Isogeny graphs of abelian varieties over finite fields

Jean Kieffer

v1.0 (December 2024)

Abstract

These are lecture notes from a graduate mini-course given at the University of Luxembourg. The first part covers mathematical background on abelian varieties over finite fields and isogenies between them (Tate's theorem, Honda-Tate theory, and work of Waterhouse; I briefly mentioned Dieudonné modules.) The second part is on recent developments in isogeny-based cryptography, namely on the notion of efficient representations of isogenies (after the SIDH attacks and work of Robert) and the recently proved equivalence between hard computational problems in this field (after Wesolowski and Page–Wesolowski.)

Important note. If you notice typos, inconsistencies, or broken English, please let me know. For updated versions, see https://members.loria.fr/JKieffer/.

Introduction

In recent years, isogeny graphs of abelian varieties over finite fields (and in particular isogeny graphs of elliptic curves) have attracted a lot of interest due to the development of isogeny-based cryptography. The whole field relies on a small number of key theorems on isogenies between abelian varieties over finite fields, such as the following.

Theorem 1 (The Deuring correspondence)**.** *Let p be a prime number, and let $B_{p,\infty}$ be the unique quaternion algebra over $\mathbb{Q}$ ramified at p and ∞.*

1. *In any isomorphism class of supersingular over $\overline{\mathbb{F}}_p$, there exists a single minimal supersingular elliptic curve defined over $\mathbb{F}_{p^2}$, up to $\mathbb{F}_{p^2}$-isomorphism. Those elliptic curves form a single isogeny class over $\mathbb{F}_{p^2}$.*

2. *The map*

$$E \mapsto \mathrm{End}(E)$$

between isomorphism classes of minimal supersingular elliptic curves $E/\mathbb{F}_{p^2}$ and maximal orders of $B_{p,\infty}$ up to conjugation, is well-defined. Each maximal order $\mathcal{O} \subset B_{p,\infty}$ has either one or two preimages; it has one preimage $E/\mathbb{F}_{p^2}$ exactly when E arises from the base change of an elliptic curve over $\mathbb{F}_p$. Otherwise, the preimages E and E' of $\mathcal{O}$ are not defined over $\mathbb{F}_p$, and E' is isomorphic to the image of E under the p-power Frobenius map.

3. *Let $E/\mathbb{F}_{p^2}$ be a minimal supersingular elliptic curve, and fix an isomorphism $\eta : \mathcal{O} \simeq \mathrm{End}(E)$ where $\mathcal{O}$ is a maximal order of $B_{p,\infty}$. Then there is a one-to-one correspondence between left $\mathcal{O}$-ideals in $B_{p,\infty}$ and isogenies with domain E; to an ideal I corresponds the isogeny $\phi_I : E \to E'$ whose kernel is*

$$\ker(\phi_I) = \bigcap_{\alpha \in I} \ker \eta(\alpha),$$

whose degree is the reduced norm of I. The endomorphism ring of E', considered as a subring of $\mathcal{O} \otimes \mathbb{Q} = B_{p,\infty}$ via

$$\beta \mapsto \frac{1}{\deg(\phi_I)} \phi_I^\vee \beta \phi_I,$$

is precisely the right order of I. Two ideals I, J are equivalent (i.e. $J = I\lambda$ for some $\lambda \in B_{p,\infty}^\times$) if and only if ϕ_I and ϕ_J have isomorphic codomains.

Theorem 2 (The CM action)**.** *Let A be a simple, ordinary abelian variety of dimension g over a finite field k. Then $F = \mathrm{End}^0(A) = \mathrm{End}(A) \otimes \mathbb{Q}$ is a CM field of degree $2g$, and we can simultaneously identify the rings $\mathrm{End}(B)$, where B runs through the k-isogeny class of A, with subrings of F in a compatible way. Let $R = \mathrm{End}(A)$, and assume further that R is a maximal order or that A is an elliptic curve. Then the subset of this isogeny class consisting of abelian varieties B such that $\mathrm{End}(B) = R$ is a principal homogeneous space under the class group $\mathrm{Cl}(R)$ of R. Invertible R-ideals act as isogenies whose degree is the norm of R, and this class group action covers all isogenies between abelian varieties with $\mathrm{End}(B) = R$.*

The aim of the first part of these notes is to open those black boxes and present some key ideas in their proofs, following the approach of Weil, Tate and Waterhouse [Tat66; Wat69] from the 1960's. We also illustrate the power of this approach through more recent applications on higher-dimensional isogeny graphs. Contrary to the classical approach of Deuring [Deu41] to Theorem 1, Tate's approach requires some knowledge on abelian varieties of higher dimensions and not only elliptic curves; rather than a red flag, we consider this a motivation to delve into this deep and beautiful subject.

The second part of the notes is focused on algorithms and cryptographic applications, and provides an introduction to the recent groundbreaking developments in isogeny-based cryptography that followed the 2022 attacks on SIDH/SIKE [CD23; MMP+23; Rob23]. We present the notion of efficient representation of isogenies after work of Robert [Rob24]. Then, we use the historical example of the Charles–Goren–Lauter hash function [CLG09] to introduce some hard algorithmic problems in isogeny graphs, and present recent results of Page–Wesolowski [PW24] showing that these problems are all equivalent under probabilistic polynomial-time reductions.

Contents

1 Mathematical background

Besides the historical papers, the major references for the theory of abelian varieties are Mumford's book [Mum70], the unfinished book by Edixhoven, van der Geer, and Moonen [EvdGM12] (Edixhoven sadly passed away in 2022), and Milne's course notes [Mil86a]. Over the complex numbers specifically, the book [BL04] is also very useful.

1.1 The geometry of abelian varieties

This section collects facts that hold true for abelian varieties over any field (but not only algebraically closed fields). Nearly all the results can be found in Mumford's book on abelian varieties [Mum70]. This material necessary for the rest of the notes, but wasn't the main focus of the class, so most proofs are not included. We assume previous exposure to algebraic geometry [Har77], including the language of schemes in some places, as well as to elliptic curves [Sil09].

We fix a base field k of characteristic p ($p = 0$ is allowed) and an algebraic closure $\overline{k}$ of k.

1.1.1 Abelian varieties and subvarieties

Definition 1.1.1. An *abelian variety* over a field k is a smooth, projective, irreducible variety over k endowed with a *group law*. The last point means that A is equipped with a multiplication morphism $m : A \times A \to A$, an inverse map $i : A \to A$, and a neutral point $e \in A(k)$ (by definition, this is the same as a morphism $e : \mathrm{Spec}(k) \to A$) such that the usual group axioms hold. In other words, the following diagrams are commutative:

$$\begin{array}{ccc} A \times A \times A & \xrightarrow{(\mathrm{id},m)} & A \times A \\ {\scriptstyle (m,\mathrm{id})}\downarrow & & \downarrow{\scriptstyle m} \\ A \times A & \xrightarrow{m} & A \end{array} \qquad \begin{array}{ccc} A & \xrightarrow{(e,id)} & A \times A \\ {\scriptstyle (id,e)}\downarrow & \searrow^{\mathrm{id}} & \downarrow{\scriptstyle m} \\ A \times A & \xrightarrow{m} & A \end{array} \qquad \begin{array}{ccc} A \times A & \xrightarrow{(\mathrm{id},i)} & A \times A \\ {\scriptstyle (i,\mathrm{id})}\downarrow & \searrow^{e} & \downarrow{\scriptstyle m} \\ A \times A & \xrightarrow{m} & A \end{array}$$

where the diagonal map in the last diagram is the compositum $A \times A \longrightarrow \mathrm{Spec}(k) \xrightarrow{e} A$.

For any k-algebra R, the set $A(R)$ of R-points of A, i.e. the set of morphisms from $\mathrm{Spec}(R)$ to A over $\mathrm{Spec}(k)$, is then equipped with a group structure that is functorial in R. (Conversely, the existence of such functorial group structures would imply the existence of m, i, and e as above by Yoneda's lemma.)

One of the most basic invariants of an abelian variety is its *dimension* (as an algebraic variety), almost always denoted by g. An *elliptic curve* is an abelian variety of dimension 1. We will see in a moment why this definition is the same as the perhaps more usual ones as Weierstrass equations [Sil09, §III.1], or smooth curves of genus 1 [Sil09, §III.3].

A fundamental fact is that abelian varieties are commutative groups: the diagram

$$\begin{array}{ccc} A \times A & \xrightarrow{\text{switch factors}} & A \times A \\ \downarrow{\scriptstyle m} & & \downarrow{\scriptstyle m} \\ A & \xrightarrow{\mathrm{id}} & A \end{array}$$

commutes, and $A(R)$ is a commutative group for each k-algebra R. Because the group law is commutative, we use the symbol $+$ for the map m, $-$ for the map i, and 0_A for e.

Remark 1.1.2. The definition of an abelian variety doesn't have to include smoothness, as it is automatic [Mum70, (i) p. 41]. We could also relax the "projective" hypothesis and write "complete" instead, but it is a fact that all abelian varieties are projective [Mum70, p. 62]. However we can't write off that word altogether: GL_n (for instance) is a group variety over k which is affine, not projective, and is not an abelian variety. (It is also not commutative for $n > 1$.)

Given an abelian variety, one can look at its subvarieties. Some of them are also subgroups:

Definition 1.1.3. An *abelian subvariety* B of A is a subvariety in the usual sense which, when endowed with the induced group law from A, becomes an abelian variety. In particular, B is irreducible and contains 0_A.

An abelian variety A over k is called *simple* if its only abelian subvarieties over k are $\{0_A\}$ and A itself. It is called *absolutely simple* if it is simple when considered as an abelian variety over an algebraic closure $\overline{k}$ of k. In particular, A is then simple.

Elliptic curves are all absolutely simple for dimension reasons.

If A and B are abelian varieties of any dimension, then $A \times B$ is an abelian variety as well: as we will see, taking products is one of the main reasons why the general setting of abelian varieties is often more powerful than that of elliptic curves alone. However $A \times B$ won't be simple unless A or B has dimension zero. For an example of a simple abelian variety that is not absolutely simple, one can consider the Weil restriction of an elliptic curve over k', if k'/k is a finite extension.

1.1.2 Morphisms, endomorphisms, subgroups, and isogenies

Definition 1.1.4. A *morphism* $\phi : A \to B$ between abelian varieties is a morphism of k-varieties respecting the group laws, i.e. satisfying $\phi(0_A) = 0_B$ and $\phi(x + y) = \phi(x) + \phi(y)$ for all points x, y on A valued in any k-algebra R. (This could also be rephrased as a commutative diagram of morphisms.)

Because abelian varieties are commutative groups, $\mathrm{Hom}(A, B)$ is always an abelian group, in other words a $\mathbb{Z}$-module, whose neutral element is the zero map $A \to B$.

Definition 1.1.5. The *endomorphism ring* of A is $\mathrm{End}(A) = \mathrm{Hom}(A, A)$ endowed with the ring structure given by addition and composition.

The ring $\mathrm{End}(A)$ is not commutative in general. We stress that throughout these notes, we only consider morphisms and endomorphisms that are defined over the chosen base field k.

Definition 1.1.6. An *isogeny* $\phi : A \to B$ between two abelian varieties is a morphism such that

1. A and B have the same dimension;
2. $\ker(\phi)$ is a finite subgroup (scheme) of A;
3. ϕ is surjective, i.e. the image of ϕ as a variety is the whole of B.

Actually, any two of these properties imply the third one [Mil86a, Prop. 8.1].

If $\phi : A \to B$ is an isogeny, then the pullback map via ϕ realizes the function field $k(B)$ as a subfield of $k(A)$, and the field extension $k(A)/k(B)$ is finite. The degree of this extension is called the *degree of* ϕ, denoted by $\deg(\phi)$.

If A and B are elliptic curves, then any nonzero morphism $\phi : A \to B$ is an isogeny: indeed the image of ϕ is connected and not $\{0_B\}$, so it has dimension 1, so it is equal to B, and thus (1) and (3) hold. This is a particular case of the following more general situation.

Proposition 1.1.7. *Let $\phi : A \to B$ be a nonzero morphism between simple abelian varieties over k. Then $\dim A = \dim B$ and ϕ is an isogeny.*

Proof. The connected component of 0_A inside $\ker(\phi)$ and the image of ϕ are abelian subvarieties of A and B respectively. Because $\phi \neq 0$ and A, B are assumed to be simple, the former subvariety must be $\{0_A\}$, and the latter B itself. Thus (2) and (3) in Definition 1.1.6 hold. □

An important example of endomorphisms of A are the *multiplications by n*, where $n \in \mathbb{Z}$. They are denoted by $[n]_A$ and are defined as follows:

- If $n \geq 0$, then $[n]_A(x) = x + \cdots + x$ (n times);
- If $n < 0$, then $[n]_A(x) = [-n]_A(-x)$.

The kernel of $[n]_A$ is called the *n-torsion subgroup* (scheme) of A, and is denoted by $A[n]$. The set of points of this subgroup scheme over $\overline{k}$ is denoted, as usual, by $A[n](\overline{k})$.

Proposition 1.1.8 ([Mum70, Prop. p. 64])**.** *Let A be an abelian variety of dimension g over k.*

1. *For any nonzero $n \in \mathbb{Z}$, the endomorphism $[n]_A$ is an isogeny of degree n^{2g}.*
2. *If n is not divisible by p, then $A[n](\overline{k}) \simeq (\mathbb{Z}/n\mathbb{Z})^{2g}$ as an abstract group.*

Proposition 1.1.9. *For every abelian varieties A, B, the group $\mathrm{Hom}(A, B)$ is torsion-free.*

Proof. Let $\phi : A \to B$ be a morphism and let $n \geq 1$ such that $n\phi = 0$. Then $\phi \circ [n]_A = 0$. However $[n]_A$ is surjective by Proposition 1.1.8, so ϕ is identically zero. □

As a consequence, the multiplication maps $[n]_A$ for $n \in \mathbb{Z}$ always form a copy of $\mathbb{Z}$ inside the endomorphism ring $\mathrm{End}(A)$.

For $n = p$ in positive characteristic, item (2) in Proposition 1.1.8 fails: as an abstract group, $A[p](\overline{k})$ is isomorphic to $(\mathbb{Z}/p\mathbb{Z})^r$ where r is an integer between 0 and g called the *p-rank* of A [Mum70, p. 147]. The situation becomes clearer if we think about finite group schemes. For every $n \neq 0$, the n-torsion subgroup $A[n]$ is a finite subgroup scheme of A of degree n^{2g}. If n is not divisible by p, then $A[n]$ is étale, so $A[n](\overline{k})$ is a group of size n^{2g}; on the other hand, $A[p]$ is not étale, and looking at its $\overline{k}$-points of $A[p]$ only reveals its étale part.

More generally, for any isogeny $\phi : A \to B$, the kernel of ϕ is a subgroup scheme of A of rank $\deg(\phi)$. If $\deg(\phi)$ is coprime to p, then ϕ is *separable*, i.e. $\#\ker(\phi)(\overline{k}) = \deg(\phi)$. These facts are part of a general correspondence between finite subgroups of A and isogenies with domain A. We first phrase the result away from the characteristic using groups of $\overline{k}$-points, then the more general result using the language of group schemes.

Proposition 1.1.10. *Let A be an abelian variety over k, and assume that k is a perfect field of characteristic p. Then the map $\phi \mapsto \ker(\phi)(\overline{k})$ realizes a one-to-one correspondence between*

- *Isogenies with $\phi : A \to B$ (where B is another abelian variety over k) such that $p \nmid \deg(\phi)$, up to postcomposition by an isomorphism $B \simeq B'$, and*

- *Finite subgroups $K \subset A(\overline{k})$ such that $p \nmid \#K$ and that are globally invariant under the absolute Galois group $G_k = \mathrm{Gal}(\overline{k}/k)$.*

The inverse map sends K to the natural quotient map $A \to A/K$, of degree $\#K$.

This proposition is roughly [Mum70, Thm. 4 p. 72], even though Mumford works over an algebraically closed field there. Proposition 1.1.10 more generally holds for separable isogenies, even if their degree is a multiple of p.

Proposition 1.1.11 ([Mum70, Cor. 1 p. 118])**.** *Let A be an abelian variety over a field k of characteristic p. Then the map $\phi \mapsto \ker(\phi)$ realizes a one-to-one correspondence between:*

- *Isogenies $\phi : A \to B$ (where B is another abelian variety over k) up to postcomposition by an isomorphism $B \simeq B'$, and*
- *Finite subgroup schemes K of A defined over k.*

The inverse map sends K to the natural quotient map $A \to A/K$ whose degree is the rank of k, i.e. the dimension of the coordinate ring of K as a k-algebra.

Considering all isogenies as quotient maps by certain subgroups is a powerful point of view. For instance, it has the following results as an immediate consequence:

Proposition 1.1.12 (Isogeny factorization)**.** *Let $\phi : A \to B$ and $\psi : A \to C$ be two isogenies, and assume that ψ is identically zero on $\ker(\phi)$, in other words $\ker(\phi) \subset \ker(\psi)$ as group schemes. Then there exists an isogeny $\psi' : B \to C$ such that $\psi = \psi' \circ \phi$.*

Proof. Since ϕ is the natural quotient map $A \to A/\ker(\phi)$, it is universal among morphisms vanishing on $\ker(\phi)$, hence the existence of ψ'. □

For every finite subgroup scheme K of A, there exists an integer $n \geq 1$ such that $K \subset A[n]$ (for instance, one can take n to be the rank of K). As a consequence, we have:

Proposition 1.1.13 (Isogenies are almost invertible)**.** *Let $\phi : A \to B$ be an isogeny. Then there exists an integer $n \geq 1$ and an isogeny $\psi : B \to A$ such that $\psi \circ \phi = [n]_A$ and $\phi \circ \psi = [n]_B$.*

Proof. Let $n \geq 1$ such that $\ker(\phi) \subset A[n]$. By Proposition 1.1.12, there exists an isogeny $\psi : B \to A$ such that $\psi \circ \phi = [n]_A$. Then $\phi(\psi(\phi(x))) = \phi(nx) = [n]_B(\phi(x))$ for every point x on A. Since ϕ is surjective as a rational map, we must have $\phi \circ \psi = [n]_B$ as well. □

In particular, being isogenous is an equivalence relation on abelian varieties defined over k. The *isogeny class* of an abelian variety A is, by definition, the set of isomorphism classes of abelian varieties B/k that are isogenous to A.

The proofs of Propositions 1.1.12 and 1.1.13 given above rely on Proposition 1.1.11 and use the language of group schemes, but one could equally write down versions using only groups of $\overline{k}$-points and Proposition 1.1.10 provided that we only consider separable isogenies, for instance isogenies whose degree is not divisible by the characteristic of k.

1.1.3 Duals, polarizations, and pairings

Constructing the *dual* of any abelian variety A is an important step in the theory, as it allows us to understand how line bundles on A behave. In particular, ample line bundles giving rise to projective embeddings of A can be used to construct *polarizations*, which are isogenies between A and its dual. The necessary emphasis on polarizations is one of the main differences between elliptic curves and higher-dimensional abelian varieties: as we will see, all elliptic curves carry a canonical principal polarization, so we can essentially forget about it.

For a general review of line bundles on smooth projective varieties, the notion of ample and very ample line bundles, their links with projective embeddings, and the definition of the Picard groups $\mathrm{Pic}(X)$ and $\mathrm{Pic}^0(X)$, we refer to [Har77, §II.5-6]. In words, $\mathrm{Pic}(X)$ is the group of (isomorphism classes of) line bundles on A (which correspond to divisors up to linear equivalence), and $\mathrm{Pic}^0(X)$ is the subgroup of line bundles that are algebraically equivalent to the trivial line bundle, i.e. the connected component of the trivial line bundle in $\mathrm{Pic}(X)$.

Theorem 1.1.14. *Let A be an abelian variety over k.*

1. $\mathrm{Pic}^0(A)$ *has a natural structure of an abelian variety defined over k, called the* dual *of A and denoted by $A^\vee$, of the same dimension as A.*
2. *For any line bundle $\mathcal{L}$ on A, the map*
$$\begin{array}{rccc} \phi_{\mathcal{L}} : & A & \to & A^\vee \\ & x & \mapsto & t_x^*\mathcal{L} \otimes \mathcal{L}^{-1}, \end{array}$$
where $t_x : y \mapsto x + y$ denotes the translation by x on A, is a morphism between A and $A^\vee$.
3. *There exists a canonical line bundle $\mathcal{P}$ on $A \times A^\vee$, the* Poincaré bundle, *with the property that $(id, \phi_{\mathcal{L}})^*\mathcal{P}$ is algebraically equivalent to $\mathcal{L}^{\otimes 2}$ for every $\mathcal{L} \in \mathrm{Pic}(A)$.*
4. *If $\mathcal{L}$ is ample, then $\phi_{\mathcal{L}} : A \to A^\vee$ is an isogeny whose degree is a perfect square.*

References for the proof. 1. [Mum70, Thm. p. 125].

2. The fact that ϕ_L is a morphism of groups is the theorem of the square [Mum70, Cor. 5 p. 131].
3. The existence of $\mathcal{P}$ is provided by [Mum70, Thm. p. 125], while the second statement appears in the proof of [Mum70, Thm. 2 p. 188].
4. If $\mathcal{L}$ is ample, then $\phi_{\mathcal{L}}$ is an isogeny by [Mum70, Thm. 1 p. 77]. The degree of $\phi_{\mathcal{L}}$ is the square of the Euler characteristic of $\mathcal{L}$ by the Riemann–Roch theorem [Mum70, p. 150]. □

Definition 1.1.15. A *polarization* on an abelian variety A is an isogeny $\lambda : A \to A^\vee$ such that $(\mathrm{id}, \lambda)^*\mathcal{P}$ is an ample line bundle on A. An equivalent definition would be: an isogeny $\lambda : A \to A^\vee$ is a polarization if and only if there exists a finite field extension k'/k and an ample line bundle $\mathcal{L}$ on A defined over k' such that $\lambda = \phi_{\mathcal{L}}$.

The dual of $A^\vee$ is canonically isomorphic to A again [Mum70, Cor. p. 132]. If $f : A \to B$ is any morphism, then the pullback of line bundles via f yields the *dual morphism* $f^\vee : B^\vee \to A^\vee$; it is an isogeny if and only if f is an isogeny, and in that case has the same degree [Mum70, Thm. 1 p. 143]. We even have the stronger property that $\ker(f^\vee)$ is the dual of $\ker(f)$ as a group scheme.

If we consider an ample line bundle $\mathcal{L}$ on A, then besides the fact that $\phi_{\mathcal{L}}$ is a polarization, a key fact is that $\mathcal{L}^{\otimes 3}$ is always very ample.

Theorem 1.1.16 (Lefschetz's theorem)**.** *Let $\mathcal{L}$ be an ample line bundle on an abelian variety A of dimension g, and let $d \geq 1$ be the integer such that $\deg \phi_{\mathcal{L}} = d^2$. Then for every integer $n \geq 3$, the line bundle $\mathcal{L}^{\otimes n}$ is very ample, and realizes A as a subvariety of $\mathbb{P}^N$ where $N = n^g d - 1$, of degree $g! n^g d$.*

Proof. Mumford proves that $\mathcal{L}^{\otimes 3}$ is very ample [Mum70, p. 163], and the same proof should work for higher n. The dimension of the space of sections of $\mathcal{L}^{\otimes n}$ is $\chi(\mathcal{L}^n) = n^g \chi(\mathcal{L}) = n^g d$ by the Riemann–Roch theorem [Mum70, p. 150]. If D is the divisor corresponding to $\mathcal{L}^{\otimes n}$, the g-fold self-intersection number of D is the degree of the image variety, and is $g! \chi(\mathcal{L}^{\otimes n})$ by the same Riemann–Roch theorem. □

Another important element on abelian varieties, closely related to ample line bundles and polarizations, is the existence of Weil pairings on certain torsion subgroups. For each $n \in \mathbb{Z}$, the dual of $[n]_A$ is $[n]_{A^\vee}$ [Mum70, (iii) p. 75], so torsion subgroups of the dual abelian variety $A^\vee$ are the duals of torsion subgroups of A. In other words, for any $n \geq 1$, there exists a canonical isomorphism

$$A^\vee[n] \simeq \mathrm{Hom}(A[n], \mathbb{G}_m)$$

(where Hom is taken in the category of k-group schemes). In other words, there exists a canonical nondegenerate pairing

$$e_n : A[n] \times A^\vee[n] \to \mu_n$$

where μ_n denotes the k-group schemes of nth roots of unity. This pairing is constructed as follows, at least when n is prime to the characteristic p of k. Let $\mathcal{L}$ be a point of $A^\vee[n](\overline{k})$, i.e. a line bundle on A such that $\mathcal{L}^{\otimes n}$ is linearly equivalent to the trivial bundle. Then the pullback line bundle $[n]_A^* \mathcal{L}$ is also trivial. If D is a divisor corresponding to $\mathcal{L}$ Let f and g be functions on A with divisors nD and $[n]_A^{-1} D$ respectively. Then for some constant $\alpha \in k^\times$, we have $f(nx) = \alpha g(x)^n$ for all $x \in A$. Let now $u \in A[n](\overline{k})$. Then the map $x \mapsto g(x)/g(x+u)$ is valued in the nth roots of unity, so must be constant on A; we declare its value to be $e_n(u, \mathcal{L})$ [Mum70, Lemma p. 184].

Proposition 1.1.17. *Assume that A is endowed with a polarization $\lambda : A \to A^\vee$, and let $n \geq 1$ prime to p. Then the pairing on $A[n] \times A[n]$ given by the formula $(x, x') \mapsto e_n(x, \lambda(x'))$ is alternating. If n is prime to the degree of λ, then this pairing is also nondegenerate.*

Proof. The fact that this pairing is alternating is [Mum70, Thm. 1 p. 186]. If n is prime to the degree of λ, then $\lambda : A[n] \to A^\vee[n]$ is bijective; because e_n is nondegenerate, the pairing $e_n(\cdot, \lambda(\cdot))$ is also nondegenerate. □

By an abuse of notation, we also denote the pairing in Proposition 1.1.17 by e_n, even though it depends on the choice of polarization on A. It is called the *Weil pairing*.

The main interest of these pairings is that they allow us to understand when isogenies are compatible with polarizations. Starting from an abelian variety A endowed with a polarization of some degree d, one can we make isogenies $A \to B$ such that B also carries a natural polarization of degree d by taking quotients by *maximal isotropic* torsion subgroups for the Weil pairing we just constructed. Recall that if R is any ring and V is an R-module endowed with an R-linear alternating form e, then a submodule W of V is called *isotropic* if e vanishes identically on $W \times W$. Further, W is called *maximal isotropic* in V if it is isotropic and the only isotropic submodule of V contaning W is W itself.

Proposition 1.1.18. *Let A be an abelian variety over k equipped with a polarization $\lambda : A \to A^\vee$ of degree d, and let $n \geq 1$ be coprime to d andp. Let $K \subset A[n](\overline{k})$ be a maximal isotropic subgroup for the Weil pairing, AND Let $\phi : A \to B = A/K$ be the quotient isogeny. Then there exists a unique polarization λ' on $B = A/K$ such that $\phi^\vee \circ \lambda' \circ \phi = n\lambda$.*

If $\lambda = \phi_{\mathcal{L}}$ and $\lambda' = \phi_{\mathcal{L}'}$, then the last equality amonts to saying that $\mathcal{L}^{\otimes n}$ and $\phi^*\mathcal{L}'$ are algebraically equivalent: see [Mum70, Proofs of Thm. 1 p. 143 and Cor. 2 p. 178]. Of course, the quotient isogeny $A \to A/K$ is k-rational if and only if the subgroup K itself is defined over k.

Remark 1.1.19. There exists a converse to Proposition 1.1.18, at least in the case where λ and λ' are principal polarizations: any isogeny $\phi : A \to B$ between principally polarized abelian varieties arises as a quotient by some maximal isotropic subgroup. However, the kernel doesn't have to be maximal isotropic in a torsion subgroup $A[n]$ for some n. In fact, the composition $\alpha = \lambda^{-1} \circ \phi^\vee \circ \lambda' \circ \phi$ defines an endomorphism of A, and the α-torsion subgroup $A[\alpha]$ carries a nondegenerate alternating pairing for which $\ker(\phi) \subset A[\alpha]$ is maximal isotropic [Mum70, (3) p. 190 and p. 231–233].

1.1.4 Examples: elliptic curves and Jacobians

An important family of abelian varieties (often the most explicitly accessible ones) arises as the Jacobians of algebraic curves: if $\mathcal{C}$ is a smooth, projective, irreducible curve $\mathcal{C}$ of genus g over a field k, the Picard group $\mathrm{Pic}^0(\mathcal{C})$ is an abelian variety of dimension g, called the *Jacobian* $\mathrm{Jac}(\mathcal{C})$ of $\mathcal{C}$ [Mil86b, Thm. 1.1].

Example 1.1.20 (The canonical principal polarization on elliptic curves)**.** Let E be an elliptic curve as in [Sil09, §III.3], i.e. a curve as above with $g = 1$ and equipped with a marked point 0_E. Let us investigate the structure of $\mathrm{Pic}^0(E)$. Line bundles on E up to isomorphism are in one-to-one correspondence with divisors on E up to linear equivalence [Har77, §II.6]. Recall that divisors on E are formal linear combinations of points of the form

$$D = \sum_{P \in E} n_P (P)$$

where $n_P \in \mathbb{Z}$ and $n_P = 0$ for almost all P, and that two divisors are called linearly equivalent if their difference is the divisor of a rational function (the formal sum of its zeroes counted with multiplicities, minus the formal sum of its poles.) The degree of such a D is by definition the sum of the coefficients n_P; divisors of functions have degree zero, so linearly equivalent divisors have the same degree.

It's easy to see that $\mathrm{Pic}^0(E)$, the set of line bundles algebraically equivalent to zero, corresponds to the set of divisors of degree zero on E: indeed, such a divisor can be algebraically transformed into the zero divisor by moving all the points P in the support to some fixed point P_0 (for instance $P_0 = 0_E$.) On the other hand, divisors algebraically equivalent to zero have degree zero because the degree is valued in the discrete set $\mathbb{Z}$.

By the Riemann–Roch theorem [Sil09, §III.3, Prop. 3.4], every degree zero divisor on E is linearly equivalent to $(-P) - (0_E)$ for a unique point P on E. The map

$$\begin{array}{rccc} \lambda : & E & \mapsto & \mathrm{Pic}^0(E) \\ & P & \mapsto & \text{the line bundle corresponding to } (0_E) - (P) \end{array}$$

is therefore a (canonical) isomorphism. In particular, E is isomorphic to its own Jacobian, and admits a group law (obtained by transporting the group law on $\mathrm{Pic}^0(E)$ under the isomorphism) with neutral element 0_E.

On the other hand, the degree one divisor (0_E) gives rise to an ample line $\mathcal{L}$ bundle on E, and one can check that $\lambda = \phi_{\mathcal{L}}$ [Con, Ex. 2.5]. Therefore, λ is a principal polarization, and elliptic curves are canonically isomorphic to their own Jacobians and their own duals.

One can show directly using Riemann–Roch (instead of Lefschetz's theorem 1.1.16) that the divisor $3(0_E)$ on E defines a very ample line bundle. One can choose sections x, y, z of this line bundle such that $(x : y : z) : E \to \mathbb{P}^2$ is an isomorphism between E and a Weierstrass equation, the point 0_E being sent to the point at infinity. (In Lefschetz's theorem, we have $g = 1, d = 1, n = 3$, so elliptic curves are indeed embedded by $3(0_E)$ as cubics in $\mathbb{P}^2$.) This provides the link with the definition in [Sil09, §III.1].

Finally, using λ as the principal polarization, one can check that the Weil pairing defined in §1.1.3 is the usual Weil pairing on the n-torsion points defined in [Sil09, §III.8].

Example 1.1.21 (The canonical principal polarization on Jacobians)**.** More generally, let $\mathcal{C}$ be a smooth curve of genus $g \geq 1$ over k. For simplicity, assume that $\mathcal{C}$ admits a k-point P_0. Then the map

$$\begin{array}{rccc} \eta_0 : & \mathcal{C} & \to & \mathrm{Jac}(\mathcal{C}) = \mathrm{Pic}^0(\mathcal{C}) \\ & Q & \mapsto & \text{the line bundle corresponding to } (P_0) - (Q) \end{array}$$

is an isomorphism between $\mathcal{C}$ and its image in $\mathrm{Jac}(\mathcal{C})$ [Mil86b, Prop. 2.3]. One can use η_0 to make a map

$$\begin{array}{rccc} \zeta : & \mathrm{Sym}^{(g-1)}(\mathcal{C}) & \to & \mathrm{Jac}(\mathcal{C}) \\ & (Q_1, \ldots, Q_{g-1}) & \mapsto & \eta_0(Q_1) + \cdots + \eta_0(Q_{g-1}). \end{array}$$

Here $\mathrm{Sym}^{(g-1)}(\mathcal{C})$, the $g-1$st symmetric power of $\mathcal{C}$, is a smooth projective variety [Mil86b, Prop. 3.1 and 3.2] and is birational to its image under ζ inside $\mathrm{Jac}(\mathcal{C})$ [Mil86b, Thm. 5.1]. This image defines a divisor called the *theta divisor*; up to algebraic equivalence, the theta divisor is independent of the choice of P_0.

It turns out that the line bundle $\mathcal{L}$ corresponding to the theta divisor is ample on $\mathrm{Jac}(\mathcal{C})$, and that the associated map $\phi_{\mathcal{L}}$ defines a principal polarization [Mil86b, Thm. 6.6]. Thus, Jacobians of curves admit a canonical principal polarization, generalizing Example 1.1.20. This polarization is independent of the choice of base point P_0, so exists as a k-rational map by descent theory even when $\mathcal{C}$ doesn't have rational points.

An important fact is that in dimension $g \leq 3$, almost all principally polarized abelian varieties are Jacobians.

Theorem 1.1.22 ([OU73])**.** *Let k be any field, and let A/k be a principally polarized abelian variety over k of dimension $g \leq 3$. Then either A is a Jacobian, i.e. there exists a smooth, projective, genus g curve $\mathcal{C}$ such that $\mathrm{Jac}(\mathcal{C}) \simeq A$ as a p.p.a.v. (perhaps defined over a finite extension of k), or A is a product of lower-dimensional p.p.a.v.'s endowed with the product polarization.*

Note that smooth curves of genus 2 are always hyperelliptic, while curves of genus 3 are either hyperelliptic or plane quartic curves. The two possibilities in Theorem 1.1.22 are mutually exclusive:

Proposition 1.1.23. *Let $\mathcal{C}$ be a smooth, projective genus g curve over any field k. Then its Jacobian is indecomposable as a p.p.a.v., in other words it is not isomorphic to a product of positive-dimensional p.p.a.v.'s endowed with the product polarization.*

Proof. On a Jacobian, the divisor associated to the polarization as constructed in Example 1.1.21 is irreducible. This would not be the case if $\mathrm{Jac}(\mathcal{C})$ was a product of two p.p.a.v.'s, but I'm not sure how that argument exactly works. See for instance [Bea13, Rem. 3.10]. □

1.1.5 Endomorphism algebras

As a rule of thumb, when studying an abelian variety A, it is desirable to know its endomorphism ring. A first step to determine $\mathrm{End}(A)$ is to understand its endomorphism algebra, the $\mathbb{Q}$-algebra one obtains from the endomorphism rings by inverting the multiplication-by-n isogenies.

Definition 1.1.24. Let A be an abelian variety. The *endomorphism algebra* of A is $\mathrm{End}^0(A) = \mathrm{End}(A) \otimes_{\mathbb{Z}} \mathbb{Q}$.

The algebra $\mathrm{End}^0(A)$ contains $\mathrm{End}(A)$ as a subring by Proposition 1.1.9. A major feature of endomorphism algebras is that they are isomorphic for isogenous abelian varieties:

Proposition 1.1.25. *Let $\phi : A \to B$ be an isogeny, and choose an integer $n \geq 1$ and an isogeny $\psi : B \to A$ such that $\psi \circ \phi = [n]_A$ and $\phi \circ \psi = [n]_B$ as in Proposition 1.1.13. Then the map*

$$\begin{array}{rccc} \eta : & \mathrm{End}^0(B) & \to & \mathrm{End}^0(A) \\ & \alpha & \mapsto & \frac{1}{n}\psi \circ \alpha \circ \phi \end{array}$$

is independent of the choices of ψ and n and is an isomorphism of $\mathbb{Q}$-algebras.

Proof. Let (ψ', n') be another choice; we can assume that $n \geq 1$ has been chosen minimal such that $\ker(\phi) \subset A[n]$ by Proposition 1.1.12. We have $\ker(\phi) \subset A[n']$ and $A[n'] \cap A[n] = A[n \wedge n']$, so $\ker(\phi) \subset A[n \wedge n']$. Because n is minimal, this shows n divides n'. Then $\psi' = (n'/n)\psi$, so (ψ, n) and (ψ', n') yield the same η. The map η is obviously a morphism of rings, and its inverse is $\beta \mapsto \frac{1}{n}\phi \circ \beta \circ \psi$, so it is an isomorphism. □

Proposition 1.1.26. *Let A be a simple abelian variety over k. Then $\mathrm{End}^0(A)$ is a* division *$\mathbb{Q}$-algebra, meaning that any nonzero element has an inverse. In particular the center of $\mathrm{End}^0(A)$ is a field of characteristic zero.*

Proof. By Proposition 1.1.7, any nonzero endomorphism ϕ of A is an isogeny. By Proposition 1.1.13, there exists $n \geq 1$ and an endomorphism $\psi : A \to A$ such that $\psi \circ \phi = \phi \circ \psi = [n]_A$; then ψ/n is the desired inverse of ϕ. □

Later on, we will prove that $\mathrm{End}^0(A)$ is always a finite-dimensional $\mathbb{Q}$-algebra; in particular, its center is a number field if A is simple.

When A is non-simple, we can still describe $\mathrm{End}^0(A)$ in terms of division $\mathbb{Q}$-algebras.

Theorem 1.1.27 (Poincaré reducibility; [Mum70, Cor. 1 p. 174])**.** *Let A be an abelian variety over k. Then there exist simple non-isogenous abelian varieties $A_1, \ldots, A_r$ over k and integers $n_1, \ldots, n_r \geq 1$ such that A is isogenous to $A_1^{n_1} \times \cdots \times A_n^{n_r}$. This decomposition is unique up to replacing each A_i by an isogenous abelian variety and permuting them. We have*

$$\mathrm{End}^0(A) \simeq \prod_{i=1}^{r} \mathrm{Mat}_{n_i \times n_i}(D_i)$$

where $D_i = \mathrm{End}^0(A_i)$ is a division $\mathbb{Q}$-algebra for each i.

In particular, we see that $\mathrm{End}^0(A)$ is division if and only if A is simple.

Example 1.1.28. Let E_1, E_2 be isogenous elliptic curves over k, and assume $\mathrm{End}^0(E_1) = \mathbb{Q}$. Then the endomorphism algebra of $E_1 \times E_2$ is $\mathrm{Mat}_{2\times 2}(\mathbb{Q})$, but the endomorphism *ring* of $E_1 \times E_2$ is the full $\mathrm{Mat}_{2\times 2}(\mathbb{Z})$ only when E_1 and E_2 are actually isomorphic.

On the endomorphism ring $\mathrm{End}(A)$ of an abelian variety A, one can consider the degree map, which is well-defined once we declare the degree of α to be zero if α is not an isogeny, and valued in $\mathbb{Z}_{\geq 0}$. It turns out that the degree map behaves like a homogeneous polynomial form degree $2g$ where $g = \dim A$ [Mum70, Thm. 2 p. 174]. For instance, it is multiplicative, and we have $\deg(n\alpha) = n^{2g}\deg(\alpha)$ for all $\alpha \in \mathrm{End}(A)$ and $n \in \mathbb{Z}$. Using homogeneity, we can extend the degree map deg to the whole algebra $\mathrm{End}^0(A)$, with values in $\mathbb{Q}$.

Theorem 1.1.29 ([Mum70, Thm. 4 p. 180])**.** *Let $\alpha \in \mathrm{End}(A)$. Then there exists a unique monic polynomial $P_\alpha \in \mathbb{Z}[X]$ of degree $2g$ such that $P_\alpha(n) = \deg([n]_A - \alpha)$ for every $n \in \mathbb{Z}$. We also have $P_\alpha(\alpha) = 0$, in particular $P_\alpha = X^{2g}$ if and only if $\alpha = 0$.*

We call P_α the *characteristic polynomial* of α; the *trace* $\mathrm{Tr}(\alpha)$ of α is defined as the usual coefficient of P_α, i.e.

$$P_\alpha = X^{2g} - \mathrm{Tr}(\alpha)X^{2g-1} + \cdots + \deg(\alpha)X^0.$$

The proof of Theorem 1.1.29 actually uses Tate modules as introduced in the next section.

1.2 Tate's isogeny theorem

As a motivation for this subsection, consider two abelian varieties A, B of respective dimensions $g, g' \geq 1$ over the field $k = \mathbb{C}$. Then there exist full rank, discrete $\mathbb{Z}$-lattices $\Lambda \subset \mathbb{C}^g$ and $\Lambda' \subset \mathbb{C}^{g'}$ such that $A \simeq \mathbb{C}^g/\Lambda$ and $B \simeq \mathbb{C}^{g'}/\Lambda'$ [Mum70, (2) p. 2]. It is then very easy to describe what $\mathrm{Hom}(A, B)$ is: as in [BL04, Prop. 1.2.1], we have

$$\mathrm{Hom}(A, B) = \{\alpha \in \mathrm{Mat}_{g'\times g}(\mathbb{C}) : \alpha\Lambda \subset \Lambda'\}.$$

By writing down how such a matrix α acts on $\mathbb{Z}$-bases of Λ and Λ', we can realize $\mathrm{Hom}(A, B)$ as a sub-$\mathbb{Z}$-module of $\mathrm{Mat}_{2g'\times 2g}(\mathbb{Z})$. In particular, $\mathrm{Hom}(A, B)$ is a free $\mathbb{Z}$-module of finite rank $r \leq 4gg'$.

The aim of this section is to derive similar conclusions from Hom modules between abelian varieties over any fields. It turns out that the *Tate modules*, constructed from torsion subgroups of the abelian varieties, are suitable replacements for the period lattices over $\mathbb{C}$. Over finite fields in particular, Tate's isogeny theorem (Theorem 1.2.11 below) asserts the existence of a bijection between $\mathrm{Hom}(A, B) \otimes \mathbb{Z}_\ell$, for all primes ℓ away from the characteristic, and Hom-sets between Tate modules.

If k has positive characteristic p, an analogue of this theory at $\ell = p$ exists and involves the *Dieudonné modules* attached to p-power torsion subgroups of the abelian varieties (seen as group schemes.) Even though Dieudonné modules are important for the theory, we will only briefly mention this more technical topic.

1.2.1 Tate modules

If A is an abelian variety of dimension g over $\mathbb{C}$ with period lattice $\Lambda \subset \mathbb{C}^g$, and if ℓ is a prime, then we naturally have

$$A[\ell] \simeq (\tfrac{1}{\ell}\Lambda)/\Lambda, \quad A[\ell^2] \simeq (\tfrac{1}{\ell}{}^2\Lambda)/\Lambda, \quad \text{etc.}$$

Consequently, we have a commutative diagram

$$\begin{array}{ccccccccc} \cdots & \xrightarrow{[\ell]_A} & A[\ell^3] & \xrightarrow{[\ell]_A} & A[\ell^2] & \xrightarrow{[\ell]_A} & A[\ell] & \xrightarrow{[\ell]_A} & \{0_A\} \\ & & \downarrow{\scriptstyle \times \ell^3} & & \downarrow{\scriptstyle \times \ell^2} & & \downarrow{\scriptstyle \times \ell} & & \downarrow \\ \cdots & \xrightarrow{\mathrm{id}} & \Lambda/\ell^3\Lambda & \xrightarrow{\mathrm{id}} & \Lambda/\ell^2\Lambda & \xrightarrow{\mathrm{id}} & \Lambda/\ell\Lambda & \longrightarrow & \{0\} \end{array}$$

Consequently, $\Lambda \otimes \mathbb{Z}_\ell$, which is the inverse limit of the quotients $\Lambda/\ell^n\Lambda$ as $n \to \infty$, is isomorphic to the inverse limit of the ℓ^n-torsion subgroups of A under the multiplication-by-ℓ maps. The crucial observation is that the latter inverse limit makes sense even outside of characteristic zero.

Definition 1.2.1. Let A be an abelian variety over any field k, and let ℓ be a prime distinct from the characteristic of k. The ℓth *Tate module* of A is

$$T_\ell(A) = \varprojlim_{n \to \infty} A[\ell^n]$$

where the inverse limit is taken with respect to the multiplication-by-ℓ maps $A[\ell^{n+1}] \to A[\ell^n]$. (Here, and in the rest of these notes, we allow ourselves to write $A[n]$ instead of $A[n](\overline{k})$ when n is prime to the characteristic: the group scheme $A[n]$ is étale so can be identified with the group of its $\overline{k}$-points.)

Since each torsion subgroup $A[\ell^n]$ is a $(\mathbb{Z}/\ell^n\mathbb{Z})$-module, and those module structures are compatible under the multiplication-by-ℓ maps, the Tate module $T_\ell(A)$ is naturally a $\mathbb{Z}_\ell$-module.

Since each torsion subgroup $A[\ell^n]$ carries an $(\mathbb{Z}/\ell^n\mathbb{Z})$-linear action of the absolute Galois group $G_k = \mathrm{Gal}(\overline{k}/k)$, the Tate module $T_\ell(A)$ also carries a natural $\mathbb{Z}_\ell$-linear action of G_k.

Proposition 1.2.2. *The Tate module $T_\ell(A)$ is a free $\mathbb{Z}_\ell$-module of rank $2g$. In other words $T_\ell(A)$ is (noncanonically) isomorphic to $\mathbb{Z}_\ell^{2g}$ as a $\mathbb{Z}_\ell$-module.*

Proof. By Proposition 1.1.8, each torsion subgroup $A[\ell^n]$ is isomorphic to $(\mathbb{Z}/\ell^n\mathbb{Z})^{2g}$. □

Recall that $\mathbb{Z}_\ell$ can be endowed with a natural topology, the *ℓ-adic topology*, in which the sets $\{\ell^n\mathbb{Z}_\ell\}_{n \geq 0}$ form a basis of neighborhoods of zero (both open and closed) and for which $\mathbb{Z}_\ell$ is a topological abelian group. Via an isomorphism $T_\ell(A) \simeq \mathbb{Z}_\ell^{2g}$, the Tate module $T_\ell(A)$ also inherits an ℓ-adic topology (which is independent of the chosen isomorphism.)

Sometimes it is useful to consider $T_\ell(A)$ as a lattice in a certain vector space. Recall that $\mathbb{Q}_\ell = \mathbb{Z}_\ell[1/\ell]$ is a field, called the field of ℓ-adic numbers; the ℓ-adic valuation on $\mathbb{Q}_\ell^\times$ is valued in $\mathbb{Z}$.

Definition 1.2.3. For each prime ℓ as above, we define $V_\ell(A) = T_\ell(A) \otimes_{\mathbb{Z}_\ell} \mathbb{Q}_\ell$. It is a $2g$-dimensional $\mathbb{Q}_\ell$-vector space and is also endowed with the ℓ-adic topology.

The Tate modules of an abelian variety A are constructed from ℓ-power torsion subgroups of A, and conversely one can describe the torsion subgroups of A (hence isogenies with domain A, thanks to Proposition 1.1.10) in terms of its Tate modules. For every $n \geq 0$, we have a canonical isomorphism

$$T_\ell(A)/\ell^n T_\ell(A) \simeq A[\ell^n]$$

If we view $T_\ell(A)$ as a lattice inside the $\mathbb{Q}_\ell$-vector space $V_\ell(A)$, we can also write

$$T_\ell(A)/\ell^n T_\ell(A) \simeq \ell^{-n} T_\ell(A)/T_\ell(A)$$

via multiplication by ℓ^n. If we write

$$A[\ell^\infty] = \bigcup_{n \geq 0} A[\ell^n],$$

noting that $V_\ell(A)$ is the reunion of all lattices $\ell^{-n} T_\ell(A)$ for $n \geq 0$, we obtain a bijection

$$V_\ell(A)/T_\ell(A) \simeq A[\ell^\infty].$$

For a finite subgroup $K \subset A[\ell^\infty]$, we denote by $\Lambda(K) \subset V_\ell(A)$ its preimage under the above isomorphism. Then we immediately have:

Proposition 1.2.4. *The association $K \mapsto \Lambda(K)$ realizes a one-to-one correspondence between finite subgroups of A of ℓ-power order and lattices in $V_\ell(A)$ containing $T_\ell(A)$. The subgroup K is canonically isomorphic to $\Lambda(K)/T_\ell(A)$ via the above isomorphism $V_\ell(A)/T_\ell(A) \simeq A[\ell^\infty]$.*

If k is perfect, then k-rational subgroups of A correspond under the bijection of Proposition 1.2.4 to overlattices of $T_\ell(A)$ that are stable under G_k.

1.2.2 Morphisms between Tate modules

Consider now a morphism $\phi \in \mathrm{Hom}(A, B)$ between two abelian varieties over k, and a prime ℓ distinct from p. For each $n \geq 0$, ϕ induces a map from $A[\ell^n]$ to $B[\ell^n]$ that is $(\mathbb{Z}/\ell^n\mathbb{Z})$-linear and Galois-equivariant (because ϕ is defined over k by assumption.) Therefore, ϕ induces a $\mathbb{Z}_\ell$-linear and G_k-equivariant map from $T_\ell(A)$ to $T_\ell(B)$, denoted by $T_\ell(\phi)$. We have thus constructed a morphism of $\mathbb{Z}$-modules

$$T_\ell : \mathrm{Hom}(A, B) \to \mathrm{Hom}_{G_k}(T_\ell(A), T_\ell(B)).$$

After tensoring with $\mathbb{Q}_\ell$, we can also consider $T_\ell(\phi)$ as a map from $V_\ell(A)$ to $V_\ell(B)$. In that case, we denote it as $V_\ell(\phi)$ to disambiguate.

Using the previous correspondence between lattices in $V_\ell(A)$ and subgroups of ℓ-power order, we can recover the ℓ-primary part of $\ker(\phi)$ from $T_\ell(\phi)$:

Proposition 1.2.5. *Let $\phi : A \to B$ be any isogeny. Then $T_\ell(\phi) : T_\ell(A) \to T_\ell(B)$ is injective with finite cokernel, and $V_\ell(\phi) : V_\ell(A) \to V_\ell(B)$ is a bijection. Under the correspondence of Proposition 1.2.4, $\ker(\phi) \cap A[\ell^\infty]$ corresponds to the sublattice $V_\ell(\phi)^{-1}(T_\ell(B))$ of $T_\ell(A)$. Via the map $V_\ell(\phi)$, we also have an isomorphism*

$$\ker(\phi) \cap A[\ell^\infty] \simeq V_\ell(\phi)^{-1}(T_\ell(B))/T_\ell(A) \simeq T_\ell(B)/V_\ell(\phi)(T_\ell(A)).$$

Proof. Left to the reader (chase through the definitions.) □

We now analyze more closely this map T_ℓ. We first prove:

Lemma 1.2.6. *The map $T_\ell : \mathrm{Hom}(A, B) \to \mathrm{Hom}_{G_k}(T_\ell(A), T_\ell(B))$ is injective. In particular $\mathrm{Hom}(A, B)$ is a torsion-free abelian group.*

Proof. Let $\phi : A \to B$ be a morphism such that $T_\ell(\phi) = 0$. This means that $\ker(\phi)$, seen as a subvariety of A, contains the torsion subgroups $A[\ell^n]$ for every $n \geq 0$. Let B be the connected component of $\ker(\phi)$ containing 0_A; it is an abelian subvariety of A, and $\ker(\phi)$ consists of finitely many translates of B in A. If $\dim B < \dim A$, then $\#\ker(\phi)[\ell^n](\overline{k}) \ll \ell^{2n(g-1)}$ as $n \to \infty$, contradicting $A[\ell^n] \subset \ker(\phi)$. Therefore $B = A$. □

In fact, one can prove a stronger result than Lemma 1.2.6. Since the codomain of T_ℓ is a $\mathbb{Z}_\ell$-module, we can also consider the following map that we also call T_ℓ by an abuse of notation

$$T_\ell : \mathrm{Hom}(A,B) \otimes_{\mathbb{Z}} \mathbb{Z}_\ell \to \mathrm{Hom}_{G_k}(T_\ell(A), T_\ell(B)).$$

Proposition 1.2.7. *The map $T_\ell : \mathrm{Hom}(A,B) \otimes_{\mathbb{Z}} \mathbb{Z}_\ell \to \mathrm{Hom}_{G_k}(T_\ell(A), T_\ell(B))$ is injective.*

Proof. We follow [Mum70, Proof of Thm. 3 p. 176] (with a small modification: we don't use Poincaré reducibility.) It is enough to prove the result when $A = B$. Indeed, we have for any two abelian varieties A, B:

$$\mathrm{End}(A \times B) = \mathrm{End}(A) \oplus \mathrm{Hom}(A,B) \oplus \mathrm{Hom}(B,A) \oplus \mathrm{End}(B).$$

If T_ℓ is injective on $\mathrm{End}(A \times B) \otimes \mathbb{Z}_\ell$, it has to be in particular injective on $\mathrm{Hom}(A,B) \otimes \mathbb{Z}_\ell$. Taking $A = B$ allows us to take advantage of the fact that $\mathrm{End}(A)$ has a ring structure, and that endomorphisms of A have characteristic polynomials (Theorem 1.1.29).

Let now M be a finitely generated subgroup of $\mathrm{End}(A)$. We can consider $\mathbb{Q}M$ as a finite-dimensional $\mathbb{Q}$-vector space in the (possibly infinite-dimensional) $\mathbb{Q}$-algebra $\mathrm{End}^0(A)$. We claim that $\mathbb{Q}M \cap \mathrm{End}(A)$ is also a finitely generated $\mathbb{Z}$-module. Let g be the dimension of A, and let $P : \mathrm{End}^0(A) \to \mathbb{Q}^{2g}$ be the map which to an endomorphism $\alpha \in \mathrm{End}^0(A)$ associates the $2g$ non-leading coefficients of its characteristic polynomial P_α. The restriction of P to $\mathbb{Q}M$ is continuous for the usual (real) topologies on $\mathbb{Q}M$ and $\mathbb{Q}^{2g}$. Let I be the open interval $(-1,1)$, and let $U = P^{-1}(I^{2g})$. It is an open neighborhood of zero in $\mathbb{Q}M$. On the hand, the characteristic polynomial of an endomorphism α has integral coefficients, and is X^{2g} if and only if $\alpha = 0$, so $U \cap \mathrm{End}(A) = \{0\}$. Thus $\mathbb{Q}M \cap \mathrm{End}(A)$ is a discrete subgroup of $\mathbb{Q}M$, which implies our claim that it is finitely generated as a $\mathbb{Z}$-module (the only discrete subgroups of $\mathbb{Q}M$ are lattices.)

Finally, we prove that $T_\ell : \mathrm{End}(A) \otimes \mathbb{Z}_\ell \to \mathrm{End}(T_\ell(A))$ is injective. It is enough to prove that $T_\ell : M \otimes \mathbb{Z}_\ell \to \mathrm{End}(T_\ell(A))$ is injective for every finitely generated $\mathbb{Z}$-module $M \subset \mathrm{End}(A)$; by the previous claim, we may also assume that $M = \mathbb{Q}M \cap \mathrm{End}(A)$ inside $\mathrm{End}^0(A)$. Let $(f_1, \dots, f_r)$ be a $\mathbb{Z}$-basis of M. Let $\lambda_1, \dots, \lambda_r \in \mathbb{Z}_\ell$ such that $T_\ell(\sum \lambda_i f_i) = 0$. Fix $n \geq 0$. We will show that the ℓ-adic numbers λ_i are zero modulo ℓ^n; this is enough to conclude as n is arbitrary. Let $b_1, \dots, b_r$ be honest integers congruent to $\lambda_1, \dots, \lambda_r$ modulo ℓ^n. Then $f = \sum b_i f_i$ is an element of $\mathrm{End}(A)$ which acts like $T_\ell(\sum \lambda_i f_i)$ on $A[\ell^n]$, in other words $A[\ell^n] \subset \ker(f)$. By Proposition 1.1.12, there exists $f' \in \mathrm{End}(A)$ such that $f = \ell^n f'$. We have $f' \in \mathrm{End}(A) \cap \mathbb{Q}M = M$, so one can also write f' as a $\mathbb{Z}$-linear combination of $f_1, \dots, f_r$, say $f' = \sum c_i f_i$. We then have $b_i = \ell^n c_i$ for every i, so $\lambda_i = 0 \bmod \ell^n$ for every i. □

Corollary 1.2.8. *For any two abelian varieties A, B over k, $\mathrm{Hom}(A,B)$ is a free abelian group of finite rank $r \leq 4 \dim(A) \dim(B)$.*

Proof. Fixing $\mathbb{Z}_\ell$-bases of $T_\ell(A)$ and $T_\ell(B)$ identifies $\mathrm{Hom}(T_\ell(A), T_\ell(B))$ with the matrix space $\mathrm{Mat}_{2\dim B, 2\dim A}(\mathbb{Z}_\ell)$, which is a free $\mathbb{Z}_\ell$-module of rank $4\dim(A)\dim(B)$. By Proposition 1.2.7, the map $T_\ell : \mathrm{Hom}(A,B) \otimes \mathbb{Z}_\ell \to \mathrm{Hom}(T_\ell(A), T_\ell(B))$ is injective, so $\mathrm{Hom}(A,B) \otimes \mathbb{Z}_\ell$ has finite rank $r \leq 4\dim(A)\dim(B)$ as a $\mathbb{Z}_\ell$-module. □

Corollary 1.2.9. *For any simple abelian variety A over k, the endomorphism algebra $\mathrm{End}^0(A)$ is a finite-dimensional division $\mathbb{Q}$-algebra. In particular, its center $Z(\mathrm{End}^0(A))$ is a number field, and $\mathrm{End}^0(A)$ is a central division (hence simple) algebra over this number field.*

For any abelian variety A over k, $\mathrm{End}^0(A)$ is a semisimple $\mathbb{Q}$-algebra, and $Z(\mathrm{End}^0(A))$ is a product of number fields. For every $\alpha \in Z(\mathrm{End}^0(A))$ and every prime ℓ distinct from the characteristic of k, the endomorphism $V_\ell(\alpha)$ of $V_\ell(A)$ is semisimple, i.e. diagonalizable over $\overline{\mathbb{Q}}_\ell$.

Proof. Everything but the very last statement on the semisimplicity of $V_\ell(\alpha)$ is a direct consequence of Proposition 1.1.26, the Poincaré reducibility theorem 1.1.27, and Corollary 1.2.8. For the last statement, we know that $Z(\mathrm{End}^0(A))$ is a product of number fields, so α has to be annihilated by some polynomial without multiple roots, and so is $V_\ell(\alpha)$. □

1.2.3 Characteristic polynomials and polarizations

The following theorem is fundamental to understand how endomorphisms may act on distinct Tate modules.

Theorem 1.2.10. *Let A be an abelian variety over any field k, and let ℓ be a prime distinct from the characteristic of k. For any $\alpha \in \mathrm{End}(A)$, the characteristic polynomial of $T_\ell(A)$, which is well-defined as a monic polynomial in $\mathbb{Z}_\ell[X]$ of degree $2g$, is precisely the polynomial P_α from Theorem 1.1.29. In particular, it has integer coefficients and is independent of ℓ; the trace and degree of α are its trace and determinant as an endomorphism of $T_\ell(A)$.*

(This result is used in the proof of Theorem 1.1.29 to make sure that P_α has integral coefficients and $P_\alpha(\alpha) = 0$. The reader can check that these properties, are not used in the following proof.)

Proof. For every isogeny $\alpha \in \mathrm{End}(A)$, the largest power of ℓ dividing $\deg(\alpha)$ is precisely $\#(\ker(\alpha) \cap A[\ell^\infty])$. By Proposition 1.2.5, this is also the cardinality of $\mathrm{Coker}(T_\ell(\beta) : T_\ell(A) \to T_\ell(A))$, which is also the largest power of ℓ dividing $\det T_\ell(\beta)$. If $|\cdot|_\ell$ denotes the ℓ-adic valuation, the equality $|\deg(\alpha)|_\ell = |\det T_\ell(\alpha)|_\ell$ (possibly ∞) remains true for all endomorphisms, not only isogenies.

Let $x_1, \dots, x_{2g}$ be the roots of P_α in $\overline{\mathbb{Q}}$; we view them as elements in $\overline{\mathbb{Q}}_\ell$ after choosing an arbitrary embedding $\overline{\mathbb{Q}} \hookrightarrow \overline{\mathbb{Q}}_\ell$. Let $y_1, \dots, y_{2g} \in \overline{\mathbb{Q}}_\ell$ be the eigenvalues of $V_\ell(\alpha)$. Then for every nonzero polynomial $F \in \mathbb{Z}[X]$ with roots $\lambda_1, \dots, \lambda_r$ and leading coefficient c, we have

$$\begin{aligned}
\left|\prod_{i=1}^{2g} F(x_i)\right|_\ell &= \left|c^{2g} \prod_{i=1}^{2g}\prod_{j=1}^{r}(x_i - \lambda_j)\right|_\ell \\
&= \left|c^{2g} \prod_{j=1}^{r} P_\alpha(\lambda_j)\right|_\ell \\
&= \left|\deg\Big(c \prod_{j=1}^{r}(\lambda_j - \alpha)\Big)\right|_\ell \\
&= |\deg F(\alpha)|_\ell \\
&= |\det T_\ell(F(\alpha))|_\ell \\
\left|\prod_{i=1}^{2g} F(x_i)\right|_\ell &= \left|\prod_{i=1}^{2g} F(y_i)\right|_\ell .
\end{aligned}$$

In the third line, we consider deg as a polynomial map $\mathrm{End}^0(A) \otimes_{\mathbb{Q}} K \to K^{2g}$ where K is a splitting field of F; this extended degree map is still multiplicative. The last equality comes from the fact that the eigenvalues of $T_\ell(F(\alpha)) = F(T_\ell(\alpha))$ over $\overline{\mathbb{Q}}_\ell$ are $F(y_1), \dots, F(y_{2g})$.

By continuity, the equality on the last line remains true if $F \in \mathbb{Z}_\ell[X]$ or $F \in \mathbb{Q}_\ell[X]$. One can then show (by carefully choosing the polynomials F) that the collections $x_1, \dots, x_{2g}$ and $y_1, \dots, y_{2g}$ have to be the same up to permutations: see [Mil86a, Lemma 12.10]. □

For later use, we note also that when A is equipped with a polarization λ, the Tate module $T_\ell(A)$ is equipped with an alternating pairing

$$e_\ell : T_\ell(A) \times T_\ell(A) \to \mathbb{Z}_\ell(1) = \varprojlim_{n \to \infty} \mu_{\ell^n}(\overline{k})$$

which reduces to the Weil pairing e_{ℓ^n} on the ℓ^n-torsion subgroups for every $n \geq 0$. The determinant of this pairing is always nonzero (thus e_ℓ is nondegenerate on $V_\ell(A) \times V_\ell(A)$ always), and is invertible when ℓ is coprime to the degree of λ.

1.2.4 Tate's isogeny theorem

If k is finite, the following very important strengthening of Proposition 1.2.7 holds.

Theorem 1.2.11. *Let A and B be abelian varieties over a finite field k, and let $G_k = \mathrm{Gal}(\overline{k}/k)$. Let ℓ be any prime distinct from the characteristic of k. Then the map*

$$T_\ell : \mathrm{Hom}(A, B) \otimes \mathbb{Z}_\ell \to \mathrm{Hom}_{G_k}(T_\ell(A), T_\ell(B))$$

is a bijection.

This theorem will let us understand morphisms, endomorphisms and isogenies between abelian varieties only in terms of their action on Tate modules, i.e. in terms of linear algebra data.

Before proving Theorem 1.2.11, we describe what $\mathrm{Hom}_{G_k}(T_\ell(A), T_\ell(B))$ is in a more hands-on way. A key point here is that when k is a finite field, the Galois group G_k is generated (as a profinite group) by the Frobenius map $x \mapsto x^q$, where q is the cardinality of k.

Definition 1.2.12. Let A be an abelian variety over k. The *Frobenius endomorphism* $\pi_A \in \mathrm{End}(A)$ is the endomorphism obtained by raising the coordinates of points of A to their qth power (in any projective embedding of A). We denote the characteristic polynomial of π_A by f_A.

Almost by definition, we then have the following result.

Proposition 1.2.13. *Let A/k be an abelian variety as above, and let ℓ be a prime distinct from the characteristic of k. The action of the Frobenius endomorphism, seen as an element of G_k, on $T_\ell(A)$ is precisely given by $T_\ell(\pi_A)$.*

Corollary 1.2.14. *Let A, B be two abelian varieties over k of dimensions g, g' respectively, and let ℓ be a prime as above. Choose arbitrary bases of $T_\ell(A)$ and $T_\ell(B)$ as $\mathbb{Z}_\ell$-modules; this identifies morphisms between Tate modules with matrices over $\mathbb{Z}_\ell$. Then we have*

$$\mathrm{Hom}_{G_k}(T_\ell(A), T_\ell(B)) = \{m \in \mathrm{Mat}_{2g' \times 2g}(\mathbb{Z}_\ell) : T_\ell(\pi_B)m = mT_\ell(\pi_A)\}.$$

We can go further and exactly compute the rank of $\mathrm{Hom}(A, B)$ in terms of the characteristic polynomials of Frobenius on A and B.

Definition 1.2.15. Let f_1 and $f_2 \in \mathbb{Z}[X]$ be two monic polynomials. Factor them into distinct irreducibles as

$$f_1 = \prod_{i=1}^{r} m_i^{a_i}, \qquad f_2 = \prod_{i=1}^{r} m_i^{b_i}.$$

where $a_i, b_i \geq 0$. Then we define

$$r(f_1, f_2) = \sum_{i=1}^{r} a_i b_i \deg(m_i).$$

Corollary 1.2.16. *In the setting of Corollary 1.2.14, we have*

$$\operatorname{rank}_{\mathbb{Z}_\ell} \operatorname{Hom}_{G_k}(T_\ell(A), T_\ell(B)) = r(f_A, f_B).$$

Proof. By Corollary 1.2.14, this rank is also the dimension over $\mathbb{Q}_\ell$ of the vector space

$$\{m \in \operatorname{Mat}_{2g' \times 2g}(\mathbb{Q}_\ell) : T_\ell(\pi_B)m = mT_\ell(\pi_A)\}.$$

We can compute this dimension after extending the base field $\mathbb{Q}_\ell$; after making such an extension, we can assume that π_A and π_B act as diagonal matrices by Corollary 1.2.9 whose diagonal coefficients are specified by the factorizations of f_A and f_B. Then, a straightforward computation (left to the reader) is enough to conclude. □

We now turn to the proof of Tate's isogeny theorem. Given that injectivity holds over any field by Proposition 1.2.7, we only have to prove surjectivity. We will see shortly that treating the case $A = B$ is enough. Roughly, the proof idea to produce elements in $\operatorname{End}(A) \otimes \mathbb{Z}_\ell$ is to use the information provided on the Tate module to construct many isogenies from A of ℓ-power degree. Because we work over a finite field, the number of possible codomains for this isogenies will be finite, so we get infinitely many isogenies of ℓ-power degree $A \to B$, where B is some fixed abelian variety. In this way, we get infinitely many cycles producing endomorphisms of A.

In Tate's original proof (the one we choose to follow in these notes), the finiteness argument is provided by the following theorem.

Theorem 1.2.17. *Let k be a finite field. Then, for every $g, d \geq 1$, the number of isomorphism classes of pairs (A, λ), where A is an abelian variety of dimension g over k and λ is a polarization on A of degree d defined over k, is finite.*

Proof. By Lefschetz's theorem (Theorem 1.1.16), any such (A, λ) gives rise to a subvariety of $\mathbb{P}^N$ for some fixed N of some fixed degree, that is well-defined up to the action of the automorphism group of $\mathbb{P}^N$ (i.e. up to changing coordinates by some matrix in $\mathrm{PGL}_N(k)$); conversely, two pairs (A, λ) whose images in $\mathbb{P}^N$ are isomorphic must be themselves isomorphic.

Next, we use the fact that subvarieties of some fixed degree in $\mathbb{P}^N$ of some fixed degree are parametrized by a *Chow variety* over k, which is a smooth projective variety of finite type [Har92, Thm. 21.2]. This Chow variety therefore has finitely many k-points, so the number of isomorphism classes of pairs (A, λ) is finite. □

Remark 1.2.18. It is true that for every $g \geq 1$, the number of abelian varieties of dimension g over k up to isomorphism is finite, regardless of polarizations. One can prove this result using Zarhin's trick: for every abelian variety A, the abelian variety $A^4 \times (A^\vee)^4$ admits a principal

polarization [EvdGM12, Thm. 11.29]. Using this stronger finiteness result somewhat simplifies the proof of Tate's isogeny theorem 1.2.11 [EvdGM12, §16.3], as one doesn't have to worry anymore about polarizations. Nevertheless, the initial proof we will follow isn't much more complicated. The notes [Lic11] are also a useful reference.

Remark 1.2.19. Theorem 1.2.11 is also true when k is a number field: this is one of the key results in Faltings' landmark paper [Fal83]. Faltings' proof also relies on a finiteness statement. However Theorem 1.2.17 would not be true if k is a number field; what is true is that the k-isogeny class of any fixed A/k is finite, because it consists of abelian varieties over k with bounded Faltings height.

In order to use Theorem 1.2.17, we equip A with some polarization of degree d (as A is projective), and we wish to construct infinitely many isogenies from A of ℓ-power degree whose kernels are maximal isotropic in certain torsion subgroups: by Proposition 1.1.18, this will be enough to guarantee that the codomain carries a polarization of the same degree as A (provided that ℓ is prime to d.) Those subgroups should further be stable under π_A for those isogenies to be defined over k. Constructing π_A-stable isotropic subgroups will be easier when π_A admits many stable lines in $A[\ell]$, in other words, when f_A factors as a product of linear polynomials mod ℓ. By the Chebotarev density theorem, this will happen for a positive density of primes ℓ, but not all of them.

Luckily, we can get away with proving Theorem 1.2.11 for one single prime ℓ, and deduce the result for all other primes. We explain this reduction in the rest of this paragraph, and defer the proof of Tate's theorem at a "nice" prime ℓ to the end of this subsection.

Proposition 1.2.20. *Assume that Theorem 1.2.11 holds when $A = B$ is any abelian variety over k. Then it holds for all pairs of abelian varieties (A, B).*

Proof. Apply Tate's theorem to $\mathrm{End}(A \times B)$ as in the proof of Proposition 1.2.7. □

Proposition 1.2.21. *Let A be an abelian variety over a finite field k and let ℓ be a prime distinct from $\mathrm{char}(k)$. Assume that the map*

$$V_\ell : \mathrm{End}(A) \otimes \mathbb{Q}_\ell \to \mathrm{End}_{G_k}(V_\ell(A))$$

is bijective. Then Theorem 1.2.11 holds for A at ℓ.

Proof. What we have to show is that the map $T_\ell : \mathrm{End}(A) \to \mathrm{End}(T_\ell(A))$ has torsion-free cokernel. Since every prime but ℓ is invertible in $\mathbb{Z}_\ell$, it is enough to prove that $\mathrm{Coker}(T_\ell)$ is free of ℓ-torsion. Let $f \in \mathrm{End}(T_\ell(A))$ and assume that $\ell f \in \mathrm{Im}(T_\ell)$. In other words, we can find endomorphisms $\phi_1, \dots, \phi_r$ of A and coefficients $\lambda_1, \dots, \lambda_r \in \mathbb{Z}_\ell$ such that $\ell f = \sum \lambda_i \phi_i$. Approximating the coefficients λ_i by integers, we find a sequence of endomorphisms $\phi_n \in \mathrm{End}(A)$ such that $T_\ell(\phi_n)$ converges to ℓf for the ℓ-adic topology. For sufficiently large n, the action of ϕ_n on $A[\ell]$ then coincides with that of ℓf, i.e. $A[\ell] \subset \ker(\phi_n)$. By isogeny factorization (Proposition 1.1.12), we can write $\phi_n = \ell\psi_n$ for some $\psi_n \in \mathrm{End}(A)$. Since the sequence $(\phi_n) = (\ell\psi_n)$ converges in $\mathrm{End}(A) \otimes \mathbb{Z}_\ell$, so does the sequence (ψ_n) (because $\mathrm{End}(A)$ is torsion-free by Lemma 1.2.6). Let $\psi \in \mathrm{End}(A) \otimes \mathbb{Z}_\ell$ be the limit. We obtain $T_\ell(\ell\psi) = \ell f$, so $T_\ell(\psi) = f$, hence $f \in \mathrm{Im}(T_\ell)$. □

Proposition 1.2.22. *Let A be an abelian variety over a finite field k. Assume that there exists a prime ℓ, distinct from $\mathrm{char}(k)$, such that $V_\ell : \mathrm{End}(A) \otimes \mathbb{Q}_\ell \to \mathrm{End}_{G_k}(V_\ell(A))$ is bijective. Then V_ℓ is bijective at all such primes ℓ, hence Theorem 1.2.11 holds for A by Proposition 1.2.21.*

Proof. By Proposition 1.2.7, we know that V_ℓ is injective at all primes ℓ distinct from $\mathrm{char}(k)$. Therefore we only have to prove an equality of dimensions over $\mathbb{Q}_\ell$ for all primes ℓ. Obviously, $\dim_{\mathbb{Q}_\ell}(\mathrm{End}(A)\otimes\mathbb{Q}_\ell) = \mathrm{rank}_{\mathbb{Z}}(\mathrm{End}(A))$ is independent of ℓ. On the other hand, we have

$$\dim_{\mathbb{Q}_\ell}(\mathrm{End}_{G_k}(V_\ell(A))) = r(f_A, f_A)$$

by Corollary 1.2.16, which is also independent of ℓ. Consequently, it is sufficient to check the equality of dimensions at a single ℓ, as claimed. □

Finally, we prove that the map

$$V_\ell : \mathrm{End}(A)\otimes\mathbb{Q}_\ell \to \mathrm{End}_{G_k}(V_\ell(A))$$

is bijective for a "nice" prime ℓ. As indicated above, we equip A with some polarization λ of degree d, and choose ℓ prime to d and such that f_A is a product of linear factors mod ℓ. Since the action of Frobenius on $V_\ell(A)$ is semisimple by Corollary 1.2.9, this implies that the $\mathbb{Q}_\ell$-algebra generated by $V_\ell(\pi_A)$ is a product of copies of $\mathbb{Q}_\ell$.

The key lemma in the proof (whose proof uses the finiteness result in Theorem 1.2.17) is the following. Recall that $V_\ell(A)$ is endowed with the nondegenerate alternating pairing e_ℓ.

Lemma 1.2.23. *Let $W \subset V_\ell(A)$ be a G_k-stable and maximal isotropic subspace. Then there exists $u \in \mathrm{End}(A)\otimes\mathbb{Q}_\ell$ such that $\mathrm{Im}(V_\ell(u)) = W$.*

Proof. Consider the following overlattices of $T_\ell(A)$:

$$X_n = (\ell^{-n}T_\ell(A)\cap W) + T_\ell(A).$$

By assumption, X_n is G_k-stable. Under the correspondence of Proposition 1.2.4, it corresponds to a subgroup $K_n \subset A[\ell^n]$ which is maximal isotropic for the Weil pairing e_{ℓ^n}. By Proposition 1.1.18, the codomain B_n of the isogeny $f_n : A \to B_n$ with kernel K_n is also equipped with a natural polarization of degree d. By Theorem 1.2.17, the abelian varieties B_n fall into finitely many isomorphism classes. Therefore, we can extract a subsequence $(f_{n_i})_{i\geq 0}$ of isogenies from A whose codomains B_{n_i} are all isomorphic to the fixed abelian variety B_{n_0}. Fix such isomorphisms $\eta_i : B_{n_0} \to B_{n_i}$. Consider now the elements

$$u_i = \ell^{n_i} f_{n_i}^{-1} \circ \eta_i \circ f_{n_0} \in \mathrm{End}^0(A),$$

and let's compute $V_\ell(u_i)(X_{n_0})$. By construction,

$$\begin{aligned} V_\ell(u_i)(X_{n_0}) &= \ell^{n_i} V_\ell(f_{n_i})^{-1} \circ V_\ell(\eta_i)(T_\ell(B_{n_0})) \\ &= \ell^{n_i} V_\ell(f_{n_i})^{-1}(T_\ell(B_{n_i})) \\ &= \ell^{n_i} X_{n_i} \subset T_\ell(A) \subset X_{n_0}. \end{aligned}$$

Therefore the elements $V_\ell(u_i)$ belong to $\mathrm{End}(X_{n_0}) \cap V_\ell(\mathrm{End}(A)\otimes\mathbb{Q}_\ell)$, which is a compact subset of $\mathrm{End}(V_\ell(A))$. Up to extracting a further subsequence, we may assume that $V_\ell(u_i)$ converges to $V_\ell(u)$, for some $u \in \mathrm{End}(A)\otimes\mathbb{Q}_\ell$.

We now prove that u satisfies the conclusion of the lemma. To show $\mathrm{Im}(V_\ell(u)) \subset W$, it is sufficient to show $V_\ell(u)(X_{n_0}) \subset W$. Let $x \in X_{n_0}$. Then

$$V_\ell(u)(x) = \lim_{i\to\infty} V_\ell(u_i)(x) \in \bigcap_{i\geq 0} \ell^{n_i} X_{n_i} = W \cap T_\ell(A) \subset W.$$

Conversely, to show $W \subset \operatorname{Im} V_\ell(u)$, it is sufficient to show that this image contains $T_\ell(A) \cap W$. Let $y \in T_\ell(A) \cap W$. By the above computation, for each $i \geq 0$, there exists $x_i \in X_{n_0}$ such that $V_\ell(u_i)(x_i) = y$. Since X_{n_0} is compact, up to extracting a further subsequence, we may assume that the sequence (x_i) converges to some $x \in X_{n_0}$. Then one can check from the definition of ℓ-adic convergence that $u(x) = y$. □

Note that in this proof, we have only used the fact that ℓ is prime to d.

We use Lemma 1.2.23 as follows. If an element $u \in \operatorname{End}(V_\ell(A))$ commutes with $V_\ell(\operatorname{End}(A) \otimes \mathbb{Q}_\ell)$, then by the lemma, it should stabilize all the maximal isotropic, G_k-stable subspaces $W \subset V_\ell(A)$. We will see that every G_k-stable line is an intersection of G_k-stable maximal isotropic subspaces, so u must also stabilize every G_k-stable line in $V_\ell(A)$. Given our assumption on ℓ, Frobenius acts as a diagonal matrix on $V_\ell(A)$, so this implies $u \in \mathbb{Q}_\ell(V_\ell(\pi_A))$. We obtain that the centralizer *of the centralizer* of $V_\ell(\operatorname{End}(A) \otimes \mathbb{Q}_\ell)$ is the centralizer of $\mathbb{Q}_\ell(V_\ell(\pi_A))$, in other words $\operatorname{End}_{G_k}(V_\ell(A))$. But this double centralizer is precisely $V_\ell(\operatorname{End}(A)\otimes)$ by the double centralizer theorem, and we are done. Let's spell out this proof in more details.

Let $D \subset \operatorname{End}(V_\ell(A))$ denote the centralizer of $V_\ell(\operatorname{End}(A) \otimes \mathbb{Q}_\ell)$. Recall that we assumed that $\mathbb{Q}_\ell(V_\ell(\pi_A))$ is a product of copies of $\mathbb{Q}_\ell$.

Lemma 1.2.24. *Let $W \subset V_\ell(A)$ be any G_k-stable isotropic subspace. Then W is D-stable.*

Proof. Note that $0 \leq \dim_{\mathbb{Q}_\ell}(W) \leq g$, where $g = \dim A$, because W is isotropic and e_ℓ is non-degenerate. We prove the lemma by descending induction on $d = \dim W$.

- If $d = g$, then W is maximal isotropic. By Lemma 1.2.23, there exists $u \in V_\ell(\operatorname{End}(A) \otimes \mathbb{Q}_\ell)$ such that $\operatorname{Im}(u) = W$. For every $v \in D$, we then have
$$v(W) = v(u(V_\ell(A))) = u(v(V_\ell(A)) \subset u(V_\ell(A)) = W.$$
- If $d < g$, we assume by induction that the lemma holds for every G_k-stable subspace W' of dimension at least $d+1$. The orthogonal $W^\perp$ of W with respect to e_ℓ satisfies $W \subset W^\perp$ and $\dim_{\mathbb{Q}_\ell} W^\perp = 2g - d \geq d + 2$. Given our assumption on $\mathbb{Q}_\ell(V_\ell(\pi_A))$, we can find lines $L_1, \ldots, L_r$ (with $r \geq 2$) in $V_\ell(A)$ that are stable under $V_\ell(\pi_A)$ (hence under G_k) and which satisfy
$$W^\perp = W \oplus L_1 \oplus L_2 \oplus \cdots \oplus L_r.$$
Let $v \in D$. By the induction hypothesis, v stabilizes $W \oplus L_1$ and $W \oplus L_2$, hence it stabilizes their intersection W. □

Note that lines in $V_\ell(A)$ are always isotropic because e_ℓ is alternating, hence Lemma 1.2.24 applies to any G_k-stable line in $V_\ell(A)$.

Lemma 1.2.25. *We have $D = \mathbb{Q}_\ell(V_\ell(\pi_A))$.*

Proof. Given our assumption on $\mathbb{Q}_\ell(V_\ell(\pi_A))$, we can decompose $V_\ell(A)$ as
$$V_\ell(A) = V_1 \oplus \cdots \oplus V_r$$
where each V_i is a nontrivial π_A-stable subspace where π_A acts as a scalar $\lambda_i \in \mathbb{Q}_\ell$, and $\lambda_1, \ldots, \lambda_r$ are distinct. Let $v \in D$. By Lemma 1.2.24, for every $1 \leq i \leq r$, v stabilizes every line in V_i. Hence v acts as a scalar μ_i on V_i as well. If F is any polynomial such that $F(\lambda_i) = \mu_i$ for each i, we then have $v = F(V_\ell(\pi_A))$. □

Proposition 1.2.26. *We have* $V_\ell(\mathrm{End}(A)\otimes\mathbb{Q}_\ell) = \mathrm{End}_{G_k}(V_\ell(A))$.

Proof. By Lemma 1.2.25, the centralizer D of $V_\ell(\mathrm{End}(A)\otimes\mathbb{Q}_\ell)$ in $\mathrm{End}(V_\ell(A))$ is $\mathbb{Q}_\ell(V_\ell(\pi_A))$, so the double centralizer of $V_\ell(\mathrm{End}(A)\otimes\mathbb{Q}_\ell)$ is $\mathrm{End}_{G_k}(V_\ell(A))$. Since $\mathrm{End}(A)\otimes\mathbb{Q}_\ell$ is a semisimple $\mathbb{Q}_\ell$-algebra, we conclude by the double centralizer theorem (see for instance [Mil20, Thm. 1.14]) that $V_\ell(\mathrm{End}(A)\otimes\mathbb{Q}_\ell) = \mathrm{End}_{G_k}(V_\ell(A))$. □

Combining Proposition 1.2.26 with the reductions from the previous paragraph, in particular Proposition 1.2.22, concludes the proof of Tate's isogeny theorem.

1.2.5 Tate's isogeny theorem at $\ell = p$

In order to have a solid foundation for the theory, knowing Tate's isogeny theorem at all primes $\ell \neq p$, where p denotes the characteristic, is not sufficient: we also need an analogue at $\ell = p$. However, the naive definition

$$T_p(A)\text{“}=\text{”}\varprojlim_{n\to\infty} A[p^n](\overline{k})$$

cannot work: for instance, if A is a supersingular elliptic curve, then this would lead to $T_p(A) = \{0\}$.

Fortunately, a solution exists: we need to consider the sequence of subgroup *schemes* (an example of *p-divisible group*)

$$\{0\} \hookrightarrow A[p] \hookrightarrow A[p^2] \hookrightarrow \cdots \hookrightarrow A[p^n] \hookrightarrow \cdots$$

and apply a certain equivalence of categories between finite group schemes of p-power and certain objects in semi-linear algebra, as follows.

Let W be the ring of Witt vectors of the finite field k: if $k = \mathbb{F}_p$ is a prime field, then $W = \mathbb{Z}_p$, and in general if $k = \mathbb{F}_{p^r}$, then W is the ring of integers in the unique unramified extension of $\mathbb{Q}_p$ of degree r. The ring W has a unique maximal ideal pW, and W/pW is isomorphic to k. There is a canonical automorphism $\sigma : W \to W$ that lifts the absolute Frobenius automorphism $x \mapsto x^p$ of k: for instance, if $W = \mathbb{Z}_p$, then σ can be constructed in terms of Teichmüller representatives. Let R be the (noncommutative!) ring $W[F, V]$, where F, V are two indeterminates subject to the relations

$$FV = VF = p, \quad \text{and for all } \alpha \in W, \quad F\alpha = \sigma(\alpha)F \text{ and } V\alpha = \sigma^{-1}(\alpha)V.$$

Theorem 1.2.27 (Dieudonné–Cartier–Oda; see [Wat69, §1.2])**.** *There exists a contravariant equivalence of categories $\mathbb{D}$ between finite commutative group schemes over k of p-power rank and R-modules of finite W-length; if G is such a group scheme of rank p^n, then $\mathbb{D}(G)$ has W-length n.*

Definition 1.2.28. We define the *Dieudonné module* of an abelian variety A over k as

$$T_p(A) = \varprojlim_{n\to\infty} \mathbb{D}(A[p^n]).$$

It is a free W-module of rank $2g$ equipped with semi-linear actions of endomorphisms F and V. (Note that the limit is an inverse one because $\mathbb{D}$ is contravariant.)

Much like what happened with Tate modules, we can recover p-torsion information on A from the Dieudonné module: for instance $A[p^n]$ is the group scheme corresponding via $\mathbb{D}$ to the quotient $T_p(A)/p^nT_p(A)$, which is an R-module of W-length p^{2ng}.

We can also consider $T_p(A)$ as a lattice inside a vector space: we have $R \otimes_{\mathbb{Z}_p} \mathbb{Q}_p = L[F, V]$, where $L = W \otimes_{\mathbb{Z}_p} \mathbb{Q}_p$ is the unramified extension of $\mathbb{Q}_p$ defined above (L is the fraction field of W).

Since p is now invertible, $L[F,V] = L[F,F^{-1}]$, so an $L[F,V]$-module is simply an L-vector space equipped with a semi-linear (but $\mathbb{Q}_p$-linear) and bijective action of F. In particular, for an abelian variety A, the linearized Dieudonné module $V_p(A) = T_p(A) \otimes_{\mathbb{Z}_p} \mathbb{Q}_p$ is a $2g$-dimensional L-vector space equipped with such an action of F.

Theorem 1.2.29. *For any abelian varieties A, B over a finite field k, the map*

$$T_p : \mathrm{Hom}(A,B) \otimes \mathbb{Z}_p \to \mathrm{Hom}_R(T_p(B), T_p(A))$$

is a bijection. For any $\alpha \in \mathrm{End}(A)$, the characteristic polynomial $P_\alpha \in \mathbb{Z}[X]$ is the characteristic polynomial of $V_p(\alpha) \in \mathrm{End}(V_p(A))$ seen as an L-vector space.

The proof of this theorem in [WM71] uses Theorem 1.2.27 as a black box and is somewhat easier (in any case shorter) then that of Theorem 1.2.11, because we already know that $\mathrm{Hom}(A,B)$ is a free $\mathbb{Z}$-module of the correct rank in terms of f_A and f_B.

In the rest of the course, we will not delve into any detailed arguments involving Dieudonné modules, as those quickly get very technical.

1.3 Honda-Tate theory

The aim of Honda-Tate theory is to describe *isogeny classes* of abelian varieties over finite fields only in terms of arithmetic information on the characteristic polynomial of Frobenius. A large part of Honda-Tate theory consists in fairly direct consequences of Tate's isogeny theorem 1.2.11.

1.3.1 Isogeny classes and the characteristic polynomial of Frobenius

Theorem 1.3.1. *Let A and B be two abelian varieties over a finite field k. The following are equivalent:*

1. *A and B are k-isogenous.*
2. *$f_A = f_B$.*
3. *For every finite extension k'/k, we have $\#A(k') = \#B(k')$.*

Proof. (1) $\implies$ (2): let ℓ be any prime distinct from $\mathrm{char}(k)$. Then $V_\ell(\phi)$, where $\phi : A \to B$ is any isogeny, realizes a Galois-equivariant isomorphism between $V_\ell(A)$ and $V_\ell(B)$, so the characteristic polynomials of Frobenius must be the same. Those are f_A and f_B by Theorem 1.1.29.

(2) $\implies$ (1): if $f_A = f_B$, then since the action of Frobenius on $V_\ell(A)$ and $V_\ell(B)$ is semi-simple (Corollary 1.2.9), $V_\ell(A)$ and $V_\ell(B)$ are isomorphic as G_k-modules. By Tate's isogeny theorem, $V_\ell(\mathrm{Hom}(A,B) \otimes \mathbb{Q}_\ell)$ contains isomorphisms between $V_\ell(A)$ and $V_\ell(B)$, hence so does $V_\ell(\mathrm{Hom}(A,B) \otimes \mathbb{Q})$. Therefore A and B are isogenous.

(2) $\iff$ (3): if k' is an extension of k of degree r, then $\#A(k')$ is the degree of $\pi_A^r - 1$. This degree has a symmetric expression in terms of the roots of f_A, hence has an expression in terms of the coefficients of f_A. Conversely, knowing sufficiently many degrees of $\pi_A^r - 1$ determines the roots of f_A: see e.g. [WM71, Thm. 7]. □

More generally, f_A divides f_B if and only if A is isogenous to some abelian subvariety of B [Tat66, §3, Thm. 1]. More can be said about isogeny classes of simple abelian varieties.

Theorem 1.3.2. *Let A be a simple abelian variety over a finite field k of characteristic p. Then:*

1. *f_A is the power of an irreducible polynomial: there exists $e \geq 1$ and an irreducible, monic polynomial $m_A \in \mathbb{Z}[X]$ such that $f_A = m_A^e$.*
2. *For every prime ℓ distinct from p, m_A is the minimal polynomial of $V_\ell(\pi_A) \in \mathrm{End}(V_\ell(A))$.*
3. *$\mathrm{End}^0(A)$ is a division algebra with center $\mathbb{Q}(\pi_A)$, the number field defined by the polynomial m_A.*
4. *We have $\dim_{\mathbb{Q}} \mathrm{End}^0(A) = e^2[\mathbb{Q}(\pi_A) : \mathbb{Q}]$ and $2 \dim A = e[\mathbb{Q}(\pi_A) : \mathbb{Q}]$.*
5. *We have $2g \leq \dim_{\mathbb{Q}} \mathrm{End}^0(A) \leq (2g)^2$.*

Proof. By Proposition 1.1.26, we know $\mathrm{End}^0(A)$ is a division algebra. By Tate's isogeny theorem, the center of $\mathrm{End}(A) \otimes \mathbb{Q}_\ell$ is $\mathbb{Q}_\ell(\pi_A)$, so the center of $\mathrm{End}^0(A)$ is $\mathbb{Q}(\pi_A)$. This shows $\mathbb{Q}(\pi_A)$ is a number field. If f_A had several prime factors, there would be zero divisors in $\mathbb{Q}(\pi_A)$, which is impossible. This proves (1). Further, m_A has to be the defining polynomial of the number field $\mathbb{Q}(\pi_A)$, so $m_A(\pi_A) = 0$, proving (2) and (3).

By Corollary 1.2.16, we have $\dim_{\mathbb{Q}} \mathrm{End}^0(A) = r(f_A, f_A)$. This value is clearly at least $2g$ (when f_A is irreducible of degree $2g$) and at most $(2g)^2$ (when $e = 2g$ and m_A is linear), proving (5).

Finally, we prove (4). We have $2 \dim A = \deg(f_A) = e \deg(m_A) = e[\mathbb{Q}(\pi_A) : \mathbb{Q}]$. Fix an ℓ such that m_A is irreducible modulo ℓ. Then $V_\ell(A)$ has the structure of a vector space of dimension e over the field $\mathbb{Q}_\ell(\pi_A)$. By Tate's isogeny theorem, $V_\ell(\mathrm{End}(A) \otimes \mathbb{Q}_\ell)$ is then identified with the space of $e \times e$ matrices over $\mathbb{Q}_\ell(\pi_A)$. Therefore

$$\dim_{\mathbb{Q}} \mathrm{End}^0(A) = \dim_{\mathbb{Q}_\ell} \mathrm{End}(A) \otimes \mathbb{Q}_\ell = e^2[\mathbb{Q}_\ell(\pi_A) : \mathbb{Q}_\ell] = e^2[\mathbb{Q}(\pi_A) : \mathbb{Q}]. \qquad \square$$

Theorem 1.3.3. *Let A be a simple abelian variety of dimension g over a finite field k of characteristic p.*

1. *The following are equivalent:*
   (a) *$\mathrm{End}(A)$ is commutative.*
   (b) *$\mathrm{End}^0(A) = \mathbb{Q}(\pi_A)$.*
   (c) *$\dim_{\mathbb{Q}} \mathrm{End}^0(A) = 2g$.*
   (d) *f_A is irreducible.*
   (e) *$e = 1$.*
   (f) *$\deg(m_A) = 2g$.*
2. *The following are equivalent:*
   (a) *$\mathbb{Q}(\pi_A) = \mathbb{Q}$.*
   (b) *$\dim_{\mathbb{Q}} \mathrm{End}^0(A) = (2g)^2$.*
   (c) *f_A is the power of a linear polynomial.*
   (d) *$e = 2g$.*
   (e) *$\deg(m_A) = 1$.*

(f) $\mathrm{End}^0(A) \simeq \mathrm{Mat}_{g\times g}(B_{p,\infty})$, *where* $B_{p,\infty}$ *denotes the unique quaternion algebra over* $\mathbb{Q}$ *ramified at* p *and* ∞.

(g) A *is isogenous to* E^g, *where* E *is an elliptic curve over* k *such that* $\mathrm{End}^0(E) = B_{p,\infty}$.

Proof. Everything is a direct consequence of Theorem 1.3.2, except deducing (f) and (g) from (a)–(e) in case (2). We postpone this until after the discussion of Brauer groups and the invariants of $\mathrm{End}^0(A)$ in the next two paragraphs. □

Definition 1.3.4. We say that A is *supersingular with all endomorphisms defined* (abbreviation: ss.def.) if (2) holds for A over k, and we say that A is *supersingular* if (2) holds for A over some finite extension k'/k.

Note that if (2) holds for A/k, it also holds for the base change of A over every finite extension k'/k, because (2) implies $\mathbb{Q}(\pi_A^r) = \mathbb{Q}$ for every $r \geq 1$.

We conclude this discussion of isogeny classes by relating the notion of supersingular abelian varieties with the perhaps usual definition in terms of p-ranks.

Proposition 1.3.5. *Let* A *be an abelian variety of dimension* g *over a finite field* k *of characteristic* p. *There exists an integer* $0 \leq r \leq g$ *such that* $A[p^n](\overline{k}) \simeq (\mathbb{Z}/p^n\mathbb{Z})^r$ *for every* $n \geq 0$. *We have* $r = 0$ *if and only if* A *is supersingular.*

Proof. This uses Tate's isogeny theorem at $l = p$ (Theorem 1.2.29) as well as an analysis of the structure of $T_p(A)$ in relation with the p-rank: see [Wat69, Cor. 4.4]. □

We say that A is *ordinary* if it has p-rank g; equivalently, half of the roots of f_A over $\overline{\mathbb{Q}}$ have p-adic valuation 0. If A is simple and ordinary, then it is in case (1) of Theorem 1.3.3 for every finite extension k'/k. The converse to this statement holds for elliptic curves (an elliptic curve is either ordinary or supersingular), but not (I think) for higher-dimensional abelian varieties.

1.3.2 Brauer groups of number fields

In order to complete Theorem 1.3.2 with the description of what exactly $\mathrm{End}^0(A)$ is in terms of f_A, we digress and discuss Brauer groups of number fields.

For now, let K be any field. We only consider finite-dimensional K-algebras. Such an algebra B is called *simple* if its only two-sided ideals are $\{0\}$ and B itself, and *central* if its center is exactly K.

Theorem 1.3.6 (Wedderburn)**.** *Let* B *be a central simple algebra over* K. *Then there exists a central division algebra* D *over* K, *unique up to isomorphism, and a unique integer* $r \geq 1$ *such that* B *is isomorphic to* $\mathrm{Mat}_{r\times r}(D)$. *More generally, if* B *is any central simple algebra over* K, *we denote by* $[B]$ *the unique Brauer class containing* B.

Wedderburn's theorem allows us to partition central simple algebras over K into equivalence classes as follows.

Definition 1.3.7. Let D be a central division algebra over K. The *Brauer class* of D consists of the isomorphism classes of the central simple K-algebras $\mathrm{Mat}_{r\times r}(D)$ for $r \geq 1$. We denote it as $[D]$.

Definition 1.3.8. The *Brauer group* of K, denoted as $\mathrm{Br}(K)$, is the group whose element are the Brauer classes of the central division algebras over K, with the following group operation: if D and D' are such division algebras, then $[D] \cdot [D'] = [D \otimes_K D']$.

The Brauer group is a well-defined abelian group, because $D \otimes_K D' \simeq D' \otimes_K D$ is a central simple algebra over K: see e.g. [Mil20, IV, Prop. 2.3].

Descibing the structure of Brauer groups of number fields is one of the key outcomes of class field theory. Before we state the main result, we describe what the (simpler) Brauer groups of local fields look like. Recall that local fields are either finite extensions of $\mathbb{Q}_p$ for some p (the non-archimedean ones), or $\mathbb{R}$ or $\mathbb{C}$ (the archimedean ones.)

Theorem 1.3.9 ([Mil20, IV, Prop. 4.3])**.** *Let K be a local field. Then there exists a canonical injective map*

$$\mathrm{inv}_K : \mathrm{Br}(K) \to \mathbb{Q}/\mathbb{Z}$$

with the following properties:

1. *If $K = \mathbb{C}$, then the image of inv_K is $\{0\}$ (in fact, the Brauer group of any algebraically closed field is trivial);*
2. *If $K = \mathbb{R}$, then the image of inv_K is $\{0, 1/2\}$ (the central division algebras over $\mathbb{R}$ are $\mathbb{R}$ and the Hamiltonian quaternions);*
3. *If K is non-archimedean, then inv_K is surjective.*

If K is a number field, then for each place v of K, we have a natural map

$$\begin{array}{ccc} \mathrm{Br}(K) & \to & \mathrm{Br}(K_v) \\ [D] & \mapsto & [D \otimes_K K_v] \end{array}$$

where K_v denotes the completion of K at v. For simplicity, we denote by inv_v the composite map $\mathrm{Br}(K) \to \mathrm{Br}(K_v) \to \mathbb{Q}/\mathbb{Z}$, and we allow central simple algebras (instead of their Brauer classes) as input of the inv functions.

Theorem 1.3.10. *Let K be a number field. Then we have an exact sequence*

$$0 \longrightarrow \mathrm{Br}(K) \longrightarrow \bigoplus_v \mathrm{Br}(K_v) \xrightarrow{\text{sum of } \mathrm{inv}_{K_v}} \mathbb{Q}/\mathbb{Z} \longrightarrow 0,$$

where the direct sum is over all places v of K. In other words:

1. *Two central division algebras D, D' over K are isomorphic if and only if $\mathrm{inv}_v(D) = \mathrm{inv}_v(D')$ for all places v of K.*
2. *If D is a central simple algebra over K, then $\mathrm{inv}_v(D) = 0$ for all places v of K but finitely many. Moreover $\sum_v \mathrm{inv}_v(D) = 0$ in $\mathbb{Q}/\mathbb{Z}$.*
3. *Conversely, let $(i_v)_v$ be a collection of elements in $\mathbb{Q}/\mathbb{Z}$ such that $i_v = 0$ for almost every v, the sum of all values i_v is 0 in $\mathbb{Q}/\mathbb{Z}$, $i_v = 0$ if v is complex, and $2i_v = 0$ if v is real. Then there exists a unique central division algebra D over K such that $\mathrm{inv}_v(D) = i_v$ for all v.*

The group $\mathrm{Br}(K)$ is torsion. If D is a central division algebra over K, then the order e of $[D]$ in $\mathrm{Br}(K)$ is the least common denominator of the invariants $\mathrm{inv}_v(D)$. We have $\dim_{\mathbb{Q}} D = e^2$.

Proof. [Mil20, VII, Thms. 7.1 and 8.1] prove everything about the exact sequence, except item (3), for which we refer to [Har13, Thm. 9.11]. For the last statements, see [Mil20, VIII, Thm. 2.6]. □

Example 1.3.11. Let p be a prime number, which we identify with a place of $\mathbb{Q}$. Let ∞ denote the real place of $\mathbb{Q}$. By Theorem 1.3.10, there exists a unique central division algebra D over $\mathbb{Q}$ such that $\mathrm{inv}_p(D) = \mathrm{inv}_\infty(D) = 1/2$ and $\mathrm{inv}_v(D) = 0$ at all other places. The dimension of D over $\mathbb{Q}$ is 4, hence it is a quaternion algebra, often denoted as $B_{p,\infty}$.

More generally, a quaternion algebra D over any number field K (i.e. a central division algebra of dimension 4) satisfies $\mathrm{inv}_v(D) \in \{0, 1/2\}$ for all places v of K. If $\mathrm{inv}_v(D) = 1/2$, i.e. $D \otimes_K K_v$ is still division, we say that D *ramifies* at v. Otherwise, i.e. if $D \otimes_K K_v \simeq \mathrm{Mat}_{2\times 2}(K_v)$, we say that D *splits* at v. Theorem 1.3.10 imples that quaternion algebras over K are characterized by the set of places at which they ramify, always finite and of even size. Hence we may refer to $B_{p,\infty}$ as the unique quaternion algebra over $\mathbb{Q}$ ramified at p and ∞ and nowhere else.

1.3.3 Invariants of endomorphism algebras

We are now ready to complete Theorem 1.3.2 with a description of the central division algebra $\mathrm{End}^0(A)$ as an element of the Brauer group of its center.

Theorem 1.3.12. *Let A be a simple abelian variety over a finite field k of characteristic p and cardinality q. Then $\mathrm{End}^0(A)$ is the unique central division algebra D over the field $\mathbb{Q}(\pi_A)$ whose invariants at all places v of $\mathbb{Q}(\pi_A)$ are the following:*

- *If v is complex, then $\mathrm{inv}_v(D) = 0$;*
- *If v is real, then $\mathrm{inv}_v(D) = 1/2$;*
- *If v lies above a prime $\ell \neq p$, then $\mathrm{inv}_v(D) = 0$;*
- *If v lies above p, then*
$$\mathrm{inv}_v(D) = \frac{\mathrm{ord}_v(\pi_A)}{\mathrm{ord}_v(q)}[K_v : \mathbb{Q}_p]$$
where K_v denotes the completion of $\mathbb{Q}(\pi_A)$ at v.

Proof. If v is complex, there is nothing to prove.

If there exists a real prime v, then by the Riemann hypothesis (Theorem 1.3.13 below), we must have $\pi_A = \pm\sqrt{q}$. There are two cases depending on whether q is a square or not, but in any case $\mathrm{End}^0(A)$ ends up not being split at v: see [Tat66, Thm. 2].

Fix now a prime $\ell \neq p$, let m_A be the minimal polynomial of Frobenius on A, and let $e \geq 1$ such that $f_A = m_A^e$ as in Theorem 1.3.2. We look at the decomposition

$$\mathbb{Q}(\pi_A) \otimes_{\mathbb{Q}} \mathbb{Q}_\ell = \prod_{v|\ell} K_v.$$

Because the action of $V_\ell(\pi_A)$ on the Tate module $V_\ell(A)$ is semisimple, $V_\ell(A)$ is a free module of rank e over $\mathbb{Q}(\pi_A) \otimes_{\mathbb{Q}} \mathbb{Q}_\ell$. In other words, we have a decomposition

$$V_\ell(A) = \prod_{v|\ell} W_v,$$

where each W_v is an e-dimensional vector space over K_v. By Tate's isogeny theorem, we have

$$D \otimes_{\mathbb{Q}} \mathbb{Q}_\ell = \mathrm{End}_{\mathbb{Q}(\pi_A)\otimes\mathbb{Q}_\ell} V_\ell(A) \simeq \prod_{v|\ell} \mathrm{Mat}_{e\times e}(K_v).$$

Therefore, $D \otimes_{\mathbb{Q}(\pi_A)} K_v$ is isomorphic to $\mathrm{Mat}_{e\times e}(K_v)$ for each place v of $\mathbb{Q}(\pi_A)$ above ℓ, in other words the invariant of D at v is zero.

Computing the invariant at v when v is p-adic uses Dieudonné modules: we omit this proof here, and refer instead to [WM71, II, Thm. 2]. □

1.3.4 The Riemann hypothesis and the Honda–Tate theorem

If A is a simple abelian variety over a finite field k and the minimal polynomial m_A of Frobenius on A is known (or alternatively, the characteristic polynomial f_A), then Theorems 1.3.2 and 1.3.12 describe the algebra $\mathrm{End}^0(A)$ uniquely. A natural question at this point is: what can we say, a priori, about this polynomial m_A?

Theorem 1.3.13 (The Riemann hypothesis for abelian varieties over finite fields)**.** *Let A be an abelian variety over a finite field k of cardinality q. Then all the complex roots of the Frobenius characteristic polynomial f_A have complex absolute value $\sqrt{q}$.*

In particular, if A is simple, then m_A is the minimal polynomial over $\mathbb{Q}$ of the *q-Weil number* $\pi_A \in \mathbb{Q}(\pi_A)$ as in the following definition.

Definition 1.3.14. An algebraic integer $\pi \in \overline{\mathbb{Q}}$ is called a *q-Weil number* if the absolute value of π in each complex embedding $\mathbb{Q}(\pi) \hookrightarrow \mathbb{C}$ is $\sqrt{q}$.

Conversely, we have:

Theorem 1.3.15 (Honda)**.** *Let q be a prime power, and let $m \in \mathbb{Z}[X]$ be the minimal polynomial of some q-Weil number $\pi \in \overline{\mathbb{Q}}$. Then there exists a simple abelian variety A over $\mathbb{F}_q$ such that $m_A = m$.*

Combining Theorems 1.3.2, 1.3.12, 1.3.13 and 1.3.15 yields what is known as *Honda-Tate theory*: there is a one-to-one correspondence between isogeny classes of simple abelian varieties over $k = \mathbb{F}_q$ and minimal polynomials of q-Weil numbers (or in other words, with q-Weil numbers up to Galois conjugation in $\overline{\mathbb{Q}}$). Further, the minimal polynomial m_A explicitly determines the endomorphism algebra of the corresponding isogeny class.

We do not prove Theorems 1.3.13 and 1.3.15. Briefly, the Riemann hypothesis stems from the fact that once we fix a polarization on A, the associated Rosati involution is a positive involution on $\mathrm{End}^0(A)$, and the division algebras with positive involutions have been classified by Albert [Mum70, Thm. 2 p. 201]. One eventually obtains that the Rosati involution corresponds to the involution $\pi \mapsto q/\pi$ on the roots of m_A [Mum70, Thm. 4 p. 206]. Honda proved his theorem by constructing abelian varieties with the required minimal polynomials, first by constructing CM abelian varieties over number fields, then by reducing those modulo primes [Hon68].

Weil numbers as in Definition 1.3.14 might seem exceedingly rare. In fact, they can be constructed easily as follows.

Proposition 1.3.16. *Let q be a prime power.*

1. *If $\pi \in \overline{\mathbb{Q}}$ is a Weil number, then $\beta = \pi + q/\pi$ is a totally real algebraic integer such that $|\beta| \leq 2\sqrt{q}$ in each embedding.*
2. *Conversely, if $\beta \in \overline{\mathbb{Q}}$ is a totally real algebraic number such that $|\beta| \leq 2\sqrt{q}$ in each embedding, then any root π of $X^2 - \beta X + q$ in $\overline{\mathbb{Q}}$ is a q-Weil number.*

Proof. Left to the reader (it's easy.) □

Note that $\mathbb{Q}(\pi)$ is usually totally imaginary when π is a q-Weil number, unless $\pi = \pm\sqrt{q}$, which is a rather special case. This remark is used in the proof of Theorem 1.3.12.

1.3.5 Example: isogeny classes of elliptic curves

To conclude, we fully spell out what Honda–Tate theory shows in the case of elliptic curves over finite fields. We continue to use notation from Theorem 1.3.2, replacing A by E. Recall the notion of ordinary and supersingular abelian varieties from Definition 1.3.4.

Proposition 1.3.17. *Let E be an elliptic curve over $k = \mathbb{F}_q$ of characteristic p. Then exactly one of the following holds:*

1. $\mathrm{End}(E)$ *is commutative;* $\mathbb{Q}(\pi_E) = \mathrm{End}^0(E)$ *is a quadratic field;* $f_E = m_E$.
2. $\mathrm{End}(E)$ *is not commutative;* $\mathbb{Q}(\pi_E) = \mathbb{Q}$; π_E *is a rational integer;* $\mathrm{End}^0(E) = B_{p,\infty}$.

If E is ordinary, then we are in case 1; if we are in case 2, then E is ss.def. and q is a square.

Proof. In Theorem 1.3.2, we have $2g = 2 = e[\mathbb{Q}(\pi_E) : \mathbb{Q}]$, so we necessarily are in one of the two cases of Theorem 1.3.3. In case 2, π_E is a rational integer of absolute value $\sqrt{q}$ by Theorem 1.3.13, so q has to be a square. □

Proposition 1.3.18. *Let $q = p^a$ be a prime power. Isogeny classes of elliptic curves over $k = \mathbb{F}_q$ are in one-to-one correspondence with integers $\beta \in \mathbb{Z}$ such that $|\beta| \leq 2\sqrt{q}$ and one of the following (mutually exclusive) conditions holds:*

1. $\beta \neq 0 \bmod p$;
2. a *is even and* $\beta = \pm 2\sqrt{q}$;
3. a *is even,* $p \neq 1 \bmod 3$, *and* $\beta = \pm\sqrt{q}$;
4. a *is odd,* $p \in \{2, 3\}$, *and* $\beta = \pm p^{(a+1)/2}$;
5. a *is odd and* $\beta = 0$;
6. a *is even,* $p \neq 1 \bmod 4$, *and* $\beta = 0$.

The isogeny class corresponding to β consists of those elliptic curves E with $f_E = X^2 - \beta X + q$, in other words $\#E(\mathbb{F}_q) = q + 1 - \beta$. Case 1 corresponds to ordinary elliptic curves, case 2 to ss.def. elliptic curves, and all the other cases to supersingular elliptic curves without all endomorphisms defined (so $\mathrm{End}(E)$ *is still commutative).*

Proof. If E is ss.def., then $\pi_E = \pm\sqrt{q}$ and q is a square; therefore $\beta = q + \pi/q$ is $\pm 2\sqrt{q}$. In all other cases, $|\beta| < 2\sqrt{q}$, so $\mathbb{Q}(\pi_E)$ is imaginary quadratic, and $m_E = f_E = X^2 - \beta X + q$.

Conversely, for such a polynomial m_E to correspond to an isogeny class of *elliptic curves*, it is necessary and sufficient that $e = 1$ in Theorem 1.3.12, in other words all the invariants of D should be zero, in other words a should divide the p-adic valuation of $P(0)$ for every irreducible factor P of $X^2 - \beta X + q$ over $\mathbb{Q}_p$. At this point, we enter a case-by-case analysis of the splitting behavior of p in $\mathbb{Q}(\sqrt{\beta^2 - 4q})$ (the splitting field of $X^2 - \beta X + q$); for details, see [Wat69, Thm. 4.1]. □

In all cases of Proposition 1.3.18 where E is supersingular, one can further say exactly over which extension of $\mathbb{F}_q$ the curve E will acquire all its endomorphisms: this extension has degree at most 4 if $p \geq 5$, and at most 6 for $p \in \{2, 3\}$ [Wat69, p. 537].

Remark 1.3.19. In particular, ss.def. elliptic curves over $\mathbb{F}_q$ (assuming q is a square) form exactly two isogeny classes: the *maximal* curves, for which $\beta = -2\sqrt{q}$, and the *minimal* curves, for which $\beta = +2\sqrt{q}$. (This terminology comes from the number of points on the curves). After base-changing to $\mathbb{F}_{q^2}$, the elliptic curves in those two classes become part of one single isogeny class (the class of minimal curves), as $\pi_E^2 = (\pm\sqrt{q})^2 = +\sqrt{q^2}$. The isogeny class of maximal ss.def. elliptic curves over $\mathbb{F}_{q^2}$ also exists, but doesn't contain any curve base-changed from $\mathbb{F}_q$.

The previous fact might be surprising to the reader used to think that all supersingular elliptic curves are defined over $\mathbb{F}_{p^2}$: this is true (as we will see), but only up to $\overline{\mathbb{F}}_p$-isomorphism. We can reconcile the two points of view by saying that each $\overline{\mathbb{F}}_p$-isomorphism class of supersingular elliptic curves over $\overline{\mathbb{F}}_p$ contains exactly one ss.def. *minimal* (or maximal) supersingular curve over $\mathbb{F}_{p^2}$, and next, work with isogenies and endomorphisms from that elliptic curve that are defined over $\mathbb{F}_{p^2}$. This is what most protocols in isogeny-based cryptography actually do.

1.4 Isomorphism classes within an isogeny class

The next step, after classifying abelian varieties over finite fields up to isogeny as in Honda–Tate theory, is to delve into the structure of one fixed isogeny class. In the case of supersingular elliptic curves and ordinary abelian varieties respectively, we would like to derive Theorem 1 (the Deuring correspondence) and Theorem 2 (the CM action) respectively. Describing the structure of all isogeny classes over finite fields in such a "purely arithmetic" way is still, in general, an unsolved (and perhaps hopeless) problem, although one that is computationally approachable: see e.g. [Mar24].

The main tactic to study one fixed isogeny class is to construct isogenies from ideals of the various endomorphism *rings* (as opposed to algebras) appearing in the isogeny class. Describing the structure of the isogeny class can then be split into two sub-problems:

1. Classifying exactly which endomorphism rings can occur in the isogeny class, and
2. Understanding how isogenies arising from ideals (or not!) connect abelian varieties with the same and/or different endomorphism rings.

At this point, this strategy is still quite vague, but it should become clearer with examples.

The two sub-problems above are solved with the help of Tate's isogeny theorem. Note that throughout, we are manipulating lattices (with extra structure) in finite-dimensional vector spaces over $\mathbb{Q}$ (namely endomorphism rings and ideals, seen inside the endomorphism algebra.) The *local-global principle* for lattices asserts that a lattice L in such a vector space V is determined by the collection of its localizations $L \otimes_{\mathbb{Z}} \mathbb{Z}_\ell \subset V \otimes_{\mathbb{Q}} \mathbb{Q}_\ell$, as ℓ runs through all prime numbers. Luckily, Tate's isogeny theorem exactly describes what the localization of these endomorphism lattices are. We can already see here that including $\ell = p$ (the characteristic of the base field k) in the picture with Theorem 1.2.29 is vital: otherwise the local-global principle doesn't work. Since we do not look closely at Dieudonné modules in this document, we will necessarily miss some key parts in the proofs. The main reference for this section is Waterhouse's classical thesis paper [Wat69].

1.4.1 Isogenies from ideals in endomorphism rings

Throughout, when we say that I is an ideal in $\mathrm{End}(A)$, we mean that I is a *left* ideal (in case $\mathrm{End}(A)$ is not commutative) and we assume that I is also a lattice in $\mathrm{End}(A)$ (i.e. it has full rank over $\mathbb{Z}$). This is equivalent to requiring that I contains an isogeny. When A is simple, this just means $I \neq \{0\}$.

Definition 1.4.1. Let A be an abelian variety over any field k, and let I be an ideal in $\mathrm{End}(A)$. We define

$$H(I) = \bigcap_{\alpha \in I} \ker(\alpha)$$

as a subgroup scheme; it is a finite subgroup scheme of A. We denote the quotient isogeny by

$$\phi_I : A \to A/H(I).$$

We also write $H_A(I)$ for $H(I)$ and $\phi_{A,I}$ for ϕ_I when disambiguation is necessary.

In general, there is no reason to assume that ϕ_I is separable. For instance, if I is a principal ideal of the form $\mathrm{End}(A)\alpha$, then $H(I) = \ker(\alpha)$ and $\phi_I = \alpha$, even if α is not separable. This is why we must formulate Definition 1.4.1 in terms of group schemes instead of groups of $\overline{k}$-points.

Conversely, starting from an isogeny with domain A (i.e. a subgroup scheme K of A via the correspondence of Proposition 1.1.11), one can look at the ideal of $\mathrm{End}(A)$

$$I = \{\alpha \in \mathrm{End}(A) : K \subset \ker(\alpha)\}.$$

If K is of the form $H(I)$, then this construction doesn't necessarily recover I (but we always recover an ideal containing I.)

Definition 1.4.2. We say that $I \subset \mathrm{End}(A)$ is a *kernel ideal* if

$$I = \{\alpha \in \mathrm{End}(A) : H(I) \subset \ker(\alpha)\}.$$

It's easy to see that every ideal I is included in some kernel ideal J such that $H(I) = H(J)$, so we don't lose much if we only consider the above construction for kernel ideals.

Since we'll be using the local-global principle everywhere for the study of these isogenies, we record how they behave in terms of Tate modules.

Lemma 1.4.3. *We have*

$$V_\ell(\phi_I)^{-1}(T_\ell(A/H(I))) = \bigcap_{\alpha \in I} V_\ell(\alpha)^{-1}(T_\ell(A)).$$

Proof. For each $\alpha \in I$, the preimage $V_\ell(\alpha)^{-1}(T_\ell(A))$ is the overlattice of $T_\ell(A)$ corresponding to the finite subgroup $\ker(\alpha)[\ell^\infty]$ of A by Proposition 1.2.5. Since

$$H(I)[\ell^\infty] = \bigcap_{\alpha \in I} \ker(\alpha)[\ell^\infty],$$

we see that the overlattice of $T_\ell(A)$ corresponding to $H(I)[\ell^\infty)$ is precisely

$$\bigcap_{\alpha \in I} V_\ell(\alpha)^{-1}(T_\ell(A)).$$

This overlattice is also $V_\ell(\phi_I)^{-1}(T_\ell(A/H(I)))$, again by Proposition 1.2.5. □

There is a similar formula for the action of ϕ_I on Dieudonné modules [Wat69, Prop. 3.8].

One convenient feature of this construction is that is works well with respect to the (compatible) multiplication of ideals. By Proposition 1.1.25, if $\phi : A \to B$ is any isogeny, we can identify $\mathrm{End}(B)$ with a subring of $\mathrm{End}^0(A)$ as follows: choose $n \geq 1$ and an isogeny $\psi : B \to A$ such that $\phi_I \circ \psi = [n]_A$ and $\psi \circ \phi_I = [n]_B$. Then we define the map

$$\begin{array}{ccccc} \eta & \mathrm{End}(B) & \to & \mathrm{End}^0(A) \\ & \alpha & \mapsto & \frac{1}{n}\psi \circ \alpha \circ \phi_I \end{array}$$

This map is independent of the choice of n and ψ.

Lemma 1.4.4. *Let I be an ideal in $\mathrm{End}(A)$, let $B = A/H(I)$, and let J be an ideal in $\mathrm{End}(B)$. Let $\eta : \mathrm{End}(B) \hookrightarrow \mathrm{End}^0(A)$ be the map constructed as above. Then $I\eta(J)$ is an ideal of $\mathrm{End}(A)$, and the composite map $\phi_{B,J} \circ \phi_{A,I}$ is precisely $\phi_{A,I\eta(J)}$.*

Proof. We use the local-global principle: for each prime ℓ, we check that the overlattice of $T_\ell(A)$ given by $\phi_{A,I}^{-1} \circ \phi_{B,J}^{-1}(T_\ell(B/H_B(J)))$ is the overlattice corresponding to $H_A(I\eta(J))$. (A similar direct computation will work for $\ell = p$, but we omit it.) We compute:

$$\begin{aligned} \phi_{A,I}^{-1} \circ \phi_{B,J}^{-1}(T_\ell(B/H_B(J))) &= \phi_{A,I}^{-1}(\bigcap_{\tau \in J} V_\ell(\tau)^{-1}(T_\ell(B))) \\ &= \bigcap_{\tau \in J} V_\ell(\phi_{A,I})^{-1} \circ V_\ell(\tau)^{-1}(T_\ell(B)) \\ &= \bigcap_{\tau \in J} V_\ell(\eta(\tau))^{-1} V_\ell(\phi_{A,I}^{-1}(T_\ell(B))) \\ &= \bigcap_{\tau \in J, \sigma \in I} V_\ell(\eta(\tau))^{-1} V_\ell(\sigma)^{-1}(T_\ell(A)) \\ &= \bigcap_{\alpha \in I\eta(J)} V_\ell(\alpha)^{-1}(T_\ell(A)). \end{aligned}$$

This is precisely the overlattice of $T_\ell(A)$ corresponding to $H_A(I\eta(J))$ by Lemma 1.4.3. □

Our stated aim is to span the isogeny class of A using isogenies of the above form. Therefore, we should study when the codomains $A/H(I)$ and $A/H(J)$ are isomorphic. We have already seen that if I is principal, then $A/H(I)$ is isomorphic to A. More generally, we call two ideals I and J *equivalent* if $I = J\lambda$ for some $\lambda \in \mathrm{End}^0(A)$; an ideal I is principal if and only if it is equivalent to the trivial ideal $\mathrm{End}(A)$. This notion is compatible with kernel ideals:

Lemma 1.4.5. *Let $I \subset \mathrm{End}(A)$ be a kernel ideal. For any $\alpha \in \mathrm{End}(A)$, $I\alpha$ is also a kernel ideal.*

Proof. Assume $H(I\alpha) \subset \ker(\lambda)$ for some $\lambda \in \mathrm{End}(A)$. Then $\ker(\alpha) \subset H(I\alpha) \subset \ker(\lambda)$, so there exists $\mu \in \mathrm{End}(A)$ such that $\lambda = \mu\alpha$. Because α is surjective, the assumption $H(I\alpha) \subset \ker(\mu\alpha)$ is equivalent to $H(I) \subset \ker(\mu)$, i.e. $\mu \in I$ because I is a kernel ideal, so $\lambda \in I\alpha$. □

Proposition 1.4.6. *1. If I and J are equivalent $\mathrm{End}(A)$-ideals, then $A/H(I)$ and $A/H(J)$ are isomorphic as abelian varieties.*

2. If $A/H(I) \simeq A/H(J)$ and if both I and J are kernel ideals, then I and J are equivalent.

Proof. 1. Let $\lambda \in \mathrm{End}^0(A)$ such that $I = J\lambda$, and fix an integer $N \geq 1$ such that $N\lambda \in \mathrm{End}(A)$. Then we have $NI = J(N\lambda)$. By Lemma 1.4.4, we can factor $\phi_{NI} = \phi_{J(N\lambda)}$ in two different ways:

$$A \xrightarrow{N} A \xrightarrow{\phi_I} A/H(I), \qquad \text{and}$$

$$A \xrightarrow{N\lambda} A \xrightarrow{\phi_J} A/H(J).$$

Thus $A/H(I)$ and $A/H(J)$, as codomains of isogenies with the same kernel, are isomorphic by Proposition 1.1.11.

2. Assume $A/H(I)$ and $A/H(J)$ are isomorphic; we identify both codomains with a fixed abelian variety B. First pick $N \geq 1$ large enough so that $H(J) \subset [N]^{-1}H(I) = \ker(N\phi_I)$. By isogeny factorization (Proposition 1.1.12), there exists an isogeny $\psi : B \to B$ such that $N\phi_I = \psi \circ \phi_J$. Next, we pick $M \geq 1$ large enough such that $M\psi \circ \phi_J = \phi_J \circ \alpha$ for some $\alpha \in \mathrm{End}(A)$: this is possible because $\phi_J^{-1} \circ \psi \circ \phi_J \in \mathrm{End}^0(A)$. Then we have $MN\phi_I = \psi_J \circ \alpha$, in other words $H(MNI) = H(J\alpha)$. Thus $MNI = J\alpha$ as both are kernel ideals by Lemma 1.4.5, hence I and J are equivalent. □

Finally, when I is a kernel ideal, we can describe the endomorphism ring of $A/H(I)$:

Proposition 1.4.7. *Let $I \subset \mathrm{End}(A)$ be a kernel ideal. Then the endomorphism ring of $A/H(I)$, seen as a subring of $\mathrm{End}^0(A)$, is precisely the right order of I.*

Proof. We again use the local-global principle: for each prime $\ell \neq p$, we check that the endomorphism ring of $A/H(I)$ equals the right order of I after tensoring with $\mathbb{Z}_\ell$. A similar proof will work at $\ell = p$. By Tate's isogeny theorem 1.2.11, the subring of $\mathrm{End}(A/H(I)) \otimes \mathbb{Z}_\ell \subset \mathrm{End}^0(A) \otimes_{\mathbb{Q}} \mathbb{Q}_\ell$ consists of those elements α such that

$$V_\ell(\alpha)\Big(\bigcap_{\tau \in I} V_\ell(\tau)^{-1}(T_\ell(A))\Big) \subset \bigcap_{\tau \in I} V_\ell(\tau)^{-1}(T_\ell(A)).$$

Equivalently, for each $\sigma \in I$, we have

$$V_\ell(\sigma\alpha)\Big(\bigcap_{\tau \in I} V_\ell(\tau)^{-1}(T_\ell(A))\Big) \subset T_\ell(A). \tag{1}$$

In particular, $V_\ell(\sigma\alpha)(T_\ell(A)) \subset T_\ell(A)$, so $\sigma\alpha \in \mathrm{End}(A)$. (We have used the local-global principle here, as this holds for all ℓ.) If α lies in the right order of I, then (1) is satisfied for each σ, as $\sigma\alpha \in I$. Conversely, if the above condition holds, then $I + I\alpha$ is an ideal of $\mathrm{End}(A)$ with $H(I + I\alpha) = H(I)$. Because I is a kernel ideal, we have $I\alpha \subset I$, i.e. α lies in the right order of I. □

1.4.2 The case of maximal orders

If the endomorphism rings of the abelian varieties we consider are maximal orders in the endomorphism algebra, then the above construction of isogenies from ideals is particularly powerful.

Proposition 1.4.8. *Let A be an abelian variety over a finite field k, and let S be any maximal order in $\mathrm{End}^0(A)$. Then there exists another abelian variety B in the isogeny class of A whose endomorphism ring is isomorphic to S.*

Proof. Since S is a lattice in $\mathrm{End}^0(A)$, there exists an integer $N \geq 1$ such that $NS \subset \mathrm{End}(A)$. Let $I = \mathrm{End}(A)NS$ be the ideal generated by NS. Then the right order of I contains S, so equals it. The endomorphism ring of $A/H(I)$ is S by Proposition 1.4.7. □

We omit the proof of the following theorem: see [Wat69, Thm. 3.15].

Theorem 1.4.9. *Let A be an abelian variety over a finite field k, and assume that $\mathrm{End}(A)$ is a maximal order. Then every ideal I of $\mathrm{End}(A)$ is a kernel ideal, the degree of ϕ_I equals the reduced norm of I, and the endomorphism ring of $A/H(I)$ is also maximal.*

The reduced norm of I is computed as follows. We know that $\mathrm{End}^0(A)$ is a direct sum of simple algebras, so a maximal order in $\mathrm{End}^0(A)$ is simply a direct sum of maximal orders in each component, and an $\mathrm{End}(A)$-ideal is a direct sum of left ideals in each component. If J is an ideal in a maximal order S of a simple algebra of dimension e^2 over its center, then $\#(S/J)$ is always an eth power, and we call its positive eth root the reduced norm of J. The reduced norm of I is then the product of the reduced norms of each of its components.

Given Theorem 1.4.9, we observe that not all isogenies are always of the form ϕ_I: for instance, if E is an elliptic curve such that $\mathrm{End}(E)$ is non-maximal, then there exists an isogeny $\phi : E \to E'$ with $\mathrm{End}(E')$ maximal, and the dual isogeny $E' \to E$ is not of the form ϕ_I.

1.4.3 The main theorems on isogeny classes

We now sketch the proofs of the two main theorems 1 and 2, starting with the latter. Recall the definition of a principal homogeneous space (PHS): we say that a set X, endowed with an action of a group G, is a *principal homogeneous space* for G if for any $x, x' \in X$, there exists a unique $g \in G$ such that $g(x) = x'$. In other words, the action of G on X is simply transitive.

Sketch of proof of Theorem 2. When $\mathrm{End}^0(A)$ is commutative, then we can identify endomorphism rings of all abelian varieties isogenous to A as subrings of $\mathrm{End}^0(A)$, simultaneously and in a compatible way.

First, we explain why the class group of $\mathrm{End}(A)$ acts freely on abelian varieties in the isogeny class with endomorphism ring R. Let I be an invertible ideal in R. By Proposition 1.4.7, the endomorphism ring of $A/H(I)$ is still R; by Proposition 1.4.6, $A/H(I)$ (up to isomorphism) does not depend on the class of I in the class group, and $A/H(I) \simeq A$ if and only if I lies in the trivial class; and by Lemma 1.4.4, we indeed have a group action.

On the other hand, the fact that there is only one orbit involves a careful analysis of lattices in Tate and (especially) Dieudonné modules, and we do not give a complete proof; see [Wat69, Thms. 4.5 and 7.2]. The similarity between the assumptions $g = 1$ and R maximal that makes the proof work is the following: if $\ell \neq p$ is a prime, then any rank 1 module over $R \otimes \mathbb{Z}_\ell$ (that is not left stable by a bigger order, when $g = 1$) is free over $R \otimes \mathbb{Z}_\ell$. □

As a complement, we can exactly describe what the possible endomorphism rings are in the ordinary case. We also omit the proof, which is heavily reliant on Dieudonné modules.

Theorem 1.4.10 ([Wat69, Thm. 7.4])**.** *Let A be simple and ordinary. Then the endomorphism rings occurring in the isogeny class of A are all orders in $\mathrm{End}^0(A)$ containing π and $q\pi^{-1}$.*

However, we can prove the following weaker statement:

Proposition 1.4.11. *Let A be an abelian variety such that $\mathrm{End}^0(A)$ is commutative. Let $R \subset \mathrm{End}^0(A)$ be any order containing π_A. Then there exists an abelian variety B in the isogeny class of A such that $\mathrm{End}(B) \otimes \mathbb{Z}_\ell = R \otimes \mathbb{Z}_\ell$ for every prime $\ell \neq p$.*

To relate this result with Theorem 1.4.10, note that the index of $\mathbb{Z}[\pi]$ inside $\mathbb{Z}[\pi, q\pi^{-1}]$ is a power of p: indeed $\pi\mathbb{Z}[\pi, q\pi^{-1}] \subset \mathbb{Z}[\pi]$, and π has norm q.

Proof. Because $\mathrm{End}(A)$ and R are lattices in the same $\mathbb{Q}$-vector space, there exist only finitely many primes ℓ such that $\mathrm{End}(A) \otimes \mathbb{Z}_\ell \neq R \otimes \mathbb{Z}_\ell$. Let ℓ be such a prime. Then there exists a lattice $\Lambda_\ell \subset V_\ell(A)$ such that the order stabilizing Λ_ℓ is exactly $R \otimes \mathbb{Z}_\ell$. (Take $\Lambda_\ell = (R \otimes \mathbb{Z}_\ell)v$ for some $v \in V_\ell(A)$.) We can also assume that Λ_ℓ contains $T_\ell(A)$. By Proposition 1.2.5, Λ_ℓ corresponds to the ℓth Tate module of the codomain B of some isogeny of ℓ-power degree $\phi : A \to B$; at all other primes, the Tate modules of A and B are the same lattices in $V_\ell(A)$. Repeating this construction at all other primes such that $\mathrm{End}(A) \otimes \mathbb{Z}_\ell \neq R \otimes \mathbb{Z}_\ell$, we find the required abelian variety B. □

In the case of supersingular elliptic curves, the only missing ingredient to prove the Deuring correspondence 1 is the following.

Theorem 1.4.12. *Let E be a ss.def. elliptic curve. Then $\mathrm{End}(E)$ is a maximal order, and every isogeny $\phi : E \to E'$ is isomorphic to ϕ_I for some ideal I of $\mathrm{End}(E)$.*

Proof. This is again proved using the local-global principle and an explicit description of Dieudonné modules (an order is maximal if and only if it is maximal everywhere locally.) At primes $\ell \neq p$, we have $\mathrm{End}^0(E) \otimes \mathbb{Q}_\ell = \mathrm{Mat}_{2\times 2}(\mathbb{Q}_\ell)$ by Tate's isogeny theorem, since the Frobenius endomorphism is a scalar. The stabilizer in $\mathrm{Mat}_{2\times 2}(\mathbb{Q}_\ell)$ of any lattice in $\mathbb{Q}_\ell^2$ is conjugate to $\mathrm{Mat}_{2\times 2}(\mathbb{Z}_\ell)$ (by choosing a basis of the lattice), hence maximal; moreover any other lattice is the stabilizer of some ideal. We omit the proof at $\ell = p$: see [Wat69, Thm. 4.5]. □

Proof of Theorem 1. The classification of isogenies from a fixed elliptic curve E in terms of ideals in its endomorphism is provided by Propositions 1.4.6 and 1.4.7 and Theorems 1.4.9 and 1.4.12.

Next, we study the map $E \mapsto \mathrm{End}(E)$. Two isomorphic maximal orders in $B_{p,\infty}$ are conjugate by the Skolem–Noether theorem [Voi21, Thm. 7.7.1], so by Theorem 1.4.12, the map $E \mapsto \mathrm{End}(E)$ (as an order in $B_{p,\infty}$ up to conjugation) is well-defined. Let E_0 be any minimal supersingular elliptic curve, and let $\mathcal{O}_0 \subset B_{p,\infty}$ be isomorphic to $\mathrm{End}(E_0)$. If $\mathcal{O} \subset B_{p,\infty}$ is any other maximal, there exists an $\mathcal{O}_0$-ideal whose right order is $\mathcal{O}$ by [Voi21, Lemma 17.4.6 and §23]. By the previous facts, the codomain of $\phi_I : E_0 \to E'$ has endomorphism ring isomorphic to $\mathcal{O}$, so the map is surjective. On the other hand, let E and E' be two supersingular, minimal elliptic curves such that $\mathrm{End}(E)$ and $\mathrm{End}(E')$ are isomorphic to the same order $\mathcal{O}$. By Theorem 1.3.1, E and E' are isogenous, so there exists an ideal I of $\mathcal{O}$ whose right order is conjugate to $\mathcal{O}$. Up to replacing I by an equivalent ideal, we can assume that the right order of I is also $\mathcal{O}$, i.e. I is a two-sided ideal. Then we know that I is an integral multiple of either (1) or the ideal $\mathfrak{p}$ corresponding to the p-Frobenius morphism $E \to E^{(p)}$ [Voi21, Thm. 18.3.6], so either $E' \simeq E$ or $E' \simeq E^{(p)}$. In particular, there exists only one isomorphism class of supersingular elliptic curves with order $\mathcal{O}$ if and only if E is the base-change of an elliptic curve over $\mathbb{F}_p$.

Finally, we prove that any isomorphism class of supersingular elliptic curves over $\overline{\mathbb{F}}_p$ contains a single minimal elliptic curve defined over $\mathbb{F}_{p^2}$. Let E be an elliptic curve defined over a finite extension k of $\mathbb{F}_p$, and let $E_0/\mathbb{F}_{p^2}$ be any minimal supersingular elliptic curve. After extending k, we can assume that E/k is ss.def. and minimal. By Theorem 1.3.1, E and E_0 are isogenous over k; by the previous results, there exists an ideal I of $\mathrm{End}(E_0)$ whose codomain is k-isomorphic to E'. But the endomorphisms of E_0 are all defined over $\mathbb{F}_{p^2}$, so ϕ_I and its codomain are defined over $\mathbb{F}_{p^2}$. Conversely, let E and E' be minimal, supersingular elliptic curves over $\mathbb{F}_{p^2}$, and choose an ideal I such that $\phi_I : E \to E'$ is an isogeny over $\mathbb{F}_{p^2}$. If E and E' become isomorphic after a finite extension $k/\mathbb{F}_{p^2}$, then applying the above theory over k shows that I is principal, hence E and E' are isomorphic over $\mathbb{F}_{p^2}$. □

1.5 Examples of isogeny graphs

Isogeny graphs of abelian varieties over finite fields are a combinatorial way of representing isogeny classes of those abelian varieties and the structure of isogenies within the class. A vertex in such a graph G typically represent a k-isomorphism class of abelian varieties over a fixed finite field k, while an edge between two vertices represented by abelian varieties A and B corresponds to an isogeny $\phi : A \to B$. Often, the isogenies appearing as edges are of a certain type only – for instance, ℓ-isogenies between elliptic curves. Thus G is, a priori, a finite directed graph where loops and multiple edges are allowed.

1.5.1 Supersingular ℓ-isogeny graphs

Definition 1.5.1. Fix distinct primes ℓ and p. The *supersingular ℓ-isogeny graph* over $\mathbb{F}_{p^2}$ is the finite graph G whose vertices are minimal, supersingular elliptic curves over $\mathbb{F}_{p^2}$ up to isomorphism, and whose edges correspond to ℓ-isogenies defined over $\mathbb{F}_{p^2}$, up to isomorphism (i.e. post-composition by an isomorphism on the codomain.)

By the Deuring correspondence (Theorem 1), the structure of the supersingular ℓ-isogeny graph is completely encoded in terms of quaternions.

In certain cases, it is convenient to consider the supersingular ℓ-isogeny graph as a regular, undirected graph. This is possible only when $p = 1 \mod 12$: in that case, the elliptic curves with j-invariant 0 and 1728 over $\overline{\mathbb{F}}_p$ are not supersingular (because p splits in both $\mathbb{Q}(i)$ and $\mathbb{Q}(\zeta_3)$), so any supersingular elliptic curve $E/\mathbb{F}_{p^2}$ satisfies $\mathrm{Aut}(E) = \{\pm 1\}$.

Proposition 1.5.2. *Assume that $p = 1 \mod 12$. Then the supersingular ℓ-isogeny graph is an undirected, $\ell + 1$-regular graph with $(p-1)/12$ vertices.*

Proof. To show that G is undirected, we check that for any two vertices E, E' of G, taking duals induces a bijection between the sets of ℓ-isogenies $E \to E'$ and $E' \to E$ up to isomorphism. But because of the above fact on automorphism groups, $\phi : E \to E'$ is isomorphic to $\psi : E \to E'$ if and only if $\psi = \pm\psi$, if and only if $\phi^\vee = \pm\psi^\vee$, if and only if $\phi^\vee$ and $\psi^\vee$ are isomorphic.

Then, the graph G is $\ell + 1$-regular because for every elliptic curve E, there are $\ell + 1$ lines in $E[\ell](\overline{k})$, hence $\ell + 1$ different ℓ-isogenies with domain E by Proposition 1.1.10. These isogenies can all be defined over $\mathbb{F}_{p^2}$ by Theorem 1.

The number of vertices comes from the mass formula in [Voi21, Thm. 25.1.1]. □

Note that for any p, the number of vertices in the supersingular ℓ-isogeny graph is $p/12 + O(1)$ by the same mass formula.

The correspondance with quaternions is a very powerful tool to study the structure of supersingular isogeny graphs. For instance, the major theorem that those graphs are Ramanujan (see §2.2.3) is proved in that way. As an easier example, one can use facts about quaternion algebras (namely the strong approximation theorem) to show that the supersingular ℓ-isogeny graph is connected:

Proposition 1.5.3 ([Voi21, Prop/ 28.4.17]). *Let $\mathcal{O} \subset B_{p,\infty}$ be a maximal order, let I be a left $\mathcal{O}$-ideal, and let $\ell \neq p$ be any prime. Then there exists $n \geq 1$ and an $\mathcal{O}$-ideal J, equivalent to I and of reduced norm ℓ^n for some $n \geq 1$. Consequently, the supersingular ℓ-isogeny graph over $\mathbb{F}_{p^2}$ is connected.*

Strong approximation can also be used to show that the 2-dimensional analogues of these graphs using superspecial abelian surfaces are connected [JZ23].

Remark 1.5.4. The supersingular ℓ-isogeny graph G' is closely related to the *Brandt class groupoid* of the quaternion algebra, defined in [Voi21, p. 19.5.4]. Define the *ℓ-Brandt graph* G' of $B_{p,\infty}$ in the obvious way: vertices in this graph correspond to maximal orders in $B_{p,\infty}$ up to conjugation, and edges correspond to ideals of reduced norm ℓ connecting (conjugates of) these orders, up to equivalence. Then, by the Deuring correspondence, there is a well-defined map $G \to G'$ under which each vertex of G' has either 1 or 2 preimages, depending on whether the corresponding elliptic curve is defined over $\mathbb{F}_p$ or not.

1.5.2 Isogeny volcanoes of ordinary elliptic curves

Next, we consider isogeny graphs of ordinary elliptic curves over any finite field $k = \mathbb{F}_q$. Fix an isogeny class V of such elliptic curves; by Theorem 1.3.1, it consists of all elliptic curves with a common number of points $q + 1 - t$ (up to isomorphism.)

Definition 1.5.5. The *ℓ-isogeny graph* G on the isogeny class V is the graph whose vertices are the elements of V, and whose edges correspond to ℓ-isogenies up to post-composition by an isomorphism.

Fix a base point $E \in V$; we are interested in the structure of the connected component G_E of E in G. Let $R = \mathrm{End}(E)$. As in §1.4.3, we can identify the endomorphism rings of the other vertices in G_E with orders in the quadratic imaginary number field $F = R \otimes \mathbb{Q}$. If R' is the endomorphism ring of another vertex E' in G_E, then $R \otimes \mathbb{Z}_{\ell'} = R' \otimes \mathbb{Z}_{\ell'}$ for any prime $\ell' \neq \ell$ by Tate's isogeny theorem. As a consequence, R' is one of the orders occurring in the tower

$$R_{\min} = R_n \subset R_{n-1} \subset \cdots \subset R_i = R \subset R_{i-1} \subset \cdots \subset R_0 = R_{\max}$$

where:

- $[R_{i-1} : R_i] = \ell$ for each $1 \leq i \leq n$,
- the conductor c_0 of R_0 is prime to ℓ,
- for each $0 \leq i \leq n$, the conductor of R_i is $c_i = \ell^i c_0$,
- n is the ℓ-adic valuation of the conductor $\mathbb{Z}[\pi_E]$ of discriminant $t^2 - 4q$,
- all the orders $R_n, \ldots, R_0$ occur as endomorphism rings in G_E.

These properties can all be derived from Theorem 1.4.10 and facts about orders in quadratic fields. We say that a vertex is *at level* i in G_E if its endomorphism ring is R_i.

The question is now: how are vertices with different endomorphism rings connected in G_E? First, Theorem 2 gives us a precise control on how many vertices in G_E have endomorphism ring $R_0 = R_{\max}$. There are three cases:

1. If ℓ is split in F, then there exists two invertible R_0-ideals $\mathfrak{l}$ and $\bar{\mathfrak{l}}$ of norm ℓ. Since $\mathfrak{l}\bar{\mathfrak{l}} = (\ell)$, the ideal $\bar{\mathfrak{l}}$ is the inverse of $\mathfrak{l}$ in the class group of R_0. Let e be the order of $\mathfrak{l}$ in this class group. By Theorem 2, there are exactly e vertices at level 0 in G_E, connected in a cycle of $\mathfrak{l}$-isogenies.

2. If ℓ is ramified in F, the situation is similar, but this time $\bar{\mathfrak{l}} = \mathfrak{l}$. There is either one or two vertices at level 0 depending on whether $\mathfrak{l}$ is principal or not, connected by a single ℓ-isogeny (which could be a loop.)

3. If ℓ is inert in F, there are no invertible R_0-ideals of norm ℓ, so there is a single vertex at level 0.

In all cases, if $i > 1$, then there exist no invertible R_i-ideals of norm ℓ, so no ℓ-isogenies between vertices at level i.

It turns out that the other edges in G_E are arranged in the shape of an *ℓ-volcano*, as in the following definition, adapted from [BJW17, Def. 1.1]. This volcano structure was first explicited in [Koh96, Prop. 23].

Definition 1.5.6. Let ℓ be a prime and $n \geq 0$. An *ℓ-volcano of depth n* is an undirected, connected graph whose vertices are partitioned into *levels* $V_0, \ldots, V_n$ such that:

- the subgraph V_0 is a finite regular graph of degree at most 2,
- for each $i < n$, there are $\ell + 1$ edges going out of each vertex in V_i,
- for each $i > 0$, each vertex in V_i has exactly one neighbor in V_{i-1}, and such edges cover all edges in the graph that are not in V_0.

Examples are shown on Figure 1.

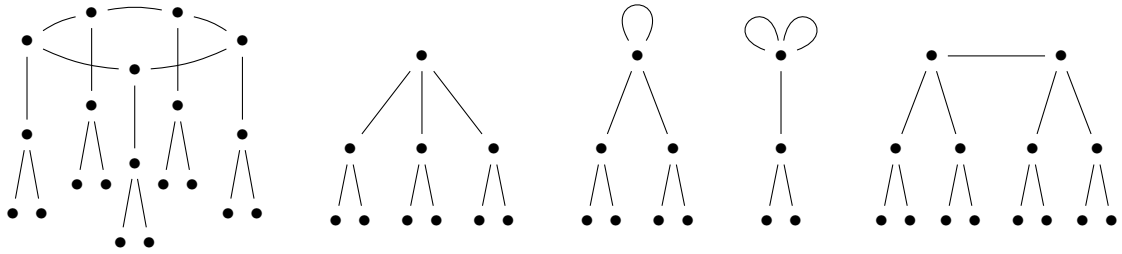

Figure 1: some 2-volcanoes of depth 2

Proposition 1.5.7. *Let $q = p^r$ be a prime power, let $\ell \neq p$ be a prime, let $E/\mathbb{F}_q$ be an ordinary elliptic curve, and let G_E be the connected component of E in the ℓ-isogeny graph defined above. Let n be the ℓ-adic valuation of the conductor of $\mathbb{Z}[\pi_E]$. Assume further that G_E does not contain any elliptic curve with j-invariant 0 or 1728. Then G_E is an ℓ-volcano of depth n, whose i-th level consists of elliptic curves with endomorphism ring R_i as defined above, for any $0 \leq i \leq n$.*

Proof. The assumption on j-invariants guarantees that the elliptic curves appearing as vertices in G_E have automorphism group $\{\pm 1\}$, so G_E is undirected. The structure of V_0 is given by Theorem 2.

Let E' be a vertex at level $i > 0$. Then we can view $T_\ell(E')$ as a lattice inside $V_\ell(E)$ whose stabilizer is the order $R_i \otimes \mathbb{Z}_\ell$, by Tate's isogeny theorem. The neighbors of E' in G_E correspond to elliptic curves whose Tate modules are overlattices of index ℓ in $T_\ell(E')$. There is exactly one such lattice whose stabilizer is $R_{i-1} \otimes \mathbb{Z}_\ell$, namely $(R_{i-1} \otimes \mathbb{Z}_\ell)T_\ell(E')$. This overlattice is Galois-invariant because the endomorphisms of E are defined over $\mathbb{F}_q$. Thus, there is at least one edge from E' to a vertex in V_{i-1}.

Next, we consider the ℓ-isogeny graph G on the whole isogeny class of E, similarly partitioned into levels $V_0', \ldots, V_n'$. Let δ be the number of outgoing edges from any vertex in V_0'. Since any vertex in G has at most $\ell + 1$ outgoing edges, we have from the previous considerations

$$\#V_1 \leq (\ell + 1 - \delta)\#V_0 \quad \text{and} \quad \#V_{i+1} \leq \ell \#V_i \text{ for each } i > 0. \tag{2}$$

By Theorem 2, we also know that for each $0 \leq i \leq n$,

$$\#V_i = \#\operatorname{Cl}(R_i).$$

However, we also know from the study of class groups of quadratic orders [Koh96, (4.2) p. 41] that

$$\#\operatorname{Cl}(R_1) = (\ell + 1 - \delta)\#\operatorname{Cl}(R_0) \quad \text{and} \quad \#\operatorname{Cl}(R_{i+1}) = \ell \#\operatorname{Cl}(R_i) \text{ for each } i > 0.$$

Hence all the inequality in (2) are equalities. Further, if $n > 0$, there is only one edge going out of every vertex in V_n because of Theorem 2, as there are no invertible R_n-ideals of norm ℓ. Thus G_E is an ℓ-volcano. □

As a consequence of Proposition 1.5.7, the ℓ-isogeny graph on constructed on the whole isogeny class of E is a union of finitely many identical ℓ-volcanoes: indeed, the shape of the level 0 vertices is completely determined by the splitting behavior of ℓ in the quadratic field F (identical across the isogeny class), and the volcano is entirely determined as a graph by its level 0 vertices and edges.

1.5.3 Cayley graphs of class groups

Another kind of isogeny graphs of ordinary abelian varieties over finite fields can be described in terms of Cayley graphs of finite groups.

Definition 1.5.8. Let H be a finite group, and $S \subset H$. The *Cayley graph* of H relative to S is the (a priori directed) graph G whose vertex set is H, with an edge from h to sh for every $h \in H$ and $s \in S$.

If S is left stable under the inverse map $s \mapsto s^{-1}$, then the Cayley graph of H relative to S. One sometimes labels the edges in the directed (resp. undirected) graph G by the corresponding

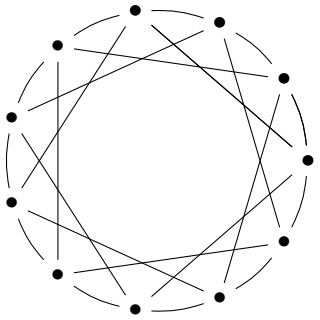

Figure 2: Cayley graph of $\mathbb{Z}/11\mathbb{Z}$ relative to $S = \{\pm 1, \pm 3\}$

elements s (resp. pairs $\{s, s^{-1}\}$). If H is an abelian group, then the Cayley graphs of H have a very regular shape, as shown in Figure 2. The Cayley graph G is connected if and only if S generates H.

Next, let $A/\mathbb{F}_q$ be an ordinary abelian variety of any dimension g, and let C denote the isogeny class of A. Then $\mathrm{End}^0(A)$ is a CM field (we are in case (1) of Theorem 1.3.3); assume that $R = \mathrm{End}(A)$ is the maximal order in $\mathrm{End}^0(A)$. Choose distinct primes $\ell_1, \ldots, \ell_r$ that are prime to the index of $\mathbb{Z}[\pi_A, q\pi_A^{-1}]$ in R. We consider the graph G whose vertex set is C and whose edges consist of isogenies of degree $d \in \{\ell_1, \ldots, \ell_r\}$.

Proposition 1.5.9. *Let $S \subset \mathrm{Cl}(R)$ be the set of classes represented by R-ideals of norm $\ell_1, \ldots, \ell_r$. Then G is isomorphic to the Cayley graph of $\mathrm{Cl}(R)$ relative to the set S.*

Note that the class group $\mathrm{Cl}(R)$ is abelian. A similar proposition would hold in the case of elliptic curves, even if their endomorphism ring is not the maximal order.

Proof. By Tate's isogeny theorem 1.2.11 and Theorem 1.4.10, we know that the endomorphism ring R' of any vertex of G satisfies $\mathbb{Z}[\pi_A, q\pi_A^{-1}] \subset R'$ and $R' \otimes \mathbb{Z}_\ell \simeq R \otimes \mathbb{Z}_\ell$ where ℓ is any prime distinct from $\ell_1, \ldots, \ell_r$. The above assumptions then force $R = R'$. The isogeny graph of G is isomorphic to the specified Cayley graph of $\mathrm{Cl}(R)$ by Theorem 2. □

1.5.4 Isogeny volcanoes in higher dimensions

We conclude part 1 of these notes with the following question: can we describe the structure of ℓ-isogeny graphs of ordinary abelian varieties A of dimension $g > 1$ over finite fields $\mathbb{F}_q$ when the index of $\mathbb{Z}[\pi_A, q\pi_A^{-1}]$ in the maximal order is divisible by ℓ?

This problem is unsolved in general, but we do again find volcano-like structures in certain specific cases, using Tate's isogeny theorem as in the proof of Proposition 1.5.7. We follow [BJW17].

The first step is to find a CM field in which certain orders of ℓ-power index in the maximal order are arranged in a tower and indexed by conductors, like in the setting of quadratic imaginary fields. This happens more generally for orders in a CM field containing the maximal order in the totally real subfield.

Proposition 1.5.10 ([BJW17, Thm. 2.1])**.** *Let K be a CM number field of degree $2g$, i.e. a totally imaginary quadratic extension of a totally real number field K^+ of degree g. Let $\mathbb{Z}_K$ (resp. $\mathbb{Z}_{K^+}$) denote the maximal order in K (resp. K^+). Then the map $f_+ \mapsto \mathbb{Z}_{K^+} + f_+\mathbb{Z}_K$ is a bijection between*

ideals in $\mathbb{Z}_{K^+}$ and orders in K containing $\mathbb{Z}_{K^+}$. The conductor of $\mathbb{Z}_{K^+} + f_+\mathbb{Z}_K$ is $f_+\mathbb{Z}_K$, and its intersection with $\mathbb{Z}_{K^+}$ is precisely f_+.

Corollary 1.5.11. *Let ℓ be a prime, and let $\mathcal{O} \subset K$ be any order such that $\mathcal{O}^+ := \mathcal{O} \cap K^+$ is maximal at ℓ, i.e. $\mathcal{O}^+ \otimes_{\mathbb{Z}} \mathbb{Z}_\ell$ is the maximal order in the $\mathbb{Q}_\ell$-algebra $K^+ \otimes_{\mathbb{Q}} \mathbb{Q}_\ell$. Then the orders $\mathcal{O}'$ in K such that $\mathcal{O}' \cap K^+ = \mathcal{O}^+$ and such that $\mathcal{O}' \otimes \mathbb{Z}_{\ell'} \simeq \mathcal{O} \otimes \mathbb{Z}_{\ell'}$ for all primes $\ell' \neq \ell$ are in bijection with ideals f_+ of $\mathbb{Z}_{K^+}$ supported at primes above ℓ, via*

$$f_+ \mapsto \text{the unique such } \mathcal{O}' \text{ satisfying } \mathcal{O}' \otimes_{\mathbb{Z}} \mathbb{Z}_\ell = (\mathbb{Z}_{K^+} + f_+\mathbb{Z}_K) \otimes_{\mathbb{Z}} \mathbb{Z}_\ell.$$

Proof. Combine Proposition 1.5.10 with the local-global principle. □

To find a tower of orders, we put ourselves in the setting of Corollary 1.5.11, and restrict attention to ideals f_+ that are powers of a single, fixed prime ideal $\mathfrak{l}$ of $\mathbb{Z}_{K^+}$ above ℓ. Thus, we fix an ordinary abelian variety A over $\mathbb{F}_q$ of dimension g, let K/K^+ be the CM field $\mathrm{End}^0(A)$, and assume that $\mathrm{End}(A) \cap K^+$ is maximal at ℓ.

Definition 1.5.12. We say that $\phi : A \to B$ is a $\mathfrak{l}$-*isogeny* if $\ker(\phi) \subset A[\mathfrak{l}]$ is a nontrivial, proper subgroup of $A[\mathfrak{l}]$ that is stable under $\mathrm{End}(A) \cap K^+$.

We know that $T_\ell(A)$ is a free $\mathbb{Z}_\ell$-module of rank $2g$, so $V_\ell(A)$ has the structure of a 2-dimensional vector space over $K^+ \otimes_{\mathbb{Q}} \mathbb{Q}_\ell$. Therefore, the kernel of an $\mathfrak{l}$-isogeny as in Definition 1.5.12 is a $\mathbb{Z}_{K^+}/\mathfrak{l}\mathbb{Z}_{K^+}$-vector space of dimension 1; this explains why $\mathfrak{l}$-isogenies are a convenient generalization of ℓ-isogenies between elliptic curves.

By considering Tate modules, we observe that when taking sequences of $\mathfrak{l}$-isogenies from A, the endomorphism rings $\mathcal{O}'$ of the abelian varieties we can reach all satisfy the properties of Corollary 1.5.11 [BJW17, Prop. 4.8]. A similar, but more technical, reasoning as in Proposition 1.5.7 leads to the following result.

Theorem 1.5.13 ([BJW17, Thm. 4.3])**.** *Assume for simplicity that $\mathbb{Z}_K^\times = \mathbb{Z}_{K^+}^\times$, in particular $\mathrm{Aut}(A) = \{\pm 1\}$. Then the connected $\mathfrak{l}$-isogeny graph G constructed from A is partitioned into levels $V_0, \dots, V_n$, where n denotes the valuation at $\mathfrak{l}$ of the conductor of $\mathbb{Z}_{K^+}[\pi_A]$: a vertex B is at level i if $\mathrm{End}(B) = \mathcal{O}_i$ where $\mathcal{O}_i$ corresponds to $f_+ = \mathfrak{l}^i$ under the bijection of Corollary 1.5.11. Moreover:*

1. *V_0 is isomorphic to the Cayley graph of the subgroup of $\mathrm{Cl}(\mathcal{O}_0)$ generated by the prime ideals of K above $\mathfrak{l}$;*
2. *There are $N(\mathfrak{l}) + 1$ edges going in and out of each vertex in V_i, for each $0 \leq i \leq n-1$;*
3. *For each $1 \leq i \leq n$, and each vertex B in V_i, there is a single edge from B to a vertex in V_{i-1}, and a single edge coming from a vertex in V_{i-1} to B; these cover all the edges outside V_0;*
4. *For each path $A \to B \to C$ in G where A, B lie at a common level i and B lies at level $i \pm 1$, we have $C \simeq A/A[\mathfrak{l}]$.*

In general, the isogeny graph G is directed. The theorem shows that G is directed if and only if it is an $N(\mathfrak{l})$-volcano of depth n as in Definition 1.5.6, if and only if $\mathfrak{l}$ is principal in $\mathcal{O}_n$. This happens, for instance, if $\mathfrak{l}$ is principal in $\mathrm{End}(A) \cap K^+$. An example where G is not a volcano appears in Figure 3, a reproduction of [BJW17, Fig. 2]: in that case, the depth is 3, the ideal $\mathfrak{l}$ has norm 2, is principal in $\mathcal{O}_0$ but not in $\mathcal{O}_1$, and is ramified in K.

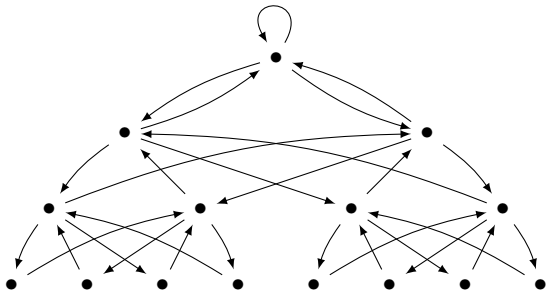

Figure 3: A directed $\mathfrak{l}$-isogeny graph

The situation becomes even more complicated when taking polarizations into account, for instance if we want to only consider principally polarized abelian varieties. In that case, we need to change the isogeny types we consider as $\mathfrak{l}$-isogenies won't preserve principal polarizability in general. For almost complete results on the graph of (ℓ, ℓ)-isogenies between ordinary p.p. abelian surfaces over finite fields, we refer to [BJW17, §6-8].

2 Introduction to isogeny-based cryptography

In this second part of this course, we jump ahead 50 years and present some recent topics in isogeny-based cryptography.

Isogeny-based cryptography is one of several branches in public-key cryptography, if arguably the most recent and least mature one. In contrast with classical elliptic-curve based cryptography, based on the computational hardness of the discrete logarithm problem on elliptic curves over finite fields, isogeny-based cryptography is supposed to be *post-quantum*; in other words the computational problem(s) on which it is based are, conjecturally, also hard for adversaries equipped with quantum computers. Compared to other families of post-quantum cryptographic protocols, particularly lattice-based cryptography, isogeny-based cryptography is less well-known, generally less efficient, but has advantages in some respects, e.g. very small key or signature sizes.

For a long time, the flagship protocol in isogeny-based cryptography was the key exchange scheme SIDH [DJP11] (Supersingular Isogeny Diffie–Hellman) proposed in 2011, and its avatar SIKE, a former candidate in the NIST standardization competition. SIDH relied on the hardness of computing a secret isogeny between known supersingular elliptic curves, but had to leak extra information on the images of torsion points through that isogeny. This ended up being too much information, and the system was decisively broken in a series of papers in 2022 [CD23; MMP+23; Rob23]. This event, however, did not end isogeny-based cryptography: instead, it revealed new ways of efficiently encoding isogenies from their action on torsion points, opening the way to a rich family of new algorithms and cryptosystems, among which [PR23; NOC+24; BDD+24]. A key feature of the SIDH attacks is that they involve isogenies between abelian varieties of higher dimensions, not only between elliptic curves.

First, we review how to encode abelian varieties and isogenies in a computer, with a focus on the notion of efficient representations of isogenies, and algorithms using them. Next, we focus on another historical example of isogeny-based scheme: the 2009 Charles–Goren–Lauter (CGL)

hash function [CLG09], and explain how it can be proven secure assuming the hardness of certain computational problems involving isogenies. Finally, we review recent work of Wesolowski [Wes22] and Page–Wesolowski [PW24] showing that these computational problems are all equivalent to the *endomorphism ring problem* for supersingular elliptic curves.

2.1 Efficient representations of isogenies

For applications to isogeny-based cryptography, it is necessary to represent isogenies in a computer in an efficient way, and to devise algorithms to manipulate them. In particular, we should be able to encode abelian varieties themselves in a computer, as they will be the domains and codomains of our isogenies, respectively. Because the SIDH attacks involve abelian varieties of any dimensions, it is necessary to work in higher dimensions as well.

2.1.1 Representing abelian varieties

Representing abelian variety A in a computer usually involves, at least in some abstract way, embedding A as a subvariety of some projective space. To simplify the discussion, we choose to work with fully explicit projective embeddings in these notes, even if writing down these embeddings is often costly (and avoidable) in practice.

Such an embedding embedding arises from some ample line bundle of A, which in turns defines a polarization $\lambda : A \to A^\vee$: thus it is far more common to encode *polarized* abelian varieties than non-polarized ones. We now present several ways of encoding abelian varieties which might be of interest depending on the case one considers.

Elliptic curves. An elliptic curve E can be written as plane cubics using a Weierstrass equation. The embedding $E \hookrightarrow \mathbb{P}^2$ is associated to the very ample line bundle whose associated divisor is $3(0_E)$. The divisor (0_E) defines a principal polarization on E, as in Example 1.1.20.

Jacobians. Beyond the case of elliptic curves, we find Jacobians of curves. It is especially interesting to focus on those when the dimension g is small because of Theorem 1.1.22. Jacobians are a special case of the next paragraph on p.p. abelian varieties, but more compact representations are usually available for them. For instance, in the case $g = 2$, smooth projective curves of genus 2 are all *hyperelliptic* of degree 5 or 6: they can be written as (the normalization of) a curve in $\mathbb{P}^2$ with an equation of the form

$$y^2 = f(x)$$

where f has degree 5 or 6 [Har77, IV, Exercise 2.2]. Points on the Jacobian can be encoded (generically) as unordered pairs of points on the curve, and the additions can be performed using Mumford coordinates [Can00]. If $g = 3$, then the most common case (in terms of dimensions in moduli) is the case where $\mathcal{C}$ is a plane quartic, and similar addition algorithms also exist [FOR08].

Principally polarized abelian varieties. More generally, if A is endowed with a principal polarization arising from some ample line bundle $\mathcal{L}$, then $\mathcal{L}^{\otimes 4}$ is a very ample line bundle on A by Lefschetz's theorem 1.1.16 and is defined over the base field by Theorem 1.1.14. One can use $\mathcal{L}^{\otimes 4}$ to construct a projective embedding of A called the *theta model* (of level 4) $\Theta : A \hookrightarrow \mathbb{P}^{4^g} - 1$.

This embedding is convenient for many reasons: knowing $\Theta(0_A)$ is enough to write down equations for the image of A using the Riemann relations [Mum66, §3]; moreover, additions of points

can be performed directly in the theta model using the same Riemann relations. The theta model also has drawbacks: the number of coordinates it uses is exponential in g rather than polynomial, and 2-torsion points on A must all be rational in order for Θ to be defined over the base field.

More generally, if we are given an abelian variety of some fixed dimension g with a polarization of some fixed degree d, then we can in principle always construct a projective embedding of A similar to the theta model, and manipulate the theta coordinates of points on A directly. However, the formulas are necessarily more complicated. This explain why almost everybody sticks to principally polarized abelian varieties.

For simplicity, we state the following theorem using the theta model for principally polarized abelian varieties, but variants would also hold for the other explicit representations discussed above.

Theorem 2.1.1. *Fix $g \geq 1$. The theta model of principally polarized abelian varieties (p.p.a.v.'s) of dimension g over finite fields has the following properties.*

1. *(Group law.) Given a p.p.a.v. A over $\mathbb{F}_q$ of dimension g embedded using its theta model $\Theta : A \hookrightarrow \mathbb{P}^N$, given $\Theta(P)$ and $\Theta(Q)$ for some $P, Q \in A(\mathbb{F}_q^r)$ for some $r \geq 1$, one can compute $\Theta(P+Q)$ and $\Theta(-P)$ in time $\mathrm{Poly}(\log q, r)$.*

2. *(Scalar multiplications.) Given (A, Θ) as above, $\Theta(P)$ for some $P \in A(\mathbb{F}_q^r)$, and $n \in \mathbb{Z}$, one can compute $\Theta(nP) \in \mathbb{P}^N(\mathbb{F}_q)$ in time $\mathrm{Poly}(\log q, \log n, r)$.*

3. *(Products.) Given (A, Θ) as above and another p.p. abelian variety (B, Θ') in the theta model of dimension at most g, one can compute $A \times B$ in the theta model in time $\mathrm{Poly}(\log q)$.*

4. *(Decompositions.) Given (A, Θ) as above that we know arises as a product of two lower-dimensional p.p. abelian varieties (B_1, Θ_1) and (B_2, Θ_2) with the product polarization, one can compute those factors as well as an explicit isomorphism $A \simeq B_1 \times B_2$ in time $\mathrm{Poly}(\log q)$.*

5. *(Quotients.) Given (A, Θ) as above, give points $\Theta(P_1), \ldots, \Theta(P_{n^2})$ where $P_1, \ldots, P_{n^2} \in A(\overline{\mathbb{F}}_q)$ are all the points of a maximal isotropic subgroup $K \subset A[n]$ for some given $n \geq 1$ prime to the characteristic of $\mathbb{F}_q$, one can compute the theta model of the quotient $B = A/K$ as well as polynomial formulas for the quotient isogeny $\phi : \Theta(A) \to \Theta(B)$, in time $\mathrm{Poly}(\log q, n)$. If K is defined over $\mathbb{F}_q$, then so are B and ϕ.*

6. *(Generating torsion.) Given (A, Θ) as above and $n \geq 1$, one can list all the points $\Theta(P)$ where $P \in A[n](\overline{\mathbb{F}}_q)$ in time $\mathrm{Poly}(\log q, n)$.*

7. *(Frobenius.) Given (A, Θ) as above over a finite field $\mathbb{F}_q$ of characteristic p, let $\pi_p : A \to A^{(p)}$ and $\pi : A \to A$ denote its p-Frobenius and q-Frobenius maps respectively. Given a point $\Theta(P)$ where $P \in A(\overline{\mathbb{F}}_q^r)$ for some $r \geq 1$, one the points $\Theta(\pi_p(P)) \in \Theta(A^{(p)})$ and $\Theta(\pi(P)) \in \Theta(A)$ in time $\mathrm{Poly}(\log q, r)$.*

Proof.

1. This uses the Riemann relations [Rob21, §I.2.7].

2. Use the double-and-add algorithm and the group law on $\Theta(A)$, which we can write down using polynomial formulas.

3. See [Mum66, Lemma 3 p. 323].

4. By the same lemma, we only have to recognized the theta model of A as the product of those of B and B' under a Segre embedding.

5. See [LR12]. Note that the points P_i are defined over an extension of $\mathbb{F}_q$ of degree $\mathrm{Poly}(n)$.

6. After computing polynomial formulas for the multiplication map $[n] : A \to A$ in the theta model (which have degree $\mathrm{Poly}(n)$), it is sufficient to compute all the roots using the resultant method, and check that they indeed give n-torsion points on A.

7. This amounts to computing the pth or qth power of the coordinates of $\Theta(P)$, which can be done in time $\mathrm{Poly}(\log q, r)$ using the square-and-multiply algorithm. □

In the rest of this section, we assume that we always manipulate isogenies between principally polarized abelian varieties in the theta model (or elliptic curves.) This is perhaps an oversimplification, as not all isogenies we manipulate in practice are of that form. Nevertheless, our discussion would still make sense if we use another kind of explicit models, provided that the analogue of Theorem 2.1.1 holds.

2.1.2 Representing isogenies

One obvious way of representing an isogeny $\phi : A \to B$ of degree d between p.p.a.v.'s, assuming we are given its domain and codomain in the theta model, is to write down polynomial formulas for the isogeny ϕ. These polynomials have degree $\mathrm{Poly}(d)$, and the cost of using them to evaluate ϕ at a given point of A, say, is at least linear in d. This is a problem: assuming that ϕ is the secret key in a cryptographic system where an honest participant should evaluate ϕ, the cost of running the protocol would be roughly the same as the cost of recovering ϕ by brute force by an attacker.

In some cases, one can do better. For instance, if ϕ arises as the composition of n isogenies $\phi_i : A_i \to A_{i+1}$ of degree $O(1)$, then we can uniquely encode ϕ by writing down polynomial representations of $\phi_1, \ldots, \phi_n$, and evaluate ϕ at a given point of A (over $\mathbb{F}_q$, say) using $\mathrm{Poly}(\log q, n) = \mathrm{Poly}(\log q, \log d)$ operations. In that case, one can hope to achieve an exponential complexity gap between the cost of evaluating ϕ and that of recovering it from scratch, a necessary feature for cryptographic applications. To formalize this discussion, we introduce the following definition, a variation on [Rob24, §2.1].

Definition 2.1.2. Fix $g \geq 1$. For each prime power q, let Φ_q be a family of isogenies between p.p.a.v.'s over $\mathbb{F}_q$ of dimension g; we require that the degrees of the isogenies in Φ_q are unbounded for infinitely many q. We define an *efficient representation* of the family of isogenies (Φ_q) with respect to some fixed algorithms Domain, Codomain, Degree, Eval to be a collection of sets $\mathrm{Isog}_q \subset \{0,1\}^*$ and maps $\mathrm{Isog}_q \to \Phi_q$ such that the following hold.

1. For each q and each $\phi \in \Phi_q$ of degree $d \geq 1$, the length of any encoding $D_\phi \in \mathrm{Isog}_q$ of ϕ has length $\mathrm{Poly}(\log q, \log d)$.

2. Given q and any encoding $D_\phi \in \mathrm{Isog}_q$ of an isogeny $\phi \in \Phi_q$ of degree d, the algorithm Domain (resp. Codomain) returns the domain (resp. codomain) of ϕ in the theta model in time $\mathrm{Poly}(\log q)$.

3. Given q and any encoding $D_\phi \in \mathrm{Isog}_q$ of an isogeny $\phi \in \Phi_q$ of degree d, the algorithm Degree returns d in time $\mathrm{Poly}(\log q, \log d)$.

4. Given q, $r \geq 1$, any encoding $D_\phi \in \mathrm{Isog}_q$ of an isogeny $\phi \in \Phi_q$ of degree d, $(A, \Theta) = \textsc{Domain}(q, D_\phi)$, $(B, \Theta') = \textsc{Codomain}(q, D_\phi)$, and any point $\Theta(P)$ where $P \in A(\mathbb{F}_{q^r})$, the algorithm Eval returns $\Theta'(\phi(P))$ in time $\mathrm{Poly}(\log q, r, \log d)$.

We call two efficient representations *equivalent* if they encode the same isogenies, and for any allowable isogeny ϕ of degree d over $\mathbb{F}_q$, an encoding of ϕ in the one representation can be transformed into an encoding of ϕ in the other representation in time $\mathrm{Poly}(\log q, \log d)$.

In the rest of this section, we review "historical" efficient representations of certain families of isogenies. Next, we discuss Kani's lemma and explain how to construct the so-called HD and CRT representations [Rob24, §2.4], which are efficient representations for the family of *all* isogenies between p.p. abelian varieties over any finite field. The existence of those representations has been a major outcome of the attacks on SIDH.

Remark 2.1.3. One could also apply the formalism used in Definition 2.1.2 to make a definition of what an *algorithmic representation of p.p.a.v.'s over finite fields* with respect to algorithms Add, Neg, etc. performing the operations listed in Theorem 2.1.1. We choose not to do this for simplicity, and continue to discuss abelian varieties in the theta model (or elliptic curves.)

2.1.3 Historical efficient representations

As indicated above, we can efficiently represent compositions of isogenies of small degree.

Proposition 2.1.4. *Fix $C \geq 1$ and $g \geq 1$. For a prime power q, let Ψ_q be the collection of all isogenies $\psi : A \to B$ where A, B are p.p. abelian varieties over $\mathbb{F}_q$ of degree at most C, and let Φ_q be the collection of all isogenies arising as (compatible) compositions of isogenies from Ψ_q. Define the set Isog_q as follows: for each compatible sequence of isogenies*

$$\psi_1 : A_0 \to A_1, \quad \ldots, \quad \psi_n : A_{n-1} \to A_n$$

in Ψ_q such that $\deg \psi_i > 1$ for each $i > 1$, add to Isog_q a bit string D_ϕ containing n, equations for $\Theta(A_0), \ldots, \Theta(A_n)$, the integers $\deg \psi_1, \ldots, \deg \psi_n$, as well as polynomial expressions representing ψ_i for each i. This bit string encodes $\phi = \psi_n \circ \cdots \circ \psi_1 \in \Phi_q$.

Then the above data is an efficient representation of the family of isogenies (Φ_q) with respect to the obvious algorithms Domain, Codomain, Degree, *and the algorithm* Eval *specified as follows: given $D_\phi \in \mathrm{Isog}_q$ as above and given $\Theta(P)$ where $P \in A(\mathbb{F}_{q^r})$, map Theta$(P)$ through the polynomial formulas defining $\psi_1, \ldots, \psi_n$ in that order.*

Proof. Everything is straightforward. For instance, the algorithm Eval performs $O(n)$ arithmetic operations on elements of $\mathbb{F}_{q^r}$ because C is bounded, so its overall cost is indeed polynomial in r, $\log q$ and $n = O(\log d)$ (because we forbid long sequences of isomorphisms.) □

In the case of elliptic curves, isogeny factorization (Proposition 1.1.12) shows that any smooth-degree isogeny can be obtain as a composition of small-degree isogenies. Formulating an analogous statement in higher dimensions is possible but more complicated, as we have to make sure that the intermediate abelian varieties are also principally polarized (we would have to talk about isotropic subgroups as in Proposition 1.1.18.) Recall that a number n is called *C-smooth* if each prime $p|n$ satisfies $p \leq C$.

Interestingly, one can also describe straightforward efficient representations for other classes of isogenies that are not necessarily of smooth degree, namely scalar multiplications and Frobenius endomorphisms. We omit the (straightforward) proofs.

Proposition 2.1.5. *Fix $g \geq 1$. For a prime power q, let Φ_q be the collection of scalar multiplication isogenies $[n] : A \to A$ where A is a p.p.a.v. of dimension g over $\mathbb{F}_q$ and $n \in \mathbb{Z}\backslash\{0\}$. Let Isog_q be the set consisting of binary representations of tuples (n, A, Θ) where $n \in \mathbb{Z}$ and (A, Θ) is such an abelian variety in the theta model.*

Then this data is an efficient representation of the family of scalar multiplication isogenies (Φ_q) with respect to the obvious algorithms Domain, Codomain, Degree, *and the algorithm* Eval *given by the double-and-add algorithm as in Theorem 2.1.1.*

Proposition 2.1.6. *Fix $g \geq 1$. For a prime power $q = p^r$, let Φ_q be the collection of q-Frobenius endomorphisms $\pi : A \to A$ and p-Frobenius maps $\pi_p : A \to A^{(p)}$ where A is a p.p.a.v. of dimension g over $\mathbb{F}_q$. Let Isog_q be the set consisting of binary representations of tuples (x, A, Θ) where $x \in \{p, q\}$ and (A, Θ) is such an abelian variety in the theta model.*

Then this data is an efficient representation of the family of Frobenius maps (Φ_q) with respect to the obvious algorithms Domain, Codomain, Degree, *and the algorithm* Eval *given by fast exponentiation as in Theorem 2.1.1.*

Moreover, if we have an efficient representation from a family of isogenies (Φ_q), then we can construct an efficient representation for the family of isogenies arising as compositions of isogenies from Φ as in Proposition 2.1.4. Consequently, we know how to efficiently represent compositions of small-degree isogenies, scalar multiplications, and Frobenius endomorphisms.

Another, completely different kind of efficient representations that have been used in isogeny-based cryptography for a long time in the special case of supersingular elliptic curves arises from the Deuring correspondence (Theorem 1): we represent isogenies in terms of the corresponding ideals connecting maximal orders in the quaternion algebra $B_{p,\infty}$. The endomorphism rings of the domain and/or codomain are necessarily part of that encoding. In order to evaluate such an isogeny at a point, one can *smoothen* the ideal, i.e. find an equivalent ideal of smooth norm (for instance of norm ℓ^n for some fixed small ℓ and some $n \geq 1$). This can be done, at least heuristically, using an algorithm by Kohel–Lauter–Petit–Tignol (KLPT) [KLP+14]. This representation in terms of ideals is at the heart of (the first version of) the SQIsign signature protocol, which predates the attacks on SIDH [DKL+20].

2.1.4 The HD and CRT representations

The idea of the HD (higher-dimensional) representation, now fundamental in isogeny-based cryptography, relies on embedding any isogeny between p.p. abelian varieties of some dimension g (often $g = 1$) as a component of an isogeny of smooth degree in dimension $2g, 4g$ or $8g$ between product abelian varieties, using Kani's lemma below as the key tool.

Recall that for an integer $n \geq 1$, an isogeny $\phi : A \to B$ between p.p. abelian varieties of dimension g is called an *n-isogeny* if $\ker(\phi) \subset A[n]$ is maximal isotropic for the Weil pairing, and B is endowed with the natural principal polarization that exists on the quotient $A/\ker\phi$ as in Proposition 1.1.18. Equivalently, we have $\lambda_A^{-1} \circ \phi^\vee \circ \lambda_B \circ \phi = [n]_A$, where λ_A and λ_B denote the principal polarizations on A and B respectively.

Theorem 2.1.7 (Kani's lemma)**.** *Let $g \geq 1$ and let $n_1, n_2 \geq 1$ be coprime integers. Consider a commutative square of isogenies between p.p. abelian varieties of dimension g over any field k:*

$$\begin{array}{ccc} A_0 & \xrightarrow{\phi_1} & A_1 \\ {\scriptstyle\phi_2}\downarrow & & \downarrow{\scriptstyle\phi_2'} \\ A_2 & \xrightarrow{\phi_1'} & A_{12} \end{array}$$

where ϕ_1, ϕ_1' are n_1-isogenies and ϕ_2, ϕ_2' are n_2-isogenies. Denote the principal polarizations on A_0, A_1, A_2, A_{12} by $\lambda_0, \lambda_1, \lambda_2, \lambda_{12}$ respectively. Then the isogeny

$$\Phi = \begin{pmatrix} \phi_1 & \lambda_1^{-1} \circ \phi_2'^{\vee} \circ \lambda_{12} \\ -\phi_2 & \lambda_0^{-1} \circ \phi_1'^{\vee} \circ \lambda_1 \end{pmatrix} : A_0 \times A_{12} \to A_1 \times A_2$$

is an $(n_1 + n_2)$-isogeny, where $A_0 \times A_{12}$ and $A_1 \times A_2$ are endowed with the product polarizations. The kernel of Φ is the following maximal isotropic subgroup in $A_0 \times A_{12}[n_1 + n_2]$:

$$\begin{aligned} \ker \Phi &= \{(\lambda_0^{-1} \circ \phi_1^{\vee} \circ \lambda_1(P), \phi_2'(P)) : P \in A_1[n_1 + n_2]\} \\ &= \{(-\lambda_0^{-1} \circ \phi_2^{\vee} \circ \lambda_2(P), \phi_1'(P)) : P \in A_2[n_1 + n_2]\} \\ &= \{(n_1 P, \phi_2' \circ \phi_1(P) : P \in A_0[n_1 + n_2]\}. \end{aligned}$$

The proof uses the following lemma.

Lemma 2.1.8. *Let $A_1, \ldots, A_r$, $B_1 \ldots, B_s$ be principally polarized abelian varieties over k, and denote their principal polarizations by $\lambda_1, \ldots, \lambda_r, \lambda_1', \ldots, \lambda_s'$. Let $\Phi : A_1 \times \cdots \times A_r \to B_1 \times \cdots \times B_s$ be any morphism representing by a matrix $M = (\Phi_{i,j})$, where $\Phi_{i,j} : A_i \to B_j$. Then the morphism $\Phi^{\vee} : B_1^{\vee} \times \cdots \times B_s^{\vee} \to A_1^{\vee} \times \cdots \times A_r^{\vee}$ is represented by the matrix $M' = (m'_{i,j})$ with $m'_{i,j} = \Phi^{\vee}_{j,i}$.*

Proof. This is obvious from the definition of duals in terms of line bundles. □

Proof of Theorem 2.1.7. By Lemma 2.1.8, we have

$$(\lambda_0 \times \lambda_{12})^{-1} \circ \Phi^{\vee} \circ (\lambda_1 \times \lambda_2) = \begin{pmatrix} \lambda_0^{-1} \circ \phi_1^{\vee} \circ \lambda_1 & -\lambda_0^{-1} \circ \phi_2^{\vee} \circ \lambda_2 \\ \phi_2' & \phi_1' \end{pmatrix} : A_1 \times A_2 \to A_0 \times A_{12}.$$

Then, a direct matrix computation yields

$$\begin{aligned} (\lambda_0 \times \lambda_{12})^{-1} \circ \Phi^{\vee} \circ (\lambda_1 \times \lambda_2) \circ \Phi &= \begin{pmatrix} \lambda_0^{-1}\phi_1^{\vee}\lambda_1\phi_1 + \lambda_0^{-1}\phi_2^{\vee}\lambda_2\phi_2 & 0 \\ 0 & \phi_2'\lambda_1^{-1}\phi_2'^{\vee}\lambda_{12} + \phi_1'\lambda_2^{-1}\phi_1'^{\vee}\lambda_{12} \end{pmatrix} \\ &= \begin{pmatrix} n_1 + n_2 & 0 \\ 0 & n_1 + n_2 \end{pmatrix} \\ &= [n_1 + n_2]_{A_0 \times A_{12}}. \end{aligned}$$

In the first line, we used that $\phi_1' \circ \phi_2 = \phi_2' \circ \phi_1$. □

Theorem 2.1.7 still holds when n_1 and n_2 are not coprime, with one exception: in that case, there is no known simple expression for $\ker \Phi$. A commutative square of isogenies as in the theorem is sometimes called an *isogeny diamond*.

We now use Kani's lemma to construct an efficient representation of all isogenies between elliptic curves over finite fields, be they of smooth degree or not. For a given isogeny $\phi : E \to E'$ of degree d, choose an integer $N = \ell_1 \cdots \ell_r$, a product of small primes coprime to d, such that $N > d$. The goal is to use Kani's lemma to embed ϕ as a component of an N-isogeny in higher dimensions. There are several possibilities. (We allow ourselves to identify elliptic curves with their duals.)

1. If $N - d = a^2$ is a perfect square, we consider the isogeny diamond

$$\begin{array}{ccc} E & \xrightarrow{[a]_E} & E \\ {\scriptstyle\phi}\downarrow & & \downarrow{\scriptstyle\phi} \\ E' & \xrightarrow{[a]_{E'}} & E' \end{array}$$

   By Kani's lemma, the isogeny $\Phi = \begin{pmatrix} a & \phi^\vee \\ -\phi & a \end{pmatrix} : E \times E' \to E \times E'$ is an N-isogeny.

2. If $N - d = a_1^2 + a_2^2$ is a sum of two squares (already a far more common occurrence), then constructing an $(N-d)$-isogeny from E is not obvious. However, the isogeny

$$\phi_1 = \begin{pmatrix} a_1 & a_2 \\ -a_2 & a_1 \end{pmatrix} : E^2 \to E^2$$

   is (by Kani's lemma) an $(a_1^2+a_2^2)$-isogeny. After doing this, we should construct a d-isogeny ϕ_2 with domain E^2 that includes ϕ as one of its components: we can simply take

$$\phi_2 = \begin{pmatrix} \phi & 0 \\ 0 & \phi \end{pmatrix} : E^2 \to E'^2,$$

   giving rise to the desired isogeny diamond. By Kani's lemma,

$$\Phi = \begin{pmatrix} a_1 & a_2 & \phi^\vee & 0 \\ -a_2 & a_1 & 0 & \phi^\vee \\ -\phi & 0 & a_1 & -a_2 \\ 0 & -\phi & a_2 & a_1 \end{pmatrix} : E^2 \times E'^2 \to E^2 \times E'^2$$

   is an N-isogeny in dimension 4.

3. Otherwise, $N - d$ will always be a sum of 4 squares $a_1^2 + a_2^2 + a_3^2 + a_4^2$. This time, we go to dimension 8: the endomorphism

$$\phi_1 = \begin{pmatrix} a_1 & a_2 & a_3 & a_4 \\ -a_2 & a_1 & -a_4 & a_3 \\ -a_3 & a_4 & a_1 & -a_2 \\ -a_4 & -a_3 & a_2 & a_1 \end{pmatrix} : E^4 \to E^4$$

   is an $(N-d)$-isogeny, allowing us to embed ϕ as a component of an N-isogeny $\Phi : E^4 \times E'^4 \to E^4 \times E'^4$ in dimension 8.

Theorem 2.1.9 (The HD representation)**.** *For a prime power q, let Φ_q be the collection of all isogenies between elliptic curves over $\mathbb{F}_q$. To an isogeny $\phi : E \to E'$ in Φ_q, we associate data D_ϕ consisting of:*

- *E, E', and $d = \deg \phi$;*
- *a collection of small distinct primes $\ell_1, \ldots, \ell_r$, prime to d, such that $N = \ell_1 \cdots \ell_r > d$;*

- *$k = 1$, 2, or 4 integers denoted by a, (a_1, a_2) or (a_1, a_2, a_3, a_4) whose squares sum to $N - d$;*
- *the following representation of the N-isogeny Φ in dimension $2k$ constructed above. We know that there exists p.p. abelian varieties $A_1, \dots, A_{r-1}$ of dimension $2k$ and isogenies*

$$\Psi_1 : E^d \times E'^d \to A_1, \quad \Psi_2 : A_1 \to A_2, \quad \dots, \quad \Psi_r : A_{r-1} \to E^4 \times E'^4$$

such that each Ψ_i is an ℓ_i-isogeny and $\Phi = \Psi_r \circ \cdots \circ \Psi_1$. The representation of Φ consists of the abelian varieties $A_1, \dots, A_{r-1}$ in the theta model together with polynomial formulas describing $\Psi_1, \dots, \Psi_r$, as in Proposition 2.1.4.

Then this data defines an efficient representation for the family of all elliptic curve isogenies (Φ_q) with respect to the obvious algorithms Domain, Codomain, *and* Degree, *and the algorithm* Eval *constructed as in Proposition 2.1.4 and followed by projection onto one of the components.*

Proof. Straightforward. □

A key feature of the HD representation is that the kernel of Φ in Kani's lemma can be constructed from the image of ϕ at ℓ-torsion points for $\ell \in \{\ell_1, \dots, \ell_r\}$. In [Rob24, §2.4], providing these images is called the "kernel version" of the HD representation, but it feels like this representation deserves its own name.

Theorem 2.1.10 (The CRT representation)**.** *Define Φ_q as in Theorem 2.1.9. To an isogeny $\phi : E \to E'$ in Φ_q, we associate data D_ϕ consisting of:*

- *E, E', and $d = \deg \phi$;*
- *a collection of small distinct primes $\ell_1, \dots, \ell_r$ prime to d such that $N = \ell_1 \cdots \ell_r > d$;*
- *for each $1 \le i \le r$, a basis (P_i, Q_i) of $E[\ell_i](\overline{\mathbb{F}}_q)$, as well as the points $\phi(P_i), \phi(Q_i)$ on E'.*

Then this data defines an efficient representation for the family of (Φ_q) of all elliptic curves isogenies that is equivalent to the HD representation.

The name "CRT" stems from the isomorphism $E[N] \simeq \prod_i E[\ell_i]$ given by the Chinese remainder theorem. We use a basis of the right hand side to encode the action of ϕ on N-torsion points.

Proof. First, we show that the CRT representation can be efficiently transformed into the HD representation. This implies that the CRT representation is efficient. Let $\phi : E \to E'$ be an isogeny of degree d. (E, E' and d are known as part of the CRT representation.) We first compute integers a_1, a_2, a_3, a_4 whose squares sum to $N - d$; this can also be done in time $\mathrm{Poly}(\log d)$ [RS85].

By Theorem 2.1.7, ϕ (or rather $-\phi$) appears as a component of an explicit N-isogeny $\Phi : E^4 \times E'^4 \to E^4 \times E'^4$. Its kernel is

$$\ker \Phi = \{((N - d)P, M\phi(P)) : P \in E^4[N]\}.$$

By isogeny factorization, Φ can be written as a composition $\Psi_1 \circ \cdots \circ \Psi_r$ where each Ψ_i is an ℓ_i-isogeny between p.p. abelian varieties of dimension 8. Moreover, the intermediate abelian varieties are principally polarizable by Proposition 1.1.18.

Next, we compute $\ker \Psi_1$ by mapping the given basis of $E^4[\ell_i]$ through the above formula. Then, we compute $\Psi_1 : E^4 \times E'^4 \to A_1$. This uses the algorithms products and quotients from

Theorem 2.1.1. The kernel of Ψ_2 is now the image under Ψ_1 of the same construction starting from $E[\ell_2]$, and we continue. In the end, we construct a p.p. abelian variety A_r that we know is isomorphic to $E^4 \times E'^4$ with its product polarization, and r isogenies $\Psi_1, \dots, \Psi_r$. We decompose A_r as a product of elliptic curves using Theorem 2.1.1 to conclude.

Conversely, the fact that the HD representation can be efficiently transformed into the CRT representation is a special case of the next theorem. □

In fact, an encoding of an isogeny ϕ in the CRT representation can be constructed as soon as we know ϕ in *any* another efficient representation: the HD representation is "universal".

Theorem 2.1.11. *Consider any efficient representation of some family of isogenies $\Phi = (\Phi_q)_q$ with respect to some algorithmic representation of elliptic curves. Then one can construct an algorithm which, given q and an encoding of some isogeny ϕ of degree d in the chosen representation, outputs an encoding of ϕ in the CRT representation in time* $\mathrm{Poly}(\log q, \log d)$.

Proof. First, we apply Degree, Domain and Codomain to obtain the two elliptic curves E, E' and the degree d of $\phi : E \to E'$. Next, we choose small distinct primes $\ell_1, \dots, \ell_r$ prime to $\deg \phi$ such that $N = \ell_1 \cdots \ell_r > \deg \phi$: this can be done in time $\mathrm{Poly}(\log d)$, with $r = O(\log d)$ and $\max_i \ell_i = O(\log d)$. For each i, we generate a basis (P_i, Q_i) of $E[\ell_i](\overline{F}_q)$ as in Theorem 2.1.1. These points are defined over extensions of $\mathbb{F}_q$ of degree $\mathrm{Poly}(\log d)$. Finally, we apply the algorithm Eval to obtain $\phi(P_i)$ and $\phi(Q_i)$ for each i. □

Remark 2.1.12. One could similarly define HD representations of isogenies in higher dimensions, but if we do it in the straightforward way, it would only cover isogenies between p.p. abelian varieties that are n-isogenies for some integer $n \geq 1$, as in the above definition. This becomes a nontrivial restriction as soon as $g > 1$.

2.1.5 Algorithms on efficient representations

To conclude this section, we discuss algorithms taking isogenies in efficient representation as input, and which output other isogenies in efficient representation. We are somewhat agnostic of what the representation really is, and only use the algorithms provided by Definition 2.1.2. To be more concrete, one can always think about isogenies in HD or CRT representation by Theorem 2.1.11.

Testing equality. We rely on the following lemma.

Lemma 2.1.13. *Let $\phi, \phi' : E \to E'$ be isogenies of the same degree d. If ϕ and ϕ' agree on $4d+1$ distinct points on E, then $\phi = \phi'$.*

Proof. The degree map on $\mathrm{Hom}(E, E')$ begaves like a quadratic form (as can be seen from Theorem 1.1.29), so $\deg(\phi - \phi') \leq 4d$. □

This lemma is tight: for every $n \in \mathbb{Z}$, $[n]_E$ and $[-n]_E$ both have degree n^2 and agree on the $4n^2$ points of $E[2n](\overline{\mathbb{F}}_p)$. The algorithm Eq testing the equality of isogenies in efficient representation then works as follows.

1. Test that the degrees, domains and codomains of ϕ and ϕ' are the same.
2. Pick a sequence of small primes $\ell_1, \dots, \ell_r$ such that $\ell_1^2 \cdots \ell_r^2$ bigger than $4d+1$; this can be done in time $\mathrm{Poly}(\log d)$.

3. For each $1 \leq i \leq r$, generate a basis (P_i, Q_i) of $E(\overline{\mathbb{F}}_p)[\ell_i]$, and check that $\phi(P_i) = \phi'(P_i)$ and $\phi(Q_i) = \phi'(Q_i)$.

The algorithm Eq runs in time $\mathrm{Poly}(\log q, \log d)$.

Negation and sum. If $\phi, \phi' : E \to E'$ are two isogenies in efficient representation, we are asked to compute efficient representations of $-\phi$ and $\phi + \phi'$. We already know their domain, codomain, and how to evaluate them efficiently (first evaluate ϕ and ϕ', then apply the group law on E'.) We also know the degree of $-\phi$, so the only nontrivial step is to be able to compute the degree of $\phi + \phi'$, as required by Definition 2.1.2.

For this, we use the Weil pairing. We already know that $\deg(\phi + \phi') \leq M = 4(\deg \phi)(\deg \phi')$. Pick small primes $\ell_1, \ldots, \ell_r$ such that $\ell_1 \cdots \ell_r > M$, and for each i, construct a basis (P_i, Q_i) of $E[\ell_i]$. Then

$$e_{\ell_i}(\phi(P_i) + \phi'(P_i), \phi(Q_i) + \phi'(Q_i)) = e_{\ell_i}(P_i, Q_i)^{\deg(\phi+\phi')}.$$

We can compute the Weil pairings and solve the discrete logarithm problem in time $\mathrm{Poly}(q, \ell_i)$, so we gain access to $\deg(\phi + \phi') \bmod \ell_i$. Finally, we use Chinese remainders to compute $\deg(\phi + \phi')$. In the end, we obtain two algorithms Neg and Sum on isogenies in efficient representation which run in polynomial time in $\log q$ and the log of the degrees.

Composition. Composing isogenies in efficient representation is easier than summing, as the degree of a composition is the product of the degrees.

The dual isogeny. Let $\phi : E \to E'$ be an isogeny of degree d. The dual isogeny $\phi^\vee : E \to E'$ also has degree d. In order to compute a CRT representation of $\phi^\vee$, we start from a CRT basis $((P_i, Q_i))_{1 \leq i \leq r}$ of $E[N]$ where $N = \ell_1 \cdots \ell_r$, on which the images of ϕ are known. Then, assuming that each prime ℓ is coprime to d, the tuple $(\phi(P_i), \phi(Q_i))_{1 \leq i \leq r}$ is a CRT basis of $E'[N]$; moreover, for every $1 \leq i \leq r$, we have

$$\phi^\vee(\phi(P_i)) = dP_i \quad \text{and} \quad \phi^\vee(\phi(Q_i)) = dQ_i.$$

This provides a CRT representation of $\phi^\vee$.

Dividing by integers. Here we are given an isogeny $\phi : E \to E'$ in CRT representation and an integer $m \geq 1$. We describe an algorithm Division which, on this input, either returns $\psi : E \to E'$ (in CRT representation) such that $\phi = m\psi$ if it exists, or a failure symbol $\perp$, in time $\mathrm{Poly}(q, \log \deg \phi, \log m)$. Simply, we can assume that the CRT modulus N is prime to m; then one can divide the images of ϕ by m (i.e. perform scalar multiplications by the inverse of m modulo ℓ_i).

This provides an efficient representation of ϕ/m if this is actually an isogeny. How can we decide whether the obtained CRT data encodes a valid isogeny? Well, we attempt to transform it into an HD representation as in Theorem 2.1.11. If the codomain of the isogeny Φ is not isomorphic to $E^4 \times E'^4$, then we output $\perp$. If it is, then the components of Φ yield finitely many honest morphisms $E \to E'$ in efficient representation, and we can check that ϕ/m is one of them with the equality algorithm.

Isogeny factorization. More generally, if we have two isogenies $\phi : E \to E'$ and $\psi : E \to E''$ in efficient representation, we wish to compute an isogeny $\eta : E' \to E''$ such that $\psi = \eta \circ \phi$ in efficient representation, or $\perp$ if such an η does not exist. In order to do this, we can simply write

$$\phi \circ \psi^\vee = d\eta$$

where $d = \deg \psi$, and apply the composition and division algorithms.

Splitting isogenies of coprime degrees. Here we assume that n_1, n_2 are coprime integers, and we are given an isogeny $\phi : E \to E'$ of degree $n_1 n_2$ in efficient representation. There exists an elliptic curve E'' as well as isogenies $\phi_1 : E \to E''$ and $\phi_2 : E'' \to E$ of degree n_1 and n_2 respectively, and such that $\phi = \phi_2 \circ \phi_1$. We are asked to compute E'' as well as the isogenies ϕ_1, ϕ_2 in efficient representation.

If $n_1 + n_2$ is smooth, then the problem is easy: by Kani's lemma, we have an isogeny diamond and we can compute $\ker \Phi$, then Φ from its kernel. We obtain E'' as one of the factors of the codomain of Φ, and the required isogenies as components of Φ.

In general, we pad ϕ with extra isogenies. The goal will be to find efficient endomorphisms ψ_1, ψ_2 of E and E' respectively, of respective degrees u, v, such that un_1 and vn_2 are coprime and $un_1 + vn_2$ is a product of small distinct primes. We can then split the isogeny $\psi_2 \circ \phi \circ \psi_1$ as $\psi_2 \circ \phi_2'$ and $\phi_1 \circ \psi_1$ using the above method in the smooth case. Finally, we apply the isogeny factorization algorithm to recover efficient representations of ϕ_2' and ϕ_1.

Note that for every $N > n_1 n_2$, there exist (efficiently computable) $u, v \geq 0$ such that $N = un_1 + vn_2$. If u and v are squares, then we can choose ψ_1 and ψ_2 to be scalar multiplications; in the worst case, u and v are sums of 4 squares, and we can find ψ_1 and ψ_2 in dimension 4 in a similar way to the HD representation.

Remark 2.1.14. This splitting algorithm can for instance be applied as follows. Let $E/\mathbb{F}_{p^2}$ be a supersingular curve with known endomorphism ring; such a curve can for instance be obtained from the reduction of a CM elliptic curve in characteristic zero. In order to generate an isogeny of some chosen degree n from E, we choose a product $N > n$ of small distinct prime numbers (coprime to n), generate an endomorphism of E of degree $n(N - n)$, and let $\phi : E \to E'$ be the first part of the splitting.

Another example is that in the setting of the CM group action, we can evaluate the action of a non-smooth ideal I as well: we only have to find two equivalent ideals J_1, J_2 of coprime norms, and split the endomorphism $\phi_{J_2}^\vee \circ \phi_{J_1}$, which we know how to efficiently evaluate.

Pushforwards of coprime degrees . Given two isogenies $\phi : E \to E_1$ and $\psi : E \to E_2$ of coprime degrees, we ask to compute an elliptic curve E' efficient representations of the isogenies $\phi' : E_2 \to E'$ and $\psi' : E_1 \to E'$ such that the diagram

$$\begin{array}{ccc} E & \xrightarrow{\phi} & E_1 \\ \downarrow{\scriptstyle \psi} & & \downarrow{\scriptstyle \psi'} \\ E_2 & \xrightarrow{\phi'} & E' \end{array}$$

commutes. This can be done as follows: applying the dual and composition algorithms, we compute an efficient representation of $\psi \circ \phi^\vee : E_1 \to E_2$. Then, we apply the splitting algorithm to compute E_3, ψ' and ϕ'.

2.1.6 Some open problems

The notion of efficient representation of isogenies is a recent one, and the following problems, suggested in [Rob24, §7], are still open.

1. Given an elliptic curve E and a point P on E of order n (say defined over the base field), how can we compute an efficient representation of the degree n isogeny $\phi : E \to E/\langle P \rangle$ in time $\mathrm{Poly}(\log q, \log n)$? The best known algorithms have complexity at least $\sqrt{n}$, so exponential in $\log n$.

2. Are there splitting and pushforward algorithms for efficient representations of isogenies of non-coprime degrees?

3. Given a supersingular elliptic curve E, can we construct an isogeny of any degree from E in efficient representation without knowing its endomorphism ring?

4. As seen above, we can evaluate the CM action for ideals of any norm, i.e. we can walk horizontally in ℓ-isogeny volcanoes, where ℓ is a large prime, in time $\mathrm{Poly}(\log \ell)$. Can we walk up and down in the volcano in time $\mathrm{Poly}(\log \ell)$ as well?

Finally, going beyond the framework of polynomial-time algorithms and Definition 2.1.2, an ongoing research effort is to make the above algorithms as efficient as possible to accelerate the cryptographic schemes they appear in.

2.2 The Charles–Goren–Lauter hash function

The CGL hash function is historically the very first cryptographic scheme based on isogenies, specifically on 2^n-isogenies between supersingular elliptic curves over a finite fields. In this section, we review its construction, and reduce its security properties to certain well-posed computational problems on isogenies.

2.2.1 Cryptographic hash functions

Before we present the CGL hash function proper, we digress and discuss cryptographic hash functions in general.

Definition 2.2.1. A family of *cryptographic hash functions* (H_λ), indexed by an integer parameter λ (the security parameter), is a family of functions

$$H_{\lambda,s} : \{0,1\}^* \to \{1, \dots, N_\lambda\}$$

where $N_\lambda \geq 1$, indexed by seeds $s \in S_\lambda$ where S_λ is a finite set, with the following properties:

1. (Efficiency.) $N_\lambda = \mathrm{Poly}(\lambda)$, and evaluating $H_{\lambda,s}$ given $s \in S_\lambda$ on a string of length n can be done in $\mathrm{Poly}(n, \lambda)$ operations.

2. (Preimage resistance.) Given random $s \in S_\lambda$ and $x \in \{1, \dots, N_\lambda\}$, it must be infeasible for any polynomial-time adversary to recover a preimage $y \in \{0,1\}^*$ of x via $H_{\lambda,s}$. More formally, if A is any polynomial-time algorithm that, given λ and random x, s, attempts to output y, the probability of success of A is $O(2^{-\lambda})$.

3. (Collision resistance.) It must be infeasible for any polynomial-time adversary to compute binary strings $y \neq y'$ such that $H_{\lambda,s}(y) = H_{\lambda,s}(y')$. More formally, if A is any polynomial-time algorithm that, given λ and a random s, attempts to output such y, y', then the probability of success of A is $O(2^{-\lambda})$.

Remark 2.2.2. Perhaps a more meaningful definition, that would make sense of the statement that a given H is a cryptographic hash function with 128 or 256 bits of security (i.e. $\lambda \to \infty$ is not allowed), would be to require (for instance, in the case of preimage resistance) that the probability of success of any algorithm A is at most $2^{-\lambda}$ times the number of binary operations performed by A. This would then require to define what a "binary operation" exactly is. I'm not sure if such a formal definition has been written down somewhere.

A typical use case of cryptographic hash functions are *digital signatures*: in order to sign a message m, the honest prover picks a random seed s, computes $H_{\lambda,s}(m)$, and uses this hash as input to the signature protocol. An adversary could perhaps be able to fake their way through the signature protocol for certain well-chosen inputs h, but preimage resistance guarantees they won't be able to find any message matching this signature. On the other hand, if adversary can construct two messages $m \neq m'$ such that $H_{\lambda,s}(m) = H_{\lambda,s}(m')$, ask an authority to sign the (completely normal) message m, and get a signature that is also valid for the (malicious) message m'. Collision resistance guarantees this does not happen.

2.2.2 Construction of the CGL hash functions

We follow the original paper [CLG09]. Pick a security parameter λ, and choose a prime number p of bit length λ such that $p = 1 \bmod 12$. We consider the isogeny class V of minimal supersingular elliptic curves over $\mathbb{F}_{p^2}$ with all endomorphisms defined; the seed space S_λ is the set V itself. We consider the supersingular ℓ-isogeny graph $G(p, \ell)$ on the vertex set V as defined in Definition 1.5.1, and take $\ell = 2$. By Proposition 1.5.2, $G(p, 2)$ is an undirected, 3-regular graph.

We also pick an arbitrary total order on the set V (for instance, by picking an arbitrary total order on $\mathbb{F}_{p^2}$, and ordering elliptic curves by their j-invariants.) In order to evaluate $H_{\lambda,s}$ on a bit string m of length n, we perform a non-backtracking walk in $G(p, 2)$ as follows:

- The starting point is $E = s$.
- For each $0 \leq i \leq n$, we look at the three neighbors of the current curve E in $G(p, 2)$, discard the curve we came from at the previous step (an arbitrary one, say the biggest, if $i = 0$) and order the two others as $E_0 < E_1$. Next, we replace the current curve E by E_{m_i}.
- Finally, we output $H_{\lambda,s}(m) = j(E) \in \mathbb{F}_{p^2}$.

This hash function is not especially pratical: for every bit of the message, we have to compute 2-isogenies between elliptic curves, which involves several arithmetic operations over $\mathbb{F}_{p^2}$. Nevertheless, it is efficient in the sense of Definition 2.2.1. Note that the generalization of this construction to higher dimensions might, perhaps surprisingly, be more efficient, mainly due to the fact that the higher number of 2-isogenies allows one to process several bits of the message at once [KMM+24].

Why should we believe that the CGL family is a family of cryptographic hash functions? Since $P \neq NP$ is still a conjecture, this can only be proven if we *assume* that certain computational problems are hard. The aim of a *security proof*, or *security reduction*, is then to show that, if an adversary was able to break the preimage resistance (say) of the CGL hash functions, then they

would also be able to solve these computational problems. We introduce these hard problems, and give the security proofs, in §2.3.

At this point, let us only mention that in the security proof of collision resistance, it is crucial that the seed s in the CGL hash function is taken (very close to) uniformly at random. How can we sample random supersingular elliptic curves? The best currently known way is to start from any vertex in the ℓ-isogeny graph, perform a random walk in the graph, and output its endpoint. It turns out that the probability distribution of this output is very close to the uniform even when the random walk is quite short (of length $O(\log p)$), because the supersingular ℓ-isogeny graphs are good *expander graphs*. In the rest of this section, we discuss expander graphs in general, and state that supersingular isogeny graphs are Ramajan (i.e. expander) with a brief sketch of proof.

2.2.3 Expander graphs

For simplicity, we only discuss d-regular, undirected graphs. Given such a graph G with vertex set V and edge set E, and an ordering $v_1, \ldots, v_n$ of v, we define the *adjacency matrix* $A(G)$ of G to be the $n \times n$ matrix whose (i, j) entry is the number of edges from v_i to v_j. It is a symmetric real matrix because G is undirected.

In fact, we can also view $A(G)$ as the matrix of a self-adjoint operator on a finite-dimensional Hilbert space, as follows. Let $L^2(V)$ be the vector space of functions $f : V \to \mathbb{C}$ endowed with the hermitian product

$$\langle f, g \rangle = \sum_{v \in V} f(v)\overline{g(v)}.$$

The *adjacency operator* associates to each $f \in L^2(V)$ the function $h : V \to \mathbb{C}$ such that for each $v \in \mathbb{C}$,

$$h(v) = \sum_{w \text{ neighbor of } v} f(w)$$

In this sum, neighbors are counted with multiplicities if multiple edges are present. There is a natural basis $(f_1, \ldots, f_n)$ of $L^2(V)$ satisfying $f_i(v_j) = \delta_{ij}$ (the Kronecker symbol). The matrix of the adjacency operator in this basis is precisely $A(G)$. As a self-adjoint operator, $A(G)$ admits n real eigenvalues, which we order as $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$. We always have $\lambda_1 = d$ and $\lambda_n \geq -d$ by the triangle inequality, and because the eigenfunction 1 is an eigenvector for the eigenvalue d.

Definition 2.2.3. Let $\varepsilon > 0$. We say that G is a (two-sided) *ε-expander graph* if $\lambda_2 \leq (1 - \varepsilon)d$ and $\lambda_n \geq -(1-\varepsilon)d$, in other words if $\lambda_1 = d$ has multiplicity one as an eigenvalue and all the other eigenvalues of $A(G)$ have absolute value at most $(1 - \varepsilon)d$.

There is an upper bound on how good expander graphs can be for a fixed d:

Theorem 2.2.4 ([Nil91])**.** *Let (G_i) be a sequence of d-regular graphs whose number of vertices goes to infinity with i. Let $\lambda_2^{(i)}$ and $\lambda_n^{(i)}$ be the second and last eigenvalues of $A(G_i)$. Then we have*

$$\max\{|\lambda_2^{(i)}|, |\lambda_n^{(i)}|\} \geq 2\sqrt{d-1} - o(1).$$

A d-regular graph for which $\max\{|\lambda_2|, |\lambda_n|\} \leq 2\sqrt{d-1}$ is called *Ramanujan*. Given the above theorem, one often says that Ramanujan graphs are "optimal" expander graphs.

There are other possible definitions of expander graphs (in term of edge expansion ratios, for instance), but the spectral definition as above is convenient to prove the following.

Proposition 2.2.5. *Let G be a d-regular ε-expander graph with n vertices. Let v be a vertex in V. For each $m \geq 1$, let δ_m denote the probability distribution on V of the endpoint of a random walk with m steps starting from v; let δ_∞ denote the uniform probability distribution on V. Then*

$$\|\delta_m - \delta_\infty\|_2 \leq (1-\varepsilon)^m.$$

In particular,

$$\sum_{v \in V} \left| \delta_m(v) - \frac{1}{n} \right| \leq \varepsilon^m \sqrt{n}.$$

Proof. We consider the operator $\frac{1}{d}A(G)$. Looking at the definition of $A(G)$ and that of a random walk, we immediately get

$$\frac{1}{d}A(G)\delta_m = \delta_{m+1} \quad \text{and} \frac{1}{d}A(G)\delta_\infty = \delta_\infty,$$

so

$$\delta_m - \delta_\infty = \frac{1}{d^m}A(G)^m(\delta_0 - \delta_\infty).$$

The function δ_∞ generates the eigenspace of $A(G)$ for the eigenvalue one. The orthogonal subspace $\delta_\infty^\perp$ is also stable under $A(G)$, and the operator norm of $1/dA(G)$ on $\delta^\perp$ is at most $1-\varepsilon$, by the expander hypothesis. But the projection of δ_0 onto $\delta_\infty^\perp$ is precisely

$$\delta_0 - \frac{\langle \delta_0, \delta_\infty \rangle}{\|\delta_\infty\|_2^2}\delta_\infty = \delta_0 - \delta_\infty.$$

Thus $\delta_m - \delta_\infty \in \delta_\infty^\perp$ for all m, and

$$\|\delta_m - \delta_\infty\|_2 \leq (1-\varepsilon)^m \|\delta_0 - \delta_\infty\|_2 \leq (1-\varepsilon)^m \|\delta_0\|_2^2 = (1-\varepsilon)^m.$$

The last inequality in Proposition 2.2.5 comes from the usual comparison between $\|\cdot\|_2$ and $\|\cdot\|_1$. □

The quantity $\sum_{v \in V} |\delta_m(v) - \frac{1}{n}|$ is called the *total variation distance* between δ_m and the uniform distribution. In the particularly strong case of Ramanujan graphs, we obtain:

Corollary 2.2.6. *Let G be a d-regular Ramanujan graph on n vertices, and let $\eta > 0$ Then the distribution of the endpoint of a random walk of length $O((\log n - \log \eta)/\log d)$ starting from any given vertex in V has statistical distance at most η from the uniform distribution.*

Proof. In that case, we have $1 - \varepsilon = O(1/\sqrt{d})$. □

2.2.4 Supersingular isogeny graphs are Ramanujan

For applications to isogeny-based cryptography, a major theorem is that (undirected) ℓ-isogeny graphs of supersingular elliptic curves are Ramanujan. This is a difficult theorem from number theory, but we attempt to sketch its proof nonetheless.

Theorem 2.2.7. *Let $p = 1 \bmod 12$ be a prime, and let $\ell \neq p$. Then the ℓ-isogeny graph G of (minimal) supersingular elliptic curves over $\mathbb{F}_{p^2}$ is an $\ell+1$-regular Ramanujan graph with $(p-1)/12$ vertices.*

The congruence assumption on p is only a technical hypothesis ensuring that G is undirected.

Proof sketch. By the Deuring correspondence (Theorem 1), we can identigy $L^2(G)$ is the space of *quaternionic modular forms* on $B_{p,\infty}$, and the adjacency operator $A(G)$ is the ℓth *Hecke operator* on this space of quaternionic modular forms.

The Jacquet-Langlands correspondence, a key example of the Langlands program, assert that $L^2(G)$ is isomorphic to the space of classical modular forms of level 2 for $\Gamma_0(p)$, and that this isomorphism preserves the action of the Hecke operators. Under this correspondence, the functions which sum to zero on G (i.e. the orthogonal of 1) correspond to cusp forms; moreover, the inner product on $L^2(G)$ corresponds to the Petersson inner product on cusp forms.

Next, we use another major theorem in number theory, namely the Ramanujan conjecture proved by Deligne: the eigenvalues of the ℓth Hecke operator on modular cusp forms of weight 2 for $\Gamma_0(p)$ have absolute values at most $2\sqrt{\ell}$. This is what we had to prove as G is $(\ell+1)$-regular. □

Remark 2.2.8. The Ramanujan property is only known for isogeny graphs of supersingular elliptic curves: for the natural generalizations of these graphs to dimension 2 (superspecial abelian surfaces), it is unknown, and even false in general [JZ23].

An important consequence of Theorem 2.2.7 is that one can efficiently sample supersingular elliptic curves with a distribution that is exponentially close to the uniform distribution: starting from any vertex, we follow a random ℓ-isogeny walk of length $O(\log p)$ and output the result. When implementing the CGL family of hash functions, it is important to generate the seed in this way for the security proof of collision resistance.

2.3 Hard problems and security proofs in isogeny-based cryptography

2.3.1 Five hard problems

Here is the list of the main computational problems that isogeny-based cryptography is based on. We fix a prime number p. Throughout, when we write *random*, we mean uniformly random, and all the supersingular elliptic curves we consider are minimal.

Problem 2.3.1 (ℓ-IsogenyPath)**.** Given two random supersingular elliptic curves $E, E'/\mathbb{F}_{p^2}$ and a prime number ℓ, find a sequence of ℓ-isogenies over $\mathbb{F}_{p^2}$ connecting E and E'.

Problem 2.3.2 (Isogeny)**.** Given two random supersingular elliptic curves $E, E'/\mathbb{F}_{p^2}$, find an isogeny $E \to E'$ in efficient representation.

Problem 2.3.3 (OneEnd)**.** Given a supersingular elliptic curve $E/\mathbb{F}_{p^2}$, find a non-scalar endomorphism $\alpha \in \mathrm{End}(E)\backslash\mathbb{Z}$ in efficient representation.

Problem 2.3.4 (EndRing)**.** Given a supersingular elliptic curve $E/\mathbb{F}_{p^2}$, find endomorphisms $\alpha_1, \alpha_2, \alpha_3, \alpha_4 \in \mathrm{End}(E)$ in efficient representation such that $\mathrm{End}(E) = \mathbb{Z}\alpha_1 + \mathbb{Z}\alpha_2 + \mathbb{Z}\alpha_3 + \mathbb{Z}\alpha_4$.

Problem 2.3.5 (MaxOrder)**.** Given a supersingular elliptic curve $E/\mathbb{F}_{p^2}$, find an order $\mathcal{O} \subset B_{p,\infty}$ (specified as $\mathbb{Z}\alpha + \mathbb{Z}\beta + \mathbb{Z}\gamma + \mathbb{Z}\delta$ for some quaternions $\alpha, \beta, \gamma, \delta$) isomorphic to $\mathrm{End}(E)$.

In the Isogeny, OneEnd, and EndRing problems, we could request the efficient representation to be the CRT or HD representations: the CRT and HD representations are universal (Theorem 2.1.11), so this would not make the problems any harder.

These computational problems have been now been publicized and studied for almost 20 years; there is a good confidence in the community that they are indeed hard problems. Therefore, one considers that a cryptographic schemes offers strong security guarantees if one can reduce its security to one of the above problems.

Remark 2.3.6. This is far from being the case for every isogeny-based scheme: their security often relies on *ad hoc* computational problems that are introduced at the same time as the scheme itself, and consequently have not already been the focus of cryptanalytic effort. A typical example of such an "alternative" problem would be to ask to compute isogenies between supersingular curves over $\mathbb{F}_{p^2}$ that aren't uniformly random, but instead are sampled from a probability distribution arising from the protocol itself, and might perhaps not be well distributed: see e.g. [BDD+24, Assumption 2]. A more extreme example was the security assumption underlying SIDH [DJP11] (computing isogenies with extra torsion information), which turned out to be an easy problem.

In the rest of this section, we present security proofs reducing the preimage-resistance of the CGL family of hash functions to the 2-IsogenyPath problem, and its collision resistance to the OneEnd problem. These security proofs presuppose that the seed s, a supersingular curve over $\mathbb{F}_{p^2}$, is chosen (very close to) uniformly at random. This suggests the question of how to construct a random supersingular elliptic curve. An efficient way of doing this is to start from a known curve E_0, compute a chain of random small-degree isogenies, and output the codomain; if we restrict to 2-isogenies, this amounts to performing a random walk in the graph $G(p, 2)$. It is then necessary to know that the distribution of the endpoint of relatively short random walks will be close to undistinguishable from the uniform distribution. This turns out to be true: isogeny graphs of supersingular elliptic curves are *Ramanujan graphs*, i.e. optimal *expander graphs*. We review this theory in the next subsection.

2.3.2 CGL security reduces to ℓ-IsogenyPath and OneEnd

Proposition 2.3.7. *Assume that there exists a polynomial-time algorithm $\mathcal{A}$ breaking the preimage-resistance property of the CGL family of hash functions. Then one can construct a polynomial-time algorithm $\mathcal{B}$ solving the ℓ-IsogenyPath problem with non-negligible probability.*

Proof. On input E and E', algorithm $\mathcal{B}$ asks the oracle $\mathcal{A}$ for a preimage of E' under the CGL hash function with seed E. By assumption $\mathcal{A}$ outputs a valid message m whose hash is E', i.e. a non-backtracking ℓ-isogeny path from E to E', with non-negligible probability. □

Proposition 2.3.8. *Assume that there exists a polynomial-time algorithm $\mathcal{A}$ breaking the collision-resistance property of the CGL family of hash functions. Then one can construct a polynomial-time algorithm $\mathcal{B}$ solving the OneEnd problem with non-negligible probability.*

Proof. On input E, the algorithm $\mathcal{B}$ asks the oracle $\mathcal{A}$ for two distinct messages with the same hash under the CGL hash function with seed E, i.e. two distinct, non-backtracking ℓ-isogeny paths $E \to E_1 \to \cdots \to E_r$ and $E \to E'_1 \to \cdots \to E'_s = E_r$ to the same elliptic curve. Let $\phi, \psi : E \to E_r$ be the composed isogenies, and onsider the endomorphism $\psi^\vee \circ \phi : E \to E$. We claim that $\psi^\vee \circ \phi \notin \mathbb{Z}$. Indeed, assume the contrary. Then there exists $n \in \mathbb{Z}$ such that $\psi^\vee \circ \phi = \pm[\ell^n]$. In particular, $\ker(\phi) \subset E[\ell^n]$, and $\ker(\psi^\vee) = \phi(E[\ell^n])$. Because the paths giving ϕ and ψ are non-backtracking, both $\ker(\phi)$ and $\ker(\psi^\vee)$ are cyclic groups. On the other hand, $E[\ell^n] \simeq (\mathbb{Z}/\ell^n\mathbb{Z})^2$, so we must have $\deg(\phi) = \deg(\psi) = \ell^n$. Then the relation $\psi^\vee \circ \phi = \pm[\ell^n]$ shows that $\psi^\vee = \pm\phi^\vee$, so $\phi = \pm\psi$, i.e. the ℓ-isogeny paths are the same; a contradiction. □

Note that in this security reduction, the assumption that $\mathcal{A}$ only has access to the seed E of the CGL hash function, and nothing else, in crucial. In fact, whoever chooses the seed in an implementation of the CGL hash function has to be a trusted authority, as they could possibly know what $\mathrm{End}(E)$ is and solve the OneEnd problem on that curve. (Conversely, generating any non-scalar endomorphism of E of norm a power of ℓ breaks collision resistance.) There is currently no known way of sampling supersingular elliptic curves at random with no information whatsoever on their endomorphism rings: the random walk method from §2.2.2 certainly doesn't achieve this, as we can track down what the endomorphism rings are along an isogeny walk.

2.3.3 The equivalence between hard problems

To conclude, we discuss the following landmark result.

Theorem 2.3.9 ([Wes22; PW24]). *Under the Generalized Riemann Hypothesis (GRH), the five computational problems ℓ-IsogenyPath, Isogeny, OneEnd, EndRing, and MaxOrder are equivalent under probabilistic polynomial-time reductions.*

This result, together with the security reductions from the last subsection, guarantees for instance that the CGL family of hash functions is cryptographically secure is EndRing is computationally hard.

Before this theorem, certain heuristic reductions had also been described. We do not present a full proof of Theorem 2.3.9 (which represents full reseach papers), but instead indicate what the easy reductions are, and sketch the proofs of some of the harder reductions. We show them as dashed arrows along with the easy reductions (thick arrows) on the following diagram:

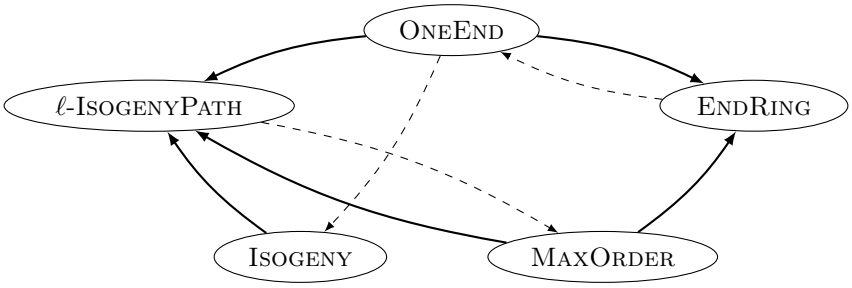

From this picture, the ℓ-IsogenyPath and EndRing problems may look, a priori, more difficult than the others. The easy reductions are as follows:

- OneEnd to EndRing: three (at least) of the basis elements in the output of EndRing do not lie in $\mathbb{Z}$, and we only need one to solve OneEnd.
- OneEnd to ℓ-IsogenyPath: choose any $\ell' \neq \ell$, and let E' be the codomain of an ℓ'-isogeny path ϕ from E. Solving ℓ-IsogenyPath provides an ℓ^n-degree isogeny $\psi : E' \to E$ for some $n \geq 1$. The composition $\psi \circ \phi$ is a non-scalar endomorphism of E.
- MaxOrder to EndRing: if we know a basis of $\mathrm{End}(E)$ in efficient representation, then we can compute the Gram matrix of the degree map (seen as a quadratic form on $\mathrm{End}^0(E)$) in this basis: if $\langle\cdot\rangle$ denotes the corresponding bilinear form, we have $\langle\alpha,\beta\rangle = \frac{1}{2}(\alpha\beta^\vee + \beta\alpha^\vee)$, and we can compute compositions and duals from the algorithms of §2.1.5. Once the Gram

matrix is known, there is a unique corresponding maximal order in $B_{p,\infty}$ up to conjugation, and this order is efficiently computable: see e.g. [Voi21, Thm. 22.1.1].

- MaxOrder to ℓ-IsogenyPath: let E be any supersingular elliptic curve, and E' an auxiliary supersingular curve with known endomorphism ring. Solving ℓ-IsogenyPath provides us with an ℓ-isogeny chain from E' to E. We can then compute the left ideal I in $\mathrm{End}(E')$ corresponding to the composite isogeny under the Deuring correspondence. The endomorphism ring of E is isomorphic to the right order of I, which is also efficiently computable.
- Isogeny to ℓ-IsogenyPath: a sequence of ℓ-isogenies from E to E' is an allowable efficient representation for the composite isogeny $E \to E'$.

The "heuristic reductions", on the other hand, are as follows:

- EndRing to OneEnd: starting from an elliptic curve E, one can construct many ℓ-power degree isogenies $\phi : E \to E'$ for several distinct small primes ℓ. One can call OneEnd on any such E' to get a non-scalar $\alpha \in \mathrm{End}(E')$; then $\phi^\vee \circ \alpha \circ \phi$ is also a non-scalar endomorphism of E. One can hope that the subrings of $\mathrm{End}(E)$ these endomorphisms generate will eventually cover the whole of $\mathrm{End}(E)$, leading to a solution of EndRing [EHL+18, §5]. However, it has been shown that this assumption is false for certain OneEnd oracles [PW24, §1.2].
- ℓ-IsogenyPath to MaxOrder: let E be any supersingular elliptic curve, and E' an auxiliary curve with known endomorphism ring. Solving MaxOrder on E provides us with a maximal order $\mathcal{O} \subset B_{p,\infty}$ isomorphic to $\mathrm{End}(E)$. Let $\mathcal{O}' \subset B_{p,\infty}$ be isomorphic to $\mathrm{End}(E')$. One can efficiently compute a connecting ideal I from $\mathcal{O}'$ to $\mathcal{O}$. The Kohel–Lauter–Petit–Tignol algorithm (KLPT) [KLP+14] then heuristically succeeds in computing another $\mathcal{O}'$-ideal J, equivalent to I, of ℓ-power norm; the right order of J is also isomorphic to $\mathrm{End}(E)$. We can then transform J into a sequence of ℓ-isogenies from E to E'.
- OneEnd to Isogeny: starting from an elliptic curve E, one can construct many ℓ-power degree isogenies $\phi : E \to E'$ for several distinct small primes ℓ. One can call Isogeny to get another isogeny $\psi : E' \to E$, and hope that the composition $\psi \circ \phi$ (for which we can compute an efficient representation) is non-scalar.

Note that the dashed arrows still turn the above diagram into a fully connected graph, i.e. our five computational problems are heuristically equivalent. Finally, we briefly discuss how the above reductions can be made rigorous.

- Given a OneEnd oracle, one can *enrich* it by composing the given endomorphisms with random walks in the ℓ-isogeny graph [PW24, Algorithm 1]. A generalized version of Theorem 2.2.7, [PW24, Thm. 1.3], then proves that the output of the enriched oracle satisfies good randomness properties [PW24, Thm. 4.2]. In particular, those properties imply that the output endomorphisms will eventually generate the endomorphism ring, perhaps with the help of the division algorithm from §2.1.5 [PW24, Thm. 7.2].
- The reduction of ℓ-IsogenyPath to MaxOrder is proved, under GRH, in [Wes22]: this technical paper provides adaptations to the KLPT algorithm that turn it into an expected polynomial-time algorithm [Wes22, Thm. 6.3]. This is the only place where GRH is used.
- Finally, one can show that the above heuristic reduction of OneEnd to Isogeny is actually valid [PW24, Thm. 8.6].

References

[BDD+24] A. Basso, P. Dartois, L. De Feo, A. Leroux, L. Maino, G. Pope, D. Robert, and B. Wesolowski. "SQIsign2D-West: the fast, the small, and the safer". In: *Advances in Cryptology – ASIACRYPT 2024*. Springer, 2024, pp. 339–370.

[Bea13] A. Beauville. "Theta functions, old and new". In: *Open Problems and Surveys of Contemporary Mathematics*. Vol. 6. Higher Education Press and International Press, 2013, pp. 99–131.

[BL04] C. Birkenhake and H. Lange. "Complex Abelian Varieties". 2nd ed. Springer, 2004.

[BJW17] E. H. Brooks, D. Jetchev, and B. Wesolowski. "Isogeny graphs of ordinary abelian varieties". *Research in Number Theory* 3 (2017), p. 28.

[Can00] D. G. Cantor. "Computing in the Jacobian of a hyperelliptic curve". *Mathematics of Computation* 48.177 (2000), pp. 95–101.

[CD23] W. Castryck and T. Decru. "An efficient key recovery attack on SIDH". In: *Advances in Cryptology – EUROCRYPT 2023*. IACR, 2023, pp. 423–447.

[CLG09] D. X. Charles, K. E. Lauter, and E. Z. Goren. "Cryptographic hash functions from expander graphs". *Journal of Cryptology* 22 (2009), pp. 93–113.

[Con] B. Conrad. *Polarizations*. URL: https://math.stanford.edu/~conrad/vigregroup/vigre04/polarization.pdf.

[DJP11] L. De Feo, D. Jao, and J. Plût. "Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies". In: Post-Quantum Cryptography. Springer, 2011, pp. 19–34.

[DKL+20] L. De Feo, D. Kohel, A. Leroux, C. Petit, and B. Wesolowski. "SQISign: compact post-quantum signatures from quaternions and isogenies". In: *Advances in Cryptology – ASIACRYPT 2020, Part I*. Springer, 2020, pp. 64–93.

[Deu41] M. Deuring. "Die Typen der Multiplikatorenringe elliptischer Functionenkörper". *Abh. Hamburg* 14 (1941), pp. 197–272.

[EvdGM12] B. Edixhoven, G. van der Geer, and B. Moonen. *Abelian varieties*. 2012. URL: http://van-der-geer.nl/~gerard/AV.pdf.

[EHL+18] K. Eisenträger, S. Hallgren, K. Lauter, T. Morrison, and C. Petit. "Supersingular isogeny graphs and endomorphism rings: reductions and solutions". In: *Advances in Cryptology – EUROCRYPT 2018*. Springer, 2018, pp. 329–368.

[Fal83] G. Faltings. "Endlichkeitssätze für abelsche Varietäten über Zahlkörper". *Invent. Math.* 73.3 (1983), pp. 349–366.

[FOR08] S. Flon, R. Oyono, and C. Ritzenthaler. "Fast addition on non-hyperelliptic genus 3 curves". In: *Algebraic Geometry and Its Applications*. Vol. 5. Ser. Number Theory Appl. World Sci. Pub., 2008, pp. 1–28.

[Har13] D. Harari. *Théorie du corps de classes*. 2013. URL: https://www.imo.universite-paris-saclay.fr/~david.harari/enseignement/corpsclass/poly.pdf.

[Har92] J. Harris. "Algebraic Geometry. A First Course". Springer, 1992.

[Har77] R. Hartshorne. "Algebraic Geometry". Springer, 1977.

[Hon68] T. Honda. "Isogeny classes of abelian varieties over finite fields". *J. Math. Soc. Japan* 20 (1968), pp. 83–95.

[JZ23] B. W. Jordan and Y. Zaytman. *Isogeny graphs of superspecial abelian varieties and Brandt matrices*. 2023. URL: http://arxiv.org/abs/2005.09031v5.

[Koh96] D. Kohel. "Endomorphism rings of elliptic curves over finite fields". University of California at Berkeley, 1996.

[KLP+14] D. Kohel, K. Lauter, C. Petit, and J.-P. Tignol. "On the quaternion ℓ-isogeny path problem". *LMS Journal of Computation and Mathematics* 17.A (2014), pp. 418–432.

[KMM+24] S. Kunzweiler, L. Maino, T. Moriya, C. Petit, G. Pope, D. Robert, M. Stopar, and Y. B. Ti. *Radical 2-isogenies and cryptographic hash functions in dimensions 1, 2 and 3*. 2024. URL: https://eprint.iacr.org/2024/1732.

[Lic11] S. Lichtenstein. *Tate's isogeny theorem for abelian varieties over finite fields*. 2011. URL: https://virtualmath1.stanford.edu/~conrad/mordellsem/Notes/L03.pdf.

[LR12] D. Lubicz and D. Robert. "Computing isogenies between abelian varieties". *Compositio Mathematica* 148.05 (2012), pp. 1483–1515.

[MMP+23] L. Maino, C. Martindale, L. Panny, G. Pope, and B. Wesolowski. "A direct key recovery attack on SIDH". In: *Advances in Cryptology – EUROCRYPT 2023*. Springer, 2023, pp. 448–471.

[Mar24] S. Marseglia. "Modules over orders, conjugacy classes of integral matrices, and abelian varieties over finite fields". In: *Sixteenth Algorithmic Number Theory Symposium (ANTS XVI)*. 2024.

[Mil86a] J. S. Milne. "Abelian varieties". In: *Arithmetic Geometry (Storrs, 1984)*. Springer, 1986, pp. 103–150.

[Mil86b] J. S. Milne. "Jacobian varieties". In: *Arithmetic Geometry (Storrs, 1984)*. Springer, 1986, pp. 167–212.

[Mil20] J. S. Milne. *Class Field Theory*. 2020. URL: https://www.jmilne.org/math/CourseNotes/cft.html.

[Mum66] D. Mumford. "On the equations defining abelian varieties. I". *Inventiones Mathematicae* 1 (1966), pp. 287–354.

[Mum70] D. Mumford. "Abelian Varieties". Oxford University Press, 1970.

[NOC+24] K. Nakagawa, H. Onuki, W. Castryck, M. Chen, R. Invernizzi, G. Lorenzon, and F. Vercauteren. "SQIsign2D-East: a new signature scheme using 2-dimensional isogenies". In: *Advances in Cryptology – ASIACRYPT 2024*. Springer, 2024, pp. 272–303.

[Nil91] A. Nilli. "On the second eigenvalue of a graph". *Discrete Math.* 91.2 (1991), pp. 207–210.

[OU73] F. Oort and K. Ueno. "Principally polarized abelian varieties of dimension 2 and 3 are Jacobian varieties". *J. Fac. Sci. Univ. Tokyo Sect. IA Math.* 20 (1973), pp. 377–381.

[PR23] A. Page and D. Robert. *Introducing Clapoti(s): evaluating the isogeny class group action in polynomial time*. 2023. URL: https://eprint.iacr.org/2023/1766.

[PW24] A. Page and B. Wesolowski. *The supersingular Endomorphism Ring and One Endomorphism problems are equivalent*. 2024. URL: https://eprint.iacr.org/2023/1399.

[RS85] M. O. Rabin and J. O. Shallit. "Randomized algorithms in number theory". *Comm. Pure Appl. Math.* 44 (1985), pp. 483–494.

[Rob21] D. Robert. *Efficient algorithms for abelian varieties and their moduli spaces*. 2021. URL: https://hal.science/tel-03498268v1.

[Rob23] D. Robert. "Breaking SIDH in polynomial time". In: *Advances in Cryptology – EUROCRYPT 2023*. Springer, 2023, pp. 472–503.

[Rob24] D. Robert. *On the efficient representation of isogenies (a survey)*. 2024. URL: https://eprint.iacr.org/2024/1071.

[Sil09] J. H. Silverman. "The Arithmetic of Elliptic Curves". 2nd ed. Springer, 2009.

[Tat66] J. Tate. "Endomorphisms of abelian varieties over finite fields". *Inventiones Mathematicae* 2 (1966), pp. 134–144.

[Voi21] J. Voight. "Quaternion Algebras". Springer, 2021.

[Wat69] W. C. Waterhouse. "Abelian varieties over finite fields". *Ann. scient. Éc. Norm. Sup.* 4.2 (1969), pp. 521–560.

[WM71] W. C. Waterhouse and J. S. Milne. "Abelian varieties over finite fields". In: *Proc. Symp. Pure Math.* 1971.

[Wes22] B. Wesolowski. "The supersingular isogeny path and endomorphism ring problems are equivalent". In: *FOCS 2021*. IEEE, 2022.