Chapter 25

Isogenies of Elliptic Curves

Isogenies are a fundamental object of study in the theory of elliptic curves. The definition and basic properties were given in Sections 9.6 and 9.7. In particular, they are group homomorphisms.

Isogenies are used in algorithms for point counting on elliptic curves and for computing class polynomials for the complex multiplication (CM) method. They have applications to cryptanalysis of elliptic curve cryptosystems. They also have constructive applications: prevention of certain side-channel attacks; computing distortion maps for pairing-based cryptography; designing cryptographic hash functions; relating the discrete logarithm problem on elliptic curves with the same number of points. We do not have space to discuss all these applications.

The purpose of this chapter is to present algorithms to compute isogenies from an elliptic curve. The most important result is Vélu's formulae, that compute an isogeny given an elliptic curve and a kernel subgroup G. We also sketch the various ways to find an isogeny given an elliptic curve and the j-invariant of an elliptic curve ℓ-isogenous to E. Once these algorithms are in place we briefly sketch Kohel's results, the isogeny graph, and some applications of isogenies. Due to lack of space we are unable to give proofs of most results.

Algorithms for computing isogenies on Jacobians of curves of genus 2 or more are much more complicated than in the elliptic case. Hence, we do not discuss them in this book.

25.1 Isogenies and Kernels

Let $E : y^2 + a_1xy + a_3 = x^3 + a_2x^2 + a_4x + a_6$ be an elliptic curve over $\Bbbk$. Recall from Section 9.6 that a non-zero isogeny $\phi : E \to \widetilde{E}$ over $\Bbbk$ of degree d is a morphism of degree d such that $\phi(\mathcal{O}_E) = \mathcal{O}_{\widetilde{E}}$. Such a map is automatically a group homomorphism and has kernel of size dividing d.

Theorem 9.7.5 states that a separable isogeny $\phi : E \to \widetilde{E}$ over $\Bbbk$ may be written in the form

$$\phi(x, y) = (\phi_1(x), cy\phi_1(x)' + \phi_3(x)), \tag{25.1}$$

where $\phi_1(x), \phi_3(x) \in \Bbbk(x)$, where $\phi_1(x)' = d\phi_1(x)/dx$ is the (formal) derivative of the rational function $\phi_1(x)$, where $c \in \overline{\Bbbk}^*$ is a non-zero constant, and where (writing $\widetilde{a}_i$ for the coefficients of $\widetilde{E}$)

$$2\phi_3(x) = -\widetilde{a}_1\phi_1(x) - \widetilde{a}_3 + (a_1x + a_3)\phi_1(x)'.$$

Lemma 9.6.13 showed that if $\phi_1(x) = a(x)/b(x)$ in equation (25.1) then the degree of ϕ is $\max\{\deg_x(a(x)), \deg_x(b(x))\}$. The kernel of an isogeny ϕ with $\phi_1(x) = a(x)/b(x)$ is $\{\mathcal{O}_E\} \cup \{P = (x_P, y_P) \in E(\overline{\Bbbk}) : b(x_P) = 0\}$. The kernel of a separable isogeny of degree d has d elements.

Let E be an elliptic curve over a field $\Bbbk$ and G a finite subgroup of $E(\overline{\Bbbk})$ that is defined over $\Bbbk$. Theorem 9.6.19 states that there is a unique elliptic curve $\widetilde{E}$ (up to isomorphism) and a separable isogeny $\phi : E \to \widetilde{E}$ over $\Bbbk$ such that $\ker(\phi) = G$. We sometimes write $\widetilde{E} = E/G$. Let ℓ be a prime such that $\gcd(\ell, \mathrm{char}(\Bbbk)) = 1$. Since $E[\ell]$ is isomorphic (as a group) to the product of two cyclic groups, there are $\ell + 1$ different subgroups of $E[\ell]$ of order ℓ. It follows that there are $\ell + 1$ isogenies of degree ℓ, not necessarily defined over $\Bbbk$, from E to other curves (some of these isogenies may map to the same image curve).

As implied by Theorem 9.6.18 and discussed in Exercise 9.6.20, an isogeny is essentially determined by its kernel. We say that two separable isogenies $\phi_1, \phi_2 : E \to \widetilde{E}$ are **equivalent isogenies** if $\ker(\phi_1) = \ker(\phi_2)$.

Exercise 25.1.1. Let $\phi : E \to \widetilde{E}$ be a separable isogeny. Show that if $\lambda \in \mathrm{Aut}(\widetilde{E})$ then $\lambda \circ \phi$ is equivalent to ϕ. Explain why $\phi \circ \lambda$ is not necessarily equivalent to ϕ for $\lambda \in \mathrm{Aut}(E)$.

Theorem 25.1.2 shows that isogenies can be written as "chains" of prime-degree isogenies. Hence, in practice one can restrict to studying isogenies of prime degree. This observation is of crucial importance in the algorithms.

Theorem 25.1.2. *Let E and $\widetilde{E}$ be elliptic curves over $\Bbbk$ and let $\phi : E \to \widetilde{E}$ be a separable isogeny that is defined over $\Bbbk$. 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 $\Bbbk$ and $\deg(\phi) = n^2 \prod_{i=1}^{k} \deg(\phi_i)$.*

Proof: Theorem 9.6.19 states that ϕ is essentially determined by its kernel subgroup G and that ϕ is defined over $\Bbbk$ if and only if G is. We will also repeatedly use Theorem 9.6.18, that states that an isogeny $\phi : E \to \widetilde{E}$ defined over $\Bbbk$ factors as $\phi = \phi_2 \circ \phi_1$ (where $\phi_1 : E \to E_1$ and $\phi_2 : E_1 \to \widetilde{E}$ are isogenies over $\Bbbk$) whenever $\ker(\phi)$ has a subgroup $G = \ker(\phi_1)$ defined over $\Bbbk$.

First, let n be the largest integer such that $E[n] \subseteq G = \ker(\phi)$ and note that $\phi = \phi' \circ [n]$ where $[n] : E \to E$ is the usual multiplication by n map. Set $i = 1$, define $E_0 = E$ and set $G = G/E[n]$. Now, let $\ell \mid \#G$ be a prime and let $P \in G$ have prime order ℓ. There is an isogeny $\phi_i : E_{i-1} \to E_i$ of degree ℓ with kernel $\langle P \rangle$. Let $\sigma \in \mathrm{Gal}(\overline{\Bbbk}/\Bbbk)$. Since $\sigma(P) \in G$ but $E[\ell] \not\subseteq G$ it follows that $\sigma(P) \in \langle P \rangle$ and so $\langle P \rangle$ is defined over $\Bbbk$. It follows that ϕ_i is defined over $\Bbbk$. Replace G by $\phi_i(G) \cong G/\langle P \rangle$ and repeat the argument. □

Exercise 25.1.3. How must the statement of Theorem 25.1.2 be modified if the requirement that ϕ be separable is removed?

Exercise 25.1.4. Let E be an ordinary elliptic curve. Let $\phi_1 : E \to E_1$ and $\phi_2 : E_1 \to E_2$ be non-zero separable isogenies over $\Bbbk$ of coprime degrees e and f respectively. Show that there is an elliptic curve $\widetilde{E}_1$ over $\Bbbk$, and a pair of non-zero separable isogenies $\psi_1 : E \to \widetilde{E}_1$ and $\psi_2 : \widetilde{E}_1 \to E_2$ of degrees f and e respectively, such that $\phi_2 \circ \phi_1 = \psi_2 \circ \psi_1$.

25.1.1 Vélu's Formulae

We now present explicit formulae, due to Vélu [617], for computing a separable isogeny from an elliptic curve E with given kernel G. These formulae work in any characteristic. As motivation for Vélu's formulae we now revisit Example 9.6.9.

Example 25.1.5. Let $E : y^2 = x^3 + x$ and consider the subgroup of order 2 generated by the point $(0, 0)$. From Example 9.2.4 we know that the translation by $(0, 0)$ map is given by

$$\tau_{(0,0)}(x, y) = \left(\frac{1}{x}, \frac{-y}{x^2} \right).$$

Hence, it follows that functions invariant under this translation map include

$$X = x + 1/x = (x^2 + 1)/x, \qquad Y = y - y/x^2 = y(x^2 - 1)/x^2.$$

One can compute that $X^3 = (x^6 + 3x^4 + 3x^2 + 1)/x^3$ and so

$$\begin{aligned} Y^2 &= y^2(x^2 - 1)^2/x^4 \\ &= (x^6 - x^4 - x^2 + 1)/x^3 \\ &= X^3 - 4X. \end{aligned}$$

It follows that the map

$$\phi(x, y) = \left(\frac{x^2 + 1}{x}, y\frac{x^2 - 1}{x^2} \right)$$

is an isogeny from E to $\widetilde{E} : Y^2 = X^3 - 4X$.

We remark that ϕ can also be written as

$$\phi(x, y) = \left(\frac{y^2}{x^2}, y\frac{x^2 - 1}{x^2} \right)$$

and can be written projectively as

$$\begin{aligned} \phi(x : y : z) &= (x(x^2 + z^2) : y(x^2 - z^2) : x^2 z) \\ &= (y(x^2 + z^2) : xy^2 - x^2 z - z^3 : xyz) \\ &= (y^2 z : y(x^2 - z^2) : x^2 z) \\ &= (xy^2 : y(y^2 - 2xz) : x^3). \end{aligned}$$

Theorem 25.1.6. *(Vélu) Let E be an elliptic curve over $\Bbbk$ defined by the polynomial*

$$F(x, y) = x^3 + a_2 x^2 + a_4 x + a_6 - \left(y^2 + a_1 xy + a_3 y \right) = 0.$$

Let G be a finite subgroup of $E(\overline{\Bbbk})$. Let G_2 be the set of points in $G - \{\mathcal{O}_E\}$ of order 2 and let G_1 be such that $\#G = 1 + \#G_2 + 2\#G_1$ and

$$G = \{\mathcal{O}_E\} \cup G_2 \cup G_1 \cup \{-Q : Q \in G_1\}.$$

Write

$$F_x = \frac{\partial F}{\partial x} = 3x^2 + 2a_2x + a_4 - a_1y \quad \textit{and} \quad F_y = \frac{\partial F}{\partial y} = -2y - a_1x - a_3.$$

For a point $Q = (x_Q, y_Q) \in G_1 \cup G_2$ define the quantities

$$u(Q) = (F_y(Q))^2 = (-2y_Q - a_1x_Q - a_3)^2$$

and

$$t(Q) = \begin{cases} F_x(Q) & \text{if } Q \in G_2 \\ 2F_x(Q) - a_1F_y(Q) & \text{if } Q \in G_1. \end{cases}$$

Note that if $Q \in G_2$ then $F_y(Q) = 0$ and so $u(Q) = 0$.

Define

$$t(G) = \sum_{Q \in G_1 \cup G_2} t(Q) \quad \textit{and} \quad w(G) = \sum_{Q \in G_1 \cup G_2} (u(Q) + x_Q t(Q))$$

and set

$$A_1 = a_1,\ A_2 = a_2,\ A_3 = a_3,\ A_4 = a_4 - 5t(G),\ A_6 = a_6 - (a_1^2 + 4a_2)t(G) - 7w(G).$$

Then the map $\phi : (x, y) \mapsto (X, Y)$ where

$$X = x + \sum_{Q \in G_1 \cup G_2} \frac{t(Q)}{x - x_Q} + \frac{u(Q)}{(x - x_Q)^2}$$

and

$$Y = y - \sum_{Q \in G_1 \cup G_2} u(Q)\frac{2y + a_1x + a_3}{(x - x_Q)^3} + t(Q)\frac{a_1(x - x_Q) + y - y_Q}{(x - x_Q)^2} + \frac{a_1u(Q) - F_x(Q)F_y(Q)}{(x - x_Q)^2}$$

is a separable isogeny from E to

$$\widetilde{E} : Y^2 + A_1XY + A_3Y = X^3 + A_2X^2 + A_4X + A_6$$

with kernel G. Further, ϕ satisfies

$$\phi^* \left(\frac{dX}{2Y + A_1X + A_3} \right) = \frac{dx}{2y + a_1x + a_3}.$$

Proof: (Sketch) The basic idea (as used in Example 25.1.5) is that the function

$$X(P) = \sum_{Q \in G} x(P + Q)$$

on E is invariant under G (in the sense that $X = X \circ \tau_Q$ for all $Q \in G$) and so can be considered as "defined on E/G". To simplify some calculations sketched below it turns out to be more convenient to subtract the constant $\sum_{Q \in G - \{\mathcal{O}_E\}} x(Q)$ from X. (Note that $x(Q) = x_Q$.) Let $t_\infty = -x/y$ be a uniformizer on E at $\mathcal{O}_E$ (one could also take $t_\infty = x/y$, but this makes the signs more messy). The function x can be written as $t_\infty^{-2} - a_1t_\infty^{-1} - a_2 - a_3t_\infty - (a_1a_3 + a_4)t_\infty^2 - \cdots$ (for more details about the expansions of x, y and ω_E in terms of power series see Section IV.1 of Silverman [564]). It follows that $X = t_\infty^{-2} - a_1t_\infty^{-1} - \cdots$ and so $v_{\mathcal{O}_E}(X) = -2$.

One can also show that $y = -t_\infty^{-3} - a_1 t_\infty^{-2} - a_2 t_\infty^{-1} - \cdots$. The function $Y(P) = \sum_{Q \in G} y(P+Q)$ is invariant under G and has $v_{\mathcal{O}_E}(Y) = -3$. One can therefore show (see Section 12.3 of Washington [626]) that the subfield $\Bbbk(X,Y)$ of $\Bbbk(x,y)$ is the function field of an elliptic curve $\widetilde{E}$ (Washington [626] does this in Lemma 12.17 using the Hurwitz genus formula). The map $\phi : (x,y) \mapsto (X,Y)$ is therefore an isogeny of elliptic curves. By considering the expansions in terms of t_∞ one can show that the equation for the image curve is $Y^2 + A_1XY + A_3Y = X^3 + A_2X^2 + A_4X + A_6$ where the coefficients A_i are as in the statement of the Theorem.

Now, let $\omega_E = dx/(2y + a_1x + a_3)$. One has $dx = (-2t_\infty^{-3} + a_1 t_\infty^{-2} + \cdots)dt_\infty$ and $2y + a_1x + a_3 = -2t_\infty^{-3} - a_1 t_\infty^{-2} + \cdots$ and so $\omega_E = (1 - a_1 t_\infty + \cdots)dt_\infty$. Similarly, $\phi^*(\omega_{\widetilde{E}}) = d(X \circ \phi)/(2Y \circ \phi + A_1 X \circ \phi + A_3) = d(t_\infty^{-2} + a_1 t_\infty^{-1} + \cdots)/(-2t_\infty^{-3} + \cdots) = (1 + \cdots)dt_\infty$. It follows that the isogeny is separable and that $\phi^*(\omega_{\widetilde{E}}) = f\omega_E$ for some function f. Further, $\mathrm{div}(\omega_E) = 0$ and $\mathrm{div}(\phi^*(\omega_{\widetilde{E}})) = \phi^*(\mathrm{div}(\omega_{\widetilde{E}})) = 0$ (by Lemma 8.5.36, since ϕ is unramified[1]) and so $\mathrm{div}(f) = 0$. It follows that f is a constant, and the power series expansions in terms of t_∞ imply that $f = 1$ as required.

Write the isogeny as $\phi(x,y) = (\phi_1(x), y\phi_2(x) + \phi_3(x))$. By Theorem 9.7.5 the isogeny is determined by $\phi_1(x)$ (for the case $\mathrm{char}(\Bbbk) = 2$ see Exercise 9.7.6). Essentially, one only has to prove Vélu's formula for $\phi_1(x)$; we do this now. First, change the definition of X to

$$X(P) = x_P + \sum_{Q \in G - \{\mathcal{O}_E\}} (x_{P+Q} - x_Q)$$

where P is a "generic point" (i.e., $P = (x_P, y_P)$ where x_P and y_P are variables) on the elliptic curve and $Q \in G - \{\mathcal{O}_E\}$. Let $F(x,y)$ be as in the statement of the Theorem and let $y = l(x)$ be the equation of the line through P and Q (so that $l(x) = \lambda(x - x_Q) + y_Q$ where $\lambda = (y_P - y_Q)/(x_P - x_Q)$). Define

$$F_1(x) = F(x, l(x)) = (x - x_Q)(x - x_P)(x - x_{P+Q}).$$

Further,

$$\frac{\partial F_1}{\partial x}(Q) = (x_Q - x_P)(x_Q - x_{P+Q})$$

and

$$\frac{\partial F_1}{\partial x} = \frac{\partial F}{\partial x} + \frac{\partial F}{\partial y}\frac{\partial l}{\partial x} = F_x + F_y \cdot \lambda.$$

Hence, $x_{P+Q} - x_Q = F_x(Q)/(x_P - x_Q) + (y_P - y_Q)F_y(Q)/(x_P - x_Q)^2$. One now considers two cases: When $[2]Q = \mathcal{O}_E$ then $F_y(Q) = 0$. When $[2]Q \neq \mathcal{O}_E$ then it is convenient to consider

$$x_{P+Q} - x_Q + x_{P-Q} - x_{-Q}.$$

Now, $x_{-Q} = x_Q$, $y_{-Q} = y_Q + F_y(Q)$, $F_x(-Q) = F_x(Q) - a_1F_y(Q)$ and $F_y(-Q) = -F_y(Q)$. The formula for $\phi_1(x)$ follows.

Now we sketch how to obtain the formula for the Y-coordinate of the isogeny in the case $\mathrm{char}(\Bbbk) \neq 2$. Note that $\phi_1(x) = x + \sum_Q [t(Q)/(x - x_Q) + u(Q)/(x - x_Q)^2]$ and so $\phi_1(x)' = 1 - \sum_Q [t(Q)/(x - x_Q)^2 + 2u_Q/(x - x_Q)^3]$. Using $\phi_3(x) = (-A_1\phi_1(x) - A_3 +$

[1] This was already discussed in Section 9.6. One can directly see that separable isogenies are unramified since if $\phi(P_1) = P_2$ then the set of pre-images under ϕ of P_2 is $\{P_1 + Q : Q \in \ker(\phi)\}$.

$(a_1x + a_3)\phi_1(x)')/2$ one computes

$$\begin{aligned} y\phi_2(x) + \phi_3(x) &= y\phi_1(x)' + \phi_3(x) \\ &= y\left(1 - \sum_Q t(Q)/(x-x_Q)^2 + 2u(Q)/(x-x_Q)^3\right) \\ &\quad -(a_1x+a_3)/2 - a_1\sum_Q [t(Q)/(x-x_Q)^2 + 2u(Q)/(x-x_Q)^3] \\ &\quad +(a_1x+a_3)/2 + (a_1x+a_3)\sum_Q [-t(Q)/(x-x_Q)^2 - 2u(Q)/(x-x_Q)^3] \\ &= y - \sum_Q \left[t(Q)\frac{y + a_1(x-x_Q) - y_Q}{(x-x_Q)^2} + u(Q)\frac{2y + a_1x + a_3}{(x-x_Q)^3} \right. \\ &\quad \left. + \frac{t(Q)((a_1x_Q + a_3)/2 + y_Q) + a_1u(Q)/2}{(x-x_Q)^2} \right]. \end{aligned}$$

It suffices to show that the numerator of the final term in the sum is equal to $a_1u(Q) - F_x(Q)F_y(Q)$. However, this follows easily by noting that $(a_1x_Q + a_3)/2 + y_Q = -F_y(Q)/2$, $u(Q) = F_y(Q)^2$ and using the facts that $F_y(Q) = 0$ when $[2]Q = \mathcal{O}_E$ and $t(Q) = 2F_x(Q) - a_1F_y(Q)$ otherwise. □

Corollary 25.1.7. *Let E be an elliptic curve defined over $\Bbbk$ and G a finite subgroup of $E(\overline{\Bbbk})$ that is defined over $\Bbbk$. Then there is an elliptic curve $\widetilde{E} = E/G$ defined over $\Bbbk$ and an isogeny $\phi : E \to \widetilde{E}$ defined over $\Bbbk$ with $\ker(\phi) = G$.*

Proof: It suffices to show that the values $t(G)$, $w(G)$ and the rational functions X and Y in Theorem 25.1.6 are fixed by any $\sigma \in \mathrm{Gal}(\overline{\Bbbk}/\Bbbk)$. □

Corollary 25.1.8. *Let $\phi : E \to \widetilde{E}$ be a separable isogeny of odd degree ℓ between elliptic curves over $\Bbbk$. Write $\phi(x,y) = (\phi_1(x), \phi_2(x,y))$, where $\phi_1(x)$ and $\phi_2(x,y)$ are rational functions. Then $\phi_1(x,y) = u(x)/v(x)^2$, where $\deg(u(x)) = \ell$ and $\deg(v(x)) = (\ell-1)/2$. Also, $\phi_2(x,y) = (yw_1(x) + w_2(x))/v(x)^3$, where $\deg(w_1(x)) \leq 3(\ell-1)/2$ and $\deg(w_2(x)) \leq (3\ell-1)/2$.*

Exercise 25.1.9. Prove Corollary 25.1.8.

Definition 25.1.10. An isogeny $\phi : E \to \widetilde{E}$ is **normalised** if $\phi^*(\omega_{\widetilde{E}}) = \omega_E$.

Vélu's formulae give a normalised isogeny. Note that normalised isogenies are incompatible with Theorem 9.7.2 (which, for example, implies $[m]^*(\omega_E) = m\omega_E$). For this reason, in many situations one needs to take an isomorphism from $\widetilde{E}$ to obtain the desired isogeny. Example 25.1.12 shows how this works.

Exercise 25.1.11. Let $\phi : E \to \widetilde{E}$ be an isogeny given by rational functions as in equation (25.1). Show that ϕ is normalised if and only if $c = 1$.

Example 25.1.12. Let $E : y^2 + xy + 3y = x^3 + 2x^2 + 4x + 2$ over $\mathbb{F}_{311}$. Then

$$E[2] = \{\mathcal{O}_E, (-1,-1), (115, 252), (117, 251)\} \subset E(\mathbb{F}_{311}).$$

Let $G = E[2]$. Applying the Vélu formulae one computes $t(G) = 8$, $w(G) = 306$, $A_1 = 1, A_2 = 2, A_3 = 3, A_4 = 275$ and $A_6 = 276$. One can check that E and

$$\widetilde{E} : Y^2 + XY + 3Y = X^3 + 2X^2 + 275X + 276$$

have the same j-invariant, but they are clearly not the same Weierstrass equation. Hence, the Vélu isogeny with kernel $E[2]$ is not the isogeny $[2] : E \to E$.

To recover the map [2] one needs to find a suitable isomorphism from $\widetilde{E}$ to E. The isomorphism will have the form $(X, Y) \mapsto (u^2X + r, u^3Y + su^2X + t)$ where we must have $u = 1/2$ to have the correct normalisation for the action of the isogeny on the invariant differential (see Exercise 25.1.13). One can verify that taking $r = 291$, $s = 233$ and $t = 67$ gives the required isomorphism from $\widetilde{E}$ to E and that the composition of the Vélu isogeny and this isomorphism is the map [2].

Exercise 25.1.13. Show that if $\phi : (x, y) \mapsto (u^2x + r, u^3y + su^2x + t)$ is an isomorphism from E to $\widetilde{E}$ then $\phi^*(\omega_{\widetilde{E}}) = \frac{1}{u}\omega_E$.

Exercise 25.1.14. Determine the complexity of constructing and computing the Vélu isogeny. More precisely, show that if $d = \#G$ and $G \subset E(\mathbb{F}_{q^n})$ then $O(dM(n, q))$ bit operations are sufficient, where $M(n, q) = M(n \log(nq))$ is the number of bit operations to multiply two degree n polynomials over $\mathbb{F}_q$.

Further, show that if d is an odd prime then $n \leq d - 1$ and so the complexity can be written as $O(d^{2+\epsilon} \log(q)^{1+\epsilon})$ bit operations.

Example 25.1.15. Consider $E : y^2 = x^3 + 2x$ over $\mathbb{F}_{37}$, with $j = 26 \equiv 1728 \pmod{37}$. We have $\#E(\mathbb{F}_{37}) = 2 \cdot 5^2$ so there is a unique point $(0, 0)$ of order 2 over $\mathbb{F}_{37}$ giving a 2-isogeny from E. Using Vélu's formulae one determines that the image of this isogeny is $E_1 : y^2 = x^3 + 29x$, which also has j-invariant 26 and is isomorphic to E over $\mathbb{F}_{37}$.

Now consider the other points of order 2 on E. Let $\alpha \in \mathbb{F}_{37^2}$ satisfy $\alpha^2 = -2$. The isogeny ϕ_2 with kernel $\{\mathcal{O}_E, (\alpha, 0)\}$ maps to $E_2 : y^2 = x^3 + 28\alpha x$, while the isogeny ϕ_3 with kernel $\{\mathcal{O}_E, (-\alpha, 0)\}$ maps to $E_3 : y^2 = x^3 - 28\alpha x$. Note that there is an isomorphism $\psi : E_2 \to E_3$ over $\mathbb{F}_{37^2}$. We have $\hat{\phi_2} \circ \phi_2 = \hat{\phi_3} \circ \phi_3 = [2]$ on E. One can also consider $\hat{\phi_3} \circ \psi \circ \phi_2$ on E, which must be an element of $\mathrm{End}(E) = \mathbb{Z}[i]$ of degree 4. One can show that it is $i[2]$ where $i(x, y) = (-x, 31y)$. Hence, [2] and $\hat{\phi_3} \circ \psi \circ \phi_2$ are equivalent isogenies.

Kohel [350] and Dewaghe [171] independently gave formulae for the Vélu isogeny in terms of the coefficients of the polynomial defining the kernel, rather than in terms of the points in the kernel. We give these formulae in Lemma 25.1.16 for the case where G has odd order (they are also given in Section 2.4 of [350]). Since a $\Bbbk$-rational subgroup of an elliptic curve can have points defined over an extension of $\Bbbk$, working with the coefficients of the polynomial can be more efficient than working with the points in G.

Lemma 25.1.16. *Let $E : y^2 + (a_1x + a_3)y = x^3 + a_2x^2 + a_4x + a_6$ be an elliptic curve over $\Bbbk$. Let G be a cyclic subgroup of $E(\overline{\Bbbk})$ of odd order $2d + 1$. Let $G_1 \subseteq G$ be such that $\#G_1 = d$ and $G = \{\mathcal{O}_E\} \cup G_1 \cup \{-Q : Q \in G_1\}$. Define*

$$\psi(x) = \prod_{Q \in G_1} (x - x_Q) = x^d - s_1x^{d-1} + s_2x^{d-2} + \cdots + (-1)^d s_d \tag{25.2}$$

where the s_i are the i-th symmetric polynomials in the roots of $\psi(x)$ (equivalently, in the x-coordinates of elements of G_1). Define $b_2 = a_1^2 + 4a_2, b_4 = 2a_4 + a_1a_3$ and $b_6 = a_3^2 + 4a_6$. Then there is an isogeny $\phi : E \to \widetilde{E}$, with $\ker(\phi) = G$, of the form $\phi(x, y) = (A(x)/\psi(x)^2, B(x, y)/\psi(x)^3)$ where $A(x)$ and $B(x, y)$ are polynomials. Indeed,

$$\frac{A(x)}{\psi(x)^2} = (2d+1)x - 2s_1 - (4x^3 + b_2x^2 + 2b_4x + b_6)(\psi(x)'/\psi(x))' - (6x^2 + b_2x + b_4)(\psi(x)'/\psi(x)).$$

The proof of Lemma 25.1.16 is given as a sequence of exercises.

Exercise 25.1.17. Let the notation be as in Lemma 25.1.16. Let $F_x(Q)$, $F_y(Q)$, $t(Q)$ and $u(Q)$ be as in Theorem 25.1.6. Show that

$$t(Q) = 6x_Q^2 + b_2 x_Q + b_4 \quad \text{and} \quad u(Q) = 4x_Q^3 + b_2 x_Q^2 + 2b_4 x_Q + b_6.$$

Exercise 25.1.18. Let the notation be as in Lemma 25.1.16. Let $F_x(Q)$, $F_y(Q)$, $t(Q)$ and $u(Q)$ be as in Theorem 25.1.6. Show that

$$\begin{aligned}
\frac{x_Q}{x - x_Q} &= \frac{x}{x - x_Q} - 1 , \\
\frac{x_Q}{(x - x_Q)^2} &= \frac{x}{(x - x_Q)^2} - \frac{1}{(x - x_Q)} , \\
\frac{x_Q^2}{(x - x_Q)} &= \frac{x^2}{x - x_Q} - x - x_Q , \\
\frac{x_Q^2}{(x - x_Q)^2} &= \frac{x^2}{(x - x_Q)^2} - \frac{2x}{(x - x_Q)} + 1 , \\
\frac{x_Q^3}{(x - x_Q)^2} &= \frac{x^3}{(x - x_Q)^2} - \frac{3x^2}{(x - x_Q)} + 2x + x_Q.
\end{aligned}$$

Exercise 25.1.19. Let the notation be as in Lemma 25.1.16. For $1 \leq i \leq 3$ define

$$S_i = \sum_{Q \in G_1} \frac{1}{(x - x_Q)^i}.$$

Show that $S_1 = \psi(x)'/\psi(x)$ and that $S_2 = -(\psi'(x)/\psi(x))' = ((\psi(x)')^2 - \psi(x)\psi(x)'')/\psi(x)^2$.

Exercise 25.1.20. Complete the proof of Lemma 25.1.16.

Exercise 25.1.21. Determine the complexity of using Lemma 25.1.16 to compute isogenies over finite fields. More precisely, show that if $G \subseteq E(\mathbb{F}_{q^n})$ is defined over $\mathbb{F}_q$ and $d = \#G$ then one can compute $\psi(x)$ in $O(d^2)$ operations in $\mathbb{F}_{q^n}$. Once $\psi(x) \in \mathbb{F}_q[x]$ is computed show that one can compute the polynomials $A(x)$ and $B(x, y)$ for the isogeny in $O(d)$ operations in $\mathbb{F}_q$.

25.2 Isogenies from j-invariants

Vélu's formulae require that one knows the kernel of the desired isogeny. But in some applications one wants to take a $\Bbbk$-rational isogeny of a given degree d (assuming such an isogeny exists) from E to another curve $\widetilde{E}$ (where $\widetilde{E}$ may or may not be known), and one does not know a specific kernel. By Theorem 25.1.2 one can restrict to the case when $d = \ell$ is prime. We usually assume that ℓ is odd, since the case $\ell = 2$ is handled by points of order 2 and Vélu's formulae.

One solution is to choose a random point $P \in E[\ell]$ that generates a $\Bbbk$-rational subgroup of order ℓ. To find such a point, compute the ℓ-division polynomial (which has degree $(\ell^2 - 1)/2$ when ℓ is odd) and find irreducible factors of it in $\Bbbk[x]$ of degree up to $(\ell - 1)/2$. Roots of such factors are points of order ℓ, and one can determine whether or not they generate a $\Bbbk$-rational subgroup by computing all points in the subgroup. Roots of factors of degree $d > (\ell - 1)/2$ cannot give rise to $\Bbbk$-rational subgroups of order ℓ. This approach is expensive when ℓ is large for a number of reasons. For a start, finding roots of degree at most $(\ell - 1)/2$ of a polynomial of degree $(\ell^2 - 1)/2$ in $\mathbb{F}_q[x]$ takes $\Omega(\ell^3 \log(\ell) \log(q))$ bit operations.

A more elegant approach is to use the ℓ-th modular polynomial. It is beyond the scope of this book to present the theory of modular functions and modular curves (some basic references are Sections 5.2 and 5.3 of Lang [366] and Section 11.C of Cox [157]). The fundamental fact is that there is a symmetric polynomial, called the **modular polynomial**[2] $\Phi_\ell(x,y) \in \mathbb{Z}[x,y]$ such that if E is an elliptic curve over a field $\Bbbk$ and $\widetilde{E}$ is an elliptic curve over $\Bbbk$ then there is a separable isogeny of degree ℓ (where $\gcd(\ell, \mathrm{char}(\Bbbk)) = 1$) with cyclic kernel from E to $\widetilde{E}$ if and only if $\Phi_\ell(j(E), j(\widetilde{E})) = 0$ (see Theorem 5, Section 5.3 of Lang [366]). The modular polynomial $\Phi_\ell(x,y)$ is a singular model for the modular curve $X_0(\ell)$ over $\mathbb{Q}$. This modular curve is a moduli space in the sense that a (non-cusp) point of $X_0(\ell)(\Bbbk)$ corresponds to a pair (E,G) where E is an elliptic curve over $\Bbbk$ and where G is a cyclic subgroup of E, defined over $\Bbbk$, of order ℓ. Note that it is possible to have an elliptic curve E together with two distinct cyclic subgroups G_1 and G_2 of order ℓ such that the image curves E/G_1 and E/G_2 are isomorphic; in this case (E,G_1) and (E,G_2) are distinct points of $X_0(\ell)$ but correspond to a repeated root of $\Phi_\ell(j(E),y)$ (it follows from the symmetry of $\Phi_\ell(x,y)$ that this is a singular point on the model). In other words, a repeated root of $\Phi_\ell(j(E),y)$ corresponds to non-equivalent ℓ-isogenies from E to some elliptic curve $\widetilde{E}$.

Since there are $\ell+1$ cyclic subgroups of $E[\ell]$ it follows that $\Phi_\ell(j(E),y)$ has degree $\ell+1$. Indeed, $\Phi_\ell(x,y) = x^{\ell+1} + y^{\ell+1} - (xy)^\ell + \cdots$ with all other terms of lower degree (see Theorem 11.18 of Cox [157] or Theorem 3 of Section 5.2 of Lang [366]). The coefficients of $\Phi_\ell(x,y)$ are large (as seen in Example 25.2.1, even when $\ell = 2$ the coefficients are large).

Example 25.2.1.

$$\begin{aligned}\Phi_2(x,y) &= x^3 + y^3 - x^2y^2 + 1488(x^2y + xy^2) - 162000(x^2 + y^2) \\ &\quad +40773375xy + 8748000000(x+y) - 157464000000000.\end{aligned}$$

Let ℓ be prime. Cohen [138] showed that the number of bits in the largest coefficient of $\Phi_\ell(x,y)$ is $O(\ell \log(\ell))$ (see Bröker and Sutherland [110] for a more precise bound). Since there are roughly ℓ^2 coefficients it follows that $\Phi_\ell(x,y)$ can be written down using $O(\ell^3 \log(\ell))$ bits, and it is believed that this space requirement cannot be reduced. Hence, one expects to perform at least $O(\ell^3 \log(\ell)) = O(\ell^{3+\epsilon})$ bit operations[3] to compute $\Phi_\ell(x,y)$. Indeed, using methods based on modular functions one can conjecturally[4] compute $\Phi_\ell(x,y)$ in $O(\ell^{3+\epsilon})$ bit operations (see Enge [194]). Using modular functions other than the j-function can lead to polynomials with smaller coefficients, but this does not affect the asymptotic complexity.

The fastest method to compute modular polynomials is due to Bröker, Lauter and Sutherland [109]. This method exploits some of the ideas explained later in this chapter (in particular, isogeny volcanoes). The method computes $\Phi_\ell(x,y)$ modulo small primes and then determines $\Phi_\ell(x,y)$ by the Chinese remainder theorem. Under the Generalized Riemann Hypothesis (GRH) the complexity is $O(\ell^3 \log(\ell)^3 \log(\log(\ell)))$ bit operations. For the rest of the chapter we abbreviate the cost as $O(\ell^{3+\epsilon})$ bit operations. The method can also be used to compute $\Phi_\ell(x,y)$ modulo p, in which case the space requirements are $O(\ell^2 \log(\ell)^2 + \ell^2 \log(p))$ bits.

The upshot is that, given an elliptic curve E over $\Bbbk$, the j-invariants of elliptic curves $\widetilde{E}$ that are ℓ-isogenous over $\Bbbk$ (where $\gcd(\ell, \mathrm{char}(\Bbbk)) = 1$) are given by the roots of

[2]The reader should not confuse the modular polynomial $\Phi_\ell(x,y)$ with the cyclotomic polynomial $\Phi_m(x)$.

[3]Recall that a function $f(\ell)$ is $O(\ell^{3+\epsilon})$ if, for every $\epsilon > 0$, there is some $C(\epsilon), L(\epsilon) \in \mathbb{R}_{>0}$ such that $f(\ell) < C(\epsilon)\ell^{3+\epsilon}$ for all $\ell > L(\epsilon)$.

[4]Enge needs an assumption that rounding errors do not affect the correctness of the output.

$\Phi_\ell(j(E), y)$ in $\Bbbk$. When E is ordinary, Theorem 25.4.6 implies that $\Phi_\ell(j(E), y)$ has either $0, 1, 2$ or $\ell + 1$ roots in $\Bbbk$ (counted with multiplicities).

Exercise 25.2.2. Given the polynomial $\Phi_\ell(x, y)$ and a value $j \in \mathbb{F}_q$ show that one can compute $F(y) = \Phi_\ell(j, y) \in \mathbb{F}_q[y]$ in $O(\ell^2(\ell \log(\ell) \log(q) + M(\log(q))))$ bit operations. Show also that one can then compute the roots $\tilde{j} \in \mathbb{F}_q$ of $F(y) = \Phi_\ell(j(E), y)$ (or determine that there are no roots) in expected time bounded by $O(\ell^2 \log(\ell) \log(q))$ field operations (which is $O(\ell^{2+\epsilon} \log(q)^3)$ bit operations).

For the rest of this section we consider algorithms to compute an ℓ-isogeny $\phi : E \to \widetilde{E}$ given an elliptic curve E and the j-invariant of $\widetilde{E}$.

Exercise 25.2.3. Let E be an elliptic curve over $\mathbb{F}_q$ and let E' over $\mathbb{F}_q$ be a twist of E. Show that there is an $\mathbb{F}_q$-rational isogeny of degree ℓ from E (to some elliptic curve) if and only if there is an $\mathbb{F}_q$-rational isogeny of degree ℓ from E'. Show that $\mathrm{End}(E) \cong \mathrm{End}(E')$ (where $\cong$ denotes ring isomorphism).

25.2.1 Elkies' Algorithm

Let $\ell > 2$ be a prime and let E be an elliptic curve over $\Bbbk$ where $\mathrm{char}(\Bbbk) = 0$ or $\mathrm{char}(\Bbbk) > \ell + 2$. Assume $j(E) \neq 0, 1728$ (for the case $j(E) \in \{0, 1728\}$ one constructs isogenies using the naive method or the methods of the following sections). Let $\tilde{j} \in \Bbbk$ be such that $\Phi_\ell(j(E), \tilde{j}) = 0$. We also assume that $\tilde{j}$ is a simple root of $\Phi_\ell(j(E), y)$ (more precisely, $(\partial\Phi_\ell(x, y)/\partial x)(j, \tilde{j}) \neq 0$ and $(\partial\Phi_\ell(x, y)/\partial y)(j, \tilde{j}) \neq 0$); see page 248 of Schoof [530] for a discussion of why this condition is not too severe.

Elkies gave a method to determine an explicit equation for an elliptic curve $\widetilde{E}$, such that $j(\widetilde{E}) = \tilde{j}$, and a polynomial giving the kernel of an ℓ-isogeny from E to $\widetilde{E}$. Elkies' original motivation (namely, algorithms for point counting) only required computing the kernel polynomial of the isogeny, but as we have seen, from this information one can easily compute the rational functions describing the isogeny. The method also works when $\ell > 2$ is composite, but that is not of practical relevance. The condition that $\mathrm{char}(\Bbbk)$ not be small (if it is non-zero) is essential.

We use the same notation as in Lemma 25.1.16: $\psi(x)$ is the polynomial of degree $(\ell - 1)/2$ whose roots are the x-coordinates of affine points in the kernel G of the isogeny and s_i are the i-th symmetric polynomials in these roots. We also define

$$p_i = \sum_{P \in G - \{\mathcal{O}_E\}} x_P^i$$

so that $p_1 = 2s_1$ and $p_2 = 2(s_1^2 - 2s_2)$ (these are Newton's formulae; see Lemma 10.7.6). While the value $\tilde{j}$ specifies the equation for the isogenous curve $\widetilde{E}$ (up to isomorphism) it does not, in general, determine the isogeny (see pages 37 and 44 of Elkies [193] for discussion). It is necessary to have some extra information, and for this the coefficient p_1 suffices and can be computed using partial derivatives of the modular polynomial (this is why the condition on the partial derivatives is needed).

The explanation of Elkies' algorithm requires theory that we do not have space to present. We refer to Schoof [530] for a good summary of the details (also see Elkies [193] for further discussion). The basic idea is to use the fact (Deuring lifting theorem) that the isogeny lifts to an isogeny between elliptic curves over $\mathbb{C}$. One can then interpret the ℓ-isogeny in terms of Tate curves[5] $\mathbb{C}^*/q^{\mathbb{Z}}$ (we have not presented the Tate curve in this

[5]The notation q here refers to $q(z) = \exp(2\pi i z)$ and not a prime power.

book; see Section C.14 of [564] or Section 5.3 of [565]) as a map from $q(z)$ to $q(z)^\ell$. As discussed on page 40 of Elkies [193], this isogeny is not normalised. There are a number of relations between the modular j-function, certain Eisenstein series, the equation of the elliptic curve (in short Weierstrass form) and the kernel polynomial of the isogeny. These relations give rise to formulae that must also hold over $\Bbbk$. Hence, one can work entirely over $\Bbbk$ and obtain the kernel polynomial.

The details of this approach are given in Sections 7 and 8 of Schoof [530]. In particular: Theorem 7.3 shows how to get j' (derivative); Proposition 7.1 allows one to compute the coefficients of the elliptic curve; Proposition 7.2 gives the coefficient p_1 of the kernel polynomial (which is a function of values specified in Proposition 7.1 and Theorem 7.3). The coefficients of the kernel polynomial are related to the coefficients of the series expansion of the Weierstrass ζ-function (see Theorem 8.3 of [530]).

The algorithm is organised as follows (see Algorithm 28). One starts with an ordinary elliptic curve $E : y^2 = x^3 + Ax + B$ over $\Bbbk$ and $j = j(E)$. We assume that $j \notin \{0, 1728\}$ and $\mathrm{char}(\Bbbk) = 0$ or $\mathrm{char}(k) > \ell + 2$. Let $\phi_x = (\frac{\partial \Phi_\ell}{\partial x})(j, \tilde{\jmath})$, $\phi_y = (\frac{\partial \Phi_\ell}{\partial y})(j, \tilde{\jmath})$, $\phi_{xx} = (\frac{\partial^2 \Phi_\ell}{\partial x^2})(j, \tilde{\jmath})$, $\phi_{yy} = (\frac{\partial^2 \Phi_\ell}{\partial y^2})(j, \tilde{\jmath})$ and $\phi_{xy} = (\frac{\partial^2 \Phi_\ell}{\partial x \partial y})(j, \tilde{\jmath})$. One computes the derivative j' and the corresponding values for E_4 and E_6. Given $\tilde{\jmath}$ one computes $\tilde{\jmath}'$ and then the coefficients $\tilde{A}$ and $\tilde{B}$ of the image curve $\tilde{E}$. Finally one computes p_1, from which it is relatively straightforward to compute all the coefficients of the kernel polynomial $\psi(x)$.

Algorithm 28 Elkies' algorithm (Source code provided by Drew Sutherland)

INPUT: $A, B \in \Bbbk$, $\ell > 2$, $j, \tilde{\jmath} \in \Bbbk$
OUTPUT: $\tilde{A}, \tilde{B}, \psi(x)$

Compute ϕ_x, ϕ_y, ϕ_{xx}, ϕ_{yy} and ϕ_{xy} ▷ Compute $\tilde{A}$ and $\tilde{B}$
Let $m = 18B/A$, let $j' = mj$, and let $k = j'/(1728 - j)$
Let $\tilde{\jmath}' = -j'\phi_x/(\ell\phi_y)$, let $\tilde{m} = \tilde{\jmath}'/\tilde{\jmath}$, and let $\tilde{k} = \tilde{\jmath}'/(1728 - \tilde{\jmath})$
Let $\tilde{A} = \ell^4 \tilde{m}\tilde{k}/48$ and $\tilde{B} = \ell^6 \tilde{m}^2 \tilde{k}/864$
Let $r = -(j'^2\phi_{xx} + 2\ell j'\tilde{\jmath}'\phi_{xy} + \ell^2\tilde{\jmath}'^2\phi_{yy})/(j'\phi_x)$ ▷ Compute p_1
Let $p_1 = \ell(r/2 + (k - \ell\tilde{k})/4 + (\ell\tilde{m} - m)/3)$
Let $d = (\ell - 1)/2$ ▷ Compute the power sums t_n of the roots of $\psi(x)$
Let $t_0 = d$, $t_1 = p_1/2$, $t_2 = ((1-10d)A - \tilde{A})/30$, and $t_3 = ((1-28d)B - 42t_1A - \tilde{B})/70$
Let $c_0 = 0$, $c_1 = 6t_2 + 2At_0$, $c_2 = 10t_3 + 6At_1 + 4Bt_0$
for $n = 2$ to $d - 1$ **do**
    Let $s = \sum_{i=1}^{n-1} c_i c_{n-i}$
    Let
$$c_{n+1} = \frac{3s - (2n-1)(n-1)Ac_{n-1} - (2n-2)(n-2)Bc_{n-2}}{(n-1)(2n+5)}$$
end for
for $n = 3$ to $d - 1$ **do**
    Let
$$t_{n+1} = \frac{c_n - (4n-2)At_{n-1} - (4n-4)Bt_{n-2}}{4n+2}$$
end for
Let $s_0 = 1$ ▷ Compute the symmetric functions s_n of the roots of $\psi(x)$
for $n = 1$ to d **do**
    Let $s_n = \frac{-1}{n} \sum_{i=1}^{n} (-1)^i t_i s_{n-i}$
end for
return $\psi(x) = \sum_{i=0}^{d} (-1)^i s_i x^{d-i}$

Exercise 25.2.4. Show that Elkies' algorithm requires $O(d^2) = O(\ell^2)$ operations in $\mathbb{k}$.

Bostan, Morain, Salvy and Schost [93] have given algorithms (exploiting fast arithmetic on power series) based on Elkies' methods. The algorithms apply when the characteristic of the field is zero or is sufficiently large compared with ℓ. There is a slight difference in scope: Elkies' starts with only j-invariants whereas Bostan et al assume that one is given elliptic curves E and $\widetilde{E}$ in short Weierstrass form such that there is a normalised isogeny of degree ℓ over $\mathbb{k}$ from E to $\widetilde{E}$. In general, one needs to perform Elkies' method before one has such an equation for $\widetilde{E}$ and so the computations with modular curves still dominate the cost. Theorem 2 of [93] states that one can compute the rational functions giving the isogeny in $O(M(\ell))$ operations in $\mathbb{k}$ when $\mathrm{char}(\mathbb{k}) > 2\ell - 1$ and when the coefficient p_1 is known. Note that Bostan et al are not restricted to prime degree isogenies. An application of the result of Bostan et al is to determine whether there is a normalised isogeny from E to $\widetilde{E}$ *without* needing to compute modular polynomials. Lercier and Sirvent [384] (again, assuming one is given explicit equations for E and $\widetilde{E}$ such that there is a normalised ℓ-isogeny between them) have showed how to achieve a similarly fast method even when the characteristic of the field is small.

A number of calculations can fail when $\mathrm{char}(\mathbb{k})$ is non-zero but small compared with ℓ, due to divisions by small integers arising in the power series expansions for the modular functions. Algorithms for the case of small characteristic will be mentioned in Section 25.2.3.

25.2.2 Stark's Algorithm

Stark [581] gave a method to compute the rational function giving the x-coordinate of an endomorphism $\phi : E \to E$ corresponding to a complex number β (interpreting $\mathrm{End}(E)$ as a subset of $\mathbb{C}$). The idea is to use the fact that, for an elliptic curve E over the complex numbers given by short Weierstrass form,

$$\wp(\beta z) = \frac{A(\wp(z))}{B(\wp(z))} \tag{25.3}$$

where A and B are polynomials and where $\wp(z) = z^{-2} + 3G_4 z^2 + \cdots$ is the Weierstrass function (see Theorem VI.3.5 of Silverman [564]). This isogeny is not normalised (since $\wp(\beta z) = \beta^{-2} z^{-2} + \cdots$ it follows that the pullback of ω_E under ϕ is $\beta\omega_E$). Stark's idea is to express $\wp$ as a (truncated) power series in z; the coefficients of this power series are determined by the coefficients of the elliptic curve E. One computes A and B by taking the continued fraction expansion of the left hand side of equation (25.3). One can apply this algorithm to curves over finite fields by applying the Deuring lifting theorem. Due to denominators in the power series coefficients of $\wp(z)$ the method only works when $\mathrm{char}(\mathbb{k}) = 0$ or $\mathrm{char}(\mathbb{k})$ is sufficiently large. Stark's paper [581] gives a worked example in the case $\beta = \sqrt{-2}$.

The idea generalises to normalised isogenies $\phi : E \to \widetilde{E}$ by writing $\wp_{\widetilde{E}}(z) = A(\wp_E(z))/B(\wp_E(z))$ where now the power series for $\wp_E(z)$ and $\wp_{\widetilde{E}}(z)$ are different since the elliptic curve equations are different. Note that it is necessary to have actual curve equations for the normalised isogeny, not just j-invariants. We refer to Section 6.2 of Bostan, Morain, Salvy and Schost [93] for further details and complexity estimates.

25.2.3 The Small Characteristic Case

As we have seen, the Elkies and Stark methods require the characteristic of the ground field to be either zero or relatively large since they use lifting to short Weierstrass forms

over $\mathbb{C}$ and since the power series expansions have rational coefficients that are divisible by various small primes. Hence, there is a need for algorithms that handle the case when $\mathrm{char}(\Bbbk)$ is small (especially, $\mathrm{char}(\Bbbk) = 2$). A number of such methods have been developed by Couveignes, Lercier, Morain and others. We briefly sketch Couveignes' "second method" [155].

Let p be the characteristic of the field. We assume that p is "small" (in the sense that an algorithm performing p operations is considered efficient). Let E and $\widetilde{E}$ be ordinary[6] elliptic curves over $\mathbb{F}_{p^m}$.

The basic idea is to use the fact that if $\phi : E \to \widetilde{E}$ is an isogeny of odd prime degree $\ell \neq p$ (isogenies of degree p are easy: they are either Frobenius or Verschiebung) then ϕ maps points of order p^k on E to points of order p^k on $\widetilde{E}$. Hence, one can try to determine the rational functions describing ϕ by interpolation from their values on $E[p^k]$. One could interpolate using any torsion subgroup of E, but using $E[p^k]$ is the best choice since it is a cyclic group and so there are only $\varphi(p^k) = p^k - p^{k-1}$ points to check (compared with $\varphi(n^2)$ if using $E[n]$). The method can be applied to any elliptic curve $\widetilde{E}$ in the isomorphism class, so in general it will not return a normalised isogeny.

Couveignes' method is as follows: First, compute points $P \in E[p^k] - E[p^{k-1}]$ and $\widetilde{P} \in \widetilde{E}[p^k] - \widetilde{E}[p^{k-1}]$ over $\overline{\mathbb{F}}_p$ and guess that $\phi(P) = \widetilde{P}$. Then try to determine the rational function $\phi_1(x) = u(x)/v(x)$ by interpolating $\phi_1(x([i]P)) = x([i]\widetilde{P})$; if this does not work then try another guess for $\widetilde{P}$. The interpolation is done as follows (we assume $p^k > 2\ell$). First, compute a polynomial $A(x)$ of degree d where $2\ell < d \leq p^k$ such that $A(x([i]P)) = x([i]\widetilde{P})$ for $1 \leq i \leq d$. Also compute $B(x) = \prod_{i=1}^{d}(x - x([i]P))$. If the guess for $\widetilde{P}$ is correct then $A(x) \equiv u(x)/v(x) \pmod{B(x)}$ where $\deg(u(x)) = \ell, \deg(v(x)) = \ell - 1$ and $v(x)$ is a square. Writing this equation as $A(x)v(x) = u(x) + B(x)w(x)$ it follows that $u(x)$ and $v(x)$ can be computed using Euclid's algorithm. The performance of the algorithm depends on the method used to determine points in $E[p^k]$, but is dominated by the fact that these points lie in an extension of the ground field of large degree, and that one expects to try around $\frac{1}{2}p^k \approx \ell$ choices for $\widetilde{P}$ before hitting the right one. The complexity is polynomial in ℓ, p and m (where E is over $\mathbb{F}_{p^m}$). When $p = 2$ the fastest method was given by Lercier [382]. For further details we refer to the surveys by Lercier and Morain [383] and De Feo [202].

25.3 Isogeny Graphs of Elliptic Curves over Finite Fields

Let E be an elliptic curve over $\mathbb{F}_q$. The $\mathbb{F}_q$-**isogeny class** of E is the set of $\mathbb{F}_q$-isomorphism classes of elliptic curves over $\mathbb{F}_q$ that are isogenous over $\mathbb{F}_q$ to E. Tate's isogeny theorem states that two elliptic curves E and $\widetilde{E}$ over $\mathbb{F}_q$ are $\mathbb{F}_q$-isogenous if and only if $\#E(\mathbb{F}_q) = \#\widetilde{E}(\mathbb{F}_q)$ (see Theorem 9.7.4 for one implication).

We have seen in Sections 9.5 and 9.10 that the number of $\mathbb{F}_q$-isomorphism classes of elliptic curves over $\mathbb{F}_q$ is roughly $2q$ and that there are roughly $4\sqrt{q}$ possible values for $\#E(\mathbb{F}_q)$. Hence, if isomorphism classes were distributed uniformly across all group orders one would expect around $\frac{1}{2}\sqrt{q}$ elliptic curves in each isogeny class. The theory of complex multiplication gives a more precise result (as mentioned in Section 9.10.1). We

[6] The restriction to ordinary curves is not a significant problem. In practice we are interested in elliptic curves over $\mathbb{F}_{p^m}$ where m is large, whereas supersingular curves are all defined over $\mathbb{F}_{p^2}$. Indeed, for small p there are very few supersingular curves, and isogenies of small degree between them can be computed by factoring division polynomials and using Vélu's formulae.

denote by π_q the q-power Frobenius map; see Section 9.10 for its properties. The number of $\mathbb{F}_q$-isomorphism classes of ordinary elliptic curves over $\mathbb{F}_q$ with $q+1-t$ points is the Hurwitz class number of the ring $\mathbb{Z}[\pi_q] = \mathbb{Z}[(t+\sqrt{t^2-4q})/2]$. This is the sum of the ideal class numbers $h(\mathcal{O})$ over all orders $\mathbb{Z}[\pi_q] \subseteq \mathcal{O} \subseteq \mathcal{O}_K$. It follows (see Remark 9.10.19) that there are $O(q^{1/2}\log(q)\log(\log(q)))$ elliptic curves in each isogeny class. For supersingular curves see Theorem 9.11.12.

Definition 25.3.1. Let E be an elliptic curve over a field $\Bbbk$ of characteristic p. Let $S \subseteq \mathbb{N}$ be a finite set of primes. Define

$$X_{E,\Bbbk,S}$$

to be the (directed) graph[7] with vertex set being the $\Bbbk$-isogeny class of E. Vertices are typically labelled by $j(\widetilde{E})$, though we also speak of "the vertex $\widetilde{E}$".[8] There is a (directed) edge $(j(E_1), j(E_2))$ labelled by ℓ for each equivalence class of ℓ-isogenies from E_1 to E_2 defined over $\Bbbk$ for some $\ell \in S$. We usually treat this as an undirected graph, since for every ℓ-isogeny $\phi : E_1 \to E_2$ there is a dual isogeny $\widehat{\phi} : E_2 \to E_1$ of degree ℓ (though see Remark 25.3.2 for an unfortunate, though rare, complication).

Remark 25.3.2. Edges in the isogeny graph correspond to equivalence classes of isogenies. It can happen that two non-equivalent isogenies from $E_1 \to E_2$ have equivalent dual isogenies from $E_2 \to E_1$. It follows that there are two directed edges in the graph from E_1 to E_2 but only one directed edge from E_2 to E_1. (Note that this does not contradict the fact that isogenies satisfy $\widehat{\widehat{\phi}} = \phi$, as we are speaking here about isogenies up to equivalence.) Such an issue was already explained in Exercise 25.1.1; it only arises if $\#\mathrm{Aut}(E_2) > 2$ (i.e., if $j(E_2) = 0, 1728$).

Definition 25.3.3. A (directed) graph is **k-regular** if every vertex has (out-)degree k (a loop is considered as having degree 1). A **path** in a graph is a sequence of (directed) edges between vertices, such that the end vertex of one edge is the start vertex of the next. We will also describe a path as a sequence of vertices. A graph is connected if there is a path from every vertex to every other vertex. The **diameter** of a connected graph is the maximum, over all pairs v_1, v_2 of vertices in the graph, of the length of the shortest path from v_1 to v_2.

There are significant differences (both in the structure of the isogeny graph and the way it is used in applications) between the ordinary and supersingular cases. So we present them separately.

25.3.1 Ordinary Isogeny Graph

Fix an ordinary elliptic curve E over $\mathbb{F}_q$ such that $\#E(\mathbb{F}_q) = q+1-t$. The isogeny graph of elliptic curves isogenous over $\mathbb{F}_q$ to E can be identified, using the theory of complex multiplication, with a graph whose vertices are ideal classes (in certain orders). The goal of this section is to briefly sketch this theory in the special case (the general case is given in Section 25.4) of the sub-graph where all elliptic curves have the same endomorphism ring, in which case the edges correspond to multiplication by prime ideals. We do not

[7] Some authors would use the name "multi-graph", since there can be loops and/or multiple edges between vertices.

[8] In the supersingular case one can label vertices as $j(\widetilde{E})$ without ambiguity only when $\Bbbk$ is algebraically closed: when $\Bbbk$ is finite then, in the supersingular case, one has two distinct vertices in the graph for $\widetilde{E}$ and its quadratic twist. For the ordinary case there is no ambiguity by Lemma 9.11.13 (also see Exercise 25.3.8).

have space to give all the details; good references for the background are Cox [157] and Lang [366].

The endomorphism ring of E (over $\overline{\mathbb{F}}_q$) is an order $\mathcal{O}$ in the quadratic imaginary field $K = \mathbb{Q}(\sqrt{t^2-4q})$. (We refer to Section A.12 for background on orders and conductors.) Let $\mathcal{O}_K$ be the ring of integers of K. Then $\mathbb{Z}[\pi_q] = \mathbb{Z}[(t+\sqrt{t^2-4q})/2] \subseteq \mathcal{O} \subseteq \mathcal{O}_K$ and if $\mathcal{O}_K = \mathbb{Z}[\theta]$ then $\mathcal{O} = \mathbb{Z}[c\theta]$ where c is the conductor of $\mathcal{O}$. The ideal class group $\mathrm{Cl}(\mathcal{O})$ is defined to be the classes of invertible $\mathcal{O}$-ideals that are prime to the conductor; see Section 7 of [157] or Section 8.1 of [366]. There is an explicit formula for the order $h(\mathcal{O})$ of the ideal class group $\mathrm{Cl}(\mathcal{O})$ in terms of the class number $h(\mathcal{O}_K)$ of the number field; see Theorem 7.24 of [157] or Theorem 8.7 of [366].

There is a one-to-one correspondence between the set of isomorphism classes of elliptic curves E over $\mathbb{F}_q$ with $\mathrm{End}(E) = \mathcal{O}$ and the set $\mathrm{Cl}(\mathcal{O})$. Precisely, to an invertible $\mathcal{O}$-ideal $\mathfrak{a}$ one associates the elliptic curve $E = \mathbb{C}/\mathfrak{a}$ over $\mathbb{C}$. An $\mathcal{O}$-ideal $\mathfrak{a}'$ is equivalent to $\mathfrak{a}$ in $\mathrm{Cl}(\mathcal{O})$ if and only if $\mathbb{C}/\mathfrak{a}'$ is isomorphic to E. One can show that $\mathrm{End}(E) = \mathcal{O}$. The theory of complex multiplication shows that E is defined over a number field (called the ring class field) and has good reduction modulo the characteristic p of $\mathbb{F}_q$. This correspondence is not canonical, since the reduction modulo p map is not well-defined (it depends on a choice of prime ideal above p in the ring class field).

Let $\mathfrak{a}$ be an invertible $\mathcal{O}$-ideal and $E = \mathbb{C}/\mathfrak{a}$. Let $\mathfrak{l}$ be an invertible $\mathcal{O}$-ideal and, interpreting $\mathfrak{l} \subseteq \mathrm{End}(E)$, consider the set $E[\mathfrak{l}] = \{P \in E(\mathbb{C}) : \phi(P) = \mathcal{O}_E \text{ for all } \phi \in \mathfrak{l}\}$. Since $\mathcal{O} \subseteq \mathbb{C}$ we can interpret $\mathfrak{l} \subseteq \mathbb{C}$, in which case

$$\begin{aligned} E[\mathfrak{l}] &\cong \{z \in \mathbb{C}/\mathfrak{a} : \alpha z \in \mathfrak{a}, \text{ for all } \alpha \in \mathfrak{l}\} \\ &\cong \mathfrak{l}^{-1}\mathfrak{a}/\mathfrak{a}. \end{aligned}$$

It follows that $\#E[\mathfrak{l}]$ is equal to the norm of the ideal $\mathfrak{l}$. The identity map on $\mathbb{C}$ induces the isogeny

$$\mathbb{C}/\mathfrak{a} \to \mathbb{C}/\mathfrak{l}^{-1}\mathfrak{a}$$

with kernel $\mathfrak{l}^{-1}\mathfrak{a}/\mathfrak{a} \cong E[\mathfrak{l}]$. The above remarks apply to elliptic curves over $\mathbb{C}$, but the theory reduces well to elliptic curves over finite fields, and indeed, every isogeny from E to an elliptic curve $\widetilde{E}$ with $\mathrm{End}(E) = \mathrm{End}(\widetilde{E})$ arises in this way. This shows that, not only do ordinary elliptic curves correspond to ideals in $\mathcal{O}$, but so do their isogenies.

Exercise 25.3.4. Show that if $\mathfrak{l} = (\ell)$ where $\ell \in \mathbb{N}$ then the isogeny $\mathbb{C}/\mathfrak{a} \to \mathbb{C}/\mathfrak{l}^{-1}\mathfrak{a}$ is $[\ell]$.

Exercise 25.3.5. Suppose the prime ℓ splits in $\mathcal{O}$ as $(\ell) = \mathfrak{l}_1\mathfrak{l}_2$ in $\mathcal{O}$. Let $\phi : E \to \widetilde{E}$ correspond to the ideal $\mathfrak{l}_1$. Show that $\widehat{\phi}$ corresponds to $\mathfrak{l}_2$.

Let ℓ be a prime. Then ℓ splits in $\mathcal{O}_K = \mathbb{Z}[\theta]$ if and only if the minimal polynomial of θ factors modulo ℓ with two linear factors. If D is the discriminant of K then ℓ splits if and only if the Kronecker symbol satisfies $(\frac{D}{\ell}) = +1$. Note that the Kronecker symbol is the Legendre symbol when ℓ is odd and

$$(\tfrac{D}{2}) = \begin{cases} 0 & D \equiv 0 \pmod 4, \\ 1 & D \equiv 1 \pmod 8, \\ -1 & D \equiv 5 \pmod 8. \end{cases} \tag{25.4}$$

Let E be an elliptic curve over $\mathbb{F}_q$ with $\mathrm{End}(E) = \mathcal{O}$ and let ℓ be coprime to the conductor of $\mathcal{O}$. There are $1 + (\frac{D}{\ell})$ prime ideals $\mathfrak{l}$ above ℓ, and so there are this many isogenies of degree ℓ from E. It follows that there are ℓ-isogenies in the isogeny graph for roughly half the primes ℓ.[9]

[9] Of course, there are still $\ell + 1$ isogenies of degree ℓ for each ℓ, but the rest of them are not to curves $\widetilde{E}$ such that $\mathrm{End}(\widetilde{E}) = \mathcal{O}$.

Let E be an elliptic curve over $\mathbb{F}_q$ corresponding to an $\mathcal{O}$-ideal $\mathfrak{a}$. Let $S \subseteq \mathbb{N}$ be a finite set of primes that are all co-prime to the conductor. Let G be the component of E in the isogeny graph $X_{E,\mathbb{F}_q,S}$ of Definition 25.3.1. Let $S' = \{\mathfrak{l}_1, \ldots, \mathfrak{l}_k\}$ be the set of classes of invertible $\mathcal{O}$-ideals above primes $\ell \in S$ and let $\langle S' \rangle$ be the subgroup of $\mathrm{Cl}(\mathcal{O})$ generated by S'. From the above discussion it follows that G can be identified with the graph whose vertices are the $\mathcal{O}$-ideal classes in the coset $\mathfrak{a}\langle S' \rangle$ and such that, for each $\mathfrak{b} \in \mathfrak{a}\langle S' \rangle$ and each $\mathfrak{l}_i \in S'$ there is an edge between $\mathfrak{b}$ and $\mathfrak{l}_i^{-1}\mathfrak{b}$. Since ideal class groups are well-understood, this correspondence illuminates the study of the isogeny graph. For example, an immediate corollary is that the graph of elliptic curves E with $\mathrm{End}(E) = \mathcal{O}$ is connected if and only if S' generates $\mathrm{Cl}(\mathcal{O})$. A well-known result of Bach states that (assuming the Riemann hypothesis for the Dedekind zeta function of K and Hecke L-functions for characters of $\mathrm{Cl}(\mathcal{O}_K)$) the group $\mathrm{Cl}(\mathcal{O}_K)$ is generated by prime ideals of norm less than $6\log(|\Delta_K|)^2$ (see page 376 of [20]) where Δ_K is the discriminant of $\mathcal{O}_K$. Another immediate corollary is that the graph is regular (i.e., every vertex has the same degree).

Remark 25.3.6. We stress that there is no canonical choice of $\mathcal{O}$-ideal $\mathfrak{a}$ corresponding to an elliptic curve E with $\mathrm{End}(E) = \mathcal{O}$. However, given a pair $(E, \widetilde{E})$ of isogenous elliptic curves with $\mathrm{End}(E) = \mathrm{End}(\widetilde{E}) = \mathcal{O}$ the ideal class corresponding to the isogeny between them is well-defined. More precisely, if E is identified with $\mathbb{C}/\mathfrak{a}$ for some $\mathcal{O}$-ideal $\mathfrak{a}$ then there is a unique ideal class represented by $\mathfrak{b}$ such that $\widetilde{E}$ is identified with $\mathbb{C}/\mathfrak{b}^{-1}\mathfrak{a}$. The only algorithm known to find such an ideal $\mathfrak{b}$ is to compute an explicit isogeny from E to $\widetilde{E}$ (using algorithms presented later in this chapter) and then determine the corresponding ideal. If one could determine $\mathfrak{b}$ efficiently from E and $\widetilde{E}$ then navigating the ordinary isogeny graph would be much easier.

Exercise 25.3.7. Let E_1 be an elliptic curve with $\mathrm{End}(E_1) = \mathcal{O}$. Let $\mathfrak{l}$ be a prime ideal of $\mathcal{O}$ above ℓ. Suppose $\mathfrak{l}$ has order d in $\mathrm{Cl}(\mathcal{O})$. Show that there is a cycle $E_1 \to E_2 \to \cdots \to E_d \to E_1$ of ℓ-isogenies.

Exercise 25.3.8. Let E be an ordinary elliptic curve over $\mathbb{F}_q$ and let E' be the quadratic twist of E. Show that the graphs $X_{E,\mathbb{F}_q,S}$ and $X_{E',\mathbb{F}_q,S}$ are identical.

Remark 25.3.9. Let ℓ split in $\mathcal{O} = \mathbb{Z}[\theta] \subseteq \mathrm{End}(E)$ and let $\mathfrak{l}_1 = (\ell, a+\theta)$ and $\mathfrak{l}_2 = (\ell, b+\theta)$ be the corresponding prime ideals. Given an isogeny $\phi : E \to \widetilde{E}$ of degree ℓ one can determine whether ϕ corresponds to $\mathfrak{l}_1$ or $\mathfrak{l}_2$ as follows: Compute (using the Elkies method if only $j(E)$ and $j(\widetilde{E})$ are known) the polynomial determining the kernel of ϕ; compute an explicit point $P \in \ker(\phi)$; check whether $[a]P + \theta(P) = \mathcal{O}_E$ or $[b]P + \theta(P) = \mathcal{O}_E$, where θ is now interpreted as an element of $\mathrm{End}(E)$. This trick is essentially due to Couveignes, Dewaghe, and Morain (see Section 3.2 of [156]; also see pages 49-50 of Kohel [350] and Galbraith, Hess and Smart [221]).

Remark 25.3.10. The ideas mentioned above show that all elliptic curves over $\mathbb{F}_q$ with the same endomorphism ring are isogenous over $\mathbb{F}_q$. Combined with the results of Section 25.4 one can prove Tate's isogeny theorem, namely that any two elliptic curves over $\mathbb{F}_q$ with the same number of points are isogenous over $\mathbb{F}_q$.

More details about the structure of the ordinary isogeny graph will be given in Section 25.4. In particular, that section will discuss isogenies between elliptic curves whose endomorphism rings are different orders in the same quadratic field.

25.3.2 Expander Graphs and Ramanujan Graphs

Let X be a finite (directed) graph on vertices labelled $\{1, \ldots, n\}$. The **adjacency matrix** of X is the $n \times n$ integer matrix A with $A_{i,j}$ being the number of edges from vertex i to vertex j. The **eigenvalues of a finite graph** X are defined to be the eigenvalues of its adjacency matrix A. For the rest of this section we assume that all graphs are un-directed. Since the adjacency matrix of an un-directed graph is real and symmetric, the eigenvalues are real.

Lemma 25.3.11. *Let X be a k-regular graph. Then k is an eigenvalue, and all eigenvalues λ are such that $|\lambda| \leq k$.*

Proof: The first statement follows since $(1, 1, \ldots, 1)$ is an eigenvector with eigenvalue k. The second statement is also easy (see Proposition 1.1.2 of Davidoff, Sarnak and Valette [166] or Theorem 1 of Murty [446]). □

Let X be a k-regular graph. We denote by $\lambda(X)$ the maximum of the absolute values of all the eigenvalues that are not equal to $\pm k$. Alon and Boppana showed that the lim inf of $\lambda(X)$ over any family of k-regular graphs (as the number of vertices goes to ∞) is at least $2\sqrt{k-1}$ (see Theorem 1.3.1 of Davidoff, Sarnak and Valette [166], Theorem 3.2 of Pizer [481] or Theorem 10 of Murty [446]). The graph X is said to be **Ramanujan** if $\lambda(X) \leq 2\sqrt{k-1}$. Define $\lambda_1(X)$ to be the largest eigenvalue that is strictly less than k.

Let G be a finite group and S a subset of G such that $g \in S$ implies $g^{-1} \in S$ (we also allow S to be a multi-set). The **Cayley graph** of G is the graph X with vertex set G and an edge between g and gs for all $g \in G$ and all $s \in S$. Murty [446] surveys criteria for when a Cayley graph is a Ramanujan graph. If certain character sums are small then X may be a Ramanujan graph; see Section 2 of [446].

Definition 25.3.12. Let X be a graph and A a subset of vertices of X. The **vertex boundary** of A in X is

$$\delta_v(A) = \{v \in X - A : \text{ there is an edge between } v \text{ and a vertex in } A\}.$$

Let E_X be the set of edges (x, y) in X. The **edge boundary** of A in X is

$$\delta_e(A) = \{(x, y) \in E_X : x \in A \text{ and } y \in X - A\}.$$

Let $c > 0$ be real. A k-regular graph X with $\#X$ vertices is a **c-expander** if, for all subsets $A \subseteq X$ such that $\#A \leq \#X/2$,

$$\#\delta_v(A) \geq c\#A.$$

Exercise 25.3.13. Show that $\delta_v(A) \leq \delta_e(A) \leq k\delta_v(A)$.

Exercise 25.3.14. Let X be a k-regular graph with n vertices that is a c-expander. Show that if n is even then $0 \leq c \leq 1$ and if n is odd then $0 \leq c \leq (n+1)/(n-1)$.

Expander graphs have a number of theoretical applications; one important property is that random walks on expander graphs reach the uniform distribution quickly.

Let X be a k-regular graph. Then

$$\#\delta_e(A) \geq \frac{k - \lambda_1(X)}{2}\#A \tag{25.5}$$

when $\#A \leq \#X/2$ (see Theorem 1.3.1 of Davidoff, Sarnak and Valette [166] or Section 4 of Murty [446][10]). Hence $\#\delta_v(A) \geq (\frac{1}{2} - \lambda_1(X)/(2k))\#A$ and so Ramanujan graphs are expander graphs. Indeed, small values for $\lambda_1(X)$ give large expansion factors. We refer to [166, 446] for further details, and references.

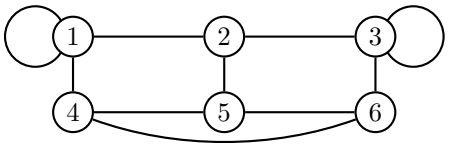

Figure 25.1: A 3-regular graph.

Exercise 25.3.15. Consider the 3-regular graph X in Figure 25.1. Determine the eigenvalues of X. Is this graph Ramanujan? Determine $\delta_v(\{1\})$, $\delta_v(\{1,2\})$ and $\delta_v(\{1,3\})$. Verify that $\#\delta_v(A) \geq \#A$ for all subsets A of vertices of X such that $\#A \leq 3$ and so X is an expander.

Exercise 25.3.16. For every $c > 0$ find an integer $n \in \mathbb{N}$, a graph X with n vertices, and a subset A of X such that $\#A \leq n/2$ but $\#\delta_v(A) < c\#A$. (Such a graph is very far from being an expander.)

Consider the ordinary isogeny graph of elliptic curves over $\mathbb{F}_q$ with $\mathrm{End}(E) = \mathcal{O}_K$, the ring of integers in $K = \mathbb{Q}(\sqrt{t^2 - 4q})$. This was shown in the previous section to be a Cayley graph for the ideal class group $\mathrm{Cl}(\mathcal{O}_K)$. Jao, Miller and Venkatesan [311] show, assuming a generalisation of the Riemann hypothesis, that the ordinary isogeny graph is an expander graph (indeed, it is "nearly Ramanujan", i.e., $\lambda_1(X) \leq O(k^\beta)$ for some $0 < \beta < 1$).

25.3.3 Supersingular Isogeny Graph

For the supersingular isogeny graph we work over $\overline{\mathbb{F}}_p$. The graph is finite. Indeed, Theorem 9.11.12 implies $p/12 - 1 < \#X_{E,\overline{\mathbb{F}}_p,S} < p/12 + 2$. Note that it suffices to consider elliptic curves defined over $\mathbb{F}_{p^2}$ (although the isogenies between them are over $\overline{\mathbb{F}}_p$ in general).

In contrast to the ordinary case, the supersingular graph is always connected using isogenies of any fixed degree. A proof of this result, attributed to Serre, is given in Section 2.4 of Mestre [419].

Theorem 25.3.17. *Let p be a prime and let E and $\widetilde{E}$ be supersingular elliptic curves over $\overline{\mathbb{F}}_p$. Let ℓ be a prime different from p. Then there is an isogeny from E to $\widetilde{E}$ over $\overline{\mathbb{F}}_p$ whose degree is a power of ℓ.*

Proof: See Corollary 78 of Kohel [350] or Section 2.4 of Mestre [419]. □

Hence, one can choose any prime ℓ (e.g., $\ell = 2$) and consider the ℓ-isogeny graph $X_{E,\overline{\mathbb{F}}_p,\{\ell\}}$ on supersingular curves over $\overline{\mathbb{F}}_p$. It follows that the graph is $(\ell+1)$-regular and connected.

Exercise 25.3.18. Let $p = 103$. Determine, using Theorem 9.11.12, the number of isomorphism classes of supersingular elliptic curves over $\mathbb{F}_p$ and over $\mathbb{F}_{p^2}$. Determine the 2-isogeny graph whose vertices are supersingular elliptic curves over $\mathbb{F}_{p^2}$.

Exercise 25.3.19. Determine the supersingular 2-isogeny graph over $\mathbb{F}_{11}$. Interpret the results in light of Remark 25.3.2.[11]
[Hint: The isomorphism classes of elliptic curves with $j(E) = 0$ and $j(E) = 1728$ are supersingular modulo 11; this follows from the theory of complex multiplication and the facts that $(\frac{-3}{11}) = (\frac{-4}{11}) = -1$.]

[10] Note that the proof on page 16 of [446] is for $\delta_e(A)$, not $\delta_v(A)$ as stated.
[11] This example was shown to me by David Kohel.

Exercise 25.3.20. Find a prime p such that the set of isomorphism classes of supersingular elliptic curves over $\mathbb{F}_p$ does not form a connected subgraph of $X_{E,\overline{\mathbb{F}}_p,\{2\}}$.

There is a one-to-one correspondence between supersingular elliptic curves E over $\overline{\mathbb{F}}_p$ and projective right modules of rank 1 of a maximal order of the quaternion algebra over $\mathbb{Q}$ ramified at p and ∞ (see Section 5.3 of Kohel [350] or Gross [268]). Pizer has exploited this structure (and connections with Brandt matrices and Hecke operators) to show that the supersingular isogeny graph is a Ramanujan graph. Essentially, the Brandt matrix gives the adjacency matrix of the graph. A good survey is [481], though be warned that the paper does not mention the connection to supersingular elliptic curves.

The supersingular isogeny graph has been used by Charles, Goren and Lauter [128] to construct a cryptographic hash function. It has also been used by Mestre and Oesterlé [419] for an algorithm to compute coefficients of modular forms.

25.4 The Structure of the Ordinary Isogeny Graph

This section presents Kohel's results on the structure of the isogeny graph of ordinary elliptic curves over finite fields. Section 25.4.2 gives Kohel's algorithm to compute $\mathrm{End}(E)$ for a given ordinary elliptic curve over a finite field.

25.4.1 Isogeny Volcanoes

Let E be an ordinary elliptic curve over $\mathbb{F}_q$ and let $\#E(\mathbb{F}_q) = q + 1 - t$. Denote by K the number field $\mathbb{Q}(\sqrt{t^2 - 4q})$ and by $\mathcal{O}_K$ the ring of integers of K. We know that $\mathrm{End}(E) = \mathrm{End}_{\overline{\mathbb{F}}_q}(E)$ is an order in $\mathcal{O}_K$ that contains the order $\mathbb{Z}[\pi_q] = \mathbb{Z}[(t+\sqrt{t^2-4q})/2]$ of discriminant $t^2 - 4q$. Let Δ_K be the discriminant of K, namely $\Delta_K = (t^2 - 4q)/c^2$ where c is the largest positive integer such that Δ_K is an integer congruent to 0 or 1 modulo 4. The integer c is the **conductor** of the order $\mathbb{Z}[\pi_q]$.

Suppose E_1 and E_2 are elliptic curves over $\mathbb{F}_q$ such that $\mathrm{End}(E_i) = \mathcal{O}_i$, for $i = 1, 2$, where $\mathcal{O}_1$ and $\mathcal{O}_2$ are orders in K containing $\mathbb{Z}[\pi_q]$. We now present some results about the isogenies between such elliptic curves.

Lemma 25.4.1. *Let $\phi : E \to \widetilde{E}$ be an isogeny of elliptic curves over $\mathbb{F}_q$. If $[\mathrm{End}(E) : \mathrm{End}(\widetilde{E})] = \ell$ (or vice versa) then the degree of ϕ is divisible by ℓ.*

Proof: See Propositions 21 and 22 of Kohel [350]. □

Definition 25.4.2. Let ℓ be a prime and E an elliptic curve. Let $\mathrm{End}(E) = \mathcal{O}$. An ℓ-isogeny $\phi : E \to \widetilde{E}$ is called **horizontal** (respectively, **ascending**, **descending**) if $\mathrm{End}(\widetilde{E}) \cong \mathcal{O}$ (respectively, $[\mathrm{End}(\widetilde{E}) : \mathcal{O}] = \ell$, $[\mathcal{O} : \mathrm{End}(\widetilde{E})] = \ell$).

Exercise 25.4.3. Let $\phi : E \to \widetilde{E}$ be an ℓ-isogeny. Show that if ϕ is horizontal (resp., ascending, descending) then $\widehat{\phi}$ is horizontal (resp., descending, ascending).

Example 25.4.4. We now give a picture of how the orders relate to one another. Suppose the conductor of $\mathbb{Z}[\pi_q]$ is 6 (e.g., $q = 31$ and $t = \pm 4$), so that $[\mathcal{O}_K : \mathbb{Z}[\pi_q]] = 6$. Write $\mathcal{O}_K = \mathbb{Z}[\theta]$. Then the orders $\mathcal{O}_2 = \mathbb{Z}[2\theta]$ and $\mathcal{O}_3 = \mathbb{Z}[3\theta]$ are contained in $\mathcal{O}_K$ and are

such that $[\mathcal{O}_K : \mathcal{O}_i] = i$ for $i = 2, 3$.

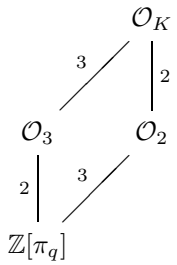

Definition 25.4.5. Let the notation be as above. If $\mathrm{End}(E) = \mathcal{O}_K$ then E is said to be on the **surface** of the isogeny graph.[12] If $\mathrm{End}(E) = \mathbb{Z}[\pi_q]$ then E is said to be on the **floor** of the isogeny graph.

By the theory of complex multiplication, the number of isomorphism classes of elliptic curves over $\mathbb{F}_q$ on the surface is equal to the ideal class number of the ring $\mathcal{O}_K$.

Theorem 25.4.6. *Let E be an ordinary elliptic curve over $\mathbb{F}_q$ as above and let $\mathcal{O} = \mathrm{End}(E)$ be an order in $\mathcal{O}_K$ containing $\mathbb{Z}[\pi_q]$. Let $c = [\mathcal{O}_K : \mathcal{O}]$ and let ℓ be a prime. Every ℓ-isogeny $\phi : E \to \widetilde{E}$ arises from one of the following cases.*

- *If $\ell \nmid c$ then there are exactly $(1+(\frac{t^2-4q}{\ell}))$ equivalence classes of horizontal ℓ-isogenies over $\mathbb{F}_q$ from E to other elliptic curves.*[13]
- *If $\ell \mid c$ then there are no horizontal ℓ-isogenies starting at E.*
- *If $\ell \mid c$ there is exactly one ascending ℓ-isogeny starting at E.*
- *If $\ell \mid [\mathcal{O} : \mathbb{Z}[\pi_q]]$ then the number of equivalence classes of ℓ-isogenies starting from E is $\ell + 1$, where the horizontal and ascending isogenies are as described and the remaining isogenies are descending.*
- *If $\ell \nmid [\mathcal{O} : \mathbb{Z}[\pi_q]]$ then there is no desending ℓ-isogeny.*

Proof: See Proposition 23 of Kohel [350]. A proof over $\mathbb{C}$ is also given in Appendix A.5 of [217]. □

Corollary 25.4.7. *Let E be an ordinary elliptic curve over $\mathbb{F}_q$ with $\#E(\mathbb{F}_q) = q + 1 - t$. Let c be the conductor of $\mathbb{Z}[\sqrt{t^2 - 4q}]$ and suppose $\ell \mid c$. Then $\ell \nmid [\mathrm{End}(E) : \mathbb{Z}[\pi_q]]$ if and only if there is a single ℓ-isogeny over $\mathbb{F}_q$ starting from E.*

Example 25.4.8. Let $q = 67$ and consider the elliptic curve $E : y^2 = x^3 + 11x + 21$ over $\mathbb{F}_q$. One has $\#E(\mathbb{F}_q) = 64 = q + 1 - t$ where $t = 4$ and $t^2 - 4q = 2^2 \cdot 3^2 \cdot (-7)$. Further, $j(E) = 42 \equiv -3375 \pmod{67}$, so E has complex multiplication by $(1 + \sqrt{-7})/2$. Since the ideal class number of $\mathbb{Q}(\sqrt{-7})$ is 1, it follows that E is the unique elliptic curve up to isomorphism on the surface of the isogeny graph.

Since 2 splits in $\mathbb{Z}[(1 + \sqrt{-7})/2]$ there are two 2-isogenies from E to elliptic curves on the surface (i.e., to E itself) and so there is only one 2-isogeny down from E. Using the modular polynomial we deduce that the 2-isogeny down maps to the isomorphism class

[12] Kohel's metaphor was intended to be aquatic: the floor represents the ocean floor and the surface represents the surface of the ocean.

[13] The symbol $(\frac{t^2-4q}{\ell})$ is the Kronecker symbol as in equation (25.4).

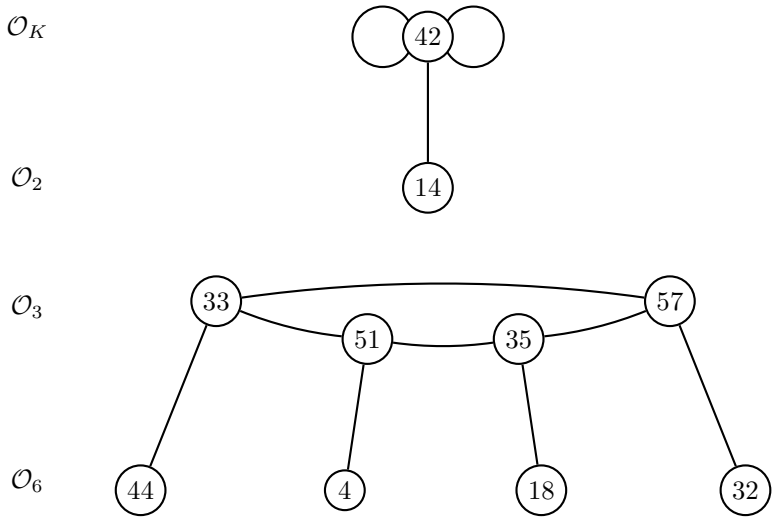

Figure 25.2: A 2-isogeny graph with two volcanoes. The symbols on the left hand side denote the endomorphism ring of curves on that level, using the same notation as Example 25.4.4.

of elliptic curves with j-invariant 14. One can verify that the only 2-isogeny over $\mathbb{F}_q$ from $j = 14$ is the ascending isogeny back to $j = 42$.

We have $(\frac{-7}{3}) = -1$ so there are no horizontal 3-isogenies from E. Hence, we expect four 3-isogenies down from E. Using the modular polynomial we compute the corresponding j-invariants to be $33, 35, 51$ and 57. One can now consider the 2-isogeny graphs containing these elliptic curves on their surfaces. It turns out that the graph is connected, and that there is a cycle of horizontal 2-isogenies from $j = 33$ to $j = 51$ to $j = 35$ to $j = 57$. For each vertex we therefore only expect one 2-isogeny down to the floor. The corresponding j-invariants are $44, 4, 18$ and 32 respectively. Figure 25.2 gives the 2-isogeny graph in this case.

Exercise 25.4.9. Draw the 3-isogeny graph for the elliptic curves in Example 25.4.8. Is $X_{E,\mathbb{F}_q,\{2,3\}}$ connected? If so, what is its diameter?

Fix a prime $\ell \mid c$ where c is the conductor of $\mathbb{Z}[\pi_q]$. Consider the sub-graph of the isogeny graph corresponding to isogenies whose degree is equal to ℓ. We call this the ℓ-isogeny graph. This graph is often not connected (for example, it is not connected when c is not a power of ℓ or when primes above ℓ do not generate $\mathrm{Cl}(\mathcal{O}_K)$). Even when ℓ splits and c is 1 or a power of ℓ, the graph is often not connected (the graph is connected only when the prime ideals above ℓ generate the ideal class group). Theorem 25.4.6 shows that each component of the ℓ-isogeny graph has a particular shape (that Fouquet and Morain [209] call a **volcano**).

We now give a precise definition for volcanoes. Let $\#E(\mathbb{F}_q) = q + 1 - t$ and let c be the conductor of $\mathbb{Z}[\pi_q]$ and suppose $\ell^m \| c$. Let $K = \mathbb{Q}(\sqrt{t^2 - 4q})$ and denote by $\mathcal{O}_K$ the maximal order in K. A **volcano** is a connected component of the graph $X_{E,\mathbb{F}_q,\{\ell\}}$. A volcano has $m + 1$ "levels" $V_0, \ldots, V_m$, being sub-graphs of the ℓ-isogeny graph; where vertices in V_i (i.e., on level i) correspond to isomorphism classes of elliptic curves $\widetilde{E}$ such that $\ell^i \| [\mathcal{O}_K : \mathrm{End}(\widetilde{E})]$. In other words, V_0 is on the surface of this component of the ℓ-isogeny graph (but not necessarily on the surface of the entire isogeny graph $X_{E,\mathbb{F}_q,S}$) and V_m is on the floor of this component of the ℓ-isogeny graph (though, again, not

necessarily on the floor of the whole isogeny graph). The surface of a volcano (i.e., V_0) is also called the **crater**. The graph V_0 is a connected regular graph with each vertex of degree at most 2. For all $0 < i \leq m$ and every vertex $v \in V_i$ there is a unique edge from v "up" to a vertex in V_{i-1}. For all $0 \leq i < m$ and every $v \in V_i$, the degree of v is $\ell + 1$. Every vertex in V_m has degree 1.

25.4.2 Kohel's Algorithm (Ordinary Case)

Kohel used the results of Section 25.4.1 to develop deterministic algorithms for computing $\mathrm{End}(E)$ (i.e., determining the level) for an elliptic curve E over $\mathbb{F}_q$. We sketch the algorithm for ordinary curves. Two facts are of crucial importance in Kohel's algorithm. The first (Corollary 25.4.7) is that one can recognise the floor when standing on it. The second fact is that if one starts a chain of ℓ-isogenies with a descending isogeny, and avoids backtracking, then all the isogenies in the chain are descending.

Before going any further, we discuss how to compute a non-backtracking chain of ℓ-isogenies. Given $j(E)$ one can compute the j-invariants of ℓ-isogenous curves over $\mathbb{F}_q$ by computing the roots of $F(y) = \Phi_\ell(j(E), y)$ in $\mathbb{F}_q$. Recall that one computes $\Phi_\ell(x, y)$ in $O(\ell^{3+\epsilon})$ bit operations and finds the roots of $F(y)$ in $\mathbb{F}_q$ in expected time bounded by $O(\ell^2 \log(\ell) \log(q))$ operations in $\mathbb{F}_q$. Let $j_0 = j(E)$ and let j_1 be one of the roots of of $F(y)$. We want to find, if possible, $j_2 \in \mathbb{F}_q$ such that there are ℓ-isogenies from E to E_1 (where $j(E_1) = j_1$) and from E_1 to E_2 (where $j(E_2) = j_2$) and such that $j_2 \neq j_0$ (so the second isogeny is not the dual of the first). The trick is to find roots of $\Phi_\ell(j_1, y)/(y - j_0)$. This process can be repeated to compute a chain $j_0, j_1, j_2, \ldots$ of j-invariants of ℓ-isogenous curves. As mentioned earlier, an alternative approach to walking in the isogeny graph is to find $\mathbb{F}_q$-rational factors of the ℓ-division polynomial and use Vélu's formulae; this is less efficient in general and the method to detect backtracking is to compute the image curve using Vélu's formulae and then compute its j-invariant.

The basic idea of Kohel's algorithm is, for each prime ℓ dividing[14] the conductor of $\mathbb{Z}[\pi_q]$, to find a chain of ℓ-isogenies from E to an elliptic curve on the floor. Suppose ℓ is a prime and $\ell^m \| c$. Kohel (on page 46 of [350]) suggests to take two non-backtracking chains of ℓ-isogenies of length at most m from E. If E is on the floor then this is immediately detected. If E is not on the surface then at least one of the initial ℓ-isogenies is descending, so in at most m steps one finds oneself on the floor. So if after m steps neither chain of isogenies has reached the floor then it follows that we must have started on the surface (and some or all of the ℓ-isogenies in the chain were along the surface). Note that, apart from the algorithm for computing roots of polynomials, the method is deterministic.

Exercise 25.4.10. Let $E : y^2 = x^3 + 3x + 6$ over $\mathbb{F}_{37}$ be an elliptic curve. Note that $\#E(\mathbb{F}_{37}) = 37 + 1 - 4$ and $4^2 - 4 \cdot 37 = -2^4 \cdot 7$. Hence the conductor is 4. We have $j(E) = 10$. Using the modular polynomial one finds the following j-invariants of elliptic curves 2-isogenous to E: $11, 29, 31$. Further, there is a single 2-isogeny from j-invariants $11, 31$ (in both cases, back to $j = 10$). But from 29 there is a 2-isogeny to $j = 10$ and two 2-isogenies to $j = 29$. What is $\mathrm{End}(E)$? Give j-invariants for a curve on the floor and a curve on the surface.

The worst case of Kohel's algorithm is when the conductor is divisible by one or more very large primes ℓ (since determining the j-invariant of an ℓ-isogenous curve is polynomial in ℓ and so exponential in the input size). Since c can be as big as $\sqrt{q}$ the above method (i.e., taking isogenies to the floor) would therefore have worst-case complexity of at least

[14] It is necessary to find the square factors of $t^2 - 4q$, which can be done in deterministic time $\tilde{O}(q^{1/6})$; see Exercise 12.5.1.

$q^{1/2}$ bit operations (indeed, it would be $O(q^{3/2+\epsilon})$ operations in $\mathbb{F}_q$ if one includes the cost of generating modular polynomials). Kohel (pages 53 to 57 of [350]) noted that when ℓ is very large one can more efficiently resolve the issue of whether or not ℓ divides the conductor by finding elements in the ideal class group that are trivial for the maximal order but non-trivial for an order whose conductor is divisible by ℓ; one can then "test" such a relation using isogenies. Using these ideas Kohel proves in Theorem 24 of [350] that, assuming a certain generalisation of the Riemann hypothesis, his algorithm requires $O(q^{1/3+\epsilon})$ bit operations. Kohel also considers the case of supersingular curves.

Bisson and Sutherland [58] consider a randomised version of Kohel's method using ideas from index calculus algorithms in ideal class groups. Their algorithm has heuristic subexponential expected complexity of $O(L_q(1/2, \sqrt{3}/2))$ bit operations. We do not present the details.

Remark 25.4.11. An important remark is that neither of the two efficient ways to generate elliptic curves over finite fields is likely to lead to elliptic curves E such that the conductor of $\mathrm{End}(E)$ is divisible by a large prime.

- When generating elliptic curves by choosing a random curve over a large prime field and counting points, then $t^2 - 4q$ behaves like a random integer and so is extremely unlikely to be divisible by the square of a very large prime
- When using the CM method then it is automatic that the curves have $q+1-t$ points where $t^2 - 4q$ has a very large square factor. It is easy to arrange that the square factor is divisible by a large prime. However, the elliptic curve itself output by the CM method has $\mathrm{End}(E)$ being the maximal order. To get $\mathrm{End}(E)$ to be a non-maximal order one can either use class polynomials corresponding to non-maximal orders or use descending isogenies. Either way, it is infeasible to compute a curve E such that a very large prime divides the conductor of $\mathrm{End}(E)$. Furthermore, Kohel's algorithm is not needed in this case, since by construction one already knows $\mathrm{End}(E)$.

Hence, in practice, the potential problems with large primes dividing the conductor of $\mathrm{End}(E)$ do not arise. It is therefore safe to assume that determining $\mathrm{End}(E)$ in practice is easy.

25.5 Constructing Isogenies Between Elliptic Curves

The **isogeny problem for elliptic curves** is: given two elliptic curves E and $\widetilde{E}$ over $\mathbb{F}_q$ with the same number of points, to compute an isogeny between them. Solving this problem is usually considered in two stages:

1. Performing a pre-computation, that computes a chain of prime-degree isogenies from E to $\widetilde{E}$. The chain is usually computed as a sequence of explicit isogenies, though one could store just the "Elkies information" for each isogeny in the chain.
2. Given a specific point $P \in E(\overline{\mathbb{F}}_q)$ to compute the image of P under the isogeny.

The precomputation is typically slow, while it is desirable that the computation of the isogeny be fast (since it might be performed for a large number of points).

An algorithm to solve the isogeny problem, requiring exponential time and space in terms of the input size, was given by Galbraith [217]. For the case of ordinary elliptic curves, an improved algorithm with low storage requirements was given by Galbraith, Hess and Smart [221]. We briefly sketch both algorithms in this section.

We now make some preliminary remarks in the ordinary case. Let c_1 be the conductor of $\mathrm{End}(E)$ and c_2 the conductor of $\mathrm{End}(\widetilde{E})$. If there is a large prime ℓ that divides c_1 but not c_2 (or vice versa), then any isogeny between E and $\widetilde{E}$ will have degree divisible by ℓ and hence the isogeny will be slow to compute. Since the conductor is a square factor of $t^2 - 4q$ it can be, in theory, as big as $q^{1/2}$. It follows that one does not expect an efficient algorithm for this problem in general. However, as mentioned in Remark 25.4.11, in practice one can ignore this bad case and assume the primes dividing the conductor are moderate.

For the rest of this section, in the ordinary case, we assume that $\mathrm{End}(E) = \mathrm{End}(\widetilde{E}) = \mathcal{O}$. (If this is not the case then take vertical isogenies from E to reduce to it.) Usually $\mathcal{O}$ is the maximal order. This is desirable, because the class number of the maximal order is typically smaller (and never larger) than the class number of the sub-orders, and so the algorithm to find the isogeny works more quickly in this case. However, for the sake of generality we do not insist that $\mathcal{O}$ is the maximal order. The general case could appear if there is a very large prime dividing the conductor of $\mathcal{O}$.

25.5.1 The Galbraith Algorithm

The algorithm of Galbraith [217] finds a path between two vertices in the isogeny graph $X_{E,\Bbbk,S}$ using a breadth-first search (see Section 22.2 of [146]). This algorithm can be used in both the ordinary and supersingular cases. More precisely, one starts with sets $X_0 = \{j(E)\}$ and $Y_0 = \{j(\widetilde{E})\}$ (we are assuming the vertices of the isogeny graph are labelled by j-invariants) and, at step i, computes $X_i = X_{i-1} \cup \delta_v(X_{i-1})$ and $Y_i = Y_{i-1} \cup \delta_v(Y_{i-1})$ where $\delta_v(X)$ is the set of vertices in the graph that are connected to a vertex in X by an edge. Computing $\delta_v(X)$ involves finding the roots in $\Bbbk$ of $\Phi_\ell(j, y)$ for every $j \in X$ and every $\ell \in S$. In the supersingular case the set S of possible isogeny degrees usually consists of a single prime ℓ. In the ordinary case S could have as many as $\log(q)$ elements, and one might not compute the whole of $\delta_v(X)$ but just the boundary in a subgraph corresponding to a (random) subset of S. In either case, the cost of computing $\delta_v(X)$ is clearly proportional to $\#X$.[15] The algorithm stops when $X_i \cap Y_i \neq \varnothing$, in which case it is easy to compute the isogeny from E to $\widetilde{E}$.

Exercise 25.5.1. Write pseudocode for the above algorithm.

Under the (probably unreasonable) assumption that new values in $\delta_v(X_i)$ behave like uniformly chosen elements in the isogeny graph, one expects from the birthday paradox that the two sets have non-empty intersection when $\#X_i + \#Y_i \geq \sqrt{\pi \#X_{E,\Bbbk,S}}$. Since the graph is an expander, we know that $\#X_i = \#X_{i-1} + \#\delta_v(X_{i-1}) \geq (1 + c)\#X_{i-1}$ when X_{i-1} is small, and so $\#X_i \geq (1 + c)^i$.

In the supersingular case we have $\#X_{E,\Bbbk,S} = O(q)$ and in the ordinary case we have $\#X_{E,\Bbbk,S} = h(\mathcal{O}) = O(q^{1/2} \log(q))$. In both cases, one expects the algorithm to terminate after $O(\log(q))$ steps. Step i involves, for every vertex $j \in X_i$ (or $j \in \delta_v(X_{i-1})$) and every $\ell \in S$, computing roots of $\Phi_\ell(j, y)$ in $\mathbb{F}_q$. One can check that if all ℓ are polynomially bounded in $\log(q)$ then the expected number of bit operations is bounded by $\sqrt{\#X_{E,\Bbbk,S}}$ times a polynomial in $\log(q)$.

In the supersingular case the algorithm performs an expected $\tilde{O}(q^{1/2})$ bit operations. In the ordinary case, by Bach's result (and therefore, assuming various generalisations of the Riemann hypothesis) we can restrict to isogenies of degree at most $6 \log(q)^2$ and so each step is polynomial-time (the dominant cost of each step is finding roots of the

[15] When all $\ell \in S$ are used at every step, to compute $\delta_v(X_i)$ it suffices to consider only vertices $j \in \delta_v(X_{i-1})$.

modular polynomial; see Exercise 25.2.2). The total complexity is therefore an expected $\tilde{O}(q^{1/4})$ bit operations. The storage required is expected to be $O(q^{1/4}\log(q)^2)$ bits.

Exercise 25.5.2. Let $m \in \mathbb{N}$. Suppose all $\ell \in S$ are such that $\ell = O(\log(q)^m)$. Let $\phi : E \to \widetilde{E}$ be the isogeny output by the Galbraith algorithm. Show, under the same heuristic assumptions as above, that one can evaluate $\phi(P)$ for $P \in E(\mathbb{F}_q)$ polynomial-time.

Exercise 25.5.3. Isogenies of small degree are faster to compute than isogenies of large degree. Hence, the average cost to compute an ℓ-isogeny can be used as a weight for the edges in the isogeny graph corresponding to ℓ-isogenies. It follows that there is a well-defined notion of shortest path in the isogeny graph between two vertices. Show how Dijkstra's algorithm (see Section 24.3 of [146]) can be used to find a chain of isogenies between two elliptic curves that can be computed in minimal time. What is the complexity of this algorithm?

25.5.2 The Galbraith-Hess-Smart Algorithm

We now restrict to the ordinary isogeny graph and sketch the algorithm of Galbraith, Hess and Smart [221]. The basic idea is to replace the breadth-first search by a random walk, similar to that used in the kangaroo algorithm.

Let H be a hash function from $\mathbb{F}_q$ to a set S of prime ideals of small norm. One starts random walks at $x_0 = j(E)$ and $y_0 = j(\widetilde{E})$ and stores ideals $\mathfrak{a}_0 = (1)$, $\mathfrak{b}_0 = (1)$. One can think of $(x_0, \mathfrak{a}_0)$ as a "tame walk" and $(y_0, \mathfrak{b}_0)$ as a "wild walk". Each step of the algorithm computes new values x_i and y_i from x_{i-1} and y_{i-1}: To compute x_i set $\mathfrak{l} = H(x_{i-1})$ and $\ell = N(\mathfrak{l})$; find the roots of $\Phi_\ell(x_{i-1}, z)$; choose the root corresponding to the ideal $\mathfrak{l}$ (using the trick mentioned in Remark 25.3.9) and call it x_i. The same process is used (with the same function H) for the sequence y_i. The ideals are also updated as $\mathfrak{a}_i = \mathfrak{a}_{i-1}\mathfrak{l}$ (reduced in the ideal class group to some short canonical representation of ideals). If $x_i = y_j$ then the walks follow the same path. We designate certain elements of $\mathbb{F}_q$ as being distinguished points, and if x_i or y_i is distinguished then it is stored together with the corresponding ideal $\mathfrak{a}$ or $\mathfrak{b}$. After a walk hits a distinguished point there are two choices: it could be allowed to continue or it could be restarted at a j-invariant obtained by taking a short random isogeny chain (perhaps corresponding to primes not in S) from E or $\widetilde{E}$.

Once a collision is detected one has an isogeny corresponding to ideal $\mathfrak{a}$ from $j(E)$ to some j, and an isogeny corresponding to ideal $\mathfrak{b}$ from $j(\widetilde{E})$ to j. Hence, the ideal $\mathfrak{a}\mathfrak{b}^{-1}$ gives the isogeny from $j(E)$ to $j(\widetilde{E})$.

Stolbunov has noted that, since the ideal class group is Abelian, it is not advisable to choose S such that $\mathfrak{l}, \mathfrak{l}^{-1} \in S$ (since such a choice means that walks remain "close" to the original j-invariant, and cycles in the random walk might arise). It is also faster to use isogenies of small degree more often than those with large degree. We refer to Galbraith and Stolbunov [230] for further details.

The remaining problem is that the ideal $\mathfrak{a}\mathfrak{b}^{-1}$ is typically of large norm. By construction, it is a product of exponentially many small primes. Since the ideal class group is commutative, such a product has a short representation (storing the exponents for each prime), but this leads to an isogeny that requires exponential computation. The proposal from [221] is to represent ideals using the standard representation for ideals in quadratic fields, and to "smooth" the ideal using standard techniques from index calculus algorithms in ideal class groups. It is beyond the scope of this book to discuss these ideas in detail. However, we note that the isogeny then has subexponential length and uses

primes ℓ of subexponential degree. Hence, the second stage of the isogeny computation is subexponential-time; this is not as fast as it would be with the basic Galbraith algorithm. The idea of smoothing an isogeny has also been used by Bröker, Charles and Lauter [108] and Jao and Soukharev [312].

Since the ordinary isogeny graph is conjecturally an expander graph, we know that a random walk on it behaves close to the uniform distribution after sufficiently many steps. We make the heuristic assumption that the pseudorandom walk proposed above has this property when the number of different primes used is sufficiently large and the hash function H is good. Then, by the birthday paradox, one expects a collision after $\sqrt{\pi h(\mathcal{O})}$ vertices have been visited. As a result, the heuristic expected running time of the algorithm is $O(q^{1/4})$ isogeny chain steps, and the storage can be made low by making distinguished elements rare. The algorithm can be distributed: using L processors of equal power one solves the isogeny problem in $\tilde{O}(q^{1/4}/L)$ bit operations.

25.6 Relating the Discrete Logarithm Problem on Isogenous Curves

The main application of the algorithms in Section 25.5 is to relate the discrete logarithm problem on curves with the same number of points. More precisely, let E and $\widetilde{E}$ be elliptic curves over $\mathbb{F}_q$ with $\#E(\mathbb{F}_q) = \#\widetilde{E}(\mathbb{F}_q)$. Let r be a large prime dividing $\#E(\mathbb{F}_q)$. A natural question is whether the discrete logarithm problem (in the subgroup of order r) has the same difficulty in both groups $E(\mathbb{F}_q)$ and $\widetilde{E}(\mathbb{F}_q)$. To study this question one wants to reduce the discrete logarithm problem from $E(\mathbb{F}_q)$ to $\widetilde{E}(\mathbb{F}_q)$. If we have an isogeny $\phi : E \to \widetilde{E}$ of degree not divisible by r, and if ϕ can be efficiently computed, then we have such a reduction.

As we have seen, if there is a very large prime dividing the conductor of $\mathrm{End}(E)$ but not the conductor of $\mathrm{End}(\widetilde{E})$ (or vice versa) then it is not efficient to compute an isogeny from E to $\widetilde{E}$. In this case one cannot make any inference about the relative difficulty of the DLP in the two groups. No example is known of elliptic curves E and $\widetilde{E}$ of this form (i.e., with a large conductor gap) but for which the DLP on one is known to be significantly easier than the DLP on another. The nearest we have to an example of this phenomenon is with elliptic curves E with $\#\mathrm{Aut}(E) > 2$ (and so one can accelerate the Pollard rho method using equivalence classes as in Section 14.4) but with an isogeny from E to $\widetilde{E}$ with $\#\mathrm{Aut}(\widetilde{E}) = 2$.

On the other hand, if the conductors of $\mathrm{End}(E)$ and $\mathrm{End}(\widetilde{E})$ have the same very large prime factors (or no large prime factors) then we can (conditional on a generalised Riemann hypothesis) compute an isogeny between them in $\tilde{O}(q^{1/4})$ bit operations. This is not a polynomial-time reduction. But, since the current best algorithms for the DLP on elliptic curves run in $\tilde{O}(q^{1/2})$ bit operations, it shows that from a practical point of view the two DLPs are equivalent.

Jao, Miller and Venkatesan [310] have a different, and perhaps more useful, interpretation of the isogeny algorithms in terms of random self-reducibility of the DLP in an isogeny class of elliptic curves. The idea is that if E is an elliptic curve over $\mathbb{F}_q$ then by taking a relatively short random walk in the isogeny graph one arrives at a "random" (again ignoring the issue of large primes dividing the conductor) elliptic curve $\widetilde{E}$ over $\mathbb{F}_q$ such that $\#\widetilde{E}(\mathbb{F}_q) = \#E(\mathbb{F}_q)$. Hence, one easily turns a specific instance of the DLP (i.e., for a specific elliptic curve) into a random instance. It follows that if there were a "large" set of "weak" instances of the DLP in the isogeny class of E then, after enough trials,

one should be able to reduce the DLP from E to one of the elliptic curves in the weak class. One concludes that either the DLP is easy for "most" curves in an isogeny class, or is hard for "most" curves in an isogeny class.