcal{O}$ are known as the **affine points** of E, and have a unique representative $(x, y, 1)$, which will be denoted by (x, y). Additionally, we note that because the only power of y occurring in the short Weierstrass equation is a square, for any point $(x, y) \in E(K)$ the point $(x, -y)$ is also in $E(K)$, and that these are the only points of E with that x-coordinate.

2.2 The group law

One of the main features of elliptic curves is the fact that their points can be made into an abelian group with the point $\mathcal{O}$ as the identity element. Addition in this group is summarized by the rule "three collinear points sum to zero" and can be more exactly be described as follows for affine points P_1 and P_2 (since the only point at infinity is the identity, this captures every nontrivial case). Let $P_1 = (x_1, y_1)$ and $P_2 = (x_2, y_2)$. We first suppose that $x_1 = x_2$ but $P_1 \neq P_2$. In this case, P_1 and P_2 must be of the form $P_1 = (a, b)$ and $P_2 = (a, -b)$ with $b \neq 0$. Since these points lie on the line $x = b$, they are collinear with $\mathcal{O}$, and the group law gives that $P_1 = -P_2$. Supposing now that $x_1 \neq x_2$, we set m to be the slope of the line between the two points as

$$m = \frac{y_1 - y_2}{x_1 - x_2}.$$

Since Equation (1) is cubic, there are exactly three points (counting multiplicity) lying on both the line $(y - y_1) = m(x - x_1)$ and E, of which P_1 and P_2 are two. We denote the third point by R, and note that by the rule we must have that $P_1 + P_2 = -R$. Writing $P_1 + P_2 = (x_3, y_3)$, we deduce that

$$x_3 = m^2 - x_1 - x_2 \text{ and}$$
$$y_3 = m(x_1 - x_3) - y_1.$$

Finally, suppose $P_1 = P_2$. In this case, we take the tangent line to E at P_1 to find the third collinear point R then negate as before. If $y_1 = 0$, then the tangent line must be vertical and $R = -R = \mathcal{O}$. If $y_1 \neq 0$, then we set m to be the slope of the tangent line to E at (a, b) given by

$$m = \frac{3x_1^2 + A}{2y_1}$$

and define x_3 and y_3 exactly as above. The group law being entirely defined, we note that inverses in the group are given by simply negating the y-coordinate. Further, one can check

that this proposed group law is in fact associative and commutative; see [Cas91, page 28] for two full proofs, [Sut19, Section 2.2.1] for an adaptation of the geometric proof in [Cas91], or [Sil86, Proposition 3.4e] for a divisor-based argument[1]. Thus, we have a genuine abelian group on the points of E. Finally, we observe that since the group law is defined entirely in terms of elementary field operations, if P_1 and P_2 both lie in $E(K)$, then their sum must as well.

2.3 The group structure of E

Having established E as an abelian group, we can begin to investigate its structure. We will focus exclusively on the case where F is a finite field and K is an algebraic extension of F, noting that the group structure of $E(K)$ in general depends greatly upon both the field K and the coefficients A and B used to construct the curve. Firstly, we observe that for any point $P \in E(K)$, we must have that $P \in E(L)$ for some finite algebraic extension L/F. Since multiples of points must remain defined over the same field and $E(L)$ is finite (the affine points being a subset of L^2), there are no points in $E(K)$ of infinite order. One important way to look at the group structure of E is through its **ℓ-torsion subgroups** $E[\ell]$, which are defined for $\ell \in \mathbb{N}$ by

$$E[\ell] = \{P \in E : \ell P = \mathcal{O}\}.$$

These subgroups are very well characterized by the following theorem.

Theorem 2.1. *Let E be an elliptic curve defined over a field F of characteristic p, and let $\ell \in \mathbb{N}$.*

1. *If ℓ is coprime to p, then*
$$E[\ell] \cong \mathbb{Z}/\ell\mathbb{Z} \times \mathbb{Z}/\ell\mathbb{Z},$$
2. *and if $\ell = p^e$, then either*
$$E[\ell] \cong \{\mathcal{O}\} \text{ for all } e$$
or
$$E[\ell] \cong \mathbb{Z}/p^e\mathbb{Z} \text{ for all } e.$$

Proof. [Sil86, Corollary 6.4] □

2.4 Isogenies

Having established both the rational point structure and the group structure on elliptic curves, it is natural to investigate maps between elliptic curves that preserve both of them, called **isogenies**. Let E_1 and E_2 be elliptic curves defined over K, with $E_1 : Y^2Z = X^3 + AZ^2 + BZ^3$. We define isogenies more precisely as maps $\phi : E_1 \to E_2$ satisfying $\phi(\mathcal{O}) = \mathcal{O}$ which can be written as

$$\phi = [\phi_X, \phi_Y, \phi_Z],$$

[1] Silverman mentions that the associative law can also be verified by direct computations, though cautions that this is a tedious and unenlightening task. We have heeded his warning, and so cannot vouch for it.

where ϕ_X, ϕ_Y, and ϕ_Z are homogeneous polynomials in $K[X, Y, Z]$ of the same degree, at least one of which is nonzero at each $P \in E_1$. To be more precise still, we take the coordinate functions to be elements of the quotient ring $K[E_1] \coloneqq K[X, Y, Z]/(Y^2Z - X^3 - AZ^2 - BZ^3)$, so that isogenies are only defined up to their behavior on E_1; this will become relevant in Section 4. The set of all isogenies from E_1 to E_2 is denoted $\mathrm{Hom}(E_1, E_2)$ and has an abelian group structure under pointwise addition. We define the **kernel** of an isogeny ϕ, denoted $\ker(\phi)$, to be its kernel as a group homomorphism; that is, all points P such that $\phi(P) = \mathcal{O}$. If we allow these isogenies to have coordinate functions in the field $K(E_1) = \operatorname{Frac} K[E_1]$ with the understanding that we clear denominators before evaluating, the following lemma gives a simpler expression for isogenies, referred to as the isogeny's **standard form**.

Lemma 2.2. *Let E_1 and E_2 be elliptic curves over a field K of characteristic not 2 or 3, and let $\phi = [\phi_X, \phi_Y, \phi_Z]$ be an isogeny from E_1 to E_2. Then ϕ can be expressed as a rational function on the affine points of E_1 as*

$$\phi(x, y) = \left(\frac{u(x)}{v(x)}, \frac{s(x)}{t(x)}y\right) \coloneqq \left[\frac{u(X)}{v(X)}, \frac{s(X)}{t(X)}Y, Z\right]$$

where $u, v, s, t \in K[x]$ and the pairs (u, v) and (s, t) share no non-unit factors.

Proof. [Sut19, Lemma 5.25] □

Suppose now that there exist isogenies $\phi_1 : E_1 \to E_2$ and $\phi_2 : E_2 \to E_1$ such that either composition of ϕ_1 with ϕ_2 is the identity. Then E_1 and E_2 are said to be **isomorphic**, and the isogenies ϕ_1 and ϕ_2 are called **isomorphisms**. It is immediate that isomorphisms of elliptic curves E_1 and E_2 induce isomorphisms between the groups $E_1(K)$ and $E_2(K)$, though note that the converse is not generally true. Since $\deg \mathrm{id}_E = 1$, for any isomorphism ϕ we must also have $\deg \phi = 1$.

2.5 Isogeny kernels

Since $P = (x_0, y_0)$ and $-P = (x_0, -y_0)$ generate the same subgroup for any P, they are either both in $\ker \phi$ or both not in $\ker \phi$. Further, since P and $-P$ are the only points of E with x-coordinate x_0, all subgroups of E (including $\ker \phi$) are uniquely determined by the x-coordinates of the points comprising that subgroup. For an isogeny ϕ, the **kernel polynomial** of ϕ is defined to be the smallest degree polynomial $h(x)$ such that $\phi(P) = \mathcal{O}$ if and only if $h(x_0) = 0$. When ϕ is written in standard form, the points in the kernel of ϕ are exactly those whose x-coordinate is a root of $v(x)$, so we always have that $h(x)$ divides $v(x)$. The **degree** of ϕ, denoted $\deg \phi$, is defined to be the maximum of the degrees of u and v. It can be shown that the composition of two isogenies $\phi_1 : E_1 \to E_2$ and $\phi_2 : E_2 \to E_3$ has degree

$$\deg(\phi_2 \circ \phi_1) = (\deg \phi_1) \cdot (\deg \phi_2);$$

see [Sut19, Corollary 6.10]. An isogeny of degree ℓ is often referred to as an **ℓ-isogeny**.

There is a relationship between the degree of ϕ and the size of its kernel. To explore this fully in the context of fields of characteristic $p \neq 0$, we make the following two definitions. We say that ϕ is **inseparable** if $(u/v)'$ is identically zero, and otherwise we say that ϕ is

separable. Inseparable isogenies can only occur over fields of positive characteristic, and it turns out that ϕ is inseparable exactly when $u(x) = f(x^p)$ and $v(x) = g(x^p)$ for some $f, g \in K[x]$, as shown in [Sut19, Lemma 6.1]. By [Sut19, Corollary 6.4], we can factor any isogeny ϕ as

$$\phi = \phi_{\text{sep}} \circ \pi^r$$

for some r, where ϕ_{sep} is separable and π is the **p-power Frobenius isogeny** taking (x, y) to (x^p, y^p). Finally, the following can be shown.

Theorem 2.3. *Let ϕ be an isogeny, and write ϕ as $\phi_{sep} \circ \pi^r$. Then $|\ker \phi| = \deg \phi_{sep}$.*

Proof. [Sut19, Theorem 6.8] □

A useful corollary to this theorem is that if $\deg \phi$ and p are coprime, the order of the kernel is simply the degree of ϕ, as π has degree p.

From the above discussion, we observe that the kernel of any isogeny must be a finite subgroup of $E(K)$. The following theorem provides a converse result to this, and will be useful in Section 8.3.

Theorem 2.4 (Vélu)**.** *Let E be an elliptic curve defined over K, and let G be any finite subgroup of $E(K)$. Then there exists an elliptic curve E' defined over some finite extension of K and a unique isogeny $\phi : E \to E'$ such that the kernel of ϕ is exactly G.*
Further, the isogeny ϕ is given explicitly by

$$\phi(x_P, y_P) = \left(x_p + \sum_{Q \in G - \{\mathcal{O}\}} (x_{P+Q} - x_Q), y_p + \sum_{Q \in G - \{\mathcal{O}\}} (y_{P+Q} - y_Q) \right).$$

We will use faster methods from [Bos+08] for computing the kernel polynomial of ϕ.

2.6 The j-invariant

Many properties of elliptic curves are invariant under isomorphism, though perhaps the most fundamental of these is the **j-invariant**. If $E : y^2 = x^3 + Ax + B$ is an elliptic curve in short Weierstrass form defined over K, the j-invariant is defined by

$$j(E) = \frac{-1728 \cdot (4A)^3}{\Delta},$$

where Δ is the discriminant of E defined earlier. The j-invariant fully characterizes isomorphism classes of elliptic curves in the following sense in [Sil86, Proposition 1.4].

- two elliptic curves E_1 and E_2 are isomorphic over $\overline{K}$ if and only if $j(E_1) = j(E_2)$, and
- for every $j_0 \in K$, there exists an elliptic curve $E(j_0)$ defined over K such that $j(E) = j_0$.

The j-invariant of elliptic curves is derived from a modular form of weight 0 for $\mathrm{SL}_2(\mathbb{Z})$ of the same name, the **j-function** $j : \mathcal{H} \to \mathbb{C}$. It turns out that for all $\ell \in \mathbb{Z}_{>0}$, the function $j(\ell\tau)$ lies in the extension field of $\mathbb{C}$ by j. The minimal polynomial of $j(\ell\tau)$ in $\mathbb{C}(j(\tau))$ is

called the **modular polynomial of level** ℓ and is usually denoted Φ_ℓ. These modular polynomials live in $\mathbb{Z}[X, Y]$ and parameterize pairs of isomorphism classes of elliptic curves together with ℓ-isogenies between them, in the sense that $\Phi_\ell(j_0, j_1) = 0$ if and only if there exists an ℓ-isogeny $\phi : E_0 \to E_1$ where $j(E_0) = j_0$ and $j(E_1) = j_1$. The modular polynomials Φ_ℓ are monic, symmetric in X and Y, and have degree $\ell + 1$ in both X and Y. for proofs of these statements, see [Lan87, Chapter 5]. The coefficients of Φ_ℓ grow very quickly, with the largest coefficient requiring $O(\ell \log \ell)$ bits to store, as shown in [Coh84].

2.7 ℓ-Isogeny graphs

We can use modular polynomials to define the **ℓ-isogeny graph** G_ℓ for a field K, a directed multigraph with vertices being elements of $\overline{K}$ and with the number of edges from j_0 to j_1 being the degree of the zero of $\Phi_\ell(j_0, j_1)$. Since Φ_ℓ is of degree $\ell + 1$ in both X and Y, the ℓ-isogeny graph is $(\ell + 1)$-regular; the vertices adjacent to j_0 are the $\ell + 1$ roots of $\Phi_\ell(j_0, Y)$, which polynomial is called the **ℓ^{th} modular polynomial instantiated at** j_0. Interpreting the vertices as isomorphism classes of elliptic curves having the specified j-invariant, paths in G_ℓ of length n correspond to sequences $(\phi_1, \ldots, \phi_n)$ of ℓ-isogenies such that $\phi_i : E_i \to E_{i+1}$. Such a sequence is referred to as a **chain**, and can be composed to obtain an ℓ^n-isogeny $\phi : E_1 \to E_{n+1}$.

To conclude our discussion of isogenies, we recall that for ℓ coprime to p, the ℓ-torsion of E is isomorphic to $\mathbb{Z}/\ell\mathbb{Z} \times \mathbb{Z}/\ell\mathbb{Z}$. If ℓ is a prime, then $E[\ell]$ has exactly $\ell + 1$ nontrivial subgroups $K_1, \ldots, K_{\ell+1}$, each of order ℓ. Using Vélu's formula, we can construct an elliptic curve E_i and a separable isogeny ϕ_i with $\phi_i : E \to E_i$ and $\ker \phi_i = K_i$ for $i = 1, \ldots, \ell + 1$. Since $\deg \phi_i = |\ker \phi_i| = \ell$, the ϕ_i are exactly the ℓ-isogenies to the neighboring vertices of $j(E)$ in G_ℓ.

3 Endomorphisms

As mentioned before, the set $\mathrm{Hom}(E, E')$ of isogenies from E to E' can be given a group structure under pointwise addition. If we restrict our attention to the case where $E = E'$, additional structure emerges: the group $\mathrm{Hom}(E, E)$ of isogenies from E to itself can be given a ring structure as well, with the product of two isogenies defined to be their composition. Elements of $\mathrm{Hom}(E, E)$ are called **endomorphisms**, and the ring of endomorphisms of E is denoted $\mathrm{End}(E)$.

We immediately see that $\mathrm{End}(E)$ is a ring with identity (this being the identity map on E), which we denote $[1]$.

3.1 Division polynomials

Defining $[n]$ to be

$$[n] \coloneqq \overbrace{[1] + \cdots + [1]}^{n \text{ times}},$$

we obtain the **multiplication by** n endomorphisms, which act on points $P \in E$ as $[n]P = nP$. Thus, the kernel of $[n]$ is the n-torsion $E[n]$. The kernel polynomial ψ_n of $[n]$ is called

the **nth division polynomial**, and can be computed using the recursive formulas

$$\psi_{2n} = \frac{1}{2y}\psi_n\left(\psi_{n+2}\psi_{n-1}^2 - \psi_{n-2}\psi_{n+1}^2\right)$$
$$\psi_{2n+1} = \psi_{n+2}\psi_n^3 - \psi_{n-1}\psi_{n+1}^3$$

as well as initial conditions for $\psi_0, \ldots, \psi_4 \in K[x, y]$ depending on E and presented in detail in [Sut19, Lecture 6]. Further, if we define

$$\phi_n = x\psi_n^2 - \psi_{n+1}\psi_{n-1}, \tag{2}$$
$$\omega_n = \frac{1}{4y}\left(\psi_{n+1}\psi_{n-1}^2 - \psi_{n-2}\psi_{n+1}^2\right), \tag{3}$$

then one can show that the standard form for $[n]$ is given by

$$[n] = \left(\frac{\phi_n(x)}{\psi_n^2(x)}, \frac{\omega_n(x,y)}{\psi_n^3(x,y)}\right)$$

where we possibly need to use the curve equation $y^2 = x^3 + Ax + B$ to eliminate some factors of y^2. Using these formulas, one can inductively prove the following theorem.

Theorem 3.1. *Let E be an elliptic curve defined over a field of characteristic p. Then for n coprime to p, the map $[n]$ is separable and has degree n^2.*

Proof. [Sut19, Theorem 6.25]. □

When n is odd, ψ_n lies in $K[x]$ and has degree $(n^2-1)/2$. Recalling that for any isogeny ϕ, and point $P = (x, y)$ is in $\ker(\phi)$ exactly when x is a root of the kernel polynomial of ϕ, we have that ψ_n completely characterizes the x-coordinates of points in the kernel of the multiplication-by-n map; that is, the points lying in the n-torsion $E[n]$.

3.2 The Frobenius endomorphism

Another important endomorphism that arises when E is defined over the finite field $\mathbb{F}_q$ is the **Frobenius endomorphism** π_E, which takes a point $P = (x, y)$ to $\pi_E P = (x^q, y^q)$. This is, in fact, an endomorphism; if (x, y) satisfy the relation $y^2 = x^3 + Ax + B$, then since A and B are in $\mathbb{F}_q$ they are fixed by the automorphism of $\overline{F_q}$ taking k to k^q, and we have that

$$y^q = (x^q)^3 + Ax^q + B.$$

The Frobenius endomorphism is noteworthy because its fixed points are exactly the $\mathbb{F}_q$-rational points of E, a fact which is the basis of modern point-counting algorithms.

3.3 The trace of an endomorphism

To define the trace of an endomorphism, we first define the **Tate module**, an algebraic object that encodes information about the ℓ^n-torsion of E for all n simultaneously for some prime $\ell \neq p$. Consider the commutative diagram given by

$$\cdots \to E[\ell^3] \to E[\ell^2] \to E[\ell] \to 0,$$

where $E[\ell^n]$ is viewed as a $\mathbb{Z}/\ell^n\mathbb{Z}$-module and the homomorphisms $E[\ell^n] \to E[\ell^{n-1}]$ are the multiplication-by-ℓ maps. The Tate module is defined to be the inverse limit of this sequence; that is, it is the unique group $T_\ell(E)$ (up to isomorphism) such that

- there exists homomorphisms $\{\psi_n : T_\ell(E) \to E[\ell^n]\}$ such that $\psi_n \circ [\ell^m] = \psi_{n-m}$, and
- for any other group G with homomorphisms $\{\varphi_n\}$ satisfying the first property, there exists a homomorphism $\pi : G \to T_\ell(E)$ such that $\varphi_n = \psi_n \circ \pi$ for all n.

Since each of the $E[\ell^n]$ is a $\mathbb{Z}/\ell^n\mathbb{Z}$-module, $T_\ell(E)$ is a module over the ℓ-adic integers $\mathbb{Z}_\ell$. As a $\mathbb{Z}_\ell$-module, the Tate module has the structure $T_\ell(E) \cong \mathbb{Z}_\ell \times \mathbb{Z}_\ell$, as shown in [Sil86, Section III.7]. We then have the following theorem.

Theorem 3.2. *Let ϕ be an endomorphism of an elliptic curve E defined over K, and let ℓ be a prime such that $\ell \neq \operatorname{char} K$. Then*

1. *ϕ induces a unique $\mathbb{Z}_\ell$-linear transformation of $T_\ell(E)$, represented by a matrix $A_\ell \in M_2(\mathbb{Z}_\ell)$,*
2. *the determinant of A_ℓ is exactly the degree of ϕ, and*
3. *for any primes $\ell_1, \ell_2 \neq \operatorname{char} K$, the matrix traces $\operatorname{tr} A_{\ell_1}$ and $\operatorname{tr} A_{\ell_2}$ are equal and lie in $\mathbb{Z} \subset \mathbb{Z}_\ell$.*

Proof. [Sil86, Proposition 2.3] □

The quantity $\operatorname{tr} A_\ell$ is called the **trace** of the endomorphism ϕ, and is denoted $\operatorname{tr} \phi$. Via a classical result from linear algebra, one then obtains the **characteristic polynomial** of an endomorphism, given by

$$\phi^2 - (\operatorname{tr} \phi)\phi + \deg \phi = 0. \tag{4}$$

This formula is central to our computations of the trace. Further, one can use Equation (4) and some additional information about the ring structure of $\operatorname{End}(E)$ to deduce the following.

Lemma 3.3. *Let $\phi \in \operatorname{End}(E)$ such that $\phi \neq [n]$ for $n \in \mathbb{Z}$. Then*

$$|\operatorname{tr}(\phi)| \leq 2\sqrt{\deg(\phi)}.$$

Proof. [Ban+19, Appendix] □

The original motivation for computing the trace of an endomorphism comes from Hasse's theorem and its close relationship to point-counting on elliptic curves; see Section 7 for more on this.

3.4 Endomorphisms as cycles in G_ℓ

As a final comment on endomorphisms generally, we provide an additional mechanism for their construction. Let ℓ be a prime (typically small), and consider the ℓ-isogeny graph G_ℓ. If we have a cycle of length n in G_ℓ beginning at j_0, this gives rise as in Section 2 to an isogeny $\phi : E_1 \to E_{n+1}$. Moreover, since $j(E_1) = j(E_{n+1})$, there exists an isomorphism $\psi : E_{n+1} \to E_1$, and thus $\psi \circ \phi : E_1 \to E_1$ is an endomorphism of degree ℓ^n. As mentioned in the Introduction, this is the only computationally feasible way to represent endomorphisms of elliptic curves defined over fields of large characteristic, and endomorphisms of this form are the primary objects of interest in this thesis.

4 Rational maps on subgroups of elliptic curves

4.1 Definitions

Let $E : y^2 = f(x)$ where $f(x) = x^3 + Ax + B$ be an elliptic curve defined over a field K, and fix an algebraic closure $\overline{K}$ of K. As discussed in Section 2.4, the behavior of isogenies $\phi : E \to E'$ on the affine points of E can be characterized by rational maps of the form

$$\phi(x, y) = (\phi_x(x, y), \phi_y(x, y)),$$

where ϕ_x and ϕ_y are elements of the fraction field of the quotient ring $K[E] = K[x, y]/(y^2 - f(x))$.

Suppose we have some finite subgroup $H < E$ with $|H|$ odd. Since $(x, y) \in H$ exactly when $(x, -y) \in H$, there exists some polynomial $h \in \overline{K}[x]$ such that $h(x_0) = 0$ if and only if there exist points (x_0, y_0) and $(x_0, -y_0)$ in H. Note that since $|H|$ is odd, H contains no points of order 2, so (x_0, y_0) and $(x_0, -y_0)$ are distinct. This polynomial h is known as the **kernel polynomial** of H, an abuse of notation justified by h being the kernel polynomial of the isogeny $\phi : E \to E/H$ given by Vélu's formula.

Since H is characterized by its kernel polynomial h, a function on the x-coordinates of points in H is determined only up to equivalence modulo $h(x)$. This is made formal by defining the **ring of regular functions** on H to be

$$\overline{K}[H] \coloneqq \overline{K}[E]/(h(x)) = \overline{K}[x, y]/(h(x), y^2 - f(x)).$$

We may then define the **field of rational functions** on H, given by $\overline{K}(H) \coloneqq \operatorname{Frac} \overline{K}[H]$.

Let E' be another elliptic curve, and consider $\operatorname{Hom}(E, E')$. We extend the canonical surjection $\overline{K}[E] \to \overline{K}[H]$ to a function

$$\operatorname{proj}_H : \operatorname{Hom}(E, E') \to \overline{K}(H)^2,$$

and call the image of this map $\operatorname{Hom}(H, E')$. Further, if H is a subgroup of E invariant under $\operatorname{End}(E)$, then we define $\operatorname{End}(H)$ to be $\operatorname{Hom}(H, E)$. If $\phi \in \operatorname{Hom}(E, E')$, we denote $\operatorname{proj}_H(\phi)$ by ϕ_H, and in the special case that $H = E[\ell]$ we denote the image by $\phi_{(\ell)}$.

A useful example of such a fixed subgroup is the ℓ-torsion $E[\ell]$ for some prime ℓ. As a non-example, the $\ell + 1$ subgroups of $E[\ell]$ are not necessarily fixed by any given isogeny, as they may be permuted.

4.2 Arithmetic in $\mathrm{Hom}(H, E')$

We wish now to describe algorithms for computing basic operations in $\mathrm{Hom}(H, E')$ similar to those given in [Sut19, Section 9.3]. First, however, we prove that elements of $\mathrm{Hom}(H, E')$ can be given in a standard form.

Lemma 4.1. *Let E and E' be elliptic curves over K, and let H be a subgroup of E with kernel polynomial $h \in \overline{K}[x]$. Then every element of $\mathrm{Hom}(H, E')$ can be written uniquely in the form*

$$\left(\frac{a(x)}{b(x)}, \frac{c(x)}{d(x)}y\right)$$

where $\gcd(a, b) = \gcd(c, d) = 1$, and h shares no roots with either b or d.

Proof. Let $\phi \in \mathrm{Hom}(E, E')$. By Lemma 2.2, we can write ϕ uniquely as

$$\phi(x, y) = \left(\frac{u(x)}{v(x)}, \frac{s(x)}{t(x)}y\right)$$

where $\gcd(u, v) = \gcd(s, t) = 1$. We now write $v(x) = v_1(x)v_2(x)$, where $v_1(x) = \gcd(v, h)$ and $v_2 = v/v_1$. Since v_2 shares no factors with h, there exist polynomials f and g such that

$$fh + gv_2 = 1.$$

Letting $\tilde{\phi} = (\frac{u}{v}, \frac{s}{t}y)$ denote the image of ϕ in $\mathrm{Hom}(H, E')$, we then have that

$$\tilde{\phi}_x = \frac{u}{v} = \frac{ug}{v_1v_2g} = \frac{ug}{v_1(1 - fh)} = \frac{ug}{v_1}.$$

Repeating this for ϕ_y gives the desired form. □

One useful consequence of this lemma is that if v and h share no roots (that is, if the kernel of ϕ has trivial intersection with H) then elements of $\mathrm{Hom}(H, E')$ can be specified by maps of the form $(a(x), b(x)y)$ where $a, b \in \overline{K}[x]/(h(x))$. In Section 8, we utilize this fact by working with subgroups that have no intersection with the kernel of the given isogeny.

We now present versions of the algorithms in [Sut19, Section 9.3], modified so as to work for any subgroup $H < E$, not just $E[\ell]$. We assume in both cases that the kernels of the input isogenies have trivial intersection with H, as this is the most relevant case for our purposes.

Algorithm 4.2. Let E be an elliptic curve defined over K, and let H be a finite subgroup of E with kernel polynomial h. Let H_1 be another finite subgroup of E satisfying $H < H_1 < E$ such that H_1 is fixed by $\mathrm{End}(E)$. Given $\phi \in \mathrm{Hom}(H, E)$ and $\psi \in \mathrm{End}(H_1)$, compute the composition $\psi \circ \phi$ in $\mathrm{Hom}(H, E)$ as follows.

1. Write $\phi(x, y) = (a(x), b(x)y)$ and $\psi(x, y) = (c(x), d(x)y)$.
2. Compute $a'(x) = c(a(x))$ and $b'(x) = d(a(x)) \cdot b(x)$ modulo $h(x)$.
3. Output $(a'(x), b'(x)y)$.

Note that if $H = H_1 = E[\ell]$ for some ℓ, Algorithm 4.2 gives multiplication in $\mathrm{End}(E[\ell])$. We now give an algorithm for computing the sum of two elements of $\mathrm{Hom}(H, E)$. We assume for this algorithm that the kernel of the sum also has trivial intersection with H. This is not guaranteed by any conditions on the input, but we will have methods for recovering from such an error whenever we have need of this algorithm in the sequel.

Algorithm 4.3. Let $E : y^2 = f(x)$ with $f(x) = x^3 + Ax + B$ be an elliptic curve and let H be a finite subgroup of E with kernel polynomial h. Let ϕ and ψ be elements of $\mathrm{Hom}(H, E)$. We compute the sum $\phi + \psi$ in $\mathrm{Hom}(H, E)$ as follows.

1. If $\phi = \psi$:
   (a) Write $\phi(x, y) = (a(x), b(x)y) = \psi(x, y)$.
   (b) Let $k(x) = 2b(x)f(x)$. If k is not relatively prime to h, **abort**; otherwise continue.
   (c) Let $m(x) = (3a(x)^2 + A) \cdot k^{-1}(x)$ modulo $h(x)$.
2. Otherwise:
   (a) Write $\phi(x, y) = (a(x), b(x)y)$ and $\psi(x, y) = (c(x), d(x)y)$.
   (b) Let $k(x) = c(x) - a(x)$. If k is not relatively prime to h, **abort**; otherwise continue.
   (c) Let $m(x) = (d(x) - b(x)) \cdot k^{-1}(x)$ modulo $h(x)$.
3. Compute $a'(x) = f(x)m(x)^2 - 2a(x)$ and $b'(x) = m(a(x) - a'(x)) - b(x)$ modulo $h(x)$.
4. Output $(a'(x), b'(x)y)$.

Finally, we present a simple algorithm for composing an element of $\mathrm{Hom}(H, E')$ with an isogeny.

Algorithm 4.4. Let E, E_1, and E_2 be elliptic curves defined over K, and let H be a finite subgroup of E with kernel polynomial h. Given $\phi \in \mathrm{Hom}(H, E_1)$ and $\psi \in \mathrm{Hom}(E_1, E_2)$, compute the image of $\psi \circ \phi$ in $\mathrm{Hom}(H, E_2)$ as follows.

1. Write $\phi(x, y) = (a(x), b(x)y)$ and
$$\psi = \left(\frac{u(x)}{v(x)}, \frac{s(x)}{t(x)}y\right).$$
2. Compute the inverses of $v(a(x))$ and $t(a(x))$ modulo $h(x)$
3. Compute the products $a'(x) = u(a(x)) \cdot v(a(x))^{-1}$ and $b'(x) = s(a(x)) \cdot t(a(x))^{-1} \cdot b(x)$ modulo $h(x)$
4. Output $(a'(x), b'(x)y)$

5 Complexity of basic operations

In analyzing the runtime of the algorithms presented in this thesis, we will use the time complexity of several more basic operations, which we present here. All complexities given in this thesis are given in terms of bit operations. In the results that follow, we use $\mathsf{M}(n)$ to indicate the time to multiply two n-bit integers, and write $\log\log n$ as $\operatorname{llog} n$. We assume that M is superlinear, and by [SS71] can take M to be in $O(n \log n \operatorname{llog} n)$ when it simplifies the complexity expressions.

5.1 Finite field operations

In this section we present some basic results on operations over $\mathbb{F}_q$ and $\mathbb{F}_q[x]$. We assume in all cases that $q = p^e$ satisfies $\log e = O(\log p)$.

First, we give the time complexities of simple operations in $\mathbb{F}_q$.

Lemma 5.1. *Let $\alpha, \beta \in \mathbb{F}_q$. Then*

1. *the sum $\alpha + \beta$ and the difference $\alpha - \beta$ can be computed in time $O(\log q)$,*
2. *the product $\alpha\beta$ can be computed in time $O(\mathsf{M}(\log q))$, and*
3. *the multiplicative inverse α^{-1} can be computed in $O(\mathsf{M}(\log q) \operatorname{llog} q)$.*

Proof. These results are proved in [Sut19, Theorem 3.17], [Sut19, Theorem 4.3], and [Sut19, Theorem 4.8] respectively. □

We next give some basic polynomial operations in $\mathbb{F}_q[x]$.

Lemma 5.2. *Let $f, g \in \mathbb{F}_q[x]$ be polynomials of degree no more than d. Then*

1. *the sum $f + g$ can be computed in $O(d \log q)$, and*
2. *the product fg can be computed in $O(\mathsf{M}(d \log q))$.*

Proof. To compute the sum of two polynomials, all we need to do is add two elements of $\mathbb{F}_q$ for each of the d coefficients, giving a total complexity of $O(d \log q)$.

See [Hv19a, Theorem 1.1] for a proof of the second claim for $q = p$. This can then be combined with Kronecker substitution to get the result for general q, per [Hv19b, Section 1.1]. □

As we will spend much time working with modular polynomial arithmetic, we present the relevant results for those operations here as well.

Lemma 5.3. *Let $f, g, h \in \mathbb{F}_q[x]$ such that f and h have degree no more than d and g has degree no more than d'. Then*

1. *the residue of f modulo h can be computed in $O(\mathsf{M}(d \log q))$,*
2. *f^{-1} modulo h can be computed in $O(\mathsf{M}(d \log q) \log d + d\ \mathsf{M}(\log q) \operatorname{llog} q)$, and*
3. *the composition $g(f)$ can be computed modulo h in $O(d'\ \mathsf{M}(d \log q))$.*

Proof. The first claim follows by using [Sut19, Algorithm 4.2] with inputs in $\mathbb{F}_q[x]$.

The second claim is proven by combining [GG13, Corollary 11.9] with the complexities of operations in $\mathbb{F}_q$ and $\mathbb{F}_q[x]$ discussed above.

To compute $g(f)$ modulo h, we use the following algorithm.

Algorithm 5.4.

1. Let $g = a_{d'} x^{d'} + \cdots + a_0$.
2. Compute f^i modulo h for $i = 1, \ldots, d'$.
3. Output $a_{d'} f^{d'} + \cdots + a_0$

In the second step, if we compute the remainder modulo h after every multiplication, then we can compute all powers of f with d' multiplications of degree d polynomials and d' reductions modulo h. The multiplications dominate, and this step takes $O(d'\ \mathsf{M}(d \log q))$. Step 3 then takes another d' multiplications, this time by the coefficients, and $O(d')$ additions. These are both subsumed in the computations for the previous step, giving the claimed total complexity. □

Finally, we give an algorithm for applying the Chinese remainder theorem to a list of polynomials and analyze its complexity. This will be utilized in Section 8.3.

Algorithm 5.5. Given polynomials $r_1, \ldots, r_n$ and $m_1, \ldots, m_n$ in $\mathbb{F}_q[x]$ of degree at most d, compute the unique degree nd polynomial f satisfying $f \equiv r_i \pmod{m_i}$ for $i = 1, \ldots, n$ as follows.

1. Initialize $m \leftarrow m_1$ and $f \leftarrow r_1$.
2. For $i = 2, \ldots, n$:
   (a) Compute $y = (r_i - c)/m$ modulo m_i.
   (b) Set $c \leftarrow c + m \cdot y$, then set $m \leftarrow m \cdot m_i$.
3. Output $c \pmod{m}$.

Proposition 5.6. *Let $r_1, \ldots, r_n$ and $m_1, \ldots, m_n$ be polynomials in $\mathbb{F}_q[x]$ of degree at most d such that the m_i are pairwise relatively prime. Assuming that $\operatorname{llog} p = O(n)$ and $d = O(n)$, Algorithm 5.5 computes the unique degree nd polynomial f satisfying $f \equiv r_i \pmod{m_i}$ for $i = 1, \ldots, n$ in time $O(n\ \mathsf{M}(nd \log q))$.*

Proof. The dominant steps in this algorithm are the polynomial multiplications and inversions. As all of the inversions are performed modulo a degree d polynomial and the multiplications are performed on polynomials of degree nd, by the assumptions on the relative sizes of the inputs the multiplications are asymptotically dominant. Since n of these multiplications must be performed, we get a total complexity of $O(n\ \mathsf{M}(nd \log q))$, as desired. □

5.2 Complexity of operations in $\mathrm{Hom}(H, E)$

In this section we analyze the complexity of the operations involving $\mathrm{Hom}(H, E)$ defined in Section 4. As Schoof's point-counting algorithm and our main algorithm for computing the trace of an endomorphism are both based upon determining the action of an endomorphism on torsion subgroups of E, we will make good use of them.

Lemma 5.7. *Let E be an elliptic curve defined over $\mathbb{F}_q$ and let H and H_1 be finite subgroups of E with kernel polynomials h and h_1 of degrees d and d_1 respectively such that $H < H_1 < E$ and H_1 is fixed by $\mathrm{End}(E)$. Then given $\phi \in \mathrm{Hom}(H, E)$ and $\psi \in \mathrm{End}(H_1)$, Algorithm 4.2 computes $\psi \circ \phi$ in time $O(d_1\ \mathsf{M}(d \log q))$.*

Proof. The dominant steps in this algorithm are the modular compositions of the degree d polynomials defining ϕ into the degree d_1 polynomials defining ψ. The four of these give a total complexity of $O(d_1\ \mathsf{M}(d \log q))$, as desired. □

As a corollary to this result, we can determine the time complexity of multiplications in $\mathrm{End}(E[\ell])$ for any integer ℓ.

Corollary 5.8. *Let E be an elliptic curve defined over $\mathbb{F}_q$ and let $\ell \in \mathbb{Z}$ be odd. For given endomorphisms $\phi, \psi \in \mathrm{End}(E[\ell])$, Algorithm 4.2 computes the product $\psi\phi \coloneqq \psi \circ \phi$ in time $O(\ell^2\ \mathsf{M}(\ell^2 \log q))$*

Proof. Since the kernel polynomial of $E[\ell]$ is the ℓ^{th} division polynomial, which for ℓ odd has degree $(\ell^2 - 1)/2$, we get $d = d_1 = O(\ell^2)$ in the previous lemma. □

We similarly determine the time complexity of additions in $\mathrm{Hom}(H, E)$ and obtain as a corollary the time to add two elements of $\mathrm{End}(E[\ell])$.

Lemma 5.9. *Let E be an elliptic curve defined over $\mathbb{F}_q$ and let $H < E$ be a finite subgroup with kernel polynomial h of degree d. Assume that $\mathrm{llog}\, q = O(\log d)$. Given $\phi, \psi \in \mathrm{Hom}(H, E)$, Algorithm 4.3 computes $\phi + \psi$ in $O(\log d\ \mathsf{M}(d \log q))$.*

Proof. Regardless of whether or not $\phi = \psi$, Algorithm 4.3 terminates after a number of polynomial multiplications, inversions, and reductions modulo h that is independent of the input size. The most expensive of these is polynomial inversion, which gives the stated complexity. □

Corollary 5.10. *Let E be an elliptic curve defined over $\mathbb{F}_q$ and let $\ell \in \mathbb{Z}$. Assume that $\mathrm{llog}\, q = O(\log \ell)$. For given endomorphisms $\phi, \psi \in \mathrm{End}(E[\ell])$, Algorithm 4.3 computes the sum $\psi + \phi$ in time $O(\log \ell\ \mathsf{M}(\ell^2 \log q))$.*

Proof. As in the proof of Corollary 5.8, we get that the ℓ^{th} division polynomial has degree $O(\ell^2)$, and substitute that into the complexity result from 5.9. □

Additionally, we can use this result to obtain the complexity of computing scalar multiples of maps using a double-and-add approach.

Corollary 5.11. *Let E be an elliptic curve defined over $\mathbb{F}_q$ and let $H < E$ be a finite subgroup with kernel polynomial h of degree d. Assume that $\mathrm{llog}\, q = O(\log \ell)$. Given $\phi \in \mathrm{Hom}(H, E)$ and $k \in \mathbb{Z}$, we can compute $k\phi$ in time $O(\log k \log d\ \mathsf{M}(d \log q))$.*

Proof. Working with the binary representation of k, we can obtain $k\phi$ with a sequence of at most $\log_2 k$ doublings or additions of ϕ. As each doubling/addition takes the time given in Lemma 5.9, we get the desired complexity. □

Finally, we can analyze the complexity of Algorithm 4.4 for computing the action of a composition of isogenies on a subgroup. This will be useful to us when handling input endomorphisms given as chains of small-degree isogenies when we use it to "walk" step by step along the chain instead of computing the whole endomorphism.

Lemma 5.12. *Let E, E_1, and E_2 be elliptic curves defined over $\mathbb{F}_q$, and let H be a finite subgroup of E with kernel polynomial h of degree d. Assume that $\mathrm{llog}\, q = O(\log d)$. Given $\phi \in \mathrm{Hom}(H, E_1)$ and $\psi \in \mathrm{Hom}(E_1, E_2)$ such that ψ is of degree d_1, Algorithm 4.4 computes the image of $\psi \circ \phi$ in $\mathrm{Hom}(H, E_2)$ in $O((d_1 + \log d)\ \mathsf{M}(d \log q))$.*

Proof. The computations in Algorithm 4.4 involve a constant number of compositions with a degree d_1 polynomial modulo h and a constant number of polynomial inversions modulo h. These give a complexity of

$$O((d_1 + \log d)\ \mathsf{M}(d \log q) + d \operatorname{llog} q\ \mathsf{M}(\log q)),$$

which we reduce to the desired complexity using the assumptions that $\operatorname{llog} q = O(\log d)$ and that M is superlinear. □

6 Supersingular elliptic curves

Let E be an elliptic curve defined over an algebraic closure K of some finite field $\mathbb{F}_p$. Considering $\mathrm{End}(E)$ as a $\mathbb{Z}$-module, it can be shown that $\mathrm{End}(E)$ has $\mathbb{Z}$-rank either 2 or 4, with ring structure either an order in an imaginary quadratic field or an order in a quaternion algebra respectively [Sil86, Corollary 9.4]. Elliptic curves with a rank 2 endomorphism ring are called **ordinary**, and those with rank 4 are called **supersingular**. Supersingular elliptic curves enjoy a number of other remarkable properties which make their study a distinct task from the ordinary case.

Firstly, all isomorphism classes of supersingular elliptic curves defined over K have j-invariant in the finite subfield $\mathbb{F}_{p^2}$ (see [Sil86, Theorem V.3.1] for a proof; Silverman attributes this result originally to [Deu41]). As all elliptic curves can be defined over the field containing their j-invariant, supersingular elliptic curves are always isomorphic to a curve defined over $\mathbb{F}_{p^2}$. The j-invariants corresponding to these isomorphism classes are known as **supersingular j-invariants**, and one can see immediately that there are only finitely many. This can be made even more precise; for a given p, there are exactly $\lfloor \frac{p}{12} \rfloor + t$ supersingular j-invariants over $\mathbb{F}_{p^2}$, where $t \in \{0, 1, 2\}$ and depends only on the residue of p modulo 12 [Deu41, Section 10].

Secondly, for every pair of supersingular elliptic curves E_1 and E_2 and every $\ell > 1$, there exists an isogeny from E_1 to E_2 of degree ℓ^n for some n [Koh96, Corollary 78]. As this isogeny can then be factored into a chain of n isogenies of degree ℓ, all supersingular j-invariants are connected in the ℓ-isogeny graph. Further, isogenies preserve the rank of the endomorphism ring, so the connected component of $G_\ell(p)$ containing the supersingular j-invariants contains *only* the supersingular j-invariants.

7 Schoof's algorithm

In this section we discuss Schoof's algorithm for determining the number of rational points on an elliptic curve defined over $\mathbb{F}_q$, originally presented in [Sch85, Section 3]. This algorithm was the first deterministic point-counting algorithm with running time polynomial in $\log q$, and though several improvements have been made (most notably those of Elkies [Elk98] and Atkin [Atk91]), Schoof's algorithm remains the foundation of the current best methods for this task. We will only be discussing the basic algorithm here; a richer discussion as well as details on the aforementioned improvements can be found in [Sch95]. A summary of Schoof's algorithm (as given in [Sut19]) is as follows.

Algorithm 7.1. Given an elliptic curve E defined over $\mathbb{F}_q$ where $q = p^r$, compute the trace of the Frobenius endomorphism of E as follows.

1. Set $M \leftarrow 1$.
2. While $M \leq 4\sqrt{q}$, for primes $\ell = 2, 3, 4, \ldots$ such that $\ell \neq p$:
   (a) Compute t_ℓ, the trace of Frobenius modulo ℓ.
   (b) Set $M \leftarrow M \cdot \ell$.
3. Use the Chinese Remainder Theorem to determine the unique integer $t \in [-2\sqrt{q}, 2\sqrt{q}]$ such that $t \equiv t_\ell \pmod{\ell}$ for all ℓ used in Step 2.
4. Output t.

The main idea powering Schoof's algorithm is that if one

(a) has an easily computable bound on the trace and

(b) can quickly compute the trace modulo many small primes,

one can use the Chinese Remainder Theorem to reconstruct the true value of the trace. Both the bound on the trace and the connection of Algorithm 7.1 to point-counting come from Hasse's Theorem.

Theorem 7.2 (Hasse)**.** *Let E be an elliptic curve defined over $\mathbb{F}_q$ and let $t = \operatorname{tr} \pi_E$ be the trace of the Frobenius endomorphism of E. Then*

1. *the number of $\mathbb{F}_q$-rational points of E is given by $|E(\mathbb{F}_q)| = q - t + 1$, and*
2. *$|t| \leq 2\sqrt{q}$.*

Proof. [Sut19, Theorem 8.3] □

The remaining piece of Schoof's algorithm that we must discuss is the polynomial time method for computing the trace of Frobenius modulo some small prime ℓ. Since any endomorphism $\phi \in \operatorname{End}(E)$ is a group homomorphism by definition, it takes $E[\ell]$ into itself, so the characteristic equation (4) holds when restricted to the ℓ-torsion; that is,

$$\phi_{(\ell)}^2 - [\operatorname{tr} \phi]_{(\ell)} \phi_{(\ell)} + [\deg \phi]_{(\ell)}. \tag{5}$$

Recalling that for all $n \in \mathbb{Z}$, we have that $[n]_{(\ell)} = n_\ell [1]_{(\ell)}$ where $n_\ell \equiv n \pmod{\ell}$, we obtain the desired efficient algorithm for computing t_ℓ. The following presentation is again from [Sut19].

Algorithm 7.3. Given an elliptic curve E defined over $\mathbb{F}_q$ where $q = p^r$ and a prime ℓ not equal to p, compute the trace of Frobenius of E modulo ℓ as follows.

1. Compute ψ_ℓ, the ℓ^{th} division polynomial of E.
2. Let π be the Frobenius enodmorphism of E. Compute the images of π, π^2, and $[q]$ in $\mathrm{End}(E[\ell])$.
3. Compute $\pi^2_{(\ell)} + [q]_{(\ell)}$
4. Compute $c\pi_{(\ell)}$ for $c = 0, \ldots, \ell - 1$ until the relation

$$\pi^2_{(\ell)} + [q]_{(\ell)} = c\pi_{(\ell)}$$

is satisfied, then output that c.

We are guaranteed that at least one value of c will satisfy the equation in Step 4, that value being t_ℓ itself. Further, the following lemma guarantees that this is in fact the *only* value of c that will satisfy Equation (5).

Lemma 7.4. *[Sut19] Let ℓ be a prime, and let ϕ be an endomorphism of E that is not the zero map on $E[\ell]$. If P be a non-zero point of E satisfying*

$$\phi^2_{(\ell)}(P) - c\phi_{(\ell)}(P) + [\deg \phi]_{(\ell)}(P)$$

for some integer c, then $c \equiv t_\ell = \operatorname{tr} \phi \pmod{\ell}$.

8 A modified Schoof's algorithm

In this section we give an explicit algorithm for computing the trace of an endomorphism α of a supersingular elliptic curve given as a chain of small-degree isogenies. Such an algorithm was originally described in [Koh96, Theorem 81], and a more detailed outline was given in [Ban+19, Appendix]. We flesh out the algorithm more fully and perform a tighter analysis of the asymptotic complexity. The main theorem that we wish to prove is as follows.

Theorem 8.1. *Let E be a supersingular elliptic curve defined over $\mathbb{F}_q$, and let $\alpha = \alpha_r \circ \cdots \circ \alpha_1$ be an endomorphism of E, where $\alpha_1, \ldots, \alpha_r$ are isogenies of degree at most d defined over $\mathbb{F}_q$. If $r = O(\log p)$ and $d = O(\log p)$ as well, then $\operatorname{tr} \alpha$ can be computed in time $\tilde{O}(\log^6 p)$. If instead $d = O(\operatorname{llog} q)$, then the computation takes $\tilde{O}(\log^5 p)$.*

This result improves upon the estimated complexity given in [Ban+19, Theorem 6.11] due partially to a tighter complexity analysis and partially to the use of a slightly different method for computing $\alpha^2_{(\ell)}$ that increases the dependence on d in exchange for lessening the dependence on r.

8.1 The algorithm

Our algorithm proceeds similarly to Algorithm 7.1 for computing the trace of Frobenius. In place of Hasse's Theorem, we instead have Lemma 3.3, which gives us the bound

$$|\operatorname{tr}\alpha| \leq 2\sqrt{\deg\alpha},$$

where the degree D of α can be determined by multiplying the degrees of the α_i. In Step 2, we iterate over small primes not dividing D. By the prime number theorem, the density of the primes is such that

$$\sum_{\text{primes } \ell \leq x} \ln \ell \approx x$$

for all x, so we get that

$$\prod_{\text{primes } \ell \leq \ln 4\sqrt{D}} \ell \approx 4\sqrt{D}.$$

More generally, if α is given as a chain of r isogenies of degree at most d, we have that $D \leq d^r$, and so the maximum size of ℓ that we need consider is $O(\ln 4\sqrt{D}) = O(r \log d)$.

Fixing a prime ℓ, we compute $t_\ell \coloneqq \operatorname{tr}\alpha \pmod{\ell}$ via a process similar to Algorithm 7.3.

Algorithm 8.2. Given an elliptic curve $E : y^2 = f(x)$ defined over $\mathbb{F}_q$, a degree D endomorphism $\alpha \in \operatorname{End}(E)$ given as a chain of isogenies $\alpha = \alpha_r \circ \cdots \circ \alpha_1$ defined over $\mathbb{F}_q$, and a small prime ℓ, compute the trace of α modulo ℓ as follows.

1. Compute ψ_ℓ, the ℓ^{th} division polynomial of E.
2. Compute $\alpha_{(\ell)}$, the image of α in $\operatorname{End}(E[\ell])$, by repeatedly applying Algorithm 4.4 to evaluate
$$\alpha_r \circ \cdots \circ \alpha_1 \circ \mathrm{id}_{E[\ell]},$$
where $\mathrm{id}_{E[\ell]}$ is the image of (x, y) in $\operatorname{End}(E[\ell])$.
3. Compute $\alpha_{(\ell)}^2$ similarly by again using Algorithm 4.4 to evaluate
$$\alpha_r \circ \cdots \circ \alpha_1 \circ \alpha_{(\ell)}.$$
4. Compute $d_{(\ell)} \coloneqq [\deg\alpha]_{(\ell)} = (\deg\alpha)\,\mathrm{id}_{E[\ell]}$ using a double-and-add procedure.
5. Compute $\alpha_{(\ell)}^2 + d_{(\ell)}$.
6. Compute $c\alpha_{(\ell)}$ for $c = 0, \ldots, \ell - 1$ until the relation
$$\alpha_{(\ell)}^2 + d_{(\ell)} = c\alpha_{(\ell)}$$
is satisfied. Then output that c.

Equation (5) still holds for α, and Lemma 7.4 extends naturally to α provided there is only trivial intersection between $\ker \alpha$ and $E[\ell]$, a condition guaranteed by only choosing ℓ coprime to the degree.

When computing $\alpha^2_{(\ell)}$, there are two natural approaches. In addition to the approach described in the above algorithm in which one simply "goes around α" twice, one can also compute $\alpha^2_{(\ell)}$ directly from $\alpha_{(\ell)}$ using Algorithm 4.2 for multiplying two elements of $\mathrm{End}(E[\ell])$. To compare the two, we analyze their complexities as follows.

Proposition 8.3. *Let E be an elliptic curve defined over $\mathbb{F}_q$, and let*

$$\alpha = \alpha_r \circ \cdots \circ \alpha_1$$

be an endomorphism of E given as a chain of r isogenies defined over $\mathbb{F}_q$ of degree no more than d. Let ℓ be an odd prime satisfying $\mathrm{llog}\, q = O(\log \ell)$. Given the ℓ^{th} division polynomial ψ_ℓ and the image $\alpha_{(\ell)}$ of α in $\mathrm{End}(E[\ell])$, we can

1. *compute $\alpha^2_{(\ell)}$ in*
$$O(r(d + \log \ell)\ \mathsf{M}(\ell^2 \log q))$$
using repeated applications of Algorithm 4.4, as described in Step 3 of Algorithm 8.2, and

2. *compute $\alpha^2_{(\ell)}$ in*
$$O(\ell^2\ \mathsf{M}(\ell^2 \log q))$$
using Algorithm 4.2 to directly multiply $\alpha_{(\ell)} \circ \alpha_{(\ell)}$ in $\mathrm{End}(E[\ell])$.

Proof. The complexity of a single application of Algorithm 4.4 follows from taking $d = \ell^2$ and $d_1 = d$ in Lemma 5.12, and is $O((d + \log \ell)\ \mathsf{M}(\ell^2 \log q))$. We then repeat this for each of the r elements of the chain to get the first claim.

The second claim is directly Corollary 5.8. □

To compare these complexities, we recall that we will need to compute $\alpha^2_{(\ell)}$ for $\ell = O(r \log d)$. Substituting this in to each of the complexities derived in the previous proposition gives a time complexity of

$$O\big(r(d + \log r))\ \mathsf{M}(r^2 \log^2 d \log q)\big)$$

for repeated composition, and for direct multiplication in $\mathrm{End}(E[\ell])$ we get a complexity of

$$O\big(r^2 \log^2 d\ \mathsf{M}(r^2 \log^2 d \log q)\big).$$

If $r = \Theta(d)$, then these two approaches differ only up to a log factor. However, if r is significantly larger than d (this is the situation in which one has a long chain of small degree isogenies), then multiplication in $\mathrm{End}(E[\ell])$ is asymptotically slower than repeated compositions. Conversely, if d is significantly larger than r, the opposite occurs.

To end this section, we present the full modified Schoof's algorithm for the computation of the trace of α.

Algorithm 8.4. Given an elliptic curve $E : y^2 = f(x)$ defined over $\mathbb{F}_q$, an endomorphism $\alpha \in \mathrm{End}(E)$ given by $\alpha = \alpha_r \circ \cdots \circ \alpha_1$ where each of the α_i is an isogeny of degree at most d defined over $\mathbb{F}_q$, compute the trace of α as follows.

1. Set $M \leftarrow 1$.
2. Compute $D = \prod_{i=1}^{r} \deg \alpha_i$, the degree of α.
3. While $M \leq 4\sqrt{D}$, for primes $\ell = 2, 3, 4, \ldots$ such that $\ell \nmid D$:
   (a) Compute t_ℓ, the trace of α modulo ℓ, via Algorithm 8.2.
   (b) Set $M \leftarrow M \cdot \ell$.
4. Use the Chinese Remainder Theorem to determine the unique integer $t \in [-2\sqrt{D}, 2\sqrt{D}]$ such that $t \equiv t_\ell \pmod{\ell}$ for all ℓ used in Step 2.
5. Output t.

8.2 Complexity analysis

In this section we analyze the complexity of Algorithm 8.4. We begin by determining the complexity of computing $\mathrm{tr}\,\alpha$ modulo ℓ for odd primes ℓ.

Lemma 8.5. *Let E be an elliptic curve defined over $\mathbb{F}_q$, and let $\alpha \in \mathrm{End}(E)$ be given as a chain of r isogenies of degree no more than d. Given an odd prime ℓ such that $\log q = O(\log \ell)$, Algorithm 8.2 computes t_ℓ in time.*

$$O\big((rd + r\log\ell + \ell\log\ell)\ \mathsf{M}(\ell^2 \log q)\big).$$

Proof. Using [Dol18, Algorithm 1], we can compute the ℓ^{th} division polynomial in time

$$O\big(\log\ell\ \mathsf{M}(\ell^2 \log q)\big).$$

Next, since the processes for computing $\alpha_{(\ell)}$ and $\alpha^2_{(\ell)}$ are nearly identical, by Proposition 8.3 we can compute both in time

$$O\big(r(d + \log\ell)\ \mathsf{M}(\ell^2 \log q)\big).$$

The computations of $d_{(\ell)} \coloneqq [D]_{(\ell)}$ and $\alpha^2_{(\ell)} + d_{(\ell)}$ are done by single scalar multiplications and additions in $\mathrm{End}(E[\ell])$, and so are dominated by the repeated additions to be considered in Step 4. When computing $d_{(\ell)} = D\,\mathrm{id}_{E[\ell]}$, we note that because $\ell\,\mathrm{id}_{E[\ell]}$ is the zero map we can simply multiply $\mathrm{id}_{E[\ell]}$ by the residue of D modulo ℓ.

If it happens that we are unable to perform the addition in Step 3 due to a nontrivial intersection of the kernel of $\alpha_{(\ell)}^2 + d_{(\ell)}$ with the ℓ-torsion of E, we may need to restart from Step 2 working modulo the kernel polynomial of this intersection. The recomputations are done with a smaller polynomial, so are dominated by the first round of computations, and we need only do this a maximum of one time. Thus, there is no effect on the asymptotic complexity if we need to restart the calculation in this way.

Finally, the computations of $c\alpha_{(\ell)}$ in Step 4 are carried out by repeatedly adding $\alpha_{(\ell)}$ to itself, and checking each addend. Thus, we will need to compute $O(\ell)$ additions on average, for a time complexity of

$$O\big(\ell \log \ell \ \mathsf{M}(\ell^2 \log q)\big).$$

This gives a total complexity of

$$O\big((rd + r\log \ell + \ell \log \ell) \ \mathsf{M}(\ell^2 \log q)\big),$$

as claimed. □

Note that the three addends in this complexity correspond roughly to scenarios where each of the inputs d, r, and ℓ are large, and that the size of q influences all of them.

This then allows us to determine the asymptotic complexity of the full modified Schoof's algorithm.

Theorem 8.6. *Let E be an elliptic curve defined over $\mathbb{F}_q$, and let $\alpha \in \mathrm{End}(E[\ell])$ be given as a chain of r isogenies of degree no more than d. Assuming that $\operatorname{llog} q = O(\log r)$, Algorithm 8.4 computes the trace of α in time*

$$O\left(\frac{r\log d}{\log r + \operatorname{llog} d}(rd + r\log r \log d) \ \mathsf{M}(r^2 \log^2 d \log q)\right).$$

Proof. The computations in Algorithm 8.4 are dominated by the computation of the trace of α modulo ℓ in Step 3a. As discussed at the beginning of Section 8, the primes ℓ will be $O(r \log d)$ in size. Thus, the computation of each t_ℓ will be performed in time

$$O\big((rd + r\log r \log d) \ \mathsf{M}(r^2 \log^2 d \log q)\big).$$

Further, the number of primes that we need to check is given by the prime number theorem as

$$O\left(\frac{r\log d}{\log r + \operatorname{llog} d}\right),$$

and multiplying these two complexities gives the desired result. □

This allows us to prove more concise asymptotics on our algorithm for some common assumptions on the relative sizes of the inputs. Recall that we take $M(n) = O(n \log n \operatorname{llog} n)$ in determining the big-$\tilde{O}$ complexities.

Corollary 8.7. *Suppose r and d are both $O(\log q)$. Then Algorithm 8.4 runs in*

$$O(\log^3 q \ \mathsf{M}(\log^3 q \operatorname{llog}^2 q)) = \tilde{O}(\log^6 q).$$

If instead $d = O(\operatorname{llog} q)$, then the runtime is

$$O(\log^2 q \operatorname{lllog}^3 q \ \mathsf{M}(\log^3 q \operatorname{llog}^2 q)) = \tilde{O}(\log^5 q).$$

8.3 Application to supersingular elliptic curves

We now fix E to be a supersingular elliptic curve defined over $\mathbb{F}_{p^2}$. It seems that one ought to be able to improve upon Algorithm 8.4 by exploiting the fact that all supersingular elliptic curves over $\overline{\mathbb{F}_p}$ are in fact defined over the field $\mathbb{F}_{p^2}$. We attempted to do this by quickly determining the factors of the ℓ^{th} division polynomial of E, computing α and α^2 restricted to the subgroups of $E[\ell]$ corresponding to these factors, then using the Chinese Remainder Theorem to reconstruct $\alpha_{(\ell)}$ and $\alpha^2_{(\ell)}$. This did not result in a significant speed-up, either asymptotic or experimental, but we present the algorithm below both as a way of cataloguing our efforts and also to illustrate the type of potential leverage points working with supersingular elliptic curves provides as compared to the ordinary case.

Fix a prime ℓ, and let $\Phi_\ell \in \mathbb{Z}[X, Y]$ be the modular polynomial of level ℓ. For an ordinary elliptic curve E' defined over $\mathbb{F}_q$, one expects $\Phi_\ell(j(E'), Y) \in \mathbb{F}_q[Y]$ to split over a degree $\ell+1$ extension of $\mathbb{F}_q$. As a consequence, one generically expects that the neighboring j-invariants to $j'(E)$ in the ℓ-isogeny graph, elliptic curves having those neighboring j-invariants, and ℓ-isogenies originating from E' are all defined over this same extension. While the ℓ^{th} division polynomial does factor into the kernel polynomials of these isogenies, working modulo these factors would incur a significant penalty in terms of bit operations, as now all field operations must happen over a much larger field.

However, for a supersingular elliptic curve E, this is no longer the case. Since all supersingular j-invariants are defined over $\mathbb{F}_{p^2}$ and all isogenies out of E are to supersingular curves, we get that the $\ell + 1$ isogenies of degree ℓ with E as their domain have kernel polynomials defined over $\mathbb{F}_{p^2}$ itself, and not an extension. Thus, field operations working modulo these kernel polynomials are no more expensive, and we can hope to achieve a speed-up.

Our approach to computing $\alpha_{(\ell)}$ and $\alpha^2_{(\ell)}$ using this factors-based approach is outlined in the following algorithm.

Algorithm 8.8. Let E be a supersingular elliptic defined over $\mathbb{F}_{p^2}$ with $j(E) \notin \{0, 1728\}$ and let α be an endomorphism of E given by $\alpha = \alpha_r \circ \cdots \circ \alpha_1$, where each of the α_i is defined over $\mathbb{F}_{p^2}$. Given an odd prime ℓ not dividing $\deg \alpha$, compute the images $\alpha_{(\ell)}$ and $\alpha^2_{(\ell)}$ of α and α^2 respectively in $\mathrm{End}(E[\ell])$ as follows.

1. Compute Φ_ℓ, the modular polynomial of level ℓ. Factor $\Phi_\ell(j(E), Y)$ completely over $\mathbb{F}_{p^2}$, and let $j_1, \ldots, j_{\ell+1}$ denote its roots. We assume that all roots obtained in this way are distinct.

2. For $i = 1, \ldots, \ell + 1$, compute the kernel polynomial of an isogeny from E to E_i, where E_i is a supersingular elliptic curve such that $j(E_i) = j_i$. This gives a list of degree $(\ell - 1)/2$ polynomials $\phi_1, \ldots, \phi_{\ell+1}$ whose product is the ℓ^{th} division polynomial of E. Let $K_1, \ldots, K_{\ell+1}$ denote the kernels of these isogenies.

3. Compute the images $\alpha_{K_i} = (a_i(x), b_i(x)y)$ and $\alpha^2_{K_i} = (c_i(x), d_i(x)y)$ of α and α^2 in $\mathrm{Hom}(K_i, E)$ for each i using repeated applications of Algorithm 4.4.

4. Use the Chinese remainder theorem to reconstruct $a \in \mathbb{F}_{p^2}[x]$ satisfying $a \equiv a_i\phi_i$ for $i = 1, \ldots, \ell + 1$. Similarly, reconstruct b, c, and d.

5. Output $\alpha_{(\ell)} = (a(x), b(x)y)$ and $\alpha^2_{(\ell)} = (c(x), d(x)y)$.

Steps 1 and 4 of Algorithm 8.8 are addressed in Section 5. The only step requiring special consideration is Step 2, for which we use an adapted version of [Gal12, Algorithm 28] (Galbraith cites [Elk98] for the original idea behind the algorithm and credits Sutherland for the implementation) together with work on the fast computation of isogenies in [Bos+08].

We now present the algorithm for Step 2.

Algorithm 8.9. Let $E : y^2 = x^3 + Ax + B$ be an elliptic curve defined over $\mathbb{F}_{p^2}$ such that $j = j(E) \notin \{0, 1728\}$, let ℓ be a positive integer, and let $\tilde{j}$ be a simple roots of $\Phi_\ell(j, \tilde{j})$. Compute the kernel polynomial ϕ_i of an ℓ-isogeny from E to some curve $\tilde{E}$ such that $j(\tilde{E}) = \tilde{j}$ as follows.

1. Compute the values f_x, f_y, f_{xy}, f_{xx}, and f_{yy}, where $f_x = \frac{\partial \Phi_\ell}{\partial x}(j.\tilde{j})$, $f_{xy} = \frac{\partial^2 \Phi_\ell}{\partial x \partial y}(j.\tilde{j})$, etc.
2. Let $m = 18B/A$, let $j' = mj$, and let $\tilde{j}' = -j' f_x/(\ell f_y)$.
3. Let $\tilde{m} = \tilde{j}'/\tilde{j}$ and let $\tilde{k} = \tilde{j}'/1728 - \tilde{j}$.
4. Compute $\tilde{A} = \ell^4 \tilde{m} \tilde{k}/48$ and compute $\tilde{B} = \ell^6 \tilde{m}^2 \tilde{k}/864$. These coefficients are such that $\tilde{E} : y^2 = x^3 + \tilde{A}x + \tilde{B}$ has j-invariant $\tilde{j}$.
5. Let $r = -((j')^2 f_{xx} + 2\ell j' \tilde{j}' f_{xy} + \ell^2 \tilde{j}' f_{yy})/(j' f_x)$.
6. Compute $\sigma = \ell \left(\frac{r}{2} + \frac{k - \ell\tilde{k}}{4} + \frac{\ell\tilde{m} - m}{3} \right)$. This is the first power sum of the kernel polynomial ϕ of the ℓ-isogeny from E to $\tilde{E}$.
7. Use [Bos+08, `fastElkies`] to compute ϕ, then output.

The assumption in Step 1 of Algorithm 8.8 is not overly restrictive for two reasons. First of all, we expect that for sufficiently large p a randomly chosen elliptic curve E will have no double edges in the ℓ-isogeny graph for any $\ell = O(r \log d)$, provided r and d are at most polylogarithmic in p. This is due to [LB20, Proposition B.3], which states that the number of elliptic curves with a non-integer endomorphism of degree M is $O(M^{3/2})$. If E has a double edge in the ℓ-isogeny graph for some ℓ, then the endomorphism defined by going out along one edge and back along the other is not an integer and has degree ℓ^2. Thus, the number of curves with double edges in the ℓ-isogeny graph for some $\ell = O(r \log d)$ is $O(r^3 \log^3 d)$, which is significantly less than the $O(p)$ supersingular elliptic curves under the above assumptions on r and d.

Secondly, even if a double edge is encountered, the algorithm is largely unaffected. Suppose that in Step 2 there were m non-distinct roots of $\Phi_\ell(j(E), Y)$ for some $m > 1$. To recover, one must insert a step between Steps 2 and 3 to find the missing factors of the ℓ^{th} division polynomial, ψ_ℓ. This can be accomplished by dividing the product of the ϕ_i successfully recovered out of ψ_ℓ, and setting the remainder to ϕ_{rem}. This polynomial will be the product of m kernel polynomials of degree ℓ isogenies, so will have degree $m(\ell - 1)/2$. Letting m range between 1 and $\ell+1$, we see that Algorithm 8.8 degenerates to the method for computing α_ℓ and α_ℓ^2 given in Algorithm 8.2. As we will see shortly, this will have virtually no effect on the asymptotic complexity.

We now restrict to the case where E is a supersingular elliptic curve and analyze the time complexities of Algorithms 8.8 and 8.9.

Proposition 8.10. *Let E be a supersingular elliptic curve defined over $\mathbb{F}_{p^2}$, let ℓ be an odd prime, and let $\tilde{j}$ be adjacent to $j(E)$ in the ℓ-isogeny graph. Assume $\operatorname{llog} p = O(\log \ell)$. Given the above inputs and Φ_ℓ, the modular polynomial of level ℓ, Algorithm 8.9 computes the kernel polynomial of an isogeny from E to a curve with j-invariant $\tilde{j}$ in time*

$$O\big(\ell^2\ \mathsf{M}(\log p)\big).$$

Proof. As Φ_ℓ is of degree $\ell + 1$ in both variables, it has $O(\ell^2)$ non-zero terms. Thus, its derivatives may be computed with $O(\ell^2)$ multiplications in $\mathbb{F}_{p^2}$, for a total complexity of $O(\ell^2\ \mathsf{M}(\log p))$. We may then evaluate these partial derivatives at $(j(E), \tilde{j})$ in $O(\ell^2\ \mathsf{M}(\log p))$ time as well.

Steps 2-6 are accomplished in a constant number of field operations in $\mathbb{F}_{p^2}$, so they are dominated by the surrounding steps.

For Step 7, [Bos+08, Theorem 2] gives that the kernel polynomial can be computed in $O(\mathsf{M}(\ell))$ field operations. A closer analysis reveals that $O(\ell)$ of these are divisions, and that some of the multiplications may be sped up with Kronecker substitution. The full complexity of this algorithm is given by $O\big(\mathsf{M}(\ell \log p) + \ell\ \mathsf{M}(\log p) \operatorname{llog} p\big)$.

Thus under the assumption that $\operatorname{llog} p = O(\log \ell)$, all steps are dominated by the cost of computing and evaluating the derivatives, $O(\ell^2\ \mathsf{M}(\log p))$. □

We may now determine the time required to compute $\alpha_{(\ell)}$ and $\alpha_{(\ell)}^2$ by first computing its restriction to the subgroups of $E[\ell]$, then using the Chinese remainder theorem to get its action on the whole ℓ-torsion.

Theorem 8.11. *Let E be a supersingular elliptic curve defined over $\mathbb{F}_{p^2}$ and let α be an endomorphism of E given as a chain of r isogenies of degree no more than d. Let ℓ be an odd prime with $\operatorname{llog} p = O(\log \ell)$ and $\log \ell = O(\log p)$. Assuming that E has no double neighbors in the ℓ-isogeny graph, Algorithm 8.8 computes $\alpha_{(\ell)}$ and $\alpha_{(\ell)}^2$ in time*

$$O\big(\ell(rd + r \log \ell)\ \mathsf{M}(\ell \log p) + \ell\ \mathsf{M}(\ell^2 \log p)\big).$$

Proof. To compute the ℓ^{th} modular polynomial Φ_ℓ, for the purposes of the complexity analysis we use the algorithm given in [BLS12, Algorithm 6.1], which runs in $O(\ell^3 \log \ell \operatorname{llog} \ell)$. We then instantiate it at $j(E)$ in $O(\ell^2\ \mathsf{M}(\log p))$. Since E is supersingular, all of its neighbors in the ℓ-isogeny graph have j-invariants in $\mathbb{F}_{p^2}$. Thus, we may factor $\Phi_\ell(j(E), Y)$ completely over $\mathbb{F}_{p^2}$, and this takes

$$O\big(\ell \log p\ \mathsf{M}(\ell \log p))\big).$$

We then use ℓ runs of Algorithm 8.9 to compute the $\ell + 1$ degree $(\ell - 1)/2$ factors of the ℓ^{th} division polynomial; this takes

$$O\big(\ell^3\ \mathsf{M}(\log p)\big).$$

We now use repeated applications of Algorithm 4.4 to compute the images of α and α^2 in $\operatorname{Hom}(K_i, E)$ for each proper subgroup $K_i < E[\ell]$; this takes

$$O\big(r\ell(d + \log \ell)\ \mathsf{M}(\ell \log p)\big)$$

in total. Finally, as shown in Lemma 5.6, we can reconstruct the images of α and α^2 in $\mathrm{End}(E[\ell])$ in

$$O\big(\ell\ \mathsf{M}(\ell^2 \log p)\big).$$

We add all of these complexities together to get that Algorithm 8.8 runs in

$$O\big(\ell(rd + r\log \ell)\ \mathsf{M}(\ell \log p) + \ell\ \mathsf{M}(\ell^2 \log p)\big),$$

as desired. □

Remark 8.12. Sutherland has pre-computed and made publicly available all modular polynomials of prime level up to up to 1000 using the methods listed in [BLS12], and in practice this is more than enough for current applications.

We can compare with the results obtained in Proposition 8.3 by again substituting in $\ell = O(r \log d)$, giving a runtime of

$$O\big(r^2 d \log d\ \mathsf{M}(r \log d \log p) + r \log d\ \mathsf{M}(r^2 \log^2 d \log p)\big).$$

From this we see that in the case that $d = O(r)$, this provides only a slight asymptotic speed-up, obtained from "moving" one factor of $r \log d$ to the outside of M. In the case that $d = O(\log r)$, there is no asymptotic benefit. Further, in practice using this approach tends to run more slowly than Algorithm 8.4, so using this algorithm as described is not recommended.

9 Conclusion

We have thus provided an explicit polynomial-time algorithm for computing of the trace of an endomorphism given by a chain of small-degree isogenies, following the work of Kohel as well as that of Banks, Camacho-Navarro, Eisenträger, Morrison, and Park. We performed a careful analysis of the asymptotic complexity of this algorithm, sharpening the bound provided in [Ban+19, Theorem 6.11] in the generic case and showing an asymptotic improvement of our algorithm in the low-degree case. Finally, we illustrated an attempt at improving the algorithm for the special case of supersingular elliptic curves. While this gave little to no asymptotic improvement, it shows the types of improvements one may be able to make in future work.

10 References

[Atk91] A.O.L. Atkin. *Unpublished manuscript*. 1991.

[Aza+20] Reza Azarderakhsh, Matthew Campagna, Craig Costello, Luca De Feo, Basil Hess, Amir Jalali, David Jao, Brian Koziel, Brian LaMacchia, Patrick Longa, Michael Naehrig, Joost Renes, Vladimir Soukharev, and David Urbanik. *Supersingular Isogeny Key Encapsulation*. 2020. URL: `https://csrc.nist.gov/CSRC/media/Projects/post-quantum-cryptography/documents/round-3/submissions/SIKE-Round3.zip`.

[Ban+19] Efrat Bank, Catalina Camacho-Navarro, Kirsten Eisenträger, Travis Morrison, and Jennifer Park. "Cycles in the supersingular ℓ-isogeny graphs and corresponding endomorphisms". In: *Research Directions in Number Theory*. Ed. by Jennifer S. Balakrishnan, Amanda Folsom, Matilde Lalín, and Michelle Manes. Vol. 19. Association for Women in Mathematics Series. Springer, Cham, 2019, pp. 41–66.

[BLS12] Reinier Bröker, Kristin Lauter, and Andrew Sutherland. "Modular polynomials via isogeny volcanoes". In: *Mathematics of Computation* 81.278 (Apr. 2012), pp. 1201–1231.

[Bos+08] Alin Bostan, François Morain, Bruno Salvy, and Éric Schost. "Fast algorithms for computing isogenies between elliptic curves". In: *Mathematics of Computation* 77.263 (July 2008), pp. 1755–1778.

[BS11] Gaetan Bisson and Andrew Sutherland. "Computing the endomorphism ring of an ordinary elliptic curve over a finite field". In: *Journal of Number Theory* 113 (2011), pp. 815–831.

[Cas91] Ian Cassels. *Lectures on elliptic curves*. Cambridge University Press, 1991.

[Coh84] Paula Cohen. "On the coefficients of the transformation polynomials for the elliptic modular function". In: *Mathematical Proceedings of the Cambridge Philosophical Society* 95 (3 1984), pp. 389–402.

[Deu41] Max Deuring. "Die Typen der Multiplikatorenringe elliptischer Funktionenkörper". In: *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 14 (Dec. 1941), pp. 197–272.

[Dol18] Javad Doliskani. "On division polynomial PIT and supersingularity". In: *Applicable Algebra in Engineering, Communication and Computing* 29 (Jan. 2018), pp. 393–407.

[Elk98] Noam Elkies. "Elliptic and modular curves over finite fields and related computational issues". In: *Computational Perspectives on Number Theory: Proceedings of a Conference in Honor of A. O. L. Atkin*. Ed. by Duncan Buell and Jeremy Teitelbaum. Vol. 7. Studies in Advanced Mathematics. American Mathematical Society and International Press of Boston, 1998, pp. 21–76.

[Gal12] Steven Galbraith. *Mathematics of Public Key Cryptography*. 2nd ed. Cambridge University Press, 2012.

[GG13] Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. 3rd ed. Cambridge University Press, 2013.

[GPS17] Steven Galbraith, Christophe Petit, and Javier Silva. "Identification protocols and signature schemes sased on supersingular isogeny problems". In: *Advances in Cryptology – ASIACRYPT 2017*. Vol. 10624. Lecture Notes in Computer Science. Springer, Cham, 2017, pp. 3–33.

[Hv19a] David Harvey and Joris van der Hoeven. "Faster polynomial multiplication over finite fields using cyclotomic coefficient rings". In: *Journal of Complexity* 54 (2019), p. 101404. ISSN: 0885-064X. DOI: https://doi.org/10.1016/j.jco.2019.03.004.

[Hv19b] David Harvey and Joris van der Hoeven. *Polynomial multiplication over finite fields in time $O(n \log n)$*. hal-02070816. 2019.

[Koh96] David Kohel. "Endomorphism rings of elliptic curves over finite fields". PhD thesis. University of California at Berkeley, 1996.

[Lan87] Serge Lang. *Elliptic Functions*. 2nd ed. Vol. 112. Graduate Texts in Mathematics. Springer-Verlag New York, 1987.

[LB20] Jonathan Love and Dan Boneh. "Supersingular curves with small non-integer endomorphisms". In: *Proceedings of the Fourteenth Algorithmic Number Theory Symposium*. Ed. by Steven D. Galbraith. Vol. 4. The Open Book Series 1. Mathematical Sciences Publishers, 2020, pp. 7–22.

[NIS17] NIST. *Post-quantum cryptography*. 2017. URL: https://csrc.nist.gov/Projects/Post-Quantum-Cryptography (visited on 04/19/2021).

[Sch85] René Schoof. "Elliptic curves over finite fields and the computation of square roots mod p". In: *Mathematics of Computation* 44.170 (Apr. 1985), pp. 483–494.

[Sch95] René Schoof. "Counting points on elliptic curves over finite fields". en. In: *Journal de Théorie des Nombres de Bordeaux* 7.1 (1995), pp. 219–254.

[Sho94] Peter Shor. "Algorithms for quantum computation: discrete logarithms and factoring". In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. IEEE, 1994, pp. 124–134.

[Sil86] Joseph Silverman. *The Arithmetic of Elliptic Curves*. 1st ed. Vol. 106. Graduate Texts in Mathematics. Springer-Verlag New York, 1986.

[SS71] Arnold Schönhage and Volker Strassen. "Schnelle Multiplikation großer Zahlen". In: *Computing* 7 (Sept. 1971), pp. 281–292.

[Sut19] Andrew Sutherland. *18.783 Elliptic Curves*. License: Creative Commons BY-NC-SA. 2019.

[Voi21] John Voight. *Quaternion Algebras*. 1st ed. Vol. 228. Graduate Texts in Mathematics. Springer International Publishing, 2021.