

Mathematics of Isogeny Based Cryptography

Luca De Feo
Université de Versailles & Inria Saclay
<http://defeo.lu/>

École mathématique africaine
May 10 – 23, 2017, Thiès, Senegal

Introduction

These lectures notes were written for a summer school on *Mathematics for Post-quantum cryptography* in Thiès, Senegal. They try to provide a guide for Masters' students to get through the vast literature on elliptic curves, without getting lost on their way to learning isogeny-based cryptography. They are by no means a reference text on the theory of elliptic curves, nor on cryptography; students are encouraged to complement these notes with some of the books recommended in the bibliography.

The presentation is divided in three parts, roughly corresponding to the three lectures given. In an effort to keep the reader interested, each part alternates between the fundamental theory of elliptic curves, and applications in cryptography. We often prefer to have the main ideas flow smoothly, rather than having a rigorous presentation as one would have in a more classical book. The reader will excuse us for the inaccuracies and the omissions.

Isogeny Based Cryptography is a very young field, that has only begun in the 2000s. It has its roots in *Elliptic Curve Cryptography* (ECC), a somewhat older branch of public-key cryptography that was started in the 1980s, when Miller and Koblitz first suggested to use elliptic curves inside the Diffie-Hellman key exchange protocol (see Section 4).

ECC only started to gain traction in the 1990s, after Schoof's algorithm made it possible to easily find elliptic curves of large prime order. It is nowadays a staple in public-key cryptography. The 2000s have seen two major innovations in ECC: the rise of *Pairing Based Cryptography* (PBC), epitomized by Joux' one-round tripartite Diffie-Hellman key exchange, and the advent of Isogeny-based cryptography, initiated by the works of Teske and Rostovtsev & Stolbunov. While PBC has attracted most of the attention during the first decade, thanks to its revolutionary applications, isogeny based cryptography has stayed mostly discrete during this time. It is only in the second half of the 2010 that the attention has partly shifted to isogenies. The main reason for this is the sudden realization by the cryptographic community of the very possibly near arrival of a *general purpose quantum computer*. While the capabilities of such futuristic machine would render all ECC and PBC suddenly worthless, isogeny based cryptography seems to resist much better to the cryptanalytic powers of the quantum computer.

In these notes, after a review of the general theory of elliptic curves and isogenies, we will present the most important isogeny-based systems, and their cryptographic properties.

L^AT_EX source code available at <https://github.com/defeo/ema2017/>.

This work is licensed under a [Creative Commons “Attribution-NonCommercial 4.0 International”](https://creativecommons.org/licenses/by-nc/4.0/) license.



Contents

I	Elliptic curves and cryptography	3
1	Elliptic curves	3
2	Maps between elliptic curves	5
3	Elliptic curves over finite fields	6
4	Application: Diffie-Hellman key exchange	7
5	Application: Elliptic curve factoring method	8
II	Isogenies and applications	10
6	Elliptic curves over \mathbb{C}	10
7	The endomorphism ring	14
8	Application: point counting	16
9	Isogeny graphs	17
10	Application: computing irreducible polynomials	19
III	Cryptography from isogeny graphs	22
11	Expander graphs	22
12	Isogeny graphs in cryptanalysis	24
13	Provably secure hash functions	25
14	Post-quantum key exchange	26
15	Further topics in isogeny based cryptography	35

Part I

Elliptic curves and cryptography

Throughout this part we let k be a field, and we denote by \bar{k} its algebraic closure. We review the basic theory of elliptic curves, and two classic applications in cryptography. The interested reader will find more details on elliptic curves in [63], and on their use in cryptography in [40, 30].

1 Elliptic curves

Elliptic curves are projective curves of genus 1 having a specified base 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 $x_i = \lambda y_i$ for all i .

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

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

By fixing arbitrarily the coordinate $x_n = 0$, we define a projective space of dimension $n - 1$, which we call the *space 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 [63, Chap. III].

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

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

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 $\mathcal{O} = (0 : 1 : 0)$.

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 $(0 : 1 : 0)$.




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

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 a 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

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

Proof. See [63, Coro. 6.4]. For the characteristic 0 case see also next part. \square

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*.

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 tracked 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 $\phi : E \rightarrow E'$ 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 [63, III, Th. 4.8]. □

Two curves are called *isogenous* if there exists an isogeny between them. We shall see in the next part 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 $E[m]$. For most elliptic curves, this is the end of the story: the only endomorphisms are the scalar multiplications. We shall however see some non-trivial endomorphisms soon.

3 Elliptic curves over finite fields

From now on we let E be an elliptic curve defined over a finite field k with q elements. 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 .

Curves over finite fields always have a non-trivial endomorphism.

Definition 8 (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 9. Let π be the Frobenius endomorphism of E . Then:

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

Corollary 10 (Hasse's theorem). Let E be an elliptic curve defined over a finite field k with q elements, then

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

Proof. See [63, V, Th. 1.1]. □

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 a special case of the famous *Weil's conjectures*, proven by Weil himself in 1949 for Abelian varieties, and more generally by Deligne in 1973.

Definition 11. 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 12. 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 [63, V, Th. 2.4]. □

We conclude with a theorem that links the isogenies between two elliptic curves with their Frobenius endomorphisms.

Theorem 13 (Sato-Tate). Two elliptic curves E, E' defined over a finite field k are isogenous over k if and only if $\#E(k) = \#E'(k)$.

4 Application: Diffie-Hellman key exchange

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 [24], 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 14 (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 , then she can compute the shared secret; the converse is not true in general, but it is 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 [40]), 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 $G = (\mathbb{Z}/p\mathbb{Z})^\times$ [40], thus requiring parameter 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 [51] and Koblitz [42] to suggest, in the 1980s, to replace $(\mathbb{Z}/p\mathbb{Z})^\times$ in the Diffie-Hellman 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 2.

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 2: The Diffie-Hellman protocol over elliptic curves

5 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 [47] 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 whose 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

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

given by the Chinese remainder theorem. The algorithm is detailed in Figure 3a. 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 3b. 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 ,

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 $y = x^e \pmod N$;
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 $b = Y^2 - X^3 - aX \pmod N$;
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 3: The $(p - 1)$ and ECM factorization algorithms

a problem that is known to be as 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 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 $E(\mathbb{F}_p) \oplus E(\mathbb{F}_q)$, and any attempt at inverting a non-invertible in $\mathbb{Z}/N\mathbb{Z}$ will result in a factorization of N .

Exercices

Exercise I.1. Prove Proposition 6.

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

Exercise I.3. Prove Proposition 9.

Exercise I.4. Using Proposition 12, devise an algorithm to effectively compute $\#E(\mathbb{F}_{q^n})$ given $\#E(\mathbb{F}_q)$.

Exercise I.5. Implement the ECDH key exchange in the language of your choice.

Exercise I.6 (Polig-Hellman algorithm). Let G be a cyclic group of order $N = pq$, generated by an element g . Show how to solve discrete logarithms in G by computing two separate discrete logarithms in the subgroups $\langle g^p \rangle$ and $\langle g^q \rangle$.

Exercise I.7. Implement the ECM factorization method in the language of your choice.




Figure 4: A complex lattice (black dots) and its associated complex torus (grayed *fundamental domain*).

Part II

Isogenies and applications

6 Elliptic curves over \mathbb{C}

Definition 15 (Complex lattice). A complex lattice Λ is a discrete subgroup of \mathbb{C} that contains an \mathbb{R} -basis.

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 $\text{Im}(\omega_1/\omega_2) > 0$; we then say that the basis has *positive orientation*. A positively oriented basis is obviously not unique, though.

Proposition 16. *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*

$$\begin{aligned}\omega'_1 &= a\omega_1 + b\omega_2, \\ \omega'_2 &= c\omega_1 + d\omega_2,\end{aligned}$$

for some matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \text{SL}_2(\mathbb{Z})$.

Proof. See [64, I, Lem. 2.4]. □

Definition 17 (Complex torus). Let Λ be a complex lattice, the quotient \mathbb{C}/Λ is called a *complex torus*.




Figure 5: Addition (left) and scalar multiplication (right) of points in a complex torus \mathbb{C}/Λ .

A convex set of class representatives of \mathbb{C}/Λ is called a *fundamental parallelogram*. Figure 4 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}/Λ , and can be graphically represented as operations on points inside a fundamental parallelogram. This is illustrated in Figure 5.

Definition 18 (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 19 (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 [64, 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 20 (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 21. *The Weierstrass 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*

$$E : y^2 = 4x^3 - g_2x - g_3$$

is an elliptic curve over \mathbb{C} . The map

$$\begin{aligned} \mathbb{C}/\Lambda &\rightarrow E(\mathbb{C}), \\ 0 &\mapsto (0 : 1 : 0), \\ z &\mapsto (\wp(z) : \wp'(z) : 1) \end{aligned}$$

is an isomorphism of Riemann surfaces and a group morphism.

Proof. See [63, 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 homotopy class of complex tori, we have a corresponding isomorphism class of elliptic curves. The converse is also true.

Theorem 22 (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 [64, 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 6a; 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

$$\begin{aligned} [m] : \mathbb{C}/\Lambda &\rightarrow \mathbb{C}/\Lambda', \\ z &\mapsto z \bmod \Lambda'; \end{aligned}$$

and we verify that it is an endomorphism because Λ and Λ' are homothetic.




Figure 6: Maps between complex tori.

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

$$\begin{aligned} \phi_\alpha : \mathbb{C}/\Lambda &\rightarrow \mathbb{C}/\Lambda', \\ z &\mapsto \alpha z \bmod \Lambda' \end{aligned}$$

that is holomorphic and also a group morphism. One example of such maps is given in Figure 6a: 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.

Theorem 23. *Let E, E' be elliptic curves over \mathbb{C} , with corresponding lattices Λ, Λ' . There is a bijection between the group of isogenies from E to E' and the group of maps ϕ_α for all α such that $\Lambda \subset \alpha\Lambda'$.*

Proof. See [63, VI, Th. 4.1]. □

Looking again at Figure 6a, 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}/Λ , 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 for any base field.

We finish this section by summarizing the most important algebraic properties of isogenies; we start with a technical definition.

Definition 24 (Degree, separability). Let $\phi : E \rightarrow E'$ 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.
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 [63, II, Th. 2.4]. □

In practice, most of the time we will be considering separable isogenies, and we can take $\deg \phi = \# \ker \phi$ as the definition of the degree. Notice that in this case $\deg \phi$ is the size of any fiber of ϕ . Separable isogenies are completely determined by their kernel, as the following proposition shows.

Proposition 25. *Let E be an elliptic curve, and let G be a finite subgroup of E . There are a unique elliptic curve E' , and a unique separable isogeny ϕ , such that $\ker \phi = G$ and $\phi : E \rightarrow E'$.*

Proof. See [63, Prop. III, 4.12]. □

The proposition justifies introducing the notation E/G for the image curve E' . We conclude with a fundamental theorem on isogenies.

Theorem 26 (Dual isogeny). *Let $\phi : E \rightarrow E'$ be an isogeny of degree m . There is a unique isogeny $\hat{\phi} : E' \rightarrow E$ 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. $\hat{\phi}$ is defined over k if and only if ϕ is;
2. $\widehat{\psi \circ \phi} = \hat{\phi} \circ \hat{\psi}$ for any isogeny $\psi : E' \rightarrow E''$;
3. $\widehat{\psi + \phi} = \hat{\psi} + \hat{\phi}$ for any isogeny $\psi : E \rightarrow E'$;
4. $\deg \phi = \deg \hat{\phi}$;
5. $\hat{\hat{\phi}} = \phi$.

7 The endomorphism ring

We have already defined an endomorphism as an isogeny from a curve to itself. If we add the multiplication-by-0, the set of all endomorphisms of E form a ring under the operations of addition and composition, denoted by $\text{End}(E)$.

We have already seen that the multiplication-by- m is a different endomorphism for any integer m , thus $\mathbb{Z} \subset \text{End}(E)$. For the case of finite fields, we have also learned about the Frobenius endomorphism π ; so certainly $\mathbb{Z}[\pi] \subset \text{End}(E)$ in this case. We shall now give a complete characterization of the endomorphism ring for any field.

Definition 27 (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 of maximal dimension.

The prototypical example of order is the ring of integers \mathcal{O}_K of a number field K , i.e., the ring of all elements of K such that their monic minimal polynomial has coefficients in \mathbb{Z} . It turns out that \mathcal{O}_K is the *maximal order* of K , i.e., it contains any other order of K .

Definition 28 (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 29 (Deuring). Let E be an elliptic curve defined over a field k of characteristic p . The ring $\text{End}(E)$ 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 the quaternion algebra ramified at p and ∞ ; in this case we say that E is supersingular.

Proof. See [63, III, Coro. 9.4] and [4]. □

In positive characteristic, a curve that is not supersingular is called *ordinary*; it necessarily has complex multiplication. We focus again on the finite field case; we have already seen that $\mathbb{Z}[\pi] \subset \text{End}(E)$. Now, Hasse's theorem can be made more precise as follows.

Theorem 30. Let E be an elliptic curve defined over a finite field. Its Frobenius endomorphism π satisfies a quadratic equation

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

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

Proof. See [63, 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 $\#E = q + 1 - t$. Now, if we let $D_\pi = t^2 - 4q < 0$, we verify that $\pi \in \mathbb{Q}[\sqrt{D_\pi}]$; so, at least in the ordinary case, we can affirm that

$$\mathbb{Z}[\pi] \subset \text{End}(E) \subset \mathcal{O}_K,$$

where $K = \mathbb{Q}[\sqrt{D_\pi}]$ is called the *endomorphism algebra* of E . The structure of the orders of K is very simple in this case.

Proposition 31. 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 d_K is the discriminant of K , the discriminant of \mathcal{O} is f^2d_K .

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

In our case, we can write $D_\pi = f^2d_K$, with d_K squarefree. Then, any order $\mathbb{Z}[\pi] \subset \mathcal{O} \subset \mathcal{O}_K$ has conductor dividing f .

8 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 [60]. 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 32 (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

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

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 [1, 2, 25, 61, 26]. 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 $\text{GL}_2(\mathbb{Z}/\ell\mathbb{Z})$, and its characteristic polynomial is exactly

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

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 $M = \begin{pmatrix} \lambda & 0 \\ 0 & \mu \end{pmatrix}$. 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 $\phi : E \rightarrow E/\langle P \rangle$, 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 $t_\pi = \lambda + q/\lambda \pmod{\ell}$.

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 $\langle P \rangle$ we 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 [61, 48, 26, 67].

9 Isogeny graphs

We now look at the graph structure that isogenies create on the set of j -invariants defined over a finite field. We start with an easy generalization of the Sato-Tate theorem 13.

Theorem 33 (Sato-Tate). *Two elliptic curves E, E' defined over a finite field are isogenous if and only if their endomorphism algebras $\text{End}(E) \otimes \mathbb{Q}$ and $\text{End}(E') \otimes \mathbb{Q}$ are isomorphic.*

An equivalence class of isogenous elliptic curves is called an *isogeny class*. In particular, we see that it is impossible for an isogeny class to contain both ordinary and supersingular curves. When we restrict to isogenies of a prescribed degree ℓ , we say that two curves are ℓ -isogenous; by the dual isogeny theorem, this too is an equivalence relation. Remark that if E is ℓ -isogenous to E' , and if E'' is isomorphic to E' , then by composition E and E'' are also ℓ -isogenous.

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.

Definition 34 (Isogeny graph). An *isogeny graph* is a (multi)-graph whose nodes are the j -invariants of isogenous curves, and whose edges are isogenies between them.

The dual isogeny theorem guarantees that for every isogeny $E \rightarrow E'$ there is a corresponding isogeny $E' \rightarrow E$ of the same degree. For this reason, isogeny graphs are usually drawn undirected. Figure 7 shows a typical example of isogeny graph, where we restrict to isogenies of degree 3.

The classification of isogeny graphs was initiated by Pizer [56, 57] and Kohel [43]; further algorithmic treatment of graphs of ordinary curves, and the now famous name of *isogeny volcanoes* was subsequently given by Fouquet and Morain [28]. We start with some generalities.

Proposition 35. *Let $E : y^2 = x^3 + ax + b$ be an elliptic curve defined over a finite field k of characteristic p , and let $\ell \neq p$ be a prime.*

1. *There are $\ell + 1$ distinct isogenies of degree ℓ with domain E defined over the algebraic closure \bar{k} .*




Figure 7: A volcano of 3-isogenies (ordinary elliptic curves, Elkies case), and the correspond tower of orders inside the endomorphism algebra.

2. There are 0, 1, 2 or $\ell + 1$ isogenies of degree ℓ with domain E defined over k .
3. If E is ordinary, there is a unique separable isogeny of degree p with domain E ; there are none if E is supersingular.
4. The map $(x, y) \mapsto (x^p, y^p)$ is a purely inseparable isogeny of degree p from E to $E^{(p)} : y^2 = x^3 + a^p x + b^p$.

There are many differences between the structure of isogeny graphs of ordinary curves and those of supersingular ones. We focus here on the ordinary case, and we leave the supersingular one for the last part.

Proposition 36 (Horizontal and vertical isogenies). *Let $\phi : E \rightarrow E'$ be an isogeny of prime degree ℓ , and let $\mathcal{O}, \mathcal{O}'$ be the orders corresponding to E, E' . Then, either $\mathcal{O} \subset \mathcal{O}'$ or $\mathcal{O}' \subset \mathcal{O}$, and one of the following is true:*

- $\mathcal{O} = \mathcal{O}'$, in this case ϕ is said to be horizontal;
- $[\mathcal{O}' : \mathcal{O}] = \ell$, in this case ϕ is said to be ascending;
- $[\mathcal{O} : \mathcal{O}'] = \ell$, in this case ϕ is said to be descending.

Proof. See [43, Prop. 21]. □

Observe that vertical isogenies can only exist for primes that divide the conductor of $\mathbb{Z}[\pi]$, so the horizontal case is the generic one. Like we did for the SEA algorithm we can further distinguish three cases, depending on the value of the Legendre symbol $\left(\frac{D}{\ell}\right)$, i.e., depending on whether π splits (Elkies case), is inert (Atkin case), or ramifies modulo ℓ . All possible cases are encoded in the following proposition.

Proposition 37. *Let E be an elliptic curve over a finite field k . Let \mathcal{O} its endomorphism ring, f its conductor, D its discriminant, π the Frobenius endomorphism, f_π the conductor of $\mathbb{Z}[\pi]$. Let ℓ be a prime different from the characteristic of k , then the types of degree ℓ isogenies with domain E are as follows:*

- If $\ell \nmid f$ and $\ell \nmid (f_\pi/f)$, there is one ascending isogeny;

- If $\ell \mid f$ and $\ell \mid (f_\pi/f)$, there is one ascending isogeny and ℓ descending ones;
- If $\ell \nmid f$ and $\ell \nmid (f_\pi/f)$, there are $1 + \left(\frac{D}{\ell}\right)$ horizontal isogenies, where $\left(\frac{D}{\ell}\right)$ represents the Legendre symbol;
- If $\ell \nmid f$ there are $1 + \left(\frac{D}{\ell}\right)$ horizontal isogenies, plus $\ell - \left(\frac{D}{\ell}\right)$ descending isogenies only if $\ell \mid (f_\pi/f)$.

Proof. See [43, Prop. 21]. □

Putting the pieces together, we see that graphs of ordinary curves have a very rigid structure: a cycle of horizontal isogenies (Elkies case), possibly reduced to one point (Atkin case), or to two points (ramified case); and a tree of descending isogenies of height $v_\ell(f_\pi)$ (the ℓ -adic valuation of the conductor of π). Such graphs are called *isogeny volcanoes* for obvious reasons (have a look at Figure 7).

The action of π on $E[\ell]$, or more generally on $E[\ell^k]$ for k large enough, can be used to determine even more precisely which isogenies are ascending, descending or horizontal. We will not give details here, but see [52, 53, 36, 21].

10 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 $\ell \nmid \#E(\mathbb{F}_q)$, thus we are no longer limited to $\ell \mid (q - 1)$; 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 [17] 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 whose 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 Exercice II.8). 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 [73].

Proposition 38 (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 $\phi : E \rightarrow E/G$, 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

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

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

Proof. See [19, §8.2]. □

Corollary 39. *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 8. 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., $(\frac{D}{\ell}) = 1$), and that each isogeny ϕ_i is horizontal, thus their composition eventually forms a cycle, the *crater* of a volcano.

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 $\ell \nmid \#E_0$ 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 $\phi_i : E_{i-1} \rightarrow E_i$;
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 $\phi_i^{-1} \circ \dots \circ \phi_1^{-1}(P)$.




Figure 8: Couveignes-Lercier algorithm to compute irreducible polynomials, and structure of the computed isogeny cycle.

Exercices

Exercice II.1. Prove Lemma 16.

Exercice II.2. Prove that y divides the m -th division polynomial ψ_m if and only if m is even, and that no division polynomial is divisible by y^2 .

Exercice II.3. Using the Sato-Tate theorem 33, prove that two curves are isogenous if and only if they have the same number of points.

Exercice II.4. Prove Propostion 35.

Exercice II.5. Prove that the dual of a horizontal isogeny is horizontal, and that the dual of a descending isogeny is ascending.

Exercice II.6. Prove that the height of a volcano of ℓ -isogenies is $v_\ell(f_\pi)$, the ℓ -adic valuation of the Frobenius endomorphism.

Exercice II.7. Let $X^2 - tX - q$ be the minimal polynomial of π , and suppose that it splits as $(X - \lambda)(X - \mu)$ in \mathbb{Z}_ℓ (the ring of ℓ -adic integers). Prove that the volcano of ℓ isogenies has height $v_\ell(\lambda - \mu)$.

Exercice II.8. Prove that $E[\ell] \subset E(\mathbb{F}_q)$ implies $\ell \mid (q-1)$.

Part III

Cryptography from isogeny graphs

11 Expander graphs

When we talk about *Isogeny Based Cryptography*, as a topic distinct from *Elliptic Curve Cryptography*, we usually mean algorithms and protocols that rely fundamentally on the structure of *large* isogeny graphs. The cryptographically interesting properties of these graphs are usually tied to their *expansion* properties.

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* 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 $v \rightarrow v_1 \rightarrow \dots \rightarrow v'$ 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 among 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 $V = \{v_1, \dots, v_n\}$ 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

$$\lambda_1 \geq \dots \geq \lambda_n.$$

It is convenient to identify functions on V with vectors in \mathbb{R}^n , and therefore also think of the adjacency matrix as a self-adjoint operator on $L^2(V)$. Then can we immediately bound the eigenvalues of G .

Proposition 40. *If G is a k -regular graph, then its largest and smallest eigenvalues λ_1, λ_n satisfy*

$$k = \lambda_1 \geq \lambda_n \geq -k.$$

Proof. See [69, Lem. 2]. □

Definition 41 (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 $\#V_i \rightarrow \infty$ 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 42 (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 $n \rightarrow \infty$. A graph such that $|\lambda_i| \leq 2\sqrt{k-1}$ for any λ_i 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 43 (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 44 (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 45. *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*

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

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 $v_1 \rightarrow \dots \rightarrow v_i$, 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 its diameter terminate on any vertex with probability close to uniform.

Proposition 46 (Mixing theorem). *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 $\#F/2\#V$.

Proof. See [38]. □

The length in the previous proposition is also called the *mixing length* of the expander graph. We conclude this section with two results on expansion in graphs of isogenies.

Theorem 47 (Supersingular graphs are Ramanujan). *Let p, ℓ be distinct primes, then*

1. *All supersingular j -invariants of curves in $\overline{\mathbb{F}}_p$ are defined in \mathbb{F}_{p^2} ;*

2. There are

$$\lfloor \frac{p}{12} \rfloor + \begin{cases} 0 & \text{if } p \equiv 1 \pmod{12} \\ 1 & \text{if } p \equiv 5, 7 \pmod{12} \\ 2 & \text{if } p \equiv 11 \pmod{12} \end{cases}$$

isomorphism classes of supersingular elliptic curves over $\bar{\mathbb{F}}_p$;

3. The graph of supersingular curves in $\bar{\mathbb{F}}_p$ with ℓ -isogenies is connected, $\ell + 1$ regular, and has the Ramanujan property.

Proof. See [63, V, Th. 4.1], [56, 57], [8]. □

Theorem 48 (Graphs of horizontal isogenies are expanders). *Let \mathbb{F}_q be a finite field and let $\mathcal{O} \subset \mathbb{Q}[\sqrt{-D}]$ be an order in a quadratic imaginary field. Let G be the graph whose vertices are elliptic curves over \mathbb{F}_q with complex multiplication by \mathcal{O} , and whose edges are (horizontal) isogenies of prime degree bounded by $(\log q)^{2+\delta}$ for some fixed $\delta > 0$. Assume that G is non-empty. Then, under the generalized Riemann hypothesis, G is a regular graph and there exists an ε , independent of \mathcal{O} and q , such that G is a one-sided ε -expander.*

Proof. See [38]. □

12 Isogeny graphs in cryptanalysis

Besides the applications to point counting mentioned in the previous part, the first application of isogenies in cryptography has been to study the difficulty of the discrete logarithm problem in 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 .

Vélu's formulas (Proposition 38) give a solution to this problem in $\tilde{O}(\#G)$ operations over the field of definition 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). Given two elliptic curves E, E' over a finite field, isogenous of known degree d , find an isogeny $\phi : E \rightarrow E'$ of degree d .

Remark that, up to automorphisms, the isogeny ϕ is typically unique. Elkies was the first to formulate the problem and give an algorithm [25, 26] whose complexity is in $O(d^3)$ in general and in $\tilde{O}(d)$ in the special context of the SEA algorithm [7, 49]. Alternate algorithms, with complexity $O(d^2)$ in general, are due to Couveignes and others [13, 14, 15, 23, 21].

Problem 3 (Isogeny path). Given two elliptic curves E, E' over a finite field k , such that $\#E = \#E'$, find an isogeny $\phi : E \rightarrow E'$ of smooth degree.




Figure 9: The meet in the middle attack in weak isogeny classes.




Figure 10: Hashing the string 010101 using an expander graph

This problem, and variations thereof, is the one that occurs most in isogeny based cryptography. It is a notoriously difficult problem, for which only algorithms exponential in $\log \#E$ are known in general. A general strategy to tackle it is by a *meet in the middle* random walk [29]: 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 [34, 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 an isogeny class of ordinary elliptic curves is $O(\sqrt{\#E})$, 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 9.

Similar ideas have been used to construct *key escrow systems* [70], and to prove random reducibility of discrete logarithms inside some isogeny classes [38].

13 Provably secure hash functions

The next application of isogeny graphs is constructing *provably secure hash functions*. 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 [8] chooses an arbitrary start vertex j_0 in an expander graph, then takes a random walk (without backtracking) 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 10.

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

Problem 5 (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 *Cayley graphs*, and was broken shortly afterwards [72, 55]. 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 worryingly, some weaknesses have recently been highlighted [44], that could potentially lead to backdoors in standardized parameters.

14 Post-quantum key exchange

We come to the last, more powerful constructions based on isogeny graphs. We present here two key exchange protocols, similar in spirit to the Diffie-Hellman protocol discussed in the first part. Both protocols are significantly less efficient than ECDH, however they are relevant because of their conjectured *quantum security*. In recent years, 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 [62] 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.

Both protocols are based on random walks in an isogeny graph. The two participants, Alice and Bob, start from the same common curve E_0 , and take a (secret) random walk to some curves E_A, E_B . After publishing their respective curves, Alice starts a new walk from E_B , while Bob starts from E_A . By repeating the “same” secret steps, they both eventually arrive on a shared secret curve E_S , only known to them. While the idea may seem simple, its realization is far from easy. Indeed, as opposed to the hash function case, we cannot be content with an arbitrary labeling of the graph edges. We must instead use the algebraic properties of the isogeny graphs to ensure that Alice and Bob’s walks “commute”.

14.1 Hard homogeneous spaces

The first protocol originates in a preprint by Couveignes [16], but was only later put into practice and popularized by Rostovtsev and Stolbunov [59, 65]. It uses random walks in graphs of ordinary curves with horizontal isogenies; in this sense, it is a direct application of Theorem 48. The protocol can be viewed as a special instance of a general construction on *Schreier graphs*, a generalization of *Cayley graphs*.

Definition 49 (Schreier graph). Let G be a group *acting freely* on a set X , in the sense that there is a map

$$\begin{aligned} G \times X &\rightarrow X \\ (\sigma, x) &\mapsto \sigma \cdot x \end{aligned}$$




Figure 11: Schreier graph of the generators of a group of order 13 under the action of $S = \{2, 3, 5, 2^{-1}, 3^{-1}, 5^{-1}\} \subset (\mathbb{Z}/13\mathbb{Z})^\times$.

such that $\sigma \cdot x = x$ if and only if $\sigma = 1$, and $\sigma \cdot (\tau \cdot x) = (\sigma\tau) \cdot x$, for all $\sigma, \tau \in G$ and $x \in X$. Let $S \subset G$ be a *symmetric* subset, i.e. one not containing 1 and closed under inversion. The *Schreier graph* of (S, X) is the graph whose vertices are the elements of X , and such that $x, x' \in X$ are connected by an edge if and only if $\sigma \cdot x = x'$ for some $\sigma \in S$.

Because of the constraints on the group action and the set S , Schreier graphs are undirected and regular, and they usually make good expanders (see exercise III.2). Note that Cayley graphs are the Schreier graphs of the (left) action of a group on itself.

As an example, take a cyclic group G of order n , then $(\mathbb{Z}/n\mathbb{Z})^\times$ acts naturally on G by the law $\sigma \cdot g = g^\sigma$ for any $g \in G$ and $\sigma \in (\mathbb{Z}/n\mathbb{Z})^\times$. This action is not free on G , but it is on the subset P of all generators of G ; we can thus build the Schreier graph (S, P) , where S is a symmetric subset that generates $(\mathbb{Z}/n\mathbb{Z})^\times$. An example of such graph for the case $n = 13$ is shown in Figure 11, where the set $S \subset (\mathbb{Z}/13\mathbb{Z})^\times$ has been chosen to contain 2, 3, 5 and their inverses.

By slightly generalizing Couveignes' work [16], we will now show how to construct a key exchange protocol based on this family of Schreier graphs. We will restrict to cyclic groups of prime order p , and we will have the cryptosystem security grow exponentially in p . Let $G = \langle g \rangle$ be a cyclic group of order p ; let $D \subset (\mathbb{Z}/p\mathbb{Z})^\times$ be a generating set such that $\sigma \in D$ implies $\sigma^{-1} \notin D$; and let $S = D \cup D^{-1}$. We call *directed route* a sequence of elements of D . A directed route $\rho \in D^*$, together with a *starting vertex* $g \in G$, defines a walk in the Schreier graph (S, G) by starting in g , and successively taking the edges corresponding to the labels in ρ . If ρ is a directed route, and $g \in G$, we write $\rho(g)$ for vertex where the walk defined by ρ and g ends. We can now define a key exchange protocol where the secrets are random directed routes, and the public data are vertices of the Schreier graph. The protocol is summarized in Figure 12.

A graphical example of this protocol with $p = 13$ and $D = \{2, 3, 5\}$ is given in Figure 13. To understand why it works, observe that if ρ is a route of length m

$$\rho = (\sigma_1, \dots, \sigma_m),$$

then

$$\rho(g) = \exp_g \left(\prod \sigma_i \right)$$

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	$\rho_A \in D^*$	$\rho_B \in D^*$
Compute public data	$g_A = \rho_A(g)$	$g_B = \rho_B(g)$
Exchange data	$g_A \longrightarrow$	$\longleftarrow g_B$
Compute shared secret	$g_{AB} = \rho_A(g_B)$	$g_{AB} = \rho_B(g_A)$

Figure 12: Key exchange protocol based on random walks in a Schreier graph.




Figure 13: Example of key exchange on the Schreier graph of Figure 11. Alice's route is represented by continuous lines, Bob's route 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.

for any $g \in G$. Hence, the order of the steps in a route does not count: what counts is only how many times each element of D appears in ρ . We immediately realize that this protocol is nothing else than the classical Diffie-Hellman protocol on the group G , presented in a twisted way.¹

For this protocol to have the same security as the original Diffie-Hellman, we need the public keys g_A, g_B to be (almost) uniformly distributed. Hence, we shall require that the graph is an expander, and that walks are longer than the mixing length; i.e., that D generates $(\mathbb{Z}/p\mathbb{Z})^\times$, and that walks have length $\sim \log p$. Since a secret route is simply defined by the number of times each element of D is present, we shall also need $\#D \sim \log p / \log \log p$ in order to have a large enough key space. If we respect all these constraints, we end up with a protocol that is essentially equivalent to the original Diffie-Hellman, only less efficient.

It is now an easy exercise to generalize to other Schreier graphs. To see how this applies to isogeny graphs, we must take a step back, and define some more objects related to elliptic curves.

Definition 50 (Fractional ideal). Let \mathcal{O} be an order in a number field K . A *fractional ideal* of \mathcal{O} is a non-zero subgroup $I \subset K$ such that

- $xI \subset I$ for all $x \in \mathcal{O}$, and
- there exists a non-zero $x \in \mathcal{O}$ such that $xI \subset \mathcal{O}$.

¹A minor difference lies in the fact that this protocol avoids non-primitive elements of G , whereas the classical Diffie-Hellman protocol may well use public keys belonging to a subgroup of G .

A fractional ideal is called *principal* if it is of the form $x\mathcal{O}$ for some $x \in K$.

Note that the ideals of \mathcal{O} are exactly the fractional ideals contained in \mathcal{O} ; however, from now on we will simply call *ideals* the fraction ideals, and we will use the name *integral ideal* for ordinary ones. An ideal I is said to be *invertible* if there is another ideal J such that $IJ = \mathcal{O}$. Invertible ideals form an Abelian group, written multiplicatively, under the operation

$$IJ = \{xy \mid x \in I, y \in J\}.$$

It is easily verified that \mathcal{O} is the neutral element of the group, and that principal ideals form a subgroup of it.

Proposition 51 (Ideal class group). *Let \mathcal{O} be an order in a number field K . Let $\mathcal{I}(\mathcal{O})$ be its group of invertible ideals, and $\mathcal{P}(\mathcal{O})$ the subgroup of principal ideals. The (ideal) class group of \mathcal{O} is the quotient*

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

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

The class group is a fundamental object in the study of number fields and their Galois theory. What is relevant to us, is the fact that the elements of $\text{Cl}(\mathcal{O})$ are *represented* by horizontal isogenies, a fact that is developed in the theory of *complex multiplication*. We only take here a small peek at the theory; see [46, 64, 18] for a detailed account.

Definition 52 (\mathfrak{a} -torsion). Let E be an elliptic curve defined over a finite field \mathbb{F}_q . Let \mathcal{O} be the endomorphism ring of E , and let $\mathfrak{a} \subset \mathcal{O}$ be an integral invertible ideal of norm coprime to q . We define the \mathfrak{a} -torsion subgroup of E as

$$E[\mathfrak{a}] = \{P \in E \mid \alpha(P) = 0 \text{ for all } \alpha \in \mathfrak{a}\}.$$

Given an ideal $\mathfrak{a} \subset \mathcal{O}$ as above, it is natural to define the (separable) isogeny $\phi_{\mathfrak{a}} : E \rightarrow E_{\mathfrak{a}}$ whose kernel is $E[\mathfrak{a}]$. This definition can be readily extended to inseparable isogenies. Since \mathfrak{a} is invertible, we can show that $\text{End}(E) \simeq \text{End}(E_{\mathfrak{a}}) \simeq \mathcal{O}$, that $E_{\mathfrak{a}}$ only depends on the class of \mathfrak{a} in $\text{Cl}(\mathcal{O})$, and that the map $(\mathfrak{a}, E) \mapsto E_{\mathfrak{a}}$ defines a group action of $\text{Cl}(\mathcal{O})$ on the set of elliptic curves with complex multiplication by \mathcal{O} .

Theorem 53. *Let \mathbb{F}_q be a finite field, and let $\mathcal{O} \subset \mathbb{Q}[\sqrt{-D}]$ be an order in a quadratic imaginary field. Denote by $\text{Ell}_q(\mathcal{O})$ the set of elliptic curves defined over \mathbb{F}_q with complex multiplication by \mathcal{O} .*

Assume that $\text{Ell}_q(\mathcal{O})$ is non-empty, then the class group $\text{Cl}(\mathcal{O})$ acts freely and transitively on it; i.e., there is a map

$$\begin{aligned} \text{Cl}(\mathcal{O}) \times \text{Ell}_q(\mathcal{O}) &\rightarrow \text{Ell}_q(\mathcal{O}) \\ (\mathfrak{a}, E) &\mapsto \mathfrak{a} \cdot E \end{aligned}$$

such that $\mathfrak{a} \cdot (\mathfrak{b} \cdot E) = (\mathfrak{a}\mathfrak{b}) \cdot E$ for all $\mathfrak{a}, \mathfrak{b} \in \text{Cl}(\mathcal{O})$ and $E \in \text{Ell}_q(\mathcal{O})$, and such that for any $E, E' \in \text{Ell}_q(\mathcal{O})$ there is a unique $\mathfrak{a} \in \text{Cl}(\mathcal{O})$ such that $E' = \mathfrak{a} \cdot E$.

A set, that is acted upon freely and transitively by a group G , is also called a *principal homogeneous space* or a *torsor* for G . An immediate consequence of the theorem above is that the torsor $\text{Ell}_q(\mathcal{O})$ has cardinality equal to the class number $h(\mathcal{O})$.

Following on from the connection between isogenies and ideals, suppose that $\ell\mathcal{O}$ splits into prime ideals as $\ell\mathcal{O} = \mathfrak{l}\bar{\mathfrak{l}}$. Set $S = \{\mathfrak{l}, \bar{\mathfrak{l}}\}$, then the Schreier graph of $(S, \text{Ell}_q(\mathcal{O}))$ is exactly the

graph of horizontal ℓ -isogenies on $\text{Ell}_q(\mathcal{O})$. More generally, if we let $S \subset \text{Cl}(\mathcal{O})$ be a symmetric subset, its Schreier graph is a graph of horizontal isogenies, and it is an expander if and only if S generates $\text{Cl}(\mathcal{O})$.

Based on this observation, we can now give a key exchange protocol based on random walks in graphs of horizontal isogenies. The general idea was already present in Couveignes' work [16], but it was Rostovtsev and Stolbunov who proposed to use isogeny computations to effectively implement the protocol [59, 65].

The protocol implicitly uses the set $\text{Ell}_q(\mathcal{O})$ of elliptic curves over \mathbb{F}_q with complex multiplication by some order \mathcal{O} ; however it never explicitly computes \mathcal{O} . Instead, it determines parameters in the following order:

1. A *large enough* finite field \mathbb{F}_q ;
2. A curve E defined over \mathbb{F}_q ;
3. The Frobenius discriminant $D_\pi = t_\pi^2 - 4q$ of E is computed through point counting, and it is verified that it contains a *large enough* prime factor;
4. A set $L = \{\ell_1, \dots, \ell_m\}$ of primes that split in $\mathbb{Z}[\pi]$, i.e., such that $\left(\frac{D_\pi}{\ell_i}\right) = 1$;
5. For each prime ℓ_i , the factorization

$$\pi^2 - t_\pi \pi + q = (\pi - \lambda_i)(\pi - \mu_i) \pmod{\ell_i}$$

is computed, and one of the roots, say λ_i , is chosen arbitrarily as *positive direction*.

The condition on the ℓ_i 's guarantees that each graph of ℓ_i -isogenies on $\text{Ell}_q(\mathcal{O})$ is 2-regular. The choice of a *positive direction* allows us to *orient* the graph, by associating to λ_i the isogeny with kernel $E[\ell_i] \cap \ker(\pi - \lambda_i)$. The key exchange now proceeds like the ordinary Diffie-Hellman protocol:

1. Alice chooses a random walk made of steps in L along the positive direction; denote the walk by $\rho_A \in L^*$, and denote by $E_A = \rho_A(E)$ the curve where the walk terminates. Note that E_A does only depends on how many times each ℓ_i appears in ρ_A , and not on their order.
2. Bob does the same, choosing a random walk ρ_B and computing $E_B = \rho_B(E)$.
3. Alice and Bob exchange E_A and E_B .
4. Alice computes the *shared secret* $\rho_A(E_B)$.
5. Bob computes the *shared secret* $\rho_B(E_A)$.

The actual computations are carried out by solving *explicit isogeny problems* (see Problem 2), in much the same way they are done in the Elkies case of the SEA algorithm (see Section 8). The protocol is summarized in Figure 14.

We conclude this section with a discussion on the security of the Rostovtsev-Stolbunov protocol. All the protocol's security rests on the isogeny path problem: given E and E_A , find an isogeny $\phi : E \rightarrow E_A$ of smooth order. To be safe against exhaustive search and meet in the middle attacks as seen in Section 12, the set $\text{Ell}_q(\mathcal{O})$ must be large. On average $\#\text{Ell}_q(\mathcal{O}) \sim \sqrt{q}$, thus we shall take $\log_2 q \approx 512$ for a security level of at most 128 bits. However, some isogeny classes are much smaller than average, this is why we also need check that D_π has a large prime factor.

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 $L = \{\ell_1, \dots, \ell_m\}$ 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 14: Rostovtsev-Stolbunov key exchange protocol based on random walks in an isogeny graph.

Furthermore, for the public and private curves to be (almost) uniformly distributed in $\text{Ell}_q(\mathcal{O})$, we need the isogeny graph to be connected; equivalently, we need the ideals $(\ell_i, \pi - \lambda_i)$ to generate $\text{Cl}(\mathcal{O})$. Theorem 48 ensures this is the case if $\#L \sim (\log q)^2$, however it is usually sufficient to take a much smaller set in practice. It is not enough to have an expander: we also need the random walks to be longer than the mixing length, that is $\sim \log q$. And, since the key space grows exponentially with $\#L$, rather than with the walk length, we shall also ask that $\#L \sim \log q / \log \log q$.

When all conditions are met, the best known attack against this cryptosystem is the meet in the middle strategy, which runs in $O(\sqrt[4]{q})$ steps. However, the real case for this system is made by looking at attacks performed on a *quantum computer*. It is well known that Shor's algorithm [62] breaks the Diffie-Hellman cryptosystem in polynomial time on a quantum computer, and thus it also breaks the protocol of Figure 12. More generally, Shor's algorithm can solve the (generalized) discrete logarithm problem in any Abelian group, and in particular in $\text{Cl}(\mathcal{O})$. However, in the Rostovtsev-Stolbunov protocol, the attacker only sees E , E_A and E_B . Since there is no canonical way to map the curves to elements of $\text{Cl}(\mathcal{O})$, it is not enough to be able to solve discrete logarithms in it.

Childs, Jao and Soukharev [9] have shown how to adapt quantum algorithms by Regev [58] and Kuperberg [45] to solve the ordinary isogeny path problem in subexponential time. Although their attack does not qualify as a total break, it makes the Rostovtsev-Stolbunov protocol even less practical. Indeed, the protocol is already very slow, mainly due to the relatively large size of the isogeny degree set L . If parameter sizes must be further enlarged to protect against quantum attacks, it seems plausible that the Rostovtsev-Stolbunov protocol may never be used in practice.

14.2 Supersingular Isogeny Diffie-Hellman

We finally come to the last cryptographic construction from isogeny graphs. Compared to the ordinary case, graphs of supersingular isogenies have two attractive features for constructing key exchange protocols. First, one isogeny degree is sufficient to obtain an expander graph; by choosing one small prime degree, we have the opportunity to construct more efficient protocols. Second, there is no action of an Abelian group, such as $\text{Cl}(\mathcal{O})$, on them; it thus seems harder to use quantum computers to speed up the supersingular isogeny path problem.

But the absence of a group action also makes it impossible to directly generalize the Rostovtsev-Stolbunov protocol to supersingular graphs. It turns out, however, that there is an algebraic structure acting on supersingular graphs. We have seen that, if E is a supersingular curve defined over \mathbb{F}_p or \mathbb{F}_{p^2} , its endomorphism ring is isomorphic to an order in the quaternion algebra $\mathbb{Q}_{p,\infty}$




Figure 15: Supersingular isogeny graphs of degree 2 (left, blue) and 3 (right, red) on \mathbb{F}_{972} .

$$\begin{array}{l}
 \ker \alpha = \langle A \rangle \subset E[\ell_A^{e_A}] \\
 \ker \beta = \langle B \rangle \subset E[\ell_B^{e_B}] \\
 \ker \alpha' = \langle \beta(A) \rangle \\
 \ker \beta' = \langle \alpha(B) \rangle
 \end{array}
 \quad
 \begin{array}{ccc}
 E & \xrightarrow{\alpha} & E/\langle A \rangle \\
 \beta \downarrow & & \downarrow \beta' \\
 E/\langle B \rangle & \xrightarrow{\alpha'} & E/\langle A, B \rangle
 \end{array}$$

Figure 16: 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.

ramified at p and at infinity. There is more: supersingular curves are in correspondence with the maximal orders of $\mathbb{Q}_{p,\infty}$, and their left ideals act on the graph like isogenies. It would be rather technical to go into the details of the theory of quaternion algebras and their maximal orders; instead, we describe the key exchange protocol using only the language of isogenies, with the *caveat* that its security can only be properly evaluated by also looking at its quaternion counterpart. The interested reader will find more details on quaternion algebras in [74, 56, 57, 43, 4, 44].

The key idea of the Supersingular Isogeny Diffie-Hellman protocol (SIDH), first proposed in [37], is to let Alice and Bob take random walks in two distinct isogeny graphs on the same vertex set. In practice, we choose a *large enough* prime p , and two *small* primes ℓ_A and ℓ_B . The vertex set is going to consist of the supersingular j -invariants defined over \mathbb{F}_{p^2} , Alice's graph is going to be made of ℓ_A -isogenies, while Bob is going to use ℓ_B -isogenies. Figure 15 shows a toy example of such graphs, where $p = 97$, $\ell_A = 2$ and $\ell_B = 3$.

Even this, though, is not sufficient to define a key exchange protocol, because there is no canonical way of labeling the edges of these graphs. We shall introduce, then, a very *ad hoc* construction leveraging the group structure of elliptic curves. Recall that a separable isogeny is uniquely defined by its kernel, and that in this case $\deg \phi = \# \ker \phi$. More precisely, a walk of length e_A in the ℓ_A -isogeny graph corresponds to a kernel of size $\ell_A^{e_A}$; and this kernel is cyclic if and only if the walk *does not backtrack*.

Hence, Alice choosing a secret walk of length e_A is equivalent to her choosing a secret cyclic subgroup $\langle A \rangle \subset E[\ell_A^{e_A}]$. If we let Alice choose one such subgroup, and Bob choose similarly 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 16.

At this point, we would like to define a protocol where Alice and Bob choose random cyclic subgroups $\langle A \rangle$ and $\langle B \rangle$ in some *large enough* torsion groups, and exchange enough information to both compute $E/\langle A, B \rangle$ (up to isomorphism), without revealing their respective secrets. We are faced with two difficulties, though:

1. The points of $\langle A \rangle$ (or $\langle B \rangle$) may not be rational. Indeed, in general they may be defined over a field extension of degree as large as $\ell_A^{e_A}$, thus requiring an exponential amount of information to be explicitly represented.
2. The diagram in Figure 16 shows no way by which Alice and Bob could compute $E/\langle A, B \rangle$ without revealing their secrets to each other.

We will solve both problems by carefully controlling the group structure of our supersingular curves. This is something that is very hard to do in the ordinary case, but totally elementary in the supersingular one, as the following proposition shows.

Theorem 54 (Group structure of supersingular curves). *Let p be a prime, and let E be a supersingular curve defined over a finite field \mathbb{F}_q with $q = p^m$ elements. Let t be the trace of the Frobenius endomorphism of E/k , then one of the following is true:*

- *m is odd and*
 - $t = 0$, or
 - $p = 2$ and $t^2 = 2q$, or
 - $p = 3$ and $t^2 = 3q$;
- *m is even and*
 - $t^2 = 4q$, or
 - $t^2 = q$, and $j(E) = 0$, and E is not isomorphic to $y^2 = x^3 \pm 1$, or
 - $t^2 = 0$, and $j(E) = 1728$, and E is not isomorphic to $y^2 = x^3 \pm x$.

The group structure of $E(\mathbb{F}_q)$ is one of the following:

- *If $t^2 = q, 2q, 3q$, then $E(\mathbb{F}_q)$ is cyclic;*
- *If $t = 0$, then $E(\mathbb{F}_q)$ is either cyclic, or isomorphic to $\mathbb{Z}/\frac{q+1}{2}\mathbb{Z} \oplus \mathbb{Z}/2\mathbb{Z}$;*
- *If $t = \mp 2\sqrt{q}$, then $E(\mathbb{F}_q) \simeq (\mathbb{Z}/(\sqrt{q} \pm 1)\mathbb{Z})^2$.*

Proof. See [74, 50]. □

Of all the cases, the only one we are concerned with is $q = p^2$, and $E(\mathbb{F}_q) \simeq (\mathbb{Z}/(p \pm 1)\mathbb{Z})^2$. Since we have full control on p , we can choose it so that $E(\mathbb{F}_q)$ contains two large subgroups $E[\ell_A^{e_A}]$ and $E[\ell_B^{e_B}]$ of coprime order. Hence, once $\ell_A^{e_A}$ and $\ell_B^{e_B}$ are fixed, we look for a prime of the form $p = \ell_A^{e_A} \ell_B^{e_B} f \mp 1$, where f is a small cofactor. In practice, such primes are abundant, and we can easily take $f = 1$. This solves the first problem: $E(\mathbb{F}_q)$ now contains $\ell_A^{e_A-1}(\ell_A + 1)$ cyclic subgroups of order $\ell_A^{e_A}$, each defining a distinct isogeny; hence, a single point $A \in E(\mathbb{F}_q)$ is enough to represent an isogeny walk of length e_A .

The second problem is solved by a very peculiar trick, which sets SIDH apart from other isogeny 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

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	$\alpha : E \rightarrow E_A = E/\langle A \rangle$	$\beta : E \rightarrow E_B = E/\langle B \rangle$
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 17: Supersingular Isogeny Diffie-Hellman key exchange protocol.

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:

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

To start the protocol, they choose random secret subgroups

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

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

$$\begin{aligned} \alpha : E &\rightarrow E/\langle A \rangle, \\ \beta : E &\rightarrow E/\langle B \rangle. \end{aligned}$$

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 $\alpha' : E/\langle B \rangle \rightarrow E/\langle A, B \rangle$, 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 it is conjectured that they do not give any advantage in computing α' 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 17, and schematized in Figure 18.

We end with a discussion on parameter sizes. It is clear that the key space of SIDH depends on the size of the subgroups $E[\ell_A^{e_A}]$ and $E[\ell_B^{e_B}]$, hence we must take $\ell_A^{e_A} \sim \ell_B^{e_B}$ so to make attacks equally hard against Alice or Bob's public data. However this puts serious constraints on the isogeny walks performed in SIDH. Indeed, we have seen that the size of the supersingular isogeny graph is $O(p)$, whereas the size of Alice's (or Bob's) public key space is only $O(\sqrt{p})$. Said otherwise, Alice and Bob take random walks *much shorter* than the diameter of the graph. At the moment, it is not clear how this affects the security of the protocol.




Figure 18: Schematics of SIDH key exchange. Quantities only known to Alice are drawn in blue, quantities only known to Bob in red.

To choose an appropriate size for p , we start by looking at attacks that only use the j -invariants published by Alice and Bob. Given curves E and E_A , connected by an isogeny of degree $\ell_A^{e_A}$, an easy variation on the meet-in-the-middle paradigm finds the secret isogeny in $O(\ell_A^{e_A/2})$ steps (and $O(\ell_A^{e_A/2})$ storage) as follows: tabulate all possible walks of length $\lfloor e_A/2 \rfloor$ starting from E , then iterate over the walks of length $\lceil e_A/2 \rceil$ starting from E_A , until a collision is found. The same collision can also be found with $O(\ell_A^{e_A/3})$ queries to a quantum oracle, using a quantum algorithm due to Tani [68]. Because the isogeny walks are shorter than the diameter, we expect to find only one collision, and that is precisely Alice's secret isogeny.

It turns out these are the best known attacks against SIDH, even taking into account the additional information passed by Alice and Bob. Hence, taking $\log_2 p = n$ offers a classical security of $\sim n/4$ bits, and a quantum security of $\sim n/6$ qubits. In conclusion, to obtain a 128-qubit and 192-bit secure system, we would have to find a 768-bit prime of the form $p = \ell_A^{e_A} \ell_B^{e_B} f \pm 1$, with $e_A \log_2 \ell_A \sim e_B \log_2 \ell_B \sim 384$. In practice, we usually take $\ell_A = 2$ and $\ell_B = 3$ for efficiency reasons, and an example of one such prime is $p = 2^{387} 3^{242} - 1$.

15 Further topics in isogeny based cryptography

We conclude these notes with a brief overview of the current research topics in isogeny based cryptography. We only focus on constructions derived from supersingular isogenies, as they currently are the most promising ones.

Efficient implementation of SIDH What makes SIDH interesting is its relatively good efficiency, especially when compared with other isogeny-based protocols. However, several optimizations are required in order to achieve a compact and fast implementation, competitive with other post-quantum key-exchange candidates. In short, one must optimize each of these levels:

- The arithmetic of \mathbb{F}_p benefits from the special form of p , especially for primes of the form $p = 2^a 3^b - 1$, as explained in [12, 41, 6];

- The arithmetic of \mathbb{F}_{p^2} benefits from the fact that -1 is not a square in \mathbb{F}_p , whenever $p \equiv -1 \pmod{2}$;
- The arithmetic of elliptic curves benefits from using Montgomery models, and optimized formulas for doublings, triplings and isogenies [22, 12, 10];
- Field inversions can be avoided using projective coordinates and *projectivized curve equations* [12];
- The full computation and evaluation of the secret isogeny from a generator of its kernel must be performed using a quasi-linear algorithm first described in [22].

Undoubtedly, the latter is the most novel and surprising of the optimizations. For lack of space, we do not describe any of them here, and we primarily address the interested reader to [22] and [12].

By putting together all the optimizations mentioned above, the SIDH scheme can be made relatively practical, as shown in [22, 12], although one of two order of magnitude slower than other post-quantum competitors. Where SIDH really excel, is in its *very short key sizes*, actually the shortest among post-quantum candidates, at the time of writing. This key size can be shrunk even more through *key compression* techniques [3, 11]. However, the size of the isogeny graph in SIDH is much larger than the size of the key space, it is thus, in principle, possible to make even shorter keys; how to do this efficiently is still an open question.

Security of SIDH We can formally state the security of SIDH as a hardness assumption on a problem called SSDDH. As mentioned previously, the best known algorithms for SSDDH have exponential complexity, even on a quantum computer.

Problem 6 (Supersingular Decision Diffie-Hellman). Let $E, \ell_A, \ell_B, e_A, e_B, P_A, Q_A, P_B, Q_B$ be the parameters of an SIDH protocol.

Given a tuple sampled with probability $1/2$ from one of the following two distributions:

1. $(E/\langle A \rangle, \phi(P_B), \phi(Q_B), E/\langle B \rangle, \psi(P_A), \psi(Q_A), E/\langle A, B \rangle)$, where
 - $A \in E$ is a uniformly random point of order $\ell_A^{e_A}$,
 - $B \in E$ is a uniformly random point of order $\ell_B^{e_B}$,
 - $\phi : E \rightarrow E/\langle A \rangle$ is the isogeny of kernel $\langle A \rangle$, and
 - $\psi : E \rightarrow E/\langle B \rangle$ is the isogeny of kernel $\langle B \rangle$;
2. $(E/\langle A \rangle, \phi(P_B), \phi(Q_B), E/\langle B \rangle, \psi(P_A), \psi(Q_A), E/\langle C \rangle)$, where A, B, ϕ, ψ are as above, and where $C \in E$ is a uniformly random point of order $\ell_A^{e_A} \ell_B^{e_B}$;

determine from which distribution the tuple is sampled.

Assuming SSDDH is hard, we can formally prove the security of the key exchange against *passive adversaries*, i.e., those adversaries who can see all messages sent between Alice and Bob, but who do not modify them. We address the interested reader to [22] for the technical details.

It is apparent that SSDDH is a very special instance of the isogeny path problem; it is thus conceivable that specially crafted algorithms could break SIDH without solving the generic isogeny path problem. As an illustration, consider the following problem.

Problem 7. Let $E, \ell_A, \ell_B, e_A, e_B, P_A, Q_A, P_B, Q_B$ be the parameters of an SIDH protocol.

Let $A \in E$ be a point of order $\ell_A^{e_A}$, and let $\phi : E \rightarrow E/\langle A \rangle$. Given $E/\langle A \rangle, \phi(P_B)$ and $\phi(Q_B)$ compute $\phi(R)$ for an arbitrary point $R \in E$ of order $\ell_A^{e_A}$.

It is easy to verify that solving this problem immediately reveals the secret $\langle A \rangle$. Indeed, $\phi(R)$ is an element of $\ker \hat{\phi}$, from which we can recover $\hat{\phi}$ and ϕ , and thus $\langle A \rangle$. An efficient solution to this problem completely breaks SIDH, without doing anything for the generic isogeny path problem.²

And indeed, although the security of SSDDH is still unblemished at the time of writing, several polynomial-time attacks have appeared against variations of SIDH. The interested reader will find more details in the following references:

- A key-recovery attack against a *static key* version of SIDH, where Alice uses a long term secret isogeny [32];
- Key-recovery attacks in various *leakage models* [32, 35, 71];
- Key recovery attacks against some *unbalanced* variants of SIDH [54].

Finally, it is worth mentioning that there is a quantum subexponential attack [5] in the case where both E and $E/\langle A \rangle$ are defined over \mathbb{F}_p .

Other protocols Key exchange is not the only public-key protocol that can be derived from isogeny graphs. It is easy, for example, to derive a public-key encryption protocol similar to El Gamal from either the Rostovtsev-Stolbunov protocol or SIDH. We illustrate the second:

- Alice’s secret key is an isogeny $\alpha : E \rightarrow E/\langle A \rangle$; her public key contains $E/\langle A \rangle$ and the evaluation of α on *Bob’s basis* $\langle P_B, Q_B \rangle$.
- To encrypt a message m , Bob chooses a random $\beta : E \rightarrow E/\langle B \rangle$, and computes the shared secret $E/\langle A, B \rangle$, which he converts to a binary string s (e.g., by hashing the j invariant of $E/\langle A, B \rangle$); he sends to Alice the message $(E/\langle B \rangle, \beta(P_A), \beta(Q_A), m \oplus s)$.
- To decrypt, Alice uses $E/\langle B \rangle, \beta(P_A), \beta(Q_A)$ to compute the shared secret $E/\langle A, B \rangle$, which she converts to s , and finally she un.masks $m \oplus s$.

In [22], it is proven that this protocol is IND-CPA secure under the SSDDH assumption. Achieving IND-CCA2 security is harder, as the attack against static keys in [32] shows, however it is possible to apply a generic transformation to obtain an IND-CCA2 secure *key encapsulation mechanism*.

One may expect that digital signatures would also generalize easily to the isogeny setting, but both Schnorr signatures and ECDSA rely on the existence of a group law on the public data, something that is missing both in the ordinary and in the supersingular case.

To our rescue, comes a zero-knowledge protocol based on the same construction shown in Figure 16. In this protocol, Alice’s secret key is an isogeny $\alpha : E \rightarrow E/\langle A \rangle$; her public key is the curve $E/\langle A \rangle$, together with a description of the action of α on $E[\ell_B^{e_B}]$, as in SIDH. To prove knowledge of α to Bob, she takes a random subgroup $\langle B \rangle \subset E[\ell_B^{e_B}]$, computes a commutative diagram as in Figure 16, and sends to Bob the curves $E/\langle B \rangle$ and $E/\langle A, B \rangle$. To verify that Alice knows the secret, Bob asks her one of two questions at random:

- either reveal the point B and its image $\alpha(B)$,
- or reveal the point $\beta(A)$.

²The converse reduction is not evident either: given an oracle solving the isogeny path problem, how can we break SIDH? A partial answer is given in [44, 32], where it is shown that, knowing the endomorphism rings of E and $E/\langle A \rangle$, an attacker can solve the isogeny path problem, and then break SIDH, in polynomial time.

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 $\langle P_B, Q_B \rangle$ of $E[\ell_B^{e_B}]$.	
Secret key	An isogeny $\alpha : E \rightarrow E/\langle A \rangle$ of degree $\ell_A^{e_A}$.	
Public key	The curve $E/\langle A \rangle$, the images $\alpha(P_B), \alpha(Q_B)$.	
	Alice	Bob
Pick random	$B \in E[\ell_B^{e_B}]$ of order $\ell_B^{e_B}$	
Compute masking isogeny	$\beta : E \rightarrow E/\langle B \rangle$	
Commit	$(E/\langle B \rangle, E/\langle A, B \rangle) \rightarrow$	
Challenge		$\leftarrow b \in \{0, 1\}$
Reveal	if $b = 0$, send $(B, \alpha(B))$ if $b = 1$, send $\beta(A) \rightarrow$	

Figure 19: Supersingular Isogeny Zero-Knowledge Identification protocol.

After receiving Alice’s answer, he accepts only if the points do define isogenies between the curves $E, E/\langle A \rangle, E/\langle B \rangle, E/\langle A, B \rangle$ as expected. The protocol is summarized in Figure 19.

Intuitively, if Alice respects the protocol, she always succeeds in convincing Bob. If she cheats, she only has one chance out of two of guessing Bob’s challenge and succeed in tricking him. Thus, by iterating the protocol a sufficient number of times, a cheater’s chance of success can be made arbitrarily small at exponential pace. The protocol is zero-knowledge because revealing B and $\alpha(B)$ does not reveal anything that Bob does not already know. Revealing $\beta(A)$ is trickier, and we need to make one more security assumption, named Decisional Supersingular Product (DSSP), to prove zero knowledge. In [22] it is proven that this protocol is secure and zero-knowledge under the SSDDH³ and DSSP assumptions.⁴

Using a generic construction, such as the Fiat-Shamir heuristic [27], it is possible to derive a signature scheme from the zero-knowledge protocol above. Alternative signature schemes based on the same construction, with different desirable properties, are presented in [33, 75]. However, all these protocols suffer from the high cost of having to iterate hundreds of times the basic building block of Figure 19. Obtaining an efficient signature scheme from isogeny assumptions is still an open problem.

More protocols can be obtained by slightly generalizing the SIDH construction. If we allow the prime to be of the form $p = \ell_A^{e_A} \ell_B^{e_B} \ell_C^{e_C} \pm 1$, we can construct a commutative cube in the same way the square of Figure 16 was constructed. Using primes of this form, Sun, Tian and Wang have proposed a *strong designated verifier* signature scheme [66]. Adding one more prime ℓ_D in the mix, Jao and Soukharev have proposed *undeniable* signatures [39]. The drawback of all these schemes is that, as we add more torsion subgroups to the base curve, the size of the primes grows, making the schemes less and less practical.

In general, isogeny graphs are much less flexible than the classical discrete logarithm problem. Many of the protocols that have been built on discrete logarithms fail to be ported to isogeny based cryptography. Devising new post-quantum protocols, retaining some of the desirable properties of classical ones, is a very active area of research in isogeny based cryptography.

³Actually, a weaker assumption named CSSI.

⁴The paper [22] also hints at a variant of the zero-knowledge protocol where Bob challenges Alice to open one out of three commitments, namely one of $B, \alpha(B), \beta(A)$. This variant is less efficient, since a cheater has 2/3 chances of success, however its security relies on the stronger isogeny walk problem, rather than on SSDDH.

Exercices

Exercice III.1. Prove Proposition 40.

Exercice III.2. Show that a Schreier graph $(S \subset G, X)$ is an ε -expander if and only if S generates G .

Exercice III.3. Derive encryption protocols *à la* El Gamal from the key exchange protocols of Section 14.

References

- [1] Arthur O. L. Atkin. The number of points on an elliptic curve modulo a prime. 1988.
- [2] Arthur O. L. Atkin. The number of points on an elliptic curve modulo a prime. <http://www.lix.polytechnique.fr/Labo/Francois.Morain/AtkinEmails/19910614.txt>, 1991.
- [3] Reza Azarderakhsh, David Jao, Kassem Kalach, Brian Koziel, and Christopher Leonardi. Key compression for isogeny-based cryptosystems. In *Proceedings of the 3rd ACM International Workshop on ASIA Public-Key Cryptography*, pages 1–10. ACM, 2016.
- [4] Juliana V. Belding. *Number Theoretic Algorithms for Elliptic Curves*. PhD thesis, University of Maryland, 2008.
- [5] 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.
- [6] Joppe W. Bos and Simon Friedberger. Fast arithmetic modulo $2^x p^y \pm 1$. Cryptology ePrint Archive, Report 2016/986, 2016. <http://eprint.iacr.org/2016/986>.
- [7] Alin Bostan, François Morain, Bruno Salvy, and Éric Schost. Fast algorithms for computing isogenies between elliptic curves. *Math. Comp.*, 77:1755–1778, September 2008.
- [8] 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.
- [9] Andrew Childs, David Jao, and Vladimir Soukharev. Constructing elliptic curve isogenies in quantum subexponential time. *Journal of Mathematical Cryptology*, 8(1):1–29, 2014.
- [10] Craig Costello and Huseyin Hisil. A simple and compact algorithm for sidh with arbitrary degree isogenies. Cryptology ePrint Archive, Report 2017/504, 2017. <http://eprint.iacr.org/2017/504>.
- [11] Craig Costello, David Jao, Patrick Longa, Michael Naehrig, Joost Renes, and David Urbanik. *Efficient Compression of SIDH Public Keys*, pages 679–706. Springer International Publishing, Cham, 2017.
- [12] Craig Costello, Patrick Longa, and Michael Naehrig. Efficient algorithms for Supersingular Isogeny Diffie-Hellman. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016: 36th Annual International Cryptology Conference*, pages 572–601. Springer Berlin Heidelberg, 2016.
- [13] Jean-Marc Couveignes. *Quelques calculs en théorie des nombres*. PhD thesis, Université de Bordeaux, 1994.
- [14] 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.
- [15] Jean-Marc Couveignes. Isomorphisms between Artin-Schreier towers. *Math. Comp.*, 69(232):1625–1631, 2000.
- [16] Jean-Marc Couveignes. Hard homogeneous spaces. <http://eprint.iacr.org/2006/291/>, 2006.

- [17] Jean-Marc Couveignes and Reynald Lercier. Fast construction of irreducible polynomials over finite fields. *To appear in the Israel Journal of Mathematics*, July 2011.
- [18] David A Cox. *Primes of the form $x^2 + ny^2$: Fermat, class field theory, and complex multiplication*, volume 34. John Wiley & Sons, 2011.
- [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 *ISSAC'13*, pages 165–172. ACM, 2013.
- [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 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.
- [24] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, 1976.
- [25] Noam D. Elkies. Explicit isogenies. 1992.
- [26] 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.
- [27] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Conference on the Theory and Application of Cryptographic Techniques*, pages 186–194. Springer, 1986.
- [28] 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.
- [29] Steven D. Galbraith. Constructing isogenies between elliptic curves over finite fields. *LMS Journal of Computation and Mathematics*, 2:118–138, 1999.
- [30] Steven D Galbraith. *Mathematics of public key cryptography*. Cambridge University Press, 2012. <https://www.math.auckland.ac.nz/~sgal018/crypto-book/crypto-book.html>.
- [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 Comput. Sci.*, pages 29–44. Springer, Berlin, 2002.

- [32] Steven D Galbraith, Christophe Petit, Barak Shani, and Yan Bo Ti. On the security of supersingular isogeny cryptosystems. <http://eprint.iacr.org/2016/859>, 2016. To appear at AsiaCrypt 2016.
- [33] Steven D. Galbraith, Christophe Petit, and Javier Silva. Signature schemes based on supersingular isogeny problems. Cryptology ePrint Archive, Report 2016/1154, 2016. <http://eprint.iacr.org/2016/1154>.
- [34] 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.
- [35] Alexandre Gélín and Benjamin Wesolowski. Loop-abort faults on supersingular isogeny cryptosystems. In *International Workshop on Post-Quantum Cryptography*, pages 93–106. Springer, 2017.
- [36] Sorina Ionica and Antoine Joux. Pairing the volcano. *Mathematics of Computation*, 82(281):581–603, 2013.
- [37] 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.
- [38] David Jao, Stephen D. Miller, and Ramarathnam Venkatesan. Expander graphs based on GRH with an application to elliptic curve cryptography. *J. Number Theory*, 129(6), 2009.
- [39] David Jao and Vladimir Soukharev. Isogeny-based quantum-resistant undeniable signatures. In *International Workshop on Post-Quantum Cryptography*, pages 160–179. Springer, 2014.
- [40] Antoine Joux. *Algorithmic cryptanalysis*. CRC Press, 2009.
- [41] Angshuman Karmakar, Sujoy Sinha Roy, Frederik Vercauteren, and Ingrid Verbauwhede. Efficient finite field multiplication for isogeny based post quantum cryptography. *Proceedings of WAIFI 2016*, 2016.
- [42] Neal Koblitz. Elliptic curve cryptosystems. *Math. Comp.*, 48(177):203–209, 1987.
- [43] David Kohel. *Endomorphism rings of elliptic curves over finite fields*. PhD thesis, University of California at Berkley, 1996.
- [44] David 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.
- [45] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM J. Comput.*, 35(1):170–188, 2005.
- [46] Serge Lang. Elliptic functions. In *Elliptic Functions*, pages 5–21. Springer, 1987.
- [47] Hendrik W. Lenstra. Factoring integers with elliptic curves. *Annals of Mathematics*, 126:649–673, 1987.
- [48] Reynald Lercier. *Algorithmique des courbes elliptiques dans les corps finis*. PhD thesis, LIX - CNRS, June 1997.

- [49] 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.
- [50] Alfred Menezes, Scott Vanstone, and Tatsuaki Okamoto. Reducing elliptic curve logarithms to logarithms in a finite field. In *STOC '91: Proceedings of the twenty-third annual ACM symposium on Theory of computing*, pages 80–89, New York, NY, USA, 1991. ACM.
- [51] 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.
- [52] Josep M. Miret, Ramiro Moreno, Ana Rio, and Magda Valls. Determining the 2-sylow subgroup of an elliptic curve over a finite field. *Mathematics of Computation*, 74(249):411–427, 2005.
- [53] Josep M. Miret, Ramiro Moreno, Daniel Sadornil, Juan Tena, and Magda Valls. An algorithm to compute volcanoes of 2-isogenies of elliptic curves over finite fields. *Applied Mathematics and Computation*, 176(2):739–750, 2006.
- [54] Christophe Petit. Faster algorithms for isogeny problems using torsion point images. Cryptology ePrint Archive, Report 2017/571, 2017. <http://eprint.iacr.org/2017/571>.
- [55] 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.
- [56] Arnold K. Pizer. Ramanujan graphs and Hecke operators. *Bull. Amer. Math. Soc. (N.S.)*, 23(1), 1990.
- [57] 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.
- [58] Oded Regev. A subexponential time algorithm for the dihedral hidden subgroup problem with polynomial space, June 2004.
- [59] Alexander Rostovtsev and Anton Stolbunov. Public-key cryptosystem based on isogenies. <http://eprint.iacr.org/2006/145/>, 2006.
- [60] René Schoof. Elliptic curves over finite fields and the computation of square roots mod p . *Math. Comp.*, 44(170):483–494, 1985.
- [61] René Schoof. Counting points on elliptic curves over finite fields. *Journal de Théorie des Nombres de Bordeaux*, 7(1):219–254, 1995.
- [62] 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.
- [63] Joseph H. Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1992.
- [64] Joseph H. Silverman. *Advanced Topics in the Arithmetic of Elliptic Curves (Graduate Texts in Mathematics)*. Springer, January 1994.

- [65] Anton Stolbunov. Constructing public-key cryptographic schemes based on class group action on a set of isogenous elliptic curves. *Adv. Math. Commun.*, 4(2), 2010.
- [66] Xi Sun, Haibo Tian, and Yumin Wang. Toward quantum-resistant strong designated verifier signature from isogenies. In *2012 Fourth International Conference on Intelligent Networking and Collaborative Systems*, 2012.
- [67] Andrew V. Sutherland. Genus 1 point counting over prime fields. Last accessed July 16, 2010. <http://www-math.mit.edu/~drew/SEArecords.html>, 2010.
- [68] Seiichiro Tani. Claw finding algorithms using quantum walk. *Theoretical Computer Science*, 410(50):5285–5297, 2009.
- [69] Terence Tao. Expansion in groups of lie type – basic theory of expander graphs. <https://terrytao.wordpress.com/2011/12/02/245b-notes-1-basic-theory-of-expander-graphs/>, 2011.
- [70] Edlyn Teske. An elliptic curve trapdoor system. *Journal of Cryptology*, 19(1):115–133, January 2006.
- [71] Yan Bo Ti. Fault attack on supersingular isogeny cryptosystems. In *International Workshop on Post-Quantum Cryptography*, pages 107–122. Springer, 2017.
- [72] 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.
- [73] Jean Vélu. Isogénies entre courbes elliptiques. *Comptes Rendus de l'Académie des Sciences de Paris*, 273:238–241, 1971.
- [74] William C. Waterhouse. Abelian varieties over finite fields. *Annales Scientifiques de l'École Normale Supérieure*, 2(4):521–560, 1969.
- [75] Youngho Yoo, Reza Azarderakhsh, Amir Jalali, David Jao, and Vladimir Soukharev. A post-quantum digital signature scheme based on supersingular isogenies. Cryptology ePrint Archive, Report 2017/186, 2017. <http://eprint.iacr.org/2017/186>.