Libraries and Learning Services

University of Auckland Research Repository, ResearchSpace

Copyright Statement

Isogenies of Abelian Varieties in Cryptography

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy in Mathematics, the University of Auckland, 2019

Doctor of Philosophy, the University of Auckland

(Mathematics)

Yan Bo TI

Co-Authorship forms are removed from digital copy and saved as a separate file on Research Space for security reasons, please email Research Space (researchspace@auckland.ac.nz) for any query.

Abstract

Isogenies of abelian varieties have been used in cryptography to create post-quantum cryptosystems. In particular, supersingular elliptic curve isogenies have been used to construct key exchange, encryption and signature protocols and hash functions. This thesis concerns itself with results relating to this cryptosystem and presents four main findings: two attacks, a reduction and a generalisation.

The two attacks on the cryptosystem are an adaptive attack and a fault attack. The adaptive attack targets instances of the cryptosystem using static keys and is able to recover the secret with close to optimal number of queries for most use cases. The fault attack targets the cryptosystem embedded in hardware and is able to recover the entire secret with one successful perturbation.

The reduction shows that breaking the cryptosystem is at most as difficult as computing endomorphism rings of supersingular elliptic curves. It relies on the equivalence of the category of supersingular elliptic curves under isogenies and the category of invertible modules under homomorphisms.

We also generalise the cryptosystem from isogenies between supersingular elliptic curves to isogenies between supersingular principally polarised abelian surfaces. In particular, we propose a genus two version of the key exchange protocol called Genus Two SIDH (G2SIDH). We perform some analysis of the security of G2SIDH by studying the isogeny graph of principally polarised abelian surfaces. A by-product of this study is that a naive generalisation of the hash function to genus two is no longer collision resistant.

Acknowledgements

First and foremost, I want express my deepest gratitude to my amazing supervisor: Steven Galbraith. In the duration of my studies, he has shared his time and attention selflessly. His aptitude and verve for research, paired with his sense of humour made this PhD a super singular period of my life.

Next, I would like to thank Barak Shani, Christophe Petit, Craig Costello, Lukas Zobernig, Chloe Martindale, Victor Flynn, and Samuel Dobson for being wonderful people and for giving me the pleasure of meeting and working with them. I would also like to thank all the great people that I had the great fortune of meeting at conferences around the world there are too many of you to list, but YOU are definitely included in the list.

Living in Auckland has been amazing in no small part due to the marvellous people I have met. I would first like to thank my group of friends, collectively known as Preseperos, and also the incredible people of the Maths department. From the early days of "Bacteria", to the golden age of "Snake", to the everlasting "25" and "OGH", there was never a want of fun. However, this thesis could have been completed in 2 years without these most welcomed distractions.

My sojourn in Auckland also precipitated in a serendipitous encounter with mi cariño, Ceci. Her constant support and encouragement was indispensable to this endeavour.

Last but not least, I have to thank my long-suffering and beloved family (Koli and Bunny included) for showing endless patience and enduring love while I was away in New Zealand and also when I was in the UK. I dedicate this thesis to them.

Introduction

The first encounter students will have with abelian varieties will most likely be the Pythagorean triples, which are integral solutions to the equation

$$a^2 + b^2 = c^2 \,.$$

The task of finding Pythagorean triples can be reduced to finding rational solutions of the Pythagorean equation. One is then usually introduced to the group law which allows for the generation of infinitely many rational solutions. Hence, this is an explicit study of abelian varieties. This gives us our first glimpse into the rich theory that is abelian varieties. The study of these varieties still remains an active area of research with the resolution of Fermat's Last Theorem as the latest landmark achievement in the discipline. We will review the necessary background on abelian varieties in Chapter 1. The aim of this chapter is to provide a concise and (hopefully) self-contained treatment of this vast topic.

And now for something completely different, we have cryptography which is the study of secure communications in the potential presence of adversaries. The first known use of ciphers dates back to around 100BC, when Julius Caesar has been recorded to have used the Caesar cipher to secure messages to his generals. Despite its early start, cryptography is a relatively young science. Cryptography has only taken its modern form with the advent of computers in the 1950's; in the process, it has also developed a distinctly number-theoretic flavour. The strongest influence number theoretic methods have had on cryptography is in public-key cryptography. The first[1] public-key cryptosystem, and possibly the earliest use of number theory in cryptography, is the Diffie–Hellman key exchange protocol [DH76]. The key exchange protocol relies on the difficulty of the Discrete Logarithm Problem (DLP) in the group of multiplicative elements of a finite field. The discrete logarithm problem in a group G is defined as: Given $g, h \in G$ such that $h = g^x$, find x.

There are many algorithms to solve a generic[2] instance of the DLP. The main algorithms are the following:

- The Baby-step Giant-step (BSGS) [Sha71] is a "space-time tradeoff" algorithm that searches for collisions that would yield a solution to the DLP by pre-computing a list and comparing sequentially generated elements against this list. The time complexity is approximately $O(\sqrt{N})$, where N is the order of the group.

- Pollard rho and Pollard Kangaroo algorithms [Pol75] are collision finding algorithms that have similar runtime complexity to the BSGS algorithm. The Pollard rho algorithm can be further improved with the use of negation maps which speeds up the original algorithm by a factor of $\sqrt{2}$ [BKL10]. Note that this technique explicitly exploits the geometric structure of elliptic curves. The two algorithms have negligible space requirements at a cost of a linear increase in complexity as compared to BSGS.

- The Pohlig–Hellman algorithm [PH78] is able to efficiently solve specific instances of the DLP when there exists a nested sequence of (small) subgroups. The aim is to solve the DLP in small prime order subgroups and lift the solution to a larger subgroup, repeating the process to the Sylow subgroup. The algorithm then combines the results using the Chinese Remainder Theorem to compute the discrete logarithm.

---

[1]Although a classified discovery has already been made in the secretive bowels of GCHQ by Williamson.
[2]Generic meaning that we only have access to the group operation and have no access to other properties/structure the group might have.

On top of these generic algorithms to solve the DLP, one could solve the DLP when more information or structure is given. In the case of the DLP over a finite field, one could employ Index Calculus methods. Index Calculus methods exploit the notion of "smoothness" and the ability to decompose arbitrary elements in the field into smooth elements. In the case of the group of multiplicative elements in a finite field, one can use non-generic methods such as Index Calculus methods to solve the DLP. The algorithm takes random powers of g, lifts them into the integers and factorises these powers. If the factors of the powers are smooth (we say that it lies in the factorbase), then the relation is kept. When factorised into elements in the factorbase, one can express the random power as a sum of the discrete logarithms of elements of the factorbase. After iterating this procedure sufficiently many times, one is able to use linear algebra methods to recover the discrete logarithms of all/most of the elements of the factorbase. Now, taking the product of h with some g^s, such that the factors of the product are all contained in the factorbase, one can recover x by solving a linear equation. This necessitated the use of larger keys to compel the adversary to devote the same amount of resources to break the cryptosystem.

The moment that abelian varieties came to the fore in cryptography was when researchers started realising that the groups of abelian varieties are particularly well-suited to modern computer arithmetic. This is due to elements of abelian varieties having a convenient representation in finite fields, and having a group law that is defined by polynomials. Furthermore, the range of parameters afforded by the Hasse interval made them more flexible in the order of the groups. The first use of abelian varieties in cryptography based a cryptosystem on the difficulty of the Elliptic Curve DLP (ECDLP).

Abelian Varieties and Cryptography

In 1985, Neal Koblitz [Kob87] and Victor Miller [Mil85] independently suggested the use of elliptic curves in cryptography, spawning Elliptic Curve Cryptography (ECC).

Besides the generic discrete logarithm algorithms to solve the DLP in any generic group, ECC has an extra geometric structure that has been exploited to solve very particular ECDLP instances. Indeed, most attacks on the ECDLP rely on transferring the DLP into another group where the DLP can (hopefully) be solved more easily.

- There exists a bilinear pairing on elliptic curves which maps into some finite field. The MOV attack [MVO91, FR94] exploits this pairing to transfer the DLP from elliptic curves to finite fields. Provided that the field is sufficiently small, one will be able to solve the DLP in the finite field and transport the solution back into elliptic curves to solve the initial ECDLP.

- Suppose L/K is a finite extension of fields, then Weil descent is a map which sends a variety over L to a variety over K. The Weil descent attack [GHS02] transports the ECDLP to the DLP in a higher dimensional variety. The attack relies on the DLP being easier in the higher dimensional variety with a smaller base field. If so, the DLP can be solved in the higher dimensional variety before bringing the solution back to solve the ECDLP.

- The anomalous curve attack [Sma99, SA98] exploits the existence of a map from the elliptic curve group into the vector space of holomorphic differentials of the curve. Since this vector space is isomorphic to the base field of the curve, the ECDLP can be transferred into the base field to be solved. This is usually extremely practical as one can employ stronger algorithms to solve the DLP in finite fields. In fact, an attacker uses the Euclidean algorithm, which is extremely efficient. This map will only exist if the characteristic of the base field divides the order of the elliptic curve group.

Despite the existence of these attacks, well-chosen elliptic curves can provide much smaller key sizes than key sizes for finite field cryptography. Elliptic curves are abelian varieties of dimension one. To generalise ECC, one can look at abelian varieties of higher dimension, and in 1989, Koblitz [Kob89] suggested the use of hyperelliptic curves in public-key cryptography. The points on the hyperelliptic curve lack a group structure; the DLP is in the group of points on the Jacobian of the curve. The main draw of hyperelliptic curve cryptography (HECC) is the hope of efficiency through computing group operations in a smaller finite field. However, it has been shown that if the dimension of the abelian variety is too large, then index calculus attacks can be more efficient than generic group algorithms [AD93, Gau00].

Hitherto, the algorithms described are known as classical algorithms; that is, they are a series of steps that are performed on a classical computer. In 1994, Peter Shor proposed a quantum algorithm to solve the DLP in a generic group. Being a quantum algorithm, it requires a quantum computer to be executed. Thus, the slow and irresistible march towards the first quantum computer capable of breaking modern cryptosystems heralds the end of cryptosystems based on the difficulty of the DLP.

Post-quantum Cryptography

This bleak landscape spurred the research into cryptosystems that are based on hard problems that cannot be efficiently solved by either classical and quantum computers. The candidates for such a cryptosystem can be grouped into 5 main categories.

- Lattice-based cryptosystems base their security on well-studied problems of lattices, namely the shortest-vector and closest-vector problems. The first lattice-based schemes began with papers by Atjai [Ajt96] in 1996, and Hoffstein, Pipher and Silverman [HPS98] in 1998. Encryption and signature protocols in lattice-based cryptosystems are generally fast, but have relatively large keys, ciphertexts, and signatures. Also, there is a long-standing problem of estimating lattice attacks that makes choosing security parameters extremely difficult.

- Code-based protocols employ error-correcting codes for encryption and key exchange and more recently, signatures. The security is based on the difficulty of decoding general linear codes as proposed by McEliece in 1978 [McE78]. The advantages of using code-based cryptography are its speed and the reliance on a problem that is known to be NP-hard. However, large key sizes have continued to deter potential users.

- Hash-based schemes use hash trees to produce signatures. The first hash based digital signature was published by Ralph Merkle in 1979 in his PhD thesis [Mer79]. Confidence in the security of these schemes come from the continued resistance to cryptanalysis of the cryptosystem since its inception. However, the lack of encryption and key exchange protocols is the main drawback for the hash-based paradigm.

- Multivariate cryptographic systems rely on the difficulty of solving systems of multivariate equations. The debut of multivariate schemes came in 1988 with a paper by Matsumoto and Imai [MI88]. These are problems that have been well-studied and can be proved to be NP-hard or NP-complete. Multivariate cryptosystems are able to build the shortest signature schemes amongst its post-quantum competitors. But it does not have efficient encryption or key exchange protocols.

- Supersingular Elliptic Curve Isogeny cryptosystems are based on the assumption that isogenies between elliptic curves are difficult to compute. A key exchange and encryption protocol has been proposed in 2011 [JD11], and signature schemes have seen improvements in their efficiency. The key exchange protocol has the shortest keys amongst its competitors, but the computational complexity of the scheme is relatively high. Confidence in this cryptosystem is still weak due to the novelty of the cryptographic assumptions.

The purpose of this thesis is to examine the security and protocols of isogeny-based cryptography. We hope to bolster confidence in supersingular elliptic curve isogeny cryptosystems through this study. In this thesis, we are most interested in the cryptanalysis of isogeny-based cryptosystems and the mathematics lurking in the background.

Isogenies of Abelian Varieties

The protagonist of isogeny-based cryptography are isogenies. An isogeny between abelian varieties is a homomorphism which is surjective and has finite kernel. It can be shown that an isogeny preserves both the algebraic and also certain geometric structures of the variety. Isogenies are classified by their degree, which is defined by the degree of the function field extension induced categorically. In the case where that extension is separable, we have that the degree coincides with the size of the kernel. We call an isogeny with degree ℓ an ℓ-isogeny.

There is a correspondence between ℓ-isogenies and subgroups of the ℓ-torsion which will play a crucial role in this thesis. Every subgroup of the torsion can be the kernel of an isogeny, and every isogeny is uniquely defined by its kernel up to composition with an isomorphism. In other words, isogenies from isomorphic varieties with the same kernels under the isomorphism will map into a pair of isomorphic varieties. The upshot of this is that it makes sense to talk about isomorphism classes of isogenies.

In the case of elliptic curves, this correspondence can be computed efficiently: given a subgroup of the ℓ-torsion, Vélu's formula allows one to compute the isogeny whose kernel is equal to the given subgroup. Conversely, given an isogeny, one can compute its kernel to obtain the subgroup of the ℓ-torsion.

A special class of isogenies that map a variety to itself are known as endomorphisms (although the converse is not always true). The set of all endomorphisms on a variety, together with the zero endomorphism, forms a ring with multiplication given by composition, and addition given by point-wise addition. Note that the zero endomorphism is a map that sends every point from a variety to the identity element. The zero endomorphism is not an isogeny, but we call it the zero endomorphism for convenience. The endomorphism ring is a $\mathbb{Z}$-module and we have the inclusion of $\mathbb{Z}$ into the endomorphism ring. The exact isomorphism type of the endomorphism ring for an arbitrary abelian variety is not known. However, the endomorphism rings of elliptic curves are well understood.

A non-trivial theorem states that endomorphism rings of elliptic curves are isomorphic to the integers, or an order of an imaginary quadratic field, or a maximal order of a quaternion algebra. A fortiori, endomorphism rings of elliptic curves over a finite field fall into the latter two categories. We call elliptic curves over a finite field with endomorphism rings isomorphic to a maximal order of a quaternion algebra, supersingular elliptic curves. If the endomorphism ring of an elliptic curve is an order of an imaginary quadratic field, we say that it is ordinary.

Now, fix two distinct primes p and ℓ, and set a field k with characteristic p. Then one can define the ℓ-isogeny graph over k to be a directed graph whose vertices are isomorphism classes of elliptic curves over k, and an edge between two vertices exists if and only if there is an ℓ-isogeny between them over k. The first interesting question one might ask is if the graph is connected. In general, the answer is no.

A deep result of Tate states that abelian varieties are isogenous over k if and only if they have the same cardinality over k. It is important to note that this result considers isogenies of all degrees and not only ℓ-isogenies. However, we can still infer from the theorem that edges can only exist between isomorphism classes of elliptic curves with the same cardinality. The result is a large number of components in the ℓ-isogeny graph. To gain a handle on the structure of the components, one must look at the endomorphism algebra of the components.

Ordinary curves form multiple ℓ-volcanoes grouped according to their endomorphism algebras. The structure of these ℓ-volcanoes will be examined in Chapter 4, where we will also see how class group actions allow us to navigate within the components.

Supersingular elliptic curves on the other hand form a single connected component. In fact, the component is an $(\ell + 1)$-regular graph. The problem of finding an isogeny between two supersingular elliptic curves can then be translated into finding a path between two vertices in the ℓ-isogeny graph. As mentioned, endomorphism rings of supersingular elliptic curves are maximal orders of a quaternion algebra. It can be shown that the quaternion algebra is the definite quaternion algebra ramified at p and infinity. Yet another important result is the equivalence of categories between the category of supersingular elliptic curves and their isogenies, and the category of invertible left modules of a maximal orders and their module homomorphisms. The arithmetic of the maximal orders of quaternion algebras allows us to translate the isogeny problem to finding invertible modules of these maximal orders.

We will examine the structure of the isogeny graph of (certain classes of) abelian surfaces in Chapter 5. This has implications on the security of cryptosystems constructed from isogenies of abelian varieties.

Isogenies of Abelian Varieties in Cryptography

Founded on the belief that finding isogenies is difficult, in 1997 Jean-Marc Couveignes presented at ENS the notion of "Hard Homogeneous Spaces". Essentially, he described a set acted transitively on by a group fulfilling certain conditions, one of which is that given two elements in the set, it is difficult to find an element in the group taking one element of the set to the other. Furthermore, he proposed the set of isomorphism classes of elliptic curves (with group action given by the class group action) as a hard homogeneous space. The paper was submitted to Crypto '97, but was not accepted, but the

manuscript was made public in 2006 [Cou06], and later published in 2007 [Cou07]. In 2004, Anton Stolbunov independently proposed a similar proposal in his Masters Thesis [Sto04] and also included a signature scheme. This culminated in a paper with Rostovtsev in 2006 [RS06] where they described a key exchange protocol that they suggested might be post-quantum secure. We will refer to this collectively as the Couveignes–Rostovtsev–Stolbunov (CRS) cryptosystem. The hope that the CRS cryptosystem might be post-quantum secure was slightly misplaced as Childs, Jao and Soukharev [CJS14] found in 2010 a subexponential quantum algorithm that is able to force users towards larger key sizes. This algorithm relies on the commutative group action of the class group, and reduces the isogeny problem to a hidden shift problem. One can then apply the Kuperberg algorithm to compute the keys.

In 2006, the use of supersingular isogenies was proposed by Kristin Lauter in "The Second Cryptographic Hash Workshop" organised by NIST. This subsequently appeared in the paper by Charles, Goren and Lauter [CLG09]. In this paper, they described a hash function that navigates the supersingular 2-isogeny graph. The input of the hash was used to deterministically generate a path from a fixed initial vertex, while the output of the hash was the endpoint of the path. Hence the security of the hash depended on the issue of finding isogenies between supersingular elliptic curves.

In 2011, David Jao and Luca De Feo [JD11] proposed a key exchange scheme which is similar to the CRS cryptosystem. The key difference is the use of supersingular elliptic curves over $\mathbb{F}_{p^2}$ which foils the quantum attack of Childs, Jao and Soukharev since the class group action does not exist on supersingular elliptic curves over $\mathbb{F}_{p^2}$. This innovation led to a different problem: that of "non-commutativity" of the operations performed by the two parties involved in key exchange. In general, isogenies will not be commutative since one can only compose isogenies if the domains and codomains are compatible. When we speak of commutative operations in the CRS framework, we speak of the representation of isogenies by the action of commuting group elements. Since this action is absent for supersingular elliptic curves over $\mathbb{F}_{p^2}$, extra torsion data is required to allow for two parties to compose their secret isogenies in a well-defined manner. The order of the compositions will not matter (hence giving us "commutativity") since quotients by subgroup do commute. By ensuring that both parties send additional information in the protocol, they were able to surmount this complication. This key exchange protocol is dubbed the Supersingular Isogeny Diffie–Hellman key exchange (SIDH).

In 2017, a proposal was submitted to the NIST standardisation process. The Supersingular Isogeny Key Encapsulation (SIKE) suite is based on the ideas of SIDH, and has a public key encryption algorithm and a key encapsulation mechanism. As of 2019, the SIKE proposal has reached the second round of the standardisation process.

In 2018, Castryck et al. delineated a supersingular isogeny cryptosystem that is able to retain the commutative action of CRS while still retaining some post-quantum security. The commutativity of this action has led the authors to name the protocol CSIDH (commutative SIDH), and pronounced as "seaside". This followed efforts from De Feo, Kieffer and Smith [DKS18] to improve the speed of the CRS cryptosystem. The key ingredient of CSIDH is the use of supersingular elliptic curves over $\mathbb{F}_p$. This imposes restrictions on the endomorphism ring which allows class group actions to once again act on the isomorphism classes of elliptic curves. However, the subexponential CJS algorithm does apply to this scheme.

We will present SIDH and the CGL hash function in Chapter 2. The treatment of CSIDH will be concise and will follow from the discussion of endomorphism rings in Chapter 4.

Since the publication of SIDH and CSIDH, numerous implementations and signature schemes have been proposed, as have some cryptanalysis on the cryptosystems. However, the task of examining the entirety of this field would be an undertaking too great for the author and indeed the reader! Hence, we will focus on the contributions that this thesis has made to the field.

Contributions

The first contribution of this thesis is to present the first attack on SIDH. The attack is an adaptive attack which means that users are no longer able to use static keys in the key exchange protocol. Static keys are used in cases where non-interactive key exchange is desired, additionally, one can obtain one-way authentication for free. This is now no longer a secure option. As mentioned, the non-commutativity of the SIDH operations necessitated the publication of auxiliary information which is absent in most Diffie–Hellman schemes. The authors noticed this oddity but did not anticipate that an adaptive attack would

be able to break the cryptosystem. Our adaptive attack, published as [GPST16], is able to break SIDH with the optimal number of queries. Indeed, by querying an oracle that returns one bit of information, the attack is able to recover the isogeny with less than $\log_2(n)$ queries, where n is approximately the number of bits of the secret. Additionally, the attack is able to evade lightweight countermeasures and can only be thwarted by countermeasures that would double the complexity of SIDH. This attack will make up the first section of Chapter 3.

The next contribution is a fault attack on SIDH [Ti17]. The fault attack targets the auxiliary information of the protocol. The crux of the analysis of the auxiliary information in [JD11] is that the auxiliary information can in no way leak information about the secret isogenies. However, the fault attack targets this very idea by changing the auxiliary information in such a way that it leaks information about the secret isogeny. In fact, a single fault will be sufficient to divulge the entire secret. The fault attack will be presented in the second section of Chapter 3.

Recall the equivalence of categories between the category of supersingular elliptic curves and isogenies, and the category of invertible left modules of a maximal order and their module homomorphism. In [KLPT14], the authors were able to show that the supersingular isogeny problem has a corresponding problem in the category of invertible modules. Furthermore, the authors were able to solve the problem in that particular category. In [GPST16] and §4.3.2, we expand on this link and show that lattice reduction methods are able to allow us to solve the supersingular isogeny problem of SIDH in the category of invertible modules. Armed with this result, we show that breaking SIDH is at most as hard as computing the endomorphism rings of an arbitrary supersingular elliptic curve.

Lastly, the final contribution of this thesis is the generalisation of elliptic curve isogeny-based cryptography. Elliptic curves are abelian varieties of dimension one, hence [FT19] and §2.3 aim to generalise this to higher dimensions. Generalisation of the cryptosystem to higher dimensions have first been proposed by Takashima in [TY09] and [Tak18]. In particular, they suggested that the CGL hash function can be generalised to abelian surfaces, where it might be more efficient. More concretely, they proposed looking at genus two hyperelliptic curves, and looking at the isogeny graph of principally polarised supersingular abelian surfaces (PPSSASs). But Takashima was wary that the isogeny graph of abelian surfaces may harbour pitfalls that would compromise on the security of the protocol. In §5.4.1 and [FT19], we show that this concern is valid as cycles in the isogeny graph mean that the hash function loses its collision resistance when generalised to abelian surfaces. Subsequently, Castryck, Decru and Smith [CDS19] have been able to tweak the genus two hash function sufficiently to avoid these collision generating cycles.

The second result of [FT19] is the generalisation of SIDH to genus two. The Genus Two SIDH (G2SIDH) is presented in §2.3. In that section, we will realise the algorithms necessary for computing isogenies for abelian surfaces and the associated protocols in G2SIDH. We will also present some security analysis in §5.4.2 using the theory developed in Chapter 5.

Notation and Conventions

A field is usually denoted by k and a finite field with q elements is denoted by $\mathbb{F}_q$. The algebraic closure of a field k is denoted by $\overline{k}$. In this thesis, we will frequently use p and ℓ to represent different primes, use E for elliptic curves, H for hyperelliptic curves, A for abelian surfaces, and X for a general abelian variety. The use of ϕ will be restricted to the discussion of isogenies. Given two abelian varieties X and Y, we write $X \sim Y$ to denote that the two are isogenous and write $X \cong Y$ if they are isomorphic. We use $\mathcal{O}$ to represent the identity of the abelian variety. We use O to mean the complexity of an algorithm, or an order of a field or algebra.

When enumerating statements in this work we will use (a), (b), ..., for independent statements; (1), (2), ..., for connected statements and conditions; and (i), (ii), ..., for equivalent statements.

Contents

Chapter 1

Abelian Varieties

Abelian varieties feature centrally in this thesis. In particular, we will focus on abelian varieties of low dimension: elliptic curves and abelian surfaces. In this chapter, we will survey some of the necessary background from algebraic geometry and state a number of results that will be used later in the thesis. We will assume familiarity with varieties.

We will divide this chapter into four sections. The first section will cover abelian varieties in total generality, in particular, the focal point will be on isogenies between abelian varieties and their properties. The second section will peel away some of the generality and specialise to elliptic curves and abelian surfaces. Here, the construction of these varieties will come to the fore and we will also discuss algorithms for dealing with these objects. We will also introduce divisors which will make explicit the group structure of these abelian varieties. The penultimate section will cover properties of the curves and surfaces, and will feature algorithms for computing the isogenies in the case of elliptic curves. In the final section, we will look to generalise the notion of supersingularity in elliptic curves to higher dimensions.

1.1   Abelian Varieties and their Isogenies

We assume general definitions and background from [Mum08]. Let k be a perfect and algebraically closed field. An *abelian variety X* is a complete algebraic variety over a field k with a group law $m : X \times X \to X$ such that m and the inverse map are both morphisms of varieties. We will use $\mathcal{O}$ to denote the identity of the abelian variety

Let X and Y be abelian varieties, then a homomorphism $\phi : X \to Y$ is called an *isogeny* if it is surjective and has a finite kernel. We say that two abelian varieties X and Y over $\mathbb{F}_q$ are *isogenous* over $\mathbb{F}_q$ if there is an isogeny $\phi : X \to Y$. The surjectivity of the isogeny induces a finite algebraic extension $\phi^*(k(Y)) \subseteq k(X)$. We ascribe the field extension the usual definition of separability. We then define the (in)separable degree of ϕ to be $[k(X) : \phi^*(k(Y))]$, the (in)separable degree of the field extensions.

The following proposition shows that the kernels of separable isogenies carry a lot of information about the isogeny.

Proposition 1.1 ([Mum08, pg.63, $(*)$]). *Let $\phi : X \to Y$ be an isogeny. We have that*

$$\# \ker \phi = \text{separable degree}(\phi).$$

In fact, we have a correspondence between finite subgroups and isogenies as presented in the next theorem. This reduces the study of morphisms between abelian varieties to the study of subgroups.

Theorem 1.2 ([Mum08, pg. 72, Thm. 4]). *Let X be an abelian variety. Then there is a 1–1 correspondence between the two sets of objects:*

  (a) finite subgroups $K \subset X$,

  (b) separable isogenies $\phi : X \to Y$, where two isogenies $\phi_1 : X \to Y_1$, $\phi_2 : X \to Y_2$, are considered equal if there is an isomorphism $\psi : Y_1 \to Y_2$ such that $\phi_2 = \psi \circ \phi_1$, which is set up by $K = \ker \phi$, and $Y = X/K$.

More is known about the kernel structure of an isogeny. Suppose that $\phi : X \to Y$ is an isogeny, then $\ker \phi$ is a finite group subscheme of X of order $\deg \phi$. Then by the structure theorem of finitely generated modules over PIDs, we have that

$$\ker \phi \cong \mathbb{Z}/a_1\mathbb{Z} \times \cdots \times \mathbb{Z}/a_n\mathbb{Z},$$

where $n \in \mathbb{Z}$ and $a_i \mid a_{i+1}$. If ϕ is separable, we say that ϕ is an $(a_1, \ldots, a_n)$-isogeny.

If n is co-prime to the characteristic of k, the *multiplication-by-n* map is a separable isogeny and will be denoted by $[n]$. Furthermore, the kernel of this map consists of points of the abelian variety of order n. The set of *n-torsion points* of an abelian variety A is denoted by $A[n] = \{P \in A(\bar{k}) \mid [n]P = \mathcal{O}\}$.

We will show in §1.2 that Pic^0 is a functor from the category of schemes to the category of abelian groups. Now, let X be an abelian variety and we want to examine $\mathrm{Pic}^0(X)$. A non-trivial discussion in [Mil86a, §9–10] shows that $\mathrm{Pic}^0(X)$ is an abelian variety and that it exists and is unique up to isomorphism. We call such an abelian variety the *dual* of X and we denote it as $X^\vee$. Since, Pic^0 is a (contravariant) functor, we can also map across isogenies $\phi : X \to Y$ to get $\phi^\vee : Y^\vee \to X^\vee$. We will say a little more about dual abelian varieties in the case of elliptic curves in §1.3.

Now, let X, Y be abelian varieties and let $\phi : X \to Y$ be an isogeny between them. We have that $\ker \phi$ is a subscheme of X, and following the construction of the dual, we have that $\ker \phi^\vee$ is a subscheme of $Y^\vee$. The next theorem shows that the functor Pic^0 applied to $\ker \phi$ would give us $\ker \phi^\vee$. In other words, the theorem shows that an isogeny and its dual have kernels which are dual to each other. In fact, we have that they are isomorphic as finite abelian groups.

Theorem 1.3. *If $\phi : X \to Y$ is an isogeny, then so is $\phi^\vee : Y^\vee \to X^\vee$. Furthermore, if ϕ is separable, then $\ker \phi$ and $\ker \phi^\vee$ are isomorphic as finite abelian groups.*

Proof. Remark (3) of [Mum08, pg. 81] proves that $\phi^\vee$ is an isogeny. To prove the next statement, we use Corollary 2 of [Mum08, pg. 74] that says that $\ker \phi$ and $\ker \phi^\vee$ are dual finite abelian groups. Hence using the result that dual finite abelian groups are isomorphic [Apo76, §6], we are done. □

The last statement of the next result that is of special interest to us: it provides us with a way to factor a large isogeny into smaller ones. One can imagine G to be the kernel of an isogeny and so the action here is really "translation-by-G"[3]. Hence the last statement can be read as saying: every isogeny whose kernel contains the kernel of another isogeny factors through the latter one.

Theorem 1.4 ([Mum08, pg. 111, Thm. 1(A)])**.** *Let G be a finite group scheme acting on a scheme X such that the orbit of any point is contained in an affine open subset of X. Then there is a pair (Y, π), where Y is a scheme and $\pi : X \to Y$ a morphism satisfying:*

(1) as a topological space, (Y, π) is the quotient of X for the action of the underlying finite group;

(2) the morphism $\pi : X \to Y$ is G-invariant, and if $\pi_*(\mathcal{O})^G$ denotes the subsheaf of $\pi_*(\mathcal{O})$ of G-invariant functions, the natural homomorphism $\mathcal{O}_Y \to \pi_*(\mathcal{O}_X)^G$ is an isomorphism.

The pair (Y, π) is uniquely determined up to isomorphism by these conditions. The morphism π is finite and surjective. Y will be denoted X/G, and it has the functorial property: $\forall$ G-invariant morphisms $f : X \to Z$, $\exists$ a unique morphism $g : Y \to Z$ such that $f = g \circ \pi$.

At this juncture, we will strip away some of the generalities and provide concrete examples to the theory we have seen thus far.

1.2 Hyperelliptic Curves and their Jacobians

The main references for this section are [Gal12, CFA+12, Har77].

Let k be a perfect field. We define a *curve C* over k to be an integral, separated scheme of finite type over k of dimension one. We call a curve *non-singular* if all the local rings are regular. Then for a curve C with genus g, there is a finite morphism $f : C \to \mathbb{P}^1$ of degree $\leq g+1$ by the Riemann–Roch theorem. We say that a curve is *hyperelliptic* if it has positive genus and there exists a map $f : C \to \mathbb{P}^1$ of degree 2. Elliptic curves with a rational point are hyperelliptic as well.

---

[3]This observation is the underlying strategy for most isogeny algorithms. Refer to §1.3.1 and §2.3.1.

Elliptic curves are the prototypical example of abelian varieties and in this section, we will set the results from the previous section in the most concrete terms: that of elliptic curves and abelian surfaces. The fact that elliptic curves are abelian varieties might confuse an inexperienced novice into thinking that abelian surfaces are hyperelliptic curves of genus two. However, hyperelliptic curves of genus > 1 are NOT abelian varieties. In this section, we will develop the theory of Jacobians and see how abelian surfaces can be constructed from hyperelliptic curves of genus two.

Elliptic curves have a geometric structure of a genus one curve and the group law is also easy to illustrate geometrically. However, the algebraic intuition of the group law of elliptic curves is not as fathomable. The group law can be made explicit once the identification between the set of points and the divisor class group is made clear. This identification can then be used to extend the group law to hyperelliptic curves of higher genus. Hence we should begin with a discussion on divisors.

1.2.1 Divisors

There are two treatments of divisors due to Cartier and Weil respectively. We will introduce Weil divisors here and omit Cartier divisors. This is because we are primarily interested in the divisors of non-singular curves and of hyperelliptic curves in particular. Non-singular curves are smooth varieties and this allows us to work exclusively with Weil divisors since the notion of Cartier and Weil divisors coincide in this setting [Har77, II.6.11].

Definition. Let C be a non-singular curve over k. A *prime Weil divisor* on C is a closed point in C. A *Weil divisor* is a finite formal sum of prime Weil divisors; more precisely, a Weil divisor is the following finite formal sum

$$D = \sum_{P \in C} n_P \cdot P, \quad n_P \in \mathbb{Z}.$$

The *degree* of D is the sum of all the n_P, i.e. $\deg(D) = \sum_{P \in C} n_P$.

We say that a Weil divisor D is k-rational if it is stable under the action of $\mathrm{Gal}(\overline{k}/k)$. Furthermore, we can define divisors of functions.

Definition. Let C be a non-singular curve, and let $f \in k(C)^*$ be a rational function on C. Then the *divisor* of f is

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

where $\mathrm{ord}_P(f)$ is the order of vanishing or order of poles of f at P (cf. [Har77, pg. 131]).

We define a relation on divisors by denoting $D \sim D'$ if and only if $D - D' = (f)$ for some f. This relation forms an equivalence relation. Divisors of functions form the trivial class, and we say that a divisor is *principal* if it is the divisor of some function.

Definition. The *divisor class group* of C is the group of divisor classes modulo linear equivalence and is denoted by $\mathrm{Pic}(C)$. We let $\mathrm{Pic}^0(C)$ denote the elements of $\mathrm{Pic}(C)$ with degree zero.

We use the notation $\mathrm{Pic}^0(C)(k)$ to denote the set of elements in $\mathrm{Pic}^0(C)$ fixed by $\mathrm{Gal}(\overline{k}/k)$.

We will now show how the group structure on the Picard group can be used to construct an abelian variety from hyperelliptic curves.

1.2.2 Jacobians

The aim of this section is to show that the Jacobian of a curve C is an abelian variety that is naturally isomorphic to $\mathrm{Pic}^0(C)$. This natural isomorphism endows the Jacobian with a group structure. We now need to give the Jacobian its geometric structure as a variety before we can call it an abelian variety. First, let C be a curve and let k be a field such that $C(k) \neq \emptyset$. Now, for any point $P \in C(\overline{k})$ define the map

$$f : C(k) \to \mathrm{Pic}^0(C)(k)$$
$$Q \mapsto Q - P.$$

This map is induced by the injective morphism $C \hookrightarrow J$, for some algebraic variety J, which has the following universal property [Mil86b, §1]:

If A is an abelian variety and $g : C \to A$ is a morphism that maps P to the identity of A, then there exists a unique homomorphism $\phi : J \to A$ of abelian varieties such that $g = \phi \circ f$, yielding the following diagram:

$$
\begin{array}{ccc}
C & \xrightarrow{\ f\ } & J \\
& {\scriptstyle g}\searrow & \downarrow{\scriptstyle \phi} \\
& & A
\end{array}
$$

The universal property uniquely characterises J, which is our Jacobian. Hence the Jacobian thus constructed is the sheafification of the Picard functor, hence bestowing it its geometric structure. The following two theorems summarise this discussion and provide us with additional properties of the Jacobian.

Theorem 1.5 ([Mil86b, Thm. 1.1])**.** *There is an abelian variety J over k and a morphism of functors $\iota : \operatorname{Pic}_C^0 \to J$ such that $\iota : \operatorname{Pic}_C^0(T) \to J(T)$ is an isomorphism of groups whenever $C(T)$ is non-empty, where T is a connected scheme over k.*

Theorem 1.6 ([HS00, Thm. A.8.1.1])**.** *Let C be a smooth projective hyperelliptic curve over k of genus $g \geq 1$ such that $C(k) \neq \emptyset$. There exists an abelian variety $\operatorname{Jac}(C)$, called the Jacobian of C, and an injection $f : C \hookrightarrow \operatorname{Jac}(C)$, called the Jacobian embedding of C, with the following properties:*

(a) *Extend f linearly to divisors on C. Then f induces a group isomorphism between $\operatorname{Pic}^0(C)$ and $\operatorname{Jac}(C)$.*

(b) *For each $r \geq 0$, define a subvariety $W_r \subseteq \operatorname{Jac}(C)$ by*

$$
W_r = \underbrace{f(C) + \cdots + f(C)}_{r \text{ copies}} \, .
$$

(By convention, $W_0 = \{0\}$.) Then

$$
\dim(W_r) = \min(r, g) \qquad \text{and} \qquad W_g = \operatorname{Jac}(C) \, .
$$

In particular, $\dim(\operatorname{Jac}(C)) = g$.

(c) *Let $\Theta = W_{g-1}$. Then Θ is an irreducible ample divisor on $\operatorname{Jac}(C)$.*

Most of the time, we will not be interested in the geometric structure of the abelian varieties, hence it will suffice to think of Jacobian varieties as just the Picard group of a hyperelliptic curve.

There are efficient methods of representing group elements and performing the group operation. The elements of the Jacobian can be represented using Mumford representations [Gal12, §10.3.1]. Using the Riemann–Roch theorem or Theorem 1.6(b), one can show that any element of the Picard group of a curve of genus g can be uniquely determined by g curve points. The Mumford representation encodes this information in a series of polynomials. The Cantor algorithms [Gal12, §10.3.2] are then able to manipulate these representations to derive the Mumford representation of the sum of two elements.

We are now able to describe abelian varieties of dimension 1 and 2. Abelian varieties of dimension 1 are simply elliptic curves, and the group structure is inherited from the isomorphism to the Jacobian as given by the results above. There is a delicate subtlety when turning to abelian surfaces. As Weil noted, the correct higher dimensional analogue of an elliptic curve is not a general abelian variety, but a principally polarised abelian variety. However, we will postpone discussion of polarisations to the next section. For now, we will treat polarisations as an unknown condition attached to abelian varieties and will round off this section by restricting our discussion to principally polarised abelian surfaces (PPASs).

The following result, which we will prove in Chapter 5, categorises all the PPASs.

Theorem 1.7. *If $A/\overline{\mathbb{F}}_p$ is a PPAS, then $A \cong J_H$ for some smooth (hyperelliptic) genus two curve H, or $A \cong E_1 \times E_2$ where E_i are elliptic curves.*

From the result, we can focus our attention on hyperelliptic curves of genus two and elliptic curves if we would like to study PPASs. Hyperelliptic curves of genus two over a field k whose characteristic is not 2, are non-singular curves with affine models of the form

$$H : y^2 = h_6 x^6 + \cdots + h_1 x + h_0$$

where $h_i \in k$. This form is unique up to fractional linear transformations of x and an associated transformation of y as given by

$$x \mapsto \frac{ax + b}{cx + d}, \quad y \mapsto \frac{ey}{(cx + d)^3},$$

where $a, b, c, d \in k$ such that $ad - bc \neq 0$ and $e \in k^*$. We say that two hyperelliptic curves of genus two, H_1 and H_2, are isomorphic if there is such a transformation between them. The isomorphism space can be parametrised by G_2-*invariants* [CNP05].

Let H be a hyperelliptic curve of genus two over a field k with the form as above with Jacobian variety J_H. Using the Riemann–Roch theorem, one can represent a point on the Jacobian as

$$P + Q - \mathcal{O}$$

where $P, Q \in H$ and $\mathcal{O}$ is described below:

(a) If $f_6 = 0$, then we let $\mathcal{O} = 2 \cdot \infty$, where ∞ is the unique point at infinity of H.

(b) If $f_6 \neq 0$, then we let $\mathcal{O} = \infty^- + \infty^+$, where $\infty^\pm$ are the two different branches of the singularity at infinity.

We will often drop the $-\mathcal{O}$ notation and write $[P + Q]$ for a point on the Jacobian.

1.3   Polarisations and Isogenies

Given an abelian variety X, recall that the dual variety $X^\vee$ exists and is unique up to isomorphism. An ample divisor $\mathcal{L}$ of X defines an isogeny $\phi_\mathcal{L} : X \to X^\vee$ known as a *polarisation* of X. If the polarisation is an isomorphism, then we say that it is *principal*.

Elliptic curves are isomorphic to their dual, and hence are canonically principally polarised. It should be noted that the dual of an isogeny between elliptic curves is closely related to the concept of dual abelian varieties. Hence the self-duality of an elliptic curve and its trivial principal polarisation lurk in the background of many of the results of elliptic curves that we use regularly.

Abelian surfaces on the other hand are not necessarily polarised, much less principally polarised. Polarisations are ample line bundles and allow for the embedding of abstract varieties into projective spaces, hence giving equations to these abstract varieties. Jacobians of hyperelliptic curve will be the main source of examples of abelian surfaces in this thesis. The theorem to follow states that every Jacobian admits a principal polarisation, we will study abelian varieties equipped with principal polarisations.

Theorem 1.8 ([Mil86b, Thm. 6.6]). *The map $f_{\mathcal{L}(\Theta)} : J \to J^\vee$ is an isomorphism, where Θ is the irreducible ample divisor as defined in Theorem 1.6. Hence every Jacobian of a hyperelliptic curve of genus two is a PPAS.*

Since we are interested in PPASs, the isogenies between them would have to preserve the additional structure of principal polarisation. There is a non-degenerate skew-symmetric bilinear pairing on a principally polarised abelian variety A over k given by

$$e_m : A[m](\overline{k}) \times A^\vee[m](\overline{k}) \to \overline{k}^*,$$

where m is co-prime to p. This is the *Weil pairing*.

Since we are working with PPASs, we can identify A and $A^\vee$ to obtain a pairing on A. Note that without a principal polarisation, this identification will not yield a non-degenerate pairing since the polarisation which sends A to its dual would have non-trivial kernel.

The following proposition shows the relation between pairings, polarisation and isogenies.

Proposition 1.9 ([Mil86a, Prop. 16.8]). *Let $\phi : A \to B$ be an isogeny of degree co-prime to $\mathrm{char}(k)$, and let $\lambda : A \to A^\vee$ be a polarisation of A. Let e^λ be defined as in [Mil86a, §16]. Then $\lambda = \phi^*(\lambda')$ for some polarisation λ' on B if and only if $\ker \phi \subset \ker \lambda$ and e^λ is trivial on $\ker \phi \times \ker \phi$.*

Remark. *Let A be a PPAS. A priori, it would seem that this proposition is not compatible with the principal polarisation of A, since if α is the principal polarisation of A, then we have that $\ker \alpha$ is trivial, and so the condition $\ker \phi \subseteq \ker \alpha$ will always be false unless ϕ is trivial.*

The subtlety in reading this result is to consider another polarisation $\lambda = [\deg(\phi)] \circ \alpha$, which is no longer a principal polarisation but we have that $\ker \phi \subseteq \ker \lambda$. Then using [Mil86a, Rem. 16.9], we have that $\deg(\lambda') = \deg(\phi)^{-2} \deg(\lambda) = \deg(\alpha) = 1$, and hence, we conclude that λ' is a principal polarisation.

This shows that a principal polarisation can induce a principal polarisation.

Hence we can see that through pairings, we are able to preserve the polarisation of abelian varieties. This is especially important to us as we need to preserve the principal polarisation. This leads naturally to the next definition.

Definition. Let A be a principally polarised abelian variety over $\mathbb{F}_q$, and let ℓ be a positive integer co-prime to q. We say a subgroup S of $A[\ell]$ is maximal ℓ-isotropic if

(1) the ℓ-Weil pairing on $A[\ell]$ restricts trivially to S, and

(2) S is not properly contained in any other subgroup of $A[\ell]$ satisfying (1).

We call the first condition the *isotropic condition*.

Definition. Let A be a PPAS over $\mathbb{F}_q$, and let ℓ be a prime co-prime to q. Then an (ℓ, ℓ)-*isogeny* is an isogeny on A such that its kernel is maximal ℓ-isotropic.

Remark. *We will mainly focus our discussion for the rest of this thesis on (ℓ, ℓ)-isogenies. Proposition 1.9 and the remark that followed clearly show that (ℓ, ℓ)-isogenies and compositions of them are the only non-trivial isogenies that will preserve principal polarisations. Hence in the study of PPASs, and of the isogenies between them, (ℓ, ℓ)-isogenies are the natural isogenies to focus our attention on.*

1.3.1 Vélu's formula

There exists an algorithm to compute ℓ-isogenies between elliptic curves with complexity $O(\ell)$. This algorithm was introduced by Vélu in 1971 [Vél71]. In the applications to follow, we will see that the isogenies we are interested in computing will have large but smooth degrees. Hence the general strategy for computing these isogenies is to factor (c.f. Theorem 1.4) them into smaller isogenies and compute with these smaller isogenies.

The algorithm takes as inputs a curve E_1 over a field k, which has the form

$$y^2 = x^3 + ax + b \,,$$

and a list of points of a finite subgroup of E_1 which we will call G. It outputs the Weierstrass model for the codomain curve E_2 of a separable isogeny, ϕ, with kernel G, and ϕ as rational maps on E_1.

The strategy of the algorithm is to represent ϕ as follows: for all $P \notin G$

$$\phi(P) = \left(x_P + \sum_{Q \in G \backslash \{\mathcal{O}\}} (x_{P+Q} - x_Q) \,,\ y_P + \sum_{Q \in G \backslash \{\mathcal{O}\}} (y_{P+Q} - y_Q) \right)$$

and for any $P \in G$, $\phi(P) = \mathcal{O}$. This representation makes explicit the invariance of ϕ under translation by elements of G and it is also clear that $G = \ker \phi$.

To generate the rational functions for ϕ, let $G^+ = (G \backslash \{\mathcal{O}\}) / \langle -1 \rangle$ be the equivalence classes of the points in G without the identity where each point is identified with its inverse. Then for each $P \in G^+$, we define the values

$$g_P^x = 3x_P^2 + a \,, \quad g_P^y = -2y_P \,, \quad v_P = 2g_P^x \,, \quad u_P = \left(g_P^y \right)^2 \,.$$

We also define

$$v = \sum_{P \in G^+} v_P \,, \quad w = \sum_{P \in G^+} u_P + x_P v_P \,.$$

Then $\phi : E_1 \to E_2$ is given by

$$\phi(x,y) = \left(x + \sum_{P \in G^+} \left(\frac{v_P}{x - x_P} - \frac{u_P}{(x - x_P)^2} \right) \,,\, y + \sum_{P \in G^+} \left(\frac{2yu_P}{(x - x_P)^3} + v_P \frac{y - y_P - g_P^x g_P^y}{(x - x_P)^2} \right) \right) \,.$$

The equation for E_2 is given by

$$y^2 = x^3 + (a - 5v)x + (b - 7w) \,.$$

1.4 Supersingularity

In this section, we will briefly touch on the concept of supersingularity. Supersingularity in elliptic curves has many equivalent definitions. In particular, we say that an elliptic curve is *supersingular* if one of the following conditions hold:

Theorem 1.10 ([Sil09, Thm. V.3.1]). *Let k be a field of characteristic p, and let E/k be an elliptic curve. Then the following are equivalent:*

(i) *$E[p^r] = 0$ for one (all) $r \geq 1$.*

(ii) *$\mathrm{End}(E)$, the endomorphism ring over the closure of k (c.f. Chapter 4), is an order in a quaternion algebra.*

To generalise the notion of supersingularity, one would hope that it would suffice for all the conditions in Theorem 1.10 to hold for a supersingular abelian variety. Unfortunately, that is not the case. In fact, for an arbitrary abelian variety, these conditions are no longer equivalent. However, when restricting to PPASs, we do have the following equivalent conditions that we can use for supersingularity:

Theorem 1.11. *Let k be an algebraically closed field of characteristic p, and let A be a PPAS, then the following are equivalent:*

(i) *A has p-rank 0, i.e. $A[p^r] = 0$ for one (all) $r \geq 1$.*

(ii) *$\mathrm{End}_k(A) \otimes \mathbb{Q} \cong M_{2 \times 2}(B_{p,\infty})$, where $B_{p,\infty}$ is the definite quaternion algebra ramified at p and infinity.*

(iii) *The characteristic polynomial of the Frobenius endomorphism is a power of a linear polynomial.*

(iv) *The Frobenius acts as a rational scalar (the centre $\mathbb{Q}[\pi] = \mathbb{Q}$).*

(v) *$[\mathrm{End}_k(A) \otimes \mathbb{Q} : \mathbb{Q}] = 16$.*

(vi) *A is k-isogenous to the square of a supersingular elliptic curve, all of whose endomorphisms are defined over k.*

(vii) *the slopes of the p-divisible group of A are $1/2$.*

Proof. Tate's theorem [Tat66, Thm. 2(d)] gives us the equivalence between (ii), (iii), (iv), (v), and (vi). The equivalence (i) $\iff$ (v) comes from [AP15]. The last equivalence (vi) $\iff$ (vii) is the result of [AP15, Thm. 1.1]. $\square$

Hence, from the theorem, we can use the following definition for supersingularity of PPASs.

Definition. Let k be a field of characteristic p, and let A/k be a PPAS. We say that A is *supersingular* if A is isogenous over $\overline{k}$ to a product of supersingular elliptic curves. We say that A is *superspecial* if A is isomorphic over $\overline{k}$ to a product of supersingular elliptic curves as PPASs.

We will also say that a hyperelliptic curve H is supersingular (resp. superspecial) if its Jacobian is supersingular (resp. superspecial).

Remark. *Note that the definition differs slightly from statement (vi) of Theorem 1.11. The definition states that A is isogenous to a **product** of supersingular elliptic curves, whereas the theorem states that A is isogenous to a **power** of a supersingular elliptic curve. A priori, one should not expect an equivalence between the two statements. However, a result in [Shi79, Thm. 3.5] (attributed to Deligne) states that A is isogenous to a product of supersingular elliptic curves if and only if it is isogenous over $\overline{k}$ to a power of a supersingular elliptic curve.*

A PPAS A is *simple* over k if it is not isogenous over k to a product of lower dimensional abelian varieties and we say that it is *absolutely simple* if it is simple over $\overline{k}$. Hence a supersingular PPAS (PPSSAS) is never absolutely simple.

Chapter 2

Isogeny-based Cryptosystems

We recall the Diffie–Hellman key exchange protocol. Let G be a cyclic group with generator g and let Alice and Bob be the two parties performing the key exchange. Alice picks a secret scalar a and publishes g^a as her public key. The main observation here is that the map $\phi_A : g \mapsto g^a$ is an endomorphism on G. Hence, the secret scalar can be seen as a secret endomorphism known only to Alice. Upon obtaining Bob's public key, she will apply her secret endomorphism on Bob's public key g^b. This allows her to compute $\phi_A : g^b \mapsto g^{ba}$ and hence obtain the shared secret.

Isogeny-based cryptosystems generalise this by using secret isogenies in place of secret endomorphisms. The objects between the maps are abelian varieties. Since the goal is to avoid the discrete logarithm problem, one should avoid using elements of the abelian varieties as objects between the maps.

This chapter introduces the schemes and protocols of isogeny-based cryptosystems. Perhaps the most important protocol is the supersingular key exchange protocol described in §2.1. First introduced by Jao and De Feo in [JD11], it is the first post-quantum key exchange protocol based on isogenies of supersingular elliptic curves. Previous attempts at using class group actions on ordinary elliptic curves by Couveignes in an unpublished manuscript in 1997 [Cou06] and Stolbunov's thesis [Sto04] in 2004 and Rostovtsev and Stolbunov [RS06] in 2006, could be attacked using a subexponential quantum algorithm [CJS14].

The first cryptographic protocol to be based on isogenies between supersingular elliptic curves is the CGL hash function proposed in 2006 [CLG09]. The hash function is described in §2.2.

In §2.3 we will see an extension of the key exchange protocol to abelian surfaces.

We have not included an exhaustive list of all schemes that employ isogeny-based assumptions, but we will included references to some of the more interesting schemes in the literature:

Signatures: [JS14], [STW14], [GPS17], [DG19]

Verifiable Delay Functions: [DMPS19]

We will survey and describe briefly some protocols that use group actions once we have studied the endomorphism ring of elliptic curves in Chapter 4. These include CSIDH [CLM$^+$18], and the Couveignes–Rostovtsev–Stolbunov (CRS) scheme mentioned above.

In general, these cryptosystems base their security upon the difficulty of the following problem:

Problem (Elliptic Curve Isogeny Problem)**.** *Let E and E' be elliptic curves such that there exists an isogeny $\phi : E \to E'$. Given E and E', find any $\psi : E \to E'$.*

2.1 Jao and De Feo Key Exchange

There are three steps in the key exchange protocol: the set-up, the key exchange and the key derivation.

In the set-up, a prime of the form $p = 2^n \cdot 3^m \cdot f - 1$ is generated where f is small and $2^n \approx 3^m$ (more generally $p = \ell_A^n \cdot \ell_B^m \cdot f \pm 1$ where ℓ_A, ℓ_B are small primes). A supersingular elliptic curve E over $\mathbb{F}_{p^2}$ is constructed, and bases P_A, Q_A and P_B, Q_B are chosen for $E[2^n]$ and $E[3^m]$. Here a "basis" means that the group $\langle P_A, Q_A \rangle$ generated by P_A and Q_A has order 2^{2n}, and similarly, $|\langle P_B, Q_B \rangle| = 3^{2m}$. The points P_A, Q_A, P_B, Q_B are defined over $\mathbb{F}_{p^2}$ and are guaranteed to exist by the choice of the prime.

In the key exchange, Alice picks random integers $0 \leq a_1, a_2 < 2^n$ (not both divisible by 2) and Bob picks random integers $0 \leq b_1, b_2 < 3^m$ (not both divisible by 3)[4]. Alice and Bob compute

$$G_A = \langle [a_1]P_A + [a_2]Q_A \rangle, \quad G_B = \langle [b_1]P_B + [b_2]Q_B \rangle$$

respectively. Using Vélu's formula, they will then be able to compute the isogenies ϕ_A and ϕ_B with respective kernels G_A and G_B. They then compute $E_A = \phi_A(E) = E/G_A$, $\phi_A(P_B)$, $\phi_A(Q_B)$ and $E_B = \phi_B(E) = E/G_B$, $\phi_B(P_A)$, $\phi_B(Q_A)$ respectively. Their respective messages in the protocol will be

$$(E_A, \phi_A(P_B), \phi_A(Q_B)), \quad (E_B, \phi_B(P_A), \phi_B(Q_A)).$$

Upon receipt of Bob's message, to derive the shared key, Alice computes

$$\langle [a_1]\phi_B(P_A) + [a_2]\phi_B(Q_A) \rangle = \langle \phi_B([a_1]P_A + [a_2]Q_A) \rangle = \phi_B(G_A).$$

Alice then computes the isogeny from E_B, with kernel equal to this subgroup. Bob will perform a similar computation and the resulting isogenies will have kernels generated by G_A and G_B (since the subgroups have a trivial intersection). Both parties will obtain a curve isomorphic to

$$E_{AB} = E/\langle G_A, G_B \rangle = E_A/\langle \phi_A(G_B) \rangle = E_B/\langle \phi_B(G_A) \rangle.$$

Note that Vélu's formula only determines codomain curves up to isomorphism, hence it is not necessary that both parties have the same curve E_{AB}. Therefore in the key derivation, the parties take the j-invariant $j(E_{AB})$ to be their shared key.

The protocol can be summarised in the following diagram, where we use the notation from above.

As mentioned, the Jao–De Feo key exchange scheme is similar to a key exchange scheme for ordinary elliptic curves proposed by Rostovtsev and Stolbunov [RS06]. The ordinary case is based on a commutative mathematical structure, however this structure enables a subexponential-time quantum algorithm [CJS14] to compromise the system. The supersingular curves variant is lacks a commutative group action and so it seems to be a promising candidate to be a post-quantum-secure cryptosystem whose best attack is (at the moment) a fully exponential algorithm. The auxiliary points included in the protocol messages allow Jao and De Feo to get around the difficulties obtaining meaningful compositions of the secret isogenies of Alice and Bob.

We stress that the isogeny problem involved here differs from a general one in several ways. On the one hand, the special primes used and the auxiliary points given to an attacker may make the supersingular isogeny problem easier than the general isogeny problem. On the other hand there is a very strong constraint imposed on the degree of the isogeny, and this might a priori make the problem harder; we discuss this issue in more detail in §4.3.2.

2.2   Hash Function

The CGL hash function is a provably collision resistant hash function constructed from the set of supersingular elliptic curves over $\mathbb{F}_{p^2}$ with ℓ-isogenies (where ℓ is a prime different from p). It was shown that computing collisions is at least as hard as computing isogenies between supersingular elliptic curves [CLG09].

We will now formally define this hash function. Let p and ℓ be two distinct prime numbers. Consider the ℓ-isogeny graph $G_{p^2,\ell}$ that we examine in §4.2.2. The supersingular elliptic curves form a component in $G_{p^2,\ell}$ which is a connected $(\ell + 1)$-regular subgraph.

---

[4]Note that this is not the true key space for Alice and Bob as we shall see later in Lemma 3.1.

The main idea for the hash function is to use the input of the hash as directions for traversing the subgraph without backtracking. A walk is a sequence of vertices, where each adjacent vertex in the sequence is connected by an edge. A walk on the isogeny graph would correspond to a sequence of ℓ-isogenies, which can be composed into a single large ℓ^n-isogeny, where n is the number of steps taken. A non-backtracking walk is walk with the condition that, on a given step, one is not allowed to return to the vertex visited on the previous step. The point of preventing backtracking is to prevent trivial collisions by inserting or appending loops to walks since walking along the same edge in different directions would result in a multiplication-by-ℓ to occur in the single large isogeny.

For the implementation of the hash function, the authors proposed a 256-bit prime p and chose $\ell = 2$. Now, pick a random vertex on the graph to be the starting vertex of the walk. The number of steps in the walk is equal to the number of bits of the input. The output of the hash will then be the final vertex of the walk. Since there are 3 isogenies emanating from each vertex, the no-backtracking condition implies that (besides the first step) there are only 2 choices for each step. The choice between the two isogenies at the n-th step is determined by the n-th input bit. This can be done deterministically as we will now describe.

Starting at an elliptic curve $E_1 : y^2 = f(x)$, the three 2-isogenies from E_1 will have kernels generated by the three Weierstrass points $P_i = (x_i, 0)$, where x_i are the three roots of the cubic $f(x)$. Fix a lexicographical order on the roots and use it to order the roots and make a choice according to the input. Without loss of generality, suppose P_1 was chosen to generate the first kernel $\phi : E_1 \to E_2$, then we have that $\phi(P_2) = \phi(P_3)$. It can be shown that $\phi(P_2)$ would generate the kernel of the dual isogeny to ϕ, and is a non-trivial 2-torsion point on E_2. Hence, one can find the abscissa of $\phi(P_2)$ and use it to find the abscissas of the two other non-trivial 2-torsion points. Using the ordering, one can make a choice on the other two points based on the next bit.

Note that for a hash to be deterministic, one must use the same implementation of Vélu's formula, since different implementations would yield different Weierstrass models. Suppose that two parties wish to use different implementations of Vélu's formula, then they would be required to compose an isomorphism after computing the isogeny to obtain curves with the same Weierstrass model.

2.3   Extensions to Genus Two

Elliptic curves are principally polarised abelian varieties of dimension one, hence we can turn to principally polarised abelian varieties of higher dimension when looking to generalise isogeny-based cryptosystems. Elliptic curves have three 2-isogenies but PPASs have fifteen $(2, 2)$-isogenies. Hence, this motivates the use of PPASs for these cryptosystems.

In this section, we will construct the key exchange protocol for genus two which we call Genus Two SIDH (G2SIDH). The scheme presented here follows the original scheme closely. Before presenting the scheme, we will review two algorithms used to select a base PPSSAS and a key from the keyspace. We will also look briefly at the isogeny algorithms employed in the scheme.

We note that the `MAGMA` implementation of the scheme is extremely slow. An example of which is presented in Appendix A.

Selecting a Base Hyperelliptic Curve

Similar to the SIDH case, we pick primes of the form $p = 2^n \cdot 3^m \cdot f - 1$.

We consider a base hyperelliptic curve given by

$$H : y^2 = x^6 + 1 \,.$$

It can be shown that the Jacobian of H is supersingular since it is the double cover of the supersingular elliptic curve $y^2 = x^3 + 1$, which is supersingular over $\mathbb{F}_p$, since $p \equiv 2 \pmod 3$ [Sil09, Eg. V.4.5]. Since a double cover induces a $(2,2)$-isogeny, this yields that the Jacobian of H is supersingular[5]. We then take a random sequence of Richelot isogenies (cf. §2.3.1) to obtain a random PPSSAS.

---

[5]We will see later in §5.4.2 that it is in fact superspecial.

Selection of Secrets

Our aim is to use scalars to encode the secret kernel to be used by the two parties of the key exchange as this allows for a compact representation of the secret.

Firstly, let $H/\mathbb{F}_q$ be a hyperelliptic curve of genus two and let J_H be its Jacobian. The secret kernels will be maximal ℓ^n-isotropic subgroups of $J_H[\ell^n]$ of order ℓ^{2n}. As we will see in §5.1, the kernels have structure $C_{\ell^n} \times C_{\ell^{n-k}} \times C_{\ell^k}$, where $0 \leq k < n/2$. Hence they should be generated by three points: Q_1, Q_2 and Q_3. Furthermore, to fulfil the condition of isotropy, we also require that the generators satisfy

$$e_{\ell^n}(Q_1, Q_2) = e_{\ell^n}(Q_1, Q_3) = e_{\ell^n}(Q_2, Q_3) = 1\,.$$

Our approach is summarised by the following steps:

Pre-computation

Step 1: Find generators for $J_H[\ell^n]$. Name them P_1, P_2, P_3, P_4.

Step 2: Find the values $\alpha_{i,j}$ such that $e_{\ell^n}(P_i, P_j) = e_{\ell^n}(P_1, P_2)^{\alpha_{i,j}}$.

Secret selection

Step 3: Pick some $(r_1, r_2, r_3, r_4) \in \{0, \ldots, \ell^n - 1\}^4$ such that they are not simultaneously divisible by ℓ.

Step 4: Pick a random[6] $0 \leq k < n/2$ and compute s_1, s_2, s_3, s_4 and t_1, t_2, t_3, t_4 by solving the two linear congruences

$$\begin{pmatrix} r_1 s_2 - r_2 s_1 + \alpha_{1,3}(r_1 s_3 - r_3 s_1) \\ +\alpha_{1,4}(r_1 s_4 - r_4 s_1) + \alpha_{2,3}(r_2 s_3 - r_3 s_2) \\ +\alpha_{2,4}(r_2 s_4 - r_4 s_2) + \alpha_{3,4}(r_3 s_4 - r_4 s_3) \end{pmatrix} \equiv 0 \mod \ell^k$$

$$\begin{pmatrix} r_1 t_2 - r_2 t_1 + \alpha_{1,3}(r_1 t_3 - r_3 t_1) \\ +\alpha_{1,4}(r_1 t_4 - r_4 t_1) + \alpha_{2,3}(r_2 t_3 - r_3 t_2) \\ +\alpha_{2,4}(r_2 t_4 - r_4 t_2) + \alpha_{3,4}(r_3 t_4 - r_4 t_3) \end{pmatrix} \equiv 0 \mod \ell^{n-k}$$

and so the Weil pairing is trivial if and only if the condition in Step 4 holds. One can check that the same would hold for the other pairing.

Step 5: Output $(s_1, \ldots, s_4, r_1, \ldots, r_4, t_1, \ldots, t_4)$ as the secret scalars which will give the generators of the kernel:

$$Q_1 = \sum [s_i] P_i\,, \quad Q_2 = \sum [r_i] P_i\,, \quad Q_3 = \sum [t_i] P_i\,.$$

Remark. *Note the following:*

(a) *Step 2 performs discrete logarithm computations modulo a 2 and 3-smooth moduli and so is extremely efficient by using the Pohlig–Hellman algorithm [Gal12, §13.2].*

(b) *In Step 4, we note that the sampling of k cannot be a uniform choice if one likes to sample the entire keyspace uniformly. This is due to the distribution of the number of subgroups of different structures. One can use the equations in Theorem 5.4 and Proposition 5.5 to find a uniform distribution of the keyspace. A quick glimpse of the equations tells us that one should select $k = 0$ with much higher probability than $k = \lfloor n/2 \rfloor$ if one aims to sample the keyspace uniformly.*

(c) *In Step 4, we pick a random solution in the solution space for r_i and t_i. It can be shown that this ensures that the isotropic condition is upheld.*

---

[6] This will not be a uniformly random choice if one wants to sample the entire keyspace. See (c) in remark to follow.

Indeed, we have that

$$e_{\ell^n}(Q_1, Q_2) = \prod_{1 \le i,j \le 4} e_{\ell^n}([s_i]P_i, [r_j]P_j)$$

$$= e_{\ell^n}(P_1,P_1)^{s_1 r_1} e_{\ell^n}(P_1,P_2)^{s_1 r_2} e_{\ell^n}(P_1,P_3)^{s_1 r_3} e_{\ell^n}(P_1,P_4)^{s_1 r_4}$$
$$e_{\ell^n}(P_2,P_1)^{s_2 r_1} e_{\ell^n}(P_2,P_2)^{s_2 r_2} e_{\ell^n}(P_2,P_3)^{s_2 r_3} e_{\ell^n}(P_2,P_4)^{s_2 r_4}$$
$$e_{\ell^n}(P_3,P_1)^{s_3 r_1} e_{\ell^n}(P_3,P_2)^{s_3 r_2} e_{\ell^n}(P_3,P_3)^{s_3 r_3} e_{\ell^n}(P_3,P_4)^{s_3 r_4}$$
$$e_{\ell^n}(P_4,P_1)^{s_4 r_1} e_{\ell^n}(P_4,P_2)^{s_4 r_2} e_{\ell^n}(P_4,P_3)^{s_4 r_3} e_{\ell^n}(P_4,P_4)^{s_4 r_4}$$
$$= e_{\ell^n}(P_1,P_2)^{\alpha_{1,2} s_1 r_2} e_{\ell^n}(P_1,P_2)^{\alpha_{1,3} s_1 r_3} e_{\ell^n}(P_1,P_2)^{\alpha_{1,4} s_1 r_4}$$
$$e_{\ell^n}(P_1,P_2)^{\alpha_{2,1} s_2 r_1} e_{\ell^n}(P_1,P_2)^{\alpha_{2,3} s_2 r_3} e_{\ell^n}(P_1,P_2)^{\alpha_{2,4} s_2 r_4}$$
$$e_{\ell^n}(P_1,P_2)^{\alpha_{3,1} s_3 r_1} e_{\ell^n}(P_1,P_2)^{\alpha_{3,2} s_3 r_2} e_{\ell^n}(P_1,P_2)^{\alpha_{3,4} s_3 r_4}$$
$$e_{\ell^n}(P_1,P_2)^{\alpha_{4,1} s_4 r_1} e_{\ell^n}(P_1,P_2)^{\alpha_{4,2} s_4 r_2} e_{\ell^n}(P_1,P_2)^{\alpha_{4,3} s_4 r_3}$$
$$= e_{\ell^n}(P_1,P_2)^{\begin{pmatrix} \alpha_{1,2} s_1 r_2 + \alpha_{1,3} s_1 r_3 + \alpha_{1,4} s_1 r_4 \\ -\alpha_{1,2} s_2 r_1 + \alpha_{2,3} s_2 r_3 + \alpha_{2,4} s_2 r_4 \\ -\alpha_{1,3} s_3 r_1 - \alpha_{2,3} s_3 r_2 + \alpha_{3,4} s_3 r_4 \\ -\alpha_{1,4} s_4 r_1 - \alpha_{2,4} s_4 r_2 - \alpha_{3,4} s_4 r_3 \end{pmatrix}} .$$

Hence, for the Weil pairing to be trivial, we require that the exponent in the last expression to be congruent to zero modulo ℓ^n. Then noting that $s_i \equiv 0 \mod \ell^{n-k}$, we have the condition in Step 4. One can repeat the same argument for Q_1 and Q_3 to obtain the second expression. Note that we have that $e_{\ell^n}(Q_2, Q_3)$ is already trivial due to the orders of Q_2 and Q_3 being ℓ^{n-k} and ℓ^k respectively.

2.3.1 Isogeny Algorithms

Computing an ℓ-isogeny between elliptic curves can be done with a complexity of $O(\ell)$ field operations. The general method to compute the codomains of this isogeny or to map points under the isogeny is to use Vélu's formula. Here, we will present algorithms for computing the codomains of $(2,2)$ and $(3,3)$-isogenies and show how we can map subgroups under these isogenies. The speed-ups come from the use of simpler representations in the computation: the use of hyperelliptic curves in the (2,2) case and the use of Kummer surfaces in the (3,3) case.

Richelot Isogenies

We will use Richelot isogenies [Ric36, Ric37] to perform $(2,2)$-isogenies as is standard in the literature. Richelot isogenies are relatively well-understood and have been implemented in various computational algebra programs. In particular, the algorithm to find the equation of the codomain curve is studied in [Smi05]. The algorithm to compute the images of points under a degree 2 isogeny is classical, and a good exposition can be found in [Smi05, Ch. 8], and also in [CF96, §9] and [BD11, §4].

Let H be a hyperelliptic curve over a field k and let $K \subset J_H[2]$ be a maximal 2-isotropic subgroup, then we can define the isogeny $\phi : J_H \to A$, where A is a PPAS, and $\ker \phi = K$. The key step of Richelot isogenies is the parametrisation of subgroups of $J_H[2]$ by quadratic splittings of the hyperelliptic polynomial. The codomain of the Richelot isogeny can be found by using the Richelot operators on the quadratic splittings. If A is the Jacobian of a hyperelliptic curve, say H', then for all H and H', we will see later that there exists a subvariety Γ of $H \times H'$ such that Γ is a *Richelot correspondence* on $H \times H'$ (cf. [Smi05, §8.4]). The isogeny can then be computed by pulling back divisor classes (Jacobian points) on H to Γ before sending it back down to H'.

More concretely, we let H have the form

$$H : y^2 = h_6 x^6 + h_5 x^5 + \cdots + h_0 .$$

If we denote the Weierstrass points of H as $(a_i, 0)$, where $i = 1, \ldots, 6$, then the points of order 2 in J_H are

$$\big[(a_i, 0) - (a_j, 0)\big]$$

for $1 \leq i < j \leq 6$, hence there are 15 of them. We denote $\left[(a_i, 0) - (a_j, 0)\right]$ by (i, j) for convenience. Then the following addition laws hold [Kuh88]:

$$(i, j) + (i, j) = 0, \quad (i, j) + (k, l) = (m, n), \quad (i, j) + (i, k) = (j, k),$$

for distinct $i, j, k, l, m, n \in \{1, 2, 3, 4, 5, 6\}$. Furthermore, the Weil pairings on these points are given by [Kuh88]

$$e_2((i, j), (i, j)) = 1, \quad e_2((i, j), (k, l)) = 1, \quad e_2((i, j), (i, k)) = -1,$$

for distinct $i, j, k, l \in \{1, 2, 3, 4, 5, 6\}$. Hence the maximal 2-isotropic subgroups of $J_H[2]$ are given by

$$\{0, (i, j), (k, l), (m, n)\},$$

i.e. disjoint pairs of the Weierstrass points. Now, notice that an element in the maximal 2-isotropic subgroup is represented by two Weierstrass points, and so can be written as a quadratic polynomial with the abscissa of the two Weierstrass points as its roots. Hence, we can represent H as

$$H : y^2 = F_1(x) F_2(x) F_3(x)$$

and we say that $\{F_1, F_2, F_3\}$ is a quadratic splitting of H. If the quadratic splittings are linearly dependent as polynomials, then we say that it is singular. In particular, the singularity can be measured by the determinant of the quadratic splitting which is given by

$$\delta = \det \begin{pmatrix} f_{0,1} & f_{1,1} & f_{2,1} \\ f_{0,2} & f_{1,2} & f_{2,2} \\ f_{0,3} & f_{1,3} & f_{2,3} \end{pmatrix},$$

where $f_{i,j}$ is the coefficient of x^i in F_j. It can be shown that singular quadratic splittings lead to Richelot isogenies that map into a product of elliptic curves [Smi05, §8.3].

Suppose that $\{F_1, F_2, F_3\}$ is a non-singular quadratic splitting of H, then applying the Richelot operator [Smi05, Def. 8.4.1] on $\{F_1, F_2, F_3\}$ would yield $\{G_1, G_2, G_3\}$, which are given by

$$G_1 = \delta^{-1} \det \begin{pmatrix} \frac{d}{dx} F_2(x) & \frac{d}{dx} F_3(x) \\ F_2(x) & F_3(x) \end{pmatrix},$$

$$G_2 = \delta^{-1} \det \begin{pmatrix} \frac{d}{dx} F_3(x) & \frac{d}{dx} F_1(x) \\ F_3(x) & F_1(x) \end{pmatrix},$$

$$G_3 = \delta^{-1} \det \begin{pmatrix} \frac{d}{dx} F_1(x) & \frac{d}{dx} F_2(x) \\ F_1(x) & F_2(x) \end{pmatrix}.$$

This yields a twist of H' that we denote as H'_d which is defined by

$$H'_d : d\tilde{y}^2 = G_1(\tilde{x}) G_2(\tilde{x}) G_3(\tilde{x})$$

for some $d \in k^*$.

Now, we are ready to describe the mapping of points under the Richelot isogeny. The covering curve, as described at the start of the section, is given by the equations

$$\Gamma_d : \left\{ \begin{array}{l} F_1(x) G_1(\tilde{x}) + F_2(x) G_2(\tilde{x}) = 0 \\ F_1(x) G_1(\tilde{x})(x - \tilde{x}) = \sqrt{d} \tilde{y} y \\ F_2(x) G_2(\tilde{x})(x - \tilde{x}) = -\sqrt{d} \tilde{y} y \end{array} \right\}.$$

To map points on J_H to $J_{H'}$, we first express an element $D \in J_H$ as the divisor class given by $D = [P - Q]$ (or maybe $[P + Q]$), where $P, Q \in H$. We denote $P = (x(P), y(P))$ and $Q = (x(Q), y(Q))$, and we need to find the points P', P'' and Q', Q'' under the following morphism:

$$\mathrm{Jac}(H) \hookrightarrow \mathrm{Jac}(\Gamma) \twoheadrightarrow \mathrm{Jac}(H').$$

Take P as an example, we use the first equation of Γ to get the abscissas of P' and P'', denoted by $x(P')$ and $x(P'')$ respectively, by solving the quadratic equation in $\tilde{x}$

$$F_1(x(P)) G_1(\tilde{x}) + F_2(x(P)) G_2(\tilde{x}) = 0.$$

We then recover the ordinate by solving for $\tilde{y}$ in the following equations

$$F_1(x(P))G_1(x(P'))(x(P) - x(P')) = \sqrt{d}y(P)\tilde{y} \quad \text{and} \quad F_1(x(P))G_1(x(P''))(x(P) - x(P'')) = \sqrt{d}y(P)\tilde{y}.$$

This gives us $P' = (x(P'), y(P'))$ and $P'' = (x(P''), y(P''))$.

Repeating the computation for Q yields us Q' and Q''. We then recover the image of D under the Richelot map by computing

$$D' = [P' + P''] + [Q' + Q''].$$

(3,3)-Isogenies over the Kummer Surface

As for (3,3)-isogenies, we note that for the purposes of genus two isogeny cryptography, we do not need to map points under the isogeny but only need to map Kummer points under the isogeny since the Jacobian points that correspond to the Kummer points generate identical subgroups.

Given an abelian surface A, the *Kummer surface* is defined by $A/\langle \pm 1 \rangle$. This is a quartic surface in $\mathbb{P}^3$. Computations of (3,3)-isogenies on the Kummer surface was the object of study of [BFT14]. We can use the formulae[7] presented in [BFT14] to compute the images of Kummer points under the isogeny. This has also been noted by Costello in [Cos18].

The methods of [BFT14] parametrises the hyperelliptic curve into the form

$$y^2 = F(x) = G_1(x)^2 + \lambda_1 H_1(x)^3 = G_2(x)^2 + \lambda_2 H_2(x)^3,$$

where $H_i, G_i, F \in k[x]$ with H_i quadratic and co-prime, G_i cubic, and $\lambda_i \in k^*$. They showed in Lemmata 3,4 and 5 that this model of the hyperelliptic curve will always have a Jacobian that admits a (3,3)-isogeny. They further showed in Theorem 6 that H_i and G_i are parametrised by 3 variables named r, s, t, hence hyperelliptic curves whose Jacobians admit a (3,3)-isogeny are parametrised as C_{rst}. Once in the parametrised form, they provided formulae to compute (3,3)-isogenies.

To use their formulae for our purposes, we need to be able to express an arbitrary hyperelliptic curve with a Jacobian that admits a (3,3)-isogeny in the C_{rst} parametrisation. This procedure is detailed in [BFT14, §4] and we will only give brief summaries of the steps of procedure. Readers interested in the intuition and the veracity of the steps are to refer to [BFT14, §4].

Firstly, given a hyperelliptic curve C, and generators for the kernel of a (3,3)-isogeny, D_1, D_2, we obtain the parametrisation

$$y^2 = F(x) = G_1(x)^2 + \lambda_1 H_1(x)^3 = G_2(x)^2 + \lambda_2 H_2(x)^3,$$

where λ_i, G_i, H_i correspond to D_i. We do this by computing the cubic G_i using D_i; noting that $2D_i = -D_i$. So we express $D_1 = [(x_1, y_1) + (x_2, y_2) - \mathcal{O}]$, and require that (x_1, y_1) satisfy $y = G_1(x)$ with multiplicity two. This gives us 4 restrictions which suffices to compute G_1. We can then compute λ_1 and H_1 from F.

Now let $k[\alpha] = k[t]/(t^3 - \lambda_2/\lambda_1)$, and we have that

$$H_1 - \alpha H_2 = LM$$

for some $L, M \in k[\alpha, x]$, and that for some $c \in k^*$,

$$
\begin{aligned}
G_2 - G_1 &= \tfrac{1}{c} \operatorname{Nm}(M), \\
G_2 + G_1 &= c\lambda_1 \operatorname{Nm}(M),
\end{aligned}
\quad \text{and} \quad
\begin{aligned}
G_1 &= \tfrac{1}{2}(c\lambda_1 \operatorname{Nm}(L) - \tfrac{1}{c}\operatorname{Nm}(M)), \\
G_2 &= \tfrac{1}{2}(c\lambda_1 \operatorname{Nm}(L) + \tfrac{1}{c}\operatorname{Nm}(M)).
\end{aligned}
$$

Furthermore, one can force L and M to take the following forms

$$L = x - u\alpha, \text{ and } M = (c_0 + c_1\alpha + c_2\alpha^2)x - (m_0 + m_1\alpha + m_2\alpha^2)$$

after a series of linear transformations.

Note that

$$(cy)^2 = (cG_1)^2 + (c^2\lambda_1)H_1^3 = (cG_2)^2 + (c^2\lambda_2)H_2^3,$$

hence we need that $c = 1$.

---

[7]The files containing the formulae can be found in `http://www.cecm.sfu.ca/~nbruin/c3xc3/`.

We remark that the series of transformations in [BFT14, §4] is incomplete. A last transformation is necessary as c has shifted away from 1 due to prior transformations. At that stage, we have the following:

$$(s, t, c_0, c_1, c_2, m_0, m_1, m_2, u) = (s', t', 1, -1, 0, -r', 0, 1, 1).$$

We need one last transformation

$$y \mapsto (4/\lambda_1)^2 y$$

and set

$$s = \lambda_1/4, \quad r = \text{Coefficient of } x \text{ in } H_1, \quad t = \text{Coefficient of } 1 \text{ in } H_1$$

to get the (r, s, t)-parametrisation of [BFT14, Thm. 6].

The key to forming the cubic formula which maps Kummer points to Kummer points under the $(3, 3)$-isogeny lies in the bi-quadratic forms on the Kummer surface [CF96, pg. 23]. Given the generators, T_1, T_2, of the maximal 3-isotropic subgroup of $J_H[3]$, the authors found two cubic forms which are each invariant under translation by T_1 and T_2 respectively. The cubic forms generate spaces of dimension 8 and intersect in dimension 4. Since Kummer surfaces can be explicitly described as quartic surfaces in $\mathbb{P}^3$, the intersection of the cubic forms fully describes the Kummer surfaces.

2.3.2 Genus Two SIDH

We will present the key exchange protocol in genus two. The astute reader will see that all the steps carry over from the scheme presented in §2.1.

Set-up

Pick a prime p of the form $p = 2^n \cdot 3^m \cdot f - 1$ where $2^n \approx 3^m$. Now, we pick a hyperelliptic curve H using the methods at the start of this section. We then generate the bases $\{P_1, P_2, P_3, P_4\}$ and $\{Q_1, Q_2, Q_3, Q_4\}$ which generate $J_H[2^n]$ and $J_H[3^m]$ respectively. The points P_i, Q_i are defined over $\mathbb{F}_{p^2}$ and are guaranteed to exist by the choice of the prime.

First Round

Alice chooses her secret scalars $(a_i)_{i=1,\ldots,12}$ using the steps outlined in at the start of this section and computes the isogeny $\phi_A : J_H \to J_A$ with kernel given by

$$\left\langle \sum_{i=1}^{4} [a_i] P_i, \ \sum_{i=5}^{8} [a_i] P_{i-4}, \ \sum_{i=9}^{12} [a_i] P_{i-8} \right\rangle.$$

She also needs to compute the points $\phi_A(Q_i)$ for $i = 1, 2, 3, 4$. She sends the tuple

$$(J_A, \phi_A(Q_1), \phi_A(Q_2), \phi_A(Q_3), \phi_A(Q_4))$$

to Bob.

At the same time, Bob chooses his secret scalars $(b_i)_{i=1,\ldots,12}$ and computes the isogeny $\phi_B : J_H \to J_B$ which has the kernel

$$\left\langle \sum_{i=1}^{4} [b_i] Q_i, \ \sum_{i=5}^{8} [b_i] Q_{i-4}, \ \sum_{i=9}^{12} [b_i] Q_{i-8} \right\rangle.$$

He computes the points $\phi_B(P_i)$ for $i = 1, 2, 3, 4$, and sends the tuple

$$(J_B, \phi_B(P_1), \phi_B(P_2), \phi_B(P_3), \phi_B(P_4))$$

to Alice.

Second Round

Alice will receive Bob's tuple and proceeds with computing the subgroup

$$\left\langle \sum_{i=1}^{4} [a_i]\phi_B(P_i), \ \sum_{i=5}^{8} [a_i]\phi_B(P_{i-4}), \ \sum_{i=9}^{12} [a_i]\phi_B(P_{i-8}) \right\rangle .$$

This is the kernel of a $(2^n, 2^{n-k}, 2^k)$-isogeny $\phi'_A : J_B \to J_{BA}$. Bob will perform a similar computation and arrive at the PPSSAS J_{AB}. But since

$$J_{AB} = J_A/\phi_A(K_B) \cong J_H/\langle K_A, K_B \rangle \cong J_B/\phi_B(K_A) = J_{BA} ,$$

they can then use the G_2-invariants of the hyperelliptic curves associated to J_{AB} and J_{BA} as their shared secret. In the unlikely event that J_{AB} is a product of elliptic curves (probability is $O(1/p)$ [CDS19, Thm. 1]), they would have to restart the protocol with a different selection of scalars.

Remark. *The method in [BFT14] allows us to find $\pm\phi_B(P_i)$ since we are working over the Kummer surface. However, we need the map*

$$(P_1, P_2, P_3, P_4) \mapsto (\phi_B(P_1), \phi_B(P_2), \phi_B(P_3), \phi_B(P_4))$$

or

$$(P_1, P_2, P_3, P_4) \mapsto (-\phi_B(P_1), -\phi_B(P_2), -\phi_B(P_3), -\phi_B(P_4))$$

to ensure that the subgroup generated by Alice in the second round is isotropic.
   To fix this problem, one can check if

$$e_{2^n}(\phi_B(P_i), \phi_B(P_j)) = e_{2^n}(P_i, P_j)^{3^m}$$

for all $1 \leq i < j \leq 4$ and negate the $\phi_B(P_i)$'s accordingly. The goal is to ensure that all the $\phi_B(P_i)$ have the same sign, as Alice's secret scalars will return the correct kernel if the signs of $\phi_B(P_i)$ are the same. Having different parities would in effect change Alice's secret scalars which would result in the wrong kernel and the wrong isogeny.
   The computation of Bob's isogeny will require the factorisation of his large isogeny into smaller isogenies. In doing so, Bob will be required to map his kernels through the smaller isogenies, hence he will encounter a similar problem with the parities of the points. To combat this issue, he would have to add a point of order 4 to the kernel points to track the parity of the points after the map. Since he has to map Alice's points in the first round, the overheads for this procedure is minimal. However, this would slow down the second round.

Chapter 3

Attack on Points

Recall that the SIDH key exchange protocol in §2.1 required both parties to send auxiliary points to complete the protocol to compose isogenies in a meaningful way. The Diffie–Hellman key exchange based on abelian groups did not require additional information to be sent; hence the security assumptions do not carry over neatly into SIDH.

In fact, we will see that the auxiliary information sent in the protocol is the target of two attacks: an adaptive attack and a fault attack.

Adaptive attacks are a standard type of attack on cryptosystems that use a static private key. These first arose in the setting of protocols based on the discrete logarithm problem, where a user can be treated as an oracle that takes as input a group element g and returns g^a for some long-term secret value a. A first kind of attack is the "small subgroup" attack of Lim and Lee [LL97]. Here a group element g of small order ℓ is sent, so that on receipt of the value g^a, one can do a search and learn $a \pmod{\ell}$. Similar ideas have been used based on "invalid curve" attacks, which involve providing a point that lies in a different group altogether (see Ciet and Joye [CJ05]). In the context of the isogeny cryptosystem, if Alice has a fixed key (a_1, a_2) then a dishonest Bob can send her (E, P, Q) and then Alice will compute an isogeny $\phi : E \to E_0$ with kernel $\langle [a_1]P + [a_2]Q \rangle$. The idea is to try to learn something about Alice's secret key (a_1, a_2) using knowledge of E_0. The possibility of such attacks is mentioned in [KLM$^+$15] and [CLN16], but neither paper presented full details of them. In §3.1, we describe a general active attack against the static-key variant of the protocol. Our attack allows to recover the whole static key with the minimum number of queries (when $\ell = 2$) and negligible computation. Our attack is not prevented by any of the validation techniques introduced in [CLN16], nor by our stronger validation technique using pairings (cf. remark in §3.1.4), nor by key compression techniques [AJK$^+$16, CJL$^+$17, ZJP$^+$18]. Our attack is prevented by the method in [KLM$^+$15], but this adds significant cost to the running time of the system.

Fault attacks exploit the leakage of sensitive information when the implementation operates under unexpected circumstances. Biehl, Meyer and Müller [BMM00] extended fault attacks on RSA cryptosystems to systems using elliptic curves. Ciet and Joye [CJ05] then refined the methods and made the attack more practical. The key insight in both papers was the absence of the a_6 elliptic curve parameter in the scalar multiplication computation. The fault changed the base point P to some P'. This meant that the output point $[\lambda]P'$, where λ is the secret, might be in a group where solving the elliptic curve discrete logarithm problem was feasible, hence allowing for the recovery of some information about λ. In §3.2, we will examine the effects of changing a point P to some random P' and attempt to recover the secret, which in this case is an isogeny ϕ. The attack would be able to recover the entire secret ϕ from a single output $\phi(P')$ with high probability. This compares well against the fault attack presented in [CJ05] where a single successful perturbation only reveals partial information of the secret. We will see the fault attack in the context of several signature schemes and key exchange protocols. The attack would work against the countermeasure proposed by Kirkwood et al. [KLM$^+$15] which is based on the Fujisaki–Okamoto transform. The main observation that underlies the attack is that users should never reveal the image of random points under the secret isogeny.

There is another attack on the auxiliary information of the key exchange protocol by Petit [Pet17]. The attack uses the auxiliary information to compute the endomorphism rings of the supersingular elliptic curves. This results in a break of the cryptosystem as we will see in §4.3.2. It should be noted

that this attack works with an "unbalanced" version of the Jao–De Feo cryptosystem where one torsion group is significantly larger than the other one. We will not be covering this attack.

3.1   Adaptive Attack

In this section, we will assume that Alice is using a static key (a_1, a_2), and that a dishonest user is playing the role of Bob and trying to learn her key. Our discussion is entirely about Alice's key and points in $E[2^n]$, but it should be clear that the same methods would work for points in $E[\ell^m]$ for any small prime ℓ (see the remark in §3.1.3 for further discussion).

There are two attack models that can be defined in terms of access to an oracle O:

(a) $O(E, R, S) = E/\langle [a_1]R + [a_2]S\rangle$.

If the scheme under attack is the key exchange scheme, this corresponds to Alice taking Bob's protocol message, completing her side of the protocol, and outputting the shared key. In the encryption protocol, this would correspond to an encryption $c = m \oplus j(E_{AB})$ *without the hash function* and Alice decrypting Bob's ciphertext and returning the plaintext m.

(b) $O(E, R, S, E')$ which returns 1 if $j(E') = j(E/\langle [a_1]R + [a_2]S\rangle)$ and 0 otherwise.

In the key exchange setting, this corresponds to Alice taking Bob's protocol message, completing her side of the protocol, and then performing some operations using the shared key that return an error message if the shared key is not the same as the j-invariant provided (e.g., the protocol involves verifying a MAC corresponding to a key derived from the session key).

In the encryption scenario [JD11, §3.2], this would correspond to Bob having access to a decryption oracle for Alice. By choosing a random ciphertext c Bob could ask for a decryption of (E_B, R, S, c) and get m such that $c = m \oplus H_k(j(E_{AB}))$. Bob can then check whether or not $c \oplus m = H_k(j(E'))$. Hence a decryption oracle for the encryption scheme gives an oracle O of this type.

Our attacks can be mounted in both models. To emphasise their power we explain them in the context of the second, weaker, model.

Before we present the attack, we make an observation through a lemma.

Lemma 3.1. *Let $P, Q \in E[2^n]$ be linearly independent generators of $E[2^n]$. Then for all $(a_1, a_2) \in \mathbb{Z}^2$ (not simultaneously even), we have that $\langle [a_1]P + [a_2]Q\rangle = \langle P + [\alpha]Q\rangle$ or $\langle [a_1]P + [a_2]Q\rangle = \langle [\alpha]P + Q\rangle$ for some $\alpha \in \mathbb{Z}$.*

Proof. If a_1 is odd, then it is invertible modulo the order of the group, so let $\theta \equiv a_1^{-1} \pmod{2^n}$, then θ must be odd, hence

$$\langle [a_1]P_A + [a_2]Q_A\rangle = \langle [\theta a_1]P_A + [\theta a_2]Q_A\rangle = \langle P_A + [\alpha]Q_A\rangle,$$

where the first equality stems from the fact that θ is co-prime to the order of the generator, and the last equality is obtained by setting $\alpha = \theta a_2$.

If a_1 is even, then a_2 must be odd, and repeating the procedure gives $(\alpha, 1)$. $\square$

This result tells us that there is no loss of generality for Alice to restrict her secret key to be $(1, \alpha)$ or $(\alpha, 1)$. However, even if Alice does not employ such a simplification, the result also tells us that there is no loss of generality for an attacker to assume the secret key is of one of these two forms. We will call this the normalised form.

3.1.1   First Step of the Attack

From Lemma 3.1, we may assume that the private key is normalised. In the following exposition, we will assume that the normalisation is $(1, \alpha)$. The case where we have $(\alpha', 1)$ where α' is even is performed in exactly the same way with some tweaks. Note that if α' is odd then it can be converted to the $(1, \alpha)$ case, so we may assume α' is even in the second case.

To differentiate between $(1, \alpha)$ and $(\alpha', 1)$ an attacker honestly generates Bob's ephemeral values $(E_B, R = \phi_B(P_A), S = \phi_B(Q_A))$ and follows the protocol to compute the resulting key E_{AB}. Then the

attacker sends $(E_B, R, S + [2^{n-1}]R)$ to Alice and tests the resulting j-invariant. Expressing this in terms of the oracle access: the attacker queries an oracle of the second type on $(E_B, R, S + [2^{n-1}]R, E_{AB})$. If the oracle returns 1 then the curve $E_B/\langle[a_1]R + [a_2](S + [2^{n-1}]R)\rangle$ is isomorphic to E_{AB} and so $\langle[a_1]R + [a_2](S + [2^{n-1}]R)\rangle = \langle[a_1]R + [a_2]S\rangle$. Hence, by the following lemma, a_2 is even and we are in the first case. If the oracle returns 0 then a_2 is odd.

Lemma 3.2. *Let $R, S \in E[2^n]$ be linearly independent points of order 2^n and let $a_1, a_2 \in \mathbb{Z}$. Then*

$$\langle[a_1]R + [a_2](S + [2^{n-1}]R)\rangle = \langle[a_1]R + [a_2]S\rangle$$

if and only if a_2 is even.

Proof. If a_2 is even then $[a_2][2^{n-1}]R = 0$ and so the result follows. Conversely, if the two groups are equal then there is some $\lambda \in \mathbb{Z}_{2^n}^*$ such that

$$\lambda([a_1]R + [a_2](S + [2^{n-1}]R)) = [a_1]R + [a_2]S.$$

Since the points are independent we have $\lambda a_2 = a_2$ and so $\lambda = 1$. Hence, since S has order 2^n, we have $a_2 2^{n-1} \equiv 0 \pmod{2^n}$ and a_2 is even. $\qquad\square$

Note that the Weil pairing

$$e_{2^n}(R, S + [2^{n-1}]R) = e_{2^n}(R, S) = e_{2^n}(P_A, Q_A)^{3^m}$$

and so the attack is not detectable using pairings (c.f. Remark in §3.1.4).

Similarly one can call the oracle on $(E_B, R + [2^{n-1}]S, S, E_{AB})$. The oracle returns 1 if and only if a_1 is even. Hence, we can determine which of the two cases we are in and determine if α is even or odd. Having recovered a single bit of α, we will now explain how to use similar ideas to recover the rest of the bits of α.

3.1.2   Continuing the Attack

We now assume that Alice's static key is of the form $(1, \alpha)$ and we write

$$\alpha = \alpha_0 + 2^1\alpha_1 + 2^2\alpha_2 + \cdots + 2^{n-1}\alpha_{n-1}.$$

The attacker will learn one bit of α for each query of the oracle. Algorithm 1 gives pseudo-code for the attack.

We now give some explanation and present the derivation of the algorithm. Suppose an attacker has recovered the first i bits of α, so that

$$\alpha = K_i + 2^i\alpha_i + 2^{i+1}\alpha',$$

where K_i is known but $\alpha_i \in \{0, 1\}$ and $\alpha' \in \mathbb{Z}$ are not known.

The attacker generates $E_B, R = \phi_B(P_A), S = \phi_B(Q_A)$ and E_{AB} as in the protocol. To recover α_i, the attacker will choose suitable integers a, b, c, d and query the oracle on

$$(E_B, [a]R + [b]S, \ [c]R + [d]S, E_{AB}).$$

The integers a, b, c, and d will be chosen to satisfy the following conditions:

(1) If $\alpha_i = 0$, then $\langle[a + \alpha c]R + [b + \alpha d]S\rangle = \langle R + [\alpha]S\rangle$.

(2) If $\alpha_i = 1$, then $\langle[a + \alpha c]R + [b + \alpha d]S\rangle \neq \langle R + [\alpha]S\rangle$.

(3) $[a]R + [b]S$ and $[c]R + [d]S$ both have order 2^n.

(4) The Weil pairing $e_{2^n}([a]R + [b]S, [c]R + [d]S)$ must be equal to

$$e_{2^n}(\phi_B(P_A), \phi_B(Q_A)) = e_{2^n}(P_A, Q_A)^{\deg\phi_B} = e_{2^n}(P_A, Q_A)^{3^m}.$$

21

The first two conditions help us distinguish the bit α_i and the latter two prevent the attack from being detected via order checking and Weil pairing validation checks respectively.

Consider the following integers:

$$a_i = 1, \quad b_i = -2^{n-i-1}K_i,$$
$$c_i = 0, \quad d_i = 1 + 2^{n-i-1}.$$

One can verify that they satisfy the third condition. To satisfy the fourth condition we need to use a scaling by θ that we will discuss later.

To show that the first two conditions are satisfied, note that $\langle [a]R + [b]S + [\alpha]([c]R + [d]S) \rangle$ is equal to

$$\langle R - [2^{n-i-1}K_i]S + [\alpha][1 + 2^{n-i-1}]S \rangle$$
$$= \langle R + [\alpha]S + [-2^{n-i-1}K_i + 2^{n-i-1}(K_i + 2^i\alpha_i + 2^{i+1}\alpha')]S \rangle$$
$$= \langle R + [\alpha]S + [\alpha_i 2^{n-1}]S \rangle$$
$$= \begin{cases} \langle R + [\alpha]S \rangle & \text{if } \alpha_i = 0, \\ \langle R + [\alpha]S + [2^{n-1}]S \rangle & \text{if } \alpha_i = 1. \end{cases}$$

By the following lemma, these two subgroups are different. Hence the response of the oracle tells us α_i.

Lemma 3.3. *Let R and S be linearly independent elements of the group $E[2^n]$ with full order, then the subgroups*

$$\langle R + [\alpha]S + [2^{n-1}]S \rangle \quad and \quad \langle R + [\alpha]S \rangle$$

are different.

Proof. The proof is very similar to the proof of Lemma 3.2. The subgroups have order 2^n, since R has order 2^n, and R and S are linearly independent. Then if the subgroups are the same, we must have some λ such that

$$[\lambda]R + [\lambda\alpha]S = R + [\alpha]S + [2^{n-1}]S.$$

By the linear independence of R and S, we can compare coefficients and conclude that $\lambda = 1$, and that $[2^{n-1}]S = \mathcal{O}$, which implies that S has order a factor of 2^{n-1}, which is a contradiction. $\square$

Algorithm 1: Adaptive attack using oracle $O(E, R, S, E')$.

Data: n, E, P_A, Q_A, P_B, Q_B, E_A, $\phi_A(P_B)$, $\phi_A(Q_B)$
Result: α

1 Set $K_0 \leftarrow 0$;
2 **for** $i \leftarrow 0$ **to** $n - 3$ **do**
3   Set $\alpha_i \leftarrow 0$;
4   Choose random (b_1, b_2);
5   Set $G_B \leftarrow \langle [b_1]P_B + [b_2]Q_B \rangle$;
6   Set $E_B \leftarrow E/G_B$ and let $\phi_B : E \rightarrow E_B$ be the isogeny with kernel G_B;
7   Set $(R, S) \leftarrow (\phi_B(P_A), \phi_B(Q_A))$;
8   Set $E_{AB} \leftarrow E_A/\langle [b_1]\phi_A(P_B) + [b_2]\phi_A(Q_B) \rangle$;
9   Set $\theta \leftarrow \sqrt{(1 + 2^{n-i-1})^{-1}} \pmod{2^n}$;
10   Query the oracle on $\left(E_B, [\theta](R - [2^{n-i-1}K_i]S), [\theta][1 + 2^{n-i-1}]S, E_{AB}\right)$;
11   **if** *Response is false* **then** $\alpha_i = 1$;
12   Set $K_{i+1} \leftarrow K_i + 2^i\alpha_i$;
13 **end**
14 Brute force $\alpha_{n-2}, \alpha_{n-1}$ using E and E_A and $K_{n-2} = \alpha \pmod{2^{n-2}}$ to find α (this requires no oracle calls);
15 Return α;

Finally, we address the fourth condition. We need that

$$e_{2^n}([a]R + [b]S, [c]R + [d]S) = e_{2^n}(R, S)^{ad-bc} = e_{2^n}(P_A, Q_A)^{3^m}.$$

The idea is that we can mask the points chosen from the attack above to satisfy the fourth condition. Recall that the points we wish to send to Alice are

$$(R', S') = (R - [2^{n-i-1}K_i]S, [1 + 2^{n-i-1}]S).$$

Computing the Weil pairing of the two points, we have

$$\begin{aligned}
&e_{2^n}(R', S')\\
&= e_{2^n}(R - [K_i 2^{n-i-1}]S, [1 + 2^{n-i-1}]S)\\
&= e_{2^n}(R, [1 + 2^{n-i-1}]S) \cdot e_{2^n}(-[K_i 2^{n-i-1}]S, [1 + 2^{n-i-1}]S)\\
&= e_{2^n}(R, S)^{1+2^{n-i-1}},
\end{aligned}$$

which is not the correct value. So we choose θ such that

$$e_{2^n}(\theta R', \theta S') = e_{2^n}(R, S)^{\theta^2(1+2^{n-i-1})} = e_{2^n}(P_A, Q_A)^{3^m} = e_{2^n}(R, S).$$

Note that $\langle [\theta]R' + [\alpha][\theta]S' \rangle = \langle [\theta](R' + [\alpha]S') \rangle = \langle R' + [\alpha]S' \rangle$ as long as θ is co-prime to the order 2^n. Hence we need θ to be the square root of $1 + 2^{n-i-1}$ modulo 2^n. The following well-known lemma by Gauss shows that such a square root exists as long as $n - i - 1 \geq 3$. Note that θ will be odd, as required.

Lemma 3.4. *If a is an odd number and $m = 8$, 16, or some higher power of 2, then a is a quadratic residue modulo m if and only if $a \equiv 1 \pmod 8$.*

The condition $n - i - 1 \geq 3$ means we may not be able to launch the attack in an undetected way for the last two bits. This is why we use a brute force method to determine these bits. To brute force the remaining bits, one uses the recovered bits to compute an isogeny from E to obtain an intermediate curve E' which is 4-isogenous to E_A, and finds the path from E' to E_A.

The attack in the case $(\alpha', 1)$ follows by swapping the roles of R and S.

3.1.3 Complexity of the Attack

The attack requires fewer than $n \approx \frac{1}{2}\log_2(p)$ interactions with Alice. This seems close to optimal for the second attack model, where the attacker only gets one bit of information at each query. One can reduce the number of queries by doing more computation (increasing the range of the brute-force search).

We now consider the attack in the context of [KLM+15] and [CLN16], i.e. in the presence of validations using Weil pairings and checking of orders, and the Fujisaki–Okamoto type approach. Due to our third and fourth conditions, the attack passes the validation steps in [CLN16], and even the stronger check of taking the degree of the isogeny into account as mentioned in the remark in §3.1.4.

The approach in [KLM+15] would be able to detect the attack. This is because the auxiliary points sent to Alice in the attack are not the correct values generated in an honest protocol run.

Remark. *We now say a few words about attacking odd prime power isogenies. Let ℓ be an odd prime such that $\ell^n \mid (p + 1)$ and $E[\ell^n] \subset E(\mathbb{F}_{p^2})$. Let P_A, Q_A be generators of $E[\ell^n]$. Alice would compute an ℓ^n-isogeny with kernel $\langle [a_1]P_A + [a_2]Q_A \rangle$ and a dishonest user Bob is trying to learn her key a_1, a_2, where a_1 and a_2 are not simultaneously divisible by ℓ. As above, we take Alice's secret key to be $(1, \alpha)$.*

The obvious generalisation for this attack is to set $R = \phi_B(P_A)$ and $S = \phi_B(Q_A)$ and to send Alice points

$$(R - [x\ell^{n-i-1}]S, [1 + \ell^{n-i-1}]S).$$

In her computation for the subgroup, Alice would compute

$$\langle R + [\alpha]S + [\ell^{n-i-1}][\alpha - x]S \rangle.$$

Since we want to compare this subgroup against $\langle R + [\alpha]S \rangle$, we need

$$(\ell^{n-i-1})(\alpha - x) \equiv 0 \pmod{\ell^n}$$

23

to ensure the subgroups computed are the same. Hence for each coefficient of a power of ℓ in the ℓ-expansion of α, we will need at most $\ell - 1$ queries to recover it.

For $\ell = 3$ this is close to optimal (at most two queries compared to $\log_2(3)$ queries), but for primes $\ell \geq 5$ this seems not optimal since one would hope that given an oracle that returns one bit of information one could learn the value with at most $\lceil \log_2(\ell) \rceil$ queries.

3.1.4 Validations and Countermeasures

The concept of "validation" is intended to prevent active attacks. In the case of protocols based on the DLP, the typical countermeasures check that g does lie in the correct group, and that the order of g is the correct value. In the context of supersingular isogeny cryptosystems the validation of (E, P, Q) should test that E really is a supersingular elliptic curve, that P and Q lie on the curve and have the correct order, and that P and Q are independent. Methods to do this are given in [CLN16].

In particular, §9 of [CLN16] presented some explicit validation steps. Their two requirements are: the points in the public key have full order and they are linearly independent. They use the Weil pairing of the two points to check independence.

Remark. *We now observe that the Weil pairing can be used to check a lot more than just independence. A standard fact is that if $\phi : E \to E'$ is an isogeny and if $P, Q \in E[N]$ then*

$$e_N(\phi(P), \phi(Q)) \; = \; e_N(P, Q)^{\deg(\phi)}$$

where the first Weil pairing is computed on E' and the second on E (for details see [Sil09, Prop. III.8.2] or [BSS05, Thm. IX.9]). This allows to validate not only that the points are independent but also that they are consistent with being the image of the correct points under an isogeny of the correct degree. Hence, a natural validation step for Alice to run in the Jao–De Feo scheme is to check

$$e_{2^n}(\phi_B(P_A), \phi_B(Q_A)) = e_{2^n}(P_A, Q_A)^{3^m}.$$

This will give her some assurance that the points $\phi_B(P_A), \phi_B(Q_A)$ provided by Bob are consistent with being the images of the correct points under an isogeny of the correct degree. However, as we have shown, this validation step is not sufficient to prevent all adaptive attacks. It will be necessary to use a much stronger protection, which we describe next.

Public key compression techniques introduced in [AJK$^+$16, CJL$^+$17, ZJP$^+$18] reduces the size of public keys at some cost in run time. We note that this use of these compression techniques does not change the oracle that we use in the adaptive attack and so the adaptive attack can still work in the presence of these compression techniques.

Kirkwood et al. introduced a method to secure the key exchange protocol of isogeny cryptosystems. This is based on the Fujisaki–Okamoto transform [FO13] which is also explained by Peikert [Pei14, §5.2] and Galbraith et al. [GPST16, §2.3]. The method allows for one party to validate the other, but for the ease of exposition, let us suppose that Alice is using a static secret and Bob needs to prove to her that he is performing the protocol correctly.

Bob would prove to Alice that he performed the protocol correctly by executing the key exchange, encrypting the random seed used to generate his private key and sending this ciphertext to Alice for her to verify that the random seed leads to the correct keys.

Applied to the Jao–De Feo protocol, we will briefly explain how Bob can prove to Alice that he has executed the protocol correctly. This is especially applicable if Alice is using a static key and Bob is potentially a malicious party.

(1) Alice computes and sends the public key $(E_A, \phi_A(P_B), \phi_A(Q_B))$.

(2) Bob receives Alice's public key $(E_A, \phi_A(P_B), \phi_A(Q_B))$.

(3) Bob obtains his random seed r_B from a random source and derives his private key using a key derivation function, KDF$_1$,

$$(b_1, b_2) = \text{KDF}_1(r_B).$$

He uses the secret key to compute $G_B = \langle [b_1]P_B + [b_2]Q_B \rangle$, and uses the Vélu formula to compute ϕ_B and $E_B = E/G_B$.

(4) Bob derives the shared secret $SS_B = j(E_{AB})$ using his private key and Alice's public key. He then computes a session key (SK) and a validation key (VK) using a key derivation function, KDF_2,

$$SK \mid VK = \mathrm{KDF}_2(j(E_{AB})).$$

(5) Bob sends his public key $(E_B, \phi_B(P_A), \phi_B(Q_A))$ and $c_B = \mathrm{Enc}_{VK}(r_B \oplus SK)$ to Alice.

(6) Using her private key and Bob's public key, Alice computes the shared secret $SS_A = j(E'_{AB})$ and derives the session and validation keys SK' and VK'. She uses these to compute

$$r'_B = \mathrm{Dec}_{VK'}(c_B) \oplus SK'.$$

She then computes Bob's secret keys from r'_B and recomputes all of Bob's operations and compares $(E'_B, \phi'_B(P_A), \phi'_B(Q_A))$ with $(E_B, \phi_B(P_A), \phi_B(Q_A))$.

If they are equal, then Alice verifies that Bob has computed the protocol correctly and proceeds to use $SK' = SK$ for future communication with Bob. Else, the protocol terminates in a non-accepting state.

This validation method can be used for both the key exchange and the encryption protocols. It also compels one party to reveal the secret used and so requires a change in secret keys after each verification. This protocol is summarised in Fig. 3.1.
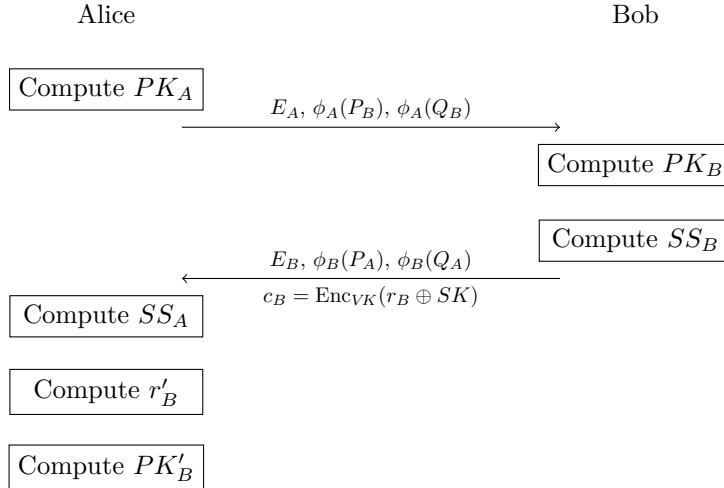

Figure 3.1: The Kirkwood et al. validation method for supersingular key exchange.

3.2 Fault Attack

Assume that the protocol under attack reveals the x-coordinate of the image of a point under the secret isogeny. The fault attack aims to force the implementation to output the image of a random point under the secret isogeny. This would allow the adversary to recover the secret. This section will illustrate how this is accomplished and review the different scenarios where the fault attack may be employed. The attack model that we are assuming supposes that an attacker is able to randomise certain values (random bit flips), in particular, the abscissa of a point to be mapped by the secret isogeny. This fault model is the same as the one proposed by Ciet and Joye [CJ05], where a fault will cause the x-coordinate of the point to be a random element of the field.

Our first observation is that computations do not involve the y-coordinate of the points. Given a curve E and a point P, a perturbation in the x-coordinate of P would result in another point P' on the same curve over a quadratic extension. Indeed, given any x, we recover the y-coordinate of P' by solving a quadratic equation which always has a solution in $\mathbb{F}_{p^2}$. In particular, any $x \in \mathbb{F}_{p^2}$ either corresponds to a point on E or a point on its quadratic twist E'. In [CLN16], for example, computations

do not distinguish between the curve E and its quadratic twist E', hence the isogeny will be evaluated correctly on any $x \in \mathbb{F}_{p^2}$. In a more general setting where the twists of the curves are treated separately, the faulted point will be on E with probability $1/2$ and on the twist with probability $1/2$. Hence the adversary may assume, after a series of faults, that a perturbed point will lie on E.

The perturbed point would be a random point on the curve. In §3.2.1, we will show how one recovers the secret isogeny given the image of the random point. This is not dissimilar to [JS14, Rem. 3.1], where Jao and Soukharev noted that a party should never disclose any information that allows an adversary to evaluate ϕ_A on $E[\ell_A^{e_A}]$. The method to recover ϕ_A given the image of a random point in $E[\ell_A^{e_A}]$ is mentioned in [DJP14, §5.1] and explained in detail in §3.2.1. In fact, we will show that a party should never reveal the image of random points under the secret isogeny.

3.2.1 Recovery of isogeny from image of random point

Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve where $p = \ell_A^{e_A} \cdot \ell_B^{e_B} \cdot f \pm 1$. Then with (P_A, Q_A), (P_B, Q_B), and (P_C, Q_C) being the generators of $E[\ell_A^{e_A}]$, $E[\ell_B^{e_B}]$, and $E[f]$ respectively, a random point $X \in E(\mathbb{F}_{p^2})$ takes the form

$$X = [u]P_A + [v]Q_A + [w]P_B + [x]Q_B + [y]P_C + [z]Q_C$$

for some $u, v, w, x, y, z \in \mathbb{Z}$.

Now suppose that we are given the image of X under the secret isogeny ϕ_A, we will show how one can use the knowledge of $\phi_A(X)$ to recover ϕ_A. Since ϕ_A is a group homomorphism and we know that X can be expressed as a linear combination of P_A, Q_A, P_B, Q_B, P_C, and Q_C, so we have

$$\begin{aligned}
\phi_A(X) &= \phi_A([u]P_A + [v]Q_A + [w]P_B + [x]Q_B + [y]P_C + [z]Q_C) \\
&= [u]\phi_A(P_A) + [v]\phi_A(Q_A) + [w]\phi_A(P_B) \\
&\quad + [x]\phi_A(Q_B) + [y]\phi_A(P_C) + [z]\phi_A(Q_C) \,.
\end{aligned}$$

Now our aim is to isolate a linear combination of $\phi_A(P_A)$ and $\phi_A(Q_A)$. To that end, we perform the operation

$$\begin{aligned}
[\ell_B^{e_B} \cdot f]\phi_A(X) &= [\ell_B^{e_B} \cdot f]([u]\phi_A(P_A) + [v]\phi_A(Q_A)) \\
&= [u']\phi_A(P_A) + [v']\phi_A(Q_A) \,,
\end{aligned}$$

and we find ourselves in the scenario described in [JS14, Remark 3.1] and [DJP14, §5.1].

Once we have $[u']\phi_A(P_A) + [v']\phi_A(Q_A)$, the subgroup generated by this point will help with the construction of the dual isogeny of ϕ_A hence recovering ϕ_A.

Lemma 3.5. *Let E_1 be a supersingular elliptic curve over $\mathbb{F}_{p^2}$, where $p = \ell_A^{e_A} \cdot \ell_B^{e_B} \cdot f \pm 1$. Suppose $\phi : E_1 \to E_2$ is an isogeny of degree $\ell_A^{e_A}$ with a cyclic kernel and let $\{P, Q\}$ be generators of $E_1[\ell_A^{e_A}]$. Then for any $X \in E_1[\ell_A^{e_A}]$, define $\psi : E_2 \to E'$ such that $\ker \psi = \langle \phi(X) \rangle$, then there exists some $\theta : E' \to E_1$ of degree ℓ_A^{ϵ}, $\epsilon \leq e_A$, such that*

$$\hat{\phi} = \theta \circ \psi \,.$$

Proof. Using Lemma 3.1, we may suppose that $\ker \phi = \langle P + [\alpha]Q \rangle$. Hence

$$\begin{aligned}
\phi(P) &= \phi(P) - \phi(P + [\alpha]Q) \\
&= -[\alpha]\phi(Q) \,.
\end{aligned}$$

Then expressing $X = [u]P + [v]Q$ for some u, v, we have

$$\langle \phi(X) \rangle = \langle [u]\phi(P) + [v]\phi(Q) \rangle = \langle [v - \alpha u]\phi(Q) \rangle = \langle [\ell_A^k]\phi(Q) \rangle \,,$$

where k is the ℓ_A-adic valuation of $(v - \alpha u)$.

Let $\psi : E_2 \to E'$ be an isogeny with kernel given by $\langle \phi(X) \rangle = \langle [\ell_A^k]\phi(Q) \rangle$. For any $Y \in E_1[\ell_A^{e_A}]$, we can write $Y = [r]P + [s]Q$ for some r, s.

If $k = 0$, then

$$\begin{aligned}
\psi \circ \phi(Y) &= \psi(\phi([r]P + [s]Q)) \\
&= \psi([s - r\alpha]\phi(Q)) \\
&= \mathcal{O} \,.
\end{aligned}$$

26

So it is clear that $E_1[\ell_A^{e_A}] \subseteq \ker(\psi \circ \phi)$. The reverse inclusion is obvious since $\ker(\psi \circ \phi)$ does not contain any non-trivial element of order co-prime to ℓ_A. So $\psi \circ \phi = [\ell_A^{e_A}]$, which implies, by the uniqueness of the dual isogeny, that $\psi = \hat{\phi}$, and $\theta : E_1 \to E_1$ is the identity isogeny.

If $k > 0$,

$$\psi \circ \phi(Y) = \psi(\phi([r]P + [s]Q))$$
$$= \psi([s - r\alpha]\phi(Q)).$$

Note that $\psi \circ \phi(Y)$ has order at most ℓ_A^k, since

$$[\ell_A^k]\psi \circ \phi(Y) = [s - r\alpha]\psi([\ell_A^k]\phi(Q)) = \mathcal{O}.$$

Now denote by $\gamma \in \mathbb{Z}_{\geq 0}$, the ℓ_A-adic valuation of $s - r\alpha$, then

$$\mathrm{ord}(\psi \circ \phi(Y)) = \mathrm{ord}(\psi([s - r\alpha]\phi(Q)))$$
$$= \ell_A^{k-\gamma}.$$

[Note that $\epsilon = k - \gamma$.]

So there exists Y such that $\gamma = 0$. Define $\theta : E' \to E_1$ such that $\ker \theta = \langle \psi \circ \phi(Y) \rangle$, then using the above argument we can see that $\theta \circ \psi = \hat{\phi}$. Furthermore, it is clear that $\deg \theta \leq \ell_A^{e_A}$. $\qquad \square$

The lemma tells us that given the image of a point in $E_1[\ell_A^{e_A}]$ under an $\ell_A^{e_A}$-isogeny, ϕ, we are able to find an isogeny ψ which is close to the dual isogeny of ϕ. To obtain the dual isogeny, one has to first recover θ. If ϵ is sufficiently small, one will be able to recover θ by brute force (just as we have done in §3.1). In fact, we will examine the size of ϵ in §3.2.2 and show that ϵ is small in most cases.

Hence we have the following algorithm to recover isogenies given the image of random points.

Algorithm 2: Recovering the dual isogeny after fault injection.

    Data: $\phi(X)$

    Result: $\hat{\phi}$

1 Set $\lambda \leftarrow \ell_B^{e_B} \cdot f$;

2 Set $T \leftarrow [\lambda]\phi(X)$;

3 Set $\psi : E_2 \to E'$ as the isogeny with kernel T;

4 **if** $\mathrm{ord}(T) = \ell_A^{e_A}$ **then**

5 | Return ψ;

6 **else**

7 | Brute force for θ;

8 Return $\theta \circ \psi$;

3.2.2   Analysis of attack

As seen in the proof of Lemma 3.5, to obtain the dual of the isogeny, we need k in the proof to be 0, or failing that, have ϵ small. But since ϵ is dependent on k, we will study k instead.

We start by fixing some $\alpha \in \mathbb{Z}/\ell_A^{e_A}\mathbb{Z}$ and suppose that u and v are selected randomly in $\mathbb{Z}/\ell_A^{e_A}\mathbb{Z}$, then we have

$$\Pr\left(\ell_A^n \text{ divides } (u - \alpha v)\right) = \frac{1}{\ell_A^n}.$$

Indeed, it is clear that we can treat $\rho = u - \alpha v$ as a single random variable, so this reduces to finding $\Pr(\ell_A^n \text{ divides } \rho)$, where ρ is randomly selected from $\mathbb{Z}/\ell_A^n\mathbb{Z}$. Since one in every ℓ_A^n elements is divisible by ℓ_A^n, we have the claim.

So $k = 0$ with probability $1 - \frac{1}{\ell_A}$. More generally, $k = \kappa$ with probability $\frac{\ell_A - 1}{\ell_A^{\kappa+1}}$. So we see that the isogeny ψ obtained from the procedure in §3.2 will be close to being the dual isogeny and brute forcing for θ is feasible. The number of neighbours one would have to check to recover θ is $3 \cdot 2^{k-1}$.

Lastly, we will address the issue of the faulted point $\phi(X)$ not having an order divisible by $\ell_A^{e_A}$. This would have the effect of decreasing the degree of ψ and so increase the degree of θ. But notice that we

can repeat the same analysis as the above to conclude that the degree of θ would be small with high probability.

Hence we have shown that Algorithm 2 has a high probability of recovering the secret isogeny.

Attack Models and their Feasibility

Interactive identification protocol and signature schemes

Signature schemes such as those proposed in [YAJ+17, JS14] employ a long term secret to generate a public key for signatures. The underlying feature of these schemes is the need to map ephemeral points under the secret isogeny. In the forthcoming exposition, we will not focus on the actual protocol, but will examine the step in the protocol which performs the map we are targeting.

The fault attack can capitalise on these protocols by injecting a fault just before the computation of the image of the ephemeral points to obtain the image of a random point under the secret isogeny. A fault would affect the validity of the signatures, however the invalidity of the signature would not prompt the signer to change their long-term secret. Hence the adversary would be able to break the signature scheme. We have to add that the compression of points [AJK+16] is an effective countermeasure that foils the attack and would also reduce the size of the responses. We will sketch the algorithm for point compression and examine their effects when examining countermeasures.

Using the notation found in [DJP14, §3.1], we see that to learn the prover's long-term secret S in the *interactive identification protocol*, the adversary needs to perturb the computation of the point $\phi(R)$. During the prover's computation, the adversary will introduce a perturbation immediately before the computation of $\phi(R)$. In particular the adversary can attempt to inject a fault into the fetching operation and cause a fault in R. This will cause the faulted point R' to be, with high probability, a point of full order. Successfully doing so would allow for the recovery of the secret isogeny ϕ. To obtain the output of the faulted point, the adversary needs the challenge bit to be 0 as described in [DJP14, §3.1]. This would happen 50% of the time and since identification schemes typically require a large number of passes, this must happen with high probability. The adversary could check the order of the points in the responses (if the challenge bit is 0) and the faulted point would have order larger than $\ell_A^{e_A}$. Using this information, the adversary would be able to use Algorithm 2 to recover S.

Due to its similarity to the identification protocol, to learn the signer's long-term secret S in the *digital signature scheme*, the steps the adversary takes are identical to the process above. The aim now is to inject a fault during the computation of $\phi(R_i)$ for some i's. A successful fault coinciding with the challenge bit being 0 would produce a point of order larger than $\ell_A^{e_A}$, so the adversary has to find that point in the signature by testing the orders of the points in the signature.

In the *undeniable signature protocol* the adversary will be able to learn the long term secret ϕ_A by inducing a fault in $\phi_M(P_C)$ before the computation of $\phi_{M,AM}(\phi_M(P_C))$ (using notation found in [JS14, §4.2]). Using $\phi_{M,AM}(X)$, the adversary would learn $\phi_{M,AM}$ and equivalently, $\phi_M(G_A)$. Since ϕ_M is computable from the message, the adversary would be able to recover G_A.

The signature scheme in [GPS17] does not output auxiliary points, hence this attack will not work on it.

Static key exchange protocol

Consider the static key exchange protocol described in §2.1. Suppose an adversary is trying to learn Alice's static secret isogeny and has the ability to cause a fault in Alice's computation. After introducing a fault in the computation of $\phi_A(P_B)$, Alice would then proceed to publish the public key tuple

$$(E_A, \phi_A(X), \phi_A(Q_B)).$$

The adversary will then be able to recover ϕ_A using Algorithm 2.

Notice that this would not be prevented by the validation method presented in §3.1.4. Since the validation method will only be able to detect misdemeanours carried out by Bob, it will not be able to prevent the fault attack. In particular, throughout the validation process the public key of Alice is only computed once and is never checked by the method. Hence the fault attack would not be detected by this validation and an adversary would be able to recover the secret isogeny as previously described.

Suppose that one party is using a static key in the key exchange protocol. An adversary would be able to recover the secret isogeny if the static public key is recomputed for each exchange. However, this is unlikely to happen since $\phi_A(P_B)$ and $\phi_A(Q_B)$ will be hardcoded for efficiency.

Now suppose that the adversary is attacking the key exchange protocol with ephemeral keys. If the secrets are not authenticated, the adversary would be able to compute $\phi_A(P_B)$, and send that in place of $\phi_A(X)$. This way, both parties would be able to derive the same shared secret. Since recovering ϕ_A from $\phi_A(X)$ can be done efficiently, and computing $\phi_A(P_B)$ is also efficient, performing the substitution before a time-out in the connection is very feasible. However, it should be noted that without authentication, it might be better to use a man-in-the-middle attack.

Remark. *The attack may also be implemented on the ephemeral key exchange protocol, but in both settings the attack would cause a failure to establish a shared secret key.*

Countermeasures

A simple countermeasure to this attack is to implement order checking before the publication of the auxiliary points. Another countermeasure that can be placed on the identification protocol and hence the signature scheme is the compression of the points $R, \phi(R)$ if the challenge bit is 0.

The compression protocol relies on an algorithm to compute canonical representations of elliptic curves in the isomorphism classes, and an algorithm to find canonical bases for an arbitrary torsion subgroup. Instead of sending $R, \phi(R)$, hence revealing $\phi(R)$ which the fault attack requires, the prover sends the scalars when representing $\phi(R)$ as a linear combination of the canonical basis. A faulted point would not be representable in this basis, and so an error should ensue. Hence the adversary would not be able to use the faulted auxiliary point to complete the attack.

Note that the compression of $\psi(S)$ will not be useful since the attack does not attack that point. Completing the fault attack on $\psi(S)$ would yield the isogeny ψ, which is not the long-term secret, hence does not achieve the goal of the attack.

Chapter 4

Endomorphism Rings of Elliptic Curves

Let X be an abelian variety, then a homomorphism $\phi : X \to X$ is called an *endomorphism*. We denote the set of all endomorphisms of X, together with the zero map (the map that sends all points to the identity), by $\mathrm{End}(X)$. We can endow the set $\mathrm{End}(X)$ with a ring structure by letting the zero and identity maps function as 0 and 1. Addition in this ring is given by point-wise addition and so the ring addition axioms are inherited from the addition axioms on the abelian variety. Multiplication in $\mathrm{End}(X)$ is given by composition of endomorphisms. We write $\mathrm{End}_k(X)$ to denote the endomorphism ring of endomorphisms of X that are defined over k, and write $\mathrm{End}(X)$ to denote the endomorphism ring with endomorphisms defined over the algebraic closure of the field X is defined over.

In this chapter, we will review the theory of the structure of the endomorphism rings of elliptic curves. A priori, it seems strange for us to turn our attention to the study of endomorphisms; especially when this is a class of isogenies that do not seem to "go anywhere". The link between the two can be seen when we show that the isogeny graph structure is determined by endomorphism rings. One can also take the view that when decomposing an endomorphism, one will obtain a sequence of isogenies which are not endomorphisms in general. Yet another connection between endomorphism rings and isogenies lies in the correspondence due to Deuring [Deu41] which establishes an equivalence of categories between the category of maximal orders of quaternion algebras with connecting ideals and the category of supersingular elliptic curves with isogenies. This equivalence of categories allows us to translate the problem from one category to another and will serve as the topic of §4.1. We will finally link it back to cryptography by showing a reduction of the SIDH problem in §2.1 to the computation of endomorphism rings.

4.1 Structure of the Endomorphism Rings of Elliptic Curves

Let E/k be an elliptic curve over k. The multiplication-by-n maps (together with the zero morphism) form a subring of $\mathrm{End}(E)$ isomorphic to $\mathbb{Z}$, and they lie in the centre of $\mathrm{End}(E)$. This begets the inclusion $\mathbb{Z} \hookrightarrow \mathrm{End}(E)$. Over a finite field, the Frobenius map is an endomorphism which is not a scalar multiplication, hence $\mathbb{Z} \subsetneq \mathrm{End}(E)$ in this instance.

Definition. We say that E/k has *complex multiplication* by O if $\mathrm{End}(E) \cong O \neq \mathbb{Z}$.

Hence, elliptic curves over a finite field will always have complex multiplication. We note that endomorphism rings are not isogeny invariant.

Example. Consider the supersingular elliptic curve $E_1 : y^2 = x^3 + x$ over $\mathbb{F}_{13^2}$, then it can be shown that $\mathrm{End}(E_1) = \mathbb{Z}\langle \pi_{E_1}, \phi \rangle$, where π_{E_1} is the Frobenius endomorphism on E_1, and $\phi(x, y) = (-x, \rho y)$, where $\rho^2 = -1$ in the field. There is an isomorphism of algebras between $\mathrm{End}(E_0) \otimes \mathbb{Q}$ and the quaternion algebra $B_{p,\infty}$ ramified at p and ∞. We give $B_{p,\infty}$ the canonical representation $B_{p,\infty} = \mathbb{Q}\langle \mathbf{i}, \mathbf{j} \rangle$, where $\mathbf{i}^2 = -1$, and $\mathbf{j}^2 = -p$, and $\mathbf{k} = \mathbf{ij}$. Then we have that $\mathrm{End}(E_1) = \mathbb{Z}\left\langle 1, \mathbf{i}, \frac{1+\mathbf{k}}{2}, \frac{\mathbf{i}+\mathbf{j}}{2} \right\rangle$, and the isomorphism of quaternion algebras is given by sending $(1, \mathbf{i}, \mathbf{j}, \mathbf{k}) \mapsto (1, \phi, \pi_{E_1}, \pi_{E_1} \circ \phi)$. One can then check that

$E_2 : y^2 = x^3 + 11x + 5$ is 2-isogenous to E_1, and also that $\text{End}(E_2)$ does not contain an element that would map into $\mathbf{i}$.

However, there is an isogeny invariant related to the endomorphism ring:

Definition. The *endomorphism algebra* of E is $\text{End}^0(E) = \text{End}(E) \otimes_{\mathbb{Z}} \mathbb{Q}$.

Theorem 4.1. *If E_1 and E_2 are isogenous elliptic curves, then $\text{End}^0(E_1) \cong \text{End}^0(E_2)$.*

Proof. Let $\phi : E_1 \to E_2$ be an isogeny between E_1 and E_2, and let $\alpha_1 \in \text{End}(E_1)$ be some endomorphism. Then we have that $\alpha_2 = \phi \circ \alpha_1 \circ \hat{\phi} \in \text{End}(E_2)$. We can obtain the degree of α_2 as

$$\deg(\alpha_2) = \deg(\phi \circ \alpha_1 \circ \hat{\phi}) = \deg(\phi)^2 \deg(\alpha_1)$$

and compute its trace as

$$\begin{aligned}
\text{tr}(\alpha_2) &= \alpha_2 + \widehat{\alpha}_2 \\
&= \phi \circ \alpha_1 \circ \hat{\phi} + \widehat{\phi \circ \alpha_1 \circ \hat{\phi}} \\
&= \phi \circ \alpha_1 \circ \hat{\phi} + \phi \circ \widehat{\alpha}_1 \circ \hat{\phi} \\
&= \phi \circ (\alpha_1 + \widehat{\alpha}_1) \circ \hat{\phi} \\
&= \deg(\phi) \, \text{tr}(\alpha_1) \, .
\end{aligned}$$

Hence $[\deg(\phi)] \circ \alpha_1$ has the same minimal polynomial as α_2 over $\mathbb{Q}$. This implies that $\mathbb{Q}(\alpha_1) \cong \mathbb{Q}(\alpha_2)$, and since this holds for any $\alpha_1 \in \text{End}(E_1)$, we have that $\text{End}^0(E_1)$ is isomorphic to a $\mathbb{Q}$-subalgebra of $\text{End}^0(E_2)$.

We can reverse the roles of α_1 and α_2 to obtain the other inclusion and conclude that the two endomorphism algebras are the same. $\qquad\square$

Furthermore, we can classify the endomorphism algebras of elliptic curves.

Theorem 4.2 ([Sil09, Cor. III.9.4]). *If E/k is an elliptic curve, then $\text{End}^0(E)$ is isomorphic to one of the following:*

(a) $\mathbb{Q}$;

(b) *an imaginary quadratic field;*

(c) *a definite quaternion algebra.*

Now that we have seen the structure of the endomorphism algebra, we will turn our attention to that of the endomorphism ring.

Theorem 4.3 ([Sil09, Cor. III.7.5]). *The endomorphism ring $\text{End}(E)$ is a free $\mathbb{Z}$-module of rank $1, 2$ or 4. In particular, it is equal to the dimension of $\text{End}^0(E)$ as a $\mathbb{Q}$-vectorspace.*

Furthermore, $\text{End}(E)$ is a subring of $\text{End}^0(E)$. Since an *order* is a lattice that is also a subring, we say that $\text{End}(E)$ is an order of $\text{End}^0(E)$. As a corollary to Theorem 4.2, we have that for any elliptic curve E over a finite field, $\text{End}(E)$ is an order of an imaginary quadratic field or an order of a definite quaternion algebra.

Recall that elliptic curves over finite fields fall into one of two categories: ordinary and supersingular. We say that an elliptic curve E is supersingular if $\text{tr}(\pi_E) \equiv 0 \pmod{p}$ and ordinary otherwise.

The following result then gives us the endomorphism algebra of ordinary elliptic curves.

Proposition 4.4. *If $E/\mathbb{F}_q$ is an ordinary elliptic curve, then $\text{End}^0(E) = \mathbb{Q}(\pi_E)$ is an imaginary quadratic field.*

Proof. Suppose for contradiction that $\pi_E \in \mathbb{Z} \subseteq \text{End}(E)$. Since $\deg(\pi_E) = q^2$, and the only integers in $\text{End}(E)$ with degree q^2 are $\pm q$, we have that $\text{tr}(\pi_E) = \pm 2q \equiv 0 \mod p$ where p is the characteristic of $\mathbb{F}_q$. This implies that E is supersingular, which is a contradiction.

So $\pi_E \notin \mathbb{Z}$, which implies that $\pi_E \notin \mathbb{Q} \subseteq \text{End}^0(E)$, because π_E is an algebraic integer. Hence either $\text{End}^0(E) = \mathbb{Q}(\pi_E)$ is an imaginary quadratic field, or $\text{End}^0(E)$ is a quaternion algebra via Theorem 4.2. In the latter case, it must contain some $\beta \in \text{End}^0(E)$ which does not commute with π_E. We will show that this does not occur.

Firstly, we will show that for all $m \geq 1$ we have $\pi_E^m = [a] \circ \pi_E + [b]$, for some $a \not\equiv 0 \pmod{p}$ and $b \equiv 0 \pmod{p}$.

We proceed by induction; noting that the base case holds with $a = 1$ and $b = 0$. For the inductive step we have:

$$
\begin{aligned}
\pi_E^{m+1} &= \pi_E \circ \pi_E^m \\
&= \pi_E \circ ([a] \circ \pi_E + [b]) \\
&= [b] \circ \pi_E + [a] \circ (\text{tr}(\pi_E) \circ \pi_E - [q]) \\
&= ([a] \circ \text{tr}(\pi_E) + [b]) \circ \pi_E - [aq] \\
&= [c] \circ \pi_E + [d],
\end{aligned}
$$

where $c = a\,\text{tr}(\pi_E) + b \not\equiv 0 \pmod{p}$, since $a\,\text{tr}(\pi_E) \not\equiv 0 \pmod{p}$ and $b \equiv 0 \pmod{p}$, and we have $d = -aq \equiv 0 \pmod{p}$, as desired.

Since $\pi_E^m = [a] \circ \pi_E + [b]$ with $a \neq 0$ and $\pi_E \notin \mathbb{Z}$ this shows that $\pi_E^m \notin \mathbb{Q}$ for any $m \geq 1$. Now given any $\alpha \in \text{End}^0(E)$, one can write α as $r\alpha = \phi$ for some $r \in \mathbb{Q}$ and $\phi \in \text{End}(E)$. The endomorphism ϕ is defined over $\mathbb{F}_{q^m}$ for some m. By writing ϕ as $\phi(x, y) = (s_1(x), s_2(x)y)$, we get

$$
(\phi \circ \pi_E^m)(x, y) = (s_1(x^{q^m}), s_2(x^{q^m})y^{q^m}) = (s_1(x)^{q^m}, s_2(x)^{q^m} y^{q^m}) = (\pi_E^m \circ \phi)(x, y),
$$

thus ϕ and therefore α commutes with π_E^m and we are done. $\qquad\square$

Remark. *Notice that $\alpha \in \mathbb{Q}(\pi_E^m) \subseteq \mathbb{Q}(\pi_E)$, hence the proof presents us with the exact quadratic field which is isomorphic to the endomorphism algebra.*

Corollary 4.5. *If $E/\mathbb{F}_q$ is an ordinary elliptic curve, then $\text{End}^0(E) \cong \mathbb{Q}(\sqrt{t - 4q})$, where $t = \text{tr}(\pi_E)$.*

Since we have that the endomorphism ring is an order, it will be contained in the unique maximal order of the quadratic field, which is the ring of integers. Also, since the Frobenius endomorphism π_E is contained in the endomorphism ring, and we have shown that it is not an integer, we have that $\mathbb{Z}[\pi_E] \subseteq \text{End}(E) \subseteq O_{\mathbb{Q}(\sqrt{D})}$. We will see later that the position of the endomorphism ring of the elliptic curve in this range will describe its place in the isogeny graph.

For now, we will turn our attention to describing the endomorphism ring of supersingular elliptic curves.

Theorem 4.6 ([Sil09, Thm. V.3.1(a)]). *Let $E/\mathbb{F}_{p^n}$ be a supersingular elliptic curve, then $\text{End}^0(E)$ is a quaternion algebra. In particular, $\text{End}(E)$ is a maximal order in $\text{End}^0(E)$.*

Remark. *Note that maximal orders of definite quaternion algebras are not unique. Hence the theorem does not imply that the endomorphism algebra determines the endomorphism ring.*

We will now explain Deuring's correspondence between the *category of supersingular elliptic curves and their isogenies* and the *category of invertible modules and their homomorphisms*. To that end, we will introduce kernel ideals which were introduced by Waterhouse in [Wat69]; they show the explicit link between isogenies and endomorphism rings. We will use them to build towards the proof as shown in §42.2 of Voight's unpublished book [Voi18]. Since the proofs of most of the results can be found readily in the book, proofs will not be provided here. We will selectively expand on some of the proofs in [Voi18] and will rely heavily on the proofs within. The focus will be on the intuition of the results leading to the correspondence.

Let E be a supersingular elliptic curve over $\overline{\mathbb{F}}_p$, and let $O = \text{End}(E)$, and $B = O \otimes \mathbb{Q}$. Let $I \subseteq O$ be a non-zero integral left O-ideal, then using [Voi18, Prop. 16.1.2], O being maximal implies that I is invertible. Define the *kernel ideal*

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

where the intersection is taken scheme-theoretically, and $E[\alpha] = \ker \alpha$ as a group scheme over $\overline{\mathbb{F}}_p$. It is sufficient to take the intersection on just the $\mathbb{Z}$-generators of I. Next, ϕ_I is defined to be an isogeny $E \to E/E[I]$ with kernel $E[I]$. We remark that $E[I]$ is finite, hence the definition makes sense.

Note that the recourse to group schemes is important for us to mention both separable and inseparable isogenies in the same breath. This allows us to reconcile with the fact that the Frobenius map has degree p but has trivial kernel. And this really comes to the fore when considering the example of $J = (\pi_E)$ which yields $E[J] = \{\mathcal{O}\}$ if we forget the group scheme structure.

The next result provides the first ingredient for proving the equivalence of categories.

Lemma 4.7. *Let $\phi_I : E \to E_I$ be an isogeny with kernel $E[I]$ as above. Then the map*

$$\phi_I^* : \mathrm{Hom}(E_I, E) \to I$$
$$\psi \mapsto \psi \circ \phi_I$$

is an isomorphism of left O-modules.

Remark. *Another way to view this result is to see it as*

$$\bigcap_{\psi \in \mathrm{Hom}(E_I, E)} \ker(\psi \circ \phi_I) = \ker \phi_I \quad or \quad \bigcap_{\psi \in \mathrm{Hom}(E_I, E)} \ker \psi = \{\mathcal{O}\}\,.$$

And to prove this, we show that

$$E[I] = \bigcap_{\alpha \in I} \ker \alpha = \bigcap_{\psi \in \mathrm{Hom}(E, E_I)} \ker(\psi \circ \phi_I)\,,$$

where the first equality is by definition and the second comes from the surjectivity of the map ϕ_I^*. Without surjectivity, we will only have a containment instead of the equality.

We identify the endomorphism ring within the endomorphism algebra with the next result:

Lemma 4.8. *The ring homomorphism*

$$\iota : \mathrm{End}(E_I) \hookrightarrow B$$
$$\beta \mapsto \phi_I^{-1} \circ \beta \circ \phi_I = \frac{1}{\deg(\phi_I)} \phi_I^\vee \circ \beta \circ \phi_I$$

is injective and $\iota(\mathrm{End}(E_I)) = \{b \in B \mid bI \subseteq I\} = O_R(I)$.

Remark. *We will see later that $O_R(I)$ is the right order of a given ideal and can be thought of as the endomorphism ring of the codomain curve. See §4.2.2.*

Proof. The map is certainly an injection since the only way that $\frac{1}{\deg(\phi_I)} \phi_I^\vee \circ \beta \circ \phi_I = 0$ is when $\beta = 0$ since all the other maps are non-zero maps. Then to show $\iota(\mathrm{End}(E_I)) = O_R(I)$, one can show the action on I by right multiplication and so $\iota(\mathrm{End}(E_I)) \subseteq O_R(I)$. We can achieve equality by using the maximality of $\mathrm{End}(E_I)$ and the fact that ι is a ring homomorphism. $\qquad\square$

Recall that isogenies are the "same" if the kernels are the same; the next result illustrates this corresponding fact which states that the ideal classes determine the isomorphism classes of the codomain curve.

Lemma 4.9. *If $J = I\beta \subseteq O$ with $\beta \in B^*$, then $E_I \cong E_J$.*

The work of Waterhouse establishes a link between isogenies and endomorphisms which we present now. Given a finite subgroup scheme $H \leq E(\overline{k})$, we define

$$I(H) = \{\alpha \in O \mid \alpha(P) = \mathcal{O}, \forall P \in H\}\,.$$

Recall the kernel ideal defined scheme-theoretically before. Now, we restrict to the case of when I contains a separable endomorphism α, this allows us to define kernel ideals in a more elementary way:

$$E[I] = \{P \in E(\overline{k}) \mid \alpha(P) = \mathcal{O}, \forall \alpha \in I\}\,.$$

We would like to think of $I(\cdot)$ as a functor from the category of subgroups to the category of ideals and that $E[\cdot]$ is the functor going the other way. The morphisms are given by inclusions in both categories. This results in $E[\cdot]$ and $I(\cdot)$ being contravariant functors. The astute reader will notice the parallels between the current exposition and the classic correspondence between algebraic sets and prime ideals in classical algebraic geometry. The two correspondences are used to prove the respective equivalence of categories. The following results that will build towards the equivalence of categories will enhance this similarity.

Lemma 4.10. *If $H_1 \subseteq H_2$ and $I(H_1) = I(H_2)$, then $H_1 = H_2$.*

Proof. We factor the isogeny through H_1 as in the following diagram

$$E \xrightarrow{\quad \phi_1 \quad} E/H_1 \xrightarrow{\quad\quad} E/H_2$$
$$\underset{\phi_I}{\underbrace{\qquad\qquad\qquad\qquad}}$$

Now, we split the proof into 2 cases: when ϕ_1 is separable, and when it is purely inseparable.

In the first case, we want to show that $H_2 \subseteq H_1$. But assume for contradiction that that isn't the case. Then there exists $\alpha \in O$ such that $\alpha(H_1) = \{\mathcal{O}\}$ but $\alpha(H_2) \neq \{\mathcal{O}\}$, so $I(H_2) \neq I(H_1)$, hence a contradiction.

In the second case, one can show that H_1 and H_2 must be the kernel of the same p^r-power Frobenius map and hence must be the same. $\qquad\square$

Proposition 4.11. *The following statements hold:*

(a) $\deg(\phi_I) = \mathrm{nrd}(I)$,

(b) $I(E[I]) = I$.

Remark. *We will make several remarks on the proof found in [Voi18, Prop. 42.2.16] and will thus use the notation found there. In the first displayed equation of Voight's proof, the equality can be presented in the following way which might be more illuminating:*

$$[\deg \beta] = \beta \circ \hat{\beta} \cong \beta\bar{\beta} = \mathrm{nrd}(\beta) = \mathrm{nrd}(I) \,.$$

Indeed, all equalities are definitions of $\deg$ and nrd and the isomorphism is given by the fact that $\hat{}$ and $\bar{}$ are involutions in their respective rings.

Next, the proof reduces the general case to the principal ideal case which has already been proven. In particular, one can use the multiplicative property of $\deg$ and the fact that II' is principal to get the divisibility on both sides. The key is the use of [Voi18, Ex. 17.5]. This proves (a), and (b) follows shortly after.

Corollary 4.12. *For every isogeny $\phi : E \to E'$, there exists a left O-ideal I and an isomorphism $\rho : E_I \to E'$ such that $\phi = \rho \circ \phi_I$. In particular, for every maximal order $O' \subseteq B$, there exists E' such that $O' \cong \mathrm{End}(E')$.*

Lemma 4.13. *For every non-zero integral left O-ideal $I, I' \subseteq O$, the natural map*

$$\mathrm{Hom}(E_I, E)\,\mathrm{Hom}(E_{I'}, E_I) \to \mathrm{Hom}(E_{I'}, E)$$

is bijective, giving a further bijection

$$\mathrm{Hom}(E_{I'}, E_I) \to (I : I')_R = I^{-1}I'$$
$$\psi \mapsto \phi_I^{-1} \circ \psi \circ \phi_{I'} \,.$$

Proof.

$$
\begin{aligned}
I\phi_I^{-1}\,\mathrm{Hom}(E_{I'}, E_I)\phi_{I'} &= \mathrm{Hom}(E_I, E)\phi_I\phi_I^{-1}\,\mathrm{Hom}(E_{I'}, E_I)\phi_{I'} \\
&= \mathrm{Hom}(E_I, E)\,\mathrm{Hom}(E_{I'}, E_I)\phi_{I'} \\
&= \mathrm{Hom}(E_{I'}, E)\phi_{I'} \\
&= I' \,.
\end{aligned}
$$

$\qquad\square$

We now have all the ingredients for the equivalence of categories to be proved. Let E_0 be a supersingular curve over $\overline{\mathbb{F}}_p$. This will serve the role as a base object. We also fix the base ring $O_0 = \text{End}(E_0)$ and the base algebra $B_0 = O_0 \otimes \mathbb{Q}$.

Remark. *One really needs a base object before looking at the equivalence of categories. While the bijection between isomorphism classes of supersingular curves and isomorphism classes of maximal orders can be done without the base objects, the morphisms between maximal orders cannot exist without the base object using the results we have developed thus far. Interested readers can refer to §42.4 of [Voi18] for a base-object free discussion.*

Theorem 4.14 ([Voi18, Thm. 42.3.2]). *Let E be a supersingular elliptic curve, then the association $E \mapsto \text{Hom}(E, E_0)$ is functorial and defines an equivalence between the category of*

supersingular elliptic curves over k, under isogenies

and

invertible left O_0-modules, under left O_0-module homomorphisms.

Remark. *The categorical view of this is that the functor $\text{Hom}(-, E_0)$ is contravariant. Hence we can define the equivalence using the association $E \mapsto \text{Hom}(E_0, E)$ to get a covariant functor with right O_0-modules.*

4.2 Structure of Isogeny Graphs of Elliptic Curves

Now that we have gained an understanding of the endomorphism rings of elliptic curves, we can put that theory to use by studying the structures of isogeny graphs. As mentioned, the isogeny graph is highly dependent on the endomorphism ring structure of elliptic curves. This will be examined fully in this section.

Definition. Let p and ℓ be distinct primes, then the ℓ-isogeny graph $G_{p^n, \ell}$ is the directed graph with vertices elements $\mathbb{F}_{p^n}$ and edges (j_1, j_2) with weights equal to the multiplicity of j_2 as a root in $\mathbb{F}_{p^n}$ of $\Phi_\ell(j_1, Y) \in \mathbb{Z}[Y]$, where $\Phi_\ell(X, Y)$ is the ℓ-modular equation [Cox13, §11.C].

The consequence of Tate's theorem [Tat66] is that $G_{p^n, \ell}$ is disconnected. Furthermore, the graph can be split into two main components: the ordinary component and the supersingular component. We will see later that the supersingular elliptic curves form a connected $(\ell + 1)$-regular component. This is a stronger claim than what Tate's theorem claims, as the isogenies now have degrees which are an ℓ-power. We will prove this in §4.2.2.

We will deal first with the ordinary components. The ordinary component can be further subdivided into subcomponents that form ℓ-*volcanoes*.

Definition ([Sut13, Def. 1]). An ℓ-*volcano* V is a connected undirected graph whose vertices are partitioned into one or more levels $V_0, \ldots, V_d$ such that the following hold:

(1) The subgraph on V_0 (the crater) is a regular graph of degree at most 2.

(2) For $i > 0$, each vertex in V_i has exactly one neighbour in level V_{i-1}, and this accounts for every edge not on the crater.

(3) For $i < d$, each vertex in V_i has degree $\ell + 1$.

(4) Each vertex in V_d has degree 1.

4.2.1 Ordinary

In this section, we will briefly review the structure of isogeny graphs for ordinary elliptic curves over a finite field. The bulk of the description of the isogeny graphs for ordinary elliptic curves was explicated by David Kohel in his PhD thesis [Koh96].

Let E/k be an elliptic curve with complex multiplication by O, an order with discriminant D in an imaginary quadratic field K. Also, let ℓ be a prime distinct from $\text{char}(k)$, and we want to consider ℓ-isogenies away from E. Let E' be an elliptic curve such that $\phi : E \to E'$, where ϕ is an ℓ-isogeny. Let $\text{End}(E') = O'$, then we have that $\text{End}^0(E) = \text{End}^0(E')$, but it is not necessarily the case that $O = O'$. The relation between O and O' is given in the following proposition:

Proposition 4.15 ([Koh96, Prop. 21]). *Let $\phi : E \to E'$ be as above, and let $\mathrm{End}(E) \cong O$ and $\mathrm{End}(E') \cong O'$, then one of the following holds:*

(a) $O = O'$,

(b) $O' \subseteq O$ and $[O : O'] = \ell$,

(c) $O \subseteq O'$ and $[O' : O] = \ell$.

We will see how this proposition dictates the structure of the graph. The following definition, while seemingly unmotivated at the moment, will become intuitively clear when seen in the context of the graph.

Definition. Using the notation from the above result:
If $O = O'$, we say that ϕ is *horizontal*.
If $O' \subseteq O$ and $[O : O'] = \ell$, we say that it is *descending*.
If $O \subseteq O'$ and $[O' : O] = \ell$, we say that it is *ascending*.
In the latter two cases, we refer to them as being *vertical*.

This leads to the main theorem of this section, which describes the ordinary components of the isogeny graph.

Theorem 4.16 ([Koh96, Prop. 23]). *Let V be an ordinary component of $G_{q,\ell}$ that does not contain the j-invariants 0 or 1728. Then V is an ℓ-volcano for which the following hold:*

(1) *The vertices in level V_i all have the same endomorphism ring O_i.*

(2) *The subgraph V_0 has degree $1 + \left(\frac{D_0}{\ell}\right)$, where $D_0 = \mathrm{disc}(O_0)$.*

(3) *If $\left(\frac{D_0}{\ell}\right) \geq 0$, then $|V_0|$ is the order of $[\mathfrak{l}]$ in $\mathrm{Cl}(O_0)$, where $\mathfrak{l}$ is an ideal above ℓ. Otherwise $|V_0 = 1|$.*

(4) *The depth of V is $d = \frac{v_\ell((t^2-4q)/D_0)}{2}$, where $t = \mathrm{tr}(\pi_E)$ for $j(E) \in V$.*

(5) *$\ell \nmid [O_K : O_0]$ and $[O_i : O_{i+1}] = \ell$ for $0 \leq i \leq d$.*

We will use the theorem to explain some of the features of the ℓ-volcanoes of ordinary components. Figures 4.1b and 4.1d show the 2 and 3-isogeny graph for an ordinary component over the field $\mathbb{F}_{37}$. The vertices are given by the j-invariants of the elliptic curves.

The first notable feature of the graph is the 4-cycle in both of the graphs. This would correspond to the crater of the volcano as given by (2) of the theorem. The 2-isogeny graph has a depth of 1, while the 3-isogeny graph has a depth of 0. One can check that the elliptic curves at the depth of 1 of the 2-isogeny graph have endomorphism rings which have conductor 2 and these vertices have degree 1. The isogenies traversing the two levels are vertical.

4.2.2 Supersingular

As mentioned, the supersingular component of the ℓ-isogeny graph forms a single connected component. However, it is a priori not clear why this would be the case. Using Tate's theorem, one is able to deduce that any two supersingular elliptic curve would be isogenous, but it does not imply that they should be connected in the ℓ-isogeny graph. To see this, we will need to invoke the equivalence of categories of Theorem 4.14. One then has to use results in Gross [Gro87] and Pizer [Piz90] which we will now retrace.

Let B be the definite quaternion algebra ramified at p and ∞. Fix some maximal order O in B, then recall that a left O-ideal is a lattice that is stable under left action by O. Associated to this left O-ideal is a *right order* of I given by $\{b \in B \mid Ib \subseteq I\}$, this is yet again another maximal order. We can define $I^{-1} = \{b \in B \mid IbI \subseteq I\}$ as the right ideal of O whose left order is the right order of I.

The left ideals are equivalent if they are associates, i.e. I and J are in the same equivalence class if $I = bJ$ for some $b \in B^*$. The number of equivalence classes is given by the *class number of B*. We can enumerate the equivalence classes of the left ideals by $\{I_1, \ldots, I_n\}$ such that $I_1 = O$ is the trivial ideal.

There is an immediate link to supersingular elliptic curves that we will now delineate. In choosing B, we have in effect chosen a prime p, over which B is ramified. In terms of supersingular elliptic curves, this serves to fix the field $\overline{\mathbb{F}}_p$, and hence fixes the set of isomorphism classes of supersingular elliptic

curves. The class number of B then corresponds to the number of isomorphism classes of supersingular elliptic curves over $\overline{\mathbb{F}}_p$.

Coming back to quaternion orders, we can associate to each ideal-class representative a right order O_i. This results in a list $\{O_1, \ldots, O_n\}$, where each conjugacy class of maximal orders is represented once or twice. The number of conjugacy classes is given by the *type number of B*.

Let n be the class number of B and let t be the type number of B, then we can note that $t < n$. The connection between the class and type numbers of B can be seen from the vantage of supersingular elliptic curves. We already stated that the class number is the number of supersingular elliptic curves over $\overline{\mathbb{F}}_p$. The type number on the other hand is the number of conjugacy classes of the endomorphisms of these supersingular elliptic curves. The disparity between the two values arises from the fact that

$$E \cong \overline{E} \iff E \text{ is defined over } \mathbb{F}_p$$

where $\cong$ represents isomorphism and $\overline{E}$ is E acted on by the Galois action, but we have that

$$\text{End}(E) \cong \text{End}(\overline{E})$$

where $\cong$ represents conjugacy. A fortiori, $\text{End}(E)$ is isomorphic to $\text{End}(\overline{E})$ as rings!

Example. This is a nice example of how non-isomorphic elliptic curves can have isomorphic endomorphism rings. Let $p = 83$, and $k = \mathbb{F}_{p^2}$ and let $E_0 : y^2 = x^3 + x$ be the curve with j-invariant $j = 1728 \equiv 68$ mod 83. Then we have the following 3-isogeny graph of supersingular elliptic curves:

Using E_0 as the base ring, we can denote the Frobenius endomorphism by i and the endomorphism $(x, y) \mapsto (-x, iy)$ by j, then we have that

$$\text{End}(E_0) = \mathbb{Z} \left\langle 1, \mathbf{j}, \frac{\mathbf{j} + \mathbf{k}}{2}, \frac{1 + \mathbf{i}}{2} \right\rangle$$

where this is over $B_{p,\infty}$, the quaternion algebra ramified at $p = 83$ and ∞, and $\mathbf{i}^2 = -83$, $\mathbf{j}^2 = -1$ and $\mathbf{k} = \mathbf{ij}$.

We write $E(j_{\text{inv}})$ for an elliptic curve whose j-invariant is given by j_{inv}. Then we can work out using methods in [GPS17, §4] that

$$\text{End}(E(50)) = \mathbb{Z} \left\langle \frac{1 + \mathbf{i}}{2}, \mathbf{i}, \frac{\mathbf{j} + 151\mathbf{k}}{18}, 9\mathbf{k} \right\rangle,$$

$$\text{End}(E(28)) = \mathbb{Z} \left\langle \frac{9 + \mathbf{i} + 59\mathbf{j} + 9\mathbf{k}}{18}, \frac{2\mathbf{i} + 37\mathbf{j} + 9\mathbf{k}}{18}, \frac{9\mathbf{j} + \mathbf{k}}{2}, \mathbf{k} \right\rangle,$$

$$\text{End}(E(17)) = \mathbb{Z} \left\langle \frac{9 + \mathbf{i} + 13\mathbf{j} + 21\mathbf{k}}{18}, \frac{\mathbf{i} + 13\mathbf{j} + 21\mathbf{k}}{9}, \frac{3\mathbf{j} + \mathbf{k}}{2}, 3\mathbf{k} \right\rangle,$$

$$\text{End}(E(\alpha)) = \text{End}(E(\bar{\alpha})) = \mathbb{Z} \left\langle \frac{9 + 3\mathbf{i} + \mathbf{j} + 133\mathbf{k}}{18}, \frac{3\mathbf{i} + \mathbf{j} + 52\mathbf{k}}{9}, \frac{\mathbf{j} + 7\mathbf{k}}{6}, 9\mathbf{k} \right\rangle.$$

Clearly, $E(\alpha) \not\cong E(\bar{\alpha})$ over k as they have different j-invariants, but as we can see, the endomorphism rings are the same.

Now, fix p and let $\{E_1, \ldots, E_n\}$ be the isomorphism classes of supersingular elliptic curves over $\overline{\mathbb{F}}_p$. Then we define $O_i = \operatorname{End}(E_i)$. This gives the isomorphism

$$I_j^{-1} I_i \cong \operatorname{Hom}(E_i, E_j)$$

as a left O_j-module and a right O_i-module.

Recall that our aim is to show that the supersingular component of $G(p^n, \ell)$ is connected. Thus far, we have shown the connection between supersingular isogenies and maximal orders of quaternion algebras. In particular, we have shown that maximal orders can be connected via left and right ideals, as a result, we are able to study the graphs made from these gadgets. The equivalence of categories allows us to transfer the isogeny graph to the graph of maximal orders, hence a connectivity result in the latter case can be translated to a connectivity result in the former.

The coup de grâce comes from a result from Pizer. Before presenting the theorem, we note that Pizer considers orders of level $N = p^2 M$ in $B_{p,\infty}$, the quaternion algebra ramified at p and ∞. This is simply a generalisation of the maximal orders. The maximal orders we are interested in are orders of level p^2.

Furthermore, the result will assert that the graph obtained is in fact connected and is *Ramanujan*. The significance of being a Ramanujan graph is that it has good mixing properties that can be used in cryptographic proofs.

Theorem 4.17 ([Piz90, Thm. 1])**.** *Let O be a maximal order in $B_{p,\infty}$ and let ℓ be a prime with $\ell < p/4$ and $\ell \nmid N$. Then the multi-graph $G(p^2, \ell)$ is defined and is a $\ell + 1$ regular connected Ramanujan graph.*

Remark. *Note that Pizer proved this theorem using properties of Brandt matrices that we have omitted. In short, for a given prime p and some integer m, Brandt matrices are matrices in $M(n, \mathbb{Z})$, where n is the number of isomorphism classes of supersingular elliptic curves over $\mathbb{F}_{p^2}$. The (i,j)-entry of the Brandt matrix is defined to be equal to the number of subgroup schemes K of order m in E_i such that $E_i/K \cong E_j$. Interested readers are encouraged to refer to [Gro87, §1–2].*

We note that all the vertices of the supersingular component can be defined over $\mathbb{F}_{p^2}$ (cf. Theorem 1.10). Also, this component is a $(\ell + 1)$-regular graph since the number of distinct subgroups of order ℓ of $E[\ell]$ is $\ell + 1$.

The supersingular component does not have the ℓ-volcano structure present in the ordinary components. This is because the endomorphism rings are all maximal orders of $B_{p,\infty}$, so there is a lack of containment of endomorphism rings that gave rise to the levels of the ℓ-volcano structure of their ordinary counterparts.

In Figure 4.1, we give an example of the isogeny graph in $\mathbb{F}_{37}$ which shows an ordinary component and the supersingular component for $\ell = 2$ and 3. In particular, one can clearly see that the out-degree of the ℓ-isogeny graphs of the supersingular components is $\ell + 1$.

4.3 Application to Cryptography

4.3.1 Group Action by Class Group

Couveignes–Rostovtsev–Stolbunov Isogeny Cryptosystem

Let p be a prime and $E/\mathbb{F}_{p^n}$ be an ordinary elliptic curve with complex multiplication by O. Recall from the discussion in §4.1 that O is an order in some imaginary quadratic field. Also, we have that if $\operatorname{Cl}(O)$ is the class group of O, and $\mathfrak{a}$ is an O-ideal, then $\mathfrak{a}$ induces an action on the set of elliptic curves with complex multiplication by O, and maps E to $E_{\mathfrak{a}}$ via the isogeny $\phi : E \to E_{\mathfrak{a}}$. The kernel of ϕ is given by $\ker \phi = \{P \in E(\overline{\mathbb{F}}_p) \mid \alpha(P) = \mathcal{O} \quad \forall \alpha \in \mathfrak{a}\}$ as we have seen with kernel ideals. This is known as the class group action on the set of isomorphism classes of elliptic curves with complex multiplication by O.

Couveignes, Rostovtsev and Stolbunov [Cou06, RS06] independently proposed a key exchange protocol based on the above action. Alice and Bob first agree on the starting curve, E_0 which is chosen from the set of isomorphism classes of elliptic curves with complex multiplication by O. Alice chooses the secret O-ideal $\mathfrak{a}$, and computes $E_{\mathfrak{a}}$ using the action above. She then publishes $E_{\mathfrak{a}}$. Bob would complete his side of the protocol and publish his public key $E_{\mathfrak{b}}$. To obtain the shared secret, Alice would compute the action of $\mathfrak{a}$ on $E_{\mathfrak{b}}$ to obtain $E_{\mathfrak{b}\mathfrak{a}}$, which will be the same as Bob's computation of $E_{\mathfrak{a}\mathfrak{b}}$, since the class group is commutative. We call this the CRS protocol.
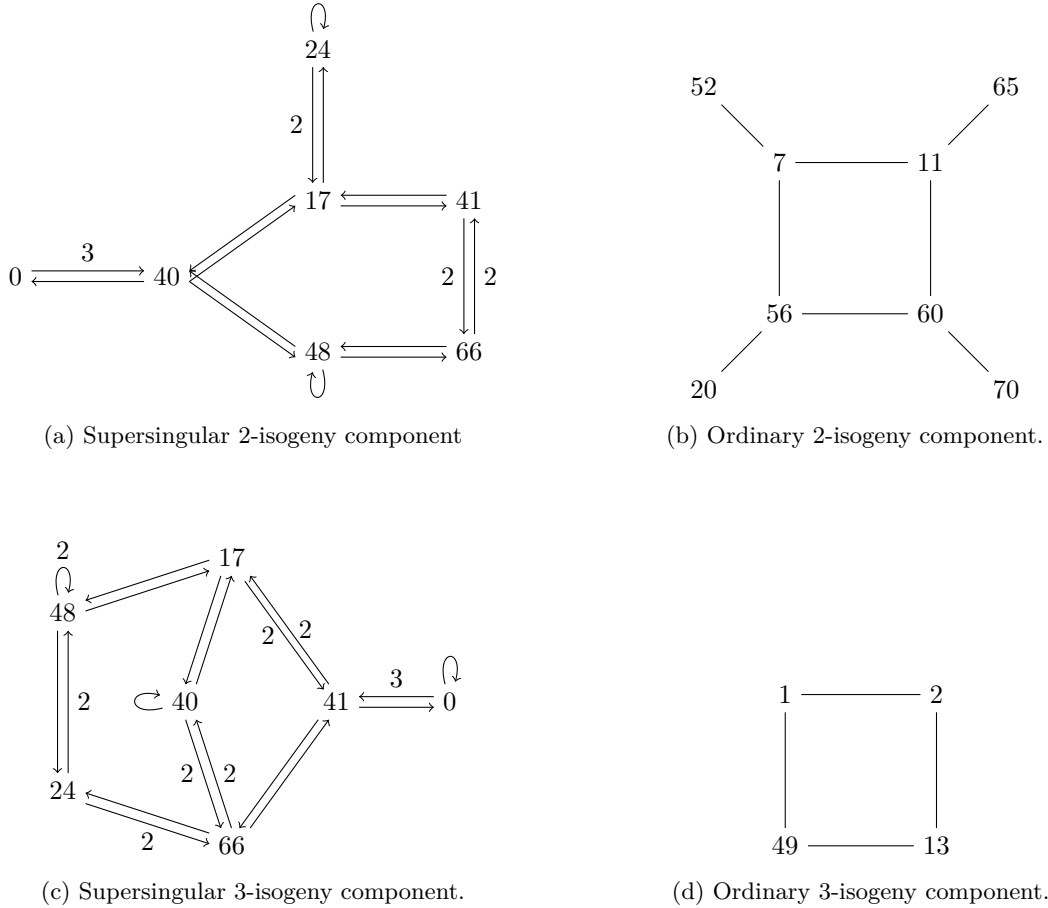
(a) Supersingular 2-isogeny component

(b) Ordinary 2-isogeny component.

(c) Supersingular 3-isogeny component.

(d) Ordinary 3-isogeny component.

Figure 4.1: Isogeny components in $\mathbb{F}_{37}$.

The hard problem here is to recover the O-ideal $\mathfrak{a}$ when given E and $E_{\mathfrak{a}}$. This is equivalent to the isogeny problem. However, it has since been discovered [CJS14] that there exists a quantum algorithm, based on the Kuperberg algorithm, that is able to solve the hard problem in sub-exponential time.

The existence of a sub-exponential algorithm did not signal the end of the cryptosystem. One could increase the parameters of the cryptosystem to thwart the attack, but this would come at the cost of the speed of the algorithm. The main bottleneck in the CRS protocol lies in the computation of the action. The computation of the action takes $O(d)$, where d is the norm of the ideal or the degree of the isogeny. Actions in the cryptosystems can factor into a sequence of smaller isogenies. However, this factorisation very much depends on the structure of the endomorphism ring of the elliptic curve. The efficiency of the cryptosystem hinges on reducing the norm of the largest factor in the factorisation of the action.

In 2018, De Feo, Kieffer, and Smith tackled the issue [DKS18] by seeking to reduce the norm of the largest factor. The difficulty lies in simultaneously controlling p and the order of the elliptic curve. The aim was to increase the size of the prime of the protocol, thereby thwarting the quantum attack, and still find an efficient method to compute the action. However, their attempts at finding ordinary elliptic curves for this cryptosystem were unfruitful.

CSIDH

An improvement to [DKS18] is CSIDH [CLM$^+$18]. CSIDH uses the same principles as the protocol in [DKS18] by maximising the number of small prime ideals where the corresponding kernels have rational points. The key idea in this case was the use of supersingular elliptic curves which allows for one to efficiently choose such an order. Recall that supersingular elliptic curves can be defined over $\mathbb{F}_{p^2}$, and $\mathrm{End}_{\overline{\mathbb{F}}_p}(E)$ is a maximal order in the quaternion algebra which is ramified at p and ∞. However, when

restricted to $\mathbb{F}_p$, one gets that $\text{End}_{\mathbb{F}_p}(E/\mathbb{F}_p)$ is an order of an imaginary quadratic field. One can choose a suitable prime and the endomorphism ring will contain suitable prime ideals. In choosing a suitable prime, such that $p+1$ contains sufficiently many small primes, we have that $\text{End}_{\mathbb{F}_p}(E/\mathbb{F}_p)$ will contain prime ideals of norm equal to the prime factors of $p+1$. Hence isogenies computed in the cryptosystem can be factored into isogenies with small prime norms which are more efficient computationally.

4.3.2 Solving the Isogeny Problem when the Endomorphism Ring is Known

In this section, we investigated the possibility of applying the equivalence of categories developed in Theorem 4.14 to the cryptosystem of Jao and De Feo. This follows [GPST16, §4] closely and is joint work with Christophe Petit.

Let $p = \ell_A^n \cdot \ell_B^m \cdot f - 1$ as in the Jao–De Feo cryptosystem, and let E and E_A be two supersingular elliptic curves such that there exists an isogeny $\phi_A : E \to E_A$ of degree ℓ_A^n between them. In this section, we will show that having the ability to efficiently compute the endomorphism rings $\text{End}(E)$ and $\text{End}(E_A)$ would result in an efficient algorithm to recover ϕ_A assuming a certain natural heuristic holds. The following result is a formal statement of this reduction which we will prove after a brief introduction.

Theorem 4.18. *Let E and E_A be supersingular elliptic curves over $\mathbb{F}_{p^2}$ such that $E[\ell_A^n] \subseteq E(\mathbb{F}_{p^2})$ and suppose there is an isogeny $\phi_A : E \to E_A$ of degree ℓ_A^n from E to E_A. Suppose there is no isogeny $\phi : E \to E_A$ of degree $< \ell_A^n$. Then, given an explicit description[8] of $\text{End}(E)$ and $\text{End}(E_A)$, there is an efficient algorithm to compute ϕ_A with high probability.*

Computing the endomorphism ring of a supersingular elliptic curve is a problem that is believed to be equivalent to computing an arbitrary isogeny between two supersingular elliptic curves. The equivalence of categories we obtained in §4.1 for supersingular elliptic curves allows us to translate the supersingular isogeny problem to the category of maximal orders. In that category, finding a connecting ideal between two maximal orders is simply linear algebra. However, we will show that finding an arbitrary isogeny is insufficient for breaking SIDH; we need to find an isogeny of the correct degree (cf. remark below). In 2014, Kohel, Petit, Lauter and Tignol [KLPT14] presented an algorithm that is able to find a connecting ideal whose norm is a prime power. However, their algorithm does not produce an isogeny that satisfies the additional constraint that it must be of small degree, as is required in SIDH ($\ell_A^n \approx p^{1/2}$). Hence the current state of affairs does not give a reduction of the form we require.

The aim of this section is to present an alternative method to [KLPT14] in this context. We use some of the notation of [KLPT14].

Remark (The Importance of the Correct Isogeny)**.** *It is not sufficient to compute an arbitrary isogeny from E to E_A if one's goal is to break SIDH. Indeed, one needs to compute an isogeny with the right properties.*

Suppose there are curves E and isogenies $\phi_A : E \to E_A$, $\phi_B : E \to E_B$ with $\ker \phi_A = G_A, \ker \phi_B = G_B$ satisfying the commutative isogeny diagram from §2.1:

$$
\begin{array}{ccc}
& E_A = \frac{E}{G_A} & \\
\phi_A \nearrow & & \searrow \\
E & & E_{AB} = \frac{E}{\langle G_A, G_B \rangle} \\
\phi_B \searrow & & \nearrow \\
& E_B = \frac{E}{G_B} &
\end{array}
$$

The protocol follows from the fact that $E/\langle G_A, G_B \rangle = E_A/\langle \phi_A(G_B) \rangle = E_B/\langle \phi_B(G_A) \rangle$, since $\ker \phi_A$ and $\ker \phi_B$ have trivial intersection, and that $\phi_A(G_B)$ and $\phi_B(G_A)$ can be computed using the auxiliary information.

However, suppose an attacker given E, E_A, E_B and the auxiliary information is able to compute some isogeny $\phi' : E \to E_A$ such that ϕ' is not equal to ϕ_A. This is the resulting picture:

---

[8]This means that one can express $\text{End}(E)$ and $\text{End}(E_A)$ in terms of a rank 4 $\mathbb{Z}$-module and give the bases of these endomorphism rings as elements of the quaternion algebra.

The attacker attempting to derive the shared secret $j(E_{AB})$ has to compute $\phi_B(\ker(\phi'))$ and will obtain an isogeny with this kernel. However, the attacker will not be able to derive the shared secret since they only have the points $\phi_B(P_A)$ and $\phi_B(Q_A)$ to work with. The attacker can only compute $\phi_B(\ker(\phi'))$ if $\ker \phi' \subseteq \langle P_A, Q_A \rangle$ which would imply that ϕ' is an isogeny of degree dividing 2^n.

This is the crux of the difficulty in giving a reduction from computing endomorphism rings to computing the secret key in the Jao–De Feo cryptosystem: known algorithms to compute an isogeny from E to E_A, given $\mathrm{End}(E)$ and $\mathrm{End}(E_A)$, are not likely to give an isogeny of the correct degree. However, we can exploit the size of the secret isogeny to recover the secret isogeny. This size restriction is absent in the general case.

We now show how the existence of a small degree isogeny actually *helps* the cryptanalysis of the Jao–De Feo cryptosystem, assuming we know (or are able to compute) the endomorphism rings of the curves in play.

Given two maximal orders O and O_A, one can compute in polynomial time an ideal I that connects them [KLPT14, Lem. 8]. Computing an isogeny of the correct degree corresponds to computing an equivalent ideal of the correct norm. In order to find such an equivalent ideal we use the following lemma.

Lemma 4.19 ([KLPT14, Lem. 5]). *Let I be a left O-ideal of reduced norm N and α an element in I. Then $I\gamma$, where $\gamma = \bar{\alpha}/N$, is a left O-ideal of norm $n(\alpha)$.*

We observe that in the context of Jao–De Feo cryptosystem, there exists by construction an element α of small norm $N\ell_A^n$ in I, corresponding via this lemma to an ideal of norm ℓ_A^n. Moreover, as Minkowski bases can be computed in polynomial time for lattices of dimension up to 4 [NS04], this element α can be efficiently recovered as long as it is in fact the smallest element in I. These observations lead to the following simple algorithm:

Algorithm 3: Computing small degree isogenies in Jao–De Feo cryptosystem given an algorithm to compute the endomorphism ring of a random supersingular elliptic curve.

Data: ℓ_A, n, E, E_A, $O = \mathrm{End}(E)$, $O_A = \mathrm{End}(E_A)$ such that E and E_A are connected by an isogeny of degree ℓ_A^n

Result: Isogeny $\varphi_A : E \to E_A$ of small degree ℓ_A^n, or *failure*

1 Compute an ideal I connecting O and O_A as in [KLPT14, Lem. 8];
2 Compute a Minkowski-reduced basis of I;
3 Let α be the non-zero element in I of minimal norm;
4 **if** $n(\alpha) \neq n(I)\ell_A^n$ **then** return *failure*;
5 Compute an ideal $I' = I\bar{\alpha}/n(I)$;
6 Compute the isogeny φ_A that corresponds to I' using methods in [GPS17, §4];
7 Return φ_A;

The probability of failure in Step 4 will be discussed after the proof. All the steps in this algorithm can be performed in polynomial time. The above discussion forms the proof of Theorem 4.18.

Proof of Theorem 4.18. Given an explicit representation of the endomorphism rings, we can translate the endomorphism rings into maximal orders of quaternion algebras. One can then find, in polynomial time, an ideal I connecting them by [KLPT14, Lem. 8].

By Lemma 4.19, it is sufficient to find an element of I of the correct norm. But given that the norm we seek is the smallest norm in the ideal, we can use lattice reduction methods to recover the smallest norm in polynomial time. Then using methods in [KLPT14], we can recover an isogeny with degree ℓ_A^n.

The isogeny recovered thus is not necessarily ϕ_A, but as the discussion that will follow, we will see that it will be ϕ_A with high probability. $\qquad\square$

In the remainder of this section, we study the success probability of this algorithm on average, and show how to use it to achieve a very large success probability.

Heuristically, we can approximate the probability that E and E_A are connected by an isogeny of degree d by estimating the probability that two randomly chosen supersingular elliptic curves are connected by an isogeny of the same degree[9].

Random pairs of supersingular elliptic curves over $\mathbb{F}_{p^2}$ are unlikely to be connected by isogenies of degrees significantly less than $\sqrt{p}$. Indeed, when $d = \prod_i p_i^{e_i}$, there are exactly

$$a(d) = \prod_i (p_i + 1) p_i^{e_i - 1}$$

isogenies of degree d from any curve E, hence any curve E is connected to at most $\sum_{d \leq D} a(\ell)$ curves E_A by an isogeny of degree at most D. A calculation given in [GPST16, App. A] shows that this sum behaves asymptotically like

$$\frac{15}{2\pi^2} D^2$$

as D tends to infinity. As there are roughly $p/12$ supersingular invariants over $\mathbb{F}_{p^2}$ we can evaluate the success probability of the above algorithm as

$$SR \approx \max\left(0,\ 1 - \frac{15 D^2}{2\pi^2} \bigg/ \frac{p}{12}\right) \approx \max\left(0,\ 1 - \frac{90}{\pi^2} \frac{\ell_A^{2n}}{p}\right).$$

For the parameters used in the Jao–De Feo cryptosystem we expect this basic attack to succeed with a probability larger than 50% if $f > \frac{180}{\pi^2} \approx 18.23$, where f is the cofactor in $p = \ell_A^n \cdot \ell_B^m \cdot f \pm 1$. The upshot is that a small cofactor would be able to defeat this attack.

However, the success rate of our attack can be easily improved in two ways. First, we can apply the algorithm separately on all curves that are at distance ℓ_A^e of E_A for some small constant e, until it succeeds for one of them. Clearly one of these curves will be connected to E by an isogeny of degree ℓ_A^{n-e}, and as a result the success rate will increase to

$$SR \approx \max\left(0,\ 1 - \frac{90}{\pi^2} \frac{\ell_A^{2(n-e)}}{p}\right).$$

With $\ell_A = 2$ and $e = 10$ this method will lead to a success rate above 99%, even when $f = 1$. Second, we can try to use the Minkowski-reduced basis computed in Step 3 of the algorithm to find an element α of the appropriate norm, even when it is not the smallest element. We explore two heuristic methods in that direction in our experiments below.

Experimental Results

We tested our algorithm in `MAGMA` with $\ell_A = 2$ and with a λ-bit prime p, a randomly selected maximal order, another random maximal order connected to the first by a path of length $\Delta = \lceil \log_{\ell_A}(p)/2 \rceil + \delta$, with $\delta \in \{-5, \ldots, 5\}$. One can traverse from the first order to the second via Δ steps in the ℓ_A-isogeny tree. This would correspond to picking a randomly selected supersingular curve over $\mathbb{F}_{p^2}$ and another supersingular curve connected to the first by an isogeny of degree ℓ_A^Δ.

The first three columns of Table 4.1 ("First basis element") correspond to the attack described in the previous section. In other words, the algorithm succeeds if the first vector output in the reduced lattice basis corresponds to the correct ℓ_A^Δ-isogeny. The next three columns ("All basis elements") correspond to a variant where instead of considering only the smallest element in Step 4 of the algorithm, we try all elements in the Minkowski-reduced basis. Finally, the last three columns ("Linear combinations") correspond to a variant where we search for α of the right norm amongst all elements of the form

---

[9]The argument is not totally accurate as E and E_A are slightly closer in the ℓ_A-isogeny graph than random pairs of curves would be. This may a priori impact the probabilities, however a significant distortion of these probabilities would reveal some unexpected properties of the graph, such as the existence of more or fewer loops of certain degrees than expected.

$\sum_{i=1}^{4} c_i \beta_i$, where $c_i \in \{-4, \dots, 4\}$ and β_i are the Minkowski-reduced basis elements. This is a naïve form of the enumeration algorithm with an arbitrary choice of 4 as a bound. Each percentage in the table corresponds to a success rate over 100 experiments, where a success is when we have found an element with correct norm using the methods mentioned above.

		First basis element			All basis elements			Linear combinations		
		λ			λ			λ		
		100	150	200	100	150	200	100	150	200
	-5	100%	99%	99%	100%	100%	99%	100%	100%	100%
	-4	93%	99%	94%	98%	99%	100%	100%	100%	100%
	-3	83%	84%	88%	92%	95%	99%	100%	100%	100%
	-2	40%	43%	45%	81%	74%	76%	100%	100%	100%
	-1	0%	2%	0%	35%	42%	35%	100%	100%	99%
δ	0	0%	0%	0%	3%	4%	3%	100%	100%	100%
	1	0%	0%	0%	1%	0%	0%	97%	99%	98%
	2	0%	0%	0%	0%	0%	0%	95%	94%	91%
	3	0%	0%	0%	0%	0%	0%	57%	68%	70%
	4	0%	0%	0%	0%	0%	0%	25%	28%	18%
	5	0%	0%	0%	0%	0%	0%	0%	3%	1%

Table 4.1: Experimental results for δ values. $\ell = 2$.

The experimental results mean that, even if the desired isogeny does not correspond to the shortest element in a reduced basis, a small exhaustive search over neighbouring curves under a 2^δ-isogeny for $1 \leq \delta \leq 5$ will with high probability lead to a curve for which the desired isogeny is the shortest element in a reduced basis.

Chapter 5

Properties of PPSSAS Isogeny Graphs

We introduced G2SIDH in Chapter 2 but omitted the security analysis. Given that the isogeny graph of elliptic curves informs us on the security of the cryptosystems on which they are based, it is imperative that we understand the isogeny graph of PPSSASs to quantify the security of G2SIDH. In this final chapter we will find the number of neighbours of an arbitrary vertex in the PPSSAS graph. We will also examine paths between two vertices in the graph.

We have seen that one can specify an isogeny by its kernel; in this chapter we explain that such an isogeny can sometimes be decomposed in more than one way. This gives rise to a peculiar substructure (ℓ-diamonds) present in the isogeny graphs of PPSSASs. We find that the ℓ-diamonds implies the existence of numerous cycles of small length in the (ℓ, ℓ)-isogeny graph of PPSSASs. All these results are given in §5.2 and §5.3. In fact, these results apply to general PPASs as we will see.

These results have implications on the cryptosystems in genus two. In particular, the number of neighbours inform us on the security of G2SIDH as we will encounter in §5.4. We will also show in §5.4 that the presence of ℓ-diamonds induces collisions in the genus two extension of the isogeny hash function.

5.1 Morphisms to Subgroups

One of the key tools in studying isogenies between abelian varieties is the correspondence between subgroups and isogenies. This subsection explains the properties a subgroup needs to have in order to correspond to an appropriate isogeny.

Recall Theorem 1.7 which allows us to restrict our attention to Jacobians of hyperelliptic curves of genus two or some reducible product of two elliptic curves. We will prove this now.

Theorem 1.7. *If $A/\overline{\mathbb{F}}_p$ is a PPAS, then $A \cong J_H$ for some smooth (hyperelliptic) genus two curve H, or $A \cong E_1 \times E_2$ where E_i are elliptic curves.*

Proof. Use [GGR05, Thm. 3.1] which says that A is isomorphic over $\mathbb{F}_{p^n}$ (for some n) to the two cases in the theorem, or to the restriction of scalars of an elliptic curve over a quadratic extension of $\mathbb{F}_{p^n}$. Since we are working over $\overline{\mathbb{F}}_p$, this is absorbed into the second case. $\square$

Note that the definition of a maximal isotropic subgroup in §1.3 includes kernels of isogenies that factor through the multiplication-by-n map. We want to focus on isogenies that do not contain these "trivial" isogenies, hence we make the following definition:

Definition. A subgroup S of $A[m]$ is *good* if $A[n] \not\subseteq S$ for any $1 < n \leq m$.

We will base our results on good maximal isotropic subgroups for the rest of this section.

The following result illustrates the preservation of principal polarisations under isogenies whose kernels are isotropic.

Proposition 5.1. *Let H be a hyperelliptic curve of genus two over $\mathbb{F}_q$. Let K be a finite, good, $\mathbb{F}_q$-rational subgroup of $J_H(\overline{\mathbb{F}}_q)$. There exists a PPAS A over $\mathbb{F}_q$, and an isogeny $\phi : J_H \to A$ with kernel K, if and only if K is a maximal m-isotropic subgroup of $J_H[m]$ for some positive integer m.*

Proof. ($\Leftarrow$) The quotient $J_H \to J_H/K$ always exists as an isogeny between abelian varieties [Ser88, III.3.12]. Since J_H is the Jacobian of a hyperelliptic curve, it has a principal polarisation λ. Now consider the polarisation $\mu = [\deg \phi] \circ \lambda$ on J_H, then we certainly have $K = \ker \phi \subseteq \ker \mu$, and since K is isotropic, we use [Mil86a, Thm. 16.8] to get a polarisation λ' on J_H/K. Using [Mil86a, Rem. 16.9], we have that $\deg \lambda' = 1$ and so J_H/K is a PPAS.

By Theorem 1.7, we have that A is the Jacobian of a hyperelliptic curve of genus two or a product of two elliptic curves.

($\Rightarrow$) This is a simple application of [Mil86a, Thm. 16.8]. $\qquad\square$

Using the results above, we can focus on the type of subgroups of the torsion group that correspond to the isogenies we would like to investigate. We will denote by C_n the cyclic group of order n, and use e to denote the group identity.

Lemma 5.2. *Let A be a PPAS. If K is a good maximal ℓ^n-isotropic subgroup, then it cannot be cyclic.*

Proof. Suppose that K is cyclic, then the pairing is trivial on K due to the alternating property of the Weil pairing. But it can then be shown that there is an isotropic subgroup isomorphic to $C_{\ell^n}^2$ that contains K. Indeed, such a subgroup must exist: Suppose that $\langle P \rangle = K$, then let $Q \in A[\ell^n]$, such that $\langle P \rangle \cap \langle Q \rangle = \{e\}$, then if $e(P, Q) = 1$, we are done. Otherwise, suppose that $e(P, Q) = \mu^d \neq 1$. One can show that there exists some Q with the properties above such that d is co-prime to ℓ^n, hence, taking $f \equiv d^{-1} \mod \ell^n$, we have that $e(P, [f]Q) = 1$, and so $K \subset \langle P, [f]Q \rangle \cong C_{\ell^n}^2$ and is isotropic. $\qquad\square$

Proposition 5.3. *Let A be a PPAS. Then the good maximal ℓ^n-isotropic subgroups of $A[\ell^n]$ are isomorphic to*

$$C_{\ell^n} \times C_{\ell^n} \quad or \quad C_{\ell^n} \times C_{\ell^{n-k}} \times C_{\ell^k}$$

where $1 \leq k \leq \lfloor n/2 \rfloor$.

Proof. We see, from Lemma 5.2 and the fact that the maximal isotropic subgroups must be good, that K must have rank 2 or 3. If K has rank 2, then it can be shown that to be maximal, K must have the structure $C_{\ell^n} \times C_{\ell^n}$ by repeated inclusions.

Let $C_{\ell^a} \times C_{\ell^b} \times C_{\ell^c} \times C_{\ell^d}$ be a subgroup of $A[\ell^n]$. To simplify notation, we write this as $[a, b, c, d]$. Without loss of generality, we can take $a \geq b \geq c \geq d$. Then we have that the dual is $[n-a, n-b, n-c, n-d]$ (since the composition with the original isogeny is multiplication-by-ℓ^n) and $n-a \leq n-b \leq n-c \leq n-d$. Hence to get the symmetry as specified by Theorem 1.3, we must have that $n - a = d$ and $n - b = c$. Since we must have that one of the indices is zero, we take $d = 0$ and the result follows. $\qquad\square$

This result narrows down the subgroups that we need to study in order to understand sequences of (ℓ, ℓ)-isogenies between PPASs.

5.2 Number of Neighbours

In this section, we will consider the structure of an (ℓ, ℓ)-isogeny graph, $\mathcal{G}_{p^n, \ell}$.

Definition. Let p and ℓ be distinct primes, then the (ℓ, ℓ)-isogeny graph $\mathcal{G}_{p^n, \ell}$ is the directed graph whose vertices are isomorphism classes of PPASs over the algebraic closure of the field, and edges (A_1, A_2) are present between two PPAS A_1, A_2 if there is an (ℓ, ℓ)-isogeny $\phi : A_1 \to A_2$.

We begin by investigating the number of neighbours that each vertex is connected to. We approach this task by choosing an arbitrary PPAS and considering isogenies emanating from this surface. Then the nascent isogeny graph is a rooted graph at the chosen surface. Our first theorem counts the number of elements n steps from the root.

Theorem 5.4. *Let A be a PPAS, ℓ be a prime different from p and $n > 2$. Then the number of good ℓ^n-maximal isotropic subgroups of $A[\ell^n]$ is*

$$\ell^{2n-3}(\ell^2+1)(\ell+1)\left(\ell^n + \ell\frac{\ell^{n-2}-1}{\ell-1}+1\right)$$

if n is even, and

$$\ell^{2n-3}(\ell^2+1)(\ell+1)\left(\ell^n + \frac{\ell^{n-1}-1}{\ell-1}\right)$$

if n is odd.

The proof of the theorem builds on the number of subgroups of the torsion subgroup (cf. Proposition 5.6). We then count the number of maximal isotropic subgroups since these form the kernels of the isogenies that preserve principal polarisations (cf. Proposition 5.5). The proof of the theorem then follows by summing the number of maximal isotropic subgroups.

Proposition 5.5. *Let A be a PPAS. Let $N(a, b, c)$ be the number of good maximal isotropic subgroups of A isomorphic to $C_{\ell^a} \times C_{\ell^b} \times C_{\ell^c}$. Then*

(a) $N(n, n-a, a) = \ell^{3n-2a-4}(\ell^2+1)(\ell+1)^2$, *where $1 \le a < n/2$;*

(b) $N(n, n, 0) = \ell^{3n-3}(\ell^2+1)(\ell+1)$;

(c) $N(2k, k, k) = \ell^{4k-3}(\ell^2+1)(\ell+1)$.

Before imposing the isotropic condition on the subgroups via pairings, it is useful to know the number of subgroups of a particular structure in the torsion group. The following proposition enumerates them.

Proposition 5.6. *Let $S(a, b, c, d)$ be the number of subgroups of $C_{\ell^n}^4$ which are isomorphic to $C_{\ell^a} \times C_{\ell^b} \times C_{\ell^c} \times C_{\ell^d}$. Then*

(a) $S(n, n-a, a, 0) = \ell^{4n-2a-6}(\ell^2+1)(\ell+1)^2(\ell^2+\ell+1)$, *where $1 \le a < n/2$;*

(b) $S(n, n, 0, 0) = \ell^{4n-4}(\ell^2+1)(\ell^2+\ell+1)$;

(c) $S(2k, k, k, 0) = \ell^{6k-5}(\ell^2+1)(\ell+1)(\ell^2+\ell+1)$.

To give a flavour of our strategy, we will prove the second case in Proposition 5.5. The proof will show us the ingredients needed for the rest of the cases.

Proposition 5.7. *Let A be a PPAS. Let $N(a, b, c)$ be the number of good maximal isotropic subgroups of A isomorphic to $C_{\ell^a} \times C_{\ell^b} \times C_{\ell^c}$. Then $N(n, n, 0) = \ell^{3n-3}(\ell^2+1)(\ell+1)$.*

Proof. Note that this is equivalent to finding a subgroup isomorphic to $C_{\ell^n}^2$ in $A[\ell^n] \cong C_{\ell^n}^4$ which satisfies the isotropic condition.

So we need to find 2 elements in $C_{\ell^n}^4$ that have full order, are isotropic under the Weil pairing and generate subgroups with trivial intersection. To make things concrete, let $\langle P_1, \ldots, P_4 \rangle = C_{\ell^n}^4$. Let us pick the first element $X \in C_{\ell^n}^4$. This involves picking a full order element in $C_{\ell^n}^4$ for which we have $\ell^{4n} - \ell^{4n-4}$ choices. Let $X = \sum [a_i] P_i$.

To pick the second element $Y \in C_{\ell^n}^4$, we need to pick a full order element but also ensure that Y is isotropic to X under the Weil pairing. If we write $Y = \sum [b_i] P_i$, then we require that

$$e_\ell(X, Y) = e_\ell(P_1, P_2)^{a_1 b_2 - a_2 b_1} \cdot e_\ell(P_1, P_3)^{a_1 b_3 - a_3 b_1} \cdot e_\ell(P_1, P_4)^{a_1 b_4 - a_4 b_1}$$
$$\cdot e_\ell(P_2, P_3)^{a_2 b_3 - a_3 b_2} \cdot e_\ell(P_2, P_4)^{a_2 b_4 - a_4 b_2} \cdot e_\ell(P_3, P_4)^{a_3 b_4 - a_4 b_3}$$
$$= 1 \,.$$

But this is a linear condition on the selection of the b_i's. Thus this gives us $\ell^{3n} - \ell^{3n-3}$ choices[10]. But we need to pick Y such that $\langle X \rangle \cap \langle Y \rangle = \{e\}$. In other words, since Y has order ℓ^n, we need that $\ell^{n-1}Y \notin \langle X \rangle$.

Given that Y has full order, we need to avoid $(\ell - 1)\ell^{3(n-1)}$ elements. Hence the total number of choices for Y is
$$\ell^{3n} - \ell^{3(n-1)} - (\ell - 1)\ell^{3(n-1)}.$$

Now, we need to divide the choices we have for X and Y by the number of generating pairs in a subgroup $C_{\ell^n}^2$. The total number of generating pairs is $(\ell^{2n} - \ell^{2(n-1)})(\ell^{2n} - \ell^{2(n-1)} - (\ell - 1)\ell^{2(n-1)})$. Hence the total number of maximal isotropic $C_{\ell^n}^2$ subgroups of $C_{\ell^n}^4$ is

$$\frac{(\ell^{4n} - \ell^{4n-4})(\ell^{3n} - \ell^{3(n-1)} - (\ell - 1)\ell^{3(n-1)})}{(\ell^{2n} - \ell^{2(n-1)})(\ell^{2n} - \ell^{2(n-1)} - (\ell - 1)\ell^{2(n-1)})} = \ell^{3n-3}(\ell^2 + 1)(\ell + 1).$$

$\square$

The over-arching strategy for counting the rest of the cases is the same as in Proposition 5.7. We first find the number of choices to pick generators for the subgroup, then we quotient that number by the number of ways the same subgroup can be generated by different generators. In other words, we need to find the following:

$$\frac{\text{Number of ways to generate the subgroup}}{\text{Number of generators generating same subgroup}}.$$

Then to find the number of maximal isotropic subgroups, we note (as we did in the proof of Proposition 5.7) that the isotropic condition is a linear constraint. This goes into reducing the count for the subsequent generators.

Using this strategy, we will prove the remaining two cases from Proposition 5.6.

Proposition 5.8. *Let $G = (C_{\ell^n})^4$. If $1 \leq a < n/2$, then the number of subgroups isomorphic to $H = C_{\ell^n} \times C_{\ell^{n-a}} \times C_{\ell^a}$ is*
$$\ell^{4n-2a-6}(\ell^2 + 1)(\ell + 1)^2(\ell^2 + \ell + 1).$$

We will first need to prove a result on the number of elements of varying orders in H.

Lemma 5.9. *With G and H as in the above proposition, we have the following:*

(a) *the number of elements of order ℓ^n in H is $\ell^{2n-1}(\ell - 1)$,*

(b) *the number of elements of order ℓ^{n-a} in H is $\ell^{2n-a-2}(\ell^2 - 1)$,*

(c) *the number of elements of order ℓ^a in H is $\ell^{3a} - \ell^{3(a-1)}$.*

Proof.   (a) The element of order ℓ^n can only come from the first component $C_{\ell^n} \subseteq H$, and there are $\ell^n - \ell^{n-1}$ of them. We then have a free choice of the second and third generators which means we have ℓ^n choices for them. Hence the result follows.

(b) The element of order ℓ^{n-a} cannot come from the last component $C_{\ell^a} \subseteq H$. So it can either come from the first or second component $C_{\ell^n}, C_{\ell^{n-a}} \subseteq H$.

Suppose that it came from the first component $C_{\ell^n} \subseteq H$, then that is a total of $\ell^{n-a} - \ell^{n-a-1}$ choices for the first generator. We then have a free choice on the second and third generators, hence the number of choices in this case is $(\ell^{n-a} - \ell^{n-a-1})\ell^{n-a}\ell^a$.

If the element of order ℓ^{n-a} came from the second component $C_{\ell^{n-a}} \subseteq H$, then we need that the rest of the components do not have order ℓ^{n-a} to avoid double counting. Hence the number of choices in this case is $\ell^{n-a-1}(\ell^{n-a} - \ell^{n-a-1})\ell^a$.

The result follows when we sum the two cases.

---

[10]To see this, note that each $e_\ell(P_i, P_j) = \mu^{\alpha_{i,j}}$, where μ is an ℓ-root of unity and $\alpha_{i,j}$ is some integer. We can express the isotropic condition as

$$b_4(\alpha_{1,4}a_1 + \alpha_{2,4}a_2 + \alpha_{3,4}a_3) \equiv \begin{array}{l} \alpha_{1,2}(a_2b_1 - a_1b_2) + \alpha_{1,3}(a_3b_1 - a_1b_3) \\ +\alpha_{2,3}(a_3b_2 - a_2b_3) + \alpha_{1,4}a_4b_1 \\ +\alpha_{2,4}a_4b_2 + \alpha_{3,4}a_4b_3 \end{array} \pmod{\ell}.$$

In the case where $(\alpha_{1,4}a_1 + \alpha_{2,4}a_2 + \alpha_{3,4}a_3) \not\equiv 0$, we have free choices for b_1, b_2, b_3 (not all divisible by ℓ) and so have $\ell^{3n} - \ell^{3n-3}$ choices.

(c) This is straightforward.

$\square$

Proof of Prop. 5.8. To count the numerator, we need to count the number of ways to choose each of the three components. The number of choices for the first component is straightforward: it is given by the number of elements of order ℓ^n in G and this is just $\ell^{4n} - \ell^{4(n-1)}$. For the second generator, we need to count all the elements of order ℓ^{n-a} that are linearly independent to the first generator. That is, denoting the first generator as A and the generator to be chosen as B, we need to ensure that $\langle A \rangle \cap \langle B \rangle = \{e\}$. In other words, since B has order ℓ^{n-a}, we need that $\ell^{n-a-1}B \notin \langle A \rangle$. Notice that $\ell^{n-a-1}B$ has order ℓ and there are only $\ell - 1$ elements of order ℓ in $\langle A \rangle$ and there are ℓ^{n-a-1} choices of B lying above $\ell^{n-a-1}B$. The last generator is chosen similarly but the condition is now $\langle C \rangle \cap \langle A, B \rangle = \{e\}$.

As for the denominator, the subtlety is getting the number of elements of the correct order in H. This has been dealt with in Lemma 5.9.

The first term is the number of elements of order ℓ^n in H which is equal to $\ell^n(\ell^n - \ell^{n-1})$.

The second term is the number of elements of order ℓ^{n-a} in H that are linearly independent to the first generator. Let our second generator be B and the first generator be A. This is similar to the numerator in that we need $\ell^{n-a-1}B \notin \langle A \rangle$. There are still $\ell - 1$ elements of order ℓ in $\langle A \rangle$, but there are now $(\ell^{n-a-1})^2\ell^a$ choices of B lying above $\ell^{n-a-1}B$. Indeed, there are ℓ^{n-a-1} choices in both the first and second components and ℓ^a choices in the last.

The last term is similar to the numerator.

Putting everything together, one can check that the count is

$$\frac{\left[\ell^{4n} - \ell^{4(n-1)}\right]\left[\ell^{4(n-a)} - \ell^{4(n-a-1)} - (p-1)p^{4(n-a-1)}\right]\left[\ell^{4a} - \ell^{4(a-1)} - (\ell^2 - 1)\ell^{4(a-1)}\right]}{\left[\ell^n(\ell^n - \ell^{n-1})\right]\left[\ell^{2n-a-2}(\ell^2 - 1) - (\ell-1)(\ell^{n-a-1})^2\ell^a\right]\left[\ell^{3a} - \ell^{3(a-1)} - (\ell^2 - 1)\ell^{3(a-1)}\right]}$$

and the result follows. $\square$

Proposition 5.10. *Let $G = (C_{\ell^{2k}})^4$, then the number of subgroups isomorphic to $C_{\ell^{2k}} \times C_{\ell^k} \times C_{\ell^k}$ is*

$$\ell^{6k-5}(\ell^2 + 1)(\ell + 1)(\ell^2 + \ell + 1).$$

Proof. The number of choices for the first generator is the number of elements with order ℓ^{2k}: there are $\ell^{4(2k)} - \ell^{4(2k-1)}$ of them.

The number of choices for the second generator is the number of elements of order ℓ^k $\left(= \ell^{4(k)} - \ell^{4(k-1)}\right)$ minus those that generate subgroups that have non-trivial intersection with the subgroup generated by the first generator $\left(= (\ell-1)\ell^{4(k-1)}\right)$.

The number of choices for the third generator is also the number of elements of order ℓ^k $\left(= \ell^{4(k)} - \ell^{4(k-1)}\right)$, minus those which are linearly dependent to the first and second generators. Indeed, fix the generators chosen for the first two components to be A and B. We need to subtract from the number of order ℓ^k elements, the number elements C of order ℓ^k such that $[\ell^{k-1}]C \in \langle A, B \rangle$. Then we note that there are $\ell^2 - 1$ elements of order ℓ in $\langle A, B \rangle$. And there are $\ell^{4(k-1)}$ elements lying over $\ell^{4(k-1)}C$.

Hence the total number of choices for the generators is:

$$\left[\ell^{4(2k)} - \ell^{4(2k-1)}\right]\left[\ell^{4(k)} - \ell^{4(k-1)} - (\ell-1)\ell^{4(k-1)}\right]\left[\ell^{4(k)} - \ell^{4(k-1)} - (\ell^2-1)\ell^{4(k-1)}\right].$$

As for the denominator, the number of elements of order ℓ^k in $C_{\ell^{2k}} \times C_{\ell^k} \times C_{\ell^k}$ is given by

$$\ell^{3k-3}(\ell-1)(\ell^2 + \ell + 1).$$

The rest of the proof then follows from the denominator taking the form

$$\left[(\ell^{2k} - \ell^{2k-1})(\ell^k)(\ell^k)\right]\left[\ell^{3k-3}(\ell-1)(\ell^2 + \ell + 1) - (\ell-1)(\ell^{3(k-1)})\right]^2.$$

$\square$

Proposition 5.11. *The number of maximal isotropic subgroups isomorphic to $C_{\ell^{2k}} \times C_{\ell^k} \times C_{\ell^k}$ is*

$$\ell^{4k-3}(\ell^2 + 1)(\ell + 1).$$

Proof. We now know how to obtain the number of subgroups isomorphic to $C_{\ell^{2k}} \times C_{\ell^k} \times C_{\ell^k}$ using the preceding result. The proof then follows from making the same observation as in Proposition 5.7 that the isotropic condition is a linear constraint on the second and third generators.

This leads to the value

$$\frac{\left[\ell^{4(2k)} - \ell^{4(2k-1)}\right]\left[\ell^{3(k)} - \ell^{3(k-1)} - (\ell-1)\ell^{3(k-1)}\right]\left[\ell^{3(k)} - \ell^{3(k-1)} - (\ell^2-1)\ell^{3(k-1)}\right]}{\left[(\ell^{2k} - \ell^{2k-1})(\ell^k)(\ell^k)\right]\left[\ell^{3k-3}(\ell-1)(\ell^2 + \ell + 1) - (\ell-1)(\ell^{3(k-1)})\right]^2},$$

which one can check is equal to

$$\ell^{3(n-1)}(\ell^2 + 1)(\ell + 1).$$

$\square$

We will omit the last proof since it follows from the above methods and the equations in Proposition 5.6.

5.3 Number of Paths Between Two Vertices

Let A be a PPAS, and let K be a finite subgroup of order ℓ^{2n}. In this section, we are interested in the paths between A and A/K in $\mathcal{G}_{\overline{\mathbb{F}}_p, \ell}$ which have length n. Hence, for the remainder of this section, paths between vertices mean paths between such an $(A, A/K)$ pairing of length n.

Now, suppose we have an isogeny which has a maximal isotropic kernel K with order ℓ^{2n}, then we can decompose this isogeny into a sequence of n (ℓ, ℓ)-isogenies:

$$A \xrightarrow{\phi_1} A_1 \xrightarrow{\phi_2} A_2 \xrightarrow{\phi_3} \ldots \xrightarrow{\phi_n} A/K.$$

As mentioned in the introduction, this decomposition of isogenies may be non-unique. The non-uniqueness arises from kernels whose structure allows for more than one subgroup isomorphic to $C_\ell \times C_\ell$. The key observation is that these subgroups form the kernels of ϕ_1. In that spirit, the next two lemmata will give properties for the kernels of the first isogeny.

Lemma 5.12. *Let A be a PPAS. Let K be a maximal isotropic subgroup of $A[\ell^n]$ which is isomorphic to $C_{\ell^n} \times C_{\ell^{n-a}} \times C_{\ell^a}$ for some $a \geq 0$. Let $\langle P, Q, R \rangle = K$ such that P, Q, R have orders $\ell^n, \ell^{n-a}, \ell^a$ respectively.*

(a) Let $P_i, Q_i, R_i \in A_i$ be elements mapped from $P = P_0$, $Q = Q_0$, $R = R_0$ under the sequence of isogenies as stated above. Then $[\ell^{n-i-1}]P_i \in \ker \phi_{i+1}$ for all $i \geq 0$.

(b) The first (ℓ, ℓ)-isogeny must have kernel

$$\langle [\ell^{n-1}]P, [\ell^{n-a-1}]Q + [k][\ell^{a-1}]R \rangle \quad \text{for } 0 \leq k \leq \ell - 1, \quad \text{or} \quad \langle [\ell^{n-1}]P, [\ell^{a-1}]R \rangle.$$

Proof. (a) One can show by contradiction that if there is a kernel not containing P_i, then we will have cyclic kernels, which cannot be a kernel of a (ℓ, ℓ)-isogeny by Lemma 5.2.

Next, let $P' \in \langle P_i \rangle$, $Q' \in \langle Q_i \rangle$, and $R' \in \langle R_i \rangle$ such that P', Q', R' all have order ℓ. Then kernels cannot be of the form $P' + Q', P' + R', Q' + R'$. Indeed, it can be shown by examining the pairing $e_\ell(P' + Q', P' + R')$ to see that one either obtains a cyclic kernel, or that the subgroup above is not isotropic.

(b) We have from the first part that $[\ell^{n-1}]P$ must be a generator of the group. The second generator must be chosen from the remaining points of order ℓ. By the isotropic condition of K, we have that they are all trivial on the pairing as well.

$\square$

Lemma 5.13. *Let $G \cong C_{\ell^n} \times C_{\ell^{n-a}} \times C_{\ell^a}$ and H be abelian groups. Let $\langle P, Q, R \rangle = K$ such that P, Q, R have orders $\ell^n, \ell^{n-a}, \ell^a$ respectively. If $\phi : G \to H$ is a group homomorphism, with*

$$\ker \phi = \left\langle [\ell^{n-1}]P, [\ell^{n-a-1}]Q + [k][\ell^{a-1}]R \right\rangle$$

for $1 \leq k \leq \ell - 1$ and $a \leq n/2$, then $H \cong C_{\ell^{n-1}} \times C_{\ell^{n-a}} \times C_{\ell^{a-1}}$.

Proof. We have that $\phi(P)$ has order ℓ^{n-1} and Q has order ℓ^{n-a}, since $[\ell^{n-a-1}]Q \notin \ker \phi$. Since the order of the kernel is ℓ^2, we must have that $H \cong C_{\ell^{n-1}} \times C_{\ell^{n-a}} \times C_{\ell^{a-1}}$. $\hfill\square$

We can now study the different isogenies that exist between two vertices on the graph. In particular, we will be counting the number of different paths between two vertices on the graph.

Consider the simple cases first, where there is only one path between two vertices, or where two vertices are separated by two (ℓ, ℓ)-isogenies.

Proposition 5.14. *Let A be a PPAS, and let $K \cong (C_{\ell^n} \times C_{\ell^{n-a}} \times C_{\ell^a})$. Let $P(n,a)$ be the number of paths from A to A/K. Then*

(a) $P(n,0) = 1$ for all n;

(b) $P(2,1) = \ell + 1$.

Proof. (a) Since kernels of (ℓ, ℓ)-isogenies cannot be cyclic, the only possible subgroup of order ℓ^2 of $C_{\ell^n} \times C_{\ell^n}$ is $C_\ell \times C_\ell$, and there is only one choice for this subgroup.

(b) Let $K = C_{\ell^2} \times C_\ell \times C_\ell$. Then from Lemma 5.12 (and using its notation) we must have that the first isogeny has kernel

$$\langle [\ell]P, Q + [k]R \rangle \quad \text{for } 0 \le k \le \ell - 1, \quad \text{or} \quad \langle [\ell]P, R \rangle.$$

There are $\ell + 1$ choices for the first kernel. Thereafter, there is only one choice for the second kernel and so we have a total of $\ell + 1$ paths.

$\hfill\square$

Now, we can prove the general case.

Theorem 5.15. *Using the notation above, where $P(n,a)$ is the number of paths in a $(C_{\ell^n} \times C_{\ell^{n-a}} \times C_{\ell^a})$-isogeny. Then $P(n,a)$ satisfies the following recursive equation:*

$$P(n,a) = 2P(n-1, a-1) + (\ell - 1)P(n-1, a),$$

where $1 \le a < n/2$, and with the following boundary conditions:

$$P(n,0) = 1, \quad P(2,1) = \ell + 1.$$

Proof. We will prove this by induction. The base cases of the induction steps are easy and the boundary conditions follow from Proposition 5.14. We will show the induction step.

Let us suppose that the recursive formula holds for $P(n-1, a-1)$ and $P(n-1, a)$. Now, suppose that our kernel is isomorphic to $C_{\ell^n} \times C_{\ell^{n-a}} \times C_{\ell^a}$. Since each (ℓ, ℓ)-isogeny has a kernel of the form $C_\ell \times C_\ell$, we have, from Lemma 5.12(2), that the first isogeny must have a kernel of the form

$$\langle [\ell^{n-1}]P, [\ell^{n-a-1}]Q + [k][\ell^{a-1}]R \rangle \quad \text{for } 0 \le k \le \ell - 1, \quad \text{or} \quad \langle [\ell^{n-1}]P, [\ell^{a-1}]R \rangle.$$

It is clear that if the kernel is given by

$$\langle [\ell^{n-1}]P, [\ell^{n-a-1}]Q \rangle \quad \text{or} \quad \langle [\ell^{n-1}]P, [\ell^{a-1}]R \rangle,$$

then the residual kernel will be of the form

$$C_{\ell^{n-1}} \times C_{\ell^{n-a-1}} \times C_{\ell^a} \text{ or } C_{\ell^{n-1}} \times C_{\ell^{n-a}} \times C_{\ell^{a-1}}$$

respectively. Otherwise, if the first kernel has the form

$$\langle [\ell^{n-1}]P, [\ell^{n-a-1}]Q + [k][\ell^{a-1}]R \rangle \quad \text{for } 1 \le k \le \ell - 1,$$

the residual kernel will be of the form $C_{\ell^{n-1}} \times C_{\ell^{n-a}} \times C_{\ell^{a-1}}$ by Lemma 5.13. Hence we are done. $\hfill\square$

Proposition 5.14 actually shows us the different paths that can exist between vertices in the graph. In particular, for kernels with rank 2, there can only be a single path between the domain and codomain. However, for kernels with rank 3, there can be a multitude of paths that exist between the domain and codomain. It can be seen that the following shapes (ℓ-diamonds) are the basic paths drawn out by kernels with group structure $C_{\ell^2} \times C_\ell \times C_\ell$ for different ℓ's.

$$\ell = 2 \qquad\qquad \ell = 3 \qquad\qquad \ell = 5 \qquad\qquad \ell = 7$$

The non-uniqueness of these paths can be seen more explicitly in this example, where we will consider kernels with order 256. The key to each example is to the find the number of $C_2 \times C_2$ subgroups of each kernel since this would correspond with the number of possible $(2,2)$-isogenies. Firstly, we note that the structure of maximal isotropic subgroups of order 256 must be $C_{16} \times C_{16}$, or $C_{16} \times C_4 \times C_4$, or $C_{16} \times C_8 \times C_2$ by Proposition 5.3. The easy case is when the kernel K_0 has the structure $C_{16} \times C_{16}$. This is because there is only one $C_2 \times C_2$ subgroup in K. Hence, there is only one isogeny path available and we have a straight line. The other isogeny graphs are given in Figure 5.1.

Now, let us consider the case when the kernel K_1 has the structure $C_{16} \times C_4 \times C_4$. We will label the isomorphism classes of the surfaces by (n), where n is a natural number. We will denote the first surface by (1).

We can represent the 3 generators of K_1 by P, Q and R, where their orders are 16, 4 and 4 respectively. There are 3 different $C_2 \times C_2$ subgroups of K given by $\langle [8]P, [2]Q \rangle$, $\langle [8]P, [2]R \rangle$ and $\langle [8]P, [2](Q+R) \rangle$ in accordance to Lemma 5.12. Hence, we can and will denote the $(2,2)$-subgroups of K by the scalar preceding Q and R. For instance, the three subgroups given here are denoted by $(2,0)$, $(0,2)$ and $(2,2)$.

These 3 subgroups lead to non-isomorphic surfaces labelled as $(2), (3)$ and (4). The edges are labelled by the subgroup corresponding to the isogeny.

Consider the vertex (2), and consider the $(2,2)$-isogeny from (2) with kernel[11] $\langle [4]P, [2]R \rangle$ and denote the codomain by (8). One can see that the isogeny from (1) to (8) has kernel $\langle [4]P, [2]Q, [2]R \rangle$.

One can also map from (3) and (4) to (8) via the kernels $(2,0)$ and $(2,0)$. Immediately, one can spot the diamonds mentioned prior to this example. Indeed, the diamonds can be seen repeatedly in the graph.

Vertices can form tips of the diamond when there is a $C_4 \times C_2 \times C_2$ subgroup in the kernel. This is best illustrated in the next example where the kernel K_2 has structure $C_{16} \times C_8 \times C_2$. Using the notation from the previous example, K_2 will be given by $\langle P', Q', R' \rangle$, where $P' = P$, $[2]Q' = Q$ and $R' = [2]R$

Starting from the vertex (1) again, we have the same 3 subgroups, which result in the same surfaces $(2), (3)$ and (4). We also have that the three surfaces will all have maps into (8) as before. However, the residual kernel at (2) is now isomorphic to $C_8 \times C_8$, hence we see that the isogeny path from (2) down to (18) is a straight line. The residual kernel at (4) on the other hand, is $C_8 \times C_4 \times C_2$, hence it contains $C_4 \times C_2 \times C_2$ as a subgroup and so, (4) forms the tip of another diamond.

Another thing to note about this case is that the moment R is in the kernel, we cannot have $C_4 \times C_2 \times C_2$ as a subgroup of the residual kernel. This can be observed from the diagonal right-to-left lines in Figure 5.1b.

Lastly, Figure 5.2 shows all the neighbours which are two $(2,2)$-isogenies away. So the top vertex is connected to each of the middle and bottom vertices by an isogeny of degree 4 and 16 respectively. The diamonds corresponding to kernels with the structure $C_4 \times C_2 \times C_2$, (though contorted) are present and their number is as predicted in Proposition 5.5.

5.4 Application to Cryptography

The discussion in the previous sections will be utilised to analyse the security of G2SIDH of §2.3 and the generalisation of the CGL hash function of §2.2 to genus two.

---

[11] Note that we actually mean $\langle [4]\phi(P), [2]\phi(R) \rangle$, where ϕ corresponds to the $(2,2)$-isogeny from (1). We will drop ϕ for ease of notation.

(a) Kernel has structure $C_{16} \times C_4 \times C_4$.　　　(b) Kernel has structure $C_{16} \times C_8 \times C_2$.

Figure 5.1: Isogeny subgraphs when the kernel has order 256.

Figure 5.2: Isogeny graph from an arbitrary vertex showing 2 layers of isogenies.

5.4.1 Collisions in the Genus Two Isogeny-based Hash Function

Recall the CGL hash function in §2.2. In the genus two case of the hash function, due to the additional isogenies available to a single vertex (15 as opposed to 3), it is hoped that one can achieve a higher security level with a smaller number of steps. In [Tak18], Takashima outlined an algorithm for obtaining a sequence of $(2, 2)$-isogenies without backtracking. He also implicitly suggested the generalisation of the CGL hash function to genus two. The genus two version of the CGL hash uses the input bits to traverse the $(2, 2)$-isogeny graph of PPSSASs. The algorithm begins at a pre-chosen PPSSAS and begins a walk based on the binary input to the algorithm. The walk on the graph is similar to the original CGL hash with a difference of an increased number of paths at each iteration.

One of the main results of [CLG09] is the proof that the CGL hash function is collision resistant. The vague intuition for this is that the supersingular elliptic curve isogeny graph is locally tree-like, i.e. there are no small cycles in a small enough subgraph. This assumption fails in the genus two case as any diamond configuration leads to a collision in the hash. An attacker can find two pairs of inputs so that the walks collide. Using the 2-diamond as an example (see Fig. 5.3), where a hash is performed by walking along the left-most path. An attacker, with the knowledge that the hash has traversed through a diamond, will be able to choose either the middle path or the right-most path to achieve a collision.

In terms of endomorphisms, the collision resistance in the CGL hash is achieved by the lack of endomorphisms of degree 2^k, where k is small, in the graph. However, as we have seen in the previous section, we might be able to find endomorphism of degree 16 (or cycles of length 4) after 2 iterations of the genus two hash.

A recent paper by Castryck, Decru and Smith [CDS19] remedied this problem by restricting the isogeny paths to not include ℓ-diamonds. Setting $\ell = 2$, they were able to work with Richelot isogenies and restrict the quadratic splittings that would result in ℓ-diamonds.

Figure 5.3: Example of a collision in the genus two hash.

5.4.2 Security and Analysis of G2SIDH

We will first show that G2SIDH takes place over the field of $\mathbb{F}_{p^2}$ as is the case in SIDH with our choice of initial hyperelliptic curve. We will then tackle the security complexity of G2SIDH by using the results of Theorem 5.4.

Base Field

Recall the procedure of choosing a base hyperelliptic curve where we performed a random walk on the $(2,2)$-isogeny graph from the vertex corresponding to $H : y^2 = x^6 + 1$. Since there is a Richelot isogeny to a product of two supersingular elliptic curves, the next result tells us that H is in fact superspecial.

Lemma 5.16 ([Oor75])**.** *Let A be an abelian variety over a field of characteristic p and of dimension $g \geq 2$, and let $E^g \to A$ be an isogeny of degree d, where E is a supersingular elliptic curve. If $p \nmid d$ then $A \cong E^g$.*

Since G2SIDH employs $(2,2)$ and $(3,3)$-isogenies over a field of characteristic co-prime to 2 and 3, the PPASs in the protocol will be superspecial. The results in a paper by Ibukiyama and Katsura [IK94] state that every principally polarised supersingular abelian surface is isomorphic to one defined over $\mathbb{F}_{p^2}$.

Hence G2SIDH with the initial selection of a base hyperelliptic curve H takes place over the field of $\mathbb{F}_{p^2}$.

Security Estimates

In this section, we will define the computational problem needed to analyse our cryptosystem.

Let p be a prime of the form $2^n \cdot 3^n \cdot f - 1$, and fix a supersingular hyperelliptic curve of genus two H over $\mathbb{F}_{p^2}$ and let J_H denote its Jacobian. Fix bases for $J_H[2^n]$ and $J_H[3^m]$, denoting them by $\{P_i\}_{i=1,2,3,4}$ and $\{Q_i\}_{i=1,2,3,4}$ respectively.

Problem (Computational Genus Two Isogeny (CG2I) Problem)**.** *Let $\phi : J_H \to J_A$ be an isogeny whose kernel is given by K. Given J_A and the images $\{\phi(Q_i)\}$, $i \in \{1, 2, 3, 4\}$, find generators for K.*

The analogue problem in genus one has been conjectured to be hard. However, due to the higher regularity of the genus two isogeny graph, we are able to perform a smaller number of isogeny computations to achieve the same security level as compared to SIDH.

Let us look at the complexities of the algorithms one can employ against the CG2I problem, where the task is to recover the isogeny $\phi_A : J_H \to J_A$ when given J_H and J_A. We note that from Theorem 5.4, we have that the number of elements in the n-sphere is approximately $\ell^{3n} \approx \sqrt{p^3}$, hence a naïve exhaustive search on the leaves of J_H has a complexity of $O(\sqrt{p^3})$. One can improve on this by considering the meet-in-the-middle search by listing all isogenies of degree ℓ^n from J_H and J_A and finding collisions in both lists. The meet-in-the-middle search has a complexity of $O(\sqrt[4]{p^3})$. [12]. This compares favourably with the genus one case which has classical security of $O(\sqrt[4]{p})$, and quantum security of $O(\sqrt[6]{p})$. An example of a prime which one can use to achieve 128-bits of security is 171-bits, whereas the genus one case requires 512-bits for the same level of security.

The additional information provided by the presence of auxiliary points has not affected the security of SIDH aside from the two attacks detailed in Chapter 3, and the attack of [Pet17]. The security of the

---

[12]We note that a recent result by Jaques and Schanck [JS19] states that the Claw finding algorithm is expensive to implement in a realistic quantum computing model, and that attackers are better off employing Grover's algorithm which has a complexity of $O(\sqrt[4]{p^3})$.

auxiliary points of G2SIDH should be inherited from the security of those points in SIDH since attacks on one should be transferable to the other.

Due to the aforementioned reasons, we conjecture the CG2I problem to be computationally infeasible.

Existing Attacks on SIDH

Due to the similarity of G2SIDH and SIDH, many of the cryptanalysis papers on SIDH [GPST16, Ti17, GW17, Pet17, EHL$^+$18] may be applicable to G2SIDH. We can group the attacks into two classes: curves and points, and computing endomorphism rings.

Attacks on curves and points include the adaptive attack [GPST16] and fault attacks [Ti17, GW17]. Attacks via the computation of endomorphism rings include the methods using auxiliary points to find a subring of the endomorphism ring [Pet17] and using the Deuring correspondence [EHL$^+$18]. The purpose of computing the endomorphism ring is due to the result in [GPST16] that showed a reduction, in most cases, that the SIDH problem is at most as difficult as computing the endomorphism ring. The key observation behind this result is that the isogenies tend to be short paths in the graph, and so a lattice reduction performed on the basis of the connecting ideal would yield an element that corresponds to the secret isogeny via results in [KLPT14].

However, we will not be able to analyse these attacks in detail and will leave this as future work.

Chapter 6

Future Directions

Following the publication of [GPST16], a countermeasure to the adaptive attack was proposed in [AJL17]. To foil the adaptive attack outlined in Chapter 3, the authors suggested that both parties perform multiple instances of the key exchange before hashing the resulting output. We will briefly sketch this idea here. To ease the exposition, we will demonstrate how Alice performs the modified key exchange with two instances; it will be obvious how one can extend it to more instances.

The set up is exactly the same as in SIDH, so let $p = 2^n \cdot 3^m \cdot f - 1$ be a prime and let E be a supersingular elliptic curve over $\mathbb{F}_{p^2}$. We fix generators for the torsion groups as such: $\langle P_A, Q_A \rangle = E[2^n]$ and $\langle P_B, Q_B \rangle = E[3^m]$.

To perform the key exchange, Alice would pick $a_1^{(i)}, a_2^{(i)}$, where $i = 1$ or 2. Since we have two instances, we will use superscripts to denote which instance we are on. Alice computes her secret subgroups

$$G_A^{(i)} = \langle [a_1^{(i)}]P_A + [a_2^{(i)}]Q_A \rangle$$

for $i = 1, 2$. These will be the kernels to her secret isogenies $\phi_A^{(i)}$. She computes the codomains $E_A^{(i)}$ and the points P_B and Q_B under both isogenies, i.e. $\phi_A^{(i)}(P_B)$ and $\phi_A^{(i)}(Q_B)$ for $i = 1, 2$. She then sends the message

$$\left(E_A^{(i)}, \phi_A^{(i)}(P_B), \phi_A^{(i)}(Q_B) \right)_{i=1,2}$$

to Bob. Bob will perform his side of the protocol and sends his public key to Alice.

To compute the shared secret, Alice would take Bob's public key and computes

$$H_{i,j} = \left\langle \left[a_1^{(i)} \right] \phi_B^{(j)}(P_A) + \left[a_2^{(i)} \right] \phi_B^{(j)}(Q_A) \right\rangle$$

for $i, j = 1, 2$. She uses this to compute $z_{i,j}$, the j-invariant of the codomain of the isogeny whose kernel is $H_{i,j}$. The shared secret will be $\mathrm{Hash}(z_{1,1}, z_{1,2}, z_{2,1}, z_{2,2})$ using some preimage resistant hash function. This multiple instance of SIDH is called k-SIDH, where k is the number of instances. The above example is 2-SIDH. The authors have proposed letting $k = 113$ for a user working in $E[2^n]$ and $k = 94$ for one working in $E[3^m]$.

An interesting question is if an extended version of the adaptive attack can be made to work in the presence of this countermeasure. If the attack is incapacitated due to this modification, one would like to know what the lowest k needed to foil the attack. A low k (such as 2) would vastly improve the efficiency of this countermeasure and would allow users to have static keys.

Moving away from genus one, there are a number of open problems to be considered in genus two. There is a problem of efficient (ℓ, ℓ)-isogeny algorithms for $\ell > 3$. This has been addressed in [CR15, LR15, LR12, CE15], but an implementation[13] based on these ideas is not yet ready for cryptographic use.

The cryptanalysis of G2SIDH that we have left open at the end of Chapter 5 is yet another avenue of research one could explore. The extension of the adaptive attack to G2SIDH has parallels to SIDH with multiple instances described above. An interesting question is if the attack on one can carry over the other.

---

[13] A library written in `MAGMA` called `AVIsog` can be found in `http://avisogenies.gforge.inria.fr/`.

The generalisation of Deuring's correspondence in Chapter 4 to the case of PPSSASs is an extremely fascinating topic. This would open up the understanding of the structure of PPAS (ℓ, ℓ)-isogeny graphs, and also the possibility of group actions on PPSSASs over the field $\mathbb{F}_p$.

Lastly, the issue of connectedness of the PPSSAS (ℓ, ℓ)-isogeny graph has not been addressed in this thesis. A conjecture in [CDS19] postulates that the superspecial (ℓ, ℓ)-isogeny graph is connected. A resolution of this hypothesis can have implications on the security of G2SIDH.

Bibliography

[AD93]   Leonard M. Adleman and Jonathan DeMarrais, *A subexponential algorithm for discrete logarithms over all finite fields*, Math. Comp. **61** (1993), no. 203, 1–15.

[AJK⁺16]   Reza Azarderakhsh, David Jao, Kassem Kalach, Brian Koziel, and Christopher Leonardi, *Key compression for isogeny-based cryptosystems*, Proceedings of the 3rd ACM International Workshop on ASIA Public-Key Cryptography, AsiaPKC@AsiaCCS, Xi'an, China, May 30 - June 03, 2016, 2016, pp. 1–10.

[AJL17]   Reza Azarderakhsh, David Jao, and Christopher Leonardi, *Post-quantum static-static key agreement using multiple protocol instances*, Selected Areas in Cryptography - SAC 2017 - 24th International Conference, Ottawa, ON, Canada, August 16-18, 2017, Revised Selected Papers, 2017, pp. 45–63.

[Ajt96]   Miklós Ajtai, *Generating hard instances of lattice problems (extended abstract)*, Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996, 1996, pp. 99–108.

[AP15]   Jeffrey D. Achter and Rachel Pries, *Superspecial rank of supersingular abelian varieties and Jacobians*, J. Théor. Nombres Bordeaux **27** (2015), no. 3, 605–624.

[Apo76]   Tom M. Apostol, *Introduction to analytic number theory*, Springer-Verlag, New York-Heidelberg, 1976, Undergraduate Texts in Mathematics.

[BD11]   Nils Bruin and Kevin Doerksen, *The arithmetic of genus two curves with $(4,4)$-split Jacobians*, Canad. J. Math. **63** (2011), no. 5, 992–1024.

[BFT14]   Nils Bruin, E. Victor Flynn, and Damiano Testa, *Descent via $(3,3)$-isogeny on Jacobians of genus 2 curves*, Acta Arith. **165** (2014), no. 3, 201–223.

[BKL10]   Joppe W. Bos, Thorsten Kleinjung, and Arjen K. Lenstra, *On the use of the negation map in the pollard rho method*, Algorithmic Number Theory, 9th International Symposium, ANTS-IX, Nancy, France, July 19-23, 2010. Proceedings, 2010, pp. 66–82.

[BMM00]   Ingrid Biehl, Bernd Meyer, and Volker Müller, *Differential fault attacks on elliptic curve cryptosystems*, Advances in Cryptology - CRYPTO 2000, 20th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2000, Proceedings, 2000, pp. 131–146.

[BSS05]   Ian F. Blake, Gadiel Seroussi, and Nigel P. Smart, *Advances in elliptic curve cryptography*, London Mathematical Society Lecture Note Series, Cambridge University Press, Cambridge, UK, 2005.

[CDS19]   Wouter Castryck, Thomas Decru, and Benjamin Smith, *Hash functions from superspecial genus-2 curves using richelot isogenies*, IACR Cryptology ePrint Archive **2019** (2019), 296.

[CE15]   Jean-Marc Couveignes and Tony Ezome, *Computing functions on Jacobians and their quotients*, LMS J. Comput. Math. **18** (2015), no. 1, 555–577.

[CF96]   J. W. S. Cassels and E. V. Flynn, *Prolegomena to a middlebrow arithmetic of curves of genus 2*, London Mathematical Society Lecture Note Series, vol. 230, Cambridge University Press, Cambridge, 1996.

[CFA+12]   Henri Cohen, Gerhard Frey, Roberto Avanzi, Christophe Doche, Tanja Lange, Kim Nguyen, and Frederik Vercauteren, *Handbook of elliptic and hyperelliptic curve cryptography, second edition*, 2nd ed., Chapman & Hall/CRC, 2012.

[CJ05]   Mathieu Ciet and Marc Joye, *Elliptic curve cryptosystems in the presence of permanent and transient faults*, Des. Codes Cryptography **36** (2005), no. 1, 33–43.

[CJL+17]   Craig Costello, David Jao, Patrick Longa, Michael Naehrig, Joost Renes, and David Urbanik, *Efficient compression of SIDH public keys*, Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part I, 2017, pp. 679–706.

[CJS14]   Andrew M. Childs, David Jao, and Vladimir Soukharev, *Constructing elliptic curve isogenies in quantum subexponential time*, J. Mathematical Cryptology **8** (2014), no. 1, 1–29.

[CLG09]   Denis Xavier Charles, Kristin E. Lauter, and Eyal Z. Goren, *Cryptographic hash functions from expander graphs*, J. Cryptology **22** (2009), no. 1, 93–113.

[CLM+18]   Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes, *CSIDH: an efficient post-quantum commutative group action*, Advances in Cryptology - ASIACRYPT 2018 - 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2-6, 2018, Proceedings, Part III, 2018, pp. 395–427.

[CLN16]   Craig Costello, Patrick Longa, and Michael Naehrig, *Efficient algorithms for supersingular isogeny diffie-hellman*, Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part I, 2016, pp. 572–601.

[CNP05]   Gabriel Cardona, Enric Nart, and Jordi Pujolàs, *Curves of genus two over fields of even characteristic*, Math. Z. **250** (2005), no. 1, 177–201.

[Cos18]   Craig Costello, *Computing supersingular isogenies on kummer surfaces*, Advances in Cryptology - ASIACRYPT 2018 - 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2-6, 2018, Proceedings, Part III, 2018, pp. 428–456.

[Cou06]   Jean-Marc Couveignes, *Hard homogeneous spaces*, IACR Cryptology ePrint Archive **2006** (2006), 291.

[Cou07]   ——— , *Quelques mathématiques de la cryptologie á clés publiques*, Nouvelles méthodes mathématiques pour la cryptographie, Société mathématique de France, 2007.

[Cox13]   David A. Cox, *Primes of the form $x^2 + ny^2$*, second ed., Pure and Applied Mathematics (Hoboken), John Wiley & Sons, Inc., Hoboken, NJ, 2013, Fermat, class field theory, and complex multiplication.

[CR15]   Romain Cosset and Damien Robert, *Computing (ℓ, ℓ)-isogenies in polynomial time on Jacobians of genus 2 curves*, Math. Comp. **84** (2015), no. 294, 1953–1975.

[Deu41]   Max Deuring, *Die Typen der Multiplikatoren ringe elliptischer Funktionenkörper*, Abh. Math. Sem. Hansischen Univ. **14** (1941), 197–272.

[DG19]   Luca De Feo and Steven D. Galbraith, *Seasign: Compact isogeny signatures from class group actions*, Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part III, 2019, pp. 759–789.

[DH76]   Whitfield Diffie and Martin E. Hellman, *New directions in cryptography*, IEEE Trans. Information Theory **22** (1976), no. 6, 644–654.

[DJP14]   Luca De Feo, David Jao, and Jérôme Plût, *Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies*, J. Mathematical Cryptology **8** (2014), no. 3, 209–247.

[DKS18]   Luca De Feo, Jean Kieffer, and Benjamin Smith, *Towards practical key exchange from ordinary isogeny graphs*, Advances in Cryptology - ASIACRYPT 2018 - 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2-6, 2018, Proceedings, Part III, 2018, pp. 365–394.

[DMPS19]   Luca De Feo, Simon Masson, Christophe Petit, and Antonio Sanso, *Verifiable delay functions from supersingular isogenies and pairings*, IACR Cryptology ePrint Archive **2019** (2019), 166.

[EHL⁺18]   Kirsten Eisenträger, Sean Hallgren, Kristin E. Lauter, Travis Morrison, and Christophe Petit, *Supersingular isogeny graphs and endomorphism rings: Reductions and solutions*, Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part III, 2018, pp. 329–368.

[FO13]   Eiichiro Fujisaki and Tatsuaki Okamoto, *Secure integration of asymmetric and symmetric encryption schemes*, J. Cryptology **26** (2013), no. 1, 80–101.

[FR94]   Gerhard Frey and Hans-Georg Rück, *A remark concerning m-divisibility and the discrete logarithm in the divisor class group of curves*, Math. Comp. **62** (1994), no. 206, 865–874.

[FT19]   E. Victor Flynn and Yan Bo Ti, *Genus two isogeny cryptography*, Post-Quantum Cryptography - 10th International Conference, PQCrypto 2019, Chongqing, China, May 8-10, 2019 Revised Selected Papers, 2019, pp. 286–306.

[Gal12]   Steven D. Galbraith, *Mathematics of public key cryptography*, Cambridge University Press, Cambridge, 2012.

[Gau00]   Pierrick Gaudry, *An algorithm for solving the discrete log problem on hyperelliptic curves*, Advances in Cryptology - EUROCRYPT 2000, International Conference on the Theory and Application of Cryptographic Techniques, Bruges, Belgium, May 14-18, 2000, Proceeding, 2000, pp. 19–34.

[GGR05]   Josep González, Jordi Guàrdia, and Victor Rotger, *Abelian surfaces of GL_2-type as Jacobians of curves*, Acta Arith. **116** (2005), no. 3, 263–287.

[GHS02]   Pierrick Gaudry, Florian Hess, and Nigel P. Smart, *Constructive and destructive facets of weil descent on elliptic curves*, J. Cryptology **15** (2002), no. 1, 19–46.

[GPS17]   Steven D. Galbraith, Christophe Petit, and Javier Silva, *Identification protocols and signature schemes based on supersingular isogeny problems*, Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I, 2017, pp. 3–33.

[GPST16]   Steven D. Galbraith, Christophe Petit, Barak Shani, and Yan Bo Ti, *On the security of supersingular isogeny cryptosystems*, Advances in Cryptology - ASIACRYPT 2016 - 22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4-8, 2016, Proceedings, Part I, 2016, pp. 63–91.

[Gro87]   Benedict H. Gross, *Heights and the special values of L-series*, Number theory (Montreal, Que., 1985), CMS Conf. Proc., vol. 7, Amer. Math. Soc., Providence, RI, 1987, pp. 115–187.

[GW17]   Alexandre Gélin and Benjamin Wesolowski, *Loop-abort faults on supersingular isogeny cryptosystems*, Post-Quantum Cryptography - 8th International Workshop, PQCrypto 2017, Utrecht, The Netherlands, June 26-28, 2017, Proceedings, 2017, pp. 93–106.

[Har77]   Robin Hartshorne, *Algebraic geometry*, Springer-Verlag, New York-Heidelberg, 1977, Gradu-ate Texts in Mathematics, No. 52.

[HPS98]   Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman, *NTRU: A ring-based public key cryp-tosystem*, Algorithmic Number Theory, Third International Symposium, ANTS-III, Portland, Oregon, USA, June 21-25, 1998, Proceedings, 1998, pp. 267–288.

[HS00]   Marc Hindry and Joseph H. Silverman, *Diophantine geometry*, Graduate Texts in Mathe-matics, vol. 201, Springer-Verlag, New York, 2000, An introduction.

[IK94]   Tomoyoshi Ibukiyama and Toshiyuki Katsura, *On the field of definition of superspecial po-larized abelian varieties and type numbers*, Compositio Mathematica **91** (1994), no. 1, 37–46 (en). MR 1273924

[JD11]   David Jao and Luca De Feo, *Towards quantum-resistant cryptosystems from supersingular el-liptic curve isogenies*, Post-Quantum Cryptography - 4th International Workshop, PQCrypto 2011, Taipei, Taiwan, November 29 - December 2, 2011. Proceedings, 2011, pp. 19–34.

[JS14]   David Jao and Vladimir Soukharev, *Isogeny-based quantum-resistant undeniable signatures*, Post-Quantum Cryptography - 6th International Workshop, PQCrypto 2014, Waterloo, ON, Canada, October 1-3, 2014. Proceedings, 2014, pp. 160–179.

[JS19]   Samuel Jaques and John M. Schanck, *Quantum cryptanalysis in the RAM model: Claw-finding attacks on SIKE*, Advances in Cryptology - CRYPTO 2019 - 39th Annual Interna-tional Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part I, 2019, pp. 32–61.

[KLM+15]   Daniel Kirkwood, Bradley C. Lackey, John McVey, Mark Motley, Jerome A. Solinas, and David Tuller, *Failure is not an option: Standardization issues for post-quantum key agree-ment*, 2015, Workshop on Cybersecurity in a Post-Quantum World.

[KLPT14]   David Kohel, Kristin Lauter, Christophe Petit, and Jean-Pierre Tignol, *On the quaternion ℓ-isogeny path problem*, LMS J. Comput. Math. **17** (2014), no. suppl. A, 418–432.

[Kob87]   Neal Koblitz, *Elliptic curve cryptosystems*, Math. Comp. **48** (1987), no. 177, 203–209.

[Kob89]   Neal Koblitz, *Hyperelliptic cryptosystems*, J. Cryptology **1** (1989), no. 3, 139–150.

[Koh96]   David Russell Kohel, *Endomorphism rings of elliptic curves over finite fields*, Ph.D. thesis, 1996, Thesis (Ph.D.)–University of California, Berkeley, p. 117.

[Kuh88]   Robert M. Kuhn, *Curves of genus* 2 *with split Jacobian*, Trans. Amer. Math. Soc. **307** (1988), no. 1, 41–49.

[LL97]   Chae Hoon Lim and Pil Joong Lee, *A key recovery attack on discrete log-based schemes using a prime order subgroupp*, Advances in Cryptology - CRYPTO '97, 17th Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 1997, Proceedings, 1997, pp. 249–263.

[LR12]   David Lubicz and Damien Robert, *Computing isogenies between abelian varieties*, Compos. Math. **148** (2012), no. 5, 1483–1515.

[LR15]   ———, *Computing separable isogenies in quasi-optimal time*, LMS J. Comput. Math. **18** (2015), no. 1, 198–216.

[McE78]   R. J. McEliece, *A Public-Key Cryptosystem Based On Algebraic Coding Theory*, Deep Space Network Progress Report **44** (1978), 114–116.

[Mer79]   Ralph C. Merkle, *Secrecy, authentication, and public key systems*, Ph.D. thesis, Stanford University, 1979.

[MI88]   Tsutomu Matsumoto and Hideki Imai, *Public quadratic polynominal-tuples for efficient signature-verification and message-encryption*, Advances in Cryptology - EUROCRYPT '88, Workshop on the Theory and Application of of Cryptographic Techniques, Davos, Switzerland, May 25-27, 1988, Proceedings, 1988, pp. 419–453.

[Mil85]   Victor S. Miller, *Use of elliptic curves in cryptography*, Advances in Cryptology - CRYPTO '85, Santa Barbara, California, USA, August 18-22, 1985, Proceedings, 1985, pp. 417–426.

[Mil86a]   J. S. Milne, *Abelian varieties*, Arithmetic Geometry (Gary Cornell and Joseph H. Silverman, eds.), Springer New York, New York, NY, 1986, pp. 103–150.

[Mil86b]   _____, *Jacobian varieties*, Arithmetic Geometry (Gary Cornell and Joseph H. Silverman, eds.), Springer New York, New York, NY, 1986, pp. 167–212.

[Mum08]   David Mumford, *Abelian varieties*, Tata Institute of Fundamental Research Studies in Mathematics, vol. 5, Published for the Tata Institute of Fundamental Research, Bombay; by Hindustan Book Agency, New Delhi, 2008, With appendices by C. P. Ramanujam and Yuri Manin, Corrected reprint of the second (1974) edition.

[MVO91]   Alfred Menezes, Scott A. Vanstone, and Tatsuaki Okamoto, *Reducing elliptic curve logarithms to logarithms in a finite field*, Proceedings of the 23rd Annual ACM Symposium on Theory of Computing, May 5-8, 1991, New Orleans, Louisiana, USA, 1991, pp. 80–89.

[NS04]   Phong Q. Nguyen and Damien Stehlé, *Low-dimensional lattice basis reduction revisited*, Algorithmic Number Theory, 6th International Symposium, ANTS-VI, Burlington, VT, USA, June 13-18, 2004, Proceedings, 2004, pp. 338–357.

[Oor75]   Frans Oort, *Which abelian surfaces are products of elliptic curves?*, Mathematische Annalen **214** (1975), no. 1, 35–47.

[Pei14]   Chris Peikert, *Lattice cryptography for the internet*, Post-Quantum Cryptography - 6th International Workshop, PQCrypto 2014, Waterloo, ON, Canada, October 1-3, 2014. Proceedings, 2014, pp. 197–219.

[Pet17]   Christophe Petit, *Faster algorithms for isogeny problems using torsion point images*, Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part II, 2017, pp. 330–353.

[PH78]   Stephen C. Pohlig and Martin E. Hellman, *An improved algorithm for computing logarithms over* $GF(p)$ *and its cryptographic significance*, IEEE Trans. Information Theory **IT-24** (1978), no. 1, 106–110.

[Piz90]   Arnold K. Pizer, *Ramanujan graphs and Hecke operators*, Bull. Amer. Math. Soc. (N.S.) **23** (1990), no. 1, 127–137.

[Pol75]   J. M. Pollard, *A Monte Carlo method for factorization*, Nordisk Tidskr. Informationsbehandling (BIT) **15** (1975), no. 3, 331–334.

[Ric36]   F. Richelot, *Essai sur une méthode générale pour déterminer les valeurs des intégrales ultra-elliptiques, fondée sur des transformations remarquables de ces transcendantes*, C. R. Acad. Sci. Paris. **2** (1836), 622–627.

[Ric37]   _____, *De transformatione integralium abelianorum primi ordinis commentatio*, J. reine angew. Math. **16** (1837), 221–341.

[RS06]   Alexander Rostovtsev and Anton Stolbunov, *Public-key cryptosystem based on isogenies*, IACR Cryptology ePrint Archive **2006** (2006), 145.

[SA98]   Takakazu Satoh and Kiyomichi Araki, *Fermat quotients and the polynomial time discrete log algorithm for anomalous elliptic curves*, Comment. Math. Univ. St. Paul. **47** (1998), no. 1, 81–92.

[Ser88]   Jean-Pierre Serre, *Algebraic groups and class fields*, Graduate Texts in Mathematics, vol. 117, Springer-Verlag, New York, 1988, Translated from the French.

[Sha71]   Daniel Shanks, *Class number, a theory of factorization, and genera*, 1969 Number Theory Institute (Proc. Sympos. Pure Math., Vol. XX, State Univ. New York, Stony Brook, N.Y., 1969), Amer. Math. Soc., Providence, R.I., 1971, pp. 415–440.

[Shi79]   Tetsuji Shioda, *Supersingular $K3$ surfaces*, Algebraic geometry (Proc. Summer Meeting, Univ. Copenhagen, Copenhagen, 1978), Lecture Notes in Math., vol. 732, Springer, Berlin, 1979, pp. 564–591.

[Sil09]   Joseph H. Silverman, *The arithmetic of elliptic curves*, second ed., Graduate Texts in Mathematics, vol. 106, Springer, Dordrecht, 2009.

[Sma99]   Nigel P. Smart, *The discrete logarithm problem on elliptic curves of trace one*, J. Cryptology **12** (1999), no. 3, 193–196.

[Smi05]   Benjamin Smith, *Explicit endomorphisms and correspondences*, Ph.D. thesis, University of Sydney, 2005.

[Sto04]   Anton Stolbunov, *Public-key encryption based on cycles of isogenous elliptic curves*, Master's thesis, Saint-Petersburg State Polytechnical University, 2004, In Russian.

[STW14]   Xi Sun, Haibo Tian, and Yumin Wang, *Toward quantum-resistant strong designated verifier signature*, IJGUC **5** (2014), no. 2, 80–86.

[Sut13]   Andrew V. Sutherland, *Isogeny volcanoes*, ANTS X—Proceedings of the Tenth Algorithmic Number Theory Symposium, Open Book Ser., vol. 1, Math. Sci. Publ., Berkeley, CA, 2013, pp. 507–530.

[Tak18]   Katsuyuki Takashima, *Efficient algorithms for isogeny sequences and their cryptographic applications*, Math. Ind. (Tokyo), vol. 29, pp. 97–114, Springer, Singapore, 2018.

[Tat66]   John Tate, *Endomorphisms of abelian varieties over finite fields*, Invent. Math. **2** (1966), 134–144.

[Ti17]   Yan Bo Ti, *Fault attack on supersingular isogeny cryptosystems*, Post-Quantum Cryptography - 8th International Workshop, PQCrypto 2017, Utrecht, The Netherlands, June 26-28, 2017, Proceedings, 2017, pp. 107–122.

[TY09]   Katsuyuki Takashima and Reo Yoshida, *An algorithm for computing a sequence of Richelot isogenies*, Bull. Korean Math. Soc. **46** (2009), no. 4, 789–802.

[Vél71]   Jacques Vélu, *Isogénies entre courbes elliptiques*, C. R. Acad. Sci. Paris Sér. A-B **273** (1971), A238–A241.

[Voi18]   John Voight, *Quaternion algebras*, `https://math.dartmouth.edu/~jvoight/quat.html`, Version 0.9.14, July 7, 2018.

[Wat69]   William C. Waterhouse, *Abelian varieties over finite fields*, Ann. Sci. École Norm. Sup. (4) **2** (1969), 521–560.

[YAJ+17]   Youngho Yoo, Reza Azarderakhsh, Amir Jalali, David Jao, and Vladimir Soukharev, *A post-quantum digital signature scheme based on supersingular isogenies*, Financial Cryptography and Data Security - 21st International Conference, FC 2017, Sliema, Malta, April 3-7, 2017, Revised Selected Papers, 2017, pp. 163–181.

[ZJP+18]   Gustavo Zanon, Marcos A. Simplício Jr., Geovandro C. C. F. Pereira, Javad Doliskani, and Paulo S. L. M. Barreto, *Faster isogeny-based compressed key agreement*, Post-Quantum Cryptography - 9th International Conference, PQCrypto 2018, Fort Lauderdale, FL, USA, April 9-11, 2018, Proceedings, 2018, pp. 248–268.

Appendix A

Implementation of Genus Two SIDH

We have implemented the key exchange scheme in `MAGMA` using p of 100-bits. This yields a classical security of 75-bits and a quantum security of 50-bits. The first round of the key exchange which required the mapping of points took 145.7 seconds for Alice and 145.41 seconds for Bob. The second round of the key exchange took 74.8 seconds for Alice and 72.29 seconds for Bob.

The implementation took parameters $e_A = 51$ and $e_B = 32$, and $f = 1$ with

$$p = 4172630516011578626876079341567 \,.$$

The base hyperelliptic curve is defined by

$$
\begin{aligned}
H : y^2 = {} & (380194068372159317574541564775i + 1017916559181277226571754002873)x^6 \\
& + (3642151710276608808804111504956i + 1449092825028873295033553368501)x^5 \\
& + (490668231383624479442418028296i + 397897572063105264581753147433)x^4 \\
& + (577409514474712448616343527931i + 1029071839968410755001691761655)x^3 \\
& + (4021089525876840081239624986822i + 3862824071831242831691614151192)x^2 \\
& + (2930679994619687403787686425153i + 1855492455663897070774056208936)x \\
& + 2982740028354478560624947212657i + 2106211304320458155169465303811
\end{aligned}
$$

where $i^2 = -1$ in $\mathbb{F}_{p^2}$.

The generators of the torsion subgroups are given by

$$
P_1 = \left(
\begin{array}{c}
x^2 + (2643268744935796625293669726227i + 1373559437243573104036867095531)x \\
+2040766263472741296629084172357i + 41148336987880572074205999666055, \\
+(2643644763015937217035303914167i + 3102052689781182995044090081179)x \\
+1813936678851222746202596525186i + 3292045648641130919333133017218
\end{array}
\right),
$$

$$
P_2 = \left(
\begin{array}{c}
x^2 + (150612007990263217492664325998i + 1228415755183185090469788608852)x \\
+510940816723538210024413022814i + 325927805213930943126621646192, \\
+(1580781382037244392536803165134i + 3887834922720954573750149446163)x \\
+167573350393555136960752415082i + 1225135781040742113572860497457
\end{array}
\right),
$$

$$
P_3 = \left(
\begin{array}{c}
x^2 + (3505781767879186878832918134439i + 1904272753181081852523334980136)x \\
+646979589883461323280906338962i + 403466470460947461098796570690, \\
+(311311346636220579350524387279i + 1018806370582980709002197493273)x \\
+1408004869895332587263994799989i + 1849826149725693312283086888829
\end{array}
\right),
$$

$$
P_4 = \left(
\begin{array}{c}
x^2 + (2634314786447819510080659494014i + 72540633574927805301023935272)x \\
+1531966532163723578428827143067i + 1430299038689444680071540958109, \\
+(3957136023963064340486029724124i + 304348230408614456709697813720)x \\
+888364867276729326209394828038i + 2453132774156594607548927379151
\end{array}
\right),
$$

$$
Q_1 = \left(
\begin{array}{c}
x^2 + (263085206348114424941031847450i + 66199700402594224448399474867)x \\
+497300488675151931970215687005i + 759563233616865509503094963984, \\
+(1711990417626011964235368995795i + 3370542528225682591775373090846)x \\
+2409246960430353503520175176754i + 1486115372404013153540282992605
\end{array}
\right),
$$

$$
Q_2 = \left(
\begin{array}{c}
x^2 + (950432829617443696475772551884i + 3809766229231883691707469450961)x \\
+1293886731023444677607106763783i + 2152044083269016653158588262237, \\
+(3613765124982997852345558006302i + 4166067285631998217873560846741)x \\
+2494877549970866914093980400340i + 3422166823321314392366398023265
\end{array}
\right),
$$

$$
Q_3 = \left(
\begin{array}{c}
x^2 + (1867909473743807424879633729641i + 3561017973465655201531445986517)x \\
+614550355856817299796257158420i + 3713818865406510298963726073088, \\
+(846565504796531694760652292661i + 2430149476747360285585725491789)x \\
+3827102507618362281753526735086i + 878843682607965961832497258982
\end{array}
\right),
$$

$$
Q_4 = \left(
\begin{array}{c}
x^2 + (2493766102609911097717660796748i + 2474559150997146544698868735081)x \\
+843886014491849541025676396448i + 2700674753803982658674811115656, \\
+(2457109003116302300180304001113i + 3000754825048207655171641361142)x \\
+2560520198225087401183248832955i + 2490028703281853247425401658313
\end{array}
\right).
$$

The secret scalars of Alice and Bob are

$$\begin{array}{llll}
\alpha_1 = 937242395764589, & \alpha_2 = 282151393547351, & \alpha_3 = 0, & \alpha_4 = 0, \\
\alpha_5 = 0, & \alpha_6 = 0, & \alpha_7 = 1666968036125619, & \alpha_8 = 324369560360356, \\
\alpha_9 = 0, & \alpha_{10} = 0, & \alpha_{11} = 0, & \alpha_{12} = 0, \\
\beta_1 = 103258914945647, & \beta_2 = 1444900449480064, & \beta_3 = 0, & \beta_4 = 0, \\
\beta_5 = 0, & \beta_6 = 0, & \beta_7 = 28000236972265, & \beta_8 = 720020678656772, \\
\beta_9 = 0, & \beta_{10} = 0, & \beta_{11} = 0, & \beta_{12} = 0,
\end{array}$$

Using their secret scalars, they will obtain the following pair of hyperelliptic curves

$$\begin{aligned}
H_A : y^2 = \ & (34047030045874958215961769650058i + 40333618126043548010579938245 9)x^6 \\
& + (30015840864247629380622762223 40i + 311047190480692260365532924751 0)x^5 \\
& + (10171993106272309835115864633 32i + 15991896986314333726508575440 71)x^4 \\
& + (24695620123390929453983656786 89i + 11545664726152368274164676245 84)x^3 \\
& + (84187423865805302301385741620 0i + 42241081564390431972913195946 9)x^2 \\
& + (35075842271804269761097720529 62i + 23312982665955694626577987360 63)x \\
& + 27298166205209051755907581870 19i + 37487040066451290004985635147 34,
\end{aligned}$$

$$\begin{aligned}
H_B : y^2 = \ & (34343946890747526635795108965 30i + 32588196103419971235766003329 54)x^6 \\
& + (33502551138208951913891435659 73i + 26818924894486594289304672201 47)x^5 \\
& + (29582988186750040620470667582 64i + 90476936207932105542507672830 9)x^4 \\
& + (27012554876080269751771810910 75i + 78703312001501214614218618255 6)x^3 \\
& + (35236758116710920224917644660 22i + 28048413535583425284080556136 9)x^2 \\
& + (32381515135507987962380525651 24i + 34378857924337731633951307005 55)x \\
& + 18293273741634100972988530687 66i + 34534895169444063163962714851 72.
\end{aligned}$$

The auxiliary points computed are the following

$$\phi_B(P_1) = \pm \begin{pmatrix} x^2 + (57696747003522438444707169185 9i + 39055912331691419936017033810 59)x \\ + 14976084511258721758524483591 37i + 26229380933247876792294133204 05, \\ (22054830267312824885077668359 20i + 18876318955336669751709604986 04)x \\ + 22704381367194868281470967681 68i + 10988930791405119751197407891 84 \end{pmatrix},$$

$$\phi_B(P_2) = \pm \begin{pmatrix} x^2 + (20028072084247624580283527344 3i + 38784721108218654809248217025 29)x \\ + 47662803181075773448874071929 0i + 29575846124545180041625195748 71, \\ (39499086219077143610718155532 77i + 63063932362073596663671832104 3)x \\ + 90159764238532415792570097688 9i + 24293023201015378212402191510 82 \end{pmatrix},$$

$$\phi_B(P_3) = \pm \begin{pmatrix} x^2 + (41331577536226942506060772314 39i + 24864103595308248650394644848 54)x \\ + 21780064637456518248306490662 6i + 12493649627329044443349026898 84, \\ (12654902465945371726616464990 03i + 21308341603491590070519744331 28)x \\ + 25802866809874256010007380109 69i + 57804661019214611469846653075 8 \end{pmatrix},$$

$$\phi_B(P_4) = \pm \begin{pmatrix} x^2 + (66011020037796840738441908308 37i + 87106350729631184785549140074) x \\ + 23303393342511305368718930390 627i + 14945115526504944791133936697 13, \\ (17063142627028927741094461459 89i + 35390744497287905908915032555 45)x \\ + 19506194536813819323291061300 08i + 68517091567074185843077492006 1 \end{pmatrix},$$

$$\phi_A(Q_1) = \begin{pmatrix} x^2 + (34640403943119329646931073486 18i + 12341214841615676111016673995 25)x \\ + 17895775393232773855271038385 i + 38568589680145916450053183269 85, \\ (24328359508557655869381466383 49i + 32674847156228220519231772140 55)x \\ + 98538613755178934876027713807 6i + 11798358869918510122340542757 35 \end{pmatrix},$$

$$\phi_A(Q_2) = \begin{pmatrix} x^2 + (36338270096097826108869629350 1i + 34995487290399225281034310547 49)x \\ + 38325125233825477164180751955 17i + 33642049662042848527625303330 38, \\ (30438171015966076121868088851 16i + 40275575671985651870961331717 34)x \\ + 40871766319171660663568861985 18i + 13271576463407603468406381463 28 \end{pmatrix},$$

$$\phi_A(Q_3) = \begin{pmatrix} x^2 + (39466841366607878818882854510 15i + 12502368537491191845026040237 17)x \\ + 33581526134833765878726676747 03i + 46725220115107638905552480947 6, \\ (15629207843681052454991327577 75i + 98792082307594685023364460049 7)x \\ + 16750057582828713370107986050 79i + 14909246691958233636017633476 29 \end{pmatrix},$$

$$\phi_A(Q_4) = \begin{pmatrix} x^2 + (16294082425577501557293307597 72i + 32352833878101392017735393736 55)x \\ + 13413806694903683434507043166 76i + 14549710227885254094961980229 605, \\ (23936759862475240326635668723 48i + 34120192049740864216160966417 02)x \\ + 18903496986565042343202833183 45i + 84169906134721523463121001207 5 \end{pmatrix}.$$

This allows for both parties to compute the final isogeny to obtain

$$\begin{pmatrix} 10550181501975738539472491986 25i + 22237138430559346779893001942 59, \\ 81906058072957201350800653723 2i + 38741924008265518316862493915 28, \\ 16588859753516044944861384828 83i + 39313544136985382924653522573 93 \end{pmatrix}$$

as their common G_2-invariants.