Isogeny Graphs in Cryptography

Luca De Feo
Université Paris Saclay – UVSQ
https://defeo.lu/

Graph Theory Meets Cryptography
July 29 – August 2, 2019, Wurzbürg, Germany

Introduction

These lectures notes were written for the summer school *Graph Theory Meets Cryptography* in Wurzbürg, Germany.

The presentation is divided in four parts, roughly corresponding to the four lectures given.

Contents

LaTeX source code available at https://github.com/defeo/wurzburg/.

Part I
Elliptic curves and isogenies

In this part, we review the basic and not-so-basic theory of elliptic curves. Our goal is to summarize the fundamental theorems necessary to understanding the foundations of isogeny based cryptography. A proper treatment of the material covered here would require more than one book, we thus skip proofs and lots of details to go straight to the useful theorems. The reader in search of a more comprehensive treatment will find more details [66, 67, 46, 55].

Throughout this part we let k be a field, and we denote by $\bar{k}$ its algebraic closure.

1 Elliptic curves

Elliptic curves are projective curves of genus 1 with a distinguished point. Projective space initially appeared through the process of adding *points at infinity*, as a method to understand the geometry of projections (also known as *perspective* in classical painting). In modern terms, we define projective space as the collection of all lines in affine space passing through the origin.

Definition 1 (Projective space)**.** The *projective space of dimension* n, denoted by $\mathbb{P}^n$ or $\mathbb{P}^n(\bar{k})$, is the set of all $(n+1)$-tuples

$$(x_0, \dots, x_n) \in \bar{k}^{n+1}$$

such that $(x_0, \dots, x_n) \neq (0, \dots, 0)$, taken modulo the equivalence relation

$$(x_0, \dots, x_n) \sim (y_0, \dots, y_n)$$

if and only if there exists $\lambda \in \bar{k}$ such that $x_i = \lambda_i y_i$ for all i.

The equivalence class of a projective point $(x_0, \dots, x_n)$ is customarily denoted by $(x_0 : \cdots : x_n)$. The set of the *k-rational points*, denoted by $\mathbb{P}^n(k)$, is defined as

$$\mathbb{P}^n(k) = \{(x_0 : \cdots : x_n) \in \mathbb{P}^n \mid x_i \in k \text{ for all } i\}.$$

By fixing arbitrarily the coordinate $x_n = 0$, we define a projective space of dimension $n-1$, which we call the *hyperplane at infinity*; its points are called *points at infinity*.

From now on we suppose that the field k has characteristic different from 2 and 3. This has the merit of greatly simplifying the representation of an elliptic curve. For a general definition, see [66, Chap. III].

Definition 2 (Weierstrass equation)**.** An *elliptic curve* defined over k is the locus in $\mathbb{P}^2(\bar{k})$ of an equation

$$Y^2Z = X^3 + aXZ^2 + bZ^3, \tag{1}$$

with $a, b \in k$ and $4a^3 + 27b^2 \neq 0$.

The point $(0 : 1 : 0)$ is the only point on the line $Z = 0$; it is called the *point at infinity* of the curve.

It is customary to write Eq. (1) in *affine form*. By defining the coordinates $x = X/Z$ and $y = Y/Z$, we equivalently define the elliptic curve as the locus of the equation

$$y^2 = x^3 + ax + b,$$

plus the point at infinity $\mathcal{O} = (0 : 1 : 0)$.

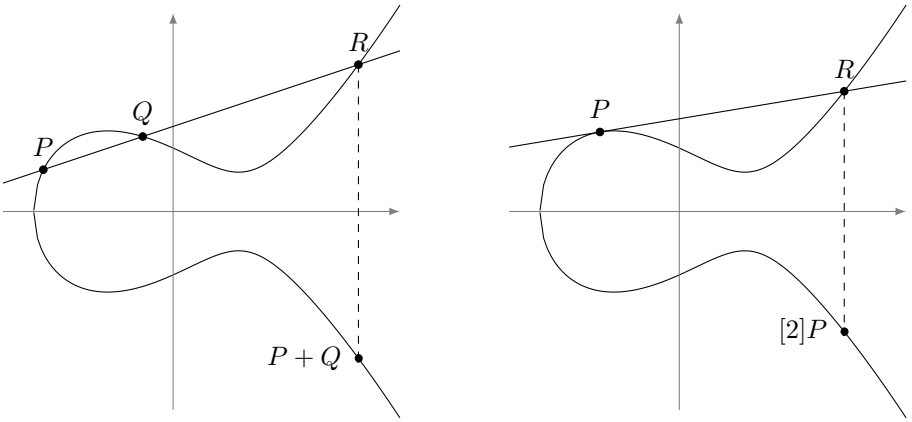

Figure 1: An elliptic curve defined over $\mathbb{R}$, and the geometric representation of its group law.

In characteristic different from 2 and 3, we can show that any projective curve of genus 1 with a distinguished point $\mathcal{O}$ is isomorphic to a Weierstrass equation by sending $\mathcal{O}$ onto the point at infinity $(0:1:0)$.

Now, since any elliptic curve is defined by a cubic equation, Bezout's theorem tells us that any line in $\mathbb{P}^2$ intersects the curve in exactly three points, taken with multiplicity. We define a group law by requiring that three co-linear points sum to zero.

Definition 3. Let $E : y^2 = x^3 + ax + b$ be an elliptic curve. Let $P_1 = (x_1, y_1)$ and $P_2 = (x_2, y_2)$ be two points on E different from the point at infinity, then we define a composition law $\oplus$ on E as follows:

- $P \oplus \mathcal{O} = \mathcal{O} \oplus P = P$ for any point $P \in E$;
- If $x_1 = x_2$ and $y_1 = -y_2$, then $P_1 \oplus P_2 = \mathcal{O}$;
- Otherwise set

$$\lambda = \begin{cases} \frac{y_2 - y_1}{x_2 - x_1} & \text{if } P \neq Q, \\ \frac{3x_1^2 + a}{2y_1} & \text{if } P = Q, \end{cases}$$

  then the point $(P_1 \oplus P_2) = (x_3, y_3)$ is defined by

$$\begin{aligned} x_3 &= \lambda^2 - x_1 - x_2, \\ y_3 &= -\lambda x_3 - y_1 + \lambda x_1. \end{aligned}$$

It can be shown that the above law defines an Abelian group, thus we will simply write $+$ for $\oplus$. The n-th scalar multiple of a point P will be denoted by $[n]P$. When E is defined over k, the subgroup of its *rational points over* k is customarily denoted $E(k)$. Figure 1 shows a graphical depiction of the group law on an elliptic curve defined over $\mathbb{R}$.

We now turn to the group structure of elliptic curves. The torsion part is easily characterized.

Proposition 4. *Let E be an elliptic curve defined over an algebraically closed field k, and let $m \neq 0$ be an integer. The m-torsion group of E, denoted by $E[m]$, has the following structure:*

- $E[m] \simeq (\mathbb{Z}/m\mathbb{Z})^2$ *if the characteristic of k does not divide m;*

- *If $p > 0$ is the characteristic of k, then*

$$E[p^i] \simeq \begin{cases} \mathbb{Z}/p^i\mathbb{Z} & \text{for any } i \geq 0, \text{ or} \\ \{\mathcal{O}\} & \text{for any } i \geq 0. \end{cases}$$

Proof. See [66, Coro. 6.4]. For the characteristic 0 case see also Section 3. □

For curves defined over a field of positive characteristic p, the case $E[p] \simeq \mathbb{Z}/p\mathbb{Z}$ is called *ordinary*, while the case $E[p] \simeq \{\mathcal{O}\}$ is called *supersingular*. We shall see an alternative characterization of supersingularity in the next section.

The free part of the group is much harder to characterize. We have some partial results for elliptic curves over number fields.

Theorem 5 (Mordell-Weil)**.** *Let k be a number field, the group $E(k)$ is finitely generated.*

However the exact determination of the rank of $E(k)$ is somewhat elusive: we have algorithms to compute the rank of most elliptic curves over number fields; however, an exact formula for such rank is the object of the *Birch and Swinnerton-Dyer conjecture*, one of the *Clay Millenium Prize Problems*.

2 Maps between elliptic curves

Finally, we focus on maps between elliptic curves. We are mostly interested in maps that preserve both facets of elliptic curves: as projective varieties, and as groups.

We first look into invertible algebraic maps, that is linear changes of coordinates that preserve the Weierstrass form of the equation. Because linear maps preserve lines, it is immediate that they also preserve the group law. It is easily verified that the only such maps take the form

$$(x, y) \mapsto (u^2x', u^3y')$$

for some $u \in \bar{k}$, thus defining an *isomorphism* between the curve $y^2 = x^3 + au^4x + bu^6$ and the curve $(y')^2 = (x')^3 + ax' + b$. Isomorphism classes are traditionally encoded by an invariant, whose origins can be traced back to complex analysis.

Proposition 6 (j-invariant)**.** *Let $E : y^2 = x^3 + ax + b$ be an elliptic curve, and define the* j-invariant *of E as*

$$j(E) = 1728\frac{4a^3}{4a^3 + 27b^2}.$$

Two curves are isomorphic over the algebraic closure $\bar{k}$ if and only if they have the same j-invariant.

Note that if two curves defined over k are isomorphic over $\bar{k}$, they are so over an extension of k of degree dividing 6. An isomorphism between two elliptic curves defined over k, that is itself not defined over k is called a *twist*. Any curve has a *quadratic twist*, unique up to isomorphism, obtained by taking $u \notin k$ such that $u^2 \in k$. The two curves of j-invariant 0 and 1728 also have *cubic*, *sextic* and *quartic twists*.

A surjective group morphism, not necessarily invertible, between two elliptic curves is called an *isogeny*. It turns out that isogenies are algebraic maps as well.

Theorem 7. *Let E, E' be two elliptic curves, and let $\phi : E \to E$ be a map between them. The following conditions are equivalent:*

1. *ϕ is a surjective group morphism,*
2. *ϕ is a group morphism with finite kernel,*
3. *ϕ is a non-constant algebraic map of projective varieties sending the point at infinity of E onto the point at infinity of E'.*

Proof. See [66, III, Th. 4.8]. □

Two curves are called *isogenous* if there exists an isogeny between them. We shall see later that this is an equivalence relation.

Isogenies from a curve to itself are called *endomorphisms*. The prototypical endomorphism is the multiplication-by-m endomorphism defined by

$$[m] \ : \ P \mapsto [m]P.$$

Its kernel is exactly the m-th torsion subgroup $E[m]$.

Since they are algebraic group morphisms, we can define addition of isogenies by $(\phi+\psi)(P) = \phi(P) + \psi(P)$, and the resulting map is still an isogeny. Thus, by including the constant map that sends every point to the point at infinity, the set of isogenies $E \to E'$ forms a group. Additionally, endomorphisms $E \to E$ support composition, distributing over addition, hence the set of all endomorphisms forms a ring, denoted by $\mathrm{End}(E)$.[1]

Since each $m \in \mathbb{Z}$ defines a different multiplication-by-m endomorphism, clearly $\mathbb{Z} \subset \mathrm{End}(E)$. But can $\mathrm{End}(E)$ be larger? We shall now give a complete characterization of the endomorphism ring for any elliptic curve.

Definition 8 (Order)**.** Let K be a finitely generated $\mathbb{Q}$-algebra. An *order* $\mathcal{O} \subset K$ is a subring of K that is a finitely generated $\mathbb{Z}$-module, and that contains a $\mathbb{Q}$-basis for K.

The prototypical example of order is the ring of integers $\mathcal{O}_K$ of a number field K. It turns out that $\mathcal{O}_K$ is the *maximal order* of K, i.e., it contains any other order of K. We shall discuss this case in depth in Section 6.

Definition 9 (Quaternion algebra)**.** A *quaternion algebra* is an algebra of the form

$$K = \mathbb{Q} + \alpha\mathbb{Q} + \beta\mathbb{Q} + \alpha\beta\mathbb{Q},$$

where the generators satisfy the relations

$$\alpha^2, \beta^2 \in \mathbb{Q}, \quad \alpha^2 < 0, \quad \beta^2 < 0, \quad \beta\alpha = -\alpha\beta.$$

Theorem 10 (Deuring)**.** *Let E be an elliptic curve defined over a field k of characteristic p. The ring* $\mathrm{End}(E)$ *is isomorphic to one of the following:*

- *$\mathbb{Z}$, only if $p = 0$;*
- *An order $\mathcal{O}$ in a quadratic imaginary field (a number field of the form $\mathbb{Q}(\sqrt{-D})$ for some $D > 0$); in this case we say that E has* complex multiplication *by $\mathcal{O}$;*
- *Only if $p > 0$, a maximal order in a quaternion algebra ramified at p and ∞; in this case we say that E is* supersingular*.*

Proof. See [66, III, Coro. 9.4] and [42]. □

In positive characteristic, a curve that is not supersingular is called *ordinary*; we shall see that it necessarily has complex multiplication.

[1] In short, isogenies are the morphisms in the Abelian category of elliptic curves.

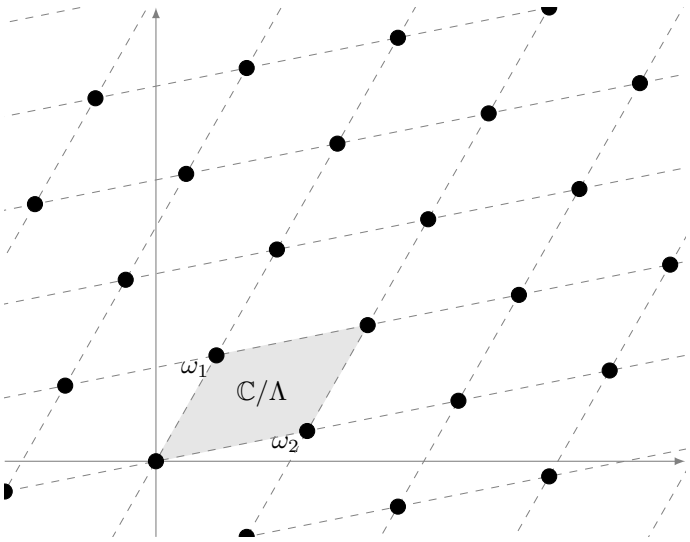

Figure 2: A complex lattice (black dots) and its associated complex torus (grayed *fundamental domain*).

3 Elliptic curves over $\mathbb{C}$

To better understand elliptic curves and their morphisms, we take a moment now to specialize them to the complex numbers.

Definition 11 (Complex lattice)**.** A *complex lattice* Λ is a discrete subgroup of $\mathbb{C}$ that contains an $\mathbb{R}$-basis of $\mathbb{C}$.

Explicitly, a complex lattice is generated by a *basis* (ω_1, ω_2), such that $\omega_1 \neq \lambda\omega_2$ for any $\lambda \in \mathbb{R}$, as

$$\Lambda = \omega_1\mathbb{Z} + \omega_2\mathbb{Z}.$$

Up to exchanging ω_1 and ω_2, we can assume that $\mathrm{Im}(\omega_1/\omega_2) > 0$; we then say that the basis has *positive orientation*. A positively oriented basis is obviously not unique, though.

Proposition 12. *Let Λ be a complex lattice, and let (ω_1, ω_2) be a positively oriented basis, then any other positively oriented basis (ω_1', ω_2') is of the form*

$$\omega_1' = a\omega_1 + b\omega_2,$$
$$\omega_2' = c\omega_1 + d\omega_2,$$

for some matrix $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in \mathrm{SL}_2(\mathbb{Z})$.

Proof. See [67, I, Lem. 2.4]. □

Definition 13 (Complex torus)**.** Let Λ be a complex lattice, the quotient $\mathbb{C}/\Lambda$ is called a *complex torus*.

A convex set of class representatives of $\mathbb{C}/\Lambda$ is called a *fundamental parallelogram*. Figure 2 shows a complex lattice generated by a (positively oriented) basis (ω_1, ω_2), together with a fundamental parallelogram for $\mathbb{C}/(\omega_1, \omega_2)$. The additive group structure of $\mathbb{C}$ carries over to $\mathbb{C}/\Lambda$, and can be graphically represented as operations on points inside a fundamental parallelogram. This is illustrated in Figure 3.

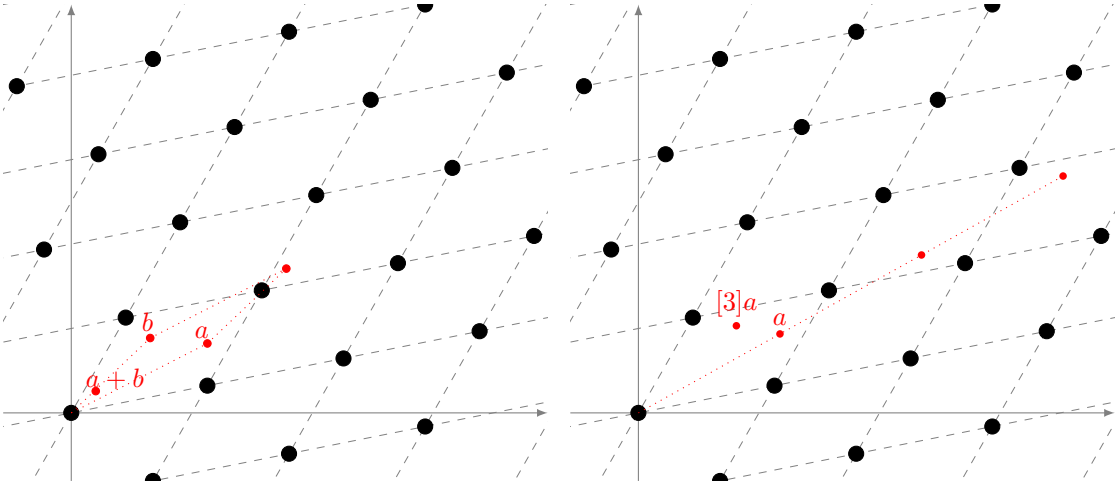

Figure 3: Addition (left) and scalar multiplication (right) of points in a complex torus $\mathbb{C}/\Lambda$.

Definition 14 (Homothetic lattices)**.** Two complex lattices Λ and Λ' are said to be *homothetic* if there is a complex number $\alpha \in \mathbb{C}^\times$ such that $\Lambda = \alpha\Lambda'$.

Geometrically, applying a homothety to a lattice corresponds to zooms and rotations around the origin. We are only interested in complex tori up to homothety; to classify them, we introduce the *Eisenstein series of weight* $2k$, defined as

$$G_{2k}(\Lambda) = \sum_{\omega \in \Lambda \setminus \{0\}} \omega^{-2k}.$$

It is customary to set

$$g_2(\Lambda) = 60G_4(\Lambda), \quad g_3(\Lambda) = 140G_6(\Lambda);$$

when Λ is clear from the context, we simply write g_2 and g_3.

Theorem 15 (Modular j-invariant)**.** *The* modular j-invariant *is the function on complex lattices defined by*

$$j(\Lambda) = 1728\frac{g_2(\Lambda)^3}{g_2(\Lambda)^3 - 27g_3(\Lambda)^2}.$$

Two lattices are homothetic if and only if they have the same modular j-invariant.

Proof. See [67, I, Th. 4.1]. □

It is no chance that the invariants classifying elliptic curves and complex tori look very similar. Indeed, we can prove that the two are in one-to-one correspondence.

Definition 16 (Weierstrass $\wp$ function)**.** Let Λ be a complex lattice, the *Weierstrass $\wp$ function* associated to Λ is the series

$$\wp(z;\Lambda) = \frac{1}{z^2} + \sum_{\omega \in \Lambda \setminus \{0\}} \left(\frac{1}{(z-\omega)^2} - \frac{1}{\omega^2} \right).$$

Theorem 17. *The Weierestrass function $\wp(z;\Lambda)$ has the following properties:*

1. *It is an* elliptic function *for* Λ, *i.e.* $\wp(z) = \wp(z+\omega)$ *for all* $z \in \mathbb{C}$ *and* $\omega \in \Lambda$.
2. *Its Laurent series around* $z = 0$ *is*
$$\wp(z) = \frac{1}{z^2} + \sum_{k=1}^{\infty} (2k+1) G_{2k+2} z^{2k}.$$
3. *It satisfies the differential equation*
$$\wp'(z)^2 = 4\wp(z)^3 - g_2\wp(z) - g_3$$
for all $z \notin \Lambda$.
4. *The curve*
$$E \ : \ y^2 = 4x^3 - g_2 x - g_3$$
is an elliptic curve over $\mathbb{C}$. *The map*
$$\begin{aligned} \mathbb{C}/\Lambda &\to E(\mathbb{C}), \\ 0 &\mapsto (0:1:0), \\ z &\mapsto (\wp(z) : \wp'(z) : 1) \end{aligned}$$
is an isomorphism of Riemann surfaces and a group morphism.

Proof. See [66, VI, Th. 3.1, Th. 3.5, Prop. 3.6]. □

By comparing the two definitions for the j-invariants, we see that $j(\Lambda) = j(E)$. So, for any homothety class of complex tori, we have a corresponding isomorphism class of elliptic curves. The converse is also true.

Theorem 18 (Uniformization theorem)**.** *Let* $a, b \in \mathbb{C}$ *be such that* $4a^3 + 27b^2 \neq 0$, *then there is a unique complex lattice* Λ *such that* $g_2(\Lambda) = -4a$ *and* $g_3(\Lambda) = -4b$.

Proof. See [67, I, Coro. 4.3]. □

Using the correspondence between elliptic curves and complex tori, we now have a new perspective on their group structure. Looking at complex tori, it becomes immediately evident why the torsion part has rank 2, i.e. why $E[m] \simeq (\mathbb{Z}/m\mathbb{Z})^2$. This is illustrated in Figure 4a; in the picture we see two lattices Λ and Λ', generated respectively by the black and the red dots. The multiplication-by-m map corresponds then to
$$\begin{aligned} [m] : \mathbb{C}/\Lambda &\to \mathbb{C}/\Lambda', \\ z &\mapsto z \bmod \Lambda'; \end{aligned}$$
or equivalently $[m] : z \mapsto mz \bmod \Lambda$, after applying the homothety $m\Lambda' = \Lambda$, as expected.

Within this new perspective, isogenies are a mild generalization of scalar multiplications. Whenever two lattices Λ, Λ' verify $\alpha\Lambda \subset \Lambda'$, there is a well defined map
$$\begin{aligned} \phi_\alpha : \mathbb{C}/\Lambda &\to \mathbb{C}/\Lambda', \\ z &\mapsto \alpha z \bmod \Lambda' \end{aligned}$$
that is holomorphic and also a group morphism. One example of such maps is given in Figure 4a: there, $\alpha = 1$ and the red lattice strictly contains the black one; the map is simply defined as reduction modulo Λ'. It turns out that these maps are exactly the isogenies of the corresponding elliptic curves.

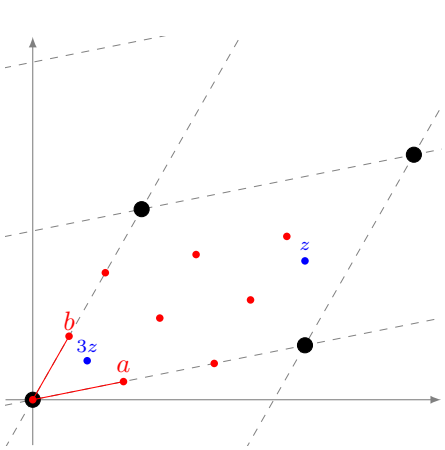

(a) 3-torsion group on a complex torus (red points), with two generators a and b, and action of the multiplication-by-3 map (blue dots).

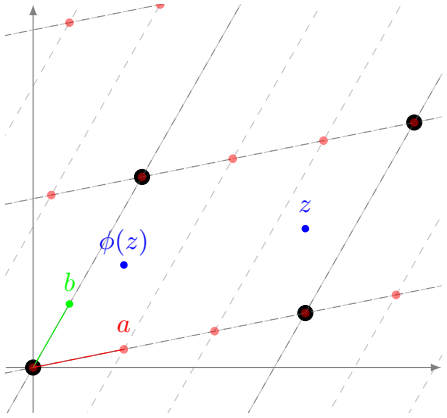

(b) Isogeny from $\mathbb{C}/\Lambda$ (black dots) to $\mathbb{C}/\Lambda'$ (red dots) defined by $\phi(z) = z \bmod \Lambda'$. The kernel of ϕ is contained in $(\mathbb{C}/\Lambda)[3]$ and is generated by a. The kernel of the dual isogeny $\hat{\phi}$ is generated by the vector b in Λ'.

Figure 4: Maps between complex tori.

Theorem 19. *Let E, E' be elliptic curves over $\mathbb{C}$, with corresponding lattices Λ, Λ'. There is a bijection between the set of isogenies from E to E' and the set of maps ϕ_α for all α such that $\alpha\Lambda \subset \Lambda'$.*

Proof. See [66, VI, Th. 4.1]. □

Looking again at Figure 4a, we see that there is a second isogeny $\hat{\phi}$ from Λ' to $\Lambda/3$, whose kernel is generated by $b \in \Lambda'$. The composition $\hat{\phi} \circ \phi$ is an endomorphism of $\mathbb{C}/\Lambda$, up to the homothety sending $\Lambda/3$ to Λ, and we verify that it corresponds to the multiplication-by-3 map. In this example, the kernels of both ϕ and $\hat{\phi}$ contain 3 elements, and we say that ϕ and $\hat{\phi}$ have *degree* 3. Although not immediately evident from the picture, this same construction can be applied to any isogeny. The isogeny $\hat{\phi}$ is called the *dual* of ϕ. Dual isogenies exist not only in characteristic 0, but also for any base field, as we shall see in Section 5.

Under which conditions does an isogeny become an endomorphism? By virtue of the last theorem, there is a one-to-one correspondence between the endomorphisms $E \to E$ and the complex numbers α such that $\alpha\Lambda \subset \Lambda$. In general, the only possible choices are given by α an integer, corresponding to scalar multiplications. For some lattices, however, something "special" happens; we shall study this case in Section 6.

4 Elliptic curves over finite fields

In this section we shift our attention to elliptic curves defined over a finite field k with q elements, which are the main objects manipulated in cryptography. Obviously, the group of k-rational points is finite, thus the algebraic group $E(\bar{k})$ only contains torsion elements, and we have already characterized precisely the structure of the torsion part of E.

For curves over finite fields, the Frobenius endomorphism plays a very special role, and governs much of their structure.

Definition 20 (Frobenius endomorphism)**.** Let E be an elliptic curve defined over a field with q elements, its *Frobenius endomorphism*, denoted by π, is the map that sends

$$(X : Y : Z) \mapsto (X^q : Y^q : Z^q).$$

Proposition 21. *Let π be the Frobenius endomorphism of E. Then:*

- $\ker \pi = \{\mathcal{O}\}$;
- $\ker(\pi - 1) = E(k)$.

Theorem 22 (Hasse)**.** *Let E be an elliptic curve defined over a finite field with q elements. Its Frobenius endomorphism π satisfies a quadratic equation*

$$\pi^2 - t\pi + q = 0,$$

for some $|t| \leq 2\sqrt{q}$.

Proof. See [66, V, Th. 2.3.1]. □

The coefficient t in the equation is called the *trace* of π. By replacing $\pi = 1$ in the equation, we immediately obtain the cardinality of E as $\#E(k) = \#\ker(\pi - 1) = q + 1 - t$.

Corollary 23. *Let E be an elliptic curve defined over a finite field k with q elements, then*

$$|\#E(k) - q - 1| \leq 2\sqrt{q}.$$

It turns out that the cardinality of E over its *base field* k determines its cardinality over any finite extension of it. This is a special case of Weil's famous conjectures, proven by Weil himself in 1949 for Abelian varieties, and more generally by Deligne in 1973.

Definition 24. Let V be a projective variety defined over a finite field $\mathbb{F}_q$, its *zeta function* is the power series

$$Z(V/\mathbb{F}_q; T) = \exp\left(\sum_{n=1}^{\infty} \#V(\mathbb{F}_{q^n}) \frac{T^n}{n}\right).$$

Theorem 25. *Let E be an elliptic curve defined over a finite field $\mathbb{F}_q$, and let $\#E(\mathbb{F}_q) = q+1-a$. Then*

$$Z(E/\mathbb{F}_q; T) = \frac{1 - aT + qT^2}{(1 - T)(1 - qT)}.$$

Proof. See [66, V, Th. 2.4]. □

5 Isogenies

We now look more in detail at isogenies of elliptic curves. We start with some basic definitions.

Definition 26 (Degree, separability)**.** Let $\phi : E \to E'$ be an isogeny defined over a field k, and let $k(E), k(E')$ be the function fields of E, E'. By composing ϕ with the functions of $k(E')$, we obtain a subfield of $k(E)$ that we denote by $\phi^*(k(E'))$.

1. The *degree* of ϕ is defined as $\deg \phi = [k(E) : \phi^*(k(E'))]$; it is always finite.

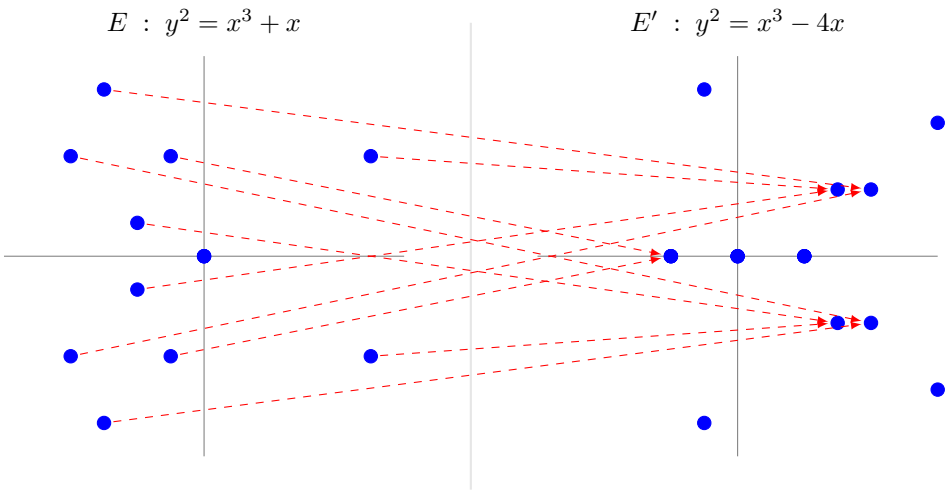

Figure 5: The isogeny $(x, y) \mapsto \left((x^2+1)/x,\ y(x^2-1)/x^2\right)$, as a map between curves defined over $\mathbb{F}_{11}$.

2. ϕ is said to be *separable*, *inseparable*, or *purely inseparable* if the extension of function fields is.
3. If ϕ is separable, then $\deg \phi = \#\ker \phi$.
4. If ϕ is purely inseparable, then $\deg \phi$ is a power of the characteristic of k.
5. Any isogeny can be decomposed as a product of a separable and a purely inseparable isogeny.

Proof. See [66, II, Th. 2.4]. □

In practice, most of the time we will be considering separable isogenies, and we can take $\deg \phi \equiv \#\ker \phi$ as the definition of the degree. Notice that in this case $\deg \phi$ is the size of any fiber of ϕ.

Example 27. The map ϕ from the elliptic curve $y^2 = x^3 + x$ to $y^2 = x^3 - 4x$ defined by

$$\begin{aligned} \phi(x, y) &= \left(\frac{x^2+1}{x}, y\frac{x^2-1}{x^2}\right), \\ \phi(0, 0) &= \phi(\mathcal{O}) = \mathcal{O} \end{aligned} \tag{2}$$

is a separable isogeny between curves defined over $\mathbb{Q}$. It has degree 2, and its kernel is generated by the point $(0, 0)$.

Plotting the isogeny (2) over $\mathbb{R}$ would be cumbersome, however, since the curves are defined by integer coefficients, we may reduce the equations modulo a prime p, then the isogeny descends to an isogeny of curves over $\mathbb{F}_p$. Figure 5 plots the action of the isogeny after reduction modulo 11. A red arrow indicates that a point of the left curve is sent onto a point on the right curve; the action on the point in $(0, 0)$, going to the point at infinity, is not shown. We observe a symmetry with respect to the x-axis, a consequence of the fact that ϕ is a group morphism; and, by looking closer, we may also notice that collinear points are sent to collinear points, also a necessity for a group morphism.

It is evident that the isogeny is 2-to-1, however we are unable to see all fibers over $\mathbb{F}_p$, because the isogeny is only surjective over the algebraic closure. This is not dissimilar from the way power-by-n maps act on the multiplicative group $k^\times$ of a field k: the map $x \mapsto x^2$, for example, is a 2-to-1 (algebraic) group morphism on $\mathbb{F}_{11}^\times$, and those elements that have no preimage, the non-squares, will have exactly two square roots in $\mathbb{F}_{11^2}$, and so on.

The most unique property of separable isogenies is that they are entirely determined by their kernel.

Proposition 28. *Let E be an elliptic curve defined over an algebraically closed field, and let G be a finite subgroup of E. There is a curve E', and a separable isogeny ϕ, such that* $\ker \phi = G$ *and $\phi : E \to E'$. Furthermore, E' and ϕ are unique up to composition with an isomorphism $E' \simeq E''$.*

Said otherwise, for any finite subgroup $G \subset E$, we have an exact sequence of algebraic groups

$$0 \longrightarrow G \longrightarrow E \xrightarrow{\phi} E' \longrightarrow 0.$$

Uniqueness up to isomorphisms justifies the notation E/G for the isomorphism class of the image curve E'. Conversely, since any non-constant morphism of elliptic curves necessarily has finite kernel, we have a bijection between the finite subgroups of a curve E and the isogenies with domain E up to isomorphisms. This correspondence is rich in consequences: it is an easy exercise to prove the following useful facts.

Corollary 29.

1. *Any isogeny of elliptic curves can be decomposed as a product of prime degree isogenies.*
2. *Let E be defined over an algebraically closed field k, let ℓ be a prime different from the characteristic of k, then there are exactly $\ell + 1$ isogenies of degree ℓ with domain E, up to isomorphism.*

Slightly more work is required to prove the following, fundamental, theorem (the difficulty comes essentially from the inseparable part, see [66, III.6.1] for a detailed proof).

Theorem 30 (Dual isogeny theorem)**.** *Let $\phi : E \to E'$ be an isogeny of degree m. There is a unique isogeny $\hat{\phi} : E' \to E$ such that*

$$\hat{\phi} \circ \phi = [m]_E, \quad \phi \circ \hat{\phi} = [m]_{E'}.$$

$\hat{\phi}$ is called the dual isogeny *of ϕ; it has the following properties:*

1. *$\hat{\phi}$ has degree m;*
2. *$\hat{\phi}$ is defined over k if and only if ϕ is;*
3. *$\widehat{\psi \circ \phi} = \hat{\phi} \circ \hat{\psi}$ for any isogeny $\psi : E' \to E''$;*
4. *$\widehat{\psi + \phi} = \hat{\psi} + \hat{\phi}$ for any isogeny $\psi : E \to E'$;*
5. $\deg \phi = \deg \hat{\phi}$;
6. $\hat{\hat{\phi}} = \phi$.

The computational counterpart to the kernel-isogeny correspondence is given by Vélu's much celebrated formulas.

Proposition 31 (Vélu [75])**.** *Let $E : y^2 = x^3 + ax + b$ be an elliptic curve defined over a field k, and let $G \subset E(\bar{k})$ be a finite subgroup. The separable isogeny $\phi : E \to E/G$, of kernel G, can be written as*

$$\phi(P) = \left(x(P) + \sum_{Q \in G \setminus \{\mathcal{O}\}} x(P+Q) - x(Q), y(P) + \sum_{Q \in G \setminus \{\mathcal{O}\}} y(P+Q) - y(Q) \right);$$

and the curve E/G has equation $y^2 = x^3 + a'x + b'$, where

$$a' = a - 5 \sum_{Q \in G \setminus \{\mathcal{O}\}} (3x(Q)^2 + a),$$
$$b' = b - 7 \sum_{Q \in G \setminus \{\mathcal{O}\}} (5x(Q)^3 + 3ax(Q) + 2b).$$

6 Complex multiplication

We conclude with one of the most power tools for the study of isogeny graphs: the theory of *complex multiplication*. Our goal is to characterize elliptic curves with endomorphism rings larger than $\mathbb{Z}$; to do so, we start from elliptic curves defined over the complex numbers. But first, we need to recall some basic definitions from algebraic number theory; for a more detailed treatment, see [47].

An *quadratic number field* is a quadratic extension K of the rationals; it is called *real* if there exists an embedding $K \subset \mathbb{R}$, *imaginary* otherwise. All such fields can be expressed as $\mathbb{Q}(\sqrt{d})$ for some integer d, the *Gaussian integers* $\mathbb{Q}(i)$ being a typical example of an imaginary one.

Definition 32 (Discriminant)**.** Let d be a square free integer, the *discriminant* of the quadratic number field $\mathbb{Q}(\sqrt{d})$ is d if $d = 1 \bmod 4$, and $4d$ otherwise.

An integer Δ that is the discriminant of a quadratic number field is called a *fundamental discriminant*.

Definition 33 (Ring of integers)**.** Let K be a number field, an *algebraic integer* of K is an element $\alpha \in K$ that is root of an irreducible monic polynomial with integer coefficients. The algebraic integers of K form a ring, called the *ring of integers* of K.

For example, $\mathbb{Z}[i]$ is the ring of integers of $\mathbb{Q}(i)$; more generally, if Δ is a fundamental discriminant, the ring of integers of $\mathbb{Q}(\sqrt{\Delta})$ is $\mathbb{Z}[\delta]$, where $\delta = (\Delta + \sqrt{\Delta})/2$. By Definition 8, an order of a quadratic field K is a subring of K that is a $\mathbb{Z}$-module of rank 2. The ring of integers $\mathcal{O}_K$ of K fits the bill: it always has $(1, \delta)$ as *integral basis*, i.e., as a set of $\mathbb{Z}$-module generators. Furthermore, it is easy to prove that any other order is contained in $\mathcal{O}_K$; for this reason we will some times call it the *maximal order* of K. More precisely, we can prove the following.

Proposition 34. *Let K be a quadratic number field, and let $\mathcal{O}_K$ be its ring of integers. Any order $\mathcal{O} \subset K$ can be written as $\mathcal{O} = \mathbb{Z} + f\mathcal{O}_K$ for an integer f, called the* conductor *of $\mathcal{O}$. If Δ_K is the discriminant of K, the* discriminant *of $\mathcal{O}$ is $f^2\Delta_K$.*

If $\mathcal{O}, \mathcal{O}'$ are two orders of discriminants Δ, Δ', then $\mathcal{O} \subset \mathcal{O}'$ if and only if $\Delta' | \Delta$.

When K is imaginary quadratic, any order $\mathcal{O} \subset K \subset \mathbb{C}$ is a complex lattice by definition. We now define a broader class of algebraic lattices, that are not necessarily rings.

Definition 35 (Fractional ideal)**.** Let $\mathcal{O}$ be an order of a number field K. A *(fractional) $\mathcal{O}$-ideal* $\mathfrak{a}$ is a finitely generated non-zero $\mathcal{O}$-submodule of K.

If $\mathfrak{a}$ is generated by a single element, then it is called *principal*. If $\mathfrak{a} \subset \mathcal{O}$, then it is called an *integral* ideal.

An $\mathcal{O}$-ideal $\mathfrak{a}$ is *invertible* if there exists another ideal $\mathfrak{a}^{-1}$ such that $\mathfrak{a}\mathfrak{a}^{-1} = \mathfrak{a}^{-1}\mathfrak{a} = \mathcal{O}$. If $\mathcal{O}$ is the maximal order of K, then any $\mathcal{O}$-ideal is invertible.

When $\mathcal{O}$ is the maximal order, we often omit specifying the order, and simply speak of (fractional) ideals of K.

Now, let K be a quadratic imaginary field. Let Λ be a complex lattice such that $\Lambda \subset K$, and define its order $\mathcal{O}_\Lambda$ to be

$$\mathcal{O}_\Lambda = \{\alpha \in K \mid \alpha\Lambda \subset \Lambda\}. \tag{3}$$

It is clear that $\mathcal{O}_\Lambda$ is a ring, and it is easy to show that it is an order of K, and thus that Λ is a fractional $\mathcal{O}_\Lambda$-ideal. Using Theorem 17 we associate to Λ a complex elliptic curve E_Λ; but then, by definition, $\mathcal{O}_\Lambda \simeq \mathrm{End}(E_\Lambda)$. Said otherwise, E_Λ *complex multiplication* by $\mathcal{O}_\Lambda$.

We have thus found a way to construct elliptic curves over the complex numbers with complex multiplication by a specified order. Conversely, every curve with complex multiplication arises this way. To show this, we look at the set of all isomorphism classes of elliptic curves with complex multiplication by a specified order $\mathcal{O}$, which we will denote by $\mathrm{Ell}(\mathcal{O})$. Because homothetic lattices give rise to isomorphic curves, fractional ideals $\mathfrak{a}$ and $c\mathfrak{a}$ will be associated to isomorphic curves $E_\mathfrak{a}$ and $E_{c\mathfrak{a}}$ as long as $c \neq 0$. This justifies looking at fractional ideals modulo principal ideals.

Definition 36 (Ideal class group)**.** Let $\mathcal{O}$ be an order of a number field K. Let $\mathcal{I}(\mathcal{O})$ be the group of invertible fractional $\mathcal{O}$-ideals, and let $\mathcal{P}(\mathcal{O})$ be the group of principal ideals.

The *ideal class group* of $\mathcal{O}$ is the quotient group

$$\mathrm{Cl}(\mathcal{O}) = \mathcal{I}(\mathcal{O})/\mathcal{P}(\mathcal{O}).$$

It is a finite Abelian group; its order is called the *class number* of $\mathcal{O}$, and denoted by $h(\mathcal{O})$.

When $\mathcal{O}$ is the maximal order, $\mathrm{Cl}(\mathcal{O})$ is also called the class group of K. The class group is a fundamental object in *class field theory*: when $\mathcal{O}$ is the maximal order, it is isomorphic to the Galois group of the maximal unramified Abelian extension of K, also called the *Hilbert class field* of K; more generally, non-maximal orders are connected to ramified Abelian extensions of K. The next theorem highlights a fundamental connection between the class group and the modular j-invariant, and thus to elliptic curves with complex multiplication by $\mathcal{O}$.

Theorem 37. *Let $\mathcal{O}$ be an order of a number field K, and let $\mathfrak{a}_1, \ldots, \mathfrak{a}_{h(\mathcal{O})}$ be representatives of $\mathrm{Cl}(\mathcal{O})$. Then:*

- *$K(j(\mathfrak{a}_i))$ is an Abelian extension of K;*
- *The $j(\mathfrak{a}_i)$ are all conjugate over K;*
- *The Galois group of $K(j(\mathfrak{a}_i))$ is isomorphic to $\mathrm{Cl}(\mathcal{O})$;*
- *$[\mathbb{Q}(j(\mathfrak{a}_i)) : \mathbb{Q}] = [K(j(\mathfrak{a}_i)) : K] = h(\mathcal{O})$;*
- *The $j(\mathfrak{a}_i)$ are integral, their minimal polynomial is called the* Hilbert class polynomial *of $\mathcal{O}$;*
- *$\mathrm{Cl}(\mathcal{O})$ acts freely and transitively on $\mathrm{Ell}(\mathcal{O})$, in particular $\#\,\mathrm{Ell}(\mathcal{O}) = h(\mathcal{O})$.*

Proof. See [67, Ch. II] and [46, Ch. 10]. □

Hence, we have completely characterized all elliptic curves with complex multiplication by an order $\mathcal{O}$, up to isomorphism; in particular, we now know that j-invariants with complex multiplication (sometimes called *singular j-invariants*) are algebraic integers. In the next part, we shall say more on how $\mathrm{Cl}(\mathcal{O})$ acts on the set $\mathrm{Ell}(\mathcal{O})$.

Example 38. Let $\mathcal{O} = \mathbb{Z}[i]$, so that $\mathcal{O}$ is the ring of integers of $\mathbb{Q}(i)$. It was already proven by Gauss that $\mathbb{Z}[i]$ is a principal ideal domain, and thus that its class group is trivial. Up to homothety, there is a unique lattice with order $\mathbb{Z}[i]$, and one such representative is $\mathbb{Z}[i]$ itself.

Recall the definition of the Eisenstein series

$$G_{2k}(\Lambda) = \sum_{\omega \in \Lambda \setminus \{0\}} \omega^{-2k}.$$

But in our case $\Lambda = \mathbb{Z}[i]$, thus $i\Lambda = \Lambda$, hence

$$G_{2k}(\Lambda) = G_{2k}(i\Lambda) = i^{-2k} G_{2k}(\Lambda) = (-1)^k G_{2k}(\Lambda).$$

In particular $G_6(\Lambda) = -G_6(\Lambda) = 0$, hence, by the definition of the modular j-invariant, $j(\mathbb{Z}[i]) = 1728$.

This shows that that the Hilbert class polynomial of $\mathbb{Z}[i]$ is $X - 1728$, and that the curve $E : y^2 = x^3 + x$ is the only curve over $\mathbb{C}$, up to isomorphism, with complex multiplication by $\mathbb{Z}[i]$. In particular, $\mathbb{Z}[i]$ contains a subgroup of units $\{\pm 1, \pm i\}$, which correspond to the four automorphisms generated by the map

$$\begin{aligned} \iota : E &\longrightarrow E, \\ (x, y) &\longmapsto (-x, iy). \end{aligned}$$

6.1 Complex multiplication for finite fields

At this point, we have a complete characterization of complex multiplication elliptic curves in characteristic 0. What happens, then, in positive characteristic p?

There are at least two ways in which we could construct elliptic curves over a finite field with endomorphism ring larger than $\mathbb{Z}$. One is to start from a complex multiplication elliptic curve E defined over a number field L, and then reduce at a place[2] $\mathfrak{p}$ over p. We write $\bar{E} = E(\mathfrak{p})$ for the reduction of E at the place $\mathfrak{p}$; if we do this carefully (for example, we must avoid singular reductions), non-trivial endomorphisms of E will descend to non-trivial endomorphisms of $\bar{E}$.

Theorem 39 (Deuring)**.** *Let E be an elliptic curve over a number field L, with complex multiplication by an order $\mathcal{O} \subset K$. Let $\mathfrak{p}$ be a place of L over p, and assume that E has non-singular reduction $\bar{E}$ modulo $\mathfrak{p}$. The curve $\bar{E}$ is supersingular if and only if p has only one prime of K above it (p ramifies or remains prime in k).*

Suppose that p splits completely in K. Let f be the conductor of $\mathcal{O}$, and write $f = p^r f_0$, where $p \nmid f_0$. Then:

- *$\bar{E}$ has complex multiplication by the order in K with conductor f_0.*
- *If $p \nmid f$, then the map $\omega \mapsto \omega(\mathfrak{p})$ defines an isomorphism of $\mathrm{End}(E)$ and $\mathrm{End}(\bar{E})$.*

Proof. See [46, Ch. 13]. □

[2] A *place* is just a fancy name for a prime ideal of L.

Note that $p > 2$ splits in K if and only if the fundamental discriminant Δ_K of K is is a square modulo p. To include the case $p = 2$, we may use the Kronecker symbol $\left(\frac{\Delta_K}{p}\right)$, which is equal to 1 if and only if p splits.

Example 40. We have seen that the elliptic curve $E/\mathbb{Q}$ defined by $y^2 = x^3 + x$ has complex multiplication by $\mathbb{Z}[i]$. Assume $p > 2$; by virtue of the theorem above, $E(p)$ is supersingular if and only if $(-4/p) = -1$, i.e., if and only if $p \equiv 3 \bmod 4$.

In particular, this implies that -1 is not a square modulo p, and thus that the automorphism $(x, y) \mapsto (-x, iy)$ does not descend to an $\mathbb{F}_p$-automorphism of $E(p)$. It does, however, descend to an $\mathbb{F}_{p^2}$-automorphism, showing that $\mathrm{End}(E(p))$ is not commutative.

Another approach is to directly construct a curve $E/\mathbb{F}_q$ so that its Frobenius endomorphism is in the desired order. Recall that the Frobenius endomorphism π satisfies a quadratic equation

$$\pi^2 - t\pi + q = 0,$$

with discriminant $\Delta_\pi = t^2 - 4q \leq 0$. Setting the case $\Delta_\pi = 0$ aside, $\mathrm{End}(E)$ necessarily contains a subring $\mathbb{Z}[\pi]$, isomorphic to an order of $\mathbb{Q}(\sqrt{\Delta_\pi})$. It turns out that these approach is essentially equivalent to the previous one, as a famous theorem shows.

Theorem 41 (Deuring's lifting theorem)**.** *Let E_0 be an elliptic curve in characteristic p, with an endomorphism ω_o which is not trivial. Then there exists an elliptic curve E defined over a number field L, an endomorphism ω of E, and a non-singular reduction of E at a place $\mathfrak{p}$ of L lying above p, such that E_0 is isomorphic to $E(\mathfrak{p})$, and ω_0 corresponds to $\omega(\mathfrak{p})$ under the isomorphism.*

Proof. See [46, Ch. 13]. □

Exercises

Exercise I.1. Prove Proposition 6.

Exercise I.2. Determine all the possible automorphisms of elliptic curves.

Exercise I.3. Prove Proposition 21.

Exercise I.4. Using Proposition 25, devise an algorithm to effectively compute $\#E(\mathbb{F}_{q^n})$ given $\#E(\mathbb{F}_q)$.

Exercise I.5. Prove Corollary 29

Exercise I.6. Let K be a complex imaginary number field, $\Lambda \subset K$ a complex lattice, and $\mathcal{O}_\Lambda$ its order as defined in Eq. (3). Prove that $\mathcal{O}_\Lambda$ is an order of K.

Exercise I.7. Let $\omega \in \mathbb{C}$ be a cube root of unity, the ring $\mathbb{Z}[\omega]$ is also known as the *Eisenstein integers*. Determine all elliptic curves with complex multiplication by $\mathbb{Z}[\omega]$.

Exercise I.8. Prove that -163 is not a square modulo all odd primes < 41. (Hint: $\mathbb{Q}(\sqrt{-163})$ has class number 1).

Part II
Isogeny graphs

We now look at isogeny graphs: graphs with isomorphisms classes of ellitpic curves for vertices, and isogenies for edges. Depending on the constraints we put on the isogenies, we will get graphs with different properties; the most important ones will be *isogeny volcanoes*, *Cayley graphs*, and *supersingular graphs*.

The classification of isogeny graphs was initiated by Mestre [53], Pizer [59, 60] and Kohel [42]; further algorithmic treatment of graphs of ordinary curves, and the now famous name of *isogeny volcanoes* was subsequently given by Fouquet and Morain [29]. We now review the different kinds of graphs.

7 Isogeny classes

We have learned previously that being isogenous is an equivalence relation,[3] it thus makes sens to speak of the *isogeny class* of an elliptic curve. Here, we are interested in characterizing these isogeny classes, and their connectivity structure. We will mostly focus on isogeny classes over finite fields, however we will occasionally mention the complex case.

We start with a theorem that links isogeny classes with the complex multiplication theory we previously learned about.

Theorem 42 (Serre-Tate)**.** *Two elliptic curves E, E' with complex or quaternionic multiplication are isogenous if and only if their* endomorphism algebras $\mathrm{End}(E) \otimes \mathbb{Q}$ *and* $\mathrm{End}(E') \otimes \mathbb{Q}$ *are isomorphic.*

In layman terms, this theorem is telling us that:

- Two curves with complex multiplications by $\mathcal{O}$ and $\mathcal{O}'$ respectively are isogenous if and only if $\mathcal{O} \subset \mathcal{O}'$ or $\mathcal{O}' \subset \mathcal{O}$; or equivalently if and only if $\mathcal{O}$ and $\mathcal{O}'$ have the same field of fractions.
- Any two supersingular curves over a field of characteristic p are isogenous.

An easy consequence for the finite field is case is the following.

Corollary 43. *Two elliptic curves E, E' defined over a finite field k are isogenous over k if and only if $\#E(k) = \#E'(k)$.*

At this stage, we are only interested in elliptic curves up to isomorphism, i.e., j-invariants. Accordingly, we say that two j-invariants are *isogenous* whenever their corresponding curves are. Like we have already done before, we shall use the notation $\mathrm{Ell}_q(\mathcal{O})$ for the subclass of elliptic j-invariants over $\bar{\mathbb{F}}_q$ with complex multiplication by an order $\mathcal{O}$.

8 Graphs

We recall some basic concepts of graph theory; for simplicity, we will restrict to undirected graphs. An undirected graph G is a pair (V, E) where V is a finite set of *vertices* and $E \subset V \times V$ is a set of unordered pairs called *edges*. Two vertices v, v' are said to be *connected by an edge*

[3] Reflexivity and transitivity are obvious, symmetry is guaranteed by the dual isogeny theorem.

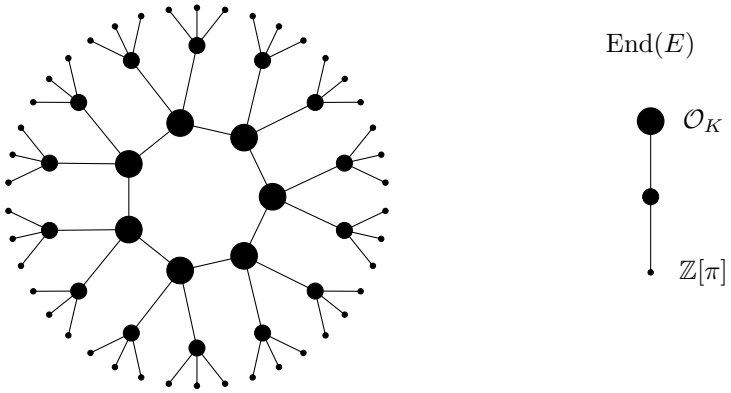

Figure 6: A volcano of 3-isogenies (ordinary elliptic curves, Elkies case), and the corresponding tower of orders inside the endomorphism algebra.

if $\{v, v'\} \in E$. The *neighbors* of a vertex v are the vertices of V connected to it by an edge. A *path* between two vertices v, v' is a sequence of vertices $v \to v_1 \to \cdots \to v'$ such that each vertex is connected to the next by an edge. The *distance* between two vertices is the length of the shortest path between them; if there is no such path, the vertices are said to be at infinite distance. A graph is called *connected* if any two vertices have a path connecting them; it is called *disconnected* otherwise. The *diameter* of a connected graph is the largest of all distances between its vertices. The *degree* of a vertex is the number of edges pointing to (or from) it; a graph where every edge has degree k is called *k-regular*. The *adjacency matrix* of a graph G with vertex set $V = \{v_1, \ldots, v_n\}$ and edge set E, is the $n \times n$ matrix where the (i, j)-th entry is 1 if there is an edge between v_i and v_j, and 0 otherwise. Because our graphs are undirected, the adjacency matrix is symmetric, thus it has n real eigenvalues

$$\lambda_1 \geq \cdots \geq \lambda_n.$$

Definition 44 (Isogeny graph)**.** An *isogeny graph* is a (multi)-graph whose vertices are the j-invariants of isogenous curves, and whose edges are isogenies between them.

The dual isogeny theorem guarantees that for every isogeny $E \to E'$ there is a corresponding isogeny $E' \to E$ of the same degree. For this reason, isogeny graphs are usually drawn undirected. Figure 6 shows a typical example of isogeny graph over a finite field, where we restrict to isogenies of degree 3.

9 ℓ-isogeny graphs

When we restrict to isogenies of a prescribed degree ℓ, we say that two curves are ℓ-isogenous; by the dual isogeny theorem, this is a symmetric relation. Remark that being ℓ-isogenous is also well defined up to isomorphism.

Let us start from the local structure: given an elliptic curve E and a prime ℓ, how many isogenies of degree ℓ have E as domain? Thanks to Proposition 28, we know this is equivalent to asking how many subgroups of order ℓ the curve has; but then we immediately know there are exactly $\ell + 1$ isogenies whenever $\ell \neq p$.

For our first example, let us consider a curve $E/\mathbb{C}$ *without* complex multiplication, i.e., such that $\mathrm{End}(E) = \mathbb{Z}$. Its ℓ-isogeny graph, i.e., the connected component of the graph of ℓ-isogenies

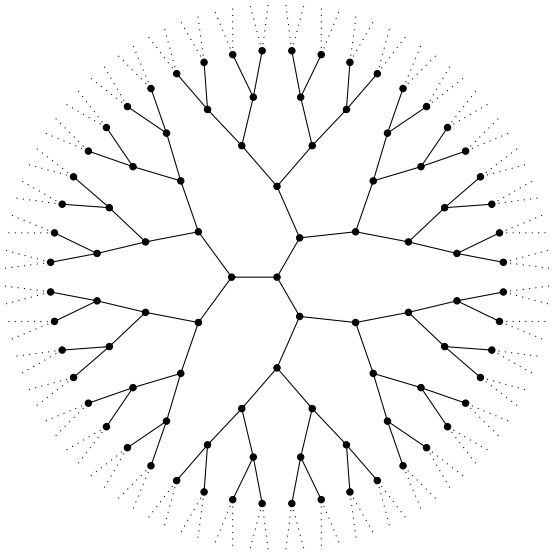

Figure 7: Infinite 2-isogeny graph of elliptic curves without complex multiplication.

containing E, is $(\ell+1)$-regular, and cannot have loops, otherwise that would provide a non-trivial endomorphism of E of degree a power of ℓ. Hence, the ℓ-isogeny graph of E is an infinite $(\ell+1)$-tree, as pictured in Figure 7.

When we think about curves over finite fields, however, some of the isogenies may only be defined in the algebraic closure, thus we would like to restrict our graphs to those isogenies that are defined over $\mathbb{F}_q$. Fortunately, we have a Swiss-army-knife to address this question: the *Frobenius endomorphism* π. Formally, an isogeny $\phi : E \to E/G$ is $\mathbb{F}_q$-rational if and only if $\pi(G) = G$, which suggests looking at the restriction of π to $E[\ell]$. Assume $\ell \neq p$, then $E[\ell]$ is a group of rank 2 and π acts on it like an element of $\mathrm{GL}_2(\mathbb{F}_\ell)$, up to conjugation. Clearly, the order of π in $\mathrm{GL}_2(\mathbb{F}_\ell)$ is the degree of the smallest extension of $\mathbb{F}_q$ where all ℓ-isogenies of E are defined. But we can tell even more by diagonalizing the matrix: π must have between 0 and 2 eigenvalues, and the corresponding eigenvectors define kernels of rational isogenies. We thus are in one of the following four cases[4]:

(0) π is not diagonalizable in $\mathbb{F}_\ell$, then E has no ℓ-isogenies.

(1.1) π has one eigenvalue of (geometric) multiplicity one, i.e., it is conjugate to a non-diagonal matrix $\left(\begin{smallmatrix}\lambda & *\\ 0 & \lambda\end{smallmatrix}\right)$; then E has one ℓ-isogeny.

(1.2) π has one eigenvalue of multiplicity two, i.e., it acts like a scalar matrix $\left(\begin{smallmatrix}\lambda & 0\\ 0 & \lambda\end{smallmatrix}\right)$; then E has $\ell+1$ isogenies of degree ℓ.

(2) π has two distinct eigenvalues, i.e., it is conjugate to a diagonal matrix $\left(\begin{smallmatrix}\lambda & 0\\ 0 & \mu\end{smallmatrix}\right)$ with $\lambda \neq \mu$; then E has two ℓ-isogenies.

Naturally, the number of eigenvalues of π depends on the factorization of its characteristic polynomial $x^2 - tx + q$ over $\mathbb{F}_\ell$, or equivalently on whether $\Delta_\pi = t^2 - 4q$ is a square modulo ℓ.

But what about the global structure? Any curve $E/\mathbb{F}_q$ can be seen as the reduction modulo p of some curve $E/\bar{\mathbb{Q}}$; thus it must inherit the connectivity structure of the isogeny graph of

[4]In the point counting literature, Case (0) is known as the *Atkin case*, and Case (2) as the *Elkies case*.

			Isogeny types		
			$\rightarrow$	$\uparrow$	$\downarrow$
$v_\ell(\Delta_\pi/\Delta_K) = 0$	$\ell \nmid [\mathcal{O}_K : \mathcal{O}]$	$\ell \nmid [\mathcal{O} : \mathbb{Z}[\pi]]$	$1 + \left(\frac{\Delta_K}{\ell}\right)$		
$v_\ell(\Delta_\pi/\Delta_K) > 1$	$\ell \nmid [\mathcal{O}_K : \mathcal{O}]$	$\ell \mid [\mathcal{O} : \mathbb{Z}[\pi]]$	$1 + \left(\frac{\Delta_K}{\ell}\right)$		$\ell - \left(\frac{\Delta_K}{\ell}\right)$
	$\ell \mid [\mathcal{O}_K : \mathcal{O}]$	$\ell \mid [\mathcal{O} : \mathbb{Z}[\pi]]$		1	ℓ
	$\ell \mid [\mathcal{O}_K : \mathcal{O}]$	$\ell \nmid [\mathcal{O} : \mathbb{Z}[\pi]]$		1	

Table 1: Number and types of ℓ-isogenies, according to splitting type of the characteristic polynomial of π.

$E/\bar{\mathbb{Q}}$. However, there is only a finite number of curves defined over $\mathbb{F}_q$, and not all isogenies will be $\mathbb{F}_q$-rational. Thus, the infinite tree must somehow "fold" or "be pruned" to fit inside $\mathbb{F}_q$.

For example, if $E/\mathbb{F}_q$ is a supersingular curve, we shall see later that its isogeny graph "folds" to a finite $(\ell+1)$-regular graph containing all supersingular curves, up to $\bar{\mathbb{F}}_q$-isomorphisms.

For the case of ordinary curves, Kohel [42] introduced a notion of "depth" in the graph. Let $E/\mathbb{F}_q$ have complex multiplication by an order $\mathcal{O}$ in a number field $K = \mathbb{Q}(\pi)$. Write $\mathcal{O}_K$ for the maximal order of K, then we know that $\mathbb{Z}[\pi] \subset \mathcal{O} \subset \mathcal{O}_K$. We have already seen that two elliptic curves are isogenous if and only if they have the same endomorphism algebra K; Kohel refined this as follows.

Proposition 45 (Kohel [42, Prop. 21])**.** *Let E, E' be elliptic curves defined over a finite field, and let $\mathcal{O}, \mathcal{O}'$ be their respective endomorphism rings. Suppose that there exists an isogeny $\phi : E \to E'$ of prime degree ℓ, then $\mathcal{O}$ contains $\mathcal{O}'$ or $\mathcal{O}'$ contains $\mathcal{O}$, and the index of one in the other divides ℓ.*

For a fixed prime ℓ, Kohel defines a curve E to be *at the surface* if $v_\ell([\mathcal{O}_K : \mathrm{End}(E)]) = 0$, where v_ℓ is the ℓ-adic valuation. E is said to be *at depth d* if $v_\ell([\mathcal{O}_K : \mathrm{End}(E)]) = d$; the maximal depth being $d_{\max} = v_\ell([\mathcal{O}_K : \mathbb{Z}[\pi]])$, curves at depth $d_{\max}$ are said to be *at the floor (of rationality)*, and $d_{\max}$ is called the *height* of the graph of E. Kohel calls then an ℓ-isogeny *horizontal* if it goes to a curve at the same depth, *descending* if it goes to a curve at greater depth, *ascending* if it goes to a curve at lesser depth.

But how many horizontal and vertical ℓ-isogenies does a given curve have? The following theorem gives a complete classification, also summarized in Table 1.

Theorem 46 (Kohel [42])**.** *Let $E/\mathbb{F}_q$ be an ordinary elliptic curve, π its Frobenius endomorphism, and Δ_K the fundamental discriminant of $\mathbb{Q}(\pi)$.*

1. *If E is not at the floor, there are $\ell + 1$ isogenies of degree ℓ from E, in total.*
2. *If E is at the floor, there are no descending ℓ-isogenies from E.*
3. *If E is at the surface, then there are $\left(\frac{\Delta_K}{\ell}\right) + 1$ horizontal ℓ-isogenies from E (and no ascending ℓ-isogenies).*
4. *If E is not at the surface, there are no horizontal ℓ-isogenies from E, and one ascending ℓ-isogeny.*

Proof. See [42, Prop. 21], or [**?**]. □

This theorem shows that, away from the surface, isogeny graphs just look like ℓ-regular complete trees of bounded height, with ℓ descending isogenies at every level except the floor. However, the surface has a more varied structure:

(0) If $\left(\frac{\Delta_K}{\ell}\right) = -1$, there are no horizontal isogenies: the isogeny graph is just a complete tree of degree $\ell + 1$ (in the graph theoretic sense) at each level but the last. We call this the *Atkin case*, as it is an extension of the Atkin case in the SEA point counting algorithm.

(1) If $\left(\frac{\Delta_K}{\ell}\right) = 0$, there is exactly one horizontal isogeny $\phi : E \to E'$ at the surface. Since E' also has one horizontal isogeny, it necessarily is $\hat{\phi}$, so the surface only contains two elliptic curves, each the root of a complete tree. We call this the *ramified case*.

(2) The case $\left(\frac{\Delta_K}{\ell}\right) = 1$ is arguably the most interesting one. Each curve at the surface has exactly two horizontal isogenies, thus the subgraph made by curves on the surface is two-regular and finite, i.e., a cycle. Below each curve of the surface there are $\ell - 1$ curves, each the root of a complete tree. We call this the *Elkies case*, again by extension of point counting.

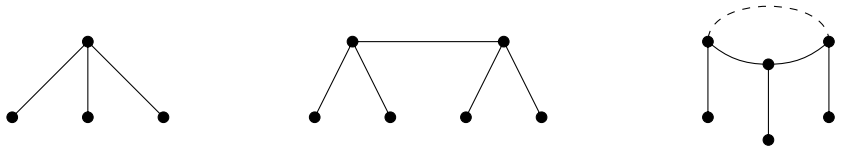

Atkin: $\left(\frac{\Delta_K}{\ell}\right) = -1$ ramified: $\left(\frac{\Delta_K}{\ell}\right) = 0$ Elkies: $\left(\frac{\Delta_K}{\ell}\right) = +1$

Figure 8: The three shapes of volcanoes of 2-isogenies of height 1.

The three cases are summarized in Figure 8. Their looks have justified the name if *isogeny volcanoes* for them [29]; in the Elkies case, we call *crater* the cycle at the surface.

We are left with one last question: how large are these graphs? Theorem 37 tells us that for any order $\mathcal{O}$ there are exactly $h(\mathcal{O})$ curves in $\mathrm{Ell}(\mathcal{O})$, thus we know exactly how many curves there are in each level of the volcano; for example we know that there will be exactly $h(\mathcal{O}_K)$ distinct trees in the Atkin case. What we do not know yet, is the number of connected components in the Elkies case. To address this question, we shall go back to complex multiplication.

10 Complex multiplication

We have already seen how the theory of complex multiplication gives a correspondence between the class group $\mathrm{Cl}(\mathcal{O})$ and the set of CM elliptic curves $\mathrm{Ell}(\mathcal{O})$. However, the (omitted) proof of Theorem 37 provides much more than a simple bijection of sets: it constructs an *action* of the group $\mathrm{Cl}(\mathcal{O})$ on the set $\mathrm{Ell}(\mathcal{O})$. We now complete our study of complex multiplication by defining the group action, and then constructing *Cayley graphs* associated to it.

We let $\mathcal{O}$ be an order in a number field K, and we assume that $\mathrm{Ell}_q(\mathcal{O})$ is non-empty. Because curves in $\mathrm{Ell}_q(\mathcal{O})$ are connected exclusively by horizontal (cyclic) isogenies, we will call it a *horizontal isogeny class*.

Let $E \in \mathrm{Ell}_q(\mathcal{O})$, let $\mathfrak{a}$ be an invertible ideal in $\mathcal{O}$, of norm coprime to q, and define the $\mathfrak{a}$-*torsion* subgroup of E as

$$E[\mathfrak{a}] = \{P \in E(\bar{\mathbb{F}}_q) \mid \sigma(P) = 0 \text{ for all } \sigma \in \mathfrak{a}\}.$$

This subgroup is the kernel of a separable isogeny $\phi_{\mathfrak{a}} : E \to E/E[\mathfrak{a}]$; it can be proven that $\phi_{\mathfrak{a}}$ is horizontal, and that its degree is the *norm* of $\mathfrak{a}$. By composing with an appropriate purely inseparable isogeny, the definition of $\phi_{\mathfrak{a}}$ is easily extended to invertible ideals of any norm.

Writing $\mathfrak{a} \cdot E$ for the isomorphism class of the image of $\phi_{\mathfrak{a}}$, we get an action $\cdot : \mathcal{I}(\mathcal{O}) \times \mathrm{Ell}_q(\mathcal{O}) \to \mathrm{Ell}_q(\mathcal{O})$ of the group of invertible ideals of $\mathcal{O}$ on $\mathrm{Ell}_q(\mathcal{O})$. It is then apparent that endomorphisms of E correspond to principal ideals in $\mathcal{O}$, and act trivially on $\mathrm{Ell}_q(\mathcal{O})$. Since the action factors through principal ideals, it natural to consider the induced action of $\mathrm{Cl}(\mathcal{O})$ on $\mathrm{Ell}_q(\mathcal{O})$. The main theorem of complex multiplication states that this action is *simply transitive*.

Theorem 47 (Complex multiplication)**.** *Let $\mathbb{F}_q$ be a finite field, $\mathcal{O} \subset \mathbb{Q}(\sqrt{-D})$ an order in a quadratic imaginary field, and $\mathrm{Ell}_q(\mathcal{O})$ the set of $\bar{\mathbb{F}}_q$-isomorphism classes of curves with complex multiplication by $\mathcal{O}$.*

Assume $\mathrm{Ell}_q(\mathcal{O})$ is non-empty, then it is a principal homogeneous space *for the class group $\mathrm{Cl}(\mathcal{O})$, under the action*

$$\begin{aligned} \mathrm{Cl}(\mathcal{O}) \times \mathrm{Ell}_q(\mathcal{O}) &\longrightarrow \mathrm{Ell}_q(\mathcal{O}), \\ (\mathfrak{a}, E) &\longmapsto \mathfrak{a} \cdot E \end{aligned}$$

defined above.

Being a principal homogeneous space (also called a *torsor*) means that, for any fixed base point $E \in \mathrm{Ell}_q(\mathcal{O})$, there is a bijection

$$\begin{aligned} \mathrm{Cl}(\mathcal{O}) &\longrightarrow \mathrm{Ell}_q(\mathcal{O}) \\ \text{Ideal class of } \mathfrak{a} &\longmapsto \text{Isomorphism class of } \mathfrak{a} \cdot E. \end{aligned}$$

This confirms what we already knew, that $\#\,\mathrm{Ell}_q(\mathcal{O}) = h(\mathcal{O})$, but also answers our question on the size of ℓ-isogeny volcanoes.

Corollary 48. *Let $\mathcal{O}$ be a quadratic imaginary order, and assume that $\mathrm{Ell}_q(\mathcal{O})$ is non-empty. Let ℓ be a prime such that $\mathcal{O}$ is ℓ-maximal, i.e., such that ℓ does not divide the conductor of $\mathcal{O}$. All ℓ-isogeny volcanoes of curves in $\mathrm{Ell}_q(\mathcal{O})$ are isomorphic. Furthermore, one of the following is true.*

(0) If the ideal (ℓ) is prime in $\mathcal{O}$, then there are $h(\mathcal{O})$ distinct ℓ-isogeny volcanoes of Atkin type, with surface in $\mathrm{Ell}_q(\mathcal{O})$.

(1) If (ℓ) is ramified in $\mathcal{O}$, i.e., if it decomposes as a square $\mathfrak{l}^2$, then there are $h(\mathcal{O})/2$ distinct ℓ-isogeny volcanoes of ramified type, with surface in $\mathrm{Ell}_q(\mathcal{O})$.

(2) If (ℓ) splits as a product $\mathfrak{l} \cdot \hat{\mathfrak{l}}$ of two distinct prime ideals, then there are $h(\mathcal{O})/n$ distinct ℓ-isogeny volcanoes of Elkies type, with craters in $\mathrm{Ell}_q(\mathcal{O})$ of size n, where n is the order of $\mathfrak{l}$ in $\mathrm{Cl}(\mathcal{O})$.

But we can extract even more information from the group action. Assume that the Frobenius endomorphism splits modulo ℓ, i.e., that

$$\pi^2 - t\pi + q = (\pi - \lambda)(\pi - \mu) \mod \ell$$

for two distinct eigenvalues λ, μ. Associate to λ and μ the prime ideals $\mathfrak{a} = (\pi - \lambda, \ell)$ and $\hat{\mathfrak{a}} = (\pi - \mu, \ell)$, both of norm ℓ; then $E[\mathfrak{a}]$ is the eigenspace of λ, and $E[\hat{\mathfrak{a}}]$ that of μ. Because $\mathfrak{a}\hat{\mathfrak{a}} = \hat{\mathfrak{a}}\mathfrak{a} = (\ell)$, the ideal classes $\mathfrak{a}$ and $\hat{\mathfrak{a}}$ are the inverse of one another in $\mathrm{Cl}(\mathcal{O})$, therefore the isogenies $\phi_{\mathfrak{a}} : E \to \mathfrak{a} \cdot E$ and $\phi_{\hat{\mathfrak{a}}} : \mathfrak{a} \cdot E \to E$ are dual to one another (up to isomorphism).

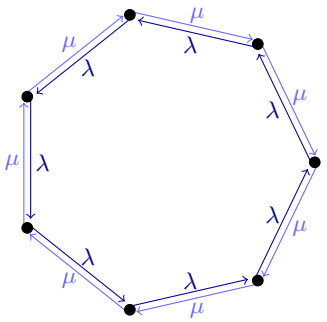

Figure 9: An isogeny cycle for an Elkies prime ℓ, with edge directions associated with the Frobenius eigenvalues λ and μ.

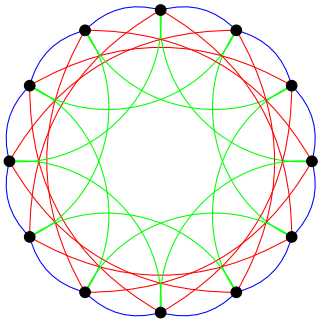

Figure 10: Graph of horizontal isogenies on 12 curves, with isogenies of three different degrees (represented in different colors).

Hence, we see that the eigenvalues λ and μ define two opposite directions on the ℓ-isogeny crater, independent of the starting curve, as shown in Figure 9. The size of the crater is the order of $(\pi - \lambda, \ell)$ in $\mathrm{Cl}(\mathcal{O})$, and the set $\mathrm{Ell}_q(\mathcal{O})$ is partitioned into craters of equal size. What we have here is a very basic example of *Cayley graph*.

Definition 49 (Cayley graph)**.** Let G be a group and $S \subset G$ be a symmetric subset (i.e., $s \in S$ implies $s^{-1} \in S$). The *Cayley graph* of (G, S) is the undirected graph whose vertices are the elements of G, and such that there is an edge between g and sg if and only if $s \in S$.

The graph in Figure 9 is isomorphic to a Cayley graph of $\mathrm{Cl}(\mathcal{O})$ for an edge set $S = \{\mathfrak{a}, \hat{\mathfrak{a}}\}$, but, unlike the Cayley graph itself, its vertex set is $\mathrm{Ell}_q(\mathcal{O})$, which is in bijection with $\mathrm{Cl}(\mathcal{O})$ only up to automorphism.[5] This graph is sometimes called the *Schreier graph* of $(\mathrm{Cl}(\mathcal{O}), S, \mathrm{Ell}_q(\mathcal{O}))$, to distinguish it from the proper Cayley graph.

In the sequel, we shall work with a larger edge set S, which will amount to "glue many isogeny craters together", as shown in Figure 10.

11 Quaternionic multiplication

Supersingular curves are generally not covered by the theory of complex multiplication. For most of them, indeed, the Frobenius endomorphism acts like an element of $\mathbb{Z}$, instead of acting like a "complex multiplier".

Supersingular curves are defined by the fact that multiplication by p is purely inseparable, i.e., $E[p]$ is trivial. This implies that the curve $E^{(p^2)}$ is isomorphic to E, and thus that both are isomorphic to a curve defined over $\mathbb{F}_{p^2}$.

When E is defined over $\mathbb{F}_p$, we can still use what we know about complex multiplication. Indeed, these curves have trace 0, and thus the Frobenius endomorphism has two distinct eigenvalue $\pm\sqrt{-p}$, implying that $\mathrm{End}_{\mathbb{F}_p}(E)$, the ring of $\mathbb{F}_p$-rational endomorphisms, is isomorphic to an order in $\mathbb{Q}(\sqrt{-p})$.

When E is defined over $\mathbb{F}_{p^2}$, however,[6] its Frobenius endomorphism must satisfy $\pi^2 - t\pi + p^2 = 0$, with t a multiple of p; hence, by Hasse's theorem, $t \in \{0, \pm p, \pm 2p\}$. The cases $t \in \{0, \pm p\}$ only happen for a very limited number of curves with j-invariant 0 or 1728; we are thus mostly

[5] Said otherwise, any vertex could correspond to 1, and we do not know which.

[6] This also applies to curves defined over $\mathbb{F}_p$, when we extend scalars.

interested in the case $t = \pm 2p$, i.e., $\pi = \pm p$. Then, the *full* endomorphism ring $\mathrm{End}(E)$ (i.e., not restricted to $\mathbb{F}_p$-rational endomorphisms) is isomorphic to a maximal order the quaternion algebra $B_{p,\infty}$ ramified at p and at infinity.

Example 50. The elliptic curve $y^2 = x^3 + x$ has supersingular reduction at all primes $p = 3 \bmod 4$. Its ring of $\mathbb{F}_p$-rational endomorphisms is generated by $\pi = \sqrt{-p}$, and it is not maximal in $\mathbb{Q}(\sqrt{-p})$.

The automorphism $\iota : (x, y) \mapsto (-x, iy)$ is only defined over $\mathbb{F}_{p^2}$, and does not commute with π. The full endomorphism ring is isomorphic to the order generated by π and ι inside the quaternion algebra $B_{p,\infty}$.

Like the CM case, isogenies are in correspondence with (left) ideals of $\mathcal{O}$. Unlike the CM case, $B_{p,\infty}$ has more than one maximal order, and there is no concept of *depth*, thus no ascending, descending or horizontal isogenies.

More precisely, let $\mathfrak{a} \subset B_{p,\infty}$ a lattice, the *left order* of $\mathfrak{a}$ is the ring $\mathcal{O}(\mathfrak{a}) = \{x \in B_{p,\infty} \mid x\mathfrak{a} \subset \mathfrak{a}\}$. Two lattices $\mathfrak{a}, \mathfrak{b}$ are said to be *right isomorphic* if $\mathfrak{a} = \mathfrak{b}x$ for some $x \in B_{p,\infty}$. If $\mathcal{O} \subset B_{p,\infty}$ is an order, $\mathfrak{a}$ is called a *left ideal* of $\mathcal{O}$ if $\mathcal{O} \subset \mathcal{O}(\mathfrak{a})$; the *left class set* $\mathrm{Cl}(\mathcal{O})$ is the set of right ideal classes of left ideals of $\mathcal{O}$. The order $\#\,\mathrm{Cl}(\mathcal{O})$ only depends on the quaternion algebra, and is called the *class number* of $B_{p,\infty}$. Analogous definitions can be given by swapping left and right; we refer to [76, Chapter 42] for more properties and definitions.

Like in the CM case, the set $\mathrm{Cl}(\mathcal{O})$ is in bijection with the vertex set of a supersingular graph.

Theorem 51. *Let $B_{p,\infty}$ be the quaternion algebra ramified at p and infinity, and let $\mathcal{O} \subset B_{p,\infty}$ be a maximal order. Let E_0/F_{p^2} be a supersingular elliptic curve with $\mathrm{End}(E_0) \simeq \mathcal{O}$.*

1. *The number of isomorphism classes of supersingular elliptic curves is equal to the class number of $B_{p,\infty}$.*
2. *There is a one-to-one correspondence $\mathfrak{a} \mapsto \mathfrak{a} \cdot E_0$ between $\mathrm{Cl}(\mathcal{O})$ and the set of isomorphism classes of supersingular elliptic curves, such that $\mathrm{End}(\mathfrak{a} \cdot E_0)$ is isomorphic to the right order of $\mathfrak{a}$.*

This theorem can be turned into an equivalence of categories, see [42, Theorem 45]. Thanks to the Eichler mass formula, we obtain the exact size of the isogeny class.

Corollary 52. *The number of isomorphism classes of supersingular elliptic curves is equal to*

$$\left\lfloor \frac{p}{12} \right\rfloor + \begin{cases} 0 & \text{if } p = 1 \mod 12, \\ 1 & \text{if } p = 5, 7 \mod 12, \\ 2 & \text{if } p = 11 \mod 12. \end{cases}$$

We thus have a bound on the size of a supersingular isogeny graph over $\mathbb{F}_{p^2}$. Since the Frobenius acts like a scalar, all isogenies are defined over $\mathbb{F}_{p^2}$, hence supersingular ℓ-isogeny graphs are necessarily $(\ell + 1)$-regular. In the next section we will learn that the supersingular ℓ-isogeny graph has a unique connected component.

12 Expander graphs from isogenies

We are finally introducing two families of isogeny graphs suitable for cryptographic use. We will want them to somehow "behave like large random graphs", while at the same time having a strong algebraic structure: the first is needed for security, the second to produce complex protocols such as key exchange.

The random-like properties of isogeny graphs are typically expressed in terms of *expansion*. An undirected graph on n vertices has n real eigenvalues $\lambda_1 \geq \cdots \geq \lambda_n$, and, if the graph is k-regular, it can be proven that $k = \lambda_1 \geq \lambda_n \geq -k$. Because of this equality, λ_1 is called the *trivial eigenvalue*. An *expander graph* is a k-regular graph such that its non-trivial eigenvalues are bounded away, in absolute value, from k. We recall here some basic facts about expanders; for an in depth review, see [34, 72].

Definition 53 (Expander graph)**.** Let $\varepsilon > 0$ and $k \geq 1$. A k-regular graph is called a (one-sided) *ε-expander* if

$$\lambda_2 \leq (1 - \varepsilon)k;$$

and a *two-sided ε-expander* if it also satisfies

$$\lambda_n \geq -(1 - \varepsilon)k.$$

A sequence $G_i = (V_i, E_i)$ of k-regular graphs with $\#V_i \to \infty$ is said to be a one-sided (resp. two-sided) *expander family* if there is an $\varepsilon > 0$ such that G_i is a one-sided (resp. two-sided) ε-expander for all sufficiently large i.

Theorem 54 (Ramanujan graph)**.** *Let $k \geq 1$, and let G_i be a sequence of k-regular graphs. Then*

$$\max(|\lambda_2|, |\lambda_n|) \geq 2\sqrt{k-1} - o(1),$$

as $n \to \infty$. A graph such that $|\lambda_j| \leq 2\sqrt{k-1}$ for any λ_j except λ_1 is called a Ramanujan graph.

The *spectral* definition of expansion is very practical to work with, but gives very little intuition on the topological properties of the graph. *Edge expansion* quantifies how well subsets of vertices are connected to the whole graph, or, said otherwise, how far the graph is from being disconnected.

Definition 55 (Edge expansion)**.** Let $F \subset V$ be a subset of the vertices of G. The *boundary of F*, denoted by $\partial F \subset E$, is the subset of the edges of G that go from F to $V \setminus F$. The *edge expansion ratio* of G, denoted by $h(G)$ is the quantity

$$h(G) = \min_{\substack{F \subset V, \\ \#F \leq \#V/2}} \frac{\#\partial F}{\#F}.$$

Note that $h(G) = 0$ if and only if G is disconnected. Edge expansion is strongly tied to spectral expansion, as the following theorem shows.

Theorem 56 (Discrete Cheeger inequality)**.** *Let G be a k-regular one-sided ε-expander, then*

$$\frac{\varepsilon}{2}k \leq h(G) \leq \sqrt{2\varepsilon}k.$$

Expander families of graphs have many applications in theoretical computer science, thanks to their *pseudo-randomness* properties: they are useful to construct *pseudo-random number generators*, *error-correcting codes*, *probabilistic checkable proofs*, and, most interesting to us, *cryptographic primitives*. Qualitatively, we can describe them as having *short diameter* and *rapidly mixing walks*.

Proposition 57. *Let G be a k-regular one sided ε-expander. For any vertex v and any radius $r > 0$, let $B(v, r)$ be the* ball *of vertices at distance at most r from v. Then, there is a constant $c > 0$, depending only on k and ε, such that*

$$\#B(v, r) \geq \min((1 + c)^r, \#V).$$

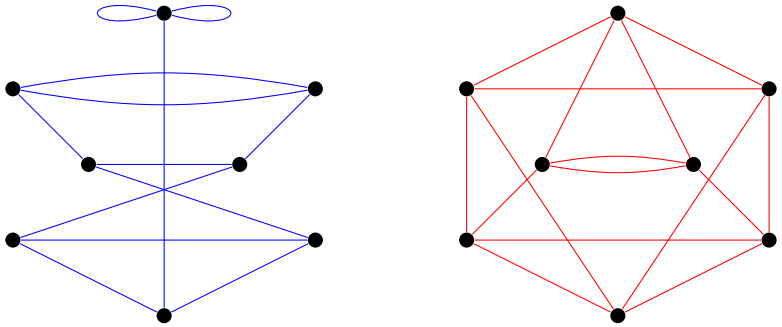

Figure 11: Supersingular isogeny graphs of degree 2 (left, blue) and 3 (right, red) on $\mathbb{F}_{97^2}$.

In particular, this shows that the diameter of an expander is bounded by $O(\log n)$, where the constant depends only on k and ε. A *random walk* of length i is a path $v_1 \to \cdots \to v_i$, defined by the random process that selects v_i uniformly at random among the neighbors of v_{i-1}. Loosely speaking, the next proposition says that, in an expander graph, random walks of length close to its diameter terminate on any vertex with probability close to uniform.

Proposition 58 (Mixing theorem ([37]))**.** *Let $G = (V, E)$ be a k-regular two-sided ε-expander. Let $F \subset V$ be any subset of the vertices of G, and let v be any vertex in V. Then a random walk of length at least*

$$\frac{\log(\#F^{1/2}/(2\#V))}{\log(1-\varepsilon)}$$

starting from v will land in F with probability at least $\#F/(2\#V)$.

The walk length in the mixing theorem is also called the *mixing length* of the expander graph.

Random regular graphs typically make good expanders, but only a handful of deterministic constructions is known, most of them based on Cayley graphs [52, 12, 34]. We just introduced Cayley graphs constructed from isogeny craters in Section 10, and, unsurprisingly, they turn out to be expanders, provided we add enough edges to them.

Theorem 59 (Jao, Miller, Venkatesan [37])**.** *Let $\mathcal{O}$ be a quadratic imaginary order, and assume that $\mathrm{Ell}_q(\mathcal{O})$ is non-empty. Let $\delta > 0$, and define the graph G on $\mathrm{Ell}_q(\mathcal{O})$ where two vertices are connected whenever there is a horizontal isogeny between them of prime degree bounded by $O((\log q)^{2+\delta})$.*

Then G is a regular graph and, under the generalized Riemann hypothesis for the characters of $\mathrm{Cl}(\mathcal{O})$, there exists an ε independent of $\mathcal{O}$ and q such that G is a two-sided ε-expander.

The theorem is readily generalized to supersingular curves and isogenies defined over $\mathbb{F}_p$.

A radically different construction of expander graphs is given by graphs of supersingular curves *defined over* $\mathbb{F}_{p^2}$ with ℓ-isogenies, for a *single* prime $\ell \neq p$. Two examples of such graphs are shown in Figure 11. This construction is related to LPS graphs [52, 51, 13], but is not isomorphic to a Cayley graph.

Theorem 60 (Mestre [53], Pizer [59, 60])**.** *Let $\ell \neq p$ be two primes. The ℓ-isogeny graph of supersingular curves in $\overline{\mathbb{F}}_p$, is connected, $(\ell+1)$-regular, and has the Ramanujan property.*

Both of these isogeny graphs will be used in the next part to build key exchange protocols. For reasons that will be apparent soon, there will only be a mild connection between the expansions properties of the graphs and the security of the protocols: the expansion theorems will mostly serve as a blueprint for devising good cryptosystems, but will have no provable impact.

Exercises

Exercise II.1. Prove Corollary 43.

Exercise II.2. Prove that the dual of a horizontal isogeny is horizontal, and that the dual of a descending isogeny is ascending.

Exercise II.3. Prove that the height of a volcano of ℓ-isogenies is $v_\ell(f_\pi)$, the ℓ-adic valuation of the Frobenius endomorphism.

Exercise II.4. Let $X^2 - tX - q$ be the minimal polynomial of π, and suppose that it splits as $(X - \lambda)(X - \mu)$ in $\mathbb{Z}_\ell$ (the ring of ℓ-adic integers). Prove that the volcano of ℓ isogenies has height $v_\ell(\lambda - \mu)$.

Exercise II.5. Prove that $E[\ell] \subset E(\mathbb{F}_q)$ implies $\ell | (q - 1)$.

Exercise II.6. Find a prime power q and an elliptic curve $E/\mathbb{F}_q$ such that the 3-isogeny volcano of E is the same as the one in Figure 6.

Part III
Key exchange

In this part we introduce our first isogeny based cryptographic protocols. We start with classic elliptic curve key exchange, then we review one of the first applications of isogeny graphs: security reduction for classic elliptic curve key exchange. Then we introduce the two best known post-quantum isogeny-based key exchange protocols: CSIDH, based on supersingular CM graphs, and SIDH, based on full supersingular graphs. We conclude with a brief discussion on the security of these protocols.

13 Diffie-Hellman key exhange

Elliptic curves are largely present in modern technology thanks to their applications in cryptography. The simplest of these application is the *Diffie-Hellman key exchange*, a cryptographic protocol by which two parties communicating over a public channel can agree on a common secret string unknown to any other party listening on the same channel.

The original protocol was invented in the 1970s by Whitfield Diffie and Martin Hellman [25], and constitutes the first practical example of *public key cryptography*. The two communicating parties are customarily called *Alice* and *Bob*, and the listening third party is represented by the character *Eve* (for *eavesdropper*). To set up the protocol, Alice and Bob agree on a set of public parameters:

- A *large enough* prime number p, such that $p-1$ has a *large enough* prime factor;
- A multiplicative generator $g \in \mathbb{Z}/p\mathbb{Z}$.

Then, Alice and Bob perform the following steps:

1. Each chooses a *secret* integer in the interval $]0, p-1[$; call a *Alice's secret* and b *Bob's secret*.
2. They respectively compute $A = g^a$ and $B = g^b$.
3. They exchange A and B over the public channel.
4. They respectively compute the *shared secret* $B^a = A^b = g^{ab}$.

The protocol can be easily generalized by replacing the multiplicative group $(\mathbb{Z}/p\mathbb{Z})^\times$ with any other cyclic group $G = \langle g \rangle$. From Eve's point of view, she is given the knowledge of the group G, the generator g, and Alice's and Bob's public data $A, B \in G$; her goal is to recover the shared secret g^{ab}. This is mathematically possible, but not necessarily *easy* from a computational point of view.

Definition 61 (Discrete logarithm)**.** Let G be a cyclic group generated by an element g. For any element $A \in G$, we define the *discrete logarithm of A in base g*, denoted $\log_g(A)$, as the unique integer in the interval $[0, \#G[$ such that

$$g^{\log_g(A)} = A.$$

It is evident that if Eve can compute discrete logarithms in G efficiently, then she can also efficiently compute the shared secret; the converse is not known to be true in general, but it is

Public parameters	Finite field $\mathbb{F}_p$, with $\log_2 p \approx 256$, Elliptic curve $E/\mathbb{F}_p$, such that $\#E(\mathbb{F}_p)$ is prime, A generator P of $E(\mathbb{F}_p)$.	
	Alice	**Bob**
Pick random secret	$0 < a < \#E(\mathbb{F}_p)$	$0 < b < \#E(\mathbb{F}_p)$
Compute public data	$A = [a]P$	$B = [b]P$
Exchange data	$A \longrightarrow$	$\longleftarrow B$
Compute shared secret	$S = [a]B$	$S = [b]A$

Figure 12: The Diffie–Hellman protocol over elliptic curves

widely believed to be. Thus, the strength of the Diffie-Hellman protocol is entirely dependent on the *hardness* of the *discrete logarithm problem* in the group G.

We know algorithms to compute discrete logarithms in a *generic* group G that require $O(\sqrt{q})$ computational steps (see [39]), where q is the largest prime divisor of $\#G$; we also know that these algorithms are *optimal for abstract cyclic groups*. For this reason, G is usually chosen so that the largest prime divisor q has size at least $\log_2 q \approx 256$. However, the proof of optimality does not exclude the existence of better algorithms for *specific* groups G. And indeed, algorithms of complexity better than $O(\sqrt{\#G})$ are known for the case $G = (\mathbb{Z}/p\mathbb{Z})^\times$ [39], thus requiring parameters of considerably larger size to guarantee cryptographic strength.

On the contrary, no algorithms better than the generic ones are known when G is a subgroup of $E(k)$, where E is an elliptic curve defined over a finite field k. This has led Miller [54] and Koblitz [41] to suggest, in the 1980s, to replace $(\mathbb{Z}/p\mathbb{Z})^\times$ in the Diffie-Helman protocol by the group of rational points of an elliptic curve of (almost) prime order over a finite field. The resulting protocol is summarized in Figure 12.

The Elliptic Curve Diffie–Hellman protocol (ECDH) is today a widely adopted standard, used for example to establish secure TLS connection, the encrypted layer of Internet. In recent years, however, the case has been made that cryptographic standards must be amended, in view of the potential threat of general purpose *quantum computers* becoming available. It is well known, indeed, that Shor's algorithm [65] would solve the factorization and the discrete logarithm problems in polynomial time on a quantum computer, thus sealing the fate of RSA, ECDH, and any other protocol based on them.

For this reason, the cryptographic community is actively seeking cryptographic primitives that would not break in polynomial time on quantum computers. In the next sections we shall present two key exchange protocols based on pseudo-random walks in certain isogeny graphs, which are conjectured to be resistant to attacks by quantum computers.

14 Isogeny graphs and discrete logarithms

Before moving to post-quantum cryptography, we quickly review one of the first applications of isogeny graphs in cryptography: hardness of discrete logarithms in families of elliptic curves. One can state several computational problems related to isogenies, both *easy* and *hard* ones. Here are some examples.

Problem 1 (Isogeny computation)**.** Given an elliptic curve E with Frobenius endomorphism π, and a subgroup $G \subset E$ such that $\pi(G) = G$, compute the rational fractions and the image curve of the separable isogeny ϕ of kernel G.

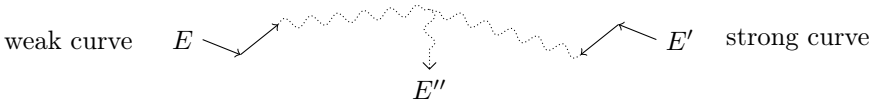

Figure 13: The meet in the middle attack in weak isogeny classes.

Vélu's formulas (Proposition 31) give a solution to this problem in $\tilde{O}(\#G)$ operations over the field of definition of E. This is nearly optimal, given that the output has size $O(\#G)$.

However in some special instances, e.g., when ϕ is a composition of many small degree isogenies, the rational fractions may be represented more compactly, and the cost may become only logarithmic in $\#G$.

Problem 2 (Explicit isogeny problem)**.** Let E be an elliptic curve, and let ℓ be an integer. Find, if it exists, an isogeny of degree ℓ with domain E.

A slightly modified version of the same problem is often found in the literature.

Problem 3. Given two elliptic curves E, E' over a finite field, isogenous of known degree ℓ, find an isogeny $\phi : E \to E'$ of degree ℓ.

Remark that, up to automorphisms, the isogeny ϕ in the latter problem is typically unique. Elkies was the first to formulate these two problems and give an algorithm [27, 28] that solves both with complexity $O(d^3)$ in general [8, 50]. Alternate algorithms, with similar complexity, are due to Couveignes and others [14, 15, 16, 24, 21].

All previous problems can be considered to be "easy" problems, because they have a polynomial time solution in the size of their output. The next problem is the blueprint for all "hard" isogeny problems used in isogeny based cryptography.

Problem 4 (Isogeny path)**.** Given two elliptic curves E, E' over a finite field k, such that $\#E = \#E'$, find an isogeny $\phi : E \to E'$ of smooth degree.

The difficulty of this problem depends on the set which E and E' are drown from, but it is in general a very difficult one, for which only algorithms exponential in $\log(\#E)$ are known. Remark that, in general, $\deg \phi$ will be exponential in $\log(\#E)$, thus the "smooth degree" constraint, which allows the output to be represented using polynomial space as a composition of small degree isogenies.

A general strategy to tackle the problem is by a *meet in the middle* random walk [30]: choose an expander graph G containing both E and E', and start a random walk from each curve. By the birthday paradox, the two walks are expected to meet after roughly $O(\sqrt{\#G})$ steps; when a collision is detected, the composition of the walks yields the desired isogeny.

The meet in the middle strategy was notoriously used to extend the power of the GHS attack on elliptic curves defined over extension fields of composite degree [33, 31]. Without going into the details of the GHS attacks, one of its remarkable properties is that only a small fraction of a given isogeny class is vulnerable to it. Finding an isogeny from an immune curve to a weak curve allows the attacker to map the discrete logarithm problem from one to the other. The average size of the isogeny class of a random ordinary elliptic curve is $O(\sqrt{\#E})$ (more on this later), thus the meet in the middle strategy yields an $O(\#E^{1/4})$ attack on any curve in the class: better than a generic attack on the discrete logarithm problem. The attack is pictured in Figure 13.

Similar ideas have been used to construct *key escrow systems* [73], and to prove random reducibility of discrete logarithms inside some isogeny classes [37].

15 Key exchange from CM graphs

The first isogeny-based protocol was introduced by Couveignes during a talk at the École Normale Superieure in 1997, although it was only published ten years later in [17]; independently, Rostovtsev and Stolbunov proposed similar protocols in [62, 68]. Couveignes' key exchange protocol was presented in a more general setting, applying to any principal homogeneous space satisfying some cryptographic properties.

Recall that a principal homogeneous space (PHS) for a group G is a set X with an action of G on X such that for any $x, x' \in X$, there is a unique $g \in G$ such that $g \cdot x = x'$. Equivalently, the map $\varphi_x : g \mapsto g \cdot x$ is a bijection between G and X for any $x \in X$. Couveignes defines a *hard homogeneous space* (HHS) to be a PHS where the action of G on X is efficiently computable, but inverting the isomorphism φ_x is computationally hard for any x.

Any HHS X for an abelian group G can be used to construct a key exchange based on the hardness of inverting φ_x: the system parameters are a HHS (G, X), and a starting point $x_0 \in X$; a secret key is a random element $g \in G$, and the associated public key is $g \cdot x_0$. If Alice and Bob have keypairs (g_A, x_A) and (g_B, x_B), respectively, then the commutativity of G lets them derive a shared secret

$$g_A \cdot x_B = g_A \cdot g_B \cdot x_0 = g_B \cdot g_A \cdot x_0 = g_B \cdot x_A.$$

The analogy with classic group-based Diffie–Hellman is evident.

Couveignes suggested to use $\mathrm{Ell}_q(\mathcal{O})$ as an instance of a HHS: the system parameters are a starting curve $E/\mathbb{F}_q$, and the associated class group $\mathrm{Cl}(\mathcal{O})$; the secret keys are random elements of $\mathrm{Cl}(\mathcal{O})$, and public keys are j-invariants of curves in $\mathrm{Ell}_q(\mathcal{O})$. However, given a generic element of $\mathrm{Cl}(\mathcal{O})$, the best algorithm [38] to evaluate its action on $\mathrm{Ell}_q(\mathcal{O})$ has subexponential complexity in q, making the protocol infeasible.

Instead, following Rostovtsev and Stolbunov [62], we may define a variant of Couveignes' HHS key exchange based on walks in a Cayley graph for G. As an example, let $G = \langle g \rangle$ be a cyclic group of order p, let $D = \{s_1, \ldots, s_n\} \subset (\mathbb{Z}/p\mathbb{Z})^\times$ be a generating set such that $\sigma \in D$ implies $\sigma^{-1} \notin D$, and let $S = D \cup D^{-1}$ so that S is symmetric as defined in 49. Then, G minus the identity is a PHS for $(\mathbb{Z}/p\mathbb{Z})^\times$ under the action

$$e \cdot g_0 = g_0^e \qquad \text{for } e \in (\mathbb{Z}/p\mathbb{Z})^\times \text{ and } g_0 \in G \setminus \{1\}.$$

We may thus define the Schreier graph of $((\mathbb{Z}/p\mathbb{Z})^\times, S, G\setminus\{1\})$, which is isomorphic to the Cayley graph $((\mathbb{Z}/p\mathbb{Z})^\times, S)$; an example for $p = 13$ is given in Figure 14.

As already seen, a random walk in this graph is a sequence of random edges starting from some vertex g_0 and ending in some vertex g_1. However we see that, because the group action of $(\mathbb{Z}/p\mathbb{Z})^\times$ is Abelian, the order in which the edges are taken from the set S does not matter for determining g_1: only matters the multiplicity of each $s \in S$. We thus define a *non-backtracking random walk* as a tuple of multiplicities $(e_1, \ldots, e_n) \in \mathbb{Z}^n$, associated to the element

$$e = \prod_{i=1}^{n} s_i^{e_i} \in (\mathbb{Z}/p\mathbb{Z})^\times,$$

defining the walk $g_0 \to e \cdot g_0$.

We can now define a key exchange protocol where the secrets are non-backtracking random walks, and the public data are vertices of the Schreier graph. The protocol is summarized in Figure 15.

Because $g_a = a \cdot g = g^a$, it is evident that this protocol is closely related to the Diffie–Hellman protocol on the group G, the only difference being that the secret exponents a, b are drawn from

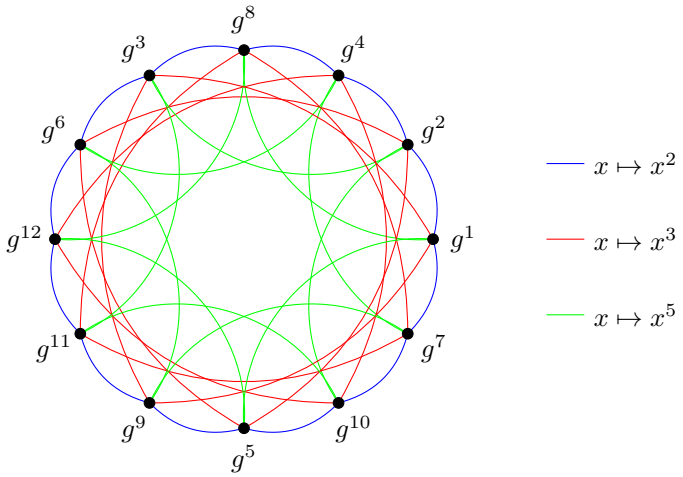

Figure 14: Schreier graph of the generators of a group of order 13 under the action of $S = \{2, 3, 5, 2^{-1}, 3^{-1}, 5^{-1}\} \subset (\mathbb{Z}/13\mathbb{Z})^\times$.

Public parameters	A group G of prime order p, A generating set $D \subset (\mathbb{Z}/p\mathbb{Z})^\times$ such that $\sigma \in D \Rightarrow \sigma^{-1} \notin D$, A generator g of G.	
	Alice	**Bob**
Pick random secret	$a = \prod_{s \in D} s^{a_i}$	$b = \prod_{s \in D} s^{b_i}$
Compute public data	$g_a = a \cdot g$	$g_b = b \cdot g$
Exchange data	$g_a \longrightarrow$	$\longleftarrow g_b$
Compute shared secret	$g_{ab} = a \cdot (g_b)$	$g_{ab} = b \cdot (g_a)$

Figure 15: Key exchange protocol based on random walks in a Schreier graph.

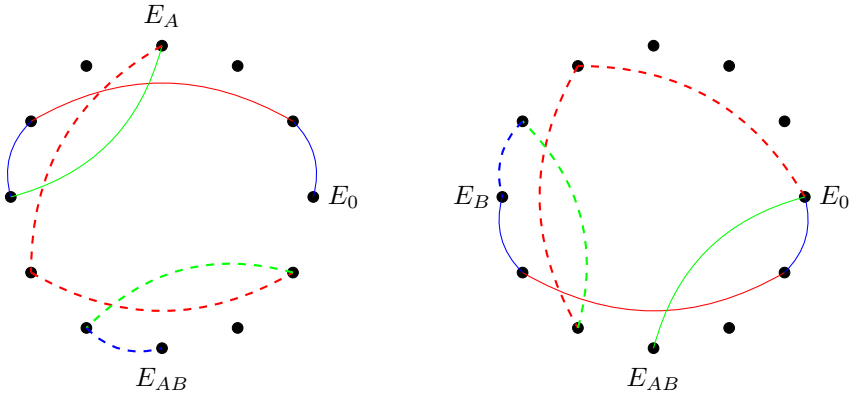

Figure 16: Example of key exchange on the isogeny graph of Figure 10. Alice's path is represented by continuous lines, Bob's path by dashed lines. On the left, Bob computes the shared secret starting from Alice's public data. On the right, Alice does the analogous computation.

a not necessarily uniform distribution. While this example instance is of no practical interest, its instantiation using a Schreier graph of the HHS $\mathrm{Ell}_q(\mathcal{O})$ yields a usable variant of Couveignes' key exchange. We fix a set S of small norm representatives of ideal classes of $\mathrm{Cl}(\mathcal{O})$, corresponding to small degree isogenies between curves in $\mathrm{Ell}_q(\mathcal{O})$. Instead of uniformly sampling secrets from $\mathrm{Cl}(\mathcal{O})$, we sample non-backtracking random walks in the Schreier graph of $(\mathrm{Cl}(\mathcal{O}), S, \mathrm{Ell}_q(\mathcal{O}))$, and exchange j-invariants as public data. The walks can be computed efficiently as a composition of small degree isogenies, and, assuming the graph is an expander and the walks are long enough, they approach the uniform distribution on $\mathrm{Ell}_q(\mathcal{O})$. The protocol is illustrated in Figure 16.

CSIDH. Even with these adjustments, the protocol is far from practical: Stolbunov managed to run a 108 bit secure implementation in around 5 minutes [69]. To understand why, let's see how a random element of $\mathrm{Cl}(\mathcal{O})$ is sampled and the group action evaluated. We have a set S of prime ideals of $\mathcal{O}$, represented as $(\pi - \lambda, \ell)$ for some Frobenius eigenvalue λ modulo a prime ℓ. A secret key corresponds to a product of ideals in S:

$$\mathfrak{s} = \prod_{\mathfrak{a}_i \in S} \mathfrak{a}_i^{e_i}. \tag{4}$$

For simplicity, we may assume that the exponents e_i are taken in a box $[-B, B]$,[7] then the size of the key space is at most $(2B+1)^{\#S}$.

On the other hand, evaluating the action of $\mathfrak{s}$ requires computing at most $\#S \cdot B$ isogenies. We see that, for a fixed set S, increasing B only increases the key space polynomially, while it also increases the running time linearly. On the other hand, for a fixed B, increasing $\#S$ exponentially increases the key space, while it only increases the running time linearly. Thus, to strike a balance between security and running time, we need to use a fairly large set S: values in the hundreds are typical for $\#S$, and all ideals in S must have different (prime) norms to avoid duplicates. Hence, evaluating the action of $\mathfrak{s}$ implies computing up to $\#S \cdot B$ isogenies of degrees as large as a few thousands!

What algorithms do we have at our disposal to compute these isogenies? We have a curve E, a prime ℓ and a *direction* $\pi - \lambda$. Without further assumptions, we have an instance of the Problem 2

[7] Negative values represent the dual direction to $(\pi - \lambda, \ell)$, associated to the ideal $(\pi - \mu, \ell)$.

Public parameters	An elliptic curve E over a finite field $\mathbb{F}_q$, D_π, the discriminant of the Frobenius endomorphism of E, A set of primes $L = \{\ell_1, \ldots, \ell_m\}$ such that $\left(\frac{D_\pi}{\ell_i}\right) = 1$, A *Frobenius eigenvalue* λ_i for each ℓ_i,	
	Alice	**Bob**
Pick random secret	$\rho_A \in L^*$	$\rho_B \in L^*$
Compute public data	$E_A = \rho_A(E)$	$E_B = \rho_B(E)$
Exchange data	$E_A \longrightarrow$	$\longleftarrow E_B$
Compute shared secret	$E_{AB} = \rho_A(E_B)$	$E_{AB} = \rho_B(E_A)$

Figure 17: CSIDH protocol, based on non-backtracking random walks in a supersingular CM graph.

above: we want to enumerate the isogenies of degree ℓ, and choose the one that is horizontal of direction $\pi - \lambda$. We are thus stuck with Elkies' or Couveignes' algorithm, both requiring $O(\ell^3)$ operations to find the isogeny. It is no surprise then that evaluating one $\mathrm{Cl}(\mathcal{O})$-action takes several minutes.

Is it possible to do better? A first is to use Vélu's formulas instead of Elikes' or Couveignes' algorithm, as first proposed in [23]. Suppose, for example, that $\pi|E[\ell]$ acts like $\left(\begin{smallmatrix} 1 & 0 \\ 0 & \mu \end{smallmatrix}\right)$, with $\mu \neq 1$. In this case, there is an easily recognizable direction associated to the eigenvalue 1: the corresponding eigenspace is the cyclic group of rational ℓ-torsion points. A point in this eigenspace can be computed by taking a random point in $E(\mathbb{F}_q)$, and multiplying it by $\#E/\ell$: there is a $(\ell - 1)/\ell$ chance that the result is not zero, and can thus be used to compute the ℓ-isogeny of direction $\pi - 1$ using Vélu's formulas.

We can do even better. Suppose that $\pi|E[\ell]$ acts like $\left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$, then both directions are recognizable: $\pi - 1$ is obtained like before, while $\pi + 1$ corresponds to the rational ℓ-torsion subgroup of a *quadratic twist*[8] of E. This constraint on π forces two conditions:

1. $q = -1 \mod \ell$,
2. $t = 0 \mod \ell$,

and this for each of the primes ℓ we want to include in the set S.

The first condition is easy to fulfill: choose a prime $q = f \cdot \prod_i \ell_i - 1$ for some cofactor f. The second one is much harder, because it essentially requires to find a curve $E/\mathbb{F}_q$ with a specific trace t. If we want E to be an ordinary curve, the best technique at our disposal consists in taking random curves $E/\mathbb{F}_q$ and computing $\#E$, until a suitable one is found.

However, if we enforce the constraint for enough primes ℓ (it is enough that $\prod \ell > 2\sqrt{q}$), then we effectively force t to be 0, and thus E to be supersingular. This was the main insight that recently lead to the key exchange protocol CSIDH [9], which is an efficient variant of the Couveignes–Rostovtsev–Stolbunov system in a supersingular CM graph.

CSIDH uses a prime p of the form $4 \cdot \prod_i \ell_i - 1$, and a supersingular curve $E/\mathbb{F}_p$ as starting point, so that $\pi|E[\ell_i] = \left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$ for all ℓ_i. By cleverly optimizing computations, CSIDH achieves a key-exchange at the 128 bits security level in only 0.1 seconds. The scheme is summarized in Figure 17.

In Section 17 we are going to present a different key exchange protocol, named SIDH, based on supersingular isogeny graphs. The graph structure will be radically different, but Vélu's formulas will still play a crucial role for its performance.

[8] A *quadratic twist* is a curve isomorphic to E over $\mathbb{F}_{q^2}$, hence it represents the same point in the isogeny graph.

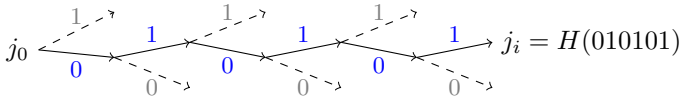

Figure 18: Hashing the string 010101 using an expander graph

16 Hash functions from Ramanujan graphs

Before presenting SIDH, we introduce a distant ancestor: a cryptographic hash function based on the Ramanujan 2-isogeny graph of supersingular elliptic curves. The mixing properties of expander graphs make them very good pseudo-random generators. For the very same reason, they can also be used to define *hash functions*. The Charles-Goren-Lauter (CGL) construction [10] chooses an arbitrary start vertex j_0 in an expander graph, then takes a non-backtracking random walk according to the string to be hashed, and outputs the arrival vertex. To fix notation, let's assume that the graph is 3-regular, then the value to be hashed is encoded as a binary string. At each step one bit is read from the string, and its value is used to choose an edge from the current vertex to the next one, avoiding the one edge that goes back. The way an edge is chosen according to the read bit need only be deterministic, but can be otherwise arbitrary (e.g., determined by some lexicographic ordering). The process is pictured in Figure 18.

For the process to be a good pseudo-random function, the walks need to be longer than the mixing length of the graph. However this is not enough to guarantee a *cryptographically strong* hash function. Indeed the two main properties of cryptographic hash functions, translate in this setting as the following computational problems.

Problem 5 (Preimage resistance)**.** Given a vertex j in the graph, find a path from the start vertex j_0 to j.

Problem 6 (Collision resistance)**.** Find a non-trivial loop (i.e., one that does not track backwards) from j_0 to itself.

Charles, Goren and Lauter suggested two types of expander graphs to be used in their constructions. One is based on LPS Cayley graphs, and was broken shortly afterwards [74, 58]. The second one is based on graphs of supersingular curves. In this context, the preimage finding problem is an instance of the isogeny path problem, while the collision finding problem is equivalent to computing a non-trivial endomorphism of the start curve j_0. In this sense, the CGL hash function on expander graphs has *provable security*, meaning that its cryptographic strength can be provably reduced to well defined mathematical problems thought to be hard.

Nevertheless, the CGL hash function has failed to attract the interest of practitioners. For one, it is considerably slower than popular hash functions such as those standardized by NIST. More concerning is the fact that, while no preimage attack is known, there is no known choice for the starting vertex that does not lead to a collision attack [43, 57, 26]. Nevertheless, the CGL hash function is a fundamental building block in isogeny basic cryptography.

17 Key exchange from supersingular graphs

With CSIDH, we have seen how supersingular curves allowed us to go from a dramatically slow protocol to a fairly efficient one. The upshot is the following: we can control the group structure of supersingular curves simply by controlling the order of the base field $\mathbb{F}_q$; this lets us choose curves with many rational points of small order, which in turn can be used to construct small

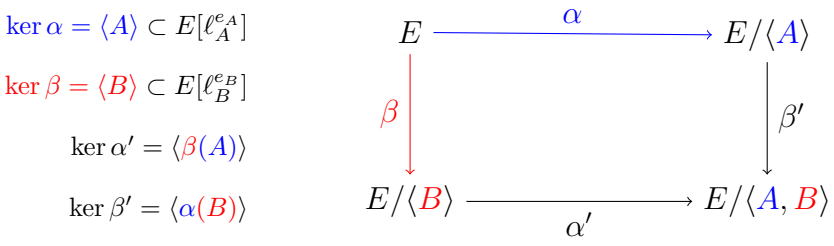

Figure 19: Commutative isogeny diagram constructed from Alice's and Bob's secrets. Quantities known to Alice are drawn in blue, those known to Bob are drawn in red.

degree isogenies via Vélu's formulas. Ultimately, specially crafted supersingular curves let us navigate their isogeny graph very efficiently.

Can we apply the same principle to the Ramanujan graphs of Theorem 60? This is the idea behind SIDH (Supersingular Isogeny Diffie–Hellman) [35, 22], historically the first practical isogeny based key exchange protocol.

SIDH uses supersingular curves $E/\mathbb{F}_{p^2}$ with trace $\pm 2p$, for a specially chosen p.[9] For these curves $\pi|E[\ell] = \pm \left(\begin{smallmatrix} p & 0 \\ 0 & p \end{smallmatrix}\right)$ for any $\ell \neq p$, and there are exactly $\ell + 1$ isogenies of degree ℓ.

Compared to the complex multiplication case, graphs of supersingular isogenies have two attractive features. First, one isogeny degree is enough to obtain an expander graph: this allows us to use isogenies of a single small prime degree, e.g., 2 or 3, instead of many small prime degrees up to the thousands. Second, there is no action of an abelian group, such as $\mathrm{Cl}(\mathcal{O})$, on them: we will see in the next section how this thwarts attacks by quantum computers.

The key idea of SIDH is to let Alice and Bob take random walks in two distinct ℓ-isogeny graphs on the same vertex set of all supersingular j-invariants defined over $\mathbb{F}_{p^2}$. We will denote by ℓ_A and ℓ_B the isogeny degrees used by Alice and Bob respectively. Figure 11 shows a toy example of such graphs, where $p = 97$, $\ell_A = 2$ and $\ell_B = 3$.

Like in CSIDH, we want to be able to evaluate ℓ-isogenies using Vélu's formulas, thus we need $p = \pm 1 \pmod{\ell}$. However, this is not enough to define a key exchange protocol, as we shall see. Instead, we will use Vélu's formulas to evaluate an isogeny of degree ℓ^e, for some large exponent e, all at once. Therefore, we select a prime of the form $p \mp 1 = \ell_A^{e_A} \ell_B^{e_B} f$, where e_A and e_B are exponents to be determined and f is a small cofactor, so that $E/\mathbb{F}_{p^2}$ contains the full subgroups $E[\ell_A^{e_A}]$ and $E[\ell_B^{e_B}]$. Typical values are $p = 2^{216}3^{137} - 1$, $p = 2^{250}3^{159} - 1$ or $p = 2^{372}3^{239} - 1$ (see [3]).

The protocol now proceeds similarly to the Couveignes–Rostovtsev–Stolbunov key exchange: Alice chooses a secret walk of length e_A in the ℓ_A-isogeny graph; this is equivalent to her choosing a secret cyclic subgroup $\langle A \rangle \subset E[\ell_A^{e_A}]$. Bob does the same in the ℓ_B-isogeny graph, choosing a secret $\langle B \rangle \subset E[\ell_B^{e_B}]$. Then, there is a well defined subgroup $\langle A \rangle + \langle B \rangle = \langle A, B \rangle$, defining an isogeny to $E/\langle A, B \rangle$. Since we have taken care to choose $\ell_A \neq \ell_B$, the group $\langle A, B \rangle$ is cyclic of order $\ell_A^{e_A} \ell_B^{e_B}$. This is illustrated in Figure 19.

After Alice and Bob have computed their respective secrets $\langle A \rangle$ and $\langle B \rangle$, we need them to exchange enough information to both compute $E/\langle A, B \rangle$ (up to isomorphism). However, publishing $E/\langle A \rangle$ and $E/\langle B \rangle$ does not give enough information to the other party, and the diagram in Figure 19 shows no way by which they could compute $E/\langle A, B \rangle$ without revealing their secrets.

We solve this problem by a very peculiar trick, which sets SIDH apart from other isogeny

[9] Note that this case includes trace zero curves $E/\mathbb{F}_p$, after extending scalars to $\mathbb{F}_{p^2}$.

Public parameters	Primes ℓ_A, ℓ_B, and a prime $p = \ell_A^{e_A}\ell_B^{e_B} f \mp 1$, A supersingular elliptic curve E over $\mathbb{F}_{p^2}$ of order $(p \pm 1)^2$, A basis $\langle P_A, Q_A\rangle$ of $E[\ell_A^{e_A}]$, A basis $\langle P_B, Q_B\rangle$ of $E[\ell_B^{e_B}]$,	
	Alice	**Bob**
Pick random secret	$A = [m_A]P_A + [n_A]Q_A$	$B = [m_B]P_B + [n_B]Q_B$
Compute secret isogeny	$\alpha : E \to E_A = E/\langle A\rangle$	$\beta : E \to E_B = E/\langle B\rangle$
Exchange data	$E_A, \alpha(P_B), \alpha(Q_B) \longrightarrow$	$\longleftarrow E_B, \beta(P_A), \beta(Q_A)$
Compute shared secret	$E/\langle A, B\rangle = E_B/\langle\beta(A)\rangle$	$E/\langle A, B\rangle = E_A/\langle\alpha(B)\rangle$

Figure 20: Supersingular Isogeny Diffie-Hellman key exchange protocol.

based protocols. The idea is to let Alice and Bob publish some additional information to help each other compute the shared secret. Let us summarize what are the quantities known to Alice and Bob. To set up the cryptosystem, they have publicly agreed on a prime p and a supersingular curve E such that

$$E(\mathbb{F}_{p^2}) \simeq (\mathbb{Z}/\ell_A^{e_A}\mathbb{Z})^2 \oplus (\mathbb{Z}/\ell_B^{e_B}\mathbb{Z})^2 \oplus (\mathbb{Z}/f\mathbb{Z})^2.$$

It will be convenient to also fix public bases of their respective torsion groups:

$$E[\ell_A^{e_A}] = \langle P_A, Q_A\rangle,$$
$$E[\ell_B^{e_B}] = \langle P_B, Q_B\rangle.$$

To start the protocol, they choose random secret subgroups

$$\langle A\rangle = \langle [m_A]P_A + [n_A]Q_A\rangle \subset E[\ell_A^{e_A}],$$
$$\langle B\rangle = \langle [m_B]P_B + [n_B]Q_B\rangle \subset E[\ell_B^{e_B}],$$

of respective orders $\ell_A^{e_A}, \ell_B^{e_B}$, and compute the secret isogenies

$$\alpha : E \to E/\langle A\rangle,$$
$$\beta : E \to E/\langle B\rangle.$$

They respectively publish $E_A = E/\langle A\rangle$ and $E_B = E/\langle B\rangle$.

Now, to compute the shared secret $E/\langle A, B\rangle$, Alice needs to compute the isogeny $\alpha' : E/\langle B\rangle \to E/\langle A, B\rangle$, whose kernel is generated by $\beta(A)$. We see that the kernel of α' depends on both secrets, thus Alice cannot compute it without Bob's assistance. The trick here is for Bob to publish the values $\beta(P_A)$ and $\beta(Q_A)$: they do not require the knowledge of Alice's secret, and we will assume that they do not give any advantage in computing $E/\langle A, B\rangle$ to an attacker. From Bob's published values, Alice can compute $\beta(A)$ as $[m_A]\beta(P_A) + [n_A]\beta(Q_A)$, and complete the protocol. Bob performs the analogous computation, with the help of Alice. The protocol is summarized in Figure 20, and schematized in Figure 21.

18 Security and quantum computers

We end this part with a quick review of the security of the key exchange protocols presented so far. The problem that is often cited as the cornerstone of isogeny based cryptography is the isogeny path problem: given isogenous curves E, E', find an isogeny of smooth degree between

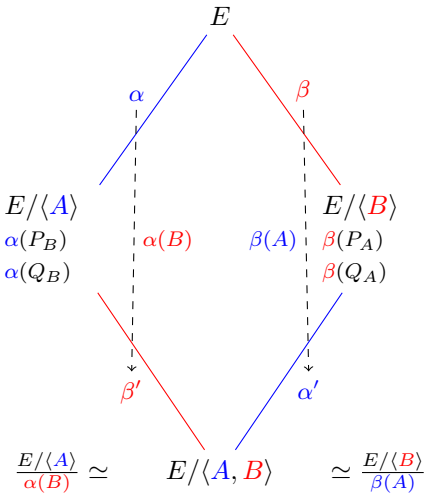

Figure 21: Schematics of SIDH key exchange. Quantities only known to Alice are drawn in blue, quantities only known to Bob in red.

them; however, each cryptosystem depends on a specific variant of the problem. Naturally, the first parameter to look at is the size of the isogeny class of E, E': too small, and we can find the isogeny by brute force.

Security of CM constructions. In the CM case, using the theory of isogeny volcanoes, we can easily find ascending paths from E and E' to two curves $\hat{E}, \hat{E}'$ with complex multiplication by the maximal order; then, we are left with the problem of finding a horizontal isogeny between $\hat{E}$ and $\hat{E}'$. Since the horizontal isogeny class of $\mathcal{O}_K$ is the smallest among all horizontal isogeny classes of curves with complex multiplication by some $\mathcal{O} \subset \mathcal{O}_K$, it makes sense to reduce to this case, as first noted by Galbraith, Hess and Smart [31, 32].

Problem 7 (Horizontal isogeny path problem)**.** Let $\mathbb{F}_q$ be a finite field, and let $\mathcal{O}_K$ be the ring of integers of a quadratic imaginary field $K = \mathbb{Q}(\sqrt{-D})$. Given two elliptic curves E, E' defined over $\mathbb{F}_q$ with complex multiplication by $\mathcal{O}_K$, find an isogeny $E \to E'$ of smooth degree.

The size of the horizontal isogeny class is $h(\mathcal{O}_K)$; it is known by the class number formula that this is in $O(\sqrt{\Delta_K} \log \Delta_K)$, and, for the typical isogeny class[10], $\Delta_K = O(q)$. The best generic attack against the Horizontal isogeny path problem is a Pollard-rho style algorithm, performing random walks from E and E' until a collision is found [31]. Its average complexity is $O(\sqrt{h(\mathcal{O}_K)})$, thus $O(q^{1/4})$ for a typical isogeny class. This justifies choosing a prime q of $4n$ bits, for a security level of 2^n, and this is indeed and what CSIDH does [9].

However, we must also ensure that the key space covers the whole $\mathrm{Ell}_q(\mathcal{O}_K)$, possibly approaching the uniform distribution. This means that isogeny walks, as in Eq. (4), must be sampled from a relatively large subset $S \subset \mathrm{Cl}(\mathcal{O}_K)$, implying that $\#S \gg \log q$. For efficiency reasons, practical instantiations take S just large enough: $\#S \sim (\log q)/2$;[11] however it will not go unnoticed that this choice is insufficient to apply Theorem 59. We may as well live with it,

[10] Including the isogeny class of trace zero supersingular curves used in CSIDH.
[11] Additional constraints in CSIDH force $\#S$ to grow as $(\log q)/(\log\log q)$.

changing our security assumptions to take into account the biased distributions given by random walks in graphs that are not known to be expanders, as it is done in [23].

Security of SIDH. Things are quite different in SIDH. We know that the supersingular isogeny graph over $\mathbb{F}_{p^2}$ has $\approx p/12$ vertices, thus in general we can find a smooth isogeny between two supersingular curves in $O(\sqrt{p})$ operations using the same kind of random walk algorithm.

However, this is not the best attack against SIDH. To understand why, we need to look at the key space. Recall that the prime in SIDH is chosen of the form $p \pm 1 = \ell_A^{e_A} \ell_B^{e_B} f$. Alice's secrets are uniformly random cyclic subgroups of $E[\ell_A^{e_A}]$; Alice's key space contains thus at most $(\ell_A + 1)\ell^{e_A - 1}$ elements. Similarly, Bob's keyspace contains at most $(\ell_B + 1)\ell^{e_B - 1}$ elements. To balance out the size of the two key spaces, we need $\ell_A^{e_A} \approx \ell_B^{e_B} \approx \sqrt{p}$. Thus, Alice's and Bob's key spaces only cover a tiny fraction of the whole supersingular graph, much less they satisfy the conditions to Theorem 60. An isogeny path in any of the two subgraphs can be found by a meet-in-the middle strategy (also called a *claw finding* algorithm) in only $O(p^{1/4})$ steps.[12]

Hence, like in the CM setting, we need to take $\log p \sim 4n$ for a security of 2^n operations. However SIDH j-invariants are elements of $\mathbb{F}_{p^2}$, thus they will typically be twice as big as j-invariants in CSIDH. Adding to that the fact that SIDH public keys contain more than a j-invariant (see Figure 21), we see that CSIDH consumes considerably less bandwidth than SIDH. This is offset by the fact that SIDH is an order of magnitude faster than CSIDH, owing to the smaller isogeny degrees.

However, we just highlighted a very important point on SIDH: it transmits more information than what we would normally feel comfortable sharing. Indeed, SIDH transmits not only the image curve $E/\langle A \rangle$, but also the image of the basis points P_B, Q_B. This would be enough information to recover the secret isogeny by interpolation techniques, however we do not know how to exploit the fact that the isogeny has smooth degree, and in fact we do not know any algorithm that takes advantage of this auxiliary information.[13]

At any rate, the security of SIDH cannot be founded on the isogeny path problem. Instead, it is necessary to make *ad hoc* assumptions taking into account the information communicated by the protocol, that go under the names of CSSI, SSCDH, SSDDH (see [35, 22] or SIDH [3]).

Quantum security. The discussion on security would not be complete without surveying quantum attacks. Indeed, the main selling point of isogeny-based key exchange protocols is their (conjectured) resistance to quantum algorithms.

Let's start with CM constructions. Couveignes' Hard Homogeneous Spaces setting is scarily similar to the Diffie–Hellman key exchange, which is indeed a special case of it. Shor's algorithm [65] solves the discrete logarithm problem in polynomial time on a quantum computer, and thus breaks the Diffie–Hellman protocol. But is there a variant of Shor's algorithm that also breaks generic HHS constructions?

Definition 62 (Hidden Subgroup Problem (HSP))**.** Let $f : G \to X$ be a function from a group G to a set X. Assume that there is a subgroup $H \subset G$ such that $f(g) = f(g')$ if and only if $g' \in gH$. The function f is said to *hide* the subgroup H, and the *hidden subgroup problem* consists in finding generators for H, given access to f.

[12] A Pollard-rho style of algorithm is not possible in this case, since its complexity would depend on the size of the whole graph. The claw finding algorithm is very memory hungry, and some argue that the RAM model is not appropriate to study its complexity. In a constant memory model, the currently best attack against SIDH is estimated to take $O(p^{3/8})$ steps [1].

[13] See [56] for a distant cousin of SIDH for which it is possible to extract useful information out of the auxiliary data.

It is well known that Kitaev's generalization of Shor's algorithm [40] solves the hidden subgroup problem in quantum polynomial time, when G is a finitely generated abelian group.

Definition 63 (Hidden Shift Problem (HShP))**.** Let $f_0, f_1 : G \to X$ be two injective functions from a group G to a set X. Assume that there is an element $s \in G$ such that $f_0(g) = f_1(gs)$ for any $g \in G$. The element s is called a *hidden shift* for f_0, f_1, and the *hidden shift problem* is to find s, given access to f_0 and f_1.

For any group G, the hidden shift problem reduces to the hidden subgroup problem for the (generalized) dihedral group $G \rtimes C_2$.[14] No generalization of Kitaev's algorithm is known for non-abelian groups, but a different family of algorithms, due to Kuperberg [44, 45] and Regev [61], solves the HShP in subexponential quantum time $\exp(\sqrt{\log \#G})$.

As first noted in [11] and then improved in [7, 5, 36], Kuperberg's algorithm can be used to solve the Horizontal isogeny path problem as follows: let E, E' be the two curves with complex multiplication by $\mathcal{O}_K$, define two functions $f_0, f_1 : \mathrm{Cl}(\mathcal{O}_K) \to \mathrm{Ell}_q(\mathcal{O}_K)$ as $f_0(\mathfrak{a}) = \mathfrak{a} \cdot E$ and $f_1(\mathfrak{a}) = \mathfrak{a} \cdot E'$, then the hidden shift defines a horizontal isogeny between E and E'.

Kuperberg's algorithm is a game changer for protocols based on complex multiplication: indeed, to ensure 2^n quantum security we need to take $\log q = O(n^2)$. The actual constant depends on the variant of Kuperberg's algorithm, and various parameters such as available quantum memory; its exact value is currently debated, but it appears that taking $\log q$ somewhere between 512 and 1024 bits grants a security of 2^{64} quantum gates [9, 23, 7, 4].

For SIDH, on the other hand, there is no group structure[15] that can be exploited by Kuperberg's algorithm. Currently, the best quantum attack against SIDH is a Grover-like claw finding algorithm due to Tani [71], requiring $O(p^{1/6})$ quantum gates (and as many qubits!). For this reason, p is typically chosen so that $\log p \sim 6n$ for a quantum security of 2^n gates, although it is debated whether Tani's algorithm actually presents an advantage over the classical claw finding attack.

[14] To reduce HShP to HSP, simply define the function f by $f(g, 1) = f_0(g)$ and $f(g, -1) = f_1(g)$, so that the hidden subgroup is generated by $(s, -1)$.

[15] Outside of the subgraph of $\mathbb{F}_p$-rational curves, but the algorithm in [6] does not impact the security of SIDH.

References

[1] Gora Adj, Daniel Cervantes-Vázquez, Jesús-Javier Chi-Domínguez, Alfred Menezes, and Francisco Rodríguez-Henríquez. On the cost of computing isogenies between supersingular elliptic curves. Cryptology ePrint Archive, Report 2018/313, 2018.

[2] Arthur O. L. Atkin. The number of points on an elliptic curve modulo a prime. Manuscript, Chicago IL, 1991.

[3] Reza Azarderakhsh, Brian Koziel, Matt Campagna, Brian LaMacchia, Craig Costello, Patrick Longa, Luca De Feo, Michael Naehrig, Basil Hess, Joost Renes, Amir Jalali, Vladimir Soukharev, David Jao, and David Urbanik. Supersingular isogeny key encapsulation, 2017.

[4] Daniel J. Bernstein, Tanja Lange, Chloe Martindale, and Lorenz Panny. Quantum circuits for the CSIDH: optimizing quantum evaluation of isogenies. Cryptology ePrint Archive, Report 2018/1059, 2018.

[5] Jean-François Biasse, Annamaria Iezzi, and Michael J. Jr. Jacobson. A note on the security of CSIDH. In *Kangacrypt 2018*, 2018.

[6] Jean-François Biasse, David Jao, and Anirudh Sankar. A quantum algorithm for computing isogenies between supersingular elliptic curves. In *International Conference in Cryptology in India*, pages 428–442. Springer, 2014.

[7] Xavier Bonnetain and André Schrottenloher. Quantum security analysis of CSIDH and ordinary isogeny-based schemes. Cryptology ePrint Archive, Report 2018/537, 2018.

[8] Alin Bostan, François Morain, Bruno Salvy, and Éric Schost. Fast algorithms for computing isogenies between elliptic curves. *Mathematics of Computation*, 77(263):1755–1778, September 2008.

[9] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. In Thomas Peyrin and Steven D. Galbraith, editors, *Advances in Cryptology – ASIACRYPT 2018*, pages 395–427. Springer International Publishing, 2018.

[10] Denis X. Charles, Eyal Z. Goren, and Kristin E. Lauter. Cryptographic hash functions from expander graphs. *Journal of Cryptology*, 22(1):93–113, January 2009.

[11] Andrew Childs, David Jao, and Vladimir Soukharev. Constructing elliptic curve isogenies in quantum subexponential time. *Journal of Mathematical Cryptology*, 8(1):1–29, 2014.

[12] Fan R.K. Chung. Diameters and eigenvalues. *Journal of the American Mathematical Society*, 2(2):187–196, 1989.

[13] Anamaria Costache, Brooke Feigon, Kristin Lauter, Maike Massierer, and Anna Puskas. Ramanujan graphs in cryptography. Cryptology ePrint Archive, Report 2018/593, 2018.

[14] Jean-Marc Couveignes. *Quelques calculs en théorie des nombres*. PhD thesis, Université de Bordeaux, 1994.

[15] Jean-Marc Couveignes. Computing ℓ-isogenies using the p-torsion. In *ANTS-II: Proceedings of the Second International Symposium on Algorithmic Number Theory*, pages 59–65, London, UK, 1996. Springer-Verlag.

[16] Jean-Marc Couveignes. Isomorphisms between Artin-Schreier towers. *Mathematics of Computation*, 69(232):1625–1631, 2000.

[17] Jean-Marc Couveignes. Hard homogeneous spaces. Cryptology ePrint Archive, Report 2006/291, 2006.

[18] Jean-Marc Couveignes and Reynald Lercier. Fast construction of irreducible polynomials over finite fields. *Israel Journal of Mathematics*, 194(1):77–105, 2013.

[19] Luca De Feo. *Algorithmes Rapides pour les Tours de Corps Finis et les Isogénies*. PhD thesis, Ecole Polytechnique X, December 2010.

[20] Luca De Feo, Javad Doliskani, and Éric Schost. Fast algorithms for ℓ-adic towers over finite fields. In *Proceedings of the 38th International Symposium on Symbolic and Algebraic Computation*, ISSAC '13, pages 165–172, New York, NY, USA, 2013. ACM.

[21] Luca De Feo, Cyril Hugounenq, Jérôme Plût, and Éric Schost. Explicit isogenies in quadratic time in any characteristic. *LMS Journal of Computation and Mathematics*, 19(A):267–282, 2016.

[22] Luca De Feo, David Jao, and Jérôme Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. *Journal of Mathematical Cryptology*, 8(3):209–247, 2014.

[23] Luca De Feo, Jean Kieffer, and Benjamin Smith. Towards practical key exchange from ordinary isogeny graphs. In Thomas Peyrin and Steven D. Galbraith, editors, *Advances in Cryptology – ASIACRYPT 2018*, pages 365–394. Springer International Publishing, 2018.

[24] Luca De Feo and Éric Schost. Fast arithmetics in Artin-Schreier towers over finite fields. In *ISSAC '09: Proceedings of the 2009 international symposium on Symbolic and algebraic computation*, pages 127–134, New York, NY, USA, 2009. ACM.

[25] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, 1976.

[26] Kirsten Eisenträger, Sean Hallgren, Kristin Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: Reductions and solutions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018*, pages 329–368. Springer International Publishing, 2018.

[27] Noam D. Elkies. Explicit isogenies. 1992.

[28] Noam D. Elkies. Elliptic and modular curves over finite fields and related computational issues. In *Computational perspectives on number theory (Chicago, IL, 1995)*, volume 7 of *Studies in Advanced Mathematics*, pages 21–76, Providence, RI, 1998. AMS International Press.

[29] Mireille Fouquet and François Morain. Isogeny volcanoes and the SEA algorithm. In Claus Fieker and David R. Kohel, editors, *Algorithmic Number Theory Symposium*, volume 2369 of *Lecture Notes in Computer Science*, pages 47–62, Berlin, Heidelberg, 2002. Springer Berlin / Heidelberg.

[30] Steven D. Galbraith. Constructing isogenies between elliptic curves over finite fields. *LMS Journal of Computation and Mathematics*, 2:118–138, 1999.

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

[32] Steven D. Galbraith and Anton Stolbunov. Improved algorithm for the isogeny problem for ordinary elliptic curves. *Applicable Algebra in Engineering, Communication and Computing*, 24(2):107–131, June 2013.

[33] Pierrick Gaudry, Florian Hess, and Niegel Smart. Constructive and destructive facets of Weil descent on elliptic curves. *Journal of Cryptology*, 15(1):19–46–46, March 2002.

[34] Oded Goldreich. *Basic Facts about Expander Graphs*, pages 451–464. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

[35] David Jao and Luca De Feo. Towards Quantum-Resistant cryptosystems from supersingular elliptic curve isogenies. In Bo-Yin Yang, editor, *Post-Quantum Cryptography*, volume 7071 of *Lecture Notes in Computer Science*, pages 19–34, Berlin, Heidelberg, 2011. Springer Berlin / Heidelberg.

[36] David Jao, Jason LeGrow, Christopher Leonardi, and Luiz Ruiz-Lopez. A polynomial quantum space attack on CRS and CSIDH. In *MathCrypt 2018*, 2018. To appear.

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

[38] David Jao and Vladimir Soukharev. A subexponential algorithm for evaluating large degree isogenies. In *ANTS IX: Proceedings of the Algorithmic Number Theory 9th International Symposium*, volume 6197 of *Lecture Notes in Computer Science*, pages 219–233, Berlin, Heidelberg, 2010. Springer.

[39] Antoine Joux. *Algorithmic cryptanalysis*. CRC Press, 2009.

[40] Alexey Yuri Kitaev. Quantum measurements and the abelian stabilizer problem. *arXiv preprint quant-ph/9511026*, 1995.

[41] Neal Koblitz. Elliptic curve cryptosystems. *Mathematics of Computation*, 48(177):203–209, 1987.

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

[43] David R. Kohel, Kristin Lauter, Christophe Petit, and Jean-Pierre Tignol. On the quaternion-isogeny path problem. *LMS Journal of Computation and Mathematics*, 17(A):418–432, 2014.

[44] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM Journal of Computing*, 35(1):170–188, 2005.

[45] Greg Kuperberg. Another Subexponential-time Quantum Algorithm for the Dihedral Hidden Subgroup Problem. In Simone Severini and Fernando Brandao, editors, *8th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2013)*, volume 22 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20–34, Dagstuhl, Germany, 2013. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[46] Serge Lang. *Elliptic Functions*, volume 112 of *Graduate texts in mathematics*. Springer, 1987.

[47] Serge Lang. *Algebraic number theory*, volume 110 of *Graduate Texts in Mathematics*. Springer-Verlag, 1994.

[48] Hendrik W. Lenstra. Factoring integers with elliptic curves. *Annals of Mathematics*, 126:649–673, 1987.

[49] Reynald Lercier. *Algorithmique des courbes elliptiques dans les corps finis*. PhD thesis, LIX - CNRS, June 1997.

[50] Reynald Lercier and Thomas Sirvent. On Elkies subgroups of ℓ-torsion points in elliptic curves defined over a finite field. *Journal de théorie des nombres de Bordeaux*, 20(3):783–797, 2008.

[51] Alexander Lubotzky. *Discrete groups, expanding graphs and invariant measures*, volume 125 of *Progress in Mathematics*. Birkhäuser Verlag, Basel, 1994.

[52] Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8(3), 1988.

[53] Jean-François Mestre. La méthode des graphes. Exemples et applications. In *Proceedings of the international conference on class numbers and fundamental units of algebraic number fields (Katata, 1986)*, Nagoya, 1986. Nagoya University.

[54] Victor S. Miller. Use of elliptic curves in cryptography. In *Lecture notes in computer sciences; 218 on Advances in cryptology–CRYPTO 85*, pages 417–426, New York, NY, USA, 1986. Springer-Verlag New York, Inc.

[55] Jürgen Neukirch. *Algebraic number theory*, volume 322. Springer Verlag, 1999.

[56] Christophe Petit. Faster algorithms for isogeny problems using torsion point images. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology – ASIACRYPT 2017*, pages 330–353. Springer International Publishing, 2017.

[57] Christophe Petit and Kristin Lauter. Hard and easy problems for supersingular isogeny graphs. Cryptology ePrint Archive, Report 2017/962, 2017.

[58] Christophe Petit, Kristin Lauter, and Jean-Jacques Quisquater. Full cryptanalysis of LPS and Morgenstern hash functions. In *Proceedings of the 6th international conference on Security and Cryptography for Networks*, SCN '08, Berlin, Heidelberg, 2008. Springer-Verlag.

[59] Arnold K. Pizer. Ramanujan graphs and Hecke operators. *Bulletin of the American Mathematical Society (N.S.)*, 23(1), 1990.

[60] Arnold K. Pizer. Ramanujan graphs. In *Computational perspectives on number theory (Chicago, IL, 1995)*, volume 7 of *AMS/IP Stud. Adv. Math.* Amer. Math. Soc., Providence, RI, 1998.

[61] Oded Regev. A subexponential time algorithm for the dihedral hidden subgroup problem with polynomial space. arXiv:quant-ph/0406151, June 2004.

[62] Alexander Rostovtsev and Anton Stolbunov. Public-key cryptosystem based on isogenies. Cryptology ePrint Archive, Report 2006/145, April 2006.

[63] René Schoof. Elliptic curves over finite fields and the computation of square roots mod p. *Mathematics of Computation*, 44(170):483–494, 1985.

[64] René Schoof. Counting points on elliptic curves over finite fields. *Journal de Théorie des Nombres de Bordeaux*, 7(1):219–254, 1995.

[65] Peter W Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on*, pages 124–134. IEEE, 1994.

[66] Joseph H. Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1992.

[67] Joseph H. Silverman. *Advanced Topics in the Arithmetic of Elliptic Curves*, volume 151 of *Graduate Texts in Mathematics*. Springer, January 1994.

[68] Anton Stolbunov. Constructing public-key cryptographic schemes based on class group action on a set of isogenous elliptic curves. *Advances in Mathematics of Communications*, 4(2), 2010.

[69] Anton Stolbunov. Cryptographic schemes based on isogenies, 2012.

[70] Andrew V. Sutherland. Genus 1 point counting over prime fields. Last accessed July 16, 2010. http://www-math.mit.edu/~drew/SEArecords.html, 2010.

[71] Seiichiro Tani. Claw finding algorithms using quantum walk. *Theoretical Computer Science*, 410(50):5285–5297, 2009.

[72] Terence Tao. Expansion in groups of Lie type – basic theory of expander graphs, 2011.

[73] Edlyn Teske. An elliptic curve trapdoor system. *Journal of Cryptology*, 19(1):115–133, January 2006.

[74] Jean-Pierre Tillich and Gilles Zémor. Collisions for the LPS expander graph hash function. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 254–269. Springer, 2008.

[75] Jacques Vélu. Isogénies entre courbes elliptiques. *Comptes Rendus de l'Académie des Sciences de Paris*, 273:238–241, 1971.

[76] John Voight. Quaternion algebras, 2018.

Part IV
Other applications

This material used to be part of the first version of these lecture notes, but we decided to discard it from the main body to focus on the more central topics.

We keep it in this appendix for historical reference, however we do not guarantee its coherence with the main material.

A Application: Elliptic curve factoring method

A second popular use of elliptic curves in technology is for factoring large integers, a problem that also occurs frequently in cryptography.

The earliest method for factoring integers was already known to the ancient Greeks: the *sieve of Eratosthenes* finds all primes up to a given bound by crossing composite numbers out in a table. Applying the Eratosthenes' sieve up to $\sqrt{N}$ finds all prime factors of a composite number N. Examples of modern algorithms used for factoring are Pollard's *Rho algorithm* and Coppersmith's *Number Field Sieve (NFS)*.

In the 1980s H. Lenstra [48] introduced an algorithm for factoring that has become known as the *Elliptic Curve Method (ECM)*. Its complexity is between Pollard's and Coppersmith's algorithms in terms of number of operations; at the same time it only requires a constant amount of memory, and is very easy to parallelize. For these reasons, ECM is typically used to factor integers having medium sized prime factors.

From now on we suppose that $N = pq$ is an integer which factorization we wish to compute, where p and q are distinct primes. Without loss of generality, we can suppose that $p < q$.

Lenstra's idea has its roots in an earlier method for factoring special integers, also due to Pollard. Pollard's $(p-1)$ *factoring method* is especially suited for integers $N = pq$ such that $p-1$ only has *small* prime factors. It is based on the isomorphism

$$\begin{aligned} \rho : \mathbb{Z}/N\mathbb{Z} &\to \mathbb{Z}/p\mathbb{Z} \times \mathbb{Z}/q\mathbb{Z}, \\ x &\mapsto (x \bmod p, x \bmod q) \end{aligned}$$

given by the Chinese remainder theorem. The algorithm is detailed in Figure 22a. It works by guessing a multiple e of $p-1$, then taking a random element $x \in (\mathbb{Z}/N\mathbb{Z})^\times$, to deduce a random element y in $\langle 1 \rangle \oplus (\mathbb{Z}/q\mathbb{Z})^\times$. If the guessed exponent e was correct, and if $y \neq 1$, the gcd of $y-1$ with N yields a non-trivial factor.

The $p-1$ method is very effective when the bound B is small, but its complexity grows exponentially with B. For this reason it is only usable when $p-1$ has small prime factors, a constraint that is very unlikely to be satisfied by random primes.

Lenstra's ECM algorithm is a straightforward generalization of the $p-1$ method, where the multiplicative groups $(\mathbb{Z}/p\mathbb{Z})^\times$ and $(\mathbb{Z}/q\mathbb{Z})^\times$ are replaced by the groups of points $E(\mathbb{F}_p)$ and $E(\mathbb{F}_q)$ of an elliptic curve defined over $\mathbb{Q}$. Now, the requirement is that $\#E(\mathbb{F}_p)$ only has small prime factors. This condition is also extremely rare, but now we have the freedom to try the method many times by changing the elliptic curve.

The algorithm is summarized in Figure 22b. It features two remarkable subtleties. First, it would feel natural to pick a random elliptic curve $E : y^2 = x^3 + ax + b$ by picking random a and b, however taking a point on such curve would then require computing a square root modulo N, a problem that is known to be has hard as factoring N. For this reason, the algorithm starts by taking a random point, and then deduces the equation of E from it. Secondly, all computations

Input: An integer $N = pq$,
a bound B on the largest prime factor of $p-1$;
Output: (p, q) or FAIL.

1. Set $e = \prod_{r \text{ prime} < B} r^{\lfloor \log_r \sqrt{N} \rfloor}$;
2. Pick a random $1 < x < N$;
3. Compute $y = x^e \mod N$;
4. Compute $q' = \gcd(y - 1, N)$;
5. **if** $q' \neq 1, N$ **then**
6. **return** $N/q', q'$;
7. **else**
8. **return** FAIL.
9. **end if**

(a) Pollard's $(p-1)$ algorithm

Input: An integer $N = pq$, a bound B;
Output: (p, q) or FAIL.

1. Pick random integers a, X, Y in $[0, N[$;
2. Compute $b = Y^2 - X^3 - aX \mod N$;
3. Define the elliptic curve $E : y^2 = x^3 - ax - b$.
4. Define the point $P = (X : Y : 1) \in E(\mathbb{Z}/N\mathbb{Z})$.
5. Set $e = \prod_{r \text{ prime} < B} r^{\lfloor \log_r \sqrt{N} \rfloor}$;
6. Compute $Q = [e]P = (X' : Y' : Z')$;
7. Compute $q' = \gcd(Z', N)$;
8. **if** $q' \neq 1, N$ **then**
9. **return** $N/q', q'$;
10. **else**
11. **return** FAIL.
12. **end if**

(b) Lenstra's ECM algorithm

Figure 22: The $(p-1)$ and ECM factorization algorithms

on coordinates happen in the projective plane over $\mathbb{Z}/N\mathbb{Z}$; however, properly speaking, projective space cannot be defined over non-integral rings. Implicitly, $E(\mathbb{Z}/N\mathbb{Z})$ is defined as the product group $E(\mathbb{F}_p) \oplus E(F_q)$, and any attempt at inverting a non-invertible in $\mathbb{Z}/N\mathbb{Z}$ will result in a factorization of N.

B Application: point counting

Before going more in depth into the study of the endomorphism ring, let us pause for a while on a simpler problem. Hasse's theorem relates the cardinality of a curve defined over a finite field with the trace of its Frobenius endomorphism. However, it does not give us an algorithm to compute either.

The first efficient algorithm to compute the trace of π was proposed by Schoof in the 1980s [63]. The idea is very simple: compute the value of $t_\pi \bmod \ell$ for many small primes ℓ, and then reconstruct the trace using the Chinese remainder theorem. To compute $t_\pi \bmod \ell$, Schoof's algorithm formally constructs the group $E[\ell]$, takes a generic point $P \in E[\ell]$, and then runs a search for the integer t such that

$$\pi([t]P) = [q]P + \pi^2(P).$$

The formal computation must be carried out by computing modulo a polynomial that vanishes on the whole $E[\ell]$; the smallest such polynomial is provided by the *division polynomial* ψ_ℓ.

Definition 64 (Division polynomial)**.** Let $E : y^2 = x^3 + ax + b$ be an elliptic curve, the *division polynomials* ψ_m are defined by the initial values

$$\begin{aligned}
\psi_1 &= 1, \\
\psi_2 &= 2y, \\
\psi_3 &= 3x^4 + 6ax^2 + 12bx - a^2, \\
\psi_4 &= (2x^6 + 10ax^4 + 40bx^3 - 10a^2x^2 - 8abx - 2a^3 - 16b^2)2y,
\end{aligned}$$

and by the recurrence

$$\begin{aligned}\psi_{2m+1} &= \psi_{m+2}\psi_m^3 - \psi_{m-1}\psi_{m+1}^3 && \text{for } m \geq 2,\\ \psi_2\psi_{2m} &= (\psi_{m+2}\psi_{m-1}^2 - \psi_{m-2}\psi_{m+1}^2)\psi_m && \text{for } m \geq 3.\end{aligned}$$

The m-th division polynomial ψ_m vanishes on $E[m]$; the multiplication-by-m map can be written as

$$[m]P = \left(\frac{\phi_m(P)}{\psi_m(P)^2}, \frac{\omega_m(P)}{\psi_m(P)^3}\right)$$

for any point $P \neq \mathcal{O}$, where ϕ_m and ω_m are defined as

$$\begin{aligned}\phi_m &= x\psi_m^2 - \psi_{m+1}\psi_{m-1},\\ \omega_m &= \psi_{m-1}^2\psi_{m+2} + \psi_{m-2}\psi_{m+1}^2.\end{aligned}$$

Schoof's algorithm runs in time polynomial in $\log \#E(k)$, however it is quite slow in practice. Among the major advances that have enabled the use of elliptic curves in cryptography are the optimizations of Schoof's algorithm due to Atkin and Elkies [2, 27, 64, 28]. Both improvements use a better understanding of the action of π on $E[\ell]$. Assume that ℓ is different from the characteristic, we have already seen that $E[\ell]$ is a group of rank two. Hence, π acts on $E[\ell]$ like a matrix M in $\mathrm{GL}_2(\mathbb{Z}/\ell\mathbb{Z})$, and its characteristic polynomial is exactly

$$\chi(X) = X^2 - t_\pi X + q \mod \ell.$$

Now we have three possibilities:

- χ splits modulo ℓ, as $\chi(X) = (X-\lambda)(X-\mu)$, with $\lambda \neq \mu$; we call this the *Elkies case*.
- χ does not split modulo ℓ; we call this the *Atkin case*;
- χ is a square modulo ℓ.

The SEA algorithm, treats each of these cases in a slightly different way; for simplicity, we will only sketch the Elkies case. In this case, there exists a basis $\langle P, Q\rangle$ for $E[\ell]$ onto which π acts as a matrix $M = \left(\begin{smallmatrix}\lambda & 0\\ 0 & \mu\end{smallmatrix}\right)$. Each of the two eigenspaces of M is the kernel of an isogeny of degree ℓ from E to another curve E'. If we can determine the curve corresponding to, e.g., $\langle P\rangle$, then we can compute the isogeny $\phi : E \to E/\langle P\rangle$, and use it to formally represent the point P. Then, λ is recovered by solving the equation

$$[\lambda]P = \pi(P),$$

and from it we recover $t_\pi = \lambda + q/\lambda \mod \ell$.

Elkies' method is very similar to Schoof's original way of computing t_π, however it is considerably more efficient thanks to the degree of the extension rings involved. Indeed, in Schoof's algorithm a generic point of $E[\ell]$ is represented modulo the division polynomial ψ_ℓ, which has degree $(\ell^2-1)/2$. In Elkies' algorithm, instead, the formal representation of $\langle P\rangle$ only requires working modulo a polynomial of degree $\approx \ell$.

The other cases have similar complexity gains. For a more detailed overview, we address the reader to [64, 49, 28, 70].

C Application: computing irreducible polynomials

In the applications seen in the first part, we have followed an old *mantra*: whenever an algorithm relies solely on the properties of the multiplicative group $\mathbb{F}_q^*$, it can be generalized by replacing $\mathbb{F}_q^*$ with the group of points of an elliptic curve over $\mathbb{F}_q$ (or, eventually, a higher dimensional Abelian variety). Typically, the generalization adds some complexity to the computation, but comes with the advantage of having more freedom in the choice of the group size and structure. We now present another instance of the same *mantra*, that is particularly remarkable in our opinion: to the best of our knowledge, it is the first algorithm where replacing $\mathbb{F}_q^*$ with $E(\mathbb{F}_q)$ required some non-trivial work with isogenies.

Constructing irreducible polynomials of arbitrary degree over a finite field $\mathbb{F}_q$ is a classical problem. A classical solution consists in picking polynomials at random, and applying an irreducibility test, until an irreducible one is found. This solution is not satisfactory for at least two reasons: it is not deterministic, and has average complexity quadratic both in the degree of the polynomial and in $\log q$.

For a few special cases, we have well known irreducible polynomials. For example, when d divides $q-1$, there exist $\alpha \in \mathbb{F}_q$ such that $X^d - \alpha$ is irreducible. Such an α can be computed using Hilbert's theorem 90, or –more pragmatically, and assuming that the factorization of $q-1$ is known– by taking a random element and testing that it has no d-th root in $\mathbb{F}_q$. It is evident that this algorithm relies on the fact that the multiplicative group $\mathbb{F}_q^*$ is cyclic of order $q-1$.

At this point our *mantra* suggests that we replace α with a point $P \in E(\mathbb{F}_q)$ that has no ℓ-divisor in $E(\mathbb{F}_q)$, for some well chosen curve E. The obvious advantage is that we now require $\ell | \#E(\mathbb{F}_q)$, thus we are no longer limited to $\ell | (q-1)$; however, what irreducible polynomial shall we take? Intuition would suggest that we take the polynomial defining the ℓ-divisors of P; however we know that the map $[\ell]$ has degree ℓ^2, thus the resulting polynomial would have degree too large, and it would not even be irreducible.

This idea was first developed by Couveignes and Lercier [18] and then slightly generalized in [20]. Their answer to the question is to decompose the map $[\ell]$ as a composition of isogenies $\hat{\phi} \circ \phi$, and then take the (irreducible) polynomial vanishing on the fiber $\phi^{-1}(P)$.

More precisely, let $\mathbb{F}_q$ be a finite field, and let $\ell \nmid (q-1)$ be odd and such that $\ell \ll q+1+2\sqrt{q}$. Then there exists a curve E which cardinality $\#E(\mathbb{F}_q)$ is divisible by ℓ. The hypothesis $\ell \nmid (q-1)$ guarantees that $G = E[\ell] \cap E(\mathbb{F}_q)$ is cyclic (see Exercise II.5). Let ϕ be the degree ℓ isogeny of kernel G, and let E' be its image curve. Let P be a point in $E'(\mathbb{F}_q) \setminus [\ell]E'(\mathbb{F}_q)$, Couveignes and Lercier show that $\phi^{-1}(P)$ is an *irreducible fiber*, i.e., that the polynomial

$$f(X) = \prod_{Q \in \phi^{-1}(P)} (X - x(Q))$$

is irreducible over $\mathbb{F}_q$.

To effectively compute the polynomial f, we need one last technical ingredient: a way to compute a representation of the isogeny ϕ as a rational function. This is given to us by the famous Vélu's formulas [75].

Proposition 65 (Vélu's formulas)**.** *Let $E : y^2 = x^3 + ax + b$ be an elliptic curve defined over a field k, and let $G \subset E(\bar{k})$ be a finite subgroup. The separable isogeny $\phi : E \to E/G$, of kernel G, can be written as*

$$\phi(P) = \left(x(P) + \sum_{Q \in G \setminus \{\mathcal{O}\}} x(P+Q) - x(Q), y(P) + \sum_{Q \in G \setminus \{\mathcal{O}\}} y(P+Q) - y(Q) \right);$$

Input: A finite field $\mathbb{F}_q$,
a prime power ℓ^e such that $\ell \nmid (q-1)$ and $\ell \ll q$;
Output: An irreducible polynomial of degree ℓ^e.

1. Take random curves E_0, until one with $\ell | \#E_0$ is found;
2. Factor $\#E_0$;
3. **for** $1 \leq i \leq e$ **do**
4. Use Vélu's formulas to compute a degree ℓ isogeny $\phi_i : E_{i-1} \to E_i$;
5. **end for**
6. Take random points $P \in E_i(\mathbb{F}_q)$ until one not in $[\ell]E_i(\mathbb{F}_q)$ is found;
7. **return** The polynomial vanishing on the abscissas of $\phi_i^{-1} \circ \cdots \circ \phi_1^{-1}(P)$.

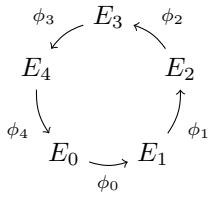

Figure 23: Couveignes-Lercier algorithm to compute irreducible polynomials, and structure of the computed isogeny cycle.

and the curve E/G has equation $y^2 = x^3 + a'x + b'$, where

$$a' = a - 5 \sum_{Q \in G \setminus \{\mathcal{O}\}} (3x(Q)^2 + a),$$
$$b' = b - 7 \sum_{Q \in G \setminus \{\mathcal{O}\}} (5x(Q)^3 + 3ax(Q) + 2b).$$

Proof. See [19, §8.2]. □

Corollary 66. *Let E and G be as above. Let*

$$h(X) = \prod_{Q \in G \setminus \{\mathcal{O}\}} (X - x(Q)).$$

Then the isogeny ϕ can be expressed as

$$\phi(X, Y) = \left(\frac{g(X)}{h(X)}, y \left(\frac{g(x)}{h(x)} \right)' \right),$$

where $g(X)$ is defined by

$$\frac{g(X)}{h(X)} = dX - p_1 - (3X^2 + a)\frac{h'(X)}{h(X)} - 2(X^3 + aX + b)\left(\frac{h'(X)}{h(X)} \right)',$$

with p_1 the trace of $h(X)$ and d its degree.

Proof. See [19, §8.2]. □

The Couveignes-Lercier algorithm is summarized in Figure 23. What is most interesting, is the fact that it can be immediately generalized to computing irreducible polynomials of degree ℓ^e, by iterating the construction. Looking at the specific parameters, it is apparent that ℓ is an *Elkies prime* for E (i.e., $\left(\frac{D}{\ell}\right) = 1$), and that each isogeny ϕ_i is horizontal, thus their composition eventually forms a cycle, the *crater* of a volcano.