COMPUTATIONAL PROBLEMS IN SUPERSINGULAR ELLIPTIC CURVE ISOGENIES

STEVEN D. GALBRAITH AND FREDERIK VERCAUTEREN

Abstract. We give a brief survey of elliptic curve isogenies and the computational problems relevant for supersingular isogeny crypto. Supersingular isogeny cryptography is attracting attention due to the fact that there are no quantum attacks known against it that are significantly faster than classical attacks. However, the underlying computational problems have not been sufficiently studied by quantum algorithms researchers, especially since there are significant mathematical preliminaries needed to fully understand isogeny crypto. The main goal of the paper is to advertise various related computational problems, and to explain the relationships between them, in a way that is accessible to experts in quantum algorithms.

1. Introduction

An isogeny is a map $\phi : E_1 \to E_2$ where E_1 and E_2 are elliptic curves. Isogenies are maps both in the sense of geometry (mapping points from one curve to another) and algebra (they are group homomorphisms). One special case of an isogeny is the multiplication by n map $[n] : E \to E$ that is the central object of study in traditional elliptic curve cryptography. The elliptic curve discrete logarithm problem is to compute n when given two points $P, Q = [n]P$ on an elliptic curve E. One can view this problem as "determining" the isogeny $\phi : E \to E$ when given two points P and $Q = \phi(P)$. As is well known, Shor's algorithm is a polynomial-time algorithm to solve this problem on a quantum computer.

Isogeny cryptosystems were first proposed by Couveignes [10] and further developed in [32, 36] (these ones were based on "ordinary curves", for some details see later sections). The "supersingular curve" case was first developed in a hash function construction by Charles, Lauter and Goren [8]. Further cryptosystems in the supersingular case were proposed by Jao and de Feo [22] and developed in subsequent research [12, 23, 9, 20].

A subexponential-time quantum algorithm for the "ordinary curve" case was discovered by Childs, Jao and Soukharev [7]. As a result, the research focus has moved entirely to the supersingular case, where only exponential-time algorithms are known. The only quantum algorithm known for the supersingular case is due to Biasse, Jao and Sankar [3], and it requires exponential time and subexponential space.

More study of supersingular isogeny crypto by experts in quantum algorithms is essential. The first aim of the paper is to give a very gentle introduction to the main mathematical ideas behind isogeny crypto (see Sections 2 to 5). The main purpose of the paper is to explain a number of inter-related computational problems, and this is done in Section 6. Progress on quantum algorithms for any of these problems would be of major significance. Finally, Section 7 surveys the current state-of-the-art for classical and quantum algorithms for these problems.

2. Elliptic curves over finite fields

General references for this section are Washington [42], Silverman-Tate [34], Silverman [33], and Sutherland [37].

Let $\mathbb{F}_q$ be a finite field. In this paper $q = p^a$ will always be a power of a large prime p, so definitely $p > 3$. An elliptic curve E over $\mathbb{F}_q$ (in short Weierstrass form) is determined by two coefficients $A, B \in \mathbb{F}_q$ and is the set of points

$$E(\mathbb{F}_q) = \{(x, y) \in \mathbb{F}_q^2 : y^2 = x^3 + Ax + B\} \cup \{\mathbf{0}_E\}$$

where $\mathbf{0}_E$ is the point $(x : y : z) = (0 : 1 : 0)$ on the projective curve $y^2 z = x^3 + Axz^2 + Bz^3$. We will just write $\mathbf{0}$ when the curve E is clear. Sometimes we also consider all the points over the algebraic closure of the field $E(\overline{\mathbb{F}}_q)$.

The set of points on an elliptic curve is an abelian group under the "chord and tangent rule". The point $\mathbf{0}$ is the identity element of the group. For any point $P = (x_P, y_P) \in E(\mathbb{F}_q)$ we have $(x_P, -y_P) \in E(\mathbb{F}_q)$ and $P + (x_P, -y_P) = \mathbf{0}$, so this is the inverse of the point. For $n \in \mathbb{N}$ and $P \in E(\mathbb{F}_q)$ we define $[n]P$ to be $P + P + \cdots + P$ (n times). For example, $[2]P = P + P$.

There are "close to q" points on an elliptic curve over $\mathbb{F}_q$. Precisely, if we write $\#E(\mathbb{F}_q)$ for the number of points and set $t = q + 1 - \#E(\mathbb{F}_q)$ then $|t| \leq 2\sqrt{q}$. An elliptic curve over $\mathbb{F}_q$ where $q = p^a$ is called **supersingular** if $p \mid t$ and is called **ordinary** otherwise. It follows that E is supersingular if $\#E(\mathbb{F}_q) \equiv 1 \pmod{p}$, and in fact for supersingular curves one has $\#E(\mathbb{F}_{q^n}) \equiv 1 \pmod{p}$ for all $n \in \mathbb{N}$. This separation of elliptic curves into supersingular and ordinary may look arbitrary and unmotivated, but we will later see how different these two classes of curves are.

For $n \in \mathbb{N}$ define $E[n] = \{P \in E(\overline{\mathbb{F}}_q) : [n]P = \mathbf{0}\}$. If $p \nmid n$ then $\#E[n] = n^2$ and, group theoretically, $E[n]$ is a direct product of two cyclic groups of order n. If E is supersingular then $E[p] = \{\mathbf{0}\}$ while if E is ordinary then $\#E[p] = p$.

An **isomorphism** of elliptic curves $f : E \to E'$ over a field $\mathbb{F}_q$ is, informally, a function described by ratios of polynomials that maps points on E to points on E', satisfies $f(\mathbf{0}_E) = \mathbf{0}_{E'}$, and has an inverse that is also described by ratios of polynomials. It follows that an isomorphism is a bijection $E(\overline{\mathbb{F}}_q) \to E'(\overline{\mathbb{F}}_q)$. It can be shown that every isomorphism is of the form

$$f(x, y) = (u^2 x + r, u^3 y + s u^2 x + t)$$

where $u, r, s, t \in \overline{\mathbb{F}}_q$. Since isomorphisms are over $\overline{\mathbb{F}}_q$ they are not necessarily maps from $E(\mathbb{F}_q)$ to $E'(\mathbb{F}_q)$. If E is an elliptic curve over $\mathbb{F}_q$ with $\#E(\mathbb{F}_q) = q + 1 - t$ then there is an elliptic curve E' over $\mathbb{F}_q$, called the **quadratic twist** of E, such that $\#E'(\mathbb{F}_q) = q + 1 + t$ and E' is isomorphic to E (the isomorphism is however not defined over $\mathbb{F}_q$).

The **j-invariant** of an elliptic curve $E : y^2 = x^3 + Ax + B$ is

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

There is an isomorphism $f : E \to E'$ if and only if $j(E) = j(E')$.

Given $j \in \overline{\mathbb{F}}_q$ with $j \neq 0, 1728$, the elliptic curve

$$E : y^2 = x^3 + \frac{3j}{1728 - j} x + \frac{2j}{1728 - j}$$

has $j(E) = j$.

We end with some final remarks about supersingular elliptic curves. First, any supersingular elliptic curve E over $\overline{\mathbb{F}}_p$ is actually defined over $\mathbb{F}_{p^2}$ (meaning, it has j-invariant in $\mathbb{F}_{p^2}$). There are about $p/12$ isomorphism classes (j-invariants) of supersingular elliptic curves in total, and $O(\sqrt{p}\log(p))$ of them have j-invariants in $\mathbb{F}_p$. When $p > 3$ then all supersingular curves E over $\mathbb{F}_{p^2}$ have $\#E(\mathbb{F}_{p^2}) = (p+1)^2$ or $(p-1)^2$, and their group structure is $C_{(p+1)}^2$ (respectively, $C_{(p-1)}^2$) where C_n denotes a cyclic group of order n.

3. Isogenies

General references for this section are Chapter 12 of Washington [42], Chapters 9 and 25 of Galbraith [17] and De Feo [13].

Let E_1, E_2 be two elliptic curves over $\mathbb{F}_q$. An **isogeny**[1] is a morphism $\phi : E_1 \to E_2$ such that $\phi(\mathbf{0}_{E_1}) = \mathbf{0}_{E_2}$. One can show that isogenies are group homomorphisms, so they are "morphisms" both in the sense of algebraic geometry and group theory. Two elliptic curves are called **isogenous** if there is an isogeny between them.

The **degree** of an isogeny is essentially the degree of polynomials describing it (see Section 12.2 of Washington [42]). The degree of an isogeny is also, in general, the number of points in the kernel (an exception is inseparable isogenies such as the Frobenius map $\phi(x, y) = (x^p, y^p)$ on elliptic curves over $\mathbb{F}_p$).

A basic example of an isogeny is the **multiplication by n map** $[n]$ on an elliptic curve E for $n \in \mathbb{N}$, which we already defined by $[n]P = P + P + \cdots + P$ (n times). This maps $\mathbf{0}$ to itself, is a group homomorphism,

[1]The word "isogeny" means "same kind" and is also used in biology and medicine.

and is described by rational functions coming from the group law. The kernel is the set of points $\{P \in E(\overline{\mathbb{F}}_q) : [n]P = \mathbf{0}\}$ which has size n^2 in general.

Example 1. *Let $E : y^2 = x^3 + x$. Then the map $[2] : E \to E$ is given by the rational function*

$$[2](x, y) = \left(\frac{(x^2-1)^2}{4(x^3+x)}, \frac{y(x^6+5x^4-5x^2-1)}{8(x^3+x)^2} \right).$$

The kernel of $[2]$ is $\mathbf{0}$ together with the three points $(x_P, 0)$ such that $x_P^3 + x_P = 0$. In other words, the kernel is the set of four points of order dividing 2.

Example 2. *Let $A, B \in \mathbb{F}_q$ be such that $B \neq 0$ and $D = A^2 - 4B \neq 0$. Consider the elliptic curve $E : y^2 = x(x^2 + Ax + B)$ over $\mathbb{F}_q$. The point $(0, 0)$ has order 2. There is an elliptic curve E' and an isogeny $\phi : E \to E'$ such that $\ker(\phi) = \{\mathbf{0}_E, (0, 0)\}$. One can verify that*

$$\phi(x, y) = \left(\frac{y^2}{x^2}, \frac{y(B-x^2)}{x^2} \right) = \left(\frac{x^2+Ax+B}{x}, y\frac{B-x^2}{x^2} \right)$$

has the desired kernel, and the image curve is $E' : Y^2 = X(X^2 - 2AX + D)$.

The **dual isogeny** to $\phi : E \to E'$ is an isogeny $\hat{\phi} : E' \to E$ such that the composition $\hat{\phi} \circ \phi : E \to E$ is simply $[\deg(\phi)]$. The dual isogeny exists for every isogeny ϕ.

A major result (often called Tate's isogeny theorem since he generalised it to Abelian varieties) is that any two elliptic curves E_1 and E_2 over $\mathbb{F}_q$ are isogenous over $\mathbb{F}_q$ (the "over $\mathbb{F}_q$" means that the isogeny is given by rational functions of polynomials in $\mathbb{F}_q[x, y]$) if and only if $\#E_1(\mathbb{F}_q) = \#E_2(\mathbb{F}_q)$. One issue that frequently causes confusion to beginners is the fact that an isogeny has a kernel and yet the two curves have the same number of points. The following example will make this clear.

Example 3. *We consider the curve $E : y^2 = x(x^2 + x + 1)$ over $\mathbb{F}_7$, which is a special case of Example 2. One can check that $\#E(\mathbb{F}_7) = 8$. Indeed $E(\mathbb{F}_7) = \{\mathbf{0}, (0, 0), (2, 0), (3, 2), (3, 5), (4, 0), (5, 1), (5, 6)\}$ and the points $(0, 0), (2, 0), (4, 0)$ all have order 2 while the points $(3, \pm 2), (5, \pm 1)$ have order 4. The isogeny ϕ given in Example 2 maps to $E' : y^2 = x(x^2 - 2x - 3)$.*

One can check that $\phi(2, 0) = \phi(4, 0) = (0, 0)$. This gives the convenient fact that if one repeats the construction of Example 2 starting from E' then one computes an isogeny to the curve $E'' : y^2 = x(x^2 + 4Ax + 16B)$ which is isomorphic to E. The composition $E \to E' \to E''$ has kernel generated by $\{(0, 0), (2, 0)\}$ and so is $E[2]$, the group of points of order 2 on E. Hence this composition is just the multiplication by 2 map. This decomposition of the multiplication by 2 map into two isogenies of degree 2 is a tool used in the proof of the Mordell-Weil theorem (see Silverman-Tate [34] for details, or any other book on elliptic curves).

One can also check that $\phi(5, 1) = \phi(3, 5) = (2, 1)$ and $\phi(5, 6) = \phi(3, 2) = (2, 6)$. Hence $\phi(E(\mathbb{F}_7))$ is a cyclic group of order 4 inside the group $E'(\mathbb{F}_7)$ of order 8. This makes sense, since we have quotiented a group of order 8 by a subgroup of order 2.

What about the other 4 points in $E'(\mathbb{F}_7)$, such as $(3, 0)$? These are the image of points on E over an extension of $\mathbb{F}_7$. Consider the point $Q = (1, \alpha) \in E(\mathbb{F}_{7^2})$ where $\alpha \in \mathbb{F}_{7^2}$ satisfies $\alpha^2 = 3$. One can check that Q has order 4, $[2]Q = (0, 0)$, and $\phi(Q) = (3, 0)$. The other "missing points" in $E'(\mathbb{F}_7)$ are similarly the image of points on E over the extension field $\mathbb{F}_{7^2}$.

The next Theorem is extremely important and useful in the subject. Every isogeny $\phi : E \to E'$ has a kernel $G = \ker(\phi)$ that is a finite subgroup of $E(\overline{\mathbb{F}}_q)$. A natural question is to what extent ϕ is uniquely defined by its kernel and which finite subgroups of $E(\overline{\mathbb{F}}_q)$ arise as a kernel of an isogeny. The answer (ignoring inseparable isogenies) is that ϕ is uniquely defined up to composition with an isomorphism by its kernel, and that every finite subgroup G of $E(\overline{\mathbb{F}}_q)$ can be the kernel of an isogeny, but the isogeny is defined over $\mathbb{F}_q$ if and only if G is defined over $\mathbb{F}_q$. The definition of "G defined over $\mathbb{F}_q$" is: If $P \in G$ and $\sigma \in \mathrm{Gal}(\overline{\mathbb{F}}_q/\mathbb{F}_q)$ then $\sigma(P) \in G$. Note that this does not imply that $G \subseteq E(\mathbb{F}_q)$.

Theorem 1. *Let E be an elliptic curve defined over $\mathbb{F}_q$ and G a finite subgroup of $E(\overline{\mathbb{F}}_q)$ that is defined over $\mathbb{F}_q$. Then there is an elliptic curve E' defined over $\mathbb{F}_q$ and a separable isogeny $\phi : E \to E'$ defined over $\mathbb{F}_q$ of degree $\#G$ with $\ker(\phi) = G$. Furthermore, if $\psi : E \to E''$ is any other separable isogeny of degree $\#G$ with $\ker(\psi) = G$ then $j(E'') = j(E')$. Hence, the image curve E' is well-defined up to isomorphism and we sometimes denote it by E/G.*

There is an explicit algorithmic proof of Theorem 1 due to Vélu [40] (for details see Silverman [33] Proposition III.4.12, Galbraith [17] Section 25.1). The algorithmic proof of this Theorem gives an explicit formula for the equation of E' and for the isogeny ϕ (as rational functions). However, the complexity of the Vélu formulae is $O(n)$ field operations to compute an isogeny of degree n, which in certain applications would be considered as exponential complexity.

A key concept that makes isogeny crypto feasible is that isogenies factor into chains. Let E and E' be elliptic curves over $\mathbb{F}_q$ and let $\phi : E \to E'$ be a separable isogeny that is defined over $\mathbb{F}_q$. Then $\phi = \phi_1 \circ \cdots \circ \phi_k \circ [n]$ where $\phi_1, \ldots, \phi_k$ are isogenies of prime degree that are defined over $\mathbb{F}_q$ and $\deg(\phi) = n^2 \prod_{i=1}^{k} \deg(\phi_i)$. What this means in practice is that an isogeny of large degree can be constructed as a composition of isogenies of small prime degree. For example, one can form a sequence of t isogenies of degree 2, and the cost to compute the composition is proportional to t, rather than the cost $O(2^t)$ of computing the composition in a single step using the Vélu formulae.

For specific crypto applications there has been a lot of nice research to speed up the computation of chains of isogenies, but we do not discuss this in this paper. See for example De Feo, Jao and Plût [12] for a taste of this.

There is one further subtlety: The Vélu algorithm outputs a particular elliptic curve in the isomorphism class, and sometimes one needs to apply a suitable isomorphism to get to the desired curve. In the key exchange protocol (see Section 6), Alice and Bob use the Vélu algorithm and they are not expected to both generate exactly the same curve; that's why the protocol works with j-invariants.

4. Endomorphisms

The general reference for this section is Section III.9 of Silverman [33] and Sutherland [37].

The endomorphism ring of E is the set of isogenies from E to itself, together with the zero map $0 : E \to E$ given by $0(P) = \mathbf{0}$. In other words

$$\mathrm{End}(E) = \{\phi : E \to E\} \cup \{0\}.$$

This is a ring where addition of isogenies is defined using elliptic curve addition as $(\phi_1 + \phi_2)(P) = \phi_1(P) + \phi_2(P)$ and multiplication is composition. Note that $\mathbb{Z} \subset \mathrm{End}(E)$ from the map $n \mapsto [n]$. Also note that this map is injective: if $n \neq m$ then we never have $[n] = [m]$, because there is always some $P \in E(\overline{\mathbb{F}}_q)$ such that $[n]P \neq [m]P$, even if $[n]P = [m]P$ for all $P \in E(\mathbb{F}_q)$.

Hence, the ring $\mathrm{End}(E)$ is a $\mathbb{Z}$-module. A non-trivial theorem (related to the fact that $\#E[n] = n^2$) is that there are only three types of ring for $\mathrm{End}(E)$: namely $\mathbb{Z}$, an order in an imaginary quadratic field, a maximal order in a quaternion algebra. Further, the case $\mathrm{End}(E) = \mathbb{Z}$ does not occur for elliptic curves over finite fields. We give examples that illustrate what is going on.

Example 4. *Let E_1, E_2 over $\mathbb{F}_{13}$ be given by $E_1 : y^2 = x^3 + x$ and $E_2 : y^2 = x^3 + 7x + 5$ Then $\#E_1(\mathbb{F}_{13}) = \#E_2(\mathbb{F}_{13}) = 20$. This is the case of ordinary curves, since $20 \not\equiv 1 \pmod{13}$. The Frobenius map $\pi(x, y) = (x^{13}, y^{13})$ is an endomorphism on E and is known to satisfy the polynomial $T^2 + 6T + 13$, meaning that*

$$\pi(\pi(P)) + [6]\pi(P) + [13]P = \mathbf{0}$$

for all points $P \in E_1(\overline{\mathbb{F}}_{13})$ (and same for $E_2(\overline{\mathbb{F}}_{13})$). It follows that End$(E_1)$ and End(E_2) contain $\mathbb{Z}[\pi]$. Since π behaves like the complex number $-3 + 2i$ it follows that the ring $\mathbb{Z}[\pi]$ is isomorphic to $\mathbb{Z}[2i]$ where $i^2 = -1$ is the usual complex number. Hence $\mathbb{Z}[\pi]$ is a subring of $\mathbb{Q}(i)$ that contains 1. In other words, it is an order.

It can be shown that End$(E_2) \cong \mathbb{Z}[2i]$. However, End$(E_1)$ is larger. The endomorphism $\psi(x, y) = (-x, iy)$ satisfies $\psi^2(x, y) = (x, -y) = [-1](x, y)$ and so we write $\psi^2 = -1$ and identify ψ with the complex number i. It follows that $\pi = -3 + 2\psi$ (assuming an appropriate choice of sign is taken when i is defined) and so End$(E_1) \cong \mathbb{Z}[i]$.

The two rings $\mathbb{Z}[i]$ and $\mathbb{Z}[2i]$ are orders in the imaginary quadratic field $\mathbb{Q}(i)$.

The general result is that an **ordinary elliptic curve** over $\mathbb{F}_q$ with $q + 1 - t$ points has Frobenius endomorphism π that satisfies $\pi(\pi(P)) - [t]\pi(P) + [q]P = \mathbf{0}$ and has an endomorphism ring that is an order in $K = \mathbb{Q}(\sqrt{d})$ with $d = t^2 - 4q < 0$ and $\mathbb{Z}[\pi] \subseteq \mathrm{End}(E) \subseteq \mathcal{O}_K$. The conductor $c = [\mathcal{O}_K : \mathbb{Z}[\pi]]$ is the largest integer such that $d/c^2 \equiv 0, 1 \bmod 4$, so there is only a finite number of possibilities for $\mathrm{End}(E)$, namely, all

the rings $\mathcal{O} = \mathbb{Z} + c'\mathcal{O}_K$ with c' a divisor of c. In the above example, $d = 6^2 - 4 \cdot 13 = -16$, $K = \mathbb{Q}(i)$, $\mathcal{O}_K = \mathbb{Z}[i]$, the conductor $c = 2$ and the order $\mathbb{Z}[i] = \mathbb{Z} + \mathcal{O}_K$, and the order $\mathbb{Z}[2i] = \mathbb{Z} + 2\mathcal{O}_K$.

Example 5. *Let $p = 11$ and consider the curve $y^2 = x^3 + x$ over $\mathbb{F}_p$. We have $\#E(\mathbb{F}_{11}) = 12$ and so E is supersingular. As in the previous example, there are endomorphisms $\psi(x, y) = (-x, iy)$ such that $\psi^2 = [-1]$ and $\psi(x, y) = (x^p, y^p)$ such that $\pi^2 = [-p]$ (this latter statement is not obvious). However a difference this time is that $p \equiv 3 \pmod 4$ and so the field element i such that $i^2 = -1$ does not lie in $\mathbb{F}_p$ but in $\mathbb{F}_{p^2}$. Hence we have $\pi \circ \psi(x, y) = \pi(-x, iy) = (-x^p, i^p y^p) = (-x^p, -iy^p)$ whereas $\psi \circ \pi(x, y) = (-x^p, iy^p)$. It follows that $\pi\psi = -\psi\pi$ and so End(E) is a non-commutative ring. In fact End(E) is now an order in the quaternion algebra $\mathbb{Q}\langle i, j\rangle$ where $i^2 = -1$ and $j^2 = -p$.*

Indeed, when E is a supersingular curve then a theorem of Deuring is that $\mathrm{End}(E)$ is a maximal order in the quaternion algebra ramified at p and infinity. It is this difference in the endomorphism rings that makes supersingular curves so different from ordinary curves.

We do not have space to give all the details of orders in imaginary quadratic fields and quaternion algebras. But suffice to say that the ordinary case has strong connections with algebraic number theory via the theory of complex multiplication (see Cox [11] and Sutherland [35]). In particular, given an elliptic curve E over $\mathbb{F}_q$ with $\mathrm{End}(E) = \mathcal{O}_K$ (the maximal order) and $\mathfrak{a}$ an $\mathcal{O}_K$-ideal, we can define the $\mathfrak{a}$-torsion subgroup as the intersection of the kernels of all elements in $\mathfrak{a}$, i.e. $E[\mathfrak{a}] = \cap_{\alpha \in \mathfrak{a}} \ker(\alpha)$ and construct an isogeny $\phi_\mathfrak{a} : E \to E_\mathfrak{a} \simeq E/E[\mathfrak{a}]$ with kernel $E[\mathfrak{a}]$. The curve $E_\mathfrak{a}$ will have the same endomorphism ring $\mathcal{O}_K$ and when $\mathfrak{a}$ is principal, $E_\mathfrak{a}$ will be isomorphic to E. For $\mathfrak{a}$ and $\mathfrak{b}$ two $\mathcal{O}_K$ ideals, we have $\phi_{\mathfrak{a}\mathfrak{b}} = \phi_\mathfrak{a} \circ \phi_\mathfrak{b}$. We thus obtain an induced action of the class group $\mathrm{cl}(O_K)$ on the set of j-invariants of elliptic curves with endomorphism ring $\mathcal{O}_K$ given by

$$[\mathfrak{a}] \star j(E) = j(E_\mathfrak{a}) .$$

This construction immediately generalizes to the case where $\mathrm{End}(E) = \mathcal{O}$ is not the maximal order.

5. Modular Polynomials and Isogeny Graphs

General references for this section are Section 11 of Cox [11], Chapter 25 of Galbraith [17], Sutherland [35], Sutherland [37], and De Feo [13].

We have seen that if $G \subseteq E(\overline{\mathbb{F}}_q)$ is a group defined over $\mathbb{F}_q$ then there is an isogeny $\phi : E \to E' = E/G$ and that this isogeny can be computed using an algorithm due to Vélu. Hence the reader might assume that in order to compute isogenies it is necessary to compute kernel points. Surprisingly there is another tool for computing isogenies that does not explicitly deal with kernel subgroups or even points on elliptic curves at all.

Let ℓ be an integer with $\ell \geq 2$. The **modular polynomial** $\Phi_\ell(x, y) \in \mathbb{Z}[x, y]$ has the following remarkable property: A pair $j, j' \in \overline{\mathbb{F}}_q$ satisfies $\Phi_\ell(j, j') = 0$ if and only if there are elliptic curves E, E' over $\overline{\mathbb{F}}_q$ with $j(E) = j$ and $j(E') = j'$ and an isogeny $\phi : E \to E'$ of degree ℓ. It follows from the dual isogeny that $\Phi_\ell(y, x) = \Phi_\ell(x, y)$.

Note that modular polynomials have high degree and very large coefficients. When ℓ is prime then $\deg_x(\Phi_\ell(x, y)) = \ell + 1$ and indeed $\Phi_\ell(x, y) = x^{\ell+1} + y^{\ell+1} + x^\ell y^\ell +$ lower terms. It requires $O(\ell^3 \log(\ell))$ bits to represent Φ_ℓ.

Hence, given an elliptic curve E over $\mathbb{F}_q$, to find the j-invariants of the ℓ-isogenous curves one simply computes the univariate polynomial $\Phi_\ell(j(E), y) \in \mathbb{F}_q[y]$ and then computes its roots in $\overline{\mathbb{F}}_q$. An algorithm due to Elkies allows to compute the kernel of the corresponding isogeny (in exponential time in ℓ) when given E and the j-invariant j' of the isogenous curve E'.

For elliptic curves over $\mathbb{F}_q$ and ℓ a prime, the **ℓ-isogeny graph** (over $\mathbb{F}_q$) is a graph (V, G) (in the sense of graph theory) whose vertices V is the set of j-invariants of elliptic curves over $\mathbb{F}_q$, i.e. is simply given by $\mathbb{F}_q$, and whose edges G are the pairs $(j(E_1), j(E_2))$ of j-invariants of ℓ-isogeneous curves.[2] At first sight the graph appears to be a directed graph, but due to the dual isogeny one can essentially think of the graph as undirected (there are two special cases involving curves with j-invariants 0 and 1728 that we don't want to discuss). Note also that the graph can be a multi-graph (two distinct edges between the same two vertices).

[2]This is the entire isogeny graph of elliptic curves. Some references only define the isogeny graph of a curve E, which is the connected component of the entire graph containing $j(E)$.

For a set $S = \{\ell_1, \ldots, \ell_k\}$ of primes $\ell_i \geq 2$, the S-isogeny graph has edge set that is the union of the edge sets of all ℓ_i-isogeny graphs for $\ell_i \in S$, and the isogeny graph is the union of all ℓ-isogeny graphs for all primes ℓ.

The definition of supersingularity implies that an elliptic curve isogeneous to a supersingular curve is itself supersingular, so connected components in the isogeny graph are either ordinary or supersingular. The structure of both components is well known: an **ordinary component** in the ℓ-isogeny graph is a so-called **ℓ-volcano** which is a connected undirected graph whose vertices can be partitioned in levels $V_0, \ldots, V_d$. V_0 is called the surface and is typically a cycle (in general a regular graph of degree ≤ 2), each vertex in V_i for $i > 0$ has exactly one neighbour in level V_{i-1} (and all edges not on the surface arise in this manner), and all vertices have degree $\ell + 1$, except for the vertices in V_d that have degree 1. All vertices in the same level V_i correspond to elliptic curves with the same endomorphism ring $\mathcal{O}_i$, and the endomorphism ring on level i has index ℓ in the endomorphism ring on level $i - 1$, i.e. $[\mathcal{O}_{i-1} : \mathcal{O}_i] = \ell$. Therefore, for the ℓ-volcano to have more than one level, it is required that $\ell \mid c$ with c the conductor. In all other cases, the ℓ-volcano only consists of its surface. If we restrict the isogeny graph to the elliptic curves with maximal endomorphism ring $\mathrm{End}(E) = \mathcal{O}_K$, then the isogeny graph is a Cayley graph for the ideal class group. Since Cayley graphs of Abelian groups (with bounded vertex degree) are not families of expander graphs, it means that the shortest path between any two vertices might be quite long.

A **supersingular component** has a totally different structure: since every j-invariant of a supersingular curve lies in $\mathbb{F}_{p^2}$, it follows that $\Phi_\ell(j(E), Y)$ for E supersingular will have $l + 1$ roots in $\mathbb{F}_{p^2}$. If we consider the ℓ-isogeny graph over $\mathbb{F}_{p^2}$, the supersingular components will all be regular graphs of degree $\ell + 1$. In fact, one can show there is only one supersingular component and this component is an expander graph. This means it has "good mixing properties" and there is a "short" path between any two vertices in the graph. Indeed it is a Ramanujan graph, which means it has essentially optimal expansion properties. We refer to Chapter 41 of Voight [41] or Pizer [28, 29] for more details (though be warned that Pizer expresses his results without mentioning elliptic curves).

Example 6. *We have*

$$\Phi_2(x,y) = x^3+y^3-x^2y^2+1488(x^2y+xy^2)-162000(x^2+y^2)+40773375xy+8748000000(x+y)-157464000000000.$$

Let E be the elliptic curve $y^2 = x^3 + x + 5$ over $\mathbb{F}_{37}$ with $\#E(\mathbb{F}_{37}) = 38$ and $j(E) = 8$. We now construct the 2-isogeny graph of E over $\mathbb{F}_{37^2}$. First $\Phi_2(j(E), y) = (y - 8)(y - \alpha)(y - \beta)$ where $\alpha, \beta \in \mathbb{F}_{37^2}$ are roots of $y^2 + 31y + 31 = 0$. Now we can consider $\Phi_2(\alpha, y) = (y - 8)(y - \beta)^2$ and $\Phi_2(\beta, y) = (y - 8)(y - \alpha)^2$.

Hence the isogeny graph of E is the multi-graph with three vertices $\{8, \alpha, \beta\}$ and three undirected edges.

6. Computational problems and relationships

This is now the main part of the paper. We want to mention some computational problems that are relevant to isogeny crypto. A quantum algorithm for any one of these problems would have major impact on the attractiveness of supersingular isogenies

The first problem is the template for the whole subject.

Definition 1. ***The general isogeny problem**: Given $j, j' \in \mathbb{F}_q$ to find an isogeny $\phi : E \to E'$, if it exists, where $j(E) = j$ and $j(E') = j'$.*

A difficulty with this problem is that a solution ϕ may require significant space to describe (in general it would be exponential in the input size). Certain special cases that arise in applications include finding a path in an isogeny graph between two elliptic curves, and in certain contexts there is a compact (polynomial-sized) description of the path. We refer to Section 6.1 for some examples of such problems.

Note that the decisional question of whether an isogeny exists or not is solvable in polynomial time: an isogeny exists if and only if $\#E(\mathbb{F}_q) = \#E'(\mathbb{F}_q)$, and computing the number of points can be done in polynomial time. If one isogeny exists then there are an infinite number of isogenies $\phi : E \to E'$. So it does not make sense to ask for a specific isogeny, unless one asks for an isogeny of minimal degree (in which case the correctness of the solution is harder to verify since there is usually no efficient way to determine whether or not there is an isogeny of smaller degree between two curves).

A variant of this problem is when one is also told the degree of ϕ. This reduces the problem space from an infinite number of isogenies to a finite number (typically one, or zero if no such isogeny exists). In some

sense, this could make the problem harder. On the other hand, knowing the degree of the isogeny could potentially make the problem easier as it could reduce the search space. An example of this problem arises from the hash function of Charles, Lauter and Goren [8].

6.1. **Supersingular Isogeny Diffie-Hellman (SIDH).** For the Jao and De Feo system [22] (also see De Feo, Jao and Plût [12]) there is a very special set-up. First choose distinct small primes ℓ_1, ℓ_2 (typically $\ell_1 = 2$ and $\ell_2 = 3$) and choose $e_1, e_2 \in \mathbb{N}$ such that $\ell_1^{e_1} \approx \ell_2^{e_2} \approx 2^\lambda$ where λ is some security parameter (typically, $\lambda = 256$). Next choose a random small integer $f \in \mathbb{N}$ until $p = \ell_1^{e_1}\ell_2^{e_2} f \pm 1$ is prime. Choose E to be a supersingular elliptic curve over $\mathbb{F}_{p^2}$ (there is an efficient algorithm to do this due to Broker [6]) such that $E(\mathbb{F}_{p^2})$ has group structure a product of two cyclic groups of order $\ell_1^{e_1}\ell_2^{e_2} f$. Fix points $R_1, S_1 \in E[\ell_1^{e_1}]$ such that the group generated by R_1 and S_1 is the whole group $E[\ell_1^{e_1}]$. Similarly, choose R_2, S_2 such that $\langle R_2, S_2 \rangle = E[\ell_2^{e_2}]$. The **SIDH public key** is (E, R_1, S_1, R_2, S_2).

The SIDH key exchange protocol (an analogue of Diffie-Hellman) works as follows: Alice chooses a secret subgroup of $E[\ell_1^{e_1}]$ by choosing an integer $0 \leq a < \ell_1^{e_1}$ and setting $T_1 = R_1 + [a]S_1$. Alice computes an isogeny $\phi_A : E \to E_A$ with kernel generated by T_1 and publishes $(E_A, \phi_A(R_2), \phi_A(S_2))$. Similarly, Bob chooses $0 \leq b < \ell_2^{e_2}$, computes $\phi_B : E \to E_B$ with kernel generated by $T_2 = R_2 + [b]S_2$ and publishes $(E_B, \phi_B(R_1), \phi_B(S_1))$. To compute the shared key, Alice computes

$$T_1' = \phi_B(R_1) + [a]\phi_B(S_1) = \phi_B(R_1 + [a]S_1) = \phi_B(T_1)$$

and then computes an isogeny $\phi_A' : E_B \to E_{AB}$ with kernel generated by T_1'. The composition $\phi_A' \circ \phi_B : E \to E_{AB}$ has kernel $\langle T_1, T_2 \rangle$. Similarly, Bob computes an isogeny $\phi_B' : E_A \to E_{AB}'$ with kernel $\langle \phi_A(R_2) + [b]\phi_A(S_2) \rangle$. The actual elliptic curve equations E_{AB} and E_{AB}' computed by Alice and Bob are not likely to be the same, but the curves are isomorphic and so $j(E_{AB}) = j(E_{AB}')$. Hence, the shared key for Alice and Bob is $j(E_{AB})$.

The protocol can be nicely expressed in terms of quotients. We can think of E_A as E/G_A for some subgroup $G_A = \langle T_1 \rangle$ of $E[\ell_1^{e_1}]$, and of E_B as E/G_B. Then $E_{AB} = E/\langle G_A, G_B \rangle$, which explains why the two parties compute the same key. Note that the protocol cannot be described purely in terms of j-invariants: One can have situations where $E/G_A \cong E/G_A'$ and $E/G_B \cong E/G_B'$ but $E/\langle G_A, G_B \rangle \ncong E/\langle G_A', G_B' \rangle$.

For more discussion of the protocol and its security we refer to [22, 12, 9].

To break this key exchange protocol is to solve a more special problem than the general isogeny problem. In particular, there is a lot of auxiliary information.

Definition 2. ***SIDH isogeny problem****: Let (E, R_1, S_1, R_2, S_2) be a SIDH public key. Let E_A be such that there is an isogeny $\phi_A : E \to E_A$ of degree $\ell_1^{e_1}$. Let $R_2' = \phi_A(R_2), S_2' = \phi_A(S_2)$. The problem is: Given $(E, R_1, S_1, R_2, S_2, E_A, R_2', S_2')$ to determine an isogeny $\phi_A : E \to E_A$ of degree $\ell_1^{e_1}$ such that $R_2' = \phi_A(R_2)$ and $S_2' = \phi_A(S_2)$.*

Notes:

(1) This problem contains exponentially much auxiliary information: Let $0 \leq x, y < \ell_2^{e_2}$ and set $T = [x]R_2 + [y]S_2$. Then $\phi_A(T) = [x]R_2' + [y]S_2'$ can be computed. Hence an attacker can compute as many pairs $(T, \phi_A(T))$ on the graph of ϕ_A as they like. A natural approach is to compute ϕ_A by solving an interpolation problem. However the difficulty is that ϕ_A has degree $\ell_1^{e_1}$ and so is described by rational functions of exponential degree. The challenge is to solve this interpolation problem using the decomposed form of ϕ_A as a sequence of e_1 isogenies of degree ℓ_1.

(2) The scheme would be totally insecure if Alice also published $R_1' = \phi_A(R_1), S_1' = \phi_A(S_1)$. An attacker would simply compute $x, y \in \mathbb{Z}$ such that $(x, y) \notin (\ell_1\mathbb{Z})^2$ but $[x]R_1' + [y]S_1' = \mathbf{0}$ (identity element on elliptic curve). This is an easy discrete log problem to solve, since the point has smooth order $\ell_1^{e_1}$. Then $[x]R_1 + [y]S_1$ is in the kernel of ϕ_A and we have likely determined the kernel exactly and hence know ϕ_A. A framework for an attack based on this idea is developed in a very recent preprint of Petit [27].

6.2. **Decisional variants.** We now describe some interesting variants of these problems.

Definition 3. ***Decisional SIDH isogeny problem****: Let (E, R_1, S_1, R_2, S_2) be a SIDH public key. Let E_A be an elliptic curve and let $R_2', S_2' \in E_A[\ell_2^{e_2}]$. Let $0 < n \leq e_1$ The problem is: Given $(E, R_1, S_1, R_2, S_2, E_A, R_2', S_2', n)$*

to determine whether or not there exists an isogeny $\phi : E \to E_A$ of degree ℓ_1^n such that $R_2' = \phi(R_2)$ and $S_2' = \phi(S_2)$.

If this problem can be solved then there is an easy way to solve the SIDH isogeny problem:[3] Let $u \in \mathbb{Z}$ be such that $u\ell_1 \equiv 1 \pmod{\ell_2}$. Given the instance $(E, R_1, S_1, R_2, S_2, E_A, R_2', S_2')$ one chooses an ℓ_1-isogeny $\psi : E_A \to E'$ and calls the decisional algorithm on $(E, R_1, S_1, R_2, S_2, E', [u]\psi(R_2'), [u]\psi(S_2'), e_1 - 1)$. If the decisional oracle says "yes", then we have correctly determined the result of the first $e_1 - 1$ steps in the path from E to E_A. Iterating this process polynomially many times solves the isogeny problem.

A strategy that does not seem to work to solve this decisional problem is to use elliptic curve pairings. The Weil pairing satisfies the compatibility condition that if $\phi : E \to E'$ and $P, Q \in E[N]$ then

$$e_N(\phi(P), \phi(Q)) = e_N(P, Q)^{\deg(\phi)}$$

where the first pairing is computed on E' and the second on E (Proposition 8.2 of Silverman). However, taking $N = \ell_2^{e_2}$, it will always be the case that $e_N([u]\psi(R_2'), [u]\psi(S_2')) = e_N(R_2, S_2)^{\ell_1^{e_1}}$ even when the curve E' does not correspond to an intermediate curve along the path of the isogeny ψ.

6.3. **Computing the endomorphism ring.** Let E be a supersingular curve such that $\mathrm{End}(E)$ is known and let E' be an arbitrary isogenous curve. In general it is believed that the problem of computing $\mathrm{End}(E')$ and the problem of computing an isogeny $\phi : E \to E'$ are broadly equivalent (see Kohel [24], Kohel, Lauter, Petit and Tignol [25]). Note that this is not true in the ordinary case; one can usually determine $\mathrm{End}(E)$ much more easily than computing isogenties (see Kohel [24], Bisson-Sutherland [4]).

In the specific SIDH cryptosystem, the base curve E is often chosen to have a special form, in which case $\mathrm{End}(E)$ is usually known. To break the cryptosystem it suffices to compute $\mathrm{End}(E_A)$. Hence another problem worthy of consideration is to compute $\mathrm{End}(E')$ for an arbitrary elliptic curve E'.

There are several possible ways one might represent $\mathrm{End}(E')$. One method is by giving explicit isogenies $\phi : E' \to E'$ as rational functions. Since the degree is usually exponential, this is typically not a useful representation. Another way is as an abstract representation as a $\mathbb{Z}$-module in a quaternion algebra. In this setting, the representation as an explicit $\mathbb{Z}$-basis with respect to the basis $\{1, \mathbf{i}, \mathbf{j}, \mathbf{k}\}$ of the quaternion algebra can have polynomial size, so this is usually what we have in mind. A thorough discussion of these issues and proofs that the endomorphism ring has a polynomial-sized representation are given by Eisentraeger, Hallgren and Morrison [15].

Definition 4. *Given a supersingular elliptic curve $E/\mathbb{F}_{p^2}$, determine an abstract representation of the maximal order $\mathrm{End}(E)$ in the quaternion algebra ramified at p and ∞.*

If one has an abstract representation of $\mathrm{End}(E)$ and $\mathrm{End}(E')$ then one has a practical description of the entire infinite set of isogenies from E to E'. In this setting, it is shown in Section 4 of [19] that one can easily find the specific isogeny required for Definition 2 using lattice reduction; because that isogeny is of particularly small degree and so corresponds to a short vector in the lattice of all isogenies.

Kohel's algorithm to compute $\mathrm{End}(E)$ works by computing paths in the isogeny graph to find several distinct isogenies $\phi : E' \to E'$. Hence the basic sub-task in this area is to compute an "arbitrary" isogeny of a curve to itself. So we single-out this problem as being worthy of research.

Definition 5. *Given a supersingular elliptic curve $E/\mathbb{F}_{p^2}$, find an isogeny $\phi : E \to E$ that is not in $\mathbb{Z}[\pi]$.*

7. What is known about algorithms

In this section we use the asymptotic notation $\tilde{O}(f(n))$ which denotes $O(f(n)\log(f(n))^k)$ for some integer $k \geq 0$. We also use the subexponential function $L_N(a, c) = \exp(c\log(N)^a \log(\log(N))^{1-a})$ such that $L_N(0, c) = \log(N)^c$ is a polynomial function while $L_N(1, c) = N^c$ is an exponential function. The cases $0 < a < 1$ are super-polynomial but also sub-exponential.

[3] This approach has been independently discovered by Thormarker [39].

Ordinary curves. The first algorithm to solve the general isogeny problem for ordinary curves is due to Galbraith [16] and proceeds in two steps:

(1) Reduce the problem to the case of elliptic curves whose endomorphism ring is maximal. Given two ordinary curves E_1 and E_2 with $\#E_1(\mathbb{F}_q) = \#E_2(\mathbb{F}_q) = q+1-t$, an algorithm due to Kohel constructs a chain of isogenies from E_i to E_i' where $\mathrm{End}(E_i') = \mathcal{O}_K$ is the maximal order of $K = \mathbb{Q}(\sqrt{t^2-4q})$. The time and space complexity of this step are $\tilde{O}(c^3)$ and $\tilde{O}(c^2)$ with c the maximum conductor of $\mathrm{End}(E_i)$, i.e. $\max_i[\mathcal{O}_K : \mathrm{End}(E_i)]$. Since c can be as large as $q^{1/2}$ in the worst case, step 1 has expected running time $\tilde{O}(q^{3/2})$ and space $\tilde{O}(q)$.
(2) Construct an isogeny between E_1' and E_2'.

Galbraith solves step (2) by constructing isogeny trees using the following procedure: pick a random prime ℓ from a well-chosen set of primes and for each vertex j in the trees, construct all ℓ-isogeneous curves by computing the roots of $\Phi_\ell(j, Y)$ in $\mathbb{F}_q$. For each root r, add it to the tree (if not yet present), and add an edge between r and j labelled with ℓ. Repeat this procedure until an edge connects both trees, at which point one has found a path of isogenies connecting E_1' and E_2'. Each ℓ-isogeny in the path can be constructed using the methods of Elkies and Vélu. This is a time-memory tradeoff algorithm based on the bi-directional search algorithm of Pohl [30]. The time and space complexity of step 2 both are $\tilde{O}(q^{1/4})$. The algorithm cannot be easily distributed or parallelised. For smooth conductor or when the E_i have maximal endomorphism rings, step 1 becomes negligible and the overall running time and storage are $\tilde{O}(q^{1/4})$. The algorithm runs in polynomial time when the class number of the endomorphism ring is small. Note that (at least, when the conductor c is small) this algorithm outputs isogeny paths of minimal length (and so has polynomial-sized output).

An improvement to step 2 is given by Galbraith, Hess and Smart [18] where instead of isogeny trees, the authors use a random walk on the isogeny graph restricted to curves whose endomorphism ring is the maximal order $\mathcal{O}_K$. This allows the algorithm to be distributed and reduces the storage costs during the first stage of the algorithm. Recall that in this case we have an action of the class group $\mathrm{cl}(O_K)$ on the set of j-invariants given by $[\mathfrak{a}] \star j(E) = j(E_\mathfrak{a})$. Here $[\mathfrak{a}]$ denotes an ideal class. Each step of the random walk will update a pair $(j, [\mathfrak{a}])$ where j is a j-invariant and $[\mathfrak{a}]$ an element of $\mathrm{cl}(\mathcal{O}_K)$. The core of the random walk consists of a deterministic (but random looking) function $f : \mathbb{F}_q \to \mathrm{cl}(O_K)$ that maps a j-invariant to an element $[\mathfrak{a}] \in \mathrm{cl}(O_K)$. The function f is used to update the pair $(j_i, [\mathfrak{a}_i])$ by defining $j_{i+1} = [f(j_i)] \star j_i$ and $[\mathfrak{a}_{i+1}] = [\mathfrak{a}_i] \cdot [f(j_i)]$. The overall algorithm then proceeds in the following way: start a first random walk with initial value $(j_0^{(1)}, [\mathfrak{a}_0^{(1)}]) = (j(E_1'), [1])$ and execute $T = O(\sqrt{h_K})$ steps with $h_K = |\mathrm{cl}(\mathcal{O}_K)|$ resulting in the pair $(j_T^{(1)}, [\mathfrak{a}_T^{(1)}])$. Then start a second random walk at $(j_0^{(2)}, [\mathfrak{a}_0^{(2)}]) = (j(E_2'), [1])$ and walk until a collision occurs, i.e. until $j_T^{(1)} = j_S^{(2)}$ for some S. The expected number of steps S is also $O(\sqrt{h_K})$ and the space requirement is clearly polynomial. The isogeny connecting E_1' and E_2' can then be represented as the class group element $[\mathfrak{a}] = [\mathfrak{a}_T^{(1)}]/[\mathfrak{a}_S^{(2)}]$. To construct the actual isogeny, the authors use an index calculus type of algorithm to find a smooth representation of $[\mathfrak{a}]$, i.e. to express $[\mathfrak{a}] = [\prod_i \mathfrak{l}^{b_i}]$ for small prime ideals $\mathfrak{l}_i$. The time complexity of step 2 remains $\tilde{O}(q^{1/4})$ but for suitable parameters the space complexity can be subexponential or even polynomial in $\log(q)$. Work of Bisson and Sutherland [5] reduces the storage requirements to find a smooth representation of an ideal within time $\tilde{O}(q^{1/4})$. The isogenies output by the algorithm are no longer necessarily of minimal length.

Galbraith and Stolbunov [21] improved the complexity of the GHS algorithm by a constant factor by modifying the random walk function so that lower-degree isogenies are used more frequently.

7.0.1. *Subexponential-time methods.* Childs, Jao and Soukharev [7] describe an improved index calculus algorithm to find a relatively compact and smooth representation of an element $[\mathfrak{a}] \in \mathrm{cl}(\mathcal{O}_K)$ that runs in sub-exponential time $L_q(1/2, \sqrt{3}/2)$ (assuming the generalised Riemann hypothesis). This algorithm can be used to speed-up the last step of the GHS algorithm above, but it also allows to evaluate the class group action $[\mathfrak{a}] \star j(E)$ for any $[\mathfrak{a}] \in \mathrm{cl}(\mathcal{O}_K)$ in sub-exponential time. Childs, Jao and Soukharev [7] also describe a quantum algorithm for step 2 above by reducing it to the abelian hidden shift problem. This problem is defined as follows: let A be a finite abelian group, T a finite set and let $f_1, f_2 : A \to T$ be black-box functions. The functions f_1, f_2 are said to hide a shift $s \in A$ if f_1 is injective and $f_2(x) = f_1(xs)$ for all $x \in A$. The goal is then to recover s by evaluating the functions f_1 and f_2. Step 2 can be easily formulated as an abelian

hidden shift problem by defining the two functions $f_b([\mathfrak{a}]) = [\mathfrak{a}] \star j(E'_b)$ for $b = 1, 2$. Indeed, let $[\mathfrak{s}]$ be the ideal class such that $[\mathfrak{s}] \star j(E'_1) = j(E'_2)$, then clearly $f_2(x) = f_1(x[\mathfrak{s}])$ for all $x \in \mathrm{cl}(\mathcal{O}_K)$. The abelian hidden shift problem can be solved using Kuperberg's algorithm [26] in $L_{|A|}(1/2)$ time, space and number of queries to f_i. Since each query takes sub-exponential time itself, the overall time and space complexity to solve step 2 on a quantum computer is $L_q(1/2)$. Remark 4.7 of [7] emphasises that there are two reasons why the time complexity is subexponential: both Kuperberg's algorithm itself, and also the classical smoothness results for computing in class groups. The output is a path in the isogeny graph of subexponential length.

Childs et al. also state that a modified algorithm due to Regev [31] allows the space complexity to be made polynomial. But this claim is incorrect in the setting of the isogeny problem, as the computation of isogenies themselves is still subexponential.

Supersingular curves. Since the ℓ-isogeny graph for supersingular curves is connected for each ℓ, it suffices to consider one ℓ only, e.g. $\ell = 2$. The meet-in-the-middle approach by Galbraith [16] can also be applied to the supersingular graph over $\mathbb{F}_{p^2}$ by building isogeny trees from E_1 and E_2 (note that step 1 can be skipped). This method will find the shortest path from E_1 to E_2, but both the time and space complexity are $\tilde{O}(p^{1/2})$ since the size of the graph is $\sim p/12$. A random walk approach as in GHS [18] would result in the same time complexity, but also the same space complexity since there is no compact representation for the path traversed from the E_i.

Delfs and Galbraith [14] study the isogeny graph restricted to supersingular curves over $\mathbb{F}_p$, which has $\tilde{O}(p^{1/2})$ nodes. The endormorphism ring of such a curve over $\mathbb{F}_p$ is, just like the ordinary case, an order in the imaginary quadratic field $K = \mathbb{Q}(\sqrt{-p})$. The $\mathbb{F}_p$-isogeny graph consists of volcanoes with depth maximum 2, hence to construct an isogeny between two supersingular curves over $\mathbb{F}_p$, one can apply the same algorithms as in the ordinary case. The resulting algorithm runs in time $\tilde{O}(p^{1/4})$ and $\tilde{O}(1)$ space when using the low memory version. The general supersingular isogeny problem can then be solved by first constructing an isogeny from E_1, E_2 to curves E'_1, E'_2 over $\mathbb{F}_p$ using self-avoiding random walks (or a depth first search through all short paths) and then running the $\mathbb{F}_p$-algorithm. Since there are $O(p)$ isomorphism classes of supersingular elliptic curves over $\mathbb{F}_{p^2}$ of which only $\tilde{O}(p^{1/2})$ are defined over $\mathbb{F}_p$, and since the isogeny graph is an expander, the expected running time of this phase will be $\tilde{O}(p^{1/2})$. So unless the curves were already defined over $\mathbb{F}_p$, the time complexity remains $\tilde{O}(p^{1/2})$, but the space complexity is $\tilde{O}(1)$. Note that the resulting isogeny does not consist of a sequence of 2-isogenies, since more primes are needed for the $\mathbb{F}_p$-isogeny graph to be connected. (If $\mathrm{End}(E_1)$ is known and simple enough then one can transform this to an isogeny of order a power of two using the ideas in [25]).

Biasse, Jao and Sankar [3] adapt both stages of the Delfs-Galbraith algorithm to the quantum setting. Firstly, the algorithm of Childs, Jao and Soukharev [7] is used to construct an isogeny between two supersingular curves over $\mathbb{F}_p$, since this case is very similar to the ordinary case. The quantum complexity of this step is sub-exponential $L_p(1/2)$. Secondly, constructing an isogeny to a curve defined over $\mathbb{F}_p$ can be done in quantum complexity $\tilde{O}(p^{1/4})$ using Grover's algorithm: since the supersingular ℓ-isogeny graph is a Ramanujan graph, it suffices to search $\tilde{O}(p^{1/2})$ paths of length $O(\log(p))$ to find a path that passes through $\mathbb{F}_p$. The overall quantum complexity of this algorithm therefore is $\tilde{O}(p^{1/4})$.

The SIDH problem given in Definition 2 is more specific than computing an isogeny between two supersingular elliptic curves in that it specifies the exact degree $\ell_1^{e_1}$ of the isogeny and also the action on the ℓ_2-torsion. This results in a faster quantum algorithm. The isogeny is composed of e_1 degree ℓ_1 isogenies and given that $\ell^{e_1} \sim p^{1/2}$ is much smaller than the size of the isogeny graph, we expect to find precisely one isogeny path from E to E_A. This path can be found by constructing two isogeny trees, starting at E and E_A, consisting of all paths of length $e_1/2$. A curve that occurs as a leaf in both trees then immediately leads to the sought isogeny. Finding a common leaf of two trees can be viewed as an instance of the claw problem: given two functions $f : A \to C$ and $g : B \to C$, find a pair (a, b) such that $f(a) = g(a)$. On a classical computer this problem can be solved in time $(|A| + |B|)$ and $O(|A|)$ space by building a hash table for $f(a)$ for $a \in A$ and comparing with $g(b)$ for all $b \in B$. Tani [38] showed that on a quantum computer this problem can be solved in quantum complexity $O((|A| \cdot |B|)^{1/3})$, resulting in a $O(p^{1/6})$ attack (since $|A| = |B| = O(p^{1/4})$). We refer to Section 5.1 of [12] for details.

A natural question is why there is a subexponential quantum algorithm for the ordinary case, but only an exponential quantum algorithm for the supersingular case. The key difference seems to be the following:

In the ordinary case, the ideal class group acts on the isogeny graph (indeed, the isogeny graph is essentially a Cayley graph). However, in the supersingular case there is no "global" algebraic object that acts on the graph. Instead, if E is an elliptic curve then every isogeny $\phi : E \to E'$ corresponds to an ideal in the maximal order $\mathrm{End}(E)$ in the quaternion algebra, but isogenies from different elliptic curves correspond to "unrelated" isogenies in an "unrelated" maximal order (in the same quaternion algebra). We refer to [24, 25, 41] for more details of the ideal-theoretic interpretation.

Acknowledgement

We thank Sean Hallgren, Christophe Petit and Drew Sutherland for discussions and comments. The second author was partly supported by the European Commission through the ICT programme under contract H2020-ICT-2014-1 645622 PQCRYPTO.

References

[1] Reza Azarderakhsh, David Jao, Kassem Kalach, Brian Koziel, and Christopher Leonardi. Key compression for isogeny-based cryptosystems. In K. Emura, G. Hanaoka and R. Zhang (eds.), AsiaPKC '16, ACM (2016) 1–10.
[2] Daniel J. Bernstein, Johannes Buchmann and Erik Dahmen. Post-Quantum Cryptography. Springer, 2009.
[3] Jean-François Biasse, David Jao, and Anirudh Sankar. A quantum algorithm for computing isogenies between supersingular elliptic curves. In W. Meier and D. Mukhopadhyay (eds), INDOCRYPT 2014, Springer LNCS 8885 (2014) 428–442.
[4] Gaetan Bisson and Andrew V. Sutherland. Computing the endomorphism ring of an ordinary elliptic curve over a finite field. Journal of Number Theory, Volume 131, Issue 5 (2011) 815–831.
[5] Gaetan Bisson and Andrew V. Sutherland. A low-memory algorithm for finding short product representations in finite groups. Designs, Codes and Cryptography, Volume 63, Issue 1 (2012) 1–13.
[6] Reinier Broker. Constructing supersingular elliptic curves. Journal of Combinatorics and Number Theory 1, Volume 1, Issue 3 (2009) 269–273.
[7] Andrew M. Childs, David Jao, and Vladimir Soukharev. Constructing elliptic curve isogenies in quantum subexponential time. J. Mathematical Cryptology **8**, no. 1 (2014) 1–29.
[8] Denis Xavier Charles, Kristin E. Lauter, and Eyal Z. Goren. Cryptographic hash functions from expander graphs. J. Cryptology, 22(1) (2009) 93–113.
[9] Craig Costello, Patrick Longa, and Michael Naehrig. Efficient algorithms for supersingular isogeny Diffie-Hellman. In M. Robshaw and J. Katz (eds.), CRYPTO 2016, Springer LNCS 9814 (2016) 572–601.
[10] Jean-Marc Couveignes. Hard homogeneous spaces. Cryptology ePrint Archive, Report 2006/291 (2006)
[11] David A. Cox. Primes of the form x2 + ny2: Fermat, class field theory, and complex multiplication. Wiley, 1997.
[12] Luca De Feo, David Jao, and Jérôme Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. J. Mathematical Cryptology, 8(3) (2014) 209–247.
[13] Luca De Feo. Mathematics of Isogeny Based Cryptography. Notes from a summer school on Mathematics for Post-quantum cryptography. `http://defeo.lu/ema2017/poly.pdf`
[14] Christina Delfs and Steven D. Galbraith. Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$. Des. Codes Cryptography 78(2) (2016) 425–440.
[15] Kirsten Eisentraeger, Sean Hallgren and Travis Morrison. On the Hardness of Computing Endomorphism Rings of Supersingular Elliptic Curves. eprint 2017/986.
[16] Steven D. Galbraith. Constructing isogenies between elliptic curves over finite fields. LMS Journal of Computation and Mathematics, 2(1) (1999) 118–13.
[17] Steven D. Galbraith. Mathematics of Public Key Cryptography. Cambridge University Press, 2012.
[18] Steven D. Galbraith, Florian Hess and Nigel P. Smart. Extending the GHS Weil Descent Attack In L. R. Knudsen (ed.), EUROCRYPT 2002, Springer LNCS 2332 (2002) 29–44.
[19] Steven D. Galbraith, Christophe Petit, Barak Shani and Yan Bo Ti. On the security of supersingular isogeny cryptosystems. In J.-H. Cheon and T. Takagi (eds.) ASIACRYPT 2016, Springer LNCS 10031 (2016) 63–91.
[20] Steven D. Galbraith, Christophe Petit and Javier Silva. Signature Schemes Based On Supersingular Isogeny Problems. To appear in Asiacrypt 2017.
Available at eprint 2016/1154.
[21] Steven D. Galbraith and Anton Stolbunov. Improved algorithm for the isogeny problem for ordinary elliptic curves. Appl. Algebra Eng. Commun. Comput., 24(2) (2013) 107–131.
[22] David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In B.-Y. Yang (ed.), PQCrypto 2011, Springer LNCS 7071 (2011) 19–34.
[23] David Jao and Vladimir Soukharev. Isogeny-based quantum-resistant undeniable signatures. In M. Mosca (ed.), PQCrypto 2014, Springer LNCS 8772 (2014) 160–179.
[24] David Kohel, Endomorphism rings of elliptic curves over finite fields, PhD thesis, University of California, Berkeley, 1996.
[25] David Kohel, Kristin Lauter, Christophe Petit and Jean-Pierre Tignol. On the quaternion ℓ-isogeny path problem. LMS Journal of Computation and Mathematics, 17A (2014) 418–432.

[26] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. SIAM J. Comput., 35(1) (2005) 170–188.
[27] Christophe Petit. Faster Algorithms for Isogeny Problems using Torsion Point Images. To appear in Asiacrypt 2017. Available at ePrint 2017/571.
[28] Arnold K. Pizer. Ramanujan graphs and Hecke operators. Bulletin of the AMS, vol. 23, No. 1 (1990) 127–137.
[29] Arnold K. Pizer. Ramanujan graphs. In D. A. Buell and J. T. Teitelbaum (eds), Computational Perspectives on Number Theory, AMS Studies in Advanced Mathematics, vol. 7 (1998) 159–178.
[30] I. Pohl, Bi-directional and heuristic search in path problems, Technical Report 104, Stanford Linear Accelerator Center, Stanford, California (1969)
[31] Oded Regev. A subexponential time algorithm for the dihedral hidden subgroup problem with polynomial space. arXiv:quant-ph/0406151, 2004.
[32] Alexander Rostovtsev and Anton Stolbunov. Public-key cryptosystem based on isogenies. Cryptology ePrint Archive, Report 2006/145, 2006.
[33] Joseph H. Silverman. The arithmetic of elliptic curves. Springer GTM 106, Springer, 1986.
[34] Joseph H. Silverman and John T. Tate. Rational Points on Elliptic Curves. Springer Undergraduate Texts in Mathematics, 1992.
[35] Andrew Sutherland. Isogeny volcanoes. In E. W. Howe and K. Kedlaya (eds.), ANTS X, The Open Book Series, Mathematical Sciences Publishers, Berkeley, 1(1) (2013) 507–530.
[36] Anton Stolbunov. Constructing public-key cryptographic schemes based on class group action on a set of isogenous elliptic curves. Advances in Mathematics of Communications, **4**, no. 2 (2010) 215–235.
[37] Andrew Sutherland. Elliptic Curves. Lecture notes from a course (18.783) at MIT, 2017. `http://math.mit.edu/classes/18.783/2017/lectures`
[38] Seiichiro Tani. Claw finding algorithms using quantum walk. Theoretical Computer Science 410 (2009) 5285–5297.
[39] Erik Thormarker. Post-Quantum Cryptography: Supersingular Isogeny Diffie-Hellman Key Exchange. Thesis, Stockholm University, 2017.
[40] Jacques Vélu. Isogénies entre courbes elliptiques. Communications de l'Académie Royale des Sciences de Paris, 273 (1971) 238–241.
[41] John Voight. Quaternion algebras. 2017. `https://math.dartmouth.edu/~jvoight/quat-book.pdf`
[42] Lawrence C. Washington. Elliptic curves: Number theory and cryptography, 2nd ed., CRC Press, 2008.
[43] Sun Xi, Haibo Tian, and Yumin Wang. Toward quantum-resistant strong designated verifier signature from isogenies. In F. Xhafa, L. Barolli, F. Pop, X. Chen and V. Cristea (eds.), International Conference on Intelligent Networking and Collaborative Systems IEEE (2012) 292–296.
[44] Youngho Yoo, Reza Azarderakhsh, Amir Jalali, David Jao and Vladimir Soukharev. A Post-Quantum Digital Signature Scheme Based on Supersingular Isogenies. To appear in Financial Crypto 2017.

Mathematics Department, University of Auckland, NZ.
E-mail address: `s.galbraith@auckland.ac.nz`

ESAT/COSIC, Electrical Engineering, KU Leuven, BE.
E-mail address: `frederik.vercauteren@esat.kuleuven.be`