Collisions in Supersingular Isogeny Graphs and the SIDH-based Identification Protocol

Wissam Ghantous[1], Federico Pintore[2] and Mattia Veroni[3]

[1] Mathematical Institute, University of Oxford, UK,
wissam.ghantous@maths.ox.ac.uk
[2] Department of Mathematics, University of Bari, IT,
federico.pintore@uniba.it
[3] NTNU - Norwegian University of Science and Technology, Trondheim, NO ,
mattia.veroni@ntnu.no

Abstract. The digital signatures that have been proposed so far in the setting of the Supersingular Isogeny Diffie-Hellman scheme (SIDH) were obtained by applying the Fiat-Shamir transform - and a quantum-resistant analogous, the Unruh transform - to an interactive identification protocol introduced by De Feo, Jao and Plût. The security of the resulting schemes is therefore deduced from that of the base identification protocol.
In this paper, we revisit the proofs that have appeared in the literature for the special soundness property of the above mentioned SIDH-based identification protocol. All such proofs consider the same extraction algorithm, which is claimed to always extract a valid witness for a statement x when given two valid transcripts, with the same commitment and different challenges, relative to x itself. We show that this is not always the case, with some explicit counterexamples. The general argument fails due to some special cycles in supersingular isogeny graphs. The existence of these special cycles not only enjoys a theoretical interest, but is generally relevant for the Isogeny-based Cryptography. We provide some theoretical results on their presence in supersingular isogeny graphs, and discuss the relevance of the obtained results for some known cryptographic applications.

Keywords: Isogeny-based Cryptography · Identification Protocol · Special Soundness · Supersingular Isogeny Graph · Digital Signature

1 Introduction

While traditional Elliptic Curve Cryptography (ECC) relies on the Elliptic Curve Discrete Logarithm Problem (ECDLP), which is easily solvable on a quantum computer using Shor's algorithm [34], Isogeny-based Cryptography is a newer field of elliptic-curve cryptosystems which rely on different assumptions - e.g., the hardness of the CSSI problem [17] - that so far have resisted quantum cryptanalysis. As a consequence, this field represents one of the options for Post-Quantum Cryptography. Nevertheless, Isogeny-based Cryptography is less studied than traditional ECC, and has only recently received broad attention to scrutinise hardness assumptions and existing primitives, and build new ones based on them.

The first isogeny-based cryptographic primitive was a Diffie-Hellman-like key exchange, usually named CRS, based on the action of ideal class groups on isomorphism classes of ordinary elliptic curves. This was designed independently by Couveignes [14] and Rostovtsev and Stolbunov [33]. Due to the inefficiency of the CRS scheme and the subsequent subexponential attack by Childs *et al.* [12], supersingular elliptic curves were suggested as a potential replacement to ordinary ones. In 2006, Charles, Lauter and Goren [11] introduced a collision resistant hash function, exploiting the supersingular isogeny graph. Subsequently, Jao *et al.* [17] constructed the Supersingular-isogeny Diffie-Hellman scheme, named SIDH, featuring small key sizes and whose security relies on a newly introduced isogeny problem, named SSDDH problem. In [23], Galbraith *et al.* showed that SIDH is vulnerable to an adaptive attack when one party uses a static key, which led to the introduction of the CCA-secure Key Encapsulation Mechanism SIKE [4]. SIKE is obtained from an encryption scheme which was proposed, together with an SIDH-based identification protocol, in [17]. Both cryptosystems are deduced from the SIDH scheme.

The only existing SIDH-based digital signatures [24, 41] were obtained by transforming the above-mentioned identification protocol into a non interactive one via the Fiat-Shamir transform [21] and a quantum-resistant analogous, the Unruh transform [36]. The UnForgeability against Chosen Message Attack (UF-CMA) of such digital signature schemes is deduced from the Honest-Verifier Zero-Knowledge (HVZK) and Special Soundness (SS) properties of the base identification protocol, and the hardness of the isogeny-based SSDDH and CSSI [17] problems. Currently, the best known classical and quantum attacks on the CSSI problem are generic claw-finding attacks, and both require exponential complexity [13, 28].

Our contribution. We revisit the proofs, provided in [17, 24, 41], for the special soundness property of the SIDH-based identification protocol. All such proofs consider the same extractor - which we denote by $\mathsf{Ex}_{\mathsf{SIDH}}$ - that, on input two valid transcripts with the same commitment and different challenges for a statement x, is claimed to output a witness w for x.

In more details, let the cyclic secret isogeny $\varphi : E_0 \to e_1$ of prime-power degree $\ell_1^{e_1}$ be a witness w for the statement $\mathsf{x} = (e_1, \varphi(P_2), \varphi(Q_2))$, where $\{P_2, Q_2\}$ is a basis of the torsion subgroup $E_0[\ell_2^{e_2}]$, with ℓ_2 a prime different from ℓ_1. The isogeny φ is uniquely identified by its kernel $\mathsf{Ker}(\varphi)$. Then, two valid transcripts with the same pair of curves (E_2, E_3) as commitment and different challenges for x give knowledge of three cyclic isogenies ϕ, ψ, ϕ' satisfying the following conditions:

- $\phi : E_0 \longrightarrow E_2$ and $\phi' : e_1 \longrightarrow E_3$ have degree $\ell_2^{e_2}$;
- $\psi : E_2 \longrightarrow E_3$ has same degree $\ell_1^{e_1}$ of φ.

Hence, the isogeny $\hat{\phi}' \circ \psi \circ \phi$ goes from E_0 to e_1, where $\hat{\phi}'$ denotes the dual isogeny of ϕ'. The extractor $\mathsf{Ex}_{\mathsf{SIDH}}$ is designed to output $\mathsf{Ker}(\hat{\phi}' \circ \psi \circ \phi) \cap E_0[\ell_1^{e_1}]$ as a witness w of x.

We depict two scenarios, detailed below, where the produced w is not a witness for x, i.e. where $\mathsf{Ex}_{\mathsf{SIDH}}$ fails in outputting the kernel defining φ. By providing some explicit examples, we demonstrate the concrete existence of such scenarios, even for some of the parameters sets considered for SIKE. Therefore, this invalidates the special-soundness proofs given in the literature, which incorrectly assumed the kernel of ψ to be always of the form $\phi(\mathsf{Ker}(\varphi))$.

In the first of the two exception scenarios, the cyclic isogeny ψ is completely unrelated to φ, i.e. $\mathsf{Ker}(\psi) \neq \phi(\mathsf{Ker}(\varphi))$, despite being an $\ell_1^{e_1}$-degree isogeny from E_0 to e_1. In this case $\mathsf{Ker}(\hat{\phi'} \circ \psi \circ \phi) \cap E_0[\ell_1^{e_1}]$ determines an isogeny of degree $\ell_1^{e_1}$ going from E_0 to an elliptic curve E' different from e_1. In the second scenario, the cyclic isogeny ψ corresponds to an isogeny φ' rather than to φ, i.e. $\mathsf{Ker}(\psi) = \phi(\mathsf{Ker}(\varphi'))$ instead of being equal to $\phi(\mathsf{Ker}(\varphi))$. The isogeny φ' is *twinned* with φ, meaning that, despite being different from φ, it goes from E_0 to e_1, has degree $\ell_1^{e_1}$ and there exists a point $R \in E_0$ of order $\ell_2^{e_2}$ such that the isogenies having kernels $\langle \varphi(R) \rangle$ and $\langle \varphi'(R) \rangle$ both go into the curve E_3. In this case $\mathsf{Ker}(\hat{\phi'} \circ \psi \circ \phi) \cap E_0[\ell_1^{e_1}]$ determines the isogeny φ' instead of the isogeny φ.

We note that both scenarios exploit *collisions* in supersingular isogeny graphs, i.e. different cyclic isogenies with the same degree, domain and codomain. In particular, in the first scenario the isogeny ψ forms a collision with the isogeny having $\phi(\mathsf{Ker}(\varphi))$ as kernel, while in the second scenario both the isogenies φ, φ' and those with kernels $\langle \varphi(R) \rangle, \langle \varphi'(R) \rangle$ form a collision. In the second case, the two collisions are tightly related, and for this reason we call them *double collisions*. The relevance of collisions in supersingular isogeny graphs has emerged also in other cryptographic contexts, as for example the adaptive attacks against SIDH [23] and the analysis of the claw-finding attack against the CSSI problem [13]. However, such collisions are usually overlooked, as they are deemed as unlikely to exist.

In this work, we study the collisions formed by cyclic isogenies of degree ℓ^e that occur in the supersingular isogeny graph $\mathcal{G}_{p^2}(\ell)$ for primes p and ℓ, whose vertices are isomorphism classes of supersingular elliptic curves over $\mathbb{F}_{p^2}$ and whose edges are isogenies of degree ℓ. By relating the total number of such collisions to traces of Brandt matrices, we express it as a sum of modified Hurwitz class numbers. We then obtain upper bounds for the total number of collisions that we can have in $\mathcal{G}_{p^2}(\ell)$. This method, however, cannot yield any meaningful lower bounds, as it would involve getting a strong handle on the behaviour of certain incomplete character sums. In fact, bounding incomplete character sums is a very difficult problem in analytic number theory, and the best known bounds are not tight enough for our purposes. Instead, we introduce a statistical model to mimic the splitting behaviour of Legendre symbols and estimate the modified Hurwitz class numbers. We then obtain heuristic lower bounds for the number of collisions, under our statistical model.

Building on the above results, we then give estimates for the number of double collisions for fixed primes p, ℓ and ℓ'. Our model suggests that the number of such double collisions is non-negligible. Moreover, our results are analysed in the context of the adaptive attacks against SIDH and the claw-finding algorithm for the CSSI problem, as they show that certain collisions are extremely rare, a fact that is implicitly assumed in [23] and [13]. We provide a rigorous proof of this fact by computing an upper bound for the corresponding collisions.

Finally, we discuss possible alternatives, from recent results on quaternion algebras, to the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$. Furthermore, we highlight that a modification of $\mathsf{Ex}_{\mathsf{SIDH}}$ which returns $\hat{\phi}' \circ \psi \circ \phi$ when $\mathsf{Ker}(\hat{\phi}' \circ \psi \circ \phi) \cap E_0[\ell_1^{e_1}]$ is not a witness of x provides the SIDH-based sigma protocol with a weaker property, called gap (or relaxed) special soundness. This property allows the extraction of a witness w which is different from the proper witness being proven within the protocol but contained in a relaxed binary relation having the proper binary relation as a subset. We stress that the digital signature schemes deduced in [24, 41] by applying the Fiat Shamir and Unruh transforms are still secure under the relaxed special soundness, since the relaxed binary relation is hard.

Related Work. In an independent work [16] recently published on eprint, De Feo *et al.* also dispute the existing proofs for the special soundness property of the SIDH-based identification protocol. In particular, they show that two valid transcripts can be easily produced for a statement x for which does not exist a corresponding witness. Consequently, not just $\mathsf{Ex}_{\mathsf{SIDH}}$, but every extractor would fail in extracting a witness for x given such transcripts. Therefore, they propose a modified version of the SIDH-based identification protocol, for which they prove special soundness. Compared to the original protocol, the modified version is less efficient both in terms of computation and bandwidth.

We note that in the literature two slightly different notions of special soundness can be found. In the first one, the extraction algorithm is required to extract a witness given two valid transcripts for a statement x in the language, i.e. there must exist a witness w for x [20, 31]. In the second one, the statement x is not required to be in the language [3, 37]. We note that the latter definition includes the former as a special case. The main reason for considering one notion instead of the other appears to be the cryptographic application for which the identification protocol which satisfies the special soundness property is considered. For some applications, the first definition is *enough*, and then there is no reason for considering the more general definition. For example, if the considered identification protocol is turned into a digital signature by applying the Fiat-Shamir transform, the first definition is sufficient to prove the UF-CMA security of the scheme.

Therefore, despite the counterexamples provided in [16], the original SIDH-based identification protocol appears to be still exploitable to construct Fiat-Shamir digital signatures. In fact, [16] shows that $\mathsf{Ex}_{\mathsf{SIDH}}$ does not provide special soundness to the protocol only when considering the second definition. Moreover, the original protocol would be the preferable choice in terms of efficiency.

It is only in this work that we show that $\mathsf{Ex}_{\mathsf{SIDH}}$ does not provide special soundness to the protocol even when considering the first definition, and therefore that the original protocol cannot be safely used for Fiat-Shamir digital signatures as it stands. In the light of this, we deem [16] and this paper to be two complementary results, both helpful in assessing the security of the SIDH-based identification protocol proposed in [17].

Roadmap. The rest of the paper is organised as follows. In Section 2, we provide some preliminaries on identification schemes, isogenies, quaternion algebras and algebraic number theory. In Section 3 we show two scenarios where the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$ fails to extract a witness for a statement x (for which a witness exists) given two valid transcripts for x. In particular, some concrete counterexamples are provided, even for some of the SIKE parameters sets. Section 4 discusses the implications of the presented counterexamples and possible countermeasures. In Section 5, we provide a deterministic upper bound on the number of collisions in an isogeny graph. We also determine heuristic lower bounds for the same quantity. Section 6 analyses the provided results on collisions in the context of the adaptive attacks against SIDH and the claw-finding algorithm for the CSSI problem. Finally, in Section 7 we draw some conclusions.

2 Preliminaries

2.1 Identification Protocols

Let X and Y be two sets whose size depends on a security parameter λ. Then $\mathcal{R} \subset X \times Y$ is a polynomially-computable binary relation on $X \times Y$ if, for any $(\mathsf{x}, \mathsf{w}) \in X \times Y$, whether (x, w) belongs to $\mathcal{R}$ can be checked in time $\mathsf{poly}(|\mathsf{x}|)$. If $(\mathsf{x}, \mathsf{w}) \in \mathcal{R}$, we call w a *witness* for the *statement* x. The *language* corresponding to $\mathcal{R}$ is $\mathcal{L}_{\mathcal{R}} = \{\mathsf{x} \in X \mid \exists\ \mathsf{w} \in Y : (\mathsf{x}, \mathsf{w}) \in \mathcal{R}\}$. Hereafter, we omit $\mathcal{R}$ from the subscript when the relation is clear from context.

An *identification protocol* ID for a polynomially-computable binary relation $\mathcal{R}$ is a special type of public-coin three-move interactive protocol between a prover and a verifier. Informally, a prover holding a pair $(\mathsf{x}, \mathsf{w}) \in \mathcal{R}$ can prove to a verifier that they possess a valid witness w for x, without revealing any information about w.

Definition 1 (Identification protocols)**.** An identification protocol ID for a binary relation $\mathcal{R}$ consists of three algorithms $(\mathsf{P} = (\mathsf{P}_1, \mathsf{P}_2), \mathsf{V})$, where V is deterministic and we assume that P_1 and P_2 are probabilistic polynomial-time (PPT) algorithms sharing states. We denote by ComSet, ChSet, and ResSet the commitment space, challenge space, and response space, respectively. Then ID has the following three-move flow:

- On input $(\mathsf{x}, \mathsf{w}) \in \mathcal{R}$, the prover runs $\mathsf{com} \leftarrow \mathsf{P}_1(\mathsf{x}, \mathsf{w})$ and sends the commitment com to the verifier.
- The verifier chooses a random challenge $\mathsf{ch} \overset{\$}{\leftarrow} \mathsf{ChSet}$, and sends ch to the prover.
- Given ch, the prover runs $\mathsf{resp} \leftarrow \mathsf{P}_2(\mathsf{x}, \mathsf{w}, \mathsf{ch})$ and returns the response resp to the verifier.
- The verifier runs $\mathsf{V}(\mathsf{x}, \mathsf{com}, \mathsf{ch}, \mathsf{resp})$ and outputs 1 if they accept, 0 otherwise.

A *transcript* $(\mathsf{x}, \mathsf{com}, \mathsf{ch}, \mathsf{resp}) \in X \times \mathsf{ComSet} \times \mathsf{ChSet} \times \mathsf{ResSet}$ of the protocol is said to be valid (relative to x) in case $\mathsf{V}(\mathsf{x}, \mathsf{com}, \mathsf{ch}, \mathsf{resp})$ outputs 1.

We require the following three properties from an identification protocol ID:

1. **Correctness.** All transcripts that are honestly generated must be valid. More precisely, for all $(\mathsf{x}, \mathsf{w}) \in \mathcal{R}$,

$$\Pr\left[\mathsf{V}(\mathsf{x}, \mathsf{com}, \mathsf{ch}, \mathsf{resp}) = 1 \;\middle|\; \begin{array}{c} \mathsf{com} \leftarrow \mathsf{P}_1(\mathsf{x}, \mathsf{w}), \\ \mathsf{ch} \xleftarrow{\$} \mathsf{ChSet}, \\ \mathsf{resp} \leftarrow \mathsf{P}_2(\mathsf{x}, \mathsf{w}, \mathsf{ch}) \end{array}\right] = 1.$$

2. **Honest-Verifier Zero-Knowledge (HVZK).** The view of an honest verifier on a protocol run can be simulated, and thus an honest verifier learns nothing on the secret witness. More formally, there exists a PPT simulator algorithm Sim that, on input a statement $\mathsf{x} \in \mathcal{L}_{\mathcal{R}}$ and challenge $\mathsf{ch} \in \mathsf{ChSet}$, outputs a commitment com and response resp such that $(\mathsf{x}, \mathsf{com}, \mathsf{ch}, \mathsf{resp})$ is a valid transcript. Moreover, the output distribution of Sim on input $(\mathsf{x}, \mathsf{ch})$ is computationally indistinguishable from the distribution of those outputs generated via an honest execution of ID conditioned on the verifier sampling ch as the challenge.

The third property which is required is special soundness. Two slightly different definitions of special soundness can be considered, which can be both found in the literature. The only difference between the two is in the set where the statement x lives.

3a. **Special Soundness.** There exists a polynomial-time extraction algorithm Ex such that, given any two valid transcripts $(\mathsf{x}, \mathsf{com}, \mathsf{ch}, \mathsf{resp})$ and $(\mathsf{x}, \mathsf{com}, \mathsf{ch}', \mathsf{resp}')$ relative to the same statement $\mathsf{x} \in \mathcal{L}_{\mathcal{R}}$, with the same commitment com and two distinct challenges $\mathsf{ch} \neq \mathsf{ch}'$, outputs w such that $(\mathsf{x}, \mathsf{w}) \in \mathcal{R}$.

3b. **Special Soundness (general).** There exists a polynomial-time extraction algorithm Ex such that, given any two valid transcripts $(\mathsf{x}, \mathsf{com}, \mathsf{ch}, \mathsf{resp})$ and $(\mathsf{x}, \mathsf{com}, \mathsf{ch}', \mathsf{resp}')$ relative to the same statement x, with the same commitment com and two distinct challenges $\mathsf{ch} \neq \mathsf{ch}'$, outputs w such that $(\mathsf{x}, \mathsf{w}) \in \mathcal{R}$.

We note that, in order for definition 3b to be verified, one can produce no more than one valid transcript relative to a statement not in the language. On the contrary, this limitation is not imposed by definition 3a. As definition 3a is contained into definition 3b, one could ask why preferring using the former over the latter. Depending on the targeted application, however, the first definition might be enough. For example, it is very common to use the Fiat-Shamir transform [21] to turn an identification protocol (with an exponentially-large challenge set) ID into a digital signature. When signing a message m, the prover (which is now the signer) first produces a commitment com by running P_1 on a pair $(\mathsf{x}, \mathsf{w}) \in \mathcal{R}$, with w acting as the secret key corresponding to the verification key x. Instead of waiting for a challenge from the verifier, the prover produces it by computing the image of a random function on the message m and the commitment

com. This allows the verifier to later replicate the computation and check that the challenge has been honestly generated. Finally, the signature consists of the commitment and the response to the commitment. The HVZK property and the special soundness defined by definition 3a of ID are sufficient to prove the UF-CMA security of the resulting digital signature scheme. In fact, the general reduction [1] to prove the digital signature secure retrieves a witness for a statement x by running an adversary against the UF-CMA game as a sub-routine, and in particular by making it output two valid transcripts relative to x. The reduction algorithm then executes the extractor Ex on the two valid transcripts. If it extracts correctly, then it knows that $\mathsf{x} \in \mathcal{L}_{\mathcal{R}}$; if it does not, then it knows that $\mathsf{x} \notin \mathcal{L}_{\mathcal{R}}$. Therefore, in this case, it is not necessary to use the more general definition of special soundness, and one can use definition 3a instead, as the possibility of producing more than one valid transcripts for a statement not in the language does not affect the reduction algorithm.

For identification protocols for binary relations defined over post-quantum algebraic structures, it is standard to consider a relaxation of the soundness notion. To be more precise, for ID it is only required the existence of an extraction algorithm which recovers a (weaker) witness in a larger binary relation $\tilde{\mathcal{R}}$ - with $\mathcal{R} \subset \tilde{\mathcal{R}}$ - rather than in $\mathcal{R}$. Such property, that can be given in two different flavours like the notion of special soundness, is formalised as follows.

Relaxed Special Soundness. There exists a polynomial-time (relaxed) extraction algorithm rEx such that, given a statement $\mathsf{x} \in \mathcal{L}_{\mathcal{R}}$ (alternatively, a statement $\mathsf{x} \in X$) and any two valid transcripts relative to x, $(\mathsf{x}, \mathsf{com}, \mathsf{ch}, \mathsf{resp})$ and $(\mathsf{x}, \mathsf{com}, \mathsf{ch}', \mathsf{resp}')$ with the same commitment com and $\mathsf{ch} \neq \mathsf{ch}'$, outputs w satisfying $(\mathsf{x}, \mathsf{w}) \in \tilde{\mathcal{R}}$, where $\tilde{\mathcal{R}}$ is a polynomially-computable binary relation on $X \times Y$ which contains $\mathcal{R}$.

It has been shown that, as long as $\tilde{\mathcal{R}}$ is still a sufficiently hard relation, then identification protocols satisfying relaxed special soundness are as useful as standard identification protocols (see, for example, [2, 5, 6, 20]).

2.2 Supersingular Elliptic Curves, Isogenies, and Hardness Assumptions

In this section we briefly recall some standard properties of isogenies between elliptic curves over finite fields. We refer the interested reader to [35] for a detailed treatment of the topic.

Let $\mathbb{F}_q$ be a finite field, where q is a power p^n of a prime $p > 5$. Given two elliptic curves E and E' defined over $\mathbb{F}_q$, an isogeny $\varphi : E \to E'$ is a non-constant regular rational map such that $\varphi(0_E) = 0_{E'}$. Every isogeny φ can be written in the form $(F_1(x)/F_2(x), yG_1(x)/G_2(x))$, where $F_1, F_2, G_1, G_2 \in \overline{\mathbb{F}}_q[x]$ ($\overline{\mathbb{F}}_q$ being the algebraic closure of $\mathbb{F}_q$), F_1 is coprime with F_2, and G_1 is coprime with G_2. If the coefficients of the four polynomials are contained in $\mathbb{F}_{q^k}$, then φ is said to be defined over $\mathbb{F}_{q^k}$, and E, E' are isogenous over $\mathbb{F}_{q^k}$. Tate's theorem states that E, E' are isogenous over $\mathbb{F}_{q^k}$ if and only if $\#E(\mathbb{F}_{q^k}) = \#E'(\mathbb{F}_{q^k})$.

An isomorphism is an isogeny that is invertible. An isogeny with the same domain and range is an endomorphism. The set $\mathsf{End}(E)$ of all endomorphisms of an elliptic curve E together with the zero map form a ring under pointwise addition and composition, called endomorphism ring of E. E is said to be ordinary if $\mathsf{End}(E)$ is commutative, supersingular otherwise. Every supersingular elliptic curve over an extension of $\mathbb{F}_p$ is isomorphic to an elliptic curve defined over $\mathbb{F}_{p^2}$.

The degree $\mathsf{deg}(\phi)$ of an isogeny φ is the maximum in $\{\deg(F_1), \deg(F_2)\}$. Two elliptic curves E and E_1 are d-isogenous if there exists an isogeny $\varphi : E \to E_1$ of degree d, and in this case we say φ is a d-isogeny. Given a natural number d and a prime power q, we denote by $\mathcal{G}_q(d)$ the graph whose vertices are $\mathbb{F}_q$-isomorphism classes of supersingular elliptic curves over $\mathbb{F}_q$ and whose edges are isogenies of degree d.

The composition of two isogenies of degrees d_1 and d_2, respectively, is an isogeny of degree $d_1 d_2$. Given an isogeny $\varphi : E \to E'$, the dual $\hat{\varphi} : E' \to E$ of φ is an isogeny - defined on the same field and having the same degree d of φ - such that $\varphi \circ \hat{\varphi} = \hat{\varphi} \circ \varphi = [d]$. If $(d, p) = 1$ and $\mathsf{Ker}(\varphi) = \langle T \rangle$, given another point $R \in E$ such that $\{T, R\}$ is a basis for $E[d]$, we have $\mathsf{Ker}(\hat{\varphi}) = \langle \varphi(R) \rangle$ [39].

Each isogeny φ has a finite kernel and, for a separable isogeny φ, it holds that $\mathsf{deg}(\varphi) = \mathsf{Ker}(\varphi)$. In the following we restrict our attention to separable isogenies. Vice versa, if H is a finite subgroup of an elliptic curve E, then there are a unique (modulo isomorphisms) elliptic curve E/H and a separable isogeny $\psi : E \to E'$ such that $\mathsf{Ker}(\psi) = H$. Both E/H and ψ can be computed with complexity $O(\#H)$ using Velu's formulas.

Given a natural number ℓ, we denote by $E[\ell]$ the ℓ-torsion subgroup $\{P \in E \mid [\ell]P = 0\}$ of E. When ℓ and p are relatively prime, $E[\ell] \simeq \mathbb{Z}_\ell \times \mathbb{Z}_\ell$. Cryptographic schemes deduced from the SIDH [17] paradigm consider primes p of the form $p = \ell_1^{e_1} \ell_2^{e_2} f \pm 1$, where e_1, e_2, f are natural numbers, ℓ_1, ℓ_2 are small primes and $\ell_1^{e_1} \approx \ell_2^{e_2}$. Under these hypotheses, every supersingular elliptic curve over $\mathbb{F}_{p^2}$ is isomorphic, over $\mathbb{F}_{p^2}$, to an elliptic curve E defined over $\mathbb{F}_{p^2}$ and such that $E(\mathbb{F}_{p^2}) = (\ell_1^{e_1} \ell_2^{e_2} f)^2$. As a consequence, $E[\ell_1^{e_1}]$, $E[\ell_2^{e_2}]$ are both contained in $E(\mathbb{F}_{p^2})$. We will denote by $\{P_A, Q_A\}$ and $\{P_B, Q_B\}$ a basis for $E[\ell_1^{e_1}]$ and $E[\ell_2^{e_2}]$, respectively.

Hardness Assumptions. The hard problems based on isogenies which we will consider in the subsequent sections are presented here. We first introduce the (arguably) most standard problem for isogenies.

Problem 1 (Explicit Isogeny Problem for Fixed Degree [15])**.** *Given a natural number d and two d-isogenous elliptic curves E_0 and E_1 the explicit isogeny problem for fixed degree d (EI_d) asks to find an isogeny $\varphi : E_0 \to E_1$ of degree d.*

The fastest algorithm to solve the above problem is based on the classical meet-in-the-middle strategy, and has computational complexity $O(d^{\frac{1}{2}})$ [17].

The following hard problems are tailored to the SIDH paradigm. In particular, the elliptic curves taken into consideration are supersingular, defined over

$\mathbb{F}_{p^2}$ for a prime p of the form $\ell_1^{e_1}\ell_2^{e_2}f \pm 1$, and with $(\ell_1^{e_1}\ell_2^{e_2})$ rational points. As the adversary is always required to determine an isogeny between two given elliptic curves, for simplicity we denote by E_0 the starting curve, and by $\{P_A, Q_A\}$, $\{P_B, Q_B\}$ two bases for $E_0[\ell_1^{e_1}]$ and $E_0[\ell_2^{e_2}]$, respectively.

Problem 2 (Computational Supersingular Isogeny (CSSI) Problem, [17])**.** *Let $\phi : E_0 \to e_1$ be an isogeny whose kernel is $\langle [m_A]P_A + [n_A]Q_A \rangle$, where m_A, n_A are uniformly sampled from $\mathbb{Z}/\ell_1^{e_1}\mathbb{Z}$ and are not both divisible by ℓ_1. Given* e$_1$*and the values $\phi(P_B), \phi(Q_B)$, determine a generator of the subgroup $\langle [m_A]P_A + [n_A]Q_A \rangle \subset E_0[\mathbb{F}_{p^2}]$.*

The fastest algorithms to solve the above problem are still based on the meet-in-the-middle strategy, both for classical and quantum attacks. The best classical attack has computational complexity $\tilde{O}(\ell_1^{e_1/2})$, and the best quantum attack has computational complexity $\tilde{O}(\ell_1^{e_1/3})$ [17].

The next two problems were originally formulated in [17] and are believed to be as hard as the computational variants. For example, the Decisional SIDH Isogeny problem reduces to its computational variant in the trivial way, and Jao and Urbanik [38] show that the converse holds as well.

Problem 3 (Decisional SIDH Isogeny Problem, [17])**.** *Given two supersingular elliptic curves E_0, E_1 and a basis $\{P'_B, Q'_B\}$ of $E_1[\ell_2^{e_2}]$, determine whether or not there exists an isogeny $\varphi : E_0 \to E_1$ such that $\mathsf{deg}(\varphi) = \ell_1^{e_1}$, and $\varphi(P_B) = P'_B$ and $\varphi(Q_B) = Q'_B$.*

Problem 4 (Decisional Supersingular Product (DSSP) Problem, [41])**.** *Let $\varphi : E_0 \to E_1$ be an isogeny of degree $\ell_1^{e_1}$. Given (E_2, E_3, ψ) sampled with probability $1/2$ from one of the following distributions, determine which distribution it is from:*

- *choose a random point $R \in E_0[\ell_2^{e_2}]$ of order $\ell_2^{e_2}$. Let $\phi : E_0 \to E_2$ and $\phi' : E_1 \to E_3$ be the isogenies with kernels $\langle R \rangle$ and $\langle \varphi(R) \rangle$, respectively. Then let $\psi : E_1 \to E_2$ be the isogeny having $\langle \phi(S) \rangle$ as kernel, where $\mathsf{deg}(\psi) = \ell_1^{e_1}$.*
- *choose E_2 randomly among all the supersingular elliptic curves defined over $\mathbb{F}_{p^2}$ having the same number of rational points as E_0. Then, choose a random point $U \in E_2$ of order $\ell_1^{e_1}$ and compute the isogeny $\psi : E_2 \to E_3$ having $\langle U \rangle$ as kernel.*

In the next section we will relate the security of the identification scheme to the DSSP problem. Since the public parameters of a protocol instance consist of the tuple $\mathsf{pp}_{\mathsf{SIDH}} = (\ell_1, \ell_2, e_1, e_2, f, p, E_0, P_1, Q_1, P_2, Q_2)$, and the hardness of the problem depends on these parameters, we denote by $\mathrm{DSSP}_{\mathsf{pp}}$ the DSSP problem on a fixed parameter set pp.

2.3 Quaternion algebras and Brandt matrices

A *quaternion algebra* B is a 4-dimensional central simple algebra over a field, which is $\mathbb{Q}$ in our case of interest. Given a basis $\{1, \mathbf{i}, \mathbf{j}, \mathbf{k}\}$ of B, with $\mathbf{ij} = -\mathbf{ji} = \mathbf{k}$, a quaternion algebra can be written as

$$B = \mathbb{Q} + \mathbb{Q}\mathbf{i} + \mathbb{Q}\mathbf{j} + \mathbb{Q}\mathbf{k}$$

with $a, b \in \mathbb{Q}^\times$ such that $\mathbf{i}^2 = a, \mathbf{j}^2 = b$. A quaternion algebra is usually denoted by the Hilbert symbol as $H_\mathbb{Q}(a, b)$ or $\left(\frac{a,b}{\mathbb{Q}}\right)$. Several examples in the literature show that a and b are not enough to uniquely determine B up to isomorphisms, so we need a new characterization.

An *order* $\mathcal{O}$ in a $\mathbb{Q}$-algebra B is a subring of B that is also a lattice (a finitely generated $\mathbb{Z}$-module) that spans B over $\mathbb{Q}$ (i.e. $\mathcal{O} \otimes_\mathbb{Z} \mathbb{Q} = B$). Given a prime p, we define $B_p := B \otimes_\mathbb{Q} \mathbb{Q}_p$ as the quaternion algebra obtained by extending the scalars of B from $\mathbb{Q}$ to $\mathbb{Q}_p$ (i.e. the p-adic completion of $\mathbb{Q}$ with respect to the p-adic norm). We extend this notation to include ∞ by setting $B_\infty := B \otimes_\mathbb{Q} \mathbb{R}$. It follows from Wedderburn's theorem that B_p is either isomorphic to the matrix algebra $M_2(\mathbb{Q}_p)$, and we say that B is *unramified* at p, or it is a division ring, in which case we say that B is *ramified* at p. A quaternion algebra is uniquely determined, up to isomorphisms, by the set of primes at which it ramifies: this set has even cardinality, and may include ∞. Conversely, any such set may arise as the set of primes at which a quaternion algebra ramifies.

Given a quaternion $\alpha = t + x\mathbf{i} + y\mathbf{j} + z\mathbf{k} \in B$, the *involution* of α is defined as $\overline{\alpha} := t - x\mathbf{i} - y\mathbf{j} - z\mathbf{k} \in B$. Via the involution we can define the *reduced trace* of α as $rd(\alpha) := \alpha + \overline{\alpha} = 2t$ and the *reduced norm* of α as $Nm(\alpha) := \alpha\overline{\alpha} = t^2 - ax^2 - by^2 + abz^2$.

Let p be a prime and B be the quaternion algebra over $\mathbb{Q}$ ramified at p and infinity. Fix a maximal order $\mathcal{O}$ of B. Since multiplication of quaternions is not commutative, we must be careful when defining the following objects which involve multiplication; their left or right analogues can be easily deduced. Two left ideals I and J of $\mathcal{O}$ are *equivalent* if there exists $\beta \in B^\times$ such that $J = I\beta$. Let $\{I_1, ..., I_n\}$ be a set of representatives of the equivalence classes of the left ideals of $\mathcal{O}$, setting $I_1 = \mathcal{O}$ by convention, and let *cl*$(\mathcal{O})$ denote the class group. The *class number* n is the same for every maximal order $\mathcal{O}$ in B. For every left ideal I_i, let $\mathcal{O}_R(I_i) := \{\beta \in B : I_i\beta \subset I_i\}$ be the right order of I_i. Then each conjugacy class of a maximal order of B is represented in the set $\{\mathcal{O}_R(I_1), ..., \mathcal{O}_R(I_n)\}$. The set $\mathit{\Gamma}_i := \mathcal{O}_R(I_i)^\times/\mathbb{Z}^\times$ is finite, as it is a discrete subgroup of the compact group $(B \otimes \mathbb{R})^\times/\mathbb{R}^\times \cong \mathrm{SO}_3(\mathbb{R})$. Let w_i be its cardinality.

We now introduce theta series and Brandt matrices, the main reference for this part being [25]. Let the inverse ideal of I_i be defined as $I_i^{-1} := \{\beta \in \mathcal{B} : I_i\beta I_i \subset I_i\}$, and $M_{ij} := I_j^{-1}I_i = \{\sum a_k b_k : a_k \in I_j^{-1}, b_k \in I_i\}$. Let $Nm(M_{ij})$ and $a \in Nm(M_{ij})$ have no common factors; the definition of *Brandt matrix* $B(m) := \left[B_{ij}(m)\right]_{1 \leq i,j \leq n}$ arises from the following definition of theta series θ_{ij}:

$$\theta_{ij}(\tau) := \frac{1}{2w_j} \sum_{a \in M_{ij}} e^{2\pi i \frac{Nm(a)}{Nm(M_{ij})}\tau} = \sum_{m \geq 0} B_{ij}(m)q^m,$$

where $q := e^{2\pi i\tau}$. The integer m is called the *degree* of $B(m)$; if $m = 0$, we have

$$B(0) = \frac{1}{2}\begin{bmatrix} \frac{1}{w_1} & \frac{1}{w_2} & \cdots & \frac{1}{w_n} \\ \frac{1}{w_1} & \frac{1}{w_2} & \cdots & \frac{1}{w_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{w_1} & \frac{1}{w_2} & \cdots & \frac{1}{w_n} \end{bmatrix},$$

while $B(1)$ is the identity matrix. We hereby list some properties of Brandt matrices, which are proved in [25, Prop. 2.7]:

1. If $m \geq 1$, then each entry $B_{ij}(m) \in \mathbb{N}$, and the sum of the entries in a row is independent of the chosen row, i.e. for all $1 \leq i \leq n$,
$$\sum_{j=1}^{n} B_{ij}(m) = \sum_{\substack{d \mid m \\ \gcd(d,p)=1}} d.$$
2. If m and m' are coprime, then $B(mm') = B(m)B(m')$.
3. If $\ell \neq p$ is a prime, then $B(\ell^k) = B(\ell^{k-1})B(\ell) - \ell B(\ell^{k-2})$ for all $k \geq 2$.
4. $w_j B_{ij}(m) = w_i B_{ji}(m)$ for all m and for all $1 \leq i, j \leq n$.

Recent advances in isogeny-based cryptography have put quaternion algebras under the spotlight, due to their intimate connection with supersingular elliptic curves via a correspondence of categories. The original result by Deuring [19] has been enriched by several later works, the last of which being [18]. In particular,

- a supersingular j-invariant over $\mathbb{F}_{p^2}$ corresponds to a maximal orders in the quaternion algebra $B_{p,\infty}$, and the set of supersingular j-invariants corresponds to the class group $cl(\mathcal{O})$;
- endomorphisms correspond to principal ideals;
- isogenies of degree ℓ correspond to left ideals of norm ℓ of maximal orders;
- composition of isogenies corresponds to multiplication of ideals

Given the further connection between quaternion algebras and Brandt matrices, we can exploit the latter to study elliptic curves. First and foremost, let us now introduce Hurwitz Class Numbers. Given an order $\mathcal{O}$ of rank 2 over $\mathbb{Z}$ of negative discriminant d, let $h(d)$ be the size of the class group $cl(\mathcal{O})$ and define $u(d) := \mathcal{O}^\times/\mathbb{Z}^\times = \mathcal{O}^\times/2$. For $D > 0$, the *Hurwitz Class Numbers* $H(D)$ is

$$H(D) = \sum_{d \cdot \mathrm{f}^2 = -D} \frac{h(d)}{u(d)}. \tag{1}$$

Let $\mathcal{O}_{-D}$ be *the* order of discriminant $-D$; we define the *modified Hurwitz Class Number* $H_p(D)$ as

$$H_p(D) := \begin{cases} 0 & \text{if } p \text{ splits in } \mathcal{O}_{-D}; \\ H(D) & \text{if } p \text{ is inert in } \mathcal{O}_{-D} \\ & \quad \text{and does not divide the conductor of } \mathcal{O}_{-D}; \\ \frac{1}{2}H(D) & \text{if } p \text{ is ramified in } \mathcal{O}_{-D} \\ & \quad \text{but does not divide the conductor of } \mathcal{O}_{-D}; \\ H(\frac{D}{p^2}) & \text{if } p \text{ divides the conductor of } \mathcal{O}_{-D}; \end{cases} \tag{2}$$

For $D = 0$, we set $H(0) = -1/12$ and $H_p(0) := \frac{p-1}{24}$.

Theorem 1. ***[25, Prop. 1.9]*** *For all $m \geq 0$,*

$$Tr(B(m)) = \sum_{\substack{s\in\mathbb{Z} \\ s^2\leq 4m}} H_p(4m - s^2).$$

Theorem 2. ***[26, Sec. 7]*** *For $m \in \mathbb{Z}$,*

$$\sum_{\substack{s\in\mathbb{Z} \\ s^2\leq 4m}} H(4m - s^2) = 2\sum_{d|m} d - \sum_{d|m} \min\{d, m/d\}.$$

3 The SIDH-based Identification Protocol and Its Special Soundness

In their seminal work [17], Jao, De Feo, and Plût proposed an identification scheme in the SIDH setting, which we will refer to as the SIDH-based identification protocol $\mathsf{ID}_{\mathsf{SIDH}}$. In this protocol, a prover persuades a verifier that, for a given pair (E_0, E_1) of supersingular elliptic curves, it knows a cyclic isogeny $\varphi : E_0 \to E_1$ - having degree $\ell_1^{e_1}$, where ℓ_1 is a prime - without revealing any information about the isogeny itself.

Later on, Yoo *et al.* [41] and Galbraith *et al.* [24] turned $\mathsf{ID}_{\mathsf{SIDH}}$ into the first SIDH-based digital signature schemes by applying the Fiat-Shamir transform [21] and the Unruh transform [36]. The Unruh transform provides security in the quantum random oracle, while the Fiat-Shamir transform generally guarantees security only in the random oracle model. The UF-CMA security of the resulting digital signatures is deduced from both the HVZK and special soundness properties of $\mathsf{ID}_{\mathsf{SIDH}}$. Proofs for the special soundness property of the protocol are given in [17, 24, 41], and all of them consider the same extractor, which we will denote by $\mathsf{Ex}_{\mathsf{SIDH}}$ in the following.

In this section, we detail two scenarios where the proposed extraction algorithm $\mathsf{Ex}_{\mathsf{SIDH}}$ fails to extract any meaningful witness for a statement $\mathsf{x} \in \mathcal{L}_{\mathcal{R}}$ when given two valid transcripts relative to x. The consequence of such failure is that, when $\mathsf{ID}_{\mathsf{SIDH}}$ is turned into a signature scheme via either the Fiat-Shamir or Unruh transform, unforgeability can no be longer argued. In the next sections, we will show some concrete examples of the two scenarios mentioned above, even for some of the SIKE parameters sets.

3.1 $\mathsf{ID}_{\mathsf{SIDH}}$ and $\mathsf{Ex}_{\mathsf{SIDH}}$

The public parameters of $\mathsf{ID}_{\mathsf{SIDH}}$ consists of a tuple $\mathsf{pp}_{\mathsf{SIDH}} = (\ell_1, \ell_2, e_1, e_2, f, p, E_0, P_1, Q_1, P_2, Q_2)$ where:

- ℓ_1 and ℓ_2 are two distinct small primes;
- e_1,e_2 and f are natural numbers, and $\ell_1^{e_1} \approx \ell_2^{e_2}$;

- p is a prime of the form $\ell_1^{e_1}\ell_2^{e_2}f \pm 1$;
- E_0 is a *starting* supersingular elliptic curve defined over the finite field $\mathbb{F}_{p^2}$;
- $\{P_1, Q_1\}$ and $\{P_2, Q_2\}$ are two bases for $E_0[\ell_1^{e_1}]$ and $E_0[\ell_2^{e_2}]$, respectively.

$\mathsf{ID}_{\mathsf{SIDH}}$ is an identification protocol for a binary relation $\mathcal{R}$ contained in $X \times Y$, where:

$$X = \{(E_1, P', Q') | E_1 \text{ sup. elliptic curve over } \mathbb{F}_{p^2} \;\wedge\; E_1[\ell_2^{e_2}] = \langle P', Q' \rangle\} \quad (3)$$

$$Y = \{\varphi \mid \varphi \text{ is a cyclic isogeny from } E_0 \text{ s.t. } \mathsf{deg}(\varphi) \in \{\ell_1^{e_1}, \ell_2^{2e_2}\ell_1^{e_1}\}\}. \quad (4)$$

The binary relation $\mathcal{R}$ is induced by the CSSI problem, and therefore it is defined as follows:

$$\mathcal{R} = \{((E_1, P', Q'), \varphi) \mid \varphi : E_0 \to E_1 \wedge \mathsf{deg}(\varphi) = \ell_1^{e_1} \wedge \varphi(P_2) = P', \varphi(Q_2) = Q'\}. \quad (5)$$

Here we note that each statement $\mathsf{x} \in \mathcal{L}_{\mathcal{R}}$ admits a unique witness w (see [38]). $\mathsf{ID}_{\mathsf{SIDH}}$ is a binary-challenge identification protocol, i.e. $\mathsf{ChSet} = \{0, 1\}$, which constits of three algorithms $((\mathsf{P}_1, \mathsf{P}_2), \mathsf{V})$ defined as:

- $\mathsf{com} \leftarrow \mathsf{P}_1((E_1, P', Q'), \varphi)$: on input $(E_1, P', Q'), \varphi) \in \mathcal{R}$, it generates two random integers m_2, n_2 in $\mathbb{Z}_{\ell_2^{e_2}}$, not both divisible by ℓ_2, and computes the point $R = [m_2]P_2 + [n_B]Q_2$ - of order $\ell_2^{e_2}$ - and the elliptic curves $E_2 = E_0/\langle R \rangle$ and $E_3 = E_1/\langle [m_2]P' + [n_2]Q' \rangle$. Then it outputs the commitment $\mathsf{com} = (E_2, E_3)$.
- $\mathsf{resp} \leftarrow \mathsf{P}_2((E_1, P', Q'), \varphi, \mathsf{ch})$: if the challenge ch is equal to 0, the algorithm sets resp to the pair (m_B, n_B). Otherwise, given the point S - of order $\ell_1^{e_1}$ - generating $\mathsf{Ker}(\varphi)$ and by the isogeny $\phi : E_0 \to E_2$ with $R = [m_2]P_2 + [n_B]Q_2$ as kernel, it sets resp to $\phi(S)$. It then returns the response resp.
- $\{0, 1\} \ni b \leftarrow \mathsf{V}((E_1, P', Q'), \mathsf{com}, \mathsf{ch}, \mathsf{resp})$: the deterministic verification algorithm checks whether
  - $(\mathsf{ch} = 0)$ given $\mathsf{resp} = (m_2, n_2)$, m_2 and n_2 are not both divisible by ℓ_2, $E_0/\langle [m_2]P_2 + [n_2]Q_2 \rangle$ is isomorphic to E_2, and $E_1/\langle [m_2]P' + [n_2]Q' \rangle$ is isomorphic to E_3;
  - $(\mathsf{ch} = 1)$ given $\mathsf{resp} = T$, $T \in E_2$ has order $\ell_1^{e_1}$ and $E_2/\langle T \rangle$ is isomorphic to E_3.

  If the conditions are fulfilled, V outputs 1 (accept), otherwise outputs 0 (reject).

It is easy to show that $\mathsf{ID}_{\mathsf{SIDH}}$ is correct. We here show that it satisfies the Honest-Verifier Zero-Knowledge property by constructing a zero-knowledge simulator Sim as follows. On input $(E_1, P', Q') \in \mathcal{L}_{\mathcal{R}}$ and $\mathsf{ch} = 0$, Sim chooses two random integers m_2, n_2 in $\mathbb{Z}_{\ell_2^{e_2}}$, not both divisible by ℓ_2, computes the point $R = [m_2]P_2 + [n_2]Q_2$ of order $\ell_2^{e_2}$, and outputs

$$(\mathsf{com} = (E_0/\langle R \rangle, E_1/\langle [m_2]P_2 + [n_2]Q_2 \rangle), \mathsf{ch} = 0, \mathsf{resp} = (m_2, n_2)).$$

This simulated transcript is distributed exactly as a real one conditioned on $\mathsf{ch} = 0$. In order to simulate a transcript on input $(E_1, P', Q') \in \mathcal{L}_{\mathcal{R}}$ and $\mathsf{ch} = 1$,

Sim chooses a random supersingular elliptic curve E_2, defined over $\mathbb{F}_{p^2}$ and with the same number of rational points as E_0. Then it chooses a random cyclic subgroup $\langle T \rangle \subset E_2[\ell_1^{e_1}]$ having order $\ell_1^{e_1}$, and outputs

$$(\mathsf{com} = (E_2, E_3 = E_2/\langle T \rangle), \mathsf{ch} = 1, \mathsf{resp} = T).$$

This is computationally indistinguishable from a valid transcript conditioned on $\mathsf{ch} = 1$, under the assumption that the $\mathrm{DSSP}_{\mathsf{pp}}$ problem (Problem 4) is hard.

In [17,24,41], the special soundness property of $\mathsf{ID}_{\mathsf{SIDH}}$ is proven by considering the same extractor $\mathsf{Ex}_{\mathsf{SIDH}}$, defined as follows. Given two valid transcripts $(\mathsf{x}, \mathsf{com}, 0, (m_2, n_2))$ and $(\mathsf{x}, \mathsf{com}, 1, T)$ - relative to the statement $\mathsf{x} = (E_1, P', Q')$, and with the same commitment $\mathsf{com} = (E_1, E_2)$ and different challenges - it outputs $\hat{\phi}(\langle T \rangle)$, where ϕ is the isogeny from E_0 with $\langle [m_2]P_2 + [n_2]Q_2 \rangle$ as kernel. We note that, under the assumption that $\langle T \rangle$ is equal to $\phi(\mathsf{Ker}(\varphi))$ - where φ is the only witness for x - $\mathsf{Ex}_{\mathsf{SIDH}}$ extracts exactly $\mathsf{Ker}(\varphi)$ (and so, equivalently, φ). Despite this assumption being made in [17,24,41], it does not appear necessary to make it in general. In the following section we detail two different scenarios where such assumption is not valid.

3.2 Scenario 1 - Single Collision

Let $\mathsf{x} = (E_1, P', Q')$ be a statement in $\mathcal{L}_{\mathcal{R}}$. Suppose there exists a point $R = [m_2]P_2 + [n_2]Q_2 \in E_0$, of order $\ell_2^{e_2}$, which defines two isogenies, $\phi : E_0 \to E_2 = E_0/\langle R \rangle$ and $\phi' = E_1 \to E_3 = E_1/\langle [m_2]P' + [n_2]Q' \rangle$, such that E_2 admits two distinct cyclic subgroups G and $\tilde{G}$ of order $\ell_1^{e_1}$ that generate two isogenies ψ, ψ' going from E_2 to E_3. The pair (ψ, ψ') forms a collision (of length e_1) in the isogeny graph $\mathcal{G}_{p^2}(\ell_1)$ We call *single collision* (of length e_1) a collision of this form. Here we assume $G = \phi(\mathsf{Ker}(\varphi))$ - where φ is the only witness for x - and we denote by T a generator for $\tilde{G}$.

Given the commitment $\mathsf{com} = (E_2, E_3)$, both $(\mathsf{x}, \mathsf{com}, \mathsf{ch} = 0, (m_2, n_2))$ and $(\mathsf{x}, \mathsf{com}, \mathsf{ch} = 1, T)$ are valid transcripts relative to x. If no extra hypotheses are made, on input such two transcripts, $\mathsf{Ex}_{\mathsf{SIDH}}$ extracts a cyclic subgroup of E_0, having order $\ell_1^{e_1}$, which defines an isogeny whose codomain E' is not $\mathbb{F}_{p^2}$-isomorphic to E_1. Therefore, $\mathsf{Ex}_{\mathsf{SIDH}}$ fails in extracting a witness for x. In order to better explain the above scenario and to show that it can actually happen in practice, we provide a concrete occurrence of such scenario.

Example 1. Consider the prime $p = (2^8)(3^5) - 1$ and the irreducible polynomial $f = x^2 + 62205x + 5 \in \mathbb{F}_p[x]$. Given $\mathbb{F}_{p^2} = \mathbb{F}_p[x]/(f)$, we denote by z a root of f, which forms a basis for $\mathbb{F}_{p^2}$ together with the identity 1. Consider as E_0 a supersingular elliptic curve with j-invariant equal to 22470. The two points $P_2 = (7077z + 32228, 17988z + 60777)$, $Q_2 = (51235z + 37453, 42878z + 1379)$ form a basis for $E_0[5^3]$. Let E_1 be the image curve of the 2^8-isogeny φ having $\langle 33446z + 46615, 52617z + 4750 \rangle$ as kernel, for which it holds $j(E_1) = 37167z + 53117$. We consider the tuple $(E_1, P' = (6505z + 32827, 20825z + 21686), Q' = (59525z + 48254, 52332z + 7163)$ as statement x. Then, for the curves $E_2 =$

$$\begin{tikzcd}
E_0 \arrow[r, "\varphi"] \arrow[d, "\phi"'] & E_1 \arrow[d, "\phi'"] \\
E_2 \arrow[r, bend left, "\psi"] \arrow[r, bend right, "\psi'"'] & E_3
\end{tikzcd}$$

Fig. 1: Scenario 1.

$E_0/\left\langle [161]P_2 + [183]Q_2 \right\rangle$, $E_3 = E_1/\left\langle [161]P' + [183]Q' \right\rangle$, there exist two distinct cyclic isogenies $\psi, \psi' : E_2 \to E_3$ of degree 2^8. ψ is the one with kernel equal to $\phi(\mathsf{Ker}(\varphi))$, where ϕ is the isogeny with kernel $\langle [161]P_2 + [183]Q_2 \rangle$. On the other hand, ψ' has kernel generated by $T = (52195z + 35063, 51186z + 33135)$, with $T \notin \phi(\mathsf{Ker}(\varphi))$. As a consequence, $\hat{\phi}(\langle T \rangle)$ defines an isogeny whose image has j-invariant equal to $55144z + 45927$, which is different from that of E_1.

An exhaustive approach has been used to produce concrete examples for Scenario 1. In particular, the simple procedure we executed goes as follows:

1. produce a prime p of the form $\ell_1^{e_1}\ell_2^{e_2}f \pm 1$, so that $\ell_1^{e_1} \approx \ell_2^{e_2}$;
2. for each vertex j_0 in the isogeny graph $\mathcal{G}_{p^2}(\ell_2)$, compute all the paths composed by e_2 steps (with no backtracking) and the corresponding arriving vertices;
3. for each arriving vertex j_2, compare all paths in the isogeny graph $\mathcal{G}_{p^2}(\ell_1)$ which originate from j_2 and composed of e_1 steps (with no backtracking);
4. an occurrence of Scenario 1 is found whenever two paths end at the same j-invariant j_3, and they are distinct in at least one step.

3.3 Scenario 2 - Double Collisions

Let $\mathsf{x} = (E_1, P', Q')$ be a statement in $\mathcal{L}_{\mathcal{R}}$. Suppose $\mathcal{L}_{\mathcal{R}}$ contains another statement $\tilde{\mathsf{x}}$, with $\tilde{\mathsf{x}} \neq \mathsf{x}$, with the same first component (modulo $\mathbb{F}_{p^2}$-isomorphisms). In other words, there exist two distinct cyclic subgroups of order $\ell_1^{e_1}$ in E_0, H and $\tilde{H}$, such that $j(E_0/H) = j(E_0/\tilde{H})$. We denote by φ and $\tilde{\varphi}$ the isogenies having kernels H and $\tilde{H}$, respectively. We further assume there exists a point $R = [m_2]P_2 + [n_2]Q_2 \in E_0$, of order $\ell_2^{e_2}$, such that $E_1/\left\langle \varphi(R) \right\rangle$ has the same j-invariant of $E_1/\left\langle \tilde{\varphi}(R) \right\rangle$. We denote by ϕ' and $\tilde{\phi}'$ the isogenies having $\langle \varphi(R) \rangle$ and $\langle \tilde{\varphi}(R) \rangle$ as kernels, respectively. The pairs $(\varphi, \tilde{\varphi})$ and $(\phi', \tilde{\phi}')$ are two collisions (of length e_1 and e_2, respectively) in the isogeny graphs $\mathcal{G}_{p^2}(\ell_1)$ and $\mathcal{G}_{p^2}(\ell_2)$, respectively. Since the second collision *originates* from the same point $R \in E_0$,

the two collisions have a tight link. For this reason we call *double collision* (of length $e_1 + e_2$) a pair of collisions of this form.
Given the commitment $\mathsf{com} = (E_2 = E_0/\langle R\rangle, E_3 = E_1/\langle\varphi(R)\rangle)$, the two transcripts $(\mathsf{x}, \mathsf{com}, \mathsf{ch} = 0, (m_2, n_2))$ and $(\mathsf{x}, \mathsf{com}, \mathsf{ch} = 1, \phi(\tilde{S}))$ - where $\tilde{S}$ generates $\tilde{H}$ - are both valid transcripts relative to x. However, on input such transcripts, $\mathsf{Ex}_{\mathsf{SIDH}}$ extracts a witness $\tilde{\mathsf{w}}$ for $\tilde{\mathsf{x}}$, rather than a witness for x.

Also for this scenario we provide a concrete counterexample.

Example 2. Let $p = (2^8)(5^3) + 1$, and $f = x^2 + 7 \in \mathbb{F}_p[x]$ be an irreducible polynomial. We denote the primitive element $\sqrt{-7}$ by z. Then we consider as E_0 a supersingular elliptic curve over $\mathbb{F}_{p^2}$ such that $j(E_0) = 80630z + 38195$. The two distinct cyclic subgroups of order 2^8 of E_0

$$H = \langle(174423z + 15317, 139167z + 27752)\rangle$$
$$\tilde{H} = \langle(279804z + 121600, 104494z + 307794)\rangle,$$

are such that $j(E_0/H) = j(E_0/\tilde{H}) = 255209z + 212204$. We denote by φ and $\tilde{\varphi}$ the isogenies with kernels H and $\tilde{H}$, respectively. Let E_1 and $\tilde{E_1}$ the images of these two isogenies. The point $R = (290744z + 184866, 22597z + 44859) \in E_0$ - having order 3^5 - is such that $E_1/\langle\varphi(R)\rangle$ and $\tilde{E_1}/\langle\tilde{\varphi}(R)\rangle$ have the same j-invariant.

Producing examples where this Scenario 2 occurs is much more difficult than it is for Scenario 1. The procedure we followed relies on two main parts. The first one consists in an exhaustive search for collisions of $\ell_1^{e_1}$-isogenies (i.e. single collisions of length e_1), and proceeds as follows:

1. produce a prime p of the form $\ell_1^{e_1}\ell_2^{e_2}f \pm 1$, so that $\ell_1^{e_1} \approx \ell_2^{e_2}$;
2. for each vertex j_0 in the isogeny graph $\mathcal{G}_{p^2}(\ell_1)$, compare all paths composed of e_1 steps (with no backtracking) and originating from the fixed vertex;
3. a single collision is found whenever two paths end at the same j-invariant (which we refer to as *colliding j-invariant*), and they are distinct in at least one step.

After a colliding j-invariant j_1 is found for a starting vertex j_0, the second part of the procedure takes the two colliding paths (which correspond to the isogenies φ and $\tilde{\varphi}$, respectively) and continues by

- constructing a supersingular elliptic curve E_0, defined over $\mathbb{F}_{p^2}$ having j_0 as j-invariant;
- comparing all pair of paths of e_2 steps in $\mathcal{G}_{p^2}(\ell_2)$ determined by the points $\varphi(R)$ and $\tilde{\varphi}(R)$, for any $R \in E_0[\ell_2^{e_2}]$ of order $\ell_2^{e_2}$.
- a collision is found whenever the arriving j-invariant is the same for both paths in a pair. In this case, the two paths are said to be *colliding*, like the arriving j-invariant.

When the two colliding paths are equal in each step, we call the whole double collision a *Florence flask*; when the two paths differ in at least one step, we call the double collision a *lemniscate*.

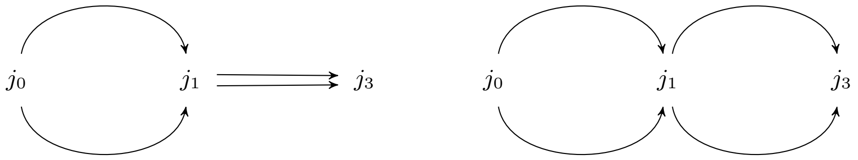

Fig. 2: A Florence flask (on the left) and a lemniscate (on the right). Recall that the paths in each collision are required to differ in at least one step, and not necessarily in all as in the above figures.

We ran some experiments for primes of small size, always setting $\ell_1 = 2$ and $\ell_2 = 3$. The results are summarized in Section 3.3, where we list the number of vertices of $\mathcal{G}_{p^2}(2)$ from which at least a single collision originates, the total number of single collisions in the graph, and the total number of double collisions (distinguishing between lemniscates and Florence Flasks). The exhaustive search conducted in both parts of the procedure becomes very expensive already with 12-bit primes and thus, from $p = 2591$ on, we restricted our analysis to some random vertices. The results show that, except for $p = 1297$, each prime we considered exhibits at least one single collision for each supersingular j-invariant taken as initial vertex.

p	**eA**	**eB**	**+1/-1**	**Initial j-invariants with collisions**	**Single collisions**	**Lemniscates**	**Florence flasks**
431	4	3	-1	37/37	183	134	286
433	4	3	+1	36/36	213	229	152
863	5	3	-1	73/73	681	246	316
1297	4	4	+1	97/108	194	127	231
2591	5	4	-1	25/25	121	44	44
2593	5	4	+1	25/25	112	77	85
15551	6	5	-1	25/25	76	16	84
62207	8	3	-1	20/20	280	14	405

Table 1: A summary of the number of single and double collisions when considering small primes p.

3.4 Scenario 1 and SIKE parameters sets

At a first glance, one might argue that the concrete examples provided in the previous subsections exist only because of the small size of the considered graphs. In order to argue that we cannot exclude the presence of similar examples also in supersingular isogeny graphs of cryptographic size, we focus on Scenario 1 for the *largest* SIKE paramaters set, called SIKEp751. This name refers to the fact that the underlying prime, denoted by p751, has size (in bits) equal to 751.

As we detail in the Appendix, the supersingular isogeny graph $\mathcal{G}_{(p751)^2}(3)$ admits two single collisions, of length 239, whose starting vertex is the (class of the) supersingular elliptic curve chosen as starting curve for SIKEp751. As the defining equation of this curve is $y^2 = x^3 + 6x^2 + x$, such curve is usually denoted by E_6. Consider one of the two collisions, and call H and $\tilde{H}$ the distinct kernels of the two colliding isogenies ψ and ψ', respectively. Let E_3 be the image curve (modulo isomorphisms) of the two isogenies.

$$\begin{tikzcd}
E_0 \arrow[r, "\varphi"] & E_1 \arrow[d, "\phi'"] \\
E_6 \arrow[u, "\overline{\phi}"] \arrow[r, bend left, "\psi"] \arrow[r, bend right, "\psi'"'] & E_3
\end{tikzcd}$$

Fig. 3: An example for Scenario 1 from a single collision originating from E_6 (the dual isogeny of $\overline{\phi}$ is ϕ).

An example of Scenario 1 can then be constructed as follows. A cyclic subgroup $K \subset E_6$ of order 2^{372} is randomly sampled, and the image curve of the isogeny $\overline{\phi}$ having K as kernel is denoted by E_0. Then, the obtained curve E_0 is set as the starting curve of $\mathsf{ID}_{\mathsf{SIDH}}$. Let $\varphi : E_0 \longrightarrow E_1$ be the isogeny with kernel $\overline{\phi}(H)$. By fixing some bases $\{P_1, Q_1\}$, $\{P_2, Q_2\}$ for $E_0[3^{239}]$ and $E_1[2^{372}]$, respectively, E_1 and φ can be turned into a statement-witness pair for $\mathsf{ID}_{\mathsf{SIDH}}$ with public parameters

$$\mathsf{pp}_{\mathsf{SIDH}} = (\ell_1 = 3, \ell_2 = 2, e_1 = 239, e_2 = 372, f = 1, p751, E_0, P_1, Q_1, P_2, Q_2).$$

In particular, the statement x is set to $(E_1, \varphi(P_2), \varphi(Q_2))$, while its corresponding witness w is set to φ. Let R be a generator of the kernel of the dual isogeny of $\overline{\phi}$, with $m_2, n_2 \in \mathbb{Z}/2^{372}\mathbb{Z}$ such that $R = [m_2]P_2 + [n_2]Q_2$, and T a generator for $\tilde{H}$. We underline that T is the generator of the kernel of ψ'. By setting $\mathsf{com} = (E_6, E_3)$, the two transcripts $(\mathsf{x}, \mathsf{com}, \mathsf{ch} = 0, (m_2, n_2))$ and $(\mathsf{x}, \mathsf{com}, \mathsf{ch} = 1, T)$ are both valid relative to x, and are an example of Scenario 1.

Remark 1. The constructed example could be deemed as not an example of Scenario 1 for SIKEp751, as the initial curve E_0 is not the prescribed one. However, if we took the same parameter set and we let the initial curve vary, this example would be perfectly acceptable. Moreover, we note that constructing an example of Scenario 1 needs a single collision of length 239 in the graph $\mathcal{G}_{(p751)^2}(3)$ (or of the length 372 in the graph $\mathcal{G}_{(p751)^2}(3)$). It appears that the only way to compute such a collision is to exploit the knowledge of the endomorphism ring of the starting vertex (as done in the Appendix by extending the procedure, based on quaternion algebras, presented in [32]). As a consequence, it seems prohibitive to obtain a single collision starting from a random vertex E_1. We believe that considering a starting curve for $\mathsf{ID}_{\mathsf{SIDH}}$ other than E_6 does not diminish the relevance of the constructed example.

4 Security implications of the two exception scenarios

In [16], De Feo *et al.* analyse the SIDH-based identification protocol $\mathsf{ID}_{\mathsf{SIDH}}$ and the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$. They show that two valid transcripts relative to a statement $\mathsf{x} \notin \mathcal{L}_{\mathcal{R}}$ can be easily produced. When considering Definition 3b for the special soundness property, this means that $\mathsf{ID}_{\mathsf{SIDH}}$ cannot be proven to enjoy such property. Therefore, they propose a modified version of $\mathsf{ID}_{\mathsf{SIDH}}$, which they prove to achieve special soundness. In such modification, P_1 appends to a commitment (E_2, E_3) also the images - trough $\psi : E_2 \to E_3$ - of a basis for $E_2[\ell_2^{e_2}]$. Such points are then involved in the verification phase for both possible challenges. Therefore, the proposed modification is less efficient in terms of computation and bandwidth with respect to the original identification scheme $\mathsf{ID}_{\mathsf{SIDH}}$.

The counterexamples provided in [16] do not affect the special soundness property of $\mathsf{ID}_{\mathsf{SIDH}}$ when considering Definition 3a. As discussed in Section 2.1, such definition can be safely considered when exploiting an identification protocol for a Fiat-Shamir digital signature scheme. As a consequence, according to the special soundness proofs provided in [17,24,41], $\mathsf{ID}_{\mathsf{SIDH}}$ could still be taken as a building block to construct Fiat-Shamir SIDH-based digital signature schemes. Moreover, not only could one use $\mathsf{ID}_{\mathsf{SIDH}}$, but this would be the preferable choice in terms of efficiency, given the lower efficiency of the modified scheme presented in [16].

However, the two exception scenarios and concrete examples described in the previous section show, for the first time, that $\mathsf{Ex}_{\mathsf{SIDH}}$ does not provide special soundness to $\mathsf{ID}_{\mathsf{SIDH}}$, even when considering Definition 3a. Therefore, the original protocol cannot be used for secure Fiat-Shamir digital signatures as it stands. In the following subsections we analyse two possible remedies to this state of affairs.

4.1 Relaxed Special Soundness

Given the sets X and Y introduced in Section 3.1, a binary relation $\tilde{\mathcal{R}}$ alternative to $\mathcal{R}$ can be considered. Such relation is induced by Problem 1, and therefore it is defined as follows:

$$\tilde{\mathcal{R}} = \{((E_1, P', Q'), \varphi) \mid \varphi : E_0 \to E_1\}. \tag{6}$$

It is clear that $\mathcal{R}$ is contained in $\tilde{\mathcal{R}}$. Furthermore, it can be shown that $\mathsf{ID}_{\mathsf{SIDH}}$ enjoys relaxed special soundness with respect to $\mathcal{R} \subset \tilde{\mathcal{R}}$.

Let $(\mathsf{x}, \mathsf{com}, \mathsf{ch}, \mathsf{resp})$, $(\mathsf{x}, \mathsf{com}, \mathsf{ch}', \mathsf{resp}')$ be two valid transcripts, relative to $\mathsf{x} \in \mathcal{L}_{\mathcal{R}}$, having the same commitment and different challenges. If $\mathsf{x} = (E_1, P', Q')$ and $\mathsf{com} = (E_2, E_3)$, the verification phase for the two challenges guarantees the knowledge of two isogenies $\phi : E_0 \to E_2$, $\phi' : E_1 \to E_3$, and of an isogeny $\psi : E_2 \to E_3$, respectively. Therefore, $\hat{\phi}' \circ \psi \circ \phi$ is an isogeny from E_0 to E_1. An extractor $\mathsf{rEx}_{\mathsf{SIDH}}$ outputting $\hat{\phi}(\mathsf{Ker}(\psi))$ or $\mathsf{Ker}(\hat{\phi}' \circ \psi \circ \phi)$ if $\hat{\phi}(\mathsf{Ker}(\psi))$ does not define an isogeny from E_0 to E_1 makes $\mathsf{ID}_{\mathsf{SIDH}}$ satisfy the relaxed special soundness property.

As we already observed, all the proofs for the Fiat-Shamir and Unruh transform, which were formally given for identification protocols satisfying special soundness, naturally extend to identification protocols that enjoy relaxed special soundness. Namely, the reduction algorithms used in the security proofs of the signature schemes based on the two kinds of identification protocols are equivalent but in the extraction of the witness. The only difference between the two cases is what we can argue about the witness extracted from two valid transcripts; in the former case, a witness for the original relation $\mathcal{R}$ is extracted, while in the latter a witness for the relaxed relation $\tilde{\mathcal{R}}$ is extracted.As long as the original relation and the relaxed relation are equally hard, which is the case for our isogeny relations (see Section 2.2), then all the proofs hold.

4.2 Alternative extractors

The two scenarios described in Section 3, which make the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$ fail, invalidate the proofs for the special soundness property (Definition 3.a) of $\mathsf{ID}_{\mathsf{SIDH}}$ based on its extractor. However, the possibility to construct an alternative special-soundness extractor for $\mathsf{ID}_{\mathsf{SIDH}}$ is not ruled out by the two scenarios.

Such an alternative extractor could be designed by taking into consideration the reduction proposed by Galbraith *et al.* in [23, Sec 4.1]. This reduction runs in polynomial time, and given an isogeny between two supersingular elliptic curves E_0 and E_1, it returns a shorter isogeny between them under the hypothesis that the endomorphism ring of E_0 is known. This tool could be used to retrieve an isogeny of degree $\ell_1^{e_1}$ from the isogeny $\widehat{\phi'} \circ \psi \circ \phi$ to . However, the reduction is probabilistic, and thus it leaves open the question whether there exist examples on which the reduction fails.

In a recent work [22], Fouotsa, Kutas and Merz presented a generalisation of the above reduction. In particular, their algorithm allows to compute a secret isogeny φ of degree N_1 between two supersingular elliptic curves E_0, E_1 given the endomorphism rings $\mathrm{End}(E_0)$, $\mathrm{End}(E_1)$ and the action of φ on a basis $\{P_1, Q_1\}$ of the torsion subgroup $E_0[N_1]$. The main idea of the new reduction is the following. Since isogenies from E_0 to E_1 form a $\mathbb{Z}$-module M of rank 4, the KLPT algorithm [29] can be used to compute a basis of M. One can then turn it in a LLL-reduced basis [30] $\{\psi_1, \psi_2, \psi_4, \psi_4\}$ and evaluate the images of $\{P_1, Q_1\}$ via these maps. Computing φ is then reduced to solving a system of linear equations to retrieve the coefficients $x_1, \dots x_4$ for which $\varphi = x_1\psi_1 + x_2\psi_2 + x_3\psi_3 + x_4\psi_4$.

As argued in [22], this solution is unique when $x_i < \frac{N_2}{2}$ and $N_1 < \frac{dN_2}{16}$, where d is the degree of the shortest isogeny between E_0 and E_1. Note that, for balanced SIDH parameters of cryptographic interest, the ratio N_1/N_2 is approximately 1, and thus this procedure works whenever the degree of the shortest isogeny between the two curves is greater than 16. Also in this case, however, the question whether there exist examples on which the algorithm fails remains open. We defer to future work a further investigation of the two above reductions, aiming at constructing an alternative special-soundness extractor for $\mathsf{ID}_{\mathsf{SIDH}}$.

5 Quantitative study of cycles in Isogeny Graphs

As detailed in Section 3.2 and Section 3.3, there are scenarios in which the existence of collisions in supersingular isogeny graphs poses a problem to the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$ previously considered to prove the special soundness property of the SIDH-based identification protocol $\mathsf{ID}_{\mathsf{SIDH}}$. We therefore are interested in quantifying these collisions and providing bounds for them.

In the rest of the section, p will be be a large prime, ℓ a small one (usually 2 or 3) and we will denote by $\mathcal{G}_{p^2}(\ell)$ the supersingular isogeny graph whose vertices correspond to supersingular elliptic curves over $\mathbb{F}_{p^2}$ - modulo isomorphisms over $\mathbb{F}_{p^2}$ - and edges correspond to isogenies of degree ℓ (up to equivalence).

We are concerned with answering two main questions about supersingular isogeny graphs, which relate to the existence of single collisions and double collisions. The problem of the existence of simple collisions is addressed in Section 5.1, while in Section 5.2 we investigate the existence of double collisions.

5.1 Simple Collisions

We begin by formally stating the first problem we focus our attention on.

Problem 5. *Fix primes $\ell < p$, and let E_0 be a uniformly random supersingular elliptic curve E_0 (i.e. a vertex) in the graph $\mathcal{G}_{p^2}(\ell)$. Determine the probability that another supersingular elliptic curve E_1 in $\mathcal{G}_{p^2}(\ell)$ is linked to E_0 by two non-equivalent isogenies φ_A and φ_B, both with cyclic kerkel and of degree* $\deg(\varphi_i) = \ell^e$ *for a fixed $e \in \mathbb{N}$, with*[4] *$p \approx \ell^{2e}$, as in Figure 4.*

The answer to this problem depends on the primes ℓ and p. It is impossible to find an accurate estimate that only depends on ℓ and works for all elliptic curves. However, we will see that we can find an upper bound for the average number of collisions that only depends on p. Finding a *good* lower bound is much more difficult. It is beyond the reach of current analytic number theory, as it involves bounds on [27, Chap. 12] that are tighter than what we expect to be provable.

[4] A lot of our work does not use this assumption. However, we do need it at the end, when giving final bounds in terms of p. Also, in Section 5.1.2, we need $p > 2\ell^e$.

$$\begin{tikzcd} E_0 \arrow[r, bend left, "\varphi_A"] \arrow[r, bend right, "\varphi_B"'] & E_1 \end{tikzcd}$$

Fig. 4: Isogenies of the same degree sending E_0 in another curve $E_1 \in \mathcal{G}_{p^2}(\ell)$.

In order to tackle Problem 5, we start by computing the expected number of pairs of same-degree isogenies mapping E_0 to some E_1, where the expectation is taken over the parameters p and E_0.

Let f and g be two isogenies with the same domain $E_0 \in \mathcal{G}_{p^2}(\ell)$. We say that $f \sim g$ if they generate the same path in $\mathcal{G}_{p^2}(\ell)$, i.e. $f = \alpha \circ g$ for some automorphism α. Given two points in $E_0[\ell^e]$, we will consider them equivalent if they generate the same torsion subgroup, i.e. if they generate the same isogeny (modulo equivalence). Let $E_0[\ell^e]_{\max}$ denote the set of point of maximal order in $E_0[\ell^e]/_\sim$. We can rephrase the above statement as follows: compute the probability of finding two non-equivalent points $P_A, P_B \in E_0[\ell^e]$ of order ℓ^e such that $E_0/\langle P_A\rangle \cong E_0/\langle P_B\rangle$.

For primes p, ℓ and some $m \in \mathbb{N}$, let $\mathscr{C}_E(\ell^m)$ denote the number of endomorphisms with no backtracking (i.e. with cyclic kernel) of *degree* ℓ^m of a vertex E in $\mathcal{G}_{p^2}(\ell)$. Let also $\mathrm{Coll}_{\ell^e}(E_0)$ be the set of collisions (f, g) of degree ℓ^e starting at E_0, such that $f \not\sim g$. Following the above discussion, we can rewrite the set $\mathrm{Coll}_{\ell^e}(E_0)$ as

$$\left\{(P_A, P_B) \in \left(E_0[\ell^e]_{\max}\right)^2 : \frac{E_0}{\langle P_A\rangle} \cong \frac{E_0}{\langle P_B\rangle}, P_A \not\sim P_B\right\} \Big/ \left((P_A, P_B) \sim (P_B, P_A)\right). \tag{7}$$

Remark 1. In Equation (7), we mod out by the relation $(P_A, P_B) \sim (P_B, P_A)$ because we regard the two pairs of isogenies (f, g) and (g, f) as the same, since they clearly generate the same collisions. Note that k isogenies between E and E' would be counted as $\binom{k}{2}$ pairs.

Let n be the number of vertices in the graph $\mathcal{G}_{p^2}(\ell)$. By definition, to obtain $\mathscr{C}_{E_i}(\ell^{2e})$, we need to subtract from $B_{ii}(\ell^{2e})$ the number of endomorphisms of E_i of degree ℓ^{2e} with scalar factors. However, this number is simply given by $B_{ii}(\ell^{2e-2})$. Indeed, an endomorphism of E_i of degree ℓ^{2e} with backtracking can be written as $[\ell] \circ f$ for an endomorphism f of E_i of degree ℓ^{2e-2}. Thus, we can write

$$\begin{aligned}\sum_{i=1}^{n}\mathscr{C}_{E_i}(\ell^{2e}) &= \frac{1}{2}\sum_{i=1}^{n}\left(B_{ii}(\ell^{2e}) - B_{ii}(\ell^{2e-2})\right)\\ &= \frac{1}{2}\mathrm{Tr}\,\left(B(\ell^{2e})\right) - \frac{1}{2}\mathrm{Tr}\,\left(B(\ell^{2e-2})\right)\\ &= \frac{1}{2}\sum_{s^2\leq 4\ell^{2e}} H_p(4\ell^{2e}-s^2) - \frac{1}{2}\sum_{s^2\leq 4\ell^{2e-2}} H_p(4\ell^{2e-2}-s^2).\end{aligned} \tag{8}$$

Note that we divided by 2 because every endomorphism on the graph $\mathcal{G}_{p^2}(\ell)$ corresponds to a path that can be taken in both directions, since the graph is undirected. In more precise terms, each endomorphism corresponds to two pairs of equivalent (as described in Remark 1) isogenies $(f,g) \sim (g,f)$ that both give us the same collision: the endomorphisms $\hat{f}\circ g$ and $\hat{g}\circ f$ that they form are dual to each other.

In addition, $H_p(0)$ appears in both sums, so we can cancel it out. Moreover, since $4\ell^{2e}$ is a square in $\mathbb{Z}$, p will always split in $\mathcal{O}_{-4\ell^{2e}}$, and $H_p(4\ell^{2e}) = 0$. Likewise, $H_p(4\ell^{2e-2}) = 0$. Hence,

$$\sum_{i=1}^{n}\mathscr{C}_{E_i}(\ell^{2e}) = \frac{1}{2}\sum_{0<s^2<4\ell^{2e}} H_p(4\ell^{2e}-s^2) - \frac{1}{2}\sum_{0<s^2<4\ell^{2e-2}} H_p(4\ell^{2e-2}-s^2). \tag{9}$$

Remark 2. As an aside, note that the above number $\sum_{i=1}^{n}\mathscr{C}_{E_i}(\ell^{2e})$ is not exactly the total number of cycles (endomorphisms with no backtracking) of degree ℓ^{2e} in $\mathcal{G}_{p^2}(\ell)$, as we are over-counting each cycle in the above sum. If we wanted to count the *exact* total number of cycles of length $2e$ in $\mathcal{G}_{p^2}(\ell)$ we would have to divide by $2e$ to get $\frac{1}{2e}\sum_{i=1}^{n}\mathscr{C}_{E_i}$.

We have now related the number of endomorphisms with no backtracking $\mathscr{C}_E(\ell^m)$ to the desired number of collisions $\mathrm{Coll}_{\ell^e}(E_0)$. This allows us to turn our attention to estimating $\mathscr{C}_E(\ell^m)$.

Lemma 1. *It is possible to bound the number $\mathscr{C}_E(\ell^r)$ using $Coll_{\ell^r}(E_0)$ as follows*

$$\mathscr{C}_{E_0}(\ell^r) \leq \#Coll_{\ell^r}(E_0) \leq \mathscr{C}_{E_0}(\ell^r) + \sum_{i=1}^{r-1}\mathscr{C}_{E_0}(\ell^i)(\ell-1)\ell^{r-1-i}. \tag{10}$$

Proof. The first inequality is clear since, by definition, every collision also constitutes and endomorphism with no backtracking. However, there are some endomorphisms with backtracking that still constitute collisions, as in Figure 5. The collisions that form endomorphisms *with* backtracking can be broken into two parts: a first part that is an endomorphism without backtracking and second part at the end that constitutes the backtracking. In Figure 5, the first part is the part drawn in black between E and E_0 that gives an endomorphism without backtracking ("abcdeffeihga") and the second part, drawn in blue, is the part that constitutes the backtracking.

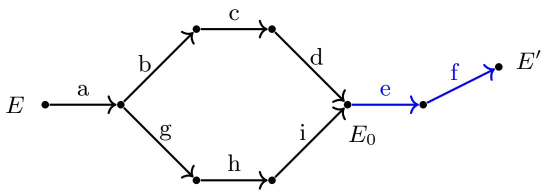

Fig. 5: Two paths (“abcdef” and “aghief”) that constitute an endomorphisms (“abcdeffeihga”) with backtracking and also a collision.

Therefore, to take the endomorphisms with backtracking into account when counting collisions, we simply need to count all endomorphisms without backtracking of different possible degrees $\mathscr{C}_{E_0}(\ell^r)$, for some $r \in \{1, ..., e-1\}$, and add an extra factor of $(\ell-1)\ell^{e-1-r}$, which is an upper bound on the number of ways we can draw a path of length $r+1$ from an elliptic curve E_0 to another one E' without backtracking, as in Figure 5. Note that this takes into account the fact that the curve E_0 already has two edges coming into it: this is why we have $(\ell-1)\ell^{e-1-r}$ instead of simply ℓ^{e-r}. □

If we were to expand the definition of H_p appearing in (9), we would see that in order to precisely compute $\sum_{i=1}^{n} \mathscr{C}_{E_i}(\ell^{2e})$, one would need to get a handle on the average splitting behaviour of $4\ell^{2e} - s^2$ in $\mathbb{F}_p$, which is beyond what is currently known. We will try to get around this issue as follows. One option is to find upper bounds (as in Section 5.1.1) for the number of collisions using the fact that $0 \leq H_p(D) \leq H(D)$. Unfortunately, using this same fact to find lower bounds would only give us a meaningless lower bond: zero. Another option (used in Section 5.1.2) is to replace the deterministic splitting of $4\ell^{2e} - s^2$ in $\mathbb{F}_p$ by a random Bernoulli process. We would essentially be modelling the Legendre symbols $\left(\frac{4\ell^{2e}-s^2}{p}\right)$ by a Bernoulli random variable. This would not give us an actual value for $\sum_{i=1}^{n} \mathscr{C}_{E_i}(\ell^{2e})$, but rather an estimate for it.

5.1.1 Upper bounds

We hereby prove an upper bound for the number of endomorphisms without backtracking, which will allow us to bound the number collisions. The expected number of endomorphisms without backtracking in $\mathcal{G}_{p^2}(\ell)$ is expressed as $\mathbb{E}_E\left[\mathscr{C}_E(\ell^{2r})\right] := \frac{1}{n}\sum_{i=1}^{n} \mathscr{C}_{E_i}(\ell^{2r})$.

Lemma 2. *The upper bound*

$$\mathbb{E}_E\left[\mathscr{C}_E(\ell^{2r})\right] \leq \frac{\ell^{2r+1}}{n(\ell-1)}$$

holds for all $r \geq 1$.

Proof. From equation (9), we have

$$\begin{align*}
\sum_{i=1}^{n} \mathscr{C}_{E_i}(\ell^{2r}) &= \frac{1}{2} \sum_{0<s^2\leq 4\ell^{2r}} H_p(4\ell^{2r}-s^2) - \frac{1}{2} \sum_{0<s^2\leq 4\ell^{2r-2}} H_p(4\ell^{2r-2}-s^2) \\
&\overset{(\diamondsuit)}{\leq} \frac{1}{2} \sum_{0<s^2\leq 4\ell^{2r}} H(4\ell^{2r}-s^2) \\
&= \frac{1}{2} \sum_{s^2\leq 4\ell^{2r}} H(4\ell^{2r}-s^2) - H(0) \\
&= \frac{1}{2}\left(2\sum_{d\mid\ell^{2r}} d - \sum_{d\mid\ell^{2r}} \min\left(d, \frac{\ell^{2r}}{d}\right)\right) + \frac{1}{12} \tag{11} \\
&= \sum_{i=0}^{2r} \ell^i - \sum_{i=0}^{r-1} \ell^i - \frac{1}{2}\ell^r + \frac{1}{12} \\
&= \sum_{i=r}^{2r} \ell^i - \frac{1}{2}\ell^r + \frac{1}{12} \\
&= \frac{\ell^{2r+1}-\ell^r}{\ell-1} - \frac{1}{2}\ell^r + \frac{1}{12}.
\end{align*}$$

Thus, putting everything together, we get that

$$\mathbb{E}_E\left[\mathscr{C}_E(\ell^{2r})\right] = \frac{1}{n}\sum_{i=1}^{n} \mathscr{C}_{E_i}(\ell^{2r}) \leq \frac{\ell^{2r+1}}{n(\ell-1)}.$$

□

Remark 3. In the proof of Theorem 2, we have the upper bound $\sum_{i=1}^{n} \mathscr{C}_{E_i}(\ell^{2r}) \leq \frac{\ell^{2r+1}-\ell^r}{\ell-1} - \frac{1}{2}\ell^r + \frac{1}{12}$. One can ask about how good it is. The only place in (11) were we use an upper bound instead of an equality is at ($\diamondsuit$), when we dropped the term $\frac{1}{2}\sum_{0<s^2\leq 4\ell^{2r-2}} H_p(4\ell^{2r-2}-s^2)$ (which has order ℓ^{2e-2}, so dropping it does not affect things much) and bounded $H_p(D)$ above by $H(D)$. Regarding the latter, we know that the definition of $H_p(D)$ depends on the splitting behaviour of p. We thus suspect that $H_p(D)$ is equal to zero half of the time and to $H(D)$ the other half, which leads us to expect that, on average, we have $H_p(D) = \frac{1}{2}H(D)$. Of course, this is not a rigorous statement and we cannot actually prove it, as the actual state of knowledge does not allow us to precisely pin down the average behaviour of the Legendre symbol $\left(\frac{4\ell^{2e}-s^2}{p}\right)$ as $s = 0, ..., 2\ell^e$. So we cannot actually say that our upper bound $\frac{\ell^{2r+1}-\ell^r}{\ell-1} - \frac{1}{2}\ell^r + \frac{1}{12}$ *is* twice the value of $\mathbb{E}_E\left[\mathscr{C}_E(\ell^{2r})\right]$, but we can imagine that is should in principal be close. We will discuss this issue further in Section 5.1.2.

We now make use of the bound we have obtained in Lemma 2 to give an estimate of the number of collisions that are expected in a graph.

Theorem 3. *We have the following upper bound for the average number of collisions of degree ℓ^e in the graph $\mathcal{G}_{p^2}(\ell)$:*

$$\mathbb{E}_E[\#Coll_{\ell^e}(E)] := \frac{1}{n}\sum_{i=1}^{n} \#Coll_{\ell^e}(E_i) \leq \frac{\ell^{2e}(\ell+1)}{n(\ell-1)}. \tag{12}$$

The average number of collisions is therefore in $O(1)$.

Proof. By Lemma 2, we know that for all $r \in \mathbb{N}$,

$$\mathbb{E}_E\left[\mathscr{C}_E(\ell^{2r})\right] \leq \frac{\ell^{2r+1}}{n(\ell-1)}.$$

Combining this with Lemma 1, we obtain

$$\begin{aligned}
\mathbb{E}_E[\#\mathrm{Coll}_{\ell^e}(E)] &\leq \mathbb{E}\left[\mathscr{C}_E(\ell^e)\right] + \sum_{r=1}^{e-1}\mathbb{E}\left[\mathscr{C}_E(\ell^r)\right](\ell-1)\ell^{e-1-r} \\
&\leq \frac{\ell^{2e+1}}{n(\ell-1)} + \sum_{r=1}^{e-1}\frac{\ell^{2r+1}}{n(\ell-1)}(\ell-1)\ell^{e-1-r} \\
&= \frac{\ell^{2e+1}}{n(\ell-1)} + \frac{\ell^e}{n}\sum_{r=1}^{e-1}\ell^r \\
&= \frac{\ell^{2e+1}}{n(\ell-1)} + \frac{\ell^e}{n}\cdot\frac{\ell^e-\ell}{\ell-1} \\
&= \frac{\ell^{2e+1}+\ell^{2e}-\ell^{e+1}}{n(\ell-1)} \\
&\leq \frac{\ell^{2e}(\ell+1)}{n(\ell-1)}.
\end{aligned}$$

□

Now that we have found an upper bound, we turn our attention to finding a lower bound.

5.1.2 Lower bounds

We would like to show that, for a fixed prime ℓ, the expected number of collisions $\mathrm{Coll}_{\ell^e}(E)$ for an arbitrary prime p and an arbitrary elliptic curve E in the graph $\mathcal{G}_{p^2}(\ell)$ is non negligible. By Lemma 1, we only need to focus on $\mathscr{C}_E(\ell^e)$ in order to bound $\mathrm{Coll}_{\ell^e}(E)$. Moreover, in order to get a handle on $\mathscr{C}_E(\ell^e)$ we need to know $\sum_{0<s^2\leq 4\ell^{2r}} H_p(4\ell^{2r}-s^2)$. The difficulty lies in relating the sums of modified Hurwitz class numbers $H_p(\cdot)$ to sums of Hurwitz class numbers $H(\cdot)$. As explained in Remark 3, our estimate is that the actual value of $\sum_{0<s^2\leq 4\ell^{2r}} H_p(4\ell^{2r}-s^2)$ is roughly $\frac{1}{2}\sum_{0<s^2\leq 4\ell^{2r}} H(4\ell^{2r}-s^2)$. However, we cannot prove this as it involves sums of Legendre symbols, and the best bounds for character sums that are known (based the bounds of Burgess in [8–10]) are far from being tight enough for what we need.

What we will do instead, in order to estimate $\mathrm{Coll}_{\ell^e}(E)$, is to model the behaviour of the Legendre symbol $\left(\frac{4\ell^{2e}-s^2}{p}\right)$ with Bernoulli events. Let

$$\varepsilon_p(D) := \begin{cases} 0 \text{ if } p \text{ splits in } \mathcal{O}_D; \\ 1 \text{ if } p \text{ is inert in } \mathcal{O}_{-D} \\ \quad \text{and does not divide the conductor of } \mathcal{O}_{-D}; \\ \frac{1}{2} \text{ if } p \text{ is ramified in } \mathcal{O}_{-D} \\ \quad \text{and does not divide the conductor of } \mathcal{O}_{-D}; \end{cases} \tag{13}$$

so that

$$H_p(4\ell^{2e} - s^2) = \varepsilon_p(4\ell^{2e} - s^2)H(4\ell^{2e} - s^2), \tag{14}$$

for all $4\ell^{2e} > s^2$. Note that we excluded the last case of the definition of *modified Hurwitz Class Numbers* in equation (2), which is when p divides the conductor of $\mathcal{O}_{s^2-4\ell^{2e}}$, because this case simply does not happen. Indeed, let $\mathfrak{f}$ denote the conductor of $\mathcal{O}_{-D}$ for $D := s^2 - 4\ell^{2e}$ with $4\ell^{2e} \geq s^2$ and $-d$ its fundamental discriminant. Then $\mathfrak{f} = \sqrt{\frac{D}{d}} < 2\ell^e < p$. Hence, p cannot divide the conductor $\mathfrak{f}$.

In the next lemma, we will replace $\varepsilon_p(4\ell^{2e} - s^2)$ with Bernoulli random variables, in order to get estimates on H_p and thus on $\mathbb{E}_E[\mathrm{Coll}_{\ell^e}(E)]$.

Let $X_0, X_1, ..., X_{2\ell^e}$ be i.i.d. Bern(1/2) random variables with values in $\{0, 1\}$ and set

$$H^*(4\ell^{2e} - s^2) = X_s \cdot H(4\ell^{2e} - s^2), \tag{15}$$

for all $s = 0, 1, ..., 2\ell^e$. Our goal is for H^*, from equation (15), to mimic H_p, from equation (14), for all $s^2 < 4\ell^{2e}$. Now that we have replaced the deterministic (but hard to pin down) behaviour of $H_p(4\ell^{2e} - s^2)$ by a probabilistic function $H^*(4\ell^{2e} - s^2)$, we can make probabilistic statements about it.

Lemma 3. *We have the following expectation for the sum* $\sum_{s^2 \leq 4\ell^{2e}} H^*(4\ell^{2e} - s^2)$*:*

$$\mathbb{E}\left[\sum_{s^2 \leq 4\ell^{2e}} H^*(4\ell^{2e} - s^2)\right] = \frac{\ell^{2e+1} - \ell^e}{\ell - 1} - \frac{1}{2}\ell^e.$$

Proof. First, let $\mathfrak{a}$ denote the sum $\sum_{s^2 \leq 4\ell^{2e}} H(4\ell^{2e} - s^2)$; then

$$\begin{aligned} \mathfrak{a} &= 2\sum_{i=0}^{2e} \ell^i - \sum_{i=0}^{2e} \min\{\ell^i, \ell^{2e-i}\} \\ &= 2\sum_{i=0}^{2e} \ell^i - 2\sum_{i=0}^{e-1} \ell^i - \ell^e \\ &= 2\frac{\ell^{2e+1} - \ell^e}{\ell - 1} - \ell^e. \end{aligned}$$

We then can easily compute the expectation $\mu := \mathbb{E}\left[\sum_{s^2\leq 4\ell^{2e}} H^*(4\ell^{2e}-s^2)\right]$ as

$$\begin{aligned}\mathbb{E}\left[\sum_{s^2\leq 4\ell^{2e}} H^*(4\ell^{2e}-s^2)\right] &= \sum_{s^2\leq 4\ell^{2e}} H(4\ell^{2e}-s^2)\mathbb{E}[X_s] \\ &= \mathfrak{a}/2 \\ &= \frac{\ell^{2e+1}-\ell^e}{\ell-1}-\frac{1}{2}\ell^e.\end{aligned} \tag{16}$$

□

Now that we know the expectation from Lemma 3, we can ask about how much the sum $\sum_{s^2\leq 4\ell^{2e}} H^*(4\ell^{2e}-s^2)$ can deviate from it's mean μ. One might initially be tempted to use the central limit theorem, and there are *triangular* versions of the CLT that allow for the setup we are in (we are *not* in the presence of a sequence of random variables, but the sequence that we are considering $\{H^*(4\ell^{2e}-s^2)\}_{s=0,\ldots,2\ell^e}$ changes for every new p). However, we would need the Lyapunov or Lindeberg condition, which is not obvious at all to prove rigorously (see Remark 4).

Instead, we need to revert to the use of *concentration inequalities* that do not use limits, such as Hoeffding's bound (see Proposition 2.5 in [40]).

Proposition 1. *Let $\varepsilon > 0$, there is a positive constant $c_\varepsilon := 1-e^{-\varepsilon^2/2}$ only depending on ε such that*

$$\mathbb{P}\left(\left|\frac{\sum_{s^2\leq 4\ell^{2e}} H^*(4\ell^{2e}-s^2)}{\mu}-1\right|\leq\varepsilon\right)\geq c_\varepsilon. \tag{17}$$

Proof. Let $\varepsilon > 0$. Using Hoeffding's bound, we can rewrite the probability

$$\mathbb{P}\left(\left|\frac{\sum_{s^2\leq 4\ell^{2e}} H^*(4\ell^{2e}-s^2)}{\mu}-1\right|\leq\varepsilon\right) =$$

$$\begin{aligned}&= \mathbb{P}\left(\left|\sum_{s^2\leq 4\ell^{2e}} H^*(4\ell^{2e}-s^2)-\mu\right|\leq\varepsilon\mu\right)\\ &= \mathbb{P}\left(\left|\sum_{s^2\leq 4\ell^{2e}} H(4\ell^{2e}-s^2)X_s-\mu\right|\leq\varepsilon\mu\right)\\ &\geq 1-\exp\left[-\frac{\varepsilon^2\mu^2}{2\cdot\sum_{s^2\leq 4\ell^{2e}}\mathrm{Var}(H(4\ell^{2e}-s^2)X_s)}\right]\\ &= 1-\exp\left[-\frac{\varepsilon^2\mu^2}{2\cdot\sum_{s^2\leq 4\ell^{2e}}(H(4\ell^{2e}-s^2))^2\mathrm{Var}(X_s)}\right]\end{aligned}$$

$$= 1 - \exp\left[-2\varepsilon^2 \frac{\left(\sum_{s^2\leq 4\ell^{2e}} H(4\ell^{2e}-s^2)\right)^2}{\sum_{s^2\leq 4\ell^{2e}} H(4\ell^{2e}-s^2)^2}\right]$$

$$\overset{(\clubsuit)}{>} 1 - \exp\left[-\frac{2\varepsilon^2\mu^2}{\mathfrak{a}^2}\right]$$

$$= 1 - e^{-\varepsilon^2/2}.$$

□

This shows the quantity $\sum_{s^2\leq 4\ell^{2e}} H^*(4\ell^{2e}-s^2)$ is close to $\mathfrak{a}/2 \approx p$ a positive proportion of the time. Indeed, if we take $\varepsilon = 0.5$, then $\sum_{s^2\leq 4\ell^{2e}} H^*(4\ell^{2e}-s^2)$ is between 0.5μ and 1.5μ with probability $\geq 11\%$ which is non-negligible. This provides more evidence to the fact that we can think of $\sum_{s^2\leq 4\ell^{2e}} H_p(4\ell^{2e}-s^2)$ as being close to $\mu = \frac{1}{2}\sum_{s^2\leq 4\ell^{2e}} H(4\ell^{2e}-s^2)$.

Note that for proving Proposition 1, we could have also used Markov's inequality, which gives slightly better bounds for certain values of ε.

Remark 4. One can actually do much better than the bound provided in Proposition 1. Indeed, in step ($\clubsuit$) of the proof of the above proposition, we bound $\frac{\left(\sum_{s^2\leq 4\ell^{2e}} H(4\ell^{2e}-s^2)\right)^2}{\sum_{s^2\leq 4\ell^{2e}} H(4\ell^{2e}-s^2)^2}$ above by 1. However, we believe that this quantity should tend to zero as p goes to infinity. Indeed, as $\sum_{s^2\leq 4\ell^{2e}} H(4\ell^{2e}-s^2) \in \Theta(p)$ (by the proof of Lemma 2), we only need to explain why $\sum_{s^2\leq 4\ell^{2e}} H(4\ell^{2e}-s^2)^2$ is on $o(p^2)$. By the Brauer-Siegel theorem [7], we have that for all $\varepsilon > 0$ there is some $c_\varepsilon > 0$ such that $h(-d) < c_\varepsilon d^{1/2+\varepsilon}$. In addition, we know that there are at most $2d^{1/4}$ square divisors of d for all d; or even better: there are $o(d^\varepsilon)$ divisors for d for all ε. This gives us a small enough bound on $H(D)$, especially when D is big, which is the case in the majority of the values $D = 4\ell^{2e} - s^2$ that we consider. It then follows that $\sum_{s^2\leq 4\ell^{2e}} H(4\ell^{2e}-s^2)^2$ is $o(\ell^{4e}) = o(p^2)$.

Now, recall from equations (8) and (9) that

$$\begin{aligned}\mathbb{E}_E\left[\mathscr{C}_E(\ell^{2e})\right] &:= \frac{1}{n}\sum_{i=1}^{n}\mathscr{C}_{E_i}(\ell^{2e}) \\ &= \frac{1}{2}\sum_{0<s^2<4\ell^{2e}} H_p(4\ell^{2e}-s^2) - \frac{1}{2}\sum_{0<s^2<4\ell^{2e-2}} H_p(4\ell^{2e-2}-s^2) \\ &= \frac{1}{2}\sum_{s^2\leq 4\ell^{2e}} H_p(4\ell^{2e}-s^2) - \frac{1}{2}\sum_{s^2\leq 4\ell^{2e-2}} H_p(4\ell^{2e-2}-s^2).\end{aligned}$$

Let us define

$$\mathscr{C}^*(\ell^{2e}) := \frac{1}{2}\sum_{s^2\leq 4\ell^{2e}} H^*(4\ell^{2e}-s^2) - \frac{1}{2}\sum_{s^2\leq 4\ell^{2e-2}} H^*(4\ell^{2e-2}-s^2)$$

which should approximate the total number of collisions $\sum_{i=1}^{n} \mathscr{C}_{E_i}(\ell^{2e})$ defined in (9). As in (16), we can compute

$$\begin{aligned}
\mathbb{E}\left[\mathscr{C}^*(\ell^{2e})\right] &= \frac{1}{2}\mathbb{E}\left[\sum_{s^2 \le 4\ell^{2e}} H^*(4\ell^{2e} - s^2)\right] - \frac{1}{2}\mathbb{E}\left[\sum_{s^2 \le 4\ell^{2e-2}} H^*(4\ell^{2e-2} - s^2)\right] \\
&= \frac{\ell^{2e+1} - \ell^e}{\ell - 1} - \frac{1}{2}\ell^e - \left(\frac{\ell^{2e-1} - \ell^{e-1}}{\ell - 1} - \frac{1}{2}\ell^{e-1}\right) \\
&= \frac{\ell^{2e+1} - \ell^e}{\ell - 1} - \frac{1}{2}\ell^e - \frac{\ell^{2e-1} - \ell^{e-1}}{\ell - 1} + \frac{1}{2}\ell^{e-1} \\
&= \frac{1}{2}\ell^{e-1}(\ell + 1)(2\ell^e - 1).
\end{aligned}$$

Using Lemma 1, we conclude the proof of the following theorem.

Theorem 4. *Using the Bernoulli model introduced in Section 5.1.2, we expect that for an arbitrary supersingular isogeny graph $\mathcal{G}_{p^2}(\ell)$, the number of collisions of degree ℓ^e is bounded below by $\frac{1}{2}\ell^{e-1}(\ell + 1)(2\ell^e - 1)$. In particular, this model predicts the number of collisions starting at a given elliptic curve to be $\Omega(1)$.*

5.2 Double Collisions

The second problem we consider is directly related with the first one.

Problem 6. *Assuming the scenario described in Problem 5, determine the probability of having two more isogenies (not necessarily distinct this time), $\psi_A, \psi_B : E_1 \longrightarrow E_2$ of degree q^f, and cyclic kernel, where q is a small prime and $f \in \mathbb{N}$, with $p \approx q^{2f}$, such that there exists $R \in E_0$ for which $R_A := \varphi_A(R)$ and $R_B := \varphi_B(R)$ are generators for the kernels of ψ_A and ψ_B, respectively.*

$$\begin{tikzcd}
E_0 \arrow[r, bend left, "\varphi_A"] \arrow[r, bend right, "\varphi_B"'] & E_1 \arrow[r, bend left, "\psi_A"] \arrow[r, bend right, "\psi_B"'] & E_2
\end{tikzcd}$$

Fig. 6: Isogenies of the same degrees sending E_0 in E_1, and E_1 in E_2.

Let us compute the average number of pairs of isogenies from E_1 to some E_2, as in Problem 6. We can rephrase this problem as follows: determine the probability of finding two points $R_A, R_B \in E_1[q^f]$ of maximal order, such that

$E_1/\langle R_A\rangle \cong E_1/\langle R_B\rangle$ and $R_i = \varphi_i(R)$ for $i = A, B$. Let $\mathbf{Coll}_{\ell^e,q^f}(E_0, E_1)$ denote the set

$$\left\{\begin{array}{c}(R_A, R_B) \in (E_1[q^f]_{\max})^2/((X,Y)\sim(Y,X)) \text{ such that}\\ E_1/\langle R_A\rangle \cong E_1/\langle R_B\rangle; R_i = \varphi_i(R) \text{ for some } R \in E_0[q^f]_{\max}\end{array}\right\}$$

of desired pairs (R_A, R_B) *up to equivalence*, where $E_1[q^f]_{\max}$ denotes the set of point of maximal order in $E_1[q^f]/_\sim$. Let also $\overline{\mathbf{Coll}}_{\ell^e,q^f}(E_1)$ denote the set:

$$\{(R_A, R_B) \in (E_1[q^f]_{\max})^2/((X,Y)\sim(Y,X)) : E_1/\langle R_A\rangle \cong E_1/\langle R_B\rangle\}.$$

Note that $\mathbf{Coll}_{\ell^e,q^f}(E_1) \subseteq \overline{\mathbf{Coll}}_{\ell^e,q^f}(E_1)$. Now, since we are allowing $R_A \sim R_B$, we can split $\mathbf{Coll}_{\ell^e,q^f}(E_1)$ and $\overline{\mathbf{Coll}}_{\ell^e,q^f}(E_1)$ into two disjoint subsets:

$$\mathbf{Coll}_{\ell^e,q^f}(E_0, E_1) = \mathbf{Coll}^{\neq}_{\ell^e,q^f}(E_0, E_1) \sqcup \mathbf{Coll}^{=}_{\ell^e,q^f}(E_0, E_1),$$

$$\overline{\mathbf{Coll}}_{\ell^e,q^f}(E_1) = \overline{\mathbf{Coll}}^{\neq}_{\ell^e,q^f}(E_1) \sqcup \overline{\mathbf{Coll}}^{=}_{\ell^e,q^f}(E_1),$$

where $\mathbf{Coll}^{\neq}_{\ell^e,q^f}(E_0, E_1)$ is the set

$$\left\{\begin{array}{c}(R_A, R_B) \in (E_1[q^f]_{\max})^2/((X,Y)\sim(Y,X)) \text{ such that}\\ E_1/\langle R_A\rangle \cong E_1/\langle R_B\rangle, R_A \neq R_B, R_i = \varphi_i(R) \text{ for some } R \in E_0[q^f]_{\max}\end{array}\right\},$$

$\mathbf{Coll}^{=}_{\ell^e,q^f}(E_0, E_1)$ is the set

$$\left\{\begin{array}{c}(R_A, R_B) \in (E_1[q^f]_{\max})^2/((X,Y)\sim(Y,X)) \text{ such that}\\ E_1/\langle R_A\rangle \cong E_1/\langle R_B\rangle, R_A = R_B, R_i = \varphi_i(R) \text{ for some } R \in E_0[q^f]_{\max}\end{array}\right\}$$
$$\cong \left\{R_A \in E_1[q^f]_{\max} \,|R_A = \varphi_A(R) = \varphi_B(R) \text{ for some } R \in E_0[q^f]_{\max}\right\},$$

$\overline{\mathbf{Coll}}^{\neq}_{\ell^e,q^f}(E_1)$ is the set

$$\left\{(R_A, R_B) \in (E_1[q^f]_{\max})^2/((X,Y)\sim(Y,X))\middle|\, E_1/\langle R_A\rangle \cong E_1/\langle R_B\rangle, R_A \neq R_B\right\}$$

and $\overline{\mathbf{Coll}}^{=}_{\ell^e,q^f}(E_1)$ is the set

$$\left\{(R_A, R_B) \in (E_1[q^f]_{\max})^2/((X,Y)\sim(Y,X))\middle|\, E_1/\langle R_A\rangle \cong E_1/\langle R_B\rangle, R_A = R_B\right\}$$
$$\cong \left\{R_A \in E_1[q^f]_{\max}\right\}.$$

It is known that the size of the torsion subgroup $E_1[q^f]$ is q^{2f} and that the number of subgroups of $E_1[q^f]$ having order q^f (i.e. the number of points with maximal order, up to equivalence) is $q^f + q^{f-1}$. We can now prove that the expected number of double collisions in a supersingular isogeny graph for a fixed initial curve is 1.

Proposition 2. *We have that*

$$\mathbb{E}[\#\boldsymbol{Coll}^{=}_{\ell^e,q^f}(E_0, E_1)] = 1.$$

Proof. First of all, notice that for $i \in \{A, B\}$ the isogeny $\varphi_i : E_0 \longrightarrow E_1$ restricts to an isomorphism $\varphi_i : E_0[q^f]_{\max} \xrightarrow{\sim} E_1[q^f]_{\max}$. Hence,

$$\mathbb{P}_{R_A, R_B \in E_1[q^f]_{\max}} \left(R_i = \varphi_i(R) \text{ for some } R \in E_0[q^f]_{\max} \right) =$$
$$= \mathbb{P}_{R_A, R_B \in E_1[q^f]_{\max}} \left(\varphi_A^{-1}(R_A) = \varphi_B^{-1}(R_B) \right).$$

So the cardinality of $\mathbf{Coll}^{=}_{\ell^e, q^f}(E_0, E_1)$ is equal to

$$\sum_{R_A, R_B \in E_1[q^f]_{\max}} \mathbb{1}_{[E_1/\langle R_A \rangle \cong E_1/\langle R_B \rangle, R_A = R_B, R_i = \varphi_i(R) \text{ for some } R \in E_0[q^f]_{\max}]}$$
$$= \sum_{R_A \in E_1[q^f]_{\max}} \mathbb{1}_{\left[\varphi_A^{-1}(R_A) = \varphi_B^{-1}(R_A)\right]}.$$

Thus,

$$\begin{aligned} \mathbb{E}[\#\mathbf{Coll}^{=}_{\ell^e, q^f}(E_0, E_1)] &= \mathbb{E} \sum_{R_A \in E_1[q^f]_{\max}} \mathbb{1}_{\left[\varphi_A^{-1}(R_A) = \varphi_B^{-1}(R_A)\right]} \\ &= \sum_{R_A \in E_1[q^f]_{\max}} \mathbb{P}\left(\varphi_A^{-1}(R_A) = \varphi_B^{-1}(R_A) \right) \\ &\overset{(\dagger)}{=} \sum_{X, Y \in E_0[q^f]_{\max}} \mathbb{P}\left(X = Y \right) \\ &= 1. \end{aligned}$$

The justification for step (†) is again, as in the first step of this proof, that $\varphi_i : E_0 \longrightarrow E_1$ restricts to an isomorphism $\varphi_i : E_0[q^f]_{\max} \xrightarrow{\sim} E_1[q^f]_{\max}$. □

We are finally ready to give an answer to Problem 6: we want to estimate how many double collisions we can have in the graph $\mathcal{G}_{p^2}(\ell)$. We have found estimates for lower bounds on the number of first collisions that could occur in subsection 5.1.2, and we have found the expected number of certain kinds of second collisions in Proposition 2. We now argue that we can combine these two results together. Indeed, the first collision has degree ℓ^e and the second has degree q^f. Thus, we can assume without loss of generality that the events of having a first collision and having a second collision are independent. This allows us to write that the expectation of this event is

$$\begin{aligned} &\mathbb{E}_{E_0}[\mathrm{Coll}_{\ell^e}(E_0)] \cdot \mathbb{E}_{E_0, E_1}\left[\mathbf{Coll}_{\ell^e, q^f}(E_0, E_1)\right] \\ &\geq \mathbb{E}_{E_0}[\mathrm{Coll}_{\ell^e}(E_0)] \cdot \mathbb{E}_{E_0, E_1}\left[\mathbf{Coll}^{=}_{\ell^e, q^f}(E_0, E_1)\right] \\ &= \mathbb{E}_{E_0}[\mathrm{Coll}_{\ell^e}(E_0)]. \end{aligned}$$

However, as amply discussed in Section 5.1.2, we cannot directly bound the quantity $\mathbb{E}_{E_0}[\mathrm{Coll}_{\ell^e}(E_0)]$ from below (without getting the trivial bound zero!). We therefore revert to using our Bernoulli model for approximating Legendre symbols. We then get the following result.

Theorem 5. *Fix an arbitrary supersingular isogeny graph as in Problems 5 and 6 and pick an arbitrary elliptic curves E_0 on it. Then, using the Bernoulli model introduced in Section 5.1.2 to estimate Legendre symbols, we expect that the number of double collisions starting at E_0 is non negligible. More precisely, it is in $\Omega(1)$.*

Remark 5. The above result uses a statistical model to approximate a deterministic event. This means that the bound is not necessarily always true, and that one could potentially hope to manufacture a prime p specifically tailored in order to ensure that the graph $\mathcal{G}_{p^2}(\ell)$ has fewer collisions. However, Bernoulli events are a very good way of approximating Legendre symbols, and therefore, on average, our bound should provide a very good indicator for what is actually happening.

6 Collisions in Isogeny-based Cryptography

When bounding the average number of collisions in Theorem 3 of Section 5, we fixed the starting elliptic curve and allowed the image curve to be any vertex on the graph. However, one can also consider the average number of collisions starting at a given elliptic curve E and colliding at another fixed elliptic curve E'. This would drastically change the order of magnitude of our final answer, and we would then get a much smaller quantity. This fact is sometimes implicitly assumed in the literature, with scarce formal justifications.

In the analysis of claw-finding algorithms for the CSSI problem of [13] - and in particular in Section 2.2. - the focus is on finding an isogeny φ from E_0 to E_1. It is claimed that, once an elliptic curve E', with isogenies $f_1 : E_0 \longrightarrow E'$ and $g_1 : E_1 \longrightarrow E'$ is found, the composition $\hat{g_1} \circ f_1$ would return φ. This argument assumes that it is impossible (or at least extremely unlikely) to have collisions $f_1, f_2 : E_0 \longrightarrow E'$ or $g_1, g_2 : E_1 \longrightarrow E'$, with $f_1 \not\sim f_2$ and $g_1 \not\sim g_2$. If such collisions happen, all the isogenies $\hat{g_i} \circ f_j$, with $i, j \in \{1, 2\}$, would be different, and the recovered isogeny would not necessarily be equal to φ.

Also in [23, Sec. 3.1], it is claimed that if two curves E/G_1 and E/G_2 are isomorphic, then one could safely assume $G_1 = G_2$. In other words, if two isogenies of the same sufficiently small degree start at the same curve and have the same codomain, then they are equivalent. This argument is mentioned in the adaptive attack developed in [23] and assumes that collisions are extremely unlikely.

We now use the calculations done in Section 5 to provide a more rigorous proof of this assumption. Let $\mathrm{Coll}_{\ell^e}(E, E')$ denote the number of collisions between E and E' of degree ℓ^e in the graph $\mathcal{G}_{p^2}(\ell)$, with $p \approx \ell^{2e}$. The number of elliptic curves that one can obtain starting from E and taking a path of length e is $\Theta(\ell^e) = \Theta(\sqrt{p})$. Therefore, we would get that $\mathrm{Coll}_{\ell^e}(E, E')$ is $O(1/\ell^e)$ which is negligible. The details of this argument are proven in the following theorem.

Theorem 6. *Given a supersingular graph $\mathcal{G}_{p^2}(\ell)$, pick uniformly at random two elliptic curves E, E' in $\mathcal{G}_{p^2}(\ell)$ and assume that there is an isogeny of degree ℓ^e with no backtracking between them. Then, the expected number of collisions of degree ℓ^e from E to E' in the graph $\mathcal{G}_{p^2}(\ell)$ is $O(1/\sqrt{p})$.*

Proof. We know from Theorem 3 that there are, on average, $O(1)$ collisions starting at E. Let m_E be the number of elliptic curves that are connected to E via a path of length e without backtracking. Let $\mathfrak{m}_E = \ell^e + \ell^{e-1}$ be the number of non equivalent isogenies of degree ℓ^e with no backtracking starting from E. Note that m_E is not exactly the same as $\mathfrak{m}_E$. This is because it possible for two non-equivalent isogenies of degree ℓ^e starting from E to arrive to the same elliptic curve E', thus forming a collision. We have shown that the number of such collisions is $O(1)$. Thus, $0 \leq \mathfrak{m}_E - m_E$ is $O(1)$, and

$$\begin{aligned} \mathbb{E}_{E,E'}[\#\mathrm{Coll}_{\ell^e}(E,E')] &:= \frac{1}{n}\sum_{i=1}^{n}\frac{1}{m_{E_i}}\sum_{E'\in\mathcal{D}}\#\mathrm{Coll}_{\ell^e}(E_i,E') \\ &\leq \frac{1}{n}\sum_{i=1}^{n}\frac{\#\mathrm{Coll}_{\ell^e}(E_i)}{m_{E_i}} \\ &\leq \frac{1}{n}\sum_{i=1}^{n}\frac{\#\mathrm{Coll}_{\ell^e}(E_i)}{\mathfrak{m}_{E_i}-(\mathfrak{m}_{E_i}-m_{E_i})} \end{aligned} \tag{18}$$

Now, from the last line of Equation (18), we know that $\mathbb{E}_E[\#\mathrm{Coll}_{\ell^e}(E)] \in O(1)$, $\mathfrak{m}_{E_i} = \ell^e + \ell^{e-1} \in \Theta(\sqrt{p})$ and $\mathfrak{m}_{E_i} - m_{E_i} \in O(1)$ for all i. Thus,

$$\mathbb{E}_{E,E'}[\#\mathrm{Coll}_{\ell^e}(E,E')] \in O(1/\sqrt{p}).$$

□

7 Conclusion

In this paper, we disputed the validity of the existing proofs for the special soundness property of the SIDH-based identification protocol $\mathsf{ID}_{\mathsf{SIDH}}$. In addition to providing concrete examples for two scenarios where the proposed extraction algorithm $\mathsf{Ex}_{\mathsf{SIDH}}$ fails, we also made a careful study of the number of collisions that occur in supersingular isogeny graphs. Our analysis shows that one can always expect the existence of certain single and double collisions, exactly those that make $\mathsf{Ex}_{\mathsf{SIDH}}$ fail. We also used our calculations to give upper bounds on the number of certain collisions in the supersingular isogeny graph $\mathcal{G}_{p^2}(\ell)$ for fixed primes p and ℓ, thus providing a rigorous justification to some informal claims which have appeared in the literature. Our calculations are general and could be exploited in many contexts within Isogeny-based Cryptography. We leave the improvement of the bounds we provided on the number of collisions for future work, as this would require developing new number theoretic tools.

Despite the fact we showed that the issue with the previous special-soundness proofs does not impact the security of the SIDH-based signature scheme that have been proposed so far, formalising the security properties of $\mathsf{ID}_{\mathsf{SIDH}}$ seems to be important to avoid any misconceptions or unnecessary flaws in existing and future isogeny-based cryptosystems.

References

[1] M. Abdalla, J. An, M. Bellare, and C. Namprempre. From identification to signatures via the fiat-shamir transform: Minimizing assumptions for security and forward-security. *IACR Cryptology ePrint Archive*, 2002:418–433, 05 2002.

[2] G. Asharov, A. Jain, A. López-Alt, E. Tromer, V. Vaikuntanathan, and D. Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In *EUROCRYPT 2012*, pages 483–501. Springer, 2012.

[3] T. Attema, R. Cramer, and L. Kohl. A compressed σ-protocol theory for lattices. IACR Cryptol. ePrint Arch. 2021 (2021): 307.

[4] R. Azarderakhsh, M. Campagna, C. Costello, L. De Feo, B. Hess, A. Jalali, D. Jao, B. Koziel, B. La Macchia, and P. Longa. Supersingular isogeny key encapsulation. Submission to the NIST Post-Quantum Standardization project, 2017.

[5] F. Benhamouda, J. Camenisch, S. Krenn, V. Lyubashevsky, and G. Neven. Better zero-knowledge proofs for lattice encryption and their application to group signatures. In *ASIACRYPT 2014*, pages 551–572. Springer, 2014.

[6] W. Beullens, S. Katsumata, and F. Pintore. Calamari and Falafl: Logarithmic (linkable) ring signatures from isogenies and lattices. In *ASIACRYPT 2020*, pages 464–492. Springer, Cham, 2020.

[7] R. Brauer. On the zeta-functions of algebraic number fields. *American Journal of Mathematics*, 69(2):243–250, 1947.

[8] D. Burgess. The character sum estimate with r= 3. *Journal of the London Mathematical Society*, 2(2):219–226, 1986.

[9] D. A. Burgess. On character sums and primitive roots. *Proceedings of the London Mathematical Society*, 3(1):179–192, 1962.

[10] D. A. Burgess. On character sums and l-series. ii. *Proceedings of the London Mathematical Society*, 3(1):524–536, 1963.

[11] D. X. Charles, K. E. Lauter, and E. Z. Goren. Cryptographic hash functions from expander graphs. *Journal of Cryptology*, 22(1):93–113, 2009.

[12] A. Childs, D. Jao, and V. Soukharev. Constructing elliptic curve isogenies in quantum subexponential time. *Journal of Mathematical Cryptology*, 8.1:1–29, 2014.

[13] C. Costello, P. Longa, M. Naehrig, J. Renes, and F. Virdia. Improved classical cryptanalysis of sike in practice. In *Public Key Cryptography 2020*, pages 505–534. Springer, 2020.

[14] J. M. Couveignes. Hard homogeneous spaces. IACR Cryptol. ePrint Arch. 2006 (2006): 29.

[15] L. De Feo. Mathematics of isogeny based cryptography. *arXiv preprint arXiv:1711.04062*, 2017.

[16] L. De Feo, S. Dobson, S. D. Galbraith, and L. Zobernig. Sidh proof of knowledge. *IACR Cryptol. ePrint Arch. 2021/1023*, 2021.

[17] L. De Feo, D. Jao, and J. Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. *Journal of Mathematical Cryptology*, 8.3:209–247, 2014.

[18] L. De 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*, pages 64–93, Cham, 2020. Springer International Publishing.

[19] M. Deuring. Die typen der multiplikatorenringe elliptischer funktionenkörper. In *Abhandlungen aus dem mathematischen Seminar der Universität Hamburg*, volume 14, pages 197–272. Springer, 1941.

[20] A. El Kaafarani and S. Katsumata. Attribute-based signatures for unbounded circuits in the rom and efficient instantiations from lattices. In *PKC 2018*, pages 89–119. Springer, 2018.

[21] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, pages 186–194. Springer, 1986.

[22] T. B. Fouotsa, P. Kutas, and S.-P. Merz. On the isogeny problem with torsion point information. Cryptology ePrint Archive, Report 2021/153, 2021. `https://ia.cr/2021/153`.

[23] S. D. Galbraith, C. Petit, B. Shani, and Y. Bo Ti. On the security of supersingular isogeny cryptosystems. In *ASIACRYPT*, pages 63–91. Springer, 2016.

[24] S. D. Galbraith, C. Petit, and J. Silva. Identification protocols and signature schemes based on supersingular isogeny problems. In *ASIACRYPT*, pages 3–33. Springer, 2017.

[25] B. Gross. Heights and the special values of L-series. In *Conference Proceedings of the CMS*, volume 7, 1987.

[26] A. Hurwitz. Ueber relationen zwischen classenanzahlen binärer quadratischer formen von negativer determinante. *Mathematische Annalen*, 25(2):157–196, 1885.

[27] H. Iwaniec and E. Kowalski. *Analytic number theory*, volume 53. American Mathematical Soc., 2004.

[28] S. Jaques and J. M. Schanck. Quantum cryptanalysis in the ram model: Claw-finding attacks on sike. In *Crypto 2019*, pages 32–61. Springer, Cham, 2019.

[29] D. Kohel, K. Lauter, C. Petit, and J.-P. Tignol. On the quaternion -isogeny path problem. *LMS Journal of Computation and Mathematics*, 17, 06 2014.

[30] A. K. Lenstra, H. W. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische annalen*, 261:515–534, 1982.

[31] F. Marc. Communication-efficient non-interactive proofs of knowledge with online extractors. In *CRYPTO*, pages 152–168. Springer, Berlin, Heidelberg, 2005.

[32] H. Onuki, Y. Aikawa, and T. Takagi. The existence of cycles in the supersingular isogeny graphs used in sike. *2020 International Symposium on Information Theory and Its Applications (ISITA)*, pages 358–362, 2020.

[33] A. Rostovtsev and A. Stolbunov. Public-key cryptosystem based on isogenies. IACR Cryptol. ePrint Arch. 2006 (2006): 145.

[34] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings of the 35th annual symposium on foundations of computer science*, pages 124–134. IEEE, 1994, November.

[35] J. H. Silverman. *The arithmetic of elliptic curves.* New York, Springer, Vol. 106, 2009.

[36] D. Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In *EUROCRYPT*, pages 755–784. Springer, 2015.

[37] D. Unruh. Post-quantum security of fiat-shamir. In *ASIACRYPY*, pages 65–95. Springer, Cham, 2017.

[38] D. Urbanik and D. Jao. Sok: The problem landscape of sidh. In *APKC*, pages 53–60. ACM, 2018.

[39] V. Vitse. Simple oblivious transfer protocols compatible with supersingular isogenies. In *International Conference on Cryptology in Africa*, pages 56–78. Springer, Cham, 2019.

[40] M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.

[41] Y. Yoo, R. Azarderakhsh, A. Jalali, D. Jao, and V. Soukharev. A post-quantum digital signature scheme based on supersingular isogenies. In *FC*, pages 163–181. Springer, 2017.

8 Appendix

In order to find a meaningful example for the scenario depicted in Section 3.2, we follow the same approach adopted by Onuki, Aikawa and Takagi [32] do determine the existence of endomorphisms of prescribed degree in supersingular isogeny graphs for SIKE parameters. The general ideal is the following:

1. since $p \equiv 15 \pmod{16}$ in our study case, we know the explicit description of $End(E_0)$, we know that there is a unique 2-isogeny that maps E_0 to a distinct curve, which we call E_6, and by Lemma 1 of [32] we know the explicit description of $End(E_6)$. In particular, any endomorphism of E_6 can be written as
$$\alpha = a + b(2\mathbf{i}) + c\frac{1+\mathbf{j}}{2} + d\frac{\mathbf{i}+\mathbf{k}}{4},$$
with $a, b, c, d \in \mathbb{Z}$ (at least one of which not divisible by ℓ) and $\mathbf{i}, \mathbf{j}, \mathbf{k}$ such that $\mathbf{i}^2 = -1, \mathbf{j}^2 = -p, \mathbf{ij} = \mathbf{k} = -\mathbf{ji}$;
2. we then consider the corresponding norm equation
$$\ell^n = \frac{1}{16}\big((4a+2c)^2 + (8b+d)^2 + (4c^2+d^2)p\big) \tag{19}$$
where ℓ^n is the norm of the endomorphism α. By Theorem 1 in [32], we can focus on integers n such that $\frac{p+1}{16} < \ell^n \leq \ell^{2e}$, and by rearranging Equation 19 and replacing $A = 4a + 2c$ and $B = 8b + d$, we reduce the problem to solving the following Diophantine equation
$$16\ell^n - (4c^2 + d^2)p = A^2 + B^2 \tag{20}$$
with c and d not both 0, and A, B such that $a = \frac{A-2c}{4} \in \mathbb{Z}, b = \frac{B-d}{8} \in \mathbb{Z}$.

The authors provide results for the first two parameter sets of SIKE, namely a collision of 3^{e_2}-isogenies for p434 and no collisions for p503. They could not examine the remaining parameter sets because of a lack of computational resources. With our implementation, we could investigate the two remaining parameter sets, running the code on a laptop mounting a Coffee Lake 2.2 GHz Intel i7 processor and 12GB of RAM. The result is that there exist no cycles (and thus no collisions) for p610 neither of degree 2^{2e_1} nor 3^{2e_2}. For p751 there are no collisions of 2^{e_1}-isogenies, while there exist two pairs of colliding 3^{e_2}-isogenies, which we now present.

Let $e_1 = 372$, $e_2 = 239$ and p751 $= 2^{e_1} \cdot 3^{e_2} - 1$. Let η be a primitive element of $\mathbb{F}_{p^2}$ with minimal polynomial $x^2 + 1$; given the $E_6[3^{e_2}]$ torsion basis

$$\begin{aligned}
P_2 = (&841279105546461804686499115213762297740963469565458990324532197\\
&170119753830282998639717315851968048398639566308956083752260055190 5\\
&715721966592107851870816915611798787999369829670862974603718692670 6\\
&8644394211287290040655663023 1, 232832339774986047278299512593543184 4\\
&906025806061286315561848048094118981177935779374110739971558302443 6
\end{aligned}$$

$$\begin{aligned}
&69052376824243054704744119836366858126425102589256068192209923162\\
&99948430865850777345758025150183257684455900332189181028 6),\\
Q_2 = (&54363521006076338401431179820675907258704379723256874298913509 5\\
&93375464906320336149438290118331236448565893730039881970952595605 7\\
&43871729893121233367553907230165153029309104381363871811519161954 3\\
&534872448539156352030073708291 2, 27468206193805043276584564703605 76\\
&39211608416213987304454325393688867446997069856205162995991758227 0\\
&79505132960707097561160222179773210530421731878375757129548644734 3\\
&9099401589351814612812541553928275189840035770968426563389 42 \cdot \eta),
\end{aligned}$$

the kernels of two distinct isogenies colliding in curves of j-invariant 36748324814 65277414370377344994980543168961447497832891959594848980770518955999 9 20954686286458329903693959432922063213725111691674155571714123246827 0 09952290798655944625158009755772764608896053794594261659270555052100 4 94910843 $\cdot\eta +$ 21591927783730499578412437897563320823963930917022966131 0 74025628978413773392103065773928886848158072953862916682620505198834 6 31526353981178245576179028287586668352682365984162211994213314615983 7 67274692829779784506037373 25459 are generated by the points

$$\begin{aligned}
K_1 = {}&19326264832755263436205782983289563837195637516764115445955224 9\\
&80550866458993667856572223600330507498160017978890 1 \cdot P_2 + 398140788\\
&12090453073616618563048908428001097524649990508670812852166946098 1\\
&866051005968073733368199418319346274 98 \cdot Q_2
\end{aligned}$$

and

$$\begin{aligned}
K_2 = {}&88315269100529870147635044900489166652883464217472613054693194 0\\
&78697241064121690098335862704399965485532642793536 6 \cdot P_2 + 398140788\\
&12090453073616618563048908428001097524649990508670812852166946098 1\\
&866051005968073733368199418319346275 03 \cdot Q_2
\end{aligned}$$

respectively. The other pair of isogenies colliding in curves of j-invariant 6679885 26030402783860739089287182477825842819805123827815659483007390842172 6 47892554819661010273680993059377869677698866009222030236698267892091 3 64176119952617805089545375532460262696634387887318564082519172094682 0 565156665988 $\cdot\eta +$ 21591927783730499578412437897563320823963930917022 96 61310740256289784137733921030657739288868481580729538629166826205051 9 88346315263539811782455761790282875866683526823659841622119942133146 1 59837672746928297797845060373732545 9 is given by the two kernels generated by

$$\begin{aligned}
K_1 = {}&19326264832755263436205782983289563837195637516764115445955224\\
&980550866458993667856572223600330507498160017978890 1 \cdot P_2 + 1036601\\
&26052076088276479166027473839647278991981771729449781337774031412
\end{aligned}$$

$$913296697444848405567396790989509467309676 9 \cdot Q_2$$

and

$$\begin{aligned} K_2 = {} & 8831526910052987014763504490048916665288346421747261305469319 4 \\ & 0786972410641216900983358627043999654855326427935366 \cdot P_2 + 103660 \\ & 126052076088276479166027473839647278991981771729449781337774031412 \\ & 913296697444848405567396790989509467309676 4 \cdot Q_2 \end{aligned}$$

respectively.