SIDH Proof of Knowledge

Luca De Feo[1], Samuel Dobson[2], Steven D. Galbraith[2], and Lukas Zobernig[2]

[1]IBM Research Europe. luca@defeo.lu
[2]Mathematics Department, University of Auckland, New Zealand.
samuel.dobson.nz@gmail.com, s.galbraith@auckland.ac.nz,
lukas.zobernig@auckland.ac.nz

February 16, 2022

Abstract

We demonstrate the soundness proof for the De Feo–Jao–Plût identification scheme (the basis for SIDH signatures) contains an invalid assumption and provide a counterexample for this assumption—thus showing the proof of soundness is invalid. As this proof was repeated in a number of works by various authors, multiple pieces of literature are affected by this result. Due to the importance of being able to prove knowledge of an SIDH key (for example, to prevent adaptive attacks), soundness is a vital property. We propose a modified sigma protocol (and hence identification scheme) fixing the issue with the De Feo–Jao–Plût scheme, and provide a proof of security of this new scheme. We also prove that a modification of this scheme allows the torsion points in the public key to be verified too. This results in a secure proof of knowledge for SIDH keys. It also avoids the SIDH identification scheme soundness issue raised by Ghantous, Pintore and Veroni (ePrint 2021/1051). In particular, this protocol provides a non-interactive way of verifying that SIDH public keys are well-formed as protection against adaptive attacks, leading to more efficient SIDH-based non-interactive key exchange (NIKE).

1 Introduction

While Supersingular Isogeny Diffie-Hellman (SIDH) [JD11, DJP14] is a fast and efficient post-quantum key exchange candidate, it has been hampered by the existence of practical adaptive attacks on the scheme—the first of these given by Galbraith et al. [GPST16] (the GPST attack), followed by other variations [FP21, UXT+22]. These attacks mean it is not safe to re-use a static key across multiple SIDH exchanges without other forms of protection. As such, various countermeasures have been proposed—though each with its unique drawbacks.

The first of these is to require one participant to use a one-time ephemeral key in the exchange, accompanied by a Fujisaki–Okamoto-type transform [HHK17] revealing the corresponding secret to the other party. This allows the recipient to verify the public key is well-formed, ensuring an adaptive attack was not used. This is what was done in SIKE [ACC+17], and converts the scheme to a secure key encapsulation mechanism (KEM). But it is of limited use in cases where both parties wish to use a long-term key.

The second countermeasure is to use many SIDH exchanges in parallel, combining all the resulting secrets into a single value, as proposed by Azarderakhsh, Jao, and Leonardi [AJL17]. This scheme is known as k-SIDH, where k is the number of keys used by each party in the exchange. The authors suggest $k = 92$ is required for a secure key exchange, as Dobson et al. [DGL+20] demonstrate how the GPST adaptive attack can be ported to $k = 2$ and above. Note that the number of SIDH instances grows as k^2, so this scheme is very inefficient. Urbanik and Jao's [UJ20] proposal attempted to improve the efficiency of this protocol by making use of the special automorphisms on curves with j-invariant 0 or 1728, but it was shown by Basso et al. [BKM+20] that Urbanik and Jao's proposal is vulnerable to a more efficient adaptive attack and actually

scales worse in efficiency than k-SIDH itself (although the public keys are approximately 4/5 of the size, it requires around twice as many SIDH instances for the same security).

Finally, adaptive attacks can also be prevented by providing a non-interactive proof that a public key is well-formed or honestly generated. While generic NIZKs would make this possible in a very inefficient manner, Urbanik and Jao [UJ20] claim a method for doing so using a similar idea to their k-SIDH improvement mentioned above. Their scheme is based on the SIDH-based identification scheme by De Feo, Jao, and Plût [DJP14], which is a fairly simple proof with single bit challenges.

Unfortunately, however, we show that the soundness of this original De Feo–Jao–Plût scheme is not rigorously proved—specifically that it does not reduce to the computational assumption they claim—and give a counterexample to this proof. Because this scheme (and proof) has since been used to build an undeniable signature by Jao and Soukharev [JS14], a signature scheme by Yoo, Azarderakhsh, Jalali, Jao, and Soukharev [YAJ+17], and also by Galbraith, Petit, and Silva [GPS20], all of these subsequent papers suffer from the same issue. Our counterexample does not apply to Urbanik and Jao's scheme, but their soundness proof nonetheless does not hold for the same reason.

Ghantous, Pintore, and Veroni [GPV21] have demonstrated that the soundness property for the De Feo–Jao–Plût scheme (and those based on it) fails for a different reason—namely the existence of multiple isogenies of the same degree between some curves. The protocols we propose in this paper are not vulnerable to the same issue, as we briefly discuss in Remark 4.

1.1 Our contributions

We present two new proofs that an SIDH public key is well-formed—meaning that, for base curve E_0 and public-key curve E_1, there is an isogeny $\phi : E_0 \to E_1$ of the correct degree (the private key or *witness*).

First, in Section 5, we propose a modification to the De Feo–Jao–Plût scheme that ensures that there is an extractor for the witness $\phi : E_0 \to E_1$. We express this in terms of a relation we call the *weak* SIDH relation. The first key idea in this protocol is the provision of bases (P_2, Q_2) for $E_2[\ell_2^{e_2}]$ and (P_3, Q_3) for $E_3[\ell_2^{e_2}]$. This allows the verifier to check that $(P_3, Q_3) = (\phi'(P_2), \phi'(Q_2))$ in the $\mathsf{chall} = 1$ case, and in the $\mathsf{chall} = 0$ case, to check that the isogenies from E_2 to E_0 and E_3 to E_1 are "parallel". The second key idea is, in the 2-special soundness proof, to view the transcript as an SIDH square where E_2 is treated as the "base curve" (instead of E_0), and where E_0 and E_3 play the roles of the participants' two public-key curves in SIDH. It then follows that there is a witness ϕ as required.

Second, in Section 6, we give a new scheme that convinces a verifier not only that there is an isogeny $\phi : E_0 \to E_1$ of the correct degree, but also that the torsion points provided in an SIDH public key are the correct images of the public parameter points under ϕ. We call this stronger relation the SIDH relation. Making this non-interactive using the Fiat-Shamir heuristic gives a secure method for proving well-formedness of SIDH public keys, which is needed if one wants to prevent adaptive attacks. This is the first such protocol in the literature and has important applications in all areas where SIDH key exchange could be used with static keys. Our scheme works with any base elliptic curve, rather than being restricted to the two curves with j-invariant 0 or 1728 as in [UJ20].

The scheme in Section 6 builds on the protocol of Section 5. However, it requires assurance that the ephemeral isogenies used in the commitments by the prover are "independent enough". To achieve this, we "double" the protocol, by essentially running two sessions of the protocol from Section 5 for each challenge bit. The prover shows that the two instances are consistent with each other by providing images of a random torsion basis in both squares, which the verifier can check are correct. The verifier also checks that the two instances are independent. This allows us to construct a 2-special soundness extractor that outputs a correct witness.

Commitments in the original De Feo–Jao–Plût scheme were just j-invariants of curves, but our new proofs require committing to various points on curves as well. This makes the proofs larger. As with the original

De Feo–Jao–Plût scheme, it is non-trivial to simulate valid protocol transcripts without knowing the witness and so we only achieve computational zero-knowledge. We introduce some new assumptions to prove the same for our schemes.

1.2 Plan of the paper

Section 2 recalls the SIDH protocol and gives some useful lemmas that are used in our soundness proofs. Section 3 presents some isogeny-based hardness assumptions and reductions, including the new decisional assumptions we need for our zero-knowledge proofs. We then recall the De Feo–Jao–Plût identification scheme in Section 4.1 and outline the issue with its proof of soundness in Section 4.2. Our first new SIDH proof is given in Section 5. We then show how the points in the SIDH public key can also be verified in Section 6. Finally, we conclude with some standard discussion on how a secure signature scheme which is a Proof of Knowledge (PoK) of an SIDH secret key can be constructed from our second scheme—the first such scheme which is sound and proves correctness of the points in the public key (a protection mechanism against adaptive attacks [GPST16, DGL+20]) in Section 7.

1.3 Acknowledgements

We thank David Jao, Jason LeGrow, and Yi-Fu Lai for useful discussion about this work. We also thank Paulo Barreto for catching some typos in this paper, and Simon-Philipp Merz for valuable comments. We thank Javad Doliskani for important observations that inspired significant improvements to this work. Finally, we would like to thank those involved with the BIRS Supersingular Isogeny Graphs in Cryptography workshop for great discussion on some questions this work raised—especially Lorenz Panny and his work analyzing SIDH squares in small fields.

2 Preliminaries

Notation. As a convention, we will use K_ϕ to denote a point which generates the kernel of an isogeny ϕ. Let $[t]$ denote the set $\{1, \ldots, t\}$.

2.1 SIDH

We now provide a brief refresher on the Supersingular Isogeny Diffie-Hellman (SIDH) key exchange protocol [JD11, DJP14] by De Feo, Jao, and Plût.

As public parameters, we have a prime $p = \ell_1^{e_1} \cdot \ell_2^{e_2} \cdot f \pm 1$, where ℓ_1, ℓ_2 are small primes, f is an integer cofactor, and $\ell_1^{e_1} \approx \ell_2^{e_2}$. We work over the finite field $\mathbb{F}_{p^2}$. Additionally we fix a base supersingular elliptic curve E and bases $\{P_1, Q_1\}$, $\{P_2, Q_2\}$ for both the $\ell_1^{e_1}$ and $\ell_2^{e_2}$-torsion subgroups of $E(\mathbb{F}_{p^2})$ respectively (such that $E[\ell_i^{e_i}] = \langle P_i, Q_i \rangle$). Typically $\ell_1 = 2$ and $\ell_2 = 3$.

It is well known that knowledge of an isogeny (up to isomorphism) and knowledge of its kernel are equivalent, and we can convert between them at will, via Vélu's formulae [Vél71]. In SIDH, the secret keys of Alice and Bob are isogenies $\phi_A : E(\mathbb{F}_{p^2}) \to E_A(\mathbb{F}_{p^2})$, $\phi_B : E(\mathbb{F}_{p^2}) \to E_B(\mathbb{F}_{p^2})$ of degree $\ell_1^{e_1}$ and $\ell_2^{e_2}$, respectively. These isogenies are generated by randomly choosing secret integers $a_i, b_i \in \mathbb{Z}/\ell_i^{e_i}\mathbb{Z}$ (not both divisible by ℓ_i) and computing the isogeny with kernel generated by $K_i = [a_i]P_i + [b_i]Q_i$. We thus unambiguously refer to the isogeny, its kernel, and such integers a, b, as "the secret key."

Figure 1 depicts the commutative diagram making up the key exchange. In order to make the diagram commute, Alice and Bob are required to not only give their image curves E_A and E_B in their respective public keys, but also the images of the basis points of the other participant's kernel on E. That is, Alice provides E_A, $P_2' = \phi_A(P_2), Q_2' = \phi_A(Q_2)$ as her public key. This allows Bob to "transport" his secret isogeny to E_A and compute ϕ_{AB} whose kernel is $\langle [a_2]P_2' + [b_2]Q_2' \rangle$. Both Alice and Bob will arrive along these transported isogenies at isomorphic image curves E_{AB}, E_{BA} (using Vélu's formulae, they will actually

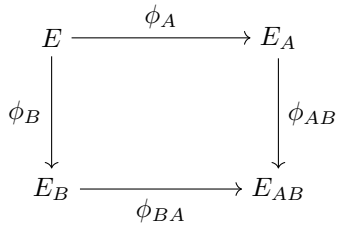

Figure 1: Commutative diagram of SIDH, where $\ker(\phi_{BA}) = \phi_B(\ker(\phi_A))$ and $\ker(\phi_{AB}) = \phi_A(\ker(\phi_B))$.

arrive at exactly the same curve [Leo20]). Two elliptic curves are isomorphic over $\overline{\mathbb{F}}_{p^2}$ if and only if their j-invariants are equal, $j(E_{AB}) = j(E_{BA})$, hence this j-invariant may be used as the shared secret of the SIDH key exchange.

Some cryptographic hardness assumptions related to isogenies and SIDH are discussed in Section 3.

2.2 Isogeny squares

We collect here some basic definitions and lemmas that we will use repeatedly throughout the paper. In the statements below, all elliptic curves are defined over a field of characteristic p.

Definition 1 (Independent points, isogenies)**.** *Let E be an elliptic curve, let $\ell \neq p$ be a prime and e an integer, let (P, Q) be a basis of $E[\ell^e]$. Let $R = [a]P + [b]Q$ and $S = [c]P + [d]Q$. The following conditions are equivalent:*

(a) (R, S) form a basis of $E[\ell^e]$.

(b) ℓ does not divide $ad - bc$, i.e., the matrix $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ is invertible modulo ℓ^e.

(c) The value of the ℓ^e-th Weil paring $w = e(R, S)$ has order ℓ^e, i.e., $w^{\ell^{e-1}} \neq 1$.

When R, S satisfy any of these, we say they are independent *of one another. Similarly, we say that two cyclic groups of order ℓ^e are independent whenever any of their generators are. Finally, we say that two isogenies of degree ℓ^e are independent if their kernels are.*

Proof. $(a) \Rightarrow (b)$: Both P, Q and R, S are bases of the same torsion subgroup $E[\ell^e]$. Hence, $A = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ is a change-of-basis from P, Q to R, S and there must be an inverse change-of-basis A^{-1} from R, S to P, Q. Then A is necessarily invertible, and therefore, so too is its determinant $ad - bc$ modulo ℓ^e.

$(b) \Rightarrow (c)$: We have that

$$w = e(R, S) = e([a]P + [b]Q, [c]P + [d]Q).$$

Then since e is bilinear, $w = e(P, Q)^{ad-bc}$. Now $e(P, Q)$ has order ℓ^e because e is surjective onto the group of ℓ^e-th roots of unity (c.f. [Sil09, Corollary III.8.1.1]), and since $\ell \nmid ad - bc$, then w must also have order ℓ^e.

$(c) \Rightarrow (a)$: Recall that $E[\ell^e] \simeq \mathbb{Z}/\ell^e\mathbb{Z} \times \mathbb{Z}/\ell^e\mathbb{Z}$ [Sil09, Corollary III.6.4b]. Thus, in order for R, S to form a basis, we must show $\langle R \rangle \cap \langle S \rangle = \{\mathcal{O}_E\}$.

Suppose $[w]R = [z]S \neq \mathcal{O}_E$ for some integers w, z. By assumption, it must be that $\ell^e \nmid w$ and $\ell^e \nmid z$. Now consider $e([w]R - [z]S, S) = 1$, since $e(\mathcal{O}_E, T) = 1$ for any T. By the bilinearity of the pairing, this gives

$$e([w]R - [z]S, S) = e(R, S)^w e(S, S)^{-z} = 1.$$

Then, because $e(S, S) = 1$, we arrive at the conclusion $e(R, S)^w = 1$, which is a contradiction since $e(R, S)$ has order ℓ^e and $\ell^e \nmid w$. Thus, there can exist no such integers w, z, and therefore $\langle R \rangle \cap \langle S \rangle = \{\mathcal{O}_E\}$. □ □

Lemma 1. *Let $\phi : E \to E/\langle R \rangle$ be an isogeny of kernel $\langle R \rangle$ and degree ℓ^e, let S be a point of order ℓ^e independent to R. Then $\phi(S)$ has order ℓ^e and generates* $\ker(\widehat{\phi})$.

Proof. Because R and S are independent (Definition 1), the subgroups generated by R and S intersect trivially. Thus, since ϕ has kernel $\langle R \rangle$, no non-trivial point in $\langle S \rangle$ is in the kernel of ϕ. Furthermore, we know that $\widehat{\phi} \circ \phi = [\ell^e]$ has kernel $E[\ell^e]$, and that $S \in E[\ell^e]$. Thus $\widehat{\phi}(\phi(S)) = \mathcal{O}$, implying $\phi(S)$ is in the kernel of $\widehat{\phi}$. The same holds for all elements $S' = [\lambda]S \in \langle S \rangle$, and since $\phi(S') \neq \mathcal{O}$ for all non-trivial S', $\phi(S)$ has order ℓ^e and generates $\ker(\widehat{\phi})$. □ □

The following lemma is the main tool we are going to use, repeatedly, to design all proofs of knowledge.

Lemma 2. *Let ℓ_1, ℓ_2 be distinct primes different from p, let e_1, e_2 be integers. Let $\phi_A : E \to E_A$ be an isogeny of degree $\ell_1^{e_1}$. Let $\phi_B : E \to E_B$ and $\phi_{AB} : E_A \to E_{AB}$ be isogenies of degree $\ell_2^{e_2}$ such that* $\ker(\phi_{AB}) = \phi_A(\ker(\phi_B))$. *Then there exists an isogeny $\phi_{BA} : E_B \to E_{AB}$ of degree $\ell_1^{e_1}$.*

Additionally, let $S \in E$ be a point of order $\ell_2^{e_2}$ such that $\ker(\phi_B)$ *and $\langle S \rangle$ are independent, let $S_B = \phi_B(S)$ and let $S_{AB} = \phi_{AB}(\phi_A(S))$. Then S_B and S_{AB} both have order $\ell_2^{e_2}$ and generate, respectively,* $\ker(\widehat{\phi_B})$ *and* $\ker(\widehat{\phi_{AB}})$. *Moreover, $\phi_{BA}(S_B) = S_{AB}$.*

This is visualized in Figure 2.

Proof. Let K_A be a generator of $\ker(\phi_A)$. Then because the degrees of ϕ_A, ϕ_B are coprime, $\phi_B(K_A)$ also has order $\ell_1^{e_1}$ and generates the kernel of some isogeny

$$\chi : E_B \to E_B/\langle \phi_B(K_A) \rangle.$$

Observe that E_{AB} is defined as the codomain of $\phi_{AB} \circ \phi_A$. We thus have that $E_{AB} \cong E/\langle K_A, K' \rangle$ for a point K' of order $\ell_2^{e_2}$ such that $\langle \phi_A(K') \rangle = \ker(\phi_{AB})$. Because $\ker(\phi_{AB}) = \phi_A(\ker(\phi_B))$, we conclude $\langle K' \rangle = \ker(\phi_B)$. Therefore, $E_B/\langle \phi_B(K_A) \rangle \cong E_{AB}$ as required.

By the conditions on S, Lemma 1 shows that $S_B = \phi_B(S)$ generates $\ker(\widehat{\phi_B})$.

One can verify that using Vélu's formula [Vél71], $\phi_{AB}(\phi_A(P)) = \phi_{BA}(\phi_B(P))$ for any point $P \in E$ (see [Leo20, Lemma 1]). Hence,

$$\begin{aligned} S_{AB} &= \phi_{AB}(\phi_A(S)) \\ &= \phi_{BA}(\phi_B(S)) = \phi_{BA}(S_B) \end{aligned}$$

Finally, because $\langle S \rangle$ and $\ker(\phi_B)$ are independent, and because $\ker(\phi_{AB}) = \phi_A(\ker(\phi_B))$, then $\langle \phi_A(S) \rangle$ and $\ker(\phi_{AB})$ must also be independent. $\phi_A(S)$ must have order $\ell_2^{e_2}$ since the degree of ϕ_A is coprime to the order of S. So by applying Lemma 1 again, we arrive at the conclusion that S_{AB} generates $\ker(\widehat{\phi_{AB}})$. □ □

The lemma above suggests an algorithm to compute the points S_B and S_{AB}, even when the isogeny ϕ_A is only known through its action on $E[\ell_2^{e_2}]$. We present such an algorithm in Figure 3.

2.3 Sigma protocols

A sigma protocol Π_Σ for a relation $\mathcal{R} = \{(X, W)\}$ is a public-coin three-move interactive proof system consisting of two parties: a verifier V and a prover P. Recall that public-coin informally means that there are no secret sources of randomness—the verifier's coin tosses are accessible to the prover. In practice this means the challenge sent by the verifier to the prover is uniformly random. For our purposes, a witness W can be thought of as a secret key, while the statement X is the corresponding public key. Thus, proving $(X, W) \in \mathcal{R}$ is equivalent to saying that X is a valid public key for which a corresponding secret key exists. We use the security parameter κ to parametrize the length of the secret keys involved.

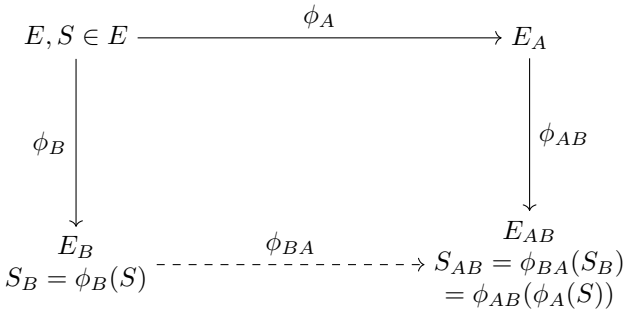

Figure 2: Lemma 2, visualized. The lemma shows that ϕ_{BA} exists and the equality on S_{AB} from both directions holds.

Input: $(E, P, Q, E_A, P_A, Q_A, \phi_B, \phi_{AB})$ such that $\langle P, Q \rangle = E[\ell_2^{e_2}]$, $\phi_B : E \to E_B$ and $\phi_{AB} : E_A \to E_{AB}$ have degree $\ell_2^{e_2}$, and for some isogeny $\phi_A : E \to E_A$ of degree $\ell_1^{e_1}$, we have $\ker(\phi_{AB}) = \phi_A(\ker(\phi_B))$, $P_A = \phi_A(P)$, and $Q_A = \phi_A(Q)$.

Output: $(S, \phi_{BA}(S))$ where $S \in E_B$ and $\phi_{BA} : E_B \to E_{AB}$ is an isogeny such that $\ker(\phi_{BA}) = \phi_B(\ker(\phi_A))$.

1: Find a point $S \in E$ such that $S_B = \phi_B(S)$ generates the kernel of $\widehat{\phi_B}$. In fact, it suffices to have either $S = P$ or Q.
2: Write $S = [a]P + [b]Q$.
3: Compute $S_A = [a]P_A + [b]Q_A$ Then $S_A = \phi_A(S)$ despite ϕ_A being unknown.
4: Compute $S_{AB} = \phi_{AB}(S_A)$.
5: Output (S_B, S_{AB}).

Figure 3: Algorithm to compute image of a single point under hidden isogeny ϕ_{BA}, as per Lemma 2.

Definition 2 (Sigma protocol)**.** *A sigma protocol Π_Σ for a family of relations $\{\mathcal{R}\}_\kappa$ parametrized by security parameter κ consists of PPT algorithms $((P_1, P_2), (V_1, V_2))$ where V_2 is deterministic and we assume P_1, P_2 share states. The protocol proceeds as follows:*

1. *Round 1: The prover, on input $(X, W) \in \mathcal{R}$, returns a commitment $\mathsf{com} \leftarrow P_1(X, W)$ which is sent to the verifier.*
2. *Round 2: The verifier, on receipt of com, runs $\mathsf{chall} \leftarrow V_1(1^\kappa)$ to obtain a random challenge, and sends this to the prover.*
3. *Round 3: The prover then runs $\mathsf{resp} \leftarrow P_2(X, W, \mathsf{chall})$ and returns resp to the verifier.*
4. *Verification: The verifier runs $V_2(X, \mathsf{com}, \mathsf{chall}, \mathsf{resp})$ and outputs either $\top$ (accept) or $\bot$ (reject).*

A transcript $(\mathsf{com}, \mathsf{chall}, \mathsf{resp})$ is said to be valid if $V_2(X, \mathsf{com}, \mathsf{chall}, \mathsf{resp})$ outputs $\top$. Let $\langle P, V \rangle$ denote the transcript for an interaction between prover P and verifier V. Relevant properties of a sigma protocol are:

Correctness: If the prover P knows $(X, W) \in \mathcal{R}$ and behaves honestly, then the verifier V accepts.

2-special soundness: There exists a polynomial-time extraction algorithm $\mathsf{Extract}$ that, given a statement X and two valid transcripts $(\mathsf{com}, \mathsf{chall}, \mathsf{resp})$ and $(\mathsf{com}, \mathsf{chall}', \mathsf{resp}')$ where $\mathsf{chall} \neq \mathsf{chall}'$, outputs a witness W such that $(X, W) \in \mathcal{R}$ with probability at least $1 - \varepsilon$ for soundness error ε.

Zero-knowledge (ZK): There exists a polynomial-time simulator Sim that, given a statement X for any

$(X, W) \in \mathcal{R}$, and for any (cheating) verifier V^*, outputs transcripts $(\mathsf{com}, \mathsf{chall}, \mathsf{resp})$ that are indistinguishable from valid interactions between a prover P and V^*.

Proof of Knowledge (PoK): There exists a polynomial-time extraction algorithm $\mathsf{Extract}$ that, given an arbitrary statement X and access to any prover P^*, outputs a witness W such that $(X, W) \in \mathcal{R}$ with probability at least $\Pr[\langle P^*, V\rangle = 1] - \varepsilon$ for knowledge error ε.

It is a known result (e.g. by Hazay and Lindell [HL10, Theorem 6.3.2]) that a correct and 2-special sound sigma protocol with challenge length t is a Proof of Knowledge with knowledge error 2^{-t}. In this paper, this will generally be a single-bit challenge sigma protocol repeated with t iterations.

3 SIDH problems and assumptions

In this section, we recall some standard isogeny-based hardness assumptions of relevance to this work. We then introduce a new decisional assumption which will be useful for the proof of zero-knowledge in Section 5. The first two are computational isogeny-finding problems.

Definition 3 (General isogeny problem)**.** *Given j-invariants $j, j' \in \mathbb{F}_{p^2}$, find an isogeny $\phi : E \to E'$ if one exists, where $j(E) = j$ and $j(E') = j'$.*

This is the foundational hardness assumption of isogeny-based cryptography, that it is hard to find an isogeny between two given curves. Note the decisional version, determining whether an isogeny exists, is easy—an isogeny exists if and only if $\#E(\mathbb{F}_{p^2}) = \#E'(\mathbb{F}_{p^2})$.

Definition 4 (Computational Supersingular Isogeny (CSSI) problem)**.** *For fixed SIDH prime p, base curve E_0, and $\ell_2^{e_2}$-torsion basis $P_0, Q_0 \in E_0$, let $\phi : E_0 \to E_1$ be an isogeny of degree $\ell_1^{e_1}$. Given an SIDH public key $(E_1, P_1 = \phi(P_0), Q_1 = \phi(Q_0))$, find an isogeny $\phi' : E_0 \to E_1$ of degree $\ell_1^{e_1}$ such that $P_1, Q_1 = \phi'(P_0)$, $\phi'(Q_0)$.*

This is problem 5.2 of [DJP14] and essentially states that it is hard to find the secret key corresponding to a given public key. This problem is also called the SIDH isogeny problem by [GV18, Definition 2].

At the heart of the adaptive attack is the problem that, given a public key (E_1, P_1, Q_1), we cannot validate that P_1, Q_1 are indeed the correct images of basis points P_0, Q_0 under the secret isogeny ϕ. The best we know how to do is to check they are indeed a basis of the correct order, and use the Weil pairing check

$$e_{\ell_2^{e_2}}(P_1, Q_1) = e_{\ell_2^{e_2}}(P_0, Q_0)^{\deg \phi}.$$

Unfortunately this holds for many different choices of basis points. Indeed, if (P_1, Q_1) are the correct images, then any pair $(aP_1 + bQ_1, cP_1 + dQ_1)$ such that $ad - bc = 1 \bmod \ell_2^{e_2}$ also passes the check. So this is not enough to uniquely determine ϕ, and, in particular, is insufficient to protect against the GPST adaptive attack.

The following decisional problem follows Definition 3 of [GV18] and is also very similar to the key validation problem of Urbanik and Jao [UJ18, Problem 3.4] (the key validation problem asks whether a ϕ of degree *dividing* $\ell_1^{e_1}$ exists). However, the previous definitions did not take the Weil pairing check into account, which would serve as a distinguisher.

Definition 5 (Decisional SIDH isogeny (DSIDH) problem)**.** *The decisional SIDH problem is to distinguish between the following two distributions:*

- $\mathcal{D}_0 = \{(E_0, P_0, Q_0, E_1, P_1, Q_1)\}$ *such that E_0 is a supersingular elliptic curve defined over $\mathbb{F}_{p^2}$, P_0, Q_0 a basis such that $E_0[\ell_2^{e_2}] = \langle P_0, Q_0\rangle$, $\phi : E_0 \to E_1$ is an isogeny of degree $\ell_1^{e_1}$, and $P_1 = \phi(P_0)$ and $Q_1 = \phi(Q_0)$.*

- $\mathcal{D}_1 = \{(E_0, P_0, Q_0, E_1, P_1, Q_1)\}$ *such that* E_0 *is a supersingular elliptic curve defined over* $\mathbb{F}_{p^2}$, P_0, Q_0 *a basis such that* $E_0[\ell_2^{e_2}] = \langle P_0, Q_0 \rangle$, E_1 *is any supersingular elliptic curve over* $\mathbb{F}_{p^2}$ *with the same cardinality as* E_0, *and* P_1, Q_1 *is a basis of* $E_1[\ell_2^{e_2}]$ *satisfying the Weil pairing check* $e_{\ell_2^{e_2}}(P_1, Q_1) = e_{\ell_2^{e_2}}(P_0, Q_0)^{\ell_1^{e_1}}$.

As shown by Galbraith and Vercauteren [GV18], Thormarker [Tho17], and Urbanik and Jao [UJ18], being able to solve this decisional problem is as hard as solving the computational (CSSI) problem, so key validation is fundamentally difficult. This is done by testing ℓ_1-isogeny neighboring curves of E_1 and learning the correct path one bit at a time.

Definition 6 (Decisional Supersingular Product (DSSP) problem)**.** *Let* $\phi : E_0 \to E_1$ *be an isogeny of degree* $\ell_1^{e_1}$. *Let* $P_0, Q_0 \in E_0[\ell_2^{e_2}]$ *be a fixed basis of the* $\ell_2^{e_2}$*-torsion subgroup. Suppose we have the following two distributions:*

- $\mathcal{D}_0 = \{(E_2, E_3, \phi')\}$ *such that there exists a cyclic subgroup* $G \subseteq E_0[\ell_2^{e_2}]$ *of order* $\ell_2^{e_2}$ *and* $E_2 \cong E_0/G$ *and* $E_3 \cong E_1/\phi(G)$, *and* $\phi' : E_2 \to E_3$ *is a degree* $\ell_1^{e_1}$ *isogeny.*
- $\mathcal{D}_1 = \{(E_2, E_3, \phi')\}$ *such that* E_2 *is a random supersingular curve with the same cardinality as* E_0, *and* E_3 *is the codomain of a random isogeny* $\phi' : E_2 \to E_3$ *of degree* $\ell_1^{e_1}$.

Let $\mathcal{O}^{\mathsf{DSSP}}$ *be an oracle which behaves as follows. On setup, with public parameters* $(E_0, P_0, Q_0, E_1, \phi(P_0), \phi(Q_0))$, *it chooses a uniformly random secret bit* $b \leftarrow \{0, 1\}$. *Each time it is queried,* $\mathcal{O}^{\mathsf{DSSP}}$ *returns a tuple from distribution* $\mathcal{D}_b$. *The DSSP problem is then, given access to such an oracle, to determine* b.

This is problem 5.5 of [DJP14] and intuitively states that it is hard to determine whether there exist valid "vertical sides" to an SIDH square given the corners and the bottom horizontal side.

3.1 A new hardness assumption

We define a new decisional isogeny assumption which will be useful for the proof of zero-knowledge in Section 5. This assumption can intuitively be seen as a "parallel" version of the DSIDH assumption above.

Definition 7 (Decisional Mirror SIDH (DMSIDH) problem)**.** *Let* $\phi : E_0 \to E_1$ *be an isogeny of degree* $\ell_1^{e_1}$. *Let* P_0, Q_0 *be a basis for the* $\ell_2^{e_2}$*-torsion subgroup* $E_0[\ell_2^{e_2}]$.

Define distributions $\mathcal{D}_0$ *and* $\mathcal{D}_1$ *as follows. Construct a random SIDH square by letting* $\psi : E_0 \to E_2$ *be a random isogeny of degree* $\ell_2^{e_2}$, *then* $\psi' : E_1 \to E_3$ *an isogeny of degree* $\ell_2^{e_2}$ *whose kernel is* $\phi(\ker(\psi))$, *and* $\phi' : E_2 \to E_3$ *an isogeny of degree* $\ell_1^{e_1}$ *whose kernel is* $\psi(\ker(\phi))$. *Construct a basis* S, T *of* $E_2[\ell_2^{e_2}]$ *with* $\langle S \rangle = \ker(\widehat{\psi})$. *Finally, the distributions are*

- $\mathcal{D}_0 = \{(\psi, \psi', S, T, \phi'(S), T')\}$ *where* $T' = \phi'(T)$
- $\mathcal{D}_1 = \{(\psi, \psi', S, T, \phi'(S), T')\}$ *where* $T' = \phi'(T + [r]S)$, *and* r *is random.*

Let $\mathcal{O}^{\mathsf{DMSIDH}}$ *be an oracle which, on setup with public parameters* $(E_0, P_0, Q_0, E_1, \phi(P_0), \phi(Q_0))$, *chooses a uniformly random secret bit* $b \leftarrow \{0, 1\}$, *then each time it is queried returns a sample from* $\mathcal{D}_b$. *The DMSIDH problem is, given access to* $\mathcal{O}^{\mathsf{DMSIDH}}$, *to determine* b. *The problem is visualized in Figure 4.*

In other words, $(E_1, \phi(P_0), \phi(Q_0))$ is an SIDH public key, and the ψ, ψ' are the vertical sides of an SIDH square. The challenge is to determine whether a point T' is the actual image of T under the hidden horizontal isogeny on the fourth (bottom) side of the SIDH square (which is guaranteed to exist).

Observe that, given an SIDH public key, one can already choose isogenies ψ, ψ' such that $\ker(\psi') = \phi(\ker(\psi))$. We can also obtain a point S and its image $\phi'(S)$ via these ψ and ψ'. This is possible due to Lemma 2 (and achieved in practice using the algorithm in Figure 3). Thus, the only additional information provided in the DMSIDH problem is a candidate image T' of one extra point T on E_2 (independent to S).

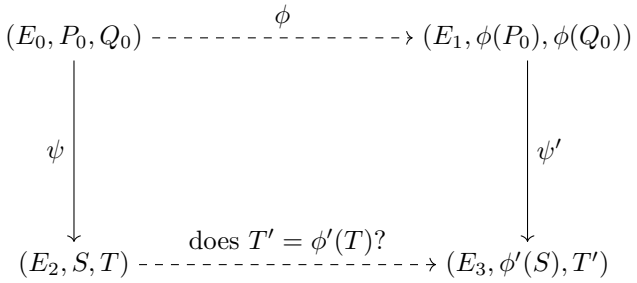

Figure 4: The Decisional Mirror SIDH (DMSIDH) problem (Definition 7) visualized. Dashed lines are secret and are not known by the adversary/distinguisher. S is such that $\langle S \rangle = \ker(\widehat{\psi})$.

3.2 Double variants

In Section 6, we propose a scheme which uses two independent SIDH squares in each round of the sigma protocol. For the zero-knowledge proof in that section, we require "double" variants of the DSSP and DMSIDH problems. We prove that the Double-DMSIDH problem is hard if the "single" version is. Hence, this definition is only needed as a tool to simplify the security proof.

The Double-DSSP problem differs from the "single" version by the introduction of two bases U_i', V_i' of the $\ell_1^{e_1}$-torsion subgroups on $E_{2,i}$, for $i \in \{0, 1\}$. As we shall see in Section 6, these extra points will be used to verify that the two independent SIDH squares in the "double" protocol both use consistent isogenies ϕ_i'. These extra points, plus the requirement that the isogenies ψ_i used in each of the two squares should be independent, mean a reduction from DSSP to the Double-DSSP problem is unlikely. We believe Double-DSSP is a hard problem.

Definition 8 (Double-DSSP Problem)**.** *For public parameters $(E_0, P_0, Q_0, E_1, \phi(P_0), \phi(Q_0))$, let $\mathcal{O}^{DSSP}$ be a DSSP instance generator oracle (with secret bit b). The double-DSSP problem is to distinguish between the following two distributions:*

- $\mathcal{D}_0 = \{(\mathsf{inst}_{i\in\{0,1\}}, U_i', V_i')\}$ *where* $\mathsf{inst}_i = (E_{2,i}, E_{3,i}, \phi_i') \leftarrow \mathcal{O}^{DSSP}$ *with* $b = 0$*, and additionally, if $\psi_i : E_0 \to E_{2,i}$ are the respective isogenies of degree $\ell_2^{e_2}$, then ψ_0 and ψ_1 are independent and $U_i', V_i' = \psi_i(U), \psi_i(V)$ where $\{U, V\}$ is a random basis of $E_0[\ell_1^{e_1}]$.*
- $\mathcal{D}_1 = \{(\mathsf{inst}_{i\in\{0,1\}}, U_i', V_i')\}$ *where* $\mathsf{inst}_i = (E_{2,i}, E_{3,i}, \phi_i') \leftarrow \mathcal{O}^{DSSP}$*, $b = 1$, and U_i', V_i' is a random basis of the $\ell_1^{e_1}$ torsion subgroup on $E_{2,i}$ such that $e_{\ell_1^{e_1}}(U_0', V_0') = e_{\ell_1^{e_1}}(U_1', V_1')$ and there is a pair (a, b) of integers such that the kernel of ϕ_i' is generated by $[a]U_i' + [b]V_i'$ for both $i \in \{0, 1\}$.*

Definition 9 (Double-DMSIDH Problem)**.** *For public parameters $(E_0, P_0, Q_0, E_1, \phi(P_0), \phi(Q_0))$, let $\mathcal{O}^{DMSIDH}$ be a DMSIDH instance generator oracle (with secret bit b). The double-DMSIDH problem is to distinguish between the following two distributions:*

- $\mathcal{D}_0 = \{\mathsf{inst}_{i\in\{0,1\}}\}$ *where* $\mathsf{inst}_i = (\psi_i, \psi_i', S_i, T_i, \phi_i'(S_i), T_i') \leftarrow \mathcal{O}^{DMSIDH}$*, $b = 0$, and ψ_0 and ψ_1 are independent.*
- $\mathcal{D}_1 = \{\mathsf{inst}_{i\in\{0,1\}}\}$ *where* $\mathsf{inst}_i = (\psi_i, \psi_i', S_i, T_i, \phi_i'(S_i), T_i') \leftarrow \mathcal{O}^{DMSIDH}$*, $b = 1$, and ψ_0 and ψ_1 are independent.*

Theorem 3. *If there exists an adversary $\mathcal{A}^{DDMSIDH}$ which makes n queries to a Double-DMSIDH oracle and guesses its bit with advantage $\mathsf{Adv}^{ddmsidh}$, then there exists an adversary that solves the DMSIDH problem (with oracle $\mathcal{O}^{DMSIDH}$) with the same advantage $\mathsf{Adv}^{ddmsidh}$, after making an expected $n(\ell_2 + 1)/\ell_2$ queries to $\mathcal{O}^{DMSIDH}$.*

Proof. Given a DMSIDH oracle $\mathcal{O}^{\mathsf{DMSIDH}}$, we simulate a Double-DMSIDH oracle as follows. Any time

$\mathcal{A}^{\mathsf{DDMSIDH}}$ asks for a sample we query $\mathcal{O}^{\mathsf{DMSIDH}}$ for $\mathsf{inst}_0 = (\psi_0, \psi_0', S_0, T_0, \phi_0'(S_0), T_0')$, then we keep querying $\mathcal{O}^{\mathsf{DMSIDH}}$ for $\mathsf{inst}_1 = (\psi_1, \psi_1', S_1, T_1, \phi_1'(S_1), T_1')$ until ψ_0 and ψ_1 are independent. Finally, we return $(\mathsf{inst}_0, \mathsf{inst}_1)$.

Write $\ker(\psi_i) = \langle [a_i]P_0 + [b_i]Q_0 \rangle$, and say that two pairs a_i, b_i ($i \in \{0,1\}$) are conjugate if $(a_0, b_0) = \lambda(a_1, b_1)$ for some invertible scalar λ. There are $\ell_2 + 1$ different such conjugacy classes of (a_i, b_i), and being in different conjugacy classes implies that $a_0'b_1' - a_1'b_0'$ is invertible. Thus, with probability $\ell_2/(\ell_2+1)$, any two random choices of ψ_i will be independent.

Therefore, if $\mathcal{A}^{\mathsf{DDMSIDH}}$ makes n queries to the Double-DMSIDH oracle, the simulation makes an expected number $n(\ell_2+1)/\ell_2$ of queries to $\mathcal{O}^{\mathsf{DMSIDH}}$. Because the simulation is perfect, whatever advantage $\mathcal{A}^{\mathsf{DDMSIDH}}$ has against Double-DMSIDH carries over to DMSIDH. □ □

4 Previous SIDH identification scheme and soundness issue

4.1 De Feo–Jao–Plût scheme

Let p be a large prime of the form $\ell_1^{e_1} \cdot \ell_2^{e_2} \cdot f \pm 1$, where ℓ_1, ℓ_2 are small primes. We start with a supersingular elliptic curve E_0 defined over $\mathbb{F}_{p^2}$ with $\#E_0(\mathbb{F}_{p^2}) = (\ell_1^{e_1}\ell_2^{e_2} f)^2$. The private key is a random point $K_\phi \in E_0(\mathbb{F}_{p^2})$ of exact order $\ell_1^{e_1}$. Define $E_1 = E_0/\langle K_\phi \rangle$ and denote the corresponding $\ell_1^{e_1}$-isogeny by $\phi : E_0 \to E_1$.

Let P_0, Q_0 be a basis of the torsion subgroup $E_0[\ell_2^{e_2}] = \langle P_0, Q_0 \rangle$. The fixed public parameters are $pp = (p, E_0, P_0, Q_0)$. The public key is $(E_1, \phi(P_0), \phi(Q_0))$. The private key is the kernel generator K_ϕ (equivalently, the isogeny ϕ). The interaction goes as follows:

1. The prover chooses a random primitive $\ell_2^{e_2}$-torsion point K_ψ as $K_\psi = [a]P_0 + [b]Q_0$ for some integers $0 \le a, b < \ell_2^{e_2}$ not both divisible by ℓ_2. Note that $\phi(K_\psi) = [a]\phi(P_0) + [b]\phi(Q_0)$. The prover defines the curves $E_2 = E_0/\langle K_\psi \rangle$ and $E_3 = E_1/\langle \phi(K_\psi) \rangle = E_0/\langle K_\psi, K_\phi \rangle$, and uses Vélu's formulae to compute the following diagram.

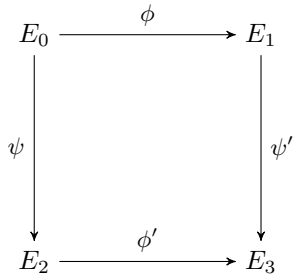

   The prover sends commitment $\mathsf{com} = (E_2, E_3)$ to the verifier.

2. The verifier challenges the prover with a random bit $\mathsf{chall} \leftarrow \{0,1\}$.

3. If $\mathsf{chall} = 0$, the prover reveals $\mathsf{resp} = (a, b)$ from which K_ψ and $\phi(K_\psi) = K_{\psi'}$ can be reconstructed. If $\mathsf{chall} = 1$, the prover reveals $\mathsf{resp} = (\psi(K_\phi) = K_{\phi'})$.

In both cases, the verifier accepts the proof if the points revealed have the correct order and generate kernels of isogenies between the correct curves. We iterate this process t times to reduce the cheating probability (where t is chosen based on the security parameter κ). Note that in an honest execution of the proof, we have

$$\widehat{\psi'} \circ \phi' \circ \psi = [\ell_2^{e_2}]\phi.$$

4.2 Issue with soundness proofs for the De Feo–Jao–Plût scheme

A core component of the security proof of the De Feo–Jao–Plût identification scheme is the soundness proof. A proof of soundness was given by multiple previous works [DJP14, YAJ+17, GPS20] based on the CSSI problem in Definition 4. A sketch of this soundness proof is as follows:

Suppose $\mathcal{A}$ is an adversary that takes as input the public key and succeeds in the identification protocol (all t iterations) with noticeable probability ϵ. Given a challenge instance $(E_0, E_1, R_0, S_0, \phi(R_0), \phi(S_0))$ for the CSSI problem, we run $\mathcal{A}$ on the tuple $(E_1, \phi(R_0), \phi(S_0))$ as the public key. In the first round, $\mathcal{A}$ outputs commitments $(E_{i,2}, E_{i,3})$ for $1 \leq i \leq t$. We then send a challenge $b \in \{0,1\}^t$ to $\mathcal{A}$ and, with probability ϵ, $\mathcal{A}$ outputs a response that satisfies the verification algorithm. Now, we use the standard replay technique: Rewind $\mathcal{A}$ to the point where it had output its commitments and then respond with a different challenge $b' \in \{0,1\}^t$. With probability ϵ, $\mathcal{A}$ outputs a valid response. This gives exactly the 2-special soundness requirement of two valid transcripts with the same commitment but different challenges.

Now, choose some index i such that $b_i \neq b'_i$. We now restrict our focus to the components (E_2, E_3) for that index, and the two responses. It means $\mathcal{A}$ sent E_2, E_3 and can answer both challenges $b = 0$ and $b = 1$ successfully. Hence $\mathcal{A}$ has provided the maps ψ, ϕ', ψ' in the following diagram.

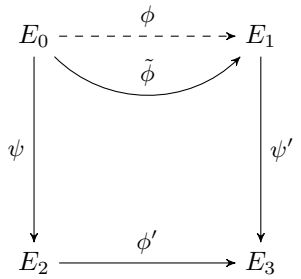

The argument proceeds as follows: We have an explicit description of an isogeny $\tilde{\phi} = \widehat{\psi'} \circ \phi' \circ \psi$ from E_0 to E_1. The degree of $\tilde{\phi}$ is $\ell_1^{e_1}\ell_2^{2e_2}$. One can determine $\ker(\tilde{\phi}) \cap E_0[\ell_1^{e_1}]$ by iteratively testing points in $E_0[\ell_1^j]$ for $j = 1, 2, \ldots$. Hence, one determines the kernel of ϕ, as desired.

However, the important issue with this argument which has so far gone unnoticed, is that it assumes $\ker(\phi) = \ker(\tilde{\phi}) \cap E_0[\ell_1^{e_1}]$. This assumption has no basis, and we will provide a simple counterexample to this argument in the following section. While we always recover an isogeny, it may not be ϕ at all—it is entirely possible the isogeny we recover does not even have codomain E_1 so this proof of 2-special soundness is not valid.

4.3 Counterexample to soundness

Fix a supersingular curve E_0 as above. Generate a random $\ell_2^{e_2}$-torsion point $K_\psi \in E_0(\mathbb{F}_{p^2})$ as $K_\psi = [a]P_0 + [b]Q_0$ for some integers $0 \leq a, b < \ell_2^{e_2}$ not both divisible by ℓ_2. Let $\psi : E_0 \to E_2$ have kernel generated by K_ψ. Then choose a random isogeny $\phi' : E_2 \to E_3$ of degree $\ell_1^{e_1}$ with kernel generated by $K_{\phi'}$. Then choose a random isogeny $\psi' : E_3 \to E_1$ of degree $\ell_2^{e_2}$. Choose points $P'_0, Q'_0 \in E_1(\mathbb{F}_{p^2})$ such that $\ker(\widehat{\psi'}) = \langle [a]P'_0 + [b]Q'_0 \rangle$. Then publish

$$(E_0, E_1, P_0, Q_0, P'_0, Q'_0)$$

as a public key. In other words, we have

$$E_0 \xrightarrow{\psi} E_2 \xrightarrow{\phi'} E_3 \xrightarrow{\psi'} E_1$$

Now there is no reason to believe that there exists an isogeny from E_0 to E_1 of degree $\ell_1^{e_1}$, yet we can respond to both challenge bits 0 and 1 in a single round of the identification scheme. Pulling back the kernel of ϕ' via ψ to E_0 will result in the kernel of an isogeny which, in general, will not have codomain E_1 (but instead

a random other curve). This is because ψ' is entirely unrelated to ψ in this case (they are not "parallel"), so we have no SIDH square.

The key observation is that a verifier could be fooled into accepting this public key by a prover who always uses the same curves (E_2, E_3) instead of randomly chosen ones. When $b = 0$ the prover responds with the pair (a, b) corresponding to the kernel of ψ and $\widehat{\psi'}$, and when $b = 1$ the prover responds with $K_{\phi'}$. The verifier will agree that all responses are correct and will accept the proof.

It is true that the verifier could test whether the commitments (E_2, E_3) are being re-used, but this has never been stated as a requirement in any of the protocol descriptions. To tweak the verification protocol we need to know how "random" the pairs (E_2, E_3) (or, more realistically, the pairs (a, b)) need to be. One may think that the original scheme seems to be secure despite the issue with the proof, as long as the commitment (E_2, E_3) is not reused every time. However, in experiments with small primes, it is entirely possible to construct instances[1] where even with multiple different commitments, a secret isogeny of the correct degree between E_0 and E_1 does not exist. We expect that this extrapolates to large primes too, although one could potentially argue that finding enough such instances is computationally infeasible.

It is also true that repeating (E_2, E_3) means the protocol is no longer zero-knowledge. We emphasize that soundness and zero-knowledge are independent security properties, which are proved separately (and affect different parties: one gives an assurance to the verifier and the other to the prover). The counterexample we have provided is a counterexample to the soundness proof. The fact that the counterexample is not consistent with the proof that the protocol is zero-knowledge is irrelevant.

Finally, one could consider basing security of the protocol on the general isogeny problem (Definition 3) because, even in our counterexample, an isogeny $E_0 \to E_1$ exists and can be extracted—it just doesn't have degree $\ell_1^{e_1}$. We find it interesting that none of the previous authors chose to do it that way. However, some applications may require using the identification/signature protocols to prove that an SIDH public key is well-formed, implying the secret isogeny has the correct degree. For such applications we need soundness to be rigorously proved.

The issue in the security proofs in the literature is not only that it is implicitly assumed that there is an isogeny of degree $\ell_1^{e_1}$ between E_0 and E_1. The key issue is that it is implicitly assumed that the pullback under ψ of $\ker(\phi')$ is the kernel of this isogeny. Our counterexample calls these assumptions into question, and shows that the proofs are incorrect as written.

To make this very clear, consider the soundness proof from De Feo, Jao, and Plût [DJP14]. The following diagram is written within the proof. It implicitly assumes that the horizontal isogeny ϕ' has kernel given by $\psi(S)$, so that the image curve is $E/\langle S, R\rangle$.

$$\begin{array}{ccc} E & & E/\langle S\rangle \\ \downarrow{\scriptstyle \psi} & & \downarrow{\scriptstyle \psi'} \\ E/\langle R\rangle & \xrightarrow[\phi']{} & E/\langle S, R\rangle \end{array}$$

This implicit assumption seems to have been repeated in all subsequent works, such as [YAJ+17] and [GPS20].

[1] Thank you to Lorenz Panny for demonstrating this.

5 First new SIDH proof

Let public parameters $pp = (p, \ell_1, \ell_2, e_1, e_2, E_0, P_0, Q_0)$ be such that $E_0(\mathbb{F}_{p^2})[\ell_2^{e_2}] = \langle P_0, Q_0 \rangle$. As before, suppose a user has a secret isogeny $\phi : E_0 \to E_1$ of degree $\ell_1^{e_1}$ with kernel $\ker(\phi) = \langle K_\phi \rangle$.

We propose a new sigma protocol to prove knowledge of this isogeny given the public key $(E_1, P_1 = \phi(P_0), Q_1 = \phi(Q_0))$. The protocol is presented in Figure 6. $\mathsf{IsogenyFromKernel}$ is a function taking a kernel point and outputting an isogeny and codomain curve with said kernel. $\mathsf{CanonicalBasis}_2$ is a deterministic function taking a curve and outputting a $\ell_2^{e_2}$-torsion basis on the given curve. $\mathsf{DualKernel}$ is a function taking an isogeny ψ and outputting a generator $K_{\widehat{\psi}}$ of the dual isogeny $\widehat{\psi}$. Figure 5 shows the commutative diagram of the sigma protocol.

Intuitively, the identification scheme follows Section 4.1, with a single bit challenge—if the challenge is 0, we reveal the vertical isogenies ψ, ψ', while if the challenge is 1, we reveal the horizontal ϕ'. The difference is the introduction of additional points on E_3 to the commitment, which force ψ, ψ' to be, in some sense "compatible" or "parallel". This restriction allows the proof of 2-special soundness to work.

We then repeat the identification scheme t times in parallel (where t is chosen based on the security parameter κ) and set com to be the concatenation of all individual $[\mathsf{com}_i]_{i\in[t]}$ for each iteration i, $\mathsf{chall} = [\mathsf{chall}_i]_{i\in[t]}$ and $\mathsf{resp} = [\mathsf{resp}_i]_{i\in[t]}$.

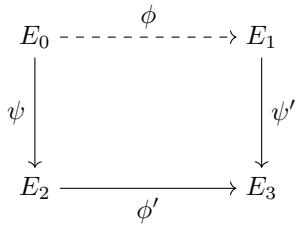

Figure 5: Commutative diagram of SIDH identification scheme

Remark 1. There are certainly improvements that can be made to improve efficiency and compress the size of signatures, but these are standard and we will not explore them here. For example, in practice the commitment information (E_3, P_3, Q_3) would be replaced with a triplet of x-coordinates, as in SIKE [ACC+17].

Theorem 4. *The sigma protocol in Figure 6 for relation*

$$\mathcal{R}_{\mathsf{weakSIDH}} = \{((E_1, P_1, Q_1), \phi) \mid \phi : E_0 \to E_1, \deg \phi = \ell_1^{e_1}\}$$

is correct, 2-special sound, and computationally zero-knowledge assuming the DMSIDH and DSSP problems are hard. Repeated with κ iterations, it is thus a Proof of Knowledge for $\mathcal{R}_{\mathsf{weakSIDH}}$ with knowledge error $2^{-\kappa}$.

Proof. We prove the three properties of Theorem 4 separately below.

Correctness: Following the protocol honestly will result in an accepting transcript. This is clear for the $\mathsf{chall} = 1$ case. For the $\mathsf{chall} = 0$ case, observe that

$$\phi'(K_{\widehat{\psi}}) = \phi'([c]P_2 + [d]Q_2) = [c]P_3 + [d]Q_3 = K_{\widehat{\psi'}},$$

thus $K_{\widehat{\psi'}}$ generates the kernel of $\widehat{\psi'}$.

round 1 (commitment)

Sample random $\ell_2^{e_2}$-isogeny kernel $\langle K_\psi \rangle \subset E_0$
Write $K_\psi = [a]P_0 + [b]Q_0 \in E_0$ for $a, b \in \mathbb{Z}/\ell_2^{e_2}\mathbb{Z}$
$K_{\psi'} \leftarrow \phi(K_\psi) = [a]\phi(P_0) + [b]\phi(Q_0) \in E_1$
$\psi, E_2 \leftarrow \mathsf{IsogenyFromKernel}(K_\psi)$
$P_2, Q_2 \leftarrow \mathsf{CanonicalBasis}_2(E_2)$
$K_{\phi'} \leftarrow \psi(K_\phi) \in E_2$
$\phi', E_3 \leftarrow \mathsf{IsogenyFromKernel}(K_{\phi'})$
$P_3, Q_3 \leftarrow \phi'(P_2), \phi'(Q_2) \in E_3$
Prover sends $\mathsf{com} \leftarrow (E_2, E_3, P_3, Q_3)$ to Verifier.

round 2 (challenge)

Verifier sends $\mathsf{chall} \leftarrow \{0, 1\}$ to Prover.

round 3 (response)

if $\mathsf{chall} = 1$ **then**
    $\mathsf{resp} \leftarrow K_{\phi'}$
else
    $K_{\widehat{\psi}} \leftarrow \mathsf{DualKernel}(\psi)$
    Write $K_{\widehat{\psi}} = [c]P_2 + [d]Q_2$ for $c, d \in \mathbb{Z}/\ell_2^{e_2}\mathbb{Z}$
    $\mathsf{resp} \leftarrow (c, d)$
Prover sends resp to Verifier.

Verification

$(E_2, E_3, P_3, Q_3) \leftarrow \mathsf{com}$
if $\mathsf{chall} = 1$ **then**
    $K_{\phi'} \leftarrow \mathsf{resp}$
    Check $K_{\phi'}$ has order $\ell_1^{e_1}$ and lies on E_2, otherwise output reject
    $P_2, Q_2 \leftarrow \mathsf{CanonicalBasis}_2(E_2)$
    $\phi', E_3' \leftarrow \mathsf{IsogenyFromKernel}(K_{\phi'})$
    Verify $E_3 = E_3'$ and $P_3, Q_3 = \phi'(P_2), \phi'(Q_2)$, otherwise output reject
else
    $(c, d) \leftarrow \mathsf{resp}$
    $P_2, Q_2 \leftarrow \mathsf{CanonicalBasis}_2(E_2)$
    $K_{\widehat{\psi}} \leftarrow [c]P_2 + [d]Q_2$
    $K_{\widehat{\psi'}} \leftarrow [c]P_3 + [d]Q_3$
    Check $K_{\widehat{\psi}}$, $K_{\widehat{\psi'}}$ have order $\ell_2^{e_2}$, otherwise output reject
    $\widehat{\psi}, E_0' \leftarrow \mathsf{IsogenyFromKernel}(K_{\widehat{\psi}})$
    $\widehat{\psi'}, E_1' \leftarrow \mathsf{IsogenyFromKernel}(K_{\widehat{\psi'}})$
    Check $E_0 = E_0'$ and $E_1 = E_1'$, otherwise output reject
Output accept

Figure 6: One iteration of the sigma protocol for our new SIDH identification scheme. The public parameters are $pp = (p, \ell_1, \ell_2, e_1, e_2, E_0, P_0, Q_0)$. The public key is (E_1, P_1, Q_1), and the corresponding secret isogeny is ϕ.

2-special soundness: Without loss of generality, suppose we obtain two transcripts $(\mathsf{com}, 0, \mathsf{resp})$, $(\mathsf{com}, 1, \mathsf{resp}')$. Then recover $(c, d) \leftarrow \mathsf{resp}$ and $K_{\phi'} \leftarrow \mathsf{resp}'$, and let ϕ' be an isogeny whose kernel is generated by $K_{\phi'}$. Applying Lemma 2, with $(\phi_A, \phi_B, \phi_{AB}) = (\phi', \widehat{\psi}, \widehat{\psi'})$, we obtain an isogeny $\chi : E_0 \to E_1$ of degree $\ell_1^{e_1}$.

The conditions of the lemma on the kernels of $\widehat{\psi}$ and $\widehat{\psi'}$ are satisfied because $\phi'(K_{\widehat{\psi}}) = K_{\widehat{\psi'}}$, as above. This shows the protocol is 2-special sound, and that it is a Proof of Knowledge of an isogeny corresponding to the given public key curve. Because this protocol does not guarantee correctness of the points P_1, Q_1 in the public key (as briefly discussed in Section 5.1), this is only a proof for the weakSIDH relation. In the next section, we will modify this protocol further to also include these torsion points in the relation.

Zero-knowledge: Proof of ZK follows as in [DJP14]. Let V^* be a cheating verifier, which shall be used as a black box by the simulator Sim. We show that Sim can generate a valid transcript for t iterations of the protocol. At each step, Sim makes a guess what the next challenge bit chall will be, and then proceeds as follows.

- If chall $= 0$, Sim simulates as per the honest protocol by choosing a random kernel $\langle K_\psi \rangle$ on E_0 of order $\ell_2^{e_2}$, writing $K_\psi = [a]P_0 + [b]Q_0$ for $a, b \in \mathbb{Z}/\ell_2^{e_2}\mathbb{Z}$, and setting $K_{\psi'} = [a]P_1 + [b]Q_1$ on E_1. Sim computes the two vertical isogenies $\psi : E_0 \to E_2, \psi' : E_1 \to E_3$ from these kernel generators respectively. The simulator then computes the corresponding dual isogenies and the canonical basis $P_2, Q_2 \leftarrow \mathsf{CanonicalBasis}_2(E_2)$. Let $K_{\widehat{\psi}}$ and $K_{\widehat{\psi'}}$ be generators of the kernels of $\widehat{\psi}$ and $\widehat{\psi'}$ respectively. The simulator writes $K_{\widehat{\psi}}$ in terms of the canonically-generated basis on E_2 as $[c]P_2 + [d]Q_2$, then chooses a torsion basis on E_3 as $P_3, Q_3 \in E_3$ in such a way that these points P_3, Q_3 are indistinguishable from points chosen in an honest protocol transcript:

  1. Obtain a point $S \in E_2$ and its image $S' = \phi'(S_2)$ via the algorithm in Figure 3 despite ϕ' being unknown (c.f. Lemma 2).
  2. Choose any $T \in E_2$ of order $\ell_2^{e_2}$ such that $E_2[\ell_2^{e_2}] = \langle S, T \rangle$.
  3. Choose a point $T' \in E_3$ such that $E_3[\ell_2^{e_2}] = \langle S', T' \rangle$, and such that $e_{\ell_2^{e_2}}(S, T)^{\ell_1^{e_1}} = e_{\ell_2^{e_2}}(S', T')$.
  4. Solve discrete logarithms of P_2, Q_2 with respect to S, T on E_2 to obtain a change-of-basis matrix, and apply the same change of basis to S', T' on E_3 to obtain points P_3, Q_3.

  Note that the above operations are efficient due to the ease of computing discrete logarithms when the group order is very smooth [Tes99].

- If chall $= 1$, the simulator chooses a random supersingular elliptic curve[2] E_2 and a random point $K_{\phi'} \in E_2$ of order $\ell_1^{e_1}$. Sim then computes an isogeny $\phi' : E_2 \to E_3$ with kernel $K_{\phi'}$. Finally, the simulator generates a canonical basis $P_2, Q_2 \leftarrow \mathsf{CanonicalBasis}_2(E_2)$, computes $P_3, Q_3 \leftarrow \phi'(P_2), \phi'(Q_2)$, and sets the commitment to (E_2, E_3, P_3, Q_3) and the response to $K_{\phi'}$.

After providing com to V^*, if the challenge V^* outputs is not the same as Sim's guess, Sim simply discards that iteration and runs again. Sim stops whenever V^* rejects or after t successful rounds. Suppose the probability of V^* not choosing the same bit as Sim's guess is noticeably different from $1/2$. Then V^* can be used as a distinguisher for the DSSP problem (in fact, an even harder problem than the DSSP where, instead of the isogeny ϕ', only its action on $E_2[\ell_2^{e_2}]$ is given). We show this below, in the chall $= 1$ case of this proof. So the probability Sim guesses correctly each round is exponentially close to $1/2$ if the DSSP problem is hard. Thus, Sim will run in polynomial-time.

To prove indistinguishability of simulated transcripts from true interactions of a prover P with V^*, it is enough to show that one round of the sigma protocol is indistinguishable (by the hybrid technique of Goldreich et al. [GMW91]).

When chall $= 0$, the choice of ψ and ψ' is done exactly as in the honest protocol, so the curves E_2, E_3 in the commitment are perfectly indistinguishable from those in honest transcripts. We show that the points P_3, Q_3 are also indistinguishable, assuming the DMSIDH problem is hard. Suppose $\mathcal{B}_0$ is a PPT adversary which can distinguish between the simulation and the real transcripts for chall $= 0$ with advantage Adv_0. Let $((E_0,$

[2] One way to do so is to take a random ℓ_2-isogeny walk from E_0. To ensure a distribution close to uniform, we take a walk of length $\gtrsim \log(p) \approx 2e_2$. However a walk of length e_2 is sufficient to get a variant of DSSP that is also believed to be hard.

$P_0, Q_0), (E_1, \phi(P_0), \phi(Q_0)), \psi, \psi', S, T, \phi'(S), T')$ be a challenge instance of the DMSIDH problem. Denote by E_2 the codomain of ψ, and E_3 the codomain of ψ'. Set $P_2, Q_2 \leftarrow \mathsf{CanonicalBasis}_2(E_2)$, and proceed as in Step 4 of the simulation to obtain points P_3, Q_3 from S', T' using a change of basis matrix $A = (a_i)$ derived from (P_2, Q_2) and (S, T):

$$\begin{pmatrix} P_2 \\ Q_2 \end{pmatrix} = \begin{pmatrix} a_0 & a_1 \\ a_2 & a_3 \end{pmatrix} \cdot \begin{pmatrix} S \\ T \end{pmatrix} \tag{1}$$

Write the kernel of $\widehat{\psi}$ as $\ker(\widehat{\psi}) = [c]P_2 + [d]Q_2$ for scalars c, d. Finally, give transcript $\mathsf{com} = (E_2, E_3, P_3, Q_3), \mathsf{chall} = 0, \mathsf{resp} = (c, d)$ to $\mathcal{B}_0$.

If $T' = \phi'(T)$ in the challenge instance of the DMSIDH problem (i.e. from distribution $\mathcal{D}_0$), then we necessarily have that $P_3 = [a_0]\phi'(S) + [a_1]\phi'(T) = \phi'(P_2)$, and similarly $Q_3 = \phi'(Q_2)$. Hence, the distribution of transcripts will be identical to the honest protocol. On the other hand, the transcript simulator selects a random T' such that $E_3[\ell_2^{e_2}] = \langle S', T' \rangle$ and $e_{\ell_2^{e_2}}(S, T)^{\ell_1^{e_1}} = e_{\ell_2^{e_2}}(S', T')$. Let $T' = [q]\phi'(T) + [r]\phi'(S) = [q]\phi'(T) + [r]S'$. The pairing condition gives $e_{\ell_2^{e_2}}(S', [q]\phi'(T) + [r]S') = e_{\ell_2^{e_2}}(S', \phi'(T))^q$ implying $q = 1$. Hence $T' = \phi'(T + [r]S)$. Then, because the transcript simulator behaves identically to the reduction in computing P_3, Q_3 (via applying the same change of basis matrix to S', T'), the transcript distribution in the reduction will be identical to the transcripts generated by the simulator. Therefore, the response from $\mathcal{B}_0$ will solve the DMSIDH problem with advantage Adv_0.

Remark 2. If there was an efficient solution to the computational version of the DMSIDH problem—that is, the problem of finding the correct image of T under the secret ϕ'—then we could obviously simulate perfectly. Moreover, if there did exist an efficient distinguisher for the DMSIDH problem, then integrating it into the verification step of the protocol in Figure 6 would be enough to prove the strong relation that we will define in Section 6. A surprising situation would only materialize if there were a gap between DMSIDH and its computational analogue, leading to an efficient, but disturbingly not zero-knowledge, protocol for both the weak and the strong relation. Our intuition tells us that such a gap should not exist, but a proof seems to be out of reach.

When $\mathsf{chall} = 1$, we consider the distribution of (E_2, E_3, ϕ'). While this distribution is not correct a priori, the DSSP computational assumption in Definition 6 implies it is computationally hard to distinguish the simulation from the real game (as in the proof in [GPS20]). Because the action of ϕ' on canonical basis $P_2, Q_2 \in E_2$ can be computed by any party who knows ϕ', the distribution of (E_2, E_3, P_3, Q_3) must also be indistinguishable between simulation and real transcripts.

Suppose $\mathcal{B}_1$ is a PPT adversary which can distinguish between the simulation and the real transcripts for $\mathsf{chall} = 1$ with advantage Adv_1. Given an instance of the DSSP problem, (E_2, E_3, ϕ'), compute $P_2, Q_2 \leftarrow \mathsf{CanonicalBasis}_2(E_2)$. Then let $P_3 = \phi'(P_2)$ and $Q_3 = \phi'(Q_2)$, and set $\mathsf{com} = (E_2, E_3, P_3, Q_3), \mathsf{chall} = 1, \mathsf{resp} = (\ker(\phi'))$. $\mathcal{B}_1$, given $(\mathsf{com}, \mathsf{chall}, \mathsf{resp})$, will then solve the DSSP with the same advantage Adv_1. It is for this same reason that a cheating verifier V^* is unable to distinguish based on com alone whether the simulator is attempting a $\mathsf{chall} = 0$ or $\mathsf{chall} = 1$ simulation with non-negligible advantage, if it cannot solve the DSSP problem with non-negligible advantage.

Hence, the scheme is computationally zero-knowledge assuming the DSSP and DMSIDH problems are hard. □ □

Remark 3. We note that the points P_1, Q_1 are not actually used in the verification algorithm, so could be omitted entirely in practice if desired. After observing just two iterations of the sigma protocol on average the verifier would be able to reconstruct (P_1, Q_1).

5.1 Why this protocol does not prove correctness of the points (P_1, Q_1)

We briefly explain why the protocol in this section does not convince a verifier that $(P_1, Q_1) = (\phi(P_0), \phi(Q_0))$. The first observation is that Figure 6 does not actually use P_1 or Q_1 anywhere, so of course, nothing is proved. But one could tweak the protocol in the $\mathsf{chall} = 0$ case to use the isogenies $\widehat{\psi} : E_2 \to E_0$ and $\widehat{\psi'} : E_3 \to E_1$ to test the points. For example, using the duals of these isogenies, one could compute integers (a, b) such that $\ker(\psi) = \langle [a]P_0 + [b]Q_0 \rangle$ and then test whether or not $\ker(\psi') = \langle [a]P_1 + [b]Q_1 \rangle$.

The problem for the verifier is that this is not enough to deduce that $(P_1, Q_1) = (\phi(P_0), \phi(Q_0))$. For example, a dishonest prover who wants to perform an attack might set $(P_1, Q_1) = (\phi(P_0), \phi(Q_0) + T)$ where T is a point of order ℓ_2. If the prover always uses integers b that are multiples of ℓ_2 then this cheating will not be detected by the verifier. Hence, the protocol needs to be changed so that the verifier can tell that the kernels of the isogenies $\widehat{\psi}$ are sufficiently independent across the executions of the protocol. This is the fundamental problem that we solve in the next section.

6 Correctness of the points in an SIDH public key

We have shown in Section 5 that successful completion of the new sigma protocol indeed proves knowledge of a degree $\ell_1^{e_1}$ isogeny from E_0 to E_1 (as per the relation $\mathcal{R}_{\mathsf{weakSIDH}}$ in Theorem 4). However, an SIDH public key (E_1, P_1, Q_1) also consists of the two torsion points, and these points are the cause of issues such as the adaptive attack [GPST16], as discussed in Section 3. In this section, we show that the choice of points P_1, Q_1 by a malicious prover is severely restricted if they must keep them consistent with "random enough" values of a, b (i.e., random choices of ψ)—preventing adaptive attacks entirely. This gives the following stronger SIDH relation:

$$\mathcal{R}_{\mathsf{SIDH}} = \left\{ ((E_1, P_1, Q_1), \phi) \;\middle|\; \begin{array}{c} \phi : E_0 \to E_1,\ \deg \phi = \ell_1^{e_1}, \\ P_1 = \phi(P_0),\ Q_1 = \phi(Q_0) \end{array} \right\}$$

Figure 7 shows the modified protocol proving this strong relation.

Let us reconsider the protocol in Figure 6 for a moment. We have that $\ker(\widehat{\psi'}) = \phi'(\ker(\widehat{\psi}))$ by the 2-special soundness of Theorem 4. Applying the algorithm in Figure 3 to $(E_2, P_2, Q_2, E_3, P_3, Q_3, E_0, E_1, \widehat{\psi}, \widehat{\psi'})$ gives us a pair $(R_0, R_1 = \chi(R_0))$ for $\chi : E_0 \to E_1$, where $\ker(\chi) = \widehat{\psi}(\ker(\phi'))$. Note that ϕ in the algorithm and Lemma 2 corresponds to ϕ' here because we have "flipped the SIDH square upside down." Because the degrees of ϕ' and $\widehat{\psi}$ are coprime, we can translate this to $\psi(\ker(\chi)) = \ker(\phi')$. Note that R_0 and R_1 will be scalar multiples (by the same scalar) of the K_ψ and $K_{\psi'}$ used by the prover in the commitment round of the protocol.

Consequently, two (honest) answers to $\mathsf{chall} = 0$ reveal two pairs of points $R_{1,0}, R_{1,1} = \phi(R_{0,0}), \phi(R_{0,1})$. If these are independent, they fix the action of ϕ on the whole $\ell_2^{e_2}$ torsion (as a basis for the $\ell_2^{e_2}$ torsion subgroups on both curves). The easiest way to enforce two such honest answers is to "double" the protocol. Thus, in each round of our new sigma protocol, we shall commit to two SIDH squares rather than just one, and require that the kernel generators of ψ in these two squares are independent from each other. We add this independence as an extra check during verification. We also require an assurance that both squares use consistent isogenies ϕ'. For this purpose we use a random $\ell_1^{e_1}$-torsion basis (U, V) on E_0 and compute the image of this basis on both curves $E_{2,i}$—if both ϕ'_i are the images of ϕ under the vertical isogenies ψ_i, then both should be representable in terms of $(\psi_i(U), \psi_i(V))$ using the same coefficients. These extra checks achieve a 2-special sound protocol for the stronger SIDH relation above. We stress that (U, V) are not made public in the commitment. In the following protocol, $\mathsf{RandomBasis}_1$ is a function taking a curve and outputting a random pair of points U, V which generate the $\ell_1^{e_1}$-torsion subgroup on the given curve. The function $\mathsf{RandomBasis}_1$ is called many times on the same curve E_0 during t rounds of the protocol and it is important that the outputs are independent and not known to the verifier in the $\mathsf{chall} = 1$ case.

round 1 (commitment)

Run **commitment** from Figure 6, giving commitment $\mathsf{com}_0 = (E_{2,0}, E_{3,0}, P_{3,0}, Q_{3,0})$. Let a_0, b_0 be the coefficients used in Line 2 and ψ_0 be the isogeny from Line 4 (of Figure 6) of this execution.

Run **commitment** from Figure 6 again, subject to one extra condition:

- If a_1, b_1 are the coefficients used in Line 2 (of Figure 6) of this execution, then require $a_0 b_1 - a_1 b_0$ invertible modulo $\ell_2^{e_2}$. Otherwise repeat Line 1 (of Figure 6).

Let $\mathsf{com}_1 = (E_{2,1}, E_{3,1}, P_{3,1}, Q_{3,1})$ be the commitment returned by this execution, and ψ_1 be the isogeny from Line 4.

$U, V \leftarrow \mathsf{RandomBasis}_1(E_0)$

for $i \in \{0, 1\}$ **do**

  Let $U_i' = \psi_i(U)$ and $V_i' = \psi_i(V)$

Output commitment $(\mathsf{com}_0, U_0', V_0', \mathsf{com}_1, U_1', V_1')$.

round 3 (response)

if $\mathsf{chall} = 1$ **then**

  Write $K_\phi = [e]U + [f]V$ for $e, f \in \mathbb{Z}/\ell_1^{e_1}\mathbb{Z}$

  Output $\mathsf{resp} \leftarrow (e, f)$

else

  for $i \in \{0, 1\}$ **do**

    $K_{\widehat{\psi}} \leftarrow \mathsf{DualKernel}(\psi)$

    Write $K_{\widehat{\psi}} = [c_i]P_2 + [d_i]Q_2$ for $c_i, d_i \in \mathbb{Z}/\ell_2^{e_2}\mathbb{Z}$

    $\mathsf{resp}_i \leftarrow (c_i, d_i)$

  Output $\mathsf{resp} \leftarrow (\mathsf{resp}_0, \mathsf{resp}_1)$.

Verification

if $\mathsf{chall} = 1$ **then**

  $(e, f) \leftarrow \mathsf{resp}$

  for $i \in \{0, 1\}$ **do**

    $(E_2, E_3, P_3, Q_3, U_i', V_i') \leftarrow \mathsf{com}_i$

    Recover $K_{\phi',i} = [e]U_i' + [f]V_i'$

    Verify $(\mathsf{com}_i, \mathsf{chall}, K_{\phi',i})$ as in Figure 6 **verification**

    If verification fails, output reject.

else

  for $i \in \{0, 1\}$ **do**

    $(E_2, E_3, P_3, Q_3, U_i', V_i') \leftarrow \mathsf{com}_i$

    $(c, d) \leftarrow \mathsf{resp}_i$

    $P_2, Q_2 \leftarrow \mathsf{CanonicalBasis}_2(E_2)$

    $K_{\widehat{\psi}} \leftarrow [c]P_2 + [d]Q_2$

    $K_{\widehat{\psi'}} \leftarrow [c]P_3 + [d]Q_3$

    Check $K_{\widehat{\psi}}$, $K_{\widehat{\psi'}}$ have order $\ell_2^{e_2}$, otherwise output reject

    $\widehat{\psi}_i, E_0' \leftarrow \mathsf{IsogenyFromKernel}(K_{\widehat{\psi}})$

    $\widehat{\psi}_i', E_1' \leftarrow \mathsf{IsogenyFromKernel}(K_{\widehat{\psi'}})$

    Check $E_0 = E_0'$ and $E_1 = E_1'$, otherwise output reject

    Choose (c', d') such that $c'd - d'c$ is invertible modulo $\ell_2^{e_2}$

    $R_{0,i} \leftarrow \widehat{\psi}_i([c']P_2 + [d']Q_2)$

    $R_{1,i} \leftarrow \widehat{\psi}_i'([c']P_3 + [d']Q_3)$

    Check there exist $a_i', b_i' \in \mathbb{Z}/\ell_2^{e_2}\mathbb{Z}$ such that, simultaneously,

      i. $R_{0,i} = [a_i']P_0 + [b_i']Q_0$,

      ii. $R_{1,i} = [a_i']P_1 + [b_i']Q_1$,

    otherwise output reject

  Check $\widehat{\psi}_0(U_0') = \widehat{\psi}_1(U_1')$ and $\widehat{\psi}_0(V_0') = \widehat{\psi}_1(V_1')$, otherwise output reject.

  Check that $a_0'b_1' - a_1'b_0'$ is invertible modulo $\ell_2^{e_2}$, otherwise output reject.

Output accept if all the above conditions hold.

Figure 7: Modification of the Sigma protocol in Figure 6 to prove the stronger relation $\mathcal{R}_{\mathsf{SIDH}}$. Round 2, and lines in gray, are unchanged from Figure 6.

Theorem 5. *For a fixed security parameter κ and SIDH public key (E, P, Q), a proof consisting of κ iterations of the sigma protocol in Figure 7 is a computationally zero-knowledge Proof of Knowledge for $\mathcal{R}_{\mathsf{SIDH}}$ with knowledge error $2^{-\kappa}$, assuming the DMSIDH and Double-DSSP problems are hard.*

Proof. Again we prove correctness, soundness, and zero-knowledge individually.

Correctness: As mentioned above, the point $R_{0,i}$ will always be an invertible scalar multiple of the point K_ψ used by the prover in the commitment round (in the i-th SIDH square) of the protocol because both K_ψ and $R_{0,i}$ are generators of the kernel of ψ in the i-th SIDH square. This implies the pair (a'_i, b'_i) is an invertible scalar multiple of (a_i, b_i). Hence, because the honest prover will use commitments such that $a_0b_1 - a_1b_0$ is invertible, then a'_i, b'_i necessarily exist such that $a'_0b'_1 - a'_1b'_0$ is invertible in line 22 of verification. Also note that because $K_{\phi',i} = [e]U'_i + [f]V'_i = [e]\psi_i(U) + [f]\psi_i(V)$ for both $i \in \{0, 1\}$, and U, V have order coprime to the degree of ψ_i, the checks involving U'_i, V'_i, e, and f will also succeed. Correctness of the rest of the protocol can also be verified in a straightforward way.

Zero-knowledge: Let V^* be a cheating verifier. Sim will generate a valid transcript for t iterations of the protocol as follows. At each step, Sim will make a guess on what the next challenge bit chall will be, and proceeds appropriately:

- If $\mathsf{chall} = 0$, Sim will behave as in the proof of Theorem 4 to generate the first SIDH square arbitrarily. The simulator will then generate a second SIDH square in almost the same way, but ensuring that the second ψ chosen uses kernel coefficients independent to those used in the first square (just like the honest prover would do in the commitment round of Figure 7). Sim will also randomly generate a basis (U, V) of the $\ell_1^{e_1}$ torsion on E_0 and compute the images $U'_i, V'_i = \psi_i(U), \psi_i(V)$ exactly as in Figure 7. The commitment and response will be formed exactly as in the honest protocol.
- When $\mathsf{chall} = 1$, the behaviour of Sim is similar to the $\mathsf{chall} = 1$ simulation in the proof of Theorem 4, but repeated twice. First, Sim will choose two random curves $E_{2,i}$, for $i \in \{0, 1\}$. Sim will then choose a random point $K_{\phi',0} \in E_{2,0}$ of order $\ell_1^{e_1}$, and a random basis $\langle U'_0, V'_0\rangle = E_{2,i}[\ell_1^{e_1}]$, and write $K_{\phi',0} = [e]U'_0 + [f]V'_0$ for integers e, f. Next, Sim will randomly generate a basis (U'_1, V'_1) of the $\ell_1^{e_1}$-torsion subgroup on $E_{2,1}$ such that $e_{\ell_1^{e_1}}(U'_0, V'_0) = e_{\ell_1^{e_1}}(U'_1, V'_1)$, and let $K_{\phi',1} = [e]U'_1 + [f]V'_1$. Let ϕ'_0, ϕ'_1 be isogenies with respective kernels $K_{\phi',0}, K_{\phi',1}$, and let $E_{3,i}$ be the codomain of ϕ'_i. Finally, the simulator generates canonical bases $P_{2,i}, Q_{2,i} \leftarrow \mathsf{CanonicalBasis}_2(E_{2,i})$, computes $P_{3,i}, Q_{3,i} \leftarrow \phi'_i(P_{2,i}), \phi'_i(Q_{2,i})$, and sets

$$\begin{aligned}\mathsf{com} &\leftarrow \left((E_{2,i}, E_{3,i}, P_{3,i}, Q_{3,i}, U'_i, V'_i)_{i\in\{0,1\}}\right),\\ \mathsf{resp} &\leftarrow (e, f).\end{aligned}$$

After providing com to V^*, if the challenge which V^* outputs is not the same as Sim's guess, Sim simply discards that iteration and runs again. Sim stops whenever V^* rejects or after t successful rounds. Suppose the probability of V^* not choosing the same bit as Sim's guess is noticeably different from $1/2$. Then V^* can be used as a distinguisher for (a harder variant of) the Double-DSSP problem, as we again show in the $\mathsf{chall} = 1$ case of this proof. So the probability Sim guesses correctly each round is exponentially close to $1/2$ if the DSSP problem is hard. Thus Sim will run in polynomial-time.

Correctness of the simulator: We first show that the simulator will successfully generate valid transcripts with the additional R_0, R_1 check in the protocol. Suppose the verifier arbitrarily chooses c', d' such that $c'd - d'c$ is invertible modulo $\ell_2^{e_2}$, where c, d were used in the response of either square $i \in \{0, 1\}$. We have that

$$R_2 = \begin{pmatrix} c' & d' \end{pmatrix} \begin{pmatrix} P_2 \\ Q_2 \end{pmatrix} = \begin{pmatrix} c' & d' \end{pmatrix} A \begin{pmatrix} S \\ T \end{pmatrix} \tag{2}$$

where the matrix A is the same change-of-basis matrix as in Equation 1. So,

$$R_0 = \widehat{\psi}(R_2) = \begin{pmatrix} c' & d' \end{pmatrix} A \begin{pmatrix} \widehat{\psi}(S) \\ \widehat{\psi}(T) \end{pmatrix} = \begin{pmatrix} c' & d' \end{pmatrix} A \begin{pmatrix} \mathcal{O}_{E_0} \\ \widehat{\psi}(T) \end{pmatrix} \tag{3}$$

because S is in the kernel of $\widehat{\psi}$. Similarly,

$$R_3 = \begin{pmatrix} c' & d' \end{pmatrix} \begin{pmatrix} P_3 \\ Q_3 \end{pmatrix} = \begin{pmatrix} c' & d' \end{pmatrix} A \begin{pmatrix} \phi'(S) \\ \phi'(T + [r]S) \end{pmatrix} \tag{4}$$

from the simulator in the proof of Theorem 4. In the case of an honest prover (or a $\mathcal{D}_0$ DMSIDH instance where $T' = \phi'(T)$), r here would be zero. Then,

$$\begin{aligned} R_1 = \widehat{\psi'}(R_3) &= \begin{pmatrix} c' & d' \end{pmatrix} A \begin{pmatrix} \widehat{\psi'}(\phi'(S)) \\ \widehat{\psi'}(\phi'(T + [r]S)) \end{pmatrix} \\ &= \begin{pmatrix} c' & d' \end{pmatrix} A \begin{pmatrix} \mathcal{O}_{E_0} \\ \widehat{\psi'}(\phi'(T)) \end{pmatrix} \end{aligned} \tag{5}$$

because again, $\phi'(S)$ is in the kernel of $\widehat{\psi'}$. Hence, we must have that $R_1 = \phi(R_0)$ in either of the two DMSIDH instance distributions (and hence also in the two Double-DMSIDH distributions). This implies that the coefficients a'_i, b'_i in each SIDH square of the protocol exist and can be used to satisfy the verification algorithm regardless of whether a simulator or honest prover has generated the transcript.

Indistinguishability of the simulator: Suppose $\mathcal{B}_0$ is a PPT adversary which can distinguish between the simulation and the real transcripts for $\mathsf{chall} = 0$ with advantage Adv_0. We show that $\mathcal{B}_0$ can then also solve the Double-DMSIDH problem with the same advantage Adv_0. Let $(\psi_i, \psi'_i, S_i, T_i, \phi'(S_i), T'_i)_{i\in\{0,1\}}$ be an instance of the Double-DMSIDH problem. For both $i \in \{0, 1\}$, we proceed as in the proof of Theorem 4 to create a transcript $\mathsf{com} = (E_{2,i}, E_{3,i}, P_{3,i}, Q_{3,i})_{i\in\{0,1\}}$, $\mathsf{chall} = 0$, $\mathsf{resp} = (c_i, d_i)_{i\in\{0,1\}}$. We also compute the images $U'_i, V'_i = \psi_i(U), \psi_i(V)$ of the random basis (U, V), exactly as above. We then provide this transcript to $\mathcal{B}_0$. This will produce an identical distribution of transcripts as those produced by the simulator because the steps are the same. Therefore, the response from $\mathcal{B}_0$ will solve the Double-DMSIDH problem with advantage Adv_0.

Now coming to the $\mathsf{chall} = 1$ case, we similarly suppose $\mathcal{B}_1$ is a PPT adversary which can distinguish between the simulation and the real transcripts for $\mathsf{chall} = 1$ with advantage Adv_1. Let $(E_{2,i}, E_{3,i}, \phi'_i, U'_i, V'_i), i \in \{0, 1\}$ be an instance of the Double-DSSP problem. As in the proof of Theorem 4, compute $P_{2,i}, Q_{2,i} \leftarrow \mathsf{CanonicalBasis}_2(E_{2,i})$, and let $P_{3,i}, Q_{3,i} = \phi'_i(P_{2,i}), \phi'_i(Q_{2,i})$. Finally, write $\ker(\phi'_i) = [e]U'_i + [f]V'_i$ and set $\mathsf{com} = (E_{2,i}, E_{3,i}, P_{3,i}, Q_{3,i}, U'_i, V'_i)_{i\in\{0,1\}}$, $\mathsf{chall} = 1$, and $\mathsf{resp} = (e, f)$, and give $(\mathsf{com}, \mathsf{chall}, \mathsf{resp})$ to $\mathcal{B}_1$. If $\mathcal{B}_1$ outputs 1, then we respond to the Double-DSSP instance with 1, and win with advantage Adv_1. For the same reason that this $\mathsf{chall} = 1$ case is hard to distinguish, a cheating verifier V^* given only com is also unable to distinguish whether Sim is attempting a $\mathsf{chall} = 0$ or $\mathsf{chall} = 1$ simulation with non-negligible probability if the Double-DSSP problem is hard.

Hence, assuming the Double-DSSP and DMSIDH problems are hard, transcripts generated by the simulator are indistinguishable from honest transcripts generated as per the protocol in Figure 7.

2-special soundness: Suppose we obtain two accepting transcripts $(\mathsf{com}, 0, \mathsf{resp})$ and $(\mathsf{com}, 1, \mathsf{resp}')$. The secret isogeny corresponding to the public key $X = (E_1, P_1, Q_1)$ can be recovered as follows, hence $\mathsf{Extract}$ can extract a valid witness W for the statement X such that $(X, W) \in \mathcal{R}_{\mathsf{SIDH}}$. From such a pair of commitments and responses, for each of the two SIDH squares committed to in Figure 7, we can recover $\phi_i : E_0 \to E_1$ of

degree $\ell_1^{e_1}$ by the proof of Theorem 4. Now,

$$\begin{aligned}\ker(\phi_0) &= \widehat{\psi}_0(\ker(\phi_0')) \\ &= \langle \widehat{\psi}_0([e]U_0 + [f]V_0)\rangle \\ &= \langle \widehat{\psi}_1([e]U_1 + [f]V_1)\rangle \\ &= \widehat{\psi}_1(\ker(\phi_1')) \quad = \ker(\phi_1).\end{aligned}$$

Therefore, we recover the same isogeny $\phi_0 = \phi_1 = \phi$ from both squares. For each of these two squares $i \in \{0, 1\}$, the verifier will choose an $R_{0,i}$ and also learn its image $R_{1,i}$ under ϕ. This follows from Lemma 2, with $S := [c']P_2 + [d']Q_2$.

Now, because the two $R_{0,i} = [a_i']P_0 + [b_i']Q_0$ are independent, $\langle R_{0,0}, R_{0,1}\rangle$ forms another basis for $\langle P_0, Q_0\rangle = E_0[\ell_2^{e_2}]$, with change-of-basis matrix

$$B = \begin{pmatrix} a_0' & b_0' \\ a_1' & b_1' \end{pmatrix}.$$

We can then see that

$$\begin{aligned}\begin{pmatrix} R_{0,0} \\ R_{0,1} \end{pmatrix} &= B \begin{pmatrix} P_0 \\ Q_0 \end{pmatrix} \\ \begin{pmatrix} \phi(R_{0,0}) \\ \phi(R_{0,1}) \end{pmatrix} = \begin{pmatrix} R_{1,0} \\ R_{1,1} \end{pmatrix} &= B \begin{pmatrix} \phi(P_0) \\ \phi(Q_0) \end{pmatrix} \\ \begin{pmatrix} R_{1,0} \\ R_{1,1} \end{pmatrix} &= B \begin{pmatrix} P_1 \\ Q_1 \end{pmatrix},\end{aligned}$$

therefore

$$B \begin{pmatrix} \phi(P_0) \\ \phi(Q_0) \end{pmatrix} = B \begin{pmatrix} P_1 \\ Q_1 \end{pmatrix},$$

and since B is invertible, we must have that $P_1 = \phi(P_0)$ and $Q_1 = \phi(Q_0)$, as required. □

□

Note that the protocol in Figure 7 essentially runs the previous protocol (in Figure 6) twice, hence the transcripts produced by this Proof of Knowledge for $\mathcal{R}_{\mathsf{SIDH}}$ will be twice the size.

Remark 4. Ghantous et al. [GPV21] discuss issues with extraction of a witness in two different scenarios. Their first scenario ("single collision") involves two distinct isogenies $\phi' : E_2 \to E_3$ in the SIDH square of the identification scheme. Neither of our new identification schemes are impacted by such collisions because the provision of points $P_3, Q_3 \in E_3$ uniquely determines the isogeny ϕ', as shown by Martindale and Panny [MP19]. Their second scenario ("double collision") involves two distinct (non-equivalent) isogenies $\phi, \tilde{\phi} : E_0 \to E_1$, both of degree $\ell_1^{e_1}$ and a point $R \in E_0$ such that

$$E_1/\langle \phi(R)\rangle \cong E_1/\langle \tilde{\phi}(R)\rangle.$$

Our second protocol, for the relation $\mathcal{R}_{\mathsf{SIDH}}$, ensures that the witness extracted is a valid witness for the public key used (including the torsion points). Hence, this second collision scenario does not have any impact on the soundness of our protocol either.

7 SIDH signatures and Non-Interactive Proof of Knowledge

We conclude with some brief, standard remarks about the use of the new protocol proposed above.

It is standard to construct a non-interactive signature scheme from an interactive protocol using the Fiat-Shamir transformation (secure in the (quantum) random oracle model [LZ19]). This works by making the challenge chall for the t rounds of the ID scheme a random-oracle output from input the commitment com and a message M. That is, for message M,

$$V_1^{\mathcal{O}}(\mathsf{com}) = \mathcal{O}(\mathsf{com} \parallel M)$$

Thus the prover does not need to interact with a verifier and can compute a non-interactive transcript. Because the sigma protocol described in the preceding sections not only proves knowledge of the secret isogeny between two curves, but also correctness of the torsion points in the public key, we obtain a signature scheme that is also a proof of knowledge of the secret key corresponding to a given SIDH public key, and proves that the SIDH public key is well-formed. For example, simply signing the public key with its own secret key using the new scheme gives a simple NIZK proof of well-formedness for the public key, which provides protection against adaptive attacks. The unforgeability of such a scheme is additionally based on the CSSI assumption.

Such a NIZK proof of knowledge of an SIDH secret key can, among other applications, be used to achieve a secure non-interactive key exchange scheme based on SIDH. Specifically, it would enable both participants to verify non-interactively that the other participant's key is honestly formed and safe to use without fear of adaptive attack.

References

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

[AJL17] Reza Azarderakhsh, David Jao, and Christopher Leonardi. Post-quantum static-static key agreement using multiple protocol instances. In *International Conference on Selected Areas in Cryptography*, pages 45–63. Springer, 2017.

[BKM+20] Andrea Basso, Péter Kutas, Simon-Philipp Merz, Christophe Petit, and Charlotte Weitkämper. On adaptive attacks against Jao-Urbanik's isogeny-based protocol. In *Progress in Cryptology - AFRICACRYPT 2020*, pages 195–213, Cham, 2020. Springer International Publishing.

[DGL+20] Samuel Dobson, Steven D. Galbraith, Jason LeGrow, Yan Bo Ti, and Lukas Zobernig. An adaptive attack on 2-SIDH. *International Journal of Computer Mathematics: Computer Systems Theory*, 5(4):282–299, 2020.

[DJP14] 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.

[FP21] Tako Boris Fouotsa and Christophe Petit. A new adaptive attack on SIDH. Cryptology ePrint Archive, Report 2021/1322, 2021. `https://ia.cr/2021/1322`.

[GMW91] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM (JACM)*, 38(3):690–728, 1991.

[GPS20] Steven D. Galbraith, Christophe Petit, and Javier Silva. Identification protocols and signature schemes based on supersingular isogeny problems. *Journal of Cryptology*, 33(1):130–175, 2020.

[GPST16] Steven D. Galbraith, Christophe Petit, Barak Shani, and Yan Bo Ti. On the security of supersingular isogeny cryptosystems. In *Advances in Cryptology – ASIACRYPT 2016*, pages 63–91. Springer Berlin Heidelberg, 2016.

[GPV21] Wissam Ghantous, Federico Pintore, and Mattia Veroni. Collisions in supersingular isogeny graphs and the SIDH-based identification protocol. Cryptology ePrint Archive, Report 2021/1051, 2021. `https://eprint.iacr.org/2021/1051`.

[GV18] Steven D. Galbraith and Frederik Vercauteren. Computational problems in supersingular elliptic curve isogenies. *Quantum Information Processing*, 17(10):1–22, 2018.

[HHK17] Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In *Theory of Cryptography Conference*, pages 341–371. Springer, 2017.

[HL10] Carmit Hazay and Yehuda Lindell. *Sigma Protocols and Efficient Zero-Knowledge*, pages 147–175. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.

[JD11] David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In *Post-Quantum Cryptography*, pages 19–34, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[JS14] David Jao and Vladimir Soukharev. Isogeny-based quantum-resistant undeniable signatures. In *PQCrypto 2014*, volume 8772 of *Lecture Notes in Computer Science*, pages 160–179. Springer, 2014.

[Leo20] Christopher Leonardi. A note on the ending elliptic curve in SIDH. Cryptology ePrint Archive, Report 2020/262, 2020. `https://ia.cr/2020/262`.

[LZ19] Qipeng Liu and Mark Zhandry. Revisiting post-quantum Fiat-Shamir. In *Annual International Cryptology Conference*, pages 326–355. Springer, 2019.

[MP19] Chloe Martindale and Lorenz Panny. How to not break SIDH. CFAIL, 2019. `https://ia.cr/2019/558`.

[Sil09] Joseph H. Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, Dordrecht, second edition, 2009.

[Tes99] Edlyn Teske. The Pohlig–Hellman method generalized for group structure computation. *Journal of symbolic computation*, 27(6):521–534, 1999.

[Tho17] Erik Thormarker. *Post-Quantum Cryptography: Supersingular Isogeny Diffie-Hellman Key Exchange*. Thesis, Stockholm University, 2017.

[UJ18] David Urbanik and David Jao. SoK: The problem landscape of SIDH. In *Proceedings of the 5th ACM on ASIA Public-Key Cryptography Workshop*, pages 53–60, 2018.

[UJ20] David Urbanik and David Jao. New techniques for SIDH-based NIKE. *Journal of Mathematical Cryptology*, 14(1):120–128, 2020.

[UXT+22] Rei Ueno, Keita Xagawa, Yutaro Tanaka, Akira Ito, Junko Takahashi, and Naofumi Homma. Curse of re-encryption: A generic power/EM analysis on post-quantum KEMs. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 296–322, 2022.

[Vél71] Jacques Vélu. Isogénies entre courbes elliptiques. *C. R. Acad. Sci. Paris Sér. A-B*, 273:A238–A241, 1971.

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