Commitment Schemes from Supersingular Elliptic Curve Isogeny Graphs

Bruno Sterner

Surrey Centre for Cyber Security, University of Surrey, UK

Abstract. In this work we present two commitment schemes based on hardness assumptions arising from supersingular elliptic curve isogeny graphs, which possess strong security properties. The first is based on the CGL hash function while the second is based on the SIDH framework, both of which require a trusted third party for the setup phrase. The proofs of security of these protocols depend on properties of non-backtracking random walks on regular graphs. The optimal efficiency of these protocols depends on the size of a certain constant, defined in the paper, related to relevant isogeny graphs, which we give conjectural upper bounds for.

Keywords: post-quantum cryptography, public-key cryptography, commitment schemes, isogeny-based cryptography, isogeny graphs

1 Introduction

Over the past several years, there has been an extensive effort to design cryptographic protocols that are believed to be resistant to quantum attacks. Thanks to Shor's celebrated breakthrough algorithm in 1994 [Sho94], modern public key cryptosystems based on the integer factorisation problem or discrete logarithm problem are not secure when presented with a quantum adversary. The field of post-quantum cryptography is the study of such protocols and it is widely considered that there are six main categories of post-quantum schemes. These are code-based, hash-based, multivariate-based, lattice-based, isogeny-based and MPC-based cryptography - each of which uses either a new problem which is conjectured to be hard to solve, or uses symmetric primitives to build new cryptosystems.

Commitment schemes [Blu83] have played a central role in the age of modern public-key cryptography. It allows a party to securely commit to particular value in such a way that other parties can be assured that it hasn't been tampered with. They have many useful applications: in secure electronic voting [CFSY96, DEG17], signature schemes [Lam79] and zero knowledge proofs [Dam98], to name a few.

One of the most important commitment schemes is the Pedersen commitment scheme [Ped92] based on the hardness of the discrete logarithm problem in a finite cyclic group. As such, it is vulnerable to Shor's algorithm which renders it insecure if a sufficiently large quantum computer is available. Therefore, one might hope to design a commitment scheme which is secure against quantum adversaries. There has been some work on constructing lattice-based commitment schemes [XXW13, BDL+18]. They use well known lattice based assumptions such as Ring-LWE, Module-LWE and Module-SIS as a basis for their security. There has also been some work on constructing code-based commitment schemes [NTWZ19] and multivariate-based commitment schemes [PBB13].

Isogeny based cryptography is one of the younger frameworks being considered as a basis for post-quantum cryptography. Supersingular isogeny graphs were introduced as a hard problem in cryptography by Charles, Goren and Lauter when they presented the CGL hash function at the NIST hash function competition in 2005 [CLG09]. Later, key exchange protocols [DFJP14] and signature schemes [GPS17, DFKL+20] were proposed based on supersingular isogeny graphs. For these cryptosystems, the underlying security depends

on the hardness of finding a path in the supersingular ℓ-isogeny graph, sometimes with extra auxiliary information.

At the time of writing, as far as we are aware, there are no published commitment schemes based on isogeny assumptions. Just as SIDH is an analogue of traditional Diffie-Hellman, one would hope that an analogue of Pedersen commitments exists in the isogeny setting. It is therefore surprising that this is not currently the case. Galbraith has declared this to be a "huge open problem" in isogeny-based cryptography [Gal20].

1.1 Contributions

In this work we present the first provably secure commitment schemes based on supersingular elliptic curve isogeny graphs. Underlying our protocols is the well-known idea of using a hash function to obtain a secure commitment scheme. In particular, we use the isogeny-based hash function from [CLG09] as a fundamental building block for our commitment scheme. We also use a trusted third party in the setup phase. This is to ensure that the endomorphism ring of the starting curve is not revealed.

Traditionally, proving the security of the resulting commitment scheme is done with the help of the random oracle to show that it is information-theoretically hiding. However, in this work we obtain such a scheme without using the random oracle model: instead we use mathematical properties of isogenies and their associated isogeny graphs to obtain a commitment scheme which is both information-theoretically hiding and computationally binding.

1.2 Outline

In Section 2 we begin with the necessary preliminaries needed for this work. This includes background on supersingular elliptic curve isogenies and we review the techniques used for computing such isogenies. We also give a formal definition of a commitment scheme and introduce the necessary security models. In Section 3 we introduce the mixing constant for any regular graph and analyse its properties. In Sections 4 and 5 we present our commitment schemes based on supersingular isogeny graphs and use the result from Section 3 to prove their security. In Section 6 we estimate the performance of our commitment schemes, both from a perspective of efficiency and size of the commitment values. We also attempt to compare our schemes to that of a lattice counterpart. Finally, in Section 7 we summarise the presented work and suggest avenues for future work.

2 Preliminaries

We refer to [Sil09] for a comprehensive background on elliptic curves and isogenies as well as [Sma15, Chapter 20] for more background on commitment schemes.

2.1 Supersingular Elliptic Curve Isogenies

An isogeny between two elliptic curves E and E' over a finite field $\mathbb{F}_q$ is non-zero rational map which maps points on $E(\overline{\mathbb{F}_q})$ to $E'(\overline{\mathbb{F}_q})$ and also defines a group homomorphism on these points. We call an isogeny $\mathbb{F}_{q^n}$-rational if the mapping is defined over $\mathbb{F}_{q^n}$. Two elliptic curves are isogenous if there is an isogeny between the two curves. The degree of an isogeny is its degree as a rational map. We call an isogeny an ℓ-isogeny if it has degree ℓ. For separable isogenies, the degree is the number of points in its kernel and, given any subgroup $G \subseteq E(\overline{\mathbb{F}_q})$, there exists a unique[1] separable isogeny $\phi : E \to E'$ whose kernel is G. If the degree of an isogeny is 1 then the map defines an isomorphism. Moreover there is an isomorphism between two elliptic curves if and only if their j-invariants are the same.

[1] Up to isomorphism.

An endomorphism of an elliptic curve E is an isogeny from E to itself. The set of all endomorphisms of E including the zero map forms a ring with addition and composition. This ring is called the endomorphism ring, denoted by $\mathrm{End}(E)$, and it is isomorphic to either an order in an imaginary quadratic field or a maximal order in a quaternion algebra. We say E is ordinary in the first case and supersingular in the second case. For example, when $p \equiv 3 \bmod 4$ the curve $E/\mathbb{F}_{p^2} : y^2 = x^3 + x$ is supersingular, has j-invariant 1728 and has an endomorphism ring $\mathrm{End}(E) = \langle 1, \iota, \frac{\iota+\pi}{2}, \frac{1+\iota\circ\pi}{2} \rangle$, where π is the Frobenius endomorphism and $\iota(x, y) = (-x, \sqrt{-1}y)$.

Given an isogeny $\phi : E \to E'$ there exists a unique isogeny $\hat{\phi} : E' \to E$ such that $\phi \circ \hat{\phi} = \hat{\phi} \circ \phi = [\deg(\phi)]$, where $[\cdot]$ is the scalar multiplication map. This isogeny $\hat{\phi}$ is called the dual isogeny.

Given primes ℓ, p where ℓ is small (typically 2) and p is large, the supersingular ℓ-isogeny graph is the graph whose vertex set is the isomorphism classes of supersingular elliptic curves[2] defined over $\mathbb{F}_{p^2}$ (labelled by the j-invariants) and two vertices are connected by a directed edge if there is an isogeny between the two curves of degree ℓ. In most circumstances this graph can be thought of as undirected since every isogeny has a dual isogeny. For each ℓ this graph is connected in the sense that any two vertices can be connected by a path in this graph, or equivalently, for each pair of supersingular elliptic curves E and E' there is an isogeny of degree ℓ^k between these curves. It is also $(\ell+1)$-regular in the sense that to each vertex there are $\ell + 1$ outgoing edges.

For such a graph, one can freely walk on the graph. Given a string $m \in \mathbb{Z}/(\ell+1)\mathbb{Z} \times (\mathbb{Z}/\ell\mathbb{Z})^{k-1}$, we denote by $\Phi_\ell(E, m)$ the supersingular elliptic curve obtained by going on a backtracking free walk in the supersingular ℓ-isogeny graph starting at E and dictated by the entries in the string m. Namely we get a sequence $E = E_0 \longrightarrow E_1 \longrightarrow \cdots \longrightarrow E_k = \Phi_\ell(E, m)$ of supersingular elliptic curves each of which is connected by an ℓ-isogeny. To get the next curve from E_{i-1} we first compute the irreducible factors (of degree up to $(\ell-1)/2$) of the ℓ-division polynomial of E_{i-1}. Roots of these factors correspond to points of order ℓ [Gal12, Section 25.2]. Up to generating the same subgroup and avoiding backtracking[3], label these roots as $P_0, P_1, \cdots, P_{l-1}$ and choose the point $Q_i := P_{m_i}$. Then compute the curve $E_{i+1} := E_i/\langle Q_i \rangle$ and its corresponding isogeny using Vélu's formula [Vél71]. Associated to $\Phi_\ell(E, m)$ is the isogeny $\phi_m : E \to \Phi_\ell(E, m)$ of degree is ℓ^k obtained as a composition of ℓ-isogenies.

Alternatively, one can walk freely on these isogeny graphs direct from a cyclic subgroup [CFL+19, Corollary 4.5]. Namely given an ℓ-power cyclic subgroup $G = \langle R \rangle \subseteq E[\ell^k] \subseteq E(\overline{\mathbb{F}_q})$ one can compute the isogeny $\phi : E \to E'$ whose kernel is G. This can be done by computing a chain of degree ℓ isogenies whereby the kernel of the n-th isogeny in this chain is equal to $\langle \ell^{k-n} R_n \rangle$ where $R_n = \phi_n(R_{n-1})$ and $R_1 = R$. Then composing these isogenies gives ϕ. This might look exactly the same the previous approach but the difference here is one can exploit optimal strategies [DFJP14, Section 4.2.2] to make this computation faster.

The former approach mentioned for computing isogenies was done by [CLG09] to achieve the isogeny-based hash function. The latter approach was done in [DFJP14] to achieve the isogeny-based key exchange SIDH and its counterpart SIKE.

2.2 Commitment Schemes

Throughout the section and the rest of this work we abbreviate "probabilistic polynomial-time" by PPT and denote any negligible function by negl.

Formally speaking, a commitment scheme consists of three algorithms: **KeyGen**(), **Commit**() and **Open**() - each of which has an implicit input 1^λ where λ is a security parameter. **KeyGen**() is a PPT algorithm that outputs the necessary public parameters needed for the protocol as well as the definition of the message space. **Commit**() is a PPT algorithm that, given the public parameters, a message m in the message space and a random $r \in \{0,1\}^\lambda$, outputs a value c which serves as the commitment to m and r. **Open**() is a deterministic polynomial-time algorithm that given the public parameters, the message m,

[2] Apriori defined over $\overline{\mathbb{F}_p}$.

[3] Which amounts to the next isogeny being the dual of the previous one.

1. $\text{PP} \leftarrow \textbf{KeyGen}()$
2. $(\mathtt{m}_0, \mathtt{m}_1) \leftarrow \mathcal{A}(\text{PP})$
3. $\mathtt{b} \in_R \{0,1\}$
4. $\mathtt{c} = \textbf{Commit}(\text{PP}, \mathtt{m}_\mathtt{b}, \mathtt{r})$
5. $\mathtt{b}' \leftarrow \mathcal{A}(\mathtt{c})$
6. `return` $\mathtt{b} == \mathtt{b}'$

1. $\text{PP} \leftarrow \textbf{KeyGen}()$
2. $(\mathtt{m}, \mathtt{m}', \mathtt{r}, \mathtt{r}', \mathtt{c}) \leftarrow \mathcal{A}(\text{PP})$
3. `return` $(\mathtt{m} \neq \mathtt{m}')$ && $(\textbf{Open}(\text{PP}, m, r, c) == \textbf{Open}(\text{PP}, \mathtt{m}', \mathtt{r}', \mathtt{c}) == 1)$

Figure 1: Hiding and binding games (resp.) for a commitment scheme.

the random r and the value c outputs a boolean value $b \in \{0,1\}$ according to whether or not c is a valid commitment to m and r.

Cryptographic applications of commitment schemes require the following two properties, known as hiding and binding. Informally, the hiding property ensures that the outputted commitment does not reveal anything about the message, while the binding property ensures that it should be hard to replicate the same commitment using a different message. We formally define these properties with aid of the games described in Figure 1. The hiding game is modelled like an indistinguishability game where the adversary is given the commitment of one of two messages and he is tasked to determine which message was used to derive the commitment. The binding game asks the adversary to find two distinct messages from your message space that gives the same commitment.

Definition 2.1. *Let $\mathcal{C}$ be a commitment scheme with a security parameter λ and $\mathcal{A}$ be an adversary. The hiding advantage for the adversary $\mathcal{A}$, denoted $\text{Adv}^{hid}_{\mathcal{C}}(\mathcal{A})$, is defined to be $2\,|\Pr[\mathcal{A}$ wins the hiding game$] - 1/2|$. More specifically, we have*

$$\text{Adv}^{hid}_{\mathcal{C}}(\mathcal{A}) = 2\left|\Pr\left[\mathcal{A}(\text{PP}, m_0, m_1, c) = 1 \;\middle|\; \begin{array}{c}\text{PP} \leftarrow \textbf{KeyGen}(),\\ m_0, m_1,\ c = \textbf{Commit}(\text{PP}, m_b, r)\end{array}\right] - \frac{1}{2}\right|.$$

We say that $\mathcal{C}$ is information-theoretically (resp. computationally) hiding if for all adversaries (resp. PPT adversaries) $\mathcal{A}$ there is a negligible function, negl*, such that the advantage of winning the hiding game is bounded above by* $\text{negl}(\lambda)$*. Furthermore we say $\mathcal{C}$ has perfect hiding if the hiding advantage is zero for any adversary.*

The following is a reformulation of the hiding advantage.

Lemma 2.2. *Given a commitment scheme $\mathcal{C}$ and an adversary $\mathcal{A}$, we have*

$$\begin{aligned}\text{Adv}^{hid}_{\mathcal{C}}(\mathcal{A}) = \Bigg|&\Pr\left[\mathcal{A}(\text{PP}, m_0, m_1, c) = 1 \;\middle|\; \begin{array}{c}\text{PP} \leftarrow \textbf{KeyGen}(),\\ m_0, m_1,\ c = \textbf{Commit}(\text{PP}, m_1, r)\end{array}\right]\\ &- \Pr\left[\mathcal{A}(\text{PP}, m_0, m_1, c) = 1 \;\middle|\; \begin{array}{c}\text{PP} \leftarrow \textbf{KeyGen}(),\\ m_0, m_1,\ c = \textbf{Commit}(\text{PP}, m_0, r)\end{array}\right]\Bigg|.\end{aligned}$$

Definition 2.3. *Let $\mathcal{C}$ be a commitment scheme with a security parameter λ and $\mathcal{A}$ be an adversary. The binding advantage for the adversary $\mathcal{A}$, denoted $\text{Adv}^{bind}_{\mathcal{C}}(\mathcal{A})$, is defined to be $\Pr[\mathcal{A}$ wins the binding game$]$. More specifically, we have*

$$\text{Adv}^{bind}_{\mathcal{C}}(\mathcal{A}) = \Pr\left[\begin{array}{c}\mathcal{A}(\text{PP}) = (m, m', r, r', c)\\ s.t.\ m \neq m' \ \&\ \textbf{Open}(\text{PP}, m, r, c) = \textbf{Open}(\text{PP}, m', r', c) = 1\end{array} \;\middle|\; \text{PP} \leftarrow \textbf{KeyGen}()\right].$$

We say that $\mathcal{C}$ is information-theoretically (resp. computationally) binding if for all adversaries (resp. PPT adversaries) $\mathcal{A}$ there is a negligible function, negl*, such that the advantage of winning the binding game is bounded above by* $\text{negl}(\lambda)$*. Furthermore we say $\mathcal{C}$ has perfect binding if the binding advantage is zero for any adversary.*

3 Walking on Regular Graphs

Let G be a graph with vertex set $V(G)$ and let $(v_k)_{k\geq 0}$ denote a random walk in G. For a positive integer d (which throughout this work will always be at least 3), we say G is d-regular if for each vertex $v \in V(G)$ the number of edges incident to the vertex[4] v is d. We say the random walk (v_k) is non-backtracking if it does not traverse on the same edge twice in a row, i.e., for each $k \geq 1$ the edges $[v_{k-1}, v_k]$ and $[v_k, v_{k+1}]$ are different.

The adjacency matrix of the d-regular graph G, A, is the matrix whose (i,j)-th entry is the number of (directed) edges at the vertex i going to the vertex j. Note that the powers of this matrix describes the number of paths (that may include backtracking paths) between two vertices of a given length. The transition matrix of G, P, is the matrix whose (i,j)-th entry is $A(i,j)/d$. Finally, a stationary distribution on $V(G)$, π, is a probability distribution on the set of vertices of G such that $\pi = \pi P$ or equivalently $\pi(y) = \sum_{x\in V(G)} \pi(x)P(x,y)$. In effect, this distribution remains unchanged as you transition through the graph. If G is strongly connected, this stationary distribution is unique [LP17, Corollary 1.17].

Given a random walk (v_k) in G, we define the worst-case total-variation distance to stationarity at time t to be

$$d(t) := \frac{1}{2} \max_{v\in V(G)} \left\{ \sum_{x\in V(G)} \left| \Pr_v(v_t = x) - \pi(x) \right| \right\}$$

where $\Pr_v$ denotes the probability given $v_0 = v$ and π is the stationary distribution on G. We define $t_{\text{MIX}}(\epsilon)$, the total-variation mixing time of (v_k) for $0 < \epsilon < 1$, as

$$t_{\text{MIX}}(\epsilon) := \min\{t : d(t) < \epsilon\}.$$

Theorem 3.1 (Rapid Mixing of Non-Backtracking Walks)**.** *Let G be a random d-regular graph with N vertices and $d \geq 3$. Let (v_k) be a non-backtracking random walk in G. Then for any fixed $\epsilon > 0$, the worst case total-variation mixing time with high probability satisfies*

$$t_{\text{MIX}}(1-\epsilon) \geq \lceil \log_{d-1}(dN) \rceil - \lceil \log_{d-1}(1/\epsilon) \rceil,$$
$$t_{\text{MIX}}(\epsilon) \leq \lceil \log_{d-1}(dN) \rceil + 3\lceil \log_{d-1}(1/\epsilon) \rceil + 4.$$

Proof. See [LS10, Theorem 2]. □

In other words, for a sufficiently small ϵ, this theorem says that the output of a non-backtracking random walk of length $O(\log_{d-1}(dN))$ on a random regular graph is indistinguishable from choosing a random vertex in the graph. It turns out that, compared to simple random walks that allow backtracking, the mixing time of non-backtracking walks is $\frac{d}{d-2}$ times smaller [LS10, Theorem 1].

In previous work including [JMV09], powers of the adjacency matrix, A^k, are used to get some mixing results on certain d-regular graph known as expander graphs. We are interested in the study of non-backtracking paths and for that we consider the following matrices: $A_1 = A$, $A_2 = A^2 - dI$ and $A_{r+1} = A_1A_r - (d-1)A_{r-1}$ for $r \geq 2$. Then A_r is the matrix whose (i,j)-th entry is equal the number of non-backtracking walks from i to j of length r [Mur03, Section 6].

[4] If the graph G is directed then we specify that the number of outgoing edges from v is d.

Lemma 3.2. *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. For any vertex i, the number of length r non-backtracking walks starting at i is precisely $d(d-1)^{r-1}$ and so we have

$$\sum_{k=1}^{\#V(G)} A_r(i,k) = d(d-1)^{r-1}.$$

Since $d \geq 3$, as $r \to \infty$ this sum tends to ∞ and hence there is some vertex j_0 such that $A_r(i, j_0) \to \infty$. For any vertex j, fix two paths (of length m_0, m_1) between $j_0 \to j$ (which can be done since the graph is connected). We ensure that the first step in these paths are different. Consider all paths $i \to j_0 \to j$ whereby we first go to j_0 and then traverse to j using one of our fixed paths making sure we avoid any backtracking. Then we have $A_r(i, j_0) \leq A_{r+m_0}(i, j) + A_{r+m_1}(i, j)$ and therefore $A_r(i, j) \to \infty$.

Hence for each i, j there is some $k(i, j)$ such that for all $k \geq k(i, j)$, we have $A_k(i, j)$ is strictly positive. Setting k_0 to be the maximum of $k(i, j)$ over all pairs of vertices (i, j) gives the result. □

Definition 3.3. *We define k_G to be the minimal k_0 such that Lemma 3.2 holds and call k_G the mixing constant for the graph G.*

The minimality of k_G means that there exists i_0, j_0 such that $A_{k_G-1}(i_0, j_0) = 0$, and for all i, j and $k \geq k_G$, $A_k(i, j) \neq 0$. Rephrasing this in the context of non-backtracking walks we obtain the following.

Corollary 3.4. *For a connected d-regular graph G (with $d \geq 3$) let k_G be the corresponding mixing constant. Then for all $k \geq k_G$ and every pair of vertices (i, j), there exists a non-backtracking path between i and j of length k.*

We now provide a lower bound on the mixing constant k_G. The following is a generalisation of the calculation done in [ACNL+19, Section 6]. There are at most $d(d-1)^{k-1}$ possible outputs to a non-backtracking walk of length k. For some large enough k this number of walks exceeds the number of vertices in G: $d(d-1)^{k-1} \geq N$. Rearranging this gives us a lower bound for the mixing constant:

Lemma 3.5. *The mixing constant k_G of a connected d-regular graph is bounded below by*

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

Theorem 3.1 hints at an upper bound for k_G. Namely for a suitably small $\epsilon > 0$ we expect that $k_G \leq t_{\text{MIX}}(\epsilon)$. In particular, if ϵ is negligibly small then the mixing constant may be at most $t_{\text{MIX}}(\epsilon)$. For instance $\epsilon = 1/dN$, then by Theorem 3.1 we get that $t_{\text{MIX}}(1/dN) \leq 4\lceil \log_{d-1}(dN) \rceil + 4$. To summarise we make the following conjecture.

Conjecture 3.6. *The mixing constant k_G of a connected d-regular graph G has the following upper bound:*

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

This upper bound can be thought of as a worst case bound among all regular graphs. Some regular graphs have faster mixing rates, such as expander graphs or Ramanujan graphs, so one would hope expect that the mixing constant would be smaller. Later we conjecture better upper bounds for this mixing constant in the context of supersingular isogeny graphs as well as providing some experimental data to support the conjecture.

4 A Commitment Scheme from Isogeny Assumptions

The idea of using Ramanujan graphs, that have optimal mixing properties [Alo86], in cryptography was first proposed by [CLG09]. More precisely they proposed to construct hash functions by going on random walks on certain Ramanujan graphs where path-finding is hard. This includes supersingular isogeny graphs which were proved by Pizer [Piz90] to be Ramanujan.

In this section we use supersingular elliptic curve isogeny graphs to construct a commitment scheme and use the graph theoretic results from Section 3 to prove its security. The idea behind our commitment scheme is, given a message m and a random r that someone wants to commit to, compute the isogeny-based hash of m concatenated by r. The output of this concatenation will be used as the commitment of the message m. Initially we present it in the supersingular 2-isogeny setting graph and later generalise it to the supersingular ℓ-isogeny with ℓ an odd prime.

4.1 Our Protocol

Let λ be a security parameter. The key generation of the commitment scheme is as follows. Choose a prime number p of 2λ bits, a supersingular elliptic curve $E/\mathbb{F}_{p^2}$ whose endomorphism ring is unknown and a positive integer k to be chosen later. Apart from a few notable exceptions, constructing a supersingular curve with unknown endomorphism ring is currently a hard problem and such a curve is commonly refereed to as a hard curve. Classical isogeny-based protocols use the j-invariant 1728 as the starting curve[5]. However, as alluded in Section 2.1 this curve has a special endomorphism ring which has an explicit form. Hence this curve cannot be used in our protocol. To get around this, we suggest that a trusted third party generates such a supersingular curve E by going on a random walk starting at a known node in the graph (for instance the j-invariant 1728 curve). As long as the trusted party does not reveal the path it took to get the curve E, the endomorphism ring of E should remain unknown. Finally, choose two random edges incident to $j(E)$ in the isogeny graph.

To commit to a message $m \in \{0,1\}^k$ first compute the curve $E_m := \Phi_2(E, m)$ (making sure the first step in the graph is one of the two edges chosen above). Then choose uniformly at random a binary string $r \in_R \{0,1\}^k$ and compute the curve $E' := \Phi_2(E_m, r)$. When you go from E_m to E', make sure to avoid any backtracking in the isogeny graph. Then return $c := j(E')$ as the commitment of the message m.

Given the message m, the random r and the commitment c, to open the commitment first compute the curve $\Phi_2(\Phi_2(E, m), r)$. Then return the boolean value $c == j(\Phi_2(\Phi_2(E, m), r))$.

Remark. *The necessity of the endomorphism ring of E remaining unknown is due to an attack by [EHL+ 18]. They are able to break the second preimage resistance of the isogeny hash function when the endomorphism ring of E is known. This will be important in the context of binding of our protocol.*

4.2 Hiding

The graph theoretic results presented in Section 3 along with the following well known result on random walks on isogeny graphs will be used here to show that the commitment scheme presented in the previous subsection is information-theoretically hiding.

Theorem 4.1. *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 and $n = \prod_i \ell_i^{e_i}$ be an integer where ℓ_i are small primes. Let $\hat{j}$ be the j-invariant reached by a random walk of degree n starting at j_0. Then for every j-invariant $\tilde{j}$ we have*

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

[5] Updated versions of SIKE use the j-invariant 287496 as the starting curve which is the neighbour of the j-invariant 1728 in the 2-isogeny graph.

Proof. See [GPS17, Theorem 1]. □

Theorem 4.2. *Let $k_{2,p}$ be the mixing constant for the supersingular 2-isogeny graph in characteristic p. Then for any $k \geq k_{2,p}$, the commitment scheme described in Section 4.1 is information-theoretically hiding.*

Proof. Fix two message strings m_0, m_1, a randomly chosen bit $b \in_R \{0,1\}$ and a resulting commitment $E' = \Phi(E_{m_b}, r)$. The goal for an adversary is to determine which message was used to get the commitment. Since the supersingular 2-isogeny graph is 3-regular, $k_{2,p}$ is well-defined. For any $k \geq k_{2,p}$, by Corollary 3.4, there is guaranteed to be a path of length k from E_{m_0} to E' and E_{m_1} to E'. Set $\alpha := \frac{3}{2\sqrt{2}} > 1$. Using Theorem 4.1 we have

$$\Pr[c = E' \mid \text{message is } m_0] - \Pr[c = E' \mid \text{message is } m_1] \leq 2\alpha^{-k}.$$

Similarly this difference is bounded below by $-2\alpha^{-k}$. Therefore the advantage of winning the hiding game is at most $2\alpha^{-k} \leq 2\alpha^{-k_{2,p}} \leq 2\alpha^{-2\lambda+\log_2(36)}$ (last inequality is a consequence of Lemma 3.5 and $N_p \geq p/12 - 1$), which proves the theorem. □

By the conjectural upper bound on the mixing constant, Conjecture 3.6, we can choose $k = 4\lceil \log_2(p) \rceil - 4$. With this choice of k we achieve information-theoretic hiding for our commitment scheme. As mentioned earlier, it could be possible to improve on this choice of k when specific graphs are used. Since supersingular isogeny graphs are Ramanujan graphs, one hopes that the mixing constant for these graphs is smaller. In particular we conjecture the following upper bound which we believe to be sharp for supersingular 2-isogeny graphs.

Conjecture 4.3. *With $k_{2,p}$ be as defined previously, we have the following upper bound*

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

In particular the constant in the big-Oh notation is at most 1.

Experimental results on this conjecture show that for every prime $p \leq 65600$ and some primes between $123000 \leq p \leq 131100$ and $234000 \leq p \leq 2^{18}$, the associated mixing constant for the supersingular 2-isogeny graph is no more than $\log_2(p) + \log_2(\log_2(p)) + \frac{3}{10}$. These mixing constants were calculated by first computing the adjacency matrix, A, for the graph and sequentially compute A_k, as defined in Section 3, until you find a value $\hat{k}$ such that the each entry of $A_{\hat{k}}$ is non-zero. We verify that $k_{2,p} = \hat{k}$ by computing $A_{\hat{k}+1}, A_{\hat{k}+2}, \cdots, A_{\hat{k}+i}$ for some small i and see if the entries in these matrices are non-zero. Since the entries of these matrices grow as we increase k, then as long as these matrices have non-zero entries, we can conclude that $k_{2,p} = \hat{k}$. Figure 2 tabulates the mixing constant in the supersingular 2-isogeny graph in characteristic p for all $p \leq 65600$.

If this conjecture is true then we can choose $k = \lceil \log_2(p) + \log_2(\log_2(p)) + 1 \rceil$ and it would significantly speed up the performance of the protocol.

4.3 Binding

We will prove that the binding of our protocol is secure under the following hard problem.

Problem 1 (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[6] of E whose degree is a prime power ℓ^e.*

[6]By non-trivial we mean an endomorphism which is not a multiplication-by-m map, $[m]$, and by cyclic we mean an endomorphism whose kernel is cyclic.

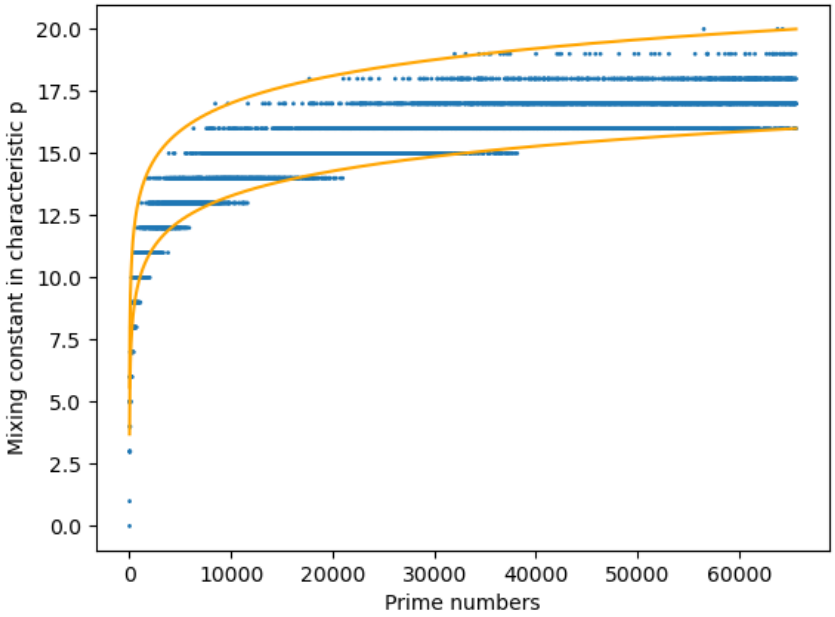

Figure 2: Mixing constant in prime characteristic p for all $p \leq 65600$. The lower bound curve is $\log_2(x)$ and the upper bound curve is $\log_2(x) + \log_2(\log_2(x))$.

A similar problem was presented in the SQISign identification protocol [DFKL+20]. The only difference is that this problem is more restrictive in the degree of the endomorphism. In their setting the degree of computed endomorphism has smooth degree.

Also, as remarked in [DFKL+20], this problem is equivalent to the endomorphism ring problem, namely compute a $\mathbb{Z}$-basis for this endomorphism ring of an elliptic curve. For more details on this equivalence see [EHL+18].

Theorem 4.4. *The commitment scheme as described in Section 4.1 is computationally binding under the Supersingular Smooth Endomorphism Problem on the curve E for the prime $\ell = 2$.*

Proof. Suppose that $\mathcal{A}$ is a PPT adversary which successfully solves the binding game for this commitment scheme. We shall construct an PPT adversary $\mathcal{A}'$ using $\mathcal{A}$ as a black box that solves the Supersingular Smooth Endomorphism Problem on the curve E.

Upon receiving the curve E, $\mathcal{A}'$ queries $\mathcal{A}$ and successfully outputs m, m', r, r' such that $m \neq m'$ and $E' = \Phi_2(E_m, r) = \Phi_2(E_{m'}, r')$. Let $\phi_m : E \to E_m$, $\phi_{m'} : E \to E_{m'}$, $\phi_r : E_m \to E'$, $\phi_{r'} : E_{m'} \to E'$ be the associated isogenies each of which has degree 2^k.

Then the composition of $\phi_m \circ \phi_r \circ \hat{\phi_{r'}} \circ \hat{\phi_{m'}}$ is an endomorphism of E whose degree is the prime power 2^{4k}. First we need to verify that this composition is non-trivial. Suppose for a contradiction that $\phi_m \circ \phi_r \circ \hat{\phi_{r'}} \circ \hat{\phi_{m'}} = [2^{2k}]$. Since the compositions $\phi_m \circ \phi_r$ and $\phi_{m'} \circ \phi_{r'}$ are isogenies from E to E' of same degree, then $\hat{\phi_{r'}} \circ \hat{\phi_{m'}}$ is the dual of $\phi_m \circ \phi_r$. Since $\widehat{\phi_{m'} \circ \phi_{r'}} = \hat{\phi_{r'}} \circ \hat{\phi_{m'}}$, we get $\phi_m \circ \phi_r = \phi_{m'} \circ \phi_{r'}$. As a result $m = m'$ and $r = r'$ – which gives the contradiction.

By removing any potential backtracking to the composition $\phi_m \circ \phi_r \circ \hat{\phi_{r'}} \circ \hat{\phi_{m'}}$ that might occur as we approach E', we get a cyclic endomorphism ψ. The adversary $\mathcal{A}'$ outputs this endomorphisms and solves the Supersingular Smooth Endomorphism Problem on E in PPT. Therefore the advantage of winning the binding game is at most the advantage of solving the above problem. □

4.4 Generalisation

In this section we generalise the above idea and construct a commitment scheme which works in the supersingular ℓ-isogeny graph for a small odd prime ℓ. Once again, key generation of the protocol is the same as described in Section 4.1.

To commit to a message $m \in \{0, 1, \cdots, \ell - 1\}^k$ first compute the curve $E_m \coloneqq \Phi_\ell(E, m)$. Then choose uniformly at random a binary string $r \in_R \{0, 1, \cdots, \ell - 1\}^k$ and compute the curve $E' \coloneqq \Phi_\ell(E_m, r)$. Once again, when you go from E_m to E', making sure to avoid any backtracking in the isogeny graph. Then return $c \coloneqq j(E')$ as the commitment of the message m.

Given the message m and the random r, to open the commitment scheme first compute the curve $\Phi_\ell(\Phi_\ell(E, m), r)$. Then return the boolean value $c == j(\Phi_\ell(\Phi_\ell(E, m), r))$.

Much like in the setting of the 2-isogeny graph, we have the following theorems proving the security of this commitment scheme.

Theorem 4.5. *Let $k_{\ell,p}$ be the mixing constant for the supersingular ℓ-isogeny graph in characteristic p. Then for any $k \geq k_{\ell,p}$, the commitment scheme described above is information-theoretically hiding.*

Proof. The proof is analogous to the proof of Theorem 4.2. □

Much like in Section 4.2, choosing $k = 4\lceil \log_\ell(p) \rceil + 8$ would be sufficient to get information-theoretic hiding.

Theorem 4.6. *The above commitment scheme is computationally binding under the Supersingular Smooth Endomorphism Problem on the curve E and the prime ℓ.*

Proof. The proof is analogous to the proof of Theorem 4.4. □

Along with this we make the following conjecture on a sharp upper bound of the mixing constant in these graphs.

Conjecture 4.7. *With $k_{\ell,p}$ as above, we have*

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

In particular the constant in the big-Oh notation is at most 1.

5 Commitments using the SIDH Approach

In this section we describe a variant of the protocol from the previous section which uses the SIDH framework. Instead of using SIDH friendly primes we use primes of the form $2^n f - 1$ and achieve the same security requirements that were achieved in the previous section. One advantage of doing this is to exploit SIDH strategies [DFJP14, Section 4.2.2] to speed up isogeny computations. (Similar ideas in the context of the hash function construction can be found here [DPB17]).

Let $p = 2^n f - 1$ be a prime with 2λ bits and f is a small integer. In the same manner as described in the previous section, choose a supersingular elliptic curve $E/\mathbb{F}_{p^2}$ whose endomorphism ring is unknown but this time we make sure that $\#E(\mathbb{F}_{p^2}) = (2^n f)^2$. This is done intentionally so that the 2^l-torsion subgroup of E entirely consists of points whose coordinates are in $\mathbb{F}_{p^2}$. Let $P_0, P_1 \in E[2^n]$ be points on E that form a basis for this 2^n-torsion subgroup of E.

Much like in the Section 4, we will go on walks in the supersingular 2-isogeny graph but instead of choosing at each step which edge to traverse, we compute the kernel subgroup and corresponding the isogeny whose kernel is this subgroup. However the longest isogeny that can be computed as a sequence of 2-isogenies using this approach has degree 2^n. So in order to attain the same security as a obtained in Section 4.2,

namely computing an isogeny of degree 2^k with $k = 4\lceil \log_2(p) \rceil \approx 4n$, we must do this isogeny computation 4 times each for the message used and the randomly generated element.

Recall that given $m_0 \in \mathbb{Z}/2^n\mathbb{Z}$ the subgroup of $E[2^n]$ defined by $\langle P + m_0 Q \rangle$ induces an isogeny $\phi_{m_0} : E \to E_{m_0}$ whose kernel is this subgroup. If we wish to extend this walk by going on another walk of degree 2^n, then we must find points P', Q' on E_{m_0} that form a basis for the respective 2^n-torsion subgroup. Also we need a procedure of computing these points in a deterministic manner. Ensuring that if replicated by another party we get the same points. We already know that $\phi_{m_0}(Q)$ has order 2^n, so set $Q' \coloneqq \phi_{m_0}(Q)$.

To deterministically compute a point P' we use techniques from [CJL+17]. We briefly summarise this method. Use the Elligator 2 method for deterministically computing points R in $E(\mathbb{F}_{p^2})$ [BHKL13], then check that $R \in E \setminus [2]E$, where $[2]E$ is the set of all 2-divisible points on E. If so then the point fR is a point of order 2^n. A check has to be made to see if this point is independent from Q'. If not then it cannot be used as a second basis element and we repeat the whole process until you compute a point P' which can be used as the second basis element. For more details on this see [CJL+17, Section 3.2].

Remark. *The choice of $Q' = \phi_{m_0}(Q)$ was done purposefully. It ensures that the isogeny induced by a kernel of the form $\langle P' + m_1 Q' \rangle$ will not result in backtracking through part of the first isogeny. This is because the kernel of the dual isogeny, $\ker(\widehat{\phi_{m_0}})$, is generated by the point Q' [NR19, Proposition 3].*

5.1 Protocol Description and Security

The key generation is as described above.

To commit to a message $m \in \mathbb{Z}/2^{4n}\mathbb{Z}$ do as follows. Compute $m_0 \coloneqq m \mod 2^n$, $m_1 \coloneqq (m - m_0)/2^n \mod 2^n$, $m_2 \coloneqq (m - m_1 - m_0 2^n)/2^{2n} \mod 2^n$ and $m_3 \coloneqq (m - m_2 - m_1 2^n - m_0 2^{2n})/2^{3n} \mod 2^n$. Notice that $m_0, m_1, m_2, m_3 \in \mathbb{Z}/2^n\mathbb{Z}$. Compute the subgroup $M_0 \coloneqq \langle P + m_0 Q \rangle$ and hence the corresponding isogeny $\phi_{m_0} : E \to E_{m_0}$ whose kernel is M_0. Compute the point $Q' \coloneqq \phi_{m_0}(Q)$ and a point P' as described above. Now compute the subgroup $M_1 \coloneqq \langle P' + m_1 Q' \rangle$ and hence the corresponding isogeny $\phi_{m_0} : E_{m_0} \to E_{m_1}$ whose kernel is M_1. Again compute the points P'', Q''. Repeat this for the integers m_2, m_3 to get isogenies $\phi_{m_i} : E_{m_{i-1}} \to E_{m_i}$ whose kernel is M_i $(i = 2, 3)$. Henceforth the curve $E_m \coloneqq E_{m_3}$ and the composition $\phi_m = \phi_{m_3} \circ \phi_{m_2} \circ \phi_{m_1} \circ \phi_{m_0}$ is the curve and isogeny obtained from the message m.

As remarked above, the choice of basis points for the 2^n-torsion subgroup is done so that we don't get any backtracking in the isogeny graph.

From here choose a random $r \in \mathbb{Z}/2^{4n}\mathbb{Z}$ and repeat the same procedure as done above. Once again you make sure that there is no backtracking through ϕ_m by making sure that you have an appropiate basis for the 2^n-torsion subgroup. The result is an isogeny $\phi_r : E_m \to E'$. Then return the curve $c \coloneqq j(E')$ as the commitment of the message m.

Much like in Section 4, given the message and the random $m, r \in \mathbb{Z}/2^{4n}\mathbb{Z}$ (resp.), to open the commitment you recompute the curve E' and return the boolean value $c == j(E')$. The deterministic nature of computing the new basis for the next 2^n-torsion subgroup means that, as long as the message m and the random r are as intended, then anyone can open the message and be assured that this is the correct message used.

Theorem 5.1. *The commitment scheme described above is information theoretically hiding and computationally binding under the Supersingular Smooth Endomorphism Problem on the curve E for the prime $\ell = 2$.*

Proof. Application of Theorem 4.2 & Theorem 4.4. □

6 Comparison

In this section we estimate the performance of these schemes, only in the setting when $\ell = 2$, and attempt to compare them to other post-quantum commitment schemes.

In the work by [DPB17], they attempted to compare the performance of the CGL hash function with a hash function that is analogous to the idea presented in Section 5. If a prime of the form $p = 2^n f - 1$ is used and k is the length of the walk you want to compute, they estimated the complexity of the CGL hash function as $kn(5.7n + 110)\mathbf{m}$ and the complexity of the SIDH variant as $kn(13.5\log(n) + 42.4)\mathbf{m}$, where $\mathbf{m}$ is the cost of performing a field multiplication. These performance timings translate to our commitment scheme constructions by choosing $k = 4\lceil \log_2(p) \rceil - 4$ with one exception. In the SIDH variant of our commitment scheme a little more work is needed then that presented above since we need to generate the basis elements for the new torsion subgroup. This requires computing one isogeny image as well as the cost of doing the Elligator 2 method to determine the second basis element. Since this is done at most 3 times, it doesn't add much to the complexity mentioned above. Approximately it adds $O(n\mathbf{m})$ to the overall complexity which is primarily dominated by the isogeny image computations.

Therefore, the performance ratio of the scheme described in Section 4 versus that of this Section 5 is approximately $(5.7n + 110)/(13.5\log(n) + 46.4 + O(1/k))$. This implies an exponential speed up in the performance of the commitment scheme presented in this Section 5 versus that described in Section 4 (especially when the prime p is of cryptographic size).

As mentioned earlier, if the validity of Conjecture 4.3 holds then the performance of these procotols will significantly speed up by up to a factor of 4.

Lets now look at the size of the commitment values in our schemes. In both variants, these values just consist of one j-invariant of a supersingular elliptic curve which is an element in $\mathbb{F}_{p^2}$. Equivalently, given a 2-dimensional representation of $\mathbb{F}_{p^2} = \mathbb{F}_p[i]$, we can express this j-invariant as two $\mathbb{F}_p$ elements. Hence, given a prime p of 2λ bits with λ a security parameter, the size of the commitment value is approximately 4λ bits or $\lambda/2$ bytes. It is worth mentioning that the size of the commitment value does not depend on the size of k. This point is consistent with most isogeny schemes, including the CGL hash function.

One can compare these commitment scheme to that of other post-quantum alternative. One clear advantage this has over other alternatives is that the size of the committed values. To target 128 bits of security, the size of the committed value in our scheme is approximately 64 B. In comparison to that of lattice based commitment schemes taken from [BDL+18, Table 2], to achieve the same level of security, the committed values is approximately 9 kB. This is much larger than that of our isogeny commitment schemes. There are a few notable drawbacks when comparing our schemes to its alternatives. First one is the performance of our schemes. Even faster variant described in Section 5 is not as fast as its lattice counterpart. This point is again consistent with most isogeny schemes. Second drawback is that it is not a homomorphic commitment scheme. Having a homomorphic commitment scheme is desirable to have since there are some strong applications that rely on the homomorphic property such as electronic voting.

7 Conclusion

In this work we presented two commitment schemes based on isogeny assumptions. This is the first provably secure commitment scheme in the isogeny literature. The scheme follows the approach of [CLG09] whereby we go on walks in supersingular isogeny graphs. We proved that this commitment scheme is secure attaining information-theoretic hiding and computational binding. We obtained information-theoretic hiding based on the existence of a mixing constant, k_G, implying that any two vertices in the graph can be connected by a non-backtracking path of fixed length k for any $k \geq k_G$. We conjectured an upper bound on this constant for both the generic setting and the specific setting of supersingular isogeny graphs. We obtained computationally binding by reducing a binding instance to a well known isogeny problem which is believed to be hard even for quantum adversaries.

We also presented a variant of this commitment scheme which is constructed through a kernel subgroup to compute the isogenies instead of going through step by step and choosing which edge to continue. Its security follows directly from the security of the previous scheme. The main advantage that this variant has over the previous commitment scheme is that of efficiency.

There are a number of open problems that arise from this work.

- Proving the explicit upper bounds for the mixing constant in both the generic setting and the special setting of the supersingular isogeny graphs.
- See how sharp we can makes these upper bounds and see if we can get close to the bound presented in Conjecture 4.3 and Conjecture 4.7 in the specific setting of supersingular isogeny graphs.
- Constructing a homomorphic commitment scheme based on isogeny assumptions. This problem would be considered a major breakthrough in this area.

Acknowledgements: The author would like to thank Robert Granger, Liqun Chen and Daniel Gardham for their comments and discussions. Also thanks to Steven Galbraith for his comments on an earlier version of the work as well as the anonymous reviewers for their comments.

References

[ACNL+19] Sarah Arpin, Catalina Camacho-Navarro, Kristin Lauter, Joelle Lim, Kristina Nelson, Travis Scholl, and Jana Sotáková. Adventures in supersingularland. *arXiv preprint arXiv:1909.07779*, 2019.

[Alo86] Noga Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.

[BDL+18] Carsten Baum, Ivan Damgård, Vadim Lyubashevsky, Sabine Oechsner, and Chris Peikert. More efficient commitments from structured lattice assumptions. In Dario Catalano and Roberto De Prisco, editors, *Security and Cryptography for Networks*, pages 368–385, Cham, 2018. Springer International Publishing.

[BHKL13] Daniel J Bernstein, Mike Hamburg, Anna Krasnova, and Tanja Lange. Elligator: Elliptic-curve points indistinguishable from uniform random strings. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 967–980, 2013.

[Blu83] Manuel Blum. Coin flipping by telephone a protocol for solving impossible problems. *ACM SIGACT News*, 15(1):23–27, 1983.

[CFL+19] Anamaria Costache, Brooke Feigon, Kristin Lauter, Maike Massierer, and Anna Puskás. Ramanujan graphs in cryptography. In *Research Directions in Number Theory*, pages 1–40. Springer, 2019.

[CFSY96] Ronald Cramer, Matthew Franklin, Berry Schoenmakers, and Moti Yung. Multi-authority secret-ballot elections with linear work. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 72–83. Springer, 1996.

[CJL+17] Craig Costello, David Jao, Patrick Longa, Michael Naehrig, Joost Renes, and David Urbanik. Efficient Compression of SIDH Public Keys. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017*, pages 679–706, Cham, 2017. Springer International Publishing.

[CLG09] Denis X. Charles, Kristin E. Lauter, and Eyal Z. Goren. Cryptographic hash functions from expander graphs. *Journal of Cryptology*, 22(1):93–113, Jan 2009.

[Cou06] Jean Marc Couveignes. Hard homogeneous spaces. *IACR Cryptol. ePrint Arch.*, 2006:291, 2006.

[Dam98] Ivan Damgård. Commitment schemes and zero-knowledge protocols. In *School organized by the European Educational Forum*, pages 63–86. Springer, 1998.

[DEG17] Ashraf Darwish and Maged M El-Gendy. A new cryptographic voting verifiable scheme for e-voting system based on bit commitment and blind signature. *Int J Swarm Intel Evol Comput*, 6(158):2, 2017.

[DFJP14] Luca De Feo, David Jao, and Jérôme Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. *Journal of Mathematical Cryptology*, 8(3):209–247, 2014.

[DFKL+20] Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski. SQISign: Compact post-quantum signatures from quaternions and isogenies. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2020*, pages 64–93, Cham, 2020. Springer International Publishing.

[DPB17] Javad Doliskani, Geovandro C. C. F. Pereira, and Paulo S. L. M. Barreto. Faster cryptographic hash function from supersingular isogeny graphs. Cryptology ePrint Archive, Report 2017/1202, 2017. https://eprint.iacr.org/2017/1202.

[EHL+18] Kirsten Eisenträger, Sean Hallgren, Kristin Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: reductions and solutions. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 329–368. Springer, 2018.

[Gal12] Steven D. Galbraith. *Mathematics of Public Key Cryptography*. Cambridge University Press, USA, 1st edition, 2012.

[Gal20] Steven D. Galbraith. Talk given at the indian institute of technology kharagpur: "similarities and differences between diffie-hellman and isogeny crypto", August 2020.

[GPS17] Steven D. Galbraith, Christophe Petit, and Javier Silva. Identification protocols and signature schemes based on supersingular isogeny problems. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology – ASIACRYPT 2017*, pages 3–33, Cham, 2017. Springer International Publishing.

[JDF11] David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In Bo-Yin Yang, editor, *Post-Quantum Cryptography*, pages 19–34, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[JMV09] David Jao, Stephen D Miller, and Ramarathnam Venkatesan. Expander graphs based on GRH with an application to elliptic curve cryptography. *Journal of Number Theory*, 129(6):1491–1504, 2009.

[Lam79] Leslie Lamport. Constructing digital signatures from a one-way function. Technical report, Citeseer, 1979.

[LP17] David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.

[LS10] Eyal Lubetzky and Allan Sly. Cutoff phenomena for random walks on random regular graphs. *Duke Mathematical Journal*, 153(3):475–510, 2010.

[Mur03] M Ram Murty. Ramanujan graphs. *Journal-Ramanujan Mathematical Society*, 18(1):33–52, 2003.

[NR19] Michael Naehrig and Joost Renes. Dual isogenies and their application to public-key compression for isogeny-based cryptography. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology – ASIACRYPT 2019*, pages 243–272, Cham, 2019. Springer International Publishing.

[NTWZ19] Khoa Nguyen, Hanh Tang, Huaxiong Wang, and Neng Zeng. New code-based privacy-preserving cryptographic constructions. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology – ASIACRYPT 2019*, pages 25–55, Cham, 2019. Springer International Publishing.

[PBB13] Albrecht Petzoldt, Stanislav Bulygin, and Johannes Buchmann. A multivariate based threshold ring signature scheme. *Applicable Algebra in Engineering, Communication and Computing*, 24(3-4):255–275, 2013.

[Ped92] Torben Pryds Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In Joan Feigenbaum, editor, *Advances in Cryptology — CRYPTO '91*, pages 129–140, Berlin, Heidelberg, 1992. Springer Berlin Heidelberg.

[Piz90] Arnold K Pizer. Ramanujan graphs and hecke operators. *Bulletin of the American Mathematical Society*, 23(1):127–137, 1990.

[RS06] Alexander Rostovtsev and Anton Stolbunov. Public-key cryptosystem based on isogenies. *IACR Cryptol. ePrint Arch.*, 2006:145, 2006.

[Sho94] 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.

[Sil09] Joseph H Silverman. *The Arithmetic of Elliptic Curves, 2nd Edition*. Graduate Texts in Mathematics. Springer, 2009.

[Sma15] Nigel P. Smart. *Cryptography Made Simple*. Springer Publishing Company, Incorporated, 1st edition, 2015.

[Vél71] Jacques Vélu. Isogénies entre courbes elliptiques. *CR Acad. Sci. Paris, Séries A*, 273:305–347, 1971.

[XXW13] Xiang Xie, Rui Xue, and Minqian Wang. Zero knowledge proofs from ring-lwe. In Michel Abdalla, Cristina Nita-Rotaru, and Ricardo Dahab, editors, *Cryptology and Network Security*, pages 57–73, Cham, 2013. Springer International Publishing.