ISOGENY GRAPH CRYPTOGRAPHY

TAE KYU KIM

Abstract. This paper is intended to provide the necessary background in elliptic curves and graph theory to understand the idea behind isogeny graph cryptography, specifically Supersingular Isogeny Diffie-Hellman. It will not cover implementation details on each cryptographic scheme nor proofs for most propositions and theorems, though a list references will be given at the end of the paper. Much of the material is taken from De Feo's "Mathematics of Isogeny Based Cryptography," an excellent paper describing elliptic curves, expander graphs, and several isogeny based protocols. This paper is part of the Cryptography class of the 2019 fall quarter at Euler Circle, and mainly serves to be another resource that the students can learn from.

1. Elliptic Curves

Before we can talk about isogenies and doing cryptography on elliptic curves, we need to first introduce elliptic curves.

Definition 1.1 (Weierstrass equation)**.** An *elliptic curve* over a field k with characteristic $\neq 2, 3$ is the locus of an equation

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

where a and b is in k and $4a^3 + 27b^2 \neq 0$. This equation is called the *Weierstrass equation.* We also include a point at infinity $\mathcal{O}$. We will denote the set of points on the elliptic curve as $E(k)$.

The condition on a and b will prevent double roots and also guarantee that the j-invariant of an elliptic curve is defined, which is an important characteristic of elliptic curves (See definition 2.1).

Elliptic curves admit a group structure, which is best illustrated over $\mathbb{Q}$. To add points P and Q on the curve, find the intersection of the line through P and Q and the Weierstrass equation (Eliminating y from the two equations gives us a cubic in x; we already have two roots, which guarantees that the third root is in our field.) Reflecting that point over the x-axis gives $P + Q$. If $P = Q$, we instead use the line tangent to the Weierstrass equation at P. We will let $\mathcal{O}$, our identity element, lie on every vertical line. That is, if P and Q have the same x-coordinate, $P + Q = \mathcal{O}$, and for all P, $P + \mathcal{O} = \mathcal{O} + P = P$.

Date: December 9, 2019.

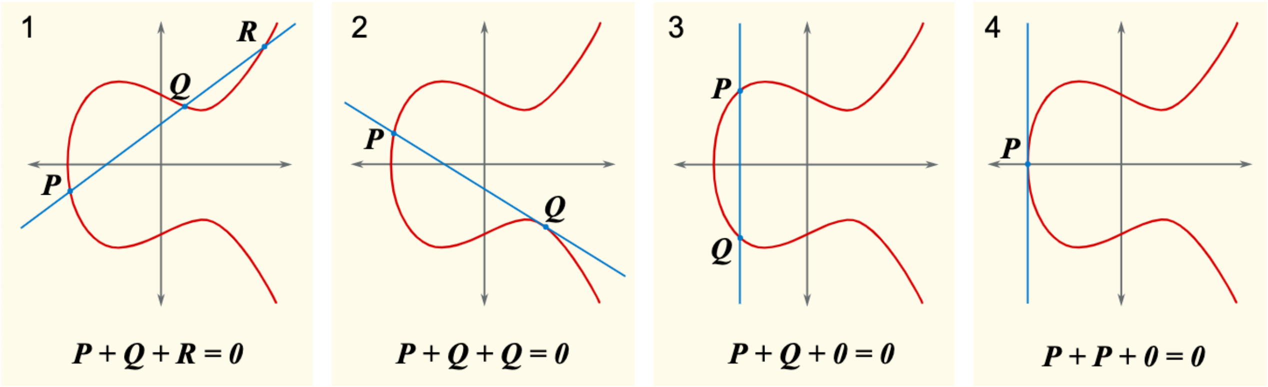

Figure 1. Adding points on $E(\mathbb{Q})$. Image from Wikimedia Commons.

Because we have addition, we can do scalar multiplication by integers, which will be denoted as $[n] : P \mapsto [n]P$. One can explicitly find the formulas to do addition and multiplication (and they are easy to compute), but we will not do so here.

Let's try to understand the group structure on $E(k)$. We can look at the torsion part (finite order) and the free part (infinite order) of the group separately. The torsion part is easily characterized:

Proposition 1.2 (m-torsion group). *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 $E[m]$, is defined as*

$$E[m] = \{[m]P = \mathcal{O} | P \in E(k)\}.$$

It 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*

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

Proof. See [Sil09, Cor 6.4]. ■

For curves over fields of positive characteristic p, the case when $E[p] \simeq \mathbb{Z}/p\mathbb{Z}$ is called *ordinary*, while the case $E[p] \simeq \mathcal{O}$ is called *supersingular*. We will come back to those terms later in the paper.

The free part of the group is more difficult to characterize. Currently we only know that it is finitely generated, and we have some algorithms to compute the ranks of most elliptic curves over number fields. However, we will not need these results in this paper.

2. Maps between elliptic curves

Proposition 2.1 (j-invariant). *Let $E : y^2 = x^3 + ax + b$ be an elliptic curve. 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 $\overline{k}$ if and only if they have the same j-invariant.

Proof. The main idea is that $E : y^2 = x^3 + ax + b$ and $E' : y^2 = x^3 + a'x + b'$ are isomorphic if and only if there exists a $c \in \overline{k}$ such that

$$a' = ac^2 \quad \text{and} \quad b' = bc^3.$$

Note that c may not necessarily be in k, so we have to find c in the algebraic closure $\overline{k}$. ■

Definition 2.2 (isogeny)**.** Let E_1 and E_2 be elliptic curves over k. An *isogeny* over k is a surjective group morphism $\phi : E_1 \to E_2$ over k such that $\phi(\mathcal{O}_{E_1}) = \mathcal{O}_{E_2}$ and that it is also an rational map.

Rational maps require that the functions that send each coordinate in E_1 to coordinates in E_2 is a ratio of polynomials in the coordinates of E_1. We need this condition so that the isogeny also respects the fact that elliptic curves are defined algebraically (i.e. are the zeros of polynomials).

More generally, the degree of an isogeny is defined using category theory. However, we will be mostly working with separable isogenies, and these have an easier notion of degree. The definition of separability requires definitions about function fields from category theory, so we omit its discussion here.

Proposition 2.3. *Let ϕ be a separable isogeny. Then* $\deg \phi = \# \ker \phi$.

Proof. See [Sil09, Thm II.2.4]. ■

Example. If E is an elliptic curve, and $C \subset E$ is a finite subgroup, then E/C is also an elliptic curve and the natural quotient map $E \to E/C$ is a (separable) isogeny. Then the degree of the isogeny is $\#C$.

If a isogeny has degree l, we will refer to it as an l-isogeny. It turns out that separable isogenies are completely determined by their kernel:

Proposition 2.4. *Let E be an elliptic curve, and let G be a finite subgroup of E. There is a unique elliptic curve E_0 and a unique separable isogeny ϕ, such that* $\ker \phi = G$ *and* $\phi : E \to E_0$.

Proof. See [Sil09, Prop III.4.12]. ■

We can have isogenies from E to itself; these are called endomorphisms (except for the zero map, which is the only non-surjective endomorphism, sending everything to 0). One important endomorphism is multiplication-by-m, denoted by

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

Another important endomorphism that isn't the multiplication map is the *Frobenius endomorphism*, which always exists in positive characteristic (base field $\mathbb{F}_p$) and is defined by

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

The set of endomorphisms under addition and composition form a ring.

Definition 2.5 (Endomorphism ring)**.** Denote by $\mathrm{End}(E)$ the set of isogenies from E to E defined over k, including the zero map; this is called the endomorphism ring of E.

A theorem due to Deuring characterizes all possibilities for the structure of the endomorphism ring, but first we need to define some terms from abstract algebra.

Definition 2.6 ($\mathbb{Q}$-algebra)**.** A $\mathbb{Q}$-*algebra* A is a vector space over $\mathbb{Q}$ equipped with a bilinear product. More generally, we can have an k-algebra (or an algebra over a k) for any field k.

A is *finitely generated* if there is a finite subset of the algebra $\{e_1, \dots, e_m\}$ such that every element in the algebra can be expressed as a polynomial in e_i with coefficients in k.

The *dimension* is the dimension of A over k as a vector space: the smallest n such that there exist independent basis elements $e_1, \dots, e_n \in A$ so that every element in A can be represented as $\sum a_i e_i$ for some $a_i \in k$. Independent basis means that a smaller subset of the basis elements cannot generate the remaining basis elements.

Example. Dimension in this context refers only to the vector space structure, not the algebra structure. For example, $\mathbb{Q}(\sqrt[3]{2})$ is generated by 1 element as a $\mathbb{Q}$-algebra, but is 3-dimensional (as a vector space) over $\mathbb{Q}$ (generated by $\{1, \sqrt[3]{2}, \sqrt[3]{4}\}$ as a $\mathbb{Q}$-vector space). As another example, the polynomial ring $\mathbb{R}[x]$ is finitely generated as an $\mathbb{R}$-algebra (generated just by x), but is an infinite-dimensional $\mathbb{R}$-vector space.

Example. The field of complex numbers $\mathbb{C}$ is an algebra over $\mathbb{R}$ as every complex number can be written as $a + bi$ where $a, b \in \mathbb{R}$, and complex multiplication $(\cdot)$ (the bilinear product) satisfies for all $x, y, z \in \mathbb{C}$ and $a, b \in \mathbb{R}$:

- Right distributivity: $(x + y) \cdot z = x \cdot z + y \cdot z$
- Left distributivity: $z \cdot (x + y) = z \cdot x + z \cdot y$
- Compatibility with scalars: $(ax) \cdot (by) = (ab)(x \cdot y)$.

Example. The field of quaternions $\mathbb{H}$ is a 4-dimensional algebra over $\mathbb{R}$; the standard basis is $(1, i, j, k)$. Note that the complex numbers and quaternions are all finitely generated over $\mathbb{R}$.

Definition 2.7 ($\mathbb{Z}$-module)**.** A $\mathbb{Z}$-*module* is like a vector space over $\mathbb{Z}$, except that a vector space requires that the underlying set of scalars is a field, while a module only requires that it is a ring. More generally, we can have an R-module where R is any ring.

Definition 2.8 (Order)**.** Let A be a finitely generated $\mathbb{Q}$-algebra. An *order* $\mathcal{O} \subset A$ is a subring of A that is a finitely generated $\mathbb{Z}$-module of maximal rank (the rank is basically the equivalent of dimension for modules).

Example. $\mathbb{Z}[i]$ is an order in $\mathbb{Q}(i)$.

Example. $\mathbb{Z}[\sqrt{-3}] \subset \mathbb{Z}[\frac{1+\sqrt{-3}}{2}]$ are both orders in $\mathbb{Q}(\sqrt{-3})$.

Theorem 2.9 (Deuring)**.** *Let E be an elliptic curve defined over a field k of characteristic p. The ring* $\operatorname{End}(E)$ *is isomorphic to one of the following:*

- *$\mathbb{Z}$, only if $p = 0$;*
- *An order $\mathcal{O}$ in an imaginary quadratic field (a field of the form $\mathbb{Q}\sqrt{-D}$ for $D > 0$); in this case we say 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 [Sil09, Cor III.9.4]. ■

For every elliptic curve, we have multiplication-by-m isogenies, so necessarily $\mathbb{Z} \subset \operatorname{End}(E)$. An elliptic curve in a positive characteristic always has the non-trivial Frobenius endomorphism, so $\mathbb{Z}[\pi] \subset \operatorname{End}(E)$; hence, E must have complex multiplication. In a positive characteristic, a curve that is not supersingular is called *ordinary.* Remember that finite fields always have positive characteristic. As we won't need ramification for our cryptography, we will not define it here.

Theorem 2.10 (Dual isogeny)**.** *Let $\phi : E \to E'$ be an isogeny of degree m. There is a unique isogeny $\hat{\phi} : E' \to E$ such that*

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

$\hat{\phi}$ is called the dual isogeny of ϕ, where $[m]_E$ and $[m]_{E'}$ refer to the multiplication-by-m isogeny on E and E'. 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 $\psi : E' \to E''$;*
(3) *$\widehat{\psi + \phi} = \hat{\psi} + \hat{\phi}$ for any isogeny $\psi : E' \to E''$;*
(4) $\deg \phi = \deg \hat{\phi}$*;*
(5) *$\hat{\hat{\phi}} = \phi$.*

Proof. See [Sil09, Thm III.6.1, Thm III.6.2]. ■

The dual isogeny is kind of an inverse of an isogeny. However, isogenies have much weaker conditions than isomorphisms. There can be isogenies between elliptic curves of different j-invariants, and the isogeny composed with its dual isogeny does not give the identity map over the elliptic curve, while an isomorphism composed with its inverse definitely does.

3. Graph Theory and Expander Graphs

We have covered a large portion of the background necessary to talk about isogenies, but we still need to talk about graphs. The following are some definitions that the reader should be familiar with.

Definition 3.1. 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* if $\{v, v'\} \in E$. The neighbors of a vertex are all vertices in V connected to it by an edge. A *path* between two vertices v and v' is a sequence of vertices $v \to v_1 \to \ldots \to v'$ such that consecutive vertices are connected by an edge. The distance between two edges is the length of the shortest path between them (If no path exists, the distance is infinite). A graph is *connected* if any two vertices have a path connecting them; it is *disconnected* otherwise. The *diameter* of a connected graph is the largest of all distances between its vertices. The *degree* of a vertex is the number of edges pointing into (or from) it; a graph where every vertex has degree k is called *k-regular*. The *adjacency matrix* of graph G with vertex set $V = \{v_1, \ldots v_n\}$ 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 graph is undirected, our adjacency matrix is symmetric. A general result is that symmetric matrices have n real eigenvalues $\lambda_1 \geq \ldots \geq \lambda_n$. For k-regular graphs, we have the following bound on the eigenvalues.

Proposition 3.2. *If G is a k-regular graph with eigenvalues $\lambda_1 \geq \ldots \geq \lambda_n$, then*

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

Proof. See [Tao11, Lem 2]. ■

Definition 3.3 (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

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

Why do we care about expander graphs? It turns out that they have a lot of applications in theoretical computer science due to their *pseudo-randomness*. What will be important for us is that they have *short diameter* and *rapidly mixing walks*. The diameter is bounded by $O(\log \# V)$, where the constant only depends on k and ε. Rapidly mixing means that any sufficiently long random walk from a vertex will land you on any other node with close to uniform probability (and the necessary length of the walk to guarantee this is relatively short).

Proposition 3.4 (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 [JMV09, Cor 1.3]. ■

We will conclude this section with a result relating graphs of supersingular curves with l-isogenies to Ramanujan graphs, which is a special type of expander graphs with greatest possible difference $\lambda_1 - \max(|\lambda_2|, |\lambda_n|)$ by having $\max(|\lambda_2|, |\lambda_n|) \approx 2\sqrt{k-1}$.

Theorem 3.5 (All 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*

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

   isomorphism classes of supersingular elliptic curves over $\overline{\mathbb{F}}_p$;
3. *The graph of supersingular curves in $\overline{\mathbb{F}}_p$ with l-isogenies is connected, $l+1$ regular, and has the Ramanujan property.*

Proof. See [Sil09, Thm V.4.1] ■

For the next theorem, an isogeny between E, E' is horizontal if $\mathrm{End}(E) \simeq \mathrm{End}(E')$.

Theorem 3.6 (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 an imaginary quadratic field. Let G be the graph which vertices are elliptic curves over $\mathbb{F}_q$ with complex multiplication by $\mathcal{O}$, and which edges are (horizontal) isogenies of prime degree bounded by $(\log q)^{2+\delta}$ for some fixed $\delta > 0$. Assume that G is nonempty. 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 [JMV09]. ■

4. Isogeny graph problems

Definition 4.1. For any prime $l \neq p$, we can construct an isogeny graph, a multi-graph (we allow multiple edges between two vertices and self-loops) in which nodes are the j-invariants of isogenous curves and edges are isogenies of degree l between them. The dual isogeny theorem implies that for every isogeny there is a corresponding reverse isogeny. For this reason, an isogeny graph is usually drawn undirected.

The following problem is considered difficult, and the security of the protocols will rely on the hardness of it.

Problem 4.2 (Isogeny path)**.** *Given two elliptic curves E, E' over a finite field K such that $\#E = \#E'$, find an isogeny $\phi : E \to E'$ of smooth degree.*

"Smooth degree" means that the degree has only small prime factors (usually we have a bound B for the "smallness" we require). This problem is considered very difficult, and the general method of attack is by having random walks from both E and E', and by the birthday paradox, the paths are expected to meet after $O\sqrt{\#G}$ steps. Intuitively, a random walk is like picking random points on the graph.

5. Hash functions and Diffie-Hellman

This section will introduce ideas of using random walks on expander graphs and isogeny graphs to encrypt information. The mixing properties of expander graphs made them good pseudo-random number generators, so they can be used to make very good hash functions.

Suppose that our graph was 3-regular and we have a binary string to encode. Start any fixed vertex. At each step, read next bit of the string and use it to determine which of the edges to traverse, avoiding the edge that goes back to the previous vertex. We also want to a deterministic way to do choose the edges at each step, so when we actually do random walks on isogeny graphs, we'll have to be selective about which isogenies we allow as edges.

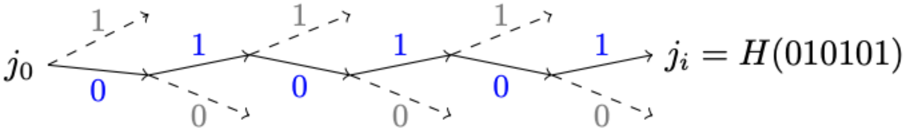

Figure 2. Hashing the string 010101 using an expander graph. Image from [Feo17].

We can describe a prototypical protocol of Diffie-Hellman using random walks. Say we have a cyclic graph generated by g with order p i.e. $G = \langle g \rangle$. The group $(\mathbb{Z}/p\mathbb{Z})^\times$ acts on G: we have $p-1$ bijections from G to G, namely taking every element to the ith power for $1 \leq i \leq p-1$.

We need to choose our bijections careful so that we can determine which edge to take when traversing our graph. Create the set $D \subset (\mathbb{Z}/p\mathbb{Z})^\times$ so that $\sigma \in D \implies \sigma^{-1} \notin D$. This will be our "forward" direction for the random walk. The "backward" direction will be D^{-1}, the set of all inverses of elements in D.

If ρ is a directed route, write $\rho(g)$ for the vertex defined by ρ and starting vertex g. If ρ is a route of length m

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

then

$$\rho(g) = g^{\prod \sigma_i}.$$

Hence, any two routes ρ_A, ρ_B commute. This allows for a Diffie-Hellman exchange:

- Public information:
  - A group G of prime order p,
  - A generating set $D \subset (\mathbb{Z}/p\mathbb{Z})^\times$
  - A generating g of G.
- Alice and Bob each pick a secret route ρ_A and ρ_B.
- Alice computes $g_A = \rho_A(g)$ and sends it to Bob.
- Bob computes $g_B = \rho_B(g)$ and sends it to Alice.
- Shared secret: $g_{AB} = \rho_A(g_B) = \rho_B(g_A)$.

The Diffie-Hellman procedure can be carried out on graphs of horizontal isogenies over elliptic curves with complex multiplication. There are a few technicalities with how to pick the edges (we need to pick a list of primes that split over $\mathbb{Z}[\pi]$), but following through with this idea gives the *Rostovtsev-Stolbunov protocol*. However, the protocol is slow as Alice and Bob must keep track of a large number of primes. Furthermore, because there is an abelian group action on our graph (the class group), the protocol may be vulnerable against quantum attacks by solving the discrete logarithm problem over the group action. Childs, Jao, and Soukharev [CJS14] have shown how to adapt quantum algorithms to solve the ordinary isogeny path problem in subexponential time. While this does not completely break the protocol, it seems less plausible as a security system that should be implemented in practice.

6. Supersingular Isogeny Key Exchange

We've looked at isogeny graphs on ordinary graphs, which have elliptic curves with endomorphism rings isomorphic to orders in imaginary quadratic fields $Q\sqrt{-D}$. Now we look at supersingular graphs, which have elliptic curves endomorphism rings isomorphic to orders in quaternion algebras.

One reason that a key exchange protocol on supersingular curves is that the endomorphism ring of supersingular curves are non-abelian, so current quantum attacks on Diffie-Hellman do not break the protocol.

The main idea of the Supersingular Isogeny Diffie-Hellman protocol (SIDH) is to let Alice and Bob take random walks in two different isogeny graphs on the same vertex set. We choose a large prime p and small primes l_A and l_B. Alice's graph is made of degree l_A isogenies, while Bob uses l_B-isogenies. The vertex set will be the j-invariants of supersingular curves defined over $\mathbb{F}_{p^2}$.

Taking a random walk in a deterministic way is difficult because there is no canonical way to label the edges. Instead, we can use the fact that separable isogenies are completely determined by their kernel and $\deg \phi = \# \ker \phi$. In addition, the degree of a composition of separable isogenies is just the product of the degrees of each isogeny. Hence, a walk of length e_A in the l_A-isogeny graph corresponds to a kernel of size $l_A^{e_A}$, and this kernel is cyclic if and only if the walk does not backtrack. Notes that separability is preserved over composition, so the kernel uniquely defines the final isogeny.

Hence, Alice taking a random walk of length e_A is equivalent to her picking a random subgroup $\langle A \rangle \subset E[l_A^{e_A}]$; similarly Bob picks his random subgroup $\langle B \rangle \subset E[l_B^{e_B}]$. Then there is a well-defined subgroup $\langle A \rangle + \langle B \rangle = \langle A, B \rangle$ that defines the isogeny to $E/\langle A, B \rangle$. As $l_A \neq l_B$, the group $\langle A, B, \rangle$ is cyclic of order $l_A^{e_A} l_B^{e_B}$.

The following theorem allows us to control the group structure of the curve

Theorem 6.1 (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 [MOV93]. ■

We're only interested in the case when $q = p^2$ and $E(\mathbb{F}_q) \simeq (\mathbb{Z}/(p \pm 1)\mathbb{Z})^2$. We can choose our p so that $E(\mathbb{F}_q)$ contains two large subgroups $E(F_q)[l_A^{e_A}]$ and $E(\mathbb{F}_q)[l_B^{e_B}]$ of coprime order by letting $p = l_A^{e_A} l_B^{e_B} f \mp 1$, where f is a small cofactor. Hence, we can choose our generator A and B to be in $E(\mathbb{F}_q)$ and not some other field extension, and one generator is sufficient to represent an isogeny walk of length e_A. Now we have an elliptic curve with

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

Then $E(\mathbb{F}_{p^2})[l_A^{e_A}] \simeq (\mathbb{Z}/l_A^{e_A}\mathbb{Z})^2$, so we need two elements in $\mathbb{Z}/l_A^{e_A}\mathbb{Z}$ for a basis of $E(\mathbb{F}_{p^2})[l_A^{e_A}]$. For convenience, this is made public for each isogeny graphs

$$E[l_A^{e_A}] = \langle P_A, Q_A \rangle,$$

$$E[l_B^{e_B}] = \langle P_B, Q_B \rangle.$$

To start, Alice and Bob choose random secret subgroups

$$\langle A \rangle = \langle [m_A]P_A + [n_A]Q_A \rangle \subset E[l_A^{e_A}],$$

$$\langle B \rangle = \langle [m_B]P_B + [n_B]Q_B \rangle \subset E[l_B^{e_B}].$$

with respective orders $l_A^{e_A}$, $l_B^{e_B}$, and compute the secret isogenies

$$\alpha : E \to E/\langle A \rangle,$$

$$\beta : E \to E/\langle B \rangle.$$

They now publish $E_A = E/\langle A \rangle$ and $E_B = E/\langle B \rangle$. Note that for $\langle A \rangle$ to have order $l_A^{e_A}$, either m_A or n_A must be coprime to l_A.

In order for Alice to compute $E/\langle A, B \rangle$, she must compute the isogeny $\alpha' : E/\langle B \rangle \to E/\langle A, B \rangle$, which kernel is generated by $\beta(A)$. Then Bob can publish the values $\beta(P_A)$ and $\beta(Q_A)$ to help Alice. It is thought that this information does not give any advantage in computing $E/\langle A, B \rangle$. The shared secret is the j-invariant of $E/\langle A, B \rangle$, which Alice computes using $E_{AB} = E_B/\langle \beta(A) \rangle$ and Bob using $E_{BA} = E_A/\langle \alpha(B) \rangle$; the j-invariants of those curves are the same.

The following proposition summarizes the described procedure for supersingular key exchange.

Proposition 6.2 (Supersingular Isogeny Diffie-Hellman (SIDH))**.** *The following procedure on an isogeny graph allows for a Diffie-Hellman key exchange between parties A and B.*

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

We end this section with a discussion on the theoretical security of this protocol. To prevent either Alice or Bob's public keys from being weaker, we want to have $l_A^{e_A} \approx l_B^{e_B}$. Theorem 3.5 (2) tells us that the size of our graph is $O(p)$, then Alice's public key space is $O(\sqrt{p})$. This means that Alice and Bob's random walks are much shorter than the diameter of the graph. However, it is not known whether this poses a weakness in the security of the protocol.

A classical attack involves taking random walks of length $l_A^{e_A/2}$ from the start and end curves to obtain a meet-in-the-middle attack of runtime $O(\sqrt[4]{p})$. This is the solution to Problem 4.2 of finding smooth isogenies between two curves. According to [Feo17], the fastest known quantum-attack is $O(\sqrt[6]{p})$.

References

[CJS14] Andrew Childs, David Jao, and Vladimir Soukharev. Constructing elliptic curve isogenies in quantum subexponential time. *Journal of Mathematical Cryptology*, 8(1):1–29, Jan 2014.

[Feo17] Luca Feo. Mathematics of isogeny based cryptography. *École Mathématique Africaine*, 11 2017.

[JMV09] David Jao, Stephen D. Miller, and Ramarathnam Venkatesan. Expander graphs based on grh with an application to elliptic curve cryptography. *Journal of Number Theory*, 129(6):1491–1504, Jun 2009.

[MOV93] A. J. Menezes, T. Okamoto, and S. A. Vanstone. Reducing elliptic curve logarithms to logarithms in a finite field. *IEEE Transactions on Information Theory*, 39(5):1639–1646, Sep. 1993.

[Sil09] Joseph H Silverman. *The Arithmetic of Elliptic Curves*, volume 106. Springer-Verlag, 2 edition, 2009.

[Tao11] Terence Tao. 254b, notes 1: Basic theory of expander graphs, Dec 2011.

Email address: `Taekyukim02@gmail.com`