UNIVERSITÀ DI PISA

DIPARTIMENTO DI MATEMATICA

Laurea Magistrale in Matematica

Abelian varieties in the theta model and applications in cryptography

Relatori:

Prof. Benjamin Smith

Prof. Davide Lombardo

Candidato:

Alessandro Sferlazza

ANNO ACCADEMICO 2023/2024

Abstract

Isogenies of principally polarised abelian varieties have been used in recent years to build cryptographic protocols that are secure against quantum computers. Though abelian varieties are classical objects in algebraic geometry, from a computational perspective they present some challenges that have been addressed only recently.

An algorithmic framework to work with abelian varieties of any dimension is provided by theta models. These are projective realisations of polarised abelian varieties defined by algebraic theta functions.

This thesis presents an introduction to the arithmetic of principally polarised abelian varieties via the theory of theta models. It concerns itself with the computation of the group law on abelian varieties and the differential addition law on their Kummer varieties, pointing out the link with the corresponding existing elliptic curve algorithms.

Then, the thesis presents some recent algorithms to efficiently compute chains of $(2, \dots, 2)$-isogenies between Kummer surfaces, and the Tate and Weil pairings on general abelian varieties, including elliptic curves and hyperelliptic Jacobians. The theoretical framework needed for pairing computation also involves the algebraic theory of biextensions. Original contributions include implementations of some of the algorithms presented.

Finally, as a cryptographic application, the recent isogeny-based digital signature scheme SQIsign2D-West is studied, with a focus on the applicability of the higher-dimensional isogeny algorithms to signature verification in small devices.

Contents

Introduction

Secure communication in internet traffic is paramount in today's digital world, and public-key cryptography is at the heart of it. Protocols for key agreement, digital signatures, encryption guarantee secrecy, authenticity and integrity for the data sent over the internet.

Elliptic Curve Cryptography (ECC) is a form of public-key cryptography that leverages the mathematical properties of elliptic curves over finite fields. Introduced in 1985 by Victor Miller [Mil85] and Neal Koblitz [Kob87] independently, it is now ubiquitous on the internet: in 2021, almost all handshakes and 25% of the certificates used by the top million websites were ECC-based [Lab21]. Its advantage consists in providing high levels of security with significantly smaller key sizes compared to traditional methods like RSA [RSA78] (still the top-used signature algorithm; its keys are 10 times larger than ECC ones at the same security level), finite-field Diffie–Hellman [DH76]. Smaller keys mean less bandwidth and often faster computations, which is particularly important in environments with constrained resources, e.g. mobile devices, smart cards, IoT devices.

Security and efficiency in ECC are rooted on a broad understanding of computational problems in number theory, mathematically interesting per se, from point counting to pairings to endomorphism rings. The search for fast pairings – one of the main algorithmic tools in ECC – has led to a whole new branch of *pairing-based* cryptography, bringing advanced functionalities like three-party Diffie–Hellman key agreement [Jou00] and identity-based encryption [BF01], without prior analogues outside of ECC.

The security of ECC relies on the Elliptic Curve Discrete Logarithm Problem (ECDLP): given a point P of large prime order on a curve E and a multiple $Q = kP$ for some integer k, find k. It is really hard to solve on classical computers: the fastest known algorithms don't do better than the theoretical upper bound on their complexity, $O(\sqrt{\#E})$, exponential in the bitsize of the input.

Despite its advantages, ECC is not considered secure in the context of quantum computing. Shor's quantum algorithm [Sho94] poses a significant threat to ECC by potentially solving the ECDLP in *polynomial* time, thus undermining the security of systems that rely on it.

This vulnerability has sparked interest in the development of *post-quantum* cryptographic protocols, which are algorithms on classical computers designed to be secure against attacks by quantum computers. Ideally, such algorithms should be drop-in replacements for current pre-quantum ones in our security protocols, so as to assure a smooth transition to a quantum-safe world. The American National Institute of Standards and Technology (NIST) has launched in 2016 a competition [NIS16] to identify and standardise post-quantum algorithms, with the first generation of standards just released in August 2024. Research for the next generations is active and ongoing.

One promising area of research in this direction is isogeny-based cryptography, a new cryptographic paradigm based on isogenies of elliptic curves. It appeared in the early 2000s [Tes06], [CLG09], [RS06], [Cou06], based on some earlier works on isogeny graphs (e.g., [Koh96]). Its security relies on the hardness of the *isogeny problem*: given two elliptic curves E, E' over $\mathbb{F}_q$ satisfying $\#E(\mathbb{F}_q) = \#E'(\mathbb{F}_q)$, Tate's theorem says there is an $\mathbb{F}_q$-isogeny connecting them; the task is finding it explicitly.

The isogeny problem is hard to solve on classical computers, like ECDLP, but no efficient *quantum* algorithm is known either. Like ECC, isogeny-based crypto stands out due to its ability to provide strong security with significantly smaller key sizes – two orders of magnitude smaller than most of other post-quantum protocols [PQS]. On the other hand, isogeny-based cryptosystems are generally considerably slower than other post-quantum paradigms, and improving their performance is an important challenge.

In the past few years, isogeny-based cryptography has underwent a major breakthrough with the attacks [CD23], [MM22], [Rob23a] on the key-exchange protocol SIDH [JD11], former NIST competition candidate. These attacks have shown that isogenies between higher-dimensional abelian varieties allow the efficient computation of any isogeny between two elliptic curves in polynomial time, overcoming the exponential complexity of the best previously known algorithms.

Abelian varieties are smooth projective algebraic groups. While the theoretical foundations of these objects in algebraic geometry are classical and well understood, their practical implementation poses significant challenges. Research to generalise ECC to higher dimensions or higher-genus hyperelliptic curves is not new (see [CFA$^+$12] for a survey). When it comes to isogenies instead, the understanding of the relative computational problems is still in its early days, which is the motivation behind this thesis. A powerful framework to study these problems is the theory of algebraic theta functions, developed by Mumford [Mum66] and adapted to the cryptographic world by more recent works (most of whose results are collected in [Rob21]). This theory allows for the design of efficient arithmetic algorithms on principally polarised abelian varieties (PPAVs) – of *any* dimension – over finite fields. Therefore, it is crucial in bridging the gap between theory and practice and generalising the well-studied case of elliptic curves (the PPAVs of dimension 1) to higher dimensions.

The aim of this thesis is to give an overview of the computational problems on PPAVs and how they can be solved efficiently using the theory of theta functions. More precisely, given a g-dimensional PPAV A over a finite field k, the thesis will present how to:

- explicitly represent A as a projective variety;
- find algorithms linked to the group law of A. We'll see *differential addition*, which takes points $P, Q, P-Q \in A$ and outputs the sum $P+Q$, multiplication of a point

P by integer scalars, and more general operations called *Riemann relations*;

- compute the Weil and Tate pairings on A;
- compute isogenies: given A and a subgroup K, compute an isogeny φ having K as a kernel, that is, compute its codomain and evaluate φ at points of A. This is nontrivial when φ has high degree [LR12]. We'll focus on the easiest case, the $(2,\dots,2)$-isogenies, where K is isomorphic to $(\mathbb{Z}/2\mathbb{Z})^g$, which turns out to be very useful in applications – in particular, it makes the aforementioned breakthroughs possible.

We will also see how to do the above only working with the Kummer variety $A/\langle -1\rangle$, whose points are points of A "up to sign".

The thesis is organised as follows: Chapter 1 provides an overview of some algorithmic aspects of elliptic curves, serving as motivation and comparison for the next chapters. Chapter 2 is an introduction to abelian varieties and the theory of algebraic theta functions, giving the necessary algebraic-geometric definitions and defining the systems of coordinates we're going to use to represent our varieties. Chapter 3 presents arithmetic algorithms: the group-law-related ones and isogeny computations. Chapter 4 presents algorithms for the computation of the Weil and Tate pairings, combining theta functions with the theory of *biextensions* of abelian varieties and cubic torsors [Bre83]. These algorithms work efficiently for *general* primes and varieties, while most optimised pairing computations today are only possible on elliptic curves with tailored parameters. Finally, Chapter 5 presents the SIDH attacks and applies the study of theta functions to the isogeny-based digital signature scheme SQIsign [DKL$^+$20], currently a NIST candidate in a call for post-quantum signature schemes. In particular, the signature verification step is studied, aiming at improving its size-efficiency for a use in memory-constrained devices.

Algorithms in Chapter 3 and 4 mostly come from the recent works [DMPR23a], [Rob24a]. Using the code attached to these two works, we provide Sagemath implementations of biextension-based pairings on hyperelliptic Jacobians, available at `https://github.com/sferl/theta-pairings-dim2`, and analyse memory-efficiency of the available algorithms for $(2^n,\dots,2^n)$-isogeny evaluations.

Acknowledgements

This work was mainly realised while the author was an intern at the GRACE team of INRIA Saclay, Palaiseau, France, and was supported by the Erasmus+ Programme funded by the European Union, the HYPERFORM consortium funded by France through Bpifrance, and also by the France 2030 program, managed by the French National Research Agency under grant agreement No. ANR-22-PETQ-0008 PQ-TLS.

The author thanks Pierrick Dartois for his clear explanations and Maria Corte-Real Santos for fruitful discussions. Many thanks to the Isogeny Club for providing such a valid learning resource and creating a network of young nice researchers.

Chapter 1

Preliminaries: Elliptic curves and Montgomery arithmetic

Elliptic curves are a fundamental object in number theory and algebraic geometry, and have found numerous applications in cryptography [CFA$^+$12]. In this chapter, we will present some algorithmic aspects of elliptic curve arithmetic. We'll work with the Montgomery model, particularly well-suited to efficiently represent elliptic curves, their Kummer lines, their points and their group law on a computer, leading to efficient algorithms at the core of many cryptographic protocols. We will highlight some properties of the Montgomery model that will motivate the study of theta models in Chapter 2.

1.1 First definitions

Although we are going to recall the basic definitions and properties on elliptic curves and their arithmetic, a little familiarity with the subject is certainly helpful. The reader interested in a thorough theoretical introduction to the subject can refer to [Sil09]. A more algorithmic approach to elliptic curves, with an explicit focus on cryptographic applications, can be found instead in [Was08].

We will also assume some familiarity with algebraic geometry, in particular notions on algebraic varieties. The reader interested in a more detailed introduction to algebraic geometry can refer to any of the standard textbooks, such as [Liu02], [Vak24].

In this first chapter, for the sake of concreteness, we will define elliptic curves not as abstract objects, but as cubic plane curves embedded in the projective plane $\mathbb{P}^2$ via some *coordinate system*, and describe their arithmetic in terms of these coordinates. We will see in Chapter 2, when introducing higher-dimensional abelian varieties, that a more abstract definition can be given, and allows for greater generalisation.

Notation 1.1. Let's first introduce some notational conventions that we will use throughout the thesis:

- By $\mathbb{F}_q$ we will denote the finite field with $q = p^r$ elements, with $p = \mathrm{char}(\mathbb{F}_q)$.

- By k we always denote a finite field $\mathbb{F}_q$ of characteristic p, where $p \neq 2, 3$ is a large prime number. In the applications, typically p is several hundred bits large, that is, $\log_2(p) \approx 128, 256, 512$.
- We will denote the algebraic closure of a field k by k^{alg}.

Definition 1.2. An *elliptic curve* E over a field k (we will write E/k) is a smooth projective algebraic curve defined by the following homogeneous equation in $\mathbb{P}^2(k)$:

$$E : Y^2Z + a_1XYZ + a_3YZ^2 = X^3 + a_2X^2Z + a_4XZ^2 + a_6Z^3, \tag{1.1}$$

with $a_1, \dots, a_6 \in k$. We write E/k to say that "E is defined over k".

Definition 1.3. For any algebraic field extension F with $k \subseteq F \subseteq k^{\text{alg}}$, we can define the set of *F-rational points* of E as

$$E(F) = \{(X : Y : Z) \in \mathbb{P}^2(F) \mid (X, Y, Z) \text{ satisfies (1.1)}\}.$$

In particular, the k^{alg}-rational points $E(k^{\text{alg}})$ are also called the *geometric points* of E. There is a distinguished point $0_E = (0 : 1 : 0) \in E(k)$ that always belongs to the curve, and is called the *point at infinity*.

Notation 1.4. In the sequel, we will often write $P \in E$ to mean that P is a geometric point in $E(k^{\text{alg}})$. For such points $P = (X_P : Y_P : Z_P)$, we write $x(P) = X_P/Z_P$ and $y(P) = Y_P/Z_P$. These are coordinates on an affine chart where the line $Z = 0$ is sent to infinity, and they are well-defined rational functions in $k^{\text{alg}}(E)$.

Definition 1.5 (Change of coordinates)**.** Let $A \in \mathbb{P}\mathrm{GL}_2(k^{\text{alg}})$ be a linear change of coordinates in the ambient space $\mathbb{P}^2$ such that $A(0 : 1 : 0) = (0 : 1 : 0)$. The transformation $(X' : Y' : Z') = A(X : Y : Z)$ maps the equation E to a different cubic homogeneous equation $E'(X' : Y' : Z')$ with the same point at infinity. We say that E' is a different equation for the *same* curve, or that E and E' are *isomorphic* curves. If A is defined over k, we say E and E' to be *isomorphic over k*.

Definition 1.6 (Short Weierstrass model)**.** Keeping the notation of the previous definition, consider the following linear change of coordinates on $\mathbb{P}^2$:

$$X' = X + \frac{a_2}{3}, \quad Y' = Y + \frac{a_1}{2}X' + \frac{a_3}{2}, \quad Z' = Z,$$

which makes sense since $\mathrm{char}(k) \neq 2, 3$. Applying this to (1.1), the equation becomes

$$E : Y'^2Z' = X'^3 + aX'Z'^2 + bZ'^3 \tag{1.2}$$

for some $a, b \in k$ that satisfy $\Delta(E) = 4a^3 + 27b^2 \neq 0$.

This is called a *short Weierstrass model* of the elliptic curve E.

Definition 1.7. Let E/k be defined by an equation of the form (1.2). The quantity $j(E) = 1728 \cdot \frac{4a^3}{4a^4+27b^2} \in k$ is called the *j-invariant* of the elliptic curve E, and is an isomorphism invariant. More generally, for any elliptic curve E defined over k, we can define its j-invariant as $j(E_W)$ for any short Weierstrass model $E_W \cong E$ (since it is isomorphism invariant, it doesn't depend on the choice of E_W).

The Group law

One fundamental fact about elliptic curves is that they are *abelian varieties*, that is, any elliptic curve E possesses a commutative group law that makes its geometric points into an abelian group. This group law has a geometric description, sketched below.

Throughout the paragraph, we consider an elliptic curve defined by an equation

$$E : BY^2Z = F(X, Z), \tag{1.3}$$

for some $B \in k$, where $F \in k[X, Z]$ is a homogeneous degree-3 polynomial with no repeated factors. The short Weierstrass model is a particular case of this form, but we will see that other types of equations, like the Montgomery model, also fit within this framework.

Lemma 1.8. *The following map*

$$[-1]\colon E(k^{\mathrm{alg}}) \to E(k^{\mathrm{alg}}), \qquad P = (X_P : Y_P : Z_P) \mapsto -P = (X_P : -Y_P : Z_P)$$

defines an involution (an order-2 automorphism) of the curve, and leaves the point at infinity $0_E = (0 : 1 : 0)$ invariant.
The line through P and $-P$ is the vertical line $X = x(P) \cdot Z$, whose third intersection point with E is 0_E.

Definition 1.9. We call $[-1]$ the *negation map* on E.

The group law of an elliptic curve is defined geometrically as follows:

Proposition 1.10. *Let $P = (X_P : Y_P : Z_P)$ and $Q = (X_Q : Y_Q : Z_Q)$ be two points in E, such that $P \neq -Q$. Let $L_{P,Q}$ be the line through P and Q in $\mathbb{P}^2$. If $P = Q$, let $L_{P,Q}$ be the tangent line to E in P. This line will have three points of intersection with E: let $R = (X_R : Y_R : Z_R)$ be the third. Then we can define $P + Q = -R$ (see also Figure 1.1). The map*

$$+\colon E \times E \to E, \qquad (P, Q) \mapsto P + Q,$$

also called as the chord-tangent law, *is a commutative algebraic group law on $E(k^{\mathrm{alg}})$, with identity element the point at infinity 0_E. It is defined over the base field k. In particular, if $P, Q \in E(k)$ are k-rational, then $P + Q \in E(k)$ is too.*

Proof. Explicit formulas for the group law can be found in [Sil09, III, Group Law Algorithm 2.3], and show associativity.

Commutativity of the group law comes from the construction, since $L_{P,Q} = L_{Q,P}$.

The group law is defined over k because the coefficients of $L_{P,Q}$ are rational functions of the coordinates of P and Q and the coefficients of E. The slope λ of $L_{P,Q}$ equals the sum of the x-coordinates of P, Q, R, so if P, Q are k-rational then $x(P+Q) = x(R)$ also is, and $y(P + Q) = -y(R)$ depends linearly on $x(R)$ and the coefficients of $L_{P,Q}$.

For the point at infinity, note that given $P \in E$ the vertical line through P intersects E at $\{P, 0_E, -P\}$, which means $P + 0_E = -(-P) = P$. □

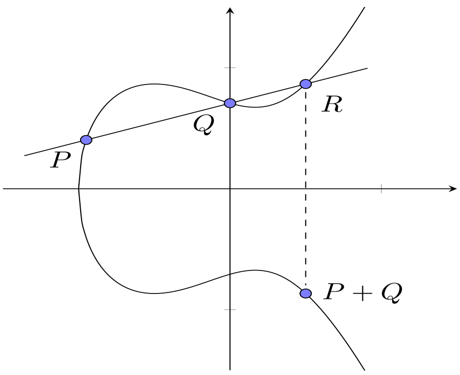

Figure 1.1: Group law on an elliptic curve (depicted over $\mathbb{R}$)

Using the group law, it is possible to define scalar multiplication:

Definition 1.11. Let $n \in \mathbb{Z}$ be an integer and $P \in E(k^{\mathrm{alg}})$ a point. The multiplication-by-n map is defined as

$$[n]\colon E(k^{\mathrm{alg}}) \to E(k^{\mathrm{alg}}), \qquad P \mapsto [n]P = \begin{cases} 0_E & \text{if } n = 0, \\ P + [n-1]P & \text{if } n > 0, \\ -[-n]P & \text{if } n < 0. \end{cases}$$

Definition 1.12. Let E/k be an elliptic curve. The *n-torsion subgroup* of E is the set

$$E[n] = \{P \in E(k^{\mathrm{alg}}) \mid [n]P = 0_E\}$$

and its points are called n-torsion points.

Proposition 1.13. *If n is an integer not divisible by p, the n-torsion subgroup $E[n]$ is a finite abelian subgroup of $E(k^{\mathrm{alg}})$ isomorphic to $\mathbb{Z}/n\mathbb{Z} \times \mathbb{Z}/n\mathbb{Z}$.*
If $n = p$, then $E[p]$ is either isomorphic to $\mathbb{Z}/p\mathbb{Z}$ or trivial. In the latter case, we say that E is supersingular*.*

Proof. See [Sil09, III, Corollary 6.4]. □

Example 1.14. Let $f(X) = F(X, 1) = (X - \alpha_1)(X - \alpha_2)(X - \alpha_3)$ where F defines the curve equation as in (1.3). The tangent line to E in $T = (\alpha_i : 0 : 1)$ is the vertical line $X = \alpha_i Z$, and its third point of intersection with E is $0_E = (0 : 1 : 0)$. Thus $T + T = 0$, that is, T is a 2-torsion point. Indeed, since $\#E[2] = 4$, this is all the 2-torsion:

$$E[2] = \{0_E, T_1 = (\alpha_1 : 0 : 1), T_2 = (\alpha_2 : 0 : 1), T_3 = (\alpha_3 : 0 : 1)\}.$$

The Montgomery model

We're interested in developing algorithms to express the group law on elliptic curves in a way that is efficient to compute. While the short Weierstrass model is quite convenient for theoretical purposes, it is not the most efficient model for computation. Instead, we can work with the *Montgomery model*, again a special case of (1.3), first introduced in [Mon87]. For more details on Montgomery arithmetic, see [CS17].

Definition 1.15. A Montgomery curve over a field k is an elliptic curve E defined by an equation of the form

$$E_{A,B}: \quad BY^2Z = X^3 + AX^2Z + XZ^2$$

with $A \in k \setminus \{\pm 2\}, B \in k \setminus \{0\}$.

For a Montgomery curve $E_{A,B}$, its j-invariant equals $j(E_{A,B}) = 256 \cdot \frac{(A^2-3)^3}{A^2-4}$.

Remark 1.16. Any elliptic curve in Montgomery form can be transformed into a short Weierstrass model via a linear change of coordinates defined over k, but the converse is not true: if E_W is a curve in Weierstrass form, there is always a k^{alg}-isomorphism from E_W to some $E_{A,B}$ in Montgomery form, but this isomorphism isn't k-rational in general. In fact, a Montgomery curve has some special structure over the base field k.

Indeed, a Montgomery curve $E_{A,B}$ always has a rational 2-torsion point. Note that $\alpha_1 = 0$ is always a root of the polynomial $F(X,1) = X^3 + AX^2 + X$. Example 1.14 shows that $T_0 = (0 : 0 : 1) \in E_{A,B}(k)$ is a point of order 2.

Remark 1.17 (Additional structure of $E_{A,B}(k)$). One can in fact show more: the curve $E_{A,B}$ has either a k-rational point of order 4, or two independent points of order 2 (that is, full k-rational 2-torsion). See [CS17, Section 2.3] for a proof.

The formulae for the group law on a Montgomery curve are quite simple (cf. the formulae for the short Weierstrass model in [Sil09]):

Proposition 1.18 (Explicit group law). *Let $P = (x_P : y_P : 1)$ and $Q = (x_Q : y_Q : 1)$ be two points on $E_{A,B}$ with $P \neq -Q$. Then $P + Q = (x_+ : y_+ : 1)$ where*

$$x_+ = B\lambda^2 - A - x_P - x_Q, \qquad y_+ = \lambda(x_P - x_Q) - y_P$$
$$\text{with} \quad \lambda = \begin{cases} \frac{y_P - y_Q}{x_P - x_Q} & \text{if } P \neq \pm Q, \\ \frac{3x_P^2 + 2Ax_P + 1}{2By_P} & \text{if } P = Q. \end{cases}$$

Though quite simple, the Montgomery group law is not as efficient as it could be: given a point $P = (X : Y : Z)$, the Y coordinate is uniquely determined (up to a sign) by the $(X : Z)$ coordinates as $Y = \sqrt{F(X,Z)/BZ}$. This means that the group law can be computed using only the $(X : Z)$-coordinates of the points, and the Y-coordinate can be recovered at the end of the computation if needed. This is the basis of the x-only arithmetic on Montgomery curves.

Consider the canonical involution $[-1]$ of Lemma 1.8 on a curve E. It is an automorphism of the curve, both as an algebraic variety and as a group. Quotienting by $\langle[-1]\rangle$ gives a double cover of E to $\mathbb{P}^1$:

$$\begin{array}{rrcll} \mathbf{x}\colon & E(k^{\mathrm{alg}}) & \to & \mathbb{P}^1(k^{\mathrm{alg}}) & \\ & P = (X:Y:Z) & \mapsto & (X:Z) = (x(P):1) & \text{if } (X,Z) \neq (0,0), \\ & 0_E = (0:1:0) & \mapsto & (1:0). & \end{array}$$

Definition 1.19. Via the cover $\mathbf{x}$, we identify the projective line $\mathbb{P}^1$ with the Kummer line $\mathcal{K}_E := E/\langle -1\rangle$ of the elliptic curve E. It is also called the x-line of E, since if we see $(1:0)$ as the point at infinity we are left with an affine line made of the x-coordinates of points of E.

The Kummer line of a Montgomery curve has another involution worth noting, that we'll see again when dealing with theta models in Chapter 2.

Proposition 1.20. *Let $T = (X_T : Y_T : Z_T)$ be a point of order 2 on E. The translation-by-T involution*

$$t_T\colon E \to E, \qquad P \to P + T$$

induces the k-automorphism of the Kummer line given by

$$t'_T : \mathbb{P}^1 \to \mathbb{P}^1, \qquad (X:Z) \mapsto (X_T X - Z_T Z : Z_T X - X_T Z).$$

In particular, if $T = T_0 = (0:0:1)$, then the transformation is just a permutation of the coordinates:

$$\mathbf{x}(P) = (X:Z) \quad \mapsto \quad \mathbf{x}(P + T_0) = (Z:X).$$

Proof. As seen in Example 1.14, $x_T = X_T/Z_T$ is a root of the polynomial $f(x) = x(x^2+Ax+1)$ and $y_T = Y_T/Z_T = f(x_T) = 0$. From the shape of the polynomial, $1/x_T$ is also a root of f, so $A = -(x_T + 1/x_T)$. The addition law 1.18 with $Q = T$ gives

$$B\lambda^2 = \frac{By_P^2}{(x_P - x_T)^2} = \frac{x_P(x_P^2 + Ax_P + 1)}{(x_P - x_T)^2}$$

and

$$\begin{aligned} x_+ = X_+/Z_+ = B\lambda^2 - (A + x_T) - x_P &= \\ &= \frac{x_P(x_P^2 + Ax_P + 1)}{(x_P - x_T)^2} - \frac{1}{x_T} - x_P \\ &= \frac{x_P x_T - 1}{x_P - x_T} = \frac{X_T X_P - Z_T Z_P}{Z_T X_P - X_T Z_P}. \end{aligned}$$

□

Arithmetic on the Montgomery Kummer line

Points on the Kummer line can be thought as points on the curve defined up to sign: the Kummer point $\mathbf{x}(P) \in \mathcal{K}_E(\mathbb{F}_q)$ is the image of two points P and $-P$ on E. They are not necessarily $\mathbb{F}_q$-rational (the Y-coordinate could be in $\mathbb{F}_{q^2}$).

The Kummer line is not an algebraic group, that is, it does not inherit an addition law from E. However, we still have a *pseudo-addition* law. Since points on the Kummer line are defined up to sign, the knowledge of $\mathbf{x}(P)$ and $\mathbf{x}(Q)$ only determines the set $\{\mathbf{x}(P+Q), \mathbf{x}(P-Q)\}$. Once we know the $\mathbf{x}$-coordinates of one of these two points, say $\mathbf{x}(P-Q) \in \mathcal{K}_E$, then we can uniquely recover the $\mathbf{x}$-coordinates of $P+Q$. Then the following is well-defined:

Definition 1.21. Let $E_{A,B}$ be a Montgomery curve, and let $P, Q \in E_{A,B}(k)$ be two points. We can define the *differential addition* operation on the Kummer line:

$$\mathsf{diff\_add}\colon \quad (\mathbf{x}(P), \mathbf{x}(Q), \mathbf{x}(P-Q)) \quad \mapsto \quad \mathbf{x}(P+Q) \in \mathcal{K}_E(k).$$

Warning: the differential addition law does not define a function on the whole $\mathcal{K}_E^3$, since it is defined only on triples of the form $(\mathbf{x}(P), \mathbf{x}(Q), \mathbf{x}(P-Q)) \in \mathcal{K}_E^3$.

In the special case where $P = Q$, then we get the doubling operation:

Definition 1.22. Let $P \in E_{A,B}(k)$ be a point. We can define the *doubling* operation on the Kummer line:

$$\mathsf{dbl}\colon \mathcal{K}_E(k) \to \mathcal{K}_E(k), \qquad \mathbf{x}(P) \mapsto \mathbf{x}(2P).$$

The above maps are not only well-defined, but also efficiently computable. If we know the $\mathbf{x}$-coordinates of two points P, Q and their difference $P - Q$, then we can compute the coordinates of $P + Q$ in a way that is faster than using the group law on E: see Algorithms 1 and 2.

Algorithm 1 x-only Montgomery differential addition

```
Input: $\mathbf{x}(P) = (X_P : Z_P), \mathbf{x}(Q) = (X_Q : Z_Q), \mathbf{x}(P-Q) = (X_- : Z_-)$ with $P, Q \in E(k)$
Output: $\mathbf{x}(P+Q) = (X_+ : Z_+)$
$U \leftarrow (X_P - Z_P) \cdot (X_Q + Z_Q)$
$V \leftarrow (X_P + Z_P) \cdot (X_Q - Z_Q)$
$X_+ \leftarrow Z_- \cdot (U+V)^2$
$Z_+ \leftarrow X_- \cdot (U-V)^2$
return $(X_+ : Z_+)$                         ▷ Complexity: 4 M, 2 S
```

We can do more, and compute scalar multiplications by any $n \in \mathbb{Z}$.

Definition 1.23. Let $P \in E_{A,B}(k)$ be a point on a Montgomery curve, n an integer. Since the elliptic involution $[-1]$ is compatible with $[n]$ (that is, $-[n]P = [n](-P)$), the following map is well defined:

$$\mathsf{ladder}\colon \mathbb{Z} \times \mathcal{K}_E(k) \to \mathcal{K}_E(k), \qquad (n, \mathbf{x}(P)) \mapsto \mathbf{x}([n]P).$$

Algorithm 2 x-only Montgomery doubling

```
Input: $\mathbf{x}(P) = (X_P : Z_P)$ with $P \in E(\mathbb{F}_q)$, and $A' = (A+2)/4$ where $A$ is the Montgomery coefficient of $E$
Output: $\mathbf{x}(2P) = (X_2 : Z_2)$
$Q \leftarrow (X_P + Z_P)^2$
$R \leftarrow (X_P - Z_P)^2$
$S \leftarrow Q - R$
$X_2 \leftarrow Q \cdot R$
$Z_2 \leftarrow S \cdot (R + A' \cdot S)$
return $(X_2 : Z_2)$                ▷ Complexity: 3 M, 2 S
```

Combining doubling and differential additions, we can get an algorithm to compute the scalar multiplication $\mathbf{x}([n]P)$ on the Kummer line in $O(\log n)$ operations.

Definition 1.24. Let $S = \{a_0, \dots, a_r\}$ be a sequence of integers such that $a_0 = 0, a_1 = 1$. It is called a *differential addition chain* if for each j in $(2, \dots, m)$ we can write $a_j = a_i + a_k$ for some $i, k < j$ and $a_i - a_k \in S$.

A differential addition chain induces an algorithm to compute $\mathbf{x}([a_r]P)$:

- At the beginning, we know $(1 : 0)$ and $\mathbf{x}(P)$.
- For each j in $(2, \dots, r)$, we write $a_j = a_i + a_l$ and compute $\mathbf{x}([a_j]P) = \mathsf{diff\_add}(\mathbf{x}([a_i]P), \mathbf{x}([a_l]P), \mathbf{x}([a_i - a_l]P))$.

Proposition 1.25. *The Montgomery ladder algorithm shown in Algorithm 3 computes the scalar multiplication* $\mathbf{x}([m]P)$ *on a Montgomery curve* $E_{A,B}$ *in* $O(\log m)$ *operations.*

Proof. Let $r = \lfloor \log m \rfloor$. We show by induction that the Montgomery ladder is induced by the following differential addition chain for m, of length $2r$:

$$S = \{\lfloor m/2^i \rfloor - 1, \lfloor m/2^i \rfloor \mid i \in (r, \dots, 0)\}.$$

Indeed, we show that the ladder keeps the following invariants:

$$R_0 = \mathbf{x}([\lfloor m/2^i \rfloor]P), \qquad R_1 = \mathbf{x}([\lfloor m/2^i \rfloor + 1]P).$$

The base case $i = r$ is clear. Suppose the invariants hold at step i, then if $m_i = 0$ we have $R_0 = \mathsf{dbl}(R_0) = \mathbf{x}([\lfloor m/2^{i-1} \rfloor])$ and $R_1 = \mathsf{diff\_add}(R_0, R_1, \mathbf{x}(P)) = \mathbf{x}([\lfloor m/2^i \rfloor + \lfloor m/2^i \rfloor + 1]P) = \mathbf{x}([\lfloor m/2^{i-1} \rfloor + 1]P)$. If $m_i = 1$, the computation is symmetric. □

Remark 1.26. In the Montgomery ladder, every time we compute $\mathsf{diff\_add}$, the difference point is P. This enhances efficiency in the algorithm, since we can re-use the quantities depending on $\mathbf{x}(P)$ at each step.

Algorithm 3 x-only Montgomery ladder

```
Input: $m = \sum_{i=0}^{r} m_i 2^i$ and $\mathbf{x}(P) = (X_P : 1)$ with $P$ in $E(\mathbb{F}_p)$
Output: $\mathbf{x}([m]P)$
(R_0, R_1) ← $(\mathbf{x}(0_E) = (1 : 0), \mathbf{x}(P))$
for $i$ in $(r, \dots, 0)$ do
    if $m_i = 0$ then
        $(R_0, R_1) \leftarrow (\mathsf{dbl}(R_0), \mathsf{diff\_add}(R_0, R_1, \mathbf{x}(P)))$
    else
        $(R_0, R_1) \leftarrow (\mathsf{diff\_add}(R_0, R_1, \mathbf{x}(P)), \mathsf{dbl}(R_1))$
return $R_0$
```

1.2 Isogenies

Definition 1.27. Let E_1 and E_2 be elliptic curves over a field k. An *isogeny* $\varphi\colon E_1 \to E_2$ is a non-constant morphism of algebraic varieties such that $\varphi(0_{E_1}) = 0_{E_2}$. An isogeny is said to be *separable* (resp. k-rational) if it is separable (resp. k-rational) as a morphism of varieties. We denote by $\deg \varphi = [k(E_1) : \varphi^* k(E_2)]$ its degree.

Remark 1.28. From now on, all isogenies will be assumed to be *separable*, unless otherwise specified.

Proposition 1.29. *Let $\varphi\colon E_1 \to E_2$ be an isogeny. On geometric points, it induces*

$$\varphi(k^{\mathrm{alg}})\colon E_1(k^{\mathrm{alg}}) \to E_2(k^{\mathrm{alg}})$$

which is a surjective map and a group homomorphism, whose kernel $K = \ker \varphi$ is a finite subgroup of $E_1(k^{\mathrm{alg}})$. [1]

If $K \leq E_1(k^{\mathrm{alg}})$ is a subgroup, then there exists a unique separable isogeny $\varphi\colon E_1 \to E_2$ such that $\ker \varphi = K$, up to post-composition with an isomorphism on the codomain. Its degree is $\deg \varphi = \#K$.

Proof. [Sil09, II, Theorem 2.3] shows surjectivity. For the group homomorphism properties, see [Sil09, III, Theorem 4.8, Corollary 4.9]. The equality $\deg \varphi = \#K$ is [Sil09, II, Proposition 2.6b]. Uniqueness of φ given K follows from the following lemma. □

Lemma 1.30 ([Sil09, III, Corollary 4.11])**.** *Let E_1 and E_2 be elliptic curves, $\varphi\colon E_1 \to E_2$ a separable isogeny, $\varphi'\colon E_1 \to E_3$ another isogeny such that $\ker \varphi \subseteq \ker \varphi'$. Then there exists a unique isogeny $\psi\colon E_2 \to E_3$ such that $\varphi = \psi \circ \varphi'$.*

Since it coincides with the degree as a morphism of varieties, the degree of an isogeny is multiplicative:

Proposition 1.31. *Let $\varphi\colon E_1 \to E_2$ and $\varphi'\colon E_2 \to E_3$ be isogenies of elliptic curves. Their degrees satisfy $\deg(\varphi' \circ \varphi) = \deg(\varphi') \deg(\varphi)$.*

[1] Technically speaking, the kernel $\ker \varphi$ is a finite subscheme of E_1. In the sequel, we restrict our attention to separable isogenies, and abuse notation by identifying $\ker \varphi$ with $\ker \varphi(k^{\mathrm{alg}})$.

Isogeny computation

Devising efficient algorithms to compute isogenies is a fundamental problem in computational number theory, with several practical applications, of which we will see some in Chapter 5. We state more precisely what *computing isogenies* means:

Problem 1.32 (Isogeny computation)**.** Let E_1 be an elliptic curve over a finite field k, $P_1, \ldots, P_r \in E_1(k)$ some points generating the kernel $K = \langle P_1, \ldots, P_r \rangle$ of a separable isogeny $\varphi \colon E_1 \to E_2 = E_1/K$. By *computing* φ, we mean running an algorithm that takes as input the points $P_1, \ldots, P_r$ and outputs:

1. Coefficients defining the codomain curve $E_2 = E_1/K$;
2. An efficient algorithm that, on input any point $R \in E_1$, outputs the image $\varphi(R)$.

If an algorithm only does part 1 it is a *codomain* algorithm, if it only does part 2 it is an *evaluation* algorithm.

In practice, since the representation of elliptic points on the Kummer line is particularly efficient, we also consider the relaxed problem:

Problem 1.33 (Kummer isogeny computation)**.** Keep the notation of Problem 1.32. An algorithm for Kummer isogeny computation takes as input the Kummer points $\mathbf{x}(P_1), \ldots, \mathbf{x}(P_r)$ and outputs:

- Those coefficients defining $E_2 = E_1/K$ that are necessary to compute the arithmetic on $\mathcal{K}_{E_2}$;
- An efficient algorithm that, on input any Kummer point $\mathbf{x}(R) \in \mathcal{K}_{E_1}$, outputs the image $\mathbf{x}(\varphi(R))$.

In view of the applications, we are mostly interested in cyclic separable isogenies:

Definition 1.34. Let E_1/k be an elliptic curve, and $P \in E_1[\ell]$ a point of order $\ell \nmid p$. Then the isogeny with kernel $K = \langle P \rangle \cong \mathbb{Z}/\ell\mathbb{Z}$ is called a cyclic *ℓ-isogeny.*

Several algorithms exist to compute ℓ-isogenies. If ℓ is prime, an ℓ-isogeny can be computed using Vélu's formulas (see [Was08, Section 12.3]) which have a complexity of $O(\ell)$ operations in k. For large ℓ, there is the more convenient VéluSqrt algorithm [BDLS20] of complexity $O(\sqrt{\ell})$. This is quite slow for cryptographic applications, where degrees of isogenies are typically several hundreds bits large. In Chapter 5, we will see that given an ℓ-isogeny $\varphi \colon E_1 \to E_2$, if we know both E_1 and E_2 and the image of a point $P \in E_1$ of order 2^n for large enough n, we can compute φ in polynomial time wirth respect to $\log \ell$ with the help of isogeny computation on abelian surfaces (the 2-dimensional analogue of elliptic curves).

When ℓ is very small instead, using Vélu's formulas is still pretty efficient: see for example the following proposition for the case $\ell = 2$ in the Montgomery model.

Proposition 1.35 ([Ren18, Section 4.2])**.** *Let E be a Montgomery elliptic curve, and $P \neq (0 : 0 : 1) \in E[2]$ be a 2-torsion point, represented as $\mathbf{x}(P) = (X_P : Z_P)$ in the Kummer variety. Then Algorithm 4 is a Kummer isogeny evaluation algorithm for $\varphi \colon E \to E/\langle P \rangle$ in the sense of Problem 1.33.*

Algorithm 4 x-only Montgomery 2-isogeny

```
Input: $\mathbf{x}(P) = (X_P : Z_P)$ where $P \in E[2]$ generates the kernel of $\varphi\colon E \to E' = E/\langle P \rangle$, any point $\mathbf{x}(R) = (X : Z)$.
Output: The point $\mathbf{x}(\varphi(R)) = (X' : Z')$.
$U \leftarrow (X + Z) \cdot (X_P - Z_P)$
$V \leftarrow (X - Z) \cdot (X_P + Z_P)$
$X' \leftarrow X \cdot (U + V)$
$Z' \leftarrow Z \cdot (U - V)$
return $(X' : Z')$                                  ▷ Complexity: 4 M
```

If ℓ is a composite number with small factors, we can compute the isogeny as a composition of (fast) smaller isogenies.

Definition 1.36. An integer N is called a *B-smooth* number if all its prime factors are smaller than a bound B. The *smoothness* of N is the largest prime factor of N.

Proposition 1.37. *Let $\varphi\colon E_0 \to E_r$ be an N-isogeny, where $N = \deg(\varphi) = \prod_{i=1}^{r} \ell_i$ is a B-smooth number, with $\ell_i \leq B$ not necessarily distinct. Then φ can be computed as a composition of smaller-degree isogenies $\varphi = \varphi_r \circ \cdots \circ \varphi_1$, where $\deg \varphi_i = \ell_i$, via Algorithm 5, with complexity $O(r(B + r \log B))$.*

Definition 1.38. An N-isogeny $\varphi = \varphi_r \circ \cdots \circ \varphi_1$ is also called an *isogeny chain* of length r.

Proof of Proposition 1.37. The decomposition of φ could be proven via a repeated application of Lemma 1.30 (see [Gal12, Theorem 25.1.2]), but we give a more explicit proof.

Let $K = \ker \varphi = \langle P \rangle$. For simplicity, we prove correctness of the algorithm when $N = 2^r$, that is, $\ell_i = 2$ for all i.

Set $K_0 = \left\langle [2^{r-1}]P \right\rangle$ and inductively define, for $i = 1, \ldots, r$:

$$\varphi_i\colon E_{i-1} \to E_i = E_{i-1}/K_{i-1}, \qquad K_i = [2^{r-1-i}]\varphi_i \circ \cdots \circ \varphi_1(K).$$

We claim that $\ker(\varphi_i \circ \cdots \circ \varphi_1) = \left\langle [2^{r-i}]P \right\rangle$ for all $i = 1, \ldots, r$.

From this, the conclusion follows: for $i = r$ we have $\ker(\varphi_r \circ \cdots \circ \varphi_1) = \langle P \rangle = \ker \varphi$, so indeed φ decomposes as $\varphi_r \circ \cdots \circ \varphi_1$ by Proposition 1.29.

Moreover, each step has indeed degree 2. In fact, $H_i := \varphi_i \circ \cdots \circ \varphi_1(K)$ is isomorphic to $\langle P \rangle / \left\langle [2^{r-i}]P \right\rangle$, which has cardinality 2^{r-i}. We then write $\# \ker \varphi_{i+1} = \# K_i = \#([2^{r-i-1}]H_i) = 2^{r-i}/2^{r-i-1} = 2$, as desired.

We prove the claim by induction on i. For $i = 1$ there's nothing to prove. Assuming the conclusion true for $i - 1$, we have

$$\begin{aligned} \ker(\varphi_i \circ \cdots \circ \varphi_1) = \{R \in E_0 \mid \varphi_{i-1} \cdots \varphi_1(R) \in K_{i-1} = \varphi_{i-1} \cdots \varphi_1(\left\langle [2^{r-i}]P \right\rangle)\} = \\ = \left\langle \ker(\varphi_{i-1} \circ \cdots \circ \varphi_1), [2^{r-i}]P \right\rangle = \\ = \left\langle [2^{r-i-1}]P, [2^{r-i}]P \right\rangle = \left\langle [2^{r-i}]P \right\rangle. \end{aligned}$$ □

Algorithm 5 Smooth isogeny chain computation

```
Input: An elliptic point $P \in E_0[N]$ with $N = \prod_{i=0}^{r} \ell_i$, a $B$-smooth integer, $p \nmid N$.
Output: The isogeny $\varphi\colon E_0 \to E_r$ with kernel $K = \langle P \rangle$, in the sense of Problem 1.32.
$P_0 \leftarrow P$
for $i$ in $(1, \dots, r)$ do                                          ▷ $r$ steps
    $R_i \leftarrow [N/\prod_{j=1}^{i} \ell_j]P_{i-1}$                  ▷ Cost: $O(\log N) = O(r \log B)$
    Compute the $\ell_i$-isogeny $\varphi_i\colon E_{i-1} \to E_i$ with kernel $K_i = \langle R_i \rangle$
                                                                        ▷ Cost: $O(\ell_i) = O(B)$
    Push the kernel through the isogeny: $P_i \leftarrow \varphi_i(P_{i-1})$    ▷ $H_i = \langle P_i \rangle$
return Compose$(\varphi_r, \dots, \varphi_1)$
```

1.3 Divisors and pairings

This section is dedicated to another computational problem on elliptic curves, that of computing pairings. In order to define them, we need to introduce the fundamental concept of divisors.

Definition 1.39. Let E/k be an elliptic curve. A *divisor* on E is a formal sum of points on E with integer coefficients:

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

with $n_P \in \mathbb{Z}$ and $n_P = 0$ for all but finitely many P.

The *support* of a divisor D is the set $\mathrm{supp}(D) = \{P \in E \mid n_P \neq 0\}$.

The *degree* of a divisor D is the sum of its coefficients: $\deg(D) = \sum_{P \in E} n_P$.

The set of all divisors on E forms a group under addition, and is denoted by $\mathrm{Div}(E)$. Its subgroup of divisors of degree 0 is denoted by $\mathrm{Div}^0(E)$.

Definition 1.40. Let $k(E)$ be the field of rational functions on the curve E. The *principal divisor* associated with a function $f \in k(E)^*$ is the divisor

$$\mathrm{div}\, f = \sum_{P \in E} \mathrm{ord}_P(f) \cdot (P)$$

where $\mathrm{ord}_P(f)$ is the order of vanishing of f at P.

The set of all principal divisors on E is denoted by $\mathrm{Princ}(E)$ and is contained in $\mathrm{Div}^0(E)$ ([Sil09, Proposition II.3.1]).

The only functions whose divisor is 0 are the constants in k^*. Therefore, if some functions $f, g \in k(E)^*$ have the same divisor, they differ by a nonzero constant in k^*.

Example 1.41. Let E be an elliptic curve in (affine) Montgomery form $By^2 = x(x^2 + Ax + 1)$ with $x = X/Z, y = Y/Z$ the affine coordinates relative to the projective coordinates $(X : Y : Z)$.

- Let $P \neq 0_E$, and $f = x - x(P)$. Then $\mathrm{div}\, f = (P) + (-P) - 2(0_E)$.

- Let $P, Q \in E$ such that $P \neq -Q$. Let $f = y - (mx + q)$ be the equation defining the line $L_{P,Q}$ through P and Q (the tangent line, if $P = Q$). Let R be the third point of intersection of $L_{P,Q}$ with E. Then $\operatorname{div} f = (P) + (Q) + (R) - 3(0_E)$.

This comes from the fact that x and y have poles of order 2 and 3 at 0_E respectively.

Definition 1.42. Define the degree-0 Picard group $\operatorname{Pic}^0(E)$ as the group of divisors of degree 0 modulo principal divisors: $\operatorname{Pic}^0(E) = \operatorname{Div}^0(E)/\operatorname{Princ}(E)$.

Two divisors that differ by a principal divisor are said to be *linearly equivalent*.

The following Theorem is the first manifestation of the *polarisations* that we'll encounter in Chapter 2 when dealing with abelian varieties.

Theorem 1.43 ([Sil09, III, Proposition 3.4])**.** *Every elliptic curve is canonically isomorphic to its degree-*0 *Picard group, via the following isomorphism:*

$$\begin{array}{rccc} \lambda_{(0_E)}\colon & E & \to & \operatorname{Pic}^0(E) \\ & P & \mapsto & [(P) - (0_E)] \end{array} \tag{1.4}$$

with inverse

$$\begin{array}{rccc} \lambda_{(0_E)}^{-1}\colon & \operatorname{Pic}^0(E) & \to & E \\ & [\sum_i n_i(P_i)] & \mapsto & \sum_i [n_i]P_i. \end{array}$$

Pairings

We are now ready to define some important pairings on elliptic curves. The interested reader can find detailed treatments in [Gal05] and [CFA$^+$12, Chapters 6, 16].

Definition 1.44. Let E/k be an elliptic curve, and let $\ell \in \mathbb{N}$ be an integer coprime with $p = \operatorname{char}(k)$. The ℓ-Weil pairing on E is a bilinear map

$$e_{W,\ell}\colon E[\ell] \times E[\ell] \to \mu_\ell$$

where μ_ℓ is the group of ℓ-th roots of unity in k^{alg}, satisfying some additional properties:

- skew-symmetry: $e_{W,\ell}(P, Q) = e_{W,\ell}(Q, P)^{-1}$ for all $P, Q \in E[\ell]$;
- non-degeneracy: $e_{W,\ell}(P, Q) = 1$ for all $Q \in E[\ell]$ if and only if $P = 0_E$.
- Galois equivariance: for all $\sigma \in \operatorname{Gal}(k^{\mathrm{alg}}/k)$, $e_{W,\ell}(\sigma(P), \sigma(Q)) = \sigma(e_{W,\ell}(P, Q))$ for all $P, Q \in E[\ell]$.
- compatibility with scalars: $e_{W,\ell\ell'}(P, Q) = e_{W,\ell}([\ell']P, Q)$ for all $P \in E[\ell\ell'], Q \in E[\ell]$.
- compatibility with isogenies: $e_{W,\ell}(\varphi(P), \varphi(Q)) = e_{W,\ell}(P, Q)^{\deg \varphi}$ for all $P, Q \in E[\ell]$ and all isogenies $\varphi\colon E \to E'$.

Concretely, we can construct the pairing via the use of divisors.

Notation 1.45. Let $f \in k(E)$ be a rational function, and $D = \sum_i n_i(P_i)$ a divisor with support disjoint from $\operatorname{div} f$. We can evaluate f on the divisor D as

$$f(D) = \prod_i f(P_i)^{n_i}.$$

Definition 1.46 (Weil pairing: operational definition)**.** Let $P, Q \in E[\ell]$ be points of order ℓ. The divisors $D_P = (P) - (0_E) = \lambda_{(0_E)}(P), D_Q = (Q) - (0_E)$ are elements of order ℓ in $\mathrm{Pic}^0(E)$. Choose $D'_Q \sim D_Q \in \mathrm{Pic}^0(E)$ with support disjoint from that of D_P. Then let $f_P, f_Q \in k(E)$ be rational functions with divisor respectively $\operatorname{div} f_P = \ell D_P = \ell(P) - \ell(0_E)$, and analogously $\operatorname{div} f_Q = \ell D'_Q$. The map

$$\begin{array}{rccc} e\colon & E[\ell] \times E[\ell] & \to & k^* \\ & (P,Q) & \mapsto & \frac{f_P(D'_Q)}{f_Q(D_P)} \end{array} \qquad D'_Q \sim D_Q \text{ in } \mathrm{Pic}^0(E)$$

is bilinear, maps $E[\ell] \times E[\ell]$ onto the group μ_ℓ of the ℓ-th roots of unity in k^*, and satisfies the additional properties listed above.

Definition 1.47. If P is a point of ℓ-torsion, a function f_P whose divisor is $\ell(P)-\ell(0_E)$ is called a *Miller function* for P.

In the algorithmic applications, another useful pairing on elliptic curves is the Tate pairing. Unlike the Weil pairing, its codomain depends on the base field k:

Definition 1.48. Let $P \in E[\ell](k), Q \in E(k)/E[\ell](k)$. Let f_P be a rational function of divisor $\operatorname{div} f_P = \ell((P) - (0_E))$, and let D_Q be a divisor that is equivalent to $(Q) - (0)$ but whose support is disjoint from $\{P, 0\}$. The map

$$\begin{array}{rccc} e_{T,\ell}\colon & E(k)[\ell] \times E[\ell] & \to & k^*/(k^*)^\ell \\ & (P,Q) & \mapsto & f_P(D_Q) \end{array}$$

is bilinear, non-degenerate, satisfies the properties of compatibility with scalars and with isogenies listed in 1.44. It is called the *ℓ-Tate pairing* on E, and its value is defined up to ℓ-th powers in k.

We want to find an algorithm to compute these pairings efficiently given P, Q, ℓ. Using the definitions above, to do so one could try and find a way to explicitly compute coefficients for the rational function f_P (and f_Q for the Weil pairing), but these functions have prohibitively large degree when ℓ is large. However, it is sufficient to find an evaluation algorithm that on input $X \in E(k)$ outputs $f_P(X)$: this can be done via Miller's algorithm.

Proposition 1.49 (Miller's algorithm)**.** *Let E/k be an elliptic curve, $P \in E[\ell](k)$ a point of order ℓ, and $D \in \mathrm{Div}(E)$ with support disjoint from $\{P, 0_E\}$. Then Algorithm 6 computes the evaluation of the Miller function f_P at the divisor D in $O(\log \ell)$ operations in k.*

In particular, the cost of computing the Tate pairing (resp. the Weil pairing, when $Q \in E[\ell]$) is $O(\log \ell)$, since it amounts to one (resp. two) Miller function evaluations.

Sketch. The key idea of Miller's algorithm is to consider the group law on E and "see it in $\mathrm{Div}(E)$". More precisely, let $P, Q \in E$, set $R = P + Q$. Then, in $\mathrm{Div}(E)$:

$$(P) - (0_E) + (Q) - (0_E) - ((R) - (0_E)) = \operatorname{div}(g_{P,Q})$$

for some $g_{P,Q} \in k(E)$: a point addition in E corresponds to a certain principal divisor in $\mathrm{Div}(E)$.

This $g_{P,Q}$ can be computed as follows. Let $L_{P,Q}$ be the line through P, Q and V_R the vertical line through $R, -R$. Then:

$$\operatorname{div} L_{P,Q} = (P) + (Q) + (-R) - 3(0_E), \qquad \operatorname{div} V_R = (R) + (-R) - 2(0_E),$$

hence we can set $g_{P,Q} = L_{P,Q}/V_R$.

Now let $f_{i,P}$ be a function with divisor $i(P) + -([i]P) - (i-1)(0_E)$, so that if $P \in E[\ell]$ then the function $f_{\ell,P}(Q)$ returned at the end is a Miller function for P evaluated at Q.

At each step of the algorithm, we have $R = [\lfloor \ell/2^{i+1} \rfloor]P =: [m_i]P$, as in a double-and-add algorithm, and we show that f equals $f_{m_i,P}$ (up to a constant). Before entering the loop (set conventionally $i = b = \lfloor \log \ell \rfloor$) this is true: $m_i = 1$ and $f = f_{1,P}$ has trivial divisor.

For the inductive step, suppose we have $(R, f) = ([m_i], f_{m_i,P})$. Then we compute $[2]R = [2m_i]P$ and $f^2 \cdot g_{R,R}(Q)$ whose divisor is

$$\begin{aligned}(2m_i(P) + -2([m_i]P) - 2(m_i - 1)(0_E)) + (2([m_i]P) - ([2m_i]P) - (0)) =\\ = 2m_i(P) + -([2m_i]P) - (2m_i - 1)(0_E) = \operatorname{div} f_{2m_i,P}.\end{aligned}$$ □

Algorithm 6 Miller's Algorithm

Input: $P \in E[\ell](k)$, a divisor $D_Q \sim (Q) - (0_E) \in \mathrm{Div}(E)$ such that $P, 0_E \notin \operatorname{supp} D_Q$, an integer $\ell = \sum_{i=0}^{b} \ell_i 2^i$

Output: The non-reduced Tate pairing $e_{T,\ell}(P, Q) = f_{\ell,P}(D_Q)$

```
$(R, f) \leftarrow (P, 1)$
for $i$ in $(b-1, \dots, 0)$ do
    $(R, f) \leftarrow ([2]R, f^2 \cdot \frac{L_{R,R}}{V_{[2]R}}(D_Q))$
    if $\ell_i = 1$ then
        $(R, f) \leftarrow (R + P, f \cdot \frac{L_{R,P}}{V_{R+P}}(D_Q))$        ▷ $R = [\lfloor \ell/2^{i-1} \rfloor]P$, $f = f_{\lfloor \ell/2^{i-1} \rfloor,P}$
return $f$
```

Chapter 2

Higher dimensions

2.1 Abelian varieties

In the last chapter, we introduced elliptic curves and some of their algorithmic aspects: an efficient coordinate representation, the group law, isogenies and pairings. The following three chapters will be devoted to generalising these aspects to general *principally polarised abelian varieties*, a higher-dimensional analogue of elliptic curves.

In this chapter, we will begin by setting a theoretical framework to work with abelian varieties, and we'll introduce convenient coordinate systems via the use of *theta structures*. Elliptic curves will fit into this framework as abelian varieties of dimension 1, and their representation will not be much different from the Montgomery model we have seen in Chapter 1.

General definitions

Let us first review some theory of abelian varieties. We will present some basic facts; for a more detailed treatment, see [Mil86a], [Mum74].

As in the case of elliptic curves, we will work over a finite field $k = \mathbb{F}_q$ of characteristic $p \neq 2, 3$. In the applications, typically we have $p \approx 2^n$ with $n \in \{128, 256, 512\}$.

Definition 2.1. A *group variety* over a field k is an algebraic variety G over k equipped with k-rational maps $m\colon G \times G \to G$, $i\colon G \to G$ and a rational point $e \in G(k)$ such that $G(k^{\mathrm{alg}})$ is a group with respect to the operations induced by the multiplication map m and the inversion map i, with e as the identity element.

Definition 2.2. An *abelian variety* A over a field k is a reduced, connected, projective group variety over k.

Remark 2.3. When working with abelian varieties abstractly, the notion of maps and points *defined over* k (or k-rational) is best formalised in the language of schemes: for example, a k-rational point in $A(k)$ is a morphism $\operatorname{Spec} k \to A$. Once we fix a projective embedding $\iota\colon A \hookrightarrow \mathbb{P}^n$, we can work with the variety as a subvariety of projective space, and the points $P \in A(k)$ defined over k are those whose coordinate representation $\iota(P) = (X_0(P) : \cdots : X_n(P))$ is defined in $\mathbb{P}^n(k)$.

Example 2.4. Let E be an elliptic curve over k. Then E, with the chord-tangent law of Proposition 1.10 as group law, the involution $[-1]$ as inversion and the point at infinity 0_E as neutral point, is an abelian variety of dimension 1.

Abelian varieties are pretty rigid objects: the fact of being projective group varieties imposes a lot of structure. For example:

Proposition 2.5 ([Mil86a, Corollary 2.2] Rigidity lemma for morphisms of abelian varieties)**.** *Let $f\colon A \to B$ be a morphism of abelian varieties over a field k. Then f is the composition of a translation map and a group homomorphism. In particular, any algebraic morphism sending the neutral point e_A to e_B is a group homomorphism.*

An immediate corollary justifies the attribute “abelian” in the name:

Corollary 2.6. *Let A be an abelian variety over a field k. Then the group law on A is commutative.*

Proof. The inversion map $i\colon A \to A$ leaves the neutral point fixed, therefore is a group homomorphism. A group is abelian if and only if inversion is a homomorphism. □

Notation 2.7. We write A/k to say A is an abelian variety over k. Since its group law is commutative, we denote it by $+$, the inversion map by $-$ (or $[-1]$), the neutral point by 0_A. We denote by $\mathcal{K}_A$ the Kummer variety $A/\langle[-1]\rangle$.

Isogenies

Definition and properties of isogenies carry over almost word by word from elliptic curves to general abelian varieties:

Definition 2.8. Let $A, B/k$ be abelian varieties, and $f\colon A \to B$ a nonconstant morphism of algebraic varieties. We say that f is an *isogeny* if it is a surjective group homomorphism with finite kernel. In that case, A is said to be *isogenous* to B.

As an algebraic morphism, f induces a field extension $[k(A) : f^*k(B)]$. We define $\deg f$ to be the degree of this extension, and we say f is separable if the extension is.

Isogenies of degree 1 of abelian varieties are invertible, and are called *isomorphisms.*

Remark 2.9. In the sequel, all the isogenies we work with will be separable. Given an isogeny f, we will often identify $\ker f$ with its geometric points $\ker f(k^{\mathrm{alg}})$.

As with elliptic curves, the simplest example of isogeny is scalar multiplication:

Example 2.10 (Scalar multiplication)**.** Let $n \in \mathbb{Z}$ be a positive integer, A an abelian variety. The multiplication-by-n map

$$[n]\colon A \to A, \qquad P \mapsto \begin{cases} 0_A & \text{if } n = 0, \\ [n-1]P + P & \text{if } n > 0, \\ -[-n]P & \text{if } n < 0, \end{cases}$$

is an isogeny of degree n^{2g}. The kernel of $[n]$ is denoted by $A[n]$, and its geometric points are called n-torsion points.

When $p \nmid n$, the isogeny $[n]$ is separable and its kernel is $A[n] \cong (\mathbb{Z}/n\mathbb{Z})^{2g}$, where g is the dimension of A.

Proof. See [Mil86a, Theorem 8.2, Remark 8.4]. □

The analogue of Lemma 1.30 holds for abelian varieties:

Lemma 2.11. *Let $f\colon A \to C$ be an isogeny of abelian varieties, and $g\colon A \to B$ a separable isogeny with $\ker g \subseteq \ker f$. Then there exists a unique isogeny $h\colon B \to C$ such that $f = h \circ g$.*

In particular, if $K = \ker g = ker f$, then B and C are isomorphic and uniquely determined (up to isomorphism) by K. In this case, we denote the codomain as A/K.

Proof. The proofs in [Sil09, III, Corollary 4.11, Proposition 4.12] apply to any abelian variety. □

Proposition 2.12. *If A, B are isogenous abelian varieties, then they have the same dimension. "Being isogenous" is an equivalence relation.*

Proof. For the first statement, see [Mil86a, Proposition 8.1]. For the second, the identity map is an isogeny, and the composition of isogenies is an isogeny. Symmetry is shown as follows. Let $f\colon A \to B$ be an isogeny. If f is separable, then $n = \deg f$ is the cardinality of $\ker f$ (as a scheme), so $\ker f \subseteq \ker([n])$. By Lemma 2.11, there exists $g\colon B \to A$ such that $[n] = g \circ f$. For the general case, see [Sil09, Theorem III.6.1]. □

2.2 Polarised abelian varieties

In this section, the theory of abelian varieties begins to differ from that of elliptic curves. First, we introduce the notion of *dual abelian variety* of an abelian variety A. Elliptic curves are canonically isomorphic to their duals, whereas this is not the case in higher dimension.

Then, we introduce *polarisations* and *principal polarisations* on abelian varieties. The latter are isomorphisms between a variety and its dual. In higher dimension they need no longer exist nor be unique; however, in order to make many constructions work properly (pairings, projective embeddings, isomorphism invariants) we need to fix the choice of a polarisation and keep track of it. This is why working with abelian varieties algorithmically is slightly more involved than with elliptic curves, and why actually our concept generalising elliptic curves will be that of *polarised* abelian varieties.

Divisors and line bundles

To be able to introduce the dual abelian variety, we first generalise the notions related to divisors on an elliptic curve to the higher-dimensional case. Definitions 1.39, 1.40, 1.42 on elliptic curves carry over to general abelian varieties by replacing points of E by codimension-1 subvarieties of A (defined over k^{alg}).

Definition 2.13 (Divisors, degree of a divisor, Picard group over k^{alg})**.**

$$\begin{aligned}
\operatorname{Div}(A) = \{\textstyle\sum_{i=1}^{m} n_i(W_i) \mid m \in \mathbb{N}, n_i \in \mathbb{Z}, W_i \subseteq A, \dim W_i = \dim A - 1 \text{ for all } i\},\\
\operatorname{div}(f) = \textstyle\sum_{W \text{ codim.-1 subvar.}} \operatorname{ord}_W(f)(W) \quad \text{for } f \in k^{\mathrm{alg}}(A)^*,\\
\operatorname{Pic}(A) = \operatorname{Div}(A)/\sim, \qquad D_1 \sim D_2 \Leftrightarrow D_1 - D_2 = \operatorname{div}(f) \text{ for some } f \in k^{\mathrm{alg}}(A)^*,\\
\deg(D) = \textstyle\sum_i n_i \quad \text{if } D = \sum_i n_i(W_i) \in \operatorname{Div}(A),\\
\operatorname{Div}^0(A) = \{D \in \operatorname{Div}(A) \mid \deg D = 0\},\\
\operatorname{Pic}^0(A) = \{D \in \operatorname{Pic}(A) \mid \deg D = 0\} = \operatorname{Div}^0(A)/\sim
\end{aligned}$$

Definition 2.14. Let A/k be a variety, $D = \sum_i n_i(W_i) \in \operatorname{Div}(A)$ a divisor. For any field extension $K \supseteq k$, we say that D is *defined* over K, and write $D \in \operatorname{Div}_K(A)$, if for all $\sigma \in \operatorname{Gal}(k^{\mathrm{alg}}/K)$, we have $\sigma(D) := \sum_i n_i(\sigma(W_i)) = D$.

A divisor class $[D] \in \operatorname{Pic}(A)$ is defined over K (we write $[D] \in \operatorname{Pic}(A)(K)$) if some divisor in the class is.[1] We use analogue notations for $\operatorname{Div}_K^0(A), \operatorname{Pic}^0(A)(K)$.

To ease our following exposition, it is useful to sometimes switch from the language of divisors to the equivalent language of invertible sheaves on A. Some familiarity with invertible sheaves is assumed, see [Vak24, Chapter 14] for an introduction.

Definition 2.15 (Invertible sheaf associated with a divisor)**.** Let A/k be a smooth variety. Any divisor on A induces an invertible sheaf $\mathcal{L} = \mathcal{L}_D = \mathcal{O}(D)$ whose local sections on a Zariski open $U \subseteq A$ can be identified with the functions $f \in k(U)^*$ such that $\operatorname{div}(f) + D = \sum_i n_i(W_i)$ with $n_i \geq 0$ for all i, and the zero section $f = 0$. Viceversa, any invertible sheaf $\mathcal{L}$ on A is of the form L_D for some D.

By abuse of terminology, we will interchangeably use the terms *invertible sheaf* and *line bundle* for $\mathcal{L}$.

In practice, we will often work with divisors and line bundles in parallel, using the language that is most convenient for the task at hand. The following table summarises a dictionary between the two languages. See [Vak24, §15.4] for more details.

Divisor D	Line bundle $\mathcal{L}_D = \mathcal{O}(D)$
Degree $\deg D$	Degree $\deg \mathcal{L}_D$
Sum $D_1 + D_2$	Tensor product $\mathcal{L}_{D_1+D_2} = \mathcal{L}_{D_1} \otimes \mathcal{L}_{D_2}$
Negation $-D$	Inverse $\mathcal{L}_{-D} = \mathcal{L}_D^{-1}$
Principal divisor $D = \operatorname{div}(f)$	Trivial line bundle $\mathcal{L}_D \cong \mathcal{O} = L_0$
Linear equivalence $D_1 \sim D_2$	Isomorphism $\mathcal{L}_{D_1} \cong \mathcal{L}_{D_2}$

[1] This definition is not entirely standard: if A is a variety over k, we could define $\operatorname{Pic}(A)(k)$ as $\operatorname{Div}_k(A)/k(A)^*$. The two definitions are not equivalent in general, but luckily they are in our case since the varieties we work with are smooth and proper.

Notation 2.16. For a smooth variety A, the set $\mathrm{Pic}_{\mathrm{LB}}(A)$ of line bundles on A up to isomorphism forms a group under tensor product, and is canonically isomorphic to $\mathrm{Pic}(A)$. In the sequel, we'll denote both groups as $\mathrm{Pic}(A)$. Likewise, the subgroup of line bundles of degree 0 is denoted by $\mathrm{Pic}^0(A)$.

Due to our algorithmic motivation, we are interested in representing *explicitly* the varieties we work with using projective coordinates. The following definition helps:

Definition 2.17. A line bundle $\mathcal{L}$ on an abelian variety A is called *ample* if some tensor power $\mathcal{L}^{\otimes n}$ is *very ample*, i.e., there are global sections $s_0, \dots, s_r \in \Gamma(A, \mathcal{L}^{\otimes n})$ such that the morphism

$$[s_0 : s_1 : \cdots : s_r] \colon A \to \mathbb{P}^r$$

is a closed embedding.

Example 2.18. Let E be an elliptic curve, and set $D = (0_E)$. The line bundle $\mathcal{L}_D^3 = \mathcal{L}_{3(0_E)}$ is very ample. Indeed, by Riemann-Roch:

$$\begin{aligned}
\dim \Gamma(\mathcal{L}_D) = 1 \quad &\rightsquigarrow \quad \Gamma(\mathcal{L}_D) = \langle Z_1 \rangle\,, \\
\dim \Gamma(\mathcal{L}_D^2) = 2 \quad &\rightsquigarrow \quad \Gamma(\mathcal{L}_D^2) = \left\langle X_2, Z_2 = Z_1^2 \right\rangle, \\
\dim \Gamma(\mathcal{L}_D^3) = 3 \quad &\rightsquigarrow \quad \Gamma(\mathcal{L}_D^3) = \left\langle X = X_2 Z_1, Y, Z = Z_1^3 \right\rangle
\end{aligned}$$

A word on the notation: here all sections Z_1, Z_2, Z are identified with the constant rational function $1 \in k(E)^*$, but they're technically different objects, being sections on different line bundles. Similarly, X_2 and X are both represented by $x \in k(E)^*$.

Using Riemann-Roch again on $\mathcal{L}_D^6$ gives

$$\dim \Gamma(\mathcal{L}_D^6) = 6 \quad \rightsquigarrow \quad \{X_2^3, X^2, XY, Y^2, XZ, YZ, Z^2\} \text{ are linearly dependent.}$$

This last linear dependence relation is equivalent to the fact that the generators X, Y, Z of $\Gamma(\mathcal{L}_D^3)$ satisfy an equation of the form (1.3), that is, they are the coordinates defining the projective embedding of E in $\mathbb{P}^2$.

If instead we use the generators of $\mathcal{L}_D^2$, we get the projective representation of the Kummer line: $(X_2 : Z_2) = (X : Z)$ induce the embedding $\mathcal{K}_E \xrightarrow{\sim} \mathbb{P}^1$. This is a more general phenomenon that we will see in the next sections.

Remark 2.19. By Lefschetz's theorem [Mum74, Section 17], if $\mathcal{L}$ is ample, then $\mathcal{L}^{\otimes n}$ is very ample when $n \geq 3$.

The dual abelian variety

Definition 2.20. Let A be an abelian variety. Then, the dual abelian variety of A is a k-variety whose K-points are $\widehat{A}(K) = \mathrm{Pic}^0(A)(K)$.

Proposition 2.21 ([Mil86a, §9, 10])**.** *Let A be an abelian variety. Its dual abelian variety $\widehat{A}$ always exists and is unique up to isomorphism. Moreover, it has the same dimension as A, and if $B = \widehat{A}$ then $A = \widehat{B}$.*

In the case of an elliptic curve E, there is an isomorphism $E \to \widehat{E}$ which, seen on geometric points, is the isomorphism $\lambda_{(0_E)}$ of Theorem 1.43.

Definition 2.22 (Dual isogeny)**.** Let $f\colon A \to B$ be a k-isogeny. Then the pull-back

$$f^*\colon \operatorname{Pic}^0(B) \to \operatorname{Pic}^0(A), \quad \mathcal{L} \mapsto f^*\mathcal{L}$$

is an isogeny. It is also denoted by $\widehat{f}\colon \widehat{B} \to \widehat{A}$ and it is called the dual isogeny of f.

Polarisations

Polarisations are isogenies between an abelian variety and its dual that look like (1.4). We will construct them by means of line bundles.

Notation 2.23 (Translation map)**.** Given a point $P \in A$, we denote by $t_P\colon A \to A$ the translation map $Q \mapsto P + Q$. It is an automorphism on A, and is k-rational if $P \in A(k)$.

Theorem 2.24 (Theorem of the square, [Mil86a, Theorem 6.7])**.** *Let $\mathcal{L}$ be an invertible sheaf on A, and P, Q be points of A. Then we have an isomorphism*

$$t^*_{P+Q}\mathcal{L} \otimes \mathcal{L} \cong t^*_P\mathcal{L} \otimes t^*_Q\mathcal{L} \tag{2.1}$$

where t_R is the translation-by-R map.

Remark 2.25. Rephrasing the theorem, if we tensor the above isomorphism by $\mathcal{L}^{-2}$, we get

$$t^*_{P+Q}\mathcal{L} \otimes \mathcal{L}^{-1} \cong (t^*_P\mathcal{L} \otimes \mathcal{L}^{-1}) \otimes (t^*_Q\mathcal{L} \otimes \mathcal{L}^{-1}).$$

In other words, the following map

$$\lambda_{\mathcal{L}}\colon A \to \operatorname{Pic}(A), \quad P \mapsto t^*_P\mathcal{L} \otimes \mathcal{L}^{-1} \tag{2.2}$$

is a group homomorphism. Moreover, $t^*_P(\mathcal{L}) \otimes \mathcal{L}^{-1}$ always has degree $\deg \mathcal{L} - \deg \mathcal{L} = 0$, so the image $\lambda_{\mathcal{L}}(A)$ lies in $\operatorname{Pic}^0(A)$.

Definition 2.26. A k-polarisation on an abelian variety A is an isogeny $\lambda\colon A \to \widehat{A}$ defined over k such that its extension of scalars to k^{alg} is of the form (2.2):

$$\lambda_{k^{\text{alg}}}\colon A_{k^{\text{alg}}} \to \widehat{A}_{k^{\text{alg}}}, \qquad P \mapsto t^*_P\mathcal{L} \otimes \mathcal{L}^{-1} \qquad \text{for some } \mathcal{L} \text{ over } A(k^{\text{alg}}).$$

An abelian variety equipped with a polarisation is called a *polarised abelian variety* (PAV). We will denote it by $(A, \mathcal{L})$, where $\mathcal{L}$ is the line bundle inducing the polarisation.

A polarisation is called *principal* if it is an isomorphism.

Remark 2.27. If A/k is an abelian variety, let $D \in \operatorname{Div}(A)(k)$ be a k-rational divisor. This defines a line bundle $\mathcal{L}_D = \mathcal{O}(D)$ on A. Then the map $\lambda_{\mathcal{L}_D}$ is a polarisation on A. If D has degree 1, it is a principal polarisation.

Definition 2.28. A pair $(A, \mathcal{L})$ where A is an abelian variety and $\lambda_{\mathcal{L}}$ is a principal polarisation on A is called a *principally polarised abelian variety* (PPAV).

Remark 2.29. First note that, if $D = \sum_i n_i(P_i)$ is a divisor, then $t_P^* D = \sum_i n_i(P_i - P)$. In particular, $t_P^*((0_E)) = (-P)$. The translation map brings a minus sign on divisors. Also, if D induces $\mathcal{L} = \mathcal{L}_D$, then $t_P^* \mathcal{L}$ is induced by $t_P^* D$.

Consider the map $\lambda_{(0_E)} \colon P \mapsto [(P) - (0)]$ defined in Theorem 1.43, and let $D = (0_E)$. The map $-\lambda_{(0_E)} \colon P \mapsto [(0_E) - (P)] = [(-P) - (0_E)] = t_P^* D - D$ is a principal polarisation induced by $\mathcal{L}_D$. The map $\lambda_{(0_E)}$ is not actually a polarisation: it would be induced by $\mathcal{L}_{-(0_E)}$, because of the sign issue above, but this invertible sheaf is not ample. Indeed, we have

$$\lambda_{\mathcal{L}_{-(0_E)}}(P) = t_P^* \mathcal{L}_{-(0_E)} \otimes \mathcal{L}_{-(0_E)}^{-1} = \mathcal{L}_{-(-P)+(0_E)} \cong \mathcal{L}_{(P)-(0_E)} = \lambda_{(0_E)}(P).$$

Definition 2.30. Let $(A, \mathcal{L}), (B, \mathcal{M})$ be PAVs and let $f \colon A \to B$ be an isogeny. We say that it respects the polarisations (or it is an isogeny of PAVs) if

$$f^* \mathcal{M} = \mathcal{L}.$$

In terms of the polarisations, this implies $f^* \lambda_{\mathcal{M}} := \widehat{f} \circ \lambda_{\mathcal{M}} \circ f = \lambda_{\mathcal{L}}$.

Remark 2.31. The reason we introduce the notion of *isogeny of PAVs* refining the notion of *isogeny* is the following. As we saw in Example 2.18, fixing a line bundle (equivalently, fixing a polarisation) gives a way to embed an abelian variety in projective space (if the line bundle is ample). In the algorithms, our varieties will always be represented via a projective embedding – since their points are written in coordinates on a computer – that is, they're equipped with a fixed polarisation. The isogenies we want to study are the ones that respect this representation, hence the polarisation.

Polarisations are also used to define pairings on an abelian variety:

Theorem 2.32 (Weil pairing). *Let A be an abelian variety and $\ell \in \mathbb{Z}$ coprime with p. There is a skew-symmetric, nondegenerate bilinear map*

$$e_m \colon A[\ell] \times \widehat{A}[\ell] \to \mu_\ell,$$

with $\mu_\ell = \{\ell\text{-th roots of unity}\}$, called the Weil pairing. *If $(A, \mathcal{L})$ is principally polarised, then we can identify A with $\widehat{A}$, obtaining a pairing on $A[\ell] \times A[\ell]$ that satisfies the properties in Definition 1.44.*

2.3 Examples

We see several examples of principally polarised abelian varieties, starting from the simplest case, elliptic curves. Where possible, since abelian varieties are projective, we see how to embed them in projective space.

Example 2.33. As we mentioned earlier, the abelian varieties of dimension 1 are exactly elliptic curves, and they're canonically principally polarised by the negative $-\lambda_{(0_E)}$ of the isomorphism of Equation (1.4).

We saw that on a given elliptic curve we can fix a projective embedding in $\mathbb{P}^2(k^{\text{alg}})$ as the projective closure of the affine curve $y^2 = f(x)$ with f a cubic polynomial having distinct roots. Polynomials of the form $f(x) = x^3 + ax + b$, $f(x) = x(x^2 + Ax + 1)$ give the short Weierstrass model, Montgomery model respectively. The group law is given by the chord-tangent law of Figure 1.1 and the neutral point is 0_E.

Moreover, the projection

$$\begin{array}{rccc} \mathbf{x}\colon & E \subseteq \mathbb{P}^2 & \to & \mathbb{P}^1 \\ & (X : Y : Z) & \mapsto & (X : Z) \end{array}$$

induces an isomorphism of the Kummer variety $\mathcal{K}_E = E/\langle \pm 1 \rangle$ to $\mathbb{P}^1$.

Example 2.34 (Products of elliptic curves)**.** Let $E_1, \ldots, E_g$ be elliptic curves. Then the product $E_1 \times \cdots \times E_g$ is an abelian variety of dimension g. The product line bundle

$$\mathcal{L} = \pi_1^* \mathcal{L}_{(0_{E_1})} \otimes \cdots \otimes \pi_g^* \mathcal{L}_{(0_{E_g})},$$

where π_i is the projection on the i-th factor, induces a principal polarisation on the product abelian variety. This is the product polarisation

$$\lambda = \lambda_{(0_{E_1})} \times \cdots \times \lambda_{(0_{E_g})} : E_1 \times \cdots \times E_g \to \widehat{E}_1 \times \cdots \times \widehat{E}_g.$$

A product of elliptic curves can be embedded in a projective space: embed the single components into $\mathbb{P}^2$, then form a product of projective planes into a single projective space (of high dimension) via a Segre embedding.

This also happens at the level of Kummer varieties. For simplicity, set $g = 2$ and let E_1, E_2 be elliptic curves. Then we have an embedding

$$\begin{array}{rcc} E_1/\langle -1 \rangle \times E_2/\langle -1 \rangle & \hookrightarrow & \mathbb{P}^3 \\ ((X_1 : Z_1), (X_2 : Z_2)) & \mapsto & (X_1X_2 : X_1Z_2 : Z_1X_2 : Z_1Z_2). \end{array}$$

Elliptic curve products are abelian variety of dimension g, but we will see that they're relatively rare among higher-dimensional abelian varieties. The following construction gives another class of abelian varieties, called *Jacobian varieties.* For more details on this, see [Mil86b].

Definition 2.35. Let C be a smooth plane projective curve of genus g defined as the unique smooth projective curve birational to the the projective closure of the affine curve

$$y^2 = f(x), \qquad f \in k[x], \quad \deg f = 2g + 1.$$

The curve C is said to be a *hyperelliptic curve.* Note that C comes equipped with a canonical involution

$$\begin{array}{rccc} \iota\colon & C & \to & C \\ & (x, y) & \mapsto & (x, -y). \end{array}$$

and a rational point $P_0 = (0 : 1 : 0)$.

There is a natural construction that attaches to a hyperelliptic curve C of genus g an abelian variety of dimension g, called the *Jacobian* of C.

Proposition 2.36 ([Mil86b, Theorem 1.1]). *Let C/k be a hyperelliptic curve given by a polynomial of degree $2g+1$. There exists a g-dimensional abelian variety, denoted by Jac_C and called the Jacobian of C, whose K-points are*

$$\mathrm{Jac}_C(K) \cong \mathrm{Pic}^0(C)(K).$$

(see Definition 2.14 for the meaning of K-rationality for divisors.)

Remark 2.37. Elliptic curves are the Jacobian varieties of themselves.

The abelian group structure on a Jacobian descends from that of the divisor group, whereas the structure of algebraic variety comes from the following:

Fact 2.38 ([Mum84, Section 2]). *For any $D \in \mathrm{Div}^0(C)$ there are points $P_1, \dots, P_r$ (unique up to reordering) with $r \le g$, $P_i \neq P_0$ and $P_i \neq \iota(P_j)$ if $i \neq j$, such that*

$$D \sim (P_1) + \dots + (P_r) - r(P_0).$$

This means that the divisor D is determined by the g-tuple of points $(P_1, \dots, P_r)$ up to permutations of the coordinates. Setting $P_{r+1} = \dots = P_g = P_0$ if $r < g$ gives a map

$$\begin{array}{rccc} \varphi_g\colon & C^g & \to & \mathrm{Jac}_C \\ & (P_1, \dots, P_g) & \mapsto & [(P_1) + \dots + (P_g) - r(P_0)]. \end{array}$$

This map is surjective, and realises Jac_C as a finite quotient of the variety C^r.

So concretely, points on a Jacobian are representable as unordered tuples of points in C. For example, when g is 2, if we have $D_1 = \varphi_2(P_1, P_2)$ and $D_2 = \varphi_2(Q_1, Q_2)$, name R_1, R_2 the two other intersection points of the cubic through P_1, P_2, Q_1, Q_2 with the curve C. Then $[D_1 + D_2] = \varphi_2(\iota(R_1), \iota(R_2))$ (see Figure 2.1). This is a higher-dimensional analogue of the chord-tangent law for elliptic curves.

Proposition 2.39 ([Mil86b, Theorem 6.6]). *Let Jac_C be the Jacobian of a genus-g hyperelliptic curve C, and define a map $\varphi_{g-1}\colon C^{g-1} \to \mathrm{Jac}_C$ as in Equation (2.38). The image $\varphi_{g-1}(C^{g-1})$ is a hypersurface $W \subseteq \mathrm{Jac}_C$, defining a divisor $\Theta = (W)$ called the* theta divisor *of C. This divisor defines a principal polarisation on Jac_C.*

Remark 2.40. When we work with a Jacobian, we always fix the principal polarisation induced by the theta divisor, being a canonical choice.

As we did with elliptic curves, we would like to represent Jacobians of genus-g hyperelliptic curves over k via some coordinate system in a projective space. In general, the representation is going to be messy: it can be shown [CF96] that a dimension-2 Jacobian can be embedded in $\mathbb{P}^{15}(k)$ as the zero locus of 72 quadratic polynomials, which is not convenient as a representation for cryptographic purposes.

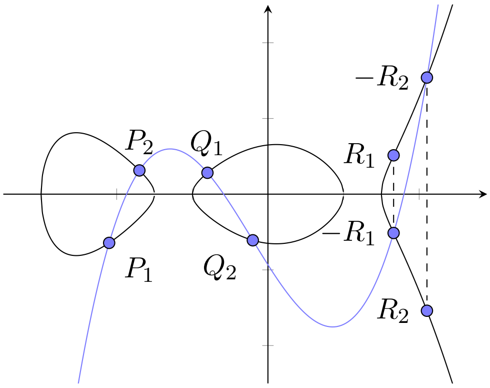

Figure 2.1: Group law on a hyperelliptic 2-dimensional jacobian (over $\mathbb{R}$)

However, we can instead embed the Kummer variety of a genus-2 curve in $\mathbb{P}^3$, as the zero locus of a single quartic polynomial, which is much more practical. To see how this is possible, we will present in the next section Mumford's formalism of theta functions, which will allow us to derive both efficient representations of abelian varieties and their Kummer varieties, and efficient algorithms for their arithmetic.

In dimension 2, the examples of this section are all the PPAVs we can encounter:

Proposition 2.41. *Let A be a 2-dimensional PPAV. Then it is isomorphic (over k^{alg}) to either the Jacobian of a curve or a product of two elliptic curves.*

Proof. See [BL04, Corollary 11.8.2]. □

2.4 Theta groups

In this section, we start to introduce Mumford's theory of algebraic theta functions, coming from [Mum66]. Theta functions are a classical tool used to study the geometry of complex abelian varieties. Mumford, that we'll follow here, gives instead a completely algebraic treatment of the theory of theta functions, that works over any field, and in particular for k of positive characteristic.

Once we set this framework, we'll be able to define *theta structures*, constructions giving convenient coordinate systems to represent our varieties.

Definition 2.42. Let $\mathcal{L}$ be an ample invertible sheaf on an abelian variety A/k. Denote by $H(\mathcal{L})$ the kernel of the polarisation $\lambda_{\mathcal{L}}$. Its geometric points $H(\mathcal{L})(k^{\text{alg}})$ are the points $P \in A$ such that $\lambda_{\mathcal{L}}(P)$ is trivial. More generally,

$$H(\mathcal{L})(F) = \{P \in A(F) \mid t_P^*\mathcal{L}_F \cong \mathcal{L}_F\} \quad \text{if } k \subseteq F \subseteq k^{\text{alg}}.$$

where $\mathcal{L}_F$ is the line bundle $\mathcal{L} \otimes_k F$ induced by the change of basis $A_F = A \times \operatorname{Spec} F$.

From now on, we assume that $\mathcal{L}$ is defined by a divisor of degree $\deg(\mathcal{L}) = d > 0$ with $p \nmid d$. Such $\mathcal{L}$ is said to be *of separable type*.

Proposition 2.43. *If $\mathcal{L}$ is an ample invertible sheaf of separable type, then $H(\mathcal{L})$ is finite and its cardinality is $\#H(\mathcal{L})(k^{\mathrm{alg}}) = \deg(\mathcal{L})^2$.*

Proof. Proved in [MFK94, Section 6.2]. □

Definition 2.44. Let $\mathcal{L}$ be an invertible sheaf of separable type on A/k. If we write

$$G(\mathcal{L})(F) = \{(P,\varphi) \mid P \in A(F), \varphi\colon \mathcal{L}_F \xrightarrow{\sim} t_P^*\mathcal{L}_F\} \qquad \text{for } k \subseteq F \subseteq k^{\mathrm{alg}},$$

then $G(\mathcal{L})$ is an algebraic group, called the *theta group* of $\mathcal{L}$, with the composition law

$$(Q,\psi) \circ (P,\varphi) = (P+Q, t_P^*\psi \circ \varphi).$$

Notation 2.45. Since we deal with sheaves of separable type, we will usually identify our algebraic groups with their geometric points, writing $H(\mathcal{L}) = H(\mathcal{L})(k^{\mathrm{alg}})$ and $G(\mathcal{L}) = G(\mathcal{L})(k^{\mathrm{alg}})$.

Notation 2.46. Let $P \in A$ be a point. For a divisor D, denote $D_P = t_P^*D - D = (t_{-P})_*D - D$. For example, if $D = (0_A)$, then $D_P = (-P) - (0_A)$.

Remark 2.47. The definition of a theta group can be rephrased in the language of divisors. If $\mathcal{L} = \mathcal{L}_D$, then $H(D) = H(\mathcal{L})$ is the group of points P such that the divisor D_P is principal, and the elements of $G(D) = G(\mathcal{L})$ are pairs (P, g_P), where g_P is a function with divisor D_P. The pair (P, g_P) is k-rational if both P and g_P are.

Proposition 2.48. *The theta group $G(\mathcal{L})$ fits into a short exact sequence of algebraic groups:*

$$1 \to \mathbb{G}_m \to G(\mathcal{L}) \to H(\mathcal{L}) \to 0.$$

Here $G(\mathcal{L}) \to H(\mathcal{L}), (P,\varphi) \mapsto P$ is the natural projection, whose kernel consists of automorphisms of $\mathcal{L}$, that is, multiplication by nonzero constants in $\mathbb{G}_m$.

Proposition 2.49. *The theta group $G(\mathcal{L})$ acts on the vector space of global sections $\Gamma(A,\mathcal{L})$, making $\Gamma(A,\mathcal{L})$ an irreducible $G(\mathcal{L})$-module, via*

$$(P,\varphi) \star s = t_{-P}^*(\varphi(s)).$$

The scalars $\mathbb{G}_m$ in $G(\mathcal{L})$ act on $\Gamma(A,\mathcal{L})$ as multiplication characters.

In the language of divisors, if g_P is above P and has $D_P = t_P^*D - D$ as divisor and s is a global section of $\mathcal{L}_D$, then $g_P \star s = (s/g_P)(\cdot - P)$, that is again a section of $\mathcal{L}_D$.

Proof. This is well-defined: if s is a global section of $\mathcal{L}$, then $\varphi(s)$ is a global section of $t_P^*\mathcal{L}$, and $t_{-P}^*(\varphi(s))$ is a global section of $\mathcal{L}$ again. It is an action because

$$(Q,\psi) \circ (P,\varphi) \star s = t_{-Q-P}^*(t_P^*\psi \circ \varphi(s)) = t_{-Q}^*\psi \circ t_{-P}^*(\varphi(s)) = (Q,\psi) \star ((P,\varphi) \star s).$$

Irreducibility of the representation is shown in [Mum66, §1, Theorem 2]. □

Example 2.50. As the simplest example, consider an elliptic curve E/k. Let $D = (0_E)$ be the degree-1 divisor defining the principal polarisation. Then $H(D)$ is trivial: indeed, $D_P = (-P) - (0_E)$ is never principal when $P \neq 0_E$. The theta group $G(D) \cong \mathbb{G}_m$ consists of just the constants.

Example 2.51. Consider again an elliptic curve E/k, say it's in Montgomery form, just to fix some projective coordinates $(X : Y : Z)$. Consider now the degree-2 divisor $D = 2(0_E)$. The associated line bundle $\mathcal{L} = \mathcal{L}_D$ has global sections generated by the x-line coordinates X, Z.

Now $H(D) = E[2]$. Indeed, $D_P - D = 2(-P) - 2(0)$ is principal if and only if $2P = 0$, that is, $P \in E[2]$.

A theta-group element over a point T of order 2 is of the form $g_T = \lambda \cdot (x - x(T))$ for some $\lambda \in k^*$, with $x = X/Z$ (see example 1.41). If we take T to be the order-2 point $(0 : 1)$ always present on a Montgomery x-line (see Remark 1.16), then we have $g_T = \lambda X/Z$.

Moreover, by Proposition 1.20 we have $(t_T^* X, t_T^* Z) = (Z, X)$ up to a constant.

The function g_T defines an isomorphism between $\mathcal{L}$ and $t_T^*\mathcal{L}$ inducing an action on the global sections $\Gamma(A, \mathcal{L}) = \langle X, Z \rangle$. Indeed we have

$$\begin{aligned}(T, g_T) \star X &= t_T^*(X/g_T) = \lambda^{-1} t_T^* Z = \lambda^{-1} X,\\ (T, g_T) \star Z &= t_T^*(Z/g_T) = \lambda^{-1}(X/Z) \cdot X.\end{aligned}$$

The action on sections in turn descends to an action on the Kummer line:

$$(T, g_T) \star \mathbf{x}(P) = ((T, g_T) \star X(P) : (T, g_T) \star Z(P)) = (Z : X)(P) = \mathbf{x}(P + T).$$

The action of g_T on a basis $\mathcal{B}$ of global sections of an ample line bundle $\mathcal{L}$ induces the translation-by-T map on the projective coordinates given by $\mathcal{B}$.

More generally, if $T = (X_T : Z_T)$ is any 2-torsion point, Proposition 1.20 gives

$$((T, g_T) \star X, (T, g_T) \star Z) = (\lambda(X_T X + Z_T Z), \lambda(Z_T X + X_T Z))$$

for some common scalar λ.

We point out another piece of structure induced by the theta group on rational torsion subgroups of an abelian variety.

Proposition 2.52 (Commutator pairing)**.** *The surjection $G(\mathcal{L}) \to H(\mathcal{L})$ induces a pairing $e_{\mathcal{L}}$ on $H(\mathcal{L})$, called the* commutator pairing*, that is skew-symmetric, nondegenerate, and has values in $\mathbb{G}_m$.*

The commutator pairing gives $H(\mathcal{L})$ a structure of symplectic group: $H(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$, one dual to the other via the pairing:

$$\begin{array}{ccc} K_2(\mathcal{L}) & \overset{\sim}{\longrightarrow} & \mathrm{Hom}(K_1(\mathcal{L}), \mathbb{G}_m) = \widehat{K_1(\mathcal{L})} \\ Q & \mapsto & e_{\mathcal{L}}(\cdot, Q) \end{array}$$

Proof. The pairing is constructed as follows. For $P, Q \in H(\mathcal{L})$, let g_P, g_Q be respective lifts in $G(\mathcal{L})$. These lifts can be chosen with a degree of freedom, that is, up to a scalar in $\mathbb{G}_m$. We can see then that the quantity

$$e_{\mathcal{L}}(P,Q) = g_P g_Q g_P^{-1} g_Q^{-1}$$

is independent of the choice of P, Q, and defines a skew-symmetric pairing $e_{\mathcal{L}} \colon H(\mathcal{L}) \times H(\mathcal{L}) \to \mathbb{G}_m$. Nondegeneracy is shown in [Mum66, §1, Theorem 1]. □

Lemma 2.53. *Let $\mathcal{L}$ be an ample separable line bundle, $n \in \mathbb{Z}$ with $\operatorname{char}(k) = p \nmid n$. Then $H(\mathcal{L}^n) = [n]^{-1}H(\mathcal{L})$ and $H(\mathcal{L}) = [n]H(\mathcal{L}^n)$.*

Proof. By the Theorem of the Square 2.24, we have

$$\lambda_{\mathcal{L}^n}(P) = t_P^*\mathcal{L}^n \otimes \mathcal{L}^{-n} = (t_P^*\mathcal{L} \otimes \mathcal{L}^{-1})^n \cong t_{[n]P}^*\mathcal{L} \otimes \mathcal{L}^{-1} = \lambda_{\mathcal{L}}([n]P).$$

It follows $H(\mathcal{L}^n) = \ker \lambda_{\mathcal{L}^n} = \ker \lambda_{\mathcal{L}} \circ [n] = [n]^{-1}H(\mathcal{L})$. In particular, $[n]H(\mathcal{L}^n) \subseteq H(\mathcal{L})$, and the converse inclusion is true because A is n-divisible (since $p \nmid n$). □

Example 2.54. Let $(A, \mathcal{L}_0)$ be a principally polarised abelian variety. Since $\lambda_{\mathcal{L}_0}$ is an isomorphism, $H(\mathcal{L}_0)$ is trivial. Then by the above lemma $H(\mathcal{L}_0^n) = A[n] \cong (\mathbb{Z}/n\mathbb{Z})^g \oplus (\mathbb{Z}/n\mathbb{Z})^g$. One can show that the commutator pairing on $\mathcal{L}_0^n$ coincides with the Weil pairing on $A[n]$.

2.5 Theta structures

We are now ready to define theta structures on PAVs $(A, \mathcal{L})$. They are a way to find a convenient coordinate system for A where the torsion points $H(\mathcal{L})$ are in a special position with lots of symmetries, and the theta group acts on these coordinates in a simple way. Later, once we bring isogenies and group arithmetic into the picture, we will see how using theta structures allows for efficient algorithms.

We are going to define an abstract analogue of $G(\mathcal{L})$, called the Heisenberg group. It is an object that summarises the group-theoretic properties of the theta group $G(\mathcal{L})$.

Definition 2.55. Let $\delta = (d_1, \dots, d_r)$ be a sequence of positive integers, such that $d_{i+1} \mid d_i, d_r > 1$.

Let $K_1(\delta) = \bigoplus_{i=1}^r (\mathbb{Z}/d_i\mathbb{Z})$ and fix a scalar product on $K_1(\delta)$ by

$$\langle (a_1, \dots, a_r) | (b_1, \dots, b_r) \rangle = \prod\nolimits_{i=1}^r \zeta_i^{a_i b_i} \mod d_i, \qquad \zeta_i \text{ a } d_i\text{-th root of unity.}$$

Define $K_2(\delta) = K_1(\delta)$, isomorphic to $\widehat{K_1(\delta)} = \operatorname{Hom}(K_1(\delta), k^{\mathrm{alg}})$ via the pairing $\langle \cdot, \cdot \rangle$. Finally let $H(\delta) = K_1(\delta) \oplus K_2(\delta)$. We have an analogue of the commutator pairing on $H(\delta)$, defined by

$$e_\delta((t_1, t_2), (t_1', t_2')) := \langle t_1 | t_2' \rangle \langle t_1' | t_2 \rangle^{-1}.$$

Let $G(\delta)$ be the group $(k^{\mathrm{alg}})^* \times K_1(\delta) \times K_2(\delta)$ equipped with the group law

$$(\lambda, t_1, t_2) \cdot (\lambda', t_1', t_2') = (\lambda\lambda' \langle t_1 | t_2' \rangle, t_1 + t_2, t_1' + t_2').$$

The group $G(\delta)$ is called the *Heisenberg group*.

Definition 2.56. Let $\mathcal{L}$ be an ample line bundle. We know there is some symplectic decomposition $H(\mathcal{L}) \cong K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$ with $K_1(\mathcal{L}) \cong K_2(\mathcal{L}) \cong \bigoplus_{i=1}^r (\mathbb{Z}/d_i\mathbb{Z})$, $d_{i+1}|d_i$, $d_r > 1$. We say that $\mathcal{L}$ is of level $\delta = (d_1, \dots, d_r)$.

Observe that the construction of the Heisenberg group mimicked the one of the theta group. Then given $\mathcal{L}$, once we choose a symplectic decomposition of $H(\mathcal{L})$, we have a commutative diagram

$$\begin{tikzcd}
1 \arrow[r] & (k^{\mathrm{alg}})^* \arrow[r] \arrow[d, equal] & G(\delta) \arrow[r] \arrow[d, "\Theta_{\mathcal{L}}"] & H(\delta) \arrow[r] \arrow[d, "\overline{\Theta}_{\mathcal{L}}"] & 0 \\
1 \arrow[r] & (k^{\mathrm{alg}})^* \arrow[r] & G(\mathcal{L}) \arrow[r] & H(\mathcal{L}) \arrow[r] & 0
\end{tikzcd}$$

where $\overline{\Theta}_{\mathcal{L}}$ is a symplectic isomorphism, meaning that $\overline{\Theta}_{\mathcal{L}}(K_i(\delta)) = K_i(\mathcal{L})$ and

$$e_\delta((t_1, t_2), (t_1', t_2')) = e_{\mathcal{L}}(\overline{\Theta}_{\mathcal{L}}(t_1, t_2), \overline{\Theta}_{\mathcal{L}}(t_1', t_2')).$$

Definition 2.57. Let $\Theta_{\mathcal{L}}$ be the central isomorphism $G(\mathcal{L}) \xrightarrow{\sim} G(\delta)$ in the above diagram. This is called a *theta structure* of level δ on A.

Remark 2.58. Given a polarised abelian variety $(A, \mathcal{L})$, different theta structures can be associated with it. Two level-δ structures $\Theta_{\mathcal{L}}, \Theta_{\mathcal{L}}'$ on A differ by an automorphism of the Heisenberg group. In practice, fixing a theta structure amounts to fixing a numbering of the d_i-torsion points of A (hence a symplectic decomposition) and, for each of these points $P = \overline{\Theta}_{\mathcal{L}}(i, j)$, a "privileged" lift $g_P = \Theta_{\mathcal{L}}(1, i, j)$ in $G(\mathcal{L})$.

Proposition 2.59 ([Mum66, §1, Proposition 3])**.** *The group $G(\delta)$ has a unique irreducible representation where scalars in $\mathbb{G}_m$ act by their multiplication character. This representation is $V(\delta) = \{g \colon K(\delta) \to k^{\mathrm{alg}}\}$ where the action is given by*

$$((\lambda, t_1, t_2) \star g)(i) = \lambda \langle i | t_2 \rangle g(i + t_1). \tag{2.3}$$

The vector space $V(\delta)$ admits a canonical basis given by the Kronecker delta functions $(\delta_i)_{i \in K(\delta)}$ on $K(\delta)$.

Both $G(\mathcal{L})$ and $G(\delta)$ have unique irreducible representations where $(k^{\mathrm{alg}})^*$ acts as multiplication. The theta structure $G(\delta) \xrightarrow{\sim} G(\mathcal{L})$ induces an isomorphism of these representations $\beta \colon V(\delta) \to \Gamma(A, \mathcal{L})$. In practice, the canonical basis of $V(\delta)$ induces a basis of $\Gamma(A, \mathcal{L})$ made of global sections of $\mathcal{L}$.

Definition 2.60. Let $(A, \mathcal{L})$ be a PAV with a theta structure $\Theta_{\mathcal{L}}$ of level δ. For $i \in K_1(\delta)$, let $\theta_i = \beta(\delta_i)$. We call θ_i an (algebraic) *theta function* on $\mathcal{L}$. The functions $(\theta_i)_{i \in K_1(\delta)}$ form a k^{alg}-basis of $\Gamma(A_{k^{\mathrm{alg}}}, \mathcal{L})$.

Pushing the action (2.3) via the isomorphism β to $\Gamma(A, \mathcal{L})$, we get the action of $G(\mathcal{L})$ on the theta functions:

$$\Theta(\lambda, t_1, t_2) \star \theta_i = \lambda \, \langle i | t_2 \rangle \, \theta_{i+t_1}.$$

Projective embeddings

Notation 2.61. If $\mathcal{L}$ is very ample and $\Theta_{\mathcal{L}}$ is a theta structure, then the theta functions determine an embedding of A in projective space. We'll denote it by

$$\Phi_{\Theta_{\mathcal{L}}} \colon A \hookrightarrow \mathbb{P}^{d_1 \cdots d_r - 1}, \qquad P \mapsto (\theta_i(P))_{i \in K(\delta)}. \tag{2.4}$$

Notation 2.62. In our cases of interest, we usually deal with PAVs of the form $(A, \mathcal{L} = \mathcal{L}_0^n)$ where $\mathcal{L}_0$ is principal and n is a small integer. In this case, we have $H(\mathcal{L}) = A[n]$, $K_1(\delta) = (\mathbb{Z}/n\mathbb{Z})^g$, and a theta structure on $\mathcal{L}$ is simply said to be of *level* n.

If $\mathcal{L}_0$ gives a principal polarisation, via the line bundle $\mathcal{L}_0^n$ we can describe A using n^g coordinates, when $n \geq 3$ (see Remark 2.19). To reduce even further the number of coordinates (in view of efficiency), we'd like to understand what happens when $n = 2$. Unfortunately, $\mathcal{L} = \mathcal{L}_0^2$ of level 2 is not very ample. However, under certain *symmetry* conditions, it lets us describe coordinates on the Kummer variety $\mathcal{K}_A = A / \langle -1 \rangle$. This means, instead of embedding the whole A in $\mathbb{P}^N$, we can represent points of A *up to sign* much more efficiently.

Definition 2.63. We say that a line bundle $\mathcal{L}$ on A is symmetric if $[-1]^*\mathcal{L} = \mathcal{L}$, and totally symmetric if there is a line bundle $\mathcal{M}$ on the Kummer variety $\mathcal{K}_A$ such that $\mathcal{L} = \pi^*\mathcal{M}$, with $\pi \colon A \to \mathcal{K}_A$ the canonical projection.

Example 2.64. Let E be an elliptic curve, let $\mathcal{L} = \mathcal{L}_{2(0_E)}$ be the line bundle whose global sections are generated the Weierstrass coordinates X, Z. Then $\mathcal{L}$ is totally symmetric, since X, Z generate the space of global sections on $\mathcal{K}_E$.

Lemma 2.65 ([Mum66, page 308])**.** *Let $\mathcal{L}$ be totally symmetric. There is an automorphism δ_{-1} of $G(\mathcal{L})$ making the following diagram commute:*

$$\begin{tikzcd}
1 \arrow[r] & (k^{\mathrm{alg}})^* \arrow[r] \arrow[d, equal] & G(\mathcal{L}) \arrow[r] \arrow[d, "\delta_{-1}"] & H(\mathcal{L}) \arrow[r] \arrow[d, "{[-1]}"] & 0 \\
1 \arrow[r] & (k^{\mathrm{alg}})^* \arrow[r] & G(\mathcal{L}) \arrow[r] & H(\mathcal{L}) \arrow[r] & 0
\end{tikzcd}$$

and an analogue D_{-1} on $G(\delta)$.

Definition 2.66. Let $\mathcal{L}$ be totally symmetric. A theta structure $\Theta_{\mathcal{L}}$ is called *symmetric* if $\Theta_{\mathcal{L}} \circ D_{-1} = \delta_{-1} \circ \Theta_{\mathcal{L}}$.

Finally, we have the necessary tools to embed Kummer varieties:

Proposition 2.67. *If* $(A, \mathcal{L}_0)$ *is a g-dimensional* geometrically simple *PPAV,* $\mathcal{L} = \mathcal{L}_0^2$ *is totally symmetric and* $\Theta_{\mathcal{L}}$ *is a symmetric theta structure of level* 2*, then the theta functions determine an embedding*

$$\Phi_{\Theta_{\mathcal{L}}} \colon \mathcal{K}_A \hookrightarrow \mathbb{P}^{2^g-1}.$$

Remark 2.68. If instead A is not simple, this does not necessarily hold. For example, if $A = E_1 \times E_2$, and we fix a level-2 theta structure on each E_i giving theta functions $(\theta_0^{E_i}, \theta_1^{E_i})$, then the Segre-like map

$$\begin{aligned}(P, Q) \mapsto (&\theta_0^{E_1}(P)\theta_0^{E_2}(Q), \theta_0^{E_1}(P)\theta_1^{E_2}(Q),\\ &\theta_1^{E_1}(P)\theta_0^{E_2}(Q), \theta_1^{E_1}(P)\theta_1^{E_2}(Q))\end{aligned}$$

is not an embedding of $A/\langle -1\rangle$ into $\mathbb{P}^3$, but of $E_1/\langle -1\rangle \times E_2/\langle -1\rangle$. We saw a similar phenomenon with Montgomery x-lines, see the end of Example 2.34. These coordinates are indeed theta functions induced by the *product theta structure*, see Lemma 3.18.

Theta-null point and torsion points

Definition 2.69. Let $\Phi \colon A \to \mathbb{P}^{n^g-1}$ be the canonical projective embedding determined by a theta structure $\Theta_{\mathcal{L}}$ of level n on A. The theta null point of A is the image of the neutral point, that is, $\Phi(0_A)$.

A good feature of theta structures is that they put $H(\mathcal{L})$ in a special position, such that once we know the theta null point, we immediately know the coordinates of all $T \in H(\mathcal{L})$.

Definition 2.70. Let $P \in A(k)$ be a rational point. The fiber $\mathcal{L}(P)$ of $\mathcal{L}$ at P is isomorphic to k, but not canonically. Fix $\Lambda_P \colon L(P) \xrightarrow{\sim} k$ an isomorphism. This is called a *rigidification* of $\mathcal{L}$ at P.

Remark 2.71. Given a section $s \in \Gamma(A, \mathcal{L})$, a rigidification Λ_P determines a value $\overline{s}(P) = \Lambda_P(s(P)) \in k$. In particular, if $\Theta_{\mathcal{L}}$ is a level-n theta structure, then Λ_P induces well-defined *affine* coordinates for P as $\overline{P} := (\overline{\theta}_i(P))_{i \in K_1(\delta)} \in k^{n^g}$. Conversely, an affine lift $\overline{P}$ of $\Phi_{\Theta_{\mathcal{L}}}(P)$ determines a rigidification Λ_P.

Proposition 2.72. *Let* $(A, \Theta_{\mathcal{L}})$ *be a PAV equipped with a theta structure of level* n*. An affine lift* $\overline{0_A}$ *of the theta null point uniquely determines affine coordinates* $\overline{T}$ *for all* $T \in H(\mathcal{L})$*, where*

$$\overline{\theta_i}(T) = \langle i|t_2\rangle^{-1}\,\overline{\theta_{i-t_1}}(0_A) \qquad \text{if } T = \overline{\Theta}_{\mathcal{L}}(t_1, t_2).$$

Proof. Fix $(T, \varphi_T) = \Theta_{\mathcal{L}}(1, t_1, t_2) \in G(\mathcal{L})$ a theta group element above T. Let Λ_0 the rigidification induced by $\overline{0_A}$. Then $\Lambda_T = \Lambda_0 \circ \varphi_T^{-1}(0_A) \colon \mathcal{L}(T) \to k$ is a well-defined rigidification of $\mathcal{L}$ at T. In coordinates, we have

$$\overline{\theta_i}(T) = \Lambda_T(\theta_i) = \Lambda_0(\varphi_T^{-1}(\theta_i)) = \Lambda_0(\Theta_{\mathcal{L}}(1, -t_1, -t_2) \star \theta_i)(0) = \langle i|t_2\rangle\,\overline{\theta_{i-t_1}}(0_A). \quad \square$$

Remark 2.73. More generally, if we consider a point $P \in A(k)$ and fix a rigidification at P, by the same reasoning we can translate it by a point in $H(\mathcal{L})$:

$$\overline{\theta_i}(P+T) = \langle i|t_2\rangle\, \overline{\theta_{i-t_1}}(P), \qquad T = \overline{\Theta_{\mathcal{L}}}(t_1, t_2).$$

2.6 More examples

Let's first see what theta structures of different levels in dimension 1 look like.

Example 2.74. Let E be an elliptic curve. The level-2 line bundle $\mathcal{L} = \mathcal{L}_{2(0_E)}$ is totally symmetric. Any symmetric $\Theta_{\mathcal{L}}$ gives an isomorphism

$$\begin{array}{ccc} E/\langle -1\rangle & \xrightarrow{\sim} & \mathbb{P}^1 \\ P & \mapsto & (\theta_0(P) : \theta_1(P)) \end{array}$$

Let $(a:b)$ be the theta null point. Then the two torsion is

$$E[2] = E[2]/\langle -1\rangle = \{0_E = (a:b), T_1 = (b:a), T_2 = (a:-b), T_1+T_2 = (-b:a)\}$$

with $K_1 = \langle T_1\rangle$, $K_2 = \langle T_2\rangle$ factors of a symplectic decomposition of $H(\mathcal{L}) = E[2]$. Note that, if the theta null point is k-rational, all the 2-torsion is.

Like in Proposition 1.20, given a point $P = (X : Z) \in \mathcal{K}_E$, the translation by a special 2-torsion point T_1 acts as permutation of the coordinates: $P + T_1 = (Z : X)$.

Definition 2.75. This model of x-line for E is called a *theta Kummer line.*

Remark 2.76 (Montgomery-theta isomorphism)*.* As the similarities suggest, a Montgomery curve and a theta-Kummer line are essentially the same model up to a linear change of variables. In a Montgomery curve, there is a canonical 4-torsion point $T' = (1 : -1)$ lying above $T = 2T' = (0 : 1)$. Let $S = (r : s) \in E[4]$ be another 4-torsion point such that (S, T) is a symplectic basis of $E[4]$. In [Rob24c, Appendix A] it is shown that the following map is an isomorphism from the x-only Montgomery model on E to a theta Kummer line having as theta null point $(r + s : r - s)$.

$$(X : Z) \in \mathbb{P}^1 \mapsto (\theta_0 = (r+s)(X-Z) : \theta_1 = (r-s)(X+Z)) \in \mathbb{P}^1.$$

Example 2.77 ([Mum66, §5.b])**.** If $(E, \mathcal{L})$ is an elliptic curve with a line bundle of level $\delta = 3$, we get a structure with marked rational 3-torsion. It determines an embedding of E as the cubic curve $X^3 + Y^3 + Z^3 - \mu XYZ$ in $\mathbb{P}^2$, for some μ with $\mu^3 \neq 1$.

Example 2.78 ([Mum66, §5.c])**.** In level $n = 4$, we get a system of 4 projective coordinates describing points on an elliptic curve. As a general fact, when $4|n$, the image of the projective embedding into $\mathbb{P}^{n^g-1}$ is an intersection of quadratic equations.

In the case of an elliptic curve E, we retrieve a Jacobi intersection model, describing the curve as intersection of two quadrics in $\mathbb{P}^3$.

Let us now turn to dimension 2.

Example 2.79 ([Mum66, §5.d])**.** Let A be a geometrically simple abelian variety of dimension $g = 2$. Up to isomorphism, it is the Jacobian of a genus-2 hyperelliptic curve C (see Proposition 2.41), with theta divisor Θ_C. Fix a symmetric theta structure of level 2 for the totally symmetric $\mathcal{L} = \mathcal{L}_{2(\Theta_C)}$. This induces a projective embedding of the Kummer surface $\mathcal{K}_A$ in $\mathbb{P}^3$, so that points on $\mathcal{K}_A$ are described by four projective coordinates $(\theta_{00}, \theta_{01}, \theta_{10}, \theta_{11})$ in $\mathbb{P}^3$.

The image of this projective embedding is a surface in $\mathbb{P}^3$ which has 16 singular points – the image points of the 2-torsion $A[2]$ – and is defined by an irreducible quartic equation.

Let $(a : b : c : d) = 0_A$ be the theta null point on the Kummer variety.

As in Example 2.74, the 2-torsion has a special form: $A[2] = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$, with

$$K_1 = \langle (b : a : d : c), (c : d : a : b) \rangle , \qquad K_2 = \langle (a : -b : c : -d), (a : b : -c : -d) \rangle$$

Like in dimension 1, by Proposition 2.72, translation by points of K_1 acts on Kummer points as a permutation of the coordinates, whereas points of K_2 act as a change of sign on some of the coordinates.

Chapter 3

Arithmetic on theta structures

In Chapter 2 we presented Mumford's theory of algebraic theta functions on a g-dimensional polarised abelian variety A over k equipped with an ample line bundle $\mathcal{L}$, and we began to see how theta structures give an explicit description of A as a projective variety. In this chapter, we will analyse more precisely the algebraic relations satisfied by the theta functions. Studying how theta functions on a line bundle $\mathcal{L}$ and its square $\mathcal{L}^{\otimes 2}$ interact with each other, we'll be able to derive algorithms to compute the group law on A. The same theory lets us describe isogenies between abelian varieties: given a theta model for $(A, \mathcal{L})$ and a subgroup $K \subseteq H(\mathcal{L})$ of the torsion underlying the theta group, we get an algorithm to compute the isogeny $\varphi\colon A \to A/K$ in the sense of Problems 1.32 and 1.33.

3.1 Isogeny theorem

Let us introduce the setting where we will work thoughout the chapter.

Let $(A, \mathcal{L})$ and $(B, \mathcal{M})$ be polarised abelian varieties of the same dimension g, where $\mathcal{L}, \mathcal{M}$ are separable ample line bundles. Let $f\colon A \to B$ be a separable isogeny of polarised abelian varieties with kernel $K = \ker(f)$. As we saw in Definition 2.30, this means that there is an isomorphism $\alpha\colon f^*\mathcal{M} \xrightarrow{\sim} \mathcal{L}$.

Definition 3.1. $K = \ker f$ is a finite group, say $K \cong \mathbb{Z}/d_1\mathbb{Z} \oplus \cdots \oplus \mathbb{Z}/d_r\mathbb{Z}$. Then we say that f is a $(d_1, \ldots, d_r)$-isogeny. In practice, in the sequel, we will be interested in $(\ell, \ldots, \ell)$-isogenies, also called ℓ-isogenies for brevity.

Remark 3.2. If $f\colon (A, \mathcal{L}) \to (B, \mathcal{M})$ is an isogeny of PAVs, then $K = \ker f$ is necessarily a subgroup of $H(\mathcal{L})$. In fact, for all $Q \in K$, we have an isomorphism

$$t_Q^*\alpha\colon t_Q^* f^*\mathcal{M} = f^*\mathcal{M} \xrightarrow{\sim} t_Q^*\mathcal{L},$$

because $f \circ t_Q = f$. Consequently,

$$\psi_Q = t_Q^*\alpha \circ \alpha^{-1}\colon \mathcal{L} \xrightarrow{\sim} t_Q^*\mathcal{L}$$

is an isomorphism, and (Q, ψ_Q) is an element of the theta group $G(\mathcal{L})$.

Proposition 3.3. *Let $f\colon (A,\mathcal{L}) \to (B,\mathcal{M})$ be an isogeny of PAVs with kernel K. There is a subgroup $\widetilde{K} \subseteq G(\mathcal{L})$ isomorphic to K via the restriction of the projection map $G(\mathcal{L}) \to H(\mathcal{L})$.*

Conversely, given a subgroup $K \subseteq H(\mathcal{L})$, if there exists a lift $\widetilde{K}$ of K in $G(\mathcal{L})$, then there exists a polarised variety $(B,\mathcal{M})$ such that $f\colon (A,\mathcal{L}) \to (B,\mathcal{M})$ is an isogeny of PAVs with kernel K. Such a lift exists if and only if K is isotropic with respect to the commutator pairing $e_{\mathcal{L}}$.

Definition 3.4. A subgroup $\widetilde{K} \subseteq G(\mathcal{L})$ such that the natural projection $G(\mathcal{L}) \xrightarrow{\pi} H(\mathcal{L})$ restricts to an isomorphism $\widetilde{K} \xrightarrow{\sim} \pi(\widetilde{K})$ is said to be a *level subgroup.*

Proof of Proposition 3.3. If f is an isogeny of PAVs and K is its kernel, the subset $\widetilde{K} = \{(Q,\psi_Q) \mid Q \in K\} \subseteq G(\mathcal{L})$, where ψ_Q was defined in Remark 3.2, is a subgroup isomorphic to K via the projection map $(Q,\psi_Q) \mapsto Q$.

The converse comes from Grothendieck's descent theory, see [Mum66, §1, Proposition 1]. The fact that K must be $e_{\mathcal{L}}$-isotropic follows from the commutativity of $\widetilde{K} \cong K$, see [Mum66, p. 293]. □

Theorem 3.5 (Mumford's Isogeny theorem)**.** *Let $f\colon (A,\mathcal{L}) \to (B,\mathcal{M})$ be an isogeny of PAVs, and $\widetilde{K}$ a corresponding level subgroup in $G(\mathcal{L})$. The following holds:*

(i) $f^{-1}(H(\mathcal{M})) \subseteq H(\mathcal{L})$.

(ii) The centraliser $Z(\widetilde{K})$ of $\widetilde{K}$ in $G(\mathcal{L})$ equals $Z' := \{(P,\varphi) \in G(\mathcal{L}) \mid f(P) \in H(\mathcal{M})\}$

(iii) There is a canonical isomorphism $G(\mathcal{M}) \cong Z(\widetilde{K})/\widetilde{K}$.

Proof. Fix an isomorphism $\alpha\colon f^*\mathcal{M} \to \mathcal{L}$.

(i) Fix $Q = f(P) \in H(\mathcal{M})$ and an isomorphism $\psi_Q\colon t_Q^*\mathcal{M} \xrightarrow{\sim} \mathcal{M}$. The composition

$$\varphi_P\colon \quad \mathcal{L} \xrightarrow{\alpha^{-1}} f^*\mathcal{M} \xrightarrow{f^*\psi_Q} f^*t_Q^*\mathcal{M} = t_P^*f^*\mathcal{M} \xrightarrow{t_P^*\alpha} t_P^*\mathcal{L}$$

is an isomorphism, so $P \in H(\mathcal{L})$.

(ii) By general descent theory, if $Q = f(P) \in B$, the line bundle $\mathcal{M}$ is defined identifying

$$\mathcal{L} \xrightarrow{\varphi} t_T^*\mathcal{L} \qquad \text{for all } (T,\varphi) \in \widetilde{K}.$$

and similarly $t_Q^*\mathcal{M}$ comes from the identification $t_P^*\mathcal{L} \xrightarrow{t_Q^*\varphi} t_{P+T}^*\mathcal{L}$. Isomorphisms $\varphi_P\colon \mathcal{L} \to t_P^*\mathcal{L}$ "descend" to $\mathcal{M}$ if and only if they commute with these identifications, that is, if for $(T,\varphi) \in \widetilde{K}$ the following commutes:

$$\begin{tikzcd}
\mathcal{L} \arrow[r, "\varphi_P"] \arrow[d, "\varphi"] & t_P^*\mathcal{L} \arrow[d, "t_P^*\varphi"] \\
t_T^*\mathcal{L} \arrow[r, "t_T^*\varphi"] & t_{P+T}^*\mathcal{L}.
\end{tikzcd}$$

This means that φ_P descends to $\mathcal{M}$, that is, induces a theta group element (Q,ψ_Q) above Q, if and only if it is in the centraliser of $\widetilde{K}$.

(iii) Given $(P, \varphi_P) \in Z'$, the isomorphism $(t_P^* \alpha)^{-1} \circ \varphi \circ \alpha \colon f^*\mathcal{M} \to f^* t_Q^* \mathcal{M}$ is of the form $f^*\psi_Q$ with $Q = f(P)$ and $(Q, \psi_Q) \in G(\mathcal{M})$, by the descent theory sketched above. Then $(P, \varphi_P) \mapsto (Q, \psi_Q)$ is a well-defined surjective map $Z' \to G(\mathcal{M})$, whose kernel is $\widetilde{K}$. □

The following lemma relates the degrees of the line bundles of isogenous PAVs, and will be important to us:

Lemma 3.6 ([Mum66, p. 291])**.** *If $f \colon (A, \mathcal{L}) \to (B, \mathcal{M})$ is an isogeny of PAVs with kernel K, then* $\deg \mathcal{L} = \#(K) \cdot \deg \mathcal{M}$.

Remark 3.7. Usually, when dealing with abelian varieties algorithmically, we'll work "in level n", meaning that every variety A will be equipped with a line bundle $\mathcal{L}$ of level n inducing a polarisation. This way A – or its Kummer variety $\mathcal{K}_A$ – will be realised projectively as a subvariety of $\mathbb{P}^{n^g-1}$ and its points can be represented as k-vectors of n^g coordinates. Usually, $\mathcal{L} = \mathcal{L}_0^n$ where $\mathcal{L}_0$ is a principal polarisation.

As a consequence of the above lemma, however, if we're given $(A, \mathcal{L})$ of level n, we cannot directly compute a nontrivial isogeny to some $(B, \mathcal{M})$ of level n. It won't be an isogeny of PAVs because it won't respect the polarisation degrees. If we want to compute an ℓ-isogenies from A to B, we should first derive a description of $(A, \mathcal{L}^\ell)$ from our knowledge of $(A, \mathcal{L})$ ("raise" the level from n to ℓn), and then compute an isogeny of PAVs $(A, \mathcal{L}^\ell) \to (B, \mathcal{M})$. The pattern is the following:

$$\begin{tikzcd}
\text{level } \ell n & (A, \mathcal{L}^\ell) \arrow[dr, "f"] & \\
\text{level } n & (A, \mathcal{L}) \arrow[u, dashed] & (B, \mathcal{M})
\end{tikzcd}$$

Isogenies and theta structures

To make the description of isogenies more explicit, we see how isogenies of PAVs relate to theta structures.

We look back at Problems 1.32 and 1.33. In our isogeny computations, we are given a theta structure of level n on a PAV $(A, \mathcal{L}, \Theta_{\mathcal{L}})$, and a subgroup $K \subseteq H(\mathcal{L})$ that will be the kernel of an isogeny f. We want to compute a theta structure $\Theta_{\mathcal{M}}$ for the codomain variety $B = A/K$ with a suitable polarisation $\mathcal{M}$ (so that f is an isogeny of PAVs), and relate the theta coordinates of A to those of B.

We need a notion of "compatibility" between two theta structures on A and B, to refine even more the notion of *isogeny of PAVs.* We start by recalling the following:

Remark 3.8. A theta structure $\Theta_{\mathcal{L}}$ induces a canonical section $s_{\mathcal{L}} \colon H(\mathcal{L}) \to G(\mathcal{L})$ that sends $P = \overline{\Theta_{\mathcal{L}}}(t_1, t_2)$ to $\Theta_{\mathcal{L}}(1, t_1, t_2)$. In particular, given any isotropic subgroup $K \subseteq H(\mathcal{L})$, there is a canonical level subgroup $\widetilde{K} = s_{\mathcal{L}}(K)$.

Definition 3.9. The level subgroup $\widetilde{K}$ defined by an isogeny of PAVs $f \colon (A, \mathcal{L}, \Theta_{\mathcal{L}}) \to (B, \mathcal{M})$ is said to be *compatible with* $\Theta_{\mathcal{L}}$ if it respects the symplectic structure induced by $s_{\mathcal{L}}$, that is, if $\widetilde{K} = (\widetilde{K} \cap \widetilde{K}_1(\mathcal{L})) \oplus (\widetilde{K} \cap \widetilde{K}_2(\mathcal{L}))$, with $\widetilde{K}_i(\mathcal{L}) = s_{\mathcal{L}}(K_1(\mathcal{L}))$.

We say that the theta structures $\Theta_{\mathcal{L}}, \Theta_{\mathcal{M}}$ are *compatible* with f if the quotient map $\alpha_f \colon Z(\widetilde{K}) \twoheadrightarrow G(\mathcal{M})$ sends $\widetilde{K}_i(\mathcal{L}) \cap Z(\widetilde{K})$ to $\widetilde{K}_i(\mathcal{M})$ for $i = 1, 2$.

Indeed, not all theta structures on $(B, \mathcal{M})$ are compatible with the isogeny f.

Theorem 3.10 (Isogeny theorem with theta structures)**.** *Let $f \colon (A, \mathcal{L}, \Theta_{\mathcal{L}}) \to (B, \mathcal{M})$ be an isogeny of PAVs with kernel $K \cong K_1 \oplus K_2$, where $\mathcal{L}$ (resp. $\mathcal{M}$) is of level $\delta_{\mathcal{L}}$ (resp. $\delta_{\mathcal{M}}$), $K_i \subseteq K_i(\mathcal{L})$. Define a level subgroup $\widetilde{K} = s_{\mathcal{L}}(K)$, so that it's automatically compatible with $\Theta_{\mathcal{L}}$. Let $K^{\perp}$ be the orthogonal complement of K in $H(\mathcal{L})$ with respect to the commutator pairing, and decompose it as $K^{\perp} = K^{\perp,1} \oplus K^{\perp,2}$. Note that $K \subseteq K^{\perp}$ since K is isotropic.*

The set of theta structures $(B, \Theta_{\mathcal{M}})$ of level $\delta_{\mathcal{M}}$ compatible with $(f, \Theta_{\mathcal{L}})$ is in bijection with the set of isomorphisms $\overline{\sigma} \colon K^{\perp,1}/K_1 \to K_1(\delta_{\mathcal{M}})$.

Fix one compatible $\Theta_{\mathcal{M}}$ and the corresponding isomorphism $\overline{\sigma}$. Also, let $\sigma \colon K^{\perp,1} \to K_1(\delta_{\mathcal{M}})$ be the projection induced by $\overline{\sigma}$. There exists a factor $\lambda \in (k^{\mathrm{alg}})^*$ such that for all $i \in K_1(\delta_{\mathcal{L}})$, we have

$$f^* \theta_i^{\mathcal{M}} = \lambda \sum_{j \in \overline{\Theta_{\mathcal{L}}}^{-1}(\sigma^{-1}(\{i\}))} \theta_j^{\mathcal{L}} \tag{3.1}$$

where $(\theta_i^{\mathcal{L}})_{i \in K_1(\delta_{\mathcal{L}})}$, $(\theta_i^{\mathcal{M}})_{i \in K_2(\delta_{\mathcal{M}})}$ are the bases of theta functions induced by $\Theta_{\mathcal{L}}$, $\Theta_{\mathcal{M}}$ respectively.

Proof. This statement is proved in [Mum66, §1, Theorem 4]. □

Example 3.11 (3-isogeny of elliptic curves in position K_2)**.** Let E be an elliptic curve equipped with a theta structure $\Theta_{\mathcal{L}}$ of level 6. Let $E[6] = \langle S_1, S_2 \rangle$ with $S_1 = \overline{\Theta}_{\mathcal{L}}(1, 0)$ and $S_2 = \overline{\Theta}_{\mathcal{L}}(0, 1)$ independent 6-torsion points induced by the theta structure.

We have $K_i(\mathcal{L}) = \langle S_i \rangle$. Let $T_2 = [2]S_2$, and let $f \colon E \to E'$ be the 3-isogeny of kernel $K = \langle T_2 \rangle = [2]K_2(\mathcal{L})$. By Proposition 3.6, the codomain curve $(E', \mathcal{M})$ must be equipped with a theta structure of level $\delta = 2$.

We want to apply the above theorem, knowing that the kernel is in position $K = [2]K_2(\mathcal{L})$, corresponding to the subgroup $\langle 0 \rangle \times \langle 2 \rangle \subseteq H(6) = (\mathbb{Z}/6\mathbb{Z})^2$. Its orthogonal via the pairing e_δ is the subgroup $\langle 3 \rangle \times \langle 1 \rangle$. More precisely, we have $K_1 = \{0_E\}, K_2 = K$, where $K_i = K \cap K_i(\mathcal{L})$, and $K^{\perp} = K^{\perp,1} \oplus K^{\perp,2} = [3]K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$.

In particular, there is a unique isomorphism

$$\sigma \colon \quad K^{\perp,1}/K_1 = [3]K_1(\mathcal{L})/\{0_E\} \cong 3\mathbb{Z}/6\mathbb{Z} \quad \xrightarrow{\sim} \quad K_1(\delta_{\mathcal{M}}) \cong \mathbb{Z}/2\mathbb{Z}$$

(namely, $i \mapsto i/3$), and this induces theta functions $(\theta_0^{\mathcal{M}}, \theta_1^{\mathcal{M}})$ such that:

$$\theta_i^{\mathcal{M}}(f(P)) = \lambda \sum_{j \in \overline{\Theta_{\mathcal{L}}}^{-1}(\overline{\sigma}^{-1}(\{i\}))} \theta_j^{\mathcal{L}} = \theta_{3i}^{\mathcal{L}} \qquad \text{for } i = 0, 1, \text{ for some } \lambda.$$

In fact, representing P and $f(P)$ in coordinates, we have:

$$P = (\theta_0^{\mathcal{L}}(P) : \cdots : \theta_5^{\mathcal{L}}(P)) = (X_0 : X_1 : X_2 : X_3 : X_4 : X_5), \quad f(P) = (X_0 : X_3).$$

Example 3.12 (3-isogeny of elliptic curves in position K_1)**.** Keep the notations of the previous example: let E be an elliptic curve equipped with a theta structure $\Theta_{\mathcal{L}}$ of level 6, with induced torsion $E[6] = \langle S_1, S_2 \rangle = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$.

This time, let $T_1 = [2]S_1$, and let $f\colon E \to E'$ be the 3-isogeny of kernel $K = \langle T_1 \rangle = [2]K_1(\mathcal{L})$. We find theta coordinates of level 2 for the codomain curve $(E', \mathcal{M})$.

Applying the isogeny theorem, using now $K = [2]K_1(\mathcal{L})$, we get

$$K_1 = K \cap K_1(\mathcal{L}) = K = [2]K_1(\mathcal{L}),$$
$$K^{\perp} = K_1(\mathcal{L}) \oplus [3]K_2(\mathcal{L}), \quad \text{so } K^{\perp,1} = K_1(\mathcal{L}).$$

Again, there is a unique isomorphism $\sigma\colon K^{\perp,1}/K_1 \cong (\mathbb{Z}/6\mathbb{Z})/(2\mathbb{Z}/6\mathbb{Z}) \xrightarrow{\sim} K_1(\delta_{\mathcal{M}}) \cong \mathbb{Z}/2\mathbb{Z}$, namely $(i \mod 6) \mapsto (i \mod 2)$. The induced theta functions $(\theta_0^{\mathcal{M}}, \theta_1^{\mathcal{M}})$ satisfy:

$$\theta_i^{\mathcal{M}}(f(P)) = \lambda \sum_{j \in \overline{\Theta_{\mathcal{L}}}^{-1}(\overline{\sigma}^{-1}(\{i\}))} \theta_j^{\mathcal{L}} = \theta_i^{\mathcal{L}} + \theta_{i+2}^{\mathcal{L}} + \theta_{i+4}^{\mathcal{L}} \qquad i = 0, 1.$$

In coordinates:

$$P^{\mathcal{L}} = (X_0 : X_1 : X_2 : X_3 : X_4 : X_5), \quad f(P)^{\mathcal{M}} = (X_0 + X_2 + X_4 : X_1 + X_3 + X_5).$$

Example 3.13. Now let's move to dimension $g = 2$.

Let $(A, \mathcal{L}_0) = \mathrm{Jac}_C$ be the Jacobian of a genus-2 hyperelliptic curve over k with its canonical principal polarisation, and suppose we're given a theta structure $\Theta_{\mathcal{L}}$ of level 4 on A. Now consider the symplectic basis $A[4] = \langle U_1, U_2, S_1, S_2 \rangle$ induced by the theta structure. Let $T_i = [2]S_i$, $i = 1, 2$, and consider the $(2,2)$-isogeny $f\colon A \to B$ having kernel $K = \langle T_1, T_2 \rangle = [2]K_2(\mathcal{L})$. As in Example 3.11, there is a unique level-2 structure $\Theta_{\mathcal{M}}$ on B compatible with the isogeny, and its coordinates satisfy the relation

$$\begin{aligned}(\theta_{00}^{\mathcal{M}}(f(P)) : \theta_{01}^{\mathcal{M}}(f(P)) : \theta_{10}^{\mathcal{M}}(f(P)) : \theta_{11}^{\mathcal{M}}(f(P))) = \\ (\theta_{00}^{\mathcal{L}}(P) : \theta_{02}^{\mathcal{L}}(P) : \theta_{20}^{\mathcal{L}}(P) : \theta_{22}^{\mathcal{L}}(P))\end{aligned}$$

where, again, $i \in (\mathbb{Z}/2\mathbb{Z})^2 \mapsto 2i \in (\mathbb{Z}/4\mathbb{Z})^2$ is the canonical embedding.

3.2 The differential addition law

In this paragraph, we will focus on the group law of an abelian variety A. We will be able to compute it explicitly *in coordinates* using theta functions. More precisely, using the isogeny theorem, we will get a *differential addition* algorithm, as in Algorithm 1: given coordinates of two points $P, Q \in A$ and their difference $P - Q$, we will be able to compute the coordinates of $P + Q$.

Symmetric theta structures

In Definition 2.66 we defined symmetric theta structures on totally symmetric line bundles, seeing that symmetric theta structures of level 2 allow us to describe Kummer varieties in projective coordinates. We see now that symmetric theta structures have a much more powerful use: they let us relate the theta functions of a line bundle $\mathcal{L}$ to those of its powers $\mathcal{L}^{2^m}$, and this will eventually let us obtain algorithms for the explicit computation of arithmetic operations on A and 2-isogenies.

As in the previous section, we need a notion of *compatibility* between different theta structures.

Let $\Theta_{\mathcal{L}} \colon G(\delta) \to G(\mathcal{L})$, with $\delta = (n, \dots, n)$, be a theta structure of level n on a PAV $(A, \mathcal{L})$. Recall that, basically by definition, it induces an isomorphism $\overline{\Theta}_{\mathcal{L}} \colon (\mathbb{Z}/n\mathbb{Z})^g \times (\mathbb{Z}/n\mathbb{Z})^g \to H(\mathcal{L}) = A[n]$. This isomorphism gives us a *numbering* of the n-torsion points, and a symplectic basis of $H(\mathcal{L})$ as follows:

$$H(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L}) = \langle S_1, \dots, S_g \rangle \oplus \langle U_1, \dots, U_g \rangle .$$

where we denote $S_i = \overline{\Theta}_{\mathcal{L}}(e_i, 0)$ and $U_i = \overline{\Theta}_{\mathcal{L}}(0, e_i)$, with e_i being the i-th standard basis vector of $(\mathbb{Z}/n\mathbb{Z})^g$. Here $K_j(\mathcal{L})$, $j = 1, 2$ are two $e_{\mathcal{L}}$-isotropic subgroups and their generators satisfy $e_{\mathcal{L}}(S_i, U_j) = \delta_{ij}$ with δ_{ij} the Kronecker delta.

Now, if we have a theta structure $\Theta_{\mathcal{L}}$ and we want another structure on $\mathcal{L}^2$ to be compatible with $\Theta_{\mathcal{L}}$, we need to impose that the numbering of the $2n$-torsion points of A induced by $\Theta_{\mathcal{L}}^2$ is compatible with the numbering of the n-torsion points induced by $\Theta_{\mathcal{L}}$. In practice, this means $\overline{\Theta}_{\mathcal{L}^2}$ inducing a symplectic basis of the $2n$-torsion $H(\mathcal{L}^2) = \langle S_1', \dots, S_g' \rangle \oplus \langle U_1', \dots, U_g' \rangle$ such that $[2]S_i' = S_i$, $[2]U_i' = U_i$.

The above paragraph tells us how two compatible theta structures on $\mathcal{L}$ and $\mathcal{L}^2$ behave on torsion points, but says nothing about their lifts in $G(\mathcal{L}), G(\mathcal{L}^2)$. In fact, we'd need a more precise notion of compatibility, but that would be too technical for our purposes: see [Mum66, §2], [Rob21, Section 2.6] for complete definitions and proofs.

We will just state some conclusions that will be useful for our next constructions.

Theorem 3.14. *Consider a PAV $(A, \mathcal{L})$ where $\mathcal{L}$ is totally symmetric of level δ.*

(i) Every symmetric theta structure $\Theta_{\mathcal{L}^2}$ induces a compatible theta structure $\Theta_{\mathcal{L}}$.

(ii) The induced $\Theta_{\mathcal{L}}$ only depends on the symplectic isomorphism $H(2\delta) \xrightarrow{\sim} H(\mathcal{L}^2)$.

(iii) Every symmetric theta structure $\Theta_{\mathcal{L}}$ is induced by some $\Theta_{\mathcal{L}^2}$, or equivalently by some symplectic isomorphism $H(2\delta) \xrightarrow{\sim} H(\mathcal{L}^2)$.

(iv) If $f \colon (A, \Theta_{\mathcal{L}}) \to (B, \Theta_{\mathcal{M}})$ is an isogeny of PAVs compatible with theta structures and $\Theta_{\mathcal{L}}$ is symmetric, then $\Theta_{\mathcal{M}}$ is also symmetric.

Remark 3.15. Throughout this chapter, we will assume $\mathcal{L}$ is a totally symmetric ample invertible sheaf on A, so that we'll be able to choose pairwise compatible theta structures on $\mathcal{L}^{\otimes 2^m}, m \in \mathbb{N}$. Suppose indeed to fix such structures $\Theta_{\mathcal{L}^{2^m}}$.

Affine coordinates

Before going on with the rest of the chapter, we want to make some remarks aimed at shifting our view towards a more computational perspective.

In this chapter, our goal is to develop algorithms giving a computational description of the arithmetic on abelian varieties. We introduced theta structures so that we have nice projective coordinate systems to represent points on abelian varieties. Indeed, recall the following:

Remark 3.16. Let $(A, \mathcal{L}_0)$ be a g-dimensional PPAV. The line bundle $\mathcal{L} = \mathcal{L}_0^n$ of level $\delta = n \geq 3$ is very ample. Any theta structure $\Theta_{\mathcal{L}}$ induces a projective embedding

$$\Phi_{\Theta_{\mathcal{L}}} \colon A \hookrightarrow \mathbb{P}^{d_1 \cdots d_r - 1}, \qquad P \mapsto (\theta_i(P))_{i \in K(\delta)}.$$

If instead $n = 2$, the line bundle $\mathcal{L} = \mathcal{L}_0^2$ is totally symmetric, A is geometrically simple and $\Theta_{\mathcal{L}}$ is a symmetric theta structure, then it induces a map $\Phi_{\Theta_{\mathcal{L}}}$ of the same form as above, that embeds $\mathcal{K}_A$ into $\mathbb{P}^{2^g-1}$.

Any point of A (or its Kummer variety, if the hypotheses above hold) can be (mathematically) represented in projective coordinates via $\Phi_{\Theta_{\mathcal{L}}}$, once we have all the necessary ingredients.

However, when representing a point $P \in A$ on a computer, we don't store in memory its *projective coordinates* $\Phi_{\Theta_{\mathcal{L}}}(P)$, but rather a certain *affine lift* $\overline{P}$ of the projective point $\Phi_{\Theta_{\mathcal{L}}}(P)$. Rescaling the affine tuple $\overline{P}$ by some nonzero scalar $\lambda \in k^*$ doesn't change the underlying projective point (that is, it still represents P) but concretely gives a different tuple when stored on a computer.

Now look back at Remark 2.71. Let $\mathcal{L}$ be a line bundle on A and $P \in A$ a point. Fixing a rigidification of $\mathcal{L}$ at P is equivalent to choosing the *value* of one (equivalently, every) global section $s \in \Gamma(A, \mathcal{L})$ at P, denoted by $\overline{s}(P)$, and in particular it is equivalent to fixing an affine lift $\overline{P} = (\overline{\theta_i}(P))_i$ of $\Phi_{\Theta_{\mathcal{L}}}(P)$.

If moreover $\mathcal{L}$ is totally symmetric, $\Theta_{\mathcal{L}}$ is symmetric and we fix a compatible structure $\Theta_{\mathcal{L}^2}$ on $\mathcal{L}^2$, then a rigidification of $\mathcal{L}$ at P canonically induces one on $\mathcal{L}^2$ at the same point, by the isomorphism chain $(\mathcal{L}^{\otimes 2})_P \cong (\mathcal{L}_P)^{\otimes 2} \cong k^{\otimes 2} \cong k$. Consequently, for example, fixing level-2 coordinates of a point $P \in A$ uniquely induces level-4 coordinates of the same point that are compatible with the level-2 ones.

These notations concerning the affine coordinates of a point $P \in A$ will be extensively used in the following sections, and ease our exposition of the arithmetic algorithms.

Later, we will see that affine coordinates in symmetric theta structures also interact well with 2-isogenies and 4-isogenies (see, e.g., Remark 3.38).

The duplication formula

We want to describe the group law of an abelian variety *in coordinates*, making use of theta structures and their theta functions. If we worked with affine varieties, the group law on the variety would translate directly to a map – the *comultiplication* map – on their coordinate rings. With projective varieties, however, this doesn't happen.

We can instead recover a coordinate description of the addition law on A studying the variety $A \times A$ directly, via the following map:

$$\begin{array}{rcl} \xi\colon A \times A & \to & A \times A \\ (P, Q) & \mapsto & (P+Q, P-Q). \end{array}$$

The map ξ is a separable isogeny, whenever $\operatorname{char} k \neq 2$. Its kernel is $\Delta(A[2])$, the 2-torsion points on the diagonal of A, and is isomorphic to $(\mathbb{Z}/2\mathbb{Z})^{2g}$.

The fact that ξ is a 2-isogeny gives us a practical way to relate theta functions on $\mathcal{L}^2$ and $\mathcal{L}$, respectively of level 2δ and δ, using the isogeny theorem we presented in Section 3.1. Eventually, we'll be able to derive algorithms for the explicit computation of the (differential) addition law using coordinates of level δ.

Theorem 3.17. *Let A be an abelian variety of dimension g, and let $\mathcal{L}$ be a symmetric invertible sheaf on A. If $\mathcal{M}$ is the invertible sheaf $\pi_1^*\mathcal{L} \otimes \pi_2^*\mathcal{L}$, where π_i is the projection on the i-th component, then $\xi^*(\mathcal{M}) \cong \mathcal{M}^{\otimes 2}$.*

Proof. By the see-saw principle [Mil86a, §5], to show the isomorphism $\xi^*(\mathcal{M}) \cong \mathcal{M}^{\otimes 2}$ it is sufficient to show that the two sheaves are isomorphic when restricted to the subschemes $A \times \{R\}$ for all $R \in A$ and $\{0_A\} \times A$. Consider for $i = 1, 2$ the morphisms

$$j_R^{(i)}\colon A \to A \times A, \quad j_R^{(1)}(P) = (P, R), \quad j_R^{(2)}(P) = (R, P).$$

Note that $\pi_i \circ j_R^{(i)} = \mathrm{id}_A$, and $\pi_i \circ j_R^{(i)}(P) = R$ for all $P \in A$ if $i \neq j$. This gives

$$(j_R^{(i)})^*(\mathcal{M}^2) \cong \mathcal{L}^2.$$

On the other hand, we have

$$\begin{aligned} (j_R^{(1)})^*\xi^*\mathcal{M} &= (j_R^{(1)})^*\xi^*(\pi_1^*\mathcal{L} \otimes \pi_2^*\mathcal{L}) = (\pi_1 \circ \xi \circ j_R^{(1)})^*\mathcal{L} \otimes (\pi_2 \circ \xi \circ j_R^{(1)})^*\mathcal{L} \\ &= t_R^*\mathcal{L} \otimes t_{-R}^*\mathcal{L} \cong \mathcal{L}^2 (\text{by the theorem of the square}) \end{aligned}$$

and

$$\begin{aligned} (j_R^{(2)})^*\xi^*\mathcal{M} &= (j_R^{(2)})^*\xi^*(\pi_1^*\mathcal{L} \otimes \pi_2^*\mathcal{L}) = (\pi_1 \circ \xi \circ j_R^{(2)})^*\mathcal{L} \otimes (\pi_2 \circ \xi \circ j_R^{(2)})^*\mathcal{L} \\ &= \mathcal{L} \otimes [-1]^*\mathcal{L} \cong \mathcal{L}^2 (\text{by symmetry}) \end{aligned}$$

as wanted. □

The following theorem can be translated "in coordinates" using theta structures. We just need to understand how to correctly set a theta structure on $A \times A$ if we have one on A. The following lemma does the work.

Lemma 3.18 (Product theta structures, [Mum66, §3, Lemma 1])**.** *Let $(A_1, \mathcal{L}_1), \ldots, (A_r, \mathcal{L}_r)$ be PAVs. Let $A = \bigoplus_{i=1}^r A_i$ with the line bundle $\mathcal{L} = \pi_1^*\mathcal{L}_1 \otimes \cdots \otimes \pi_r^*\mathcal{L}_r$ defining the product polarisation (where $\pi_i\colon A \to A_i$ is the i-th projection).*

- *$H(\mathcal{L}) = \bigoplus_{i=1}^r H(\mathcal{L}_i)$ and the decomposition is orthogonal w.r.t. $e_{\mathcal{L}}$.*
- *$G(\mathcal{L}) \cong \left(\bigoplus_{i=1}^r G(\mathcal{L})\right) / \{(\lambda_1, \dots, \lambda_r) \in (k^*)^r \mid \lambda_1 \cdots \lambda_r = 1)\}$*
- *Suppose each A_i is equipped with a theta structure $\Theta_{\mathcal{L}_i}$ of level δ_i. Then we can naturally define a theta structure $\Theta_{\mathcal{L}}$ of level $\delta = (\delta_1, \dots, \delta_r)$, such that the canonical theta functions are*

$$\theta_i^{\mathcal{L}} = \pi_1^* \theta_{i_1}^{\mathcal{L}_1} \cdots \pi_r^* \theta_{i_r}^{\mathcal{L}_r} \qquad \text{for all } i \in K_1(\delta_1) \times \cdots \times K_r(\delta_r).$$

For simplicity, we will denote $\theta_i^{\mathcal{L}} = \pi_1^* \theta_{i_1}^{\mathcal{L}_1} \cdots \pi_r^* \theta_{i_r}^{\mathcal{L}_r}$ as $\theta_{i_1}^{\mathcal{L}_1} * \cdots * \theta_{i_r}^{\mathcal{L}_r}$.

Remark 3.19. Let's look at the above lemma in coordinates, by means of Remark 3.15.

Let $P_i \in A_i(k)$ be rational points for $i = 1, \dots, r$. Fix a rigidification of the line bundles $\mathcal{L}_i$ at P_i, inducing affine values $\overline{\theta_j}^{\mathcal{L}_i}(P_i) \in k$. The lemma reads:

$$\overline{\theta_i}^{\mathcal{L}}(P_1, \dots, P_r) = \overline{\theta_{i_1}}^{\mathcal{L}_1}(P_1) \cdots \overline{\theta_{i_r}}^{\mathcal{L}_r}(P_r) \qquad \text{for all } i \in K_1(\delta_1) \times \cdots \times K_r(\delta_r).$$

Now that we know how to equip a product invertible sheaf with a theta structure, look back at Theorem 3.17. Applying the isogeny theorem to the sheaf isomorphism $\xi^* \mathcal{M} \cong \mathcal{M}^{\otimes 2}$, we get the following corollary.

Notation 3.20. Let $i = (i_1, \dots, i_g) \in (\mathbb{Z}/n\mathbb{Z})^g$ be a multi-index. We will denote by ℓi the multi-index $(\ell i_1, \dots, \ell i_g) \in (\mathbb{Z}/\ell n\mathbb{Z})^g$, that is the image of i via the canonical embedding $(\mathbb{Z}/n\mathbb{Z})^g \hookrightarrow (\mathbb{Z}/\ell n\mathbb{Z})^g$. Viceversa, if $j \in (\mathbb{Z}/\ell n\mathbb{Z})^2$ is in the image (we write $j \equiv 0 \pmod{\ell}$) then $i = j/\ell$ is its unique preimage.

Corollary 3.21 (Duplication formula)**.** *Let $(\theta_i^{\mathcal{L}})_{i \in K_1(\delta)}, (\theta_i^{\mathcal{L}^2})_{i \in K_1(2\delta)}$ be the theta functions coming from compatible symmetric theta structures on A, and say $\mathcal{L}$ has level $\delta = n$. Remember that $K_1(\delta) = (\mathbb{Z}/n\mathbb{Z})^g, K_1(2\delta) = (\mathbb{Z}/2n\mathbb{Z})^g$.*

Let $P, Q \in A(k)$ be two rational points. Then there is a $\lambda \in k^*$ such that for all $i_1, i_2 \in K_1(\delta)$ and rigidifications of $\mathcal{L}$ at P, Q and of $\mathcal{L}^2$ at $P \pm Q$ the following holds:

$$\overline{\theta_{i_1}}^{\mathcal{L}}(P+Q) \overline{\theta_{i_2}}^{\mathcal{L}}(P-Q) = \lambda \sum_{\substack{j_1, j_2 \in (\mathbb{Z}/2n\mathbb{Z})^g \\ j_1 + j_2 = 2i_1 \\ j_1 - j_2 = 2i_2}} \overline{\theta_{j_1}}^{\mathcal{L}^2}(P) \overline{\theta_{j_2}}^{\mathcal{L}^2}(Q).$$

Proof (sketch). Let $\mathcal{M} = \pi_1^* \mathcal{L} \otimes \pi_2^* \mathcal{L}$. By Theorem 3.17, we have $\xi^* \mathcal{M} \cong \mathcal{M}^{\otimes 2} \cong \pi_1^* \mathcal{L}^2 \otimes \pi_2^* \mathcal{L}^2$. We fix on the line bundle $\mathcal{N} := \pi_1^* \mathcal{L}^2 \otimes \pi_2^* \mathcal{L}^2$ the product theta structure given by $\Theta_{\mathcal{L}^2}$ on both factors, and we call it $\Theta_{\mathcal{N}}$. Similarly, we fix the product theta structure $\Theta_{\mathcal{M}}$ on $\mathcal{M}$ given by $\Theta_{\mathcal{L}}$ on both factors. We have

$$\begin{aligned} H(\mathcal{N}) &= K_1(\mathcal{L}^2)^{\oplus 2} \oplus K_2(\mathcal{L}^2)^{\oplus 2} && \text{with } A[2n] = H(\mathcal{L}^2) = K_1(\mathcal{L}^2) \oplus K_2(\mathcal{L}^2), \\ H(\mathcal{M}) &= K_1(\mathcal{L})^{\oplus 2} \oplus K_2(\mathcal{L})^{\oplus 2} && \text{with } A[n] = H(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L}). \end{aligned}$$

We claim that, by compatibility of the theta structures on $\mathcal{L}$ and $\mathcal{L}^2$, the structures $\Theta_{\mathcal{N}}$ and $\Theta_{\mathcal{M}}$ are compatible with the isogeny ξ. Since we didn't give a proper definition of compatibility, we omit this verification.

Now, we apply the isogeny theorem to the isogeny ξ. Its kernel is $K = \Delta(A[2]) = \{(P,P) \mid P \in A[2]\}$. Writing $A[2] = [n]K_1(\mathcal{L}^2) \oplus [n]K_2(\mathcal{L}^2)$, we get the symplectic decomposition $K = \Delta([n]K_1(\mathcal{L}^2)) \oplus \Delta([n]K_2(\mathcal{L}^2))$. Via the theta structure, we have the following correspondences to subgroups of $(\mathbb{Z}/2n\mathbb{Z})^{4g}$ (where the components are rearranged so that the first half is $K_1(2\delta)$ and the second half $K_2(2\delta)$):

$$\begin{aligned}\overline{\Theta}_{\mathcal{N}}^{-1}(K) &= \{(t_1,t_1,t_2,t_2) \mid 2t_1 = 0, 2t_2 = 0\},\\ \overline{\Theta}_{\mathcal{N}}^{-1}(K_1) &= \{(t_1,t_1,0,0) \mid 2t_1 = 0, 2t_2 = 0\},\\ \overline{\Theta}_{\mathcal{N}}^{-1}(K_2) &= \{(0,0,t_2,t_2) \mid 2t_1 = 0, 2t_2 = 0\}.\end{aligned}$$

Similarly, $K_1^{\perp}$ corresponds to the subgroup of elements orthogonal to K_2 (being contained in the isotropic $K_1(\delta)$, it is already orthogonal to K_1):

$$\begin{aligned}\{(x,x',0,0) \mid \langle x|t_2\rangle \langle x'|t_2\rangle = \langle x+x'|t_2\rangle = 0 \text{ for all } t_2 \in (n\mathbb{Z}/2n\mathbb{Z})^g\}\\ = \{(x,x',0,0) \mid x + x' \equiv 0 \pmod 2\}\end{aligned}$$

With some more work, one can show that the map $\sigma\colon K_1^{\perp} \to K_1(\delta)$ is defined by $\overline{\Theta_{\mathcal{N}}}(x,x',0,0) \mapsto ((x+x')/2, (x-x')/2) \in (\mathbb{Z}/n\mathbb{Z})^{2g}$. Finally, by the isogeny theorem, fixing indices $i_1, i_2 \in (\mathbb{Z}/n\mathbb{Z})^g$, we get

$$\xi^*(\theta_{i_1}^{\mathcal{L}} * \theta_{i_2}^{\mathcal{L}}) = \sum_{(j_1,j_2)\in\overline{\Theta}_{\mathcal{N}}^{-1}(\sigma^{-1})\{(i_1,i_2)\}} \theta_{j_1}^{\mathcal{L}^2} * \theta_{j_2}^{\mathcal{L}^2}.$$

If we fix rigidifications of $\mathcal{L}$ at $P, Q, P-Q, P+Q$ (which are uniquely determined up to a scalar) we get the desired formula. □

The differential addition algorithm

The duplication formula of Corollary 3.21 finally gives us an algorithm for the differential addition on an abelian variety A.

Notation 3.22. We will use the following change of variables in the sequel. Let $(\theta_i^{\mathcal{L}})_{i\in(\mathbb{Z}/n\mathbb{Z})^g}$ be theta functions of *even* level $2n$ attached to some $\Theta_{\mathcal{L}}$ on A.

Let $u \in (\mathbb{Z}/2n\mathbb{Z})^g$ and $t' \in (\mathbb{Z}/2\mathbb{Z})^g$. We denote

$$U_{t',u}^{\mathcal{L}} = \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \langle t|t'\rangle\, \theta_{u+nt}^{\mathcal{L}} \qquad \text{with } \langle t|t'\rangle = (-1)^{t_1t_1'+\cdots+t_gt_g'}. \tag{3.2}$$

If we fix affine evaluations $\overline{\theta_i}^{\mathcal{L}}(P)$ on some P, then $\overline{U}_{t',u}^{\mathcal{L}}(P) \in k$ is well-defined as well.

Theorem 3.23 (Differential addition). *Let $(A,\mathcal{L})$ be a PAV, with $\mathcal{L}$ of level n, and fix a theta structure $\Theta_{\mathcal{L}}$ on it. Let $P, Q \in A(k)$ be two rational points. For $R \in \{P, Q, P-Q, 0_A\}$, fix affine coordinates $\overline{R} = (\overline{\theta_i})_{i\in(\mathbb{Z}/n\mathbb{Z})^g}$.*

If $\overline{U}^{\mathcal{L}^2}_{i,0}(0_A)$ and $\overline{\theta}^{\mathcal{L}}_i(P-Q)$ are nonzero for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$, there is an algorithm

$$\mathsf{diff\_add}\colon (\overline{P}, \overline{Q}, \overline{P-Q}, \overline{0_A}) \mapsto \overline{P+Q}$$

that returns (deterministically) affine theta coordinates of $P+Q$ of level n given those of $P, Q, P-Q, 0_A$, with a complexity of $O(n^g)$ field operations.Algorithm 7 describes the computations in level $n = 2$.

Definition 3.24. Specialising $\mathsf{diff\_add}$ with $P = Q$, we get an algorithm

$$\mathsf{dbl}\colon \overline{P} \mapsto \overline{2P}$$

called the *doubling algorithm.* See Algorithm 8 for the computations in level 2.

Lemma 3.25. *Let $(A, \mathcal{L}, \Theta_{\mathcal{L}})$ be of level n. Fix level-n coordinates $\overline{P}, \overline{Q}, \overline{P-Q}, \overline{0_A}$. They determine coordinates $\overline{P+Q}$ satisfying the following equations:*

$$\overline{\theta}^{\mathcal{L}}_{(u+v)/2}(P+Q)\overline{\theta}^{\mathcal{L}}_{(u-v)/2}(P-Q) = \sum_{t' \in (\mathbb{Z}/2\mathbb{Z})^g} \overline{U}^{\mathcal{L}^2}_{t',u}(P)\overline{U}^{\mathcal{L}^2}_{t',v}(Q) \tag{3.3}$$

$$\overline{U}^{\mathcal{L}^2}_{t',u}(P)\overline{U}^{\mathcal{L}^2}_{t',v}(Q) = \frac{1}{2^g} \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \langle t|t'\rangle \, \overline{\theta}^{\mathcal{L}}_{(u+v)/2+t}(P+Q)\overline{\theta}^{\mathcal{L}}_{(u-v)/2+t}(P-Q) \tag{3.4}$$

with $u, v \in (\mathbb{Z}/2n\mathbb{Z})^g$, $u \equiv v \mod 2$, $t' \in (\mathbb{Z}/2\mathbb{Z})^g$.

Proof. We verify (3.3). The following chain of equalities holds:

$$\begin{aligned}
\sum_{t'} \overline{U}^{\mathcal{L}^2}_{t',u}(P)\overline{U}^{\mathcal{L}^2}_{t',v}(Q) &= \sum_{t',t_1,t_2} \langle t_1|t'\rangle \, \overline{\theta}^{\mathcal{L}^2}_{u+2t_1}(P) \, \langle t_2|t'\rangle \, \overline{\theta}^{\mathcal{L}^2}_{v+2t_2}(Q) \\
&= \sum_{t_1,t_2} \overline{\theta}^{\mathcal{L}^2}_{u+2t_1}(P)\overline{\theta}^{\mathcal{L}^2}_{v+2t_2}(Q) \sum_{t'} \langle t_1+t_2|t'\rangle \\
\big(\text{using } \textstyle\sum_{t'} \langle t_1+t_2|t'\rangle = 2^g\delta_{t_1t_2}\big) &= 2^g \sum_{t} \overline{\theta}^{\mathcal{L}^2}_{u+2t}(P)\overline{\theta}^{\mathcal{L}^2}_{v+2t}(Q) \\
&= 2^g \overline{\theta}^{\mathcal{L}}_{(u+v)/2}(P+Q)\overline{\theta}^{\mathcal{L}}_{(u-v)/2}(P-Q)
\end{aligned}$$

where all the summation indices are in $(\mathbb{Z}/2\mathbb{Z})^g$. The last equality follows from the duplication formula and the fact that the pairs (j_1, j_2) that satisfy $(j_1+j_2)/2 = (u+v)/2$ and $(j_1-j_2)/2 = (u-v)/2$ are exactly those of the form $(u+t, v+t)$ with $t \in (\mathbb{Z}/2\mathbb{Z})^g$. Replacing $\overline{\theta}^{\mathcal{L}}_i(P+Q)$ with $(1/2^g)\overline{\theta}^{\mathcal{L}}_i(P+Q)$ gives the desired result.

Equation (3.4) is obtained from (3.3) similarly. □

Proof of Theorem 3.23. Specialise the equations of the above lemma as follows. Use (3.4) with $u = v = 0$, $P = Q = 0_A$ to get

$$\overline{U}^{\mathcal{L}^2}_{t',0}(0_A)^2 = (1/2^g) \textstyle\sum_t \langle t|t'\rangle \, \overline{\theta}^{\mathcal{L}}_t(0_A)^2. \tag{3.5}$$

Then (3.4) with $v = 0, Q = 0_A$ and $i \in (\mathbb{Z}/2\mathbb{Z})^g, u = 2i$ gives:

$$\overline{U}^{\mathcal{L}^2}_{t',2i}(P)\overline{U}^{\mathcal{L}^2}_{t',0}(0_A) = (1/2^g)\sum_t \langle t|t'\rangle\, \overline{\theta}^{\mathcal{L}}_{i+t}(P)^2. \tag{3.6}$$

Finally, use (3.3) with $u = 2i, v = 0$:

$$\begin{aligned} \overline{\theta}^{\mathcal{L}}_i(P+Q)\overline{\theta}^{\mathcal{L}}_i(P-Q) &= \sum_{t'} U^{\mathcal{L}^2}_{t',2i}(P)U^{\mathcal{L}^2}_{t',0}(Q) \\ &= \sum_{t'} \frac{(\overline{U}^{\mathcal{L}^2}_{t',2i}(P)\overline{U}^{\mathcal{L}^2}_{t',0}(0_A))\cdot(\overline{U}^{\mathcal{L}^2}_{t',0}(Q)\overline{U}^{\mathcal{L}^2}_{t',0}(0_A))}{\overline{U}^{\mathcal{L}^2}_{t',0}(0_A)^2} \end{aligned} \tag{3.7}$$

The last term is computable from $\overline{0_A}, \overline{P}, \overline{Q}, \overline{P-Q}$ given the equations above. □

Remark 3.26. We would need to make sure that $U^{\mathcal{L}^2}_{t',0}(0)$ is always nonzero for all t'. Usually, this is the case. To see how to treat the case when this does not hold, see [DMPR23a, Remark 5] or [Dar24, Appendix A].

Notation 3.27. In the algorithms, we use level-2 coordinates. A point R is represented as a vector $x = (x_i)_{i\in(\mathbb{Z}/2\mathbb{Z})^g}$ with $x_i = \overline{\theta_i}(R)$. Rescaling a vector x by a scalar λ is denoted as λx or $\lambda \cdot x$.

The expressions of the form $\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} x_t$ are implemented as a matrix-vector product with the Hadamard matrix $\mathcal{H}$ defined by $\mathcal{H}_{i,j} = \langle i|j\rangle$. The sum in (3.2) is then computed as $\mathcal{H}\cdot x$. Its cost is $g \cdot 2^g$ additions, as explained in [Dar24, Appendix D.2].

Given two vectors x, y, we denote the component-wise product as $x * y$, component-wise inversion as $\mathcal{I}(x)$ and component-wise squaring as $\mathcal{S}(x)$.

Algorithm 7 Theta differential addition on a Kummer variety

```
Input: Affine level-2 theta coordinate vectors $\overline{P}, \overline{Q}, \overline{P-Q}, \overline{0_A}$.
Output: Affine level-2 theta coordinates $\overline{P+Q}$.
  $P' \leftarrow \mathcal{H}\circ\mathcal{S}(\overline{P})$          ▷ $(P')_{t'} = 2^g\overline{U}^{\mathcal{L}^2}_{t',0}(P)\overline{U}^{\mathcal{L}^2}_{t',0}(0_A)$ using (3.6)
  $Q' \leftarrow \mathcal{H}\circ\mathcal{S}(\overline{Q})$          ▷ Same but with $Q$
  $R' \leftarrow \mathcal{H}\circ\mathcal{S}(\overline{0_A})$        ▷ $(R')_{t'} = 2^g\overline{U}^{\mathcal{L}^2}_{t',0}(0_A)^2$ using (3.5)
  $R_\pm \leftarrow \mathcal{H}(P' * Q' * \mathcal{I}(R'))$   ▷ $(R_\pm)_i = 2^g\overline{\theta}^{\mathcal{L}}_i(P+Q)\overline{\theta}^{\mathcal{L}}_i(P-Q)$ using (3.7)
                                   ▷ We use: $U^{\mathcal{L}^2}_{t',2i} = \langle i|t'\rangle\, U^{\mathcal{L}^2}_{t',0}$ if $i$ is in $(\mathbb{Z}/2\mathbb{Z})^n$.
  $\overline{P+Q} \leftarrow (1/2^g)\cdot R_\pm * \mathcal{I}(\overline{P-Q})$   ▷ Divide by the coordinates of $P-Q$
  return $\overline{P+Q}$
```

3.3 Riemann relations

Beyond the duplication formula, there are more general algebraic relations that are satisfied by theta functions of level-n theta structures. In particular, theta functions

Algorithm 8 Theta doubling on a Kummer variety

```
Input: Affine level-2 theta coordinate vectors $\overline{P}, \overline{0_A}$.
Output: The point $\overline{2P} = \mathsf{dbl}(\overline{P})$.
  $P' \leftarrow \mathcal{H} \circ \mathcal{S}(\overline{P})$
  $R' \leftarrow \mathcal{H} \circ \mathcal{S}(\overline{0})$
  $R'' \leftarrow \mathcal{H}(\mathcal{S}(P') * \mathcal{I}(R'))$
  $\overline{2P} \leftarrow (1/2^g) \cdot R'' * \mathcal{I}(\overline{0})$
  return $\overline{2P}$
```

satisfy quadratic equations called the *Riemann relations*. Their name is due to the classical complex theory of theta functions. These relations give information on the arithmetic of the variety, like the duplication formula did, and let us derive more general algorithms.

Again, we suppose to have a fixed symmetric theta structure $\Theta_{\mathcal{L}}$ of level δ on A. This means $\mathcal{L}$ is totally symmetric; in particular, $2|\delta$.

Theorem 3.28 (Riemann relations)**.** *Let $P_1, P_2, P_3, P_4 \in A(k)$ be rational point and R a geometric a point satisfying $P_1 + \cdots + P_4 = 2R$. Now let $Q_i = R - P_i$ for $i = 1, \dots, 4$.*

Let $i_1, i_2, i_3, i_4 \in K_1(\delta)$, fix m such that $\sum_{l=1}^{4} i_l = 2m$, let $j_l = m - i_l$. There are rigidifications of $\mathcal{L}$ at the P_i, Q_j such that for all $t' \in (\mathbb{Z}/2\mathbb{Z})^g$ the following equality holds:

$$\begin{aligned}(3.8)\qquad &\Big(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \langle t|t'\rangle\, \overline{\theta}_{i_1+t}(P_1)\overline{\theta}_{i_2+t}(P_2)\Big)\Big(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \langle t|t'\rangle\, \overline{\theta}_{i_3+t}(P_3)\overline{\theta}_{i_4+t}(P_4)\Big) = \\ &\Big(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \langle t|t'\rangle\, \overline{\theta}_{j_1+t}(Q_1)\overline{\theta}_{j_2+t}(Q_2)\Big)\Big(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \langle t|t'\rangle\, \overline{\theta}_{j_3+t}(Q_3)\overline{\theta}_{j_4+t}(Q_4)\Big)\end{aligned}$$

Proof. The proof is a straightforward but tedious computation. The interested reader can find it in [Rob10, Theorem 4.4.6]. □

Notation 3.29. If (3.8) holds, we say that the tuple $(\overline{P_1}, \dots, \overline{P_4}, \overline{Q_1}, \dots, \overline{Q_4})$ is in Riemann position.

Remark 3.30. Let $(\overline{P_1}, \dots, \overline{P_4}, \overline{Q_1}, \dots, \overline{Q_4})$ be a tuple in Riemann position. The Riemann relations (3.8) say that, if we rescale seven out of these eight coordinate vectors by a nonzero scalar, the eighth one is uniquely determined. We can retrieve these coordinates algorithmically, as we are going to see.

Lemma 3.31. *Let $\Theta_{\mathcal{L}}$ be a symmetric theta structure on A. Its canonical theta functions satisfy*

$$[-1]\theta_i^{\mathcal{L}} = \theta_{-i}^{\mathcal{L}}.$$

Proof (sketch). Since $\Theta_{\mathcal{L}}$ is symmetric, the isogeny $[-1]\colon (A, \Theta_{\mathcal{L}}) \to (A, \Theta_{\mathcal{L}})$ is compatible with the theta structures. We can apply the isogeny theorem in coordinates, to get $[-1]^*\theta_i^{\mathcal{L}} = \lambda\theta_{-i}^{\mathcal{L}}$ for some $\lambda \in k^*$. Since $[-1]$ is an involution, we get $\lambda^2 = 1$. In [Mum66, p. 331] Mumford shows $\lambda = 1$ when $\mathcal{L}$ is very ample. □

In what follows, we work on $(A, \Theta_{\mathcal{L}})$ an abelian variety with $\Theta_{\mathcal{L}}$ symmetric.

Example 3.32. A tuple of the form $(\overline{P}, \overline{P}, \overline{-Q}, \overline{Q}; \overline{0_A}, \overline{0_A}, \overline{P+Q}, \overline{P-Q})$ is in Riemann position. The Riemann relations in this case yield exactly the duplication formula, hence the algorithms $\mathsf{diff\_add}$ and dbl.

Iterating doublings and differential additions as in the Montgomery ladder, we get:

Proposition 3.33. *Let $m \in \mathbb{Z}$ be a positive integer and $P, Q \in A(k)$ be two rational points. Fix level-n affine theta coordinates $\overline{P}, \overline{Q}, \overline{P+Q}, \overline{0_A}$. Given this data, Algorithm 9 computes affine theta coordinates for $mP+Q$ in time $O(\log m)$:*

$$\mathsf{ladder}\colon (m, \overline{P}, \overline{Q}, \overline{P+Q}, \overline{0_A}) \mapsto \overline{mP+Q}.$$

If $\overline{0_A}$ is fixed, we may omit it from the expression for compactness. Specialising with $Q = 0_A$, we get an algorithm $\mathsf{mult}\colon (m, P) \mapsto \overline{mP} = \mathsf{ladder}(m, \overline{P}, \overline{0_A}, \overline{P}, \overline{0_A})$.

Algorithm 9 Theta ladder and scalar multiplication

```
Input: Affine level-$n$ theta coordinates $\overline{P}, \overline{Q}, \overline{P+Q}, \overline{0_A}$, a positive $m = 1+\sum_{i=0}^{r} m_i 2^i$.
Output: The point $\overline{mP+Q} = \mathsf{ladder}(m, \overline{P}, \overline{Q}, \overline{P+Q}, \overline{0_A})$.
  $(R_0, R_1, R_2) \leftarrow (\overline{0_A}, \overline{P}, \overline{P+Q})$          ▷ Invariant: $R_1 = R_0 + P, R_2 = R_1 + Q$
  for $i$ in $(r, \dots, 0)$ do
    if $m_i = 0$ then                                ▷ for simplicity, omit input $\overline{0_A}$
      $(R_0, R_1, R_2) \leftarrow (\mathsf{dbl}(R_0), \mathsf{diff\_add}(R_0, R_1, \overline{P}), \mathsf{diff\_add}(R_2, R_0, \overline{P+Q}))$
    else
      $(R_0, R_1) \leftarrow (\mathsf{diff\_add}(R_1, R_0, \overline{P}), \mathsf{dbl}(R_1), \mathsf{diff\_add}(R_2, R_1, \overline{Q}))$
  return $R_2$
```

Proposition 3.34. *The tuple $(P+Q+R, P, Q, R; 0_A, Q+R, P+R, P+Q)$ is in Riemann position. It follows that fixing affine coordinates of (the null point,) the points P, Q, R and their pairwise sums determines coordinates for $P+Q+R$. This can be computed via an algorithm*

$$\mathsf{3WayAdd}\colon (\overline{P}, \overline{Q}, \overline{R}; \overline{Q+R}, \overline{P+R}, \overline{P+Q}; \overline{0_A}) \mapsto \overline{P+Q+R}.$$

of complexity $O(n^g)$, shown in Algorithm 10 for $n = 2$. Its computations descend directly from (3.8) *applied to the above Riemann tuple.*

Remark 3.35. More generally, iterating the construction with r points $P_1, \dots, P_r$, we can fix affine theta coordinates of all points P_i and pairwise sums $P_i + P_j$, and get canonically some affine theta coordinates of $P_1 + \cdots + P_r$.

Notation 3.36. Let $\overline{P} = (\overline{\theta_i}(P))_i \in k^{n^g}$ be affine coordinates of a rational point $P \in A(k)$, and $\lambda \in k^*$ a nonzero scalar. We denote by $\lambda\overline{P}$ the tuple $(\lambda\overline{\theta_i}(P))_i \in k^{n^g}$.

Algorithm 10 Theta three-way addition

```
Input: Affine level-2 theta coordinates $\overline{P}, \overline{Q}, \overline{R}, \overline{Q+R}, \overline{P+R}, \overline{P+Q}, \overline{0_A}$
Output: Affine level-2 theta coordinates $\overline{P+Q+R}$.
      ▷ Note: we assume that $\mathcal{I}$ is always applied to tuples with all nonzero entries.
  $S_1 \leftarrow \mathcal{H}(\overline{Q} * \overline{R})$
  $S_2 \leftarrow \mathcal{H}(\overline{0_A} * \overline{Q+R})$
  $S_3 \leftarrow \mathcal{H}(\overline{P+R} * \overline{P+Q})$
  $S' \leftarrow \mathcal{H}(S_2 * S_3 * \mathcal{I}(S_1))$
  $S \leftarrow (1/2^g) \cdot S' * \mathcal{I}(\overline{P})$
  return $S$
```

The algorithms we presented so far take as input some affine coordinates of points on the abelian variety. If for some input point P we rescale its affine coordinates $\overline{P}$ by some scalar $\lambda \in k^*$, the output of the algorithms changes as well. We describe this behaviour in the following lemma.

Lemma 3.37 (Homogeneity relations)**.**

(i) *Suppose* $(\overline{P_1}, \dots, \overline{P_4}, \overline{Q_1}, \dots, \overline{Q_4})$ *are in Riemann position. Then* $(\lambda_1\overline{P_1}, \dots, \lambda_4\overline{P_4}, \mu_1\overline{Q_1}, \dots, \mu_4\overline{Q_4})$ *is in Riemann position if and only if* $\prod_{i=1}^r \lambda_i = \prod_{i=1}^r \mu_i$.

(ii) $\mathsf{diff\_add}(\lambda_P\overline{P}, \lambda_Q\overline{Q}, \lambda_-\overline{P-Q}, \lambda_0\overline{0_A}) = \frac{\lambda_P^2\lambda_Q^2}{\lambda_-\lambda_0^2} \cdot \mathsf{diff\_add}(\overline{P}, \overline{Q}, \overline{P-Q}, \overline{0_A})$.

(iii) $\mathsf{3WayAdd}(\lambda_P\overline{P}, \lambda_Q\overline{Q}, \lambda_R\overline{R}, \lambda_{QR}\overline{Q+R}, \lambda_{PR}\overline{P+R}, \lambda_{PQ}\overline{P+Q}, \lambda_0\overline{0_A}) = \frac{\lambda_{QR}\lambda_{PR}\lambda_{PQ}\lambda_0}{\lambda_P\lambda_Q\lambda_R} \cdot \mathsf{3WayAdd}(\overline{P}, \overline{Q}, \overline{R}, \overline{Q+R}, \overline{P+R}, \overline{P+Q}, \overline{0_A})$

(iv) $\mathsf{ladder}(m, \lambda_P\overline{P}, \lambda_Q\overline{Q}, \lambda_+\overline{P+Q}, \lambda_0\overline{0_A}) = \frac{\lambda_+^m\lambda_P^{m(m-1)}}{\lambda_Q^{m-1}\lambda_0^{m(m-1)}} \cdot \mathsf{ladder}(m, \overline{P}, \overline{Q}, \overline{P+Q}, \overline{0_A})$.

(v) $\mathsf{mult}(m, \lambda\overline{P}) = \lambda^{m^2}\mathsf{mult}(m, P)$.

Proof. We prove some items of the lemma. The others are similar.

(i) Both sides of Equation (3.8) are linear in the coordinates of each point. If we rescale $\overline{P_i} \mapsto \lambda_i\overline{P_i}$, the left hand side of the equation gets multiplied by $\prod_{i=1}^4 \lambda_i$. Analogously, the right hand side gets multiplied by $\prod_{i=1}^4 \mu_i$ by the rescaling of the $\overline{Q_i}$. The equation is preserved if and only if these two products are equal.

(ii) This follows from (3.7) similarly: both the coordinates of P and Q are squared in the numerator and those of 0_A are squared in the denominator. To retrieve the coordinates of the $P+Q$, one must also divide by the coordinates of $P-Q$.

(iv) We prove this by induction on m. If $m = 1$, the statement is trivial. Suppose it holds for $j \leq m-1$. At step i (where i goes from r down to 0), the ladder algorithm 9 computes the points $R_0 = \overline{nP}$, $R_1 = \overline{(n+1)P}$, $R_2 = \overline{(n+1)P+Q}$ with $n = \lfloor (m-1)/2^i \rfloor$. By the inductive hypothesis, the rescaling of the input points multiplies each R_i by a scalar λ_{R_i} with

$$\lambda_{R_0} = \frac{\lambda_P^{n^2}}{\lambda_0^{n^2}}, \quad \lambda_{R_1} = \frac{\lambda_P^{(n+1)^2}}{\lambda_0^{(n+1)^2}}, \quad \lambda_{R_2} = \frac{\lambda_+^{n+1}\lambda_P^{(n+1)n}}{\lambda_Q^n\lambda_0^{(n+1)n}}.$$

In the final step, if the last digit is even we have R_0, R_1, R_2 as above with $m-1 = 2n$. The final result is computed as $R_2' = \mathsf{diff\_add}(\lambda_{R_2} R_2, \lambda_{R_0} R_0, \lambda_+ \overline{P+Q}, \lambda_0 \overline{0})$. By (ii), this equals $\mu \cdot \overline{mP+Q}$ with

$$\mu = \frac{\lambda_{R_0}^2 \lambda_{R_2}^2}{\lambda_+ \lambda_0^2} = \frac{\lambda_P^{2n^2}}{\lambda_0^{2n^2}} \frac{\lambda_+^{2(n+1)} \lambda_P^{2(n+1)n}}{\lambda_Q^{2n} \lambda_0^{2(n+1)n}} \frac{1}{\lambda_+ \lambda_0^2} = \frac{\lambda_+^{2n+1} \lambda_P^{(2n+1)(2n)}}{\lambda_Q^{2n} \lambda_0^{(2n+1)(2n)}}$$

as wanted. □

3.4 The 2-isogeny algorithm

As another immediate application of the duplication formula, we describe an algorithm to compute $(2,2)$-isogenies between principally polarised abelian *surfaces*. In Chapter 5, we will extend this to the computations of chains of $(2,2)$-isogenies of length m.

Let $(A, \mathcal{L}), (B, \mathcal{M})$ be dimension-2 PAVs equipped with line bundles of level 2. Let $f\colon A \to B$ be a $(2,2)$-isogeny of PAVs, and suppose given a level-2 symmetric theta structure $\Theta_{\mathcal{L}}$ on A such that $\ker f = K_2(\mathcal{L}) \cong (\mathbb{Z}/2\mathbb{Z})^2$, where $A[2] = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$ is the symplectic decomposition of the 2-torsion induced by the theta structure. Fix a symmetric theta structure $\Theta_{\mathcal{L}^2}$ that is compatible with $\Theta_{\mathcal{L}}$.

Via the duplication formula, we are able to describe $\Theta_{\mathcal{L}^2}$ knowing $\Theta_{\mathcal{L}}$. We saw in Example 3.13 that, if we have theta coordinates of level 4 on A, we can immediately derive a level-2 theta structure $\Theta_{\mathcal{M}}$ on the codomain B. The isogeny theorem gives the following relation between theta coordinates:

$$f^* \theta_i^{\mathcal{M}} = \theta_{2i}^{\mathcal{L}^2} \qquad \text{for all } i \in K_1(2,2) = (\mathbb{Z}/2\mathbb{Z})^2. \tag{3.9}$$

By the isogeny theorem (see also Example 3.13), if the kernel of our 2-isogeny f is $K = [2]K_2(\mathcal{L})$, then we'll have

$$K_1(\mathcal{M}) \cong K_1(\mathcal{L}), \quad K_2(\mathcal{M}) \cong K_2(\mathcal{L}^2)/K_2(\mathcal{L}). \tag{3.10}$$

Fix the notations above for the rest of the section. As discussed in Problem 1.33, we will describe:

- A *codomain* algorithm that given some information on A finds level-2 coordinates of the theta null point $0_B \in B$. By "some information", we mean that a list of generators of the kernel is sufficient; however, if we know some higher torsion points lying above the kernel generators, the algorithm becomes more efficient.
- An *evaluation* algorithm that takes as input level-2 coordinates $\overline{0_B}, \overline{P}$ for any $P \in A$ and computes level-2 coordinates of $f(P) \in B$ with respect to the theta structure found by the codomain algorithm.

The codomain algorithm

We begin by a simple observation:

Remark 3.38. Let $(\overline{\theta}_i^{\mathcal{L}}(P))_{i\in(\mathbb{Z}/2\mathbb{Z})^2}$ be affine coordinates of a rational point of A. By Equation (3.9), this induces affine coordinates on the codomain $\overline{\theta}_i^{\mathcal{M}}(f(P)) = \overline{\theta}_{2i}^{\mathcal{L}^2}(P)$. Applying the change of variable (3.2) we get $\overline{U}_{t',0}^{\mathcal{M}}(f(P)) = \overline{U}_{t',0}^{\mathcal{L}^2}(P)$. In particular, this applies to $P = 0_A, f(P) = 0_B$.

We want to compute $\overline{0_B}$ given our knowledge of A. We first show a relatively trivial algorithm: suppose we know an affine coordinate vector $\overline{0_A}$ of the theta null point of A. Without even needing the kernel generators explicitly, we can use (3.5) to get the squared values $(\overline{U}_{t',0}^{\mathcal{L}^2}(0_A)^2)_{t'} = (\overline{U}_{t',0}^{\mathcal{L}^2}(0_B)^2)_{t'}$, that is, we get the vector $(\alpha^2, \beta^2, \gamma^2, \delta^2)$ with $\mathcal{H}(\overline{0_B}) = (\alpha, \beta, \gamma, \delta)$. One of these four values is nonzero, say α. Up to rescaling, we can set $\alpha = 1$ and extract square roots of the others to get $(1, \beta, \gamma, \delta) = \overline{0_B}$ as wanted. Each square root extraction has a sign ambiguity, but all eight sign choices give isomorphic theta structures on the same variety B, by [Rob24c, Example B.3].

However, extracting square roots in finite fields is relatively expensive if compared to additions, multiplications and inversions. If we have some extra torsion information on A lying above the kernel of the isogeny, we can compute 0_B without extracting any square roots.

Theorem 3.39 (2-isogeny codomain algorithm on Kummer surfaces)**.** *Consider the setup above. Suppose known affine level-*2 *theta coordinate vectors* $\overline{0_A}, \overline{S_1}, \overline{S_2}$ *where* $S_i, i = 1, 2$ *are* 8*-torsion points satisfying* $[4]S_i = T_i$ *and such that* $\langle T_1, T_2\rangle = K_2(\mathcal{L})$ *is the kernel of the isogeny* f*. Let* $\overline{0_B}$ *be an affine lift of the theta null point of* B*, where the theta structure on* B *is induced by* $\Theta_{\mathcal{L}^2}$*. If all the components of* $\mathcal{H}(\overline{0_B})$ *are nonzero, then Algorithm 11 computes this vector, renormalised as* $R = (1, \beta, \gamma, \delta)$*, using a constant number of additions, multiplications and inversions.*

Proof of correctness. Since the point S_i is of 8-torsion and lies above a kernel point of order 2, for $i = 1, 2$ we have that $f(S_i)$ lies in $B[4]$. More precisely, $[2]S_i$ lies in $K_2(\mathcal{L}^2)$, so by (3.10) we have that $R_i = [2]f(S_i) \in K_2(\mathcal{M})$ is a point underlying the theta group of $\mathcal{M}$. By Remark 2.73, translating by these points R_i induces an action (as a sign character) on the coordinates of the points of B: if $\overline{P} = (x, y, z, w)$, then we get canonically

$$\overline{R_1 + P} = (x, -y, z, -w), \qquad \overline{R_2 + P} = (x, y, -z, -w).$$

We also know that $f(S_i) + R_i = -f(S_i)$ since $f(S_i)$ is a point of order 4. Since level-2 coordinates are invariant by point negation (Lemma 3.31, they're projective coordinates on $\mathcal{K}_A$), the coordinates of $f(S_i)$ assume the following form:

$$\begin{aligned} \overline{f(S_1)} &= (x_1, 0, z_1, 0), & \overline{f(S_2)} &= (x_2, y_2, 0, 0), \\ \mathcal{H}(\overline{f(S_1)}) &= (x', x', y', y'), & \mathcal{H}(\overline{f(S_2)}) &= (z', w', z', w') \end{aligned}$$

with $(x', y') = (x_1 + z_1, x_1 - z_1)$ and $(z', w') = (x_2 + y_2, x_2 - y_2)$.

Now, we choose a rigidification of $\mathcal{M}$ at 0_B so that we can write $\mathcal{H}(\overline{0_B}) = (1, \beta, \gamma, \delta)$. We can do this normalisation since $H(\overline{0_B})$ has all nonzero entries by assumption. Apply Equation (3.6) to the points $f(S_i)$:

$$\mathcal{H} \circ \mathcal{S}(S_1) = \left(\sum\nolimits_t \langle t|t'\rangle \, \overline{\theta}_t^{\mathcal{L}}(P)^2\right)_{t'} = \left(\overline{U}_{t',0}^{\mathcal{L}^2}(f(S_1))\overline{U}_{t',0}^{\mathcal{L}^2}(0_B)\right)_{t'} = (x', x'\beta, y'\gamma, y'\delta)$$

and similarly $\mathcal{H} \circ \mathcal{S}(S_2) = (z', w'\beta, z'\gamma, w'\delta)$. The algorithm finally recovers $1, \beta, \gamma, \delta$ from the relations between $\mathcal{H} \circ \mathcal{S}(S_1)$ and $\mathcal{H} \circ \mathcal{S}(S_2)$. □

Algorithm 11 Codomain computation of a 2-isogeny of Kummer surfaces

```
Input: Affine level-2 theta coordinates $\overline{0_A}, \overline{S_1}, \overline{S_2}$ where $\langle [4]S_1, [4]S_2 \rangle = K_2(\mathcal{L})$ is the
  kernel of an isogeny of PPAVs $f\colon A \to B$.
Output: Affine level-2 theta coordinates $(1, \beta, \gamma, \delta) = \mathcal{H}(\overline{0_B})$.
  $P_1 = (x', x'\beta, y'\gamma, y'\delta) \leftarrow \mathcal{H} \circ \mathcal{S}(\overline{S_1})$
  $P_2 = (z', w'\beta, z'\gamma, w'\delta) \leftarrow \mathcal{H} \circ \mathcal{S}(\overline{S_2})$
  $\beta \leftarrow (x'\beta)/x'$
  $\gamma \leftarrow (z'\gamma)/z'$
  $\delta \leftarrow (y'\delta) \cdot \gamma/(y'\gamma)$
  return $(1, \beta, \gamma, \delta)$.
```

The evaluation algorithm

Given the codomain algorithm, we can now describe an evaluation algorithm for the 2-isogeny f. We want to compute the affine coordinates of $f(P)$ for any $P \in A$. As in the rest of the section, fix compatible theta structures $\Theta_{\mathcal{L}}, \Theta_{\mathcal{L}^2}, \Theta_{\mathcal{M}}$, where $\mathcal{L}$ and $\mathcal{M}$ have level 2.

Theorem 3.40 (2-isogeny evaluation algorithm on Kummer surfaces)**.** *In the notations above, suppose given level-2 affine theta coordinates $\mathcal{H}(\overline{0_B}) = (1, \beta, \gamma, \delta)$ (all nonzero) and $\overline{P}$ for a rational point $P \in A(k)$. Algorithm 12 computes some level-2 affine theta coordinates $\overline{f(P)}$, in a constant number of additions, multiplications and inversions.*

Proof of correctness. Let $\overline{0_A}$ be the tuple of affine coordinates of 0_A rescaled in such a way that the induced coordinate vector $\overline{0_B}$ on the codomain satisfies $\mathcal{H}(\overline{0_B}) = (1, \beta, \gamma, \delta)$. Apply Equation (3.6) to get (up to a scalar factor that we'll incorporate in the coordinates of $f(P)$):

$$\begin{aligned}\overline{U}_{t',0}^{\mathcal{M}}(f(P))\overline{U}_{t',0}^{\mathcal{M}}(0_B) &= \overline{U}_{t',0}^{\mathcal{L}^2}(P)\overline{U}_{t',0}^{\mathcal{L}^2}(0_A) \\ &= \sum\nolimits_t \langle t|t'\rangle \, \theta_t^{\mathcal{L}^2}(P)^2 = \left(\mathcal{H} \circ \mathcal{S}(\overline{P})\right)_{t'}.\end{aligned}$$

Dividing out by the coordinates $\overline{U}_{t',0}^{\mathcal{M}}(0_B) = \mathcal{H}(\overline{0_B})_{t'}$, nonzero by assumption, we get the "twisted" coordinates $\mathcal{H}(\overline{f(P)})$. To get coordinates for $\overline{f(P)}$, apply the inverse of the Hadamard transform, that is $(1/2^g)\mathcal{H}$. We can again incorporate the scalar factor $1/2^g$ in the coordinates of $f(P)$: rescaling its coordinates still gives a valid result (it just corresponds to a different choice of rigidification of $\mathcal{M}$ at the point $f(P)$). □

Algorithm 12 Evaluation algorithm of a 2-isogeny of Kummer surfaces

```
Input: Affine level-2 theta coordinates $\overline{P}$ and $R_0 = \mathcal{H}(\overline{0_B}) = (1, \beta, \gamma, \delta)$.
Output: Affine level-2 theta coordinates $\overline{f(P)}$.
  $P' = \mathcal{H} \circ \mathcal{S}(\overline{P})$
  $Q' = P' * \mathcal{I}(R_0)$
  $\overline{f(P)} = \mathcal{H}(Q')$
  return $\overline{f(P)}$.
```

Remark 3.41 (Special cases). In our algorithms, we assumed $\mathcal{H}(\overline{0_B})$ had all nonzero coordinates. Luckily, this is usually not a problem, as it generically holds for a principally polarised abelian surface B. The only case where it does not happen is when B is a product of elliptic curves (which are rare among 2-dimensional PPAVs) with a non-product theta structure. The codomain and evaluation algorithms can be adapted to handle this special case, as done in [DMPR23a, Algorithms 8, 9].

Remark 3.42 (Generalisations). The algorithms we showed in this section work on Kummer varieties of principally polarised abelian *surfaces* (PPAVs of dimension 2). We'll see in Chapter 5 that for the cryptographic applications it is useful to have a 2-isogeny algorithm in dimension $4, 8$ as well. This has been recently done in [Dar24].

Attempts for 4-isogenies

In dimension 1, working with elliptic curve arithmetic in the Montgomery model, one can optimise chains of 2-isogenies using the fact, shown in [CH17], that 4-isogenies in the Montgomery model have a faster formula with respect to the chaining of two consecutive 2-isogenies. In practice, one performs two 2-isogenies at once. One would like to see if something similar is possible in dimension 2. We'll show a codomain algorithm that, given points of order 16 above the kernel points of a 4-isogeny (that is, two 2-isogenies chained) output the codomain of both the 2-isogenies at once. We'll see a possible application in Section 5.5.

We start from an abelian variety $(A, \mathcal{L})$ with a symmetric level-2 theta structure $\Theta_{\mathcal{L}}$, and a 4-isogeny $f\colon (A, \mathcal{L}) \to (B, \mathcal{M})$ with $\mathcal{M}$ of level 2. The goal is to find a compatible level-8 structure $\Theta_{\mathcal{L}^4}$ on A so that

$$\ker f = [2]K_2(\mathcal{L}^4) = [2]^{-1}K_2(\mathcal{L}), \qquad f^*\theta_i^{\mathcal{M}} = \theta_{4i}^{\mathcal{L}^4}.$$

By a generalisation of the duplication formula, proved by Koizumi in [Koi76], see also [Rob21, Theorem 2.7.1], one can obtain the following:

Lemma 3.43. *Let P be a point of $A(k)$, and $\overline{P}$ its affine coordinates in the theta structure $\Theta_{\mathcal{L}}$, $\overline{2P} = \mathsf{dbl}(P)$. For some scalar $\lambda \in k$, the following holds:*

$$\overline{\theta}^{\mathcal{L}}_i(2P) = \lambda \sum_{t\in\mathbb{Z}(2)} \overline{\theta}^{\mathcal{L}^4}_{i+4t}(P), \qquad i \in (\mathbb{Z}/8\mathbb{Z})^2.$$

Definition 3.44. Let $(A, \mathcal{L})$ be a PAV with an attached theta structure $\Theta_{\mathcal{L}}$, and fix $\overline{0_A}$ an affine lift of its theta null point. Let $T \in A[\ell]$ be an ℓ-torsion point. We say that an affine theta coordinate vector $\overline{T}^*$ is a *canonical ℓ-torsion lift* if $\mathsf{mult}(\ell, \overline{T}^*) = \overline{0_A}$. Note that for a general $\overline{T}$, we only know $\mathsf{mult}(\ell, \overline{T}) = \lambda\overline{0_A}$ for some nonzero λ.

Proposition 3.45. *Let T_1, T_2 be 4-torsion points of A generating $\ker f = \langle T_1, T_2\rangle = K_2(\mathcal{L}^2)$. Fix an affine lift of the theta null point $\overline{0_A}$, and suppose we're given level-2 canonical 4-torsion lifts $\overline{T_1}$, $\overline{T_2}$, $\overline{T_1+T_2}$. Then the theta null point of B is given by*

$$\overline{0_B} = \mathcal{H}^{-1}(\overline{\theta}^{\mathcal{L}}_{00}(\overline{0}), \overline{\theta}^{\mathcal{L}}_{00}(\overline{T_1}), \overline{\theta}^{\mathcal{L}}_{00}(\overline{T_2}), \overline{\theta}^{\mathcal{L}}_{00}(\overline{T_1+T_2}))$$

and the theta null point of the "intermediate" codomain $C = A/[2]\ker f$ is given by

$$\overline{0_C} = (\psi(\overline{0_A}), \psi(\overline{T_1}), \psi(\overline{T_2}), \psi(\overline{T_1+T_2})), \qquad \psi(R) = \sum\nolimits_j (\mathcal{H}\circ\mathcal{S}(R))_j.$$

Proof. Consider 8-torsion points S_1, S_2 with $[2]S_i = T_i$. We have that S_i are torsion points underlying the theta group $G(\mathcal{L}^4)$, defining an action on level-8 coordinates:

$$\begin{aligned}
(\overline{\theta}^{\mathcal{L}^4}_{i_1,i_2}(S_1+P))_{(i_1,i_2)\in(\mathbb{Z}/8\mathbb{Z})^2} &= (\zeta_8^{i_1}\overline{\theta}^{\mathcal{L}^4}_{i_1,i_2}(P))_{i_1,i_2}, \qquad \zeta_8 \text{ primitive 8-th root of 1},\\
(\overline{\theta}^{\mathcal{L}^4}_{i_1,i_2}(S_2+P))_{(i_1,i_2)\in(\mathbb{Z}/8\mathbb{Z})^2} &= (\zeta_8^{i_2}\overline{\theta}^{\mathcal{L}^4}_{i_1,i_2}(P))_{i_1,i_2}.
\end{aligned}$$

In particular, restricting our view to indices $i \equiv 0 \pmod 4$, writing

$$\overline{f(P)} = (\overline{\theta}^{\mathcal{L}^4}_i(P))_{i\in\{00,04,40,44\}} = (x, y, z, w) \tag{3.11}$$

we obtain

$$\begin{aligned}
(\overline{\theta}^{\mathcal{L}^4}_i(S_1+P))_{i\in\{00,04,40,44\}} &= (x, -y, z, -w),\\
(\overline{\theta}^{\mathcal{L}^4}_i(S_2+P))_{i\in\{00,04,40,44\}} &= (x, y, -z, -w),\\
(\overline{\theta}^{\mathcal{L}^4}_i(S_1+S_2+P))_{i\in\{00,04,40,44\}} &= (x, -y, -z, w)
\end{aligned} \tag{3.12}$$

Now set $P = 0_A$ in Lemma 3.43: if $(a, b, c, d) = (\overline{\theta}^{\mathcal{L}^4}_i(0_A))_{i\in\{00,04,40,44\}} = \overline{0_B}$, then equations (3.11) and (3.12) imply that there exist scalars $\lambda_1, \ldots, \lambda_4$ making the following hold:

$$\begin{aligned}
\overline{\theta}^{\mathcal{L}}_{00}(0_A) &= \lambda_1(a+b+c+d),\\
\overline{\theta}^{\mathcal{L}}_{00}(T_1) &= \lambda_2(a-b+c-d),\\
\overline{\theta}^{\mathcal{L}}_{00}(T_2) &= \lambda_3(a+b-c-d),\\
\overline{\theta}^{\mathcal{L}}_{00}(T_1+T_2) &= \lambda_4(a-b-c+d).
\end{aligned}$$

In [LR12] it is shown that all the λ_i are equal to 1 when the affine lifts $\overline{T_j}$ are canonical 4-torsion lifts, hence

$$\mathcal{H}(\overline{0_B}) = \mathcal{H}(a,b,c,d) = (\overline{\theta}_{00}^{\mathcal{L}}(\overline{0}), \overline{\theta}_{00}^{\mathcal{L}}(\overline{T_1}), \overline{\theta}_{00}^{\mathcal{L}}(\overline{T_2}), \overline{\theta}_{00}^{\mathcal{L}}(\overline{T_1 + T_2}))$$

as wanted.

The last part of the proposition follows from [Rob21, Corollary 2.10.9]. □

We're now left with the task of computing canonical 4-torsion lifts required as input for Proposition 3.45:

Proposition 3.46. *Let $S \in A[16]$ be a torsion point, $\overline{S}$ a corresponding affine level-2 theta coordinate vector and $g_{[8]S} \in G(\mathcal{L})$ the privileged theta group element induced by the 2-torsion point $[8]S$, inducing an action on affine level-2 coordinates. Define $R_3 = g_{[8]S} \star \mathsf{mult}(3, \overline{S}), R_5 = \mathsf{mult}(5, \overline{S})$. Both R_3 and R_5 are affine coordinate vectors for $[5]S$, so we can define their ratio as $r = (R_5)_i/(R_3)_i$ for any nonzero coordinate i. The point*

$$\overline{T}^* = r \cdot \mathsf{mult}(4, \overline{S})$$

is a canonical lift of the 4-torsion point $T = [4]S$.

Proof. Let $\overline{T} = \overline{4S} = \mathsf{mult}(4, \overline{S})$. We are looking for canonical lifts $\overline{T}^* = \mu_4 \overline{T}$ and $\overline{S}^* = \mu_{16}\overline{S}$. The final goal is to find the value of μ_4.

From the homogeneity relations (Lemma 3.37) we have $\mu_4 \cdot \overline{T} = \overline{T}^* = [4]\overline{S}^* = \mu_{16}^{4^2}[4]\overline{S} = \mu_{16}^{16}\overline{T}$, which implies $\mu_4 = \mu_{16}^{16}$.

We can compute μ_{16}^{16} as follows. Since $\overline{S}^*$ is a canonical 16-torsion lift, it satisfies $g_{[8]S} \star [3]\overline{S}^* = [11]\overline{S}^* = -[5]\overline{S}^*$. Note that level-2 coordinates are invariant by point negation. Again by homogeneity (note that the action by $g_{[8]S}$ is linear), we have:

$$\mu_{16}^{3^2} \cdot (g_{[8]S} \star (\overline{\theta}_i(\overline{3S}))_i) = \mu_{16}^{5^2} \cdot (\overline{\theta}_i(\overline{5S}))_i.$$

Taking ratios, we recover $\mu_{16}^{25-9} = \mu_{16}^{16} = \mu_4$ as desired. □

Chapter 4

Pairings

This chapter is devoted to the algorithmic computation of pairings on Jacobian varieties, with a particular focus on elliptic curves. Pairings were introduced in Section 1.3, where we presented Miller's algorithm to compute the Weil and Tate pairings on elliptic curves. Here, we will present some recent algorithms introduced in [Sta08], [LR10], [LR15], that are more efficient than the standard Miller algorithm, simpler to implement and allow for generalisation to higher dimensions. Like Miller's algorithm, these algorithms work on all abelian varieties, whereas the efficiency of most state-of-the-art implementations of pairings is limited to curves with specific parameters over specific finite fields.

These algorithms make use of two ingredients: the arithmetic algorithms in the theta model developed in Sections 3.2, 3.3, and the theory of biextensions of abelian groups (in particular, of abelian varieties) introduced in [Mum68], [Gro72].

We review some computational theory of biextensions we're going to need, then use it for pairing computations and provide an implementation in dimension 2 as example.

4.1 Biextensions

Biextensions of algebraic groups

Definition 4.1. Let H, G be algebraic groups. An *extension* of H by G is an algebraic group T fitting in an exact sequence

$$0 \longrightarrow G \longrightarrow T \longrightarrow H \longrightarrow 0. \tag{4.1}$$

Example 4.2. Let $(A, \mathcal{L})$ be a polarised abelian variety with $\mathcal{L}$ of separable type, and let $H(\mathcal{L})$ be defined as in Definition 2.42. Then the theta group $G(\mathcal{L})$ is an extension of $H(\mathcal{L})$ by $\mathbb{G}_m$.

Definition 4.3. Let A_1, A_2, G be algebraic groups. A *biextension* X of $A_1 \times A_2$ by G is an algebraic group equipped with projection maps $\pi_i \colon X \to A_i$ such that for all $a_1 \in A_1$, the fiber $\pi_1^{-1}(a_1) \subseteq X$ is an extension of A_2 by G, and similarly for all $a_2 \in A_2$ the fiber $\pi_2^{-1}(a_2)$ is an extension of A_1 by G.

For $i = 1, 2$, we say that $x \in B$ lies above $a_i \in A_i$ if $\pi_i(x) = a_i$.

Extensions are naturally linked to the concept of torsors.

Definition 4.4. Let G, A be algebraic groups. The *trivial G-torsor* over A is the algebraic group $G \times A$ equipped with the action of G on $G \times A$ given by $g \cdot (h, a) = (gh, a)$.

A *G-torsor* over A is an algebraic group T equipped with a projection map $\pi \colon T \twoheadrightarrow A$ and an action $G \times T \to T$ compatible with the projection on A (that is, $\pi(g \cdot t) = \pi(t)$ for all $g \in G, t \in T$), such that the map $G \times T \to T \times T$ given by $(g, t) \mapsto (t, g \cdot t)$ is an isomorphism, and such that T is locally trivial: for all Zariski-open $U \subseteq A$, the preimage $\pi^{-1}(U)$ is the trivial G-torsor over U.

Example 4.5. Indeed, an extension T of A by G is a G-torsor over A, where the projection map is given by the short exact sequence (4.1), and the action $G \times T \to T$ is just multiplication (seeing G as a subgroup of T). Likewise, in a biextension X of $A_1 \times A_2$ by G, the fibers $\pi_1^{-1}(a_1)$ and $\pi_2^{-1}(a_2)$ are G-torsors over A_2 and A_1 respectively.

Since the fibers of a biextension above $a_i \in A_i$ are extensions over the other group (in particular, they're groups themselves), a biextension is always equipped with two *partial group laws.*

Definition 4.6. Let X be a biextension of $A_1 \times A_2$ by G. Denote by $x_{a,b}$ a general biextension element lying above (a, b).

We define on B a *partial* group law $\star_1$ that multiplies two elements x_{a,a_2}, x_{a',a_2} of the fiber $\pi_2^{-1}(a_2)$, for some $a_2 \in A_2$, into an element $x_{a,a_2} \star_1 x_{a',a_2} = x_{aa',a_2}$ of the same fiber. Similarly, a partial group law $\star_2$ is defined, acting on fibers of the form $\pi_1^{-1}(a_1)$.

Biextensions of abelian varieties

We now specialise the construction of biextensions to the case of abelian varieties. Let A be an abelian variety, and let $D \in \operatorname{Pic}(A)$ be an ample divisor inducing a polarisation λ_D via the line bundle $\mathcal{L}_D$. For the rest of the section, suppose $\mathcal{L}_D$ is totally symmetric, so that we can work with level-2 theta coordinates.

The description of biextensions of abelian varieties comes from general facts:

Fact 4.7 ([Gro72, Remarque 2.9.6])**.** *Let A, B be abelian varieties. The set of biextensions of $A \times B$ by $\mathbb{G}_m$ is in 1-to-1 correspondence with the set of isogenies $f \colon A \to \widehat{B} = \operatorname{Pic}^0(B)$. In particular, if $A = B$, the biextensions of $A \times A$ by $\mathbb{G}_m$ are in 1-to-1 correspondence with polarisations of A.*

Before stating the next theorem, we recall Notation 2.46: for any $P \in A$ and $D \in \operatorname{Div}(A)$, we denote $D_P = t_P^*(D) - D$.

Theorem 4.8. *Let A be an abelian variety, D an ample divisor inducing a polarisation λ_D. The biextension X_D of $E \times E$ by $\mathbb{G}_m$ corresponding to λ_D can be described as* *follows. The elements of X_D are tuples $(P, Q, g_{P,Q})$, where $g_{P,Q}$ is a rational function on A with divisor*

$$\operatorname{div} g_{P,Q} = t_P^* D_Q - D_Q = D_{P+Q} - D_P - D_Q.$$

The projection maps are the natural ones: $\pi_1\colon (P, Q, g_{P,Q}) \mapsto P$, $\pi_2\colon (P, Q, g_{P,Q}) \mapsto Q$. The partial group law $\star_1$ is given by $(P_1, Q, g_{P_1,Q}) \star_1 (P_2, Q, g_{P_2,Q}) = (P_1+P_2, Q, g_{P_1+P_2,Q})$ where

$$g_{P_1+P_2,Q} = g_{P_1,Q}(\cdot) g_{P_2,Q}(\cdot + P_1) = g_{P_1,Q}(\cdot) g_{P_2,Q}(\cdot) \frac{g_{P_1,P_2}(\cdot + Q)}{g_{P_1,P_2}(\cdot)}, \tag{4.2}$$

for any g_{P_1,P_2} above (P_1, P_2). The result is independent of the choice of g_{P_1,P_2} since any other choice differs by a nonzero scalar factor that cancels out in the fraction.

The partial group law $\star_2$ is obtained symmetrically, exchanging the roles of the Ps and the Qs.

Proof. The result is proved in [Rob24a, Theorem 3.6]. □

By induction, using (4.2), one can prove the following fact, that we are going to use in the next sections.

Corollary 4.9. *Let $(A, \mathcal{L})$ be a principally polarised abelian variety, where $\mathcal{L}$ is induced by a divisor D. Let $P \in A$ be a point, $m \in \mathbb{N}$ an integer. One inductively shows*

$$g_{mP,Q} = g_{P,Q}^{\star_1, m} = g_{P,Q}^m(\cdot) \frac{f_{m,P}(\cdot + Q)}{f_{m,P}(\cdot)}$$

where $f_{m,P} = g_{P,P} \cdot g_{2P,P} \cdot \dots \cdot g_{(m-1)P,P}$ is a function with divisor

$$\operatorname{div} f_{m,P} = D_{mP} - mD_P.$$

This fact will be useful for pairing computations: what we are going to present will be generalisations of Miller's algorithm for elliptic curves (see Proposition 1.49). Note, indeed, that if E is an elliptic curve, $D = (0_E)$ gives the principal polarisation and $P \in E$ is a point of ℓ-torsion, then we have $\operatorname{div} f_{\ell,P} = D_{\ell P} - \ell D_P = \ell((0_E) - (-P)) \sim \ell((P) - (0_E))$. Up to sign conventions in the defintions, $f_{\ell,P}$ is a Miller function for P.

4.2 Rigidifications and cubical representations

Cubical representations

We are now going to represent biextension elements in a more concrete way, using the affine coordinates we discussed in Section 3.2. More precisely, once we fix an ample totally symmetric divisor D with the associated line bundle $\mathcal{L}$ and biextension X_D, we also fix a theta structure $\Theta_{\mathcal{L}}$ of level n. Fixing a rigidification of $\mathcal{L}$ at P is equivalent to fixing a tuple of affine theta coordinates $\overline{P}$.

Recall that a biextension element $g_{P,Q}$ has divisor $t^*_{P+Q}D + D - t^*_P D - t^*_Q D = D_{P+Q} - D_P - D_Q$. The invertible sheaf

$$\mathcal{L}_{\mathrm{sq}} = t^*_{P+Q}\mathcal{L} \otimes \mathcal{L} \otimes t^*_P\mathcal{L}^{-1} \otimes t^*_Q\mathcal{L}^{-1}$$

is isomorphic to the trivial sheaf $\mathcal{O}_A$, by the theorem of the square (Theorem 2.24). If $\mathcal{L} = \mathcal{L}_D$, then the above sheaf is canonically isomorphic to $\mathcal{L}_{D_{P+Q}-D_P-D_Q}$

Fix a rigidification at 0_A of the line bundle $\mathcal{L}_{\mathrm{sq}}$, or equivalently, fix affine coordinates $\overline{0_A}, \overline{P}, \overline{Q}, \overline{P+Q}$. This induces a local trivialisation around 0_A, but since the whole sheaf is trivial it extends to a global trivialisation. In particular, a unique biextension element is induced by this trivialisation, that is, the rational function $g_{P,Q}$ of divisor $D_{P+Q} - D_P - D_Q$ corresponding to 1 via the isomorphism.

Notation 4.10. We denote the biextension element induced by the affine coordinate vectors $\overline{0_A}, \overline{P}, \overline{Q}, \overline{P+Q}$ as $g_{P,Q} = [\overline{0_A}, \overline{P}, \overline{Q}, \overline{P+Q}]$. This is called a *cubical point* in [Rob24a], or a *cubical representation* of $g_{P,Q}$, because the properties of the torsors attached to our biextension come from the Theorem of the Cube [Bre83].

Remark 4.11. The cubical representation is redundant: the same $g_{P,Q}$ is induced by other quadruples

$$g_{P,Q} = [\lambda_0\overline{0_A}, \lambda_P\overline{P}, \lambda_Q\overline{Q}, \lambda_+\overline{P+Q}], \qquad \lambda_0\lambda_+ = \lambda_P\lambda_Q.$$

Indeed, if we rescale the rigidifications at $0, P, Q, P+Q$ by scalars as above, the condition $\lambda_0\lambda_+ = \lambda_P\lambda_Q$ ensures that these scalars cancel each other out in the isomorphism $\mathcal{L}_{P+Q} \otimes \mathcal{L}_{0_A} \otimes \mathcal{L}_P^{-1} \otimes \mathcal{L}_Q^{-1} \cong \mathcal{O}_A$, so the isomorphism stays the same and so does $g_{P,Q}$.

The cubical representation works well with the partial group laws of the biextension:

Proposition 4.12. *Consider a biextension on a polarised abelian variety A, and let $g_{P_1,Q}, g_{P_2,Q}$ be two biextension elements respectively represented by $[\overline{0}, \overline{P_1}, \overline{Q}, \overline{P_1+Q}]$ and $[\overline{0}, \overline{P_2}, \overline{Q}, \overline{P_2+Q}]$. Let $\overline{P_1+P_2}$ be any affine coordinate vector of the point P_1+P_2, and let $\overline{P_1+P_2+Q} = \mathsf{3WayAdd}(\overline{P_1}, \overline{P_2}, \overline{Q}, \overline{P_2+Q}, \overline{P_1+Q}, \overline{P_1+P_2})$.*

A cubical representation of their $\star_1$-product $g_{P_1,Q} \star_1 g_{P_2,Q} = g_{P_1+P_2,Q}$ is given by

$$g_{P_1+P_2,Q} = [\overline{0}, \overline{P_1+P_2}, \overline{Q}, \overline{P_1+P_2+Q}].$$

For a given integer m and a point $P \in A$, let $\overline{mP} = \mathsf{mult}(m, P)$ and $\overline{mP+Q} = \mathsf{ladder}(m, \overline{P}, \overline{Q}, \overline{P+Q}, \overline{0_A})$. The expression

$$g^{\star_1,m}_{P,Q} = g_{mP,Q} = [\overline{0}, \overline{mP}, \overline{Q}, \overline{mP+Q}]$$

is a cubical representation of the exponentiation of $g_{P,Q}$ in the partial group law $\star_1$.

Proof. The result is proved in [Rob24a, Theorem 4.16]. □

We can now see how to *evaluate* a biextension function via the cubical representation.

Proposition 4.13. *Let $(A, \mathcal{L})$ be a polarised abelian variety with $\mathcal{L} = \mathcal{L}_D$ symmetric. Let $X \in \Gamma(A, \mathcal{L})$ be a section with zero divisor D. Let $g_{P,Q}$ be a biextension element represented by $[\overline{0_A}, \overline{P}, \overline{Q}, \overline{P+Q}]$. For any point $R \in A$ such that does not lie in the support of $D + \operatorname{div} g_{P,Q}$, fix affine coordinates $\overline{P+R}, \overline{Q+R}$ and compute $\overline{P+Q+R} =$ $\mathsf{3WayAdd}(\overline{P}, \overline{Q}, \overline{R}, \overline{Q+R}, \overline{P+R}, \overline{P+Q})$. The evaluation of $g_{P,Q}$ at R is given by*

$$g_{P,Q}(R) = \frac{\overline{X}(R)\overline{X}(P+Q+R)}{\overline{X}(P+R)\overline{X}(Q+R)}.$$

Proof. This is proved in [Rob24a, Theorem 4.28]. □

When D is an effective divisor, there is indeed a section X as in the proposition. For example, if $A = E$ is an elliptic curve and D is $2(0_E)$ or $3(0_E)$, then we can take X to be the projective coordinate Z.

Finally, we remember that the theta group $G(D) = G(\mathcal{L})$ acts on sections of $\mathcal{L}$, hence on the affine coordinates giving the cubical representation. We state how this action works:

Proposition 4.14 ([Rob24a, Lemma 4.18])**.** *Let $(A, \mathcal{L})$ be a PAV with $\mathcal{L} = \mathcal{L}_D$ symmetric. Let $T \in H(D)$ be an element underlying the theta group. For all $Q \in A$, the point T induces a privileged biextension element $g_{T,Q}$. More precisely, the map*

$$s_Q \colon T \mapsto g_{T,Q} = g_T(\cdot + Q)/g_T, \qquad (T, g_T) \in G(D) \text{ any element above } T$$

is a group homomorphism, and

$$T \star g_{P,Q} = s_Q(T) \star_1 g_{P,Q}$$

is a group action of $H(D)$ on the biextension.

If $g_{P,Q}$ is represented by $[\overline{0_A}, \overline{P}, \overline{Q}, \overline{P+Q}]$, then we have

$$T \star g_{P,Q} = [\overline{0_A}, g_T \star \overline{P}, \overline{Q}, g_T \star \overline{P+Q}] = [\overline{0_A}, \overline{T+P}, \overline{Q}, \overline{T+P+Q}].$$

The expressions above don't depend on the choice of g_T by Remark 4.11: a different choice of g_T would rescale the second and fourth point by the same scalar.

4.3 Pairing computations

In this section, we are finally going to see how to compute pairings using biextensions and the arithmetic of cubical representations.

We refer the reader back to Section 1.3 for the definitions of the Weil and Tate pairings on elliptic curves. We can generalise definitions as follows for a g-dimensional PAV A.[1] As in Definition 2.42, if D is an ample divisor on A we denote by $H(D)$ the kernel of the polarisation λ_D induced by D.

[1] The following is not the standard treatment of the Weil and Tate pairings. Its equivalence to our definitions is due to [Rob23b]

Definition 4.15 (Tate, Weil pairing on PAVs)**.** Let $(A, \mathcal{L}_D)$ be a PAV over k.

- Consider a point $P \in H(\ell D)$. Let P' be a formal sum of points $P' = \sum_i n_i(P_i)$ that satisfies $\sum_i n_i = 0$, $\sum_i n_i P_i = P$. For example, one can take $P' = (P) - (0)$. A *Miller function* for ℓ, P, D is a rational function f_{ℓ, D'_P} on A with divisor $\ell D_{P'} := \ell \sum_i n_i t^*_{P_i} D$. By the theorem of the square, $D_{P'}$ is linearly equivalent to D_P, so $\ell D_{P'} \sim \ell D_P$ is indeed principal. In fact, P lies in the kernel of $\lambda_{\ell D}$ by assumption and $\ell D_P = \lambda_{\ell D}(P) = 0$.
- Let Q be a rational point in $A(k)$. Consider a formal sum of points Q' as above that does not intersect the support of $D_{P'}$. The D-Tate pairing is a nondegenerate bilinear map
$$e_{T,D,\ell} \colon H(\ell D)(k) \times A(k)/[\ell]A(k) \to k^*/(k^*)^\ell$$
that can be computed as
$$e_{T,D,\ell}(P, Q) = f_{\ell, D_{P'}}(Q') = \prod\nolimits_i f_{\ell, D_{P'}}(Q_i)^{n_i} \pmod{(k^*)^\ell}.$$
- Let $P, Q \in H(\ell D)$. Consider two formal sums of points P', Q' as above having disjoint supports. The D-Weil pairing is the nondegenerate bilinear map
$$e_{W,D,\ell} \colon H(\ell D) \times H(\ell D) \to \mu_\ell, \qquad (P, Q) \mapsto f_{\ell, D_{P'}}(Q')/f_{\ell, D_{Q'}}(P').$$
Note that while the definition of the Tate pairing depends on k (as both its domain and codomain do), the Weil pairing is geometric: it is defined over geometric points, and its output lies in k^{alg}.
One can show that the expressions for $e_{T,D,\ell}$ and $e_{W,D,\ell}$ are independent of the choice of P', Q' by Weil reciprocity [Lan83, §VI, Proposition 4].

These definitions are more general than the ones given for elliptic curves, since the divisor D comes into play.

Remark 4.16. On an elliptic curve E, if we take the divisor $D = (0_E)$, then the generalised Weil pairing $e_{W,D,\ell}$ coincides with the standard Weil pairing $e_{W,\ell}$ (likewise, $e_{T,D,\ell}$ is $e_{T,\ell}$). This happens in general for principally polarised abelian varieties, for which the standard Weil pairing $e_{W,\ell}$ was defined in Theorem 2.32.

Lemma 4.17 ([Rob21, p. 56])**.** *Let D_0 be an ample divisor on an abelian variety A. If $D = mD_0$ for some positive integer m, then we have*
$$e_{T,D,\ell} = e^m_{T,D_0,\ell}, \qquad e_{W,D,\ell} = e^m_{W,D_0,\ell}.$$

Odd-order pairing computations

Having defined our pairings, we can now look at how to compute them.

Lemma 4.18. *Let ℓ be an integer, $P, Q \in A$ two points, such that P is a point of ℓ-torsion. If $g_{P,Q}$ is a biextension element above (P, Q), then its $\star_1$-exponentiation is given by*
$$g_{\ell P, Q} = g_{P,Q}(\cdot)^\ell \cdot \frac{f_{\ell,D,P}(\cdot)}{f_{\ell,D,P}(\cdot + Q)}$$

where $f_{\ell,D,P} = (g_{P,P} \cdots g_{(\ell-1)P,P})^{-1}$ is a Miller function for ℓ, P, D.

Proof. The result follows immediately from Corollary 4.9, observing that the divisor of $f_{\ell,D,P}^{-1}$ is $D_{\ell P} - \ell D_P = -\ell D_P$. □

Throwing in cubical representations and theta functions, we get an algorithm for order-ℓ pairings, when ℓ is odd.

Theorem 4.19. *(Squared Tate pairing on PPAVs) Let $(A, \mathcal{L}_0)$ be a PPAV, equipped with $\mathcal{L} = \mathcal{L}_0^2$ totally symmetric of level 2 and a symmetric theta structure $\Theta_{\mathcal{L}}$. Let ℓ be a positive integer, $P \in A[\ell](k)$ a k-rational torsion point, $Q \in A(k)$ another rational point, $g_{P,Q}$ a rational biextension element. Define $g_{\ell P,Q} = g_{P,Q}^{\star_1,\ell}$.*

The evaluation of its inverse $g_{\ell P,Q}(R)^{-1}$ at any point R, modulo ℓ-th powers in k^*, equals the squared Tate Pairing $e_{T,\ell}(P,Q)^2$ between points P and $Q \mod [\ell]A(k)$.

In particular, if we fix a cubical representation $\overline{0_A}, \overline{P}, \overline{Q}, \overline{P+Q}$ made of affine level-2 coordinates (using theta functions on $\mathcal{L}$) and we define

$$\overline{\ell P} = \mathsf{mult}(\ell, \overline{P}), \qquad \overline{\ell P + Q} = \mathsf{ladder}(\ell, \overline{P}, \overline{Q}, \overline{P+Q})$$

then we can compute

$$g_{\ell P,Q}(0)^{-1} = \frac{\overline{\theta_i}(Q)}{\overline{\theta_i}(\ell P + Q)} \frac{\overline{\theta_j}(\ell P)}{\overline{\theta_j}(0_A)},$$

where i, j are such that θ_i and θ_j don't vanish at Q and 0_A respectively.

Before the proof, we note that, even if the theorem only allows us to compute the squared pairing, when ℓ is odd it is easy to then recover the non-squared pairing $e_{T,\ell}(P,Q)$: if $r = 1/2 \pmod{\ell}$, we can compute $(e_{T,\ell}(P,Q)^2)^r \equiv e_{T,\ell}(P,Q) \pmod{k^\ell}$.

Proof of Theorem 4.19. Let $D_0, D = 2D_0$ the divisors corresponding respectively to $\mathcal{L}_0, \mathcal{L}$.

First of all, we observe that the divisor of $g_{\ell P,Q}$ is $D_{\ell P+Q} - D_{\ell P} - D_Q$, which equals 0 since $\ell P = 0_A$. That is, $g_{\ell P,Q}$ is a constant, and so is its inverse. By Lemma 4.18, up to ℓ-th powers its value is $g_{\ell P,Q}(0)^{-1} \equiv f_{\ell,D,P}((Q) - (0_A))$ where $f_{\ell,D,P}$ is a Miller function[2] for ℓ, P, D. By definition, the Tate pairing is exactly $e_{T,D,\ell}(P,Q) = f_{\ell,D,P}((Q) - (0_A))$. Finally, Lemma 4.17 yields that $e_{T,D,\ell} = e_{T,D_0,\ell}^2 = e_{T,\ell}^2$ is the squared pairing as desired.

The expression involving the cubical representation comes from Propositions 4.12 and 4.13: we have indeed

$$g_{\ell P,Q}(0) = \frac{\overline{X}(\ell P + Q)\overline{X}(0_A)}{\overline{X}(Q)\overline{X}(\ell P)}$$

for some projective coordinate $X \in \Gamma(A, \mathcal{L})$ that has D as a divisor of zeroes. Note however that ℓP and 0 are the same point on the abelian variety (and so are $\ell P + Q$

[2] Technically, to avoid zeroes and poles for the sake of well-definedness, $(Q) - (0_A)$ could need to be replaced by some equivalent combination of points Q', as in Definition 4.15. However, we say right below how to make $g_{\ell P,Q}(0)$ always well-defined.

and Q), and their affine coordinate vectors are a rescaling of one another: they satisfy $\overline{\ell P} = \lambda \overline{0_A}$ for some λ, so that we have

$$\overline{\theta_i}(\ell P)/\overline{\theta_i}(0_A) = \lambda = \overline{X}(\ell P)/\overline{X}(0_A)$$

for all i. Choosing i that makes the expression well-defined, the conclusion follows. □

A similar discussion can be carried out for the Weil pairing, giving the following

Corollary 4.20. *Consider a PPAV $(A, \mathcal{L}_0)$ and $\mathcal{L}, D, \Theta_{\mathcal{L}}$ as above. Let $P, Q \in A[\ell]$ be ℓ-torsion points. The ratio of biextension functions $g_{P,\ell Q}(R)/g_{\ell P,Q}(R)$ evaluated at any point R gives the squared Weil pairing $e_{W,\ell}(P,Q)^2$ and can be computed as a ratio of theta coordinates*

$$e_{W,\ell}(P,Q)^2 = \frac{\overline{\theta_i}(\ell Q + P)\overline{\theta_j}(0_A)}{\overline{\theta_i}(P)\overline{\theta_j}(\ell Q)} \frac{\overline{\theta_{i'}}(Q)\overline{\theta_{j'}}(\ell P)}{\overline{\theta_{i'}}(\ell P + Q)\overline{\theta_{j'}}(0_A)}$$

If ℓ is odd, we can easily recover the non-squared pairing $e_{W,\ell}(P,Q) = e_{W,\ell}(P,Q)^{2r}$, with $r = 2^{-1} \mod \ell$.

Even-order pairing computations

The above theorems give fast recipes to compute the squared Tate and Weil pairing, and the non-squared ones when the integer ℓ is odd. When instead $\ell = 2m$ is even, some more work is needed.

Let $D = 2D_0$ be a totally symmetric ample divisor of degree 2, like before. Let $P \in A[\ell](k)$ be a rational point of ℓ-torsion and $Q \in A(k)$ another rational point. Note that we have $A[\ell] = H(\ell D_0) = H(mD)$. Using this, we will compute

$$\begin{aligned} e_{T,\ell}(P,Q) &= e_{T,D_0,2m}(P,Q) \\ \text{(by compatibility of the Tate pairing with integers)} \quad &= e_{T,D_0,m}(P,Q)^2 \\ \text{(by Lemma 4.17)} \quad &= e_{T,D,m}(P,Q), \end{aligned}$$

where the last term makes sense since P lies in the kernel $H(mD)$, and can be computed via the cubical arithmetic, using coordinates that are sections of $\mathcal{L}_D$.

We have an analogue of Lemma 4.18 to find a Miller function in this case:

Lemma 4.21. *Let m be a positive integer, D a divisor on A, and $P, Q \in A$ two points with $P \in H(mD)$. Let $g_{P,Q}$ be a biextension element above (P,Q) and $g_{mP} \in G(mD)$ an element of the theta group lying above P (that is, a function with divisor D_{mP}). If we write*

$$g_{mP,Q} = g_{P,Q}^{\star_1 m} = g_{P,Q}(\cdot)^m \frac{f'_{m,D,P}(\cdot)}{f'_{m,D,P}(\cdot + Q)}, \qquad f'_{m,D,P} = (g_{P,P} \cdots g_{(m-1)P,P})^{-1},$$

then the function $f'_{m,D,P} g_{mP}$ is a Miller function for m, D, P.

Proof. The product $f'_{m,D,P}g_{mP}$ has the required divisor: indeed $\operatorname{div} f'_{m,D,P}g_{mP} = (mD_P - D_{mP}) + D_{mP} = mD_P$. □

Using this, we can compute nonsquared Tate pairings also in the case of even ℓ:

Theorem 4.22 (Nonsquared Tate pairing on PPAVs)**.** *Let $(A, \mathcal{L}_0)$ be a PPAV, equipped with $\mathcal{L} = \mathcal{L}_0^2$ totally symmetric of level 2 and a symmetric theta structure $\Theta_{\mathcal{L}}$. Denote by D_0, D the divisors corresponding to $\mathcal{L}_0, \mathcal{L}$ respectively. Let $\ell = 2m$ be a positive integer, $P \in A[\ell]$ a torsion point, $Q \in A$ another point.*

If a biextension function $g_{P,Q} \in X_D$ is the tensor square of a biextension function $g'_{P,Q}$ in X_{D_0}, then the function $g_{mP,Q}^{-1}\frac{g_{-mP}(\cdot)}{g_{-mP}(\cdot+Q)}$ is a constant and equals the Tate pairing $e_{T,\ell}(P,Q)$, with $g_{mP,Q} = g_{P,Q}^{\star 1 m}$ and $g_{-mP} \in G(mD)$ any function above $-mP$ (since we take ratios, the constant is independent of the choice of g_{-mP}).

If we fix a cubical representation $\overline{0_A}, \overline{P}, \overline{Q}, \overline{P+Q}$ made of affine level-2 coordinates (using theta functions on $\mathcal{L}$) for $g_{P,Q}$, and we define

$$\overline{mP} = \mathsf{mult}(m, \overline{P}), \qquad \overline{mP+Q} = \mathsf{ladder}(m, \overline{P}, \overline{Q}, \overline{P+Q})$$

then we can compute the pairing using the action of the theta group on global sections:

$$e_{T,\ell}(P,Q) = \frac{\overline{\theta_i}(Q)}{(g_{-mP} \star \overline{\theta_i})(mP+Q)} \frac{(g_{-mP} \star \overline{\theta_j})(mP)}{\overline{\theta_j}(0_A)} \tag{4.3}$$

where i, j are such that θ_i and θ_j don't vanish at Q and 0_A respectively.

On input $m, \overline{0_A}, \overline{P}, \overline{Q}, \overline{P+Q}$, Algorithm 13 computes the Tate pairing $e_{T,\ell}(P,Q)$ in time $O(\log \ell)$ using the above formulas.

Proof. The fact that the function $h = g_{mP,Q}^{-1}\frac{g_{-mP}(\cdot)}{g_{-mP}(\cdot+Q)}$ is constant can be seen by looking at its divisor. To see that this constant is exactly the Tate pairing, we observe that this function, evaluated at 0_A, is equal to $s((Q) - (0_A))g_{P,Q}^m$, where $s = f'_{m,D,P}g_{-mP}^{-1}$, with the notations of Lemma 4.21, is a Miller function, and we can ignore $g_{P,Q}^m$ because by assumption $g_{P,Q}$ is already the square of a rational element, so $g_{P,Q}^m$ is the ℓ-th power of a rational element.

Finally, to make h computable in practice, we show it's a biextension element and give a cubical representation:

$$h^{-1} = g_{mP,Q}\frac{g_{-mP}(\cdot+Q)}{g_{-mP}(\cdot)} = (-mP) \star g_{mP,Q}.$$

If $g_{mP,Q}$ is represented by $[\overline{0_A}, \overline{mP}, \overline{Q}, \overline{mP+Q}]$, then $(-mP) \star g_{mP,Q}$ is represented by $[\overline{0_A}, (-mP) \star \overline{mP}, \overline{Q}, (-mP) \star \overline{mP+Q}]$, which gives exactly Equation (4.3) when evaluated at 0_A and inversed. □

The Weil pairing is computed similarly, and is shown in Algorithm 14.

Algorithm 13 Tate pairing using cubical level-2 representation

```
Input: Affine level-2 coordinates $\overline{0_A}, \overline{P}, \overline{Q}, \overline{P+Q}$ with $P \in A[\ell](k)$, $Q \in A(k)$, a positive integer $\ell$
Output: A representative of the Tate pairing $e_{T,\ell}(P,Q)$ in $k$
  if $\ell$ is odd then $n \leftarrow \ell$ else $n \leftarrow \ell/2$
  $R \leftarrow \mathsf{mult}(n, \overline{P})$                              ▷ $R = \overline{nP}$
  $R' \leftarrow \mathsf{ladder}(n, \overline{P}, \overline{Q}, \overline{P+Q}, \overline{0_A})$          ▷ $R' = \overline{nP+Q}$
  if $\ell$ is even then    ▷ Compute the theta action by the 2-torsion point $-nP = nP$
      $R \leftarrow (nP) \star R$                                    ▷ $R = (nP) \star \overline{nP}$
      $R' \leftarrow (nP) \star R'$                                 ▷ $R' = (nP) \star \overline{nP+Q}$
  $\lambda_1 = (\overline{Q})_i / R'_i$ for the first $i$ such that $R'_i \neq 0$
  $\lambda_2 = (\overline{0_A})_i / R_i$ for the first $i$ such that $R_i \neq 0$
  $e = \lambda_1 / \lambda_2$
  if $\ell$ is odd then
      $e \leftarrow e^{(\ell+1)/2}$
  return $e$
```

Algorithm 14 Weil pairing using cubical level-2 representation

```
Input: Affine level-2 coordinates $\overline{0_A}, \overline{P}, \overline{Q}, \overline{P+Q}$ with $P, Q \in A[\ell]$, an integer $\ell \leq 0$
Output: The Weil pairing $e_{W,\ell}(P,Q)$
  $e_{PQ} \leftarrow e_{T,\ell}(P,Q)$ as computed in Algorithm 13
  $e_{QP} \leftarrow e_{T,\ell}(Q,P)$ as computed in Algorithm 13
  return $e_{PQ}/e_{QP}$
```

Implementation and remarks

The above pairing algorithms have been implemented by Robert in Magma [BCP97] for hyperelliptic Jacobians in the AVIsogenies library [BCR11] and in SageMath [S$^+$24], in the case of elliptic curves (both in the Montgomery model and the theta model), in [Rob23c]. We combined the latter with the SageMath code of [DMPR23b], which provides an interface for theta Kummer surfaces, in order to give a proof-of-concept computations the biextension pairing algorithms on hyperelliptic genus-2 Jacobians, available at `https://github.com/sferl/theta-pairings-dim2`. SageMath currently doesn't have an interface like AVIsogenies to work with general theta PAVs – Tate and Weil pairing algorithms are even lacking in dimension 2, though public efficient implementations of Miller's algorithm on genus-2 Jacobians exist (see [AFK24] among the latest ones). As our proof of concept shows, if such an interface did exist, [Rob23c] could be easily adapted to it, becoming an efficient and general alternative to Miller's algorithm.

A point worth remarking, from the implementor's side, is that in the elliptic curve case biextension pairings only need Montgomery ladders and a few edits on differential additions and doublings, that is just a few edits close to what any library on elliptic curve computation already has.

Chapter 5

Applications to cryptography: SQIsign2D

5.1 Cryptography

The main motivation for the algorithms on abelian variety we went through in the last chapters comes from isogeny-based cryptography.

Cryptography is the study of secure communication between two parties. Its protocols guarantee authenticity, integrity and secrecy of data sent across insecure channels. It is central in today's internet communication, where it is used to secure transactions, messages and personal data.

Public-key cryptography is a type of cryptography where each party owning a private key, which is kept secret, also possesses a corresponding public key, which can be shared with anyone. One can build several types of protocols through this system: e.g., digital signatures, key exchanges and encryption. We'll focus in particular on digital signatures, where the owner of a private key uses it to sign a message, and anyone else can verify the signature using the signer's public key. The security of the system relies on the fact that it is computationally infeasible to derive the private key from the public key.

Diffie–Hellman Key exchange

The simplest example of public-key cryptography is the Diffie–Hellman key exchange protocol [DH76]. In this protocol, the two parties, customarily called Alice and Bob, agree on a large prime number p and a generator g of the multiplicative group $\mathbb{F}_p^\times$, and execute the following protocol (see also Figure 5.1) to get a shared secret in the end, secret to everyone but Alice and Bob.

- Alice chooses a secret integer a and sends $g^a \mod p$ to Bob.
- Bob chooses a secret integer b and sends $g^b \mod p$ to Alice.
- Both parties can then compute the shared secret $g^{ab} \mod p$.

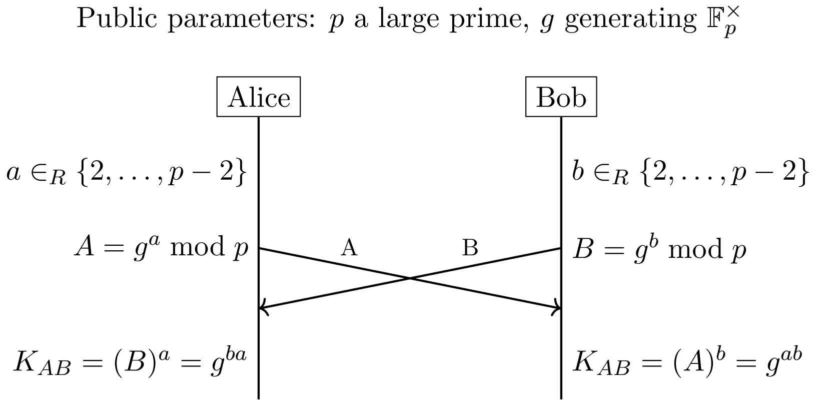
Public parameters: p a large prime, g generating $\mathbb{F}_p^\times$
Alice
Bob
$a \in_R \{2, \dots, p-2\}$
$b \in_R \{2, \dots, p-2\}$
$A = g^a \bmod p$
A
B
$B = g^b \bmod p$
$K_{AB} = (B)^a = g^{ba}$
$K_{AB} = (A)^b = g^{ab}$

Figure 5.1: The Diffie–Hellman key exchange.

Here, g^a (resp. g^b) $\bmod p$ is the public key of Alice (resp. Bob), and a (resp. b) is the private key. The security of the protocol relies on the difficulty of computing discrete logarithms $g^a \mapsto a$ in $\mathbb{F}_p^\times$.

Remark 5.1. In the DH protocol, $\mathbb{F}_p^\times$ can be replaced by any cyclic group where discrete logarithms are hard. Indeed, a variant of the Diffie–Hellman protocol is the *Elliptic Curve Diffie–Hellman* (ECDH) protocol, where the group $\mathbb{F}_p^\times$ is replaced by the group of points of an elliptic curve E over a finite field $\mathbb{F}_p$.

The Diffie–Hellman protocol (as well as its elliptic curve variant) is secure against classical computers, but it is vulnerable to quantum computers. In 1994, Peter Shor showed that a quantum computer could factor large integers and solve the discrete logarithm problem over finite abelian groups in polynomial time [Sho94]. This means that, if a large enough quantum computer is built, it could break the security of the Diffie–Hellman protocol and most other public-key cryptosystems currently in use. This is a major concern for the security of the internet, and that's why the cryptographic community is now looking for new *post-quantum* cryptographic protocols that would not break in polynomial time under quantum attacks.

The American National Institute of Standards and Technology (NIST), the de-facto worldwide reference for standardisation in cryptography, has launched in 2016 the Post-Quantum Cryptography Standardization Project [NIS16], a competition aimed at indentifying and standardising quantum-secure cryptographic protocols that run on classical computers. Ideally, they'd be drop-in replacements of the (pre-quantum) protocols currently in use, so that when quantum computers become a reality, a smooth transition to quantum-secure protocols will have already happened. The first round of standards has been just released[1] (August 2024), but the search for post-quantum protocols is still ongoing and active, looking for diversity in the mathematical constructions underlying the protocols and for advanced functionalities.

[1] `https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards`, accessed: September 2024

5.2 Isogeny-based cryptography

Isogeny-based cryptography is a branch of post-quantum cryptography whose constructions are based on walks in the graph of isogenies between elliptic curves. The study of isogenies is classical in algebraic geometry, but their algorithmic investigation is relatively recent. Work on isogeny graphs can be dated back to the late 1980s [Mes86] (see also the work on isogeny volcanoes in [Koh96]). The first applications in cryptography came in the early 2000s with works by Teske [Tes06], Charles, Goren and Lauter [CLG09], Couveignes [Cou06] and Rostovtsev and Stolbunov [RS06]. Since then, many other protocols have been proposed: key exchanges (such as SIDH [JD11], CSIDH [CLM$^+$18]), encryption schemes (such as FESTA [BMP23]), digital signatures (we'll examine SQISign [DKL$^+$20] in the next sections), and frameworks providing more advanced features [Bas24].

Hard problems

The security of isogeny-based cryptography relies on the hardness of the

Problem 5.2 (Isogeny problem)**.** Given two elliptic curves $E, E'/\mathbb{F}_q$ isogenous to each other over $\mathbb{F}_q$, compute a possible $\varphi\colon E \to E'$ over $\mathbb{F}_q$ (in the sense of Problem 1.32).

Solving the isogeny problem is hard for classical computers: the best known algorithms have exponential complexity in $\log q, \log \#E$ (see [Gal24] for the most recent advances). The isogeny problem is believed to be hard even for quantum computers: see [GV18] for a discussion.

Remark 5.3. The corresponding decisional problem, i.e., deciding whether two elliptic curves $E, E'/\mathbb{F}_q$ are $\mathbb{F}_q$-isogenous, is instead easy: it is equivalent to determining whether $E(\mathbb{F}_q)$ and $E'(\mathbb{F}_q)$ have the same cardinality, which can be done in polynomial time with respect to $\log q$ using Schoof's algorithm [Sch95].

The class of *supersingular* elliptic curves (see Definition 1.13) has some nice properties that make them particularly suitable to work with both for the construction of cryptographic protocols and for their security. In fact, supersingular isogeny graphs don't have any structure that current quantum attacks can exploit.[2] The following facts hold:

Proposition 5.4. *Let E be a supersingular elliptic curve over a finite field $\mathbb{F}_q$ of characteristic p.*

(i) The endomorphism ring of E is an order in the quaternion algebra $B_{p,\infty}$.

(ii) The j-invariant $j(E)$ belongs to $\mathbb{F}_{p^2}$. Therefore, up to isomorphism, E can be assumed to be defined over $\mathbb{F}_{p^2}$.

[2] Ordinary isogeny graphs, as the one in [RS06], admit the action of an abelian group, the class group, exploitable by a sub-exponential quantum algorithm by Kuperberg [Kup05], while the supersingular graph admits no such group action.

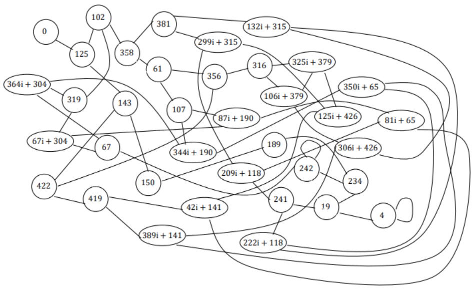

Figure 5.2: The 2-isogeny graph over $\mathbb{F}_{431^2}$ (reproduced from [Cos19])

(iii) The number of $\mathbb{F}_p^{\text{alg}}$*-isomorphism classes of supersingular elliptic curves is given by* $\lfloor p/12 \rfloor + \delta$*, with* $\delta \in \{0, 1, 2\}$*.*

(iv) Suppose p *is greater than* 3 *and* E *is defined over* $\mathbb{F}_p$*. We have* $\#E(\mathbb{F}_p) = p + 1$ *and* $\#E(\mathbb{F}_{p^2}) = (p + 1)^2$*. More precisely,* $E(\mathbb{F}_{p^2}) = E[p + 1] \cong (\mathbb{Z}/(p + 1)\mathbb{Z})^2$*.*

(v) Let ℓ *be a prime distinct from* p*. The graph of* ℓ*-isogenies between supersingular elliptic curves over* $\mathbb{F}_{p^2}$ *is connected,* $\ell + 1$*-regular, and is a Ramanujan graph. In particular, it has a* rapid mixing *property: starting from any* E_0 *and taking a walk* $E_0 \rightsquigarrow E$ *of length* $O(\log p)$ *in the graph, the curve* E *at the end of the walk is close-to-uniformly distributed.*

Proof. The first statements are standard in textbooks on elliptic curves, see for example [Sil09, V, Theorems 3.1, 4.1, 2.3.1]. For the last one, see [CGL09, Theorem 4.2]. □

By Proposition 1.37, a walk in the ℓ-isogeny graph is the same as a ℓ^n-isogeny. This leads to the more concrete hard problem:

Problem 5.5 (Isogeny path)**.** Given two elliptic curves $E, E'/\mathbb{F}_q$ such that there exists a *smooth-order* $\varphi\colon E \to E'$ defined over $\mathbb{F}_q$ (there can be many), compute one such φ.

Other natural hard problems can be defined over supersingular elliptic curves:

Problem 5.6 (Endomorphism ring problem)**.** Given a random supersingular elliptic curve $E/\mathbb{F}_q$, find a $\mathbb{F}_p^{\text{alg}}$-basis of its endomorphism ring.

Problem 5.7 (One endomorphism ring problem)**.** Given a random supersingular elliptic curve $E/\mathbb{F}_q$, find a non-scalar endomorphism of E.

Some recent work showed that all these problems are computationally equivalent to the supersingular isogeny problem: see [PW23], [HW23].

The SIDH protocol

As a first example of isogeny-based protocol, we present SIDH (Supersingular Isogeny Diffie–Hellman). It is a key exchange protocol that mimics the Diffie–Hellman protocol, but uses isogenies between supersingular elliptic curves instead of exponentiation in a finite field. It was first introduced in [JD11]. Though it has be shown not to be secure, both this protocol and the attacks to it led to important developments in isogeny-based cryptography, as we'll see.

As public parameters, suppose fixed a (publicly known) supersingular elliptic curve E_0 defined over $\mathbb{F}_{p^2}$. First consider the following (flawed) protocol:

Protocol 5.8.

(i) Alice chooses a secret walk $\varphi_A \colon E_0 \to E_A = E_0/\langle P_A \rangle$ in the 2-isogeny graph, or equivalently, chooses a 2^a-torsion point P_A. She sends E_A to Bob.

(ii) Symmetrically, Bob chooses a walk $\varphi_B \colon E_0 \to E_B/\langle P_B \rangle$, $P_B \in E[3^b]$. He sends E_B to Alice.

(iii) Alice and Bob try and find the shared secret $E_{AB} = E_A/\langle \varphi_A(P_B) \rangle = E_B/\langle \varphi_B(P_A) \rangle$.

If we state the protocol this way, its security relies on the hardness of the isogeny problem, that is, the difficulty of computing the isogenies φ_A and φ_B from the curves E_A and E_B. However, for the protocol to work, Alice must be able to compute $\varphi_B(P_A)$ while φ_B remains secret (Alice knows nothing about φ_B except for the codomain), and vice-versa for Bob.

Therefore, the protocol must be modified as in Figure 5.3: first, fix (S_A, T_A) a deterministic $(\mathbb{Z}/2^a\mathbb{Z})$-basis of $E_0[2^a]$, and (S_B, T_B) a deterministic $(\mathbb{Z}/3^b\mathbb{Z})$-basis of $E_0[3^b]$. Then replace steps (i) and (ii) by:

(i') Alice chooses $P_A = S_A + r_A T_A$ for a secret integer r_A (hence determines φ_A), and sends out the action of φ_A on S_B, T_B along with $E_A = E_0/\langle P_A \rangle$.

(ii') Bob defines r_B, P_B, φ_B, E_B analogously, and sends $\varphi_B(S_A)$, $\varphi_B(S_B)$, E_B.

Remark 5.9. The isogeny computation done by Alice and Bob can be made very efficient just by choosing the prime p wisely. Indeed, let $p = f \cdot 2^a 3^b - 1$: by Proposition 5.4 (iv), choosing E_0 over p we know that the 2^a-torsion and the 3^b torsion are fully $\mathbb{F}_p^2$-rational. Therefore, the corresponding chains of 2-isogenies (resp. 3-isogenies) can be computed using Proposition 1.37, directly over $\mathbb{F}_{p^2}$ without going to an extension field, and this is key to efficiency.

More generally, if $\varphi \colon E_0 \to E'$ is an isogeny defined over $\mathbb{F}_{p^2}$, then the codomain also satisfies[3] $E'(\mathbb{F}_{p^2}) = E'[p+1]$.

The SIDH attacks

In this final form, the SIDH protocol can be executed efficiently, at the cost of publicly exchanging some more pieces of information: but this added information is susceptible of being exploited to retrieve the secret. Indeed, the final protocol is not secure. Recent attacks [CD23] [MM22], [Rob23a] have shown that knowing the action of a

[3] For the readers with a bit of background on elliptic curves: here's why. The p^2-Frobenius endomorphism F_{E_0} of E_0 satisfies $F_{E_0} + [p] = 0$, so it acts as the identity on $p+1$-torsion points. This is why $E_0[p+1]$ is $\mathbb{F}_{p^2}$-rational. If $\varphi \colon E_0 \to E'$ is defined over $\mathbb{F}_{p^2}$, then the same relation holds for $F_{E'}$ as well, therefore $E'[p+1]$ is $\mathbb{F}_{p^2}$-rational too.

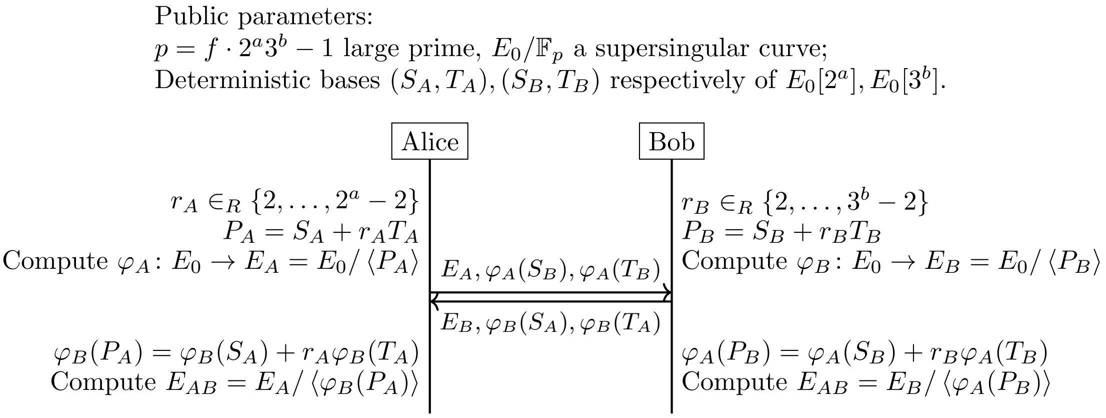

Public parameters:
$p = f \cdot 2^a 3^b - 1$ large prime, $E_0/\mathbb{F}_p$ a supersingular curve;
Deterministic bases $(S_A, T_A), (S_B, T_B)$ respectively of $E_0[2^a], E_0[3^b]$.

Alice

Bob

$r_A \in_R \{2, \dots, 2^a - 2\}$
$P_A = S_A + r_A T_A$
Compute $\varphi_A \colon E_0 \to E_A = E_0/\langle P_A \rangle$

$r_B \in_R \{2, \dots, 3^b - 2\}$
$P_B = S_B + r_B T_B$
Compute $\varphi_B \colon E_0 \to E_B = E_0/\langle P_B \rangle$

$E_A, \varphi_A(S_B), \varphi_A(T_B)$

$E_B, \varphi_B(S_A), \varphi_B(T_A)$

$\varphi_B(P_A) = \varphi_B(S_A) + r_A \varphi_B(T_A)$
Compute $E_{AB} = E_A/\langle \varphi_B(P_A) \rangle$

$\varphi_A(P_B) = \varphi_A(S_B) + r_B \varphi_A(T_B)$
Compute $E_{AB} = E_B/\langle \varphi_A(P_B) \rangle$

Figure 5.3: The SIDH key exchange.

secret isogeny on the N-torsion points for some smooth N (like $N = 2^a, 3^b$) allows to efficiently retrieve the isogeny itself.

The attack, inspired by [Kan97], is made possible by the following lemma.

Notation 5.10. Let $f \colon A_1 \to A_2$ be an isogeny of PPAVs with fixed principal polarisations λ_i on A_i. We denote by $\widetilde{f} \colon A_2 \to A_1$ the isogeny $\lambda_1^{-1} \circ \widehat{f} \circ \lambda_2$.

Lemma 5.11 (Kani, [Kan97, §2]). *Consider a commutative diagram of PPAVs*

$$\begin{tikzcd}
A \arrow[r, "\psi"] \arrow[d, "\varphi"] & A' \arrow[d, "\psi'"] \\
B \arrow[r, "\varphi'"] & B'
\end{tikzcd}$$

where φ is a separable isogeny of degree m, and ψ is an isogeny of degree n. Let $N = m + n$. The isogeny of PPAVs

$$\Phi = \begin{pmatrix} \psi & \widetilde{\psi'} \\ -\varphi & \widetilde{\varphi'} \end{pmatrix} : A \times B' \to A' \times B$$

is a N-isogeny. If m, n are coprime, we can write

$$\begin{aligned} \ker \Phi = \{(\widetilde{\psi}(P) + \widetilde{\varphi}(Q), \psi'(P) + \varphi'(Q)) \mid P \in A'[N], Q \in B[N]\} = \\ = \{([m]P,\ \varphi' \circ \varphi(P)) \mid P \in A[N]\}. \end{aligned}$$

The idea of the attack is that an unknown, large-degree isogeny φ of elliptic curves can be "embedded" into a 2-dimensional isogeny Φ of degree N (where N is chosen depending on $\deg \varphi$), once we know the action of φ on N-torsion points. Using algorithms for isogenies on higher-dimensional abelian varieties, the isogeny Φ can be computed efficiently when N is smooth, and then φ can be extracted from it. In particular, if $N = 2^n$, we can compute Φ as a chain of 2-isogenies, which we saw how to compute in Section 3.4.

Corollary 5.12 (Embedding lemma, [Rob22b, Lemma 2.4]). *Let $\varphi\colon E_1 \to E_2$ be a degree-m isogeny of elliptic curves, and suppose there exists some $N = 2^n$ such that $N - m = a^2$ is a square,* $\gcd(a, m) = 1$*. The isogeny*

$$\Phi = \begin{pmatrix} [a] & -\widetilde{\varphi} \\ \varphi & [a] \end{pmatrix} : E_1 \times E_2 \to E_1 \times E_2$$

is an N-isogeny. Its kernel is $\ker \Phi = \{([a]P, -\varphi(P)) \mid P \in E_1[N]\}$ *and is computable once we know a and the action of φ on the N-torsion of E_1.*

In particular, given any point $X \in E_1$, we can compute $\varphi(X) = \pi_2(\Phi(-X, 0))$.

More generally, if $N - m$ is a sum of $r \in \{2, 4\}$ squares, the isogeny φ can be embedded into a $2r$-dimensional isogeny $E_1^r \times E_2^r$ using the same technique. Recall that all positive integers can be written as a sum of 4 *squares.*

We saw in Section 1.2 that $\mathbb{F}_q$-isogenies of degree m can be computed from their kernel generators in time $O(\sqrt{\ell})$, where ℓ is the largest prime dividing m. If moreover we know the action of the isogeny on a point of large enough 2^n-torsion, the embedding lemma lets us instead compute isogenies of degree $m < 2^n$ in *polynomial time* with respect to $\log m$ (By [Rob23a], it is $O(\log^2 m)$), regardless of how m is factored. If we know nothing about φ, this is of little help: computing this last action requires $O(\sqrt{2^n} \log^2 q)$ – still exponential in $\log m$ – according to [Rob22a]; if we know some torsion-point information instead, the computation of φ becomes really efficient.

Indeed, in the last years, after being used to break the security of SIDH, the embedding lemma has been used constructively in new protocols, such as [BMP23]. Several pre-existing isogeny-based protocols that used 1-dimensional isogenies were now updated so as to make use of the embedding lemma, turning into "higher-dimensional" versions (e.g., SCALLOP-HD [CLP24], SQIsignHD [DLRW24]), which often turn out to be more efficient than their 1-dimensional counterparts. We refer the reader to [Rob24b] for a survey on the HD-novelties introduced by the SIDH attacks.

5.3 Digital signatures and ZK identification schemes

In the following paragraphs, we'll present an isogeny-based digital signature scheme called SQIsign2D-West, a variant of the signature scheme SQIsign that uses higher-dimensional isogenies. Before going into the details, we want to state what is a digital signature scheme in general. For a complete treatment, see [MvOV01, Chapter 11].

As usual, consider two parties Alice and Bob. Alice wants to send a message m to Bob, together with a *signature*, that is, a proof that the message is authentic and Alice is indeed the author of the message. Given m and a signature for it, anyone can verify that the signature is correct, but only Alice can produce it.

Definition 5.13. A digital signature scheme is the datum of three algorithms:

- **Key generation** $\mathsf{KeyGen}() \mapsto (\mathsf{sk}, \mathsf{pk})$. An algorithm that generates a key pair $(\mathsf{sk}, \mathsf{pk})$ for Alice, such that Alice keeps the private key sk secret at all times, while pk is public and linked to Alice's identity.

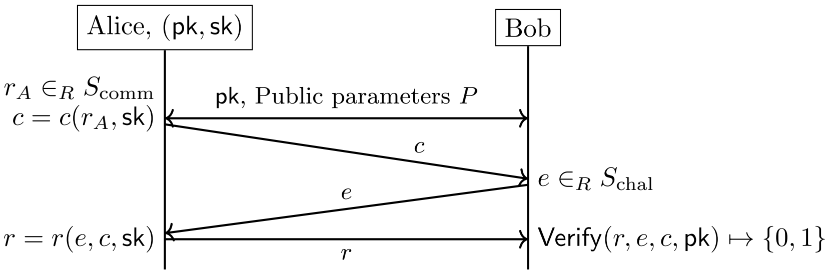
Alice, $(\mathsf{pk}, \mathsf{sk})$
Bob
$r_A \in_R S_{\text{comm}}$
$c = c(r_A, \mathsf{sk})$
pk, Public parameters P
c
$e \in_R S_{\text{chal}}$
e
$r = r(e, c, \mathsf{sk})$
r
$\mathsf{Verify}(r, e, c, \mathsf{pk}) \mapsto \{0, 1\}$

Figure 5.4: A sigma protocol.

- **Signature** $\mathsf{Sign}(m, \mathsf{sk}) \mapsto \sigma$. An algorithm that takes a message and Alice's secret data and output a signature σ for this message.
- **Verification** $\mathsf{Verify}(m, \sigma, \mathsf{pk}) \mapsto b \in \{0, 1\}$. An algorithm that takes a message m, a corresponding signature σ produced by the owner of pk and returns 1 if the signature is valid, 0 otherwise.

Remark 5.14. A digital signature should satisfy at least the following properties, in order to be used to sign documents and messages in the real world:

- *Correctness*: $\mathsf{Verify}(m, \mathsf{Sign}(m, \mathsf{sk}), \mathsf{pk}) = 1$: if a signature is produced correctly, the verification algorithms always considers it valid.
- *Unforgeability*: given a new message m for which we don't know any signature, it is computationally hard to forge a signature σ on m under the public key pk without prior knowledge of the private key sk.

There is a way to build digital signature schemes (in particular, the Sign algorithm) from the more general construction of a *zero-knowledge identification scheme*, where a prover Alice tries to prove her identity to a verifier Bob. In practice, Alice proves to Bob that she knows her secret key sk, but *without* revealing it. Bob should gain no information about sk at the end of the protocol, hence the name *zero-knowledge.*

Protocol 5.15 (Sigma protocol)**.** A zero-knowledge identification scheme usually follows the following pattern, called a *sigma* (Σ) *protocol* because of Figure 5.4.

(i) Alice and Bob agree on public parameters; Alice has a private-public keypair $(\mathsf{sk}, \mathsf{pk})$ and Bob knows that the public key pk belongs to Alice.
(ii) Alice sends a first message c, called a *commitment*, based on her secret data masked by some randomness r_A, which is also kept secret.
(iii) Bob sends a *challenge* e to Alice, chosen uniformly at random from some finite challenge space.
(iv) Alice sends a *response* r to Bob's challenge, based on e, c, using the knowledge of her secret sk.
(v) Bob accepts r if it is a valid response to his challenge, otherwise rejects it.

Remark 5.16. A sigma protocol must satisfy three properties:

- *Completeness.* If Alice knows the secret, then Bob always accepts the response (that is, Alice can always play her role in the protocol according to the rules).
- *Soundness.* If Alice doesn't know the secret, then Bob accepts the response with negligible probability.
- *Zero-knowledge.* Bob learns nothing about the secret, except that Alice knows it.

Definition 5.17 (Hash function)**.** A cryptographically secure length-n hash function is a function $H\colon \{0,1\}^* \to \{0,1\}^n$ that takes any string s and outputs a fixed-length, uniformly random-looking string $H(s)$, in such a way that the following tasks are computationally infeasible without further knowledge:

- *Preimage resistance*: given $H(s)$, find s.
- *Second preimage resistance*: given s, find $s' \neq s$ such that $H(s) = H(s')$.
- *Collision resistance*: find s_1, s_2 such that $H(s_1) = H(s_2)$.

An interactive sigma protocol like the one above can be turned into the a non-interactive protocol by means of the *Fiat–Shamir transform*. The idea is that the challenge e generated by Bob is "randomness Alice can't control". To make the protocol non-interactive, we replace e by $H(c)$ (that is, the hash of the data publicly exchanged before the challenge step) that Alice can compute by herself. Indeed, this is randomness Alice can't control: the output of H is random-looking, and Alice cannot choose c to make e assume a value of her choice, by preimage resistance.

We can then transform the protocol into the signing step of a signature scheme: if Alice wants to sign a message m, she includes the message within the hashed data, that is, let $e = H(c||m)$ be the challenge (where $c||m$ is the string concatenation of c and m), and send (m, c, r) as signature in the end. Intuitively, the protocol proves that the author of m knows sk, so it must be Alice.

Then, if another person (say Bob) wants to verify the signature, then he can recompute e using the same hash function H, and check that the response is correct.

One can show [PS96] that, if the identification scheme we start with has its desired properties, then good properties on the resulting digital signature scheme also hold.

5.4 SQIsign

The cryptographic protocol motivating the study of the present thesis is SQIsign [DKL+20], alongside with its higher-dimensional version. SQIsign is a digital signature scheme based on isogenies of supersingular elliptic curves, derived from the SQIsign interactive identification scheme that we are going to describe below.

The protocol involves two parties Alice and Bob, as usual. At the beginning, Alice and Bob agree on public parameters:

- A large prime $p \equiv 3 \pmod 4$, of size $\log p \approx \lambda$ bits, where λ is a security parameter.[4] The prime p is chosen so that all the isogenies that need to be

[4] This generally means that the best known algorithms to break the security of the protocol have a complexity of approximately 2^λ operations.

$$\begin{tikzcd}
& E_0 \arrow[r, "\varphi_{\mathrm{sk}}"] \arrow[dl, "\varphi_{\mathrm{comm}}"'] & E_{\mathrm{pk}} \arrow[dr, "\varphi_{\mathrm{chal}}"] & \\
E_{\mathrm{comm}} \arrow[rrr, dashed, "\varphi_{\mathrm{resp}}"'] & & & E_{\mathrm{chal}}
\end{tikzcd}$$

Figure 5.5: Isogenies in the SQIsign identification scheme

computed in the protocol are defined over $\mathbb{F}_{p^2}$ to maximise efficiency.

- A supersingular elliptic curve E_0 defined over $\mathbb{F}_{p^2}$ with known endomorphism ring. In practice, E_0 is chosen to be the curve with j-invariant $j(E_0) = 1728$, which is indeed supersingular when $p \equiv 3 \pmod 4$.

Alice is assigned a key pair formed by a public supersingular curve E_{pk} and a secret isogeny $\varphi_{\mathrm{sk}} \colon E_0 \to E_{\mathrm{pk}}$. The interactive protocool is the following. When we say *random isogeny* in the commitment and challenge phases, we mean uniformly random within a set $S_{\mathrm{comm}}, S_{\mathrm{chal}}$ respectively.

Protocol 5.18 (The SQIsign identification scheme (sketch))**.**

- **Commitment.** Alice generates a random commitment isogeny $\varphi_{\mathrm{comm}} \colon E_0 \to E_{\mathrm{comm}}$ and sends E_{comm} to Bob, keeping φ_{comm} secret.
- **Challenge.** Bob generates a random challenge isogeny $\varphi_{\mathrm{chal}} \colon E_{\mathrm{pk}} \to E_{\mathrm{chal}}$ and sends the pair $(\varphi_{\mathrm{chal}}, E_{\mathrm{chal}})$ to Alice.
- **Response.** Alice computes an isogeny $\varphi_{\mathrm{resp}} \colon E_{\mathrm{comm}} \to E_{\mathrm{chal}}$ and sends φ_{resp} to Bob. The isogeny is chosen so that no secrets are leaked: something more complicated than outputting $\varphi_{\mathrm{resp}} = \varphi_{\mathrm{chal}} \circ \varphi_{\mathrm{sk}} \circ \widehat{\varphi_{\mathrm{comm}}}$ must be done. For the same reason, the diagram in Figure 5.5 is *not* commutative.
- **Verification.** Bob checks that φ_{resp} is indeed an isogeny $E_{\mathrm{comm}} \to E_{\mathrm{chal}}$. If the verification is successful, Bob accepts Alice's proof of identity.

The protocol satisfies the three properties of a sigma protocol of Remark 5.16.

The soundness of this protocol follows from the hardness of the supersingular one-endomorphism problem 5.7, that is as hard as the supersingular isogeny problem:

Proposition 5.19. *Protocol 5.18 is sound.*

Proof. Suppose Bob accepts Alice's response with probability close to 1. We want to show that Alice knows an isogeny $E_0 \to E_{\mathrm{pk}}$.

Run the protocol twice and suppose Alice uses twice the same random commitment. The two executions of the protocol give the following two transcripts: $(E_{\mathrm{comm}}, \varphi_{\mathrm{chal}}, E_{\mathrm{chal}}, \varphi_{\mathrm{resp}})$, $(E_{\mathrm{comm}}, \varphi'_{\mathrm{chal}}, E'_{\mathrm{chal}}, \varphi'_{\mathrm{resp}})$. The isogeny $\widehat{\varphi'_{\mathrm{chal}}} \circ \varphi'_{\mathrm{resp}} \circ \varphi_{\mathrm{resp}} \circ \widehat{\varphi_{\mathrm{chal}}}$ is a non-trivial endomorphism of E_{pk}. This means Alice can solve Problem 5.7. As we said when introducing this problem, solving it is equivalent to solving the isogeny problem, that is, retrieving a secret key linked to Alice's public key E_{pk}. □

Showing completeness and the zero-knowledge property needs instead a more detailed description of the steps of the protocol.

SQIsign2D-West

The SQIsign protocol as originally conceived in [DKL$^+$20] is quite slow, though relatively fast if compared to many other isogeny-based protocols. The use of the Embedding Lemma 5.12 gave birth in 2024 to new versions of the protocol [DLRW24], [BDD$^+$24], [NO24], [DF24], which appear to be faster. We will examine the SQIsign2D-West protocol [BDD$^+$24], which bears the most similarities to the SQIsign sigma protocol we described above. The reference diagram is now in Figure 5.6.

Protocol 5.20 (The SQIsign2D-west identification scheme (sketch)).

Parameter selection. Let λ be a security parameter. The prime p is chosen so that $p \equiv 3 \pmod 4$, and $p = f \cdot 2^e - 1$ with $e \approx 2\lambda$ and f a small cofactor. The curve E_0 is chosen to be the Montgomery curve curve $Y^2Z = X^3 + XZ^2$ with j-invariant $j(E_0) = 1728$, that is supersingular and defined over $\mathbb{F}_p$.

This choices imply that for all supersingular curves $E/\mathbb{F}_p^{\text{alg}}$, we can assume E to be defined over $\mathbb{F}_{p^2}$ and $\mathbb{F}_{p^2}$-isogenous to E_0, so that all the $(p+1)$-torsion is $\mathbb{F}_{p^2}$-rational (see Remark 5.9). In particular, the 2^e-torsion of E is defined over $\mathbb{F}_{p^2}$, and so are all the cyclic 2^e-isogenies.

We also fix a large integer $N_{\text{comm}} \geq 2^{4\lambda}$ and parameters $e_{\text{resp}} \approx \lambda$, $e_{\text{chal}} \leq e - e_{\text{resp}}$ such that $2^{e_{\text{resp}}}$ (resp. $2^{e_{\text{chal}}}$) is the number of possible responses (resp. challenges).

Commitment generation. The commitment isogeny φ_{comm} is a random isogeny of degree N_{comm}.

Challenge generation. The challenge isogeny φ_{chal} is an isogeny of degree 2^e defined by a challenge integer $\mathsf{chal} < 2^{e_{\text{chal}}}$. More precisely, let m be the message to be signed and H a length-e_{chal} cryptographically secure hash function, and define $\mathsf{chal} = H(\text{pk}, E_{\text{comm}}, m)$. Find deterministically a torsion basis $(P_{\text{chal}}, Q_{\text{chal}})$ of $E_{\text{pk}}[2^e]$. The isogeny φ_{chal} is the 2^e-isogeny with kernel $\langle P_{\text{chal}} + [\mathsf{chal}]Q_{\text{chal}} \rangle$.

Response generation. The response isogeny φ_{resp} is an isogeny of (any) degree $N_{\text{resp}} < 2^{e_{\text{resp}}}$. By techniques outside of the scope of this thesis, an isogeny φ_{resp} is found so that it is random among the isogenies $E_{\text{comm}} \to E_{\text{chal}}$ of degree bounded by $2^{e_{\text{resp}}}$. Alice sends Bob a convenient representation of φ_{resp}.

Verification. Bob checks that φ_{resp} is indeed an isogeny $E_{\text{comm}} \to E_{\text{chal}}$.

Proposition 5.21. *Protocol 5.20 is complete, sound and is conjectured to be zero-knowledge.*

Proof. The paper [BDD$^+$24], where the protocol was introduced, proves these properties. Soundness is shown as in 5.19. Proving completeness amounts to specifying a working algorithm for Alice's signature and for Bob's verification. In the sections that follow, we're going to detail Bob's verification algorithm. □

5.5 Application: towards verification in small devices

While the first generation of post-quantum standards was being selected by NIST, new calls for proposals were published. One of them [NIS22] looks for digital signature

$$\begin{tikzcd}
& E_0 \arrow[rr, "\varphi_{\text{sk}}"] \arrow[dl, "\varphi_{\text{comm}}"'] & & E_{\text{pk}} \arrow[dr, "\varphi_{\text{chal}}"] & \\
E_{\text{comm}} \arrow[rr, "\varphi^{(\text{odd})}_{\text{resp}}"] \arrow[dd, dashed, "\varphi_{\text{aux}}"'] & & E_{\text{int}} \arrow[rr, "\varphi^{(\text{even})}_{\text{resp}}"] \arrow[dd, dashed, "\varphi'_{\text{aux}}"] & & E_{\text{chal}} \\
& \circlearrowleft & & & \\
E_{\text{aux}} \arrow[rr, dashed, "\varphi'"'] & & E'_{\text{aux}} & &
\end{tikzcd}$$

Figure 5.6: Isogenies in the SQIsign2D-West scheme

algorithms, and SQIsign is a current candidate for it. Besides general-purpose signature schemes, the call also looks for signature schemes with small signatures and fast verification, criteria respected by SQIsign and its variants.

The reason behind this is that oftentimes signatures are produced once (as an example, new software is produced and distributed, and the signature is created alongside the software itself) then verified multiple times by different devices. In particular, embedded systems and Internet-of-Things devices are likely to have to perform cryptographic verification algorithms, while having particularly limited memory resources.

The SQIsign protocol already prioritises fast verification over signing by design (see also [CEMR24] for improvements in this direction), but it is significantly slower than its higher-dimensional analogues. We are therefore interested in studying whether the protocol is still compatible with verification being run in small devices once we bring 2-dimensional isogenies into the picture.

The verification algorithm

To describe the verification algorithm, we must first specify what's the "convenient representation" of φ_{resp} mentioned above, such that Bob can efficiently compute it to run verification. In other words, we need to say what Alice's signature contains.

Remark 5.22. Writing $N_{\text{resp}} = 2^{n_{\text{even}}} \cdot m$ with m odd, by Proposition 1.37 we can factor the response isogeny as $\varphi_{\text{resp}} = \varphi^{(\text{even})}_{\text{resp}} \circ \varphi^{(\text{odd})}_{\text{resp}}$, where $\varphi^{(\text{odd})}_{\text{resp}} : E_{\text{comm}} \to E_{\text{int}}$ is a m-isogeny and $\varphi^{(\text{even})}_{\text{resp}} : E_{\text{int}} \to E_{\text{chal}}$ is a $2^{n_{\text{even}}}$-isogeny.[5]

Recall that a $2^{n_{\text{even}}}$-isogeny is efficiently computable as a chain of 2-isogenies, again by Proposition 1.37. The odd part instead needs the Embedding Lemma 5.12: Alice finds an auxiliary curve E_{aux} and an isogeny $\varphi_{\text{aux}} : E_{\text{comm}} \to E_{\text{aux}}$ that fits in the lower

[5] It can happen that the last part of the 2-isogeny path determined by $\varphi^{(\text{even})}_{\text{resp}}$ coincides with the last part of the path determined by φ_{chal}. The authors of [BDD+24] introduce some technical modifications to handle this case and optimise computations. For simplicity of exposition, we will assume this case doesn't happen.

commutative square of the diagram in Figure 5.6, such that

$$\Phi = \begin{pmatrix} \varphi_{\mathrm{resp}}^{(\mathrm{odd})} & -\widehat{\varphi'_{\mathrm{aux}}} \\ \varphi_{\mathrm{aux}} & \widehat{\varphi'} \end{pmatrix} : E_{\mathrm{comm}} \times E'_{\mathrm{aux}} \to E_{\mathrm{int}} \times E_{\mathrm{aux}}$$

is a 2-dimensional $2^{e_{\mathrm{resp}}-n_{\mathrm{even}}}$-isogeny "embedding" the 1-dimensional $\varphi_{\mathrm{resp}}^{(\mathrm{odd})}$.

Proposition 5.23 ([BDD+24], Algorithm 7)**.** *Alice has an algorithm to compute a response isogeny φ_{resp}, decomposed in an even and an odd part as above, together with the following data:*

(i) a 2-dimensional $2^{e_{\mathrm{resp}}-n_{\mathrm{even}}}$-isogeny

$$\Phi = \begin{pmatrix} \varphi_{\mathrm{resp}}^{(\mathrm{odd})} & -\widehat{\varphi'_{\mathrm{aux}}} \\ \varphi_{\mathrm{aux}} & \widehat{\varphi'} \end{pmatrix} : E_{\mathrm{comm}} \times E'_{\mathrm{aux}} \to E_{\mathrm{int}} \times E_{\mathrm{aux}};$$

(ii) the curve E_{aux} and a deterministic torsion basis $\langle P_{\mathrm{aux}}, Q_{\mathrm{aux}}\rangle = E_{\mathrm{aux}}[2^{e_{\mathrm{resp}}}]$;

(iii) a torsion basis $\langle P_{\mathrm{chal}}, Q_{\mathrm{chal}}\rangle = E_{\mathrm{chal}}[2^{e_{\mathrm{resp}}}]$ such that:

(a) the points $[2^{e_{\mathrm{resp}}-n_{\mathrm{even}}}]P_{\mathrm{chal}}, [2^{e_{\mathrm{resp}}-n_{\mathrm{even}}}]Q_{\mathrm{chal}}$ generate $\ker(\widehat{\varphi_{\mathrm{resp}}^{(\mathrm{even})}})$;

(b) the points $P_{\mathrm{int}} = \widehat{\varphi_{\mathrm{resp}}^{(\mathrm{even})}}(P_{\mathrm{chal}}), Q_{\mathrm{int}} = \widehat{\varphi_{\mathrm{resp}}^{(\mathrm{even})}}(Q_{\mathrm{chal}})$ generate $E_{\mathrm{int}}[2^{e_{\mathrm{resp}}-n_{\mathrm{even}}}]$;

(c) $\ker \Phi = \langle (P_{\mathrm{int}}, [2^{n_{\mathrm{even}}}]P_{\mathrm{aux}}), (Q_{\mathrm{int}}, [2^{n_{\mathrm{even}}}]Q_{\mathrm{aux}})\rangle$.

Alice's response is then $\sigma = (E_{\mathrm{aux}}, n_{\mathrm{even}}, P_{\mathrm{chal}}, Q_{\mathrm{chal}})$. This allows Bob to compute the dual isogeny $\widehat{\varphi_{\mathrm{resp}}}$ and recover E_{comm} from the signature and the public data. If he actually recovers E_{comm}, the signature is considered valid. The verification algorithm is then outlined in Algorithm 15.

In practice, Bob's task amounts to performing some 2-isogeny chains, of which two are in dimension 1, and another is in dimension 2.

2-isogeny chains

In Section 3.4 we presented a generic algorithm to compute a $(2, 2)$-isogeny step $\varphi_i \colon A \to A'$ that works when both A and A' are irreducible abelian varieties equipped with level-2 theta coordinates.

In order to implement SQIsign2D verification, we need something more: we want to compute a chain of $(2, 2)$-isogenies between abelian varieties, having products of elliptic curves both as the domain and the final codomain. With high probability, the intermediate steps of the chain involve irreducible abelian varieties, so the algorithms of Section 3.4 can be used there. The authors of [DMPR23a] complete the picture, by describing the following algorithms handling the first and the last step, where one of the abelian varieties is a product of elliptic curves:

- a *change of basis* algorithm that takes two Montgomery elliptic curves E_1, E_2 along with bases of their 2^n-torsion, and finds a suitable level-2 theta structure $\Theta_{\mathcal{L}}$ on $A = E_1 \times E_2$ such that $A[2] = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$ is a symplectic decomposition and $K_2(\mathcal{L}) = \ker(\varphi_0)$ is the kernel of the first step φ_0 of the isogeny φ.

Algorithm 15 Verification algorithm in SQIsign2D

```
Input: $E_{\mathrm{pk}}$, $E_{\mathrm{comm}}$, $\mathsf{chal}$, $\sigma = (E_{\mathrm{aux}}, n_{\mathrm{even}}, P_{\mathrm{chal}}, Q_{\mathrm{chal}})$, a message $m$
Output: 1 if $\sigma$ is a valid signature for $m$, 0 otherwise.
Compute torsion bases $\langle P_{\mathrm{aux}}, Q_{\mathrm{aux}}\rangle = E_{\mathrm{aux}}[2^{e_{\mathrm{resp}}}]$ and $\langle P_{\mathrm{pk}}, Q_{\mathrm{pk}}\rangle = E_{\mathrm{pk}}[2^e]$ deterministically (using an algorithm shared with Alice).
Retrieve challenge:
    $\mathsf{chal} \leftarrow H(\mathrm{pk}, E_{\mathrm{comm}}, m)$
    Compute $E_{\mathrm{chal}} \leftarrow E_{\mathrm{pk}}/\langle P_{\mathrm{pk}} + [\mathsf{chal}]Q_{\mathrm{pk}}\rangle$
Verify even part of the response:
    $P' \leftarrow [2^{e_{\mathrm{resp}}-n_{\mathrm{even}}}]P_{\mathrm{chal}}, Q \leftarrow [2^{e_{\mathrm{resp}}-n_{\mathrm{even}}}]Q_{\mathrm{chal}}$
    Compute $\widehat{\varphi_{\mathrm{resp}}^{(\mathrm{even})}}\colon E_{\mathrm{chal}} \to E_{\mathrm{int}}$ from its kernel $\langle P', Q'\rangle$.
    $P_{\mathrm{int}} \leftarrow \widehat{\varphi_{\mathrm{resp}}^{(\mathrm{even})}}(P_{\mathrm{chal}})$, $Q_{\mathrm{int}} \leftarrow \widehat{\varphi_{\mathrm{resp}}^{(\mathrm{even})}}(Q_{\mathrm{chal}})$
Apply the Embedding Lemma:
    $P_{2D} \leftarrow (P_{\mathrm{int}}, [2^{n_{\mathrm{even}}}]P_{\mathrm{aux}})$, $Q_{2D} \leftarrow (Q_{\mathrm{int}}, [2^{n_{\mathrm{even}}}]Q_{\mathrm{aux}})$
    Compute $\widetilde{\Phi}\colon E_{\mathrm{int}} \times E_{\mathrm{aux}} \to F_1 \times F_2$ with kernel $\langle (P_{2D}, Q_{2D})\rangle$.
Verify $F_2 \cong E_{\mathrm{comm}}$.
return 1 if the last check was successful, 0 otherwise.
```

- the algorithms we studied in the previous chapters: general *codomain* and *evaluation* isogeny algorithms taking as input an abelian variety A with level-2 theta structure $\Theta_{\mathcal{L}}$, where $K_2(\mathcal{L}) = \ker(\varphi')$ is the kernel of a $(2,2)$-isogeny $\varphi'\colon A \to A'$ where neither A nor A' is a product, and outputting a representation of the codomain A' and a recipe to evaluate φ on any $Q \in A$.
- Special codomain and evaluation algorithms when either A or A' is a product of elliptic curves.
- A *splitting* algorithm that takes an input an A that is a product of elliptic curves (with any theta structure, not necessarily the *product theta structure* of Remark 2.68) and outputs the two factors E_1, E_2 in Montgomery form.

Memory usage

The steps of SQIsign2D verification (Algorithm 15) are executed in sequence, so the overall amount of memory needed to run the algorithm is roughly equal to the memory used by the most memory-intensive step, that is, the 2-dimensional isogeny Φ'. Indeed, the first step (generating a torsion basis) amounts to generating two independent random points on an elliptic curve; the last one (verifying an isomorphism between two elliptic curves) is just an equality check between their j-invariants. We're left with two 2-isogeny chains in dimension 1 and one in dimension 2, and the latter inherently uses more space than the former. Indeed, points on a surface are represented by more coordinates than points on a curve, and the kernel of a 2-isogeny in dimension g needs g generators to be represented.

Modelling a typical setting of resource-constrained devices (see e.g. [GHK+21] in the context of verification of post-quantum signatures), we assume that the stack memory available to the verification algorithm is limited to 8 kB. In real-life situations, cryptographic verification is usually performed as a subroutine of a larger program, so the stack memory is shared with other tasks and the available space may be even smaller.

According to [BDD+24], depending on the security level λ, the bitsize of the prime p can assume the values

$$t = \lceil \log_2 p \rceil \in \{248, 376, 500\}.$$

A Kummer point over $\mathbb{F}_{p^2}$ (stored as four $\mathbb{F}_{p^2}$ elements, that is, eight $\mathbb{F}_p$ elements) then takes up $8t$ bits $= t$ bytes.

This means there's space for $8192/t \in [16, 33]$ Kummer points in memory to safely run the verification algorithm. One can check by inspection that the steps of the $(2, 2)$-isogeny chain listed in Section 5.5 need less than 4 auxiliary Kummer points at the same time to be executed (see also [Dar24, Appendix B.1] for a change of basis using even less memory); as above, the steps are run sequentially, so the auxiliary memory of one step can be erased and rewritten in the subsequent step. This suggests implementing SQIsign2D verification in small devices should largely be possible; a more precise analysis would need a reference implementation of the protocol and is subject to unpredictable variations due to the specific architecture of the device hardware.

Strategies

The current reference implementation of $(2, 2)$-isogeny chains in SageMath [DMPR23b] applies a space-time tradeoff (suggested in [DJP11] for 1-dimensional chains of 2-isogenies) to save time at the expense of memory to perform the computation of the chain. Given a $(2, 2)$-isogeny chain of length n, say $\varphi_n \circ \cdots \circ \varphi_1$, let $\mathcal{B}_0 = (S_0, T_0)$ be such that S_0, T_0 are 2^n-torsion points with $\langle [4]S_0, [4]T_0 \rangle = \ker(\varphi_n \cdots \varphi_1)$, and $\ker \varphi_i$ is given by the points of $[2^{n-i}]\mathcal{B}_{i-1}$ with $\mathcal{B}_i = \varphi_i(\mathcal{B}_{i-1})$, $i = 1, \ldots, 4$. At each step we need to compute the kernel generators of φ_i via some isogenies and doublings (or better, the 8-torsion points above them), that is, we need to compute the leaves of the tree in Figure 5.7.

The naive approach implied by Proposition 1.37 would require computing all the $\mathcal{B}_i$ from $\mathcal{B}_{i-1}$, then their multiples $[2^{n-i}]\mathcal{B}_i$ from $\mathcal{B}_i$ at each step, which would have a total cost of $O(n^2)$ point doublings in the variety. This corresponds to walking along the tree following the path in figure 5.8a.

However, we can do better and walk instead along the tree of Figure 5.7, which has $n \log_2 n$ nodes, at the cost of storing the branching points to reuse them to optimise computations. At a given time, this strategy requires $2 \cdot \lceil \log_2 n \rceil$ points in memory.

Coming back to the memory analysis above, at the lowest security level, with $t = 248$, considering $n = \lceil \log_2 p \rceil$, the number of branching points to store is $2 \cdot \lceil \log_2 248 \rceil = 16$, which could fit within the available storage space. At higher security levels, this becomes inapplicable.

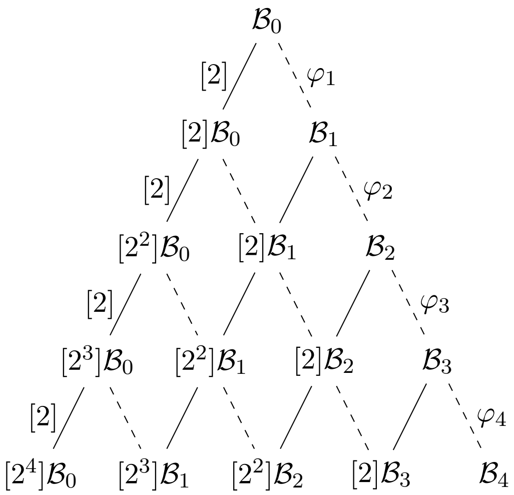

Figure 5.7: Computation tree for a 2^n-isogeny, $n = 4$, in dimension $g = 2$.
$B_0 = (S_0, T_0)$ is a pair of 2^{n+2}-torsion points such that $([4]S_0, [4]T_0)$ is the kernel of $\varphi_n \circ \cdots \circ \varphi_1$.

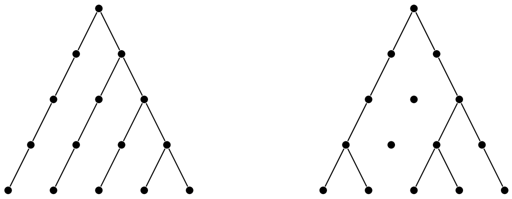

(a) Naive strategy (b) Balanced strategy

An intermediate solution could be to use an algorithm that splits the chain into consecutive 4-isogenies instead of 2-isogenies (see Section 3.4). In this case, we'd have to store three kernel points instead of two: $S_0, T_0, S_0 + T_0$; however, the number of branching nodes to store would be now $\lceil \log_4 t \rceil$, that is, it'd be halved, and this could be a good compromise between time and memory efficiency.

Bibliography

[AFK24] Mónica P. Arenas, Georgios Fotiadis, and Elisavet Konstantinou, *Special TNFS-secure pairings on ordinary genus 2 hyperelliptic curves*, Progress in Cryptology - AFRICACRYPT 2024 (Berlin, Heidelberg), Springer-Verlag, 2024, pp. 285–310.

[Bas24] Andrea Basso, *POKE: A framework for efficient PKEs, split KEMs, and OPRFs from higher-dimensional isogenies*, Cryptology ePrint Archive, Paper 2024/624, 2024, `https://eprint.iacr.org/2024/624`.

[BCP97] Wieb Bosma, John Cannon, and Catherine Playoust, *The Magma algebra system. I. The user language*, J. Symbolic Comput. **24** (1997), no. 3-4, 235–265, Computational algebra and number theory (London, 1993). MR MR1484478

[BCR11] Gaetan Bisson, Romain Cosset, and Damien Robert, *AVIsogenies (Version 0.3)*, `https://www.math.u-bordeaux.fr/~damienrobert/avisogenies/`, 2011, Accessed: September 2024.

[BDD$^+$24] Andrea Basso, Luca De Feo, Pierrick Dartois, Antonin Leroux, Luciano Maino, Giacomo Pope, Damien Robert, and Benjamin Wesolowski, *SQIsign2D-West: The Fast, the Small, and the Safer*, Cryptology ePrint Archive, Paper 2024/760, 2024, `https://eprint.iacr.org/2024/760`.

[BDLS20] Daniel J. Bernstein, Luca De Feo, Antonin Leroux, and Benjamin Smith, *Faster computation of isogenies of large prime degree*, ANTS-XIV - 14th Algorithmic Number Theory Symposium, Proceedings (Auckland, New Zealand) (Steven D. Galbraith, ed.), vol. 4, Mathematical Sciences Publishers, June 2020, pp. 39–55.

[BF01] Dan Boneh and Matt Franklin, *Identity-based encryption from the Weil pairing*, Advances in Cryptology — CRYPTO 2001, Lecture notes in computer science, Springer Berlin Heidelberg, Berlin, Heidelberg, 2001, pp. 213–229.

[BL04] Christina Birkenhake and Herbert Lange, *Complex abelian varieties*, 2nd, augmented ed., Grundlehren der mathematischen Wissenschaften; vol. 302, Springer, Berlin, 2004.

[BMP23] Andrea Basso, Luciano Maino, and Giacomo Pope, *FESTA: Fast Encryption from Supersingular Torsion Attacks*, Advances in Cryptology – ASIACRYPT 2023 (Singapore) (Jian Guo and Ron Steinfeld, eds.), Springer Nature Singapore, 2023, pp. 98–126.

[Bre83] Lawrence Breen, *Fonctions thêta et théorème du cube*, 1983 ed., Lecture notes in mathematics, Springer, Berlin, Germany, March 1983 (fr).

[CD23] Wouter Castryck and Thomas Decru, *An efficient key recovery attack on SIDH*, Advances in Cryptology – EUROCRYPT 2023 (Berlin, Heidelberg), no. 5, Springer-Verlag, 2023, p. 423–447.

[CEMR24] Maria Corte-Real Santos, Jonathan Komada Eriksen, Michael Meyer, and Krijn Reijnders, *AprèsSQI: Extra fast verification for SQIsign using extension-field signing*, Lecture Notes in Computer Science, Springer Nature Switzerland, Cham, 2024, pp. 63–93.

[CF96] John W. S. Cassels and Victor Flynn, *Prolegomena to a middlebrow arithmetic of curves of genus 2*, London Mathematical Society Lecture Note Series, Cambridge University Press, 1996.

[CFA$^+$12] Henri Cohen, Gerhard Frey, Roberto Avanzi, Christophe Doche, Tanja Lange, Kim Nguyen, and Frederik Vercauteren, *Handbook of elliptic and hyperelliptic curve cryptography, second edition*, 2nd ed., Chapman & Hall/CRC, 2012.

[CGL09] Denis Charles, Eyal Goren, and Kristin Lauter, *Families of Ramanujan graphs and quaternion algebras*, Groups and symmetries, CRM Proceedings Lecture notes, vol. 47, American Mathematical Society, Providence, RI, 2009.

[CH17] Craig Costello and Huseyin Hisil, *A simple and compact algorithm for SIDH with arbitrary degree isogenies*, Advances in Cryptology – ASIACRYPT 2017 (Cham) (Tsuyoshi Takagi and Thomas Peyrin, eds.), Springer International Publishing, 2017, pp. 303–329.

[CLG09] Denis Charles, Kristin Lauter, and Eyal Goren, *Cryptographic hash functions from expander graphs*, J. Cryptology **22** (2009), no. 1, 93–113, first appeared in Jan 2006.

[CLM$^+$18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes, *CSIDH: An efficient post-quantum commutative group action*, Advances in Cryptology – ASIACRYPT 2018 (Cham) (Thomas Peyrin and Steven D. Galbraith, eds.), Springer International Publishing, 2018, pp. 395–427.

[CLP24] Mingjie Chen, Antonin Leroux, and Lorenz Panny, *SCALLOP-HD: Group action from 2-dimensional isogenies*, Public-Key Cryptography – PKC 2024 (Cham) (Qiang Tang and Vanessa Teague, eds.), Springer Nature Switzerland, 2024, pp. 190–216.

[Cos19] Craig Costello, *Supersingular isogeny key exchange for beginners*, Cryptology ePrint Archive, Paper 2019/1321, 2019, `https://eprint.iacr.org/2019/1321`.

[Cou06] Jean-Marc Couveignes, *Hard homogeneous spaces*, Cryptology ePrint Archive, Paper 2006/291, 2006, `https://eprint.iacr.org/2006/291`.

[CS17] Craig Costello and Benjamin Smith, *Montgomery curves and their arithmetic*, Journal of Cryptographic Engineering **8** (2017), 227 – 240.

[Dar24] Pierrick Dartois, *Fast computation of 2-isogenies in dimension 4 and cryptographic applications*, Cryptology ePrint Archive, Paper 2024/1180, 2024, `https://eprint.iacr.org/2024/1180`.

[DF24] Max Duparc and Tako Boris Fouotsa, *SQIPrime: A dimension 2 variant of SQISignHD with non-smooth challenge isogenies*, Cryptology ePrint Archive, Paper 2024/773, 2024, `https://eprint.iacr.org/2024/773`.

[DH76] Whitfield Diffie and Martin Hellman, *New directions in cryptography*, IEEE Transactions on Information Theory **22** (1976), no. 6, 644–654.

[DJP11] Luca De Feo, David Jao, and Jérôme Plût, *Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies*, Cryptology ePrint Archive, Paper 2011/506, 2011, `https://eprint.iacr.org/2011/506`.

[DKL$^+$20] Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski, *SQISign: Compact post-quantum signatures from quaternions and isogenies*, Advances in Cryptology – ASIACRYPT 2020 (Cham) (Shiho Moriai and Huaxiong Wang, eds.), Springer International Publishing, 2020, pp. 64–93.

[DLRW24] Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski, *SQIsignHD: New dimensions in cryptography*, Advances in Cryptology – EUROCRYPT 2024 (Cham) (Marc Joye and Gregor Leander, eds.), Springer Nature Switzerland, 2024, pp. 3–32.

[DMPR23a] Pierrick Dartois, Luciano Maino, Giacomo Pope, and Damien Robert, *An algorithmic approach to* $(2,2)$*-isogenies in the theta model and applications to isogeny-based cryptography*, Cryptology ePrint Archive, Paper 2023/1747, 2023, `https://eprint.iacr.org/2023/1747`.

[DMPR23b] ———, *Two-isogenies*, `https://github.com/ThetaIsogenies/two-isogenies/`, 2023, Accessed: September 2024.

[Gal05] Steven D. Galbraith, *Pairings*, Advances in Elliptic Curve Cryptography (Ian F. Blake, Gadiel Seroussi, and Nigel P.Editors Smart, eds.), London Mathematical Society Lecture Note Series, Cambridge University Press, 2005, p. 183–214.

[Gal12] ———, *Mathematics of public key cryptography*, Cambridge University Press, Cambridge, England, June 2012.

[Gal24] ———, *Climbing and descending tall volcanos*, Cryptology ePrint Archive, Paper 2024/924, 2024, `https://eprint.iacr.org/2024/924`.

[GHK$^+$21] Ruben Gonzalez, Andreas Hülsing, Matthias J. Kannwischer, Juliane Krämer, Tanja Lange, Marc Stöttinger, Elisabeth Waitz, Thom Wiggers, and Bo-Yin Yang, *Verifying post-quantum signatures in 8 kB of RAM*, Post-Quantum Cryptography, Lecture notes in computer science, Springer International Publishing, Cham, 2021, pp. 215–233.

[Gro72] Alexander Grothendieck, *Biextensions de faisceaux de groupes*, Groupes de Monodromie en Géométrie Algébrique (Berlin, Heidelberg), Springer Berlin Heidelberg, 1972, pp. 133–217 (fr).

[GV18] Steven D. Galbraith and Frederik Vercauteren, *Computational problems in supersingular elliptic curve isogenies*, Quantum Inf. Process. **17** (2018), no. 10, (265) 1–22.

[HW23] Arthur Herlédan-Le Merdy and Benjamin Wesolowski, *The supersingular endomorphism ring problem given one endomorphism*, Cryptology ePrint Archive, Paper 2023/1448, 2023.

[JD11] David Jao and Luca De Feo, *Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies*, Post-Quantum Cryptography (Berlin, Heidelberg) (Bo-Yin Yang, ed.), Springer Berlin Heidelberg, 2011, pp. 19–34.

[Jou00] Antoine Joux, *A one round protocol for tripartite Diffie–Hellman*, Lecture Notes in Computer Science, Springer Berlin Heidelberg, Berlin, Heidelberg, 2000, pp. 385–393.

[Kan97] Ernst Kani, *The number of curves of genus two with elliptic differentials.*, Journal für die reine und angewandte Mathematik **1997** (1997), no. 485, 93–122.

[Kob87] Neal Koblitz, *Elliptic curve cryptosystems*, Mathematics of computation **48** (1987), no. 177, 203–209.

[Koh96] David Kohel, *Endomorphism rings of elliptic curves over finite fields*, Ph.D. thesis, University of California, Berkeley, 1996.

[Koi76] Shoji Koizumi, *Theta relations and projective normality of abelian varieties*, American Journal of Mathematics **98** (1976), no. 4, 865–889.

[Kup05] Greg Kuperberg, *A subexponential-time quantum algorithm for the dihedral hidden subgroup problem*, SIAM J. Comput. **35** (2005), no. 1, 170–188.

[Lab21] F5 Labs, *The 2021 TLS telemetry report*, `https://www.f5.com/labs/articles/threat-intelligence/the-2021-tls-telemetry-report`, 2021, Accessed: September 2024.

[Lan83] Serge Lang, *Abelian varieties*, 1 ed., Springer, New York, NY, September 1983.

[Liu02] Qing Liu, *Algebraic geometry and arithmetic curves*, Oxford Graduate Texts in Mathematics, vol. 6, Oxford University Press, Oxford, 2002, Translated from the French by Reinie Erné, Oxford Science Publications.

[LR10] David Lubicz and Damien Robert, *Efficient pairing computation with theta functions*, ANTS-IX - 9th Algorithmic Number Theory Symposium, Proceedings (Guillaume Hanrot, François Morain, and Emmanuel Thom, eds.), Lecture Notes in Computer Science, vol. 6197, Springer–Verlag, 7 2010.

[LR12] ———, *Computing isogenies between abelian varieties*, Compositio Mathematica **148** (2012), no. 5, 1483–1515.

[LR15] ———, *A generalisation of Miller's algorithm and applications to pairing computations on abelian varieties*, J. Symbolic Comput. **67** (2015), 68–92.

[Mes86] Jean-François Mestre, *La méthode des graphes. Exemples et applications*, Proceedings of the international conference on class numbers and fundamental units of algebraic number fields (Katata, 1986), Nagoya Univ., Nagoya, 1986, pp. 217–242. MR 891898

[MFK94] David Mumford, John Fogarty, and Frances Kirwan, *Geometric invariant theory*, Ergebnisse der Mathematik und Ihrer Grenzgebiete, 3 Folge/A Series of Modern Surveys in Mathematics Series, Springer Berlin Heidelberg, 1994.

[Mil85] Victor S. Miller, *Use of elliptic curves in cryptography*, Conference on the theory and application of cryptographic techniques, Springer, 1985, pp. 417–426.

[Mil86a] James S. Milne, *Abelian varieties*, Arithmetic Geometry (Gary Cornell and Joseph H. Silverman, eds.), Springer New York, New York, NY, 1986, pp. 103–150.

[Mil86b] ______, *Jacobian varieties*, Arithmetic Geometry (Gary Cornell and Joseph H. Silverman, eds.), Springer New York, New York, NY, 1986, pp. 167–212.

[MM22] Luciano Maino and Chloe Martindale, *An attack on SIDH with arbitrary starting curve*, Cryptology ePrint Archive, Paper 2022/1026, 2022, `https://eprint.iacr.org/2022/1026`.

[Mon87] Peter L. Montgomery, *Speeding the Pollard and elliptic curve methods of factorization*, Mathematics of Computation **48** (1987), 243–264.

[Mum66] David Mumford, *On the equations defining abelian varieties, I.*, Inventiones mathematicae **1** (1966), 287–354.

[Mum68] ______, *Bi-extensions of formal groups*, Algebraic Geometry, International Colloquium (Bombay), Tata Institute of Fundamental Research, 1968.

[Mum74] ______, *Abelian varieties*, Studies in mathematics, Oxford University Press, 1974.

[Mum84] ______, *Tata Lectures on Theta II. Jacobian theta functions and differential equations. With the collaboration of C. Musili, M. Nori, E. Previato, M. Stillman and H. Umemura.*, vol. 43, p. 272, Birkhäuer Verlag, Boston, 1984.

[MvOV01] Alfred J. Menezes, Paul C. van Oorschot, and Scott A. Vanstone, *Handbook of applied cryptography*, CRC Press, 2001.

[NIS16] NIST, *Post quantum cryptography standardization*, `https://csrc.nist.gov/Projects/post-quantum-cryptography/post-quantum-cryptography-standardization`, 2016, Accessed: September 2024.

[NIS22] ______, *Standardization of additional digital signature schemes*, `https://csrc.nist.gov/Projects/pqc-dig-sig/standardization`, 2022, Accessed: September 2024.

[NO24] Kohei Nakagawa and Hiroshi Onuki, *SQIsign2D-East: A new signature scheme using 2-dimensional isogenies*, Cryptology ePrint Archive, Paper 2024/771, 2024, `https://eprint.iacr.org/2024/771`.

[PQS] PQShield, *Post-quantum signatures zoo*, `https://pqshield.github.io/nist-sigs-zoo/`, Accessed: September 2024.

[PS96] David Pointcheval and Jacques Stern, *Security proofs for signature schemes*, Advances in Cryptology — EUROCRYPT 1996 (Ueli Maurer, ed.), Springer Berlin Heidelberg, Berlin, Heidelberg, 1996, pp. 387–398.

[PW23] Aurel Page and Benjamin Wesolowski, *The supersingular endomorphism ring and one endomorphism problems are equivalent*, Cryptology ePrint Archive, Paper 2023/1399, 2023.

[Ren18] Joost Renes, *Computing isogenies between Montgomery curves using the action of (0, 0)*, Post-Quantum Cryptography - 9th International Conference, PQCrypto 2018, Fort Lauderdale, FL, USA, April 9-11, 2018, Proceedings (Tanja Lange and Rainer Steinwandt, eds.), Lecture Notes in Computer Science, vol. 10786, Springer, 2018, pp. 229–247.

[Rob10] Damien Robert, *Theta functions and cryptographic applications*, Ph.D. thesis, Université Henri-Poincaré, Nancy 1, France, July 2010.

[Rob21] ———, *Efficient algorithms for abelian varieties and their moduli spaces*, June 2021, HDR thesis, `http://www.normalesup.org/~robert/pro/publications/academic/hdr.pdf`.

[Rob22a] ———, *Evaluating isogenies in polylogarithmic time*, Cryptology ePrint Archive, Paper 2022/1068, 2022, `https://eprint.iacr.org/2022/1068`.

[Rob22b] ———, *Some applications of higher dimensional isogenies to elliptic curves (overview of results)*, Cryptology ePrint Archive, Paper 2022/1704, 2022, `https://eprint.iacr.org/2022/1704`.

[Rob23a] ———, *Breaking SIDH in polynomial time*, Advances in Cryptology – EUROCRYPT 2023 (Cham) (Carmit Hazay and Martijn Stam, eds.), Springer Nature Switzerland, 2023, pp. 472–503.

[Rob23b] ———, *The geometric interpretation of the Tate pairing and its applications*, Cryptology ePrint Archive, Paper 2023/177, 2023, `https://eprint.iacr.org/2023/177`.

[Rob23c] ———, *Kummer-line*, `https://gitlab.inria.fr/roberdam/kummer-line`, 2023, Accessed: September 2024.

[Rob24a] ———, *Fast pairings via biextensions and cubical arithmetic*, Cryptology ePrint Archive, Paper 2024/517, 2024, `https://eprint.iacr.org/2024/517`.

[Rob24b] ———, *On the efficient representation of isogenies (a survey)*, Cryptology ePrint Archive, Paper 2024/1071, 2024, `https://eprint.iacr.org/2024/1071`.

[Rob24c] ———, *Some notes on algorithms for abelian varieties*, Cryptology ePrint Archive, Paper 2024/406, 2024, `https://eprint.iacr.org/2024/406`.

[RS06] Alexander Rostovtsev and Anton Stolbunov, *Public-key cryptosystem based on isogenies*, IACR Cryptol. ePrint Arch. **2006** (2006), 145.

[RSA78] Ronald L. Rivest, Adi Shamir, and Leonard Adleman, *A method for obtaining digital signatures and public-key cryptosystems*, Communications of the ACM **21** (1978), no. 2, 120–126.

[S$^+$24] William A. Stein et al., *Sage Mathematics Software (Version 10.3)*, The Sage Development Team, 2024, `http://www.sagemath.org`.

[Sch95] René Schoof, *Counting points on elliptic curves over finite fields*, Journal de théorie des nombres de Bordeaux **7** (1995), no. 1, 219–254. MR 1413578

[Sho94] Peter W. Shor, *Algorithms for quantum computation: discrete logarithms and factoring*, Proceedings 35th Annual Symposium on Foundations of Computer Science, 1994, pp. 124–134.

[Sil09] Joseph H. Silverman, *The arithmetic of elliptic curves*, Graduate Texts in Mathematics, Springer New York, 2009.

[Sta08] Katherine E. Stange, *Elliptic nets and elliptic curves*, Ph.D. thesis, Brown University, 2008.

[Tes06] Edlyn Teske, *An elliptic curve trapdoor system*, J. Cryptology **19** (2006), no. 1, 115–133, first appeared in Jan 2003.

[Vak24] Ravi Vakil, *MATH 216: Foundations of Algebraic Geometry*, 2024, `https://math.stanford.edu/~vakil/216blog/`.

[Was08] Lawrence C. Washington, *Elliptic curves: Number theory and cryptography*, 2nd ed., Chapman & Hall/CRC, 2008.