EXPLOITING HIGHER DIMENSIONS IN ISOGENY-BASED CRYPTOGRAPHY: LECTURE NOTES FOR THE 2026 ARIZONA WINTER SCHOOL

SARAH ARPIN AND CHLOE MARTINDALE

The main characters of this lecture series are elliptic curves and abelian varieties over finite fields. Elliptic curves fit in a "sweet spot" in terms of balancing arithmetic complexity and computational efficiency. The points of an elliptic curve form a group, whereas the points on curves of genus 0 or > 1 do not immediately form a group. We can efficiently work with elliptic curves over finite fields, developing fast arithmetic using efficient isomorphism class representatives.

Elliptic curves first came to the attention of cryptographers in the mid-1980s [47, 39] as a drop-in group replacement for $\mathbb{F}_p^\times$ in the discrete log problem. The index calculus algorithm, which provides a subexponential algorithm solving the discrete log problem in $\mathbb{F}_p^\times$, simply does not apply to the group of points of an elliptic curve. Thus elliptic curves over finite fields became standard objects in cryptography. Elliptic curves over finite fields come in two flavors: supersingular and ordinary. However, the elliptic curve discrete log problem is not equally difficult for all elliptic curves. In fact, supersingular elliptic curves are particularly susceptible to the MOV attack: If $E/\mathbb{F}_p$ is supersingular, then $\#E(\mathbb{F}_p) = p + 1$ and if $P \in E(\mathbb{F}_p)$ is a point of exact order N then $N|(p + 1)$ which implies that the embedding degree of N in $\mathbb{F}_p$ is only 2. This small embedding degree allows us to translate the discrete log problem on $E(\mathbb{F}_p)$ to a related discrete log problem in the group $\mathbb{F}_{p^2}^\times$ which can be attacked via index calculus. While supersingular elliptic curves are not suitable for DLP applications, they are perfectly suitable for cryptographic protocols based on other hard problems.

In particular, supersingular elliptic curves are at the core of isogeny-based cryptography. First used in a hash function [18], isogeny graphs of supersingular elliptic curves support cryptographic protocols with advanced functionality, many of which we will survey in this lecture series. The most modern of these protocols use isogenies of higher-dimension abelian varieties, where by "higher-dimension" we explicitly mean dimensions 2, 4, and 8. The key reason higher-dimensional abelian varieties appear in modern isogeny-based cryptography is that they allow us to represent and verify isogenies that are infeasible to compute directly in dimension one.

The sections of these notes correspond to the lectures we will give at the Arizona Winter School 2026.

- We introduce the mathematics of supersingular elliptic curves in Section 1, and give a first look at an isogeny-based digital signature.
- We will discuss the computational tools for isogeny-based cryptography, in particular efficient representations of isogenies using abelian varieties in Section 2.
- With these computational tools in hand, we describe higher-dimension variants of SQIsign in Section 3.
- In Section 4 we look at other new isogeny-based cryptographic protocols which use the arithmetic geometry of abelian varieties constructively.

Section 5 provides the project descriptions for the students of AWS who are selected to work in our project group. We kindly request that researchers outside of the group of graduate students selected for this project group please not read these project proposals with too much interest.

Acknowledgements. The authors extend their deepest thanks to the organizers of the 2026 Arizona Winter School for the opportunity to develop this course: Thank you Anna Medvedovsky, Anthony Várilly-Alvarado, and Isabel Vogt (main program) with Renee Bell, Daniel Erman, Brandon Levin, Padma Srinivasan, and Hang Xue.

Version. These notes are Version 1.0, compiled on January 22, 2026. Comments welcome.

1. Introduction to the mathematics of isogeny-based cryptography

An **elliptic curve** E defined over a field K is a smooth, projective, algebraic curve of genus one with a marked point ∞_E. If the characteristic of K is not 2 or 3, then E can be described as the set of solutions to an affine short Weierstrass equation:

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

with $a, b \in K$ and $4a^3 + 27b^2 \neq 0$, together with a unique point ∞_E at infinity (denoted ∞ when E is understood). The points on an elliptic curve satisfy a group law and that group law is algebraic, which gives E the structure of an **abelian variety**. The point ∞ is the identity element of the group of points.[1]

Since the points of E satisfy a group law, we have $\mathbb{Z} \subseteq \mathrm{End}_K(E)$, where $\mathrm{End}_K(E)$ denotes the **endomorphism ring** of E[2]. When the subscript K is omitted, $\mathrm{End}(E)$ denotes the set of all endomorphisms of E defined over an algebraic closure $\overline{K}$ of K: $\mathrm{End}(E) = \mathrm{End}_{\overline{K}}(E)$, which is sometimes called the **geometric endomorphism ring** of E. The structure of $\mathrm{End}_K(E)$ reveals a lot about the curve E, and is of particular computational interest when working over finite fields. This leads us to our first computational hard problem:

Problem 1 (Endomorphism Ring Problem)**.** *Given a supersingular elliptic curve E over a field of characteristic p, compute a basis for* $\mathrm{End}(E)$.

When the field characteristic p is large and no extra information is given about the curve E, this problem is believed to be computationally infeasible. This computationally hard problem is equivalent [46] to the **isogeny problem**:

Problem 2 (Isogeny Problem)**.** *Given two supersingular elliptic curves E and E', find an isogeny $\varphi : E \to E'$.*

This problem is also computationally infeasible for random supersingular elliptic curves. It should be noted that there are certain curves for which these problems are in fact easy: for example, the existence of extra automorphisms makes it easier to compute the endomorphism ring - see Example 19.

Many properties of elliptic curves are defined and studied via their coordinate rings and function fields. The **affine coordinate ring** of an elliptic curve $E : y^2 = x^3 + ax + b$ over K is the ring $K[x, y]/(y^2 - x^3 - ax - b)$, often denoted $K[E]$. The **function field** of E, denoted $K(E)$, is the field of fractions (quotient field) of the coordinate ring of E.

1.1. Isogenies.

An **isogeny** $\varphi : E_1 \to E_2$ of elliptic curves E_1, E_2 is a rational map which is defined on all points of E_1 which satisfies $\varphi(\infty_{E_1}) = \infty_{E_2}$. In particular, isogenies are morphisms of elliptic curves which respect the group law:

$$\varphi(P + Q) = \varphi(P) + \varphi(Q).$$

Most of the time when we talk about an isogeny $\varphi : E_1 \to E_2$ we will mean a nonzero isogeny. Please forgive any omission of the qualifier "nonzero."

Every nonzero isogeny gives an injection of function fields:

$$\varphi^* : \overline{K}(E_2) \to \overline{K}(E_1)$$

[1]The point at infinity is less mysterious when working with projective equations: $E : Y^2Z = X^3 + AXZ^2 + BZ^3$ contains the point $\infty = [0 : 1 : 0]$.

[2]The multiplication-by-m maps $[m] : E \to E$ are nonconstant for every $m \in \mathbb{Z} \setminus \{0\}$ [70, Prop. III.4.2.a]

$$f \mapsto f \circ \varphi.$$

The **degree** of φ is the degree of the field extension $\overline{K}(E_1)/\varphi^*(\overline{K}(E_2))$, and for the zero isogeny $[0]$ we set $\deg([0]) = 0$. A nonzero isogeny φ is (separable, inseparable, purely inseparable) if the corresponding field extension is (separable, inseparable, purely inseparable). Every isogeny $\varphi : E_1 \to E_2$ is associated with a unique **dual isogeny** $\widehat{\varphi} : E_2 \to E_1$ satisfying

$$\varphi \circ \widehat{\varphi} = [\deg \varphi]_{E_2} \text{ and } \widehat{\varphi} \circ \varphi = [\deg \varphi]_{E_1},$$

where $[m]_{E_i}$ denotes the multiplication-by-m endomorphism of E_i.

Proposition 3. *The degree map is multiplicative:* $\deg(\varphi \circ \psi) = (\deg \varphi)(\deg \psi)$.

Proof. Follows from the corresponding property for degrees of field extensions. □

Lemma 4. *Let $\varphi : E_1 \to E_2$ be a separable isogeny. Then, $\# \ker \varphi = \deg \varphi$.*

Proof. This comes from the fact that separable isogenies are unramified morphisms. See [70, Thm. III.4.10] for details. □

Lemma 5. *Take $[m] : E \to E$ to be the multiplication-by-m map for some nonzero $m \in \mathbb{Z}$. Then,* $\deg([m]) = m^2$.

Proof. We provide a start: By definition, $\deg([m]) = [K(E) : [m]^*(K(E))]$. Suppose E is given by the equation $y^2 = x^3 + ax + b$. Let $x : E \to \mathbb{P}^1$ denote the x-coordinate map. Let $x_m = x \circ [m] : E \to \mathbb{P}^1$. Notice the relation:

$$[K(E) : (x_m)^*(K(E))] = [K(E) : [m]^*(K(\mathbb{P}^1))] \cdot [[m]^*(K(\mathbb{P}^1)) : (x_m)^*(K(E))]$$

and use this to prove the desired result. □

The m-torsion points of E, denoted by $E[m](K)$, is the subgroup of points $P \in E(K)$ such that $[m](P) = \infty$. We denote by $E[m]$ the group of m-torsion points of E defined over $\overline{K}$.

Theorem 6. *Let E be an elliptic curve defined over a field K. Let m be an integer coprime to the characteristic of the field K. The m-torsion points of E are isomorphic to the group $\mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}$:*

$$E[m] \cong \mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}.$$

Proof. This result is actually true for abelian varieties of arbitrary dimension. We will revisit this result, see Theorem 28. □

Theorem 7. *Let $G \subseteq E(K)$ be a finite subgroup of points of an elliptic curve E with order coprime to the characteristic of K. Then, G defines an isogeny $\varphi : E \to E/G$ with $\ker \varphi = G$ which is unique up to post-composition with an automorphism.*

Proof. Given a Weierstrass equation for E and explicit coordinates for the points of G, Vélu [74] provided explicit formulae for the rational map φ. There have been a number of updates to this original work ([8, 55, 20]); however, in the following lectures we will focus instead on a different set of techniques for effectively computing isogenies. □

Isogenies give a first glimpse into understanding the theory underlying elliptic curves via their function fields. This is part of a broader categorical equivalence:

Theorem 8 (Rem. 2.5 [70], Cor. 6.12 [33])**.** *The following categories are equivalent:*

(1) smooth curves defined over K, and surjective morphisms defined over K;
(2) finitely generated field extensions L/K of transcendence degree one with $L \cap \overline{K} = K$, field injections fixing K.

- $\mathrm{Hom}_K(E_1, E_2)$ denote the collection of isogenies from E_1 to E_2 which are defined over K. $\mathrm{Hom}_K(E_1, E_2)$ is a group via the group laws on E_1 and E_2 and the fact that isogenies respect group operations.
- $\mathrm{End}_K(E)$ denotes the collection of isogenies from E to E which are defined over K. $\mathrm{End}_K(E)$ is a ring, where the addition map is inherited from the fact $\mathrm{End}_K(E) = \mathrm{Hom}_K(E, E)$ and the multiplication map is given by composition.
- $\mathrm{Aut}_K(E)$ denotes the collection of degree-1 isogenies from E to E which are defined over K. $\mathrm{Aut}_K(E)$ is a group under the composition operation.

Proposition 9 (Prop. III.4.2.b [70])**.** *The group* $\mathrm{Hom}(E_1, E_2)$ *is a torsion-free* $\mathbb{Z}$*-module, and* $\mathrm{End}(E)$ *is a ring of characteristic* 0*.*

Proof. Let $\varphi \in \mathrm{Hom}(E_1, E_2)$ be nonzero and suppose $[m] \circ \varphi = [0]$ for some integer m. Then $\deg([m])\deg(\varphi) = 0$ by Proposition 3. Since $\mathbb{Z}$ is torsion-free, this implies $m = 0$. Since $\mathrm{End}(E) = \mathrm{Hom}(E, E)$, this also shows that the ring $\mathrm{End}(E)$ has characteristic 0. □

Proposition 10 (Prop. III.4.2.c [70])**.** *The ring* $\mathrm{End}(E)$ *has no nonzero zero divisors.*

Proof. Again use the property that the degree map is multiplicative (Proposition 3). □

We remark that every endomorphism $\varphi \in \mathrm{End}(E)$ has a minimal polynomial $m_\varphi(t) \in \mathbb{Z}[t]$ such that $m_\varphi(\varphi) = [0]$.

There are two ways to approach their theory and computation of isogenies of elliptic curves. The first method being direct computation by kernels and explicit maps (see Lemma 4 and Theorem 7). This method works generally for all separable isogenies of elliptic curves. However, there is a subset of isogenies which correspond to ideals of the endomorphism ring. In the case where the endomorphism ring of the elliptic curve is isomorphic to an imaginary quadratic order, one can immediately see these isogenies as realizing the action of the class group of $\mathrm{End}(E)$ on a particular subset of elliptic curves. The following result is a step in this direction, which we will build upon in Section 1.5.

Proposition 11. $\mathrm{Hom}_K(E_1, E_2)$ *is a left* $\mathrm{End}_K(E_2)$*-module and a right* $\mathrm{End}_K(E_1)$*-module, where the actions are given by post- and pre-composition, respectively.*

Proof. The only property to check is well-definedness, which is immediate. □

Theorem 12 ([70])**.** *Let E/K be an elliptic curve. Then End(E) is one of:*

- $\mathbb{Z}$*,*
- *an order in a quadratic number field, or*
- *a maximal order in a quaternion algebra.*

Proof. This result follows from showing that the ring $\mathrm{End}(E)$ of endomorphisms of an elliptic curve E is a characteristic 0 ring with no zero divisors with rank at most four as a $\mathbb{Z}$-module, and that $\mathrm{End}(E)$ has an anti-involution $\widehat{\cdot}$ which satisfies the following properties:

$$\widehat{\alpha + \beta} = \widehat{\alpha} + \widehat{\beta},\ \widehat{\alpha\beta} = \widehat{\beta}\widehat{\alpha},\ \widehat{\widehat{\alpha}} = \alpha,$$

for all $\alpha, \beta \in \mathrm{End}(E)$, and that $\widehat{a} = a$ for any endomorphism $a \in \mathbb{Z} \subseteq \mathrm{End}(E)$. It is further required that $\alpha\widehat{\alpha} \geq 0$ and $\alpha\widehat{\alpha} = 0$ if and only if $\alpha = 0$. □

1.2. **Elliptic Curves in Characteristic $p > 0$.** The above results hold for elliptic curves over arbitrary fields. When we restrict to elliptic curves defined over a finite field $\mathbb{F}_q$ of characteristic p, additional structure arises that is inherited from the arithmetic of the base field. The Galois group $\mathrm{Gal}(\mathbb{F}_q/\mathbb{F}_p)$ is generated by the p-power Frobenius field automorphism:

$$\pi_p : \mathbb{F}_q \to \mathbb{F}_q$$
$$\pi_p(x) = x^p.$$

Let $E : y^2 = x^3 + ax + b$ be an elliptic curve defined over $\mathbb{F}_q$. The field automorphism π_p motivates the definition of the **p-power Frobenius isogeny**:

$$\pi_p : E \to (E^{(p)} : y^2 = x^3 + a^p x + b^p)$$
$$\pi_p(x, y) = (x^p, y^p).$$

The dual of the Frobenius isogeny is the **Verschiebung** isogeny.

Theorem 13 (Hasse-Weil bound)**.** *Let $E/\mathbb{F}_q$ be an elliptic curve over a finite field $\mathbb{F}_q$. The number of $\mathbb{F}_q$-rational points on E satisfies the following inequality:*

$$q + 1 - 2\sqrt{q} \leq \#E(\mathbb{F}_q) \leq q + 1 + 2\sqrt{q}.$$

Proof. See [70, Thm. V.1.1]. The basic idea is to use the q-power Frobenius morphism of the curve to define the set of $\mathbb{F}_q$-points:

$$E(\mathbb{F}_q) = \{P \in E : \pi_q(P) = P\} = \ker([1] - \pi_q).$$

We can show that the map $[1] - \pi_q$ is separable, so the size of its kernel is equal to its degree. Compute the degree directly by $([1] - \pi_q)([1] - \widehat{\pi}_q)$ and apply the Cauchy-Schwarz inequality. □

Elliptic curves over fields of characteristic p are either **supersingular** or **ordinary**. Both types have been used in cryptography, each with its own strengths and weaknesses. The use of supersingular elliptic curves can render the discrete logarithm problem vulnerable to the MOV attack. The use of ordinary elliptic curves in isogeny-based protocols is also insecure. The distinct properties associated with these two types of curves are important to study closely.

Theorem 14 (Thm V.3.1 [70])**.** *Let $E/\mathbb{F}_q$ be an elliptic curve defined over a finite field $\mathbb{F}_q$ of characteristic p. Then, the following are equivalent:*

*(1) E is **supersingular**;*
(2) $E[p^r] = \{\infty_E\}$ for one (and thus all) integers $r \geq 1$;
(3) the Verschiebung isogeny $\widehat{\pi}_{p^r}$ is purely inseparable for one (and thus all) $r \geq 1$;
(4) the multiplication-by-p map $[p]$ is purely inseparable and $j(E) \in \mathbb{F}_{p^2}$;
(5) the ring $\mathrm{End}(E)$ is a maximal order in a quaternion algebra.

If E is not supersingular, then the following equivalent conditions hold:

*(1) E is **ordinary**;*
(2) $E[p^r] \cong \mathbb{Z}/p^r\mathbb{Z}$ for all integers $r \geq 1$;
(3) $\mathrm{End}(E)$ is an order in an imaginary quadratic field.

Proof. See [70] for details. The main ideas come from the fact that $[p] = \widehat{\pi}_p \circ \pi_p$ and the injection $\mathrm{End}(E) \hookrightarrow \mathrm{End}(T_p(E))$, where $T_p(E)$ denotes the p-adic Tate module of E. □

Remark 15. The reader accustomed to working with elliptic curves as schemes will note that the p-torsion subgroups of elliptic curves over finite fields have more structure than what is observed from the affine/projective picture. If E is supersingular, the group scheme $E[p]$ is isomorphic to the unique non-split extension of α_p by itself (often denoted M_2). If E is ordinary, the group scheme $E[p]$ is isomorphic to $\mu_p \times \mathbb{Z}/p\mathbb{Z}$. Here, α_p is the kernel of the p-power Frobenius map on $\mathbb{G}_a$ and μ_p is the kernel of the p-power Frobenius map on $\mathbb{G}_m$. See [38, Thm. 2.9.3] for a long (but worthwhile) journey to truth.

1.3. **Supersingular elliptic curves.** Supersingular elliptic curves are defined by the properties listed in Theorem 14. For the purposes of this section fix $E/\mathbb{F}_q$ a supersingular elliptic curve over the field $\mathbb{F}_q$ with $\operatorname{char} \mathbb{F}_q = p$.

By Theorem 14, the j-invariant of a supersingular elliptic curve lies in $\mathbb{F}_{p^2}$, so in particular there are finitely many isomorphism classes of supersingular elliptic curves over $\overline{\mathbb{F}}_p$.

Theorem 16 (Thm. V.4.1 [70])**.** *Let $p \geq 5$. The number of supersingular elliptic curves (up to $\overline{\mathbb{F}}_p$-isomorphism) over $\overline{\mathbb{F}}_p$ is equal to*

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

Proof. Using the Legendre form of an elliptic curve $E_\lambda : y^2 = x(x-1)(x-\lambda)$ over $\overline{\mathbb{F}}_p$, one can define a polynomial $H_p(x)$ which vanishes at $x = \lambda \in \overline{\mathbb{F}}_p$ if and only if E_λ is supersingular. The result follows by counting roots of this polynomial $H_p(x)$. □

The property of being supersingular is an isogeny-invariant. To see this, note that an isogeny $\varphi : E \to E'$ induces an isomorphism of group schemes $E[p] \to E'[p]$. The structure of $E[p]$ determines whether or not E is supersingular, and this structure is unchanged by isogeny.

Before going into more detail on the geometric endomorphism ring of a supersingular elliptic curve as a quaternion order, we remark that the ring of $\mathbb{F}_p$-rational endomorphisms of a supersingular elliptic curve over $\mathbb{F}_p$ is actually commutative:

Theorem 17 (Prop. 2.5 [26])**.** *Let $p > 3$ be a prime and set $K = \mathbb{Q}(\sqrt{-p})$. Let $\mathcal{O}_K$ denote the ring of integers of K. Then, there is a one-to-one correspondence between the set of isomorphism classes of supersingular elliptic curves defined over $\mathbb{F}_p$ and isomorphism classes of elliptic curves over $\mathbb{C}$ whose endomorphism ring is isomorphic to either $\mathbb{Z}[\sqrt{-p}]$ or $\mathcal{O}_K$.*

Proof. Proof idea: By CM theory, the class number of an imaginary quadratic order $\mathcal{O}$ in a CM field K dictates how many isomorphism classes of elliptic curves over $\mathbb{C}$ have endomorphism ring isomorphic to $\mathcal{O}$. Taking the reduction of these curves modulo p give supersingular elliptic curves defined over $\mathbb{F}_p$. See [26, Prop. 2.5] and [69, p. II.4] for details. □

The quaternion algebra $B_{p,\infty} = \operatorname{End}(E) \otimes_{\mathbb{Z}} \mathbb{Q}$ is ramified at p and ∞, meaning that the completions $B_{p,\infty} \otimes_{\mathbb{Q}} \mathbb{Q}_p$ and $B_{p,\infty} \otimes_{\mathbb{Q}} \mathbb{R}$ are division algebras (meaning every nonzero element has a two-sided multiplicative inverse). The **discriminant** of $B_{p,\infty}$ is p (more generally, the discriminant is the product of the finite primes at which the quaternion algebra ramifies). The endomorphism ring $\operatorname{End}(E)$ is a maximal order in $B_{p,\infty}$. Orders in $B_{p,\infty}$ are rank-4

lattices over $\mathbb{Z}$ which are also subrings of $B_{p,\infty}$. A maximal order in $B_{p,\infty}$ has discriminant p. See [75] for a complete treatment of quaternion algebras.

Recall from the introduction that the **endomorphism ring problem** (Problem 1) asks to compute the (geometric) endomorphism ring of a given supersingular elliptic curve E. To "compute" $\mathrm{End}(E)$ usually means to give a basis of four endomorphisms $\alpha_1, \alpha_2, \alpha_3, \alpha_4$ of E which generate $\mathrm{End}(E)$ as a $\mathbb{Z}$-module. This can mean as rational maps, as kernels, or as some oracle which is able to evaluate the maps α_i at points of E. For random supersingular elliptic curve in characteristic p for p large enough, this computational problem is believed to be intractable. In fact, it is even difficult to find a single non-scalar endomorphism – this problem is known as the **one endomorphism problem**. The endomorphism ring problem and the one endomorphism problem were shown to be computationally equivalent [58].

Problem 18 (One Endomorphism Problem)**.** *Given a supersingular elliptic curve E over a field of characteristic p, compute one nonscalar endomorphism in* $\mathrm{End}(E)$*.*

Example 19. It is not *always* hard to compute the endomorphism ring of an elliptic curve. For a very special case, consider the elliptic curve $E : y^2 = x^3 + x$ defined over $\mathbb{F}_p$ with $p \equiv 3 \pmod 4$. This curve has an automorphism of order four:

$$[i] : E \to E$$

$$(x, y) \mapsto (-x, iy),$$

where $i \in \mathbb{F}_{p^2} \setminus \mathbb{F}_p$ such that $i^2 = -1$. Since this curve is defined over $\mathbb{F}_p$, we also have the p-power Frobenius isogeny $\pi_p \in \mathrm{End}(E)$. Already, the endomorphisms $[1], [i], \pi_p$, and $\pi_p \circ [i]$ generate a quaternion order of discriminant $4p$ – very close to maximal! To complete the computation, note

$$E[2] = \{(0, 0), (i, 0), (-i, 0), \infty\}.$$

As an exercise, find an endomorphism divisible by [2] by finding an endomorphism which vanishes on $E[2]$. This computation was simple – make sure you understand why this is the exception and not the rule. We will use this elliptic curve with known endomorphism ring again in the sections which follow.

While it is generically difficult to compute the endomorphism ring of a random supersingular elliptic curve, it is theoretically possible to list all of the possibilities as abstract orders in a quaternion algebra $B_{p,\infty}$. The following theorem of Pizer [62] gives an explicit basis for a maximal order in the quaternion algebra $B_{p,\infty}$, whose generators are also described.

Theorem 20 (Prop. 5.1 & 5.2 [62])**.** *Fix a prime p. The (unique up to isomorphism) quaternion algebra $B_{p,\infty}$ ramified precisely at p and ∞ is given by:*

$$B_{p,\infty} = \begin{cases} \left(\frac{-1,-1}{\mathbb{Q}}\right) & \text{if } p = 2 \\ \left(\frac{-1,-p}{\mathbb{Q}}\right) & \text{if } p \equiv 3 \pmod 4 \\ \left(\frac{-2,-p}{\mathbb{Q}}\right) & \text{if } p \equiv 5 \pmod 8 \\ \left(\frac{-p,-q}{\mathbb{Q}}\right) & \text{if } p \equiv 1 \pmod 8, \end{cases}$$

where q is a prime with $q \equiv 3 \pmod 4$ and p is not a quadratic residue modulo q, and $\left(\frac{a,b}{\mathbb{Q}}\right)$ denotes the quaternion algebra generated over $\mathbb{Q}$ by the elements $1, i, j, ij$ with $i^2 = a, j^2 =$

$b, ij = -ji$. In each of these quaternion algebras, we can write down an explicit $\mathbb{Z}$-basis for a maximal order in $B_{p,\infty}$:

$$\begin{cases} \frac{1}{2}(1+i+j+ij), i, j, ij & \text{if } p = 2 \\ 1, i, \frac{1+j}{2}, \frac{i+ij}{2} & \text{if } p \equiv 3 \pmod 4 \\ \frac{1}{2}(1+j+ij), \frac{1}{4}(i+2j+ij), j, ij & \text{if } p \equiv 5 \pmod 8 \\ \frac{1}{2}(1+j), \frac{1}{2}(i+ij), \frac{1}{q}(j+aij), k & \text{if } p \equiv 1 \pmod 8, \end{cases}$$

where $a \in \mathbb{Z}$ such that $q|(a^2p+1)$.

From a given maximal order $\mathcal{O} \subseteq B_{p,\infty}$, the other maximal orders are precisely the right orders of the various left ideal classes of $\mathcal{O}$. This can be explicitly computed using Step 3 of the main algorithm in [62]. Thus the difficulty of computing $\mathrm{End}(E)$ rests solely on what we will refer to as the *effective Deuring correspondence*.

Remark 21 ($\mathrm{End}_{\mathbb{F}_p}(E)$)**.** Let $E/\mathbb{F}_p$ be a supersingular elliptic curve. Then $\mathrm{End}(E)$ is an order in a quaternion algebra, but $\mathrm{End}_{\mathbb{F}_p}(E)$, the ring of endomorphisms which are defined over $\mathbb{F}_p$, is actually an order in an imaginary quadratic field. See [77, Ch. 6]. This 'nicer' endomorphism ring structure has been leveraged to solve the isogeny problem: see the work of [26] and many follow-up optimizations. Not every supersingular elliptic curve is defined over $\mathbb{F}_p$, of course, so we cannot always leverage this nice structure. However, the theory of orientations will give us some way of extending the commutative endomorphism algebra picture to more general classes of supersingular elliptic curves. We'll see a bit of this in the last lecture, Section 4.1.1.

1.4. **Deuring Correspondence.** By Theorem 14, the (geometric) endomorphism rings of supersingular elliptic curves are maximal orders in quaternion algebras. Given a prime p, we can enumerate the finite list of isomorphism classes of supersingular elliptic curves by the supersingular j-invariants in $\mathbb{F}_{p^2}$. We can also enumerate the maximal orders in the quaternion algebra $B_{p,\infty}$, which we identify with $B_{p,\infty} = \mathrm{End}(E) \otimes_{\mathbb{Z}} \mathbb{Q}$ for a given supersingular elliptic curve $E/\mathbb{F}_{p^2}$. It is natural to match the supersingular elliptic curves with the maximal orders by computing the image of their endomorphism rings embedded in $B_{p,\infty}$. This association is one-to-one for curves with $j \in \mathbb{F}_p$ and two-to-one for curves with $j \in \mathbb{F}_{p^2} \setminus \mathbb{F}_p$. Such an association is useful – we can study supersingular elliptic curves by studying maximal orders in $B_{p,\infty}$ – but an association is not a categorical equivalence.

In [28], Deuring proved a categorical equivalence (although not phrased in this language at the time) between supersingular elliptic curves and quaternion objects by fixing a basepoint in each category. For a modern phrasing of this result, we paraphrase [75, Thm. 42.3.2]:

Theorem 22. *Let $\mathcal{S}$ denote the category whose objects are supersingular elliptic curves over $\overline{\mathbb{F}}_p$ and whose morphisms are isogenies. Fix $E \in \mathcal{S}$. Let $\mathcal{O}$ denote the endomorphism ring of E and let $B_{p,\infty} = \mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Q}$ denote the quaternion algebra ramified at p and ∞. Let $\mathcal{B}$ denote the category whose objects are left $\mathcal{O}$-modules and whose morphisms are nonzero left $\mathcal{O}$-module homomorphisms. Define a contravariant functor $F : \mathcal{S} \to \mathcal{B}$ as follows:*

$$E' \in \mathcal{S} \mapsto \mathrm{Hom}(E', E)$$

$$(\varphi : E_1 \to E_2) \mapsto (\mathrm{Hom}(E_2, E) \to \mathrm{Hom}(E_1, E) \text{ via } \psi \mapsto \psi \circ \varphi).$$

Then, F is an equivalence of categories.

This equivalence of categories is a useful computational tool, when it is computable. As discussed, the endomorphism ring problem (Problem 1) is generally computationally infeasible. However, with knowledge of the endomorphism ring of a supersingular elliptic curve, translating isogenies to and from the quaternion category is both possible and computationally advantageous.

Definition 23. Let E be a supersingular elliptic curve over $\overline{\mathbb{F}}_p$, and let $\mathrm{End}(E)$ denote its endomorphism ring. Let I denote a proper integral left-ideal of $\mathrm{End}(E)$ with norm coprime to p. Define a subgroup

$$E[I] = \bigcap_{\alpha \in I} \ker \alpha$$

and an isogeny $\varphi_I : E \to E/E[I]$ with $\ker \varphi_I = E[I]$.

For further details, see [75, p. 42.2].

There are extensions of the Deuring correspondence to higher dimension abelian varieties. Deligne [27] provides the first such generalization, for ordinary abelian varieties. Work of Centeleghe and Stix [17, 16] extended this to simple abelian varieties with commutative endomorphism ring.

Theorem 24 (Thm. 1.1 [35])**.** *In particular, fix a supersingular elliptic curve $E/\mathbb{F}_q$ with p-power Frobenius isogeny π_p and suppose that either:*

- *$\mathbb{F}_q = \mathbb{F}_p$ and $R = \mathrm{End}_{\mathbb{F}_q}(E) \cong \mathbb{Z}[\pi_p]$, or*
- *$\mathbb{F}_q = \mathbb{F}_{p^2}$ and $R = \mathrm{End}_{\mathbb{F}_q}(E)$ is a maximal order in a quaternion algebra.*

There is an equivalence of categories between the category of finitely presented torsion-free left-R modules and the category of abelian varieties isogenous to a power of E. Under this correspondence, the rank of the R-module corresponds to the power of E.

This theorem unifies several categorical equivalences already in the literature, including that of Deuring. The authors [35] provide a comprehensive treatment, and the introduction of this work cites several other places where the result appeared in less generality. In Lecture 4 (Section 4), we will discuss an extension of this equivalence of categories which has been used to create isogeny-based cryptographic protocols which take advantage of both higher dimension abelian varieties and group actions.

The reader interested in further equivalences of categories for elliptic curves and abelian varieties may also be interested in [3, 7], and many more.

1.5. **Group actions and the CSIDH protocol.** The first isogeny-based cryptographic primitives to use a group action were based on ordinary elliptic curves, due to the existence of a class group action coming from the ideal class group of an imaginary quadratic order [19, 67]. In particular, given an ordinary elliptic curve $E/\mathbb{F}_q$, the endomorphism ring of E is an imaginary quadratic order and the ideal class group of $\mathrm{End}(E)$ acts on the collection of elliptic curves whose endomorphism ring is isomorphic to $\mathrm{End}(E)$ [77, Thm. 4.5]. The geometric endomorphism rings of supersingular elliptic curves are noncommutative quaternion orders, and do not have a class group of ideals[3]. However, if E is a supersingular elliptic curve defined over $\mathbb{F}_p$, then $\mathrm{End}_{\mathbb{F}_p}(E)$ is isomorphic to an imaginary quadratic order, and again one can use the class group of this order to act on the set of elliptic curves whose $\mathbb{F}_p$-endomorphism

[3]There is a class group of two-sided ideals for maximal orders in the quaternion algebra $B_{p,\infty}$, but it is always either trivial or isomorphic to $\mathbb{Z}/2\mathbb{Z}$.

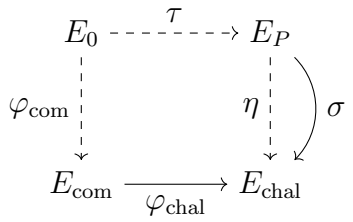

FIGURE 1. Isogenies involved in the SQIsign [25] protocol described in Section 1.6. Dashed isogenies are only known to the prover, and all other information is publicly known and communicated. The isogeny η makes the diagram commute, whereas the isogeny σ does not make the diagram commute.

ring is isomorphic to $\mathrm{End}_{\mathbb{F}_p}(E)$. The CSIDH protocol [15] relies on the difficulty of inverting this class group action. The details of this protocol were described in the PAWS lecture notes [41], but we recall the basics of the class group action here.

Definition 25. Let $E/\mathbb{F}_q$ be an elliptic curve defined over a field of characteristic p with a commutative endomorphism ring R[4]. Let I be a proper, integral ideal of R with norm coprime to p. Define a subgroup

$$E[I] = \bigcap_{\alpha \in I} \ker \alpha$$

and an isogeny $\varphi_I : E \to E/E[I]$ with $\ker \varphi_I = E[I]$. The action

$$I \star E = E/E[I]$$

defines a free action of the class group of R on the set of elliptic curves with endomorphism ring isomorphic to R. This action is transitive if p is split or ramified in $R \otimes_{\mathbb{Z}} \mathbb{Q}$, and has two orbits if p is ramified in $R \otimes_{\mathbb{Z}} \mathbb{Q}$.

This is essentially the same correspondence between ideals and isogenies that we have for supersingular elliptic curves and their geometric endomorphism rings in $B_{p,\infty}$ that we saw in Definition 23. However, there is no "class group" of left ideals of a quaternion algebra, so this ideal-isogeny correspondence is more powerful when working with commutative rings. The theory of orientations gives another way to use a class group action even when working inside the noncommutative geometric endomorphism ring of a supersingular elliptic curve. We visit this perspective in more detail in Section 4.1.1.

1.6. **A first isogeny-based digital signature.** A physical signature marks a particular document and certifies the signer's acknowledgement of the document. It is verifiable by checking the handwriting. A **digital signature** serves the same purpose but for digital documents: it certifies the signer's acknowledgement of a particular digital document, and it is verifiable by some publicly available information. Just as we hope our physical signatures are unforgeable, we want digital signatures to be unforgeable as well.

Definition 26 (Digital Signature Scheme)**.** A digital signature scheme consists of a public verification key `pk`, a private signing key `sk`, and two algorithms:

[4] R can be $\mathrm{End}_{\overline{\mathbb{F}}_p}(E)$ (if E is ordinary) or $\mathrm{End}_{\mathbb{F}_p}(E)$ (as in CSIDH) – the main point is we assume this ring is commutative.

- `Sign`: A signing algorithm which takes input a document D and a private signing key `sk` and produces a signature σ_D for document D.
- `Verify`: A verification algorithm which takes input a document D, signature σ_D for D, and public verification key `pk`, and returns `True` if σ_D is a valid signature associated with document D an `False` otherwise.

A digital signature should be *unforgeable* in that it should be difficult to produce a signature which verifies using `pk` without knowledge of `sk`.

For more information about the security of digital signatures, see [71, p. 11.7].

The first isogeny-based digital signature scheme was proposed in 2020 [25]. This protocol relies exclusively on dimension-1 abelian varieties. Although we will describe the current state of the art in Lecture 3 (see Section 3), presenting the original SQIsign protocol here allows us to apply the machinery developed in this lecture and provides a simpler model of an isogeny-based digital signature.

SQIsign is a digital signature scheme obtained via the Fiat–Shamir transform [32] applied to an identification protocol. For our purposes, it suffices to describe SQIsign at the level of the underlying sigma protocol. We will delay a description of the Fiat-Shamir transform to Section 3

A **sigma protocol** is a three-round interactive proof of knowledge in which a prover convinces a verifier that they know a secret, without revealing any information about the secret beyond its existence. In the case of SQIsign, this secret is an isogeny of supersingular elliptic curves. See Figure 1 for an overview of the isogenies involved. Let λ denote the security parameter of the protocol (meaning that the probability of an attacker breaking the protocol is at most $2^{-\lambda}$). The public parameters for this protocol with security parameter λ are:

- a prime $p \approx 2^{2\lambda}$;
- a supersingular elliptic curve $E_0/\overline{\mathbb{F}}_p$ with known endomorphism ring;
- $D_{\text{chal}} \approx 2^\lambda$ an odd integer;
- $D_{\text{res}} = 2^e$ for $e \approx \log(p)$.

The prover takes a random isogeny $\tau : E_0 \to E_P$, with $\deg \tau = D_\tau$ a large prime. Their public key is the curve E_P (which should appear to be random) and their secret is the isogeny τ. The prover needs to convince the verifier that they know τ, without revealing any information about τ. The three rounds proceed as follows:

- **Key generation:** The secret isogeny $\tau : E_0 \to E_P$ is sampled using knowledge of the endomorphism ring of E_0. Concretely, a random left-ideal I_τ of norm D_τ. In the ideal class $[I_\tau]$ contains an integral ideal J of smooth norm, which is found using the KLPT algorithm [40]. The ideal J_τ corresponds to an isogeny φ_{J_τ} whose codomain is the same as the codomain of the isogeny τ. Concretely, we have selected I_τ and J_τ from the same ideal class, s0 that $\varphi_{J_\tau} : E_0 \to E_P$ and $(\varphi_{I_\tau} = \tau) : E_0 \to E_P$. The public key E_P is computed using φ_{J_τ}.
- **Commitment:** The prover generates another secret isogeny $\varphi_{\text{com}} : E_0 \to E_{\text{com}}$, and publicly commits to the codomain curve E_{com}.
- **Challenge:** The verifier computes an isogeny $\varphi_{\text{chal}} : E_{\text{com}} \to E_{\text{chal}}$ with cyclic kernel and $\deg \varphi_{\text{chal}} = D_{\text{chal}}$ and communicates φ_{chal} to the prover.
- **Response:** Define $\eta = (\varphi_{\text{chal}} \circ \varphi_{\text{com}} \circ \widehat{\tau}) : E_P \to E_{\text{chal}}$. Using η, the prover computes a new isogeny $\varphi_{\text{res}} : E_P \to E_{\text{chal}}$ with $\deg \varphi_{\text{res}} = D_{\text{res}}$ to send as a response.

The verifier checks that the response isogeny φ_{res} is indeed an isogeny from $E_P \to E_{\text{chal}}$. At this stage, it is not clear how the prover computes the response isogeny, so we now explain this step in more detail. Note that if the prover had simply revealed the isogeny $\eta = (\varphi_{\text{chal}} \circ \varphi_{\text{com}} \circ \widehat{\tau}) : E_P \to E_{\text{chal}}$, the verifier would learn the secret isogeny τ[5]. At a high-level, the prover uses knowledge of $\text{End}(E_P)$ and η to translate η to a left ideal class $[I_\eta]$ of $\text{End}(E_P)$. To avoid revealing η (and thus revealing τ!), the prover chooses another ideal $J \in [I_\eta]$ such that the reduced norm of J is D_{res}. Then, the isogeny σ is revealed with $\ker \sigma = E[J]$. This process is completed in the `SigningKLPT` algorithm of [25]. This algorithm generalizes the `KLPT` algorithm [40].

For a sigma protocol to result in a secure digital signature, it must satisfy certain properties. The original SQIsign proposal did not contain a complete proof of security, as details were omitted and heuristics not rigorously investigated. In 2025, the authors of [1] provided a complete proof of security for the version of SQIsign submitted to round 2 of the NIST standardization process. This variant of SQIsign does not contain optimizations using higher-dimension abelian varieties, so the precise proof of security provided in [1] does not apply to SQIsign2D-West, the variant of SQIsign which we will discuss in-depth in Lecture 3 Section 3. However, these security conditions remain relevant to SQIsign2d-West, so we will provide a brief discussion here and then revisit these conditions in 3.3.

1.6.1. *Completeness.* An identification protocol is **complete** if an honest verifier is convinced by an honest prover. The completeness of SQIsign comes from the fact that the isogeny $\sigma : E_P \to E_{\text{chal}}$ produced by the prover is indeed an isogeny of the correct degree. This follows from the proof of the `SigningKLPT` algorithm in [25, Prop. 9].

1.6.2. *Soundness.* An identification protocol is **sound** if a dishonest prover (who does not know τ) is unable to convince the verifier that they know τ, except with negligible probability. The justification for soundness provided in the original SQIsign paper [25] argued that an adversary who is able to break the soundness of the protocol would be able to efficiently solve the endomorphism ring problem (Problem 1). In [1], the authors note that the underlying problem is actually the one endomorphism problem (Problem 18). The two hard problems are equivalent [58], but the equivalence requires changing the starting curve. This leads to a complete proof of soundness provided in [1], where the assumptions are adjusted slightly in their "hint-assisted" model.

1.6.3. *Zero Knowledge.* An identification protocol is **zero knowledge** if the verifier learns no information about the prover's secret through the interaction. In SQIsign, this means that the verifier doesn't learn any information about the secret isogeny $\tau : E_0 \to E_P$ from the response isogeny $\sigma : E_P \to E_{\text{chal}}$. The original KLPT algorithm would have revealed information about the endomorphism ring of E_P, thus breaking this critical security property. The algorithm `SigningKLPT` introduced in [25] avoids this information leakage, but there are other serious issues in proving zero knowledge.

To prove a protocol satisfies the zero knowledge property, one must show that there exists a simulator which produces accepting transcripts at the same rate as actual executions of the protocol, without access to the secret information. The transcript (tuple of commitment, challenge, and response data) can be set up by the simulator in any order, so usually this

[5]Make sure you can see why this is! Note that learning $\widehat{\tau}$ is equivalent to learning τ.

involves selecting a random challenge, selecting a response, and then computing the commitment that forces the transcript to pass the verification step. In SQIsign, it is not known how to sample efficiently from the algorithmically defined set of isogenies of degree-2^x from E_P without knowledge of the secret key. If the endomorphism ring of E_P is known, it is possible to efficiently sample uniformly from all isogenies of degree-2^x, but this is notably not the same set actually sampled from in the protocol. In [1], the authors overcome this security barrier by building a "hint-assisted" simulator.

2. Efficient representation and evaluation of isogenies

2.1. **Abelian varieties over finite fields.** In this section, we formally introduce abelian varieties, their isogenies, and their polarizations, spiced with plenty of intuition and motivation. A good textbook reference for the theory of abelian varieties over finite fields is the book in preparation by Edixhoven, van der Geer, and Moonen, available on van der Geer's website [31]. At the end of this section, we cover the classification of principal polarizations of supersingular abelian surfaces of Ibukiyama, Katsura, and Oort [34], giving a first look at how to work with simple explicit representations in this area. We start this with intuition and motivation for the following definition:

Definition 27. An **abelian variety** (of which an elliptic curve is an example) is a smooth projective algebraic variety that is also an algebraic group.

'Projective' means that it can be embedded into projective space, and 'algebraic' means that the image of this embedding can be expressed as algebraic equations. When working with abelian varieties in a computational setting, we fix an embedding into projective space by choosing a 'polarization'. A **polarized abelian variety** refers to the data (A, φ), where φ is the polarization, which allows A to be represented by equations as φ fixes the choice of embedding into projective space. Before formally defining polarizations, let us motivate the last part of the definition of 'abelian variety', namely that it is 'also an algebraic group'. This just means that there is a well-defined group law on the set of rational points which can be represented in terms of rational functions, which should be familiar from the elliptic curve setting. The following theorem gives the structure of this group of rational points:

Theorem 28. *Let A be an abelian variety of dimension g over a field k. Let m be an integer coprime to the characteristic of the field k. The m-torsion points of A are isomorphic to $(\mathbb{Z}/m\mathbb{Z})^{2g}$:*

$$A[m](\overline{k}) \cong (\mathbb{Z}/m\mathbb{Z})^{2g}.$$

Proof. Over the complex numbers, an abelian variety X is isomorphic to $\mathbb{C}^{2g}/\Lambda$, where Λ is a full rank lattice in $\mathbb{C}^{2g}$. The m-torsion points are the points x of $\mathbb{C}^{2g}$ such that $mx \in \Lambda$, so $X[n] \cong \left(\frac{1}{m}\Lambda\right)/\Lambda \cong \left(\left(\frac{1}{m}\mathbb{Z}\right)/\mathbb{Z}\right)^{2g} \cong (\mathbb{Z}/m\mathbb{Z})^{2g}$. For an algebraic explanation over an arbitrary algebraically closed field, see the proof in [51, Ch. II.6], which uses properties of finite group schemes. □

Just like the elliptic curve case, for an abelian variety defined over a field of characteristic p, the p-torsion behaves differently to the m-torsion of Theorem 28. However, unlike the elliptic curve case, there are more than two options. To introduce these, we must first generalize the notions of isogeny and isomorphism covered already in the elliptic curve case: An **isogeny** is a morphism of abelian varieties that is finite as a morphism of varieties and is surjective (after covering polarizations more precisely, we will introduce the more restrictive but more explicit notion of a 'polarization-respecting' isogeny). The **degree** of an isogeny is its degree as a morphism of varieties. An **isomorphism** is an isogeny of degree one.

Given an abelian variety A defined over a field k of characteristic p, we say that A is **supersingular** if there exists an isogeny over $\overline{k}$ to a product of supersingular elliptic curves $E_1 \times \cdots \times E_g$. We say that A is **superspecial** if there exists an isomorphism over $\overline{k}$ to a product of supersingular elliptic curves $E_1 \times \cdots \times E_g$. In fact, all superspecial abelian varieties are isomorphic (note that we don't yet require these isomorphisms to be polarization-preserving,

else this wouldn't be true), which turns out to be a crucial factor in having nice simple explicit representations of the polarizations to work with (see Section 2.1.1). Superspecial abelian varieties of a given dimension are currently considered to be the most appropriate generalization of supersingular elliptic curves for the purposes of isogeny-based cryptography. This is for several reasons, first and foremost because the isogeny graph is the most similar to the elliptic curve case [14], but also crucially because we don't have yet the tools to even write down supersingular abelian surfaces that are not superspecial, let alone compute isogenies between them!

We now formally define polarizations, before turning to examples and tools to work explicitly with polarizations in the superspecial case. We will see that a polarization is an isogeny between an abelian variety A and its dual $\widehat{A}$. However, the definition of the dual of an abelian variety over an arbitrary field k is a somewhat complex side-note to this lecture series (see for example [50, Chapter 6]), especially as a polarization is only explicitly identified as such by lifting to characteristic zero, as we will see below. Hence, we define the notion of dual for characteristic zero only and refer to [50] for further details in positive characteristic. The **Picard group** of A over an aribitrary field k, written as $\mathrm{Pic}(A)$, is the group of isomorphism classes of line bundles on A.[6]

Proposition 29. *For an abelian variety A over an arbitrary field k and a line bundle $\mathcal{L}$ on A, the map defined by*

$$\begin{array}{rccc} \varphi_{\mathcal{L}}: & A(k) & \to & \mathrm{Pic}(A) \\ & x & \mapsto & [T_x^*\mathcal{L} \otimes \mathcal{L}^{-1}], \end{array}$$

where T_x denotes translation by x and $[-]$ denotes the equivalence class of $-$ in $\mathrm{Pic}(A)$, is a homomorphism.

Proof. See [31, Corollary 2.10]. □

If k is algebraically closed, we define $\mathrm{Pic}^0(A)$ to be the subgroup of $\mathrm{Pic}(A)$ consisting of classes of line bundles $\mathcal{L}$ such that $\varphi_{\mathcal{L}} = 0$; then $\mathrm{Pic}^0(A)$ carries a canonical structure of an abelian variety over k [51, Chapter III, Corollary 5], which we define to be the **dual abelian variety** $\widehat{A}$. If $\mathcal{L}$ is an ample line bundle then $\varphi_{\mathcal{L}}$ defines an isogeny $A \to \widehat{A}$ [31, Theorem 6.18].

Definition 30. For an abelian variety A over an arbitrary field k, we define a **polarization** to be an isogeny

$$\varphi : A \longrightarrow \widehat{A}$$

over k such that there exists an ample line bundle $\mathcal{L}$ of the base change $A \times \overline{k}$ for which $\varphi \times \overline{k} = \varphi_{\mathcal{L}}$. We call this polarization **principal** if φ is an isomorphism (or equivalently if $\mathcal{L}$ is very ample).

We refer to the pair (A, φ), where φ is a (principal) polarization of A, as a **(principally) polarized abelian variety**. See Section 2.2 for examples of two-dimensional principally polarized abelian varieties.

Remark 31. Choosing a very ample divisor on an abelian variety can be thought of as the same thing as choosing an embedding into projective space. This connection will be

[6] If you don't link bundles, you can replace every instance of 'line bundle' with 'divisor' in your head (and replace tensors with additions) without losing the meaning of the story here.

made explicit in the section on theta coordinates, which is the name given to the associated embedding; see Section 2.2.1.

When working with the group law on abelian varieties explicitly, we will always work via a (principal) polarization so that we can compute using the coordinates in projective space, as mentioned above. In isogeny-based cryptography, we of course need to be able to compute isogenies between abelian varieties, not just use the group law on a given abelian variety, which leads us to the necessity of 'polarization-respecting isogenies'.

Definition 32. Let (A, φ) and (A', φ') be principally polarized abelian varieties of dimension g. Let $f : A \to A'$ be an isogeny of degree N^g. We say that f is **polarization-respecting** (or **polarization-preserving** or **polarized**)[7] if the following diagram commutes:

$$\begin{array}{ccc} A & \xrightarrow{f} & A' \\ {\scriptstyle \varphi \circ [N]}\downarrow & & \downarrow{\scriptstyle \varphi'} \\ \widehat{A} & \xleftarrow[\widehat{f}]{} & \widehat{A}'. \end{array}$$

Such an isogeny is often referred to as a N**-isogeny** or an $(N, \ldots, N)$**-isogeny** (where the tuple has length g). This does not define all possible polarization-respecting isogenies; in some cases there exist such isogenies of degree $< N^g$, for example see [44, Definition 1.5.3], but we leave this for further reading as the focus of this document is those isogenies that are currently used in cryptographic primitives. See Section 2.3 for more on computing (N, N)-isogenies in dimension two.

2.1.1. *Principally polarized superspecial abelian varieties.* We briefly focus on a special subcase of the above: principally polarized superspecial abelian varieties, as these are a focus of all isogeny-based primitives making use of dimension > 1. They are also in many ways beautifully simple to work with thanks to a classification due to Ibukiyama, Katsura, and Oort [34].

Let p be a prime. In any field k of characteristic p, there is only one superspecial abelian surface up to non-polarization-preserving $\overline{k}$-isomorphism. Hence, the set of all principally polarized abelian varieties (up to isomorphism) can be represented by the set of all non-isomorphic principal polarizations on any given principally polarized superspecial abelian variety, for example E^g, where E is the supersingular curve of known endomorphism ring, i.e. $\mathcal{O} := \mathrm{End}(E)$ is isomorphic to one of the orders in Theorem 20). This gives rise to the following:

Theorem 33. *Let p be a prime. Let $\mathcal{O}$ be the quaternion order defined by appropriate choice in Theorem 20 for the given prime p. Let A be a superspecial abelian variety of dimension g defined over $\mathbb{F}_{p^2}$. Then the principal polarizations of A are in bijection with the set*

$$\{M \in GL_g(\mathcal{O}) : M = \overline{M}^T,\ M \text{ positive definite}\}.$$

Proof. See [34, Proposition 2.8]. □

[7]The literature has many names for this concept, including omitting it all together and just saying 'isogeny', in which case it is up to the reader to notice that the author doesn't mean just an isogeny of abelian varieties, which is weaker.

This gets even more explicit when we restrict to dimension two (which we will, a lot of the time, in practice):

Corollary 34. *Let p be a prime and let $\mathcal{O}$ be the quaternion order defined as in Theorem 20. Let A be a superspecial abelian surface defined over $\mathbb{F}_{p^2}$. Then the principal polarizations of A are in bijection with the set*

$$\mathcal{S} = \left\{ \begin{pmatrix} s & r \\ \overline{r} & t \end{pmatrix} \in M_2(\mathcal{O}) : s, t, \in \mathbb{Z}_{>0}, st - r\overline{r} = 1 \right\}.$$

Proof. See [34, Corollary 2.9]: Note that by the full statement of the previous theorem in [34] the image being positive definite ensures that the divisor C is ample, and $d = 1$ ensures that the polarization is principal. □

With the representation of Corollary 34, we could choose to notate a principally polarized superspecial abelian surface by (E^2, M), for some $M \in \mathcal{S}$, rather than (A, φ). A benefit of this notation is that it gives us a very simple description of (polarization-preserving) isogenies and their duals:

Given two principally polarized abelian surfaces (E^2, M) and (E^2, M'), an (N, N)-isogeny between them is a matrix $\Gamma \in M_2(\mathcal{O})$ such that

$$NM = \overline{\Gamma}^T M' \Gamma. \tag{1}$$

The dual isogeny is given by $\overline{\Gamma}^T$.

To compare (1) with the previous definition of (N, N)-isogeny (Definition 32), note that $\overline{M}^T = M$, so the dual abelian surface of (E^2, M) is canonically isomorphic to (E^2, M), and similarly for M'; then we see that (1) is the commutative diagram of Definition 32 but with this new notation plugged in:

$$\begin{array}{ccc} (E^2, M) & \xrightarrow{\Gamma} & (E^2, M') \\ {\scriptstyle [N]}\downarrow & & \downarrow{\scriptstyle [1]} \\ (E^2, M) & \xleftarrow[\overline{\Gamma}^T]{} & (E^2, M'). \end{array}$$

Finally, we make a small comment on the different endomorphism rings that appear in the literature with the aid of the Ibukiyama-Katsura-Oort notation. There are three different sets that are of interest for different applications.

- **The full ring of (possibly unpolarized) endomorphisms**. For any superspecial abelian surface (E^2, M), this is isomorphic to $M_2(\mathcal{O})$. It is the same for every choice of M as all superspecial abelian surfaces are isomorphic as unpolarized abelian varieties.
- **The set of principally polarized endomorphisms**. What appears to be the most interesting thing to study, from a cryptographic and/or computational perspective, would be the set of (N, N)-endomorphisms for a given (E^2, M). That is, the set

  $$\{\Gamma \in M_2(\mathcal{O}) : \exists N \in \mathbb{Z} \text{ s.t. } NM = \overline{\Gamma}^T M \Gamma\}.$$

  However, this set isn't even a ring! Exercise: prove this.[8]

[8]To the best knowledge of the authors, no statement of this kind appears in the literature so far, although it is undoubtedly known to experts. We (re)discovered this during discussions with Luciano Maino and Lorenz Panny.

- **The symmetric endomorphism ring**. This is just another name for the set of principal polarizations described in Theorem 33.

2.2. **Explicit computations for genus 2 curves.** We now turn our attention to dimension two, that is, to principally polarized (superspecial) abelian surfaces. The following corollary to Torelli's theorem classifies principally polarized abelian surfaces:

Theorem 35 (Dimension-2 Torelli)**.** *Every principally polarized abelian surface is either a Jacobian of a genus-2 curve of a product of elliptic curves.*

Proof. See [78, Satz 2], [48, Sec. 12&13]. □

We first give a brief introduction to (Jacobians) of genus 2 curves, mostly following Chapters 1 and 2 of Cassels and Flynn [12]. We will, following Cassels and Flynn, typically write a genus 2 curve in a canonical form

$$\mathcal{C} : y^2 = f(x),$$

where $f(x) \in k[x]$ is polynomial of degree 6 with no multiple factors.[9] Unlike with elliptic curves, we can't just projectify this affine model using the map $(x, y) \mapsto (X/Z, Y/Z^3)$, as this results in a singularity at $Z = 0$. However, we do have a complete nonsingular model of

$$\mathcal{C} : y^2 = f(x) = f_0 + f_1 x + \cdots + f_6 x^6,\ f_i \in k \tag{2}$$

in $\mathbb{P}^4(k)$ given by the following equations:

$$\begin{aligned} Y^2 &= f_0 X_0^2 + f_1 X_0 X_1 + f_2 X_1^2 + f_3 X_1 X_2 + f_4 X_2^2 + f_5 X_2 X_3 + f_6 X_3^2, \\ 0 &= X_0 X_2 - X_1^2, \\ 0 &= X_0 X_3 - X_1 X_2, \\ 0 &= X_1 X_3 - X_2^2. \end{aligned}$$

To recover (2), set $X_0 = 1$, set $x^j = X_j$ and $y = Y$. There are then two points not captured in the affine picture (think of the 'point at infinity' on an elliptic curve), namely

$$[X_0 : X_1 : X_2 : X_3 : Y] = [0 : 0 : 0 : 1 : \pm\sqrt{f_6}],$$

these are referred to as the **points at infinity** and are denoted $\infty^{\pm}$.

The **Jacobian** of a genus 2 (or g) curve is given by $\mathcal{J}(\mathcal{C}) = \mathrm{Pic}^0(\mathcal{C})$. This comes with a canonical principal polarization; the Jacobian is therefore a principally polarized abelian surface. In light of the previous section, how do we get our hands on this 'canonical principal polarization'? From a computational perspective we consider a canonical principal polarization to be a canonical embedding into projective space, which in this case is $\mathbb{P}^{15}$. This is given explicitly in [12, Chapter 2, Appendix II], but is gigantic; it is also implemented in various computer algebra packages such as SageMath [**sage**] and Magma [10]. A much more manageable approach is to use 'theta coordinates': see Section 2.2.1 for how to, given an abelian variety A/k of dimension g endowed with a principal polarization $\mathcal{L}$, define a canonical embedding $A \hookrightarrow \mathbb{P}_k^{2^g-1}$.

[9] It is in some cases possible to reduce the degree of f to 5 by means of a birational transformation [12, Chapter 1 Section 1], but this leads to many case distinctions in all that follows, so we prefer to stick to this canonical form.

However, the group of rational points on $\mathcal{J}(\mathcal{C})$ is blessedly easy to describe; see [12, Chapter 1] for proofs of the following: Let $\mathcal{C}$ be a genus 2 curve defined over a field k. A **k-rational point** on $\mathcal{J}(\mathcal{C})$ is given by an unordered pair

$$\{P_1, P_2\}$$

where either

- $P_1 = \infty^+$ and $P_2 = \infty^-$ (this will later play the role of the group identity), or
- $P_1, P_2 \in \mathcal{C}(k)$, or
- there exists a quadratic extension $K \subseteq k$ with a canonical involution $\bar{\cdot}$ such that $(x, y) = P_1$, $(\overline{x}, \overline{y}) = P_2 \in \mathcal{C}(K)$.

Additionally, we identify any points of the form $\{(x, y), (x, -y)\}$, together with $\{\infty^+, \infty^-\}$.

Example 36. Let $\mathcal{C}/\mathbb{F}_3 : y^2 = x^6 + 1$ and define $\mathbb{F}_3(i) = \mathbb{F}_3[x]/(x^2 + 1)$. Then

$$\mathcal{C}(\mathbb{F}_3) = \{(0, \pm 1), \infty^{\pm}\},$$

$$\mathcal{C}(\mathbb{F}_3(i)) = \{(0,1), (0,-1), (1,i), (1,-i), (-1,i), (-1,-i), (i,0), (-i,0), \infty^{\pm}\},$$

and

$$\begin{aligned}\mathcal{J}(\mathcal{C})(\mathbb{F}_3) = \{&\{\infty^+, \infty^-\} = \{(0,1),(0,-1)\} = \{(1,i),(1,-i)\} = \{(-1,i),(-1,-i)\},\\ &\{\infty^+, \infty^+\},\\ &\{\infty^-, \infty^-\},\\ &\{(0,1), \infty^+\},\\ &\{(0,1), \infty^-\},\\ &\{(0,-1), \infty^+\},\\ &\{(0,-1), \infty^-\},\\ &\{(0,1),(0,1)\},\\ &\{(0,-1),(0,-1)\},\\ &\{(i,0),(-i,0)\}\}.\end{aligned}$$

As a sanity check, we can check that we have the right number of points against the generalization of the point-counting formulae for elliptic curves [61]:

$$\begin{aligned}\#\mathcal{C}(\mathbb{F}_p) &= 1 - t + p,\\ \#\mathcal{C}(\mathbb{F}_{p^2}) &= 1 - t^2 + 4p + 2s + p^2,\\ \#\mathcal{J}(\mathcal{C})(\mathbb{F}_p) &= 1 - t + 2p + s - tp + p^2.\end{aligned}$$

In our example we have $\#\mathcal{C}(\mathbb{F}_3) = 4$, so $t = 0$, and $\#\mathcal{C}(\mathbb{F}_9) = 10$, so $s = -6$, giving $\#\mathcal{J}(\mathcal{C})(\mathbb{F}_3) = 10$, which indeed matches our calculations.

In a similar fashion to elliptic curves, when $k = \mathbb{R}$ the group law of the Jacobian $\mathcal{J}$ of a genus 2 curve $\mathcal{C}/\mathbb{R}$ has a beautiful geometric description. For the sake of drawing a nice picture, we describe the simplest case that can easily be represented on the affine model.

- The identity of the group is the point $\{\infty^+, \infty^-\}$, together with all points identified with this.

- If $P = (x, y)$ is an affine point on $\mathcal{C}$, we define $-P = (x, -y)$. If $P = \infty^{\pm}$ is a point at infinity on $\mathcal{C}$, we define $-P = \infty^{\mp}$. Then, negation in $\mathcal{J}(\mathcal{C})$ is defined by
$$\ominus\{P, Q\} = \{-P, -Q\}.$$
- Given two points $\{P_1, P_2\}$ and $\{P_3, P_4\} \in \mathcal{J}(\mathcal{C})(\mathbb{R})$, define g to be the unique cubic passing through all four points (possibly with multiplicities). This then intersects $\mathcal{C}$ in exactly two more points (possibly with multiplicities), call these points P_5 and $P_6 \in \mathcal{C}(\mathbb{R})$; we define
$$\{P_1, P_2\} \oplus \{P_3, P_4\} = \ominus\{P_5, P_6\}.$$

To compute with all the points on the Jacobian in the same way, these computations need to take place in the projective model, but as with elliptic curves, we can work directly with the affine model and consider the calculations with $\infty^{\pm}$ on a case-by-case basis. We draw an example of such an affine picture below: TODO.

2.2.1. *Theta coordinates.* Let A be an abelian variety defined over a field k. Let $\mathcal{L}$ be a very ample line bundle on A. Then by e.g. [73, Exercise 16.2.A] there exists $m \in \mathbb{Z}$ such that there is an embedding
$$A \hookrightarrow \mathbb{P}_k^m.$$
In this section, culminating in Definition 41, we describe a canonical way of constructing such an embedding with $m = n^g - 1$, for your choice of $n \geq 2$. This construction is primarily due to Mumford [51] and generalized to characteristic p by Robert [66]; we will follow Vuille's thesis [76, Section 2.1] as this contains an excellent description of Mumford's work in modern notation, as well as taking inspiration from [24].

This embedding is induced by a 'theta structure', which is an isomorphism between a 'theta group' and a 'Heisenberg group', all of which we now define.

Recall that an ample line bundle $\mathcal{L}$ on an abelian variety A has an associated isogeny
$$\begin{array}{rccc} \varphi_{\mathcal{L}}: & A & \to & A^\vee \\ & x & \mapsto & [T_x^*\mathcal{L} \otimes \mathcal{L}^{-1}]. \end{array}$$
It is often simpler computationally to work with divisor notation rather than line bundle notation, so let us take D to be the divisor associated to[10] $\mathcal{L}$; then $\varphi_D := \varphi_{\mathcal{L}} : x \mapsto T_x^* D - D$.

Definition 37. Let A be an abelian variety defined over k and let D be a divisor of A such that φ_D defines a principal polarization on A. The **theta group of level** n on (A, φ_D) is given by
$$\mathcal{G}(nD) := \{(x, \psi_x) \in A[n] \times k(A)^\times : \operatorname{div}(\psi_x) = T_x^* nD - nD\}$$
under the group law
$$(x, \psi_x) * (y, \psi_y) = (x + y, T_x^* \psi_y \circ \psi_x).$$

Definition 38. Let $g \in \mathbb{Z}_{\geq 1}$ and let k be a field. The **Heisenberg group**[11] $\mathcal{H}(n)$ associated to g and k is given by
$$\mathcal{H}(n) = k^\times \times (\mathbb{Z}/n\mathbb{Z})^g \times \operatorname{Hom}((\mathbb{Z}/n\mathbb{Z})^g, k^\times),$$
under the group law
$$(a_1, x_1, \chi_1) * (a_2, x_2, \chi_2) = (a_1 a_2 \chi_2(x_1), x_1 + x_2, \chi_1 + \chi_2).$$

[10] See [73, Section 15.6] for the meaning of 'associated to'.

[11] Mumford introduced this group, rather than Heisenberg, but he named it after Heisenberg in his paper after observing some fun connections to the Heisenberg group that comes up in physics.

Definition 39. Let g, k, A, and D be as above. Let $\mathcal{H}(n)$ be the Heisenberg group associated to g and k and let $\mathcal{G}(nD)$ be the theta group of level n on (A, φ_D). The δ_{-1} **operator** is given on $\mathcal{H}(n)$ by

$$\begin{array}{rccc} \delta_{-1}: & \mathcal{H}(n) & \mapsto & \mathcal{H}(n) \\ & (a,x,\chi) & \mapsto & (a,-x,1/\chi) \end{array}$$

and on $\mathcal{G}(nD)$ by

$$\begin{array}{rccc} \delta_{-1}: & \mathcal{G}(nD) & \mapsto & \mathcal{G}(nD) \\ & (x,\psi_x) & \mapsto & (-x,[-1]^*\psi_x). \end{array}$$

Definition 40. Let $g \in \mathbb{Z}_{\geq 1}$ and let k be a field. Let A be an abelian variety defined over k and let D be a symmetric divisor ($[-1]^*D = D$) on A such that $\varphi_D = \varphi_{\mathcal{L}}$ is a principal polarization of A. Let $\mathcal{H}(n)$ be the Heisenberg group associated to g and k and let $\mathcal{G}(nD)$ be the theta group of level n of (A, φ_D).

A **symmetric theta structure** $\Theta^{\mathcal{L}}$ of type n on (A, φ_D) is an isomorphism

$$\Theta^{\mathcal{L}} : \mathcal{H}(n) \tilde{\to} \mathcal{G}(nD)$$

that commutes with δ_{-1} and which induces the identity on the natural embedding of $k^\times$ in both groups.

Now, we can use the theta structure to construct a basis for the global sections $\Gamma(A, nD)$, which will give us our projective embedding, as follows. There is an action of $\mathcal{H}(n)$ on the vector space $V(n)$ of k-valued functions on $(\mathbb{Z}/n\mathbb{Z})^g$ given by

$$\begin{array}{ccccccc} \mathcal{H}(n) & \times & V(n) & \to & V(n) \\ (a,x,\chi) & , & f(\bullet) & \mapsto & a\chi(\bullet)f(x+\bullet). \end{array}$$

This is an irreducible representation called the **Schrödinger representation** and is unique up to isomorphism [49, Proposition 6.2].

There is an action of $\mathcal{G}(nD)$ on space of global sections $\Gamma(A, nD)$ given by

$$\begin{array}{ccccccc} \mathcal{G}(nD) & \times & \Gamma(A,nD) & \to & \Gamma(A,nD) \\ (x,\psi_x) & , & s & \mapsto & T_{-x}^*(\psi_x \circ s). \end{array}$$

As nD is ample, as long as $\mathrm{char}(k) \neq 2$, this defines an irreducible representation [76, Proposition 2.13]. Therefore $\Gamma(A, nD)$ is an irreducible representation of the Heisenberg group $\mathcal{H}(n)$, via $\Theta^{\mathcal{L}}$, so is isomorphic to the Schrödinger representation $V(n)$ (at least when $\mathrm{char}(k) \neq 2$). Let

$$r : V(n) \to \Gamma(A, nD) \tag{3}$$

be the corresponding $\mathcal{H}(n)$-equivariant isomorphism. There is a canonical basis

$$\{\delta_i : (\mathbb{Z}/n\mathbb{Z})^g \to k\}_{i \in (\mathbb{Z}/n\mathbb{Z})^g}$$

of $V(n)$ defined by the Kronecker delta function

$$\delta_i(j) = \begin{cases} 1 & i = j \\ 0 & \text{otherwise.} \end{cases}$$

Together with r, this gives us a canonical basis of the global sections – we have our projective embedding:

Definition 41. Let $(A, \mathcal{L})$ be a principally polarized abelian variety of dimension g defined over a field k. Let D be the divisor corresponding to $\mathcal{L}$, let r be the isomorphism of (3), and let δ_i be the Kronecker delta function. Then the **theta coordinates** define an embedding

$$A \hookrightarrow \mathbb{P}_k^{n^g-1}$$

and are defined by

$$x \mapsto [\theta_i^{\Theta_{\mathcal{L}}}(x) := (r \circ \delta_i)(x)]_{i \in (\mathbb{Z}/n\mathbb{Z})^g}.$$

Example 42. TODO

We will leave our foray into theta functions here for now. For more proofs, generalizations, and general merriment, we refer the enthusiastic reader to [76, Chapter 2].

2.3. **Isogenies in dimension two.** As we are talking here about isogeny-based cryptography, we of course need to be able to compute isogenies between principally polarized abelian varieties. Recall the definition of a polarization-respecting isogeny from Definition 32; all the isogenies in this section will respect polarizations. For most applications, we have everything we need from dimensions one and/or two, so we focus here on dimension two. In fact, for many applications we can even control the situation to the extent that we only care about computing chains of $(2, 2)$-isogenies, in which case the classical **Richelot isogenies** suffice (almost, we need to be a bit careful with products of elliptic curves). These are nice because they are conceptually simple, using only the curve equations, so we focus on this case in these notes. However, using theta coordinates is the most efficient method to date, see [24] for the state-of-the-art on $(2, 2)$-isogenies with theta coordinates and [79] for the state-of-the-art on (ℓ, ℓ)-isogenies with theta coordinates. (We may add some more on computing with theta functions in the final version of these notes.)

2.3.1. *Richelot isogenies.* Richelot's original paper [63] from 1837 gave elegantly simple formulae to compute a $(2, 2)$-isogeny between the Jacobians of two genus 2 curves. A good modern reference for Richelot isogenies is Benjamin Smith's PhD thesis [72, Section 8]. When working with these explicit models (i.e. not with theta coordinates), there are four case distinctions for computation:

(1) $f : E_1 \times E_2 \to E_3 \times E_4$; in this case f is just a 2×2 matrix of elliptic curve isogenies that can be computed by Vélu's formulae (see Theorem 7),
(2) $f : E_1 \times E_2 \to \mathcal{J}(C)$; this is referred to as a **gluing** isogeny,
(3) $f : \mathcal{J}(C) \to E_1 \times E_2$; this is referred to as a **splitting** isogeny, and
(4) $f : \mathcal{J}(C) \to \mathcal{J}(C')$; this is the most generic case.

Consider first the most generic case, that is, case (4). Let

$$C/k : y^2 = f(x)$$

be a hyperelliptic curve of degree 6. Let $f(x) = f_0(x)f_1(x)f_2(x)$ be a factorization of $f(x)$ into quadratic factors $f_i(x)$. For each quadratic factor $f_i(x)$, the roots $\alpha_i, \overline{\alpha_i}$ of f_i define an order two point

$$x_i = \{(\alpha_i, 0), (\overline{\alpha_i}, 0\} \in \mathcal{J}(C)[2]$$

and by definition $x_0 + x_1 + x_2 = 0$, so

$$\langle x_0, x_1 \rangle \cong \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$$

is the kernel of a $(2,2)$-isogeny. Write $f_i(x) = f_{i0} + f_{i1}x + f_{i2}x^2$ and define

$$\Delta = \begin{vmatrix} f_{00} & f_{01} & f_{02} \\ f_{10} & f_{11} & f_{12} \\ f_{20} & f_{12} & f_{22} \end{vmatrix}.$$

Then, if $\Delta \neq 0$, the codomain of the $(2,2)$-isogeny with kernel $\langle x_0, x_1 \rangle$ is given by the Jacobian of the curve C', where

$$C'/k : y^2 = g_0(x)g_1(x)g_2(x)$$

with

$$g_i(x) = \Delta^{-1}(f_j'(x)f_k(x) - f_k'(x)f_j(x))$$

for each cyclic permutation of (i, j, k).

The case $\Delta = 0$ is the split case, i.e., case (3), which we do not cover here. We refer the reader to [56] for a detailed description of how to compute $(2,2)$-isogeny formulae, including the gluing and splitting, and with accompanying code.

2.4. **Connecting dimensions.** In Lecture 1 (Section 1), we introduced supersingular elliptic curves and isogeny based cryptography. In this lecture, we have thus far focused on abelian varieties over finite fields. The following theorem of Kani provides a bridge between these two topics. In particular, this theorem has two consequences which are used again and again in isogeny-based cryptography: It allows us to factor isogenies, see Section 2.4.1, and it allows for the efficient representation of isogenies, see Section 2.4.2.

Theorem 43 (Kani's Reducibility Criterion [37, Thm. 2.3])**.** *Let f, A, and B be pairwise coprime integers such that $B = f + A$ and let E_1, E_2, E_3, E_4 be elliptic curves such that there exist isogenies $\varphi_f, \varphi_A, \varphi, g_A$, and g_f (resp.) of degrees f, A, fA, A, and f (resp.) for which the diagram in Figure 2 commutes. Then, the isogeny*

$$\Phi = \begin{pmatrix} \varphi_f & -\widehat{\varphi}_A \\ g_A & \widehat{g}_f \end{pmatrix} : E_1 \times E_2 \to E_3 \times E_4$$

is a (B, B)-isogeny with respect to the product polarizations on $E_1 \times E_2$ and $E_3 \times E_4$ with kernel given

$$\ker \Phi = \{([A]P, \varphi(P)) : P \in E_1[B]\}.$$

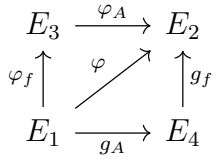

FIGURE 2. Kani's Reducibility Criterion

Proof. The proof proceeds by direct computation. See [43, Sec. 2.3] for details. □

2.4.1. *Consequence 1: Factoring isogenies.* Let f, A, and B be pairwise coprime integers such that $B = f + A$ and let $\varphi : E_1 \to E_2$ be an isogeny of elliptic curves of degree fA. Suppose that we can compute the action of φ on $E_1[B]$. That is, we know the blue values in the following commutative diagram:

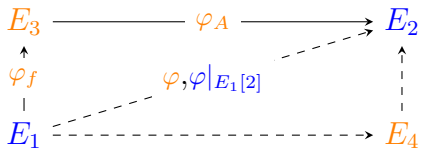

Then Kani's Reducibility criterion gives us the tool we need to find the orange values, since

$$\Phi = \begin{pmatrix} \varphi_f & -\widehat{\varphi_A} \\ * & * \end{pmatrix} : E_1 \times E_2 \to E_3 \times E_4$$

is a (B, B)-isogeny, and

$$\ker(\Phi) = \{([A]P, \varphi(P)) : P \in E_1[B]\}.$$

In particular, if we have managed to control the degrees so that $B = A + f$ is smooth, we can compute the elliptic curves E_3 and E_4 via which φ factors, as well as the corresponding isogenies.

2.4.2. *Consequence 2: Efficient representation of isogenies.* Let $\psi : E \to E'$ be an isogeny. Robert pointed out in [65] that Kani's Reducibility Criterion can also be used to store ψ efficiently – in principle φ can have huge degree and equations to match.

(1) If ψ has smooth degree, we factor ψ into small-degree parts and store each of these.
(2) If ψ has non-smooth degree but which can factor as $\deg(\psi) = fA$ such that $f+A = B$ is smooth and f, A, B are pairwise coprime, we fix $\langle P_B, Q_B \rangle = E_1[B]$. and store ψ as
$$\psi(P_B), \psi(Q_B).$$
Then for any $Q \in E_1$, we compute $\psi(Q)$ using Φ–note that, setting $\varphi = \psi$, we have exactly the blue information as in the above picture.
Exercise: Write an algorithm to compute $\psi(Q)$.
(3) If ψ has non-smooth degree which doesn't factor as $\deg(\psi) = fA$ such that $f + A = B$ with B smooth and f, A, B are pairwise coprime but we know $\mathrm{End}(E_1)$, we can use a trick introduced in QFESTA [52, Alg. 2]: Let $\deg(\psi) = f < B$, with $\gcd(f, B) = 1$ and B smooth (typically a a power of 2). Sample a degree-$(f(B - f))$ endomorphism $\theta \in \mathrm{End}(E_1)$ using [25, FullRepresentInteger alg.]. Now, in the picture for Consequence 1, set $E_2 = E_1$, $\varphi = \theta$, and $\varphi_f = \psi$; for $Q \in E_1$, we compute $\psi(Q)$ as the first coordinate of $\Phi((Q, \mathcal{O}))$.
We can store ψ efficiently as
$$(B - f, \theta(P_B), \theta(Q_B)),$$
where P_B, Q_B is basis for $E_1[B]$. If $\gcd(f, B)$¿1, the gcd is necessarily smooth, so we can combine this with (1) to get an efficient representation for the whole thing in this way.
(4) If φ has non-smooth degree which doesn't factor as $\deg(\varphi) = fA$ such that $f + A = B$ is smooth and the endomorphism ring of E_1 is unknown, we can instead use a generalization of Kani's Reducibility Criterion to 4/8 dimensions instead of 1/2 dimensions due to Robert [64]. We leave this for further reading.

This consequence has far reaching implications: it allows us not only to efficiently compute images under isogenies of arbitrary degree but allows us to store the information in polynomial memory! One such application is in the current incarnation of SQIsign, based on SQIsign2D-West [6], SQIsign2D-East [53], and SQIPrime [30].

2.4.3. *Consequence 3: IdealToIsogeny for arbitrary norm.* The final consequence we discuss here is, given a supersingular elliptic curve E_0 with known endomorphism ring $\mathrm{End}(E_0)$, the possibility to turn a left-ideal I of $\mathrm{End}(E_0)$ (including one of non-smooth norm) directly into an efficient representation for the isogeny (in the sense of Consequence 2) corresponding to that ideal (c.f. Definition 23). This idea takes inspiration from the class group action framework Clapoti [57], which we will describe further in the last lecture, Section 4, and appeared for the first time publicly in [6].[12]

Given a left-$\mathrm{End}(E_0)$ ideal I, the goal is to obtain an efficient representation of $\varphi_I : E_1 \to E_I$ by computing $\varphi_I(P_0)$ and $\varphi_I(Q_0)$, where (P_0, Q_0) is a fixed basis for $E_0[B]$, and B is smooth as above.

The first step is to search within the equivalence class of the ideal I to find two ideals I_1, I_2 of coprime norms $\mathrm{nrd}(I_1)$ and $\mathrm{nrd}(I_2)$ (resp.), each expected to be approximately $\sqrt{p}$ but chosen as small as possible. The ideals I_1, I_2 are resampled until a solution (u, v) is found to the equation:

$$u \cdot \mathrm{nrd}(I_1) + v \cdot \mathrm{nrd}(I_2) = 2^e. \tag{4}$$

By a pigeon hole principle argument, one can see that a sufficient condition for a solution to exist is $\mathrm{nrd}(I_1) \cdot \mathrm{nrd}(I_2) < B$, however solutions may exist even if this does not hold. The expected size of the solutions u and v is also approximately $\sqrt{p}$.

Using the FixedDegreeIsogeny method (point (3) in Section 2.4.2), one can efficiently represent isogenies φ_u, φ_v of E_I of degrees u and v (resp.), and say $\varphi_u : E_I \to E_u$, $\varphi_v : E_I \to E_v$. The authors provide options for optimizing this step of the subroutine, and indeed this step and the previous step have been the focus of follow-up work providing optimizations of the SQIsign2D-West algorithm [60].

Let φ_1, φ_2 denote the isogenies corresponding to the ideals I_1, I_2, respectively. Note that these isogenies map from E_0 to E_I, since they are in the same ideal class as I. Since a solution to equation (4) has been found, the next step is to apply Kani's reducibility criterion to the isogeny diamond formed by $\varphi_1 \circ \widehat{\varphi}_u : E_u \to E_I$ and $\varphi_v \circ \widehat{\varphi}_2 : E_I \to E_v$, pictured here:

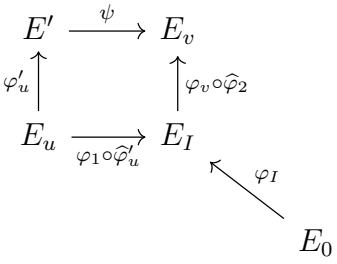

[12]This algorithm represents one of the most significant updates from the original version of SQIsign and of SQIsignHD: previous versions of the protocol imposed restrictions on the norm of the starting ideal. From a security perspective, this was a significant hitch, as strong security assumptions come from uniformly sampling a random distribution, and putting restrictions on the norm immediately violates this uniformity.

Applying Theorem 43, we obtain an isogeny $\Phi : E_u \times E_v \to E_I \times E'$ with

$$\Phi = \begin{pmatrix} \varphi_1 \circ \widehat{\varphi}_u & \varphi_2 \circ \widehat{\varphi}_v \\ -\varphi'_u & \widehat{\psi} \end{pmatrix}.$$

Knowledge of the isogeny Φ allows us to evaluate φ_1, φ_2 on arbitrary points, and so we can obtain a representation of φ_1 on the 2^e-torsion basis P_0, Q_0: $(\varphi_1(P_0), \varphi_1(Q_0))$.

The last step is to leverage φ_1 (really, I_1) to obtain an efficient representation of φ_I. There exists $\beta \in I$ such that $I_1 = I\overline{\beta}/\mathrm{nrd}(I)$, so that $\widehat{\varphi}_1 \circ \varphi_I = \beta$, and equivalently $\varphi_I = [1/d_1]\varphi_1 \circ \beta$. The basic idea is to use the fact that $\widehat{\varphi}_1 \circ \varphi_I$ is an endomorphism of E_0, but some care must be taken using the ideals I_1, I because the product $I\overline{I}_1$ is only defined up to conjugation of $\overline{I}_1$. We refer the reader to [6, Lemma 9] for details.

3. Isogeny-based digital signatures

There are now several variants of the original SQIsign protocol described in Section 1.6. The efficient representations of isogenies discussed in the previous lecture (Section 2 play a key role in all of these variants. SQIsignHD appeared first, and the other three were developed concurrently by distinct research groups. In particular:

- SQIsignHD [23]: This was the first variation after Kani's reducibility criterion was first used in isogeny-based cryptography. The main improvement of this protocol over the original SQIsign is the representation of isogenies of supersingular elliptic curves by images of torsion points allowing for the reconstruction of the isogeny. The authors work with abelian varieties of dimensions 4 ("SQIsignHDFast") and 8 ("SQIsignHDRigorous") to avoid imposing additional constraints on the isogenies they can represent[13]. Abelian surfaces (dimension 2) would be preferable for efficiency, but the authors note that, with their approach, this case imposes too many restrictive conditions on the isogenies. The authors use the higher-dimensional representation of an isogeny through images of torsion points throughout the protocol, and only evaluate a higher-dimension isogeny of abelian varieties in the verification step. As a result, the verification step was slower in this protocol.
- SQIPrime [30]: The authors of this protocol provide dimension-4 and dimension-2 variants. In spirit, this protocol is closest to SQIsignHD, with the main difference appearing the challenge round. Their goal was to keep the computations in as low a dimension as possible.
- SQIsign2D-East [53]: The authors improve the verification step of SQIsignHD by evaluating a dimension 2 isogeny instead of dimension 4 or 8. The authors provide a new algorithm for computing the commitment, they propose different constraints on the degree of the response isogeny, and they compute an auxiliary isogeny which becomes part of the signature. This protocol was simultaneously weakened and repaired in [13]. As of the writing of this document, the proof of unforgeability contains an error.
- SQIsign2D-West [6]: This variant attempts to balance the improvements offered by higher-dimension abelian varieties with the efficiency of staying in low-dimension. They work exclusively with two-dimensional isogeny representations, overcoming the previous limitations by broadening the class of isogenies that can be represented in that setting.

SQIsign2D-West is advancing in the National Institute of Standards and Technology (NIST) standardization process, so we will focus on the mathematical and cryptographic properties of this protocol.

3.1. Description of SQIsign2D-West sigma protocol.

The setup follows the same skeleton as the original version of SQIsign: it is a sigma protocol which uses a variant of the Fiat-Shamir transform to obtain a digital signature.

We again let λ denote the security parameter for the protocol. The public parameters for SQIsign2D-West are as follows:

- integer $e \approx 2\lambda$;

[13]Offering two variants reflects the current need for flexibility in our cryptographic standards. It is a positive mark to be able to balance speed, security, and signature size for different use-cases.

- integer e_{res} such that $2^{e_{\text{res}}} > 2\sqrt{2p}/\pi$;
- integer e_{chal} such that $2^\lambda < 2^{e_{\text{chal}}} \ll 2^{2\lambda}$ is the size of the challenge space;
- $e_{\text{chal}} + e_{\text{res}} \leq e$[14];
- a prime $p \approx 2^{2\lambda}$ of the form $p = c2^e - 1$, where c is a small cofactor;
- a supersingular elliptic curve $E_0/\overline{\mathbb{F}}_p$ with known endomorphism ring, in particular the paper references $E_0 : y^2 = x^3 + x$;
- (P_0, Q_0): a basis for $E_0[2^e]$;
- $D_{\text{sk}} \approx 2^{4\lambda}$ odd;
- $D_{\text{com}} = \ell_{\text{com}}^n \geq 2^{4\lambda}$, for some large prime $\ell_{\text{com}} > 2^{e_{\text{res}}}$ and some $n > 0$[15];

The steps of the sigma protocol proceed as follows:

- **Key generation (Alice):**
  - Sample a random left-ideal I_{sk} of $\mathcal{O}_0$ of norm D_{sk}.
  - Compute the secret isogeny $\varphi_{\text{sk}} : E_0 \to E_{\text{pk}}$.
  - Compute a basis $\{P_{\text{pk}}, Q_{\text{pk}}\}$ of $E_{\text{pk}}[2^e]$ and, via φ_{sk}, a basis $\{\beta_1, \beta_2, \beta_3, \beta_4\}$ of $\text{End}(E_{\text{pk}})$.
  - Publish E_{pk} and $\{P_{\text{pk}}, Q_{\text{pk}}\}$ and store the secrets φ_{sk} and $\beta_i(P_{\text{pk}})$ and $\beta_i(Q_{\text{pk}})$.
- **Commitment (Alice):**
  - Sample a random left-ideal I_{com} of $\mathcal{O}_0$ of norm D_{com}.
  - Let $\varphi_{\text{com}} = \varphi_{I_{\text{com}}}$ in the sense of Definition 23. Using Consequence 3 of Section 2.4.2, compute $E_{\text{com}} = \varphi_{\text{com}}(E_0)$ and store the secret $\varphi_{\text{com}}|_{E_0[2]}$.
  - Publish E_{com}.
- **Challenge (Bob):**
  - Sample a random integer chal in $[0, 2^{e_{\text{chal}}}]$.
  - Send chal.
    This will define an isogeny φ_{chal} of degree 2^e with kernel $\langle P_{\text{pk}} + [\text{chal}]Q_{\text{pk}} \rangle$.
- **Response (Alice):**
  - Compute I_{chal} such that $\varphi_{\text{chal}} = \varphi_I$ (c.f. Definition 23).
  - Compute $J = \overline{I}_{\text{com}} \cdot I_{\text{sk}} \cdot I_{\text{chal}}$.
  - Compute a uniformly distributed ideal I_{res} equivalent to J of norm $< 2^{e_{\text{res}}}$. Then $\varphi_{\text{res}} := \varphi_{I_{\text{res}}}$ fits in a SQIsign-style non-commutative diagram:

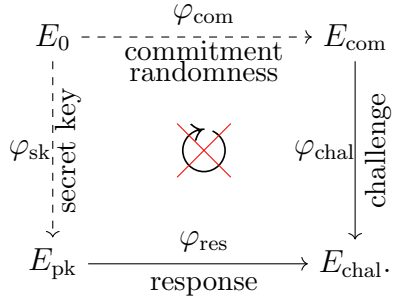

    However, instead of computing and sharing φ_{res} directly à-la-SQIsign, we use Kani's Reducibility Criterion to save time and memory.
    - Define $\varphi_{\text{res}} := \varphi_{I_{\text{res}}}$. Then $\varphi_{\text{res}} = \varphi_{\text{res}}^{\mathsf{bt}} \circ \varphi_{\text{res}}^{\mathsf{cyc}}$, where $\varphi_{\text{res}}^{\mathsf{bt}}$ is the part of φ_{res} which backtracks along $\widehat{\varphi}_{\text{chal}}$ and $\varphi_{\text{res}}^{\mathsf{cyc}}$ is the part of φ_{res} which doesn't

[14] required in the proof of soundness
[15] required in the proof of zero knowledge

backtrack. The isogeny $\varphi_{\text{res}}^{\mathsf{bt}}$ has degree 2^{bt}, where $\mathsf{bt} \geq 0$; note that the integer bt uniquely determines $\varphi_{\text{res}}^{\mathsf{bt}}$.

* Let q' be the odd part of $\deg(\varphi_{\text{res}})$ and let $f = 2^{e_{\text{res}}-\mathsf{bt}} - q'$. Choose a left ideal I_f of $\text{End}(E_{\mathsf{com}})$ of norm f.
* Using Consequence 3 of Section 2.4.3, compute the isogeny $\varphi_{\text{aux}} = \varphi_{I_f}$ (c.f. Definition 23). Define $E_{\text{aux}} = \varphi_{\text{aux}}(E_{\text{com}})$ and compute a deterministic basis $\{P_{\text{aux}}, Q_{\text{aux}}\}$ of $E_{\text{aux}}[2^{e_{\text{res}}-\mathsf{bt}}]$. Let ψ be the isogeny that makes the following diagram commute:

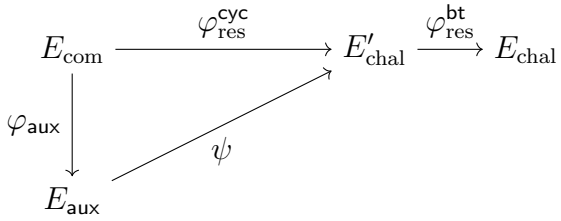

  Using Consequence 3 of Section 2.4.2, compute the action of $\varphi_{\text{res}}^{\mathsf{cyc}}$ on $E_{\text{com}}[2^{e_{\text{res}}-\mathsf{bt}}]$ and deduce $P_{\text{chal}} := \psi(P_{\mathsf{aux}})$ and $Q_{\text{chal}} := \psi(Q_{\mathsf{aux}})$.
* Let $2^k = \deg(\varphi_{\text{res}})/q'$. Publish

$$E_{\text{aux}}, P_{\text{chal}}, Q_{\text{chal}}, \mathsf{bt}, k.$$

- **Verification (Bob):**
  - *Compute E'_{chal} from bt*: Compute the isogeny $E_{\text{pk}} \to E'_{\text{chal}}$ with kernel

$$\langle [2^{\mathsf{bt}}]P_{\text{pk}} + [2^{\mathsf{bt}}\text{chal}]Q_{\text{pk}} \rangle.$$

  - *If $k > 0$, factor out the even part*: define φ_{2^k} and $\varphi_{q'}$ to be the factors

$$\varphi_{\text{res}}^{\mathsf{cyc}} = \varphi_{2^k} \circ \varphi_{q'}$$

    of degree 2^k and q' respectively. Then $\ker(\widehat{\varphi}_{2^k}) = \langle [2^{e_{\text{res}}-k-\mathsf{bt}}]P_{\text{chal}}, [2^{e_{\text{res}}-k-\mathsf{bt}}]Q_{\text{chal}} \rangle$. Reset

$$E'_{\text{chal}} \leftarrow \widehat{\varphi}_{2^k}(E'_{\text{chal}}),\ P_{\text{chal}} \leftarrow \widehat{\varphi}_{2^k}(P_{\text{chal}}),\ Q_{\text{chal}} \leftarrow \widehat{\varphi}_{2^k}(Q_{\text{chal}}).$$

    Then reset ψ to be the isogeny that makes the following diagram commute:

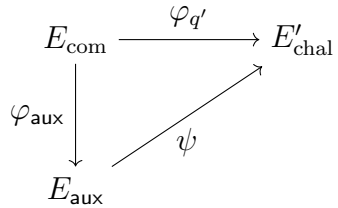

  - *Verify knowledge of secrets by computing $\varphi_{q'}$*: Compute the deterministic basis $\{P_{\mathsf{aux}}, Q_{\mathsf{aux}}\}$ of $E_{\mathsf{aux}}[2^{e_{\text{res}}-\mathsf{bt}}]$. Using Consequence 1 of Section 2.4.1 with $E_1 = E_{\mathsf{aux}}$, $E_2 = E'_{\text{chal}}$, $E_3 = E_{\text{com}}$, and $\varphi = \psi$, noting that $P_{\text{chal}} = \psi(P_{\mathsf{aux}})$ and $Q_{\text{chal}} = \psi(Q_{\mathsf{aux}})$ determines $\psi|_{E_{\mathsf{aux}}[2]}$, deduce $\varphi_{q'}$; check whether or not the image is E'_{chal}.[16]

[16]Note that in applying Consequence 1 you compute a 2-dimensional isogeny which may end in a Jacobian if the verification is failing. If this happens, return False.

In the key generation and commitment phases, it is necessary to sample a "random" isogeny of a fixed norm. Both times this is achieved through an algorithm called RandomFixedNormIdeal. In this fairly straightforward subroutine, knowledge of a basis for an endomorphism ring $\mathcal{O} \cong \mathrm{End}(E)$ is used to sample cyclic isogenies of E of fixed degree n, where the only restriction on n is that it is coprime to p. Cyclic isogenies correspond to *primitive* ideals, namely ideal I such that $I \not\subseteq m\mathcal{O}$ for any $m \in \mathbb{Z}_{>1}$.

3.2. **Fiat-Shamir transform.** While the original SQIsign [25] used a straightforward application of the Fiat-Shamir transform [32], SQIsign2D-West leverages the *commitment-recoverable* property to apply a modified Fiat-Shamir transform and obtain an even smaller signature.

The Fiat-Shamir transform turns an (interactive) sigma protocol into a digital signature. Recall that a digital signature is not interactive (and see Section 1.6 for details): we want the signer to be able to `Sign` a document D producing a signature σ_D and a verifier to be able to `Verify` that signature σ_D on the document D using the public key information associated to the signer. The sigma protocol requires back-and-forth communication between the prover (signer) and verifier, so this interactivity must be removed in order to obtain a digital signature. This process proceeds as follows:

- The signer produces follows the usual key generation and commitment processes, producing the public key $\mathrm{pk} = (E_{\mathrm{pk}}, P_{\mathrm{pk}}, Q_{\mathrm{pk}})$ and commitment φ_{com}.
- The signer uses a cryptographic hash function H to produce a challenge φ_{chal} from φ_{com} and the document D. Explicitly, $\varphi_{\mathrm{chal}} = H(\varphi_{\mathrm{com}}, \mathrm{pk}, D)$.
- The signer produces the response φ_{res} to the challenge φ_{chal}.
- The digital signature consists of an efficient representation of the transcript $(\varphi_{\mathrm{com}}, \varphi_{\mathrm{res}})$ corresponding to the public key $(E_{\mathrm{pk}}, P_{\mathrm{pk}}, Q_{\mathrm{pk}})$.
- To verify, the verifier computes the corresponding challenge φ_{chal} using the cryptographic hash function H, and then verifies that $(\varphi_{\mathrm{com}}, \varphi_{\mathrm{chal}}, \varphi_{\mathrm{res}})$ is a valid transcript corresponding to the document D and public key $(E_{\mathrm{pk}}, P_{\mathrm{pk}}, Q_{\mathrm{pk}})$.

The above process is actually repeated κ times in parallel, so that the final transcript consists of a tuple of κ transcripts. The reason for the repetition is as follows: even without knowing the secret isogeny φ_{sk}, a dishonest prover could still "guess" and convince the verifier to accept the digital signature with probability 2^{-1}. With repetition, this probability decreasts to $2^{-\kappa}$.

3.3. **Security.** To show that the sigma protocol described above will result in a secure digital signature via the Fiat-Shamir transform, we show that the sigma protocol satisfies the necessary conditions: completeness, soundness, and zero knowledge.

3.3.1. *Completeness.* Similarly to the original SQIsign, completeness is a direct consequence of the correctness of the algorithms for the steps of the sigma protocol: The algorithms provided are correct, and in particular the response will be an isogeny between the correct curves, $\varphi_{\mathrm{res}} : E_{\mathrm{pk}} \to E_{\mathrm{chal}}$.

3.3.2. *Soundness.* As a reminder, an identification protocol is sound if a dishonest prover (who does not know φ_{sk}) is unable to convince the verifier that they know φ_{sk}, except with negligible probability. The SQIsign2D-West protocol has **2-special soundness**, meaning that if a prover can produce two valid responses to two separate challenges corresponding

to the same commitment, then that prover must know the secret. Succinctly, the secret φ_{sk} can be computed from the two valid transcripts

$$(\text{com}, \text{chal}, \text{res}), (\text{com}, \text{chal}', \text{res}').$$

The idea is to use these two transcripts to construct a non-scalar endomorphism $\eta \in \text{End}(E_{\text{pk}})$. From the transcript information, one can construct isogenies $\sigma : E_{\text{com}} \to E_{\text{chal}}$ and $\sigma' : E_{\text{com}} \to E'_{\text{chal}}$ of degrees $\leq 2^{e_{\text{res}}}$. (If the prover is honest, the isogenies σ and σ' are equivalent to the honest response isogenies φ_{res} and φ'_{res} from the protocol, but all we can assume in this proof of soundness is that the transcripts provided pass verification. This is also the reason why we do not draw the secret isogeny φ_{sk} in the diagram below.) See the diagram below to summarize the isogenies obtained from these transcripts.

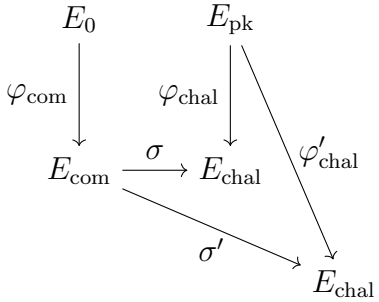

As seen in the diagram, this information yields an endomorphism $\eta \in \text{End}(E_{\text{pk}})$, namely

$$\eta = \widehat{\varphi}'_{\text{chal}} \circ \sigma' \circ \widehat{\sigma} \circ \varphi_{\text{chal}}.$$

It remains only to show that $\eta \notin \mathbb{Z}$, but this follows from the degree restrictions imposed in the protocol. See [6, Thm. 17] for details.

3.3.3. *Zero Knowledge.* As a reminder, to prove a protocol satisfies the zero knowledge property, one must show that there exists a simulator which produces accepting transcripts at the same rate as actual executions of the protocol, without access to the secret information. The SQIsign2D-West protocol is shown to have the honest-verifier zero knowledge property[17] in [6, Thm. 22], under two new oracle models, namely UTO and FIDIO.

A **uniform target oracle** (UTO) produces a random isogeny $\varphi : E \to E'$, given a supersingular elliptic curve $E/\mathbb{F}_{p^2}$ and an integer N to upper-bound the degree of φ. By "uniformly random isogeny φ", we mean that simultaneously:

- the codomain E' of φ is uniformly random with respect to the distribution of all supersingular elliptic curves over $\mathbb{F}_{p^2}$, and
- φ is uniformly random with respect to the distribution of all isogenies from the curve E of degree $\leq N$.

If $N > 2\sqrt{2p}/\pi$, then such an oracle is guaranteed to exist: For $N > 2\sqrt{2p}/\pi$, the number of left ideals of norm $\leq N$ of a fixed maximal order in a quaternion algebra whose right order is any given maximal order $\mathcal{O}'$ in that quaternion algebra is approximately the same for all possible choices of $\mathcal{O}'$. This comes from a version of Minkowski's bound for quaternion ideal classes, see [75, Prop. 17.5.6].

A **fixed degree isogeny oracle** (FIDIO) takes in a supersingular elliptic curve $E/\mathbb{F}_{p^2}$ and outputs a uniformly random isogeny $\varphi : E \to E'$ of degree N.

[17] a weaker property than zero knowledge, as it assumes the verifier is executing the protocol honestly.

The simulator proceeds to create transcripts which are statistically indistinguishable from transcripts resulting in actual executions of the protocol as follows:

- Generate a challenge isogeny $\varphi_{\text{chal}} : E_{\text{pk}} \to E_{\text{chal}}$ according to the protocol.
- Give the UTO input $(E_{\text{chal}}, 2^{e_{\text{res}}})$ and call the output $\widehat{\varphi}_{\text{res}} : E_{\text{chal}} \to E_{\text{com}}$.
- Decompose the isogeny φ_{res} into its odd-degree part $\varphi_{\text{res}}^{(1)}$ of degree q', and 2-power degree part ψ of degree 2^k. Further decompose the isogeny ψ into a part of degree 2^{bt} which backtracks along $\widehat{\varphi}_{\text{chal}}$ and the nonbacktracking part of degree $2^{k-\mathsf{bt}}$. Let $r' = k - \mathsf{bt}$.
- Give the FIDIO input $(E_{\text{com}}, 2^{e_{\text{res}},-r'} - q')$ and call the output $\varphi_{\mathsf{aux}} : E_{\text{com}} \to E_{\mathsf{aux}}$.

Honest executions of the protocol are computed using an equivalent procedure to the above, but in a different order of operations (commitment, challenge, response). This is the essence of the proof of honest verifier zero knowledge.

4. Overview of class group actions in higher-dimensional isogeny-based cryptography

Class group actions and Kani's Reducibility Criterion (Theorem 43) both provide effective tools in isogeny-based cryptography. A new wave of protocols seeks to combine these tools to create new flexible, efficient cryptographic primitives. In particular, one can use higher-dimension abelian varieties to evaluate the class group action on oriented supersingular elliptic curves. This framework, introduced by [57], gave the first polynomial time algorithm to compute the class group action for any given class group. See the blog post [59] for a history of the problem that [57] solves. The methods used are inspired by the work of [52], who gave an algorithm RanIsogImages for evaluating isogenies of non-smooth degree using Kani'a reducibility criterion. Although the "full description" promised in the abstract for Clapoti(s) has not yet appeared, other groups of authors have published extensions and improvements to the public version Clapoti(s) framework[18], which we will overview in Section 4.2.

4.1. The Clapoti Framework.

The framework for efficient evaluation of the class group action for isogeny-based cryptography was reworked in [57]. This work begat a variety of improvements in this direction, many of which we will visit in subsequent sections.

As this framework applies in the setting of oriented supersingular elliptic curves, we begin by recalling the main theory and definitions here.

4.1.1. *Oriented elliptic curves.* Orientations provide a tool for using commutative orders inside of an elliptic curve's endomorphism ring, even when the endomorphism ring itself is noncommutative. Such a strategy also works for ordinary elliptic curves and supersingular elliptic curves defined over $\mathbb{F}_p$, whose endomorphism rings are commutative. We will present this theory in its full generality, but keep in mind the special case where $\mathrm{End}(E)$ is noncommutative.

Definition 44 (Orientations)**.** Fix an elliptic curve E over a finite field $\mathbb{F}_q$. Let K be an imaginary quadratic field which embeds into the endomorphism algebra of E. Namely, there exists an embedding

$$\iota : K \hookrightarrow \mathrm{End}(E) \otimes_{\mathbb{Z}} \mathbb{Q}.$$

The pair (E, ι) is a **K-oriented elliptic curve**, or that E is K-oriented with orientation ι. Let $\mathcal{O}$ be the order of K satisfying $\iota(K) \cap \mathrm{End}(E) = \iota(\mathcal{O})$. We say that E is **primitively $\mathcal{O}$-oriented**.

Two K-oriented elliptic curves $(E, \iota), (E', \iota')$ are **isomorphic** if there exists an isomorphism $\eta : E \to E'$ such that $\iota'(-) = \eta \circ \iota(-) \circ \widehat{\eta}$.

For the remainder of this section, we fix K to be an imaginary quadratic field and $\mathcal{O}$ to be an order in K. We fix a prime p, and all elliptic curves are defined over a field of characteristic p. Let $\mathrm{SS}_{\mathcal{O}}^{pr}$ denote the set of isomorphism classes of primitively $\mathcal{O}$-oriented supersingular elliptic curves over $\overline{\mathbb{F}}_p$.

The orientation allows the translation between imaginary quadratic ideals of $\mathcal{O}$ and isogenies of E. This theory gives a general framework for the group action used in CSIDH [15] (see the PAWS2025 notes [41] for an exposition).

[18] which we will from here on refer to as *Clapoti*.

Isogenies of the underlying elliptic curves induce **K-oriented isogenies** in the following way: If (E, ι) is a primitively $\mathcal{O}$-oriented elliptic curve and $\varphi : E \to E'$ is an isogeny, then φ induces a K-orientation $\varphi_*\iota$ on E' defined:

$$\varphi_*\iota(-) = \frac{1}{[\deg \varphi]}\varphi \circ \iota(-) \circ \widehat{\varphi}.$$

The induced K-oriented elliptic curve $(E', \varphi_*\iota)$ is primitively $\mathcal{O}'$-oriented for an order $\mathcal{O}'$ of K for which exactly one of the following holds:

- $\mathcal{O}' = \mathcal{O}$, in which case we say φ is **horizontal**;
- $[\mathcal{O}' : \mathcal{O}] = \deg \varphi$, in which case we way φ is **ascending**;
- $[\mathcal{O} : \mathcal{O}'] = \deg \varphi$, in which case we way φ is **descending**.

Definition 45 (Class group action)**.** Let (E, ι) be a primitively $\mathcal{O}$-oriented elliptic curve. Let $\mathfrak{a}$ be an integral ideal of $\mathcal{O}$ coprime to the conductor of $\mathcal{O}$ and of norm coprime to p. The intersection:

$$E[\mathfrak{a}] = \bigcap_{\alpha \in \mathfrak{a}} \ker(\iota(\alpha))$$

defines a finite subgroup of E. This group in turn defines an isogeny $\varphi_\mathfrak{a} : E \to E/E[\mathfrak{a}]$ with kernel $E[\mathfrak{a}]$. Furthermore, we obtain a K-oriented isogeny

$$\varphi_\mathfrak{a} : (E, \iota) \to (E/E[\mathfrak{a}], (\varphi_\mathfrak{a})_*\iota)$$

Remark 46. We briefly justify why the above action is of the class group of $\mathcal{O}$: First, note that all ideals in the class group of $\mathcal{O}$ are by definition coprime to the conductor of $\mathcal{O}$, and from each ideal class we can always choose an integral representative of norm coprime to p. If $\mathfrak{a}$ is principal, say $\mathfrak{a} = (\alpha)$, it is necessary to show that $(\varphi_\mathfrak{a})_*\iota = \iota$, but this follows from the commutativity of the elements of $\mathcal{O}$.

4.1.2. *Class groups and supersingular elliptic curves.* We restrict our attention to the supersingular case by considering this class group action on $\mathrm{SS}^{pr}_\mathcal{O}$.

Theorem 47. *The class group of the order $\mathcal{O}$ acts freely on the set $SS^{pr}_\mathcal{O}$. Moreover:*

- *this action is transitive if p is ramified in $\mathcal{O}$,*
- *this action has two orbits if p is inert in $\mathcal{O}$.*

Proof. See [54, Thm. 3.4], which requires some CM-theory for elliptic curves. This result extends classical work of Waterhouse [77, Thm. 4.5], who considered the case where the ring of $\mathbb{F}_q$-endomorphisms of an elliptic curve was isomorphic to an imaginary quadratic order $\mathcal{O}$. As noted by Schoof in [68, Proof of Thm. 4.5], there was a slight omission in the original statement of Waterhouse, which is corrected in a later theorem of the same paper, namely [77, Thm. 5.3]. □

Remark 48. In Theorem 47, the case where p is split in $\mathcal{O}$ is omitted – in this case, the isogeny class of elliptic curves with a primitive $\mathcal{O}$-orientation is ordinary, not supersingular.

In the setup of Clapoti [57], we drop the orientation notation (E, ι) and focus explicitly on the elliptic curves. To further condense notation, let $E_\mathfrak{a}$ denote the codomain of the isogeny $\varphi_\mathfrak{a}$, in particular: $E_\mathfrak{a} = E/E[\mathfrak{a}]$. The following proposition (an application of Theorem 43) provides the necessary foundation to unconditionally compute the action of $[\mathfrak{a}] \in \mathrm{Cl}(\mathcal{O})$ on a supersingular elliptic curve $E \in \mathrm{SS}^{pr}_\mathcal{O}$ in polynomial time.

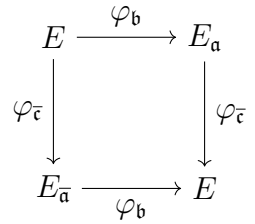

FIGURE 3. Clapoti isogeny diamond, see Proposition 49.

Proposition 49 (Prop. 2.1 [57])**.** *Let $\mathfrak{a}, \mathfrak{b}, \mathfrak{c}$ be integral ideals of $\mathcal{O}$ coprime to the conductor of $\mathcal{O}$. Suppose $\mathfrak{a}, \mathfrak{b}, \mathfrak{c}$ all belong to the same class in $\mathrm{Cl}(\mathcal{O})$, and suppose the norms $N(\mathfrak{b})$, $N(\mathfrak{c})$ of $\mathfrak{b}$ and $\mathfrak{c}$ are coprime.*

Then, there is an $N = (N(\mathfrak{b}) + N(\mathfrak{c}))$-isogeny $\Phi : E \times E \to E_{\mathfrak{a}} \times E_{\bar{\mathfrak{a}}}$ with kernel

$$\ker \Phi = \{([N(\mathfrak{b})]P, (\varphi_{\mathfrak{b}} \circ \varphi_{\bar{\mathfrak{c}}})P : P \in E[N]\}.$$

This isogeny is efficiently computable when the N-torsion group $E[N]$ is accessible.

Proof. We immediately remark that $[\mathfrak{b}] = [\mathfrak{a}]$, so the corresponding isogenies $\varphi_{\mathfrak{a}}$ and $\varphi_{\mathfrak{b}}$ from the elliptic curve E have the same codomain. Likewise, $[\mathfrak{c}] = [\mathfrak{a}]$ so $[\bar{\mathfrak{c}}] = [\bar{\mathfrak{a}}]$ and the codomains of the isogenies $\varphi_{\bar{\mathfrak{a}}}$ and $\varphi_{\bar{\mathfrak{c}}}$ of E are also the same.

The accessiblity of the group $E[N]$ means we can find generators P, Q defined over a small field. In the application in mind, $p = c \cdot 2^e - 1$ so that supersingular elliptic curves over $\mathbb{F}_p$ necessarily have $p + 1$ points (think back to Theorem 13 and Theorem 14 to recall why this is true). Then, $\#E(\mathbb{F}_p) = c \cdot 2^e$, so if e is large we have a lot of 2-torsion points accessible and we can take $N = 2^e$[19]. The goal becomes to find two ideals $\mathfrak{b}$ and $\mathfrak{c}$ equivalent to $\mathfrak{a}$ and satisfying

$$N(\mathfrak{b}) + N(\mathfrak{c}) = N.$$

Once this is achieved, the corresponding isogenies fit into an isogeny diamond (see Figure 3), and Theorem 43 (Kani's reducibility criterion) applies.

The resulting isogeny $\Phi : E \times E \to E_{\mathfrak{a}} \times E_{\bar{\mathfrak{a}}}$ is given by the matrix

$$\Phi = \begin{pmatrix} \varphi_{\mathfrak{b}} & \widehat{\varphi_{\bar{\mathfrak{c}}}} \\ -\varphi_{\bar{\mathfrak{c}}} & \widehat{\varphi_{\mathfrak{b}}} \end{pmatrix},$$

and kernel as above, given by Theorem 43. The computation of the kernel depends on the accessbility of the points in $E[N]$, and the isogenies $[N(\mathfrak{b})]$ and $(\varphi_{\mathfrak{b}} \circ \varphi_{\bar{\mathfrak{c}}})$ are efficient to compute. □

In Proposition 49, it is essential to have an efficient method for finding the ideals $\mathfrak{b}$ and $\mathfrak{c}$ in the same ideal class as $\mathfrak{a}$. Given $\mathfrak{b}$ in the same class as $\mathfrak{a}$, the proof of Proposition 2.1 [57] describes how to explicitly obtain a third equivalent ideal $\mathfrak{c}$. If we already have $\mathfrak{b}$, then the norm of this target ideal $\mathfrak{c}$ must satisfy $N(\mathfrak{c}) = N - N(\mathfrak{b})$. Explicitly, this means that $\mathfrak{c} = (\overline{\gamma}_c / N(\mathfrak{a}))\mathfrak{a}$, for some $\gamma_c \in \mathfrak{a}$ satisfying $N(\gamma_c) = N(\mathfrak{c})N(\mathfrak{a})$.

[19]Most applications take $N = 2^e$, although the idea is general.

4.2. **Beyond Clapoti.** Various research groups have optimized and extended the Clapoti framework in several ways. In the following list, we highlight a selection of the these advances (current as of January 2026). The list is chronological, based on the first public appearance of the corresponding paper.

- In SQIsign2D-West ([6], as described in Section 3), authors attribute the inspiration for their IdealToIsogeny algorithm to the Clapoti authors. The IdealToIsogeny algorithm extends the framework of Clapoti from quadratic ideals to quaternion ideals.
- In KLaPoTi [60], the authors use the KLPT algorithm to solve a key equation in Clapoti framework and work with 2-dimensional $(2^e, 2^e)$-isogenies.
- In [21] PEGASIS, the authors provide a refinement of the techniques in the Clapoti framework and give an algorithm to compute isogenies coming from a class group action using isogenies of dimension-4 abelian varieties.
- In [9] Qlapoti, the authors address the question of finding the quaternion equivalents of the ideals $\mathfrak{b}, \mathfrak{c}$ as in Proposition 49. In particular, the algorithm takes on input a quaternion ideal I and outputs two equivalent ideals I_1, I_2 satisfying
$$N(I_1) + N(I_2) = 2^e.$$
This optimization improves the speed of the IdealToIsogeny algorithm in [6].
- In [22] qt-PEGASIS, the authors provide a simplification of the algorithm described in PEGASIS using a refinement of a technique described in Qlapoti. The qt-PEGASIS algorithm uses quadratic ideals instead of quaternion ideals, and avoids a complicated rerandomization technique which was necessary for PEGASIS.

5. Project Descriptions

(1) **The Spine of the HD Isogeny Graph** The structure of the $\mathbb{F}_p$-subgraph of the full supersingular isogeny graph in dimension 1, the *spine*, was studied in [4]. The best known generic attack [26] on the isogeny problem (Problem 2) first computes a path to the spine in the supersingular isogeny graph. The complexity of this algorithm is dependent on the structure of the spine within the full graph.

Now, suppose you are given a supersingular elliptic curve $E/\mathbb{F}_{p^2}$. Then the Weil restriction of E is an abelian surface over $\mathbb{F}_p$, which lies in the natural generalization of the spine to two dimensions. This doesn't give an obvious attack (see [45, Section 3.3]), but it does show the importance of properly understanding the structure of the spine of the two-dimensional isogeny graph. One could even take this further: in general, what are the possible graph-theoretic structures for the spine of the g-dimensional isogeny graph?

(2) **Computing 2D cyclic isogenies** Thus far, the literature on computing isogenies in higher dimensions has focussed primarily on $(\ell, \dots, \ell)$-isogenies. That is, isogenies of degree ℓ^g which are maximally isotropic with respect to the ℓ-adic Tate pairing.

For simplicity, we now restrict to $g = 2$. The kernel of an (ℓ, ℓ)-isogeny is isomorphic as a group to $\mathbb{Z}/\ell\mathbb{Z} \times \mathbb{Z}/\ell\mathbb{Z}$ and is generated by an ℓ-torsion point. However, these are not the only isogenies which respect polarizations. There has also been some work on computing *cyclic isogenies*. Crucially, in some cases there exist isogenies whose degree does not grow with the dimension. In particular, let (A, ξ) be a principally polarized abelian surface with maximal real multiplication by a degree 2 real number field K_0. Suppose that ℓ splits into principal ideals $\mu\mathcal{O}_{K_0}$ and $\overline{\mu}\mathcal{O}_{K_0}$ in the maximal order $\mathcal{O}_{K_0}$ of K_0. A *μ-torsion* point of (A, ξ) is then a point P such that $\mu(P) = 0$; recall that $\mu \in \mathcal{O}_{K_0}$ and $\mathcal{O}_{K_0}$ is isomorphic to a subring of $\mathrm{End}(A)$ so μ is an endomorphism of A. If μ is totally positive and of prime norm and P is nontrivial, then P generates the kernel of an isogeny from (A, ξ) of prime degree ℓ which respects both polarizations and real multiplication, called a *cyclic isogeny* in the literature.

These cyclic isogenies are notoriously difficulty to compute, despite their potential for outperforming $(\ell, \dots, \ell)$-isogenies as the dimension grows. There exists a method for computation [29] but no working implementation. However, Luciano Maino proposed a new method for computing examples of these cyclic isogenies in his PhD thesis [42, Chapter 7] making use of Kani's Reducibility Criterion. While this method requires computing the $(\ell, \dots, \ell)$-isogenies lying above the cyclic isogenies, it does have potential to give us the first concrete examples of cyclic isogenies and increase our understanding of these objects.

This project proposes to fully work out the idea presented in [42, Chapter 7] for computing cyclic isogenies in dimension two and get a first implementation of some examples.

(3) **Non-standard graph structures in HD** To date, the only isogeny graph structures that have received (public) attention are the graphs of

- principally polarized supersingular or superspecial abelian varieties (see [2, 36]);
- principally polarized g-dimensional abelian varieties with locally maximal real multiplication by a given totally real number field of degree g over $\mathbb{Q}$ (see [44, Chapter 3], [11]);

- abelian varieties of dimension-g with commutative, Bass endomorphism rings (see [5]).

However, there are many more possibilities for graph structures. Understanding more of the possibilities could open the door for graph-theoretic cryptographic algorithms. Some interesting cases to start off with:

- Let $E/\mathbb{F}_{p^2}$ be supersingular and $E'/\mathbb{F}_{p^k}$ be ordinary. What are the possible structures of isogeny graphs containing $E \times E'$?
- Formalize the dimension-g CM case as outlined thus far only in talks such as `https://www.martindale.info/talks/Budapest.pdf`.
- Does anything special happen when the endomorphism algebra is very special, e.g. a cyclotomic field?

References

[1] Marius A. Aardal, Andrea Basso, Luca De Feo, Sikhar Patranabis, and Benjamin Wesolowski. "A Complete Security Proof of SQIsign". In: *Advances in Cryptology – CRYPTO 2025*. Ed. by Yael Tauman Kalai and Seny F. Kamara. Cham: Springer Nature Switzerland, 2025, pp. 190–222. ISBN: 978-3-032-01887-8.

[2] Yusuke Aikawa, Ryokichi Tanaka, and Takuya Yamauchi. "Isogeny graphs on supersingular abelian varieties: eigenvalues and connection to Bruhat-Tits buildings". In: *Canad. J. Math.* 76.6 (2024), pp. 1891–1916. ISSN: 0008-414X,1496-4279. DOI: 10.4153/S0008414X23000676. URL: https://doi-org.ezproxy.lib.vt.edu/10.4153/S0008414X23000676.

[3] Sarah Arpin. "Adding level structure to supersingular elliptic curve isogeny graphs". In: *J. Théor. Nombres Bordeaux* 36.2 (2024), pp. 405–443. ISSN: 1246-7405,2118-8572. DOI: 10.5802/jtnb.1283. URL: https://doi-org.ezproxy.lib.vt.edu/10.5802/jtnb.1283.

[4] Sarah Arpin, Catalina Camacho-Navarro, Kristin Lauter, Joelle Lim, Kristina Nelson, Travis Scholl, and Jana Sotáková. "Adventures in supersingularland". In: *Exp. Math.* 32.2 (2023), pp. 241–268. ISSN: 1058-6458,1944-950X. DOI: 10.1080/10586458.2021.1926009.

[5] Sarah Arpin, Stefano Marseglia, and Caleb Springer. *Isogeny graphs of abelian varieties and singular ideals in orders*. 2025. arXiv: 2508.03570 [math.NT]. URL: https://arxiv.org/abs/2508.03570.

[6] Andrea Basso, Pierrick Dartois, Luca De Feo, Antonin Leroux, Luciano Maino, Giacomo Pope, Damien Robert, and Benjamin Wesolowski. "SQIsign2D–West". In: *Advances in cryptology—ASIACRYPT 2024. Part III*. Vol. 15486. Lecture Notes in Comput. Sci. Springer, Singapore, 2025, pp. 339–370. ISBN: 978-981-96-0890-4; 978-981-96-0891-1. DOI: 10.1007/978-981-96-0891-1_11.

[7] Jonas Bergström, Valentijn Karemaker, and Stefano Marseglia. *Abelian varieties over finite fields with commutative endomorphism algebra: theory and algorithms*. 2025. arXiv: 2409.08865 [math.NT]. URL: https://arxiv.org/abs/2409.08865.

[8] Daniel J. Bernstein, Luca De Feo, Antonin Leroux, and Benjamin Smith. "Faster computation of isogenies of large prime degree". In: *ANTS XIV—Proceedings of the Fourteenth Algorithmic Number Theory Symposium*. Vol. 4. Open Book Ser. Math. Sci. Publ., Berkeley, CA, 2020, pp. 39–55. ISBN: 978-1-935107-08-8; 978-1-935107-07-1. DOI: 10.2140/obs.2020.4.39. URL: https://doi-org.ezproxy.lib.vt.edu/10.2140/obs.2020.4.39.

[9] Giacomo Borin, Maria Corte-Real Santos, Jonathan Komada Eriksen, Riccardo Invernizzi, Marzio Mula, Sina Schaeffler, and Frederik Vercauteren. "Qlapoti: Simple and Efficient Translation of Quaternion Ideals to Isogenies". In: *Advances in Cryptology – ASIACRYPT 2025*. Ed. by Goichiro Hanaoka and Bo-Yin Yang. Singapore: Springer Nature Singapore, 2026, pp. 174–205. ISBN: 978-981-95-5113-2.

[10] Wieb Bosma, John Cannon, and Catherine Playoust. "The Magma algebra system. I. The user language". In: *J. Symbolic Comput.* 24.3-4 (1997). Computational algebra and number theory (London, 1993), pp. 235–265. ISSN: 0747-7171. DOI: 10.1006/jsco.1996.0125. URL: http://dx.doi.org/10.1006/jsco.1996.0125.

[11] Ernest Hunter Brooks, Dimitar Jetchev, and Benjamin Wesolowski. "Isogeny graphs of ordinary abelian varieties". In: *Res. Number Theory* 3 (2017), Paper No. 28, 38.

ISSN: 2522-0160,2363-9555. DOI: 10.1007/s40993-017-0087-5. URL: https://doi-org.ezproxy.lib.vt.edu/10.1007/s40993-017-0087-5.

[12] J. W. S. Cassels and E. V. Flynn. *Prolegomena to a Middlebrow Arithmetic of Curves of Genus 2*. London Mathematical Society Lecture Note Series. Cambridge University Press, 1996.

[13] Wouter Castryck, Mingjie Chen, Riccardo Invernizzi, Gioella Lorenzon, and Frederik Vercauteren. *Breaking and Repairing SQIsign2D-East*. Cryptology ePrint Archive, Paper 2024/1453. 2024. URL: https://eprint.iacr.org/2024/1453.

[14] Wouter Castryck, Thomas Decru, and Benjamin Smith. "Hash functions from superspecial genus-2 curves using Richelot isogenies". In: *Journal of Mathematical Cryptology* 14.1 (2020), pp. 268–292. DOI: doi:10.1515/jmc-2019-0021. URL: https://doi.org/10.1515/jmc-2019-0021.

[15] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. "CSIDH: an efficient post-quantum commutative group action". In: *Advances in cryptology—ASIACRYPT 2018. Part III*. Vol. 11274. Lecture Notes in Comput. Sci. Springer, Cham, 2018, pp. 395–427. ISBN: 978-3-030-03332-3; 978-3-030-03331-6. DOI: 10.1007/978-3-030-03332-3_15. URL: https://doi-org.ezproxy.lib.vt.edu/10.1007/978-3-030-03332-3_15.

[16] Tommaso Giorgio Centeleghe and Jakob Stix. "Categories of abelian varieties over finite fields II: abelian varieties over $\mathbb{F}_q$ and Morita equivalence". In: *Israel J. Math.* 257.1 (2023), pp. 103–170. ISSN: 0021-2172,1565-8511. DOI: 10.1007/s11856-023-2536-2. URL: https://doi-org.ezproxy.lib.vt.edu/10.1007/s11856-023-2536-2.

[17] Tommaso Giorgio Centeleghe and Jakob Stix. "Categories of abelian varieties over finite fields, I: Abelian varieties over $\mathbb{F}_p$". In: *Algebra Number Theory* 9.1 (2015), pp. 225–265. ISSN: 1937-0652,1944-7833. DOI: 10.2140/ant.2015.9.225. URL: https://doi-org.ezproxy.lib.vt.edu/10.2140/ant.2015.9.225.

[18] Denis X. Charles, Kristin E. Lauter, and Eyal Z. Goren. "Cryptographic hash functions from expander graphs". In: *J. Cryptology* 22.1 (2009), pp. 93–113. ISSN: 0933-2790,1432-1378. DOI: 10.1007/s00145-007-9002-x. URL: https://doi-org.ezproxy.lib.vt.edu/10.1007/s00145-007-9002-x.

[19] Jean-Marc Couveignes. *Hard Homogeneous Spaces*. Cryptology ePrint Archive, Report 2006/291. https://ia.cr/2006/291. 2006.

[20] Thinh Hung Dang and Dustin Moody. "New types of formula for isogenies between elliptic curves". In: *Des. Codes Cryptogr.* 93.10 (2025), pp. 4525–4543. ISSN: 0925-1022,1573-7586. DOI: 10.1007/s10623-025-01692-y. URL: https://doi-org.ezproxy.lib.vt.edu/10.1007/s10623-025-01692-y.

[21] Pierrick Dartois, Jonathan Komada Eriksen, Tako Boris Fouotsa, Arthur Herlédan Le Merdy, Riccardo Invernizzi, Damien Robert, Ryan Rueger, Frederik Vercauteren, and Benjamin Wesolowski. "PEGASIS: Practical Effective Class Group Action using 4-Dimensional Isogenies". In: *Advances in Cryptology – CRYPTO 2025*. Ed. by Yael Tauman Kalai and Seny F. Kamara. Cham: Springer Nature Switzerland, 2025, pp. 67–99. ISBN: 978-3-032-01855-7.

[22] Pierrick Dartois, Jonathan Komada Eriksen, Riccardo Invernizzi, and Frederik Vercauteren. *qt-Pegasis: Simpler and Faster Effective Class Group Actions*. Cryptology ePrint Archive, Paper 2025/1859. 2025. URL: https://eprint.iacr.org/2025/1859.

[23] Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski. "SQIsignHD: new dimensions in cryptography". In: *Advances in cryptology—EUROCRYPT 2024. Part I*. Vol. 14651. Lecture Notes in Comput. Sci. Springer, Cham, [2024] ©2024, pp. 3–32. ISBN: 978-3-031-58715-3; 978-3-031-58716-0. DOI: 10.1007/978-3-031-58716-0_1. URL: https://doi-org.ezproxy.lib.vt.edu/10.1007/978-3-031-58716-0_1.

[24] Pierrick Dartois, Luciano Maino, Giacomo Pope, and Damien Robert. "An Algorithmic Approach to (2, 2)-Isogenies in the Theta Model and Applications to Isogeny-Based Cryptography". In: *Advances in Cryptology – ASIACRYPT 2024*. Ed. by Kai-Min Chung and Yu Sasaki. Singapore: Springer Nature Singapore, 2025, pp. 304–338. ISBN: 978-981-96-0891-1.

[25] Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski. "SQISign: Compact Post-quantum Signatures from Quaternions and Isogenies". In: *Advances in Cryptology – ASIACRYPT 2020*. Ed. by Shiho Moriai and Huaxiong Wang. Cham: Springer International Publishing, 2020, pp. 64–93. ISBN: 978-3-030-64837-4.

[26] Christina Delfs and Steven D. Galbraith. "Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$". In: *Des. Codes Cryptogr.* 78.2 (2016), pp. 425–440. ISSN: 0925-1022,1573-7586. DOI: 10.1007/s10623-014-0010-1.

[27] Pierre Deligne. "Variétés abéliennes ordinaires sur un corps fini". In: *Invent. Math.* 8 (1969), pp. 238–243. ISSN: 0020-9910,1432-1297. DOI: 10.1007/BF01406076. URL: https://doi-org.ezproxy.lib.vt.edu/10.1007/BF01406076.

[28] Max Deuring. "Die Typen der Multiplikatorenringe elliptischer Funktionenkörper." In: *Abh. Math. Sem. Hansischen Univ.* 14 (1941), pp. 197–272.

[29] Alina Dudeanu, Dimitar Jetchev, Damien Robert, and Marius Vuille. "Cyclic isogenies for abelian varieties with real multiplication". In: *Mosc. Math. J.* 22.4 (2022), pp. 613–655. ISSN: 1609-3321,1609-4514. DOI: 10.17323/1609-4514-2022-22-4-613-655.

[30] Max Duparc and Tako Boris Fouotsa. *SQIPrime: A dimension 2 variant of SQISignHD with non-smooth challenge isogenies*. Cryptology ePrint Archive, Paper 2024/773. 2024. URL: https://eprint.iacr.org/2024/773.

[31] Bas Edixhoven, Gerard van der Geer, and Ben Moonen. *Abelian varieties*. (Draft in preparation but public). 2026. URL: http://van-der-geer.nl/~gerard/AV.pdf.

[32] Amos Fiat and Adi Shamir. "How to prove yourself: Practical solutions to identification and signature problems". In: *Conference on the theory and application of cryptographic techniques*. Springer. 1986, pp. 186–194.

[33] Robin Hartshorne. *Algebraic geometry*. Vol. No. 52. Graduate Texts in Mathematics. Springer-Verlag, New York-Heidelberg, 1977, pp. xvi+496. ISBN: 0-387-90244-9.

[34] Tomoyoshi Ibukiyama, Toshiyuki Katsura, and Frans Oort. "Supersingular curves of genus two and class numbers". In: *Compositio Mathematica* 57 (1986), pp. 127–152. URL: https://www.numdam.org/item/CM_1986__57_2_127_0/.

[35] Bruce W. Jordan, Allan G. Keeton, Bjorn Poonen, Eric M. Rains, Nicholas Shepherd-Barron, and John T. Tate. "Abelian varieties isogenous to a power of an elliptic curve". In: *Compos. Math.* 154.5 (2018), pp. 934–959. ISSN: 0010-437X,1570-5846. DOI: 10.1112/S0010437X17007990. URL: https://doi-org.ezproxy.lib.vt.edu/10.1112/S0010437X17007990.

[36] Bruce W. Jordan and Yevgeny Zaytman. *Isogeny graphs of superspecial abelian varieties and Brandt matrices*. 2025. arXiv: 2005.09031 [math.NT]. URL: https://arxiv.org/abs/2005.09031.

[37] Ernst Kani. "The number of curves of genus two with elliptic differentials." In: *Journal für die reine und angewandte Mathematik* 1997.485 (1997), pp. 93–122. DOI: doi:10.1515/crll.1997.485.93.

[38] Nicholas M. Katz and Barry Mazur. *Arithmetic moduli of elliptic curves*. Vol. 108. Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, 1985, pp. xiv+514. ISBN: 0-691-08349-5; 0-691-08352-5. DOI: 10.1515/9781400881710. URL: https://doi.org/10.1515/9781400881710.

[39] Neal Koblitz. "Elliptic curve cryptosystems". In: *Math. Comp.* 48.177 (1987), pp. 203–209. ISSN: 0025-5718,1088-6842. DOI: 10.2307/2007884. URL: https://doi-org.ezproxy.lib.vt.edu/10.2307/2007884.

[40] David Kohel, Kristin Lauter, Christophe Petit, and Jean-Pierre Tignol. "On the quaternion ℓ-isogeny path problem". In: *LMS J. Comput. Math.* 17 (2014), pp. 418–432. ISSN: 1461-1570. DOI: 10.1112/S1461157014000151. URL: https://doi-org.ezproxy.lib.vt.edu/10.1112/S1461157014000151.

[41] Sabrina Kunzweiler. *Introduction to mathematical cryptography*. Lecture notes for the Preliminary Arizona Winter School 2025. Preliminary Arizona Winter School / Southwest Center for Arithmetic Geometry. 2025. URL: https://swc-math.github.io/aws/2026/PAWSKunzweiler/2025PAWSKunzweilerNotes.pdf.

[42] Luciano Maino. "Factoring Isogenies in Higher Dimension and Applications". Available at https://research-information.bris.ac.uk/en/studentTheses/factoring-isogenies-in-higher-dimension-and-applications. PhD thesis. University of Bristol, School of Computer Science, Mar. 2025.

[43] Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, and Benjamin Wesolowski. "A direct key recovery attack on SIDH". In: *Advances in cryptology—EUROCRYPT 2023. Part V*. Vol. 14008. Lecture Notes in Comput. Sci. Springer, Cham, 2023, pp. 448–471. ISBN: 978-3-031-30588-7; 978-3-031-30589-4. DOI: 10.1007/978-3-031-30589-4_16.

[44] Chloe Martindale. "Isogeny graphs, modular polynomials, and applications". Universiteit Leiden and Bordeaux University. PhD thesis. 2018. URL: https://www.martindale.info/research/Thesis.pdf.

[45] Chloe Martindale and Lorenz Panny. *How to not break SIDH*. Cryptology ePrint Archive, Paper 2019/558. 2019. URL: https://eprint.iacr.org/2019/558.

[46] Arthur Herlédan Le Merdy and Benjamin Wesolowski. *Unconditional foundations for supersingular isogeny-based cryptography*. Cryptology ePrint Archive, Paper 2025/271. 2025. URL: https://eprint.iacr.org/2025/271.

[47] Victor S. Miller. "Use of elliptic curves in cryptography". In: *Advances in cryptology—CRYPTO '85 (Santa Barbara, Calif., 1985)*. Vol. 218. Lecture Notes in Comput. Sci. Springer, Berlin, 1986, pp. 417–426. ISBN: 3-540-16463-4. DOI: 10.1007/3-540-39799-X_31. URL: https://doi-org.ezproxy.lib.vt.edu/10.1007/3-540-39799-X_31.

[48] J. S. Milne. "Jacobian varieties". In: *Arithmetic geometry (Storrs, Conn., 1984)*. Springer, New York, 1986, pp. 167–212. ISBN: 0-387-96311-1.

[49] D. Mumford. "On the equations defining abelian varieties. I". In: *Invent. Math.* 1 (1966), pp. 287–354. ISSN: 0020-9910,1432-1297. DOI: 10.1007/BF01389737. URL: https://doi.org/10.1007/BF01389737.

[50] D. Mumford, J. Fogarty, and F. Kirwan. *Geometric invariant theory*. Vol. 34. Ergebnisse der Mathematik und ihrer Grenzgebiete (2) [Results in Mathematics and Related Areas (2)]. Springer-Verlag, Berlin, 1994, pp. xiv+292. ISBN: 3-540-56963-4.

[51] David Mumford. *Abelian varieties*. Vol. 5. Tata Institute of Fundamental Research Studies in Mathematics. With appendices by C. P. Ramanujam and Yuri Manin, Corrected reprint of the second (1974) edition. Tata Institute of Fundamental Research, Bombay; by Hindustan Book Agency, New Delhi, 2008, pp. xii+263. ISBN: 978-81-85931-86-9; 81-85931-86-0.

[52] Kohei Nakagawa and Hiroshi Onuki. "QFESTA: efficient algorithms and parameters for FESTA using quaternion algebras". In: *Advances in cryptology—CRYPTO 2024. Part V*. Vol. 14924. Lecture Notes in Comput. Sci. Springer, Cham, [2024] ©2024, pp. 75–106. ISBN: 978-3-031-68387-9; 978-3-031-68388-6. DOI: 10.1007/978-3-031-68388-6_4. URL: https://doi-org.ezproxy.lib.vt.edu/10.1007/978-3-031-68388-6_4.

[53] Kohei Nakagawa and Hiroshi Onuki. *SQIsign2D-East: A New Signature Scheme Using 2-dimensional Isogenies*. Cryptology ePrint Archive, Paper 2024/771. 2024. URL: https://eprint.iacr.org/2024/771.

[54] Hiroshi Onuki. "On oriented supersingular elliptic curves". In: *Finite Fields Appl.* 69 (2021), Paper No. 101777, 18. ISSN: 1071-5797,1090-2465. DOI: 10.1016/j.ffa.2020.101777. URL: https://doi-org.ezproxy.lib.vt.edu/10.1016/j.ffa.2020.101777.

[55] Saki Otsuki, Hiroshi Onuki, and Tsuyoshi Takagi. "Improvement of the square-root Vélu's formulas for isogeny-based cryptography". In: *JSIAM Lett.* 15 (2023), pp. 61–64. ISSN: 1883-0609,1883-0617. DOI: 10.14495/jsiaml.15.61. URL: https://doi-org.ezproxy.lib.vt.edu/10.14495/jsiaml.15.61.

[56] Rémy Oudompheng and Giacomo Pope. *A Note on Reimplementing the Castryck-Decru Attack and Lessons Learned for SageMath*. Cryptology ePrint Archive, Paper 2022/1283. 2022. URL: https://eprint.iacr.org/2022/1283.

[57] Aurel Page and Damien Robert. *Introducing Clapoti(s): Evaluating the isogeny class group action in polynomial time*. Cryptology ePrint Archive, Paper 2023/1766. 2023. URL: https://eprint.iacr.org/2023/1766.

[58] Aurel Page and Benjamin Wesolowski. "The Supersingular Endomorphism Ring and One Endomorphism Problems are Equivalent". In: *Advances in Cryptology – EUROCRYPT 2024*. Ed. by Marc Joye and Gregor Leander. Cham: Springer Nature Switzerland, 2024, pp. 388–417. ISBN: 978-3-031-58751-1.

[59] Lorenz Panny. *CSI-FiSh really isn't polynomial-time*. Accessed: 2026-01-09. 2023. URL: https://yx7.cc/blah/2023-04-14.html.

[60] Lorenz Panny, Christophe Petit, and Miha Stopar. *KLaPoTi: An asymptotically efficient isogeny group action from 2-dimensional isogenies*. Cryptology ePrint Archive, Paper 2024/1844. 2024. URL: https://eprint.iacr.org/2024/1844.

[61] J. Pila. "Frobenius maps of abelian varieties and finding roots of unity in finite fields". In: *Math. Comp.* 55.192 (1990), pp. 745–763. ISSN: 0025-5718,1088-6842. DOI: 10.2307/2008445. URL: https://doi.org/10.2307/2008445.

[62] Arnold Pizer. "An algorithm for computing modular forms on $\Gamma_0(N)$". In: *J. Algebra* 64.2 (1980), pp. 340–390. ISSN: 0021-8693. DOI: 10.1016/0021-8693(80)90151-9. URL: https://doi-org.ezproxy.lib.vt.edu/10.1016/0021-8693(80)90151-9.

[63] Fried. Jul. Richelot. In: *Journal für die reine und angewandte Mathematik* 1837.16 (1837), pp. 221–284. DOI: doi:10.1515/crll.1837.16.221. URL: https://doi.org/10.1515/crll.1837.16.221.

[64] Damien Robert. "Breaking SIDH in polynomial time". In: *Advances in cryptology—EUROCRYPT 2023. Part V*. Vol. 14008. Lecture Notes in Comput. Sci. Springer, Cham, 2023, pp. 472–503. ISBN: 978-3-031-30588-7; 978-3-031-30589-4. DOI: 10.1007/978-3-031-30589-4_17.

[65] Damien Robert. "On the Efficient Representation of Isogenies". In: *Number-Theoretic Methods in Cryptology*. Ed. by Andrzej Dabrowski, Josef Pieprzyk, and Jacek Pomykała. Cham: Springer Nature Switzerland, 2025, pp. 3–84. ISBN: 978-3-031-82380-0.

[66] Damien Robert. "Theta functions and cryptographic applications". PhD thesis. L'université Henri Poincaré — Nancy 1, 2010. URL: https://www.math.u-bordeaux.fr/~damienrobert/pro/publications/academic/phd.pdf.

[67] Alexander Rostovtsev and Anton Stolbunov. *PUBLIC-KEY CRYPTOSYSTEM BASED ON ISOGENIES*. Cryptology ePrint Archive, Paper 2006/145. 2006. URL: https://eprint.iacr.org/2006/145.

[68] René Schoof. "Nonsingular plane cubic curves over finite fields". In: *J. Combin. Theory Ser. A* 46.2 (1987), pp. 183–211. ISSN: 0097-3165,1096-0899. DOI: 10.1016/0097-3165(87)90003-3. URL: https://doi-org.ezproxy.lib.vt.edu/10.1016/0097-3165(87)90003-3.

[69] Joseph H. Silverman. *Advanced topics in the arithmetic of elliptic curves*. Vol. 151. Graduate Texts in Mathematics. Springer-Verlag, New York, 1994, pp. xiv+525. ISBN: 0-387-94328-5. DOI: 10.1007/978-1-4612-0851-8. URL: https://doi.org/10.1007/978-1-4612-0851-8.

[70] Joseph H. Silverman. *The arithmetic of elliptic curves*. Second. Vol. 106. Graduate Texts in Mathematics. Springer, Dordrecht, 2009, pp. xx+513. ISBN: 978-0-387-09493-9. DOI: 10.1007/978-0-387-09494-6. URL: https://doi.org/10.1007/978-0-387-09494-6.

[71] Nigel P. Smart. *Cryptography made simple*. Information Security and Cryptography. Springer, Cham, 2016, pp. xii+481. ISBN: 978-3-319-21935-6; 978-3-319-21936-3. DOI: 10.1007/978-3-319-21936-3. URL: https://doi-org.ezproxy.lib.vt.edu/10.1007/978-3-319-21936-3.

[72] Benjamin Smith. "Explicit endomorphisms and correspondences". PhD thesis. University of Sydney, 2005.

[73] Ravi Vakil. *The rising sea—foundations of algebraic geometry*. Princeton University Press, Princeton, NJ, [2025] ©2025, pp. xxiii+662. ISBN: 978-0-691-26866-8; 978-0-691-26867-5; 978-0-691-26868-2.

[74] Jacques Vélu. "Isogénies entre courbes elliptiques". In: *C. R. Acad. Sci. Paris Sér. A-B* 273 (1971), A238–A241.

[75] John Voight. *Quaternion algebras*. Vol. 288. Graduate Texts in Mathematics. Springer, Cham, [2021] ©2021, pp. xxiii+885. ISBN: 978-3-030-56692-0; 978-3-030-56694-4. DOI: 10.1007/978-3-030-56694-4. URL: https://doi-org.ezproxy.lib.vt.edu/10.1007/978-3-030-56694-4.

[76] Marius Vuille. "Computing Cyclic Isogenies between Principally Polarized Abelian Varieties over Finite Fields". PhD thesis. École polytechnique fédéral de Lausanne, 2020. URL: https://infoscience.epfl.ch/server/api/core/bitstreams/afae03ce-cfe5-4aea-97c6-159edcdfddd9/content.

[77] William C. Waterhouse. "Abelian varieties over finite fields". In: *Ann. Sci. École Norm. Sup. (4)* 2 (1969), pp. 521–560. ISSN: 0012-9593. URL: http://www.numdam.org/item?id=ASENS_1969_4_2_4_521_0.

[78] André Weil. "Zum Beweis des Torellischen Satzes". In: *Nachr. Akad. Wiss. Göttingen. Math.-Phys. Kl. IIa.* 1957 (1957), pp. 33–53.

[79] Ryo Yoshizumi, Hiroshi Onuki, Ryo Ohashi, Momonari Kudo, and Koji Nuida. "Efficient Theta-Based Algorithms for Computing (ℓ, ℓ)-Isogenies on Kummer Surfaces for Arbitrary Odd ℓ". In: *Post-Quantum Cryptography*. Ed. by Ruben Niederhagen and Markku-Juhani O. Saarinen. Cham: Springer Nature Switzerland, 2025, pp. 3–37. ISBN: 978-3-031-86602-9.