Ryan Rueger

Computing the Isogeny Class-Group Action on Ordinary Elliptic Curves by going into Higher Dimensions

A thesis submitted to the ETH Zurich
for the degree of *master* in mathematics.
Supervised by *Dr. Luca De Feo* of IBM Research,
and *Prof. Ueli Maurer* of ETH Zurich

Spring 2024

Abstract

We assess the feasibility of using the new higher-dimensional methods of Clapoti [2] to evaluate the isogeny class-group action on ordinary elliptic curves; this action is (conjecturally) an instance of a post-quantum cryptographic group action from which many primitives can be built.

By specialising the theory to ordinary elliptic curves and then giving an almost complete Sage implementation for toy examples, we conclude that the only bottleneck to practical usage of Clapoti lies in the fact that we are considering *ordinary* elliptic curves. As such, we suspect Clapoti to not only be asymptotically polynomial time, but concretely efficient when evaluating the isogeny class-group action on *supersingular* elliptic curves.

The bottleneck for ordinary elliptic curves arises as follows. To obtain a cryptographically secure instance of the isogeny class-group action, we require the (ordinary) elliptic curves in question to have endomorphism rings with large class-groups; on the other hand, to utilise currently available higher-dimensional tools, we also require the curves to have large 2^n-torsion. These requirements appear to be fundamentally at odds with each other for ordinary elliptic curves with current knowledge.

Contents

1 Introduction **1**

1.1 A brief history of isogeny-based cryptography . . . 2
1.2 Cryptography from group actions . . . 4
1.3 Isogeny problems and hardness assumptions . . . 5
1.4 The isogeny class-group action . . . 7
1.5 Clapoti and computing the class-group action in polynomial time . . . 8
1.6 Goals of this work . . . 9

2 Higher-dimensional Isogenies **12**

2.1 Isogenies between elliptic curves . . . 12
2.2 Isogenies between principally polarised Abelian varieties . . . 16
2.3 Embedding isogenies: Kani's Lemma and Zarhin's Trick . . . 21
2.4 Computing higher-dimensional Isogenies . . . 26
2.5 Application: Breaking SIDH by embedding isogenies . . . 29
2.6 Application: Computing the isogeny class-group action . . . 30
2.7 Algorithms for Clapoti . . . 33
2.8 Consequences for isogeny-based cryptography . . . 34

3 Class-groups of Orders in Imaginary Quadratic Number Fields **35**

3.1 Orders . . . 35
3.2 The ideal class-group of an order . . . 38
3.3 Explicit representation of ideals . . . 43

4 Some experiments: Clapoti for ordinary elliptic curves **45**

4.1 Clapoti with 2-dimensional 2-isogenies . . . 45
4.2 Rational points and field extensions . . . 47
4.3 The CM method for constructing ordinary elliptic curves . . . 48
4.4 Large Class-Groups from walking down the Volcano . . . 50
4.5 Implementing the Class-Group . . . 51
4.6 Lattice points to isogenies . . . 53
4.7 Detecting split isogenies . . . 55
4.8 Results, Conclusion and Outlook . . . 56

References **59**

Acknowledgements

I am extremely grateful for Luca De Feo's extraordinary commitment to supervising this thesis. The countless hours of discussion and incredible patience for my questions have made this thesis, and my time at IBM Research, the absolute highlight of my entire degree. I also thank Ueli Maurer for being my ETH advisor and facilitating this externally written thesis.

Besides Luca, I would like to thank Sina Schaeffler, Giacomo Borin and Andrea Basso for including me in the isogeny group at IBM as one of their own; being in such close proximity to cutting edge research has been very inspiring and I have learned a great deal more from our weekly group meetings than I could have ever imagined. I could not have hoped for a more enjoyable and supportive working environment.

I am also grateful for the help of Pierrick Dartois throughout my thesis, especially for all the explanations and intuitions on higher-dimensional isogenies.

I also thank the wider Foundations of Cryptography group at IBM for immediately welcoming me so warmly and for always making interesting conversation over lunch. In particular Sebastian Faller and Simon Rastikian, who could always be convinced to take a coffee break.

I thank IBM Research for providing me an office to work from and for paying my attendance fee to EuroCrypt 2024 hosted here in Zürich. Visiting this conference and the preceding Isogeny Brainstorm Days workshop introduced me to the wider community, giving a face to many of the authors of papers I had been reading throughout my thesis.

I thank Kenny Paterson for igniting my general interest in cryptography, and introducing me to isogeny based cryptography in the first place. I do not think my studies would have taken the direction they did without the semester project on isogenies suggested by Kenny.

Finally, I thank my family and friends for their unconditional support for all my endeavours and their delightful company throughout my studies.

1
Introduction

Since Shor's 1994 discovery of polynomial time quantum algorithms for integer factorisation and computing discrete logarithms [71], there has been an effort within public-key cryptography to develop schemes whose security does not rely on the hardness of these two problems. Indeed, until then, almost all widely used schemes relied on these two hardness assumptions; most notably RSA and the Diffie-Hellman key exchange[1]. This effort is known as *Post-Quantum* cryptography.

Cryptography based on the hardness of finding an *isogeny* between elliptic curves is known as *isogeny-based* cryptography and is conjectured to be resistant against attacks from quantum computers. A main selling point of isogeny-based cryptography relative to other fields in post-quantum cryptography (lattices, codes, multivariate, hash-based), is that isogeny-based schemes are very compact. For example, at the time of publishing in 2020, the combined size of public key and signature of the SQISign [47] signature was an order of magnitude smaller than other post-quantum candidates in [46]; with SQISign's signature sizes almost being halved again in SQISign-HD [18] for the same security level, the SQISign family continues to be the most compact post-quantum signature schemes. Conversely, isogeny-based schemes are generally slower than their counterparts.

A particularly powerful construction is one of a *cryptographic group action*. In a sense, this generalises usual group exponentiation. Indeed, suppose G is a cyclic group of order N. Then $E = (\mathbb{Z}/N\mathbb{Z})^\times$ is a multiplicative group that acts on G by exponentiation: for e in E and g in G we define $e \cdot g = g^e$. Now replacing E with an *arbitrary* group (not tied to G in any way) and G with an arbitrary set, we can obtain a general group action. We can formulate analogous hardness problems to the discrete logarithm problem and the computational Diffie-Hellman for a general group action to obtain the notion of *cryptographic* group actions. From these, a whole host of primitives can be built [25]. Importantly, the study of isogenies gives us a conjectured cryptographic group action via the *isogeny class-group action*. The groups in the cryptographic group action framework are commutative, and so are susceptible to the sub-exponential time attacks of Childs-Jao-Soukharev [44] which employ Kuperberg's algorithm to solve linear hidden shifts [81]. Consequently, these schemes to have larger-than-expected security parameters.

Until Page and Robert's CLAPOTI [2][2] (CLass group Action in POlynomial TIme), a polynomial time evaluation of the isogeny class-group action has remained just out of reach, despite receiving a lot of attention from the community [41, 43, 48, 49]. We stress *polynomial time*, since the standard strategy has traditionally been to perform sub-

[1] More precisely, the security of the DH key-exchange relies on the computational Diffie-Hellman problem. However, there exists a general quantum reduction between these two problems [63–65], and a *de facto* general classical reduction [66, 67].

[2] Stylised with SMALL-CAPS in this way by Page and Robert.

exponential pre-computation steps to arrive at instantiations of an isogeny class-group action, which are then themselves practically efficient. These pre-computation methods, involving computing the structure of class-groups of imaginary quadratic number fields, cannot be scaled to higher security parameters as may be required by improvements of attacks in the style of [44]; they also do not change the fact that, regardless of practical efficiency, we did not have an asymptotically polynomial time algorithm for evaluation of the isogeny class-group action before the advent of Clapoti.

1.1 A brief history of isogeny-based cryptography

In 1997 Couveignes [38], and then independently in 2006 Rostovstev-Stolbunov [39], discovered that a key-exchange could be derived from the isogeny class-group action on a set of (isomorphism classes of) *ordinary* elliptic curves. Also in 2006, Charles-Goren-Lauter [40] discovered a hash function from walks through *supersingular* isogeny graphs, further cementing the value of studying isogenies for cryptography and marking the beginning of a new field of post-quantum cryptography.

Unfortunately, the endomorphism rings, and therefore the class-groups, of ordinary elliptic curves are commutative and so schemes derived from their actions are susceptible to quantum sub-exponential time attacks by reducing the problem of inverting the group action to a hidden linear shift in the class-group [44]. In fact, by the ideal-to-isogeny correspondence, these methods apply more generally to the *isogeny problem* between ordinary elliptic curves. This inspired De Feo-Jao to construct a key-exchange *SIDH* (Supersingular Isogeny Diffie-Hellman) [26] from isogenies between supersingular elliptic curves; because the endomorphism rings of supersingular curves are indeed non-commutative (they are quaternion algebras).

In the next decade, a lot of work was published on isogney-based cryptography with a second flagship construction *CSIDH* (Commutative-SIDH, pronounced *Sea Side*[3]) emerging in 2018 [41]. It is an effective commutative group action on (isomorphism classes of) supersingular elliptic curves, from which a key-exchange can immediately be instantiated; and in following works, signature schemes *SeaSign* [42] and *CSI-FiSh* [43].

Being a commutative group action, the Childs-Jao-Soukharev sub-exponential quantum attacks still apply to the CSIDH action; however, CSIDH is efficient enough to remain practical even at large parameter choices. More concretely, the proof-of-concept `C` implementation in [41] of the non-interactive Diffie-Hellman key-exchange derived from CSIDH with 128 classical bits of security and 40 bits of quantum security [90, 91] takes around 80 milliseconds for an exchange. The current state-of-the-art (constant time) implementation of CSIDH is in CTIDH [70]. It is worth noting that exact estimates on the bits of quantum security are hotly debated, due to the different time/memory tradeoffs and which simplifying assumptions are made on the type of quantum gates used. In any case, both CSIDH and CTIDH are significantly faster than a Couveignes-Rostovstev-Stolbunov type action, despite the speedups from De Feo-Kieffer-Smith in [72].

[3] Prompting many ocean related names in isogeny-based cryptography.

In parallel to CSIDH, SIDH continued to be worked on and advanced to round 3 of the NIST post-quantum standardisation process in 2020 as an alternate candidate for key encapsulation (now under the name SIKE) [45]. Independently from the frameworks of SIDH and CSIDH, the isogeny-based signature scheme *SQISign* [47] was developed, and then submitted to the NIST post-quantum standardisation process as a candidate for digital signatures in 2023 [46].

In 2022, work by Castryck-Decru led to the (re)discovery of *Kani's Lemma* from [20, Proof of Th. 2.6, pg. 100]. Together with some knowledge of the endomorphisms of the domain curve, use of this lemma resulted in a devastating attack on SIDH [15] by the same authors, and also independently by Maino et al. in [29]. Robert realised the underlying structure of these ideas, and combined Kani's Lemma and *Zahrin's trick* into a generic framework for efficiently representing and computing non-smooth degree isogenies by embedding them into smooth-degree isogenies between products of elliptic curves [16]; most notably without knowledge of any additional endomorphisms on the starting curve. We call this framework *Robert's embedding lemma* in this work.

Products of elliptic curves somehow have "higher dimension", and so these isogenies between them are often referred to as *higher dimensional* isogenies. As noted in the discussion on follow-up works in [15], the work of Robert dashed any hopes of recovering any security of SIDH by tweaking parameters, changing starting curves or performing SIDH in higher dimensions.

We stress that Robert's embedding lemma forms an *attack* only against SIDH and its variants; however, since it is efficient, it can be used *constructively* in other contexts, to create new schemes or improve existing ones. Notable "HD-fications" of existing schemes are *SQISignHD* [18] (and its newer 2-dimensional variants [53–55]) and *SCALLOP-HD* [49] (an effective commutative group action on supersingular elliptic curves); examples of new applications are in the *FESTA* PKE [50] and more recently, Basso's *POKE* framework which allows one to instantiate a range of primitives (PKEs, Split KEMs, and OPRFs) [51] from isogeny hardness assumptions. From this small sample of new work, we conclude that the SIDH attacks, also known as *torsion-point attacks*, have ultimately led to a sustainable growth in the isogeny-based cryptography community.

Using the same techniques, but with a different philosophical approach, Page and Robert have recently published Clapoti [2], the first generic algorithm for computing the class-group action of both ordinary and supersingular elliptic curves in polynomial time, with no pre-computation requirements. This removes the scalability problems of CSIDH and its variants. This will be discussed in subsection 1.5.

In subsection 1.6 we discuss the main goal of this thesis, namely to asses whether Clapoti, an asymptotically polynomial time algorithm, can be leveraged to concretely instantiate an efficient algorithm that evaluates the isogeny class-group action.

1.2 Cryptography from group actions

The isogeny class-group action arises very naturally in isogeny-based cryptography and is a promising candidate for a conjectural *cryptographic group action*. That is, a group action with certain hardness properties, from which we can build many primitives [25], in particular a non-interactive key exchange which we present here.

Let G be a commutative group acting on the set X, with both X, G finite. We say G acts *effectively* on X, if given some encodings of X, G there exist polynomial time (in $\log(|X| + |G|)$) algorithms for

(i) Membership testing for both X, G. That is, given a bit string, determine whether it constitutes a valid encoding of an element of X or G.
(ii) Equality testing in G. That is, given two encodings, decide whether they encode the same element.
(iii) Unique representations for X. That is, given an element x, compute a unique bit string encoding of x. We note that this implies efficient equality testing in X.
(iv) Random sampling in G. That is, to sample from G randomly according to a given a distribution.
(v) Computing the group law in G. That is, for any pair of elements g, h compute gh, and given g compute g^{-1}.
(vi) Computing the group action $G \curvearrowright X$. That is, given two elements g, x compute gx.

For example, a vector space acts efficiently on itself by translations if addition of vectors is efficient. It is exactly this prototypical example, that inspired Couveignes' naming of the following problems, posed under the assumption that the algorithms of (i)-(vi) are available.

Problem (Vectorisation)**.** *Given x, y in X find g in G (if it exists) so that $y = gx$.*

Problem (Parallelisation)**.** *Given x, gx, y in X compute $z = gy$.*

When the action is free and transitive we say that X is a *G-homogeneous space* or *G-torsor* and solutions to the vectorisation problem exist and are unique. Therefore we henceforth assume that G acts freely and transitively. Note that, inspired by Stolbunov's thesis [82], the vectorisation problem is known as the *Group Action Inverse Problem* in the CSI-FiSH paper [43].

These problems, vectorisation and parallelisation, are group-action analogues to the discrete logarithm and computational Diffie-Hellman problem respectively. The vectorisation problem trivially reduces to the parallelisation problem. Indeed given a vectorisation oracle, we can recover g from a (x, gx, y) parallelisation challenge to compute $z = gy$ in one oracle call and one group action computation. The converse reduction is given in the quantum setting by [63–65]. Since we are designing schemes for the *post-quantum* domain, we consider these problems to be computationally equivalent in the following.

When the parallelisation problem (and therefore vectorisation) is hard (i.e. there exists no known algorithm to solve it, that runs in time polynomial in $\log(|G| + |X|)$), we say that X is a *hard* G-homogeneous space, language coined by Couveignes in [38].

We now present a Diffie-Hellman style non-interactive key-exchange from an effective group action on a hard homogeneous space. Suppose Alice and Bob would like to communicate over an insecure channel. Let G, X and x in X be public along with all algorithms to compute points (i)-(vi). Alice and Bob randomly sample a, b respectively from G and call these their private keys. They then compute ax, bx to form their public keys. Finally, they (non-interactively) derive a shared secret by computing $k = b(ax) = a(bx)$ from the other party's respective public key. Note that the group must be commutative!

We see that for this primitive, we need at least: algorithm (iv) to sample the private keys; algorithm (vi) to compute the shared secret; algorithm (iii) to put the shared secret into a common canonical form; algorithm (i) for public key validation. On the latter: since the group action is transitive, every element of X is a valid public key and so to verify that Bob's public key is well-formed, Alice must (can) only verify that bx lies in X.

A direct attack on the shared secret requires computing $k = abx = bax$ from (x, ax, y), which is exactly the parallelisation problem. From our choice to regard the parallelisation and vectorisation problem to be equivalent, it is now justified to investigate attacks on the vectorisation problem for cryptanalysis of such a key-exchange.

As outlined in Sec 3.1 of [25], the best classical attacks against the vectorisation problem (called *one-wayness* in the paper) are meet-in-the-middle attacks, with the best due to Galbraith, Hess, and Smart [68] having complexity $O(\sqrt{G})$. This means, to achieve λ bits of security, the group must be of size $2^{2\lambda}$. Whilst the main feature of [44] is to reduce the *isogeny class-group* action to the linear hidden shift problem, their methods (Lemma 5.1) can be applied to any commutative group action and have complexity $O(\exp(\sqrt{\log(G)}))$, because the underlying tool is Kuperberg's (linear hidden shift) algorithm [81].

1.3 Isogeny problems and hardness assumptions

To understand the hardness assumptions underpinning cryptographic schemes based on isogenies, we should first describe what an isogeny is. A more precise treatment of elliptic curves and isogenies will appear in Subsection 2.2 when we generalise to higher dimensions, so, for now, we collect only the most important features without explicit references. All results can be found in Silverman III.4 [28] or Sutherlands' Lectures [7].

An *elliptic curve* E over k (in short Weierstrass form) is a subset of $\mathbb{P}^2(\overline{k})$ described by the points $(X : Y : Z)$ satisfying $ZY^2 = X^3 + aXZ^2 + bZ^3$ with a, b in k and $4a^3 + 27b^2 \neq 0$. Every elliptic curve over k has a group structure whose defining maps are given by polynomials with coefficients in k. This makes the k-*rational* points $E(k)$, namely the points on E with coordinates in k, a subgroup of E.

A map $\varphi\colon E \to E'$ between two elliptic curves is called an *isogeny* if it is a non-zero group morphism whose components $\varphi = (\varphi_X : \varphi_Y : \varphi_Z)$ are described by polynomials in X, Y, Z. We say that φ is *defined over* k if each component is a polynomial with coefficients in k.

These components can be put into a standard form $\varphi = (\psi_1\psi_4 : Y\psi_3\psi_2 : \psi_2\psi_4)$ whereby ψ_i are polynomials in only X and ψ_1, ψ_2 and ψ_3, ψ_4 are each coprime pairs. From this form, we can attach an integer constant called the *degree* of φ, and it is equal to the degree of ψ_2 as a polynomial (this is also the degree of ψ_4).

Moreover, we associate to every isogeny $\varphi\colon E \to E'$ its *reverse*[4] $\widetilde{\varphi}\colon E' \to E$, an isogeny which satisfies $\widetilde{\varphi}\varphi = [\deg(\varphi)]_E$ and $\varphi\widetilde{\varphi} = [\deg(\varphi)]_{E'}$. Here, $[N]_E$ denotes the *multiplication-by-N* isogeny on E, namely the map that takes a point e in E to $Ne = e + \cdots + e$.

The *kernel* of an isogeny is equal to its kernel as a group morphism, that is $\ker(\varphi) = \{e \in E \mid \varphi(e) = 0_{E'}\}$. Isogenies have finite kernels, and conversely for every finite subgroup $G \subseteq E$ there exists an isogeny $\varphi\colon E \to E'$ with kernel G. We denote the kernel of $[N]_E$ by $\ker([N]_E) = E[N]$ and call it the *N-torsion*. It comprises all the points in E whose order divides N.

If $\varphi(e) = 0_{E'}$ then $\widetilde{\varphi}\varphi(e) = \deg(\varphi)e = 0_E$ and so the kernel of a degree-d isogeny φ must lie in $E[d]$. Moreover, $E[d] \cong \mathbb{Z}/d\mathbb{Z} \times \mathbb{Z}/d\mathbb{Z}$ when d is prime to $\mathrm{char}(k)$, and so the kernel of any isogeny is spanned by at most two generators.

Finally, an isogeny is *separable* if the size $|\ker(\varphi)|$ of its kernel equals its degree $\deg(\varphi)$. This means every separable isogeny is uniquely defined by its kernel up to isomorphism.

The existence of the reverse map $\widetilde{\varphi}$ makes the notion of *isogenous* an equivalence relation; indeed, if there exists an isogeny $E \to E'$, the reverse ensures the existence of an isogeny $E' \to E$. This leads us to consider the first isogeny

Problem. *Given two elliptic curves E, E' defined over a finite field k, decide whether they are isogenous over k.*

This problem can be solved in (classical) polynomial time using two results. Indeed, firstly by Theorem 1.(c) of [59], two elliptic curves are isogenous by isogenes defined over k if and only if they have the same number of k-rational points; and secondly the Schoof-Elkies-Atikin (SEA) point counting algorithm [60] counts k-rational points on elliptic curves in polynomial time, giving us an exact tool for *deciding* this isogeny problem.

So we would like to convert this problem into a *computational* one; essentially a problem that asks to produce an isogeny between E and E'. Since an isogeny is a function, it will be described by a program for which there may not exist a canonical form. In the literature, we speak of an *isogeny representation* for such a program [89].

A naïve representation is simply by the polynomial maps $\varphi = (\varphi_X : \varphi_Y : \varphi_Z)$, perhaps in standard form $\varphi = (\psi_1\psi_4 : Y\psi_3\psi_2 : \psi_2\psi_4)$. These can be evaluated in the time it takes

[4] This is usually referred to as the *dual* in the literature. We will see why we used the term *reverse* later.

to evaluate a polynomial of degree $\deg(\psi_2) = \deg(\varphi)$. That is, linearly in the degree of the isogeny.

Conversely, recall that every separable isogeny is uniquely defined by its kernel up to isomorphism. Vélu's formulae [61] give an algorithm to compute a (separable) isogeny from its kernel, and so another way to represent an isogeny is by a description of its kernel. Vélu's formulae also have complexity $O(\deg(\varphi))$ to evaluate an isogeny.

In special cases, the `VeluSqrt` methods of [69] yield an isogeny representation of complexity of $O(\sqrt{\deg(\varphi)})$. Namely, when $\ker(\varphi) = \langle P \rangle$ is generated by a point P of prime order in $E(F_q)$. Then `VeluSqrt` can evaluate φ on points in $E(F_q)$ in the advertised time and recover (the equation for) E'.

This discussion prompts us to formalise the second isogeny hardness

Problem. *Given two elliptic curves E, E' defined over a finite field k, find an isogeny $E \to E'$ that is defined over k and return an isogeny representation of this isogeny.*

An up-to-date resource for current isogeny problems and the status of the best algorithms to solve them is maintained at [58].

1.4 The isogeny class-group action

Due to Deuring [28, Cor. 9.4], it is well-known that the endomorphism ring $\mathrm{End}(E)$ of an elliptic curve E/F_q, comprising isogenies from $E \to E$ (and the zero map on E), is isomorphic to either (i) an order within an imaginary quadratic field or (ii) a maximal order within a quaternion algebra. In the first case, we say that E has *complex multiplication* by $\mathcal{O} \cong \mathrm{End}(E)$ and call the curve *ordinary*; and in the second case, we say that E has *quaternionic multiplication* and call E *supersingular*. A precise description of orders in quadratic imaginary number fields will come in Subsection 3.1.

For ordinary elliptic curves, we can form a naturally arising action in the following way. Let E be an ordinary elliptic curve, $\mathrm{End}(E)$ its endomorphism ring, isomorphic to $\mathcal{O}$ an order inside the imaginary quadratic field $\mathbb{Q}(\sqrt{D})$. For every ideal I of $\mathcal{O}$ we can write down the intersection $H(I)$ of all kernels of endomorphisms in I. As the intersection of finite subgroups, $H(I)$ is a finite subgroup and so gives rise to an isogeny $\varphi_I \colon E \to E'$ with kernel $H(I)$ where E' is uniquely defined by $H(I)$ up to isomorphism. Notably, if $I = \alpha\mathcal{O}$ is a principal ideal, then $H(I) = \ker(\alpha)$ (because $\mathcal{O}$ contains 1), so $\ker(\varphi_I) = \ker(\alpha)$ and $E' \cong E$ (because we know that $\alpha \colon E \to E$).

This suggests to define an action of $G = \{\text{ideals of } \mathcal{O}\} / \{\text{principal ideals of } \mathcal{O}\}$ on isomorphism classes of elliptic curves by mapping the pair $[I], [E]$ to the class of the codomain $[E']$ of the arising isogeny $\varphi_I \colon E \to E'$. We write $[I] \cdot [E] = [E']$. Since principal ideals correspond to endomorphisms, the (isomorphism class of the) target curve does not depend on the representative of I. Alas, there are still two obvious hurdles to making this suggestion precise. Firstly, how we make G into a group; and secondly, determining whether the (isomorphism class of the) target curve E' of an isogeny $\varphi_I \colon E \to E'$ also has an endomorphism ring $\mathrm{End}(E') \cong \mathcal{O}$, so that the action of G is still well-defined on

E'. After all, we must define the set on which G acts: if $\mathrm{End}(E') \ncong \mathrm{End}(E) = \mathcal{O}$, it is not clear how an ideal class $[J]$ in G should act on $[I] \cdot [E] = [E']$.

The first problem, of turning G into a group is solved by using the language of the *ideal class-group*. First, we only consider *invertible* $\mathcal{I}(\mathcal{O})$ ideals, namely ideals $I \subseteq \mathcal{O}$ for which there exist another ideal $J \subseteq \mathcal{O}$ so that $IJ \subseteq \mathcal{O}$ is principal. On $\mathcal{I}(\mathcal{O})$, we define the "group law"[5] by usual ideal multiplication. Clearly every principal ideal is invertible by this definition and the set of principal ideals $\mathcal{P}(\mathcal{O}) \subseteq \mathcal{I}(\mathcal{O})$ form a subgroup. Finally, to make this into a group proper (with unique inverses), we define the *ideal class-group* as the quotient $\mathrm{Cl}(\mathcal{O}) = \mathcal{I}(\mathcal{O})/\mathcal{P}(\mathcal{O})$.

The second problem, of determining the endomorphism ring of E' when $\varphi_I \colon E \to E'$ is an isogeny with kernel $H(I)$ from an invertible ideal I is answered by Waterhouse in [73, Th. 4.5]. It tells us that indeed, $\mathrm{End}(E) \cong \mathrm{End}(E')$. Isogenies between elliptic curves with isomorphic endomorphism ring are called *horizontal* in general. Not only this, it tells us that the class-group of $\mathcal{O}$ acts freely and transitively on $\mathrm{Ell}_q(\mathcal{O})$, the set of isomorphism classes of elliptic curves E' over F_q with endomorphism rings $\mathrm{End}(E') \cong \mathcal{O}$. In particular $\left|\mathrm{Ell}_q(\mathcal{O})\right| = |\mathrm{Cl}(\mathcal{O})|$.

This action can be extended to supersingular curves using the language of *orientations*. The endomorphism ring of a supersingular elliptic curve is a quaternion algebra of rank-4 viewed as a $\mathbb{Z}$-module. An *$\mathcal{O}$-orientation* of a supersingular elliptic curve by an order $\mathcal{O}$ of an imaginary quadratic number field is an injection $\iota \colon \mathcal{O} \hookrightarrow \mathrm{End}(\mathcal{O})$.

Stated by Colò and Kohel [74], and then proved by Onuki [75], the isogeny class-group action of $\mathcal{O}$ on the set of $\mathcal{O}$-oriented supersingular elliptic curves $\mathrm{Ell}_q^{\mathrm{SS}}(\mathcal{O})$ is free and transitive; making $\mathrm{Ell}_q^{\mathrm{SS}}(\mathcal{O})$ a homogeneous space for $\mathrm{Cl}(\mathcal{O})$.

Due to Siegel [79], it is known that $|\mathrm{Cl}(\mathcal{O}_D)| \sim \sqrt{-D}$. So to achieve 128 bits of classical security, we require D to have 512 bits: firstly the class-group $\mathrm{Cl}(\mathcal{O}_D)$ has size 256 bits, and the attack of Galbraith, Hess and Smart [68] halve the bits again.

The vectorisation problem of the isogeny class-group action is closely related to the isogeny problem. Indeed, given an oracle to compute the vectorisation problem, we can use Page and Robert's Clapoti to compute an efficient isogeny representation of an isogeny between the two challenge curves.

1.5 Clapoti and computing the class-group action in polynomial time

The CSIDH group action is actually a *restricted* cryptographic group action (in the language of [25]), meaning the action can only be computed efficiently on a generating subset $g_1, \ldots, g_n$. Computing gx for arbitrary $g = g_1^{e_1} \cdots g_n^{e_n}$ is expensive for large exponents e_i.

A method to remedy this problem, first employed by CSI-FiSh, is to use relations in G to rewrite $g = g_1^{e_1} \cdots g_n^{e_n}$ as a somehow "smaller" product $g = g_1^{e_1'} \cdots g_n^{e_n'}$ using lattice

[5] This is not a well-defined group, because there is not a unique neutral element for this group law, and so inverses are not unique either.

reduction techniques on the lattice of relations satisfied by elements in $G = \mathrm{Cl}(\mathcal{O})$.

This is accounted more precisely by Panny in [57]. We can view every element $g = g_1^{e_1} \cdots g_n^{e_n}$ in G as an integer vector $(e_1, \ldots, e_n)$. By computing the (structure of the) class-group, we obtain the *relation lattice* $\Lambda \subseteq \mathbb{Z}^n$, comprising vectors $(e_1, \ldots, e_n)$ so that $g_1^{e_1} \cdots g_n^{e_n} = e$ is the neutral element.

Now doing a sub-exponential lattice reduction (e.g. via the BKZ algorithm), we obtain a shorter basis of Λ, and given an element $g = g_1^{e_1} \cdots g_n^{e_n}$ we can find a shorter representation $g = g_1^{e'_1} \cdots g_n^{e'_n}$ by solving the closest vector problem in Λ. Indeed, if $(\lambda_1, \ldots, \lambda_n)$ in Λ is the closest vector to $(e_1, \ldots, e_n)$ we obtain $g = g_1^{(e_1-\lambda_1)} \cdots g_n^{(e_n-\lambda_n)}$ (because $g_1^{\lambda_1} \cdots g_n^{\lambda_n} = e$). We remark that solving the closest vector problem can be done in polynomial time, and the distance to the closest Λ vector is dictated by the quality of the reduced basis of Λ.

From this we conclude that the better the quality of the reduced basis of Λ (i.e. the shorter the basis vectors), the smaller $e_i - \lambda_i$ will be, and the cheaper it is to evaluate g. In other words, there is a tradeoff between finding a good quality basis of Λ and evaluating the action: the more time spent on lattice reduction, the cheaper the evaluation and vice versa. Even with a quantum computer, there is no simultaneous polynomial time solution to this conflict.

Nevertheless, for concrete instantiations (and not interest in asymptotic complexity), we see that it is advantageous allocate as many resources as possible to the class-group computation and lattice reduction steps, as these are merely pre-computations and only need to be done once. CSI-FiSh [43] did exactly this for the CSIDH-512 parameter set, and pre-computed the structure of $\mathrm{Cl}(\mathcal{O})$ with a reduced basis, where $\mathcal{O}$ is an order in $\mathbb{Q}(\sqrt{D})$ with D of 512-bits. This took around 50 core-years. Computing larger class-groups to obtain better security appears to be out of reach at the current time since this has classical sub-exponential cost (and we do not have sufficiently powerful quantum computers to execute the polynomial time quantum algorithms). We stress that for asymptotic analysis, we still have the non-polynomial lattice reduction step to consider.

Page and Robert's Clapoti is the first truly polynomial-time evaluation of the isogeny class-group action. It does not require any (significant) pre-computation and so can be scaled to any parameters of choice. More precisely, Clapoti does not need to know the structure of the class-group to operate and is an (efficient) cryptographic group action.

1.6 Goals of this work

The goal of this thesis is to asses the practical feasibility of employing Page and Robert's Clapoti [2] algorithm to compute the isogeny class-group action on ordinary elliptic curves. Their main result was to deliver an *asymptotic* polynomial time algorithm to compute the isogeny class-group action on any elliptic curve (ordinary or supersingular), but by demanding new tools for computing higher-dimensional isogenies and specific algorithms for manipulating ideals in the ideal class-group, their paper leaves open the

question of whether CLAPOTI can be used for *concrete*, efficient evaluation of the isogeny class-group action.

The general technique is of CLAPOTI is to embed isogenies arising from the (isogeny) class-group action into isogenies of higher dimension. Indeed, if $\varphi_I \colon E \to E_I$ is the isogeny arising from the action of the class $[I]$ in the class-group $\mathrm{Cl}(\mathcal{O})$, CLAPOTI will embed this into a new isogeny $\Phi_I \colon E \times \cdots \times E \to (E_I = E'_1) \times \cdots \times E'_n$ of dimension n which it hopes is easier to compute. The source curves are all E, but the target curves $E'_1, \ldots, E'_n$ may be different. We see that intuitively, φ_I was embedded into the "first component" of Φ_I.

Fortunately, the dimension n of the higher-dimensional isogeny Φ_I into which is φ_I embedded remains modest: n is only ever $2, 4$ or 8. Unfortunately, at the time of writing, the only available implementation for higher-dimensional isogenies can compute isogenies of dimension $n = 2$. This places restrictions on which isogenies φ_I can be computed, and requires careful manipulation of the ideals in the class-group (e.g. finding other ideals J, representing the same class $[I] = [J]$ of whom we hope that the isogeny $\varphi_J \colon E \to E_J \cong E_I$ has "better" properties).

Using the library of [3], that computes 2-dimensional isogenies, I wrote an almost complete Sage implementation to compute the isogeny class-group action on ordinary elliptic curves. Falling just short of a complete evaluation, the work of this thesis justifies the claim that CLAPOTI cannot be used to practically evaluate the isogeny class-group action on ordinary elliptic curves, due to limitations in the framework used to compute higher-dimensional isogenies.

More precisely, regardless of the implementation, current understanding of practically computing higher-dimensional isogenies between products of elliptic curves $E_1 \times \cdots \times E_n \to E'_1 \times \cdots \times E'_n$ requires the source curves to have a specific structure: it requires the E_i to have sufficiently large rational 2^n-torsion. This requirement is fundamentally at odds with also requiring the (ordinary) source curves E_i to have endomorphism rings $\mathcal{O}_i$ with cryptographically large class-groups $\mathrm{Cl}(\mathcal{O}_i)$.

Recall that we begin with an ordinary elliptic curve E that has endomorphism ring $\mathcal{O}$. Then $\mathrm{Cl}(\mathcal{O})$ acts on $\mathrm{Ell}_q(\mathcal{O})$ via isogenies $\varphi_I \colon E \to E_I$ from the ideals I representing elements $[I]$ of the class-group. For this to be cryptographically useful, $\mathrm{Cl}(\mathcal{O})$ must be large. Now CLAPOTI tries to embed φ_I into e.g. 2-dimensions $\Phi_I \colon E \times E \to E_I \times E'_I$. To compute this, we require E to have rational 2^n-torsion, demonstrating that both requirements (endomorphism ring has large class-group, and sufficient 2^n-torsion) are necessary.

From this, however, we may draw the following positive conclusion. The methods of CLAPOTI should be practically efficient when acting on supersingular elliptic curves. They are not burdened with the conflict between having sufficient 2^n-torsion and having an endomorphism ring with a large class-group. We explain why in Subsection 4.8.

A more detailed description of the results and conclusions of this work is found in Subsection 4.8. There we list the implementations of this thesis and give examples of their capabilities. The source code of these implementations is available at

`github.com/rrueger/IsogenyClassGroupComputation`

In Section 2, we introduce the necessary tools from arithmetic geometry to formulate the CLAPOTI argument and specialise its use to ordinary elliptic curves. We demonstrate how Kani's Lemma (Lemma 2.4), a central part of the higher-dimensional theory, arises naturally when generalising isogenies between elliptic curves to isogenies between products of elliptic curves.

In Section 3, we supply the necessary theory to write algorithms for the ideal class-group. In particular, we give a computationally advantageous explicit representation of ideals in the class-group.

In Section 4, we combine the theory of the previous two sections and provide methods for finding suitable ordinary elliptic curves with the desired features of sufficient rational 2^n-torsion and endomorphism ring with large class groups (in an cryptographically unsafe way). On these toy examples, test our implementations.

2
Higher-dimensional Isogenies

In this chapter we describe how isogenies between elliptic curves can be computed efficiently by going into higher dimensions.

We give a high-level overview in the following. The tools we use to *compute* a (separable) isogeny (in any dimension r) generally take as input the kernel as a set of generators and output equations of the codomain (up to isomorphism) and an explicit formula to compute the isogeny. The best known algorithms of this type are variations on *Vélu's formulae* [61] and run in asymptotic linear time in the degree $d^r = |\ker(\varphi)|$. That is, exponential in the size $r\log(d)$ of the degree.

Therefore, traditionally (without higher-dimensional methods), the only computable isogenies between elliptic curves of large degree were those with *smooth* degree. Indeed, if $d = p_1 \cdots p_n$ with $p_i \leq B$ possibly repeated, then the d-isogeny φ can be decomposed into n isogenies $\varphi = \varphi_1 \cdots \varphi_n$ of degree p_i. Each of these can be computed in time linear in p_i and so the total computation is of cost $O(p_1 + \cdots + p_n) = O(nB) = O(\log(d)B)$.

The insight of Kani's Lemma is that an isogeny of arbitrary (read *non-smooth*) degree d can be embedded into a *smooth* degree isogeny in higher dimensions in such a way that the values and codomain of the original isogeny can be recovered after evaluating the higher-dimensional isogeny.

This gives us the slogan: smooth-degree isogenies can be computed efficiently, but non-smooth-degree isogenies cannot be computed efficiently; however embedding a non-smooth-degree isogeny into higher dimensions allows us to compute the starting isogeny efficiently.

2.1 Isogenies between elliptic curves

We first give a short primer on projective varieties. We begin with affine space. Let k be a field and $\overline{k}$ its algebraic closure. This distinction is important, and to be explicit we will always write $\overline{k}$ instead of simply requiring k to be algebraically closed.

We denote by $\mathbb{A}^n(\overline{k})$ the *affine* $\overline{k}$-space, it is a topological space whose underlying set is $\overline{k}^n$ which is endowed with the *Zariski* topology. This is the topology generated by the *closed* sets $Z(f) \stackrel{\text{def.}}{=} \{x \in \mathbb{A}^n(\overline{k}) \mid f(x) = 0\}$ where f is a polynomial in $\overline{k}[x_1, \ldots, x_n]$. Closed subsets of $\mathbb{A}^n(\overline{k})$, that is, arbitrary intersections of $Z(f_i)$, are called *affine varieties*.

We define the *projective* $\overline{k}$-space $\mathbb{P}^n(\overline{k})$ to be the equivalence classes of $\mathbb{A}^{n+1}(\overline{k}) - \{0\}$ modulo the equivalence $x \sim y$ if there exists a λ in $\overline{k}$ such that $\lambda x = y$, that is, modulo scalar equivalence. The points in $\mathbb{P}^n(\overline{k})$ are usually 0-indexed and written as $(x_0 : x_1 : \ldots : x_n)$. Since points are only defined up to scalars, we can only define the zero-locus $Z(f) = \{x \in \mathbb{P}^n(\overline{k}) \mid f(x) = 0\}$ of *homogeneous* polynomials in $\overline{k}[x_0, x_1, \ldots, x_n]$, namely polynomials which satisfy $f(\lambda x) = \lambda^d f(x)$ for some fixed d and all λ in $\overline{k}$. Else, we encounter problems: the polynomial $f(x_0, x_1) = x_0^2 + x_0$ is zero evaluated at $(1, -1)$

but not $(2,-2)$ even though these represent the same point in $\mathbb{P}^1(\overline{k})$. The constant d attached to f in the equation $f(\lambda x) = \lambda^d f(x)$ is the *degree* of f. The *Zariski* topology on $\mathbb{P}^n(\overline{k})$ is the topology generated by the closed sets $Z(f)$. Irreducible closed subsets of $\mathbb{P}^n(\overline{k})$ are called *projective varieties*. Recall that a topological space is *irreducible* if it cannot be written as the union of non-empty proper closed sets (a stronger notion than *connectedness*).

Although *a priori* every affine (projective) variety is the *arbitrary* intersection of some zero-loci $Z(f_i)$, Hilbert's *basis theorem* and *Nullstellensatz* together tell us that every (projective) variety V is *cut out* by finitely many polynomials f_i, that is $V = Z(f_1) \cap \cdots \cap Z(f_n)$.

The full generality of *varieties* (and *schemes*) comes in the language of *locally ringed spaces*. These are topological spaces endowed with a sheaf of rings, whose stalks are local rings. In the following, we show that affine varieties $V = Z(f_1, \dots, f_m) \subseteq \mathbb{A}^n(\overline{k})$ are indeed locally-ringed by their *structure sheaf*, usually denoted $\mathcal{O}_V$ or Γ_V.

This sheaf is given on open sets $U \subseteq V$ as the ring $\mathcal{O}_V(U)$ of *regular functions* $U \to \overline{k}$, namely the ring of functions $U \to \overline{k}$ which can be locally written as the quotient of two polynomials in $\overline{k}[x_1, \dots, x_n]$; addition and multiplication on this ring are given by pointwise addition and multiplication. This assignment $U \mapsto \mathcal{O}_V(U)$ satisfies the *restriction* and *gluing* axioms that a sheaf must satisfy (i) the restriction $f|_W$ of a regular function f in $\mathcal{O}_V(U)$ to an open subset W is again a regular function in $\mathcal{O}_V(W)$; and (ii) regular functions f_i in $\mathcal{O}_V(W_i)$ defined on an open cover W_i of U, compatible on the intersections $W_i \cap W_j$, glue uniquely to a regular function f on V. The *coordinate* ring of V is given by $\mathcal{O}_V(V) \cong \overline{k}[x_1, \dots, x_n]/(f_1, \dots, f_m)\overline{k}[x_1, \dots, x_n]$ and usually denoted by $k[V]$. When V is irreducible, $k[V] = \mathcal{O}_V(V)$ is an integral domain and we define the *field of functions* on V to be $k(V) = \mathrm{Frac}(k[V])$. Note that in $k[V]$ and $k(V)$, the "k" is simply notation, and does not reflect the field $\overline{k}$ over which the variety V is defined (and is why it is not written as $\overline{k}[V]$).

As for any sheaf, the *stalk* $\mathcal{O}_{V,p}$ of $\mathcal{O}_V$ at a point p in V is defined as the limit $\mathcal{O}_{V,p} = \lim_{p \in U} \mathcal{O}_V(U)$. Less abstractly, this is simply the ring of equivalence classes $[(U, f)]$ comprising pairs of open neighbourhoods $U \subseteq V$ of p and regular functions $f\colon U \to \overline{k}$, modulo the equivalence $(U, f) \sim (U', f')$ if there exists an open subset $W \subseteq U \cap U'$ so that $f|_W = f'|_W$. The stalk of an affine variety V at the point p is equal to the localisation of the coordinate ring $\mathcal{O}_V(V)$ by the maximal ideal of regular functions in $\mathcal{O}_V(V)$ vanishing at p [9, 3.19]. Hence $\mathcal{O}_{V,p}$ is a local ring (it is a localisation) and V is locally-ringed by $\mathcal{O}_V$.

In general, a morphism of a locally-ringed spaces $(X, \mathcal{O}_X) \to (Y, \mathcal{O}_Y)$ is given by a pair (φ, φ^*) comprising a continuous map $\varphi\colon X \to Y$ and a *pullback map* φ^* which induces a morphism of rings $\varphi^*_U\colon \mathcal{O}_Y(U) \to \mathcal{O}_X(\varphi^{-1}(U))$ on every open set $U \subseteq Y$ subject to the following conditions (i) the pullback commutes with restrictions, that is, for all open $V \subseteq U \subseteq Y$ we have $\varphi^*_U(f|_V) = \varphi^*_U(f)|_{\varphi^{-1}(V)}$; and (ii) the induced map on stalks $\varphi^*_p\colon \mathcal{O}_{Y,\varphi(p)} \to \mathcal{O}_{X,p}; [(U, f)] \mapsto [(\varphi^{-1}(U), \varphi^*_U(f))]$ is *local*, that is, it sends the maximal ideal of $\mathcal{O}_{Y,\varphi(p)}$ to that of $\mathcal{O}_{X,p}$. A morphism (φ, φ^*) is an isomorphism if φ has

a continuous inverse, and if φ^* induces isomorphisms $\varphi_U^*\colon \mathcal{O}_Y(U) \xrightarrow{\sim} \mathcal{O}_X(\varphi^{-1}(U))$.

With this terminology, we define a *variety* V over k to be a locally-ringed space with a sheaf of $\overline{k}$-algebras $\mathcal{O}_V$ that is covered by *affine patches*. These patches are open sets $U_i \subseteq V$ that are isomorphic to affine varieties $W_i \subseteq \mathbb{A}^n(\overline{k})$ over $\overline{k}$ as locally-ringed spaces. By this definition, projective varieties in $\mathbb{P}^n(\overline{k})$ are in fact varieties. Indeed, the open sets $U_i = \mathbb{P}^n(\overline{k}) - Z(X_i) = \{(x_0 : x_1 : \cdots : x_{i-1} : 1 : x_{i+1} : \cdots : x_n) \mid x_i \in \overline{k}\}$ form an open cover of $\mathbb{P}^n(\overline{k})$, and are each isomorphic to $\mathbb{A}^n(\overline{k})$. Hence for any projective variety $V \subseteq \mathbb{P}^n(\overline{k})$ (a closed subset) we obtain the affine patches $U_i \cap V$ by restricting the isomorphisms $U_i \cong \mathbb{A}^n(\overline{k})$ to $U_i \cap V$. This justifies overloading the previous term "projective variety" to mean a variety that is isomorphic to a closed subset of $\mathbb{P}^n(\overline{k})$.

Since the structure sheaf of a variety is not only a sheaf of rings, but also a sheaf of $\overline{k}$-valued functions (regular functions) any continuous map $\varphi\colon X \to Y$ automatically satisfies the conditions (i) and (ii) that a morphism of locally-ringed spaces must satisfy if $f \circ \varphi$ lies in $\mathcal{O}_X(\varphi^{-1}(U))$ for every f in $\mathcal{O}_Y(U)$ and $U \subseteq Y$ open. Moreover, a morphism of locally-ringed spaces between varieties $\varphi\colon X \to Y$ is automatically an isomorphism if φ has a continuous inverse φ^{-1} and $\varphi^{-1}\colon Y \to X$ is a morphism of locally-ringed spaces.

This immediately implies that every morphism of affine varieties $V \subseteq \mathbb{A}^n(\overline{k}) \to W \subseteq \mathbb{A}^m(\overline{k})$ is given by a map $\varphi = (\varphi_1, \ldots, \varphi_n)$ whose components are regular functions $V \to \overline{k}$. This is why morphisms of varieties are often also called *regular maps* (regular *map*s have regular *functions* as their components), a convention we will adopt for its brevity.

In the projective case we can be a little more concrete. Here a map $V \subseteq \mathbb{P}^n(\overline{k}) \to W \subseteq \mathbb{P}^m(\overline{k})$ is regular if and only if $\varphi = (\varphi_1 : \ldots : \varphi_m)$ with φ_i (homogeneous) polynomials φ_i in $\overline{k}[x_1, \ldots, x_n]$, not arbitrary regular functions. Indeed, we show that every map $\eta = (\eta_1 : \ldots : \eta_m)\colon V \subseteq \mathbb{P}^n(\overline{k}) \to W \subseteq \mathbb{P}^m(\overline{k})$ whose components are regular functions $V \to \overline{k}$ is equal to some $\varphi = (\varphi_1 : \ldots : \varphi_m)$ with φ_i polynomials. Essentially this is a consequence of projective varieties being irreducible by definition, and that we can multiply through by the denominators of the regular functions η_i without changing the value of the function η in projective space.

Following this outline, we first show that a regular function $\zeta\colon V \to \overline{k}$ defined on an projective variety $V \subseteq \mathbb{P}^n(\overline{k})$ is given by the quotient of two polynomials in $\overline{k}[x_0, x_1, \ldots, x_n]$ on the whole of V. Recall that a function $\zeta\colon V \to \overline{k}$ is regular if a every point p in V there exists a neighbourhood $U \subseteq V$ of p and polynomials f_p, g_p in $\overline{k}[x_0, x_1, \ldots, x_n]$ such that $\zeta|_{U_p} = f_p/g_p$. (We previously summarised this as "locally being the quotient of polynomials"). We remark that for this quotient to be well-defined, these polynomials must be homogeneous of the same degree and g_p must be non-zero on U_p. Now consider another point q in V, a neighbourhood U_q thereof and f_q, g_q polynomials in $\overline{k}[x_0, x_1, \ldots, x_n]$ so that $\zeta|_{U_q} = f_q/g_q$. Hence $f_p/g_p = f_q/g_q \iff f_p g_q = f_q g_p$ on the open set $U_p \cap U_q$ (non-empty because V is irreducible). However $f_p g_q = f_q g_p$ holds exactly on $Z(f_p g_q - f_q g_p) \subseteq \mathbb{P}^n(\overline{k})$ a closed set containing $U_p \cap U_q$. Note that $f_p g_q - f_q g_p$ is a homogeneous polynomial, because each of f_p, g_p, f_q, g_q are homogeneous and each pair f_p, g_p and f_q, g_q are each of the same degree. However, since V is irreducible, the intersection

of any two non-empty open sets is non-empty dense. Hence $Z(f_p g_q - f_q g_p)$ contains the closure of $U_p \cap U_q$, which is the whole of V. Consequently $\zeta = f_p/g_p = f_q/g_q$ on the entirety of V. This quotient is not just a local representation.

So we know that $\eta = (\eta_1 : \ldots : \eta_n)$ has components each given by some $\eta_i = f_i/g_i$ with f_i, g_i homogeneous polynomials in $\overline{k}[x_0, x_1, \ldots, x_n]$. Moreover, since η is well-defined, the g_i are non-zero on V and so we can multiply every component η_i by $g_1 \cdots g_n$ to obtain a well defined map $\eta' = ((g_1 \cdots g_n)\eta_1 : \ldots : (g_1 \cdots g_n)\eta_n)$. However, since η is a map into projective space, $\eta = \eta'$. Hence $\eta = \varphi = (\varphi_1 : \ldots : \varphi_n)$ with $\varphi_i = g_1 \cdots g_n \eta_1$ homogeneous polynomials in $\overline{k}[x_0, x_1, \ldots, x_n]$ as claimed.

Finally, we note that it is customary to denote x_0, x_1, x_2 by X, Y, Z when discussing $\mathbb{P}^2(\overline{k})$. This is a little unfortunate, for it mixes the use of Z as an indeterminate variable, and $Z(\cdot)$ as the *zero-locus*. From context, however, there should not be cause for confusion.

We now recall the following definition of an elliptic curve (there are many equivalent ones [11, pg. 1]). An *elliptic curve* E over the base field k (characteristic $\neq 2, 3$) in *short Weierstrass form* is a projective variety cut out from $\mathbb{P}^2(\overline{k})$ by an equation $Y^2Z = X^3 + aXZ + bZ^3$ where a, b in k satisfy $4a^3 + 27b^2 \neq 0$. The coefficients a, b are called the *Weierstrass invariants* of E. A general *elliptic curve* E over k (characteristic $\neq 2, 3$) is then a projective variety that is isomorphic to an elliptic curve in short Weierstrass form.

We stress that the curve is defined using data in k but described with points in the algebraic closure $\overline{k}$. This difference is plainly clear when $k = F_q$ is finite (our case of interest) because then the set of solutions $Y^2Z = X^3 + aXZ + bZ^3$ would be finite, and we would have no "geometry" to speak of. That $4a^3 + 27b^2 \neq 0$ ensures *non-singularity* of E. This can be understood as a geometric statement: it essentially means that the curve has no "kinks".

Choosing $0_E = (0 : 1 : 0)$ to be the neutral element of E, we obtain a unique commutative group law on E [11, Rem. 1.3] given by a regular map $+ \colon E \times E \to E$ defined over k [28, III.2 Alg. 2.3]. That is to say, that the components of $+$ have coefficients in k. Here it was important that the a, b defining E lie in k. In fact, any choice of neutral element $0_E = (X_0 : Y_0 : Z_0)$ with X_0, Y_0, Z_0 in k results in a group law given by a regular map defined over k by [11, Cor 1.2]. An important consequence of this, is that the set of k-rational points $E(k) \subseteq E$, namely the points $(X : Y : Z)$ in E with X, Y, Z in k, forms a subgroup of E.

Further, we note that $0_E = (0 : 1 : 0)$ is the only point on E for which $Z = 0$. So, for all other points $P = (X : Y : Z) \neq 0_E$ we have $Z \neq 0$ and we can normalise $P = (X/Z : Y/Z : 1)$. This gives rise to the *affine coordinates* of E and we may write $E = Z(y^2 - (x^3 + ax + b)) \cup \{0_E\}$ where 0_E is an abstract point called *the point at infinity*. It is common to write projective coordinates in upper case X, Y, Z and affine coordinates in lower case x, y.

Since elliptic curves are groups and projective varieties, it is natural to define a *morphism* of elliptic curves $\varphi \colon E \to E'$ to be a morphism of both projective varieties

(regular map) and a morphism of groups. In the literature (especially that of Abelian varieties) these morphisms are often called *homomorphisms*.

For elliptic curves there exists a more explicit *standard form* [7, Lecture 4] for any isogeny $\varphi\colon E \to E'$, written in affine form as

$$\varphi(x:y:1) = \left(\frac{u(x)}{v(x)} : \frac{s(x)}{t(x)}y : 1\right) = (u(x)t(x) : v(x)s(x)y : t(x)v(x))$$
$$\text{and} \qquad \varphi(0:1:0) = (0:1:0)$$

with u, v and s, t each coprime. We note that, since φ is a map into E' cut out by $ZY^2 = X^3 + a'XZ^2 + b'Z^3$, that

$$y^2 s(x)^2 v(x)^3 = t(x)^2\left(u(x)^3 + a'u(x)^2 + b'\right)$$

and so $s(x) = 0$ if and only if $t(x) = 0$. Moreover, the kernel of φ is given by the zeros of $s(x)$. Hence, if φ is separable, its degree is exactly the degree of $s(x)$.

Together with the results from our informal introduction to isogenies between elliptic curves in Subsection 1.3, this lays the technical foundations for generalising to higher dimensions.

2.2 Isogenies between principally polarised Abelian varieties

In this subsection we introduce principally polarised Abelian varieties, isogenies between them and discuss how they generalise elliptic curves to higher dimensions.

A short remark on the literature and presentation. At the time of writing, there appear to be three standard references for (not necessarily complex) Abelian varieties, Milne [11], Edixhoven-Moonen-Gerard [31] and Mumford [27]. Unlike the others, Milne takes an approach amenable to working with varieties in the sense of algebraic spaces as opposed to full generality of schemes; this is advantageous because it generalises the usual approach to elliptic curves more directly. Conversely, Milne's book is not as comprehensive as the other references. We try to stay with Milne's presentation, and only cite the others when necessary.

Consider the product $E \times E'$ of two elliptic curves E, E'. As the product of groups, it is a group and as the product of irreducible varieties it is an irreducible variety ([12, Prop. 5.20]); but it is not an elliptic curve for it somehow has "dimension" 2. In general, a variety together with a group law is known as an *Abelian variety*, but it turns out, that Abelian varieties alone are too general and we require them to be *principally polarised*. This involves the notion of the dual (Abelian) variety. To successfully convince ourselves that these new objects truly are a generalisation of elliptic curves, and to understand what features they have, we give a brief introduction in the following.

With our definition of an elliptic curve from the previous subsection in mind, we now give what we intuitively would like the definition of an Abelian variety to be. An *Abelian variety* A over the field k is a non-singular irreducible projective variety A over $\overline{k}$ defined using data in k, equipped with a commutative group law $A \times A \to A$ and an inverse $A \to A$ whose formulae are given by regular maps defined over k. We denote the

law of A by $+$, the inverse by $-$ and the (k-rational) neutral element by 0_A or just 0. We denote such an Abelian variety by A/k like we do for elliptic curves.

As with any algebro-geometric object, the language of schemes can be used for rigorous formalisation and analysis. Since Abelian varieties are still varieties, they do not require the full generality of schemes and an intermediate formalisation, namely that of *algebraic spaces* can be used [10, Ch. 11].

We recall that the study of "classical"[6] varieties lives in a world in which all fields k are algebraically closed, and attaches to every finitely generated k-algebra a geometric object which we call an *affine variety*[7]; the study of schemes generalises this by assigning every arbitrary ring a geometric object which we call an *affine scheme*; the study of algebraic spaces apply the construction of affine schemes not to any arbitrary ring, but again to finitely generated k-algebras where k is *not* necessarily algebraically closed. Here, the notion of defining a variety with data over k, but viewing it in $\mathbb{P}^n(\overline{k})$ (or $\mathbb{A}^n(\overline{k})$, the non-projective *affine* space) is formalised by *extending scalars*. Milne notes that extending scalars generally behaves well when the base field k is perfect (e.g. $k = F_q$ finite) [11, pg. 168].

Nevertheless, we will forgo all of these formalisms, and content ourselves with the following definition. We recall that a projective variety over $\overline{k}$ is a (locally-ringed space that is isomorphic to some) closed subset of $\mathbb{P}^n(\overline{k})$, which in turn is cut out by some homogeneous polynomials f_i in $\overline{k}[X_1, \dots, X_n]$. To capture the notion, that our Abelian variety A is defined using data in k, we require A to be a projective variety in which each of the f_i have a scalar λ_i so that $\lambda_i f_i$ lies in $k[X_1, \dots, X_n] \subseteq \overline{k}[X_1, \dots, X_n]$. Since scaling the polynomials does not change the zero-locus, this effectively equivalent to asking that the f_i lie in $k[x_1, \dots, x_n]$ directly. We recall that a regular map $V \subseteq \mathbb{P}^n(\overline{k}) \to W \subseteq \mathbb{P}^m(\overline{k})$ of (irreducible projective) varieties over $\overline{k}$ is a map $\varphi = (\varphi_1, \dots, \varphi_m)$, whose components φ_i are be (homogeneous) polynomials in $\overline{k}[X_1, \dots, X_n]$. We say that such a regular map is *defined over* an extension k' of k, if there exists a constant λ in $\overline{k}$ (universal for all $i = 1, \dots, n$) so that the $\lambda\varphi_i$ lie in $k'[X_1, \dots, X_n] \subseteq \overline{k}[X_1, \dots, X_n]$. When $k' = k$, we simply say φ is defined over k. This tells us what it means for the group law and inverse to be regular *over* k as we stipulated in our informal definition. In particular, the set of k-rational points $A(k)$ on A then form a subgroup.

We note that under this definition, elliptic curves are Abelian varieties. Indeed, an elliptic curve E/k is a projective variety furnished with a group law and inverse defined by regular maps over k. We also note that under this definition, the product of (finitely many) Abelian varieties over k is an Abelian variety over k. Indeed, the product of projective varieties is again a projective variety, a fact shown by the *Segre embedding* $\mathbb{P}^n \times \mathbb{P}^m \hookrightarrow \mathbb{P}^{(m+1)(n+1)-1}$ taking $((x_0 : \cdots : x_n), (y_0 : \cdots : y_m)) \mapsto (x_i y_j)_{i,j}$. This map has integer coefficients, so the defining equations and the group laws (defined over k) transferred under the embedding are still defined over k.

[6] I quote "classical", since Milne overloads the term "algebraic variety" when he introduces algebraic spaces to now allow non-algebraically closed fields.

[7] When introducing varieties, we did this in reverse. We began with a geometric object (points in some space satisfying some polynomials) and defined on this a sheaf of rings, which was a $\overline{k}$-algebra.

We note that Abelian varieties are automatically non-singular. We recall that a point P on an affine variety V is *non-singular* or *regular* if the stalk at P is regular, or equivalently if P lies in exactly one irreducible component W, and if the dimension of the Zariski tangent space $T_p(V)$ at P has dimension (as a vector space) equal to the dimension of W (as a topological space) [12, Cor. 4.45]. This is analogous to being non-singular on a manifold. Importantly, being (non-)singular is a local condition. We now note that the non-singular locus N of A (or indeed any variety) is open and dense [12, Th. 4.37]. However, the translation isomorphisms $t_a \colon A \to A; x \mapsto x + a$ cover the entirety of A with the non-singular patch N, and so the whole of A is non-singular. In more detail, let us consider the point P on A outside of N, which is *a priori* singular and pick a non-singular point Q in N. The translation map t_{Q-P} maps P to Q, and in particular establishes an isomorphism between any neighbourhood of P to a neighbourhood of Q. Since being non-singular is a local condition, this suffices to conclude that P is also non-singular. In essence, it is the existence of the regular group law that enforces non-singularity.

The *dimension* of an Abelian variety A (or any variety) is its dimension as a topological space. That is, the supremum over n for which there exists a strictly increasing chain $Z_0 \subsetneq \cdots \subsetneq Z_n$ of (non-empty) irreducible closed subsets of A. For example, k viewed as an affine variety over k (i.e. endowed with the Zariski topology) has dimension 1 because the Zariski topology on k is the cofinite topology: the only closed sets are finite sets (and the whole of k). Note that k is not an Abelian variety for it is not projective. We can equivalently define the dimension of an Abelian variety A (or any irreducible variety) to be the transcendence degree $\mathrm{trdeg}_k(k(A))$ over k of its field of functions $k(A)$ [12, Sec. 5.j]. In fact, since $k(U) \cong k(A)$ for any open dense subset U of A, we see that Elliptic curves have dimension 1: the function field $k(U) = k(x,y)/(x^3 + ax + b - y^2)k(x,y)$ of the affine patch $U = E - Z(Z)$[8] on E has transcendence degree 1 over k. Consequently elliptic curves over k are Abelian varieties over k of dimension 1. A hard(er) fact to prove, is that in fact elliptic curves are *all* of the Abelian varieties of dimension 1. Finally, given an Abelian variety, its n-fold product A^n has dimension $n \dim(A)$.

A regular map $A \to B$ between two Abelian varieties that preserves the group structures is called a *homomorphism*. A surjective homomorphism between two Abelian varieties of the same dimension is called an *isogeny*. In analogy to results on elliptic curves, we note that a regular map must only assign $0_A \mapsto 0_B$ for it to preserve the whole group structure. This is a consequence of the important *Rigidity theorem* [11, Th. 1.1]. In particular, isogenies between elliptic curves are isogenies between Abelian varieties. We say an isogeny between r-dimensional Abelian varieties is an *r-dimensional* isogeny. This gives meaning to the statement embedding an isogeny into "higher dimension". Finally we define the *degree* $\deg(f)$ of an isogeny $f \colon A \twoheadrightarrow B$ to be the degree of the induced field extension $f^* \colon k(B) \hookrightarrow k(A)$ of the function fields $k(A), k(B)$ of A, B (i.e. the degree of f as a regular map); and call f *separable* if the extension is separable. This

[8] That is, the zero-locus of the variable Z in $k[X,Y,Z]$.

is compatible with our definition of the degree of an isogeny between elliptic curves.

We define the kernel $\ker(f)$ of an isogeny $f\colon A \to B$ to be its kernel as a group morphism, that is, the set of points in A sent to 0_B by f. The kernel of any isogeny is finite (in fact, this is equivalent to an isogeny being surjective [11, Prop. 7.1]), and forms a subgroup of A (it is the kernel of a group morphism). Conversely, for every finite subgroup K of A there exists an isogeny $A \to B$, unique up to compositions with isomorphisms, with kernel K [11, Th. 8.10, Rem. 8.12]. Therefore, we call B the *quotient* of A by K and write $B = A/K$. It is defined uniquely up to isomorphism.

As with elliptic curves, the prototypical example of an isogeny of Abelian varieties is the endomorphism of *scalar multiplication* or *multiplication-by-N*, defined as mapping $a \mapsto Na$ for some integer N. These maps are surely regular because the group laws are regular, and send 0 to 0. Proving that they are surjective (except when $N = 0$) is a more involved [11, Th. 7.2]. Nevertheless, they are isogenies. We usually denote multiplication-by-N on A by $[N]_A$ or just $[N]$ since the multiplication-by-N map is well-defined on every Abelian variety and acts in the same way. We write $A[N] = \ker([N]_A)$ and call this the *N-torsion* on A. Moreover, since the group laws are defined over k, $[N]_A$ is a well-defined group morphism on $A(k)$.

An isogeny between Abelian varieties defined over k is separable when its degree is coprime to $\operatorname{char} k$ [14, Prop. 2.20]; and, as with elliptic curves, $\deg(f) = |\ker(f)|$ for separable isogenies. As such, multiplication-by-N for N coprime to the characteristic of k satisfies $|\ker([N]_A)| = \deg([N]_A) = N^{2\dim(A)}$ with the last equality due to [11, Th. 7.2].

For every degree-d isogeny $f\colon A \to B$ there exists another isogeny $g\colon B \to A$ so that $gf = [d]_A$ and $fg = [d]_B$ [31, Prop 5.12]. This gives us the *isogenous* equivalence relation as we had it for elliptic curves. In fact, the theorem of Tate [59, Theorem 1.(c)] that we cited to tell us whether elliptic curves are k-isogenous, also applies for Abelian varieties. That is, two Abelian varieties A, B over k are isogenous via an isogeny defined over k if and only if $|A(k)| = |B(k)|$.

However, if $r = \dim(A) = \dim(B) \geq 2$ and d prime to $\operatorname{char} k$ then $\deg(g)\deg(f) = d^{2r}$ and consequently $\deg(g) = d^{2r-1}$. We do not see this ugly imbalance $\deg(f) \neq \deg(g)$ for elliptic curves because $d^{2r-1} = d$ when the dimension $r = 1$. Fortunately, there is a remedy for this, giving us the *reversed* map $\widetilde{f}\colon B \to A$ with $\deg(f) = \deg(\widetilde{f})$ and $\widetilde{f}f = [d']_A$. Unfortunately, d' is no longer the degree of f, and perhaps even more unfortunately, it requires more sophisticated tools from our algebro-geometric toolbox, namely the *dual Abelian variety*.

For every Abelian variety A, there exists another Abelian variety $A^\vee$ of dimension $\dim(A)$ which we call the *dual* of A; and for every homomorphism $f\colon A \to B$, there exists a homomorphism $f^\vee\colon B^\vee \to A^\vee$ which we call the *dual* of f. If f is an isogeny, then $f^\vee$ is also an isogeny of degree $\deg(f)$ which we call the *dual isogeny* [11, Th. 9.1; 31, Th. 6.18, Def. 6.19]. The dual variety (and so dual isogeny) is uniquely defined up to unique isomorphism by a universal property (which we do not state) [11, Rem. 8.7]. The dual is reflexive, that is $(f^\vee)^\vee = f$ and $(A^\vee)^\vee = A$. In other words, the map $A \to A^\vee$ is a *dualising* contravariant functor in the category of Abelian varieties. Finally,

$(A \times A)^\vee \cong A^\vee \times A^\vee$. That is, the dual commutes with the product.

Every divisor D (an irreducible subvariety of codimension 1) of an Abelian variety A induces an isogeny $\lambda_D \colon A \to A^\vee$. We call such isogenies *polarisations* of A. We call the pair (A, λ_D) a polarised Abelian variety (PAV). When λ_D is an isomorphism, we call it a *principal* polarisation and call the pair (A, λ_D) a *principally* polarised Abelian variety (PPAV). Elliptic curves are canonically principally polarised. This (correctly) suggests that PPAVs are the "true" generalisation of elliptic curves to higher dimensions.

Milne notes that A and its dual $A^\vee$ are isogenous, but not usually isomorphic [11, pg. 3, Abelian varieties as generalizations of elliptic curves]. Conversely, A is always isogenous to a PPAV, which may not be the dual [27, Cor. 1, pg. 234; 31, Cor. 11.26]. See [30] for an explicit example of an Abelian variety that is not isomorphic to its dual. Moreover, even if an Abelian variety *is* isomorphic to its dual, this isomorphism might not be canonical, and so we must choose one explicitly.

Once a polarisation $\lambda_D \colon A \to A^\vee$ is chosen, we often denote it by λ_A. There is no notational ambiguity because A is not a divisor of A (it has codimension 0). Hence it makes sense to speak of (A, λ_A) being a (P)PAV.

When (A, λ_A) is a (P)PAV, $A \times A$ is a (P)PAV endowed with the (principal) *product polarisation* $\lambda_A \times \lambda_A \colon A \times A \to (A \times A)^\vee \cong A^\vee \times A^\vee$. We will denote the product polarisation of the n-fold product of (A, λ_A) by $\lambda_A^{(n)}$.

A homomorphism $f \colon (A, \lambda_A) \to (B, \lambda_B)$ between PPAVs can be *reversed* by defining $\widetilde{f} = \lambda_A^{-1} f^\vee \lambda_B \colon B \to A$. If f is an isogeny, then so is $\widetilde{f}$. Since forming the dual does not change the degree, $(\deg(f^\vee) = \deg(f))$, and the polarisations are principal (i.e. isomorphisms), we see that $\deg(f) = \deg(\widetilde{f})$ as promised. We also note that the reverse is also reflexive, that is the reverse of $\widetilde{f}$ is again f. Moreover, the reverse of a composition is the reversed composition of the reversals, that is, $\widetilde{fg} = \widetilde{g}\widetilde{f}$. These features are important, and illustrate why we require the target (B, λ_B) to be *principally* polarised. A priori, for $\widetilde{f} = \lambda_A^{-1} f^\vee \lambda_B$ to be well-defined, we only need λ_A to be invertible. We also require B to be a PPAV to form the reverse of the reverse.

Since elliptic curves are canonically principally polarised, we get the reversal of isogenies between elliptic curves for free. Somewhat regrettably, the reversal of an isogeny between elliptic curves is called the *dual* in most literature, most notably in [28]; of course it is only *somewhat* regrettable because, after all, elliptic curves are canonically isomorphic to their dual. That is to say, the distinction between the reversal and dual is less important for isogenies between elliptic curves, since we almost always consider elliptic curves as PPAVs together with their canonical principal polarisation.

If $\widetilde{f}f = [d_1]_A$ and $f\widetilde{f} = [d_2]_B$ are scalar multiplications on A and B respectively, the scalars d_1, d_2 must be the same $N = d_1 = d_2$ and we call f an *N-isogeny*. Indeed, $[d_2]_B f = f\widetilde{f}f = f[d_1]_A = [d_1]_B f$ so $d_1 = d_2$ because f is surjective. By symmetry, $\widetilde{f}$ is an N-isogeny if and only f is. Analogously, we call an endomorphism that is an N-isogeny, an N-endomorphism. As with isogenies of elliptic curves, the composition of an N-isogeny and an M-isogeny is a NM-isogeny.

We note that isogenies between elliptic curves are all automatically N-isogenies,

where N coincides with the degree. This is not true for higher dimensional isogenies: $f\widetilde{f}$ is not scalar multiplication for every (higher-dimensional) isogeny $f\colon A \to B$ (see example shortly after Lemma 2.3), let alone scalar multiplication by its degree $\deg(f)$. What is true, however, is that the degree of an N-isogeny $A \to B$ of r-dimensional PPAVs is N^r. This demonstrates that the new nomenclature for isogenies between (potentially higher dimensional) Abelian varieties is consistent with that of isogenies between elliptic curves where we called degree-d isogenies *d-isogenies*.

We also make two easy observations about the kernel of an N-isogeny $f\colon A \to B$. First, that $\ker(f) = \widetilde{f}(B[N])$. Indeed, $f(\widetilde{f}(B[N])) = NB[N] = \{0\}$ so $\widetilde{f}(B[N]) \subseteq \ker(f)$. Conversely, for every x in $\ker(f)$ there exists some y in B so that $\widetilde{f}(y) = x$ because isogenies are surjective. Then $0 = f(x) = f\widetilde{f}(y) = Ny$ and so y lies in $B[N]$ and $x = \widetilde{f}(y)$ in $\widetilde{f}(B[N])$. Hence $\ker(f) \subseteq \widetilde{f}(B[N])$ and $\ker(f) = \widetilde{f}(B[N])$. The second observation follows immediately from the first, and is simply that $\ker(f) \subseteq A[N]$.

2.3 Embedding isogenies: Kani's Lemma and Zarhin's Trick

In what follows, we describe how isogenies between PPAVs can be embedded into an isogenies of higher dimensions. The prototypical example (and the case we are most interested in) is to embed isogenies of elliptic curves (1-dimensional) into higher dimensions (more precisely: 1, 2, or 4-dimensions).

The key property of such an embedding, is that a (lower-dimensional) isogeny of *arbitrary* degree can be embedded into a higher-dimensional isogeny of *smooth* degree. Using methods to compute higher dimensional isogenies, we hope to efficiently compute the smooth-degree higher-dimensional isogeny and efficiently recover the values of the lower-dimensional isogeny

We first begin with an elementary construction. Let $f\colon (A, \lambda_A) \to (B, \lambda_B)$ be an isogeny between PPAVs. Then $f^{(n)} = (f, \dots, f)\colon A^n \to B^n$ is an isogeny between $n\dim(A) = n\dim(B)$ dimensional PPAVs. Notably, its reversed map $\widetilde{f^{(n)}}\colon B^n \to A^n$ is given by $(\widetilde{f}, \dots, \widetilde{f})$ because the product polarisations on A^n, B^n are just the n-fold products $\lambda_A^{(n)}, \lambda_B^{(n)}$ of the polarisations λ_A, λ_B on A, B respectively. In other words, the reversal of n-copies are n-copies of the reversal $\widetilde{f^{(n)}} = (\widetilde{f})^{(n)}$. Hence the resulting isogeny $f^{(n)}$ is also an N-isogeny, but now has degree $\deg(f)^n$.

For isogenies between products of elliptic curves we have a stronger converse in the following

Lemma 2.1. *An isogeny $f\colon E_1 \times \cdots \times E_n \to E_1' \times \cdots \times E_n'$ is given by a matrix $(f_{j,i})_{i,j}$ of isogenies (or zero-maps) $f_{i,j}\colon E_i \to E_j'$ acting by matrix-multiplication.*

Proof. To prove this, we simply define $f_{i,j} = \pi_j f \zeta_i$ where ζ_i is the inclusion $E_i \hookrightarrow E_1 \times \cdots \times E_n; x_i \mapsto (0_{E_1}, \dots, 0_{E_{i-1}}, x_i, 0_{E_{i+1}}, \dots, 0_{E_n})$ and π_j is the projection $E_1' \times \cdots \times E_n' \twoheadrightarrow E_j'$. The maps $f_{i,j}$ are surely regular and send 0_{E_i} to $0_{E_j'}$. So, by [8, Prop. II.6.8], $f_{i,j}$ are either constantly zero, or surjective. Hence $f_{i,j}$ is either a (non-zero) isogeny, or the zero-map.

We must now only verify that $f = (f_{j,i})_{i,j}$. Since f is a group morphism

$$f(x) = f(x_1, \dots, x_n) = \sum_{j=1}^{n} f(0_{E_1}, \dots, 0_{E_{j-1}}, x_j, 0_{E_{j+1}}, \dots, 0_{E_n}) = \sum_{j=1}^{n} f\zeta_j(x_j),$$

so $\pi_i f(x) = \pi_i f\zeta_1(x_1) + \cdots \pi_i f\zeta_n(x_n) = ((f_{j,i})_{i,j}(x))_i = \pi_i((f_{j,i})_{i,j}(x))$ and we conclude that $f = (f_{j,i})_{i,j}$. □

This inspires us to look at matrices of isogenies. We return to the more general setting of Abelian varieties. If $A_1, \dots, A_n, B_1, \dots, B_n$ are r-dimensional PPAVs and $f_{i,j}\colon A_i \to B_j$ isogenies, then the matrix $F = (f_{j,i})_{i,j}$ is a homomorphism $A_1 \times \cdots \times A_n \to B_1 \times \cdots \times B_n$ when acting my matrix multiplication. We verify that this is well-defined, in particular that we must "transpose" the maps: the i-th element of the image vector lies in B_i and is given by $b_i = F_{i,1}(a_1) + \cdots + F_{i,n}(a_n) = f_{1,i}(a_1) + \cdots + f_{1,i}(a_n)$, which is indeed a sum of elements in B_i.

Even though such a matrix is a homomorphism between equal dimensional Abelian varieties, it must not be surjective, and so, is not necessarily an isogeny. An obvious example is $(1, 1; 1, 1)\colon A \times A \to A \times A$. A partial result is found in [31] where it is proven that a non-singular integer matrix is an isogeny (Prop. 11.28; note: to speak of the degree of a homomorphism implies that it is surjective, and so an isogeny). This demonstrates that [8, Prop. II.6.8], an important ingredient to the proof of Lemma 2.1, does not generalise to higher dimensions. Nevertheless, we have defined the duals and reversals of a homomorphism and so can try to determine when F is an N-isogeny. All that is required to be shown is that $\widetilde{F}F = [N]$ and $F\widetilde{F} = [N]$, because then F (and $\widetilde{F}$) is surjective, and so isogenies. For this, we need the following

Lemma 2.2. *The dual of the homomorphism $F = (f_{j,i})_{i,j}$ is given by $F^\vee = (f_{i,j}^\vee)_{i,j}$.*

That is, the dual of a matrix of homomorphisms is the transpose of the entry-wise duals. This can be proven using the usual strategy for objects defined with a universal property: pick the object which intuitively should be correct, show that it satisfies the universal property and conclude that the intuition was correct. In our case $F^\vee$ is a map $B_1^\vee \times \cdots \times B_n^\vee \to A_1^\vee \times \cdots \times A_n^\vee$, and since F comes from a matrix, we expect its dual $F^\vee$ to also come from a matrix. Hence $(F^\vee)_{i,j}$ is a map $B_j^\vee \to A_i^\vee$ for which only the candidate $f_{i,j}^\vee$ remains. We stress that the above is only a proof *idea*, because we did not state the universal property of the dual, so we could not have verified it. The result is correct nevertheless. It is proven for integer matrices [31, 11.28], and the reversal version (Cor. 2.3 below) is proven for 2×2 matrices of isogenies in [20, pg. 100] (the result from which Kani's Lemma is taken).

Since the product polarisations act componentwise, $\lambda_A^{(n)}, \lambda_B^{(n)}$ act like scalar matrices and $\widetilde{F} = (\lambda_A^{(n)})^{-1} F^\vee \lambda_B^{(n)} = (\lambda_A^{-1})^{(n)} F^\vee \lambda_B^{(n)}$. From this, we conclude the

Corollary 2.3. *The reversal of the homomorphism $F = (f_{j,i})_{i,j}$ is given by $\widetilde{F} = (\widetilde{f_{i,j}})_{i,j}$.*

So like the dual, the reversal of a matrix of isogenies is the transpose of the entry-wise reversals.

This immediately tells us that even when a matrix of isogenies is an isogeny, it is likely not an N-isogeny. That is, $F\widetilde{F}$ is likely not of the form $[N] = N\mathrm{id}_{n\times n}$. We can easily construct such an example: let $F = (1, 2; 3, 4)$ (this is an isogeny because it is a non-singular matrix), then $F\widetilde{F} = FF^t = (5, 11; 11, 25)$ is not a scalar multiplication.

We will find criteria to ensure a 2×2 matrix F is in fact an N-isogeny. To that end, let $f_{i,j}\colon A_i \to B_j$ be $d_{i,j}$-isogenies between the PPAVs A_1, A_2, B_1, B_2 and set $F = (f_{j,i})_{i,j}\colon A_1 \times A_2 \to B_1 \times B_2$.

We compute $\widetilde{F}F\colon A_1 \times A_2 \to A_1 \times A_2$ to obtain

$$\begin{aligned}\begin{pmatrix}\widetilde{f_{1,1}} & \widetilde{f_{1,2}}\\ \widetilde{f_{2,1}} & \widetilde{f_{2,2}}\end{pmatrix}\begin{pmatrix}f_{1,1} & f_{2,1}\\ f_{1,2} & f_{2,2}\end{pmatrix} &= \begin{pmatrix}\widetilde{f_{1,1}}f_{1,1} + \widetilde{f_{1,2}}f_{1,2} & \widetilde{f_{1,1}}f_{2,1} + \widetilde{f_{1,2}}f_{2,2}\\ \widetilde{f_{2,1}}f_{1,1} + \widetilde{f_{2,2}}f_{1,2} & \widetilde{f_{2,1}}f_{2,1} + \widetilde{f_{2,2}}f_{2,2}\end{pmatrix}\\ &= \begin{pmatrix}[d_{1,1}+d_{1,2}]_{A_1} & \widetilde{f_{1,1}}f_{2,1} + \widetilde{f_{1,2}}f_{2,2}\\ \widetilde{f_{2,1}}f_{1,1} + \widetilde{f_{2,2}}f_{1,2} & [d_{2,1}+d_{2,2}]_{A_2}\end{pmatrix}.\end{aligned}$$

This is a scalar multiplication if and only if

$$\text{(i) } d_{1,1} + d_{1,2} = d_{2,1} + d_{2,2}\ ,\quad \text{(ii) } \widetilde{f_{1,1}}f_{2,1} + \widetilde{f_{1,2}}f_{2,2} = 0 \text{ and } \text{(iii) } \widetilde{f_{2,1}}f_{1,1} + \widetilde{f_{2,2}}f_{1,2} = 0.$$

Point (i) delivers important insight. If F is an N-isogeny, then $N = d_{1,1} + d_{2,1}$. This illustrates behaviour that was previously impossible: obtaining from some isogenies a new isogeny whose degree is a *sum* of degrees, and not the *product* of degrees as one obtains through composition.

We also note that these points together imply that $d_{1,1} = d_{2,2}$ and $d_{1,2} = d_{2,1}$. Indeed, (ii) and (iii) both imply $d_{1,1}d_{2,1} = d_{1,2}d_{2,2}$, so solving (i) for $d_{2,1}$ and substituting this in we obtain $0 = d_{1,1}d_{2,1} - d_{1,2}d_{2,2} = d_{1,1}(d_{1,1}+d_{1,2}-d_{2,2}) - d_{1,2}d_{2,2} = (d_{1,1}+d_{1,2})(d_{1,1}-d_{2,2})$; for which only $d_{1,1} = d_{2,2}$ is possible because degrees are positive.

Points (ii) and (iii) are equivalent, since one equation is the reversal of the other. Writing down the commutative diagram corresponding to point (iii) we obtain

$$\begin{array}{ccc} A_1 & \xrightarrow{f_{1,1}} & B_1 \\ {\scriptstyle -f_{1,2}}\downarrow & & \downarrow{\scriptstyle \widetilde{f_{2,1}}} \\ B_2 & \xrightarrow[\widetilde{f_{2,2}}]{} & A_2 \end{array} \tag{2.1}$$

From this, we conclude that a 2×2 matrix $F = (f_{j,i})_{i,j}$ of $d_{i,j}$-isogenies $f_{i,j}\colon A_i \to B_j$ is a N-isogeny $A_1 \times A_2 \to B_1 \times B_2$ if and only if the isogenies $f_{i,j}$ fit into the commutative diagram (2.1) with $d_{1,1} = d_{2,2}$ and $d_{1,2} = d_{2,1}$. In this case $N = d_{1,1} + d_{2,1} = d_{2,1} + d_{2,2}$.

This is usually stated in reverse (starting with a square) and is known as Kani's

Lemma 2.4 (Kani, [17, Lem 2.1]). *Let*

$$\begin{array}{ccc} A & \xrightarrow{g_1} & C \\ {\scriptstyle h_1}\downarrow & & \downarrow{\scriptstyle g_2} \\ C' & \xrightarrow[h_2]{} & B \end{array}$$

be a (not necessarily commuting) diagram of isogenies between principally polarised

Abelian varieties such that g_1, h_2 are d_1-isogenies and h_1, g_2 are d_2-isogenies. In particular, the degrees of parallel isogenies in the diagram are equal.

The resulting Kani map

$$K = \begin{pmatrix} g_1 & \widetilde{g_2} \\ -h_1 & \widetilde{h_2} \end{pmatrix} : A \times B \to C \times C'$$

is a $d_1 + d_2$ isogeny if and only if the square commutes.

In such cases we call the square a Kani square.

At least within the isogeny-based cryptography community, this Lemma is attributed to Kani, for it is an observation taken from the proof of Kani's *Reducibility Theorem* [20, Th. 2.6, pg. 100].

Since isogenies are uniquely defined by their kernels (up to isomorphism) and our computational tools will use the kernel to compute isogenies, we would like to compute the kernel of K as a function of the g_i, h_j. Since we are interested in separable isogenies, we assume that $d = d_1 + d_2$ is prime to the characteristic of k over which the PPAVs A, C, C', B are defined.

We give the description of the kernel in an extended version of Kani's

Lemma 2.5. *Let*

$$\begin{array}{ccc} A & \xrightarrow{g_1} & C \\ {\scriptstyle h_1}\downarrow & & \downarrow{\scriptstyle g_2} \\ C' & \xrightarrow[h_2]{} & B \end{array}$$

be a Kani square and $K = (g_1, \widetilde{g_2}; -h_1, \widetilde{h_2}) \colon A \times B \to C \times C'$ the corresponding Kani map. Then

$$\ker(K) = \left\{ \left(\widetilde{g_1}(v) - \widetilde{h_1}(w), g_2(v) + h_2(w) \right) \mid v \in C[d], w \in C'[d] \right\}.$$

If the kernels of $\widetilde{g_1}$, g_2 intersect trivially we say the square is orthogonal *and*

$$\ker(K) = \{ (\widetilde{g_1}(v), g_2(v)) \mid v \in C[d] \}.$$

Finally, if d_1 is prime to d_2 we say the square is minimal *and*

$$\ker(K) = \{ (d_1 x, g_2 g_1(x)) \mid x \in A[d] \}.$$

Proof. To prove the lemma, we recall that $\ker(f) = \widetilde{f}(B[N])$ for any N-isogeny $f \colon A \to B$ and compute

$$\begin{aligned} \ker(K) &= \widetilde{F}(C[d] \times C'[d]) \\ &= \begin{pmatrix} \widetilde{g_1} & -\widetilde{h_1} \\ g_2 & h_2 \end{pmatrix} (C[d] \times C'[d]) \\ &= \left\{ \left(\widetilde{g_1}(v) - \widetilde{h_1}(w), \widetilde{g_2}(v) + \widetilde{h_2}(w) \right) \mid v \in B_1[d], w \in B_2[d] \right\}. \end{aligned}$$

Giving the first kernel description of the lemma. This is not particularly enlightening. However, when the kernels of $\widetilde{g_1}, g_2$ intersect trivially, $\widetilde{K}$ is injective on $C \times \{0\}$ and

we obtain the second, simpler, description of the kernel

$$\ker(K) = \widetilde{K}(C[d] \times C'[d]) \overset{(*)}{=} \widetilde{K}(C[d] \times \{0\}) = \{(\widetilde{g_1}(v), g_2(v)) \mid v \in C'[d]\}.$$

When the kernels $\widetilde{g_1}, g_2$ intersect trivially, $\widetilde{K}$ is injective because $\widetilde{K}(v,0) = (\widetilde{g_1}(v), g_2(v)) = (0,0)$ if and only if v lies in $\ker(\widetilde{g_1}) \cap \ker(g_2)$; and $(*)$ holds because $\widetilde{K}(C[d] \times \{0\})$ is contained in $\widetilde{K}(C[d] \times C'[d])$, but

$$\begin{aligned} d^{2r} = |C[d] \times \{0\}| &\overset{\widetilde{K}\text{ inj.}}{=} \left|\widetilde{K}(C[d] \times \{0\})\right| \\ &\leq \left|\widetilde{K}(C[d] \times C'[d])\right| = |\ker(K)| \overset{K\text{ sep.}}{=} \deg(K) = d^{2r} \end{aligned}$$

shows that these nested sets have the same (finite) cardinality. Hence they must be equal.

Finally, we note that when d_1 is prime to d_2 (in particular then, the kernels of $\widetilde{g_1}, g_2$ intersect trivially), the map $g_1 \colon A[d] \to C[d]$ is a bijection and we yield the final, even simpler, description of the kernel

$$\begin{aligned} \ker(K) &= \{(\widetilde{g_1}(v), g_2(v)) \mid v \in C[d]\} \\ &= \{(\widetilde{g_1}g_1(x), g_2g_1(x)) \mid x \in A_1[d]\} \\ &= \{(d_1x, g_2g_1(x)) \mid x \in A_1[d]\}. \end{aligned}$$

Indeed, g_1 is bijective on $A[d]$ because on the one hand, it is injective: $\ker(g_1) \cap A[d] \subseteq A[d_1] \cap A[d] = \{0\}$ because d_1 is prime to d_2 and therefore also to $d = d_1 + d_2$; and on the other hand $A[d], C[d]$ are finite sets of the same cardinality, so injectivity implies surjectivity. □

This lemma has deep consequences for computing isogenies. It immediately suggests that, if $A[d]$ is accessible (e.g. d is power-smooth), and $f = g_2g_1 = h_2h_1$ is known on (a basis of) $A[d]$, then we can compute f on all A efficiently using algorithms for computing the d-isogeny K in higher dimensions.

Prompted by our previous discussion, we ask whether it is possible to construct N-isogenies emanating from A for arbitrary N. Considering only N-endomorphisms, we answer this question positively with the following

Lemma 2.6 (Zarhin's Trick)**.** *Let A be an Abelian variety and m a non-negative integer. There exists an m-endomorphism α of A^s, whereby $s = 1, 2$ or 4 depending on whether m can be written as the sum of 1, 2 or 4 squares. Finding this endomorphism takes randomised expected time $O((\log(m))^2(\log\log(m))^{-1})$.*

We immediately note that *every* non-negative integer is the sum of at most four squares by Jacobi's four-square theorem. If A in this lemma is an elliptic curve (i.e. 1-dimensional), this lemma guarantees the existence of an endomorphism of A^s of *arbitrary* degree.

Proof of the lemma. We give a direct construction. When $m = m_1^2$ is a square, we may simply write $\alpha = [m_1]$. When $m = m_1^2 + m_2^2$, we can write down the matrix $\alpha = (m_1, m_2; -m_2, m_1)$ as an endomorphism of A^2. A brief computation shows that

$\alpha\widetilde{\alpha} = \alpha\alpha^t = (m_1^2 + m_2^2)\,\mathrm{id}_2$ and so α is an m-endomorphism. The same can be done in 4 dimensions when $m = m_1^2 + m_2^2 + m_3^2 + m_4^2$. Here we write

$$\alpha = \begin{pmatrix} m_1 & m_2 & m_3 & m_4 \\ -m_2 & m_1 & m_4 & -m_3 \\ -m_3 & -m_4 & m_1 & m_2 \\ -m_4 & m_3 & -m_2 & m_1 \end{pmatrix}$$

for which a computation shows that $\alpha\widetilde{\alpha} = \alpha\alpha^t = (m_1^2 + m_2^2 + m_3^2 + m_4^2)\,\mathrm{id}_4$ and so α is an m-endomorphism.

In any of these cases, we must find m_1, m_2, m_3, m_4 (some possibly zero), such that $m = m_1^2 + m_2^2 + m_3^2 + m_4^2$. This can be done in the advertised time [34]. In practice, sage has a fast implementation.

We note that α in no way depended on A. That is to say, the same $s \times s$ matrix α induces an m-endomorphism on all s-copies of PPAVs simultaneously.

Finally, we remark that the choices of the structures of the 2 and 4-dimensional matrices did not result from blind luck, but reflect ways to write complex numbers and quaternions as matrices so that their products are respected by matrix multiplication. That is, writing $M_z = (\mathrm{re}(z), \mathrm{im}(z); -\mathrm{im}(z), \mathrm{re}(z))$ for any complex number, we obtain $M_{zw} = M_z M_w = M_w M_z$. □

This lemma is sometimes referred to as *Zarhin's Trick*, for it is exactly this construction that is allows one to construct a principal polarisation on $A^4 \times (A^\vee)^4$ for *any* (not necessarily polarised) Abelian variety [31, Theorem. 11.29].

2.4 Computing higher-dimensional Isogenies

In this subsection, we will describe some algorithms to compute isogenies in higher dimensions. In analogy to Vélu's formulae, these tools take a model of an Abelian variety (c.f. Weierstrass or Montgomery forms for elliptic curves) together with a kernel to compute both what the target Abelian variety is and a formula for evaluating the isogeny corresponding to the kernel.

The computational cost of the formulae for computing an N-isogeny $f\colon A \to B$ will depend linearly on the degree of f. Recall, the degree of an r-dimensional N-isogeny is N^r.

We now assume that N is prime to the characteristic of the field over which A, B are defined, making f a separable isogeny. The best-known generic method of computing an N-isogeny f with composite $N = p_1^{e_1} \cdots p_n^{e_n}$ is to factor f into ${p_i}^{e_i}$-isogenies; and then compute these successively as a composition of p_i-isogenies. We can optimise the latter step using different *strategies*, entirely in analogy to [26, Sec. 4.2.2]. The procedure for factoring N-isogenies $A \to B$ is given by the proof of Proposition 13 in [18].

Since every N-isogeny $A \to B$ can be factored into the composition of smaller degree, we turn to computing p-isogenies for p prime. In [3] Dartois, Maino, Pope and Robert give explicit algorithms and implementations for computing 2-isogenies between

2-dimensional Abelian varieties. They use the nomenclature of "$(2,2)$-isogenies" for isogenies of this type, because their kernels are isomorphic to $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$. They give methods for effectively computing $(2^n, 2^n)$-isogenies as chains of $(2,2)$-isogenies.

We briefly remark on what this means practically. We recall that, ultimately, our goal remains to compute isogenies between *elliptic curves*, i.e. 1-dimensional isogenies. We will see in the next subsection, that this can be achieved by embedding 1-dimensional isogenies into an N-isogeny f of dimension $r = 2, 4$ or 8. In fact, this will be a *split* isogeny, that is, the domain and codomain of f will be products of elliptic curves. Moreover, we will see that we have almost complete freedom on N (in fact, we will choose N freely, and conclude from this what the dimension of embedding r will be). This tells us that a working implementation for $(2^n, 2^n)$-isogenies is a significant step in computing isogenies between elliptic curves. We will comment on the limitations of embedding to dimension $r = 2$ later.

Although the r-dimensional N-isogenies obtained from our embedding procedure will be between products of elliptic curves, the intermediate PPAVs A_i we obtain after factoring $f = f_n \cdots f_1 \colon A \to B$ into $f_i \colon A_i \to A_{i+1}$ will (very) likely not be. That is, they will be generic Abelian varieties of dimension r and not r-fold products of elliptic curves.

Every 2-dimensional Abelian variety is either the product of two elliptic curves or the Jacobian of a divisor [33, Sec. 2, Theorem of Weil]. Whilst there exist computational methods to compute on Jacobians [6; 32, Th. 3], Dartois et. al. use Mumford's generic *theta coordinates* because Weil's theorem does not generalise to higher dimensions: higher dimensional ($r \geq 4$) Abelian varieties must no longer be Jacobians or products of elliptic curve (although the converse is true, higher-dimensional Jacobians and products of elliptic curves are Abelian varieties) [18, App. F]. Using theta coordinates starting in dimension 2, although alternative computational methods exist, provides a consistent framework for the higher dimensions 4, 8 later on. Whilst SQISignHD does compute isogenies in 4-dimensions, the current library is only implemented for isogenies of a very specific form. Certain subroutines are hard-coded for this form and so there is no non-trivial generalisation of this library. However, this is to be remedied by Dartois, who has a new library for computation of generic 2-isogenies in dimension 4 using the theta coordinates framework [52][9]. Moreover, there have already been efficiency improvements to [3] in [5].

A fundamental challenge that the algorithms of [3, 19] must overcome is, that they take as input the kernel of the isogeny (and the domain) and so they only have enough information to compute the isogeny up to isomorphism. That is, given the kernel of an isogeny $f \colon A \to B$, they compute $g \colon A \to C$ with $\ker(g) = \ker(f)$ and $B \cong C$.

Fortunately, the library of [3] (which computes 2^n-isogenies $E_1 \times E_2 \to E_1' \times E_2'$ from their kernel) computes a final *splitting* isogeny for us, so that the map it evaluates g is of the form $g \colon E_1 \times E_2 \to E_1'' \times E_2'' \cong E_1' \times E_2'$ (and not $g \colon E_1 \times E_2 \to C$ where C is some arbitrary 2-dimensional PPAV).

[9] This was published days before the end of this thesis.

From Lemma 2.1, we know that every isogeny $f\colon E_1' \times E_2' \to E_1'' \times E_2''$ comes from a matrix $(f_{ji})_{ij}$ of isogenies $f_{ij}\colon E_i' \to E_j''$. In particular every isomorphism $\varphi\colon E_1' \times E_2' \xrightarrow{\sim} E_1'' \times E_2''$ is of this form. However, since degrees are non-negative and

$$\begin{pmatrix} \widetilde{\varphi_{1,1}} & \widetilde{\varphi_{1,2}} \\ \widetilde{\varphi_{2,1}} & \widetilde{\varphi_{2,2}} \end{pmatrix} \begin{pmatrix} \varphi_{1,1} & \varphi_{2,1} \\ \varphi_{1,2} & \varphi_{2,2} \end{pmatrix} = \begin{pmatrix} \widetilde{\varphi_{1,1}}\varphi_{1,1} + \widetilde{\varphi_{1,2}}\varphi_{1,2} & \widetilde{\varphi_{1,1}}\varphi_{2,1} + \widetilde{\varphi_{1,2}}\varphi_{2,2} \\ \widetilde{\varphi_{2,1}}\varphi_{1,1} + \widetilde{\varphi_{2,2}}\varphi_{1,2} & \widetilde{\varphi_{2,1}}\varphi_{2,1} + \widetilde{\varphi_{2,2}}\varphi_{2,2} \end{pmatrix} \overset{(!)}{=} [1]_E$$

we deduce that either $\deg(\varphi_{1,1}) = 1$ and so $\deg(\varphi_{1,2}) = 0, \deg(\varphi_{2,1}) = 0, \deg(\varphi_{2,2}) = 1$; or $\deg(\varphi_{1,1}) = 0$, and so $\deg(\varphi_{1,2}) = 1, \deg(\varphi_{2,2}) = 0, \deg(\varphi_{2,1}) = 1$. In other words, every isomorphism $\varphi\colon : E_1' \times E_2' \xrightarrow{\sim} E_1'' \times E_2''$ is of the form

$$\begin{pmatrix} \varphi_{1,1} & 0 \\ 0 & \varphi_{2,2} \end{pmatrix} \quad \text{or} \quad \begin{pmatrix} 0 & \varphi_{2,1} \\ \varphi_{1,2} & 0 \end{pmatrix}$$

with $\varphi_{ij}\colon E_i' \to E_j''$ isomorphisms. As in the one-dimensional case, an isogeny from computing a kernel is unique up to isomorphism, but now in the higher dimensional case (split isogenies between products of elliptic curves) we may also lose track of the ordering of the curves.

As a result, we know that $E_1' \cong E_1''$ and $E_2' \cong E_2''$ or $E_1' \cong E_2''$ and $E_2' \cong E_1''$. What this means in practice, is that instead of computing $f = (f_1, f_2; f_3, f_4)$ the library [3] may possibly compute $g = (\varphi_{1,1}, 0; 0, \varphi_{2,2})(f_1, f_2; f_3, f_4)$ or $g = (0, \varphi_{2,1}; \varphi_{1,2}, 0)(f_1, f_2; f_3, f_4)$ instead.

Using the Weil pairing we can recover which of these two isogenies the library computed. Let l be a small prime not dividing $\deg(f_1) = \deg(f_4)$ or $\deg(f_2) = \deg(3)$ (these degrees are known, because the Kani square is known) and let $m = (m_1, m_2; m_3, m_4)\colon E_1 \times E_2 \to E_1'' \times E_2''$ be the mystery isogeny computed by the library. Taking a basis P, Q of $E_1[l]$ and computing $m((P,0)) = (m_1(P), m_3(P))$ and $m((Q,0)) = (m_1(Q), m_3(Q))$ we get the images $m_1(P), m_1(Q)$ in E_1''. The Weil pairing on E_1 then tells us that

$$e_l^{E_1''}(m_1(P), m_1(Q)) = e_l^{E_1}(P,Q)^{\deg(m_1)}.$$

On the other hand, we can compute $e_l^{E_1}(P,Q)^{\deg(f_1)}$ and $e_l^{E_1}(P,Q)^{\deg(f_2)}$ manually. If

$$e_l^{E_1''}(m_1(P), m_1(Q)) = e_l^{E_1}(P,Q)^{\deg(f_1)}, \tag{2.2}$$

then we know that $\deg(m_1) \equiv \deg(f_1) \pmod{l}$, and since $\deg(f_1) \not\equiv \deg(f_2) \pmod{l}$, we conclude that $m_1 = \varphi_{1,1} f_1$. In particular $E_1' \cong E_1''$ (and $E_2 \cong E_2''$). Conversely, if (2.2) does not hold, then $E_1' \cong E_2''$ (and $E_2 \cong E_1''$) by the same argument over degrees.

This procedure can be done in any dimension, and we will call it the *Weil pairing trick*.

Finally, we note that for general degree isogenies there exist algorithms to compute r-dimensional l-isogenies in time $O((nl)^r)$ using theta coordinates of *level* n, where n is even and l prime to n [19, Theorem 1.3]. In Remark 4.8, the authors discuss how the condition of coprimality of n, l can be relaxed, but still require l to be odd. Whilst [3] also uses theta coordinates to compute isogenies, they work with level 2 (which is not coprime to the degree of the isogenies computed).

2.5 Application: Breaking SIDH by embedding isogenies

Zarhin's trick together with Kani's lemma delivers Robert's embedding

Lemma 2.7 (Robert's embedding lemma)**.** *Every N-isogeny $f\colon A \to B$ between r-dimensional principally polarised Abelian varieties can be embedded into an M-isogeny*

$$K = (f^{(s)}, \widetilde{\alpha_B}; -\alpha_A, \widetilde{f}^{(s)})\colon A^s \times B^s \to B^s \times A^s$$

in dimension $2sr$, for any $M > N$. Here α_A, α_B are $(M-N)$-endomorphisms of A^s, B^s respectively, both given by the same $s \times s$ integer matrix obtained from Lemma 2.6. Hence $s = 1, 2$ or 4 depending on whether $M - N$ is the sum of $1, 2$ or 4 squares.

Proof. Because $f^{(s)}, \widetilde{f}^{(s)}$ are s-copies of $f, \widetilde{f}$, they act like scalar matrices and commute with everything. As such,

$$\begin{array}{ccc} A^s & \xrightarrow{f^{(s)}} & B^s \\ {\scriptstyle \alpha_A}\downarrow & & \downarrow{\scriptstyle \alpha_B} \\ A^s & \xrightarrow[f^{(s)}]{} & B^s \end{array}$$

is a Kani square and the Kani map K is a $N + (M - N) = M$ isogeny.

If M is prime to N, then $M - N$ prime to N and the square is minimal. As such, the kernel of K is given by

$$\ker(K) = \left\{ (Nx, \alpha_B f^{(s)}(x) \mid x \in A^s[M] \right\}.$$

□

Corollary 2.8 ([17, Prop. 2.9])**.** *Let $f\colon A \to B$ be an N-isogeny. If f is known on (a basis of) $A[l_i^{e_i}]$ for sufficiently many primes l_i such that $M = \prod_{i=1}^{n} l_i^{e_i} > N$ is coprime to N, then the embedding map K can be efficiently computed.*

Proof. We use Robert's embedding lemma to embed f into an M-isogeny in dimension $2sr$ where s is the number of squares required to write $M - N$ as. Because we know f on $A[l_i^{e_i}]$, we know f on $A[M]$ and can evaluate f on $A[M]$ in $O(\log(N)\log\log(N))$ [16, Lem. 3.3] operations on A. As such, the kernel of K can be computed efficiently. □

We now illustrate how these embedding lemmas can be used to break the SIDH [26, Sec. 3.2] key-exchange.

Very briefly accounted, the exchange occurs as follows. A public supersingular elliptic curve E_0 is chosen, a basis P_A, Q_A of $E[l_A^{e_A}]$ and a basis P_B, Q_B of $E[l_B^{e_B}]$ with l_A, l_B coprime primes. Alice and Bob choose random invertible elements a, b from $\mathbb{Z}/l_A^{e_A}\mathbb{Z}$ and $\mathbb{Z}/l_B^{e_B}\mathbb{Z}$ respectively. They then each compute an isogeny $\varphi_A\colon E_0 \to E_A, \varphi_B\colon E_0 \to E_B$ from the kernels $K_A = \langle P_A + aQ_A \rangle$ and $K_B = \langle P_B + bQ_B \rangle$. Finally, Alice and Bob publish the triples $(E_A, \varphi_A(P_B), \varphi_A(Q_B))$ and $(E_B, \varphi_B(P_A), \varphi_B(Q_A))$ respectively.

To recover a shared secret with Bob, Alice writes down $K_{AB} = \langle \varphi_B(P_A) + a\varphi_B(Q_A) \rangle$ and computes the image E_S of the isogeny $\varphi_{AB}\colon E_B \to E_S$ with kernel K_{AB}. Since

l_A is prime to l_B, the points $\varphi_B(P_A), \varphi_B(Q_A)$ form a basis of the $l_A^{e_A}$-torsion on E_B. Consequently, the isogeny $\varphi_{AB}\varphi_B \colon E_0 \to E_S$ has kernel $\ker(\varphi_A) + \ker(\varphi_B)$. This symmetry means that when Bob writes down $K_{BA} = \langle\, \varphi_A(P_B) + b\varphi_A(Q_B) \,\rangle$ and computes $\varphi_{BA} \colon E_A \to E_{S'}$, the curve $E_{S'}$ is isomorphic to E_S. Therefore, choosing the shared secret to be the j-invariant of $E_S, E_{S'}$, Bob and Alice will obtain the same secret value.

Now, given Bob's torsion point information $\varphi_B(P_A), \varphi_B(Q_A)$, we can recover Bob's secret isogeny φ_B. We assume that $l_A^{e_A} > l_B^{e_B}$. Let α be the $s \times s$ matrix given by Lemma 2.6 such that α induces a $(l_A^{e_A} - l_B^{e_B})$-endomorphism on E_0 and E_B. Embedding φ_B into the Kani map $K = (\varphi_B^{(s)}, \widetilde{\alpha_{E_B}}; -\alpha_{E_0}, \varphi_B^{(s)})$ gives a $l_A^{e_A}$-isogeny in dimension $r = 2s$ whose kernel is given by

$$\begin{aligned}\ker(K) &= \left\{ (l_A^{e_A} x, \alpha_{E_B}\varphi_B^{(s)}(x) \mid x \in A[l_A^{e_A}] \right\} \\ &= \left\langle\, (l_A^{e_A} P_A, \alpha_{E_B}\varphi_B^{(s)}(P_A)), (l_A^{e_A} Q_A, \alpha_{E_B}\varphi_B^{(s)}(Q_A)) \,\right\rangle\end{aligned}$$

We note that the Kani square corresponding to K is indeed minimal. Since P_A, Q_A are a basis of $E_0[l_A^{e_A}]$ we can compute the kernel immediately using the Bobs's torsion point information. Moreover, since K is a smooth isogeny, given its kernel, we can efficiently recover the values of φ_B.

2.6 Application: Computing the isogeny class-group action

In this subsection we describe Page and Robert's CLAPOTI method for computing the isogeny class-group action in polynomial time [2].

Until now, we have required some torsion knowledge. That is, to compute an isogeny f (or its embedding into higher dimension K), we needed to know the image of f on some smooth torsion. We wish to forgo this requirement. We will see that it is still possible to compute the class-group action using the same techniques, but without torsion information.

As before, we begin with a d_f-isogeny $f \colon A \to B$ between PPAVs. Suppose we can find another d_g-isogeny $g \colon A \to B$ with d_g prime to d_f. Then by forming the reversals $\widetilde{f}, \widetilde{g} \colon B \to A$, we can then build the commuting square

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ {\scriptstyle j}\downarrow & & \downarrow{\scriptstyle \widetilde{g}} \\ C & \xrightarrow[\widetilde{h}]{} & A \end{array}$$

by reversing the pushout square of $\widetilde{f}$ and $\widetilde{g}$

$$\begin{array}{ccc} B & \xrightarrow{\widetilde{f}} & A \\ {\scriptstyle \widetilde{g}}\downarrow & & \downarrow{\scriptstyle j} \\ A & \xrightarrow[h]{} & C \end{array}$$

It is a minimal Kani square, and so its Kani map $K = (f, g; -j, h)$ is a $N = d_f + d_g$ isogeny with kernel $\ker(K) = \left\{ (d_f x, \widetilde{g} f(x) \mid x \in A[N] \right\}$.

We immediately note that $\widetilde{g} f \colon A \to A$ is an endomorphism. We hope, that given a suitable representation, that this should be easier to compute than a generic isogeny. Indeed, if A is an elliptic curve, computing $\widetilde{g} f$ should be easy if the endomorphism ring is known, reducing the problem of computing the kernel of K to the N-torsion being accessible.

Unfortunately, there is no reason why the N-torsion should be accessible, for it is only the sum of two integers d_f, d_g. In a similar fashion to before, we fix this problem by composing f and g with a, b-endomorphisms α, β of A such that the new Kani map K is an $N = ad_f + bd_g$ isogeny. Since d_f, d_g are coprime we hope to achieve any $N > d_f d_g$ with this method and so can choose an arbitrary N so that $A[N]$ is accessible.

Without any specific knowledge on the endomorphism ring of A, we again use Lemma 2.6 to obtain α, β as $s \times s$ integer matrices yielding the isogenies $f^{(s)}\alpha, g^{(s)}\beta \colon A^s \to B^s$. We write down the (not necessarily minimal) Kani square

$$\begin{array}{ccc} A^s & \xrightarrow{f^{(s)}\alpha} & B^s \\ {\scriptstyle J}\downarrow & & \downarrow{\scriptstyle \widetilde{g^{(s)}\beta}} \\ C^s & \xrightarrow[\widetilde{H}]{} & A^s \end{array}$$

obtained by reversing, pushing out, and reversing again, as before. The resulting Kani map $K = (f^{(s)}\alpha, g^{(s)}\beta; -J, H)$ is an N-isogeny by construction. For now, we are not too interested in what the maps J, H actually are, only that they exist.

If $\gcd(a, b) = g > 1$, we can replace N by N/g, a by a/g, and b by b/g because if the N-torsion is accessible, then so is the N/g-torsion. So we may assume that a, b are coprime, making ad_f prime to bd_g, the Kani square minimal and the kernel of K easy to write down

$$\ker(K) = \left\{ \left(ad_f x, \widetilde{g^{(s)}\beta} f^{(s)}\alpha(x) \right) \mid x \in A[N] \right\}.$$

Moreover, since $\widetilde{g^{(s)}\beta} = \widetilde{\beta}\,\widetilde{g^{(s)}} = \widetilde{\beta}(\widetilde{g})^{(s)}$ and $f^{(s)}, (\widetilde{g})^{(s)}$ are scalar matrices they commute with α, β and we can rewrite the endomorphism as

$$\widetilde{g^{(s)}\beta} f^{(s)}\alpha = \beta\alpha(\widetilde{g}f)^{(s)}.$$

where we stress that $(\widetilde{g}f)^{(s)}$ is a s-diagonal copy of an endomorphism on A.

We still have not solved the problem of finding the d_g-isogeny $g \colon A \to B$ with d_g coprime to d_f.

We tackle this problem in the context of elliptic curves. Namely, given an isogeny $f \colon A \to B$ between elliptic curves, find another isogeny $g \colon A \to B$ with degree $\deg(g)$ coprime to $\deg(f)$. Translating this problem to the world of ideals using the ideal-to-isogeny correspondence we obtain a partial solution that can be used to compute the isogeny class-group action.

Let $\mathcal{O}$ be an order inside the ring of integers $\mathcal{O}_D$ of the imaginary quadratic number field $\mathbb{Q}(\sqrt{D})$, where $D < 0$ squarefree. We recall that $\mathrm{Ell}_q(\mathcal{O})$ denotes the set of (ordinary) elliptic curves E over F_q with complex multiplication by $\mathcal{O}$, that is $\mathrm{End}(E) \cong \mathcal{O}$; and that $[I]$ in $\mathrm{Cl}(\mathcal{O})$ acts on E in $\mathrm{Ell}_q(\mathcal{O})$ yielding E_I, where E_I is the codomain of the isogeny φ_I whose kernel is given by the intersection of kernels $\ker(\sigma)$ for all σ in I.

If J is another representative of $[I]$, then $E_J \cong E_I$. Moreover, if the norm of J is prime to the norm of I, we almost have what we want: $\varphi_I \colon E \to E_I$ and $\varphi_J \colon E \to E_J$ are isogenies of coprime norm, with the same domain and *almost* the same codomain.

We note at this point, that E_I, E_J are only well defined up to isomorphism anyway, but when writing down the square

$$\begin{array}{ccc} E & \xrightarrow{\varphi_I} & E_I \cong E_J \\ \downarrow & & \downarrow{\scriptstyle \widetilde{\varphi_J}} \\ * & \longrightarrow & E \end{array}$$

and computing the kernel of the arising Kani map, we need to know what the isomorphism $E_I \xrightarrow{\sim} E_J$ is. This can be done efficiently, because every isomorphism of ordinary elliptic curves is given by an element in F_q^d where d is an extension of degree at most 6.

Page-Robert fix this problem without extensions, by writing down the minimal Kani square

$$\begin{array}{ccc} E & \xrightarrow{\varphi_I} & E_I \\ {\scriptstyle \varphi'_{\bar{J}}}\downarrow & & \downarrow{\scriptstyle \varphi_{\bar{J}}} \\ E_{\bar{J}} & \xrightarrow[\varphi'_I]{} & E \end{array} \tag{2.3}$$

Because $\mathcal{O}$ acts by horizontal isogenies, the maps φ maps in the diagram above are all well-defined. That is, since $\mathcal{O} \cong \mathrm{End}(E) \cong \mathrm{End}(E_I) \cong \mathrm{End}(E_{\bar{J}})$, we can pick ideals $I, J \subseteq \mathcal{O}$ and form the maps φ_I emanating on any of the curves $E, E_I, E_{\bar{J}}$. The square commutes because $\varphi_I\varphi_J = \varphi_{IJ}$ and $\mathcal{O}$ is commutative. Finally, the square is a minimal because $\deg(\varphi_I) = \mathrm{N}(I)$ is prime to $\deg(\varphi_{\bar{J}}) = \mathrm{N}(\bar{J}) = \mathrm{N}(J)$.

The corresponding Kani map $K = (\varphi_I, \widehat{\varphi_{\bar{J}}}; -\varphi'_{\bar{J}}, \widehat{\varphi'_I})$ has kernel

$$\ker(K) = \left\{ (\mathrm{N}(I)x, \varphi_{\bar{J}}\varphi_I(x)) \mid x \in E[N] \right\}.$$

where $N = \mathrm{N}(I) + \mathrm{N}(J)$. The endomorphism $\varphi_{\bar{J}}\varphi_I = \varphi_{\bar{J}I}$ corresponds to φ_γ where $\bar{J}I = \gamma\mathcal{O}$ and γ lies in $\mathcal{O}$. Indeed, since $[I] = [J]$, there exists by definition a μ in $\mathrm{Frac}(\mathcal{O})$ so that $J^{-1}I = \mu\mathcal{O}$; hence $\bar{J}I = \mathrm{N}(J)J^{-1}I = \mathrm{N}(J)\mu\mathcal{O}$ which is an integral ideal because I, J are. Setting $\gamma = \mathrm{N}(J)\mu$ proves the claim. Moreover, γ can be computed efficiently by finding the smallest element in the lattice $\bar{J}I$ by way of a generalised gcd computation, known as the *Lagrange-Gauss Algorithm* [22, Alg. 23].

As before, the $N = \mathrm{N}(I) + \mathrm{N}(J)$ torsion may not be available. So we lift the square (2.3) into higher dimensions with the $s \times s$ integer matrix a, b-endomorphisms α, β until the $N = a\mathrm{N}(I) + b\mathrm{N}(J)$ torsion is accessible on E.

There is one subtlety to keep in mind however, namely that in dimension 4 endomorphisms do not commute (in dimension 1 and 2 they do). To remedy this, we can perform a *metacommutation* (viewing the 4×4 matrices as quaternions) and find matrices α', β' so that $\beta\alpha = \alpha'\beta'$ and that α' is a b-endomorphism, and β' an a-endomorphism. If $s = 1$ or 2, we simply write $\alpha' = \beta$ and $\beta' = \alpha$. Then

$$\begin{array}{ccc} E^s & \xrightarrow{\varphi_I^{(s)}\alpha} & E_I^s \\ {\scriptstyle \varphi'^{(s)}_{\bar{J}}\beta'}\downarrow & & \downarrow{\scriptstyle \beta\varphi_{\bar{J}}^{(s)}} \\ E_{\bar{I}}^s & \xrightarrow[\alpha'\varphi'^{(s)}_I]{} & E^s \end{array}$$

is a commuting minimal Kani square. Indeed,

$$\beta\varphi_{\bar{J}}^{(s)}\varphi_I^{(s)}\alpha = \beta\varphi_{\bar{J}I}^{(s)}\alpha = \beta\alpha\varphi_{\bar{J}I}^{(s)} = \alpha'\beta'\varphi_{\bar{J}I}^{(s)} = \alpha'\varphi_{JI}^{(s)}\beta' = \alpha'\varphi_{JI}'^{(s)}\beta' = \alpha'\varphi_{\bar{J}}'^{(s)}\varphi_I'^{(s)}\beta'$$

and because α' is a b-endomorphism and β' an a-endomorphism, the parallel degrees match. Note that in dimensions $s = 2, 4$, $\varphi_*^{(s)}$ is a scalar matrix and commutes with everything; when $s = 1$, $\varphi_{\bar{J}I}$ is an endomorphism, and so commutes with the other endomorphisms $\alpha, \beta, \alpha', \beta'$

Unfortunately, metacommutation exists only in the *Hurwitz* quaternions. That is, α', β' may have half-integer entries. This can be fixed by replacing N with $4N$, a with $2a$, b with $2a$, α with 2α and β with 2β. Then we must compute the 4 torsion lying above $4N$, which is a one-time computation. We continue writing N for simplicity.

The Kernel of the Kani map

$$K = \left(\varphi_I^{(s)}\alpha, \widetilde{\beta\varphi_{\bar{J}}^{(s)}}; -\varphi_{\bar{J}}'^{(s)}\beta', \widetilde{\alpha'\varphi_I'^{(s)}}\right) = \left(\varphi_I^{(s)}\alpha, \widetilde{\varphi_{\bar{J}}}^{(s)}\widetilde{\beta}; -\varphi_{\bar{J}}'^{(s)}\beta', \widetilde{\varphi_I'}^{(s)}\widetilde{\alpha'}\right)$$

is given by

$$\ker(K) = \left\{(Nx, \beta\alpha\varphi_{I\bar{J}}^{(s)}(x)) \mid x \in E^s[N]\right\} = \left\{(Nx, \beta\alpha\varphi_Y^{(s)}(x)) \mid x \in E^s[N]\right\}. \quad (2.4)$$

Using the algorithms of [19], we can now compute K in polynomial time, and using the Weil pairing trick, obtain an equation for E_I.

Using the library of [3], we do not need to deal with metacommutations to obtain E_I. Indeed, the description of the kernel (2.4) does not need to know what the other two maps (left and bottom) in the Kani square are. We know that they must exist, by the existence of pushouts.

2.7 Algorithms for Clapoti

In this subsection we present more explicitly the algorithms required to compute the Clapoti method in polynomial time.

Our first challenge is to, given an invertible integral ideal I, find an equivalent integral ideal J with coprime norm. Finding such an ideal is the content of the following

Lemma 2.9 ([2, Lem. 2.5]). *Given an invertible integral ideal $I \subseteq \mathcal{O}$ it is possible to find an integral ideal $J \subseteq \mathcal{O}$ equivalent to I but with coprime norm. This algorithm runs in randomised polynomial time $O(\log(\mathrm{N}(I)))$.*

Proof. We recall that $I^{-1} = \bar{I}\mathcal{O}/\mathrm{N}(I)$ and so for any α in I^{-1} we get an integral ideal αI. If the norm of αI is prime to that of I, we are done. We can simplify this procedure a little, by choosing α in $\bar{I}$, and verifying that $\gcd(\mathrm{N}(I), \mathrm{N}(\alpha I)) = \mathrm{N}(I)^2$ instead.

The assignment $\alpha \to \mathrm{N}(\alpha I)$ is a (positive definite) quadratic form in 2 variables. More precisely, if we let $I = m[a, b + fw_D]$, then $\bar{I} = m[a, -(b + D) + fw_D]$ and writing $\alpha = \alpha(x, y) = xma + ym(-(b + D) + fw_D)$ we obtain $\alpha(x, y) \mapsto \mathrm{N}(\alpha I)$, which is a positive definite quadratic form in 2 variables x, y.

Using [18, Lem. 4.2.1] we can sample randomly from the set

$$\{(x, y) \in \mathbb{Z} \times \mathbb{Z} \mid \mathrm{N}(\alpha(x, y)I) \leq \rho\}$$

in time polynomial in $\log(\rho)$ and $\log(\mathrm{N}(I))$. Picking $\rho = \mathrm{poly}(\mathrm{N}(I))$, we expect to find a norm coprime to that of I in polynomial time of $\log(\mathrm{N}(I))$. □

Let I be an invertible integral ideal of the ideal class-group $\mathrm{Cl}(\mathcal{O})$. The following algorithm computes the isogeny $E \to E_I$ in polynomial time.

(i) Use Lemma 2.9 to find an equivalent integral invertible ideal J of coprime norm.
(ii) Choose $N \geq \mathrm{N}(I)\mathrm{N}(J)$ coprime to $\mathrm{N}(I)\mathrm{N}(J)$; so that the N-torsion $E[N]$ is accessible. For example, we may choose N powersmooth.
A more concrete example is to choose $N = p^l$ where p is the smallest prime not dividing $\mathrm{N}(I)\mathrm{N}(J)$, and l the smallest l so that $p^l \geq \mathrm{N}(I)\mathrm{N}(J)$.
(iii) Since $\mathrm{N}(I), \mathrm{N}(J)$ are coprime there exist integral a, b so that $a\mathrm{N}(I) + b\mathrm{N}(J) = N$. Moreover, since $N \geq \mathrm{N}(I)\mathrm{N}(J)$ we may choose $a, b \geq 0$.
(iv) Let $g = \gcd(a, b)$. If $g > 1$, replace N with N/g, a with a/g and b with b/g. Now a, b are coprime, and in particular $a\mathrm{N}(I)$ is coprime to $b\mathrm{N}(J)$.
(v) Apply Lemma 2.6 to find a, b endomorphisms α, β
(vi) Use our construction to obtain an isogeny $f\colon E^s \times E^s \to E_I^s \times E_{\overline{J}}^s$.
(vii) By methods of [3, 19], the isogeny f can be evaluated in polynomial time, in particular equations for the codomains are given. We can use the Weil pairing trick to to distinguish E_I and $E_{\overline{J}}$.

2.8 Consequences for isogeny-based cryptography

We conclude this section by discussing the impact that Kani's lemma and its applications have on isogeny-based cryptography.

We stress that Robert's embedding lemma (Lemma 2.7) and the attacks stemming from it fundamentally require the knowledge of $f\colon A \to B$ on some torsion $A[d]$. This cannot be circumvented. From that perspective, one could say that the SIDH protocol was incredibly unlucky, because it stumbled upon an alternative efficient representation of isogenies.

Since the whole Lemma relies on this information, adaptations of SIDH have been proposed that mask the torsion information of f [76].

We also stress that efficient computation of the isogeny-class-group using Clapoti does not impact the security of a key-exchange derived from this action. The computation requires knowledge of the secret ideal class, and so does not constitute an attack.

3

Class-groups of Orders in Imaginary Quadratic Number Fields

We know from Deuring, that the endomorphism of an ordinary elliptic curve over a finite field is isomorphic to an *order* inside an imaginary quadratic field. By the ideal-to-isogeny correspondence, every ideal of the endomorphism ring $\mathrm{End}(E)$ of a curve E corresponds to an isogeny. The composition of isogenies corresponds to the product of ideals. This gives a natural monoid structure on the set of ideals. Since the principal ideals correspond to endomorphisms it is natural to form the quotient by principal ideals. We will see that a more careful treatment leads to the *ideal class-group*.

In this section we introduce the ideal class-group, and give some algorithms for the ideals of imaginary quadratic number fields. In general, an ideal class group can be attached to any domain. We will present some general theory, but focus on the cases that interest us.

3.1 Orders

A finite (and therefore algebraic) extension of the rationals is called a *number field*. We define the *ring of integers* $\mathcal{O}_k$ of a number field k to be the integral closure of the integers in k. That is, $\mathcal{O}_k$ comprises the elements of k that satisfy a (univariate) monic polynomial with integer coefficients.

We note that the ring of integers $\mathcal{O}_k$ in k play the role of the integers $\mathbb{Z}$ in $\mathbb{Q}$ in the sense that $\mathrm{Frac}(\mathcal{O}_k) = k$. Indeed, let x in k satisfy $a_n x^n + a_{n-1}x^{n-1} + \cdots + a_0 = 0$ with $a_n \neq 0$ and a_i in $\mathbb{Q}$. We may assume that the a_i are in fact integers, otherwise we multiply through by the product of all their denominators. Multiplying through by a_n^{n-1} we obtain

$$a_n^n x^n + a_n^{n-1}a_{n-1}x^{n-1} + \cdots + a_n^{n-1}a_0 = (a_n x)^n + a_{n-1}(a_n x)^{n-1} + \cdots + a_n^{n-1}a_0 = 0$$

and so $a_n x$ satisfies a monic polynomial with integer coefficients. Hence $a_n x$ lies in $\mathcal{O}_k$ and so $x = (a_n x)/a_n$ must lie in $\mathrm{Frac}(\mathcal{O}_k)$. The converse $\mathrm{Frac}(\mathcal{O}_k) \subseteq k$ is clear. We will see more properties later on which $\mathbb{Z}$ and $\mathcal{O}_k$ share.

An *order* of a number field k is a subring that is generated as a $\mathbb{Z}$-module by a $\mathbb{Q}$-basis of k. Equivalently, it is a subring $\mathcal{O}$ that is finitely generated as a free $\mathbb{Z}$-module for which we obtain k when we extend the scalars to $\mathbb{Q}$; in formulaic terms: $\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Q} = k$. We note that automatically the $\mathbb{Z}$-rank of $\mathcal{O}$ is $\dim_{\mathbb{Q}}(k)$. Indeed, elements of $\mathcal{O}$ are $\mathbb{Z}$-linearly independent if and only if they are $\mathbb{Q}$-linearly independent. Since $\mathcal{O}$ is a finitely generated $\mathbb{Z}$-module, so is any ideal I of $\mathcal{O}$; hence I is also a finitely generated $\mathcal{O}$-module. Consequently, any order $\mathcal{O}$ is a *Noetherian* ring.

In showing that $\mathrm{Frac}(\mathcal{O}_k) = k$, we actually showed that $\mathcal{O}_k \otimes \mathbb{Q} = k$ because the conclusion $x = (a_n x)/a_n$ lying in k only required dividing by integer a_n, not a generic

element in $\mathcal{O}_k$. Since $\mathcal{O}_k$ is finitely generated as a $\mathbb{Z}$-module [13, Ch. 2], we conclude that $\mathcal{O}_k$ is an order.

It turns out that we can classify the orders of number fields very compactly.

Theorem 3.1 ([7, Th. 13.26, Th. 13.27])**.** *Let k be a number field. Then*

(i) every order $\mathcal{O}$ of k is contained in $\mathcal{O}_k$; and
(ii) every order $\mathcal{O}$ of k can be written as $\mathcal{O} = \mathbb{Z} + f\mathcal{O}_k$.

We call $|f|$ the conductor *of $\mathcal{O} = \mathbb{Z} + f\mathcal{O}_k$. It is equal to the index $[\mathcal{O}_k : \mathcal{O}]$.*

This justifies calling $\mathcal{O}_k$ *the* maximal order in k.

Proof. To show (i), we simply demonstrate that every element α in $\mathcal{O}$ is integral over $\mathbb{Z}$. To that end, consider the $\mathbb{Z}$-submodule Ω of $\mathcal{O}$ generated by powers of α. It is finitely generated, because $\mathcal{O}$ is and let $\beta_1, \dots, \beta_l$ be a $\mathbb{Z}$-generating set. Each of the generators β_i is an element of Ω, so a $\mathbb{Z}$-linear combination of power of α. Let N be the largest power of α that appears in any β_i. Then α^{N+1} is contained in Ω and equal to a linear $\mathbb{Z}$-linear combination $\lambda_1\beta_1 + \cdots + \lambda_l\beta_l$. As such, $\alpha^{N+1} - (\lambda_1\beta_1 + \cdots + \lambda_l\beta_l) = 0$, and replacing each α^n with x^n we obtain a polynomial in x with integer coefficients which is satisfied by α. Hence α is integral over $\mathbb{Z}$, and lies in $\mathcal{O}_k$.

Point (ii) now follows rather swiftly. Since $\mathcal{O} \subseteq \mathcal{O}_k$, and both are free $\mathbb{Z}$-modules of finite rank, the index $f = [\mathcal{O}_k : \mathcal{O}]$ is finite. Because $\mathcal{O}$ is a $\mathbb{Z}$-submodule of $\mathcal{O}_k$, the map $\mathcal{O}_k \to \mathcal{O}_k; x \mapsto fx$ preserves $\mathcal{O}$ and descends to a well-defined map on the quotient $\mathcal{O}_k/\mathcal{O}$. In fact, it descends to the zero-map, and so we conclude that $f\mathcal{O}_k \subseteq \mathcal{O}$. Then $\mathbb{Z} + f\mathcal{O}_k \subseteq \mathcal{O}$. However, $\mathbb{Z} + f\mathcal{O}_k$ also has index f in $\mathcal{O}_k$, so we must have equality $\mathcal{O} = \mathbb{Z} + f\mathcal{O}_k$. □

When k is a 2-dimensional extension of $\mathbb{Q}$, we say k is a *quadratic* number field. Every quadratic number field can be uniquely realised as $\mathbb{Q}(\sqrt{D})$ where D is a squarefree integer, not 0 or 1. For this reason, we write $\mathcal{O}_D$ for the ring of integers of $k = \mathbb{Q}(\sqrt{D})$. When $D > 0$ we call $\mathbb{Q}(\sqrt{D})$ a *real* quadratic number field, else $\mathbb{Q}(\sqrt{D})$ is called *imaginary*. The *conjugate* off an element $z = a + b\sqrt{D}$ is denoted $\overline{z} = a - b\sqrt{D}$. Using Theorem 3.1, we immediately see that that orders are stable under conjugation. Indeed, α and $\overline{\alpha}$ satisfy the same polynomials, provided they have real coefficients, and so the ring of integers $\mathcal{O}_k$ is closed under conjugation.

Finally, for $\alpha = a + b\sqrt{D}$ in $k = \mathbb{Q}(\sqrt{D})$, we define the *trace* of α to be $\mathrm{tr}(\alpha) = \alpha + \overline{\alpha} = 2a$ and the *norm* to be $\mathrm{N}(\alpha) = \alpha\overline{\alpha} = a^2 - b^2D$. We note that this definition of the trace and the norm *a priori* depend on the way we represent elements in k. More precisely, they appear to depend on the basis $k/\mathbb{Q}$ we choose. Fortunately, this is not the case. The multiplication-by-α map on k is $\mathbb{Q}$-linear, and so, may be represented by a 2×2 matrix $M_\alpha = (a, bD; b, a)$ with rational entries. A short computation verifies this. Then the trace of α is equal to the trace of M_α and the norm of α, the determinant of M_α. This perspective dispels any concerns, because the trace and determinant of matrices are famously basis-independent. This also tells us immediately, that the norm is multiplicative, and that $\mathrm{N}(c\alpha) = c^2\mathrm{N}(\alpha)$ for any rational c.

We can describe the rings of integers of quadratic number fields rather compactly.

Theorem 3.2. *The ring of integers of the quadratic number field $\mathbb{Q}(\sqrt{D})$ is given by*

$$\mathcal{O}_D = \begin{cases} \mathbb{Z}[\sqrt{D}] = \mathbb{Z} + \sqrt{D}\mathbb{Z} & D \not\equiv 1 \pmod 4 \\ \mathbb{Z}[(1+\sqrt{D})/2] = (1/2)\mathbb{Z} + (\sqrt{D}/2)\mathbb{Z} & D \equiv 1 \pmod 4 \end{cases}$$

writing $\Delta_D = D$ when $D \equiv 1 \pmod 4$ and $\Delta_D = 4D$ else, we obtain

$$\mathcal{O}_D = \mathbb{Z}\left[(\Delta_D + \sqrt{\Delta_D})/2\right].$$

The number Δ_D is called the *discriminant* of $\mathbb{Q}(\sqrt{D})$ and uniquely defines the field since $\mathbb{Q}(\sqrt{D}) = \mathbb{Q}(\sqrt{\Delta_D})$.

Before proceeding to the proof, we note that an element $\alpha = a + b\sqrt{D}$ in k is integral over $\mathbb{Z}$ if and only if $\mathrm{tr}(\alpha)$ and $\mathrm{N}(\alpha)$ are integers. Indeed, the minimal polynomial of α is $x^2 - (\alpha + \overline{\alpha})x + \alpha\overline{\alpha} = x^2 - \mathrm{tr}(\alpha)x + \mathrm{N}(\alpha)$.

Proof. Suppose $\alpha = a + b\sqrt{D}$ lies in $\mathcal{O}_D$. Then $2a$ is an integer and we perform a case distinction to whether a is an integer or a half-integer.

Case 1a Let a be an integer. Then b^2D is an integer and writing $b = \lambda/\mu$ in simplest form we see that μ^2 must divide D. However, D is squarefree and so $\mu = 1$. Consequently, we have shown that if a is an integer, then so is b. Hence $\mathcal{O}_D \subseteq \mathbb{Z} + \sqrt{D}\mathbb{Z} = \mathbb{Z}[\sqrt{D}] \subseteq \mathbb{Z}[(1+\sqrt{D})/2]$.

Case 2a Now let a be a half integer $\gamma/2$ with γ odd. Then $\gamma^2/4 - b^2D = \zeta$ is an integer, and so $4b^2D$ is. Again writing $b = \lambda/\mu$ in simplest form we see that μ^2 must divide $4D$. Since D is squarefree $\mu = 1$ or 2. If $\mu = 1$, then b is an integer, and so $\mathrm{N}(\alpha) = \gamma^2/4 - b^2D$ being integer implies $\gamma^2/4$ is an integer. This can only be if γ is even. A contradiction. Hence $\mu = 2$, and λ odd. Now we have $\gamma^2 - 4b^2D = \gamma^2 - D\lambda^2 = 4\zeta$. Since γ and λ are odd, we have $\gamma^2 \equiv \lambda^2 \pmod 4$ and so $D \equiv 1 \pmod 4$ must be true. In this case, we have shown that $\mathcal{O}_D \subseteq \mathbb{Z}[(1+\sqrt{D})/2]$.

In conclusion $\mathcal{O}_D \subseteq \mathbb{Z}[\sqrt{D}]$ when $D \not\equiv 1 \pmod 4$ and $\mathcal{O}_D \subseteq \mathbb{Z}[(1+\sqrt{D})/2]$ when $D \equiv 1 \pmod 4$. For the converse, we also perform a case distinction.

Case 1b If α lies in $\mathbb{Z}[\sqrt{D}]$, then $\alpha = a + b\sqrt{D}$ with a, b integers and so α lies in $\mathcal{O}_D$ because the norm and trace are integers. Together with Case 1a we have equality $\mathcal{O}_D = \mathbb{Z}[\sqrt{D}]$ when $D \not\equiv 1 \pmod 4$.

Case 2b If $\alpha = a + b\sqrt{D}$ lies in $\mathbb{Z}[(1+\sqrt{D})/2]$ with half-integers a, b, then the trace $2a$ is clearly an integer, and the norm $a^2 - D^2b^2$ is an integer when $D \equiv 1 \pmod 4$. Indeed $a = \gamma/2, b = \lambda/2$ with γ, λ odd, so $\gamma^2, \lambda^2 \equiv 1 \pmod 4$ and so $\gamma^2 - D\lambda^2 \equiv 0 \pmod 4$ when $D \equiv 1 \pmod 4$; we conclude that $\gamma^2 - D\lambda^2 = 4\zeta$ so $\mathrm{N}(\alpha) = a^2 - Db = \zeta$ is an integer. Together with Case 2a, we have equality $\mathcal{O}_D = \mathbb{Z}[(1+\sqrt{D})/2]$.

This shows our first description of $\mathcal{O}_D$. The second $\mathcal{O}_D = \mathbb{Z}[(\Delta_D + \sqrt{D})/2]$ follows immediately by definition. □

If we write $w_D = (\Delta_D + \sqrt{\Delta_D})/2$ we get $\mathcal{O}_D = \mathbb{Z}[w_D]$. Moreover, $w_D^2 - \Delta_D w_D + (\Delta_D^2 -$

$\Delta_D)/4 = 0$ and so $\mathbb{Z}[w_D] = \mathbb{Z} + w_D\mathbb{Z}$. Notably, we see that $\mathrm{N}(w_D) = (\Delta_D^2 - \Delta_D)/4$ is an integer. This also follows from the fact that $\Delta_D = 0, 1 \pmod 4$.

We can translate Theorem 3.1 to tell us that every order $\mathcal{O}$ of a quadratic number field k can be written as $\mathcal{O} = \mathbb{Z} + fw_D\mathbb{Z}$. An immediate consequence is that the set of orders $\mathbb{Z} + fw_D\mathbb{Z}$ is exactly equal to the set of all $\mathbb{Z}[\alpha]$ with $a + fw_D$, a, f integers and $f \neq 0$. Indeed,

$$\mathbb{Z}[\alpha] = \mathbb{Z} + \mathbb{Z}\alpha = \mathbb{Z} + \mathbb{Z}(a + fw_D) = \mathbb{Z} + fw_D\mathbb{Z}.$$

That is, every order is principal (generated by one element as a ring) inside $\mathcal{O}_D$.

Finally, we make a note on the discriminant Δ_D. We have just seen that every order $\mathcal{O} = \mathbb{Z} + fw_D\mathbb{Z}$ can be written as $\mathbb{Z} + \alpha\mathbb{Z}$. Defining the discriminant $\Delta(\mathcal{O})$ of $\mathcal{O} = \mathbb{Z} + \alpha\mathbb{Z}$ to be the discriminant of the minimal polynomial of α, we see that $\Delta(\mathcal{O}_D) = \Delta_D^2 - (\Delta_D^2 - \Delta_D) = \Delta_D$ and $\Delta(\mathbb{Z} + fw_D\mathbb{Z}) = f^2\Delta_D$. As the discriminant uniquely defined the field, the discriminant $f^2\Delta_D$ of an order uniquely defines the order $\mathbb{Z} + fw_D\mathbb{Z}$ and the field $\mathbb{Q}(\sqrt{D})$. Indeed, given the discriminant $n = f^2\Delta_D$ of an order (f, Δ_D unknown) we extract the squarefree part of n to obtain D. From this we compute $\Delta_D = D$ or $\Delta_D = 4D$ depending on whether $D \equiv 1 \pmod 4$ or not and set $f = \sqrt{n/\Delta_D}$.

As an aside, we note that recovering the squarefree (or square) component of a given integer is difficult, and conjectured to be as hard as factoring [62].

3.2 The ideal class-group of an order

Suppose D is a domain and F its field of fractions. A non-zero[10] D-module I of F is said to be a *fractional ideal* of D if there exists a non-zero d in D so that $dI \subseteq D$. One can think of d as a "universal denominator", killing all the denominators of the fractions of the elements in I. To emphasise the distinction between fractional ideals $I \subseteq F$ and "usual" ideals $J \subseteq D$, we call the ideals $J \subseteq D$ *integral* ideals. In both cases, we also use the short hand *I is a fractional D-ideal*, or *J is an integral D-ideal*. We immediately see the fractional ideals are exactly D-modules of the form αJ where α lies in F and $J \subseteq D$ is an integral ideal. Finally, we remark that when D is Noetherian, i.e. every ideal of D is finitely generated (like in any order of a number field), then the fractional ideals of D are exactly finitely generated D-modules contained in F.

The product of two fractional ideals is a fractional ideal again, so the fractional ideals form a monoid with multiplication. For any fractional ideal I we define $I^{-1} = \{\alpha \in F \mid \alpha I \subseteq D\}$. This is again a fractional ideal [37, Sec. 10.1], and we call I (and I^{-1}) *invertible* if $II^{-1} = I^{-1} = D$. Using this terminology, we must be a little careful: just because I^{-1} exists (it always does), does not mean that I is invertible. Clearly every *principal* fractional ideal $I = \alpha D$ for α in F is invertible.

From this, we define the *class-group* $\mathrm{Cl}(\mathcal{O})$ of $\mathcal{O}$ to be the quotient $\mathcal{I}(\mathcal{O})/\mathcal{P}(\mathcal{O})$ where $\mathcal{I}(\mathcal{O})$ denotes the invertible fractional ideals of $\mathcal{O}$ and $\mathcal{P}(\mathcal{O})$ the principal fractional ideals. Notably, every class in the class-group is represented by an integral ideal.

[10] Not all authors follow this convention. Jacobson, Cox, Milne, Sutherland and Neukirch do; whilst Atiyah, Macdonald do not.

Before we continue to classify the invertible ideals of D, we recall some more general features of integral extensions and closures. Let $R \subseteq E$ be rings. An element e of E is said to be *integral over* R if e satisfies a monic (univariate) polynomial with coefficients in R. Clearly every element in R is integral over R. If every element in E is integral over R, we say that E is an *integral extension* of R. The *integral closure* of R in E is the set of all elements in E that are integral over R. The integral closure forms a ring [35, Cor. 5.3]. Moreover, we say that R is *integrally closed* in E if the integral closure of R is R itself. Finally, the integral closure of R in E is integrally closed in E [35, Cor. 5.5]. Consequently, $\mathcal{O}_k$ is integrally closed.

Now, whether or not other ideals of D are invertible is governed in part by the following

Proposition 3.3 (Dedekind Domains [13, Rem. 3.25; 23, Def. 3.2])**.** *Let D be a Noetherian domain. The following are equivalent*

(i) D is integrally closed and non-zero prime ideals are maximal;
(ii) every non-zero properly contained ideal of D can be uniquely written as the product of prime ideals of D; and
(iii) every fractional ideal of D is invertible.

Any D satisfying any, and therefore all, of these properties is called a *Dedekind* domain. The integers are clearly Dedekind, or any PID for that fact. Although suggested by (ii), being UFD is not enough. Indeed, $k[x, y]$ is UFD for any field k, and the ideal $xk[x, y]$ generated by x is prime yet not maximal.

We suspect $\mathcal{O}_k$ is a Dedekind domain, for it is already Noetherian and integrally closed. The only thing left to prove is that prime ideals are maximal. This is very well described by the content of the *Cohen-Seidenberg* theorems: let E be an integral extension of R, then the *restriction* $P \cap R$ of a prime ideal $P \subseteq E$ is maximal in R if and only if P is maximal in E [35, Cor. 5.8]. Applying this to $\mathbb{Z} \subseteq \mathcal{O}_k$, we see that every prime ideal of $\mathcal{O}_k$ is maximal. Indeed, the restriction $P \cap \mathbb{Z}$ of a prime non-zero ideal P of $\mathcal{O}_k$ is a prime ideal of $\mathbb{Z}$ and so also a maximal. We conclude that the ring of integers of any number field is a Dedekind ring. As such, all fractional ideals of $\mathcal{O}_k$ are invertible.

Conversely, we see that there are non-invertible fractional ideals of every non-maximal order in a quadratic number field $k = \mathbb{Q}(\sqrt{D})$ because non-maximal orders $\mathcal{O} = \mathbb{Z} + fw_D\mathbb{Z}$ are not integrally closed. Indeed, because w_D is integral over $\mathbb{Z}$ (it even *generates* the integral closure of $\mathbb{Z}$ in k), it is integral over $\mathcal{O}$, but it is patently not contained in $\mathcal{O}$.

Sutherland gives us an explicit example of a non-invertible ideal in the following. Let $D = -1$, then $w_D = 2+\sqrt{-1} = 2+\iota$ and $\mathcal{O}_D = \mathbb{Z}+w_D\mathbb{Z} = \mathbb{Z}+\iota\mathbb{Z}$ comprises the *Gaussian integers*. Now $\mathcal{O} = \mathbb{Z} + 2w_D\mathbb{Z} = \mathbb{Z} + 2\iota\mathbb{Z}$ is an order in $\mathcal{O}_D$ and then $I = 2\mathbb{Z} + 2\iota\mathbb{Z}$ an integral ideal thereof. Writing down $I^{-1} = \{\alpha \in \mathrm{Frac}(\mathcal{O}) = \mathbb{Q} + \iota\mathbb{Q} \mid \alpha I \subseteq I\} = \mathbb{Z} + \iota\mathbb{Z}$, we see that $II^{-1} = I \neq \mathcal{O}$ and conclude that I is not invertible. We did not even need to construct a "difficult" fractional ideal to arrive at a non-example.

To get a better understanding of which ideals are invertible in a non-maximal order we have introduce to some more ingredients that end up being especially useful for imaginary quadratic fields, which is our case of interest.

First, we generalise the norm of *elements* to norms of *ideals* in an *imaginary* (a condition that we have not used yet) quadratic number fields. Let $\alpha = a + bfw_D$ lie in the order $\mathcal{O} = \mathbb{Z} + fw_D\mathbb{Z}$. Then

$$\begin{aligned}\alpha\mathcal{O} &= (a + bfw_D)\mathbb{Z} + (afw_D + bf^2w_D^2)\mathbb{Z} \\ &= (a + bfw_D)\mathbb{Z} + (afw_D + bf^2(\Delta_D w_D - (\Delta_D^2 - \Delta_D)/4))\mathbb{Z} \\ &= (a + bfw_D)\mathbb{Z} + (-bf^2(\Delta_D^2 - \Delta_D)/4 + (a + bf\Delta_D)fw_D)\mathbb{Z}).\end{aligned}$$

Now, for any free module $F = e_1\mathbb{Z}+e_2\mathbb{Z}$, the submodule $G = (ae_1+be_2)\mathbb{Z}+(ce_1+de_2)\mathbb{Z}$ with $ad - bc \neq 0$ has index $|ad - bc|$ in F [21, Ex. 7.15]. That is, $|F/G| = |ad - bc|$. Therefore, using $e_1 = 1, e_2 = fw_D$, we obtain

$$\left|\mathcal{O}/\alpha\mathcal{O}\right| = a(a + bf\Delta_D) + b^2f^2(\Delta_D^2 - \Delta_D)/4.$$

This is exactly the norm of α. Indeed, writing $\overline{\alpha} = a + bf\overline{w_D} = a + bf(\Delta_D - w_D) = a + bf\Delta_D - bfw_D$ we compute

$$\begin{aligned}\mathrm{N}(\alpha) &= (a + bfw_D)(a + bf\Delta_D - bfw_D) \\ &= a(a + bf\Delta_D) - abfw_D + bfw_D(a + bf\Delta_D) - b^2f^2(\Delta_D w_D - (\Delta_D^2 - \Delta_D)/4) \\ &= a(a + bf\Delta_D) + (\Delta_D^2 - \Delta_D)/4).\end{aligned}$$

This inspires us to define the *norm* $\mathrm{N}(I)$ of an integral ideal $I \subseteq \mathcal{O}$ to be its index $[\mathcal{O} : I] = |\mathcal{O}/I|$. To ensure this is finite, we first remark that every ideal $I \subseteq \mathcal{O}$ contains an integer. Indeed, if α lies in $\mathcal{O}$, then by description of orders (Theorem 3.1), and using that the ring of integers $\mathcal{O}_D$ is stable under conjugation, we see that $\overline{\alpha}$ lies in $\mathcal{O}$. Hence $\mathrm{N}(\alpha) = \alpha\overline{\alpha}$ also lies in I. This is an integer because α is integral over $\mathbb{Z}$ (α lies in $I \subseteq \mathcal{O} \subseteq \mathcal{O}_D$). Now $\mathrm{N}(\alpha)\mathcal{O} \subseteq I \subseteq \mathcal{O}$ and so $\mathcal{O}/I \subseteq \mathcal{O}/\mathrm{N}(\alpha)\mathcal{O}$. However, since $\mathcal{O}$ is an order, and so a finitely generated free $\mathbb{Z}$-module, the quotient $\mathcal{O}/m\mathcal{O}$ is finite. This verifies that the index of any ideal of $\mathcal{O}$ is finite.

We now generalise this definition to fractional ideals. Let I be a fractional ideal of an order $\mathcal{O} \subseteq \mathcal{O}_k$, so that $I = \alpha J$ for α in k and $J \subseteq \mathcal{O}$ an integral ideal. We write $\alpha = (a + bfw_D)/(c + dfw_D)$ with integers a, b, c, d and notice that $\alpha = (a + bfw_D)\overline{(c + dfw_D)})/\mathrm{N}(c+dfw_D)$. Since $J' = (a+bfw_D)\overline{(c + dfw_D)})J$ is an integral ideal again, we can assume that α lies in $\mathbb{Q}$. In fact we can go further to assume that $\alpha = 1/n$ for some non-zero integer n. We call $J \subseteq \mathcal{O}$ *primitive* if J/m is not integral for any integer $m > 1$. We now define the *norm* of $I = \alpha J$ to be $\mathrm{N}(I) = \mathrm{N}(\alpha)\mathrm{N}(J)$ if J is a primitive $\mathcal{O}$-ideal and α in $\mathbb{Q}$.

We note that $(\alpha\mathcal{O})(\overline{\alpha}\mathcal{O}) = \mathrm{N}(\alpha)\mathcal{O}$ and see that this generalises to non-principal ideals

Lemma 3.4 ([21, Lem. 7.14])**.** *Let $\mathcal{O}$ be an order in an imaginary quadratic field and I, J invertible ideals of $\mathcal{O}$. Then $\mathrm{N}(IJ) = \mathrm{N}(I)\mathrm{N}(J)$; and $I\,\overline{I} = \mathrm{N}(I)\mathcal{O}$.*

This lemma is very powerful, because it immediately tells us how to compute the inverse of an ideal quickly. If I is invertible, then $I\bar{I}/\mathrm{N}(I) = \mathcal{O}$. By uniqueness of inverses in monoids, we conclude that $I^{-1} = \bar{I}/\mathrm{N}(I)$. This can be used in the converse, if $I\bar{I}/\mathrm{N}(I) \neq \mathcal{O}$, then I is not invertible. We also see, that in our previous example with Gaussian integers, our computation $(2\mathbb{Z} + 2\iota\mathbb{Z})^{-1} = \mathbb{Z} + \iota\mathbb{Z}$ was correct.

Sutherland gives us a counterexample to the multiplicativity of norms for non-invertible ideals. Let us return to the setting of Gaussian integers. Here $D = -1$, $\Delta_D = -4$, $w_D = 2 + \sqrt{-1} = 2 + \iota$. Let $\mathcal{O} = \mathbb{Z} + 2w_D\mathbb{Z} = \mathbb{Z} + 2\iota\mathbb{Z}$. The ideal $I = 2\mathbb{Z} + 2\iota\mathbb{Z}$ of $\mathcal{O}$ has norm 2. On the other hand, we compute $I^2 = (2(\mathbb{Z} + \iota\mathbb{Z}))^2 = 4(\mathbb{Z} + \iota\mathbb{Z})^2 = 4\mathbb{Z} + 4\iota\mathbb{Z}$, which however has norm 8 and not 4. By using the formula $I^{-1} = J = \bar{I}/\mathrm{N}(I)$ we compute $I^{-1} = \mathbb{Z} + \iota\mathbb{Z}$ to verify that I is indeed not invertible, $II^{-1} = 2(\mathbb{Z} + \iota\mathbb{Z})^2 = 2\mathcal{O} \neq \mathcal{O}$.

We also have the following, rather immediate, corollary of 3.4. Namely, that given an invertible integral ideal $I \subseteq \mathcal{O}$, the set of all equivalent integral ideals $J \subseteq \mathcal{O}$ is given by

$$\left\{\bar{\beta}I/\mathrm{N}(I) \mid \beta \in I\right\}.$$

Indeed, any two fractional ideals $\mathcal{I}, \mathcal{J} \subseteq \mathrm{Frac}(\mathcal{O})$ are equivalent if and only if there exists a γ in $\mathrm{Frac}(\mathcal{O})$ such that $\gamma\mathcal{I} = \mathcal{J}$. Moreover, with γ in $\mathrm{Frac}(\mathcal{O})$, the ideal $\gamma\mathcal{I}$ is integral if and only if γ lies in $\mathcal{I}^{-1}$. Consequently, an integral ideal J is equivalent to I if and only if $J = \gamma I$ with γ in I^{-1}. Since I is invertible, $I^{-1} = \bar{I}/\mathrm{N}(I)$ (by the lemma) and we conclude that the integral ideals J equivalent to I are exactly the ideals of the form $\bar{\beta}I/\mathrm{N}(I)$ with β in I. These ideals have norm $\mathrm{N}(\beta)/\mathrm{N}(I)$.

Writing $\beta = xma + ym(b + fw_D)$ in $I = m[a, b + fw_D]$ for integers x, y, we see that

$$\begin{aligned}\frac{\mathrm{N}(\beta)}{\mathrm{N}(I)} &= \frac{1}{m^2a}m^2\left(x^2a^2 + xya(2b + f\Delta_D) + y^2\mathrm{N}(b + fw_D)\right)\\ &= x^2a + xy(2b + f\Delta_D) + y^2\mathrm{N}(b + fw_D)/a.\end{aligned}$$

We remark that, since I is an ideal a divides $\mathrm{N}(b + fw_D)$ and so $\mathrm{N}(\beta)/\mathrm{N}(I)$ is a binary quadratic form as a function of x, y. Moreover, it has discriminant $(2b + f\Delta_D)^2 - 4\mathrm{N}(b + fw_D) = f^2\Delta_D < 0$ and so is positive definite (since we assume $a \geq 0$ when writing ideals in the standard form $m[a, b + fw_D]$).

From this discussion we conclude the

Corollary 3.5. *Let $I \subseteq \mathcal{O}$ be an invertible integral ideal. The set of all equivalent integral ideals $J \subseteq \mathcal{O}$ is given by $\{\bar{\beta}I/\mathrm{N}(I) \mid \beta \in I\}$. Moreover, the ideal $I = m[a, b + fw_D]$ has an equivalent integral ideal J of norm n if and only if*

$$x^2a + xy(2b + f\Delta_D) + y^2\mathrm{N}(b + fw_D)/a = n$$

has integer solutions x, y.

We now have developed the language to formulate the

Lemma 3.6 ([21, Lem. 7.18])**.** *Let $\mathcal{O}$ be an order of a quadratic number field, I an integral thereof and l an integer. Then*

(i) I is prime to $l\mathcal{O}$ if and only if $\mathrm{N}(I)$ is prime to $\mathrm{N}(l\mathcal{O}) = l^2$; and
(ii) I is invertible if $\mathrm{N}(I)$ is coprime to the conductor f of $\mathcal{O}$.

Proof. We recall that two ideals I, J of a ring R are coprime if $I + J = R$.

Since I is an ideal, the multiplication-by-l map $m_l \colon \mathcal{O} \to \mathcal{O}; x \mapsto mx$ induces a well-defined map $\overline{m_l} \colon \mathcal{O}/I \to \mathcal{O}/I$ on the quotients. Now the condition $I + l\mathcal{O} = \mathcal{O}$ means nothing more than $l\mathcal{O} = \mathcal{O} \pmod{I}$, which in turn is equivalent to $\overline{m_l}$ being surjective. Since $\mathcal{O}/I$ is finite (pg. 40, definition of norm), this is equivalent to $\overline{m_l}$ being bijective. Moreover, $\overline{m_l}$ is a morphism of groups, so in fact this is equivalent to $\overline{m_l}$ being an automorphism of groups.

Since $\mathcal{O}/I$ is a finite Abelian group, we can write $\mathcal{O}/I$ as a product of cyclic groups $\mathbb{Z}/n_1\mathbb{Z} \times \cdots \mathbb{Z}/n_r\mathbb{Z}$. Then $\overline{m_l}$ is an automorphism on $\mathcal{O}/I$ if and only if the induced maps on each $\mathbb{Z}/n_i\mathbb{Z}$ are isomorphisms for all n_i, which this is the case if and only if l is prime to all n_i. This is true if and only if l is prime to $|\mathcal{O}/I| = \mathrm{N}(I)$. This proves (i).

To prove (ii), let $k = \mathbb{Q}(\sqrt{D})$ with squarefree D be the quadratic number field of which $\mathcal{O} \subseteq \mathcal{O}_D$ is an order. We note that the set $\mathcal{O}_I = \{x \in k \mid xI \subseteq I\}$ is contained in $\mathcal{O}_D$. Some authors call $\mathcal{O}_I$ the *order of* I [1, Def. 42]. Now let x lie in $\mathcal{O}_I$. Since $\mathrm{N}(I)$ is prime to l, we we know by part (i) that $\mathcal{O} = I + f\mathcal{O}$ and we have

$$x\mathcal{O} = x(I + f\mathcal{O}) = xI + xf\mathcal{O} \subseteq I + f\mathcal{O}_D \subseteq I + \mathcal{O} \subseteq \mathcal{O}.$$

This demonstrates that $\mathcal{O}_I \subseteq \mathcal{O}$. The converse $\mathcal{O} \subseteq \mathcal{O}_I = \{x \in k \mid xI \subseteq I\}$ is always true because I is an ideal.

Ideals for which $\mathcal{O}_I = \mathcal{O}$ are called *proper* and, in quadratic imaginary fields, they are exactly the invertible ideals [21, Prop. 7.4], completing the proof of (ii). More generally, the proper ideals of an order of any (not necessarily quadratic) number field, form a group under multiplication [80, Prop. 4.11]. □

Unfortunately (ii) of this lemma cannot be improved. That is, there exist invertible ideals whose norm are not prime to the conductor. We construct an example in the following. Let $D = -1, \Delta_D = -4, w_D = 2 + \sqrt{-1} = 2 + \iota$. Let $\mathcal{O} = \mathbb{Z} + 6w_D\mathbb{Z} = \mathbb{Z} + 6\iota\mathbb{Z}$ and $I = 2\mathbb{Z} + 6\iota\mathbb{Z}$. Verify that I is an ideal by computing $(2\alpha + 6\beta\iota)(\gamma + 6\iota\delta) = (2\alpha\gamma - 36\beta\delta) + 6(\beta\gamma + 2\alpha\delta)\iota$. Then I has norm 2 in $\mathcal{O}$. We compute $I^{-1} = \bar{I}/2 = \mathbb{Z} + 3\iota\mathbb{Z}$ and verify $II^{-1} = (2\mathbb{Z} + 6\iota\mathbb{Z})(\mathbb{Z} + 3\iota\mathbb{Z}) = \cdots = \mathbb{Z} + 6\iota\mathbb{Z} = \mathcal{O}$. Hence I is invertible, even though its norm is not prime to the conductor.

However, there is a fix to this imprecision in the following. We first note that the integral ideals prime to $f\mathcal{O}$ are invertible, and so lie in $\mathcal{I}(\mathcal{O})$. Moreover, the product of integral ideals prime to $f\mathcal{O}$ are again prime to $f\mathcal{O}$ because the product of their norms is again coprime to f. So we can define $\mathcal{I}(\mathcal{O}, f)$ to be the subgroup inside $\mathcal{I}(\mathcal{O})$ generated by integral ideals that are prime to $f\mathcal{O}$. We can do the same for principal ideals and denote this subgroup $\mathcal{P}(\mathcal{O}, f)$ and it turns out [21, Prop. 7.19], that

$$\mathcal{I}(\mathcal{O}, f)/\mathcal{P}(\mathcal{O}, f) \cong \mathrm{Cl}(\mathcal{O})$$

3.3 Explicit representation of ideals

To perform the CLAPOTI algorithm outlined in Subsection 2.7, we must be able to compute with ideals. We have seen that orders are free $\mathbb{Z}$-modules; to use a more compact notation, we will write $[a_1, \dots, a_n]$ for the $\mathbb{Z}$-module $a_1 e_1 \mathbb{Z} + \cdots a_2 e_n \mathbb{Z}$. More concretely, in this section $[z, w] = z\mathbb{Z} + w\mathbb{Z}$.

Proposition 3.7 ([24, Prop. 3.7; 36, Sec. 5.2]). *Let $\mathcal{O} = \mathbb{Z} + f\mathbb{Z}w_D$ be an order inside a quadratic number field. Then*

(i) every subgroup $M \subseteq \mathcal{O}$ can be written as $M = [n, c + mfw_D]$ with unique integers $0 \leq n, m$ and $0 \leq c \lneq n$; and

(ii) the subgroup $M = [n, c + mfw_D]$ is an ideal if and only if

$$m \mid n, \quad m \mid c, \quad \text{and} \quad n \mid m\mathrm{N}(c/m + fw_D);$$

consequently, we may write an ideal as $m[a, b + fw_D]$ with unique integers $0 \leq m, a$ and $0 \leq b \lneq a$.

Proof. To prove (i), let $H = \{\zeta \mid \mu + \zeta f w_D \in M\}$. This is a subgroup of $\mathbb{Z}$ and so can be written as $H = m\mathbb{Z}$ for some unique non-negative integer m. Moreover, $M \cap \mathbb{Z} \subseteq \mathbb{Z}$ is also a subgroup and so $M \cap \mathbb{Z} = n\mathbb{Z}$ for some unique non-negative integer n. Now, since m lies in H, there exists a c in $\mathbb{Z}$ so that $c + mfw_D$ lies in M. If $c + mfw_D$ lies in M, then so does $c + \lambda n + mfw_D$. As such, we may choose $0 \leq c \lneq n$ uniquely.

Now, let $x = a + bfw_D$ lie in M. Then b lies in H and is $b = \lambda m$ for some integer λ. Then $x - \lambda(c + mfw_D) = a + bfw_D - \lambda c - \lambda mfw_D = a - \lambda c$ lies in $\mathbb{Z} \cap M$. As such $a - \lambda c = \mu n$ for some integer μ. In conclusion $x = \mu n + \lambda(c + mfw_D)$ is in $[n, c + mfw_D]$. We note that n, m and c did not depend on x, and so we conclude $M \subseteq [n, c + mfw_D]$. The converse $[n, c + mfw_D] \subseteq M$ is clear.

Now let us prove (ii). Since I is a group by definition, it is an ideal if and only if it is closed by multiplication from $\mathcal{O}$. Since $\mathcal{O} = [1, fw_D]$, we must verify that $1 \cdot I \subseteq I$ and $fw_D \cdot I \subseteq I$ (i.e. verify on generators). The first condition is trivially true. For the second, we compute $fw_D \cdot n$ and $fw_D \cdot (c + mfw_D)$ and find exact conditions that ensure they lie in I.

First, nfw_D lies in I if and only if $nfw_D = \lambda n + \mu(c + mfw_D)$ for integers μ, λ. That is, $n = \lambda\mu$ and $\lambda n + \mu c = 0$. Writing this more compactly, we see nfw_D lies in I if and only if $m \mid n$ and $m \mid c$. So let us write $I = m[a, b + fw_D]$ where $n = ma, c = mb$ from now on.

The second generator

$$\begin{aligned} fw_D \cdot (c + mfw_D) &= fw_D \cdot (ma + mfw_D) \\ &= m(bfw_D + f^2 w_D^2) \\ &= m(bfw_D + f^2(\Delta_D w_D - (\Delta_D^2 - \Delta_D)/4)) \\ &= m(-f^2(\Delta_D^2 - \Delta_D)/4 + (b + f\Delta_D)fw_D) \end{aligned}$$

lies in I if and only if there exist integers λ, μ so that

$$\lambda ma + \mu mb = -mf^2(\Delta_D^2 - \Delta_D)/4 \qquad \text{and} \qquad \mu m = m(b + f\Delta_D).$$

Substituting μm, we obtain $\lambda ma + mb(a + f\Delta_D) = -mf^2(\Delta_D^2 - \Delta_D)/4$ which is equivalent to

$$-\lambda a = b(b + f\Delta_D) + f^2(\Delta_D^2 - \Delta_D)/4 = b^2 + bf\Delta_D + f^2(\Delta_D^2 - \Delta_D)/4 = \mathrm{N}(b + fw_D).$$

In other words, $fw_D \cdot (c + mfw_D)$ lies in I if and only if $a \mid \mathrm{N}(b + fw_D)$; or equivalently $n \mid m\mathrm{N}(c/m + fw_D)$.

The final statement, that b can now be chosen uniquely $0 \leq b \not\leq a$ follows immediately. □

We can now compute the norms of ideals very easily.

Proposition 3.8. *The index of $M = [n, c + mfw_D]$ in $\mathcal{O} = \mathbb{Z} + f\mathbb{Z}w_D$ is $|\mathcal{O}/M| = nm$.*

Consequently, if $I = m[a, b + fw_D]$ is an ideal, then $\mathrm{N}(I) = |\mathcal{O}/I| = m^2 a$.

Proof. We prove this by showing that every element $x = a + bfw_D$ in $\mathcal{O}/M$ can be represented by $y = a' + b'fw_D$ where $0 \leq ab' \lneq n$ and $0 \leq b' \lneq m$. Let $b = \lambda m + b'$ so that $0 \leq b' \lneq m$. Then $x - \lambda(c + mfw_D) = a - \lambda c + b'fw_D$. Now choosing $a - \lambda c = \mu n + a'$ so that $0 \leq a' \lneq n$ we have $x - \lambda(c + mfw_D) - \mu n = a' + b'fw_D$ concluding the proof. □

Lemma 3.9. *There exists an invertible ideal I in $\mathcal{O} = [1, fw_D]$ with odd norm n prime to $f^2\Delta_D$ if and only if $f^2\Delta_D$ is a square modulo $4n$.*

Proof. We cite the well-known result, that the map $ax^2 + bxy + cy^2 \mapsto [a, -(b + \Delta_D)/2 + fw_D]$ induces a bijection between the form class-group of discriminant $f^2\Delta_D$ and the ideal class-group of $\mathcal{O}$ (e.g. [21, Th. 7.7]). We note that $[a, -(b + f\Delta_D)/2 + fw_D]$ is in fact an integral ideal: since $b^2 - 4ac = f^2\Delta_D$ we see that $b^2 \equiv f^2\Delta_D \pmod 4$ and so $b \equiv f\Delta_D \pmod 2$ because $\Delta_D \equiv 0, 1 \pmod 4$. Moreover the norm of this ideal is a.

So there exists an ideal in $\mathcal{O}$ with norm a if and only if there exists a binary quadratic form $Q = ax^2 + bxy + cy^2$ with discriminant $f^2\Delta_D$. Clearly Q represents a. It is now a standard result of binary quadratic forms, that odd n are represented by a form with discriminant $f^2\Delta_D$ if and only if $f^2\Delta_D$ is a square modulo $4n$ (e.g. [21, Lemma 2.5]).

Conversely, if there exists a form Q that represents a, then there exist b, c so that $ax^2 + bxy + cy^2$ is equivalent to Q. □

4
Some experiments: Clapoti for ordinary elliptic curves

4.1 Clapoti with 2-dimensional 2-isogenies

Since [3] is currently the only available library for computing generic higher-dimensional isogenies we will investigate computing the isogeny class-group action on ordinary curves in two-dimensions. Moreover, using this library means that the induced Kani map K must be a 2^n-isogeny and so the endomorphisms α, β of the general CLAPOTI algorithm must be given by 1×1 matrices (i.e. scalars), or just endomorphisms of the curve. Finally, the library requires the 2^{n+2}-torsion to be accessible on the starting curve.

Using the *CM method*, we can efficiently construct ordinary elliptic curves E with sufficient torsion; however the class-groups of the endomorphism rings $\mathrm{End}(E)$ of curves found this way will be very small. Conversely, picking random Weierstrass invariants a, b in F_q and instantiating an elliptic curve from them will produce an ordinary elliptic curve with large class-group $\mathrm{Cl}(\mathrm{End}(E))$ with overwhelming probability, but we have no control over the rational torsion. In essence, with current knowledge, there is a tension between finding ordinary elliptic curves with targeted rational torsion (e.g. 2^n-torsion) *and* curves with large class-groups $\mathrm{Cl}(\mathrm{End}(E))$. We will circumvent this problem in a way that is not cryptographically secure, but allows us to compute examples.

In either case, fortunately, the group structure of $E(F_q)$ is invariant under horizontal isogenies. In a formula $E(F_q) \cong \mathcal{O}/(\pi - 1)\mathcal{O}$ where $\mathcal{O} \cong \mathrm{End}(E)$. Consequently, we must only construct one ordinary curve E with satisfactory torsion and class-group $\mathrm{Cl}(\mathrm{End}(E))$; and then all curves in $\mathrm{Ell}_q(\mathrm{End}(E))$ will also have sufficient torsion.

We now recall the algorithm presented in 2.7 and adapt it to the two-dimensional case. We assume that $\mathrm{Ell}_q(\mathcal{O})$ comprises (isomorphism classes of) elliptic curves with sufficient 2^n-torsion. What "sufficient" means in practice is discussed in the section on results. We also assume that $\mathrm{Cl}(\mathcal{O})$ is not the trivial group, so that the resulting action is non-trivial.

Now, given an ideal class $[H]$, we must find other representatives $[I] = [J] = [H]$ for which there exist endomorphisms α, β of E so that the *degree equation*

$$\deg(\alpha)\mathrm{N}(I) + \deg(\beta)\mathrm{N}(J) = 2^m \qquad \text{with} \qquad \gcd(\deg(\alpha)\mathrm{N}(I), \deg(\beta)\mathrm{N}(J)) = 1$$

is satisfied for some $m \leq n$. From here, we can compute the endomorphism μ such that $I\bar{J} = \mu\mathcal{O}$ and compute the kernel

$$\ker(K) = \{(\deg(\alpha)\mathrm{N}(I)x, \mu\alpha\beta(x) \mid x \in E[2^n]\}$$

of the Kani map $K\colon E \times E \to E_I \times E_{\bar{I}}$, a 2^n-isogeny in 2-dimensions. This kernel can then be passed to the library of [3] to recover equations for $E_I, E_{\bar{I}}$ which we distinguish using the Weil pairing trick.

We explain different ways in which one can find the matching endomorphisms α, β.

We briefly recall that the norm of any endomorphism $\alpha = x+yfw_D$ in $\mathcal{O} = [1, fw_D] \subseteq \mathcal{O}_D$ is given by

$$\deg(\alpha) = x^2 + xyf\Delta_D + y^2f^2(\Delta_D^2 - \Delta_D)/4 \tag{4.1}$$

This inspires us to investigate two cases: $y = 0$ and $y \neq 0$.

When $y = 0$, the degrees $\deg(\alpha), \deg(\beta)$ achieved are simply squares. The degree equation $a^2\mathrm{N}(I) + b^2\mathrm{N}(J) = 2^m$ has solutions if and only if the binary quadratic form with coefficients $(\mathrm{N}(I), 0, \mathrm{N}(J))$ represents 2^m. Solutions to this can be found with Cornacchia's algorithm [78]. This suggests the following "reverse" approach.

We propose the following algorithm to find pairs I, J of equivalent ideals with coprime norm inside an order $\mathcal{O} \subseteq \mathcal{O}_D$ of conductor f, so that there exist a, b with $a^2\mathrm{N}(I) + b^2\mathrm{N}(J) = 2^n$.

(i) Using Lemma 3.9 we write down the set of possible norms of ideals representing any class in $\mathrm{Cl}(\mathcal{O})$. This set is unbounded, but by Minkowski's bound, we know that every class is represented by some ideal of norm at most $\sqrt{-\Delta/3}$ (where Δ is the discriminant of $\mathcal{O}$). We consider the subset of possible norms bounded by some (small) multiple of the Minkowski bound and call it B.

(ii) We enumerate pairs of coprime norms $n_I \leq n_J$ in B, and find solutions a, b to $\mathrm{N}(I)x^2 + \mathrm{N}(J)y^2 = 2^n$ if they exist.

(iii) We now iterate over ideals I with norm n_I, and test whether there exists an ideal J that is *equivalent* to I with norm n_J. Then we have found an endomorphism pair $\alpha = a, \beta = b$ for the class represented by I. This can be verified using Corollary 3.5.

(iv) Repeating this process until all classes in $\mathcal{O}$ have been exhausted gives us a very crude, yet concrete method to find matching (scalar) endomorphisms for every class in the class group. Clearly this can only apply to small toy examples.

When we allow $y \neq 0$, the degrees $\deg(\alpha), \deg(\beta)$ are simply values given by Equation 4.1. Here we present the approach detailed in Clapoti. Given a class $[H]$ in $\mathrm{Cl}(\mathcal{O})$ we perform the following steps

(i) Sample random elements γ, δ in H to form the equivalent ideals $I = \overline{\gamma}H/\mathrm{N}(H), J = \overline{\delta}H/\mathrm{N}(H)$. We do not use any advanced sampling technique like in [18, Lem. 4.2.1], we just choose γ, δ from a ball with small enough radius.

(ii) If $\gcd(\mathrm{N}(J), \mathrm{N}(I)) \neq 1$, re-sample the ideals

(iii) Find Bèzout coefficients a, b so that $a\mathrm{N}(I)+b\mathrm{N}(J) = 2^n$ and $\gcd(a\mathrm{N}(I), b\mathrm{N}(J)) = 1$.

(iv) Use Cornacchia's algorithm to find endomorphisms α, β in $\mathcal{O}$ with $\deg(\alpha) = a, \deg(\beta)$ using Equation 4.1, if they exist.

To use the library to compute a 2^n-isogeny $f = (f_1, f_2; f_3, f_4)\colon E_1 \times E_1 \to E_1' \times E_2'$, we must supply the equations of E_1, E_2 over F_q and generators P_1, P_2 on E_1 and Q_1, Q_2 on E_2 of the kernel G of f. From this, it is able to compute the codomain curves E_1', E_2', and returns a function which one can evaluate on arbitrary points of $E_1 \times E_2$. The *worked example* of the `README` [4] is very clear. We note that the implementation requires knowledge of $E_1[2^{n+2}], E[2^{n+2}]$, and so must to compute the additional field extensions

if these points are not already rational. We ensure the rationality requirement is met by construction.

4.2 Rational points and field extensions

We have seen that to embed an isogeny into higher dimensions or to compute the isogeny class-group action, we must have access to some torsion points. More precisely, when evaluating the induced minimal Kani map $A \times * \to * \times *$ (an N-isogeny) we computed its kernel from the points in the N-torsion $A[N]$ on A.

However, although the multiplication-by-N maps are rational, their kernels must not be. A simple counterexample is when $C = \left|A(F_q)\right|$, then $A[N]$ cannot be rational ($A[N] \subseteq A(F_q)$) for any $N \geq C$. Since the N-torsion on any Abelian variety A is finitely generated (in fact, we know it is generated by $2\dim(A)$ elements), we require at most a finite extension of F_q to ensure $A[N] \subseteq A(F_q{}^d)$.

The exact degree of extension required to contain the N-torsion points has been studied carefully in the context of point counting. The SEA point counting algorithm [60, Sec. 5] in particular uses *N-division polynomials* to compute the order of k-rational points, a technique we now describe in broad strokes.

We recall that $\left|E(F_q)\right| = q + 1 - \mathrm{tr}(\pi_q)$ [28, V.2 Th. 2.3.1], so computing $\left|E(F_q)\right|$ is equivalent to computing the trace of Frobenius. The idea of SEA is to compute the trace of π_q when restricted to small torsion subgroups $E[l_i]$ for primes $l_1, \dots, l_n$ until $L = l_1 \cdots l_n$ is large enough (that is, $L > 4\sqrt{q}$) and then to piece this information together to arrive at the full trace. The key insight is that π_q restricted to $E[l]$ is a F_l-linear map on a vector space $E[l] \cong F_l^2$ satisfying $\pi_q^2 - t'\pi_q + q = 0$ (c.f. Sec. 3.1) for some $0 \leq t' < l$, and that $t' \equiv \mathrm{tr}(\pi_q) \pmod{l}$. The l-division polynomial ψ_l is the univariate polynomial whose roots are the x-coordinates of points in $E[l]$. Hence evaluating π_q on $E[l]$ to find t' can be done by noting $\pi_q^2 - t'\pi_q + q = 0$ on $E[l]$ if and only if

$$(x^{q^2}, y^{q^2}) -_E t'(x^q, y^q) +_E q'(x, y) \equiv 0 \pmod{\psi_l, y^2 - (x^3 + ax + b)}$$

where $ZY^2 = X^3 + aXZ^2 + bZ^3$ is the defining equation of E, and we have used affine coordinates in the equation and $0 \leq q' < l$ is the representative of q modulo l. It is important that the addition, subtraction $-_E, +_E$ and scalar multiplication occur in E, that is, defined by the polynomial group laws of E.

Since the l-division polynomial ψ_l has roots exactly on the x-coordinates on the non-zero points in $E[l]$, (i) $\deg(\psi_l) = (l^2 - 1)/2$ when l is prime to q, (ii) $\deg(\psi_l) = (l - 1)/2$ when $l = q$ and E ordinary, (iii) $\deg(\psi_l) = 0$ when $l = q$ and E supersingular [28, V.3 Th. 3.1]. The factor of half arises, because if (x, y) has order l then so does $(x, -y)$. Hence, *a priori*, the degree d of the extension required to make the l-torsion rational is the gcd of the degrees of the irreducible factors. Alas, this is not the case, and the degree of the extension is at most $2\deg(\psi_l)$ [77, Cor. 3.3].

This answers the question for primes l. For a product of primes, we obtain an extension of degree at most the product of the degrees. This gives us the notion of *accessible* torsion, we say that the N-torsion is *accessible* if N is *B-powersmooth*, that

is with $N = p_1^{e_1} \cdots p_n^{e_n}$ we have $p_i^{e_i} \leq B$. This is used by Robert in the proof that CLAPOTI is polynomial time. In practice however, we will simply construct curves so that our desired N-torsion is in fact rational.

4.3 The CM method for constructing ordinary elliptic curves

From our discussion in the previous subsection, we see the need for constructing elliptic curves with prescribed rational N-torsion. The order $C = \left|E(F_q)\right|$ of a curve with rational N-torsion must be a multiple of N^2. We can find curves with a prescribed number of F_q-rational points using the folklore *CM method* outlined below [7, Sec. 22.4]. Then, given such a curve, we can test whether $E[N]$ is rational by computing $E(F_q) \cong \mathbb{Z}/m_1\mathbb{Z} \times \mathbb{Z}/m_2\mathbb{Z}$. When $E(F_q)$ is very smooth (which we will enforce by design $N = 2^n$), this can be done efficiently in Sage.

The current implementation must factor the order of $E(F_q)$ (which is easy, if it is smooth); perform computations of cost $\Theta(\sqrt{\ell})$ for every prime ℓ which has rational ℓ^∞ torsion [87]; and finally performs computations of polylogarithmic cost in q.

Let E/F_q be an ordinary elliptic curve with Frobenius endomorphism π and endomorphism ring $\mathcal{O} \subseteq \mathcal{O}_D$ of conductor f and discriminant $\Delta = f^2\Delta_D$. Since π is an endomorphism of E, it is an element of $\mathcal{O}_D$ and so satisfies a monic polynomial with integer coefficients $\pi^2 - \mathrm{tr}(\pi)\pi + \mathrm{N}(\pi) = 0$ (c.f. Sec. 3.1). It is well-known [7, Lec. 7; 28, V.2 Th. 2.3.1] that

$$\mathrm{N}(\pi) = q \qquad \text{and} \qquad \mathrm{tr}(\pi) = q + 1 - C \quad \text{where} \quad C = \left|E(F_q)\right|.$$

We will denote $\mathrm{tr}(\pi)$ by t in the following.

We write $\pi = (t + v\sqrt{\Delta})/2$ in $\mathcal{O}$ for some integer v and conclude from $\mathrm{N}(\pi) = q$ that $4q = t^2 - v^2\Delta$. This equality is known as the *norm equation* (of Frobenius). We rewrite this as a function of C by substituting $t = q + 1 - C$ and obtain

$$\begin{aligned} 0 &= (q+1)^2 - (2C+4)(q+1) + C^2 - v^2\Delta + 4 \\ &= q^2 - 2(C+1)q + (C-1)^2 - v^2\Delta \end{aligned} \tag{4.2}$$

to conclude

$$q = C + 1 \pm \sqrt{4C + v^2\Delta}.$$

The celebrated theory of *complex multiplication* together with Deuring's *reduction theorem* gives us a converse result [1, Th. 72, Th. 74]. Given integers q, C, v, Δ satisfying (4.2) (where q is a prime power, and Δ a discriminant) there exists an elliptic curve E/F_q with endomorphism ring of discriminant Δ and $\left|E(F_q)\right| = C$. We obtain the curve E (or more precisely, its j-invariant) *up to twist* as a root of the Δ-th *Hilbert Polynomial* $H_\Delta(X)$ over F_q and *a priori*, from [1, Th. 72] we only have that $\mathrm{End}(E) \cong \mathcal{O}$. However, we will see that from our prescribed values of q, v, C, Δ satisfying Equation (4.2) we can augment E to ensure that $C = \left|E(F_q)\right|$ holds.

A *twist* of an elliptic curve E/F_q is another elliptic curve E'/F_q isomorphic to E via an isomorphism that is not defined over F_q. The only isomorphisms of elliptic

curves in the Weierstrass model $E \to E'$ are of the form $(x, y) \mapsto (u^2x, u^3y)$ with u in $\overline{F_q}$ [7, Proof of Th. 14.13; 28, III.1, pg.45]. The new curve E' has Weierstrass invariants $a' = u^4a, b' = u^6b$, where a, b are the Weierstrass invariants of E (recall, these invariants are the constants from the defining equation of the respective curves). When $j(E) = 1278 \cdot 4a^3/(4a^3 + 27b^3) \neq 0, 1728$ we see that u^2 lies in F_q since a', b' lie in F_q.

Such an isomorphism (u^2 in F_q, but u not in F_q) is called a *quadratic* twist and importantly curves E, E' connected via a quadratic twist satisfy $|E| + |E'| = 2(q + 1) = 2(t + C)$ [7, Sec 8.5]. This can be seen by observing that for every x in F_q the value $z = x^3 + ax + b$ is either (i) a square $z = y^2$ and (x, y) is a point on E; or (ii) not a square, but u^3y is a square for some u in $\overline{k}$ and so (x, y) is a point on the quadratic twist E'.

The final ingredient, is to determine when the resulting curve E (obtained from a root of the Hilbert Polynomial) is *ordinary*. Deuring's theory of reduction tells us that the curve E is ordinary if and only if $((\Delta_D/q)) = 1$, where $((\cdot/\cdot))$ is the *Kronecker symbol* [86, 13.4 Th.12].

In conclusion, obtaining E from the root of $H_\Delta(X)$ over F_q results in an ordinary elliptic curve with complex multiplication by an order $\mathcal{O}$ with discriminant Δ, if $((\Delta/q)) = 1$. We count the number of points on $E(F_q)$, and if it is not C, we know its quadratic twist E' *will* have C F_q-rational points.

Since we are hunting for curves E with a given rational 2^n-torsion, the cardinality of $C = \left|E(F_q)\right|$ must be a multiple of 2^{2n}. We are now presented with an algorithm. For a given discriminant Δ, begin with small cofactors $m = 1, 2, ...$, set $C = m2^{2n}$ and find solutions to (4.2) such that q is a prime power using Cornacchia's algorithm [78]. Then compute the Hilbert polynomial H_D, factor it and obtain a j-invariant of an ordinary curve with CM by $\mathcal{O}$ of conductor f, verifying first that the q found satisfies $((\Delta/q)) \neq 1$ for the given discriminant.

Sutherland notes that computing the Hilbert polynomial for large discriminants is prohibitively expensive [7, Rem 22.11; 85]. Given that the size of the class-group grows with the square root of the discriminant, and that the security (against classical attacks) of the derived action grows with the square root of the size of the class-group, we see that it is infeasible to use the CM method alone to construct ordinary elliptic curves with CM by orders with cryptographically large class-groups.

Even so, this method, together with an unoptimised sage implementation, we achieve elliptic curves with reasonably large available 2^n-torsion (but small class-groups). For example, we compute an ordinary curve with rational 2^{1024}-torsion and endomorphism ring of class number 1.

```
sage: E = EllipticCurve(GF(2**1024 * (2**1028 * 37 - 39) + 1), (0,
5121092947330662515151619982475874094566226753276844093738976787870750133519585767421321426392770947688733563538632973688209855713862038297106368015766965593933289723766553866427647543256243972844687345707642874784928675987713500883108953390975603715661616866070446843621382992061122675768504249400103602254608456139867502235017936745282023943611664506292257158333497914122536673366301109662413780908007288187217860650572926436828760576365095328133967699729996503407719593938443703329736741932047126853875775703179850452982555955326466701429024110206858946306530532211249464416126167612119801911533762422093580424246210 8, 0, 1, 0))
sage: E.torsion_basis(2**1024)
```

```
# Successfully computes 2^1024-torsion basis
[...]
sage: E.torsion_basis(2**1025)
[...]
ValueError: curve does not have full rational 2**1025-torsion
```

This computation took $\sim$ 28 hours on a laptop.

4.4 Large Class-Groups from walking down the Volcano

We can obtain larger class-groups by walking down the volcano. If a curve E has rational ℓ^m torsion, taking m vertical ℓ-isogeny steps to arrive at a curve E_{ℓ^m} we obtain an endomorphism ring $\mathrm{End}(E_{\ell^m}) \cong \mathcal{O}_{\ell^m}$ of conductor ℓ^m and so a class-group $\mathrm{Cl}(\mathcal{O}_{\ell^m})$ of size [21, Th. 7.24]

$$|\mathrm{Cl}(\mathcal{O}_{\ell^m})| = \ell^{m-1}\left(1 - \left(\left(\frac{\Delta_D}{\ell}\right)\right)\frac{1}{\ell}\right)$$

where $((\cdot/\cdot))$ is the *Jacobi symbol*. Doing this, we forfeit some rational torsion.

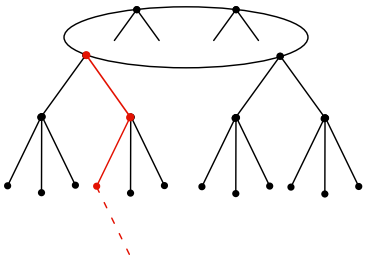

The isogeny volcano in a picture. Each node represents an (isomorphism class of an) elliptic curve, and two nodes are connected by a line if they are ℓ-isogenous. A "walk" constitutes choosing a different ℓ-isogeny every step of the way. The resulting isogeny, obtained by composing each of these ℓ-isogenies is cyclic if and only if this walk does not backtrack.

The circle at the top of the volcano is called the *crater*. The curves on the crater each have endomorphism ring $\mathcal{O}_D$, and so there are $|\mathrm{Cl}(\mathcal{O}_D)|$ many of them. Recall, we called isogenies connecting elliptic curves with isomorphic endomorphism ring *horizontal*; we now see the origins of this terminology: these isogenies are horizontal on the isogeny volcano. We note that in our case, the class number $|\mathrm{Cl}(\mathcal{O})|$ is 1 and so our crater collapses to what looks more like a "peak".

We recall that the security of a key-exchange derived from the class-group action generally depends on the size of the class-group. Although walking down the volcano in small ℓ-steps in this way does increase the size of the class-group, it *does not* increase the security of the scheme [84].

We note that SCALLOP [48] instantiates a large class-group from a relatively small discriminant and large discriminant. However, crucially, in their case, the conductor is a large prime, and so this can be viewed as one "large" step, with no intermediate steps in between.

4.5 Implementing the Class-Group

Unfortunately, at the time of writing Sage does not have an ideal class-group implementation for non-maximal orders.

```
sage: version()
'SageMath version 10.4, Release Date: 2024-07-19'
sage: QuadraticField(-3).order_of_conductor(1).class_group()
class-group of order 1 of Number Field in a with defining polynomial x^2 + 3
sage: QuadraticField(-3).order_of_conductor(2).class_group()
...
NotImplementedError: non-maximal orders are not yet supported
```

Using the theory developed in Subsection 3.3, we can write a class to perform calculations on ideals of non-maximal orders in a very explicit way.

Initialisation The general interface takes as input an order $\mathcal{O}$ with conductor f and integer coefficients m, a, b to construct the ideal $I = m[a, b + f w_D] \subseteq \mathcal{O}$. Using Proposition 3.7 the constructor of the class can verify whether the supplied values indeed define an ideal.

Conjugation The conjugate of $I = m[a, b + f w_D]$ is simply $I = m[a, -(b + f\Delta_D) + f w_D]$ because $\mathrm{tr}(w_D) = w_D + \overline{w_D} = \Delta_D$. Indeed,

$$\overline{b + f w_D} = b + f\overline{w_D} = b + f(\Delta_D - w_D) = b + f\Delta_D - f w_D = -\left(-(b + f\Delta_D) + f w_D\right)$$

so the conjugate of I is given by

$$\begin{aligned}
\overline{m[a, b + f w_D]} &= \overline{ma\mathbb{Z} + m(b + f w_D)\mathbb{Z}} \\
&= ma\mathbb{Z} + m\overline{(b + f w_D)}\mathbb{Z} \\
&= ma\mathbb{Z} - m(-(b + f\Delta_D) + f w_D)\mathbb{Z} \\
&= ma\mathbb{Z} + m(-(b + f\Delta_D) + f w_D)\mathbb{Z} \\
&= m[a, -(b + f\Delta_D) + f w_D].
\end{aligned}$$

Ideal multiplication Given two ideals $I = [a, b + f w_D]$ and $J = [c, d + f w_D]$ the product IJ is generated as a $\mathbb{Z}$-module by the 4 pairwise products of $\mathbb{Z}$-generators of I and J, namely $ac, ad + af w_D, cb + cf w_D$, and

$$(b + f w_D)(d + f w_D) = bd - f^2(\Delta_D^2 - \Delta_D)/4 + (b + d + f\Delta_D) f w_D.$$

Proposition 3.7 tells us that IJ must be of the form $z[x, y + f w_D]$. Instead of finding a closed formula for z, x, y as a function of a, b, c, d we resort to an algorithm.

To that end, we write the submodule $M = [\alpha + \beta f w_D, \gamma + \delta f w_D]$ as a matrix in row-form in the basis $(f w_D, 1)$ and get $(\beta, \alpha; \delta, \gamma)$. It is now exactly now the content of Proposition 3.7 to say that M is an ideal of $\mathcal{O}$ if and only if there exist integer-coefficient row-reductions to write $(\beta, \alpha; \delta, \gamma)$ as an upper-triangular matrix $T = (z, zy; 0, zx)$ with $0 \leq y \lneq x$ and $0 \lneq z$. Sage[11] says matrices of this form are in *Hermite normal form* (HNF).

[11] This differs from e.g. Cohen's definition [36, Def. 2.4.2] in that Sage requires $T_{j,i} < T_{i,i}$ for $j < i$ and Cohen requires $T_{i,j} < T_{i,i}$ for $j > i$.

To use Sage's inbuilt algorithm, we must set up the following 2×4 matrix (since the operations are done on rows, not columns)

$$G = \begin{pmatrix} 0 & ac \\ a & ad \\ c & cb \\ b + d + f\Delta_D & bd - f^2(\Delta_D^2 - \Delta_D)/4 \end{pmatrix}$$

and call

```
G.hermite_form(include_zero_rows=False)
```

to yield some 2×2 matrix $(Z, Y; 0, X)$ from which we recover $z = Z, y = Y/z, x = X/z$ so that $IJ = z[x, y + fw_D]$.

Reduced Bases Recall the *norm* of an element $\alpha = x + yfw_D$ of an element in $\mathcal{O}$ is given by

$$\mathrm{N}(\alpha) = \alpha\overline{\alpha} = x^2 + xyf\Delta_D + y^2f^2(\Delta_D^2 - \Delta_D)/4.$$

We say that a basis β_1, β_2 of a lattice $L = \sigma_1\mathbb{Z} + \sigma_2\mathbb{Z}$ is *(Lagrange-Gauss) reduced* if $\mathrm{N}(\beta_1) \leq \mathrm{N}(\beta_2) \leq \mathrm{N}(\alpha)$ for all non-zero α in L. Such a basis always exists [22, Lem. 17.1.5].

As a direct application: an ideal I is principal if and only if $I = \beta_1\mathcal{O}$, whereby β_1, β_2 is a reduced basis of I. Indeed, when $I = \alpha\mathcal{O}$ is principal, then $\mathrm{N}(\alpha) \leq \mathrm{N}(\alpha\omega) = \mathrm{N}(\alpha)\mathrm{N}(\omega)$ for all non-zero ω in $\mathcal{O}$ because $(x, y) \mapsto \mathrm{N}(x + yfw_D)$ is a positive-definite binary quadratic form. The converse is clear.

We can compute a reduced basis we employ Algorithm 23 of [22]. It is essentially a generalised gcd computation and is also very similar to the Gram-Schmidt orthogonalisation process. This algorithm computes the scalar product of vectors in the lattice, but describes vectors using integer coefficients in the starting basis. In our case, the starting basis is $(1, fw_D)$, and so the scalar product is given on $\mathbb{Z}^2$ by

$$\begin{aligned}
&\langle (x, y), (w, z) \rangle_{(1, fw_D)} \\
&= \frac{1}{4} \left\langle \left(2x + yf\Delta_D, yf\sqrt{-\Delta_D}\right), \left(2w + zf\Delta_D, zf\sqrt{-\Delta_D}\right) \right\rangle_{(1,1)} \\
&= \frac{1}{4} \left((2x + yf\Delta_D)(2w + zf\Delta_D) - yzf^2\Delta_D\right) \\
&= \frac{1}{4} \left(4xw + 2(yw + xz)f\Delta_D + yzf^2\Delta_D^2 - yzf^2\Delta_D\right) \\
&= \frac{1}{4} \left(4xw + 2(yw + xz)f\Delta_D + yzf^2(\Delta_D^2 - \Delta_D)\right) \\
&= xw + (yw + xz)f\Delta_D/2 + yzf^2(\Delta_D^2 - \Delta_D)/4
\end{aligned}$$

which coincides with the norm when $(x, y) = (w, z)$. Here $\langle \cdot, \cdot \rangle_{(1,1)}$ denotes the usual scalar product on $\mathbb{C}$, viewing $\sqrt{\Delta_D}$ as $\iota\sqrt{-\Delta_D}$. Notably, the scalar product may result in half-integer values due to the middle term (recall that $\Delta_D \equiv 0, 1 \pmod 4$, so $(\Delta_D^2 - \Delta_D)/4$ is an integer).

With this scalar product (compatible with our norm), we can implement Algo-

rithm 23 of [22] to obtain a reduced basis of any ideal.

Ideal equivalence Two ideals I, J are equivalent in the class-group if and only if $\bar{I}J$ is principal. Using our conjugation and basis reduction algorithms, we can verify this.

Representatives of the class-group Given the *class number* $|\mathrm{Cl}\,\mathcal{O}|$ of an order $\mathcal{O}$, we can produce non-equivalent ideals until we have exhausted the whole class-group. In fact, the class-group $\mathrm{Cl}(\mathcal{O}_D)$ of the ring of integers $\mathcal{O}_D$ of an imaginary quadratic number field is almost cyclic most of the time. More precisely, the maximal subgroup of odd order of $\mathrm{Cl}(\mathcal{O}_D)$ is cylic 97.7575% of the time, when $D < 0$ is picked at random [83, 9.I (C1)]. This suggests to produce one ideal inside $\mathcal{O}_D$ and power this until the whole class-group is exhausted. This results in ideals of exponentially large norm.

Another approach is to very naïvely manually search for ideals with small norm. Given a target norm n, we factor n into $n = ma$ and then search for $0 \leq b \lneq a$ so that $a \mid \mathrm{N}(b + fw_D)$. The resulting submodule $I = m[a, b + fw_D]$ is an ideal thanks to Proposition 3.7. Performing this on $n = 1, 2, \ldots$ and verifying whether the newly found ideal is equivalent to any of the previously found ideals gives a slow, but workable approach to computing class groups. My Sage implementation can compute class-groups of size 2^{13} this way in reasonable time ($\sim$ 4 hours on a laptop).

Note that every class in $\mathrm{Cl}(\mathcal{O})$ is represented by an ideal of norm bounded by $\sqrt{-\Delta/3}$ (where $\Delta = f^2\Delta_D$, and $\mathcal{O} = [1, fw_D]$). This is a special case of *Minkowski's bound*, but is more directly seen as a consequence of [36, Lem 5.3.4] when recalling that the form class-group is isomorphic to the ideal-class-group via $ax^2 + bxy + cy^2 \mapsto [a, -(b + \Delta_D)/2 + fw_D]$ [21, Th. 7.7].

4.6 Lattice points to isogenies

Recalling the Clapoti construction, we see that we must evaluate the isogeny corresponding to the endomorphism γ defined by $I\bar{J} = \gamma\mathcal{O}$. Our implementation of ideals gives us γ in the form $\gamma = a + bfw_D$. So the challenge remains to compute fw_D as an isogeny. We present Robert's *division algorithm* in [17] section 2.3 *Evaluating endomorphisms*; and adapt it more favourably to our cases by first computing w_D and then relating this to fw_D by "*lollipopping*" on the volcano.

Let E be an elliptic curve with endomorphism ring $\mathcal{O}_D = [1, w_D]$. We relate $w_D = (\Delta_D + \sqrt{\Delta_D})/2$ to the Frobenius endomorphism π of E. We combine the equalities

$$4q = t^2 - v^2\Delta_D \quad \text{and} \quad \pi^2 - t\pi + q = 0$$

to conclude $2\pi = t + v\sqrt{\Delta_D}$ and so $\sqrt{\Delta_D} = (2\pi - t)/v$. Here we have implicitly chosen a sign for how π is embedded into $\mathcal{O}$. Moreover, we get an equation for

$$w_D = \frac{1}{2}\left(\Delta_D + \sqrt{\Delta_D}\right) = \frac{1}{v}\left(\frac{v\Delta_D - t}{2} + \pi\right).$$

Now, $t = q + 1 - \left|E(F_q)\right|$ is even by choice (the 2^n-torsion is rational $E[2^n] \subseteq E(F_q)$); and since $4q - t^2 = -v^2\Delta_D$ is even, we obtain an integer expression for $(v\Delta_D - t)/2$.

This means that w_D can be evaluated on points P with order N prime to v very

efficiently. Indeed, if v is prime to N, $[v]$ is invertible on $E[N]$. Then $w_D = v^{-1}((v\Delta_D - t)/2 + \pi)$ is scalar multiplication and evaluations of the Frobenius on $E[N]$ (v^{-1} is the inverse of v modulo N). The same is true for evaluating fw_D on points P with order prime to v/f.

We recall that in our application of CLAPOTI, we instantiate a curve with rational $N = 2^n$ torsion and then evaluate endomorphisms on the N-torsion. We now describe the relationship between N and v.

We begin by claiming that if the N-torsion is rational on an elliptic curve E defined over F_q, then $N \mid (q-1)$. Indeed, the Weil pairing $e_N\colon E[N] \times E[N]$ takes values of N-th roots of unity in $\overline{F_q}$; but at the same time, it is Galois invariant [28, 8.1.(d)] and so if $E[N]$ is F_q-rational, the image of e_N must lie in $F_q \subseteq \overline{F_q}$. Hence there is a point in F_q of multiplicative order N, and so N divides $q-1$.

On the other had, when the N-torsion is rational, then $N^2 \mid \left|E(F_q)\right|$ simply because $\mathbb{Z}/N\mathbb{Z} \times \mathbb{Z}/N\mathbb{Z} \cong E[N] \subseteq E(F_q)$.

We now turn back to Equation (4.2) from our CM method. We rewrite it as a quadratic in $q-1$ to obtain

$$v^2\Delta_D = q^2 - 2(C+1)q + (C-1)^2 = (q-1)^2 - 2C(q-1) + C(C-4)$$

where $C = \left|E(F_q)\right|$ and $f \mid v$. Hence N^2 divides $v^2\Delta_D$.

In our concrete examples, we chose $D < 0$ to have class number 1, for example $D = -7$ and $\Delta_D = -7$; moreover we picked $N = 2^n$ and $f = \ell^m$ where ℓ is some small odd prime (by first finding a curve with $2^n\ell^m$-rational torsion, and then walking m steps down the ℓ-isogeny volcano). This illustrates that, in our cases, $N = v$ and so we have no hope of evaluating $w_D = (v\Delta_D - t + 2\pi)/2v$ on points of order N using this method.

An alternative approach is to use Vèlu's formulae. Indeed, fw_D is an isogeny of degree $d_f = f^2(\Delta_D^2 - D)/4$. We can iter the d_f-torsion to find the kernel of fw_D. This becomes exponentially expensive as f grows, so we compute fw_D by "lollipopping". Since we are trying to compute fw_D as an endomorphism on a curve F obtained by walking down the ℓ-volcano from a curve E at the crater[12], we have an isogeny connecting $\nu\colon E \to F$ such that $fw_D = \nu w_D \widehat{\nu}$. That is, fw_D as an endomorphism on F is the composition $\nu w_D \widehat{\nu}$, where w_D is viewed as an endomorphism on E.

This technique is called "lollippoping" because, we are computing the *"lollipop"* endomorphism as inspired by this picture

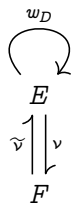

Now, computing w_D as an endomorphism on E requires only itering points in the $d = d_1 = (\Delta_D^2 - \Delta_D)/4$ torsion of E. When $N = v$ is prime to d, we can evaluate

[12]Since we are in the class-number 1 case, the crater looks more like a "peak"

$w_D = (v\Delta_D - t + 2\pi)/2v$ on points in the d-torsion, because v is invertible on $E[d]$.

4.7 Detecting split isogenies

We recall that an Abelian surface (a two-dimensional Abelian variety) is either isomorphic to some product of elliptic curves or not. Let A, A' be the latter and E_i, E'_j elliptic curves. Then we say an isogeny $f\colon E_1 \times E_2 \to A$ is a *glueing* isogeny; an isogeny $f\colon A \to E_1 \times E_2$ a *splitting* isogeny; and finally, an isogeny $f\colon E_1 \times E_2 \to E'_1 \times E'_2$ a *split* isogeny. (We apparently do not require a term for non-splitting, non-glueing isogenies $f\colon A \to A'$).

These differences are of practical importance, because the library of [3] assumes that the first step in the decomposition chain of a 2-dimensional 2^n-isogeny is a glueing. This means, to call the library successfully, we must detect, and then manually compute, any split isogenies that come first.

In pictures. We first have the expected case of [3]

$$E \times E \overset{\text{glueing}}{\longrightarrow} A_1 \longrightarrow \cdots \longrightarrow A_{n-1} \overset{\text{splitting}}{\longrightarrow} E_I \times E_{\bar{I}}$$

but it is also possible to have

$$E \times E \overset{\text{split}}{\longrightarrow} E_1 \times E'_1 \overset{\text{split}}{\longrightarrow} E_2 \times E'_2 \overset{\text{glueing}}{\longrightarrow} A_3 \longrightarrow \cdots \longrightarrow A_{n-1} \overset{\text{splitting}}{\longrightarrow} E_I \times E_{\bar{I}}$$

or other combinations of splits, glueings and splittings.

We can easily detect whether a 2-isogeny is split or not by looking at its kernel as described by the following

Lemma 4.1. *Let E_1, E_2 be elliptic curves and $f\colon E_1 \times E_2 \to A$ be a 2-isogeny between PPAV with kernel $K = \langle (P_1, P_2), (Q_1, Q_2) \rangle$. Then f is split if and only if*

(i) $P_2 = 0,\ Q_1 = 0$ *or* *(ii)* $P_1 = 0,\ Q_2 = 0$ *or* *(iii)* $\varphi(P_1) = P_2,\ \varphi(Q_1) = Q_2$

where $\varphi\colon E_1 \xrightarrow{\sim} E_2$ is an isomorphism.

Proof. Let $f\colon E_1 \times E_2 \to A \cong E'_1 \times E'_2$ be a split isogeny. By Lemma 2.1, we know that $f = (f_1, f_2; f_3, f_4)$ is a matrix of isogenies (or zero-maps). Letting $d_i = \deg(f_i)$ we now distinguish all three possible cases for the degrees

$$\text{(i) } d_1 = d_4 = 2,\ d_2 = d_3 = 0 \qquad \text{(ii) } d_1 = d_4 = 0,\ d_2 = d_3 = 2$$
$$\text{(iii) } d_1 = d_4 = 1,\ d_2 = d_3 = 1$$

These correspond to

$$(i)\ \begin{pmatrix} f_1 & 0 \\ 0 & f_4 \end{pmatrix}\ \textit{“diagonal”} \quad (ii)\ \begin{pmatrix} 0 & f_2 \\ f_3 & 0 \end{pmatrix}\ \textit{“anti-diagonal”} \quad (iii)\ \begin{pmatrix} f_1 & f_2 \\ f_3 & f_4 \end{pmatrix}\ \textit{“full”}$$

We treat only cases (i) and (iii) for (i) and (ii) are virtually identical.

In the “full” case (ii), Kani tells us that the kernel of f is given by $\ker(f) = \{(x, f_2^{-1} f_1(x)) \mid x \in E_1[2]\}$ (Lemma 2.5). Indeed, since $\deg(f_i) = 1$, the f_i are all isomorphisms and $\widetilde{f_i} = f_i^{-1}$. Moreover, $f_2^{-1} f_1$ is an isomorphism of $E_1 \to E_2$, proving the claim.

In the "diagonal" case, the kernel is clearly given by $\ker(f_1) \times \ker(f_4)$. Since f_1, f_4 are 2-isogenies, they are cyclic and $\ker(f_1) = \langle P_1 \rangle, \ker(f_4) = \langle Q_2 \rangle$ for some points P_1, Q_2 in the 2-torsion of E_1 and E_2 respectively, proving the claim.

So we have shown that every split isogeny has a kernel that satisfies one of the three cases. The converse follows immediately from our construction.

Given a kernel $K = \langle (P_1, 0), (0, Q_2) \rangle$ as in case (i), we simply write down the isogenies f_1, f_4 with kernels $\langle P_1 \rangle, \langle Q_2 \rangle$ to form $f = (f_1, 0; 0, f_4)$. The same is done in case (ii).

In the full case, we recover the isomorphism φ from the coordinates, since the only isomorphisms are of the form $(x, y) \mapsto (u^2 x, u^3 y)$ for u in $\overline{F_q}$ [7, Proof of Th. 14.13; 28, III.1, pg.45]. Now we form the split isogeny $f = (\varphi, 1; -\varphi, 1)$. We verify that f is indeed a 2-isogeny

$$f\widetilde{f} = \begin{pmatrix} \varphi & 1 \\ -\varphi & 1 \end{pmatrix} \begin{pmatrix} \varphi^{-1} & -\varphi^{-1} \\ 1 & 1 \end{pmatrix} = \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix} = [2]$$

and that its kernel is

$$\begin{aligned} \ker(f) &= \{(P_1, P_2) \in E_1[2] \times E_2[2] \mid \varphi(P_1) + P_2 = 0 \text{ and } -\varphi(P_1) + P_2 = 0\} \\ &= \{(P_1, P_2) \in E_1[2] \times E_2[2] \mid \varphi(P_1) = P_2\} \end{aligned}$$

where we used that $P_i = -P_i$ since these points lie in $E_i[2]$. Finally, we note that f is indeed split, since $f\colon E_1 \times E_2 \to E_2 \times E_2$.

In all cases (i)-(iii), we could recover a split isogeny with kernel $K = \langle (P_1, P_2), (Q_1, Q_2) \rangle$, and so have proven that an isogeny with such a kernel is split, completing the proof. □

We note that the very first isogeny in the decomposition $E \times E \to E_I \times E_{\overline{I}}$ has $E_1 = E_2$ in the language of the lemma. Here it is even easier to detect whether we are in the full case or not. Indeed, if $j(E) \neq 0, 1728$, then the only possible automorphisms $E \cong E_1 \xrightarrow{\sim} E \cong E_2$ are $\varphi = \pm 1$.

4.8 Results, Conclusion and Outlook

Piecing together the previous subsections, I have written an almost complete Sage implementation to compute the isogeny class-group action on ordinary elliptic curves. Alas, successfully evaluating the action on our toy examples with this code still remains out of reach due to what I suspect are small bugs, and not fundamental problems with the methodology

From this, and having specialised the theory of Clapoti [2] explicitly to ordinary elliptic curves, I conclude that the only remaining "true" hurdle to obtain a cryptographic group action from the isogeny class-group action on ordinary elliptic curves is to find ordinary elliptic curves with rational 2^n torsion with endomorphism ring $\mathcal{O}$ of large prime discriminant (being prime ensures no volcano-style attacks apply).

This suggests that Clapoti can be successfully leveraged to compute the isogeny class-group on *supersingular* curves, since for these it is much easier to find curves with the prescribed torsion and endomorphism rings. For example, one can produce a large

prime of the form $p = 4l_1 \cdots l_n - 1$ with l_i prime and then consider the supersingular curve described by $y^2 = x^3 + x$ over F_{p^2}. Theorem 2 of [88] then tells us that the $l_1 \cdots l_n$-torsion is rational on E. This similar to how CSIDH [41] finds suitable curves. The corresponding (F_p-rational) endomorphism ring $\mathrm{End}_{F_p}(E)$ of E has discriminant $-4p$, yielding a class-group $\mathrm{Cl}(\mathrm{End}_{F_p}(E))$ of order $\sim \sqrt{p}$.

We summarise Sage implementations presented by this thesis and what their capabilities are in the following.

The source code is available at

`github.com/rrueger/IsogenyClassGroupComputation`

Curve finding Given a target (prime power) torsion, my implementation produces an ordinary elliptic curve in reasonable time. For example, producing a curve with rational 2^{64}-torsion takes ~ 1 second, with rational 2^{512}-torsion in ~ 8 hours and with rational 2^{1024}-torsion in ~ 28 hours. This is implemented by the function `find_E`.

Recalling Minkowski's bound from Subsection 4.5, telling us that every element of the class-group $\mathrm{Cl}(\mathcal{O})$ of an order $\mathcal{O}$ with discriminant Δ is represented by an ideal of norm at most $\sqrt{-\Delta/3}$; and recalling Siegel's theorem stating $|\mathrm{Cl}(\mathcal{O})| \sim \sqrt{-\Delta}$ we see that at least the $\sim 2^{512}$-torsion must be rational if the degree equation $\deg(\alpha)\mathrm{N}(I)+\deg(\beta)\mathrm{N}(J) = 2^n$ is to have solutions, and the class-group has size 2^{256} (to achieve 128 bits of classical security). One interpretation of this heuristic, is that my implementation yields curves that have sufficient 2^n-torsion for cryptographic purposes.

Moreover, using the volcano walking techniques discussed in Subsection 4.4, I have written an implementation that yields an ordinary elliptic curve with endomorphism ring of prescribed (prime power) conductor and prescribed (prime power) rational-torsion. For example, it can produce a curve with rational 2^{64}-torsion and endomorphism ring of conductor 3^{27} in ~ 1 second.

Ideals and the class-group I written have a crude implementation of ideals in the $m[a, b + fw_D] \subseteq [1, fw_D] = \mathcal{O}$ form. This allows for easily constructing ideals of specific norms, and using the Minkowski bound, exhaustively search for representatives of elements in the class-group $\mathrm{Cl}(\mathcal{O})$ of smallest possible norm. I could compute class toy examples of ideal class-groups of size $\sim 2^{13}$ in ~ 4 hours using this implementation. This is implemented by the class `Cideal` and the method `find_class_group`.

Endomorphism finding for degree equations Given an ordinary elliptic curve E with attached class-group $\mathrm{Cl}(\mathcal{O})$ of size $\sim 2^{10}$, I could compute matching endomorphism pairs α, β in $\mathrm{End}(E) \cong \mathcal{O}$ to satisfy the degree equations $\deg(\alpha)\mathrm{N}(I) + \deg(\beta)\mathrm{N}(J) = 2^n$ ($n \leq 64$) for every class $[I] = [J]$ in $\mathrm{Cl}(\mathcal{O})$ in ~ 10 hours. This was implemented using naïve random sampling from a small enough radius, so that the resulting equivalent ideals had sufficiently small norms and the degree equation had solutions. This is implemented by the method `find_pairs_full`. There is also a version, which searches only for integer endomorphisms implemented by `find_pairs_squares`.

Detecting split isogenies From the theory developed in Subsection 4.7, I wrote an implementation for detecting, and then computing split isogenies that appear at the beginning of a 2-isogeny chain in two dimensions. The source for this idea is in the

`experiments.py` file.

Unfortunately, these routines pieced together are not robust enough to successfully compute a full evaluation, but it appears that with more careful analysis the algorithms should execute fully.

As an outlook, we should investigate whether it is possible to forgo the requirement of rational torsion. This obstacle insurmountable on a first glance, since the kernel of a 2^n isogeny in Kani's Lemma (Lemma 2.5) is patently given by evaluations of endomorphisms on the 2^n-torsion. If this is not rational, we cannot hope to pass a kernel to a library of type [3] in the way we currently do. Solving this problem would resolve the tension between finding ordinary elliptic curves with rational 2^n-torsion.

Conversely, it is certainly worth looking at how we construct ordinary elliptic curves with prescribed rational torsion. The folklore CM method is slightly more general than what we want to achieve: it finds curves with prescribed (rational) *order*, not rational *torsion*. As discussed in Subsection 4.6, requiring the N-torsion to be rational on a curve E/F_q immediately implies that $N \mid (q-1)$, a condition that does not appear in the usual CM method. This additional constraint may make parameter searching faster, but is unlikely to yield a breakthrough since we still rely on factoring the Hilbert Class polynomial. Nevertheless, a closer look at these methods may prove fruitful. According to Sutherland's lectures of 2021, the 2012 paper [85] is still the state of the art. Perhaps new efforts a decade later could yield new results.

The implementations of this thesis can be improved in various places. For example, the random sampling of equivalent ideals can be done as it is done in SQISignHD [18, Lem. 4.2.1]. The current implementation simply samples integers x, y from a small enough ball to hope that the resulting equivalent ideal $J = \overline{\beta} I / \mathrm{N}(I)$ with $\beta = max + y(b + fw_D) \subseteq I = m[a, b + fw_D]$ has norm $\mathrm{N}(I)$ is small enough so the degree equation has solutions.

Pierrick Dartois' paper [52] (published only days ago!) provides the first generic algorithms to compute 4-dimensional 2^n-isogenies (as opposed to the implementation of SQISignHD [18] which could only evaluate 4-dimensional 2^n-isogenies of very specific form). Working towards a fully-featured (e.g. robustly handling split isogenies) implementation of these algorithms would certainly be very beneficial for any higher-dimensional isogeny-based research, and in particular computing the isogeny class-group action in the Clapoti way.

References

[1] *Mathematics of Isogeny Based Cryptography*. Luca De Feo (2017). From arxiv.org/abs/1711.04062.

[2] *Introducing Clapoti(s): Evaluating the isogeny class-group action in polynomial time*. Aurel Page and Damien Robert (2023). From eprint.iacr.org/2023/1766.

[3] *An Algorithmic Approach to* $(2,2)$*-isogenies in the Theta Model and Applications to Isogeny-based Cryptography*. Pierrick Dartois, Luciano Maino, Giacomo Pope, and Damien Robert (2023). From eprint.iacr.org/2023/1747.

[4] *Theta Isogenies*. Pierrick Dartois, Luciano Maino, Giacomo Pope, and Damien Robert (2023),
Webpage github.com/ThetaIsogenies/two-isogenies

[5] *A Note on* $(2,2)$*-isogenies via Theta Coordinates*. Jianming Lin, Saiyu Wang, and Chang-An Zhao (2024). From eprint.iacr.org/2024/971.

[6] *Efficient Computation of* $(3^n, 3^n)$*-Isogenies*. Thomas Decru and Sabrina Kunzweiler (2023). From eprint.iacr.org/2023/376.

[7] *Elliptic Curves*. Andrew Sutherland (2021). From ocw.mit.edu/courses/18-783-elliptic-curves-spring-2021.

[8] *Algebraic Geometry*. Robin Hartshorne (1977), Springer.

[9] *Algebraic Geometry*. Andreas Gathmann (2022). From agag-gathmann.math.rptu.de/en/alggeom.php.

[10] *Algebraic Geometry (v5.10)*. James S. Milne (2008), 234 pp. From jmilne.org/math.

[11] *Abelian Varieties (v2.00)*. James S. Milne (2008), 166+vi pp. From jmilne.org/math.

[12] *Algebraic Geometry (v6.02)*. James S. Milne (2017), 221 pp. From jmilne.org/math.

[13] *Algebraic Number Theory (v3.08)*. James S. Milne (2020). From jmilne.org/math.

[14] *Fields and Galois Theory (v5.10)*. James S. Milne (2022). From jmilne.org/math.

[15] *An efficient key recovery attack on SIDH*. Wouter Castryck and Thomas Decru (2022). From eprint.iacr.org/2022/975.

[16] *Breaking SIDH in Polynomial Time*. Damien Robert (2022). From eprint.iacr.org/2022/1038.

[17] *Some applications of higher dimensional isogenies to elliptic curves (overview of results)*. Damien Robert (2023). From eprint.iacr.org/2022/1704.

[18] *SQISignHD: New Dimensions in Cryptography*. Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski. From eprint.iacr.org/2023/436.

[19] *Fast change of level and applications to isogenies*. David Lubicz and Damien Robert (2022). From www.normalesup.org/~robert/pro/publications/articles/change_level.pdf.

[20] *The Number of Curves of Genus Two with Elliptic Differentials*. Ernst Kani (1997). From doi.org/10.1515/crll.1997.485.93. Available from Kani's homepage mast.queensu.ca/~kani/papers/numgenl.pdf, with page numbers begin at 1 instead of 93.

[21] *Primes of the form $x^2 + ny^2$: Fermat, class field theory, and complex multiplication*. David A. Cox (2013), Vol. 34, John Wiley & Sons.

[22] *Mathematics of Public Key Cryptography. Version 2.0*. Steven D. Galbraith (2018). From math.auckland.ac.nz/~sgal018/crypto-book/crypto-book.html.

[23] *Algebraic Number Theory*. Jürgen Neukirch (1992).

[24] *Algebraic Number Theory*. Franz Lemmermeyer (2005).

[25] *Cryptographic Group Actions and Applications*. Navid Alamati, Luca De Feo, Hart Montgomery, and Sikhar Patranabis (2020). From eprint.iacr.org/2020/1188.

[26] *Towards Quantum-Resistant Cryptography from Supersingular Elliptic Curve Isogenies*. Luca De Feo, David Jao, and Jérôme Plût (2011). From eprint.iacr.org/2011/506.

[27] *Abelian Varieties*. David Mumford (1970). Note: Page numbers are consistent with 1970 edition.

[28] *The arithmetic of elliptic curves*. Joseph H. Silverman (2009), Vol. 106, Springer.

[29] *A Direct Key Recovery Attack on SIDH*. Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, and Benjamin Wesolowski (2023). From eprint.iacr.org/2023/640.

[30] *Regarding the dual of an abelian variety*. Ariyan Javanpeykar (2017), Webpage math.stackexchange.com/q/2111737. Ariyan Javanpeykar is stackexchange user math.stackexchange.com/users/61397/ariyan-javanpeykar

[31] *Abelian Varieties*. Bas Edixhoven, Ben Moonen, and Gerard van der Geer. From van-der-geer.nl/~gerard/AV.pdf.

[32] *Computing functions on Jacobians and their quotients*. Jean-Marc Couveignes and Tony Ezome (2015).

[33] *Abelian surfaces and jacobian varieties over finite fields*. Hans-Georg Rück (1990).

[34] *Finding The Four Squares in Lagrange's Theorem*. Paul Pollack and Enrique Treviño (2018). From campus.lakeforest.edu/trevino/finding4squares.pdf.

[35] *Introduction to Commutative Algebra*. Michael Atiyah and Ian G. Macdonald (1969).

[36] *A course in computational algebraic number theory*. Henri Cohen (1993).

[37] *Basic Algebra II*. Nathan Jacobson (1989).

[38] *Hard Homogeneous Spaces*. Jean-Marc Couveignes (1997). From eprint.iacr.org/2006/291.

[39] *Public-key Cryptosystem Based on Isogenies*. Alexander Rostovstev and Anton Stolbunov (2006). From eprint.iacr.org/2006/145.

[40] *Cryptographic hash functions from expander graphs*. Denis Charles, Eyal Goren, and Kristin Lauter (2006). From eprint.iacr.org/2006/021.

[41] *CSIDH: An Efficient Post-Quantum Commutative Group Action*. Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes (2018). From eprint.iacr.org/2018/383.

[42] *SeaSign: Compact isogeny signatures from class-group actions*. Luca De Feo and Steven D. Galbraith (2018). From eprint.iacr.org/2018/824.

[43] *CSI-FiSh: Efficient Isogeny based Signatures through class-group Computations*. Ward Beullens, Thorsten Kleinjung, and Frederik Vercauteren (2019). From eprint.iacr.org/2019/498.

[44] *Constructing elliptic curve isogenies in quantum subexponential time*. Andrew M. Childs, David Jao, and Vladimir Soukharev (2010). From arxiv.org/abs/1012.4019. The original version is from December 2010. In April 2018 there was a revision.

[45] *Post-Quantum Cryptography: Round 3*. NIST (2020), Webpage csrc.nist.gov/Projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions

[46] *Post-Quantum Cryptography: Digital Signature Schemes: Round 1 Additional Signatures*. NIST (2023), Webpage csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures

[47] *SQISign: compact post-quantum signatures from quaternions and isogenies*. Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski (2020)

[48] *SCALLOP: scaling the CSI-FiSh*. Luca De Feo, Tako Boris Fouotsa, Péter Kutas, Antonin Leroux, Simon-Philipp Merz, Lorenz Panny, and Benjamin Wesolowski (2023). From eprint.iacr.org/2023/058.

[49] *SCALLOP-HD: group action from 2-dimensional isogenies*. Mingjie Chen, Antonin Leroux, and Lorenz Panny (2023). From eprint.iacr.org/2023/1488.

[50] *FESTA: Fast Encryption from Supersingular Torsion Attacks*. Andrea Basso, Luciano Maino, and Giacomo Pope (2023). From eprint.iacr.org/2023/660.

[51] *POKE: A Framework for Efficient PKEs, Split KEMs, and OPRFs from Higher-dimensional Isogenies*. Andrea Basso (2024). From eprint.iacr.org/2024/624.

[52] *Fast computation of 2-isogenies in dimension 4 with the Theta model and cryptographic applications*. Pierrick Dartois (2024). From eprint.iacr.org/2024/1180.

[53] *SQIsign2D-West: The Fast, the Small, and the Safer*. Andrea Basso, Pierrick Dartois, Luca De Feo, Antonin Leroux, Luciano Maino, Giacomo Pope, Damien Robert, and Benjamin Wesolowski (2024). From eprint.iacr.org/2024/760.

[54] *SQIsign2D-East: A New Signature Scheme Using 2-dimensional Isogenies*. Kohei Nakagawa and Hiroshi Onuki (2024). From eprint.iacr.org/2024/771.

[55] *SQIPrime: A dimension 2 variant of SQISignHD with non-smooth challenge isogenies*. Max Duparc and Tako Boris Fouotsa (2024). From eprint.iacr.org/2024/773.

[56] *SQIAsignHD: SQIsignHD Adaptor Signature*. Péter Kutas and Farzin Renan (2024). From eprint.iacr.org/2024/561.

[57] *CSI-FiSh really isn't polynomial-time*. Lorenz Panny (2023), Webpage yx7.cc/blah/2023-04-14.html

[58] *Is SIKE Broken Yet?*. Luca De Feo (2023), Webpage issikebrokenyet.github.io

[59] *Endomorphisms of Abelian Varieties over Finite Fields*. John Tate (1966). From doi.org/10.1007/BF01404549.

[60] *Counting points on elliptic curves over finite fields*. René Schoof (1995). From www.jstor.org/stable/43972442. Also available from Schoof's homepage at www.mat.uniroma2.it/~schoof/ctg.pdf.

[61] *Isogénies entre courbes elliptiques*. Jacques Vélu (1971).

[62] *Importance of determining whether a number is squarefree, using geometry*. Bill Dubuque (2010). From math.stackexchange.com/q/2249 (Date 2021-05-14). Bill Dubuque is stackexchange user math.stackexchange.com/users/242/bill-dubuque.

[63] *Quantum Equivalence of the DLP and CDHP for Group Actions*. Steven Galbraith, Lorenz Panny, Benjamin Smith, and Frederik Vercauteren (2018). From eprint.iacr.org/2018/1199.

[64] *Full Quantum Equivalence of Group Action DLog and CDH, and More*. Hart Montgomery and Mark Zhandry (2022). From eprint.iacr.org/2022/1135.

[65] *A Simpler and More Efficient Reduction of DLog to CDH for Abelian Group Actions*. Steven Galbraith, Yi-Fu Lai, and Hart Montgomery (2024). From eprint.iacr.org/2024/191.

[66] *The Relationship Between Breaking the Diffie-Hellman Protocol and Computing Discrete Logarithms*. Ueli Maurer and Stefan Wolf (1999). From crypto.ethz.ch/publications/files/MauWol99b.pdf.

[67] *Dlog is Practically as Hard (or Easy) as DH – Solving Dlogs via DH Oracles on EC Standards*. Alexander May and Carl Richard Theodor Schneider (2023). From eprint.iacr.org/2023/539.

[68] *Extending the GHS Weil descent attack*. Steven G. Galbraith, Florian Hess, and Nigel P. Smart (2001). From eprint.iacr.org/2001/054.

[69] *Faster computation of isogenies of large prime degree*. Daniel J. Bernstein, Luca De Feo, Antonin Leroux, and Benjamin Smith (2020). From eprint.iacr.org/2020/341.

[70] *CTIDH: faster constant-time CSIDH*. Gustavo Banegas, Daniel J. Bernstein, Fabio Campos, Tung Chou, Tanja Lange, Michael Meyer, Benjamin Smith, and Jana Sotáková (2021). From https://eprint.iacr.org/2021/633.

[71] *Algorithms for quantum computation: discrete logarithms and factoring*. Peter Shor (1994). From ieeexplore.ieee.org/document/365700.

[72] *Towards practical key exchange from ordinary isogeny graphs*. Luca De Feo, Jean Kieffer, and Benjamin Smith (2018). From eprint.iacr.org/2018/485.

[73] *Abelian varieties over finite fields*. William C. Waterhouse (1969). From www.numdam.org/articles/10.24033/asens.1183.

[74] *Orienting supersingular isogeny graphs*. Leonardo Colò and David Kohel (2020). From eprint.iacr.org/2020/985.

[75] *On Oriented Supersingular Elliptic Curves*. Hiroshi Onuki (2020). From arxiv.org/pdf/2002.09894.

[76] *M-SIDH and MD-SIDH: countering SIDH attacks by masking information*. TakoBoris Fouotsa, Tomoki Moriya, and Christophe Petit (2023). From eprint.iacr.org/2023/013.

[77] *Computing the degree of the field of n-torsion points*. Adam Van Tuyl (2001). From ms.mcmaster.ca/~vantuyl/papers/ntorsion_2001.pdf.

[78] *On Cornacchia's algorithm for solving the diophantine equation* $u^2 + dv^2 = m$. F. Morain and J.-L. Nicolas (1990). From www.lix.polytechnique.fr/Labo/Francois.Morain/Articles/cornac.pdf.

[79] *Über die Classenzahl quadratischer Zahlkörper*. Carl Siegel (1935). From eudml.org/doc/205054.

[80] *Introduction to the arithmetic theory of automorphic functions*. Goro Shimura (1971).

[81] *A Subexponential-Time Quantum Algorithm for the Dihedral Hidden Subgroup Problem*. Greg Kuperberg (2005). From doi.org/10.1137/S009753970343634.

[82] *Cryptographic Schemes Based on Isogenies*. Anton Stolbunov (2012). From ntnuopen.ntnu.no/ntnu-xmlui/bitstream/handle/11250/262577/529395_FULLTEXT01.pdf.

[83] *Heuristics on class-groups of number fields*. Henri Cohen and H. W. Lenstra Jr. (1983), Number Theory Noordwijkerhout.

[84] *On the security of OSIDH*. Pierrick Dartois and Luca De Feo (2021). From eprint.iacr.org/2021/1681.

[85] *Accelerating the CM method*. Andrew Sutherland (2012). From arxiv.org/pdf/1009.1082.

[86] *Elliptic Functions*. Serge Lang (1987), Springer.

[87] *Sage Documentation* (2024), Webpage doc.sagemath.org/html/en/reference/arithmetic_curves/sage/schemes/elliptic_curves/ell_finite_field.html

[88] *Deuring for the People: Supersingular Elliptic Curves with Prescribed Endomorphism Ring in General Characteristic*. Jonathan Komada Eriksen, Lorenz Panny, Jana Sotáková, and Mattia Veroni (2023). From eprint.iacr.org/2023/106.

[89] *On the efficient representation of isogenies (a survey)*. Damien Robert (2024). From eprint.iacr.org/2024/1071.

[90] *Quantum Security Analysis of CSIDH*. Xavier Bonnetain and André Schrottenloher (2018). From https://eprint.iacr.org/2018/537.

[91] *Quantum circuits for the CSIDH: optimizing quantum evaluation of isogenies*. Daniel J. Bernstein, Tanja Lange, Chloe Martindale, and Lorenz Panny (2018). From eprint.iacr.org/2018/1059.