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

Wissam Ghantous[1], Shuichi Katsumata[2], Federico Pintore[3] and Mattia Veroni[4]

[1] Mathematical Institute, University of Oxford, UK,
wissam.ghantous@maths.ox.ac.uk
[2] National Institute of Advanced Industrial Science and Technology (AIST), JP,
shuichi.katsumata@aist.go.jp
[3] Department of Mathematics, University of Bari, IT,
federico.pintore@uniba.it
[4] NTNU - Norwegian University of Science and Technology, Trondheim, NO ,
mattia.veroni@ntnu.no

Abstract. The digital signature schemes 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 analog, 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 aforementioned SIDH-based identification protocol. All such proofs consider the same extraction algorithm, which is claimed to always extract the 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, which we call collisions, in supersingular isogeny graphs. We provide some theoretical results on their existence, and discuss their impact on the security of the SIDH-based digital signatures. Relying on the Generalised Riemann Hypothesis, we also introduce an alternative extractor for which we rigorously prove the special soundness property.

Keywords: Post-quantum Cryptography · 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 [Sho94], isogeny-based cryptography is a newer field of elliptic-curve cryptosystems which rely on different assumptions — e.g., the hardness of the CSSI problem [DFJP14] — 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 — because it was independently designed by Couveignes [Cou06] and Rostovtsev and Stolbunov [RS06] — based on the action of ideal class groups on isomorphism classes of ordinary elliptic curves. Due to the inefficiency of the CRS scheme and the subsequent subexponential attack by Childs *et al.* [CJS14], supersingular elliptic curves were considered as a potential replacement to ordinary ones. In 2006, Charles, Lauter and Goren [CLG09] introduced a collision-resistant hash function based on supersingular isogeny graphs. Subsequently, Jao, De Feo and Plût [DFJP14] 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 [GPSBT16], Galbraith *et al.* showed that the key exchange SIDH is vulnerable to an adaptive attack when a party uses a static key, which led to the introduction of the CCA-secure Key Encapsulation Mechanism SIKE [ACC$^+$17]. SIKE is obtained from an encryption scheme which was proposed, together with an identification protocol which we will denote by $\mathsf{ID}_{\mathsf{SIDH}}$, in [DFJP14]. Both cryptosystems are deduced from the SIDH scheme.

The only existing SIDH-based digital signatures [YAJ$^+$17, GPS17] were obtained by turning the aforementioned identification protocol $\mathsf{ID}_{\mathsf{SIDH}}$[5] into a non-interactive one via the Fiat-Shamir transform [FS86] or a quantum-resistant analog, namely the Unruh transform [Unr15]. The UnForgeability against Chosen Message Attack (UF-CMA) of such digital signature schemes is claimed under the Honest-Verifier Zero-Knowledge (HVZK) and special soundness properties of the base identification protocol $\mathsf{ID}_{\mathsf{SIDH}}$.

Different proofs that $\mathsf{ID}_{\mathsf{SIDH}}$ enjoys special-soundness have been provided [DFJP14, YAJ$^+$17, GPS17]. All such proofs consider the same deterministic extractor - which we denote by $\mathsf{Ex}_{\mathsf{SIDH}}$ - that, on input two valid transcripts relative to a statement x, with the same commitment and different challenges, is claimed to output the witness w for x.

In more details, let the cyclic secret isogeny $\varphi : E_0 \longrightarrow E_1$ of prime-power degree $\ell_1^{e_1}$ be the witness w for the statement $\mathsf{x} = (E_1, \varphi(P_2), \varphi(Q_2))$, where E_0 is a fixed supersingular elliptic curve, $\{P_2, Q_2\}$ 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 $\ker(\varphi)$. Then, two valid transcripts relative to x, with the same pair of curves (E_2, E_3) as commitment and different challenges, 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 $\ell_2^{e_2}$;

[5] To be more precise, the protocol obtained by running multiple parallel executions of $\mathsf{ID}_{\mathsf{SIDH}}$.

- $\psi : E_2 \longrightarrow E_3$ has the 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 $\ker(\hat{\phi}' \circ \psi \circ \phi) \cap E_0[\ell_1^{e_1}]$ as the witness w of x.

Our contribution. In this work, we first revisit the proofs for the special soundness property of $\mathsf{ID}_{\mathsf{SIDH}}$. In particular, we depict two scenarios where the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$, given two valid transcripts relative to a statement x, with the same commitment and different challenges, produces an output which is not the witness for x, i.e. $\mathsf{Ex}_{\mathsf{SIDH}}$ fails in outputting the kernel defining the witness φ. By providing some explicit examples, we demonstrate the concrete occurrence of such scenarios, even for some of the parameter 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 $\phi(\ker(\varphi))$.

In the first of the two exception scenarios, the cyclic isogeny ψ is completely unrelated to φ, i.e. $\ker(\psi) \neq \phi(\ker(\varphi))$, despite being an $\ell_1^{e_1}$-degree isogeny from E_2 to E_3. In this case, $\ker(\hat{\phi}' \circ \psi \circ \phi) \cap E_0[\ell_1^{e_1}]$ determines an isogeny of degree $\ell_1^{e_1}$ non-equivalent with φ (its image might not even be isomorphic to E_1). In the second scenario, the cyclic isogeny ψ corresponds to an isogeny φ' rather than to φ, i.e. $\ker(\psi) = \phi(\ker(\varphi'))$ instead of being equal to $\phi(\ker(\varphi))$. The isogeny φ' is *twinned* with φ, meaning that, despite being non-equivalent with φ, it goes from E_0 to E_1, has degree $\ell_1^{e_1}$ and the kernel $\langle R \rangle$ of ϕ is 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, $\ker(\hat{\phi}' \circ \psi \circ \phi) \cap E_0[\ell_1^{e_1}]$ determines the isogeny φ' instead of the isogeny φ.

Both scenarios exploit *collisions* in supersingular isogeny graphs, i.e. non-equivalent cyclic isogenies with same degree, domain and codomain. In particular, in the first scenario the isogeny ψ forms a collision with the isogeny having $\phi(\ker(\varphi))$ 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*. We will show that the second scenario is a special case of the first one, and prove that the instances that make the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$ fail all fall within the first scenario. As a consequence, studying the existence of collisions in isogeny graphs appears to be of primary importance to determine the failure rate of the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$, and thus the security of $\mathsf{ID}_{\mathsf{SIDH}}$ and the SIDH-based digital signatures.

The relevance of collisions in supersingular isogeny graphs has emerged also in other cryptographic contexts. We can mention, for example, the adaptive attacks against SIDH [GPSBT16] and the analysis of the claw-finding attack against the isogeny-based CSSI problem [CLN+20], which is presumed to be equivalent to the SSDDH problem underlying the security of SIDH [DFJP14]. However, such collisions are usually overlooked, as they are deemed unlikely to exist.

In order to fill this gap, we conduct a rigorous analysis on the existence of collisions in supersingular isogeny graphs. We consider any large prime p and small prime ℓ, and we study the collisions formed by cyclic isogenies of degree

ℓ^e — with $e \in \mathbb{N}$ — that occur in the supersingular isogeny graph $\mathcal{G}_{p^2}(\ell)$, whose vertices are isomorphism classes of supersingular elliptic curves over $\mathbb{F}_{p^2}$ and whose edges are isogenies of degree ℓ. We express the total number of such collisions as a sum of modified Hurwitz class numbers, by relating it to traces of Brandt matrices. 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 bound, 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 [Bur62, Bur63, Bur86] 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.

The lower bounds we obtain prove that collisions do occur in isogeny graphs. On the other hand, our upper bounds suggest that, under two assumptions on the distribution of such collisions, their occurrence is too low to affect the security of the digital signature schemes obtained from $\mathsf{ID}_{\mathsf{SIDH}}$ provided the verification algorithm are modified in such a way to require λ different commitments for each signature. In fact, when considering a signature forgery for a uniformly random statement x, the probability that a collision between E_2 and E_3 exists for at least $\lambda - \lambda_0$ distinct commitments (E_2, E_3) — where λ is the security parameter of the scheme and $\lambda_0 = \log(\lambda)\log(\log(\lambda))$ — is negligible.

As current knowledge on supersingular isogeny graphs makes it hard to formally assess the two distributional assumptions that we make, we provide an alternative argument to show that the SIDH-based digital signatures remain secure despite the existence of inputs in which $\mathsf{Ex}_{\mathsf{SIDH}}$ fails. In particular, we design a new extractor $\mathsf{NEx}_{\mathsf{SIDH}}$ which, on input two valid transcripts relative to a statement x, with the same commitment and different challenges, produces an output which is *always* the witness for x. The new extractor runs in expected polynomial time under the Generalised Riemann Hypothesis and combines together slight variations of two algorithms presented in recent papers relative to isogeny-based cryptography [Wes21, FKM21]. $\mathsf{NEx}_{\mathsf{SIDH}}$ is similar to $\mathsf{Ex}_{\mathsf{SIDH}}$ but works over quaternion algebras to retrieve the witness whenever $\mathsf{Ex}_{\mathsf{SIDH}}$ fails.

As a final contribution, we use our results on the occurrence of collisions and one of the distributional assumptions to provide a formal proof for the implicit assumptions made in the context of the adaptive attacks against SIDH [GPSBT16] and the claw-finding algorithm for the CSSI problem [CLN$^+$20].

Related Work. In a recent independent work [DFDGZ21], De Feo *et al.* also dispute the existing proofs for the special soundness property of the SIDH-based identification protocol $\mathsf{ID}_{\mathsf{SIDH}}$. In particular, they show that two valid transcripts with different challenges 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 the witness for x given such transcripts. Therefore, they propose an amended version of $\mathsf{ID}_{\mathsf{SIDH}}$ at the cost of slightly more computations and heavier bandwidth.

We note that in the literature two similar notions of special soundness can be found. In the first one, the extraction algorithm is required to extract a witness given two valid transcripts relative to a statement x in the language, i.e. there must exist a witness w for x [EK18, Fis05]. In the second one, the statement x is not required to be in the language [ACK21, Unr17]. We note that the latter definition includes the former as a special case. The main reason for considering one notion instead of the other is the cryptographic application for which the identification protocol is considered. For some applications, the first definition is *enough*, and then there is no need to consider the more general definition. For example, if the considered identification protocol is turned into a digital signature scheme by applying the Fiat-Shamir transform, the first definition (together with the HVZK property) is sufficient to prove the UF-CMA security of the scheme.

As the counterexamples provided in [DFDGZ21] are for statements x not in the language, they only invalidate the special soundness proofs when considering the second definition and do not say anything about the security of the signature schemes deduced from $\mathsf{ID}_{\mathsf{SIDH}}$. Consequently, in order to construct digital signatures, there would be no reason to consider the less efficient protocol proposed in [DFDGZ21] instead of the original identification protocol.

It is only in this work that we show that $\mathsf{Ex}_{\mathsf{SIDH}}$ does not provide special soundness to the $\mathsf{ID}_{\mathsf{SIDH}}$ protocol even when considering the first definition. This implies that the proofs of security of the signature schemes deduced from the protocol needed the revision that we perform in this paper. The two arguments we provide to prove the security of the existing SIDH-based signatures — the small failure probability of $\mathsf{Ex}_{\mathsf{SIDH}}$ and the new extractor $\mathsf{NEx}_{\mathsf{SIDH}}$ — allow to target optimisation efforts on them instead of shifting to signatures based on the more costly variation of $\mathsf{ID}_{\mathsf{SIDH}}$. In the light of this, we deem [DFDGZ21] and this paper to be two complementary results, both helpful in assessing the security of the SIDH-based identification protocol $\mathsf{ID}_{\mathsf{SIDH}}$ proposed in [DFJP14].

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 depict the two scenarios in which the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$ fails and we provide some concrete counterexamples, even for some of the SIKE parameter sets. Section 4 discusses some cryptographic implications of the presented counterexamples. 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 builds upon the presented bounds to show that, under two assumptions on the distribution of collisions in supersingular isogeny graphs, the probability that the transcripts within a forged digital signature lead to an extractor failure is negligible. In Section 7 we introduce a new extraction algorithm $\mathsf{NEx}_{\mathsf{SIDH}}$ which extracts the witness of a statement x even in the cases that make $\mathsf{Ex}_{\mathsf{SIDH}}$ fail. Finally, in Section 8 we draw some conclusions.

2 Preliminaries

2.1 Identification Protocols

Let X and Y be two sets whose sizes depend 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* $\mathsf{ch} \xleftarrow{\$} \mathsf{ChSet}$*, 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.

In the following, the statement x will be sometimes removed from transcripts when it is clear from the context. 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 $\mathsf{x} \in \mathcal{L}$ 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 lies.

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 $\mathsf{x} \in \mathcal{L}$, 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 should be able to produce no more than one valid transcript relative to a statement not in the language $\mathcal{L}$. On the contrary, this limitation is not imposed by Definition 3a. As Definition 3a is contained into Definition 3b, one could ask why choose the former over the latter. The reason is that, depending on the targeted application, the first definition might suffice. For example, it is very common to use the Fiat-Shamir transform [FS86] to turn an identification protocol ID into a digital signature. If the challenge set is not exponentially large in the security parameter, one must run multiple parallel executions of the base identification protocol to achieve a negligible soundness error. Then, 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 [AABN02] 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. Since the reduction is guaranteed to be given $\mathsf{x} \in \mathcal{L}$ as its challenge, Definition 3a suffices to argue that the reduction obtains the corresponding witness w. However, since this contradicts the hardness of the language, we can conclude that no efficient adversary can break the UF-CMA security. 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 transcript relative to a statement not in the language does not affect the reduction algorithm.

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 [Sil09] for a detailed treatment of the topic.

Let $\mathbb{F}_q$ be a finite field, where q is a power of a prime $p \geq 5$. Given two elliptic curves E and E' defined over $\mathbb{F}_q$, an isogeny $\varphi : E \longrightarrow E'$ 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 the endomorphism ring of E. If $\mathsf{End}(E)$ is commutative, then E is said to be ordinary, supersingular otherwise. Every supersingular elliptic curve defined over an extension of $\mathbb{F}_p$ is isomorphic to an elliptic curve defined over $\mathbb{F}_{p^2}$.

The degree $\deg(\varphi)$ of an isogeny φ is the maximum in $\{\deg(F_1), \deg(F_2)\}$. Two elliptic curves E and E' are d-isogenous if there exists an isogeny $\varphi : E \longrightarrow E'$ of degree d, and in this case we say φ is a d-isogeny. Given a prime number ℓ and a prime power $q = p^n$ with $p \neq \ell$, we denote by $\mathcal{G}_q(\ell)$ the graph whose vertices are $\mathbb{F}_q$-isomorphism classes of supersingular elliptic curves defined over $\mathbb{F}_q$ and whose edges are equivalence classes of ℓ-isogenies, where two isogenies are equivalent if they have the same kernel. We say that two non-equivalent ℓ^e-isogenies with cyclic kernels, isomorphic domains and isomorphic codomains form a *collision* of length $2e$ in $\mathcal{G}_q(\ell)$.

The composition of two isogenies of degrees d_1 and d_2, respectively, is an isogeny of degree d_1d_2. Given an isogeny $\varphi : E \longrightarrow E'$, the dual $\hat{\varphi} : E' \longrightarrow E$ 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 $\ker(\varphi) = \langle T \rangle$, given another point $R \in E$ such that $\{T, R\}$ is a basis for $E[d]$, we have $\ker(\hat{\varphi}) = \langle \varphi(R) \rangle$ [Vit19].

Each isogeny φ has a finite kernel and, for a separable isogeny φ (i.e. an isogeny that induces a separable extension of function fields), it holds that $\deg(\varphi) = \# \ker(\varphi)$. 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 $\psi : E \longrightarrow E'$ such that $\ker(\psi) = H$. Both E/H and ψ can be computed with complexity $O(\#H)$ using Velu's formulas. When $\ker(\varphi)$ is a cyclic group, we say that φ is a cyclic isogeny.

Given a natural number ℓ, we denote by $E[\ell]$ the ℓ-torsion subgroup $\{P \in E \mid [\ell]P = 0_E\}$ of E. When ℓ and p are relatively prime, $E[\ell] \simeq (\mathbb{Z}/\ell\mathbb{Z}) \times (\mathbb{Z}/\ell\mathbb{Z})$.

Cryptographic schemes deduced from the SIDH paradigm [DFJP14] use primes $p = \ell_1^{e_1}\ell_2^{e_2}f \pm 1$, 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 defined 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, $E[\ell_1^{e_1}]$ and $E[\ell_2^{e_2}]$ are both contained in $E(\mathbb{F}_{p^2})$. We denote by pp the tuple $(\ell_1, \ell_2, e_1, e_2, f, p, E_0, P_1, Q_1, P_2, Q_2)$ of public parameters for an SIDH-based scheme, where E_0 is a fixed supersingular elliptic curve defined over $\mathbb{F}_{p^2}$ with $\#E_0(\mathbb{F}_{p^2}) = (\ell_1^{e_1}\ell_2^{e_2}f)^2$, and $\{P_1, Q_1\}$ and $\{P_2, Q_2\}$ are bases for $E_0[\ell_1^{e_1}]$ and $E_0[\ell_2^{e_2}]$, respectively.

Hardness Assumptions. We now present the hard problems which the security of $\mathsf{ID}_{\mathsf{SIDH}}$ is based on. They are tailored to the SIDH paradigm, and therefore we explicit their dependence on the tuple of SIDH parameters pp.

Problem 1 (Computational Supersingular Isogeny (CSSI) Problem, [DFJP14]). Let $\varphi : E_0 \longrightarrow E_1$ be an isogeny whose kernel is $\langle [m_1]P_1 + [n_1]Q_1 \rangle$, where m_1, n_1 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 $\varphi(P_2), \varphi(Q_2)$, the computational supersingular isogeny problem $\mathrm{CSSI}_{\mathsf{pp}}$ requires to determine a generator of the subgroup $\langle [m_1]P_1 + [n_1]Q_1 \rangle$.

The fastest algorithms to solve the above problem are 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 $\tilde{O}(\ell_1^{e_1/3})$ [DFJP14].

The next problem was originally formulated in [DFJP14] and it is believed to be computationally infeasible as well.

Problem 2 (Decisional Supersingular Product (DSSP) Problem, [YAJ+17]). Let $\varphi : E_0 \longrightarrow E_1$ be an isogeny of degree $\ell_1^{e_1}$ and kernel $\langle S \rangle$. Given (E_2, E_3, ψ) sampled with probability 1/2 from one of the following distributions, the decisional supersingular product problem $\mathrm{DSSP}_{\mathsf{pp}}$ requires to 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 $\phi : E_0 \longrightarrow E_2$ and $\phi' : E_1 \longrightarrow E_3$ be the isogenies with kernels $\langle R \rangle$ and $\langle \varphi(R) \rangle$, respectively. Then let $\psi : E_1 \longrightarrow E_2$ be the isogeny having $\langle \phi(S) \rangle$ as kernel, where $\deg(\psi) = \ell_1^{e_1}$.
- 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 $\psi : E_2 \longrightarrow E_3$ having $\langle U \rangle$ as kernel.

2.3 Quaternion Algebras and Brandt Matrices

A *quaternion algebra* B over a field $\mathbb{K}$ is a 4-dimensional central simple algebra over $\mathbb{K}$. Throughout this work, we restrict our attention to the case $\mathbb{K} = \mathbb{Q}$. Each quaternion algebra B over $\mathbb{Q}$ admits a basis $\{1, \mathbf{i}, \mathbf{j}, \mathbf{k}\}$, with $\mathbf{ij} = -\mathbf{ji} = \mathbf{k}$ and $\mathbf{i}^2 = a, \mathbf{j}^2 = b$ for some $a, b \in \mathbb{Q}^\times$. Given a quaternion $\alpha = t + x\mathbf{i} + y\mathbf{j} + z\mathbf{k} \in B$, the *canonical involution* of α is defined as $\overline{\alpha} := t - x\mathbf{i} - y\mathbf{j} - z\mathbf{k} \in B$. Via the canonical involution, we can define the *reduced trace* of α as $trd(\alpha) := \alpha + \overline{\alpha} = 2t$ and the *reduced norm* of α as $nrd(\alpha) := \alpha\overline{\alpha} = t^2 - ax^2 - by^2 + abz^2$.

For the rest of this section, let p be a fixed prime. 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 $\mathbb{Q}_p$ (the 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_ν — where ν is a place, i.e. a prime or ∞ — is either a division ring or it isomorphic to the matrix algebra $M_2(\mathbb{Q}_\nu)$. In the first case we say that B is *ramified* at ν. A quaternion algebra is uniquely determined, up to isomorphisms, by the set of places ν at which it ramifies; this set has even cardinality. Conversely, given any finite set of places of even cardinality, there is a quaternion algebra which ramifies precisely at these places.

A *fractional ideal* I in a quaternion algebra B is a (free) $\mathbb{Z}$-lattice of rank four, and it is represented as $I = \alpha_1\mathbb{Z} + \alpha_2\mathbb{Z} + \alpha_3\mathbb{Z} + \alpha_4\mathbb{Z}$ for a basis $\{\alpha_1, \alpha_2, \alpha_3, \alpha_4\}$ of B. The *norm* of a fractional ideal is defined as $n(I) = \gcd(\{nrd(\alpha) : \alpha \in I\})$. An *order* $\mathcal{O}$ in a quaternion algebra B is a subring of B that is also a fractional ideal. Since all ideals we will work with are fractional, from now on we simply call them *ideals*. For every ideal I in B, we can define the *left order* $\mathcal{O}_L(I) := \{\beta \in B : \beta I \subset I\}$ and the *right order* $\mathcal{O}_R(I) := \{\beta \in B : I\beta \subset I\}$. Saying that I is a *left* $\mathcal{O}$*-ideal* means that $\mathcal{O} = \mathcal{O}_L(I)$, and saying it is a *right* $\mathcal{O}'$*-ideal* means that $\mathcal{O}' = \mathcal{O}_R(I)$.

For the rest of this paper, we will only consider the quaternion algebra $B_{p,\infty}$ over $\mathbb{Q}$ ramified at p and ∞, since the endomorphism ring of any supersingular elliptic curve defined over $\mathbb{F}_{p^2}$ is isomorphic to a maximal order in $B_{p,\infty}$. Two ideals I, J in $B_{p,\infty}$ are *equivalent* if there exists $\beta \in B_{p,\infty}^\times$ such that $J = I\beta$. Given a maximal order $\mathcal{O}$ in $B_{p,\infty}$, the class group $cl(\mathcal{O}) = \{I_1, I_2, ...\}$ is the set of representatives of the equivalence classes of the left $\mathcal{O}$-ideals, where $I_1 = \mathcal{O}$ by convention. This set is finite, its cardinality n is called the *class number* of $B_{p,\infty}$, as it is the same for every maximal order $\mathcal{O}$ in $B_{p,\infty}$. The set $\Gamma_i := \mathcal{O}_R(I_i)^\times/\mathbb{Z}^\times$ is finite, as it is a discrete subgroup of the compact Lie group $(B_{p,\infty} \otimes \mathbb{R})^\times/\mathbb{R}^\times \cong \mathrm{SO}_3(\mathbb{R})$; let w_i be its cardinality.

We now introduce theta series and Brandt matrices; the main reference for this part is [Gro87]. Let the inverse ideal of I_i be defined as $I_i^{-1} := \{\beta \in B_{p,\infty} : I_i\beta I_i \subset I_i\}$, and $M_{ij} := I_j^{-1}I_i = \{\sum_{k=1}^{N} a_k b_k : N \in \mathbb{N}, a_k \in I_j^{-1}, b_k \in I_i\}$. We define the reduced norm $nrd(M_{ij})$ of M_{ij} to be the unique positive rational number such that all quotients $nrd(a)/nrd(M_{ij})$, for $a \in M_{ij}$, are coprime integers. The definition of *Brandt matrix* $B(m) := \left[B_{ij}(m)\right]_{1\leq i,j\leq n}$, for $m \in \mathbb{Z}$, is obtained from the following definition of theta series θ_{ij}:

$$\theta_{ij}(\tau) := \frac{1}{2w_j} \sum_{a \in M_{ij}} q^{\frac{nrd(a)}{nrd(M_{ij})}} = \sum_{m \in \mathbb{Z}} B_{ij}(m) q^m,$$

where $q := e^{2\pi i \tau}$ and $\tau \in \{z \in \mathbb{C} : \mathrm{Im}(z) > 0\}$. We hereby list some properties of Brandt matrices, which are proved in [Gro87, Proposition 2.7]:

1. If $m \geq 1$, then $B_{ij}(m) \in \mathbb{N}$, 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 [Deu41] has been enriched by several later works, the last of which being [DFKL$^+$20]. In particular,

- any supersingular j-invariant (the j-invariant of any supersingular elliptic curve) over $\mathbb{F}_{p^2}$ corresponds to a maximal order in the quaternion algebra $B_{p,\infty}$, and the set of supersingular j-invariants over $\mathbb{F}_{p^2}$ is in bijection with the class group $cl(\mathcal{O})$, for any maximal order $\mathcal{O}$ in $B_{p,\infty}$;
- endomorphisms of a supersingular elliptic curve E over $\mathbb{F}_{p^2}$ correspond to ideals of the form $\mathcal{O}a$ in $\mathcal{O} \simeq \mathrm{End}(E)$ for some $a \in \mathcal{O}$;
- ℓ-isogenies from a supersingular elliptic curve E over $\mathbb{F}_{p^2}$ correspond to left $\mathcal{O}$-ideals of norm ℓ, where $\mathrm{End}(E) \simeq \mathcal{O}$, which are integral, i.e. are contained in $\mathcal{O}$; 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 $\mathcal{O}$ in an imaginary quadratic extension of $\mathbb{Q}$, let d be its (negative) discriminant, $h(d)$ the size of the class group $cl(\mathcal{O})$ and $u(d) := \#(\mathcal{O}^\times/\mathbb{Z}^\times)$. Fix $D > 0$ and let $\mathcal{O}_{-D}$ be *the* order of discriminant $-D$. The *Hurwitz Class Number* $H(D)$ is

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

and the *modified Hurwitz Class Number* $H_p(D)$, for a prime p, is

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

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

Note that $H_p(D) \leq H(D)$. This is trivially true for the first three cases of (Equation (2)). When p divides the conductor of $\mathcal{O}_{-D}$, expanding the definition of $H(\frac{D}{p^2})$ shows that $H_p(D) = H(\frac{D}{p^2}) \leq H(D)$. We conclude this section with the following two results on Hurwitz class numbers.

Theorem 1. ***[Gro87, Prop. 1.9]*** *For all integers $m \geq 0$,*

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

Theorem 2. ***[Hur85, §7]*** *For all integers $m \geq 1$,*

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

where the sum runs over the positive divisors of m.

3 The SIDH-based Identification Protocol and Its Special Soundness

In their seminal work [DFJP14], 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 proves to a verifier that, for a given pair (E_0, E_1) of supersingular elliptic curves, it knows a cyclic isogeny $\varphi : E_0 \longrightarrow E_1$ - having degree $\ell_1^{e_1}$, where ℓ_1 is a prime - without revealing any information about the isogeny itself.

Later on, Galbraith *et al.* [GPS17] and Yoo *et al.* [YAJ$^+$17] turned $\mathsf{ID}_{\mathsf{SIDH}}$ into the first SIDH-based digital signature schemes by applying the Fiat-Shamir transform [FS86] and the Unruh transform [Unr15], respectively. The Unruh transform provides security in the quantum random oracle model, while the Fiat-Shamir transform generally guarantees security only in the classical 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 this protocol are given in [DFJP14, YAJ$^+$17, GPS17], 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 the unforgeability proofs of the signature schemes obtained by applying either the Fiat-Shamir or Unruh transform on $\mathsf{ID}_{\mathsf{SIDH}}$ need to be reviewed. In the following, we show some concrete examples of the two scenarios mentioned above, even for some of the SIKE parameter sets.

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

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

- ℓ_1 and ℓ_2 are two distinct small primes;
- e_1, e_2 are natural numbers such that $\ell_1^{e_1} \approx \ell_2^{e_2}$;
- $f \in \mathbb{N}$ is a small cofactor;
- p is a prime of the form $\ell_1^{e_1}\ell_2^{e_2} f \pm 1$;
- E_0 is a fixed supersingular elliptic curve defined over $\mathbb{F}_{p^2}$;
- $\{P_1, Q_1\}$ and $\{P_2, Q_2\}$ are two bases for $E_0[\ell_1^{e_1}]$ and $E_0[\ell_2^{e_2}]$, respectively.

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

$$X =\{(E_1, P', Q') \mid E_1 \text{ is a supersing. elliptic curve over } \mathbb{F}_{p^2}, E_1[\ell_2^{e_2}] = \langle P', Q'\rangle\}$$
$$Y =\{\varphi \mid \varphi \text{ is a cyclic isogeny from } E_0 \text{ s.t. } \deg(\varphi) = \ell_1^{e_1}\}.$$

This relation is induced by the CSSI problem and is therefore defined as follows:

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

Here we note that each statement $\mathsf{x} \in \mathcal{L}_{\mathcal{R}}$ admits a unique witness w (see [UJ18]). $\mathsf{ID}_{\mathsf{SIDH}}$ 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 $\mathbb{Z}/\ell_2^{e_2}\mathbb{Z}$, not both divisible by ℓ_2, and computes the point $R = [m_2]P_2 + [n_2]Q_2$ - 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 be the pair (m, n). Otherwise, given the point S — of order $\ell_1^{e_1}$ — generating $\ker(\varphi)$ and the isogeny $\phi : E_0 \longrightarrow E_2$ with kernel generated by $R = [m_2]P_2 + [n_2]Q_2$, it sets resp to be $\phi(S)$. It then returns resp.
- $\{0, 1\} \ni b \leftarrow \mathsf{V}((E_1, P', Q'), \mathsf{com}, \mathsf{ch}, \mathsf{resp})$: the deterministic verification algorithm works as follows:
  - if $\mathsf{ch} = 0$, and thus $\mathsf{resp} = (m, n)$, it checks that m and n are not both divisible by ℓ_2, that $E_0/\langle [m]P_2 + [n]Q_2\rangle$ is isomorphic to E_2, and that $E_1/\langle [m]P' + [n]Q'\rangle$ is isomorphic to E_3;
  - if $\mathsf{ch} = 1$, and thus $\mathsf{resp} = T$, it checks that $T \in E_2$ has order $\ell_1^{e_1}$ and that $E_2/\langle T\rangle$ is isomorphic to E_3.

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

Remark 1. We have given a description of $\mathsf{ID}_{\mathsf{SIDH}}$ in full generality. In particular, we did not specify any conditions on ℓ_1 and ℓ_2, and E_0 can be any supersingular elliptic curve over $\mathbb{F}_{p^2}$. This general setting will be considered until Section 7, where stricter conditions will be posed for ℓ_1, ℓ_2 and E_0. Such conditions reflect the design choices which led to the different SIKE parameters sets.

It is easy to show that $\mathsf{ID}_{\mathsf{SIDH}}$ is correct. We now prove 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 $\mathbb{Z}/\ell_2^{e_2}\mathbb{Z}$, 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

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

Simulated transcripts are distributed exactly as the real ones 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 (2) is hard.

In [DFJP14, YAJ+17, GPS17], 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 $(\mathsf{x}, \mathsf{com}, 0, (m, n))$ 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 $\ker(\hat{\phi}' \circ \psi \circ \phi) \cap E_0[\ell_1^{e_1}]$. Here ϕ is the isogeny from E_0 with $\langle [m]P_2 + [n]Q_2\rangle$ as kernel, ψ is the isogeny from E_2 with kernel $\langle T\rangle$ and ϕ' is the isogeny from E_1 having $\langle [m]P' + [n]Q'\rangle$ as kernel.

Remark 2. Equivalently, the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$ can be defined to output $\hat{\phi}(\langle T\rangle)$, as in [DFJP14].

We note that, under the assumption that $\langle T\rangle$ is equal to $\phi(\ker(\varphi))$ - where φ is the only witness for x - $\mathsf{Ex}_{\mathsf{SIDH}}$ extracts exactly $\ker(\varphi)$ (and so, equivalently, φ). Despite this assumption being made in [DFJP14, YAJ+17, GPS17], it does not appear to hold in general, as we show considering two different scenarios detailed below.

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 $R = [m_2]P_2 + [n_2]Q_2$ of order $\ell_2^{e_2}$ such that $E_2 = E_0/\langle R\rangle$ admits two distinct cyclic subgroups, G and $\tilde{G}$ of order $\ell_1^{e_1}$, that generate two isogenies $\psi, \tilde{\psi}$ both from E_2 to $E_3 = E_1/\langle [m_2]P' + [n_2]Q'\rangle$. The pair $(\psi, \tilde{\psi})$ forms a collision of length e_1 in the isogeny graph $\mathcal{G}_{p^2}(\ell_1)$. Here we denote by T a generator of $\tilde{G}$, by ϕ the isogeny from E_0 with kernel $\langle R\rangle$ and by ϕ' the isogeny from E_1 with kernel $\varphi(R)$, where φ is the only witness for x. We also assume that $G = \phi(\ker(\varphi))$.

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. On input such two transcripts, $\mathsf{Ex}_{\mathsf{SIDH}}$ extracts a cyclic subgroup of E_0, having order $\ell_1^{e_1}$, which defines

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

Fig. 1: Scenario 1.

an isogeny non-equivalent to the only witness φ for x. Therefore, $\mathsf{Ex}_{\mathsf{SIDH}}$ fails in extracting the 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 it.

Example 1. Consider the prime $p = (2^8)(3^5) - 1$ and the irreducible polynomial $g = x^2 + 62205x + 5 \in \mathbb{F}_p[x]$. Then $\mathbb{F}_{p^2} = \mathbb{F}_p[x]/(g)$ and we denote by z a root of g, which forms, together with the identity 1, a basis for $\mathbb{F}_{p^2}$ over $\mathbb{F}_p$. Let E_0 be 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 $E_0[3^5]$. Let E_1 be the image curve of the 2^8-isogeny φ having $\langle 33446z + 46615, 52617z + 4750\rangle$ as kernel, for which $j(E_1) = 37167z + 53117$. We consider the tuple $(E_1, P' = (6505z + 32827, 20825z + 21686), Q' = (59525z + 48254, 52332z + 7163)$ as statement x. Then, for the curves $E_2 = E_0/\langle [161]P_2 + [183]Q_2\rangle$, $E_3 = E_1/\langle [161]P' + [183]Q'\rangle$, there exist two distinct cyclic isogenies $\psi, \tilde{\psi} : E_2 \longrightarrow E_3$ of degree 2^8. The isogeny ψ is the one with kernel equal to $\phi(\ker(\varphi))$, where ϕ is the isogeny with kernel $\langle [161]P_2 + [183]Q_2\rangle$. On the other hand, $\tilde{\psi}$ has kernel generated by $T = (52195z + 35063, 51186z + 33135)$, with $T \notin \phi(\ker(\varphi))$. As a consequence, $\hat{\phi}(\langle T\rangle)$ defines an isogeny $\tilde{\varphi}$ whose image has j-invariant equal to $55144z + 45927$, which is different from that of E_1. Therefore $\tilde{\varphi}$ differs from the witness φ for x.

To produce the above concrete example for Scenario 1, the following procedure to perform an exhaustive search was used:

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 $\mathcal{G}_{p^2}(\ell_2)$

(a) compute all the paths of length e_2 (with no backtracking) and the corresponding arriving vertices;

(b) 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 of length e_1 (with no backtracking). 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, suppose there exist two distinct cyclic subgroups H and $\tilde{H}$ of order $\ell_1^{e_1}$ in E_0 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 $E_1/\langle\varphi(R)\rangle$ has the same j-invariant of $E_1/\langle\tilde{\varphi}(R)\rangle$. 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 $(\phi', \tilde{\phi}')$ form 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) in $(\mathcal{G}_{p^2}(\ell_1), \mathcal{G}_{p^2}(\ell_2))$ 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 the witness $\tilde{\mathsf{w}}$ for $\tilde{\mathsf{x}}$, rather than the witness for x. As a side note, we can use ϕ to push forward the kernels of φ and $\tilde{\varphi}$ and obtain the isogenies ψ and $\tilde{\psi}$, which form a collision of length e_1 in $\mathcal{G}_{p^2}(\ell_1)$. Below, we provide a concrete counterexample for this scenario as well.

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

Fig. 2: Scenario 2.

Example 2. Let $p = (2^8)(5^3) + 1$, and $g = x^2 + 7 \in \mathbb{F}_p[x]$ be an irreducible polynomial. We denote by z a root of g, and therefore $\{1, z\}$ is a basis over $\mathbb{F}_p$ of the extension field $\mathbb{F}_{p^2} = \mathbb{F}_p[z]/(g)$. Let E_0 be 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

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

determine the two isogenies $\varphi : E_0 \longrightarrow E_1 = E_0/H$ and $\tilde{\varphi} : E_0 \longrightarrow \tilde{E_1} = E_0/\tilde{H}$ such that $j(E_1) = j(\tilde{E_1}) = 255209z + 212204$. 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.

Remark 3. The second scenario we described above is actually a special case of the first scenario, where E_1 is isomorphic to E' over $\mathbb{F}_{p^2}$ (see Figure 1). Nevertheless, we decided to treat it separately as it allows for an alternative procedure to construct concrete examples where $\mathsf{Ex}_{\mathsf{SIDH}}$ fails in outputting the correct witness. This alternative procedure is detailed in the following lines.

The procedure we followed to produce examples where Scenario 2 occurs relies on two main parts. The first one consists in an exhaustive search for collisions of length e_1 in $\mathcal{G}_{p^2}(\ell_1)$, 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 $\mathcal{G}_{p^2}(\ell_1)$
   (a) compare all paths of length e_1 (with no backtracking) originating from the vertex j_0;
   (b) a 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;
– for any $R \in E_0[\ell_2^{e_2}]$ of order $\ell_2^{e_2}$, comparing all pairs of paths of length e_2 in $\mathcal{G}_{p^2}(\ell_2)$ determined by the points $\varphi(R)$ and $\tilde{\varphi}(R)$. A collision is found whenever the arriving j-invariant is the same for both paths in a pair. This second collision, together with the one produced in the first part, form a double collision.

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

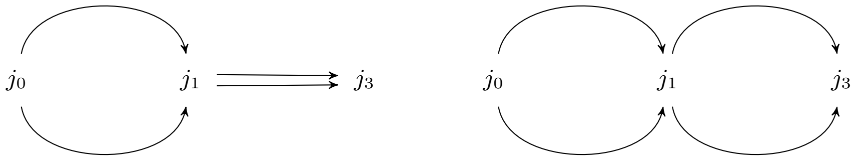

Fig. 3: A Florence flask (on the left) and a lemniscate (on the right). We stress that two non-equal colliding paths differ in at least one step, and not necessarily in all of them as one might deduce from the above figures.

We ran some experiments by considering primes p of the form $\ell_1^{e_1}\ell_2^{e_2} \pm 1$ of small size, always setting $\ell_1 = 2$ and $\ell_2 = 3$. The results are summarized in Table 1, where we list the number of vertices of $\mathcal{G}_{p^2}(2)$ from which at least a collision of length e_1 originates, the total number of collisions of length e_1 in the graph $\mathcal{G}_{p^2}(2)$, and the total number of double collisions of length (e_1, e_2) in $(\mathcal{G}_{p^2}(2), \mathcal{G}_{p^2}(3))$ (distinguishing between lemniscates and Florence Flasks). The exhaustive search conducted in both parts of the above procedure becomes very expensive already with 12-bit primes and thus, from $p = 2591$ on, we restricted our analysis to a random subset of vertices. The results show that, except for $p = 1297$, for each prime p we considered and for each vertex j in $\mathcal{G}_{p^2}(2)$ (or in the randomly-selected subgraph) exhibits at least one collision of length e_1 originating from j.

p	$\mathbf{e_1}$	$\mathbf{e_2}$	**+1/-1**	**Initial j-invariants with collisions**	**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 collisions and double collisions for the considered primes p.

3.4 Scenario 1 and SIKE Parameter Sets

One might object that the concrete examples provided in the previous subsections only occurred because of the small sizes of the considered isogeny 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 paramater set, called SIKEp751. This name reflects the fact that the underlying prime $p751 = 2^{372}3^{239} - 1$ has bit-length equal to 751. As we detail in Appendix A, the supersingular isogeny graph $\mathcal{G}_{(p751)^2}(3)$ admits two collisions of length 239, whose starting vertex is the (isomorphism 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 $\tilde{\psi}$, respectively. Let E_3 be the image curve (modulo isomorphisms) of the two isogenies. 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 the public parameters for $\mathsf{ID}_{\mathsf{SIDH}}$. Let $\varphi : E_0 \longrightarrow E_1$ be the isogeny with kernel $\overline{\phi}(H)$. By fixing bases $\{P_1, Q_1\}$, $\{P_2, Q_2\}$ for $E_0[3^{239}]$ and $E_0[2^{372}]$ respectively, $(E_1, P' = \varphi(P_2), Q' = \varphi(Q_2))$ and φ are a statement-witness pair (x, w) for $\mathsf{ID}_{\mathsf{SIDH}}$ with public parameters

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

Let R generate 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 let T generate $\tilde{H} = \ker(\tilde{\psi})$. 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.

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

Fig. 4: An example of Scenario 1 from a collision originating from E_6 (we denote by ϕ the dual isogeny of $\overline{\phi}$).

Remark 4. The constructed example could be deemed as an invalid example of Scenario 1 for SIKEp751, as the initial curve E_0 is not the one prescribed by

SIKEp751. However, if we considered the same parameter set and let the initial curve not necessarily be E_6, this example would be perfectly acceptable. Moreover, we note that constructing an example of Scenario 1 requires a collision of length 239 in the graph $\mathcal{G}_{(p751)^2}(3)$ (or of length 372 in the graph $\mathcal{G}_{(p751)^2}(2)$). It appears that the only way to compute such a collision is to exploit the knowledge of the endomorphism ring of the starting curve (as done in Appendix B by extending the procedure, based on quaternion algebras, presented in [OAT20]). As a consequence, it seems prohibitive to obtain a 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 cryptographic relevance of the constructed example.

4 Security Implications of the Two Exception Scenarios

In [DFDGZ21], 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. This means that $\mathsf{ID}_{\mathsf{SIDH}}$ does not enjoy special soundness as defined in Definition 3b. 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 — through the isogeny $\psi : E_2 \longrightarrow E_3$ — of a basis of $E_2[\ell_2^{e_2}]$. 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 [DFDGZ21] do not affect the existing proofs for 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 turning an identification protocol into a digital signature scheme using the Fiat-Shamir transform. As a consequence, according to the special soundness proofs provided in [DFJP14, YAJ$^+$17, GPS17], $\mathsf{ID}_{\mathsf{SIDH}}$ 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, given the lower efficiency of the modified scheme presented in [DFDGZ21].

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. We wonder how this weakness affects the security of the SIDH-based digital signatures. In order to answer this question, we demonstrate in the following Theorem 3 that Scenario 1 detailed in the previous section (which admits Scenario 2 as a special sub-case) is the only one that makes the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$ fail. Then, in Section 5, we will determine the probability of such a scenario occurring.

This probability relies on counting the number of single collisions of given length e in the graph $\mathcal{G}_{p^2}(\ell)$, which is interesting in its own right within the theory of isogeny graphs. Next, in Section 6.1 we will exploit the obtained results on the existence of single collisions to show that, under two general assumptions on the distribution of such collisions, for a uniformly random statement x, the probability that a single collision between E_2 and E_3 exists for at least $\lambda - \lambda_0$ different commitments (E_2, E_3) — where λ is the security parameter of the scheme and $\lambda_0 = \log \lambda \log(\log \lambda)$ — is negligible. In light of this, the security proofs for the SIDH-based digital signatures remain reliable despite the existence of inputs in which $\mathsf{Ex}_{\mathsf{SIDH}}$ fails, provided the verification algorithm are modified in such a way to require λ different commitments for each signature.

However, current knowledge on supersingular isogeny graphs makes it hard to formally assess the two distributional assumptions we introduce. To further corroborate the security of the SIDH-based digital signatures, we provide an alternative extractor $\mathsf{NEx}_{\mathsf{SIDH}}$, which runs in expected polynomial time under the Generalised Riemann Hypothesis. On input two valid transcripts relative to a statement x, with the same commitment and different challenges, the new extractor always outputs the witness for x. We will detail $\mathsf{NEx}_{\mathsf{SIDH}}$ in Section 7.

Theorem 3. *Given a statement* $\mathsf{x} = (E_1, P', Q') \in \mathcal{L}_{\mathcal{R}}$ *and a commitment* $\mathsf{com} = (E_2, E_3)$*, if* $\mathsf{t}_1 = (\mathsf{x}, \mathsf{com}, \mathsf{ch} = 0, (m, n))$ *and* $\mathsf{t}_2 = (\mathsf{x}, \mathsf{com}, \mathsf{ch} = 1, T)$ *are valid transcripts relative to* x *that make* $\mathsf{Ex}_{\mathsf{SIDH}}$ *fail, then they form an instance of Scenario 1.*

Proof. Let $\phi : E_0 \longrightarrow E_2$ be the isogeny with kernel $\langle [m]P_2 + [n]Q_2 \rangle$, where $E_0[\ell_2^{e_2}] = \langle P_2, Q_2 \rangle$. Since x belongs to the language $\mathcal{L}_{\mathcal{R}}$, there exists a cyclic $\ell_1^{e_1}$-isogeny $\varphi : E_0 \longrightarrow E_1$ such that $\varphi(P_2) = P'$ and $\varphi(P_2) = Q'$. Therefore, for any valid transcript $\mathsf{t}_1 = (\mathsf{x}, \mathsf{com}, \mathsf{ch} = 0, (m, n))$, the transcript $\mathsf{t}_2' = (\mathsf{x}, \mathsf{com}, \mathsf{ch} = 1, \phi(\ker(\varphi))$ is also valid. If $T = \phi(\ker(\varphi))$ (i.e. $\mathsf{t}_2 = \mathsf{t}_2'$) the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$ does not fail on input t_1 and t_2. Thus, T must be different from $\phi(\ker(\varphi)$, and the two cyclic $\ell_1^{e_1}$-isogenies $\psi, \psi' : E_2 \longrightarrow E_3$ having T and $\phi(\ker(\varphi)$ as respective kernels form a collision of length e_1 in $\mathcal{G}_{p^2}(\ell_1)$. In other words, t_1 and t_2 are an instance of Scenario 1.

5 Quantitative Study of Cycles in Isogeny Graphs

As detailed in Section 3.2 and in Section 3.3, there are cases in which the existence of collisions in supersingular isogeny graphs poses a problem to the success of 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 are therefore interested in quantifying these collisions. Our approach builds upon the techniques used in [Gha21], but here a different kind of collisions is considered. Furthermore, we deepen the analysis in order to provide heuristic lower bounds in addition to deterministic upper bounds.

In the rest of the section, p will 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 begin by formally stating the problem we focus our attention on, which is related to the existence of single collisions.

Problem 3. Fix primes $\ell < p$, and let E_0 be a uniformly random supersingular elliptic curve (i.e. a vertex) in the graph $\mathcal{G}_{p^2}(\ell)$. Given $e \in \mathbb{N}$, determine the expected number of single collisions of length e in $\mathcal{G}_{p^2}(\ell)$ starting from E_0.

We will be able to provide an upper bound for the expectation the above problem requires to compute. Finding a *good* lower bound, on the other hand, appears to be beyond the reach of current analytic number theory, as it would involve bounds on incomplete character sums [IK04, Chap. 12] that are tighter than both known and conjectured results.

Given two points in $E_0[\ell^e]$, we will consider them equivalent w.r.t. $\sim$ if they generate the same torsion subgroup, i.e. if they generate equivalent isogenies. Let $E_0[\ell^e]_{\max}$ denote the classes of points of maximal order in $E_0[\ell^e]/_\sim$. We can rephrase the above statement as follows: compute the expected number of points P_A, P_B in different classes of $E_0[\ell^e]_{\max}$ such that $E_0/\langle P_A\rangle \cong E_0/\langle P_B\rangle$.

Let $\mathscr{C}_{E_0}(\ell^{2e})$ denote the number of endomorphisms with no backtracking (i.e. with cyclic kernel) of degree ℓ^{2e} originating from E_0 in $\mathcal{G}_{p^2}(\ell)$. Let also $\mathsf{Coll}_{\ell^e}(E_0)$ be the cardinality of the set of single collisions (f,g) in $\mathcal{G}_{p^2}(\ell)$ of length e and starting at E_0. Following the above discussion, we see that $\mathsf{Coll}_{\ell^e}(E_0)$ is the cardinality of the set

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

Remark 5. In the above formula, we mod out by the equivalence relation $\sim$ because we regard (f,g) and (g,f) as the same single collision. Note that k isogenies between E and E' would be counted as $\binom{k}{2}$ single collisions.

Let n be the number of vertices in the graph $\mathcal{G}_{p^2}(\ell)$ and $\{E_{(i)}\}_{i=1,\dots,n}$ be representatives of the vertices of $\mathcal{G}_{p^2}(\ell)$. By definition, to obtain $\mathscr{C}_{E_{(i)}}(\ell^{2e})$, we need to subtract from $B_{ii}(\ell^{2e})$ — the entry of the Brandt matrix corresponding to all endomorphisms of $E_{(i)}$ of degree ℓ^{2e} — the number of endomorphisms of $E_{(i)}$ of degree ℓ^{2e} with non-cyclic kernel, i.e. isogenies corresponding to paths with backtracking. This number is simply equal to $B_{ii}(\ell^{2e-2})$. Indeed, since composing an isogeny of degree ℓ with its dual gives the scalar isogeny $[\ell]$, 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,

$$\begin{aligned}\sum_{i=1}^{n} \mathscr{C}_{E_{(i)}}(\ell^{2e}) &= \frac{1}{2}\sum_{i=1}^{n}\left(B_{ii}(\ell^{2e}) - B_{ii}(\ell^{2e-2})\right)\\ &= \frac{1}{2}\mathrm{Tr}\left(B(\ell^{2e})\right) - \frac{1}{2}\mathrm{Tr}\left(B(\ell^{2e-2})\right) \qquad (3)\\ &= \frac{1}{2}\sum_{s^2 \le 4\ell^{2e}} H_p(4\ell^{2e} - s^2) - \frac{1}{2}\sum_{s^2 \le 4\ell^{2e-2}} H_p(4\ell^{2e-2} - s^2).\end{aligned}$$

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. In addition, $H_p(0)$

appears in both sums, so we can cancel it out. Hence,

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

Remark 6. As an aside, note that the above number $\sum_{i=1}^{n} \mathscr{C}_{E_{(i)}}(\ell^{2e})$ is not the total number of cycles (closed paths, with no backtracking) of length $2e$ 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 now relate the number of endomorphisms with no backtracking $\mathscr{C}_{E_0}(\ell^{2e})$ to the desired number of single collisions $\mathsf{Coll}_{\ell^e}(E_0)$. This allows us to turn our attention to estimating $\mathscr{C}_{E_0}(\ell^{2e})$.

Lemma 1. *It is possible to bound the number* $\mathsf{Coll}_{\ell^e}(E_0)$ *using* $\mathscr{C}_E(\ell^{2e})$ *as follows*

$$\mathscr{C}_{E_0}(\ell^{2e}) \leq \mathsf{Coll}_{\ell^e}(E_0) \leq \mathscr{C}_{E_0}(\ell^{2e}) + \sum_{r=1}^{e-1} \mathscr{C}_{E_0}(\ell^{2r})(\ell-1)\ell^{e-1-r}. \qquad (5)$$

Proof. The first inequality is clear since, by definition, every single collision also constitutes an endomorphism with no backtracking. However, there are some endomorphisms with backtracking that still constitute single collisions as in Figure 5. The single collisions that form endomorphisms *with* backtracking can be

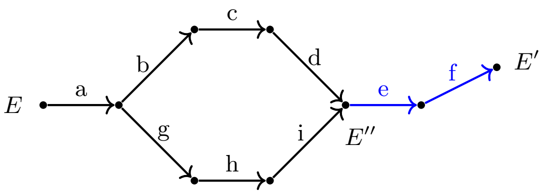

Fig. 5: Two paths ("abcdef" and "aghief") that constitute an endomorphism ("abcdeffeihga") with backtracking and also a single collision.

decomposed into two parts: a first part that is an endomorphism without backtracking ("abcdihga" in Figure 5) and second part at the end that constitutes the backtracking ("effe" in Figure 5).

Therefore, to take the endomorphisms with backtracking into account when counting single collisions, we simply need to count all endomorphisms without backtracking of different possible degrees $\mathscr{C}_{E_0}(\ell^{2r})$, 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 $e-r$ from an elliptic curve E'' to another one E' without backtracking, as in Figure 5 where $r = 4$. Note that this takes into account the fact that the curve E'' already has two edges coming into it: this is why we have $(\ell - 1)\ell^{e-1-r}$ instead of simply ℓ^{e-r}.

In order to compute the bounds in (5) using Equation (4), by definition of H_p, one would need to determine the average splitting behaviour of $4\ell^{2e} - s^2$ in $\mathbb{F}_p$. One option when computing upper bounds (as in Section 5.1) for the number of collisions is to use 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 bound: zero. To remedy this issue, the deterministic splitting of $4\ell^{2e} - s^2$ in $\mathbb{F}_p$ is replaced by a random Bernoulli process in Section 5.2. In particular, the Legendre symbols $\left(\frac{4\ell^{2e}-s^2}{p}\right)$ are modelled as Bernoulli random variables. This does not give an exact value for $\sum_{i=1}^{n} \mathscr{C}_{E_{(i)}}(\ell^{2e})$, but rather an estimate for it.

5.1 Upper Bounds

We hereby prove an upper bound for the number of endomorphisms in $\mathcal{G}_{p^2}(\ell)$ of degree ℓ^{2e} and without backtracking, which will allow us to bound the number of single collisions of length e in the same graph.

Lemma 2. *The following upper bound holds:*

$$\sum_{i=1}^{n} \mathscr{C}_{E_{(i)}}(\ell^{2e}) \leq \frac{\ell^{2e+1}}{\ell - 1}.$$

Proof. From equation (4), we have

$$\begin{aligned}
\sum_{i=1}^{n} \mathscr{C}_{E_{(i)}}(\ell^{2e}) &= \frac{1}{2} \sum_{s^2 < 4\ell^{2e}} H_p(4\ell^{2e} - s^2) - \frac{1}{2} \sum_{s^2 < 4\ell^{2e-2}} H_p(4\ell^{2e-2} - s^2) \\
&\overset{(\diamondsuit)}{\leq} \frac{1}{2} \sum_{s^2 < 4\ell^{2e}} H_p(4\ell^{2e} - s^2) \\
&\overset{(\heartsuit)}{\leq} \frac{1}{2} \sum_{s^2 < 4\ell^{2e}} H(4\ell^{2e} - s^2) \\
&= \frac{1}{2} \sum_{s^2 \leq 4\ell^{2e}} H(4\ell^{2e} - s^2) - H(0) \qquad (6) \\
&= \frac{1}{2} \left(2 \sum_{d \mid \ell^{2e}} d - \sum_{d \mid \ell^{2e}} \min\left(d, \frac{\ell^{2e}}{d}\right) \right) + \frac{1}{12} \\
&= \sum_{i=0}^{2e} \ell^i - \sum_{i=0}^{e-1} \ell^i - \frac{1}{2}\ell^e + \frac{1}{12} \\
&= \frac{\ell^{2e+1} - \ell^e}{\ell - 1} - \frac{1}{2}\ell^e + \frac{1}{12}.
\end{aligned}$$

Thus,

$$\sum_{i=1}^{n} \mathscr{C}_{E_{(i)}}(\ell^{2e}) \leq \frac{\ell^{2e+1}}{\ell - 1} \in O(\ell^{2e}).$$

Remark 7. The only places were we use an upper bound instead of an equality in (6) are ($\diamondsuit$) and ($\heartsuit$). At ($\diamondsuit$), we drop a term which is $O(\ell^{2e-2})$. At ($\heartsuit$), we bound $H_p(D)$ above by $H(D)$. The definition of $H_p(D)$ depends on the splitting behaviour of p. It is thus reasonable to expect that $H_p(D)$ is equal to zero half of the time and to $H(D)$ the other half, which would mean that $H(D) = 2H_p(D)$ on average; we discuss this issue further in Section 5.2. Therefore, one can expect the R.H.S. of ($\heartsuit$) to be, on average, twice as big as the L.H.S.

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

Theorem 4. *The number of single collisions of length e in the graph $\mathcal{G}_{p^2}(\ell)$ is bounded above as follows:*

$$\sum_{i=1}^{n} \mathsf{Coll}_{\ell^e}(E_{(i)}) \leq \frac{\ell^{2e}(\ell+1)}{\ell - 1}. \tag{7}$$

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

$$\sum_{i=1}^{n} \mathscr{C}_{E_{(i)}}(\ell^{2r}) \leq \frac{\ell^{2r+1}}{\ell - 1}.$$

Combining this with Lemma 1, we obtain

$$\begin{aligned}
\sum_{i=1}^{n} \mathsf{Coll}_{\ell^e}(E_{(i)}) &\leq \sum_{i=1}^{n} \mathscr{C}_{E_{(i)}}(\ell^{2e}) + \sum_{r=1}^{e-1}\sum_{i=1}^{n} \mathscr{C}_{E_{(i)}}(\ell^{2r})(\ell-1)\ell^{e-1-r} \\
&\leq \frac{\ell^{2e+1}}{\ell-1} + \ell^e \sum_{r=1}^{e-1} \ell^r \\
&= \frac{\ell^{2e+1} + \ell^{2e} - \ell^{e+1}}{\ell - 1} \\
&\leq \frac{\ell^{2e}(\ell+1)}{\ell-1}.
\end{aligned}$$

Corollary 1. *The expected number of single collisions of length e starting at a uniformly random vertex E of $\mathcal{G}_{p^2}(\ell)$ is bounded above as follows:*

$$\mathbb{E}_E[\mathsf{Coll}_{\ell^e}(E)] := \frac{1}{n}\sum_{i=1}^{n} \mathsf{Coll}_{\ell^e}(E_{(i)}) \leq \frac{\ell^{2e}(\ell+1)}{n(\ell-1)}. \tag{8}$$

Consequently, if $p \approx \ell^{2e}$, the expectation $\mathbb{E}_E[\mathsf{Coll}_{\ell^e}(E)]$ is $O(1)$.

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

5.2 Lower Bounds

Let us fix a prime ℓ and pick an arbitrary prime p such that $p \approx \ell^{2e}$ for some $e \in \mathbb{N}$. We want to show that $\mathbb{E}_E\left[\mathsf{Coll}_{\ell^e}(E)\right]$ is non negligible. By Lemma 1, we only need to focus on $\mathscr{C}_E(\ell^{2e})$ in order to bound $\mathsf{Coll}_{\ell^e}(E)$. Estimating the expected value of $\mathscr{C}_E(\ell^{2e})$ requires the computation of $\sum_{s^2<4\ell^{2e}} H_p(4\ell^{2e}-s^2)$, in light of Equation (4). The difficulty lies in relating the sums of modified Hurwitz class numbers $H_p(\cdot)$ to the sums of Hurwitz class numbers $H(\cdot)$. As noted in Remark 7, our hypothesis is that the actual value of $\sum_{s^2<4\ell^{2e}} H_p(4\ell^{2e}-s^2)$ is roughly $\frac{1}{2}\sum_{s^2<4\ell^{2e}} H(4\ell^{2e}-s^2)$. However, we cannot prove this as it involves sums of Legendre symbols, and the best bounds for incomplete character sums that are known (based on the bounds of Burgess in [Bur62, Bur63, Bur86]) are far from being tight enough for what we need.

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

$$\varepsilon_p(D) := \begin{cases} 0 \text{ if } p \text{ splits in } \mathcal{O}_{-D}; \\ 1 \text{ if } p \text{ is inert in } \mathcal{O}_{-D}; \\ \frac{1}{2} \text{ if } p \text{ is ramified in } \mathcal{O}_{-D} \\ \quad \text{but does not divide the conductor of } \mathcal{O}_{-D}; \end{cases} \tag{9}$$

so that

$$H_p(4\ell^{2e}-s^2) = \varepsilon_p(4\ell^{2e}-s^2)H(4\ell^{2e}-s^2), \tag{10}$$

for all $4\ell^{2e} > s^2$. Note that in Equation (9) 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 $D := 4\ell^{2e}-s^2$ with $4\ell^{2e} > s^2$ 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 replace $\varepsilon_p(4\ell^{2e}-s^2)$ with Bernoulli random variables, in order to get estimates on H_p and thus on $\mathbb{E}_E[\mathsf{Coll}_{\ell^e}(E)]$. 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

$$H^*(4\ell^{2e}-s^2) = X_s \cdot H(4\ell^{2e}-s^2), \tag{11}$$

for all $s = 0, 1, ..., 2\ell^e$. Our goal is for H^* to mimic H_p 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*

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

Proof. We can easily compute

$$\begin{aligned}\mathbb{E}\left[\sum_{s^2\le 4\ell^{2e}} H^*(4\ell^{2e}-s^2)\right] &= \sum_{s^2\le 4\ell^{2e}} H(4\ell^{2e}-s^2)\mathbb{E}[X_s] \\ &= \sum_{i=0}^{2e}\ell^i - \frac{1}{2}\sum_{i=0}^{2e}\min\{\ell^i,\ell^{2e-i}\} \qquad (12) \\ &= \frac{\ell^{2e+1}-\ell^e}{\ell-1} - \frac{1}{2}\ell^e.\end{aligned}$$

The next proposition investigates how much the sum $\sum_{s^2\le 4\ell^{2e}} H^*(4\ell^{2e}-s^2)$ can deviate from its expectation, obtained in Lemma 3. We make use of *concentration inequalities*, such as Hoeffding's bound (see Proposition 2.5 in [Wai19])[6].

Proposition 1. *Let* $\mu := \mathbb{E}\left[\sum_{s^2\le 4\ell^{2e}} H^*(4\ell^{2e}-s^2)\right]$. *Given* $\varepsilon > 0$, *the following inequality holds*

$$\mathbb{P}\left(\left|\frac{\sum_{s^2\le 4\ell^{2e}} H^*(4\ell^{2e}-s^2)}{\mu} - 1\right| \le \varepsilon\right) \ge 1 - \exp\left(-\varepsilon^2/2\right). \qquad (13)$$

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

$$\begin{aligned}&\mathbb{P}\left(\left|\frac{\sum_{s^2\le 4\ell^{2e}} H^*(4\ell^{2e}-s^2)}{\mu} - 1\right| \le \varepsilon\right) \\ &= \mathbb{P}\left(\left|\sum_{s^2\le 4\ell^{2e}} H(4\ell^{2e}-s^2)X_s - \mu\right| \le \varepsilon\mu\right) \\ &\ge 1 - \exp\left(-\frac{2\varepsilon^2\mu^2}{\sum_{s^2\le 4\ell^{2e}} H(4\ell^{2e}-s^2)^2}\right) \\ &\overset{(\clubsuit)}{>} 1 - \exp\left(-\frac{2\varepsilon^2\mu^2}{\left(\sum_{s^2\le 4\ell^{2e}} H(4\ell^{2e}-s^2)\right)^2}\right) \\ &= 1 - \exp\left(-\varepsilon^2/2\right).\end{aligned}$$

This shows the quantity $\sum_{s^2\le 4\ell^{2e}} H^*(4\ell^{2e}-s^2)$ is close to

$$\mu = \frac{1}{2}\sum_{s^2\le 4\ell^{2e}} H(4\ell^{2e}-s^2) \approx p \qquad (14)$$

[6] There exists a version of the central limit theorem — for a triangular array of random variables — that allows for the setup we are in. However, this version requires the Lyapunov or Lindeberg condition, which is not obvious to prove (see Remark 8).

(see Lemma 3) with a positive probability. For example, when ε is equal to 0.5, $\sum_{s^2 \leq 4\ell^{2e}} H^*(4\ell^{2e} - s^2)$ is between 0.5μ and 1.5μ with probability $\geq 11\%$. Therefore, Proposition 1 provides more evidence to the fact that we can estimate $\sum_{s^2 \leq 4\ell^{2e}} H_p(4\ell^{2e} - s^2)$ with $\mu = \frac{1}{2} \sum_{s^2 \leq 4\ell^{2e}} H(4\ell^{2e} - s^2)$.

Remark 8. The bound in Proposition 1 appears not to be tight. In fact, using Markov's inequality instead of Hoeffding's inequality gives slightly better bounds for certain values of ε (but slightly worse bounds for other values of ε). Furthermore, in step (♣) of the proof, we bound $\frac{\sum_{s^2 \leq 4\ell^{2e}} H(4\ell^{2e}-s^2)^2}{\left(\sum_{s^2 \leq 4\ell^{2e}} H(4\ell^{2e}-s^2)\right)^2}$ above by 1. However, this quantity tends to zero as p goes to infinity. Indeed, by Equation (14), we only need to explain why $\sum_{s^2 \leq 4\ell^{2e}} H(4\ell^{2e}-s^2)^2$ is in $o(p^2) = o(\ell^{4e})$. We prove that in Appendix B.

Our goal is to give a lower bound for the number $\sum_{i=1}^{n} \mathsf{Coll}_{\ell^e}(E_{(i)})$ of single collisions of length e in $\mathcal{G}_{p^2}(\ell)$, using the Bernoulli model, as in Equation (11). To do so, we turn our attention to bounding the number of endomorphisms $\sum_{i=1}^{n} \mathscr{C}_{E_{(i)}}(\ell^{2e})$ of degree ℓ^{2e} from below (see Lemma 1). Let us define

$$\mathscr{C}^*(\ell^{2e}) := \frac{1}{2} \sum_{s^2 < 4\ell^{2e}} H^*(4\ell^{2e} - s^2) - \frac{1}{2} \sum_{s^2 < 4\ell^{2e-2}} H^*(4\ell^{2e-2} - s^2) \tag{15}$$

which should approximate the total number of endomorphisms $\sum_{i=1}^{n} \mathscr{C}_{E_i}(\ell^{2e})$ expressed in (Equation (4)). As in (Equation (12)), we can compute

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

We summarize the above results with the following theorem.

Theorem 5. *Using the Bernoulli model introduced in Equation (11), the expected number of single collisions of length e in any given supersingular isogeny graph $\mathcal{G}_{p^2}(\ell)$ is bounded below by $\frac{1}{4}\ell^{e-1}(\ell+1)(2\ell^e - 1)$. In particular, this model predicts the expected number of single collisions of length e starting at a uniformly random vertex in $\mathcal{G}_{p^2}(\ell)$ to be bounded below by $\frac{1}{4n}\ell^{e-1}(\ell + 1)(2\ell^e - 1)$, which is $\Omega(1)$.*

6 Security of the SIDH-based Digital Signatures

The goal of this Section is twofold. We will start by studying the single collisions occurring between two fixed vertices in the graph $\mathcal{G}_{p^2}(\ell)$. Thereafter we will, under some mild assumptions, study the probability of having more collisions than a certain value between a subset of vertices of $\mathcal{G}_{p^2}(\ell)$. This is directly related to the failure rate of the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$. Secondly, as an application, we provide a formal justification to a claim made in the literature about the unlikely existence of certain collisions.

6.1 The failure rate of $\mathsf{Ex}_{\mathsf{SIDH}}$

When bounding the expected number of collisions in Section 5, we considered a uniformly random vertex as the starting point, and allowed the image curve to be any vertex on the graph. However, one can also consider the expected number of collisions between a fixed vertex and a uniformly random vertex of the graph. This drastically changes the order of magnitude of the expectation, as proved in the following Theorem 6. Such result relies on the assumption that the collisions of length e in $\mathcal{G}_{p^2}(\ell)$ are evenly distributed among the vertices of the graph. In particular, for any vertex $E \in \mathcal{G}_{p^2}(\ell)$, this would mean that:

$$\mathsf{Coll}_{\ell^e}(E) = \frac{1}{n}\sum_{i=1}^{n} \mathsf{Coll}_{\ell^e}(E_{(i)}). \tag{16}$$

By considering Theorem 4, we obtain the following reformulation.

Assumption 1 *Single collisions of length e in $\mathcal{G}_{p^2}(\ell)$ are evenly distributed among the vertices of the graph, i.e. for any $E \in \mathcal{G}_{p^2}(\ell)$,*

$$\mathsf{Coll}_{\ell^e}(E) \leq \frac{1}{n}\frac{\ell^{2e}(\ell+1)}{\ell-1}. \tag{17}$$

Let $\mathcal{D}_\ell(E;e)$ be the set of distinct elliptic curves that are connected to a vertex E via a path of length e without backtracking in $\mathcal{G}_{p^2}(\ell)$. Let also $\mathsf{Coll}_{\ell^e}(E_2, E_3)$ denote the number of single collisions of length e between $E_2, E_3 \in \mathcal{G}_{p^2}(\ell)$.

Theorem 6. *Let $E_2 \in \mathcal{G}_{p^2}(\ell)$ and $E_3 \in \mathcal{D}_\ell(E_2;e)$. Then, under Assumption 1, when $n = \#\mathcal{G}_{p^2}(\ell) \approx \ell^{2e}$, we have*

$$\mathbb{E}_{E_3}[\mathsf{Coll}_{\ell^e}(E_2, E_3)] \in O(1/\sqrt{p}).$$

Proof. Denote by m_{E_2} the cardinality of $\mathcal{D}_\ell(E_2;e)$. Let $\mathfrak{m}_{E_2} = \ell^e + \ell^{e-1}$ be the number of non-equivalent cyclic isogenies of degree ℓ^e starting from E_2. By Assumption 1, $\mathfrak{m}_{E_2} - m_{E_2} \leq \frac{1}{n}\frac{\ell^{2e}(\ell+1)}{\ell-1}$. The expectation of $\mathsf{Coll}_{\ell^e}(E_2, E_3)$ over

$E_3 \in \mathcal{D}_\ell(E_2; e)$ can be bounded above as follows:

$$\begin{aligned}
\mathbb{E}_{E_3}[\mathsf{Coll}_{\ell^e}(E_2, E_3)] &:= \frac{1}{m_{E_2}} \sum_{E_3 \in \mathcal{D}_\ell(E_2;e)} \mathsf{Coll}_{\ell^e}(E_2, E_3) \\
&= \frac{1}{m_{E_2}} \mathsf{Coll}_{\ell^e}(E_2) \\
&\leq \frac{1}{m_{E_2}} \frac{1}{n} \frac{\ell^{2e}(\ell+1)}{\ell-1} \\
&\leq \frac{1}{\ell^e + \ell^{e-1} - \frac{1}{n}\frac{\ell^{2e}(\ell+1)}{\ell-1}} \frac{1}{n} \frac{\ell^{2e}(\ell+1)}{\ell-1} \\
&= \frac{\ell^{2e}}{n\ell^{e-1}(\ell-1) - \ell^{2e}}.
\end{aligned} \tag{18}$$

Therefore, when $n \approx \ell^{2e}$, we conclude that $\mathbb{E}_{E_3}[\mathsf{Coll}_{\ell^e}(E_2, E_3)] \in O(1/\sqrt{p})$

We recall that in the sigma protocol $\mathsf{ID}_{\mathsf{SIDH}}$, once E_0 and E_2 are fixed, the curves E_3 and E_1 uniquely determine each other. This allows us to define the following random variable on E_1 in terms of the corresponding E_3:

$$\chi_{E_2}(E_1) := \mathbb{1}_{[\exists \text{ collision from } E_2 \text{ to } E_3]}.$$

To assess the security of a digital signature Π deduced from $\mathsf{ID}_{\mathsf{SIDH}}$ by applying the Fiat-Shamir transform (or a quantum-secure variant), we focus on bounding $\mathbb{P}_{E_1}(\sum_{E_2 \in \mathcal{D}_{\ell_2}(E_0;e_2)} \chi_{E_2}(E_1) \geq \lambda - \lambda_0)$, where $\lambda = \Theta(\log(p))$ is the security parameter and $\lambda_0 = \log(\lambda)\log(\log(\lambda))$. In the following lines we explain our motivation, and our reasoning holds provided that the verification algorithms of the digital signatures are slightly modified in order to accept as valid only signatures with distinct commitments. This does not introduce any significant modification, as honest signers would consider different commitments with overwhelming probability.

In the usual security reduction, which is given in the Random Oracle Model, an adversary $\mathcal{A}$ in the UF-CMA game for Π is run as a subroutine by an adversary $\mathcal{B}$ against the CSSI problem. As such, $\mathcal{A}$ receives as input a triplet $\mathsf{x} = (E_1, P', Q')$, where E_1 is the image of an $\ell_1^{e_1}$-isogeny $\varphi : E_0 \longrightarrow E_1$ having $\langle [m_1]P_1 + [n_1]Q_1 \rangle$ as kernel — for uniformly random $m_1, n_1 \in \mathbb{Z}/\ell_1^{e_1}\mathbb{Z}$ not both divisible by ℓ_1 — and $P' = \varphi(P_2)$, $Q' = \varphi(Q_2)$. Here, we recall that $\{P_1, Q_1\}$ and $\{P_2, Q_2\}$ are bases of $E_0[\ell_1^{e_1}]$ and $E_0[\ell_2^{e_2}]$, respectively. Let σ be a valid forgery produced by $\mathcal{A}$. Then $\sigma = (\overline{\mathsf{com}}, \overline{\mathsf{ch}}, \overline{\mathsf{resp}})$ where $\overline{\mathsf{com}} = (\mathsf{com}_1, \ldots, \mathsf{com}_\lambda)$, $\overline{\mathsf{resp}} = (\mathsf{resp}_1, \ldots, \mathsf{resp}_\lambda)$ and $\overline{\mathsf{ch}} = (\mathsf{ch}_1, \ldots, \mathsf{ch}_\lambda) \in \{0,1\}^\lambda$ is the output of the Random Oracle on input the signed message m and $\overline{\mathsf{com}}$. We note that, for each $i \in \{1, \ldots, \lambda\}$, $(\mathsf{x}, \mathsf{com}_i, \mathsf{ch}_i, \mathsf{resp}_i)$ is a valid transcript relative to x.

We denote by T the number of the transcripts within σ that were obtained by guessing the corresponding challenges, i.e. without exploiting any knowledge of the witness w for x. In other words, these transcripts were constructed by guessing T zero entries of $\overline{\mathsf{ch}}$. Then $T \in O(\log(\lambda))$, since $\mathcal{A}$ is a polynomial-time adversary whose advantage in non-negligible. The adversary is then run a second time in order to obtain a second forgery $\sigma' = (\overline{\mathsf{com}}, \overline{\mathsf{ch}}', \overline{\mathsf{resp}}')$. With some probability, the Hamming distance between $\overline{\mathsf{ch}}$ and $\overline{\mathsf{ch}}'$ is non-zero. The transcripts corresponding to one of the positions where $\overline{\mathsf{ch}}$ and $\overline{\mathsf{ch}}'$ differ make $\mathsf{Ex}_{\mathsf{SIDH}}$ extract the solution of the CSSI problem, unless the two transcripts form an instance of Scenario 1, i.e. there is a collision between E_2 and E_3, where $\mathsf{com} = (E_2, E_3)$ is the common commitment of the two transcripts.

In Theorem 7, we prove that $\mathbb{P}_{E_1}(\sum_{E_2 \in \mathcal{D}_{\ell_2}(E_0;e_2)} \chi_{E_2}(E_1) \geq \lambda - \lambda_0)$ is negligible when $\lambda_0 = \log(\lambda)\log(\log(\lambda))$. In other words, out of the λ pairs of transcripts, with overwhelming probability, at most $\lambda - \lambda_0$ make the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$ fail. Therefore, the reduction will extract the witness w from the two forgeries σ, σ' with overwhelming probability, since at least $\lambda_0 - T$ *good* commitments that allow to extract the exact witness will exist. In fact, $\lambda_0 \in \omega(\log(\lambda))$ and therefore $\lambda_0 - T$ is larger than 0 for any large enough lambda.

Remark 9. The security argument described above remains valid if the function λ_0 is replaced by any other function in $\omega(\log \lambda)$ for which Theorem 7 still holds.

Let $\bar{\chi}(E_1) := \sum_{E_2 \in \mathcal{D}_{\ell_2}(E_0;e_2)} \chi_{E_2}(E_1)$ be the number of commitments for the statement $(E_1, \varphi(P), \varphi(Q))$ forming an instance of Scenario 1.

Theorem 7. *Let $p = \ell_1^{e_1}\ell_2^{e_2} f \pm 1$ a prime (with ℓ_1, ℓ_2 small primes, and $f \in \mathbb{N}$ a small factor coprime to ℓ_1 and ℓ_2) such that $\ell_1^{e_1} \approx \ell_2^{e_2}$, $\lambda \in \Theta(\log(p))$ and $\lambda_0 = \log(\lambda)\log(\log(\lambda))$. If the random variables $\{\chi_{E_2} : E_2 \in \mathcal{D}_{\ell_2}(E_0; e_2)\}$ are independent then, under Assumption 1, the quantity $\mathbb{P}_{E_1}(\bar{\chi} \geq \lambda - \lambda_0)$ is negligible.*

Proof. Let m_{E_2}, m_{E_1} be the number of curves in $\mathcal{D}_{\ell_2}(E_0, e_2), \mathcal{D}_{\ell_1}(E_0, e_1)$ respectively. Given $E_2 \in \mathcal{D}_{\ell_2}(E_0, e_2)$, we define $\alpha_{E_2} := \#\{E_1 \in \mathcal{D}_{\ell_1}(E_0, e_1) : \chi_{E_2}(E_1) = 1\}$, and $\alpha := \max_{E_2}\{\alpha_{E_2}\}$. We have that $\alpha \leq \frac{1}{n}\frac{\ell_1^{2e_1}(\ell_1+1)}{\ell_1 - 1} \in O(1)$ by Assumption 1. We can write

$$\begin{aligned}
\mathbb{P}_{E_1}(\bar{\chi} \geq \lambda - \lambda_0) &= \mathbb{P}_{E_1}\left(e^{\bar{\chi}} \geq e^{\lambda - \lambda_0}\right) \\
&\overset{(\dagger)}{\leq} e^{-(\lambda-\lambda_0)}\mathbb{E}_{E_1}\left[e^{\bar{\chi}}\right] \\
&= e^{-\lambda+\lambda_0}\mathbb{E}_{E_1}\left[\exp\left(\sum_{E_2} \chi_{E_2}\right)\right] \\
&= e^{-\lambda+\lambda_0}\mathbb{E}_{E_1}\left[\prod_{E_2} \exp\left(\chi_{E_2}\right)\right] \\
&\overset{(\ddagger)}{=} e^{-\lambda+\lambda_0}\prod_{E_2}\mathbb{E}_{E_1}\left[\exp\left(\chi_{E_2}\right)\right]
\end{aligned}$$

$$
\begin{aligned}
&= e^{-\lambda+\lambda_0} \prod_{E_2} \frac{1}{m_{E_1}} \sum_{E_1} e^{\chi_{E_2}(E_1)} \qquad (19)\\
&= e^{-\lambda+\lambda_0} \prod_{E_2} \frac{1}{m_{E_1}} \left(\alpha_{E_2} e + (m_{E_1} - \alpha_{E_2})\right)\\
&\leq e^{-\lambda+\lambda_0} \prod_{E_2} \left(1 + \frac{\alpha}{m_{E_1}}(e-1)\right)\\
&= e^{-\lambda+\lambda_0} \left(1 + \frac{\alpha}{m_{E_1}}(e-1)\right)^{m_{E_2}}.
\end{aligned}
$$

In Equation (19), we used Markov's inequality at (†), and the independence of the χ_{E_2}'s at (‡). Now recall that by hypothesis, $m_{E_1}, m_{E_2} \in \Theta(\sqrt{p})$, $\lambda \in \Theta(\log(p))$ and $\lambda_0 = \log(\lambda)\log(\log(\lambda))$. We get that

$$
\mathbb{P}_{E_1}\left(\bar{\chi} \geq \lambda - \lambda_0\right) \in O\left(\frac{\log(p)^{\log(\log(\log(p)))}}{p}\right)
$$

because $\lim_{p\to\infty}\left(1 + \frac{\alpha}{\sqrt{p}}(e-1)\right)^{\sqrt{p}} = e^{\alpha(e-1)} < \infty$. Thus, $\mathbb{P}_{E_1}\left(\bar{\chi} \geq \lambda - \lambda_0\right)$ is negligible.

In Theorem 7, the assumption that the random variables $\{\chi_{E_2} : E_2 \in \mathcal{D}_{\ell_2}(E_0; e_2)\}$ are independent can be replaced by a better tailored assumption, defined as follows.

Assumption 2 *Fix E_0 and E_1 in $\mathcal{G}_{p^2}(\ell)$, then*

$$
\sum_{E_2 \in \mathcal{D}_{\ell_2}(E_0; e_2)} \mathsf{Coll}_{\ell_1^{e_1}}(E_2, E_3) \leq \frac{1}{\#\mathcal{D}_{\ell_1}(E_0; e_1)} \sum_{\substack{E_1' \in \mathcal{D}_{\ell_1}(E_0; e_1)\\ E_2 \in \mathcal{D}_{\ell_2}(E_0; e_2)}} \mathsf{Coll}_{\ell_1^{e_1}}(E_2, E_3'), \quad (20)
$$

where E_3' on the right hand side is the elliptic curve corresponding to E_1' as in the $\mathsf{ID}_{\mathsf{SIDH}}$ protocol.

Assumption 2 combined with Theorem 6 (which works under Assumption 1) tells us that there is a bound $\mathfrak{b} \in O(1)$ such that $\sum_{E_2} \mathsf{Coll}_{\ell^e}(E_2, E_3) \leq \mathfrak{b}$. Note that $\mathfrak{b}$ does not depend on E_3 (or equivalently on E_1).

Theorem 8. *Under Assumption 1 and Assumption 2, $\mathbb{P}_{E_1}\left(\bar{\chi} \geq \lambda - \lambda_0\right)$ is negligible.*

Proof. As in the proof of Theorem 7, let m_{E_2}, m_{E_1} be the number of curves in $\mathcal{D}_{\ell_2}(E_0, e_2), \mathcal{D}_{\ell_1}(E_0, e_1)$ respectively. Let also $\alpha_{E_2} = \#\{E_1 \in \mathcal{D}_{\ell_1}(E_0, e_1) : \chi_{E_2}(E_1) = 1\}$, and $\alpha = \max_{E_2}\{\alpha_{E_2}\}$. We know that $\alpha \leq \frac{1}{n}\frac{\ell_1^{2e_1}(\ell_1+1)}{\ell_1 - 1} \in O(1)$ by

Assumption 1. We can write

$$\begin{aligned}\mathbb{P}_{E_1}\left(\bar{\chi} \geq \lambda - \lambda_0\right) &= \mathbb{P}_{E_1}\left(e^{\bar{\chi}} \geq e^{\lambda - \lambda_0}\right)\\ &\overset{(\dagger)}{\leq} e^{-\lambda+\lambda_0}\mathbb{E}\left[e^{\bar{\chi}}\right]\\ &= e^{-\lambda+\lambda_0}\mathbb{E}\left[\exp\left(\sum_{E_2\in\mathcal{D}_{\ell_2}(E_0,e_2)} \chi_{E_2}\right)\right]\\ &= e^{-\lambda}e^{\lambda_0}\mathbb{E}\left[\prod_{E_2\in\mathcal{D}_{\ell_2}(E_0,e_2)} \exp\left(\chi_{E_2}\right)\right].\end{aligned} \tag{21}$$

Fix a labeling $\{E_2^{(1)}, E_2^{(2)}, ..., E_2^{(m_{E_2})}\}$ and $\{E_1^{(1)}, E_1^{(2)}, ..., E_2^{(m_{E_1})}\}$ of the elements in $\mathcal{D}_{\ell_2}(E_0, e_2)$ and $\mathcal{D}_{\ell_1}(E_0, e_1)$ respectively. Consider the matrices $A = [a_{ij}]_{i,j} := \left[\chi_{E_2^{(i)}}(E_1^{(j)})\right]_{i,j}$ and $B := [e^{ta_{ij}}]_{i,j}$. Then, $\alpha_{E_2^{(i)}} = \sum_{j=1}^{\kappa_1} a_{ij}$. The matrix A has entries in $\{0, 1\}$ and has less than $\alpha \in O(1)$ occurrences of the value 1 on each row. Let $\beta_{E_1^{(j)}} := \sum_{i=1}^{\kappa_2} a_{ij}$ and $\beta := \max_j\{\beta_{E_1^{(j)}}\}$. Assumptions 1 and 2 tell us that $\beta \in O(1)$.

Going back to (21), we see that

$$\begin{aligned}\mathbb{P}_{E_1}\left(\bar{\chi} \geq \lambda - \lambda_0\right) &\leq e^{-\lambda}e^{\lambda_0}\mathbb{E}\left[\prod_{E_2\in\mathcal{D}_{\ell_2}(E_0,e_2)} \exp\left(\chi_{E_2}\right)\right]\\ &= e^{-\lambda}e^{\lambda_0}\frac{1}{m_{E_1}}\sum_{E_1\in\mathcal{D}_{\ell_1}(E_0,e_1)}\prod_{E_2\in\mathcal{D}_{\ell_2}(E_0,e_2)} \exp\left(\chi_{E_2}\right).\end{aligned} \tag{22}$$

The quantity $\sum_{E_1}\prod_{E_2}\exp\left(\chi_{E_2}\right)$ is simply the sum of the products of all the entries of each row of B.

Assume for now that $\alpha_{E_2^{(i)}} = 1\forall i$, i.e. the matrix A has entries in $\{0, 1\}$ and has exactly one 1 on each row. Then the quantity $\sum_{E_1}\prod_{E_2}\exp\left(\chi_{E_2}\right)$ is maximized exactly when all the 1's line up on one column (and is minimized when all the 1's of each row lie on different columns). However, we can have at most β occurrences of 1's on each column. So, $\sum_{E_1}\prod_{E_2}\exp\left(\chi_{E_2}\right)$ is maximized if we have β occurrences of 1's on as many columns as possible (i.e. $\lfloor m_{E_2}/\beta\rfloor$ columns). This also means that the remaining $m_{E_1} - \lfloor m_{E_2}/\beta\rfloor$ columns of A are full of zeroes. Thus, we have $\sum_{E_1}\prod_{E_2}\exp\left(\chi_{E_2}\right) \leq e^{\beta}m_{E_2}/\beta + m_{E_1} - m_{E_2}/\beta$. So we can go back to (22) and write

$$\begin{aligned}\mathbb{P}_{E_1}\left(\bar{\chi} \geq \lambda - \lambda_0\right) &\leq e^{-\lambda}e^{\lambda_0}\frac{1}{m_{E_1}}\sum_{E_1}\prod_{E_2}\exp\left(\chi_{E_2}\right)\\ &\leq e^{-\lambda}e^{\lambda_0}\frac{1}{m_{E_1}}\left(e^{\beta}\frac{m_{E_2}}{\beta} + m_{E_1} - \frac{m_{E_2}}{\beta}\right).\end{aligned}$$

We thus obtain that $\mathbb{P}_{E_1}\left(\bar{\chi} \geq \lambda - \lambda_0\right)$ is negligible because $\frac{1}{m_{E_1}}(e^{\beta}\frac{m_{E_2}}{\beta} + m_{E_1} - \frac{m_{E_2}}{\beta}) \in O(1)$, $\lambda_0 = \log(\lambda)\log(\log(\lambda))$ and $e^{-\lambda} \in O(1/p)$.

Finally, if we remove the assumption that $\alpha_{E_2^{(i)}} = 1 \ \forall i$, we obtain

$$\sum_{E_1}\prod_{E_2} \exp\left(\chi_{E_2}\right) \leq e^{-\lambda} e^{\lambda_0} \frac{\alpha}{m_{E_1}} \left(e^{\beta} \frac{m_{E_2}}{\beta} + m_{E_1} - \frac{m_{E_2}}{\beta} \right).$$

And as $\alpha \in O(1)$, $\mathbb{P}_{E_1}\left(\bar{\chi} \geq \lambda - \lambda_0\right)$ is still negligible.

6.2 Collisions in isogeny-based cryptography

Under Assumption 1, we will now provide a formal justification to the fact that collisions in supersingular isogeny graphs occur with a negligible probability. In the analysis of claw-finding algorithms for the CSSI problem in [CLN+20], the focus is on finding a cyclic $\ell_1^{e_1}$-isogeny φ between two given vertices, E_0 and E_1, of $\mathcal{G}_{p_2}(\ell_1)$. In [CLN+20, Sec. 2.2] it is claimed that, once an elliptic curve E' and two cyclic $\ell_1^{e_1/2}$-isogenies $f_1 : E_0 \longrightarrow E'$ and $g_1 : E_1 \longrightarrow E'$ are found, the composition $\hat{g_1} \circ f_1$ would return φ. This argument holds because the authors assume that there exists a unique cyclic $\ell_1^{e_1}$-isogeny between E_0 and E_1. By regarding E_1 as a uniformly random curve in the set $\mathcal{D}_{\ell_1}(E_0; e_1)$, Theorem 6 formally proves that there is no lack of generality in considering a collision-free definition of the CSSI problem as in [CLN+20].

Similarly in [GPSBT16, Sec. 3.1], it is implicitly assumed that if two curves E_B/G_1 and E_B/G_2 are isomorphic, with G_1 and G_2 cyclic subgroups of order $\ell_1^{e_1}$ in $E_B[\ell_1^{e_1}]$, then G_1 must be equal to G_2. In other words, if two cyclic isogenies of the same degree start at the same curve and have the same image, then they are equivalent. This is exactly the argument we have put under the spotlight in this work. Since, in the above assumption, E_B is a fixed vertex of the graph $\mathcal{G}_{p^2}(\ell_1)$, while E_B/G_1 is a uniformly random curve among those connected to E_B by a cyclic $\ell_1^{e_1}$-isogeny, we can again exploit Theorem 6 to affirm that the probability that $E_B/G_2 \cong E_B/G_1$ with $G_1 \neq G_2$ is negligible.

7 A New Extractor

Under the two assumptions made in Section 6.1 on supersingular isogeny graphs, the extractor $\mathsf{Ex}_{\mathsf{SIDH}}$ can still be exploited to prove the security of the SIDH-based digital signatures despite it granting no special soundness to $\mathsf{ID}_{\mathsf{SIDH}}$ because of the counterexamples described in Section 3. However, one could still prove the special soundness of $\mathsf{ID}_{\mathsf{SIDH}}$ by relying on alternative extractors unaffected by such counterexamples. In this section we present a new extractor $\mathsf{NEx}_{\mathsf{SIDH}}$ which makes the identification protocol $\mathsf{ID}_{\mathsf{SIDH}}$ satisfy Definition 3a. The new extractor initially operates as $\mathsf{Ex}_{\mathsf{SIDH}}$, but then deviates from it in checking if the extracted isogeny is the correct witness. By Theorem 3, the latter happens only if the two valid transcripts in input form an instance of Scenario 1 (see Section 3.2). The core idea behind $\mathsf{NEx}_{\mathsf{SIDH}}$ is to overcome this problem by shrinking the long isogeny $\rho = \hat{\phi}' \circ \psi \circ \phi : E_0 \longrightarrow E_1$, determined by the input, into the correct witness. The resize of ρ is based on two algorithms recently presented in the context of isogeny-based cryptography. The first one, due to Wesolowski [Wes21], reduces the problem of computing the maximal order of a given supersingular elliptic curve to the problem of computing an ℓ^e-isogeny between two supersingular elliptic curves over $\mathbb{F}_{p^2}$ for some integer $e \in \mathbb{N}$ and a prime ℓ coprime with p. The second one, due to Fouotsa et al. [FKM21], computes an isogeny between two supersingular curves of known endomorphism rings from torsion point information.

Since variants of both algorithms are employed for $\mathsf{NEx}_{\mathsf{SIDH}}$, we briefly recall the original algorithms, we motivate the need for and present the two variants, and we combine them in the design of the new extractor $\mathsf{NEx}_{\mathsf{SIDH}}$.

Remark 10. The many algorithms in this section work in similar settings, and nevertheless present some crucial differences in parameter choices. In particular, the reader might wonder when E_0 has to be intended as a specific elliptic curve, and when it can be any supersingular elliptic curve. In Section 7.1, the curve E_0 is a specific elliptic curve whose defining equation depends on whether $p \equiv 3 \mod 4$, $p \equiv 1 \mod 8$ or $p \equiv 5 \mod 8$. In each of the three cases, one picks a special maximal order $\tilde{\mathcal{O}}$ in $B_{p,\infty}$ according to [Wes21, Lemma 2.3], and the corresponding curve $\tilde{E}$ is then computed according to [Wes21, Lemma 2.5]. For example, $\tilde{E} : y^2 = x^3 + x$ when $p \equiv 3 \mod 4$. In Section 7.2 the setting remains the same, thus Algorithm 1 works with $\tilde{E}$. In Section 7.3 we extend an algorithm from Fouotsa et al. [FKM21], for which E_0 can be any supersingular elliptic curve connected to $\tilde{E}$ by an isogeny of suitable smooth degree. In Section 7.4 we shift our attention to concrete instantiations of $\mathsf{ID}_{\mathsf{SIDH}}$. In this case, the fixed starting curve E_0 is set to be that defined by $y^2 = x^3 + 6x^2 + x$, which is the only curve non-isomorphic and 2-isogenous to $\tilde{E}$.

7.1 Wesolowski's algorithm and its building blocks

Given a supersingular elliptic curve E over $\mathbb{F}_{p^2}$, Algorithm 3 in [Wes21] reduces the problem of computing a maximal order $\mathcal{O} \simeq \mathrm{End}(E)$ in $B_{p,\infty}$ to the problem of computing a path between two vertices in $\mathcal{G}_{p^2}(\ell)$. The new extractor $\mathsf{NEx}_{\mathsf{SIDH}}$ will make use of Algorithm 1, a variation of Wesolowski's result, to determine a maximal order isomorphic to $\mathsf{End}(E_1)$ given a smooth-degree isogeny from $\tilde{E}$ to a non-isomorphic curve E_1.

In order to better explain the modification to Algorithm 3 in [Wes21] that we propose, we briefly recall its main building blocks.

The first one is an algorithm from Galbraith et al. [GPS17, Algorithm 2] that allows to translate an ideal to its matching isogeny under the Deuring correspondence, provided it originates from a curve whose endomorphism ring is isomorphic to $\tilde{O}$. We denote such algorithm by `IdealToIsogeny`.

Lemma 4. *[GPS17, Lemma 5] On input a left $\tilde{\mathcal{O}}$-ideal I of norm $N = \prod_i \ell_i^{e_i}$, the* `IdealToIsogeny` *algorithm outputs its corresponding isogeny $\varphi_I : \tilde{E} \longrightarrow E$ in running time that is polynomial in $\log p$ and $\max_i\{\ell_i^{e_i}\}$ under the Generalised Riemann Hypothesis.*

The second building block is an algorithm, denoted by `IsogenyToIdeal` in the following, that does the opposite of `IdealToIsogeny`, i.e. it computes the ideal corresponding to an isogeny of composite degree under the Deuring correspondence.

Lemma 5. *[Wes21, Lemma 7.1] On input an isogeny $\varphi : \tilde{E} \longrightarrow E$ of composite degree $\prod_i \ell_i^{e_i}$, the algorithm* `IsogenyToIdeal` *outputs the corresponding left $\tilde{\mathcal{O}}$-ideal I_φ in running time that is polynomial in $\log p$ and $\max_i\{\ell_i^{e_i}\}$ under the Generalised Riemann Hypothesis.*

The third algorithm, which we will denote by $\texttt{EquivIdeal}_c$, computes an equivalent ideal to a given one such that it satisfies some constraints on the norm. The algorithm results from [Wes21, Theorem 6.4], which we recall here below.

Theorem 9. *[Wes21, Theorem 6.4] There exists an integer $c \in \mathbb{N}$ and an algorithm $\texttt{EquivIdeal}_c$ such that, on input a left $\tilde{\mathcal{O}}$-ideal I, it outputs an equivalent ideal of $(\log p)^c$-powersmooth norm. The algorithm runs in expected polynomial time in $\log p$ and $\log n(I)$ under the Generalised Riemann Hypothesis.*

7.2 Computing maximal orders from composite isogenies

Given $\tilde{\mathcal{O}} \subset B_{p,\infty}$, $\tilde{E}$ over $\mathbb{F}_{p^2}$, and any other supersingular elliptic curve E_* over $\mathbb{F}_{p^2}$, Wesolowski [Wes21, Algorithm 3] reduces the problem of computing a maximal order $\mathcal{O}_* \simeq \mathrm{End}(E_*)$ to the one of computing a path from $\tilde{E}$ to E_* in $\mathcal{G}_{p^2}(\ell)$, for a prime ℓ different from p.

In our specific use-case, given a composite-degree isogeny $\omega = \eta \circ \rho \circ \nu : \tilde{E} \longrightarrow E_*$, $\mathsf{NEx}_{\mathsf{SIDH}}$ needs to compute a maximal order of $B_{p,\infty}$ isomorphic to $\mathrm{End}(E_*)$. Here, ν is an isogeny from $\tilde{E}$ to E_0 (where E_0 is the fixed starting curve considered for $\mathsf{ID}_{\mathsf{SIDH}}$), $\rho = \hat{\phi}' \circ \psi \circ \phi : E_0 \longrightarrow E_1$, and $\eta : E_1 \longrightarrow E_*$ has degree ℓ_1^t (for some t which does not depend on the characteristic p of the base field). To this end, we extend Wesolowski's algorithm [Wes21, Algorithm 3] to allow the input isogeny to have composite degree. As we argue in the proof of Theorem 10, our modification does not alter the correctness nor the efficiency of the algorithm.

Theorem 10. *Algorithm 1 is correct and runs in expected polynomial time in $\log p$ under the Generalised Riemann Hypothesis.*

Proof. The first iteration of the outer **for**-loop corresponds exactly to the application of [Wes21, Algorithm 3] to ω_1. When we compose the isogeny $\omega_{2,1}$ with (the new) ψ_0, one of small degree ℓ_2 and the other of $(\log p)^c$-powersmooth degree, we do not change the complexity of `IsogenyToIdeal`, which still runs in polynomial time in $\log p$ (Lemma 5). Its output is still $(\log p)^c$-powersmooth, and thus $\texttt{EquivIdeal}_c$ and `IdealToIsogeny` still run in polynomial time in $\log p$ (Theorem 9 and Lemma 4). Correctness and efficiency of the remaining steps necessary to "exhaust" ω_2 follow directly from Wesolowski's result. Repeating

Algorithm 1 Compute a maximal order from a composite-degree isogeny.

```
Require: An isogeny $\omega : \tilde{E} \longrightarrow E_*$ of composite degree $\prod_{i=1}^{n} \ell_i^{e_i}$, the $\ell_i$ not all
    necessarily distinct, with $\max_i\{\ell_i\}$ and $\sum_i e_i$ polynomial in $\log p$.
Ensure: A basis of $\mathcal{O}_* \simeq \mathrm{End}(E_*)$.
let $\omega = \omega_n \circ \cdots \circ \omega_2 \circ \omega_1$, with $\deg(\omega_i) = \ell_i^{e_i}$;
let $\psi_0 : \tilde{E} \longrightarrow \tilde{E}$ be the identity isogeny on $\tilde{E}$;
let $c \in \mathbb{N}$ as per Theorem 9;
for $i = 1, 2, \dots, n$ do
    let $\omega_i = \omega_{i,e_i} \circ \omega_{i,e_i-1} \circ \cdots \circ \omega_{i,1}$ with $\deg(\omega_{i,j}) = \ell_i$;
    for $j = 1, 2, \dots, e_i$ do
        $I_{i,j} \longleftarrow$ IsogenyToIdeal$(\omega_{i,j} \circ \psi_{j-1})$;
        $J_{i,j} \longleftarrow$ EquivIdeal$_c(I_{i,j})$ of $(\log p)^c$-powersmooth norm;
        $\psi_j \longleftarrow$ IdealToIsogeny$(J_{i,j})$;
    end for
    $\psi_0 \longleftarrow \psi_{e_i}$;
end for
$\mathcal{O}_* \longleftarrow \mathcal{O}_R(J_{n,e_n})$;
return a basis of $\mathcal{O}_*$.
```

the same argument for all the remaining components ω_i of ω concludes the proof of correctness.

We are now left to prove the efficiency of the algorithm as a whole. By hypothesis, $\sum_i e_i$ is polynomial in $\log p$, and thus the whole algorithm is polynomial in $\log p$.

7.3 Extending Foutsa et al.'s algorithm

On input a long isogeny $\rho : E_0 \longrightarrow E_1$ and $\nu : \tilde{E} \longrightarrow E_0$, $\mathsf{NEx}_{\mathsf{SIDH}}$ first computes a set $H = \{E_{1,i} \mid i = 1, \dots, k\}$ of elliptic curves *close* to E_1 (as we will explain below). It then iterates our extension of Wesolowski's algorithm, and finally applies an algorithm recently introduced by Fouotsa et al. [FKM21] to determine the witness φ.

The latter algorithm takes as input two supersingular elliptic curves E, E' over $\mathbb{F}_{p^2}$ of known endomorphism rings and corresponding maximal orders $\mathcal{O} \simeq \mathsf{End}(E)$ and $\mathcal{O}' \simeq \mathsf{End}(E')$. Let $N_1, N_2 \in \mathbb{N}$ be two coprime integers, $\varphi : E \longrightarrow E'$ an N_1-isogeny and $\varphi(P), \varphi(Q)$ the images of a basis $\{P, Q\}$ of the N_2-torsion subgroup $E[N_2]$. If the shortest isogeny between E and E' has degree $d > \frac{16N_1}{N_2}$, Algorithm 3 from [FKM21] - which we rename `IsogenyFromTorsion` - can recover φ from $E, E', \mathcal{O}, \mathcal{O}', \varphi(P), \varphi(Q)$ with complexity polynomial in $\log p$ under the Generalised Riemann Hypothesis.

Given the condition on d, one needs to rule out the existence of any isogeny between E and E' of degree up to $\frac{16N_1}{N_2}$ in order to ensure a successful application of `IsogenyFromTorsion`. In the specific context of $\mathsf{NEx}_{\mathsf{SIDH}}$, the goal is to compute the secret isogeny $\varphi : E_0 \longrightarrow E_1$ of degree N_1 by knowing its behaviour on the N_2-torsion points. Despite $N_1 = \ell_1^{e_1} \approx N_2 = \ell_2^{e_2}$ by design of $\mathsf{ID}_{\mathsf{SIDH}}$, the approximation might hide a discrepancy between the two values N_1 and N_2, as it happens when considering the SIKE parameter sets [7]. Therefore, if the new extractor applied `IsogenyFromTorsion` directly to φ, it would fail with a certain probability.

Our idea to remove any failure probability of `IsogenyFromTorsion` when used within $\mathsf{NEx}_{\mathsf{SIDH}}$ is to brute-force a small portion of the secret isogeny and apply `IsogenyFromTorsion` to recover the remaining part. In details, let t be the smallest integer such that $\frac{16N_1'}{N_2} = \frac{16N_1}{\ell_1^t N_2} < 1$, where $N_1' := N_1/\ell_1^t$. Then φ can be factored as $\varphi = \hat{\eta} \circ \varphi'$, where $\hat{\eta}$ is an isogeny of degree ℓ_1^t. If $\hat{\eta}$ is correctly guessed, φ' can be computed with `IsogenyFromTorsion` without any failure probability. In fact, φ' would have degree $N' = \frac{N_1}{\ell_1^t}$ with $\frac{16N'}{N_2} < 1$. The correct $\hat{\eta}$ is guessed by computing all ℓ_1^t-isogenies from E_1, which explains why the dual of η appears in the factorisation of φ. The algorithm iterated by $\mathsf{NEx}_{\mathsf{SIDH}}$ is depicted in Algorithm 2.

Remark 11. The guessing strategy implies running `IsogenyFromTorsion` on two elliptic curves which might not be connected by any isogeny of the prescribed degree N_1'. This scenario is not considered in the original formulation of the algorithm, which implicitly fails whenever this hypothesis is not met. More specifically, one of the steps in `IsogenyFromTorsion` consists in solving a system of linear equations to express the secret isogeny in terms of a basis of $\mathrm{Hom}(E_0, E_{1,i})$, which is the 4-rank lattice of all isogenies from E_0 to $E_{1,i}$ (see [FKM21, Algorithm 2] for more details). For all but the correct η, the algorithm will be unable to find the unique solution to the system, whose construction is not affected by the lack of any isogeny of degree N_1'. For this reason, we modify the output space of `IsogenyFromTorsion`, setting the output to $\perp$ whenever the algorithm fails in solving the system of linear equations.

Theorem 11. *Assume that the set H of 2^t-isogenies from E_1 can be computed in polynomial time in $\log p$. Then Algorithm 2 is correct and runs in expected polynomial time in $\log p$, when t does not depend on p and under the Generalised Riemann Hypothesis.*

Proof. Correctness and efficiency of Step 1. follow directly from Theorem 10. For each guess η, the algorithm computes the image $E_* = E_1/\ker(\eta)$, the maximal order $\mathcal{O}_* \simeq \mathrm{End}(E_*)$ and the images $\varphi'(P)$ and $\varphi'(Q)$. E_* is obtained by applying Vélu's formulae, and $\mathcal{O}_*$ using Algorithm 1. Since $\varphi(P) = \hat{\eta} \circ \varphi'(P)$, then $\eta \circ \varphi(P) = \eta \circ \hat{\eta} \circ \varphi'(P) = [\ell_1^t]\varphi'(P)$. We can thus compute $\varphi'(P) = [(\ell_1^t)^{-1} \bmod N_2]\,\eta(\varphi(P))$, and $\varphi'(Q)$ in a similar fashion. These images can be computed for each $\eta \in H$, but only the correct one will ensure that the output of step 8 is not $\perp$. The only η leading to a proper solution will be the correct guess, and we can then reconstruct $\varphi = \hat{\eta} \circ \varphi'$.

[7] For the parameter sets SIKEp434, SIKEp503, SIKEp610 and SIKEp751, $\lfloor \frac{16N_1}{N_2} \rfloor$ is 35, 64, 9 and 1790 respectively.

Algorithm 2 Compute an isogeny from torsion points and maximal orders.

```
Require: A balanced prime $p = N_1 N_2 f \pm 1 = \ell_1^{e_1} \ell_2^{e_2} f \pm 1$; two supersingular elliptic
    curves $E_0, E_1$ over $\mathbb{F}_{p^2}$; an isogeny $\nu : \tilde{E} \longrightarrow E_0$ of degree $\prod_{i=1}^{n'} {\ell'_i}^{e'_i}$ with $\max_i\{\ell'_i\}$
    and $\sum_i e'_i$ polynomial in $\log p$; the images $\varphi(P), \varphi(Q)$ of a basis $\{P, Q\}$ of the
    $N_2$-torsion subgroup $E_0[N_2]$; an isogeny $\rho : E_0 \longrightarrow E_1$ of degree $\prod_{i=1}^{n} \ell_i^{e_i}$ with
    $\max_i\{\ell_i\}$ and $\sum_i e_i$ polynomial in $\log p$.
Ensure: The $N_1$-isogeny $\varphi : E_0 \longrightarrow E_1$.
 1: $\mathcal{O}_0 \longleftarrow$ Algorithm 1 $(\nu)$;
 2: let $t$ be the smallest integer s.t. $(16N_1)/(\ell_1^t N_2) < 1$
 3: let $H$ denote the set of $2^t$-isogenies from $E_1$;
 4: for each $\eta \in H$ do
 5:     let $E_* = E_1 / \ker(\eta)$;
 6:     $\mathcal{O}_* \longleftarrow$ Algorithm 1 $(\eta \circ \rho \circ \nu)$;
 7:     let $\varphi'(P) = [(\ell_1^t)^{-1} \mod N_2]\, \eta(\varphi(P))$;
 8:     let $\varphi'(Q) = [(\ell_1^t)^{-1} \mod N_2]\, \eta(\varphi(Q))$;
 9:     $\varphi' \longleftarrow$IsogenyFromTorsion$(E_0, E_*, \mathcal{O}_0, \mathcal{O}_*, \varphi'(P), \varphi'(Q))$;
10:     if $\varphi' \neq \perp$ then
11:         break
12:     end if
13: end for
14: return $\varphi = \hat{\eta} \circ \varphi'$.
```

7.4 The new extractor

We conclude this section by giving a detailed description of the new extractor $\mathsf{NEx}_{\mathsf{SIDH}}$ in Algorithm 3 and arguing its efficiency.

Theorem 12. *Algorithm 3 is correct and runs in expected polynomial time in* $\log p$ *under the Generalised Riemann Hypothesis.*

Proof. The correctness follows from that of Algorithm 2, and from Theorem 3. For the efficiency of the algorithm, we focus on line 7 and on SIDH/SIKE sets of parameters. Therein, the starting elliptic curve E_0 is 2-isogenous to $\tilde{E}$. In addition, the primes ℓ_1, ℓ_2 are small constants (tipically 2 and 3) for efficiency reasons, thus independent on p; as such, they satisfy the conditions on the ℓ_i's of Algorithm 2. The degree of ρ is either $\ell_1^{2e_1} \ell_2^{e_2}$ or $\ell_1^{e_1} \ell_2^{2e_2}$, which are both approximately $p^{3/2}$. Therefore, also the condition on the e_i's is met. The $\ell_1^{e_1}$-torsion subgroup of E_0 is defined over $\mathbb{F}_{p^2}$ for efficiency reasons, so is $E_1[\ell_1^t]$, and thus Theorem 11 holds. In fact, t (defined in Section 7.3) does not depend on p. The latter can be assumed without loss of generality as SIDH and SIKE (and then $\mathsf{ID}_{\mathsf{SIDH}}$) work over balanced primes of the form $p = N_1 N_2 f \pm 1 = \ell_1^{e_1} \ell_2^{e_2} f \pm 1$ with f small natural number. It is therefore reasonable to translate the balance condition to a formal statement of this fashion: there exists a constant $t \leq 15$ such that $(16N_1)/(\ell_1^t N_2) < 1$.

Algorithm 3 The knowledge extractor $\mathsf{NEx}_{\mathsf{SIDH}}$

```
Require: A statement $\mathsf{x} = (E_1, \varphi(P), \varphi(Q))$, two valid transcripts $(\mathsf{x}, \mathsf{com}, 0, (m, n))$
    and $(\mathsf{x}, \mathsf{com}, 1, T)$ relative to the statement $\mathsf{x}$, on the same commitment $\mathsf{com} =$
    $(E_2, E_3)$ and distinct challenges.
Ensure: The $\ell_1^{e_1}$-isogeny witness $\varphi : E_0 \longrightarrow E_1$.
let $\phi : E_0 \longrightarrow E_2$ be the isogeny of kernel $\langle mP + nQ \rangle$ (where $\langle P, Q \rangle = E_0[\ell_2^{e_2}]$);
let $\psi : E_2 \longrightarrow E_3$ be the isogeny of kernel $\langle T \rangle$;
let $\phi' : E_1 \longrightarrow E_3$ be the isogeny of kernel $\langle m\varphi(P) + n\varphi(Q) \rangle$;
let $\overline{\varphi} : E_0 \longrightarrow \overline{E_1}$ be the isogeny of kernel $\ker(\hat{\phi}' \circ \psi \circ \phi) \cap E_0[\ell_1^{e_1}]$;
if $j(\overline{E_1}) \neq j(E_1) \ \vee \ \overline{\varphi}(P) \neq \varphi(P) \ \vee \ \overline{\varphi}(Q) \neq \varphi(Q)$ then
    $\varphi \longleftarrow$ Algorithm 2 $(E_0, E_1, \eta, \varphi(P), \varphi(Q), \hat{\phi}' \circ \psi \circ \phi)$;
end if;
return $\varphi$
```

$$\begin{tikzcd}
 & & E_{1,1} & \\
 & & E_{1,2} & \\
 & & \vdots & \\
\tilde{E} \arrow[r, bend left, "\nu"] & E_0 \arrow[l, bend left, "\hat{\nu}"] \arrow[r, "\varphi'"] \arrow[uuur, dotted] \arrow[uur, dotted] \arrow[ddr, dotted] \arrow[ddd, "\phi"'] & E_* & E_1 \arrow[l, "\eta"'] \arrow[uuul, dotted] \arrow[uul, dotted] \arrow[ddl, dotted] \arrow[ddd, "\phi'"] \\
 & & \vdots & \\
 & & E_{1,k} & \\
 & E_2 \arrow[rr, "\psi"] & & E_3
\end{tikzcd}$$

Fig. 6: The new extractor $\mathsf{NEx}_{\mathsf{SIDH}}$.

8 Conclusion

In this paper, we have 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 extraction algorithm $\mathsf{Ex}_{\mathsf{SIDH}}$ — considered by existing proofs — fails, we have also carefully studied the number of collisions that occur in supersingular isogeny graphs. Our analysis shows that the special cycles that make $\mathsf{Ex}_{\mathsf{SIDH}}$ fail do exist in supersingular isogeny graphs. However, if we assume that such cycles are evenly distributed over the vertex set, their existence does not affect the security of the SIDH-based digital signatures, provided the verification algorithms are modified in such a way to require λ different commitments for each signature. Our calculations are general and can be exploited in other contexts within isogeny-based cryptography. We leave for future work to further improve the bounds we provided on the number of collisions and to assess the assumptions on their distribution, as doing so would require developing new number theoretic tools. Furthermore, we have introduced a new extraction algorithm $\mathsf{NEx}_{\mathsf{SIDH}}$ which makes $\mathsf{ID}_{\mathsf{SIDH}}$ enjoy the special soundness property under the Generalised Riemann Hypothesis. This further corroborate our conclusion that the security proofs of the digital signature schemes deduced from $\mathsf{ID}_{\mathsf{SIDH}}$ are reliable.

References

[AABN02] Michel Abdalla, Jee An, Mihir Bellare, and Chanathip 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.

[ACC$^+$17] Reza Azarderakhsh, Matthew Campagna, Craig Costello, Luca De Feo, Basil Hess, Amir Jalali, David Jao, Brian Koziel, Brian La Macchia, and Patrick Longa. Supersingular isogeny key encapsulation. Submission to the NIST Post-Quantum Standardization project, 2017.

[ACK21] Thomas Attema, Ronald Cramer, and Lisa. Kohl. A compressed σ-protocol theory for lattices. In *CRYPTO 2021*, pages 549–579. Springer, Cham, 2021.

[Bra47] Richard Brauer. On the zeta-functions of algebraic number fields. *American Journal of Mathematics*, 69(2):243–250, 1947.

[Bur62] David A Burgess. On character sums and primitive roots. *Proceedings of the London Mathematical Society*, 3(1):179–192, 1962.

[Bur63] David A Burgess. On character sums and l-series. ii. *Proceedings of the London Mathematical Society*, 3(1):524–536, 1963.

[Bur86] DA Burgess. The character sum estimate with r= 3. *Journal of the London Mathematical Society*, 2(2):219–226, 1986.

[CJS14] Andrew Childs, David Jao, and Vladimir Soukharev. Constructing elliptic curve isogenies in quantum subexponential time. *Journal of Mathematical Cryptology*, 8.1:1–29, 2014.

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

[CLN$^+$20] Craig Costello, Patrick Longa, Michael Naehrig, Joost Renes, and Fernando Virdia. Improved classical cryptanalysis of sike in practice. In *Public Key Cryptography 2020*, pages 505–534. Springer, 2020.

[Cou06] Jean Marc Couveignes. Hard homogeneous spaces. IACR Cryptol. ePrint Arch. (2006): 29, 2006.

[Deu41] Max Deuring. Die typen der multiplikatorenringe elliptischer funktionenkörper. In *Abhandlungen aus dem mathematischen Seminar der Universität Hamburg*, volume 14, pages 197–272. Springer, 1941.

[DFDGZ21] Luca De Feo, Samuel Dobson, Steven D. Galbraith, and Lukas Zobernig. Sidh proof of knowledge. *IACR Cryptol. ePrint Arch. 2021/1023*, 2021.

[DFJP14] Luca De Feo, David Jao, and Jérôme Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. *Journal of Mathematical Cryptology*, 8.3:209–247, 2014.

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

[EK18] Ali El Kaafarani and Shuichi Katsumata. Attribute-based signatures for unbounded circuits in the rom and efficient instantiations from lattices. In *PKC 2018*, pages 89–119. Springer, 2018.

[Fis05] Marc Fischlin. Communication-efficient non-interactive proofs of knowledge with online extractors. In *CRYPTO*, pages 152–168. Springer, Berlin, Heidelberg, 2005.

[FKM21] Tako Boris Fouotsa, Péter Kutas, and Simon-Philipp Merz. On the isogeny problem with torsion point information. Cryptology ePrint Archive, Report 2021/153, 2021. https://ia.cr/2021/153.

[FS86] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, pages 186–194. Springer, 1986.

[Gha21] Wissam Ghantous. Loops, multi-edges and collisions in supersingular isogeny graphs. *arXiv preprint arXiv:2101.08761*, 2021.

[GPS17] Steven D Galbraith, Christophe Petit, and Javier Silva. Identification protocols and signature schemes based on supersingular isogeny problems. In *ASIACRYPT*, pages 3–33. Springer, 2017.

[GPSBT16] Steven D. Galbraith, Christophe Petit, Barak Shani, and Yan Bo Ti. On the security of supersingular isogeny cryptosystems. In *ASIACRYPT*, pages 63–91. Springer, 2016.

[Gro87] B. Gross. Heights and the special values of L-series. In *Conference Proceedings of the CMS*, volume 7, 1987.

[Hur85] Adolf Hurwitz. Ueber relationen zwischen classenanzahlen binärer quadratischer formen von negativer determinante. *Mathematische Annalen*, 25(2):157–196, 1885.

[IK04] Henryk Iwaniec and Emmanuel Kowalski. *Analytic number theory*, volume 53. American Mathematical Soc., 2004.

[OAT20] Hiroshi Onuki, Yusuke 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.

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

[Sho94] Peter 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.

[Sil09] Joseph H. Silverman. *The arithmetic of elliptic curves.* New York, Springer, Vol. 106, 2009.

[UJ18] David Urbanik and David Jao. Sok: The problem landscape of sidh. In *APKC*, pages 53–60. ACM, 2018.

[Unr15] Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In *EUROCRYPT*, pages 755–784. Springer, 2015.

[Unr17] Dominique Unruh. Post-quantum security of fiat-shamir. In *ASIACRYPY*, pages 65–95. Springer, Cham, 2017.

[Vit19] Vanessa Vitse. Simple oblivious transfer protocols compatible with supersingular isogenies. In *International Conference on Cryptology in Africa*, pages 56–78. Springer, Cham, 2019.

[Wai19] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.

[Wes21] Benjamin Wesolowski. The supersingular isogeny path and endomorphism ring problems are equivalent. Cryptology ePrint Archive, Report 2021/919, 2021. https://ia.cr/2021/919.

[YAJ$^+$17] Youngho Yoo, Reza Azarderakhsh, Amir Jalali, David Jao, and Vladimir Soukharev. A post-quantum digital signature scheme based on supersingular isogenies. In *FC*, pages 163–181. Springer, 2017.

Appendix A

In order to produce an example of Scenario 1 (see Section 3.2) for the largest set of SIKE parameters, namely SIKEp751, we relied on the approach adopted by Onuki, Aikawa and Takagi [OAT20] to determine endomorphisms of prescribed degree in supersingular isogeny graphs for SIKE parameters. We sketch the aforementioned approach in the following.

1. Since the prime p in each SIKE parameter set is congruent to 15 modulo 16, the explicit description of the endomorphism ring of $E_0 : y^2 = x^3 + x$ is known. There exists a 2-isogeny that maps E_0 to the curve $E_6 : y^2 = x^3 + 6x^2 + x$, while the others map to E_0 itself. Lemma 1 of [OAT20] gives an explicit description of $\mathrm{End}(E_6)$ via the maximal order $\mathcal{O}_6$ isomorphic to $\mathrm{End}(E_6)$. In particular, any endomorphism of E_6 corresponds to a quaternion in $\mathcal{O}_6 \subset B_{p,\infty} = \left(\frac{-1,-p}{\mathbb{Q}}\right)$ that can be written as
$$\alpha = a + 2b\mathbf{i} + c\frac{1+\mathbf{j}}{2} + d\frac{\mathbf{i}+\mathbf{k}}{4}, \tag{23}$$
with $a, b, c, d \in \mathbb{Z}$ (not all divisible by ℓ). Recall that $\{1, \mathbf{i}, \mathbf{j}, \mathbf{k}\}$ is the canonical base for $B_{p,\infty}$ such that $\mathbf{i}^2 = -1, \mathbf{j}^2 = -p, \mathbf{ij} = \mathbf{k} = -\mathbf{ji}$.
2. In order to find a quaternion of norm ℓ^n, where $\ell^n \in \{\ell_1^{2e_1}, \ell_2^{2e_2}\}$, the following norm equation is obtained by equating the norm of the right-hand side of Equation (23) to ℓ^n:
$$\ell^n = \frac{1}{16}\big((4a+2c)^2 + (8b+d)^2 + (4c^2+d^2)p\big). \tag{24}$$
By rearranging Equation (24) and replacing $A = 4a + 2c$ and $B = 8b + d$, the equation is reduced to the following Diophantine equation
$$16\ell^n - (4c^2 + d^2)p = A^2 + B^2 \tag{25}$$
with c and d not both 0.

In [OAT20], the authors applied the above approach on the first two SIKE parameter sets only, namely SIKEp434 (where $e_1 = 216$ and $e_2 = 137$) and SIKEp503 (where $e_1 = 250$ and $e_2 = 159$). The results for these two cases show that there exists only one endomorphism of degree 3^{274} and no endomorphisms of degree 2^{432} when considering SIKEp434, and no endomorphisms of degrees 2^{500} or 3^{318} for SIKEp503. The authors could not apply their approach to SIKEp610 and SIKEp751 in [OAT20].

In this paper, we fill the gap by investigating the two remaining parameter sets. Our code was run on a laptop mounting a Coffee Lake 2.2 GHz Intel i7 processor and 12GB of RAM. The results of our tests show that for SIKEp610 (where $e_1 = 305$ and $e_2 = 192$) there exists no endomorphism of degree 2^{610} nor 3^{384}. For SIKEp751 (where $e_1 = 372$ and $e_2 = 239$), there is no endomorphism of degree 2^{744}, but there are two endomorphisms of degree 3^{478}. These two endomorphisms determine two pairs of colliding isogenies of degree 3^{239}, 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

$$\begin{aligned}
P_2 = (&841279105546461804686499115213762297740963469565458990324532197\\
&170119753830282998639717315851968048398639566308956083752260055190\\
&571572196659210785187081691561179878799936982967086297460371869267\\
&068644394211287290040655663023 1, 232832339774986047278299512593543 1\\
&844906025806061286315561848048094118981177935779374110739971558302\\
&443669052376824243054704744119836366858126425102589256068192209923\\
&162999484308658507773457580251501832576844559003321891810286),\\
Q_2 = (&543635210060763384014311798206759072587043797232568742989135095\\
&933754649063203361494382901183312364485658937300398819709525956057\\
&438717298931212333675539072301651530293091043813638718115191619543\\
&534872448539156352030073708291 2, 274682061938050432765845647036057 6\\
&392116084162139873044543253936888674469970698562051629959917582270\\
&795051329607070975611602221797732105304217318783757571295486447343\\
&909940158935181461281254155392827518984003577096842656338942 \cdot \eta),
\end{aligned}$$

the following points K_1 and K_2 generate two distinct cyclic isogenies whose images have j-invariant j:

$$\begin{aligned}
K_1 &= 193262648327552634362057829832895638371956375167641154459552249\\
&\quad 805508664589936678565722236003305074981600179788901 \cdot P_2 + 398140788\\
&\quad 120904530736166185630489084280010975246499905086708128521669460981\\
&\quad 866051005968073733368199418319346274 98 \cdot Q_2,\\
K_2 &= 883152691005298701476350449004891666528834642174726130546931940\\
&\quad 786972410641216900983358627043999654855326427935366 \cdot P_2 + 398140788\\
&\quad 120904530736166185630489084280010975246499905086708128521669460981\\
&\quad 866051005968073733368199418319346275 03 \cdot Q_2,\\
j &= 367483248146527741437037734499498054316896144749783289195959484\\
&\quad 898077051895599992095468628645832990369395943292206321372511169167\\
&\quad 415557171412324682700995229079865594462515800975577276460889605379\\
&\quad 459426165927055505210049491084 3 \cdot \eta + 2159192778373049957841243789 75\\
&\quad 633208239639309170229661310740256289784137733921030657739288868481\\
&\quad 580729538629166826205051988346315263539811782455761790282875866683\\
&\quad 526823659841622119942133146159837672746928297797845060373732545 9.
\end{aligned}$$

The other colliding isogenies have kernels generated by H_1 and H_2 respectively, and both their images have j-invariant equal to j':

$$\begin{aligned}
H_1 = {} & 19326264832755263436205782983289563837195637516764115445955224 \\
& 9805508664589936678565722236003305074981600179788901 \cdot P_2 + 1036601 \\
& 26052076088276479166027473839647278991981771729449781337774031412 \\
& 9132966974448484055673967909895094673096769 \cdot Q_2, \\
H_2 = {} & 88315269100529870147635044900489166652883464217472613054693194 \\
& 0786972410641216900983358627043999654855326427935366 \cdot P_2 + 1036601 \\
& 26052076088276479166027473839647278991981771729449781337774031412 \\
& 9132966974448484055673967909895094673096764 \cdot Q_2, \\
j' = {} & 667988526030402783860739089287182477825842819805123827815659483 \\
& 0073908421726478925548196610102736809930593778696776988660092222030 \\
& 2366982678920913641761199526178050895453755324602626966343878887318 \\
& 5640825191720946820565156665988 \cdot \eta + 215919277837304995784124378975 \\
& 6332082396393091702296613107402562897841377339210306577392888868481 \\
& 5807295386291668262050519883463152635398117824557617902828758666683 \\
& 5268236598416221199421331461598376727469282977978450603737325459.
\end{aligned}$$

Appendix B

Lemma 6. *In the context of Remark 4, $\sum_{s^2 \le 4\ell^{2e}} H(4\ell^{2e} - s^2)^2$ is in $o(\ell^{4e})$.*

Proof. Let $\varepsilon > 0$. By the Brauer-Siegel theorem [Bra47], there exists $c_\varepsilon > 0$ such that $h(-d) < c_\varepsilon d^{\varepsilon + 1/2}$.
Let $\mathcal{N}_s$ be the number of square divisors of $4\ell^{2e} - s^2$ and $\mathcal{N} := \max_{s^2 \le 4\ell^{2e}} \mathcal{N}_s$. Then, we can write

$$\begin{aligned}
\sum_{s^2 \le 4\ell^{2e}} H(4\ell^{2e} - s^2)^2 &= \sum_{s^2 \le 4\ell^{2e}} \left(\sum_{d \cdot \mathfrak{f}^2 = 4\ell^{2e} - s^2} \frac{h(-d)}{u(d)} \right)^2 \\
&\le \sum_{s^2 \le 4\ell^{2e}} \left(\sum_{d \cdot \mathfrak{f}^2 = 4\ell^{2e} - s^2} c_\varepsilon d^{\varepsilon + 1/2} \right)^2 \\
&\le c_\varepsilon^2 \sum_{s^2 \le 4\ell^{2e}} \left(\mathcal{N} \cdot (4\ell^{2e})^{\varepsilon + 1/2} \right)^2 \\
&\le c_\varepsilon^2 \sum_{s^2 \le 4\ell^{2e}} \mathcal{N}^2 \cdot (4\ell^{2e})^{1 + 2\varepsilon} \\
&\le c_\varepsilon^2 \cdot 4\ell^e \cdot \mathcal{N}^2 \cdot (4\ell^{2e})^{1 + 2\varepsilon} \\
&= k_\varepsilon \cdot \mathcal{N}^2 \cdot (\ell^e)^{3 + 4\varepsilon},
\end{aligned}$$

where $k_\varepsilon := 16^{1+\varepsilon} c_\varepsilon^2$. So, taking $\varepsilon = \frac{1}{8}$

$$\begin{aligned} \lim \frac{\sum_{s^2 \leq 4\ell^{2e}} H(4\ell^{2e} - s^2)^2}{\ell^{4e}} &\leq \lim \frac{k_{\frac{1}{8}} \cdot \mathcal{N}^2 \cdot (\ell^e)^{3.5}}{\ell^{4e}} \\ &\leq k_{\frac{1}{8}} \lim \frac{\mathcal{N}^2}{\ell^{e/2}}. \end{aligned} \tag{26}$$

But now, it is known that the number of divisors of α is in $o(\alpha^\delta)$ for all $\delta > 0$, i.e. for all $\delta > 0$,

$$\lim \frac{\#\{\text{divisors of } \alpha\}}{\alpha^\delta} = 0.$$

Thus,

$$\lim \frac{\mathcal{N}_s}{4\ell^{2e\delta}} = 0, \forall \delta > 0. \tag{27}$$

Whence, going back to (26), and taking $\delta = 1/4$ in (27), we get

$$\lim \frac{\sum_{s^2 \leq 4\ell^{2e}} H(4\ell^{2e} - s^2)^2}{\ell^{4e}} = 0.$$

In other words, $\sum_{s^2 \leq 4\ell^{2e}} H(4\ell^{2e} - s^2)^2$ is in $o(\ell^{4e})$.