university of groningen

faculty of science and engineering

Isogeny-based commitment schemes

Master's Project Mathematics

November 2025

Student: V. E. M. Klijn

First supervisor: Prof. dr. J. Top

Second supervisor: Prof. J. S. Müller

Supervisor TNO Groningen: S. E. Bootsma

Supervisor TNO Den Haag: T. Attema

Acknowledgements

This thesis was written as an internship at TNO. I would like to thank TNO for providing the opportunity to work on such an interesting topic, and in particular to Sven Bootsma for setting up the internship.
Thank you to Jaap for all of the useful and informative meetings. Thank you for the encouragement and for giving me the time I needed when things were not going well.
Many thanks also to Sven and Thomas for supervising me at TNO. Our weekly meetings brought me new understanding and sometimes forced me to reconsider what I knew about things I thought I understood already. I am grateful for the understanding and space I was given during difficult times.
Thanks to all of the amazing people I met at TNO. From the many hours spent playing chess to the lunch breaks playing table tennis and the many fun and thoughtful conversations, it made my time at TNO very enjoyable.
Thank you to all of my friends and family for their support through all of the difficulties of the past year and a half. In particular, thank you to Fedel for being the most supportive partner I could ever ask for.

Abstract

The field of post-quantum cryptography studies alternatives to currently existing cryptographic methods, as current methods are unsafe if a suitably powerful quantum computer is constructed. One such alternative is based on isogenies between elliptic curves, in a field known as isogeny-based cryptography. While there are many isogeny-based protocols, until 2021 there was no isogeny-based commitment scheme. Moreover, in 2024 a modification to the isogeny-based commitment scheme was proposed that removes the need for a trusted third party during the setup phase. We discuss both of these protocols and study the theory behind them, which allows us to slightly generalise the proposed modification.

Contents

1 Introduction

The field of cryptography is crucial to the everyday functioning of modern society. It forms the backbone of all forms of online communication, from sending text messages to making bank transactions. If the security of currently used cryptographic systems were compromised, it would have major worldwide consequences. Quantum computers form a threat to this security. A suitably powerful quantum computer has yet to be built, but we have known how they work in theory for decades. Shor's algorithm was proposed in 1994 [56][57] and an implementation on a powerful enough quantum computer would break widely-used cryptographic protocols, such as RSA and Elliptic Curve Diffie-Hellman.
To prepare for the event of the creation of a powerful enough quantum computer (known as Q-day), the field of Post-Quantum Cryptography (PQC) was born. In PQC, researchers study and construct cryptographic protocols that are thought to be resilient against quantum computers. They do this by basing their protocols on mathematical problems that are thought not to have efficient methods to solve them, both on current computers and on quantum computers. One approach in PQC bases itself on isogenies between elliptic curves. This is known as isogeny-based cryptography. The security comes from the fact that it is easy to construct some isogeny from a subgroup of an elliptic curve, but given just the domain and codomain of an isogeny, it is difficult to exactly recover the isogeny.

While isogeny-based cryptography has a rich history with many interesting protocols having been proposed over the years, it took until 2021 to propose the first commitment scheme based on isogenies [62]. A commitment scheme allows one to prove to a verifier that they have committed to a specific message, without actually revealing the content of the message. A commitment scheme is an example of a cryptographic primitive, which means that it forms a building block to construct more complicated cryptographic protocols. For example, commitment schemes are used in the construction of zero-knowledge proofs, where one can prove to a verifier that they have certain information without actually revealing the information. It is therefore an important addition to the world of isogeny-based cryptography to have a commitment scheme. This thesis studies the proposed isogeny-based commitment schemes and discusses the theory behind it. We discuss the potential of a homomorphic isogeny-based commitment scheme. We also study a modification, proposed in 2024 [53], that removes the need for a trusted third party in the setup phase of the commitment scheme. By looking at the mathematical background behind this proposal, we use similar strategies to prove original results that allow us to extend this modification to work over fields with different characteristics than the ones originally required.

In Section 2, we go over the necessary background in the theory of elliptic curves, though the reader is assumed to be familiar with this material. In Section 3, we introduce isogeny graphs, which have properties that will form the backbone of the commitment schemes that we study. We also discuss some more facts about elliptic curves that are particularly relevant. In Section 4, we go over the necessary background in cryptography, with the goal of defining a commitment scheme. We also briefly discuss other approaches within PQC that are not based on isogenies. In Section 5, we discuss the history of isogeny-based cryptography and we introduce a few notable protocols. This serves as a motivation for isogeny-based commitment schemes. In Section 6, we study both the 2021 isogeny-based commitment scheme [62] and the proposed 2024 modification [53] in detail. We also add some new results that provide us with more options for the characteristic of the field that we define our elliptic curves over. Finally, we discuss the potential of a homomorphic isogeny-based commitment scheme.

2 Preliminaries

In this chapter, we go over the required background on elliptic curves and isogenies. We also discuss endomorphism rings. Familiarity with introductory algebraic geometry is assumed.
In what follows, k is a perfect field.

2.1 Elliptic Curves and isogenies

In this section, we discuss some basics on elliptic curves, and later, specifically elliptic curves over a finite field. No proofs will be given in this section, refer to [59] for more details.

2.1.1 Definition and group structure

Definition 2.1. *An elliptic curve is a smooth, projective, algebraic curve of genus 1, with a specified base point $\mathcal{O}$, defined over k.*

Any elliptic curve defined over a field k can be written in Weierstrass form:

$$Y^2Z + a_1XYZ + a_3YZ^2 = X^3 + a_2X^2Z + a_4XZ^2 + a_6Z^3,$$

or affinely,

$$y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$$

for coefficients $a_1, a_2, a_3, a_4, a_6 \in k$. In Weierstrass form, the base point $\mathcal{O}$ is the point at infinity, so we have $\mathcal{O} = (0 : 1 : 0)$. If $\mathrm{char}(k) \neq 2, 3$, we can reduce the above form to the short Weierstrass form:

$$y^2 = x^3 + ax + b.$$

There are many values associated to these coefficients, most notably the discriminant Δ and the j-invariant. In the short Weierstrass form, we define $\Delta := -16(4a^3 + 27b^2)$. A Weierstrass equation is an elliptic curve if and only if $\Delta \neq 0$. If $\Delta \neq 0$, we define $j := -1728\frac{(4a)^3}{\Delta} = 1728\frac{4a^3}{4a^3+27b^2}$ to be the j-invariant of an elliptic curve.
There are many other forms in which the equation of an elliptic curve can be written, some of which are more convenient for certain purposes. One such form is called the Montgomery form. This form is given as follows:

$$By^2 = x^3 + Ax^2 + x.$$

In the Montgomery form, we have $\Delta = A^2 - 4$ and $j = \frac{256(A^2-3)^3}{A^2-4}$. Notably, these values are independent of B. The article [14] provides an introduction to this form.
The Montgomery form has the benefit of providing efficient computation of the addition of points on the curve. This addition is the one given by the group law that exists on all elliptic curves.

Definition 2.2. *Let P_1, P_2 be two points on an elliptic curve E. One defines the following composition law on E: let L be the line through P_1 and P_2 (if $P_1 = P_2$, take the tangent line to E at P_1). Then, since E is cubic, by Bézout's theorem there exists a third point R on E that also intersects L. Then, take the line L' through R and $\mathcal{O}$. The third point of intersection of this line with E is denoted by $P_1 \oplus P_2$.*

Proposition 2.3. *The composition law defined in Definition 2.2 satisfies the properties of an abelian group law, where $\mathcal{O}$ is the zero element.*

Proof. See Proposition III.2.2 of [59]. □

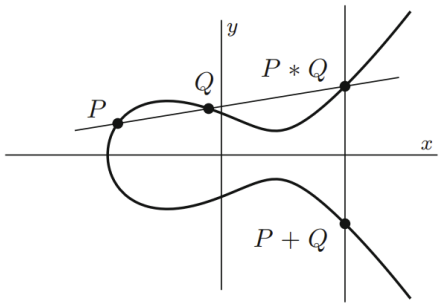

Figure 1: The group law of an elliptic curve, viewed over $\mathbb{R}$.

For convenience, we will denote the group law by $+$ (and $-$) instead of using $\oplus$. We also use the notation $[m]P$ to mean a point P added to itself m times. We call $[m]$ the *multiplication-by-m map*. This is an example of an isogeny, which we define in Definition 2.8. For an elliptic curve defined over a field k, we denote by $E(k)$ the subgroup of k-rational points.

Proposition 2.4. *Let $P_1 = (x_1, y_1)$, $P_2 = (x_2, y_2)$ be two points on an elliptic curve $E : y^2 = x^3 + ax + b$. We have the following formulas:*

- *$P + \mathcal{O} = \mathcal{O} + P = P$ for any $P \in E$.*
- *If $x_1 = x_2$ and $y_1 = -y_2$, then $P_1 + P_2 = \mathcal{O}$.*
- *Otherwise, set*

$$\lambda = \begin{cases} \frac{y_2 - y_1}{x_2 - x_1} & \text{if } P_1 \neq P_2, \\ \frac{3x_1^2 + a}{2y_1} & \text{if } P_1 = P_2, \end{cases}$$

then the point $P_1 + P_2 = (x_3, y_3)$ is given by

$$\begin{aligned} x_3 &= \lambda^2 - x_1 - x_2 \\ y_3 &= -\lambda x_3 - y_1 + \lambda x_1 \end{aligned}$$

Proof. See Group Law Algorithm III.2.3 of [59]. □

Such addition formulas can be given more generally for curves in long Weierstrass form. Additionally, with elliptic curves written in other forms, such as the Montgomery form, there are different formulas.

Next, we move on to the torsion structure of elliptic curves.

Definition 2.5. *Let E/k be an elliptic curve, and $m \in \mathbb{Z} \setminus \{0\}$. The m-torsion subgroup of E, denoted $E[m]$, is defined as the kernel of the multiplication-by-m map $[m]$.*

In other words, we can think of $E[m]$ as the set of points of order dividing m on E. The structure of this subgroup is easy to classify:

Proposition 2.6. *Let E/k be an elliptic curve, and $m \in \mathbb{Z} \setminus \{0\}$. The subgroup $E[m]$ has the following structure:*

- *$E[m] \cong \mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}$ if $\mathrm{char}(k) \nmid m$.*

- *If $char(k) = p > 0$, we have*

$$E[p^i] \cong \begin{cases} \mathbb{Z}/p^i\mathbb{Z} \text{ for all } i \geq 0 \text{ or} \\ \{\mathcal{O}\} \text{ for all } i \geq 0 \end{cases}$$

Over fields with positive characteristic p, curves containing a point of order p are called **ordinary**, while curves without any such point are called **supersingular**. Throughout this chapter, we will encounter equivalent definitions of these notions.
Note that not all m-torsion points may be defined over the field k. They may only exist over field extensions.

2.1.2 Isogenies

We would like to be able to define maps from one elliptic curve to the other. The easiest way to do this is with an isomorphism. An **isomorphism** of elliptic curves is an algebraic morphism of elliptic curves that has an inverse. For elliptic curves in short Weierstrass form, an isomorphism can be written as a change of variables

$$x = u^2x', y = u^3y'$$

for some $u \in \bar{k}^*$, where $\bar{k}$ denotes an algebraic closure of k. Isomorphism classes are given by the j-invariant.

Proposition 2.7. *Two elliptic curves are isomorphic over $\bar{k}$ if and only if they have the same j-invariant.*

Proof. See Proposition III.1.4b of [59]. □

Over a field that is not algebraically closed, it only holds that two isomorphic curves will have the same j-invariant. The other direction does not necessarily hold over non-algebraically closed fields.
Instead of isomorphisms, one can look more generally at isogenies between elliptic curves, which will be an important topic of study in this thesis.

Definition 2.8. *An isogeny between two elliptic curves E, E' is a morphism of algebraic curves $\phi : E \to E'$ such that $\phi(\mathcal{O}_E) = \mathcal{O}_{E'}$.*

As it turns out, this definition is enough to show that an isogeny is not just a morphism of algebraic curves, it is also a group homomorphism from $E(k)$ to $E'(k)$ for any field k. Moreover, if ϕ is not trivial, it is surjective if k is algebraically closed, since it is a morphism of curves.
Two curves are *isogenous* is there exists a nontrivial isogeny between them. Two curves are *isogenous over k* if there exists a nontrivial isogeny defined over k between them. Being isogenous is an equivalence relation.
An important property of isogenies is their degree, which we define as follows:

Definition 2.9. *The degree of an isogeny ϕ is its degree as an algebraic map, that is, $[k(E) : \phi^*(k(E'))]$. An isogeny is separable, inseparable or purely inseparable if the extension of function fields is separable, inseparable or purely inseparable.*

Additionally, the kernel of an isogeny is a finite group. This is because the kernel of a morphism of smooth curves is at most of size the degree, and because the kernel is a subgroup of the group of points since an isogeny is a homomorphism.

Proposition 2.10. *Let ϕ be an isogeny. Then,*

1. *If ϕ is separable, $\deg(\phi) = \#\ker(\phi)$.*
2. *If ϕ is purely inseparable, then $\deg(\phi)$ is a power of the characteristic of k.*
3. *Any isogeny can be written as the product of a separable and a purely inseparable isogeny.*

In practice, we will mostly focus on separable isogenies, since our main method to generate isogenies only generates separable ones. Therefore, we can easily determine the degree of an isogeny by looking at its kernel. When the kernel of a separable isogeny is cyclic, we call the isogeny a *cyclic isogeny*.
For an elliptic curve E given by the equation $y^2 = f(x)$ for some cubic polynomial $f(x)$, and an isogeny $\phi : E \to E'$, we can explicitly write out ϕ as sending a point (x, y) to $\left(\frac{u_1(x)}{v_1(x)}, y\frac{u_2(x)}{v_2(x)}\right)$, where $u_1(x)$ and $v_1(x)$ are coprime. With this, we can say that an isogeny ϕ is defined over a field k if the coefficients of u_1, v_1, u_2, v_2 lie in k. Additionally, this allows for a more explicit formulation of the degree of an isogeny.

Lemma 2.11. *Let ϕ be an isogeny defined by $\phi(x, y) = \left(\frac{u_1(x)}{v_1(x)}, y\frac{u_2(x)}{v_2(x)}\right)$ with $u_1(x)$ and $v_1(x)$ coprime. The degree of ϕ is given by $\deg(\phi) = \max\{\deg(u_1), \deg(v_1)\}$. If the derivative $\left(\frac{u_1}{v_1}\right)'(x)$ is not identically 0, then ϕ is separable.*

Proof. See Lemma 6.2 in [61]. □

The simplest example of an isogeny is the multiplication-by-m map, denoted by $[m]$, which sends a point P to $[m]P$. The kernel of this map is the m-torsion subgroup $E[m]$. The map $[m]$ goes from an elliptic curve to itself, that is, it is an *endomorphism*.
Two isogenies with the same domain and codomain can be added by $(\phi + \psi)(P) = \phi(P) + \psi(P)$. This means the set $\mathrm{Hom}(E, E')$, defined as the set of isognies from E to E', is a group. The zero element is given by the trivial isogeny that sends everything to the point at infinity. Moreover, if the domain and codomain are the same, we have endomorphisms. We can compose endomorphisms, and thus we obtain a ring structure. The endomorphism ring, $\mathrm{End}(E)$, will be studied in detail in Section 2.2. Knowing that $\mathrm{End}(E)$ is a ring, one can see $\mathrm{Hom}(E, E')$ as a left $\mathrm{End}(E)$-module, or a right $\mathrm{End}(E')$-module.
The following proposition tells us that the isomorphism class of the codomain of an isogeny is determined by its kernel. We will make heavy use of this fact in this thesis, as it allows us to create isogeny graphs. More on isogeny graphs in Section 3.

Proposition 2.12. *Let E be an elliptic curve and let G be a finite subgroup of E. Then, there exists an elliptic curve E' and a separable isogeny ϕ such that $\ker(\phi) = G$ and $\phi : E \to E'$. Furthermore, E' and ϕ are unique up to isomorphism.*

Proof. See Proposition III.4.12 of [59]. □

The curve E' in the proposition is usually denoted by E/G. This notation is motivated by the uniqueness of the codomain up to isomorphism, making it resemble the traditional theory of taking a quotient of a group by a (normal) subgroup. For this reason, we may also talk about taking the quotient of E by G.
There exist formulas that, given a kernel, computes a representative codomain and the isogeny. These were originally found by Vélu [65].

Proposition 2.13. *Let $E : y^2 = x^3 + ax + b$ be an elliptic curve defined over some field k and let $G \subset E(\bar{k})$ be a finite subgroup. Let ϕ be the separable isogeny such that $\ker(\phi) = G$. For any point $P \notin G$, we have*

$$\phi(P) = \left(x(P) + \sum_{Q \in G \setminus \{\mathcal{O}\}} x(P+Q) - x(Q), \quad y(P) + \sum_{Q \in G \setminus \{\mathcal{O}\}} y(P+Q) - y(Q) \right),$$

where the resulting points lie on the curve with equation $y^2 = x^3 + a'x + b'$ with

$$a' = a - 5 \sum_{Q \in G \setminus \{\mathcal{O}\}} \left(3x(Q)^2 + a\right),$$
$$b' = b - 7 \sum_{Q \in G \setminus \{\mathcal{O}\}} \left(5x(Q)^3 + 3ax(Q) + 2b\right).$$

This curve is a representative for E/G.

Another important concept is the dual isogeny. The dual isogeny is generally a useful tool to have when working with isogenies, and some of its basic properties can be very useful.

Theorem 2.14. *Let $\phi : E \longrightarrow E'$ be an isogeny of degree m. There exists an isogeny $\widehat{\phi} : E' \longrightarrow E$ called the dual isogeny such that*

$$\widehat{\phi} \circ \phi = [m]_E, \phi \circ \widehat{\phi} = [m]_{E'},$$

which satisfies the following properties:

1. *$\widehat{\phi}$ is unique.*
2. *$\widehat{\phi}$ is defined over k if and only if ϕ is.*
3. *Let $\lambda : E' \longrightarrow E''$ be another isogeny. Then, $\widehat{\lambda \circ \phi} = \widehat{\phi} \circ \widehat{\lambda}$.*
4. *Let $\psi : E \longrightarrow E'$ be another isogeny. Then, $\widehat{\psi + \phi} = \widehat{\psi} + \widehat{\phi}$.*
5. $\deg \phi = \deg \widehat{\phi}$.
6. $\widehat{\widehat{\phi}} = \phi$.

Proof. See Theorem III.6.2 of [59]. □

2.2 Endomorphisms and endomorphism rings

2.2.1 Quaternion algebras

The structure of the endomorphism ring $\mathrm{End}(E)$ will be a crucial tool for studying isogeny-based cryptography. However, before we can classify endomorphism rings of elliptic curves, we first need to recall some theory about quaternion algebras. While quaternion algebras can be defined over any field, we will restrict our focus to the case of $\mathbb{Q}$ as this is the relevant case for us. Recall that a $\mathbb{Q}$-algebra is simply a $\mathbb{Q}$-vector space equipped with a compatible ring structure.

Definition 2.15. *A quaternion algebra over $\mathbb{Q}$ is a $\mathbb{Q}$-algebra of the form*

$$\mathbb{Q} + i\mathbb{Q} + j\mathbb{Q} + k\mathbb{Q},$$

where i, j, k satisfy the relations $i^2 = a, j^2 = b, k = ij = -ji$ for $a, b \in \mathbb{Q}$. Such an algebra is denoted by $\left(\frac{a,b}{\mathbb{Q}}\right)$.

An element α of a quaternion algebra $\left(\frac{a,b}{\mathbb{Q}}\right)$ can be written as $t + xi + yj + zk$, with $t, x, y, z \in \mathbb{Q}$. We call t the real part, and $xi + yj + zk$ the imaginary part. The conjugate $\bar{\alpha}$ works similarly to the complex numbers, in that $\bar{\alpha} = t - xi - yj - zk$. This also gives us the reduced norm and reduced trace of α, with $Nrd(\alpha) := \alpha\bar{\alpha} = t^2 - ax^2 - by^2 + abz^2$ and $Trd(\alpha) := \alpha + \bar{\alpha} = 2t$ respectively. The regular norm and trace also exist, but these are not used in the context of quaternion algebras. A quaternion algebra K is *split* at a prime p if $K \otimes_{\mathbb{Q}} \mathbb{Q}_p \cong M_2(\mathbb{Q}_p)$. Put differently, K is split if it is a matrix algebra when viewed as a $\mathbb{Q}_p$-algebra. Equivalently, viewed as a $\mathbb{Q}_p$-algebra, there exists a nontrivial element with norm zero. If p is not split, we call it *ramified*.
A quaternion algebra K is called *ramified* at ∞ if there does not exist a nontrivial element of norm zero in $K \otimes_{\mathbb{Q}} \mathbb{R}$. This is equivalent to having a and b both negative. The *reduced discriminant* of a quaternion algebra is the product of the primes over which it ramifies.

We denote by $B_{p,\infty}$ the unique (up to isomorphism) quaternion algebra that ramifies at p and ∞. This quaternion algebra has reduced discriminant p. Depending on the value of p, we pick the following representatives:

Proposition 2.16. *Let p be a prime. We choose the following representatives of $B_{p,\infty}$:*

1. $B_{p,\infty} \cong \left(\frac{-1,-1}{\mathbb{Q}}\right)$ *if* $p = 2$;
2. $B_{p,\infty} \cong \left(\frac{-1,-p}{\mathbb{Q}}\right)$ *if* $p \equiv 3 \bmod 4$;
3. $B_{p,\infty} \cong \left(\frac{-2,-p}{\mathbb{Q}}\right)$ *if* $p \equiv 5 \bmod 8$;
4. $B_{p,\infty} \cong \left(\frac{-r,-p}{\mathbb{Q}}\right)$ *if* $p \equiv 1 \bmod 8$, *where* $r \equiv 3 \bmod 4$ *is a prime that is not a square modulo* p.

These representatives will be useful when dealing with elliptic curves over a finite field with characteristic satisfying one of these modular conditions.
We will now give the definition of an order in a quaternion algebra, and introduce some relevant properties of orders.

Definition 2.17. *An order in a quaternion algebra is a rank 4 $\mathbb{Z}$-module that is also a subring of said algebra. An order is maximal if it is not contained in any other order.*

Proof. See [20]. □

If we have a $\mathbb{Z}$-basis $\{\alpha_1, \alpha_2, \alpha_3, \alpha_4\}$ of an order $\mathcal{O}$, we define the *reduced discriminant* to be $\mathrm{discrd}(\mathcal{O}) := |\det(\mathrm{Tr}(\alpha_i \bar{\alpha_j}))_{1 \le i,j \le 4}|^{1/2}$.

Lemma 2.18. *An order in a quaternion algebra is maximal if and only if it has reduced discriminant equal to the reduced discriminant of the quaternion algebra.*

Therefore, within $B_{p,\infty}$, we can verify that an order is maximal by checking whether its reduced discriminant equals p. For a non-maximal order, we can consider the reduced discriminant divided by p as the index of the maximal order over the non-maximal one.

2.2.2 Endomorphism rings

We have now discussed the necessary background to understand the upcoming result about the classification of endomorphism rings. As previously discussed, the endomorphism $[m]$ always exists, where a point is sent to m times itself. This means that there will always be an injection $\mathbb{Z} \hookrightarrow \mathrm{End}(E)$. However, there might be more endomorphisms than this. Deuring's result, known as the *Deuring correspondence*, classifies endomorphism rings of elliptic curves over any field [20].

Theorem 2.19 (Deuring correspondence)**.** *Let E be an elliptic curve defined over a field k with $\mathrm{char}(k) = p$. The endomorphism ring* $\mathrm{End}(E)$ *is isomorphic to one of the following:*

- $\mathbb{Z}$.
- *An order $\mathcal{O}$ in an imaginary quadratic field $\mathbb{Q}(\sqrt{-D})$ for some squarefree $D > 0$. In this case, we say that E has complex multiplication (CM) by $\mathcal{O}$.*
- *A maximal order in the quaternion algebra $B_{p,\infty}$.*

If $p = 0$, only the first two cases are possible.

The Deuring correspondence gives us an equivalent way to define supersingular and ordinary curves. If $\mathrm{char}(k) > 0$, a curve E is supersingular if $\mathrm{End}(E)$ is isomorphic to a maximal order in a quaternion algebra. It is ordinary if it is not supersingular.
Associated to the endomorphism ring $\mathrm{End}(E)$ is the endomorphism algebra, which is defined as $\mathrm{End}(E) \otimes \mathbb{Q}$. Looking at the Deuring correspondence, in the first case the endomorphism algebra has rank 1, in the second case it has rank 2, and in the third case it has rank 4. We can thus also say that an elliptic curve is supersingular if its endomorphism algebra has rank 4. Furthermore, we have that the endomorphism algebra is invariant under taking isogenies, as characterised by the following theorems.

Theorem 2.20. *Two elliptic curves E, E' with complex multiplication are isogenous (over the algebraic closure) if and only if their endomorphism algebras $\mathrm{End}(E) \otimes \mathbb{Q}$ and $\mathrm{End}(E') \otimes \mathbb{Q}$ are isomorphic.*

Theorem 2.21. *Any two supersingular elliptic curves defined over a field of characteristic p are isogenous (over the algebraic closure).*

This indeed tells us that if we take an isogeny of any elliptic curve, including those with endomorphism ring $\mathbb{Z}$, the endomorphism algebra is unchanged.
The Deuring correspondence also gives us an easy way to compute the degree of an endomorphism. The degree of an endomorphism is equal to the norm of the corresponding element in the endomorphism ring.

2.2.3 The Frobenius endomorphism

In the case of elliptic curves over finite fields, the Deuring correspondence tells us that there will always be endomorphisms that are not the multiplication-by-m map, which we will refer to as nontrivial endomorphisms. We will now introduce one such endomorphism.
In what follows, let E be an elliptic curve over a finite field $\mathbb{F}_q$ with $q = p^n$ elements, where $n \geq 1$. For any such elliptic curve, there exists what is known as the *Frobenius endomorphism.*

Definition 2.22. *The Frobenius endomorphism, denoted by F, is the map that sends a point (x, y) on E to the point (x^q, y^q).*

There are some quick observations that one can make about the Frobenius endomorphism.

Proposition 2.23. *Let F be the Frobenius endomorphism of $E/\mathbb{F}_q$. Then:*

- $\ker(F) = \{\mathcal{O}\}$*;*
- $\ker(F - 1) = E(\mathbb{F}_q)$*, where 1 is the identity map on E.*

The second point in particular says that the fixpoints of F are precisely the $\mathbb{F}_q$-rational points. Thus, points with coordinates in field extensions of $\mathbb{F}_q$ are never fixed. By Lemma 2.11, the $\mathbb{F}_q$-Frobenius map has degree q.
The Hasse bound gives an easy estimate for the amount of $\mathbb{F}_q$-rational points that an elliptic curve can have.

Theorem 2.24 (Hasse)**.** *Let $E/\mathbb{F}_q$ be an elliptic curve. Then,*

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

Theorem 2.25. *Let $E/\mathbb{F}_q$ be an elliptic curve, and let F be its Frobenius endomorphism. Then, F satisfies the equation*

$$F^2 - tF + q = 0$$

for some $t \in \mathbb{Z}$ such that $|t| \leq 2\sqrt{q}$.

We call t the *trace* of F. More specifically we have $t = 1 + q - \#E(\mathbb{F}_q)$. The trace gives us another equivalent definition of supersingularity: E is supersingular if t is congruent to zero mod p. Over a field $\mathbb{F}_p$ with p a prime not equal to 2 or 3, this is equivalent to $t = 0$ due to the Hasse bound.
From Theorem 2.25, we can see that the Frobenius endomorphism will never be in $\mathbb{Z}$, with one exception for the case of a supersingular curve defined over $\mathbb{F}_{p^n}$, with n even. This is shown as follows: note that the equation $F^2 - tF + q$ has solutions $F = \frac{t \pm \sqrt{t^2 - 4q}}{2}$. Because $|t| \leq 2\sqrt{q}$, we have $t^2 - 4q \leq 0$. If we suppose $F \in \mathbb{Z}$, we must have $t^2 - 4q = 0$, and so we obtain $2\sqrt{q} = \pm t$. Since $t \in \mathbb{Z}$, this means that q must be a square (an even power of p) and t is 0 mod p, which implies that our curve is supersingular.
The endomorphism ring of an elliptic curve defined over a finite field contains $\mathbb{Z}[F]$, which, outside of the case mentioned above, is strictly bigger than $\mathbb{Z}$. The Deuring correspondence already told us that the endomorphism ring was bigger than $\mathbb{Z}$, but it is particularly nice to have $\mathbb{Z}[F]$, as the map F is easy to write down. This is useful since it is not always easy to find the full endomorphism ring.
Another equivalent definition for supersingularity also uses the Frobenius endomorphism: a curve is supersingular if the dual of the Frobenius endomorphism is purely inseparable.

The invariance of endomorphism algebras under isogenies reduces to a very simple condition for the case of finite fields.

Theorem 2.26. *Two elliptic curves E, E' defined over $\mathbb{F}_q$ are isogenous over $\mathbb{F}_q$ if and only if $\#E(\mathbb{F}_q) = \#E'(\mathbb{F}_q)$.*

To finish this chapter, we will collect all the equivalent definitions of supersingularity mentioned throughout this chapter.

Proposition 2.27. *Let E be a supersingular elliptic curve defined over a field k with $char(k) = p > 0$. Then the following are equivalent:*

1. *$E[p^r] = 0$ for all $r \geq 0$;*
2. *The endomorphism algebra of E has rank 4;*
3. *The dual of the Frobenius map F is purely inseparable;*
4. *The trace of Frobenius t is congruent to 0 modulo p.*

3 Isogeny graphs

In this chapter, we go through the necessary background on graph theory. After this, we are able to talk about isogeny graphs and their properties. We will mostly focus on isogeny graphs of supersingular elliptic curves.

3.1 Graphs

Before we can talk about isogeny graphs and their properties, we need to discuss some concepts from graph theory. This will be a very brief introduction and no proofs will be given. For a more detailed background on expander graphs and other applications of them, see [31].
Let G be a graph. A graph consists of vertices and edges that connect the vertices. We denote the vertices as a set $V(G)$ and the edges as a multiset $E(G)$, where $E(G)$ contains ordered pairs of elements of $V(G)$.

Definition 3.1. *The degree of a vertex is the number of outgoing edges of this vertex. For a positive integer d we call a graph d-regular if every vertex of G has degree equal to d.*

A graph is *undirected* if for any two vertices v, v', there are as many edges from v to v' as from v' to v.

Definition 3.2. *The adjacency matrix A of a graph G with n vertices is the $n \times n$ matrix whose (i, j)-th entry is the number of edges from the vertex i to the vertex j.*

If G is undirected, the adjacency matrix A is symmetric and thus the eigenvalues are real. For an undirect d-regular graph G it holds that the eigenvalues $\lambda_1 \geq \ldots \geq \lambda_n$ of A are such that $|\lambda_i| \leq d$ for all i.
One thing that can be studied about a graph is its expansion. The expansion of a graph captures how "well-connected" a graph is. It is formally defined as follows:

Definition 3.3 (Expander graph)**.** *Let $\varepsilon > 0$ and $d \geq 1$. A d-regular undirected graph is called a one-sided ε-expander if*

$$\lambda_2 \leq (1 - \varepsilon)d,$$

and a two-sided ε-expander if it also satisfies

$$\lambda_n \geq (-1 - \varepsilon)d.$$

The boundary of a subset of vertices of a graph are all the vertices that are connected via an edge to a vertex in the subset. Intuitively, we can think of good expander graphs as graphs that have low degree (that is, all vertices have low degree), but every subset of vertices has a large boundary. For our purposes, a relevant subset of expander graphs are Ramanujan graphs.

Definition 3.4 (Ramanujan graph)**.** *Let G be a d-regular undirected graph. We call G a Ramanujan graph if for all eigenvalues λ_i of the adjacency matrix of G it holds that either $|\lambda_i| = d$ or $|\lambda_i| \leq 2\sqrt{d-1}$.*

A *random walk* on a graph is a process where we start on some vertex of a graph and select one of the edges at random. After this, we go to the other vertex that this edge is connected to. We repeat this k times, where k determines the *length* of the walk. A random walk is *non-backtracking* if in each step, we do not traverse the edge that was used in the previous step. Non-backtracking walks will be very important in the commitment schemes that we study in Section 6.
The (i, j)-th entry of the adjacency matrix A can be thought of as the amount of walks of length 1 that exist between vertices i and j. Similarly, powers of the adjacency matrix A^k have the amount of walks of length k on each entry. To include only non-backtracking walks, we use the matrices A_i (see Section 8 of [40]). For a d-regular graph, these are defined as follows: $A_1 = A$, $A_2 = A^2 - dI$, and $A_{r+1} = A_1 A_r - (d-1)A_{r-1}$ for $r \geq 2$. Now, A_r is the matrix whose (i, j)-th entry equals the number of non-backtracking walks of length r between vertices i and j. We can say something useful about these matrices:

Lemma 3.5. *Let G be a connected d-regular graph with $d \geq 3$. Then there exists some positive integer k_0 such that for all $k \geq k_0$, A_k has entries which are all non-zero.*

Proof. See [62], Lemma 3.2. □

In other words, if we perform a sufficiently long non-backtracking random walk on a regular graph, we will be able to reach any vertex starting from any other vertex. This lemma allows us to define the mixing constant:

Definition 3.6 (Mixing constant)**.** *The mixing constant k_G of a graph G is defined as the smallest value k_0 such that A_k does not have any entry equal to zero.*

The mixing constant is useful for giving a lower bound on the required length of a random walk whose distribution of possible endpoints is hard to distinguish from a uniform distribution. This has applications in cryptography. For a d-regular graph, we can fairly easily provide a lower bound for the mixing constant:

Lemma 3.7. *The mixing constant k_G of a connected d-regular graph G with N vertices is bounded below by*

$$k_G \geq \log_{d-1}(N) - \log_{d-1}(d) + 1.$$

Proof. From a given starting vertex, there are $d(d-1)^{k-1}$ non-backtracking walks of length k (for the first step, there are d options, while for the other steps there are $d-1$ options as we cannot backtrack). In the best case scenario, these walks all end at different vertices. The mixing constant must be a walk with a length that can end on at least N different vertices, in other words:

$$d(d-1)^{k_G-1} \geq N.$$

Some simple rearrangement then gives the desired inequality. □

There are no proven upper bounds for the mixing constant. However, in his work on an isogeny-based commitment scheme (which will be discussed in detail in Section 6), Sterner [62] also gives a conjectural upper bound for the mixing constant based on a known result on non-backtracking random walks on d-regular graphs.

Conjecture 3.8. *The mixing constant k_G of a connected d-regular graph G with N vertices has the following upper bound:*

$$k_G \leq 4\lceil \log_{d-1}(dN) \rceil + 4.$$

3.2 Isogeny graphs

In this section, we introduce isogeny graphs. Recall from Proposition 2.12 that isogenies are uniquely defined by their kernel. Since the kernel is preserved if an isogeny is post-composed with an isomorphism of curves, it makes sense to consider such a map equivalent to the original isogeny. This forms an equivalence relation on isogenies. With this, we can formulate what it means to be an isogeny graph.

Definition 3.9. *An isogeny graph is a graph that has j-invariants of isogenous curves as its nodes, and equivalence classes of isogenies between these invariants as its edges.*

Notice that since every isogeny has a dual isogeny, we generally draw isogeny graphs as undirected graphs. There is a notable exception to this, namely the case where the domain of the isogeny has a nontrivial automorphism group (that is, an automorphism group larger than $\{\pm 1\}$). This occurs only at j-invariants 0 and 1728.
In this situation, given some isogeny with this domain, we can generate a non-equivalent isogeny with the same codomain by pre-composing with the nontrivial automorphism, as this may result in a different kernel.

However, both of these non-equivalent isogenies will have an equivalent dual isogeny. Let $\phi : E \to E'$ be an isogeny, and let σ be a nontrivial automorphism on E. Then, $\phi \circ \sigma$ is a non-equivalent isogeny. However, the dual of ϕ is $\widehat{\phi}$, while the dual of $\phi \circ \sigma$ is $\widehat{\sigma} \circ \widehat{\phi}$, which is equivalent to $\widehat{\phi}$ because it is post-composed by an isomorphism. Thus, in this case, we need to draw our edges in a directed manner. A case where this occurs is described in the example at the end of this section, see Figure 2.
Notice that such behaviour can only occur with a nontrivial automorphism, since the automorphism $[-1]$ preserves the kernel of an isogeny because the kernel is a finite subgroup.
In isogeny-based cryptography, we generally restrict our attention to ℓ-isogeny graphs, where ℓ is a prime number. An ℓ-isogeny graph only has ℓ-isogenies (isogenies of degree ℓ) as its edges. Since the degree of a (separable) isogeny is the same as the cardinality of its kernel, there will always be $\ell + 1$ isogenies of degree ℓ, provided $\ell \neq p$. This follows from the fact that the group $(\mathbb{Z}/l\mathbb{Z})^2$ has precisely $\ell+1$ subgroups of order ℓ.

While ordinary elliptic curves will not be used in this thesis, we will briefly introduce their structure here for completeness. Isogeny graphs of ordinary elliptic curves are quite rigid. They are in the form of what is called a *volcano*. They were first studied by Kohel [30]. The following proposition tells us how isogenies of ordinary elliptic curves work:

Proposition 3.10. *Let $\phi : E \longrightarrow E'$ be an isogeny of prime degree ℓ, and let $\mathcal{O}, \mathcal{O}'$ be the respective endomorphism rings of E, E'. Then either $\mathcal{O} \subset \mathcal{O}'$ or $\mathcal{O}' \subset \mathcal{O}$, and one of the following holds:*

- *$\mathcal{O} = \mathcal{O}'$, in which case we call ϕ horizontal.*
- *$[\mathcal{O}' : \mathcal{O}] = \ell$, in which case we call ϕ ascending.*
- *$[\mathcal{O} : \mathcal{O}'] = \ell$, in which case we call ϕ descending.*

Since endomorphism algebras are invariant under isogenies 2.20, we know that $\mathcal{O}$ and $\mathcal{O}'$ are orders in the same endomorphism algebra. The fact that they differ by an ℓ-isogeny ensures that one is always contained in the other.

3.3 Supersingular isogeny graphs

In this section, we restrict our attention to supersingular isogeny graphs and their properties. The case of supersingular isogeny graphs is different from the case of ordinary isogeny graphs. All supersingular $\overline{\mathbb{F}}_p$-isomorphism classes, which we call the full supersingular isogeny graph, do not have a structure as rigid as the isogeny volcanoes from the ordinary case in Proposition 3.10. However, there is still a lot to say about supersingular isogeny graphs. We will denote the supersingular ℓ-isogeny graph over a field of characteristic p by $G_\ell(p)$.

Ordinary elliptic curves can have their j-invariant exist over any extension field. On the other hand, supersingular j-invariants do not have as many possibilities.

Proposition 3.11. *Let E be a supersingular elliptic curve defined over a field k of characteristic $p > 0$. Then $j(E) \in \mathbb{F}_{p^2}$.*

Proof. (Adapted from [59], Theorem V.3.1a.) For simplicity of the proof we assume $p \neq 2, 3$, but the statement also holds for those p. We know E is defined by a Weierstrass equation $y^2 = x^3 + Ax + B$ for $A, B \in k$. For a prime power q of p, let $E^{(q)}$ denote the elliptic curve defined by $y^2 = x^3 + A^q x + B^q$. Let $F^{(p)}$ denote the p-power Frobenius isogeny from E to $E^{(p)}$ that sends a point (x, y) to (x^p, y^p) (the general q-power Frobenius map is a more general version of the Frobenius endomorphism). Because E is supersingular, by Proposition 2.27 we know that the dual of $F^{(p)}$ is purely inseparable, with inseparable degree p. So, we can write $\widehat{F^{(p)}} = \widehat{F^{(p)}}_{\text{sep}} \circ \widehat{F^{(p)}}_{\text{insep}}$ by Proposition 2.10. Furthermore, Corollary II.2.12 of [59] tells us that

$\widehat{F^{(p)}}_{\text{insep}} = F^{(p)}$. Since $\widehat{F^{(p)}}$ and $F^{(p)}$ both have degree p, we have that $\widehat{F^{(p)}}_{\text{sep}}$ has degree 1 and is therefore an isomorphism. As $\widehat{F^{(p)}} \circ F^{(p)} = [p]$, we obtain

$$[p] = \widehat{F^{(p)}} \circ F^{(p)} = \widehat{F^{(p)}}_{\text{sep}} \circ (F^{(p)})^2$$

and $\widehat{F^{(p)}}_{\text{sep}}$ is an isomorphism from $E^{(p^2)}$ to E. This means that

$$j(E) = j(E^{p^2}) = j(A^{p^2}, B^{p^2}) = j(A, B)^{p^2} = j(E)^{p^2}.$$

So, we conclude that $j(E)$ is fixed by the map $x \mapsto x^{p^2}$ of k, which means that $j(E) \in \mathbb{F}_{p^2}$. □

It holds that in any field, given some element a, we can construct an elliptic curve with j-invariant equal to a (see Proposition III.1.4c in [59]). So in particular, any supersingular curve defined over a field of characteristic p is isomorphic (not necessarily over $\mathbb{F}_{p^2}$) to a curve defined over $\mathbb{F}_{p^2}$. Furthermore, for every supersingular j-invariant, there exists a curve defined over $\mathbb{F}_{p^2}$ with Frobenius trace equal to $-2p$. So, while we have previously defined supersingular isogeny graphs with $\overline{\mathbb{F}}_p$-isomorphism classes, we can show that supersingular ℓ-isogeny graphs are actually fully defined over $\mathbb{F}_{p^2}$. This is given by the following proposition:

Proposition 3.12. *The supersingular ℓ-isogeny graph is isomorphic to the graph of $\mathbb{F}_{p^2}$-isomorphism classes of elliptic curves with trace $-2p$, with isogenies defined over $\mathbb{F}_{p^2}$.*

Proof. Let E be a supersingular curve defined over $\mathbb{F}_{p^2}$. Theorem 4.6 from [54] tells us that depending on the congruence class of p, there are either zero or two $\mathbb{F}_{p^2}$-isomorphism classes for traces 0 and $\pm p$. We will show that for trace 0, if there are two $\mathbb{F}_{p^2}$-isomorphism classes, both of these isomorphism classes have j-invariant 1728.

If $t = 0$, we know that the $\mathbb{F}_{p^2}$-Frobenius endomorphism F satisfies $F^2 + p^2 = 0$ (see Theorem 2.25). Let $F^{(p)}$ denote the p-power Frobenius isogeny (as seen in the proof of Proposition 3.11). Notice that $(F^{(p)})^2 = F$. In the proof of Proposition 3.11, we show that $[p] = \widehat{F^{(p)}}_{\text{sep}} \circ (F^{(p)})^2$. Denoting the inverse of $\widehat{F^{(p)}}_{\text{sep}}$ as ϕ, we obtain $F = \phi \circ [p]$.

We thus have $0 = F^2 + p^2 = p^2\phi^2 + p^2 = p^2(\phi^2 + 1)$. This means that ϕ is such that $\phi^2 = -1$. Therefore, ϕ is an automorphism of order 4, which is equivalent to having $j(E) = 1728$.

A similar argument holds for the cases of traces $\pm p$, but here we show that for such traces an automorphism of order 6 exists. This is equivalent to $j(E) = 0$.

From this, we conclude that if $j(E) \notin \{0, 1728\}$, the trace is not equal to 0 or $\pm p$. But E is supersingular, which means we must have trace equal to $2p$ or $-2p$. If the trace of E is $2p$, the quadratic twist of E will have trace $-2p$.

For the cases of $j(E) \in \{0, 1728\}$, which is supersingular if $p \equiv 3 \bmod 4$ for $j(E) = 1728$ and $p \equiv 2 \bmod 3$ for $j(E) = 0$, we have a different approach. We know that the curve given by $y^2 = x^3 + x$ has j-invariant 1728, while the curve given by $y^2 = x^3 + 1$ has j-invariant 0. Both of these curves are defined over $\mathbb{F}_p$, and the Hasse bound tells us that any supersingular elliptic curve defined over $\mathbb{F}_p$ must have trace 0. In other words, the $\mathbb{F}_p$-Frobenius map $F^{(p)}$ satisfies $(F^{(p)})^2 + p = 0$, so $(F^{(p)})^2 = -p$. Now, the $\mathbb{F}_{p^2}$-Frobenius F is the $\mathbb{F}_p$-Frobenius squared, so the F is the map $[-p]$. Therefore we obtain $p^2 + tp + p^2 = 0$, and so we have $t = -2p$.

In any case, for any supersingular j-invariant, we can find a curve defined over $\mathbb{F}_{p^2}$ with trace $-2p$ with that j-invariant. It follows that these curves all have $(p+1)^2$ $\mathbb{F}_{p^2}$-points. Because the $\mathbb{F}_{p^2}$-Frobenius is $[-p]$, and we know the Frobenius map fixes all $\mathbb{F}_{p^2}$-points, we have that $(\mathbb{Z}/(p+1)\mathbb{Z})^2$ is the structure of the $\mathbb{F}_{p^2}$-points. Since these curves all have the same amount of $\mathbb{F}_{p^2}$-points, Theorem 2.26 tells us that all of these curves are also $\mathbb{F}_{p^2}$-isogenous.

Thus, we conclude that the supersingular ℓ-isogeny graph is isomorphic to the graph of $\mathbb{F}_{p^2}$-isomorphism classes of curves with trace $-2p$ with isogenies defined over $\mathbb{F}_{p^2}$. □

By taking twists, the graph is also isomorphic to the graph of $\mathbb{F}_{p^2}$-isomorphism classes of curves with trace $2p$. In practice, we will consider the supersingular isogeny graph as the graph with curves of trace $-2p$, since the well-known equations for curves with j-invariants 0 and 1728 already have this trace.
Now that we know that every supersingular j-invariant has a curve with trace $-2p$ in it, we can show that every supersingular j-invariant has a curve in Montgomery form with $B = 1$.

Lemma 3.13. *Let $j' \in \mathbb{F}_{p^2}$ be a supersingular j-invariant. Then there exists a Montgomery curve $E : y^2 = x^3 + Ax^2 + x$ such that $j(E) = j'$.*

Proof. See the proof of Proposition 5.10 in [39]. □

An idea of the proof is as follows. The cases $j' = 0$ and $j' = 1728$ follow by noticing that the j-invariant of a Montgomery curve is given by $\frac{256(A^2-3)^3}{A^2-4}$. For $j' = 0$ we need that A is a square root of 3. Such an A lies in $\mathbb{F}_{p^2}$, so we use this value of A to define our Montomery curve with j-invariant 0. For $j =' 1728$ we simply take $A = 0$ as $y^2 = x^3 + x$ has j-invariant 1728.
For other j-invariants, we know that there exists a curve E' in Weierstrass form such that $j(E') = j'$ and the 4-torsion is fully contained in $E'(\mathbb{F}_{p^2})$. By making use of the *division polynomial* (which we will not define here), one can show that the conditions to be able to define an isomorphism to a Montgomery curve $By^2 = x^3+Ax^2+x$ are met. Then, this j-invariant also contains the Montgomery curve $y^2 = x^3+Ax^2+x$.

Besides the knowledge about the j-invariant, we can also count the amount of supersingular j-invariants easily.

Proposition 3.14. *The number of $\overline{\mathbb{F}}_p$-isomorphism classes of supersingular elliptic curves in characteristic $p > 3$ is equal to*

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

Proof. See [59], Theorem V.4.1c. □

We can view this statement as saying that we always have at least $\lfloor \frac{p}{12} \rfloor$ isomorphism classes, adding an isomorphism class for j-invariants 0 or 1728, should those j-invariants be supersingular. For $j = 0$ this occurs if p is 2 mod 3, and for $j = 1728$ this occurs if p is 3 mod 4. For the cases $p = 2$ and $p = 3$, we have one j-invariant, being $j = 0$ and $j = 1728$, respectively.

Remark 3.15. *When we look at ℓ-isogeny graphs and potential endomorphisms, we are generally interested in cyclic isogenies and cyclic endomorphisms. Recall that these are isogenies with a cyclic kernel. These isogenies are interesting because they correspond precisely with the non-backtracking walks that one can take on a supersingular isogeny graph.*
Because it is cyclic, a cyclic kernel contains precisely one subgroup of order ℓ, which decides which direction to walk to from the starting point. After this, the points of order ℓ^2 in the cyclic kernel will be turned into points of order ℓ when taken through the first isogeny, and it once again decides which isogeny to take. This process continues until we arrive at the end of our walk. Notice that this automatically prevents backtracking (that is, instantly going back along the dual of an isogeny that we just took). This is because taking the dual would equate to a multiplication-by-ℓ map, meaning the entire ℓ-torsion would be in the kernel and it is no longer cyclic.

The supersingular isogeny graph turns out to have a powerful property:

Theorem 3.16 (Pizer)**.** *The supersingular ℓ-isogeny graph is a Ramanujan graph.*

This result was proven by Pizer [49][50], but the proof is far beyond the scope of this thesis. Note that the supersingular ℓ-isogeny graph is directed if $p \neq 1 \bmod 12$ because of the presence of j-invariants 0 or 1728, as explained after Definition 3.9. This means that we can not directly apply the theory from Section 3.1. However, the eigenvalues of the supersingular isogeny graph are still real, which is enough for the result to be proven.
Pizer's result tells us that supersingular isogeny graphs are a good candidate for potential use in cryptographic protocols, because Ramanujan graphs are thought to have lower mixing constants than general d-regular graphs. Sterner [62] gives a conjectural upper bound for the mixing constant of the 2-isogeny graph based on experimental results:

Conjecture 3.17. *Let $k_{2,p}$ be the mixing constant for the supersingular 2-isogeny graph in characteristic p. We have the following upper bound:*

$$k_{2,p} \leq \log_2(p) + \log_2(\log_2(p)) + 1.$$

Later, he generalises this to a conjectural upper bound for the mixing constant of general ℓ-isogeny graphs:

Conjecture 3.18. *Let $k_{\ell,p}$ be the mixing constant for the supersingular ℓ-isogeny graph in characteristic p. We have the following upper bound:*

$$k_{\ell,p} \leq \log_\ell(p) + \log_\ell(\log_\ell(p)) + 1.$$

The following theorem gives concrete bounds on the probability of the outcome of a random walk.

Theorem 3.19. *Given a prime number p, let j_0 be a supersingular j-invariant in characteristic p, N_p be the number of supersingular j-invariants in characteristic p (see Proposition 3.14), and let $n = \prod_i \ell_i^{e_i}$ be an integer where ℓ_i are primes. Let $\widehat{j}$ be the j-invariant reached by a random walk of degree n starting at j_0. Then for every j-invariant $\overline{j}$ we have*

$$\left| Pr[\widehat{j} = \overline{j}] - \frac{1}{N_p} \right| \leq \prod_i \left(\frac{2\sqrt{\ell_i}}{\ell_i + 1} \right)^{e_i}.$$

Proof. See Theorem 1 of [25]. □

We will consider a small example to get a better idea of how supersingular isogeny graphs work. For our example, we take $p = 127$ and $\ell = 2$. Since -1 is not a square in $\mathbb{F}_{127}$, we will consider $\mathbb{F}_{127^2}$ as $\mathbb{F}_{127}[i]$, where i is such that $i^2 = -1$. In Figure 2, we can see what the 2-isogeny graph over this field looks like.

We observe that there are 11 vertices in the graph, which is expected since $\lfloor 127/12 \rfloor = 10$, and we have $p \equiv 3 \bmod 4$ but not $p \equiv 2 \bmod 3$.
On j-invariants 77 and 126, we can see endomorphisms of degree ℓ. The endomorphism on j-invariant 77 is there because this is the curve with j-invariant 1728 (as $1728 \equiv 77 \bmod 127$), which has a non-trivial automorphism group. This causes the two non-equivalent isogenies that go from j-invariant 77 to j-invariant 95 to have equivalent dual isogenies, as explained after Definition 3.9. In fact, j-invariant 1728 always goes to the same j-invariant in the 2-isogeny graph. This is proven in Proposition 6.13.
On the other hand, j-invariant 126 has a trivial automorphism group. However, we can still predict that this endomorphism will occur via the *modular polynomial*. We will avoid technical details, but in short, the modular polynomial of degree N parametrizes the j-invariants of elliptic curves that are related via an N-isogeny. Normally, this polynomial has two variables, but if we consider both variables the same, solutions to that polynomial provide precisely the j-invariants where endomorphisms of degree N occur. For the degree 2 modular polynomial, $j = 8000$ is a solution, which reduces to 126 mod 127. See Example 1 in Section 2.5 of [37] for a more complete explanation using modular polynomials.

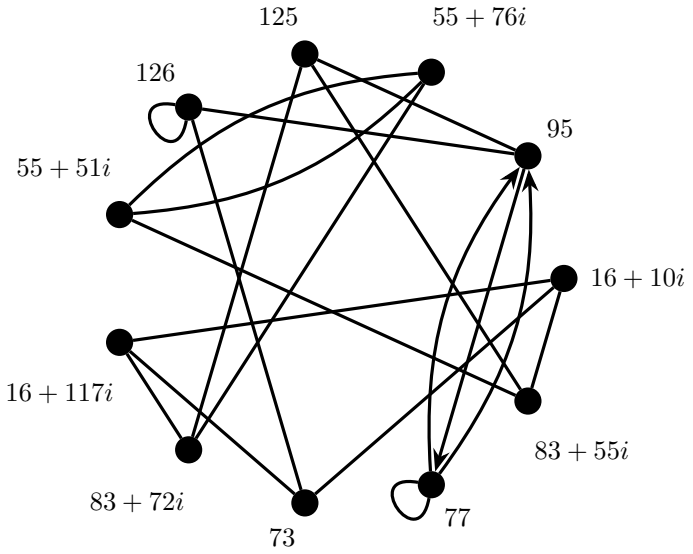

Figure 2: $G_2(127)$

We can also see that between j-invariants $55 + 51i$ and $55 + 76i$, there are two non-equivalent isogenies. This can also be explained by the modular polynomial: these two values of j provide a root of multiplicity 2. Since this is connected to the modular polynomial, we know that this behaviour does not occur often, so when working with cryptographically large p this is not an issue.
One final thing of note is the fact that there are two non-equivalent isogenies going from j-invariant 77 to j-invariant 95, while only one isogeny goes the other way. As was explained earlier, this is due to the fact that j-invariant 77 has a non-trivial automorphism group, which causes the two non-equivalent isogenies to have equivalent duals.

4 Cryptography

In this section, we discuss some important notions in cryptography that will be used in the rest of this thesis. For a more comprehensive treatment of the field of cryptography, an interested reader should refer to sources such as [28].

4.1 Basics of cryptography

4.1.1 General background

The need to communicate in a secure manner, without allowing unwanted eavesdroppers to know what is being communicated, is ancient. The Romans already used ciphers to encrypt military tactics by shifting each letter of their message by a previously agreed upon amount of letters in the alphabet. This process is known as the Caesar cipher. Much has happened since, but the need for secure communication remains. Nowadays, we would like to be able to communicate securely over the internet with, for example, banks or other people. The Caesar cipher is not a safe way to do this, since it is easy to simply brute force all possible options. There are modern and widely used methods to encrypt data that are assumed to be secure. The most important one among them is the Advanced Encryption Standard (AES). It is an example of a symmetric-key algorithm. A **key**, in the context of cryptography, is something that can encode or decode data. A **symmetric-key algorithm** is an algorithm that uses the same secret key for the encryption of plaintext and the decryption of ciphertext.
This means that with a safe symmetric-key algorithm, anyone with the secret key will be able to securely send and receive information over public channels, where anyone can see what is being sent, by only sending encrypted data. While this method works well, and AES is widely used to securely communicate, a problem still arises: if two parties want to securely communicate using a secret key, how can they both safely obtain a copy of this key? In ancient times, one party could come up with a key and send a messenger to the other party with a copy of the key. On the internet, this cannot be done this easily, as communication happens over public channels which can be read by malicious parties. The solution is to send information over public channels in such a way that the public information by itself is useless to determine the secret key, while both parties can use the public information to privately determine the secret key. This is known as key exchange.

Definition 4.1. *A **key exchange** is a method to safely exchange cryptographic keys between two parties over an insecure channel.*

A simple and famous example of a public key exchange is the Diffie-Hellman key exchange.

Example 4.2 (Diffie-Hellman key exchange)**.** *First, Alice and Bob agree on some public parameters, namely a cyclic group G of order n and a generator $g \in G$. Then, the following steps happen:*

1. *Alice and Bob respectively choose random integers $a, b \in \{1, \ldots, n-1\}$. We call a Alice's secret key, and b is Bob's secret key.*
2. *Alice computes her public key $A := g^a$, while Bob computes his public key $B := g^b$.*
3. *Alice and Bob exchange A and B over a public channel.*
4. *Alice computes $S := B^a = g^{ba}$, while Bob computes $A^b = g^{ab} = g^{ba} = S$. Thanks to commutativity, Alice and Bob have computed their shared secret key.*

The current widely-used implementation is called Elliptic Curve Diffie-Hellman (ECDH). This uses a cyclic subgroup of the group of points on an elliptic curve over some large finite field as its group G.
An eavesdropper (whom we will call Eve) only knows the generator g and the public keys A and B. The only

known way for Eve to recover Alice and Bob's private key is by finding the discrete logarithm of either of their public keys. For certain cyclic groups, such as those used in ECDH, this is assumed to be computationally infeasible to do. This is known as the Discrete Logarithm Problem (DLP). While this computational difficulty is currently the case, this might not remain true in the future, as will be discussed in more detail in Section 4.2.

There exist very few cryptographic algorithms that can be proven to be safe. One notable exception is the one-time pad. The one-time pad uses a single-use key that is of the same size as the message that is encrypted. By combining each bit of the message with a bit of the key, we obtain a provably unbreakable encryption system. The downside of the one-time pad is that it is extremely inefficient, making it unfeasible for practical use. For this reason, almost all cryptographic protocols rely on the assumption that some problem is difficult to solve.
Besides Diffie-Hellman being based on the difficulty of solving DLP, there are many other cryptographic algorithms based on solving other mathematical problems. We make the assumption that such a problem cannot be solved efficiently, that is, that the problem is *hard* to solve. The term *efficiently* refers to solving such a problem in polynomial time. Such an assumption is called a **computational hardness assumption**. Other than DLP, the most famous example of a problem assumed to be hard is that of prime factorisation. The best known application of this hardness assumption is RSA. There is no publicly known efficient method of efficiently factorising a number into prime factors on a classical computer. The story is different for quantum computers, see Section 4.2.

We will now discuss some important cryptographic concepts that will motivate our study of isogeny-based commitment schemes. The formal definitions of these terms are rather technical, so as to not broaden the scope of this thesis too much, we will introduce these terms in a somewhat less formal way. See [28] for more details.
Symmetric-key algorithms and key exchanges solve a lot of security concerns that arise from trying to communicate digitally. However, some issues remain. One such issue is that with key exchange, it is impossible to verify that the message Bob receives is actually from Alice. An adversary could intercept Alice's message and send their own message to Bob, pretending to be Alice. This would allow them to obtain a shared key with Bob, while Bob thinks he is making a shared key with Alice. This issue is resolved by the concept of a digital signature.

Definition 4.3. *A **digital signature** is a method to verify the authenticity of a digital message.*

A digital signature works by adding a signature made by some private key. A public key is then distributed to everyone, which allows anyone who receives the message to verify that the sender of the message is indeed who they are claiming to be.
Another important notion is a hash function.

Definition 4.4. *A **hash function** is a function that takes as input some message of arbitrary length and outputs something of fixed length.*

Hash functions are used in a wide variety of contexts, but in the context of cryptography, **cryptographic hash functions** are studied. Cryptographic hash functions have some specific properties that make them useful for cryptographic applications, including digital signatures:

- Given just the output, it is hard to find some input of the function that produces said output. This is called *pre-image resistance.*
- Given an input and its output, it should be difficult to find another input that produces the same output. If one finds such an input, this is called a collision, and the difficulty of finding a collision is called *collision resistance.*

We will see an example of a cryptographic hash function based on isogenies in Section 6.1.1.

We now discuss some more terms from cryptography that are relevant for us.
If we have an algorithm, such as a hash function or a symmetric-key algorithm, an important property to look at is its **time complexity**. This measures how fast an algorithm runs. More formally, it measures the amount of steps that an algorithm needs to take depending on the size of the input. If an algorithm is in *polynomial time*, it means that the required amount of operations can be bounded by a polynomial. This is usually considered fast enough for practical implementations, though very large polynomial time algorithms can still be too slow. Algorithms can also be in *exponential time*, meaning the amount of required steps increases exponentially if the input increases in size. This is considerably worse and are these algorithms are generally considered to be too slow for practical implementations, but similar to the case of polynomial time, especially fast exponential time algorithms can be practically used. Algorithms can have other time complexities, but these two are the most important ones.
An algorithm is *probabilistic* if its output is based on some probability distribution. Crucially, it means that giving it the same input multiple times does not guarantee the same outcome. An algorithm that, given the same input, always provides the same output is called deterministic. We can consider deterministic algorithms as a subset of probabilistic algorithms, where the probability distribution is one that picks one outcome with probability 1.
A cryptographic hash function is a candidate for a **one-way function**. One-way functions are functions that are easy to compute on any input, but given the image of an input, it is hard to compute an element of the preimage. We can think of the terms easy and hard as the computations being in polynomial time and exponential time, respectively. One-way functions are conjectured to exist, but there is no known proof of their existence. However, they are used extensively in cryptography. A specific class of one-way functions are **trapdoor functions**. These are one-way functions with the extra property that there exists some secret information t that makes the preimage easy to compute. An example of this is in RSA, in which a secret x is encrypted as $x^e \bmod n$, with n the product of two large primes. It is hard to compute x given just x^e, but with knowledge of the factorisation of n, it becomes easy to find x.
If we have a cryptographic scheme, we are interested in how secure it is. One measurement of this is the **security parameter**. The security parameter is commonly denoted by λ. It is the value that dictates how many bits of security will be required. To achieve λ bits of security, an attacker should require 2^λ operations to break the scheme. We typically measure the security of a cryptographic scheme in terms of the security parameter λ. We can then vary λ according to the required level of security.

4.1.2 Commitment schemes

Before we can talk about commitment schemes, we first need to define the notion of negligibility. We can intuitively think of something as *negligible* if it is so small that it can be safely ignored. There is a more formal way to define it, which we will need, since we need negligibility in some proofs.

Definition 4.5. *A **negligible function** is a function $\mu : \mathbb{N} \to \mathbb{R}$ such that for every polynomial p with $p(x) > 0$ for all x, there exists a positive integer N_p such that $|\mu(x)| < \frac{1}{p(x)}$ for all $x > N_p$.*

If we say that some function is negligible, it means that there exists a negligible function in the same variables as the function that is always larger than the function.

We are now ready to introduce commitment schemes. Roughly speaking, a commitment scheme allows a sender to commit to a certain value, without having to reveal it. After the value is revealed, a verifier is able to check that the commitment matches the value that was revealed. Formally, we define a commitment scheme as follows.

Definition 4.6 (Commitment Scheme)**.** *A commitment scheme consists of three algorithms, called* ***KeyGen****,* ***Commit****, and* ***Open****. The first algorithm,* ***KeyGen****, is a probabilistic polynomial time algorithm that takes as input the security parameter λ and outputs the public parameters PP and the message space $\mathcal{M}$ used in the protocol. The second algorithm,* ***Commit****, is a deterministic algorithm that takes as input the public parameters as well as a message $m \in \mathcal{M}$ and some randomness r, where r is drawn at random from some randomness space $\mathcal{R}$. The output is a value c, which we call the commitment to m. The third algorithm,* ***Open****, is a deterministic algorithm that takes as input the public parameters, m, r and c, and returns a boolean $b \in \{0, 1\}$ depending on whether c is a valid commitment to m given r.*

As an analogy to understand what a commitment scheme is, we can think about the situation where Alice and Bob would like to remotely play a game of rock-paper-scissors. Alice decides on a move and sends it to Bob. Upon receiving Alice's move, Bob makes his own move. However, Alice does not believe that Bob will play fair, as he could simply look at Alice's choice and respond with the winning move. On the other hand, if Alice tells Bob she made a choice while keeping the choice a secret, Bob does not trust Alice, as after Bob makes his choice, she could pretend that her secret choice was the move that beats Bob. To solve this, Alice puts her choice (rock, paper or scissors) on a piece of paper in a locked box, and sends the box to Bob. After Bob states his move to Alice, she can send him the key to the box, and Bob can verify Alice's original move. This prevents cheating from both sides.
A commitment scheme is an example of a cryptographic primitive. We can think of cryptographic primitives as the building blocks that make up larger cryptographic protocols. For example, commitment schemes are used in the construction of zero knowledge proofs, which is a protocol where one convinces the truth of some statement to a verifier, without revealing anything about the statement itself.

To make a commitment scheme useful in a cryptographic setting, it needs to satisfy two properties, called the **hiding property** and the **binding property**. The hiding property ensures that anyone who receives just the commitment c is unable to derive any information about the message m that c is a commitment of. The hiding property can be thought of in terms of the hiding game. In the hiding game, an adversary $\mathcal{A}$ gets to pick any two messages after seeing the public parameters. Then, one of these two messages is chosen at random and is committed to. The adversary wins if it can correctly guess which of the two messages was committed to. A commitment scheme is hiding if no adversary can win the hiding game with (significantly) better odds than random guessing. More formally, we have the following:

Algorithm 1 Hiding Game

1: $PP \leftarrow \mathbf{KeyGen}(\lambda)$
2: $(m_0, m_1) \leftarrow \mathcal{A}(PP)$
3: $b \in_R \{0, 1\}$
4: $c = \mathbf{Commit}(PP, m_b, r)$
5: $b' \leftarrow \mathcal{A}(c)$
6: **return** $b == b'$

This formulation of the hiding property highlights the need for the randomness r in the commitment scheme, as without it, an adversary could simply compute the resulting commitments from both messages and determine which message is correct, easily winning the hiding game.
The binding property ensures that it is difficult to find two messages that have the same commitment. This can be thought of in terms of the binding game. In the binding game, an adversary $\mathcal{A}$ has to choose two messages along with two instances of randomness. It wins the game if the messages are distinct and they both commit to the same value. A commitment scheme is binding if no adversary has a (significant) probability to win the binding game. More formally, we have the following:

Algorithm 2 Binding Game

1: $PP \leftarrow \mathbf{KeyGen}(\lambda)$
2: $(m, m', r, r', c) \leftarrow \mathcal{A}(PP)$
3: **return** $(m \neq m')$ && $(\mathbf{Open}(PP, m, r, c) == \mathbf{Open}(PP, m', r', c) == 1)$

Returning to our locked box analogy, we can see how this process informally satisfies both properties. Because Bob cannot see inside the box and he cannot open it, he does not obtain any information about the content of the box. This gives us the hiding property. Since Alice cannot change the contents of the box without Bob noticing, we also have the binding property.
We can formalise how strong the hiding and binding properties of a commitment scheme are with the hiding advantage and the binding advantage.

Definition 4.7. *Given a commitment scheme C with security parameter λ and an adversary $\mathcal{A}$, we define the hiding advantage of $\mathcal{A}$ as*

$$Adv^{hid}_{C,\lambda}(\mathcal{A}) := 2\left|Pr(\mathcal{A} \text{ wins the hiding game}) - \frac{1}{2}\right|,$$

and we define the binding advantage of $\mathcal{A}$ as

$$Adv^{bind}_{C,\lambda}(\mathcal{A}) := Pr(\mathcal{A} \text{ wins the binding game}).$$

We say that C is information-theoretically hiding if for all adversaries $\mathcal{A}$ there exists a negligible function μ such that $Adv^{hid}_{C,\lambda}(\mathcal{A}) \leq \mu(\lambda)$. We call C computationally hiding if the same holds, but only for probabilistic polynomial time adversaries, not all adversaries. A commitment scheme has perfect hiding if the hiding advantage is zero for any adversary. Analogous definitions hold for the binding property.

We can prove the following reformulation of the hiding advantage:

Lemma 4.8. *Given a commitment scheme C and an adversary $\mathcal{A}$, we have*

$$Adv^{hid}_{C,\lambda}(\mathcal{A}) = \left|Pr(\mathcal{A} \text{ picks } b' = 1 \mid b = 1) - Pr(\mathcal{A} \text{ picks } b' = 1 \mid b = 0)\right|.$$

Proof. Using Definition 4.7, we have

$$\begin{aligned}
\mathrm{Adv}^{hid}_{C,\lambda}(\mathcal{A}) &= 2\left|Pr(\mathcal{A} \text{ wins the hiding game}) - \frac{1}{2}\right| \\
&= 2\left|Pr(b' = b \text{ and } b = 1) + Pr(b' = b \text{ and } b = 0) - \frac{1}{2}\right| \\
&= 2\left|Pr(b = 1)Pr(b' = b \mid b = 1) + Pr(b = 0)Pr(b' = b \mid b = 0) - \frac{1}{2}\right| \\
&= 2\left|\frac{1}{2}Pr(b' = b \mid b = 1) + \frac{1}{2}Pr(b' = b \mid b = 0) - \frac{1}{2}\right| \\
&= |Pr(b' = 1 \mid b = 1) + Pr(b' = 0 \mid b = 0) - 1| \\
&= |Pr(b' = 1 \mid b = 1) + 1 - Pr(b' = 1 \mid b = 0) - 1| \\
&= |Pr(b' = 1 \mid b = 1) - Pr(b' = 1 \mid b = 0)| \\
&= \left|Pr(\mathcal{A} \text{ picks } b' = 1 \mid b = 1) - Pr(\mathcal{A} \text{ picks } b' = 1 \mid b = 0)\right|.
\end{aligned}$$

□

It would be very useful if a commitment scheme has both perfect hiding and perfect binding. However, this is impossible.

Lemma 4.9. *There does not exist a commitment scheme that is both perfectly hiding and perfectly binding.*

Proof. Suppose there exists a commitment scheme that is both perfectly hiding and perfectly binding. Let c be a commitment of the pair (m, r). Because the scheme is perfectly hiding, there must exist (m', r') with $m \neq m'$ that also produces c as its commitment. If not, then given enough time, an adversary could always win the hiding game by choosing one of the messages to be m and trying every possible value of r. However, the existence of two pairs $(m, r), (m', r')$ with $m \neq m'$ that both produce the commitment c goes against the perfect binding property, so we have a contradiction. □

The above lemma also holds if one of the binding and hiding properties are not perfect, but information-theoretic. The reasoning is similar, because information-theoretic binding and hiding allow for computationally unbounded adversaries. However, given perfect binding, computational hiding is still achievable. Similarly we can still have computational binding, given perfect hiding.

4.1.3 Constructing a commitment scheme from a generic hash function

Any hash function allows one to construct a commitment scheme. This is done as follows: let H be a hash function. To commit to a message m, one generates a string r that satisfies the security parameter. The commitment is given by $c := H(m, r)$. Intuitively, we obtain binding because a hash function should be collision-resistant, meaning that it should be hard to find another input for H that produces the same result. The preimage resistance of a hash function says that the output of a hash function should not tell you anything about its input, which gives us the hiding property.

One way to prove binding and hiding of this generic construction is by making use of the so-called *Random Oracle Model* (ROM). In this model, the hash function H is replaced by an oracle $\mathcal{O}$, which is a function chosen uniformly at random from all functions from the domain to the codomain of H. Essentially, this turns H into a truly random function. In security proofs, the adversary is allowed to query the oracle as often as they want. With this approach, we can prove binding and hiding of this generic construction.
However, the ROM is not reflective of reality. We can not be sure that security in the ROM translates to security in the real world, when using an actual hash function. Therefore, when we give a concrete instantiation of this construction, we make use of the underlying properties specific to the hash function used. This allows us to prove binding and hiding more concretely for specific instantiations. The commitment schemes that we study in Section 6 are concrete instantiations of this generic construction.

4.2 Post-quantum cryptography

In this section, we introduce the field of post-quantum cryptography and the various approaches that exist within it.
Many popular protocols used in modern cryptographic applications are variants of ECDH or RSA. However, in 1994, Peter Shor published algorithms [56][57] for a quantum computer that break both the DLP and prime factorisation problem in polynomial time. Thus, the existence of a sufficiently fast quantum computer would compromise digital security worldwide.
Luckily, as far as we know, sufficiently fast quantum computers do not yet exist. Research and development is rapid, however, so in 2014, the National Institute for Standards and Technology (NIST) launched a competition to create and standardise algorithms that are also safe against quantum computers [44]. The study of quantum-safe algorithms is also known as Post-Quantum Cryptography (PQC). In August 2024, three standards were finalised after multiple selection rounds in the NIST competition, with a fourth standard still in development. Of course, research into PQC is still ongoing.
In PQC, there are various hardness assumptions. These assumptions have led to the development of many

cryptographic protocols that are thought to be safe against quantum computers. We will briefly go over the most popular approaches to PQC by giving intuition about the trapdoor function used.

Lattice-based cryptography is based on the use of lattices. A lattice in $\mathbb{R}^n$ is the subgroup generated by any set of n linearly independent vectors. A lattice has more than one basis. Hard problems for lattices are the Shortest Vector Problem (SVP), which asks to find the shortest non-zero element of a given lattice, and the Closest Vector Problem (CVP), which asks, given an arbitrary point in $\mathbb{R}^n$, to give the closest lattice point. The CVP trapdoor function is the computation of a point using some basis of a lattice and applying a small error. This is easy to do, but given just the point, solving the CVP is hard. However, given a 'good' basis, where good means that the basis vectors are short and as orthogonal as possible, this becomes very easy. This good basis is the private key. The public key is another basis, but this one is bad, which makes the CVP much more difficult to solve, and using certain algorithms may even lead to an incorrect solution. Lattice-based methods are popular and are regarded as very promising for cryptography, mainly due to their quick computation times. They are also versatile, allowing for many different primitives. On the other hand, key sizes are relatively large, and the security has been reduced consistently via increasingly efficient attacks, forcing parameters to be increased. Two lattice-based algorithms, Crystals-KYBER and Crystals-DILITHIUM, have been selected by NIST as some of the first standardised post-quantum algorithms. These were turned into standards FIPS 203 and FIPS 204[45].

Code-based cryptography makes use of error-correcting codes. These are codes that are commonly used in communication on noisy channels, allowing for the receiver to be able to read the contents even if some parts of the code are corrupted. The hardness in code-based systems is based on the decoding problem. The decoding problem asks to reconstruct a codeword from a noisy codeword. This functions as the trapdoor. The secret information is the knowledge of an efficient decoding algorithm.
The first code-based cryptosystem is also the one that is still most popular today, and it is called the McEliece cryptosystem. In the original proposal [35], McEliece suggests the use of a secret Goppa code, obscured by other codes to produce a public key. A Goppa code is a certain type of error-correcting code. A message is then encoded by this public key, and a small error is added. Only the person with knowledge of the secret key is able to efficiently decode the ciphertext back to the original message. While this system has resisted decades of attempts to break it, a big problem is that key sizes are massive: while most methods provide key sizes of a few kilobytes, with some going as low as a few hundred bytes, the McEliece system produces keys that are multiple megabytes in size. Therefore, for some applications this system is impractical to use. Besides McEliece, another code-based protocol is HQC[26], which is due to be standardised[43].

Multivariate cryptography uses multivariate polynomials over a finite field. The trapdoor function involves a composition of some secret maps, which are easy to invert individually. However, their composition amounts to solving a system of multivariate polynomials, which is proven to be NP-complete. This allows for construction of digital signature schemes that are considered to be safe against quantum computers, such as the Unbalanced Oil and Vinegar scheme[29].

Hash-based cryptography uses the security of hash functions. Its trapdoor functionality logically comes from the hash functions used: given just the output, the input is difficult to find, so the input needs to be kept secret. Currently, the main use case is to construct digital signature schemes. The digital signature scheme SPHINCS+[5] has been standardised by NIST for digital signatures, under the standard FIPS 205[45]. Hash-based schemes work for any hash function, which is advantageous as it allows easy swapping of hash functions in case a hash function that is in use turns out to be unsafe.

Finally, there is **Isogeny-based cryptography**. The hardness assumption is usually some variation of the

Isogeny Path Problem: given two isogenous elliptic curves, provide an isogeny between them. In general, isogenies are easy to compute given their kernel. Without this information, this becomes a more difficult task, so this serves as the secret information of the trapdoor function. While isogeny-based cryptography has had a rocky history, it is nevertheless a promising field of research. Besides that, even with other approaches becoming more standardised, it is good to research alternatives in case of one of the approaches being broken. The rest of this thesis will focus on isogeny-based cryptography.

5 Isogeny-based Cryptography

One might wonder why we would be interested in isogeny-based cryptography. Within post-quantum cryptography, lattice-based cryptography receives most of the attention because of its promising efficiency. Isogeny-based cryptography is not considered to be very efficient, making it unappealing for widespread adoption. However, compared to long-established assumptions like prime factorisation and DLP, the post-quantum assumptions have not had the same level of scrutiny applied to them. Therefore, it makes sense to study various approaches. If one of them ends up being unsafe, we can switch to another approach. Besides this, studying isogeny-based cryptography more could lead to developments to improve efficiency.
In this chapter, we will sketch the landscape of the field of isogeny-based cryptography. After a brief history in Section 5.1, we discuss some of the most significant protocols in isogeny-based crypto in Section 5.2 (CRS), Section 5.3 (SIDH) and Section 5.4 (CSIDH). In Section 5.5, we discuss the most important hardness assumptions in isogeny-based cryptography. Finally, we consider recent developments and future directions for the field of isogeny-based cryptography in Section 5.6.

5.1 A brief history

Isogeny-based cryptography is a relatively young field, only starting to gain steam in the late 2010s. Despite this, its history could be considered somewhat tumultuous. In this section, we discuss some history and we give an overview of the most important work done in the field to date. A good introduction to isogeny-based cryptography can be found in the lecture notes of De Feo [22].
For decades, cryptography was mostly focused on long-established methods, including RSA and Elliptic Curve Cryptography (ECC). The first instances of isogeny-based cryptography can be found in works of Couveignes [15] from 1997 and Rostovtsev-Stolbunov [52] from 2006. However, not much interest was shown until the threat of quantum computers brought the motivation to study PQC.
Among the submitted algorithms of the NIST PQC competition, there was one that used isogenies, called SIKE (Supersingular Isogeny Key Exchange) or SIDH (Supersingular Isogeny Diffie-Hellman), initially proposed in 2011 by Jao and de Feo [27]. More specifically, SIDH used isogenies of supersingular elliptic curves. The key exchange worked similarly to the classic Diffie-Hellman key exchange. This algorithm made it to round 4 of the NIST competition. However, in 2022, Castryck and Decru [9] published a devastating attack that completely broke SIDH on a classical computer. The attack allows one to recover the isogenies used to compute the private key, thereby giving access to the private key. Independently, Maino and Martindale also published an attack on SIDH [33], and both attacks were later improved upon and generalised by Robert [51].
In the meantime, other algorithms were developed that make use of isogenies. The most important examples are SQISign [24] and CSIDH [10]. Fortunately, both of these protocols are unaffected by the attack on SIDH, since that attack is based on specific public information that was deemed safe to share. Furthermore, in the years following the SIDH attack, the method used in the attack has been turned into a method to create new and more efficient algorithms, including SQISignHD [17] and (Q)FESTA [3][41].

5.2 The Couveignes-Rostovtsev-Stolbunov protocol

The first idea of a protocol based on isogeny graphs was given by Couveignes [15] in 1997. However, his work was only circulated in private circles until the same idea was independently found in 2006 by Rostovtsev and Stolbunov [52]. After Rostovtsev and Stolbunov published their works, Couveignes published his original notes from 1997. The key idea from these works is to define a group action of the class group of a given imaginary quadratic order on the set of elliptic curves with this order as its endomorphism ring. Since the full endomorphism ring of an ordinary elliptic curve is an imaginary quadratic order, these curves were an obvious candidate to do this with. This protocol is commonly referred to as the Couveignes-Rostovtsev-Stolbunov (CRS) protocol. We will briefly go over the mathematics behind the CRS protocol.

Let E be an elliptic curve defined over a finite field $\mathbb{F}_q$, with $\mathcal{O}$ its endomorphism ring. Let $\mathrm{Cl}(\mathcal{O})$ denote the class group of $\mathcal{O}$. Given an ideal $I \in \mathrm{Cl}(\mathcal{O})$ of norm coprime to q, we define the I-torsion subgroup of E as

$$E[I] = \{P \in E : \alpha(P) = 0 \text{ for all } \alpha \in I\}.$$

By Proposition 2.12 we can define a separable isogeny $\phi_I : E \to E_I$ with kernel $E[I]$. Restricting now to the case of ordinary elliptic curves, we can define a group action of $\mathrm{Cl}(\mathcal{O})$ on the set $\mathrm{Ell}_q(\mathcal{O})$, the set of elliptic curves with endomorphism ring equal to $\mathcal{O}$. For an element I of $\mathrm{Cl}(\mathcal{O})$, the action is given by constructing the isogeny with the I-torsion subgroup as its kernel, and taking the codomain of this isogeny. This action is transitive and free. This means that for each pair $E, E' \in \mathrm{Ell}_q(\mathcal{O})$ there exists $I \in \mathrm{Cl}(\mathcal{O})$ such that $I \cdot E = E'$. Also, besides the identity element, no element of $\mathrm{Cl}(\mathcal{O})$ fixes any element of $\mathrm{Ell}_q(\mathcal{O})$.
Suppose we have a prime ℓ that splits in $\mathcal{O}$, that is, we can write $\ell\mathcal{O} = \mathfrak{l}\bar{\mathfrak{l}}$, where $\mathfrak{l}$ is a prime ideal in the class group. Having $\mathfrak{l}$ and $\bar{\mathfrak{l}}$ act on an elliptic curve in $\mathrm{Ell}_q(\mathcal{O})$ gives us two distinct elliptic curves, where the isogenies are horizontal (recall Proposition 3.10). We thus obtain a graph of degree 2. We can arbitrarily pick one of these two isogenies to be the 'positive' direction. This process can be repeated for any other prime that splits.
We are now ready to start defining the protocol itself. Define $L = \{\ell_1, \ldots, \ell_m\}$ a set of primes that split in $\mathbb{Z}[F]$. The key idea is that if a prime splits in $\mathbb{Z}[F]$, it splits in $\mathcal{O}$ since $\mathbb{Z}[F] \subset \mathcal{O}$. Thus, there is no need to explicitly compute the endomorphism ring, as this might take a long time to do. The protocol now proceeds as follows:

1. Both Alice and Bob start on the same starting curve E.
2. Alice takes a random walk of steps in L along the positive direction, denoted by $\rho_A \in L^*$. The walk terminates on the curve $E_A := \rho_A(E)$. Due to commutativity of the class group, the order of the steps does not matter.
3. Bob does the same, performing a walk ρ_B, which terminates on the curve $E_B := \rho_B(E)$.
4. Alice and Bob exchange E_A and E_B.
5. Alice computes the shared secret E_{AB} as $E_{AB} = \rho_A(E_B)$ while Bob computes $E_{AB} = \rho_B(E_A)$.

In spirit, this protocol resembles the traditional Diffie-Hellman protocol by making use of commutativity. While the CRS protocol is interesting from a theoretical perspective, the practical side leaves much to be desired. Computations take too long to ever be useful. A 2018 paper [23] uses the CRS protocol as its basis while implementing more efficient parameters and algorithms. This provides some improvement, but not enough to make it useful in practice. However, with CRS, it was clear that isogeny-based cryptography had potential.

5.3 SIDH

The SIDH protocol was initially proposed in 2011 by Jao and De Feo [27], with various improvements made over the years, until in 2016 it was submitted in the NIST competition. The basic idea is as follows: Alice and Bob both start on the same isomorphism class of some elliptic curve E. From here, they both perform a walk on an l_A- and l_B-isogeny graph respectively, where l_A and l_B are primes. They send the resulting curve E_A or E_B to the other party, along with some additional information. Using this additional information, the other party can use the same isogeny that they used in their first walk. The result is that Alice and Bob both land on the same elliptic curve E_{AB}. The j-invariant of this curve is then used as their private key.
Associated to Alice is her prime l_A. Alice picks a torsion group $E[l_A^{e_A}]$ with basis $\langle P_A, Q_A \rangle$. She makes this

information public. She now picks a secret subgroup A of order $l_A^{e_A}$ of $E[l_A^{e_A}]$. Using Vélu's formulas, this gives her an isogeny α with kernel A. Analogously, Bob has a torsion group $E[l_B^{e_B}]$ with basis $\langle P_B, Q_B \rangle$. He now picks a secret subgroup B, giving an isogeny β with kernel B.
The problem with SIDH is the fact that to be able to reach E_{AB} from, say, E_A, Bob cannot directly apply his isogeny. After all, this isogeny goes from E to E_B. To solve this, besides E_A, Alice also sends Bob $\alpha(P_B)$ and $\alpha(Q_B)$. From this, Bob is able to compute $\alpha(B)$. To obtain E_{AB}, he computes the isogeny from E_A with kernel $\alpha(B)$. Commutativity ensures that Alice will also end up on E_{AB} from E_B with an isogeny with $\beta(A)$ as its kernel.

However, the solution introduced above also allows for the aforementioned attack. The following fact follows from the Cauchy-Schwarz inequality:

Proposition 5.1. *Let $\varphi : E \to E'$ be an isogeny of degree d. The images of $4d+1$ points on E under φ uniquely determine φ.*

Note that, in particular, the images of the basis points of the torsion subgroup in SIDH satisfy the condition of knowing at least $4d+1$ images. The attack makes use of isogenies between abelian surfaces. This allows us to construct an algorithm to recover α from the images of $\alpha(P_B)$ and $\alpha(Q_B)$. With knowledge of α, one can reach the secret key without much trouble. The attack is considered a total break since it can be performed very quickly even on a classical computer. This has rendered SIDH and variants that also use torsion point information unsafe to use. For an extensive discussion of the mathematical aspects of this attack, we refer to a master's thesis on the topic [34].

5.4 CSIDH

From its inception, there were concerns about the torsion-point information that SIDH publishes. However, the use of supersingular elliptic curves was very enticing. It is incredibly easy to control the size of the group of rational points, since for a prime $p > 3$, a supersingular elliptic curve defined over $\mathbb{F}_p$ will have $p+1$ $\mathbb{F}_p$-rational points. Because of this, efforts were made to design an implementation with supersingular curves that did not make use of torsion-point information.
The reason that the CRS protocol works is because it makes use of the class groups of the endomorphism rings of ordinary elliptic curves, which in this case is commutative. However, in the supersingular case, the full endomorphism ring is a maximal order in a quaternion algebra. This is not commutative, which provides a problem. In the CRS protocol, the class group of the endomorphism ring has to be abelian, since the key exchange relies on this commutativity. In 2018, Castryck, Lange, Martindale, Panny and Renes found a way to get around this problem. Instead of considering the full endomorphism ring, they would consider the ring of $\mathbb{F}_p$-rational endomorphisms. This ring is an order in a quadratic imaginary field, which is commutative. Thus, using the class group of this ring, they could once again define a group action, working with supersingular elliptic curves this time.
A major benefit of using supersingular curves is that for $p > 3$, a supersingular elliptic curve E defined over $\mathbb{F}_p$ will have $\#E(\mathbb{F}_p) = p+1$. A large hurdle in efficiency with CRS is the fact that we want a curve of highly composite group order for the rational points. This is because having small prime factors in this group order allows for isogenies of small degree, which are more efficient to compute. With CSIDH, you simply pick a prime of the form $4 \cdot \ell_1 \cdot \ldots \cdot \ell_n - 1$, where the ℓ_i are small distinct odd primes. The factor 4 that is present here is to ensure that $p \equiv 3 \bmod 4$, which ensures that the curve $E(1728) : y^2 = x^3 + x$ is supersingular with $\mathrm{End}_{\mathbb{F}_p}(E(1728)) = \mathbb{Z}[F]$. We can use $E(1728)$ as the starting curve in CSIDH.
We are interested in curves that have $\mathbb{F}_p$-endomorphism ring isomorphic to $\mathbb{Z}[F]$, and for the right choice of p the authors prove a very useful statement:

Proposition 5.2. *Let $p \geq 5$ be a prime such that $p \equiv 3 \bmod 8$, and let $E/\mathbb{F}_p$ be a supersingular elliptic*

curve. Then $\mathrm{End}_{\mathbb{F}_p}(E) = \mathbb{Z}[F]$ *if and only if there exists* $A \in \mathbb{F}_p$ *such that* E *is* $\mathbb{F}_p$*-isomorphic to the curve* $E_a : y^2 = x^3 + Ax^2 + x$*. Moreover, if such an* A *exists then it is unique.*

This allows for two things: first, any isomorphism class can be stored as a single value of A, which is efficient for key storage. Second, the form given above is the Montgomery form, which allows for very speedy computations. The mathematical theory of the group action also works without the Montgomery form, but on the practical side the use of the Montgomery form is what makes CSIDH viable to implement.
There is much more to be said about CSIDH and its design choices, as well as the mathematical concepts behind it. An extensive analysis of this can be found in a master's thesis on this topic [18].

5.5 Hardness assumptions in isogeny-based cryptography

There are a number of hardness assumptions that are commonly used in isogeny-based cryptography. Generally, these problems rely on the difficulty of computing an unknown isogeny, even with knowledge of the domain and codomain.
The problem known as the Computational Supersingular Isogeny Problem is the one that was used to claim security of SIDH. It asks to compute an isogeny of a certain degree, given knowledge of the action of the isogeny on a torsion subgroup of similar size as the degree.

Problem 5.3 (Computational Supersingular Isogeny Problem)**.** *Let* p *be a prime and let* E *and* E' *be two supersingular elliptic curves over* $\mathbb{F}_{p^2}$*. Let* A *and* B *be two integers and let* ϕ *be an isogeny from* E *to* E' *of degree* A*. Given* p*,* E*,* E'*,* A*,* B*, and the action of* ϕ *on a torsion subgroup of* E *of order* B*, compute* ϕ*.*

In the original variant of this problem, A and B were powers of small primes. However, the formulation was generalised later. This problem famously has a polynomial time solution from the Castryck-Decru attack [9] and should thus never be used for real applications.
However, removing the part that provides torsion information, we arrive at a problem that (along with variants) is often used, called the ℓ-Isogeny Path Problem. This problem asks for a path from one curve to another in the ℓ-isogeny graph:

Problem 5.4 (ℓ-Isogeny Path Problem)**.** *Given a prime* p *and two supersingular elliptic curves* E *and* E' *over* $\mathbb{F}_{p^2}$*, find a path from* E *to* E' *in the* ℓ*-isogeny graph.*

A general strategy to solve this problem is a meet-in-the-middle random walk, see Section 3 of [1]. The idea is to go on a random walk from both E and E'. We expect the two paths to collide in roughly $1.5 \cdot 2^N$ steps, where N is the size of the isogeny graph.
Another important problem is the Endomorphism Ring Problem. This problem asks, given a random supersingular curve, to compute its endomorphism ring.

Problem 5.5 (Endomorphism Ring Problem)**.** *Given a prime* p *and a supersingular elliptic curve* E *over* $\mathbb{F}_{p^2}$*, find four endomorphisms of* E *that generate* $\mathrm{End}(E)$ *as a lattice.*

In [66], it is proven that Problems 5.4 and 5.5 are equivalent. That is, either problem can be efficiently reduced to the other.
There are many small variations of the above problems that work in specific contexts. There also exist other problems. However, the problems mentioned here are what most isogeny-based protocols base themselves on.

5.6 The future

After Castryck and Decru's attack was published, there was much uncertainty in the world of isogeny-based cryptography. The protocol considered to be state of the art, one of the big contenders in the post-quantum

cryptography standardization process, was completely destroyed. Not all hope is lost, however. The algorithm initially constructed as a means to attack SIDH, was found to be usable to strengthen existing protocols, or to construct entirely new protocols.
One example is the improvement of SQISign [24]. This is a digital signature algorithm based on isogenies. A big drawback of this protocol was the fact that the isogenies involved were required to be of very big degree. This made the protocol very inefficient. Using the SIDH algorithm, it is possible to choose significantly lower degrees. This allows for very small signature sizes, which is useful for cases where memory is limited. Additionally, computations are sped up considerably. This has resulted in multiple improvements over time, such as SQISignHD [17], SQISign2D-West [2] and SQISign2D-East [42].
A novel algorithm based on the attack is the public-key encryption protocol FESTA [3]. It uses the techniques from the SIDH attack to construct a new trapdoor mechanism. A refined version with improved parameters was then proposed with QFESTA [41].
Besides this, there are plenty of protocols left unaffected by the SIDH attack, including CSIDH. These protocols are still secure and improvements are still made. So, the field of isogeny-based cryptography still has a promising future ahead.

As research continues, more and more cryptographic primitives are finding an isogeny-based counterpart. Until 2021, there existed a gap in the literature in the form of an isogeny-based commitment scheme. In order to fill this gap, in 2021 Sterner proposed an isogeny-based commitment scheme [62]. Not much research on the topic has been done since, except for a modification of Sterner's scheme from 2024 [53] that removes the need for a trusted third party. These two works will be the focus of the remainder of this thesis.

6 Isogeny-based commitment schemes

In this chapter, we discuss isogeny-based commitment schemes. In Section 6.1, we introduce Sterner's proposal for an isogeny-based commitment scheme. Next, we go over the proposed modification that removes the need for a trusted third party in Section 6.2. In Section 6.3, we go over some new results that extend the characteristics over which we can apply the method that removes a trusted third party. Finally, we briefly discuss homomorphic commitments in Section 6.4.

6.1 Isogeny-based commitment schemes

In 2021, Sterner proposed the first (publicly known) commitment schemes based on isogeny assumptions [62]. Commitment schemes that are lattice-based [67][4][38], code-based [46] and multivariate-based [48] already existed. Sterner's work introduced a new post-quantum commitment scheme. The two commitment schemes Sterner proposed work similarly: perform a walk in a full supersingular isogeny graph. However, the method to perform this walk differs. The first is based on the CGL hash function and is discussed in Section 6.1.2, while the second uses an approach similar to how isogenies in SIDH are computed and is discussed in Section 6.1.3.
Sterner's CGL hash construction is an instantiation of the generic hash construction described in Section 4.1.3, making use of the CGL hash function. The SIDH construction is another hash function, as explained in [21], which makes this another instantiation of the generic construction. The main difference between the generic construction and Sterner's proposals is that we are now working with isogenies, and as such we do not need to rely on the ROM for our security proofs. Instead, we are able to prove security based on isogeny assumptions.

6.1.1 The CGL hash function

To be able to talk about Sterner's first proposal for an isogeny-based commitment scheme, we need to introduce the CGL hash function. The CGL hash function [11] is an isogeny-based hash function proposed in 2006. The idea is to make use of the supersingular ℓ-isogeny graph and perform a walk. The security comes from the fact that this graph is an expander graph, as well as the hardness of computing isogenies between supersingular curves.
Recall from Section 3.3, given primes ℓ and p with $\ell \neq p$, the supersingular ℓ-isogeny graph has vertices corresponding to j-invariants in $\mathbb{F}_{p^2}$ and it is $(\ell+1)$-regular and connected. On this graph, we can freely perform a walk. From a given vertex corresponding to a curve E, each edge is an isogeny of degree ℓ. Thus, the kernels of these $\ell+1$ isogenies correspond precisely with the $\ell+1$ subgroups of order ℓ of the curve E. The CGL hash function is described in Algorithm 3.

We will write the CGL hash function as a function $CGL(E, \ell, P_0, m)$ that returns $j(E_k)$.
Note that if we choose the point P_i at any point during our walk, we end up with an isogeny $\varphi_i \circ \varphi_{i+1}$ from E_{i-1} of degree ℓ^2 with two independent points of order ℓ in its kernel, namely P_{i-1} and the point that generates $\ker \varphi_i$. As these two points generate all of $E[\ell]$, we conclude that the isogeny of degree ℓ^2 is the multiplication-by-ℓ map, which means that the isogeny with kernel P_i is the dual of φ_i. Therefore, by excluding P_i, we guarantee that the walk does not backtrack, which is important to get proper mixing properties as discussed in Section 3.1.
While the CGL hash function works for any ℓ, we generally choose ℓ to be small to prevent having to work over large extensions of $\mathbb{F}_{p^2}$. For example, in the case $\ell = 2$, all points of order 2 already have coordinates in $\mathbb{F}_{p^2}$. For larger ℓ, the points of order ℓ can exist in larger extensions of $\mathbb{F}_{p^2}$. The larger ℓ becomes, the larger this extension becomes and the more inefficient the computations get.

Algorithm 3 The CGL hash function

1: Fix a prime p and a starting curve $E_0/\mathbb{F}_{p^2}$. Fix a prime $\ell \neq p$ and a point $P_0 \in E_0[\ell]$.
2: Define an ordering of the generators of the $\ell + 1$ subgroups of order ℓ that works on any supersingular curve. It does not matter how this order is defined, as long as it works consistently on all curves in the graph.
3: Define the message space $\mathcal{M} = \{0, 1, \ldots, \ell - 1\}^*$.
4: Let $m = (m_1, \ldots, m_k) \in \mathcal{M}$.
5: For $i = 1, 2, \ldots, k$, do the following:
Discard the point P_{i-1} and order the ℓ remaining generators of the order ℓ subgroups, labelling them 0 through $\ell - 1$.
Set S_{i-1} to be the point labelled m_i (recall that $m_i \in \{0, \ldots, \ell - 1\}$).
Compute the isogeny φ_i with kernel $\langle S_{i-1} \rangle$, define the codomain of φ_i as E_i.
Set $P_i = \varphi_i(P_{i-1})$
6: After this is done for all i, we have an isogeny $\varphi = \varphi_k \circ \ldots \circ \varphi_1 : E_0 \to E_k$. Return $j(E_k)$.

The above discussion talks about the CGL hash function as if it specifically uses isogenies, but it is worth noting that the original proposal of the CGL hash function [11] is to construct a hash function based on walks on expander graphs. The motivation for this is that expander graphs in general have good mixing properties. They then propose two instantiations of this construction for specific families of expander graphs, one of them being the family of supersingular ℓ-isogeny graphs. The other proposal is to use the family of Lubotzky-Phillips-Sarnak (LPS) expander graphs. This other proposal was broken and is not considered to be safe hash function [63].
There are other methods to construct expander graphs. In theory, any such construction should provide a different instantiation of the CGL hash function, which will lead to a similar commitment scheme to Sterner's scheme that we introduce in Section 6.1.2. All we require to get a hash function from an expander graph is that we are able to label the edges from a given vertex in a consistent way, so that a given input will always lead to the same output.
A commitment scheme constructed from such a hash function is secure as long as the construction of the expander graph allows for a collision-resistant instantiation of the CGL hash function. If this is the case, we are able to prove the binding property via this collision resistance. Collision resistance of such a hash function based on walks in an expander graph comes from the difficulty of finding a cycle in the expander graph. The difficulty of this problem must be proven separately for each instantiation of the hash function. For the isogeny-based version, this is explained in Section 6.1.5.
The hiding property will always follow from the expander properties of the graph, similar to Section 6.1.4. The mixing constant of the expander graph might vary depending on the construction, so the minimum required length of a walk could differ.

6.1.2 Sterner's commitment scheme using the CGL hash function

Now that we have explained how the CGL hash function works, Sterner's commitment scheme is constructed as an instantiation from the generic commitment scheme construction using a hash function. More specifically, the commitment scheme is given as follows:

Algorithm 4 Sterner's commitment scheme using the CGL hash function

1: **KeyGen**(1^λ): Given a security parameter λ, fix a prime p and a positive integer k_r (see Section 6.1.6 for specific parameter choices). Fix a small prime ℓ, a supersingular elliptic curve $E/\mathbb{F}_{p^2}$ and $P_0 \in E[\ell]$. The public parameters are $PP := \{p, \ell, k_r, E, P_0\}$ and the message space is $\mathcal{M} = \{0, \dots, \ell - 1\}^*$.
2: **Commit**(PP, m, r): Given PP, a message $m \in \mathcal{M}$, and a random $r \in \{0, \dots, \ell - 1\}^{k_r}$, return $c = CGL(E, \ell, P_0, m||r)$, with $m||r$ the concatenation of m and r.
3: **Open**(PP, m, r, c): Given PP, a message $m \in \mathcal{M}$, $r \in \{0, \dots, \ell - 1\}^{k_r}$ and a commitment c, return $c == CGL(E, \ell, P_0, m||r)$.

The binding property of this commitment scheme comes from the difficulty of finding a collision in the supersingular ℓ-isogeny graph. This is because it is hard to compute a path between two curves in the graph. The hiding property comes from the mixing properties of the graph. If our walk is long enough, it becomes hard to say anything about the starting point of the walk.

To get a starting curve, a trusted third party is needed. Known methods for randomly generating a supersingular curve all provide knowledge about the endomorphism ring of the curve that is generated. This is a problem, as a dishonest user could break the binding property using this knowledge. This requires us to use a trusted third party, who generates a starting curve for us and then discards the information about the endomorphism ring of the starting curve. There are currently no known methods to randomly sample a supersingular curve from the full supersingular isogeny graph in an efficient way that do not reveal any information about the endomorphism ring. So far, attempted methods have not yielded any results, as explained in the papers [7] and [39].
Sterner suggests that the third party starts at the curve $E(1728)$ (that is, the curve with j-invariant 1728 given by the equation $y^2 = x^3 - x$) and goes on a walk from here, for example, by using the CGL hash function. This enforces the use of a prime $p \equiv 3 \bmod 4$, as this is equivalent to $E(1728)$ being supersingular. The curve at the end of the walk will be the starting curve E of the commitment scheme. The trusted third party makes E public and discards the information about the path taken.
With knowledge of the path taken and knowledge of the endomorphism ring at the start of the path, one can compute the endomorphism ring of the resulting curve. However, without knowledge of the path, there are no known ways to efficiently compute the endomorphism ring of the resulting curve. Thus, as long as the path taken is kept secret, it should be computationally infeasible to find the endomorphism ring of the starting curve, which prevents the binding property from being broken.

With the starting curve in place, we commit to a message $m \in \{0, \dots, \ell - 1\}^*$ (that is, m can be of any length) by first taking a uniformly random string $r \in \{0, \dots, \ell - 1\}^{k_r}$. The value of k_r must be sufficiently large to make the commitment scheme hiding. This will be discussed in more detail in the proof of the hiding property, see Theorem 6.3. We then concatenate m and r into one and we hash the concatenation $m||r$ into a commitment c, beginning from the starting curve E. By concatenating m and r, we ensure that no backtracking takes place. To open the commitment c, a verifier needs m, r and E to again compute the hash function with starting curve E.

Remark 6.1. *As mentioned above, knowledge of the endomorphism ring of the starting curve can break the binding property. This is achieved as follows. Recall that to win the binding game, we need to find two different messages $m \neq m'$ along with randomness r, r' such that the commitment of m and r is the same as the commitment of m' and r'. With knowledge of the endomorphism ring, we can compute an endomorphism of degree ℓ^e, where e is some positive integer. We can split this endomorphism into e isogenies of degree ℓ. Now we can compose the first $e - n$ isogenies, as well as the last n isogenies. The first composition and the dual of the second composition now provide two different paths from our starting curve to the same resulting curve, that is, the same commitment, breaking the binding property.*

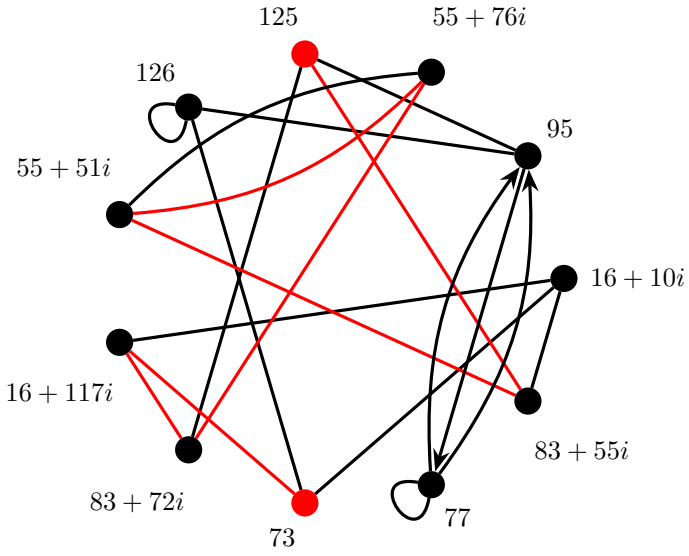

Figure 3: The CGL hash function from j-invariant 95 with input 101001.

Example 6.2. *To demonstrate how the commitment scheme works, we will work out a small example, revisiting the isogeny graph $G_2(127)$ from Figure 2. Due to its scale, this example is not safe for practical applications. However, it is useful to get a better idea of what is happening.*

Our trusted third party has computed the CGL hash of some secret input starting from the curve $E(1728)$, providing us with the starting j-invariant of 125. The input of their hash is then discarded. Suppose that we would like to commit to the message $m := 101$. We also need to add some randomness, which we randomly draw to be $r := 001$. Concatenating m and r, we obtain the string 101001, which we will put into the CGL hash function. We will now proceed through the isogeny graph without backtracking. At each step, we label the two edges that we did not arrive from with 0 and 1 according to our predefined method of ordering the generators.

To start, we need a representative of the isomorphism class of j-invariant 125. One such curve is the one given by the equation $y^2 = x^3 + x + 21$. The 2-torsion points of this curve are the points $(30, 0), (112 + 26i, 0), (112 + 101i, 0)$. We arbitrarily pick the point $(112 + 101i, 0)$ as our point P_0, that is, the point that will ensure that we do not backtrack. We now label them: the point $(30, 0)$ is labelled 0, and the point $(112 + 26i, 0)$ is labelled 1. Since the first bit of our message is a 1, we construct the isogeny with kernel $\mathcal{O}, (112 + 26i, 0)$. Using a computer algebra package such as MAGMA, we can compute the isogeny, as well as the destination curve (though this can be done by hand fairly easily in the case of 2-isogenies). This tells us that we arrive at the curve $y^2 = x^3 + (30 + 16i)x + (125 + 110i)$, which has j-invariant $83 + 55i$. The isogeny sends the point P_0 to the point $P_1 := (30 + 75i, 0)$.

We repeat this procedure for the remaining bits. We consider the 2-torsion points of the curve $y^2 = x^3 + (30 + 16i)x + (125 + 110i)$ and we take out the point P_1. Our next bit is a 0, so the 2-torsion point labelled 0 will be the kernel of our next isogeny. In our case, this is the point $(36 + 121i, 0)$. This gets us to a curve with j-invariant $55 + 51i$. After this, our walk takes us, in order, to the j-invariants $55 + 76i, 83 + 72i, 16 + 117i$ and finally 73. Thus, we end up with our final commitment $c = 73$. The path taken is highlighted in figure 3.

We can also see what happens if we do not use a trusted third party. As will be shown in Proposition 6.7, the curve with j-invariant 95 has endomorphism ring $\mathbb{Z} \oplus \mathbb{Z}(2i) \oplus \mathbb{Z}(\frac{1+j}{2}) \oplus \mathbb{Z}(\frac{i+k}{4})$, where $i^2 = -1$, $j^2 = -p$

and $k = ij = -ji$. We can compute that an element $\alpha = a + 2bi + c\frac{1+j}{2} + d\frac{i+k}{4}$ of this ring has degree $\frac{1}{16}\left((4a+2c)^2 + (8b+d)^2 + (4c^2+d^2)p\right)$. Taking $a = b = c = 0$ and $d = 1$ gives us an endomorphism of degree 8. We can see in the graph that indeed, we can find a non-backtracking path of length 3: starting at 95, we travel to 126, then we take the endomorphism on 126, after which we travel back to 95. A slightly more interesting path is the endomorphism of degree 64 by taking $a = b = 0$, $c = 1$, $d = 2$. We can see a loop of length 6 in the graph that starts at 95, then goes to 126, 73, $16 + 117i$, $83 + 72i$, 125 and finally back to 95. We can split this up into a path from 95 to $16 + 117i$ and a path from $16 + 117i$ to 95. The dual of the latter along with the first path gives us two distinct paths of length 3 to the node $16 + 117i$.

6.1.3 Sterner's commitment scheme using the SIDH approach

The CGL hash function approach is one way to perform a walk on an isogeny graph. However, there are other ways to do this. The approach given in this section is inspired by the way that isogenies in SIDH are computed. It should be noted that while Sterner's work is from before the fall of SIDH, this approach is not compromised by the attacks on SIDH. The SIDH attacks make use of torsion information that is made public, but here, no torsion information is revealed. This method also focuses on the case of degree 2 isogenies, although in theory other prime degrees could be used.
Besides the way the walk is performed, nothing else is different about this commitment scheme. The proofs for hiding and binding (See Sections 6.1.4 and 6.1.5) work in the same way. However, this approach is more efficient to compute, see Section 6.1.6.

The SIDH approach works with primes of the form $p = 2^n f - 1$, where f is a small integer. Note that by making use of a prime $\ell^n f - 1$, this method could theoretically work for any degree isogeny. However, it is most efficient to work with $\ell = 2$, to prevent having to work over large extensions of $\mathbb{F}_{p^2}$. Therefore, we will focus on this case here. Our starting curve is again given with unknown endomorphism ring. This curve is generated in the same way as Section 6.1.2. This time, we make sure that $\#E(\mathbb{F}_{p^2}) = (2^n f)^2 = (p+1)^2$ and therefore $E(\mathbb{F}_{p^2}) \cong (\mathbb{Z}/(p+1)\mathbb{Z})^2$. This ensures that the 2^l-torsion for $l \leq n$ has only points that have coordinates in $\mathbb{F}_{p^2}$. We take $P, Q \in E[2^n]$ to be points on E that form a basis of the 2^n-torsion.
As with the previous method, we will go on a walk in the ℓ-isogeny graph. The key difference is that instead of choosing at each vertex which edge to traverse next, we have a kernel subgroup of order 2^n that decides the route taken. This subgroup produces an isogeny of degree 2^n. In the ℓ-isogeny graph, this translates to a walk of length n. The length of this walk is not enough to achieve proper security, as will be discussed in Section 6.1.6. However, we can repeat this computation four times to achieve the same level of security as the CGL method. The protocol is given in Algorithm 5.

Having introduced Algorithm 5, it is important to note that while we specify the size of m in the algorithm, m can be of any size. However, a larger m might require more isogenies of degree 2^n to be computed.
Some work is required to ensure that we can compute multiple degree 2^n isogenies in a row. In the algorithm, we need to generate points on the intermediate curves E_{m_i} and E_{r_i}. These points need to be found in a deterministic manner to ensure that anyone can repeat the procedure and obtain the same outcome. We will go through this procedure for the curve E_{m_0} that is reached after the first step.

We need to find points P', Q' that form the basis of $E_{m_0}[2^n]$. We can easily find one point: we know that $\phi_{m_0}(Q)$ will have order 2^n, so we will set $Q' := \phi_{m_0}(Q)$. This choice is deliberate because it prevents backtracking from occurring. Backtracking is prevented because $\widehat{\phi}$ has its kernel generated by Q'. To see this, notice that $\widehat{\phi}(\phi(Q)) = [2^n]Q = 0$, so $\phi(Q) \in \ker(\widehat{\phi})$, and since $\phi(Q)$ is of order 2^n it must generate the entire kernel of $\widehat{\phi}$. Recall that $\widehat{\phi}$ is a cyclic isogeny from Remark 3.15. The first step of the dual is given by the 2-isogeny with kernel $\langle [2^{n-1}]Q' \rangle$. An isogeny with kernel $\langle P' + m_1 Q' \rangle$ for some m_1 will have its first step given by the 2-isogeny with kernel $\langle [2^{n-1}]P' \rangle$ or the 2-isogeny with kernel $\langle [2^{n-1}](P' + Q') \rangle$. In either case,

Algorithm 5 Sterner's commitment scheme using the SIDH approach

1: **KeyGen**(1^λ): Given a security parameter λ, fix a prime $p = 2^n f - 1$, with f a small integer. Fix a supersingular elliptic curve $E/\mathbb{F}_{p^2}$ such that $\#E(\mathbb{F}_{p^2}) = (2^n f)^2$. Fix two points $P, Q \in E[2^n]$ such that $E[2^n] = \langle P, Q\rangle$. The public parameters are $PP := \{p, E, P, Q\}$ and the message space is $\mathcal{M} = [0, 2^{4n} - 1]$.

2: **Commit**(PP, m, r): Given PP, a message $m \in \mathcal{M}$, and a random $r \in [0, 2^{4n} - 1]$, compute $m_0 := m \bmod 2^n$, $m_1 := \frac{m-m_0}{2^n} \bmod 2^n$, $m_2 := \frac{m-m_0-m_1 2^n}{2^{2n}} \bmod 2^n$, $m_3 := \frac{m-m_0-m_1 2^n - m_2 2^{2n}}{2^{3n}} \bmod 2^n$.
Compute $M_0 := \langle P + m_0 Q\rangle$ and compute $E_{m_0} := E/\langle M_0\rangle$ given by isogeny ϕ_{m_0}. Compute P' and $Q' := \phi_{m_0}(Q)$.
Compute $M_1 := \langle P' + m_1 Q'\rangle$ and compute $E_{m_1} := E_{m_0}/\langle M_1\rangle$ given by isogeny ϕ_{m_1}. Compute P'' and $Q'' := \phi_{m_1}(Q')$.
Compute $M_2 := \langle P'' + m_2 Q''\rangle$ and compute $E_{m_2} := E_{m_1}/\langle M_2\rangle$ given by isogeny ϕ_{m_2}. Compute P''' and $Q''' := \phi_{m_2}(Q'')$.
Compute $M_3 := \langle P''' + m_3 Q'''\rangle$ and compute $E_{m_3} := E_{m_2}/\langle M_3\rangle$ given by isogeny ϕ_{m_3}. Compute P_m and $Q_m := \phi_{m_3}(Q''')$. Define $E_m := E_{m_3}$.
Compute $r_0 := r \bmod 2^n$, $r_1 := \frac{r-r_0}{2^n} \bmod 2^n$, $r_2 := \frac{r-r_0-r_1 2^n}{2^{2n}} \bmod 2^n$, $r_3 := \frac{r-r_0-r_1 2^n - r_2 2^{2n}}{2^{3n}} \bmod 2^n$.
Compute $R_0 := \langle P_m + r_0 Q_m\rangle$ and compute $E_{r_0} := E_m/\langle R_0\rangle$ given by isogeny ϕ_{r_0}. Compute P'_m and $Q'_m := \phi_{r_0}(Q_m)$.
Compute $R_1 := \langle P'_m + r_1 Q'_m\rangle$ and compute $E_{r_1} := E_{r_0}/\langle R_1\rangle$ given by isogeny ϕ_{r_1}. Compute P''_m and $Q''_m := \phi_{r_1}(Q'_m)$.
Compute $R_2 := \langle P''_m + r_2 Q''_m\rangle$ and compute $E_{r_2} := E_{r_1}/\langle R_2\rangle$ given by isogeny ϕ_{r_2}. Compute P'''_m and $Q'''_m := \phi_{r_2}(Q''_m)$.
Compute $R_3 := \langle P'''_m + r_3 Q'''_m\rangle$ and compute $E_{r_3} := E_{r_2}/\langle R_3\rangle$ given by isogeny ϕ_{r_3}. Return $c = j(E_{r_3})$.

3: **Open**(PP, m, r, c): Given PP, a message $m \in \mathcal{M}$, $r \in [0, 2^{4n} - 1]$ and a commitment c, compute $j(E_{r_3})$ in the same manner as in **Commit**. Return $c == j(E_{r_3})$.

the first step will never be the first step of the dual and so it will never backtrack.

To produce a point P', Sterner makes use of the Elligator 2 method [6] with some predetermined parameters, as explained in [12]. Recall from Lemma 3.13 that a curve E defined over $\mathbb{F}_{p^2}$ is always isomorphic to a curve in Montgomery form given by $y^2 = f(x) = x^3 + Ax^2 + x$ for some $A \in \mathbb{F}_{p^2}$. We will assume our curves to be of this form. For a point (x, y) on this curve to be in $E(\mathbb{F}_{p^2})$, we need that $f(x)$ is a square in $\mathbb{F}_{p^2}$. In Elligator 2, we let u be any non-square element in $\mathbb{F}_{p^2}$. For any $r \in \mathbb{F}_{p^2}$, write

$$v = -\frac{A}{1 + ur^2} \text{ and } v' = \frac{A}{1 + ur^2} - A.$$

Notice that $v' = -v - A$ and $v' = vur^2$. We thus obtain that $f(v') = v'^3 + Av'^2 + v' = v'(v'(v' + a) + 1) = v'((-v - A) - v + 1) = ur^2 v(v^2 + Av + 1) = ur^2 f(v)$. So, $f(v)$ and $f(v')$ differ by a non-square element. Since the Legendre symbol is multiplicative, this ensures that exactly one of $f(v)$ and $f(v')$ is a square. Therefore, exactly one of v and v' is an x-coordinate of a point in $E(\mathbb{F}_{p^2})$, after which we can also compute a y-coordinate.
In [12], as well as in Sterner's case, we have $p \equiv 3 \bmod 4$, so we write $\mathbb{F}_{p^2} = \mathbb{F}_p[i]$ with $i^2 = -1$. Fix a non-square element u of $\mathbb{F}_{p^2}$. The authors of [12] use a predetermined value of p, so in their case they pick $u = i + 4$. A table is precomputed consisting of the values $\frac{-1}{1+ur^2} \in \mathbb{F}_{p^2}$ where r^2 ranges from 1 to 10. This table is public and anyone can use it to generate points on a desired Montgomery curve by combining it with the corresponding value of A. The final step is to generate points, starting with the first value of r^2 in the table, until we find a point R such that $R \in E \setminus E[2]$. If this is the case, then fR is a point of order 2^n. We need to check if this point is independent from Q'. If this is true, set $P' := fR$. Otherwise, continue generating points until a suitable candidate is found. The probability of failure after trying all points in the

table is very low.

6.1.4 Hiding property

Sterner's isogeny-based commitment scheme has information-theoretic hiding (recall Definition 4.7). The proofs for both the hiding and the binding property function almost the same for both the CGL hash version and the SIDH version of Sterner's scheme. The proofs, originally given by Sterner, work for both versions of the scheme.

Theorem 6.3. *Let $k_{\ell,p}$ be the mixing constant for the supersingular ℓ-isogeny graph in characteristic p. Fix λ a security parameter. Then for any $k \geq k_{\ell,p}$, the commitment schemes described in Sections 6.1.2 and 6.1.3 are information-theoretically hiding.*

Proof. Fix two messages m_0, m_1, choose a random bit $b \in \{0,1\}$ and commit to the message m_b, resulting in a commitment c, represented by a curve E'. An adversary must determine which message was used to get the commitment. The supersingular ℓ-isogeny graph is $\ell+1$-regular, so the mixing constant $k_{\ell,p}$ is well-defined. By definition of the mixing constant (see Definition 3.6), for any $k \geq k_{\ell,p}$, there exists a path of length k from E_{m_0} to E' and from E_{m_1} to E'. Using Theorem 3.19, we have for $i \in \{0,1\}$ that

$$\left| Pr[c = j(E') \mid \text{message is } m_i] - \frac{1}{N_p} \right| \leq \left(\frac{2\sqrt{\ell}}{\ell+1} \right)^k.$$

Subtracting the m_1 case from the m_0 case, we obtain

$$|Pr[c = j(E') \mid \text{message is } m_0] - Pr[c = j(E') \mid \text{message is } m_1]| \leq 2 \left(\frac{2\sqrt{\ell}}{\ell+1} \right)^k.$$

Notice that the left hand side of this inequality is the best case scenario for an adversary trying to win the hiding game. Even if an adversary knew the exact probabilities, they could never pick between m_0 and m_1 with more certainty than the above inequality. Thus, the hiding advantage as formulated in Lemma 4.8 is bounded above by this.
Since $\left(\frac{2\sqrt{\ell}}{\ell+1}\right) < 1$, we have that

$$2 \left(\frac{2\sqrt{\ell}}{\ell+1} \right)^k \leq 2 \left(\frac{2\sqrt{\ell}}{\ell+1} \right)^{k_{\ell,p}} \leq 2 \left(\frac{2\sqrt{\ell}}{\ell+1} \right)^{\log_\ell(p-12) - \log_\ell(12(\ell+1))+1}.$$

The last inequality follows from Lemma 3.7 and Proposition 3.14, using that $N_p \geq \frac{p}{12} - 1$. As we will explain in Section 6.1.6, p is a prime of size approximately 2λ bits, where λ is the security parameter. Therefore, $\log_\ell(p-12)$ will be at most 2λ, decreasing as ℓ increases. We will write $\log_\ell(p-12) = n\lambda$ for some n.
We can thus see that as a function of λ, $2\left(\frac{2\sqrt{\ell}}{\ell+1}\right)^{n\lambda - \log_\ell(12(\ell+1))+1}$ is negligible. Therefore, for any adversary, the hiding advantage is bounded by a negligible function, which proves information-theoretic hiding. □

6.1.5 Binding property

Sterner's scheme is computationally binding assuming the hardness of the following problem:

Problem 6.4 (Supersingular Smooth Endomorphism Problem)**.** *Given a prime p, a supersingular elliptic curve E over $\mathbb{F}_{p^2}$ and a small prime ℓ, compute a non-trivial cyclic endomorphism on E of degree a power of ℓ.*

This problem is similar to Problem 1 presented in [24]. In [24] it is also remarked that this problem is equivalent to Problem 5.5.
The following proof is adapted from Sterner's proof [62].

Theorem 6.5. *The commitment schemes described in Sections 6.1.2 and 6.1.3 are computationally binding under the assumption that the Supersingular Smooth Endomorphism Problem is hard.*

Proof. Suppose we have a PPT adversary $\mathcal{A}$ that solves the binding game for this commitment scheme. We will now construct a PPT adversary $\mathcal{A}'$ that uses $\mathcal{A}$ to solve the Supersingular Smooth Endomorphism Problem on E.
Let $c(m, r)$ denote the commitment to a message m and randomness r, either via the CGL hash approach or the SIDH approach. Given E, our adversary $\mathcal{A}'$ now queries $\mathcal{A}$ to obtain values m, m', r, r' such that $m \neq m'$, $r \neq r'$ and E' such that $j(E') = c(m, r) = c(m', r')$. From computing the commitments, we get two isogenies $\phi_{m||r} : E \longrightarrow E'$ and $\phi_{m'||r'} : E \longrightarrow E'$, both of degree ℓ^{2k}. The map $\widehat{\phi_{m'||r'}} \circ \phi_{m||r} : E \longrightarrow E$ is an endomorphism of E with degree ℓ^{4k}. Suppose this endomorphism was trivial, that is, $\widehat{\phi_{m'||r'}} \circ \phi_{m||r} = [\ell^{4k}]$. Since our maps are of the same degree and the dual isogeny is unique, this tells us that $\widehat{\phi_{m'||r'}} = \widehat{\phi_{m||r}}$. This also means that $\phi_{m'||r'} = \phi_{m||r}$, which would imply $m = m'$ and $r = r'$, which is a contradiction.
We can write our map as a composition of isogenies of degree ℓ. By taking out any backtracking that occurs here from the composition, we end up with a cyclic endomorphism ψ that solves Problem 6.4.
Therefore, the advantage of winning the binding game is no larger than the advantage of solving the Supersingular Smooth Endomorphism Problem. Thus, assuming this is a hard problem to solve, our commitment scheme is computationally binding. □

6.1.6 Parameter choices and performance

In this section, we go over the parameter choices required to make this commitment scheme binding and hiding according to the desired level of security. We also analyse the performance of this scheme and we compare it to a lattice-based commitment scheme.
Sterner's scheme is information-theoretically hiding, while it is computationally binding. For this reason, attacks on the binding property are more relevant, as it relies only on the hardness of another problem that we assume to be difficult.

We first consider the binding property. The best known attack on Problem 6.4 is one by Delfs and Galbraith [19], which requires $\tilde{O}(p^{1/2})$ operations, where p is the characteristic used. It is not immediately obvious that this is an attack on Problem 6.4, as it claims to be an attack that finds a path between two curves in an ℓ-isogeny graph. However, Problem 6.4 is equivalent to the Endomorphism Ring Problem (given a supersingular curve, compute a $\mathbb{Z}$-basis for its endomorphism ring) as explained in Remark 8 of [24]. This problem, in turn, was proven to be equivalent to the ℓ-isogeny path problem (given two supersingular curves, find a path between them in the ℓ-isogeny graph) in [66]. Finally, the ℓ-isogeny path problem is the problem that the attack from Delfs and Galbraith solves.
To achieve λ bits of security, we need to pick a p that requires 2^λ operations. The attack from Delfs and Galbraith means that we need p such that $p^{1/2} = 2^\lambda$. So, we need $p \approx 2^{2\lambda}$. In other words, p should be of size approximately 2λ bits. The commitment in both schemes is a j-invariant, that is, it is an element of $\mathbb{F}_{p^2}$. This value can be written as $a + bi \in \mathbb{F}_p[i]$, so we can store a commitment as two elements of $\mathbb{F}_p$, which means we need 4λ bits, or $\frac{\lambda}{2}$ bytes to store our commitment. The size only depends on the value of p, and it is independent of k, which is the length of the walk that we take. So, no matter the length of the message, the commitment will be of the same size. For example, to obtain 128 bits of security, the commitment will only be 512 bits or 64 bytes, which is very small.
Next, to achieve the hiding property, the only thing we need is to pick a value for k_r that is larger than the mix-

ing constant of the supersingular isogeny graph. As there are no proven upper bounds for the mixing constant, we rely on conjectural bounds. Under Conjecture 3.8, we can pick $k_r = 4\lceil \log_\ell((\ell+1)N_p\rceil + 4$. Furthermore, if Sterner's Conjecture 3.18 holds, we can pick $k_r = \lceil \log_\ell(p) + \log_\ell(\log_\ell(p)) + 1\rceil$, which would allow ever better performance. For the $\ell = 2$ case, these bounds are $k_r = 4\lceil \log_2(p)\rceil - 4$ and $k_r = \lceil \log_2(p) + \log_2(\log_2(p)) + 1\rceil$. In the SIDH variant, we restrict ourselves to the case $\ell = 2$. The highest degree isogeny that we can compute is 2^n. As we use a prime of the form $p = 2^n f - 1$ with f a small integer, this means that to attain hiding, under Conjecture 3.8 we would need to compute enough isogenies of degree 2^n to get a walk longer than the bound. By computing the bound for our p, we get a bound of approximately $4n$, which tells us that we need to compute four isogenies of degree 2^n. Under Conjecture 3.17, we can again do the computation for our p. This tells us that we would only need to compute two isogenies of degree 2^n to have a long enough path to achieve hiding.

For computational performance, we will look at the setting $\ell = 2$. If we use a prime of the form $p = 2^n f - 1$, and k is the length of the message, the work of [21] estimates the CGL hash function to require approximately $k_r n(5.7n + 110)\mathbf{m}$, where $\mathbf{m}$ is the cost of performing a single field multiplication. In [21], there is also a hash function that works similarly to the SIDH variant of Sterner's scheme, which is estimated to require approximately $k_r n(13.5\log(n) + 42.4)\mathbf{m}$. This performance estimation translates to Sterner's scheme by picking $k_r = 4\lceil \log_2(p)\rceil - 4$, with the added caveat that we need to also spend some time generating new basis elements for the 2^n-torsion on the intermediate curves in the SIDH method. This is done at most three times, and each time it involves computing an isogeny image and performing the Elligator 2 method. It approximately adds $O(n\mathbf{m})$ to the complexity, which is not significant compared to the cost of the rest of the computation. We conclude that the SIDH method is much faster than the CGL hash function method.

There exist commitment schemes with other post-quantum alternatives. For example, there are lattice-based commitment schemes. Sterner's isogeny-based method provides one clear advantage, which is the size of the commitment. For 128 bits of security, Sterner's scheme produces commitments of only 64 bytes. This is much better than, for example, lattice-based commitment schemes. In [4], a commitment with 128 bits of security is 9 kilobytes. However, the downside of the isogeny-based method is in its computational efficiency. While the works on lattice-based commitment schemes use different methods to measure computational efficiency than Sterner, which makes direct comparisons difficult, it does hold true that lattice-based methods are much faster to compute. Even the SIDH variant is not as fast as the lattice-based method.

6.2 An isogeny-based commitment scheme without a trusted third party

In this section, we go over a proposed modification to Sterner's commitment scheme that removes the need for a trusted third party. This is achieved by making use of Theorem 6.6. We will introduce and prove this theorem in Section 6.2.1. After this, we describe the modified commitment scheme, prove binding and hiding, and go over the parameter choices. Finally, we briefly go over an alternate modification to remove the need for a trusted third party.

6.2.1 Avoiding a trusted third party

An issue that is not explicitly pointed out in Sterner's original proposal is the requirement of a trusted third party in the setup phase of the scheme. During the setup, we use a trusted third party to generate a starting curve with unknown endomorphism ring. This is done in other isogeny-based protocols too. As explained in Remark 6.1, it is crucial to not know the endomorphism ring, as with knowledge of it one can obtain an endomorphism of degree a power of a small prime. Splitting this endomorphism into two, we obtain two messages with the same commitment, breaking the binding property. In 2024, Saah, Fouotsa, Fouotsa and Nkuimi-Jugnia [53] proposed a way to get around the requirement of a trusted third party. If one chooses the endomorphism ring carefully, it is possible to prove that endomorphisms of certain degree do not exist at all.

If we use the curve that has this endomorphism ring as the starting curve, the above method to break the binding property does not work, so knowledge of the endomorphism ring of the starting curve is irrelevant. The idea is to use the curve $E_6 : y^2 = x^3 + 6x^2 + x$ as the starting curve, defined over $\mathbb{F}_{p^2}$ where $p \equiv 15 \bmod 16$. This curve is the unique 2-isogenous neighbour of j-invariant 1728 in the supersingular isogeny graph. This will be proven in Proposition 6.13. In a security analysis of SIDH/SIKE, it was shown by Onuki, Aikawa and Takagi that in this setting, cyclic endomorphisms of certain degrees do not exist at all [47], as described in the following theorem:

Theorem 6.6. *Let ℓ be a prime number that does not split in $\mathbb{Z}[\sqrt{-1}]$ (that is, $\ell \not\equiv 1 \bmod 4$), and let $\varphi = (\varphi_1, \ldots, \varphi_n)$ and $\psi = (\psi_1, \ldots, \psi_m)$ be two distinct paths of respective lengths n and m from E_6 to the same curve E in $G_\ell(p)$ without backtracking. Then one of the following holds:*

- $\ell^{n+m} \geq \frac{p+1}{16}$;
- $\ell = 2$ *and either φ or ψ has a form $\varphi' \circ \varphi_0$ where $\varphi_0 : E_6 \longrightarrow E(1728)$ is of degree 2 and $E(1728)$ has j-invariant 1728.*

This theorem puts a lower bound on the degree of endomorphisms on E_6, depending on the characteristic p that we are working over. So, if we perform two walks that are short enough, they are guaranteed to not finish at the same ending curve. If they did, they would form an endomorphism of a degree lower than the bound. So, we need to increase our p to allow our walks to be of useful length. This makes our scheme perfectly binding and makes it so that we do not need a trusted third party, since knowledge of the endomorphism ring is no longer a concern.

Theorem 6.6 allows for a modified version of Sterner's commitment scheme, which will be constructed in Section 6.2.2. Before we do so, we will prove the theorem. To be able to prove this, we need two lemmas. The first lemma, Lemma 6.8, gives the structure of the endomorphism ring of E_6, given that $p \equiv 15 \bmod 16$. The second lemma, Lemma 6.9, tells us that an endomorphism in the supersingular ℓ-isogeny graph can always be shortened so that it does not contain any multiple of ℓ.
To prove Lemma 6.8, we need to know the structure of the endomorphism ring of the curve $E(1728)$, which is the curve $y^2 = x^3 - x$ with j-invariant 1728. This structure is well-known, but a proof is usually omitted. For completeness, we include a proof here.

Proposition 6.7. *Let $E(1728)$ be the curve $y^2 = x^3 - x$ with j-invariant 1728 and let p be a prime such that $p \equiv 3 \bmod 4$. We have that*

$$\mathrm{End}(E(1728)) \cong \mathbb{Z} + \mathbb{Z}i + \mathbb{Z}\frac{1+j}{2} + \mathbb{Z}\frac{i+k}{2}$$

where the right hand side is a maximal order in the quaternion algebra $B_{p,\infty} \cong \left(\frac{-1,-p}{\mathbb{Q}}\right)$ as in Proposition 2.16.

Proof. We have chosen the curve $y^2 = x^3 - x$ because it has trace 0 as a curve over $\mathbb{F}_p$. As a result, the $\mathbb{F}_p$-Frobenius endomorphism $F^{(p)}$ satisfies $(F^{(p)})^2 = [-p]$. Besides this, on supersingular curves with j-invariant 1728 there exists the automorphism $\iota : (x, y) \to (-x, iy)$ of order 4 with $\iota^2 = [-1]$. So, we have $F^{(p)}$ and ι in $\mathrm{End}(E(1728))$ and they are independent from each other. Hence, we can define an isomorphism from the endomorphism algebra of $E(1728)$ to the quaternion algebra $B_{p,\infty} \cong \left(\frac{-1,-p}{\mathbb{Q}}\right)$ that sends ι to i and $F^{(p)}$ to j.
Now observe that $y^2 = x^3 - x$ has 2-torsion points $(0, 0), (-1, 0), (1, 0)$, so its full 2-torsion group is defined over $\mathbb{F}_p$. This means that $F^{(p)}$ fixes the 2-torsion, and trivially the identity map 1 does as well. So, the map $F^{(p)} + 1$ contains the 2-torsion in its kernel, meaning that we can divide this map by 2. Therefore, the

map $\frac{F^{(p)}+1}{2}$ lies in $\mathrm{End}(E(1728))$. Going back to our map to $B_{p,\infty}$, we can see that $\mathrm{End}(E(1728))$ (which is contained in the endomorphism algebra) is sent to an order containing at least i and $\frac{1+j}{2}$. The order generated by i and $\frac{1+j}{2}$ is

$$\mathbb{Z} + \mathbb{Z}i + \mathbb{Z}\frac{1+j}{2} + \mathbb{Z}\frac{i+k}{2}.$$

We can calculate (see after Definition 2.17) that this order is maximal, so by the Deuring correspondence (Theorem 2.19) we know that this must be the full endomorphism ring of $E(1728)$. □

From this endomorphism ring, it is not too hard to construct a different maximal order whose corresponding j-invariant will be 2-isogenous to the curve $E(1728)$. In [32], Lemma 4.2, the authors prove that, given two maximal orders, the index of either maximal order over the intersection of the two is equal to the smallest degree isogeny between the two corresponding j-invariants. That is, if we want to find a curve that is 2-isogenous to $E(1728)$, it suffices to provide a maximal order that has index 2 over the intersection of said order with $\mathrm{End}(E(1728))$. This is achieved by multiplying one of the basis elements by 2, while dividing another by 2. This gives us a candidate in the maximal order given by $\mathbb{Z} + \mathbb{Z}(2i) + \mathbb{Z}\frac{1+j}{2} + \mathbb{Z}\frac{i+k}{4}$. However, we need to ensure that all of the basis elements have integer norm, otherwise we do not have a maximal order. A quick computation shows that the element $\frac{i+k}{4}$ has norm $\frac{p+1}{16}$. This enforces the use of a prime $p \equiv 15 \bmod 16$.
In this setting for p, we have found a maximal order different from $\mathrm{End}(E(1728))$ which is 2-isogenous to j-invariant 1728. Since the unique 2-isogenous neighbour of j-invariant 1728 is E_6, this means that the maximal order described here must be $\mathrm{End}(E_6)$.

Lemma 6.8. *If $p \equiv 15 \bmod 16$, then we have*

$$\mathrm{End}(E_6) \cong \mathbb{Z} + \mathbb{Z}(2i) + \mathbb{Z}\frac{1+j}{2} + \mathbb{Z}\frac{i+k}{4}$$

Proof. The proof of this lemma is roughly the reasoning given above, but backwards. □

The second lemma that we need to prove Theorem 6.6 says that any non-integer endomorphism can be shortened to not contain any factor that is multiplication-by-ℓ.

Lemma 6.9. *Let E and E' be supersingular elliptic curves, and $(\varphi_1, \ldots, \varphi_n)$ and $(\psi_1, \ldots, \psi_m)$ distinct paths from E to E' without backtracking in $G_\ell(p)$. Then the composition*

$$\widehat{\psi}_1 \circ \cdots \circ \widehat{\psi}_m \circ \varphi_n \circ \cdots \circ \varphi_1$$

is a non-integer endomorphism on E. Furthermore, if $m \leq n$ there exist integers $0 \leq m' \leq m$ and $1 \leq n' \leq n$ such that

$$\widehat{\psi}_0 \circ \cdots \circ \widehat{\psi}_{m'} \circ \varphi_{n'} \circ \cdots \circ \varphi_1$$

is in $\mathrm{End}(E) \backslash \ell\,\mathrm{End}(E)$, where ψ_0 is the identity map on E.

The proof of this lemma is quite heavy in notation. An intuitive way to think about this proof is as follows: if our composition is in $\ell\,\mathrm{End}(E)$, the only place where we could find ℓ is at the end of the paths, because the paths are non-backtracking. This occurs if the paths are the same for some amount of steps at the end. We then cut off the ends of the paths that are the same to obtain an endomorphism that is not in $\ell\,\mathrm{End}(E)$.
The proof written here is more explicit than the original proof in [47].

Proof of Lemma 6.9. Without loss of generality, assume $m \leq n$. We define $\xi_i = \varphi_i$ for $i \in \{1, \ldots, n\}$ and $\xi_i = \widehat{\psi}_{m+n+1-i}$ for $i \in \{n+1, \ldots, n+m\}$. This allows us to write the composition $\widehat{\psi}_1 \circ \cdots \circ \widehat{\psi}_m \circ \varphi_n \circ \cdots \circ \varphi_1$

as $\xi := \xi_1 \circ \cdots \circ \xi_{n+m}$. Denote the domain and codomain of ξ_i as E_{i-1} and E_i, respectively.
Suppose ξ is in $\ell \operatorname{End}(E)$. This means that $E[\ell] \subset \ker \xi$. Then, there exists some j such that $\ker \xi_j \circ \cdots \circ \xi_1$ does not contain $E[\ell]$, but $\ker \xi_{j+1} \circ \cdots \circ \xi_1$ does. As we are working with cyclic kernels (see Remark 3.15), we let $P \in E$ be a generator of $\ker \xi_j \circ \cdots \circ \xi_1$. Let Q be a point on $E[\ell]$ such that $Q \notin \ker \xi_j \circ \cdots \circ \xi_1$. The point $\xi_{j-1} \circ \cdots \circ \xi_1(P)$ is in $E_{j-1}[\ell]$, and $\xi_{j-1} \circ \cdots \circ \xi_1(Q)$ is a point in $E_{j-1}[\ell]$ independent of it, since otherwise $Q \in \ker \xi_j \circ \cdots \circ \xi_1$. So, $\xi_{j-1} \circ \cdots \circ \xi_1(P)$ and $\xi_{j-1} \circ \cdots \circ \xi_1(Q)$ generate $E_{j-1}[\ell]$.
Since $E[\ell] \subset \ker \xi_{j+1} \circ \cdots \circ \xi_1$, we have $\xi_{j+1} \circ \cdots \circ \xi_1(P) = \xi_{j+1} \circ \cdots \circ \xi_1(Q) = 0_{E_{j+1}}$. This means that $E_{j-1}[\ell] \subset \ker \xi_{j+1} \circ \xi_j$. Furthermore, because $\#E_{j-1}[\ell] = \ell^2$ and $\xi_{j+1} \circ \xi_j$ is separable of degree ℓ^2, we have $E_{j-1}[\ell] = \ker \xi_{j+1} \circ \xi_j$. In other words $\xi_{j+1} \circ \xi_j = [\ell]$, which means that $\widehat{\xi_{j+1}}$ and ξ_j are equivalent.
The paths $(\varphi_1, \ldots, \varphi_n)$ and $(\psi_1, \ldots, \psi_m)$ are without backtracking, so we must have $j = n$. Therefore, ϕ_n and ψ_m are equivalent. We can take out ϕ_n and $\widehat{\psi_m}$ from our composition to obtain the endomorphism

$$\widehat{\psi}_1 \circ \cdots \circ \widehat{\psi}_{m-1} \circ \varphi_{n-1} \circ \cdots \circ \varphi_1.$$

If this endomorphism is in $\ell \operatorname{End}(E)$ again, we repeat the above process. Eventually, we obtain the endomorphism

$$\alpha := \widehat{\psi}_0 \circ \cdots \circ \widehat{\psi}_{m'} \circ \varphi_{n'} \circ \cdots \circ \varphi_1,$$

with m', n' integers such that $0 \leq m' \leq m$ and $1 \leq n' \leq n$. The situation in which m' could be zero is if the path $(\psi_1, \ldots, \psi_m)$ is contained in the path $(\varphi_1, \ldots, \varphi_n)$. For example, we could have φ_1 be an endomorphism of degree ℓ, after which $\varphi_2 = \psi_1, \ldots, \varphi_n = \psi_m$.
On the other hand, φ_1 can not be taken out of the composition, so n is at least 1. This could only happen if $m = n$, but then this could only occur if the paths are equal. Since the paths are distinct, φ_1 never vanishes and so α is a non-integer endomorphism. Therefore, the composition

$$\widehat{\psi}_1 \circ \cdots \circ \widehat{\psi}_m \circ \varphi_n \circ \cdots \circ \varphi_1 = \ell^{\frac{n-n'+m-m'}{2}} \alpha$$

is also a non-integer endomorphism. □

With Lemmas 6.8 and 6.9, we are now ready to prove Theorem 6.6.

Proof of Theorem 6.6. Without loss of generality, we assume $n \geq m$. By Lemma 6.9, we get

$$\alpha := \widehat{\psi_0} \circ \ldots \circ \widehat{\psi_{m'}} \circ \varphi_{n'} \circ \ldots \circ \varphi_1$$

for integers $0 \leq m' \leq m$, $1 \leq n' \leq n$, with $\alpha \in \operatorname{End}(E_6) \backslash \ell \operatorname{End}(E_6)$. In particular, α has a cyclic kernel in E_6, so it is a walk on the ℓ-isogeny graph. By Lemma 6.8, we can write

$$\alpha = a + b2i + c\frac{1+j}{2} + d\frac{i+k}{4}$$

where $a, b, c, d \in \mathbb{Z}$ and at least one of a, b, c, d is not divisible by ℓ, since α does not contain any multiples of ℓ. The degree of the endomorphism α is equal to the norm of the corresponding element in the quaternion algebra. Doing the computation, we obtain

$$\deg \alpha = \frac{1}{16}\left((4a + 2c)^2 + (8b + d)^2 + (4c^2 + d^2)p\right).$$

As long as at least one of c or d is non-zero, we have

$$\ell^{n+m} \geq \deg \alpha \geq \frac{p+1}{16}$$

In the case $\ell \neq 2$, we have $\alpha \notin \mathbb{Z} + 2i\mathbb{Z}$. This is because $\deg \alpha$ is some power of ℓ and we know ℓ does not split in $\mathbb{Z}[\sqrt{-1}]$. So, in $\mathbb{Z}[\sqrt{-1}]$, the only way to have an element of norm a power of ℓ is by taking an element of

norm 1 and multiplying it with a power of ℓ. However, since ℓ is odd, none of these elements can exist in $\mathbb{Z}[2i]$. Thus, we indeed have that at least one of c or d is non-zero, and the above inequality holds.
Suppose $\ell = 2$ and $c = d = 0$. We know $\alpha \notin 2\,\mathrm{End}(E_6)$. Furthermore, we require that $\deg \alpha = a^2 + 4b^2 = 2^k$ for some k. We thus obtain that a must be even, which implies that b must be odd, otherwise $\alpha \in 2\,\mathrm{End}(E_6)$. Looking at our options modulo 16, we see that $a^2 + 4b^2$ will be equal to either 4 or 8 modulo 16. Therefore, the only powers of two that the degree can actually be is 4 or 8. This gives us $a = \pm 2, b = \pm 1$, $a = 0, b = \pm 1$ and $a = \pm 2, b = 0$. The latter is the multiplication-by-2 map, which we exclude. So, up to signs we have two options for α, namely $\alpha = 2 + 2i$ or $\alpha = 2i$.
The map $2 + 2i$ is given by first going to $E(1728)$, then taking the endomorphism $1 + i$, and then going back to E_6 again. The map $2i$ is given by first going to $E(1728)$, then taking the automorphism i, and then going back to E_6 again. In both cases, the point of order 2 in the kernel is the one that generates the kernel of the map going to $E(1728)$. Thus, we know that the path φ must start with the unique 2-isogeny from E_6 to $E(1728)$. This completes the proof. □

6.2.2 An isogeny-based commitment scheme with starting curve E_6

We are now ready to describe the modified version of Sterner's commitment scheme. The idea is to use the curve E_6 as the starting curve, while working over a characteristic $p \equiv 15 \bmod 16$ and picking k_m and k_r such that $k_m + k_r < \frac{1}{2}\log_\ell(\frac{p+1}{16})$. The commitment scheme is given as follows:

Algorithm 6 No trusted third party commitment scheme

1: **KeyGen**(1^λ): Given a security parameter λ, fix a prime $p \equiv 15 \bmod 16$ and positive integers k_m, k_r such that $k_m + k_r < \frac{1}{2}\log_\ell(\frac{p+1}{16})$ (see Section 6.2.5 for specific parameter choices). Fix ℓ a small prime equal to 2 or congruent to 3 mod 4 and $P_0 \in E_6[\ell]$. If $\ell = 2$, P_0 is such that $j(E_6/\langle P_0\rangle) = 1728$. The public parameters are $PP := \{p, \ell, k_m, k_r, E_6, P_0\}$ and the message space is $\mathcal{M} = \{0, \ldots, \ell - 1\}^{k_m}$.
2: **Commit**(PP, m, r): Given PP, a message $m \in \mathcal{M}$, and a random $r \in \{0, \ldots, \ell - 1\}^{k_r}$, return $c = CGL(E_6, \ell, P_0, m||r)$.
3: **Open**(PP, m, r, c): Given PP, a message $m \in \mathcal{M}$, $r \in \{0, \ldots, \ell - 1\}^{k_r}$ and a commitment c, return $c == CGL(E_6, \ell, P_0, m||r)$.

The only change compared to Sterner's commitment scheme is that this one has a fixed starting curve E_6, as opposed to a starting curve provided by a trusted third party. Besides this, the rest of the commitment scheme is exactly the same. No longer needing a trusted third party makes the scheme more secure. Another benefit to using E_6 as the starting curve is that we obtain perfect binding, see Section 6.2.3. The downside of this change is that we are required to use a very large value for p, which decreases efficiency, as we will discuss in Section 6.2.5.
It is crucial to note that specifically for the case $\ell = 2$, we ensure that the first step of a commitment does not go to $E(1728)$, which is not hard to do by specifying which point of order 2 we exclude on the first step of the walk.

6.2.3 Binding property

The commitment scheme described above has perfect binding. Since we picked k_m, k_r such that $k_m + k_r < \frac{1}{2}\log_\ell(\frac{p+1}{16})$, two walks of this length will have degree less than $\frac{p+1}{16}$. Theorem 6.6 tells us that an endomorphism of such degree can only exist if $\ell = 2$ and the first step of one of the walks is the isogeny from E_6 to $E(1728)$. We have ensured that this situation will not occur. Therefore, no two pairs $(m, r), (m', r')$ with $m \neq m'$ will produce the same commitment, rendering the binding game unwinnable. This implies that this commitment scheme has perfect binding.

6.2.4 Hiding property

By Lemma 4.9, this commitment scheme can never be information-theoretically hiding, because it is perfectly binding. This is different from Sterner's scheme, which has information-theoretic hiding. The reason this is different is because the hiding property in Sterner's scheme depends on k_r being larger than the mixing constant. This gives us information-theoretic hiding. However, because k_r is so constrained due to the upper bound of the degree being $\frac{p+1}{16}$, it is not possible to make k_r large enough to even reach the theoretical lower bound of the mixing constant of a supersingular isogeny graph. Furthermore, we can no longer use the mixing constant as a basis for proving security of the hiding property. The mixing constant tells us that if we take a long enough walk, we can move between any two curves in the ℓ-isogeny graph. This necessarily means that two messages can return the same commitment, which would contradict the concept of perfect binding.
To formalise this, the authors introduce the following set for each tuple (d, E, ψ), where d is a positive integer and $\psi : E_6 \to E$ is a cyclic isogeny. After this, the set is used to prove a lemma stating what was described above.

$$\mathrm{Isog}_d(E,\psi) := \left\{ \begin{array}{c} j(E'/\mathbb{F}_{p^2});\ \text{there exists a cyclic isogeny } \phi : E \to E' \text{ of degree } d \\ \text{such that } \phi \circ \psi \text{ is a cyclic isogeny} \end{array} \right\}$$

Lemma 6.10. *Let $\phi_0 : E_6 \to E_0$ and $\phi_1 : E_6 \to E_1$ be two non-equivalent cyclic isogenies of degree ℓ^{k_m} and let k_r be an integer such that $k_m + k_r < \frac{1}{2}\log_\ell\left(\frac{p+1}{16}\right)$. Then $\mathrm{Isog}_{\ell^{k_r}}(E_0,\phi_0) \cap \mathrm{Isog}_{\ell^{k_r}}(E_1,\phi_1) = \varnothing$ in the following cases:*

1. *$\ell \neq 2$;*
2. *$\ell = 2$ and neither ϕ_0 nor ϕ_1 is a path containing the curve of j-invariant 1728.*

Proof. Suppose $\mathrm{Isog}_{\ell^{k_r}}(E_0,\phi_0) \cap \mathrm{Isog}_{\ell^{k_r}}(E_1,\phi_1) \neq \varnothing$ and let E be a curve such that $j(E) \in \mathrm{Isog}_{\ell^{k_r}}(E_0,\phi_0) \cap \mathrm{Isog}_{\ell^{k_r}}(E_1,\phi_1)$. Then there exist two distinct isogenies $\psi_0 : E_0 \to E$ and $\psi_1 : E_1 \to E$ such that $\psi_0 \circ \phi_0$ and $\psi_1 \circ \phi_1$ are cyclic isogenies. This gives two distinct paths from E_6 to E in the ℓ-isogeny graph, both of length $k_m + k_r$. We know that $2(k_m + k_r) < \log_\ell\left(\frac{p+1}{16}\right)$. This is a contradiction with Theorem 6.6. □

Instead of information-theoretic hiding, we can still prove computational hiding. To do this, the authors of [53] introduce a different problem that, provided it is hard, gives computational hiding for their commitment scheme.

Problem 6.11. *Let $k_m, k_r \in \mathbb{N}*$ such that $k_m + k_r < \frac{1}{2}\log_\ell\left(\frac{p+1}{16}\right)$. Find two cyclic isogenies $\phi_0 : E_6 \to E_0$ and $\phi_1 : E_6 \to E_1$ of degree ℓ^{k_m} and a probabilistic polynomial time distinguisher which, given some $E' \in \mathrm{Isog}_{\ell^{k_r}}(E_b,\phi_b)$ for an unknown $b \in \{0,1\}$, determines the value of b.*

This problem is somewhat similar to Problem 5.4 since it asks whether there exists an isogeny of degree ℓ^{k_r} between E' and E_0, or between E' and E_1. We can use this problem to prove computational hiding, as originally proven in [53]:

Theorem 6.12. *The commitment scheme described in Section 6.2.2 is computationally hiding under the hardness of Problem 6.11.*

Proof. Suppose we have a probabilistic polynomial time adversary $\mathcal{A} = (A_1, A_2)$ that solves the hiding game for this commitment scheme. First, A_1 returns two messages m_0 and m_1, after which one of these messages is committed to at random. Upon receiving the commitment, A_2 distinguishes which of m_0 and m_1 has been committed to. We will now construct a probabilistic polynomial time adversary $\mathcal{A}'$ that uses $\mathcal{A}$ to solve Problem 6.11.
First, $\mathcal{A}'$ queries A_1 and uses m_0 and m_1 to create cyclic isogenies ϕ_{m_0}, ϕ_{m_1} of degree ℓ^{m_k}. After this, given some E' in one of $\mathrm{Isog}_{\ell^{k_r}}(E_0,\phi_0)$ or $\mathrm{Isog}_{\ell^{k_r}}(E_1,\phi_1)$ it queries A_2 with E' to determine whether it was the

result of a commitment to m_0 or to m_1. It then returns the corresponding set $\text{Isog}_{\ell^{k_r}}(E_b, \phi_b)$ for $b \in \{0, 1\}$. Therefore, the advantage of winning the hiding game is no larger than the advantage of solving Problem 6.11. Thus, assuming this is a hard problem to solve, the commitment scheme is computationally hiding.

□

6.2.5 Parameter choices and performance

The major benefit to this scheme over Sterner's scheme is the fact that we no longer need a trusted third party. However, this scheme is much less efficient compared to Sterner's scheme because we have to use a much larger value of p. Having a larger value of p increases the size of the commitment as well as the computational efficiency. We will go over the requirements to make this scheme secure to understand why.

To obtain perfect binding as described in the commitment scheme from Section 6.2.2, we need to work over a field of characteristic p, where p is a prime such that $p \equiv 15 \bmod 16$ and, for a message of length k_m and randomness of length k_r, we have $k_m + k_r < \frac{1}{2}\log_\ell(\frac{p+1}{16})$. We also require that ℓ, the degree of the isogenies in our isogeny graph, does not split in $\mathbb{Z}[\sqrt{-1}]$.
To obtain the hiding property, we need a sufficiently large value of k_r. The best known attack on problems like Problem 6.11 is to attempt to compute an isogeny of degree ℓ^{k_r} from each of E_0 and E_1. This can be done with an algorithm like vOW [13], which requires approximately $\frac{N^{3/2}}{w^{1/2}}$ isogeny computations, where $N \approx \ell^{k_r/2}$ and w is $\log_2$ of the size of the memory in bits. This can be parallelised with multiple processors. There also exists a meet-in-the-middle attack, which has faster running time of approximately $2N$, but worse space complexity. This attack is described in section 3 of [1]. In [1], sections 5.4 to 5.6, generous estimates for the amount of memory and the amount of processors are given. This provides us with good bounds for k_r given a security parameter λ. For 128 bits of security, in the $\ell = 2$ case, we require $k_r \geq 216$, while for $\ell = 3$, we require $k_r \geq 137$. However, to be conservative, we assume a perfectly running meet-in-the-middle attack. With running time of about $2\ell^{k_r/2}$ operations, for a security level λ we require $k_r = 2\lambda$ for the case $\ell = 2$. This number can decrease if ℓ increases, but this comes at the cost of making the isogeny computations slower, as a smaller ℓ results in faster computations. The lower value of k_r does not weigh up to the decrease in computation speed, so picking $\ell = 2$ is fastest.

The length of the message k_m will be $n\lambda$ for some rational n. Together with the size of k_r, we require a prime p such that $\ell^{k_m+k_r} = \ell^{(2+n)\lambda} < \sqrt{\frac{p+1}{16}}$. In other words, our prime is such that $\log_\ell(p+1) \approx 2(n+2)\lambda + \log_\ell(16)$. So, as the size of our messages increases, so does our value of p. In terms of bits, our prime will be of size $2(n+2)\lambda + 4$ bits. In Sterner's scheme, p was fixed to be 2λ bits, regardless of the message size. This represents a major loss in efficiency for this method, since an increased size of p affects both the size of the commitment, as well as the efficiency of the isogeny computations.

6.2.6 Using a uniformly random starting curve

A clear problem arises from the proposal to use E_6 as the starting curve: to have perfect binding and remove the need for a trusted third party, we need to use very large values for p, leading to major efficiency loss. To somewhat counteract this, the authors propose another method. We will briefly go over this method, but we will not discuss the mentioned algorithms in depth. The full discussion can be found in Section 5 of [53].

In this method, instead of a fixed starting curve, a publicly generated uniformly random curve is used as the starting curve. While this curve will have known endomorphism ring, the authors introduce another computational assumption by arguing that with current known algorithms, there are lower bounds on the degrees of efficiently computable prime power degree endomorphisms. Therefore, as long as we ensure that

the degree of our walk is below the square root of this lower bound, we obtain computational binding.
According to the discussion in Section 5.1 of [53], when working over characteristic p, the lower bound of the degree of an efficiently computable endomorphism is $p^{5/2}$. To be conservative, we turn this into p^2. We thus require k_m and k_r such that $\ell^{2(k_m+k_r)} < p^2$. We still require $k_r = 2\lambda$ and we have $k_m = n\lambda$. This gives us $\log_\ell(p) \approx (2+n)\lambda$. This method would approximately halve the size of p in bits.
To generate a uniformly random curve, we start with a known curve such as $E(1728)$, and we perform a CGL hash using a random string as its input. This is one of the only known methods to generate a uniformly random curve, and it gives information about the endomorphism ring. There are no known methods to generate a uniformly random curve without revealing the endomorphism ring. This is a major open problem [7][39].

This scheme has computational binding under a problem similar to Problem 6.4. In this variant, we ask for a cyclic endomorphism of specified degree ℓ^e for some fixed $e \leq 2\log_\ell(p)$. Because current algorithms are unable to efficiently compute endomorphisms of degree below $p^{5/2}$, we obtain computational binding.
Similar to the earlier method that uses E_6 as the starting curve, this scheme also has computational hiding. The proof uses a problem similar to Problem 6.11.

6.3 Starting on j-invariant 54000

In this section, we use strategies analogous to those used in Theorem 6.6 to prove a similar but new result, now for j-invariant 54000. This result is given by Theorem 6.18. This allows us to extend the commitment scheme without a trusted third party from Section 6.2: by using j-invariant 54000 as the starting curve, we can work over characteristics p where p is of the form 11 modulo 12. In this case, the bound of the endomorphism degree is also slightly improved, allowing us to use a lower value of p. However, in practice, it does not improve the efficiency of the scheme. At best, it might decrease the size of p by a single bit.

The proposal to specifically use the curve E_6 as the starting curve is not very surprising. The work by Onuki, Aikawa and Takagi was done in the study of SIDH/SIKE, which was the most well-known isogeny-based protocol before it was broken in 2022. Generally, SIDH/SIKE already made use of primes of the form 3 mod 4 and E_6 was a common starting curve because it is always supersingular for these primes, so its properties were well-studied. This does not mean that this is the only way to go about it, however.
The theorem by Onuki, Aikawa and Takagi uses the fact that the curve $E(1728)$ can only go to the curve E_6 via a 2-isogeny. This can actually be proven fairly easily with some knowledge of the automorphism group of $E(1728)$.

Proposition 6.13. *Let $E(1728)$ be an elliptic curve with j-invariant 1728 defined over a field with characteristic $p \equiv 3 \bmod 4$. A 2-isogeny with domain $E(1728)$ is either an endomorphism, or it has its codomain in the isomorphism class of E_6.*

Proof. Since there are three subgroups of order 2 on any elliptic curve, we know that there exist three isogenies of degree 2.
The curve $E(1728)$ has a nontrivial automorphism group generated by the map $\iota : (x, y) \mapsto (-x, iy)$. We can see that $\iota^2 = [-1]$. We thus have the inclusion $\mathbb{Z}[\iota] \subset \mathrm{End}(E(1728))$. We know that when regarding the endomorphism ring as a maximal order, the degree of an endomorphism is the same as the norm of this element in the maximal order. The element $1 - i$ has norm 2, and hence the map $1 - \iota$ has degree 2. This is the first 2-isogeny, and its codomain is $E(1728)$. We denote the kernel as $\langle T \rangle$, that is, it is the subgroup generated by the point T of order 2.
Next, we take another point of order 2 not equal to T, which we will call P. Note that since the structure of the 2-torsion subgroup is $(\mathbb{Z}/2\mathbb{Z})^2$, the final point of order 2 can be written as $P + T$. One way to view $\ker(1 - \iota)$ is that it describes the points that are fixed by the map ι. Since it has degree 2, the kernel has 2

elements, namely $\mathcal{O}$ and T. In particular, this tells us that $\iota(P) = P + T$ and $\iota(P + T) = P$.
Let ϕ be the isogeny with kernel $\langle P \rangle$. This map has codomain E'. Since we know $\iota(P + T) = P$, this tells us that the map $\phi \circ \iota$ has kernel $\langle P + T \rangle$. Therefore, the isogeny with kernel $\langle P + T \rangle$ has the same codomain as ϕ.
To show that this codomain is in the isomorphism class of E_6, we do an explicit calculation. An explanation on how to do this is given in [60], Proposition 3.7. We take the elliptic curve $E : y^2 = x^3 - x$. We will now take the quotient by the point $P = (-1, 0)$. We first redefine the coordinates such that we can write P as $(0, 0)$, so we get the curve $E : y^2 = x^3 - 3x^2 + 2x$. The codomain of the quotient by $(0, 0)$ is then given by $E' : y^2 = x^3 + 6x^2 + x$, which is the curve E_6. From this, it follows that any isogeny from a curve with j-invariant 1728 has its codomain in the same isomorphism class, since we can just apply isomorphisms before and after the isogeny. □

The above statement can be proven in various different ways. However, the method above gives good intuition as for why this happens. This proof makes use of the existence of non-trivial automorphisms. This tells us that something like this likely also happens with j-invariant 0. Here, we have an automorphism group that contains a third root of unity. Therefore, the procedure we follow in the previous proof for 2-isogenies can be translated to the case of 3-isogenies. The proof requires slightly more work, but the method is ultimately the same.

Proposition 6.14. *Let $E(0)$ be an elliptic curve with j-invariant 0 defined over a field with characteristic $p \equiv 5 \bmod 6$. A 3-isogeny with domain $E(0)$ is either an endomorphism, or it has its codomain in the isomorphism class of j-invariant -12288000 (modulo p).*

Proof. Since there are four subgroups of order 3 on any elliptic curve, we know that there exist 4 isogenies of degree 3.
The curve $E(0)$ has a nontrivial automorphism group generated by the map $\omega : (x, y) \mapsto (\omega x, y)$. We have that $\omega^2 + \omega + 1 = 0$ and $\omega^3 = 1$. We thus have the inclusion $\mathbb{Z}[\omega] \subset \mathrm{End}(E(0))$. In the ring $\mathbb{Z}[\omega]$, the norm N of an element $a + b\omega$ is $N(a + b\omega) = a^2 - ab + b^2$. From here, we see that the element $1 - \omega$ has norm 3, and hence the corresponding map has degree 3. This is the first 3-isogeny, with codomain $E(0)$. We denote its kernel by $\langle T \rangle$, where T is a point of order 3.
Let P, Q and R denote three points of order 3 that generate the other three subgroups of order 3. Note that we can write $Q = T + P$ and $R = 2T + P$, since T and P generate the 3-torsion subgroup. Since we know $\ker(1 - \omega) = \langle T \rangle$, we know that the elements of $\langle T \rangle$ are also precisely the ones fixed by ω. This also tells us that the map ω^2 has $\langle T \rangle$ as its fixpoints.
We now look at the where P is sent to under the map ω. Suppose $\omega P = 2P$. Then, $\omega^3 P = \omega(\omega(\omega P)) = 2P$, but $\omega^3 = 1$, so this cannot happen. If $\omega P = 2Q$, we obtain that $\omega Q = \omega(T+P) = T+2Q = T+2T+2P = 2P$. Thus, $\omega^3 P = \omega(\omega(2Q)) = \omega(P) = 2Q$, which again is impossible. Similarly, we cannot have $\omega P = 2R$. Thus, the only options are that $\omega P = Q$ or $\omega P = R$. If $\omega P = Q$, then $\omega Q = \omega(T + P) = T + Q = 2T + P = R$ and $\omega R = \omega(2T + P) = 2T + Q = P$. In the case that $\omega P = R$, we get $\omega R = Q$ and $\omega Q = P$.
Let ϕ be the isogeny with kernel $\langle P \rangle$. This map has codomain E'. In the case that $\omega P = Q$, we obtain that $\ker(\phi \circ \omega) = \langle R \rangle$, and $\ker(\phi \circ \omega^2) = \langle Q \rangle$. In the case that $\omega P = R$, we get $\ker(\phi \circ \omega) = \langle Q \rangle$, and $\ker(\phi \circ \omega^2) = \langle R \rangle$. Thus, in either case, the isogenies with kernels $\langle Q \rangle$ and $\langle R \rangle$ have the same codomain as ϕ.
To show that the j-invariant of the codomain will always be -12288000, one can do an explicit computation, taking the elliptic curve $E : y^2 = x^3 - 1$ as the starting curve. We will not write out the computations, but explicit formulas can be found in Section 3 of [64] or Theorem 7.1 of [8]. □

We can obtain a similar result for the 2-isogeny case and j-invariant 0. This allows for a slight improvement to the bound given by Theorem 6.6.

Proposition 6.15. *Let $E(0)$ be an elliptic curve with j-invariant 0 defined over a field with characteristic $p \equiv 5 \bmod 6$. A 2-isogeny with domain $E(0)$ has its codomain in the isomorphism class of j-invariant* 54000

(modulo p).

Proof. The three points of order 2 on $E(0)$ are not fixed by ω, since we already know the fixpoints of the map $1-\omega$ are points of order 3. Therefore, the three points of order 2 are sent to each other in a cyclic manner. Let P, Q, R be the points of order 2 such that $\omega(P) = Q$ and $\omega(Q) = R$. Let ϕ be the isogeny with kernel $\langle P \rangle$. Then the isogeny $\phi \circ \omega$ has kernel $\langle R \rangle$ and $\phi \circ \omega^2$ has kernel $\langle Q \rangle$. Therefore, the three isogenies of degree 2 all have the same codomain.
To show the j-invariant, we do an explicit calculation. We take the curve $E : y^2 = x^3 + 1$. Redefining the coordinates so we can write the point $(-1, 0)$ as $(0, 0)$, we obtain the curve $E : y^2 = x^3 - 3x^2 + 3x$. Taking the quotient by $(0, 0)$ gives us an isogeny with codomain $E' : y^2 = x^3 + 6x^2 - 3x$. This curve has j-invariant 54000 over $\mathbb{Q}$, hence it has j-invariant 54000 mod p when reduced to a curve over $\mathbb{F}_p$. □

In Section 3.2 of [36], McMurdy provides us with a representative of $B_{p,\infty}$ for the case of $p \equiv 5 \bmod 6$. With this representative we can describe the structure of the endomorphism ring of the curve with j-invariant 0.

Proposition 6.16. *Let $E(0)$ be the curve $y^2 = x^3 - 1$ with j-invariant 0 and let p be a prime such that $p \equiv 5 \bmod 6$. We have that*

$$\mathrm{End}(E(0)) \cong \mathbb{Z} + \mathbb{Z}\left(\frac{-1+i}{2}\right) + \mathbb{Z}j + \mathbb{Z}\left(\frac{3+i+3j+k}{6}\right)$$

where the right hand side is a maximal order in the quaternion algebra $B_{p,\infty} \cong \left(\frac{-3,-p}{\mathbb{Q}}\right)$.

We proceed in the same way as Lemma 6.8: we know the curve with j-invariant 54000, which we will denote by $E(54000)$, is the only curve that is 2-isogenous to $E(0)$. We can once again divide one of the basis elements by 2, while multiplying another by 2, to obtain a candidate maximal order given by $\mathbb{Z} + \mathbb{Z}(-1+i) + \mathbb{Z}j + \mathbb{Z}\left(\frac{3+i+3j+k}{12}\right)$. To ensure that all basis elements have integer norm, we need a prime $p \equiv 11 \bmod 12$, since the element $\frac{3+i+3j+k}{12}$ has norm $\frac{p+1}{12}$.

Lemma 6.17. *If $p \equiv 11 \bmod 12$, then we have*

$$\mathrm{End}(E(54000)) \cong \mathbb{Z} + \mathbb{Z}(-1+i) + \mathbb{Z}j + \mathbb{Z}\left(\frac{3+i+3j+k}{12}\right)$$

Proof. The proof is analogous to the proof of Lemma 6.8. □

This lemma, along with Lemma 6.9, allows us a slightly higher lower bound on the degree of endomorphisms in the ℓ-isogeny graph.

Theorem 6.18. *Let ℓ be a prime number that does not split in $\mathbb{Z}[\frac{-1+\sqrt{-3}}{2}]$, and let $\phi = (\phi_1, \ldots, \phi_n)$ and $\psi = (\psi_1, \ldots, \psi_m)$ be two distinct paths of respective lengths n and m from $E(54000)$ to the same curve E in $G_\ell(p)$ without backtracking. Then one of the following holds:*

- $\ell^{n+m} \geq \frac{p+1}{12}$
- *$\ell = 2$ and either φ or ψ has a form $\varphi' \circ \varphi_0$ where $\varphi_0 : E(54000) \to E(0)$ is of degree 2 and $E(0)$ has j-invariant 0.*
- *$\ell = 3$ and either φ or ψ has a form $\varphi' \circ \varphi_0$ where $\varphi_0 : E(54000) \to E(54000)$ is of degree 3.*

Proof. The proof resembles the proof of Theorem 6.6. We again start with Lemma 6.9 to obtain

$$\alpha := \widehat{\psi_0} \circ \ldots \circ \widehat{\psi_{m'}} \circ \varphi_{n'} \circ \ldots \circ \varphi_1$$

for integers $0 \leq m' \leq m$, $1 \leq n' \leq n$, with $\alpha \in \mathrm{End}(E(54000)) \backslash \ell \,\mathrm{End}(E(54000))$. In particular, α has a cyclic kernel in $E(54000)$. By Lemma 6.17, we can write

$$\alpha = a + b(-1+i) + cj + d\left(\frac{3+i+3j+k}{12}\right)$$

where $a, b, c, d \in \mathbb{Z}$ and at least one of a, b, c, d is not divisible by ℓ. The degree of α is equal to the norm of the corresponding element in the quaternion algebra. We obtain

$$\deg \alpha = \frac{1}{12}(12a^2 - 24ab + 6ad + 48b^2 + d^2 + p(12c^2 + 6cd + d^2)).$$

As long as at least one of c or d is non-zero, we have

$$\ell^{n+m} \geq \deg \alpha \geq \frac{p+1}{12}.$$

In the case $\ell \notin \{2, 3\}$, we have $\alpha \notin \mathbb{Z}[\sqrt{-3}]$. This is because $\deg \alpha$ is some power of ℓ and we know ℓ is inert in $\mathbb{Z}[\frac{-1+\sqrt{-3}}{2}]$. So, in $\mathbb{Z}[\frac{-1+\sqrt{-3}}{2}]$, the only way to have an element of norm a power of ℓ is by taking an element of norm 1 and multiplying it with a power of ℓ. However, since ℓ is odd, none of these elements can exist in $\mathbb{Z}[\sqrt{-3}]$. Thus, we indeed have that at least one of c or d is non-zero, and the above inequality holds.
Suppose $\ell = 2$ and $c = d = 0$. We know 2 does not split in $\mathbb{Z}[\frac{-1+\sqrt{-3}}{2}]$. We have that $\deg \alpha = a^2 - 2ab + 4b^2$ and $\alpha \notin 2\,\mathrm{End}(E(54000))$. If we check modulo 8, we find that our only option to obtain a power of 2 is if $a = 0, b = \pm 1$, $a = 2, b = 1$ or $a = -2, b = -1$. These are (up to sign) precisely the endomorphisms obtained by first taking the 2-isogeny back to $j = 0$, then applying either the automorphism ω or ω^2, and then going back to $j = 54000$. In any case, if $c = d = 0$, one of the paths starts by going to j-invariant 0.
Suppose $\ell = 3$ and $c = d = 0$. We again have that $\deg \alpha = a^2 - 2ab + 4b^2$ and $\alpha \notin 3\,\mathrm{End}(E(54000))$. If we check modulo 9, we find that our only option to obtain a power of 3 is if $a = 1, b = 1$ or $a = -1, b = -1$. Up to sign, this corresponds with one endomorphism of degree 3 on j-invariant 54000.

□

This result allows us to extend the commitment scheme without trusted third party by also allowing the use of j-invariant 54000 as a starting curve, if a prime of the form 11 mod 12 is used. In this version of the commitment scheme, extra care needs to be taken compared to Algorithm 6.2.2. In the same way, if $\ell = 2$, we avoid going to j-invariant 0 in the first step. But in this version of the scheme, we also need to ensure that if $\ell = 3$, we do not take the endomorphism on $E(54000)$ as the first step.
While the bound of the endomorphism degree is slightly improved in Theorem 6.18 compared to Theorem 6.6, allowing us to use a lower value of p, in practice it does not improve the efficiency of the scheme. At best, p might be a single bit smaller in size.

We only use the fact that ℓ is not split in one part of the proofs of Theorems 6.6 and 6.18, the part where we show that we cannot have endomorphisms of degree a power of ℓ. If we work carefully, we could also include split primes. In Theorem 6.6, we are working with $\mathbb{Z}[i]$. A prime in $\mathbb{Z}[i]$ is split if and only if it is equivalent to 1 mod 4. If ℓ is split, we can write it as a product of an element α and its conjugate, both of norm ℓ. In the endomorphism ring, this translates to two non-equivalent endomorphisms of degree ℓ. Since any power of ℓ in $\mathbb{Z}[i]$ factorises uniquely to powers of α and its conjugate, this tells us that the only way to have an endomorphism of small degree is by making use of these two degree ℓ endomorphisms. Therefore, if we avoid taking either of these endomorphisms as the first step of our isogeny paths, we will be in the case of $\ell^{n+m} \geq \frac{p+1}{16}$ again. This requires some extra attention at the start, as it only allows for $\ell - 1$ options for the first step of our walk. After this, we are again free to choose ℓ options.

Remark 6.19. *Similar results could be obtained for j-invariant -3375. This is a complex multiplication curve over $\mathbb{Q}$, given by the curve $y^2 = x^3 - 35x + 98$, with its endomorphism ring being equal to $\mathbb{Z}[\frac{1+\sqrt{-7}}{2}]$ by*

[58]. This curve has two endomorphisms of degree 2, meaning it only has one 2-isogenous neighbour, which is the curve $y^2 = x^3 - 595x + 5586$ with j-invariant 16581375. Moreover, the kernel of this isogeny to its neighbour is defined over $\mathbb{Q}$, so we can reduce it to $\mathbb{F}_p$. From there, we can once again say something about the minimum degree of endomorphisms on this j-invariant. However, it seems unlikely that these results will give significant improvements to the value of p, but it could allow for an extension to more characteristics.
We could also work out a result for the 3-isogenous neighbour of j-invariant 0, which is j-invariant -12288000, as was shown in Proposition 6.14.
Moreover, along with the 2-isogenous neighbours of j-invariants 1728 and 0 that we saw in Sections 6.2 and 6.3, respectively, j-invariants 16581375 and -12288000 are the only j-invariants for which we can state such a result. As such, these curves are the only ones that can be used as a starting curve of a commitment scheme in the same manner as Section 6.2. This is because these are the only CM curves defined over $\mathbb{Q}$ that have ℓ-isogenous neighbours that are also CM curves defined over $\mathbb{Q}$, for some ℓ. This information is captured in the table of section A.4 in [55]. When reduced modulo p with p satisfying the right congruence condition, we obtain supersingular curves defined over $\mathbb{F}_p$. For specific values of p, one could do computations to find curves with specific endomorphism rings, but the curves listed here are the only ones that can be written down for any p.

6.4 Homomorphic commitments

We have now discussed the existing isogeny-based commitment schemes in depth. In this section, we introduce homomorphic commitment schemes and we comment on the potential existence of an isogeny-based commitment scheme that is homomorphic.
The most well-known commitment scheme in use today is the Pedersen commitment scheme. Its security is based on the discrete logarithm problem (DLP). The scheme works as follows: Alice chooses a ring with p elements, along with two multiplicative generators g and h. To commit to a secret value $x \in \{0, \dots, p-1\}$, Alice takes this x along with a random $r \in \{0, \dots, p-1\}$. The commitment is defined as $c = g^x h^r$. The randomness ensures that the scheme is hiding, while binding is ensured assuming the DLP is hard to solve.
An interesting property of the Pedersen commitment scheme is that it is *homomorphic*. This essentially means that the commitment function C acts as a homomorphism. That is, given two pairs of messages and randomness, $(x, r), (x', r')$, we have $C(x, r) \cdot C(x', r') = (g^x h^r) \cdot (g^{x'} h^{r'}) = g^{x+x'} h^{r+r'} = C(x+x', r+r')$. Being homomorphic can be a desirable property for a commitment scheme to have. It allows multiple commitments to be combined into a single commitment, still without revealing any information. An application of this is for secure electronic voting [16].
In the world of isogenies, we have found analogues of the original Diffie-Hellman protocol in the form of SIDH and later CSIDH. The original Diffie-Hellman protocol relies on the DLP for its security, while SIDH and CSIDH use isogeny assumptions. Since the Pedersen commitment also relies on the DLP for its security, it is a natural question to ask whether there also exists an isogeny-based analogue for the Pedersen commitment. While it is difficult to give a definitive answer to this question, it seems difficult to believe that such an analogue exists.

A homomorphic commitment scheme requires the commitments to be elements of a group, so that we can add them. In isogeny-based cryptography, the secret is usually an isogeny, while the commitment is an elliptic curve (or the j-invariant of one). There is no inherent structure on a set of elliptic curves that would allow us to define a group operation on this set. It therefore seems unlikely that current constructions of isogeny-based commitment schemes could yield a homomorphic commitment scheme. If there does exist a way to construct a homomorphic commitment scheme based on isogenies, it would need to involve a novel approach and it would be a major breakthrough.
The reason that we are able to find an isogeny-based analogue for Diffie-Hellman, but not for the Pedersen commitment scheme, is because Diffie-Hellman has less stringent requirements than the Pedersen commitment.

In Diffie-Hellman, all we need is a commutative group action, so that the order in which Alice and Bob act with their secret keys does not matter. The secret key that is constructed can be an element of a set, since there is no need to add two keys to each other. This is why CSIDH works as an isogeny-based analogue to the original Diffie-Hellman protocol: we are able to define a group action on the set of elliptic curves, and this is enough to construct a key exchange protocol. However, the Pedersen commitment scheme also asks for structure on the set of commitments, which current isogeny-based commitment schemes do not have.

7 Conclusion and further research

7.1 Conclusion

In this thesis, we studied isogeny-based commitment schemes. We started by discussing the necessary background in elliptic curves in Section 2. In Section 3 we introduced relevant concepts in graph theory to be able to talk about isogeny graphs, and in particular, supersingular isogeny graphs. We also introduced the relevant background in cryptography in Section 4 to be able to talk about commitment schemes formally.
In Section 5, we gave a short overview of the field of isogeny-based cryptography. In itself, this serves as an introduction to the most relevant concepts within this field. We discuss some famous protocols and we go over the main hardness assumptions. In addition to this, this section serves to motivate our study of isogeny-based commitment schemes.

In Section 6, we first studied Sterner's [62] original proposals for isogeny-based commitment schemes. The first method uses the CGL hash function [11] to go on a walk in the supersingular ℓ-isogeny graph, while the second method computes kernels in a way similar to SIDH to perform a walk in the supersingular ℓ-isogeny graph. We give an explicit example of the CGL hash function variant of the commitment scheme. We also prove the hiding and binding properties of these commitment schemes and we discuss the required parameters to make this commitment scheme secure.
Next, we studied Saah, Fouotsa, Fouotsa and Nkuimi-Jugnia's [53] proposed modification to Sterner's commitment scheme. This modification involves using the curve E_6 as the starting curve. To make this scheme secure, the characteristic p that we work over is drastically increased. This makes this version of the scheme less efficient than Sterner's, but it has the benefit of removing the need for a trusted third party. In this section, we also wrote out the proof of Theorem 6.6 in detail, and we worked out the proofs of the two lemmas needed to prove this theorem, including some extra details that were omitted in the original work.

In Section 6.3, we present original results by adapting the strategies used in the proof of Theorem 6.6 to a different setting. This allows us to state a similar but new result in Theorem 6.18. With this, we are able to construct an adaptation of the isogeny-based commitment scheme without trusted third party that works over fields of characteristics $p \equiv 11 \bmod 12$.

7.2 Further research

There are multiple potential avenues of research that can still be done within the realm of isogeny-based commitment schemes.

The most obvious work is to find different constructions for an isogeny-based commitment scheme. This could result in potential improvements in efficiency. Since the original proposal predates the SIDH attack, it might be possible to construct a scheme that implements efficiency gains using techniques similar to the attack. With a novel approach, it might also be possible to find a homomorphic commitment scheme.

As was already discussed in Remark 6.19, a similar result to Theorems 6.6 and 6.18 about the minimum degrees of endomorphisms could be obtained for the 2-isogenous neighbour of j-invariant -3375 and the 3-isogenous neighbour of j-invariant 0. This will further expand the potential congruence classes that the characteristic of our base field could be in. It may also be possible to expand this characteristic in another way, by finding a more generalised version of the endomorphism ring of the neighbouring curves. Currently, we have the endomorphism ring of E_6, the 2-isogenous neighbour of $j = 1728$, but only for primes 15 modulo 16, while the endomorphism ring of $j = 1728$ is known for primes 3 modulo 4. The method to find $\mathrm{End}(E_6)$ is

quite simple, but with some more work it might be possible to write down a form of $\mathrm{End}(E_6)$ that works for primes 3 modulo 4. From there, a similar result to Theorem 6.6 could again be proven about the minimum degrees of endomorphisms on E_6.
It is also possible to extend Theorem 6.6 to work with split primes, as explained above Remark 6.19. This would allow the commitment scheme without trusted third party to work with more degrees of isogenies. Care needs to be taken to avoid taking either of the two degree ℓ endomorphisms that now exist on the starting curve, but after this, everything works the same.

An experimental observation that was made during the creation of this thesis is that the smallest degree endomorphisms usually ended up being of a specific form. This form was to take the shortest route from the starting curve to a curve with a degree ℓ endomorphism (excluding the direct neighbour for cases $\ell = 2$ and $\ell = 3$), then taking this endomorphism, and finally taking the exact same route back. When splitting this up into two isogenies, they are technically two distinct paths, so it satisfies all the conditions of Theorems 6.6 and 6.18. With modular polynomials, we know before starting the random walk which j-invariants will have such an endomorphism. This leads to the question whether it is possible to account for these j-invariants by ignoring endomorphisms and ensuring we do not take an endomorphism as a step on a walk. For the case $\ell = 2$, this would effectively amount to ignoring this curve altogether, so it would require an extra step in the random walk to achieve the same security. However, this might lead to an improvement to the minimum endomorphism degree.

References

[1] G. Adj, D. Cervantes-Vázquez, J. Chi-Domínguez, A. Menezes, and F. Rodríguez-Henríquez. On the Cost of Computing Isogenies Between Supersingular Elliptic Curves. In C. Cid and M. J. J. Jr., editors, *Selected Areas in Cryptography - SAC 2018 - 25th International Conference, Calgary, AB, Canada, August 15-17, 2018, Revised Selected Papers*, volume 11349 of *Lecture Notes in Computer Science*, pages 322–343. Springer, 2018. doi:10.1007/978-3-030-10970-7_15.

[2] A. Basso, P. Dartois, L. D. Feo, A. Leroux, L. Maino, G. Pope, D. Robert, and B. Wesolowski. SQIsign2D-West - The Fast, the Small, and the Safer. In K. Chung and Y. Sasaki, editors, *Advances in Cryptology - ASIACRYPT 2024 - 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9-13, 2024, Proceedings, Part III*, volume 15486 of *Lecture Notes in Computer Science*, pages 339–370. Springer, 2024. doi:10.1007/978-981-96-0891-1_11.

[3] A. Basso, L. Maino, and G. Pope. FESTA: Fast Encryption from Supersingular Torsion Attacks. In *Advances in Cryptology – ASIACRYPT 2023: 29th International Conference on the Theory and Application of Cryptology and Information Security, Guangzhou, China, December 4–8, 2023, Proceedings, Part VII*, page 98–126, Berlin, Heidelberg, 2023. Springer-Verlag. doi:10.1007/978-981-99-8739-9_4.

[4] C. Baum, I. Damgård, V. Lyubashevsky, S. Oechsner, and C. Peikert. More Efficient Commitments from Structured Lattice Assumptions. In D. Catalano and R. D. Prisco, editors, *Security and Cryptography for Networks - 11th International Conference, SCN 2018, Amalfi, Italy, September 5-7, 2018, Proceedings*, volume 11035 of *Lecture Notes in Computer Science*, pages 368–385. Springer, 2018. doi:10.1007/978-3-319-98113-0_20.

[5] D. J. Bernstein, C. Dobraunig, M. Eichlseder, S. Fluhrer, S.-L. Gazdag, A. Hülsing, P. Kampanakis, S. Kölbl, T. Lange, M. M. Lauridsen, F. Mendel, R. Niederhagen, C. Rechberger, J. Rijneveld, and P. Schwabe. SPHINCS+, Stateless hash-based signatures. URL: https://sphincs.org/index.html.

[6] D. J. Bernstein, M. Hamburg, A. Krasnova, and T. Lange. Elligator: elliptic-curve points indistinguishable from uniform random strings. In A. Sadeghi, V. D. Gligor, and M. Yung, editors, *2013 ACM SIGSAC Conference on Computer and Communications Security, CCS'13, Berlin, Germany, November 4-8, 2013*, pages 967–980. ACM, 2013. doi:10.1145/2508859.2516734.

[7] J. Booher, R. Bowden, J. Doliskani, T. B. Fouotsa, S. D. Galbraith, S. Kunzweiler, S. Merz, C. Petit, B. Smith, K. E. Stange, Y. B. Ti, C. Vincent, J. F. Voloch, C. Weitkämper, and L. Zobernig. Failing to Hash Into Supersingular Isogeny Graphs. *Comput. J.*, 67(8):2702–2719, 2024. doi:10.1093/COMJNL/BXAE038.

[8] S. Bootsma. Mordell's Theorem Over Rational Function Fields Via Descent by 3-Isogeny. Bachelor's thesis, Rijksuniversiteit Groningen, 2020.

[9] W. Castryck and T. Decru. An Efficient Key Recovery Attack on SIDH. In C. Hazay and M. Stam, editors, *Advances in Cryptology – EUROCRYPT 2023*, pages 423–447, Cham, 2023. Springer Nature Switzerland.

[10] W. Castryck, T. Lange, C. Martindale, L. Panny, and J. Renes. CSIDH: An Efficient Post-Quantum Commutative Group Action. In *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*, page 395–427, Berlin, Heidelberg, 2018. Springer-Verlag. doi:10.1007/978-3-030-03332-3_15.

[11] D. X. Charles, K. E. Lauter, and E. Z. Goren. Cryptographic Hash Functions from Expander Graphs. *J. Cryptol.*, 22(1):93–113, 2009. doi:10.1007/S00145-007-9002-X.

[12] C. Costello, D. Jao, P. Longa, M. Naehrig, J. Renes, and D. Urbanik. Efficient Compression of SIDH Public Keys. In J. Coron and J. B. Nielsen, editors, *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*, volume 10210 of *Lecture Notes in Computer Science*, pages 679–706, 2017. doi:10.1007/978-3-319-56620-7_24.

[13] C. Costello, P. Longa, M. Naehrig, J. Renes, and F. Virdia. Improved Classical Cryptanalysis of SIKE in Practice. In A. Kiayias, M. Kohlweiss, P. Wallden, and V. Zikas, editors, *Public-Key Cryptography - PKC 2020 - 23rd IACR International Conference on Practice and Theory of Public-Key Cryptography, Edinburgh, UK, May 4-7, 2020, Proceedings, Part II*, volume 12111 of *Lecture Notes in Computer Science*, pages 505–534. Springer, 2020. doi:10.1007/978-3-030-45388-6_18.

[14] C. Costello and B. Smith. Montgomery curves and their arithmetic. *Journal of Cryptographic Engineering*, 8:227–240, 2018.

[15] J.-M. Couveignes. Hard Homogeneous Spaces. Cryptology ePrint Archive, Paper 2006/291, 2006. URL: https://eprint.iacr.org/2006/291.

[16] R. Cramer, M. K. Franklin, B. Schoenmakers, and M. Yung. Multi-Autority Secret-Ballot Elections with Linear Work. In U. M. Maurer, editor, *Advances in Cryptology - EUROCRYPT '96, International Conference on the Theory and Application of Cryptographic Techniques, Saragossa, Spain, May 12-16, 1996, Proceeding*, volume 1070 of *Lecture Notes in Computer Science*, pages 72–83. Springer, 1996. doi:10.1007/3-540-68339-9_7.

[17] P. Dartois, A. Leroux, D. Robert, and B. Wesolowski. SQIsignHD: New Dimensions in Cryptography. In M. Joye and G. Leander, editors, *Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part I*, volume 14651 of *Lecture Notes in Computer Science*, pages 3–32. Springer, 2024. doi:10.1007/978-3-031-58716-0_1.

[18] A. de Bruijn. Design choices in CSIDH. Master's thesis, Rijksuniversiteit Groningen, 2025.

[19] C. Delfs and S. D. Galbraith. Computing isogenies between supersingular elliptic curves over $\mathbb{F}_{\mathrm{p}}$. *Des. Codes Cryptogr.*, 78(2):425–440, 2016. doi:10.1007/S10623-014-0010-1.

[20] M. Deuring. Die Typen der Multiplikatorenringe elliptischer Funktionenkörper. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 14:197–272, 1941.

[21] J. Doliskani, G. C. C. F. Pereira, and P. S. L. M. Barreto. Faster Cryptographic Hash Function from Supersingular Isogeny Graphs. In B. Smith and H. Wu, editors, *Selected Areas in Cryptography - 29th International Conference, SAC 2022, Windsor, ON, Canada, August 24-26, 2022, Revised Selected Papers*, volume 13742 of *Lecture Notes in Computer Science*, pages 399–415. Springer, 2022. doi:10.1007/978-3-031-58411-4_18.

[22] L. D. Feo. Mathematics of Isogeny Based Cryptography, 2017. arXiv:1711.04062.

[23] L. D. Feo, J. Kieffer, and B. Smith. Towards Practical Key Exchange from Ordinary Isogeny Graphs. In T. Peyrin and S. D. Galbraith, editors, *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*, volume 11274 of *Lecture Notes in Computer Science*, pages 365–394. Springer, 2018. doi:10.1007/978-3-030-03332-3_14.

[24] L. D. Feo, D. Kohel, A. Leroux, C. Petit, and B. Wesolowski. SQISign: Compact Post-quantum Signatures from Quaternions and Isogenies. In S. Moriai and H. Wang, editors, *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7-11, 2020, Proceedings, Part I*, volume 12491 of *Lecture Notes in Computer Science*, pages 64–93. Springer, 2020. doi:10.1007/978-3-030-64837-4_3.

[25] S. D. Galbraith, C. Petit, and J. Silva. Identification Protocols and Signature Schemes Based on Supersingular Isogeny Problems. *Journal of Cryptology*, 33:130–175, 2020.

[26] HQC Team. Hamming Quasi-Cyclic (HQC), 2025. URL: https://pqc-hqc.org/doc/hqc_specifications_2025_08_22.pdf.

[27] D. Jao and L. De Feo. Towards Quantum-Resistant Cryptosystems from Supersingular Elliptic Curve Isogenies. In B.-Y. Yang, editor, *Post-Quantum Cryptography*, pages 19–34, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[28] J. Katz and Y. Lindell. *Introduction to Modern Cryptography, Second Edition*. Chapman & Hall/CRC, 2nd edition, 2014.

[29] A. Kipnis, J. Patarin, and L. Goubin. Unbalanced Oil and Vinegar schemes. *EUROCRYPT 1999*, LNCS 1592:206–222, 1999.

[30] D. R. Kohel. *Endomorphism rings of elliptic curves over finite fields*. PhD thesis, University of California at Berkeley, 1996. URL: https://www.i2m.univ-amu.fr/perso/david.kohel/pub/thesis.pdf.

[31] E. Kowalski. *An introduction to expander graphs*. Société Mathématique de France, Marseilles, 2019.

[32] J. Love and D. Boneh. Supersingular curves with small noninteger endomorphisms. *Open Book Series*, 4:7–22, 12 2020. doi:10.2140/obs.2020.4.7.

[33] L. Maino, C. Martindale, L. Panny, G. Pope, and B. Wesolowski. A Direct Key Recovery Attack on SIDH. In *Advances in Cryptology – EUROCRYPT 2023: 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part V*, page 448–471. Springer-Verlag, 2023. doi:10.1007/978-3-031-30589-4_16.

[34] R. P. Mamachan. The mathematical aspects of the Castryck-Decru key recovery attack on SIDH. Master's thesis, Università degli Studi di Padova, 2023.

[35] R. J. McEliece. A Public-Key Cryptosystem Based on Algebraic Coding Theory. *DSN Progress report*, 44:114–116, 1978.

[36] K. McMurdy. Explicit Representation of the Endomorphism Rings of Supersingular Elliptic Curves, 2014. URL: https://pages.ramapo.edu/~kmcmurdy/research/McMurdy-ssEndoRings.pdf.

[37] J.-F. Mestre. La méthode des graphes. Exemples et applications. Class numbers and fundamental units of algebraic number fields, Proc. Int. Conf., Katata/Jap. 1986, 217-242 (1986)., 1986.

[38] H. Miyaji, Y. Wang, and A. Miyaji. Lattice-Based Commitment Scheme for Low Communication Costs. *IEEE Access*, 12:111400–111410, 2024. doi:10.1109/ACCESS.2024.3421995.

[39] M. Mula, N. Murru, and F. Pintore. On Random Sampling of Supersingular Elliptic Curves. *Annali di Matematica Pura ed Applicata (1923 -)*, 204:1293–1335, 2025.

[40] M. R. Murty. Ramanujan Graphs: An Introduction. *Indian Journal of Discrete Mathematics*, 6(2):91–127, 2020.

[41] K. Nakagawa and H. Onuki. QFESTA: Efficient Algorithms and Parameters for FESTA Using Quaternion Algebras. In L. Reyzin and D. Stebila, editors, *Advances in Cryptology - CRYPTO 2024 - 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2024, Proceedings, Part V*, volume 14924 of *Lecture Notes in Computer Science*, pages 75–106. Springer, 2024. doi:10.1007/978-3-031-68388-6_4.

[42] K. Nakagawa, H. Onuki, W. Castryck, M. Chen, R. Invernizzi, G. Lorenzon, and F. Vercauteren. SQIsign2D-East: A New Signature Scheme Using 2-Dimensional Isogenies. In K. Chung and Y. Sasaki, editors, *Advances in Cryptology - ASIACRYPT 2024 - 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9-13, 2024, Proceedings, Part III*, volume 15486 of *Lecture Notes in Computer Science*, pages 272–303. Springer, 2024. doi:10.1007/978-981-96-0891-1_9.

[43] National Institute of Standards and Technology. NIST selects HQC as Fifth Algorithm for Post-Quantum Encryption. URL: https://www.nist.gov/news-events/news/2025/03/nist-selects-hqc-fifth-algorithm-post-quantum-encryption.

[44] National Institute of Standards and Technology. Post-Quantum Cryptography. URL: https://csrc.nist.gov/projects/post-quantum-cryptography.

[45] National Institute of Standards and Technology. NIST Releases First 3 Finalized Post-Quantum Encryption Standards, 2024. URL: https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards.

[46] K. Nguyen, H. Tang, H. Wang, and N. Zeng. New Code-Based Privacy-Preserving Cryptographic Constructions. In S. D. Galbraith and S. Moriai, editors, *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8-12, 2019, Proceedings, Part II*, volume 11922 of *Lecture Notes in Computer Science*, pages 25–55. Springer, 2019. doi:10.1007/978-3-030-34621-8_2.

[47] H. Onuki, Y. Aikawa, and T. Takagi. The Existence of Cycles in the Supersingular Isogeny Graphs Used in SIKE. In *International Symposium on Information Theory and Its Applications, ISITA 2020, Kapolei, HI, USA, October 24-27, 2020*, pages 358–362. IEEE, 2020.

[48] A. Petzoldt, S. Bulygin, and J. Buchmann. A multivariate based threshold ring signature scheme. *Appl. Algebra Eng. Commun. Comput.*, 24(3-4):255–275, 2013. doi:10.1007/S00200-013-0190-3.

[49] A. K. Pizer. Ramanujan graphs and Hecke operators. *Bulletin of the American Mathematical Society (New Series)*, 23(1):127–137, 1990.

[50] A. K. Pizer. Ramanujan Graphs. *Computational perspectives on number theory (Chicago, IL, 1995)*, 7:159–178, 1998.

[51] D. Robert. Breaking SIDH in Polynomial Time. In C. Hazay and M. Stam, editors, *Advances in Cryptology – EUROCRYPT 2023*, pages 472–503, Cham, 2023. Springer Nature Switzerland.

[52] A. Rostovtsev and A. Stolbunov. Public-key cryptosystem based on isogenies. Cryptology ePrint Archive, Paper 2006/145, 2006. URL: https://eprint.iacr.org/2006/145.

[53] G. T. Saah, T. B. Fouotsa, E. Fouotsa, and C. N. Jugnia. Avoiding trusted setup in isogeny-based commitments. *Des. Codes Cryptogr.*, 93(8):3207–3225, 2025. doi:10.1007/S10623-025-01633-9.

[54] R. Schoof. Nonsingular plane cubic curves over finite fields. *Journal of Combinatorial Theory, Series A*, 46(2):183–211, 1987. URL: https://www.sciencedirect.com/science/article/pii/0097316587900033.

[55] J.-P. Serre. *Lectures on the Mordell-Weil Theorem*, volume E15 of *Aspects of Mathematics*. Vieweg+Teubner Verlag Wiesbaden, 1st edition, 1989.

[56] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 124–134, 1994. doi:10.1109/SFCS.1994.365700.

[57] P. W. Shor. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Journal on Computing*, 26(5):1484–1509, Oct. 1997. doi:10.1137/s0097539795293172.

[58] J. H. Silverman. *Advanced Topics in the Arithmetic of Elliptic Curves*. Graduate Texts in Mathematics. Springer New York, NY, 1st edition, 1994.

[59] J. H. Silverman. *The Arithmetic of Elliptic Curves*, volume 106 of *Graduate Texts in Mathematics*. Springer New York, NY, 2nd edition, 2009.

[60] J. H. Silverman and J. T. Tate. *Rational Points on Elliptic Curves*. Undergraduate Texts in Mathematics. Springer Cham, 2nd edition, 2015.

[61] M. Soeten. Hasse's Theorem on Elliptic Curves. Master's thesis, Rijksuniversiteit Groningen, 2013.

[62] B. Sterner. Commitment schemes from supersingular elliptic curve isogeny graphs. *Mathematical Cryptology*, 1(2):40–51, 2021.

[63] J.-P. Tillich and G. Zémor. Collisions for the LPS Expander Graph Hash Function. In N. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, pages 254–269, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.

[64] J. Top. Descent by 3-isogeny and 3-rank of quadratic fields. *Advances in Number Theory*, pages 303–317, 1991.

[65] J. Vélu. Isogénies entre courbes elliptiques. *Comptes Rendus de l'Académie des Sciences de Paris*, 273:238–241, 1971. URL: https://gallica.bnf.fr/ark:/12148/bpt6k56191248/f52.item.

[66] B. Wesolowski. The supersingular isogeny path and endomorphism ring problems are equivalent. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 1100–1111. IEEE, 2021. doi:10.1109/FOCS52979.2021.00109.

[67] X. Xie, R. Xue, and M. Wang. Zero Knowledge Proofs from Ring-LWE. In M. Abdalla, C. Nita-Rotaru, and R. Dahab, editors, *Cryptology and Network Security*, pages 57–73, Cham, 2013. Springer International Publishing.