SQISignHD: New Dimensions in Cryptography

Pierrick Dartois[1,2], Antonin Leroux[3,4], Damien Robert[1,2] and Benjamin Wesolowski[5]

[1] Univ. Bordeaux, CNRS, Bordeaux INP, IMB, UMR 5251, F-33400, Talence, France
[2] INRIA, IMB, UMR 5251, F-33400, Talence, France
[3] DGA-MI, Bruz, France,
[4] IRMAR - UMR 6625, Université de Rennes, France
[5] ENS de Lyon, CNRS, UMPA, UMR 5669, Lyon, France

Abstract. We introduce SQISignHD, a new post-quantum digital signature scheme inspired by SQISign. SQISignHD exploits the recent algorithmic breakthrough underlying the attack on SIDH, which allows to efficiently represent isogenies of arbitrary degrees as components of a higher dimensional isogeny. SQISignHD overcomes the main drawbacks of SQISign. First, it scales well to high security levels, since the public parameters for SQISignHD are easy to generate: the characteristic of the underlying field needs only be of the form $2^f 3^{f'} - 1$. Second, the signing procedure is simpler and more efficient. Third, the scheme is easier to analyse, allowing for a much more compelling security reduction. Finally, the signature sizes are even more compact than (the already record-breaking) SQISign, with compressed signatures as small as 116 bytes for the post-quantum NIST-1 level of security. These advantages may come at the expense of the verification, which now requires the computation of an isogeny in dimension 4, a task whose optimised cost is still uncertain, as it has been the focus of very little attention.

Acknowledgements. We thank Luca De Feo for his advice all along this project and for suggesting the title of this paper. This project was supported by ANR grant CIAO (ANR-19-CE48-0008) and PEPR PQ-TLS (the France 2030 program under grant agreement ANR-22-PETQ-0008 PQ-TLS).

1 Introduction

Isogeny-based cryptography has been a promising area of research in post-quantum cryptography since Couveignes, Rostovtsev and Stolbunov introduced the first key exchange using ordinary isogenies [Cou06; RS06]. Schemes from this family often distinguish themselves by their compactness, in particular with respect to key sizes. It is notably the case of the digital signature scheme SQISign [DKLPW20; DLW22], the most compact post-quantum signature scheme by a decent margin. However, efficiency has been a recurring challenge for isogeny-based schemes, and indeed, SQISign is much slower than other post-quantum signatures.

In this paper, we introduce SQISignHD, a new digital signature scheme derived from SQISign. As in [GPS16], SQISign uses the Deuring correspondence between supersingular elliptic curves and quaternion orders. This Deuring correspondence is a powerful tool to construct cryptosystems because it is one way: it is easy to turn an order into the corresponding elliptic curve, but the converse direction is the presumably hard *supersingular endomorphism ring problem* [EHLMP18; Wes22]. In SQISign, the signer's public key is a supersingular elliptic curve, and a signature effectively proves that the signer knows the associated quaternion order. This requires algorithms to translate between orders (and ideals in these orders) and elliptic curves (and isogenies from these curves). This translation is costly, and crucially requires the ideals (or isogenies) to have smooth norms (or degrees). The original methods have been improved upon [DLW22], but that remains the bottleneck of SQISign. Another issue with SQISign is its scalability to higher security levels. Indeed, to set public parameters, one needs to find a prime p such that p^2-1 has a very large smooth factor. Searching for such primes p becomes harder as the security level grows, and is still an active area of research [CMN21; BSC+22; Ahr23]. Besides, the security of SQISign relies on the fact that signatures are computationally indistinguishable from random isogenies of fixed powersmooth degrees. There is no known formal proof of this *ad hoc* heuristic assumption.

The new scheme SQISignHD follows a similar outline as SQISign, but resolves its main drawbacks by fundamentally reforging the computational approach. The main ingredient is the ground-breaking technique that has recently led to the downfall of SIDH [CD22; MMPPW23; Rob22a]. Namely, these attacks use a lemma due to Kani [Kan97] combined with Zahrin's trick, which allows one to "embed" any isogeny into an isogeny of higher dimension. As remarked in [Rob22b], this technique allows one to describe an isogeny by listing only the image of a few well-chosen points; from this description, one can efficiently evaluate the isogeny on any other point, regardless of the factorisation pattern of the underlying isogeny. This newly gained freedom on usable isogenies allows SQISignHD to overcome the main drawbacks of SQISign.

Our contribution. We introduce the digital signature scheme SQISignHD. It leverages recent algorithmic breakthroughs [CD22; MMPPW23; Rob22a] to overcome the main drawbacks of SQISign. It has the following advantages:

- SQISignHD scales well to high security levels. Indeed, while SQISign requires a search for primes p with strong constraints, the primes used in SQISignHD may be of the form $c2^f3^{f'}-1$, where c is some (preferably small) cofactor. Such primes, already used in SIDH [JD11], are easy to find, and allow for efficient field arithmetic.
- The signing procedure of SQISignHD is simpler and more efficient than SQISign. Let us stress that no high dimensional isogeny needs to be computed when signing. Preliminary implementation results show that key generation and signature times are both significantly better than SQISign. The

implementation, and a precise performance analysis, will soon be made available.
- SQISignHD is easier to analyse, allowing for a much more compelling security reduction to the supersingular endomorphism problem. Unlike in SQISign, our proof of the zero-knowledge property in SQISignHD relies on simple and plausible heuristic assumptions. In fact, we propose two variants of SQISign, one of which is less efficient but benefits from a heuristic-free analysis. In both cases, the zero-knowledge property is based on a simulator which is given access to a non-standard oracle. We carefully discuss the impact of this oracle on the supersingular endomorphism problem.
- SQISignHD signatures are even more compact than SQISign, as they are only 6.5λ bits long, for λ bits of security. In particular, they are as small as 116 bytes for the NIST-1 security level. SQISign already had the most compact signature and public keys combined of all post-quantum signature schemes, and SQISignHD breaks this record.

These advantages may come at the expense of the verification, which now requires the computation of a chain of 2-isogenies in dimension 4 (or 8 in the less efficient variant). We provide an algorithm for the verification, and an estimate of its complexity. Its implementation is left for future work, hence the cost of verification is still uncertain. The verification in SQISign also requires the computation of a (longer!) chain of 2-isogenies, but only in dimension 1.

1.1 A modular overview of SQISignHD

We will present two versions of SQISignHD, optimised in different directions. FastSQISignHD is optimised for speed, while RigorousSQISignHD is optimised for the security proof. Note that the security proof applies to both: the difference lies in the proof being unconditional for RigorousSQISignHD when given access to an oracle, but requiring additional heuristics for FastSQISignHD (see Section 6.2 and Section 6.3). Under the hood, FastSQISignHD relies on isogenies of dimension 4, while RigorousSQISignHD relies on isogenies of dimension 8. The reader may sense the parallel with the heuristic (dimension 4) and rigorous (dimension 8) variants of the algorithms of [Rob22a].

We present here the main algorithmic building blocks of SQISignHD to give a modular overview of the protocol. Those algorithms will be presented in detail in the course of the paper for the fast and rigorous version of SQISignHD.

Public set-up. We choose a prime p and a supersingular elliptic curve $E_0/\mathbb{F}_{p^2}$ of known endomorphism ring $\mathcal{O}_0 \cong \mathrm{End}(E_0)$ such that E_0 has smooth torsion defined over a small extension of $\mathbb{F}_{p^2}$ (of degree 1 or 2). In practice, one may use the curve $E_0 : y^2 = x^3 + x$ (and $p \equiv 3 \mod 4$).

Key generation. The prover generates a random secret isogeny $\tau : E_0 \longrightarrow E_A$ of fixed smooth degree D_τ. Then, the prover publishes E_A. Knowing τ, only the prover can compute the endomorphism ring $\mathrm{End}(E_A)$. In the fast method $\mathsf{FastKeyGen}$, the isogeny τ has degree $D_\tau = \Theta(p)$, which is heuristically sufficient to ensure that the distribution of E_A is computationally indistinguishable from uniform. In the alternate method $\mathsf{RigorousKeyGen}$, the degree is chosen a bit larger to make the distribution of E_A statistically close to uniform.

Commitment. The prover generates a random isogeny $\psi : E_0 \longrightarrow E_1$ of smooth degree D_ψ coprime to D_τ, and returns E_1 to the verifier (ψ being secret). The resulting distribution for E_1 is as close as possible to the uniform distribution in the supersingular isogeny graph. As in the key generation, we propose a fast procedure $\mathsf{FastCommit}(E_0)$ in Section 4.3 resulting in a distribution heuristically indistinguishable from uniform, and a slower variant $\mathsf{RigorousCommit}(E_0)$ in Section 4.4 which guarantees statistical closeness to uniform.

Challenge. The verifier generates a random isogeny $\varphi : E_1 \longrightarrow E_2$ of smooth degree D_φ sufficiently large for φ to have high entropy. Then, φ is sent to the prover. The $\mathsf{Challenge}$ procedure is described in Section 4.2.

Response. The prover generates an *efficient representation* of an isogeny $\sigma : E_A \longrightarrow E_2$ of small degree $q \simeq \sqrt{p}$ in the sense of the following definition and returns it to the verifier.

Definition 1.1.1. An *efficient representation* of an isogeny $\varphi : E \longrightarrow E'$ defined over a finite field $\mathbb{F}_q$ is given by a couple $(D, \mathscr{A})$ where:

(i) D is some data of size polynomial in $\log(\deg(\varphi))$ and $\log(q)$ determining the isogeny φ in a unique way.

(ii) $\mathscr{A}$ is a universal algorithm independent of φ returning $\varphi(P)$ as input D and $P \in E(\mathbb{F}_{q^k})$ in polynomial time in $k \log(q)$ and $\log(\deg(\varphi))$.

There always exists an efficient representation of a smooth degree isogeny. For instance, it can be written as a chain of small degree isogenies. Until the recent attacks on SIDH [CD22; MMPPW23; Rob22a], we did not know how to efficiently represent isogenies with non-smooth degrees without revealing the endomorphism ring of the domain. For that reason, the original version of SQISign uses smooth degree isogenies for the signature. These smooth degree isogenies are found with a variant of the KLPT algorithm [KLPT14] and have very big degree $\simeq p^{15/4}$. This not only hurts efficiency, but also security: the isogeny σ is so carefully crafted that it is hard to simulate, and as a result, the zero-knowledge property of SQISign is very *ad hoc*.

Now, the methods from [CD22; MMPPW23; Rob22a] give much more freedom on the isogenies that can be efficiently represented. This allows SQISignHD to improve both efficiency (using isogenies σ of degree as low as $\simeq \sqrt{p}$), and security (the isogenies σ are now nicely distributed, hence simulatable).

The idea is to "embed" σ into an isogeny of higher dimension — and that only requires knowing the image of a few points through σ. As in the attacks against

SIDH, such an isogeny can have dimension 2, 4 or 8. We shall see that dimension 2 has little interest compared to the original SQISign protocol from an efficiency and security point of view. In SQISignHD, we propose a response procedure FastRespond to represent σ in dimension 4, and an alternative procedure RigorousRespond based on an isogeny computation in dimension 8. The procedure FastRespond is fast, and its security analysis relies on reasonable heuristics. On the other hand, RigorousRespond is much slower (though still polynomial time), but allows for a rigorous analysis.

In either case, for efficiency reasons, the prover does not actually compute higher dimensional isogenies but only images of some points through σ (we explain how these points are evaluated in the course of the paper). Those points provide an efficient representation of σ (along with $\deg(\sigma)$) and this data is sent to the prover who can then compute higher dimensional isogenies representing σ.

Verification. The verifier checks that the response returned by the prover (points of E_2) correctly represents an isogeny $\sigma : E_A \longrightarrow E_2$. We propose two procedures FastVerify and RigorousVerify computing isogenies embedding σ in dimension 4 or 8. The efficiency of that task remains to be determined: isogeny computations in dimension 4 has been the subject of very little literature, and no implementation suited to the requirements of SQISignHD is presently available. We refer to Appendix C for an estimate of the number of operations required for the verification. We expect an optimized implementation to compare favourably to the original SQISign procedure.

$$\begin{tikzcd}
E_0 \arrow[r, "\tau"] \arrow[d, "\psi"'] & E_A \arrow[d, "\sigma"] \\
E_1 \arrow[r, "\varphi"] & E_2
\end{tikzcd}$$

— Public

— Prover's secret

Fig. 1. General principle of the SQISign/SQISignHD identification protocol.

1.2 Constructing a signature with the Fiat-Shamir transform

We transform our identification protocol into a signature scheme using the Fiat-Shamir transform [FS87] as in the original SQISign protocol.

Decomposing the degree of the challenge into primes $D_\varphi := \prod_{i=1}^{r} \ell_i^{e_i}$ and setting $\mu(D_\varphi) := \prod_{i=1}^{r} \ell_i^{e_i-1}(\ell_i + 1)$, we define a secure hash function in the supersingular $\{\ell_1, \cdots, \ell_r\}$-isogeny graph mapping a supersingular elliptic curve

E and an integer $s \in [\![1 \,;\, \mu(D_\varphi)]\!]$ to a cyclic D_φ-isogeny $\Phi(E, s)$. Such a hash function has been constructed in [DDF+21, § 3.1], which is a generalization of [CLG09]. We also use another secure hash function $H : \{0,1\}^* \longrightarrow [\![1 \,;\, \mu(D_\varphi)]\!]$.

Signature. To sign a message m with a secret key $\tau : E_0 \longrightarrow E_A$, generate a random commitment $\psi : E_0 \longrightarrow E_1$, let $s := H(j(E_1), m)$ and $\varphi := \Phi(E_1, s) : E_1 \longrightarrow E_2$. From the knowledge of τ, φ and ψ, construct an efficient representation $R = (\sigma(P_1), \sigma(P_2), q)$ given by the image of torsion points by a response isogeny $\sigma : E_A \longrightarrow E_2$ and return (E_1, R) as a signature.

Verification. A verifier receiving a signature (E_1, R) associated to the message m and public key E_A computes $s = H(j(E_1), m)$ and then $\varphi = \Phi(E_1, s) : E_1 \longrightarrow E_2$. The verifier finally checks that R represents correctly an isogeny $\sigma : E_A \longrightarrow E_2$ by computing a higher dimensional isogeny, as explained previously.

Once it is established that the SQISignHD identification protocol is complete, sound, and honest verifier zero-knowledge, and assuming the hardness of the endomorphism ring problem, we obtain a universal unforgeable signature against chosen message attacks in the random oracle model [VV15, Theorem 7]. Completeness will be clear by construction: a honest verifier always accepts a honest execution of the protocol. Other security assumptions (especially the zero-knowledge property which is the less trivial) will be justified in Section 6 for both FastSQISignHD and RigorousSQISignHD.

1.3 Contents

The rest of this paper is organized as follows. In Section 2, we give mathematical background and recall some algorithms for the effective Deuring correspondence already introduced in previous versions of SQISign [DKLPW20; DLW22]. In Section 3, we present the core idea of our paper: how to embed signature/response isogenies in higher dimension with Kani's lemma. Section 4 introduces algorithms for key generation, commitment and challenge whereas Section 5 presents the response and verification phase for both FastSQISignHD and RigorousSQISignHD. A security analysis of both versions of the SQISignHD identification protocol is conducted in Section 6. Finally, we discuss the expected performance of the digital signature scheme derived from FastSQISignHD in Section 7.

Some proofs of our results are deferred to Appendix A. A slight difficulty in the dimension 8 case is that for the security proof some coprimality conditions may not be assumed, the response and verification to handle this case are treated in Appendix B. Finally Appendix C details the verification algorithm when using the theta model to compute isogenies in dimension 4 and 8, and in particular gives an algorithm to compute a 2^e-isogeny and the corresponding number of arithmetic operations.

2 Preliminaries

2.1 Abelian varieties and their isogenies

An abelian variety A over a field k is a connected projective k-variety with an algebraic group law (which is then automatically abelian by rigidity). Abelian varieties are generalizations of elliptic curves in any dimension. In particular, elliptic curves are abelian varieties of dimension 1 and products of elliptic curves are abelian varieties.

A morphism of abelian varieties is an algebraic map $\varphi : A \to B$ which is a group homomorphism; by rigidity it suffices to check that $\varphi(0_A) = 0_B$. It is an *isogeny* if it is surjective and has finite kernel. The degree of an isogeny is its degree as a rational map. If φ is a *separable* isogeny, then $\deg(\varphi) = \# \ker(\varphi)$. If $\deg(\varphi)$ is coprime to the characteristic p of the base field k, then φ is automatically separable. As in elliptic curves, the multiplication by n map in an abelian variety $[n]_A : A \longrightarrow A$ is an isogeny of degree n^{2g} with $g := \dim(A)$ and its kernel $A[n]$ is isomorphic to $(\mathbb{Z}/n\mathbb{Z})^{2g}$ when n is coprime to p.

To any abelian variety A, we associate its *dual abelian variety* $\widehat{A}$, which has the same dimension. The dual defines a contravariant functor: any isogeny $\varphi : A \longrightarrow B$ induces a *dual isogeny* $\widehat{\varphi} : \widehat{B} \longrightarrow \widehat{A}$ which has the same degree. A *polarization* is an isogeny $\lambda : A \longrightarrow \widehat{A}$ induced by an ample line bundle. A polarized abelian variety (A, λ) is *principally polarized* if λ is an isomorphism.

If $n \in \mathbb{N}^*$ is coprime to p, then we have a non-degenerate pairing $e_n : A[n] \times \widehat{A}[n] \longrightarrow \overline{k}^*$ called the *Weil pairing*. Given a polarization $\lambda : A \longrightarrow \widehat{A}$, the Weil pairing yields a non-degenerate antisymmetric pairing $e_n^\lambda : A[n] \times A[n] \longrightarrow \overline{k}^*$.

Morally, a polarization can be seen as "a way to represent an abelian variety". Indeed, in plain generality, we do not have nice analogues for the Weierstrass model for elliptic curves, but every abelian variety can be described by a theta model [Mum66; Mum67a; Mum67b]. We refer to the notes of Milne [Mil86] or the book of Mumford [Mum74] for a complete introduction to abelian varieties.

2.2 The Deuring correspondence

Quaternions, orders and ideals. Let $\mathcal{B}_{p,\infty}$ be the quaternion algebra over $\mathbb{Q}$ ramifying at p and ∞. By [Piz80, Proposition 5.1], there exists a $\mathbb{Q}$-basis $(1, i, j, k)$ of $\mathcal{B}_{p,\infty}$ with $j^2 = -p$, $k = ij = -ji$ and $i^2 = -1, -2, -q$, when $p \equiv 3 \mod 4$, $p \equiv 5 \mod 8$ and $p \equiv 1 \mod 8$ respectively, q being a prime such that $(-p/q) = 1$. $\mathcal{B}_{p,\infty}$ has a *conjugation* $\alpha := x + iy + jz + kt \longmapsto \overline{\alpha} := x - iy - jz - kt$. For all $\alpha \in \mathcal{B}_{p,\infty}$, we define the *reduced norm* $\operatorname{nrd}(\alpha) := \alpha\overline{\alpha}$ and *trace* $\operatorname{Tr}(\alpha) := \alpha + \overline{\alpha}$.

A *fractional ideal* $I \subset \mathcal{B}_{p,\infty}$ is a $\mathbb{Z}$-lattice of rank 4. We also define the *reduced norm* of I as $\operatorname{nrd}(I) := \gcd\{\operatorname{nrd}(\alpha) \mid \alpha \in I\}$ and the conjugation $\overline{I} := \{\overline{\alpha} \mid \alpha \in I\}$. If $I \subset J$ are two fractional ideals, then $[J : I] = \operatorname{nrd}(I)^2 / \operatorname{nrd}(J)^2$. If $(\alpha_1, \cdots, \alpha_4)$ is a basis of I, then $|\det(\operatorname{Tr}(\alpha_i\overline{\alpha_j}))_{1 \leq i,j \leq 4}|^{1/2}$ does not depend on the basis. This invariant is called the the *reduced discriminant* of I and denoted by $\operatorname{discrd}(I)$.

An *order* of $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ is a fractional ideal which is stable by multiplication and contains 1. We say it is *maximal* if it is maximal for the inclusion. If I is a fractional ideal, we define its *left order* $O_L(I) := \{\alpha \in \mathcal{B}_{p,\infty} \mid \alpha I \subseteq I\}$ and its *right order* $O_R(I) := \{\alpha \in \mathcal{B}_{p,\infty} \mid I\alpha \subseteq I\}$. We say that I is a left (respectively right) $\mathcal{O}$-ideal when $\mathcal{O} \subseteq O_L(I)$ (respectively $\mathcal{O} \subseteq O_R(I)$). We also say that I *connects* $\mathcal{O}$ and $\mathcal{O}'$ when $\mathcal{O} = O_L(I)$ and $\mathcal{O}' = O_R(I)$. I is *integral* if $I \subset O_L(I)$. In this case, $I \subset O_R(I)$ and both $O_L(I)$ and $O_R(I)$ are maximal. **In the following, we shall only consider integral ideals and simply refer to them as ideals.** Two fractional ideals $I \sim J$ are equivalent if there exists $\beta \in B_{p,\infty}^*$ such that $J = I\beta$. In that case, $O_L(I) = O_L(J)$ and $O_R(I) = \beta O_R(J)\beta^{-1}$.

The Deuring correspondence. The Deuring correspondence due to Max Deuring [Deu41] draws a parallel between the world of quaternions and the world of supersingular elliptic curves. Indeed, if $E/\mathbb{F}_{p^2}$ is a supersingular elliptic curve, then its endomorphism ring $\mathrm{End}(E)$ is isomorphic to a maximal order $\mathcal{O} \subset \mathcal{B}_{p,\infty}$.

Example 2.2.1. If $p \equiv 3 \mod 4$, the elliptic curve $E_0 : y^2 = x^3 + x$ defined over $\mathbb{F}_p$ is supersingular and has a very explicit endomorphism ring $\mathrm{End}(E_0)$ isomorphic to $\mathcal{O}_0 := \langle 1, i, (i+j)/2, (1+k)/2\rangle$, where j corresponds to the Frobenius endomorphism $(x, y) \in E_0 \longmapsto (x^p, y^p) \in E_0$ and i corresponds to the automorphism $(x, y) \in E_0 \longmapsto (-x, \zeta y) \in E_0$ (with $\zeta \in \mathbb{F}_{p^2}$ such that $\zeta^2 = -1$).

This is one of the very few examples where $\mathrm{End}(E_0)$ can be easily and explicitly computed. Computing $\mathrm{End}(E)$ is difficult in general.

Let E be a supersingular elliptic curve and $\mathcal{O} := \mathrm{End}(E)$. An isogeny $\phi : E \longrightarrow E'$ has a *kernel ideal* $I_\phi := \{\alpha \in \mathcal{O} \mid \forall P \in \ker(\phi), \quad \alpha(P) = 0\}$, which is a left $\mathcal{O}$-ideal of norm $\mathrm{nrd}(I_\phi) = \deg(\phi)$. Conversely, any left $\mathcal{O}$-ideal I defines an isogeny $\phi_I : E \longrightarrow E_I$ of kernel $E[I] := \{P \in E \mid \forall \alpha \in I, \quad \alpha(P) = 0\}$ and degree $\deg(\phi_I) = \mathrm{nrd}(I)$. We have $I_{\phi_I} = I$ and $\phi_{I_\phi} = \phi$ so this correspondence is one to one.

The Deuring correspondence between ideals and isogenies satisfies the following properties: two equivalent ideals $I \sim J$ have isomorphic codomains $E_I \simeq E_J$, the endomorphism ring of the codomain E_I is $\mathrm{End}(E_I) \cong O_R(I)$, the conjugate $\overline{I}$ corresponds to the dual isogeny $\widehat{\phi_I}$, the kernel ideal of the composite $\phi \circ \psi$ isogeny is $I_{\phi\circ\psi} = I_\psi \cdot I_\phi$, a principal ideal corresponds to an endomorphism. For a thorough presentation of quaternions and the Deuring correspondence, we recommend the book of Voight [Voi20].

Accessible torsion to make the Deuring correspondence effective. Making the Deuring correspondence effective means computing the isogeny $\phi_I : E \longrightarrow E_I$ associated to an ideal I and conversely, computing the kernel ideal I_ϕ of a known isogeny $\phi : E \longrightarrow E'$ when $\mathrm{End}(E)$ is known. Until recently, this could be done in polynomial time only when the ideal norm/the degree is smooth and the necessary $\mathrm{nrd}(I)$-torsion to do these computations is "accessible" in the following sense.

Definition 2.2.2. Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve and $T := \prod_{i=1}^r \ell_i^{e_i}$ be an integer, where the ℓ_i are distinct primes. We say that E has *accessible* T-torsion if $E[\ell_i^{e_i}]$ is defined over an extension of $\mathbb{F}_p$ of degree polynomial in $\log(p)$ for all $i \in [\![1\ ;\ r]\!]$.

Usually, in cryptographic protocols we choose p to ensure the T-torsion is defined over $\mathbb{F}_{p^2}$ or $\mathbb{F}_{p^4}$ to optimize T-isogeny computations when T is smooth. In general, if T is B-powersmooth with B polynomial in $\log(p)$, the T-torsion is always accessible and we can compute T-isogenies in polynomial time as a product of low degree isogenies as in [EHLMP18, Proposition 4]. Until recently, those were the only way to compute isogenies and make the Deuring correspondence effective. In this work, we propose a method to compute isogenies of non-smooth degree (see Section 3). We specialize it for SQISignHD but this could be easily generalized.

2.3 Algorithms for effective Deuring correspondence

In this section, we recall already known polynomial time algorithms making the Deuring correspondence effective. Those algorithms are used as ingredients of SQISignHD. They were mainly introduced for previous versions of SQISign [DKLPW20; DLW22].

Pushing the endomorphism ring through an isogeny. In this paragraph, we introduce an algorithm to compute a basis of $\mathrm{End}(E)$ that we can easily evaluate when we know an isogeny $E_0 \longrightarrow E$ and a basis of $\mathrm{End}(E_0)$ (in practice, E_0 is the elliptic curve of Example 2.2.1).

Definition 2.3.1. Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve and $\mathcal{O}$ be a maximal quaternion order in $\mathcal{B}_{p,\infty}$ isomorphic to $\mathrm{End}(E)$. An *eval-basis* of $\mathrm{End}(E)$ is the data of a basis $(\alpha_1, \cdots, \alpha_4)$ of $\mathcal{O}$ and an isomorphism $\varepsilon : \mathcal{O} \overset{\sim}{\longrightarrow} \mathrm{End}(E)$ such that the $\varepsilon(\alpha_i)$ can be evaluated at any point of E in polynomial time in $\log(p)$. We say it is an N*-eval-basis* if we can only evaluate the $\varepsilon(\alpha_i)$ on points of order coprime to N.

For any left-ideal $I \subseteq \mathcal{O}$, we define an (N-)eval-basis of I in a similar way. Such a basis can be obtained from an (N-)eval-basis of $\mathrm{End}(E)$.

Assume that we know an eval-basis $((\alpha_1, \cdots, \alpha_4), \varepsilon)$ of $\mathrm{End}(E_0)$. Let $\psi : E_0 \longrightarrow E_1$ be an isogeny of degree N with an efficient representation. Here, we explain how to use ψ to compute an N-eval-basis of $\mathrm{End}(E_1)$.

By [Voi20, Lemma 42.2.9], the map

$$\begin{aligned} \iota : \mathrm{End}(E_1) &\longrightarrow \mathcal{B}_{p,\infty} \\ \phi &\longmapsto \frac{1}{N}\varepsilon^{-1}(\widehat{\psi} \circ \phi \circ \psi) \end{aligned}$$

```
Algorithm 1: PushEndRing [EHLMP18, Algorithm 4]
Data: An eval-basis $((\alpha_1, \cdots, \alpha_4), \varepsilon)$ of $\mathrm{End}(E_0)$ and an isogeny $\psi : E_0 \longrightarrow E_1$
      of degree $N$ with known kernel ideal $I_\psi$.
Result: An $N$-eval-basis $(\beta_1, \cdots, \beta_4)$ of $\mathrm{End}(E_1)$.
  Compute a $\mathbb{Z}$-basis $(\beta_1, \cdots, \beta_4)$ of $\mathcal{O}_1 := 1/N\overline{I_\psi}I_\psi$;
  Write $\beta_i := 1/N \sum_{j=1} c_{i,j}\alpha_j$, with $c_{i,j} \in \mathbb{Z}$, for all $i = 1, \cdots, 4$;
  Let $\varepsilon' : \mathcal{O}_1 \xrightarrow{\sim} \mathrm{End}(E_1), \beta_i \longmapsto 1/N \sum_{j=1}^4 c_{i,j}\psi \circ \varepsilon(\alpha_j) \circ \widehat{\psi}$;
  Return $(\beta_1, \cdots, \beta_4)$ and $\varepsilon'$;
```

```
Algorithm 2: KernelToIdeal_$D$
Data: A point $P \in E_1$ of smooth order $D$ and an eval-basis $((\beta_1, \cdots, \beta_4), \varepsilon')$
      of $\mathrm{End}(E_1) \cong \mathcal{O}_1$ that can be evaluated on $E_1[D]$.
Result: The ideal $I(\langle P \rangle) \subseteq \mathcal{O}_1$ associated to $\langle P \rangle$.
  Compute $Q_i := \varepsilon'(\beta_i)(P)$ for $i = 1, \cdots, 4$;
  Find $i, j$ such that $(Q_i, Q_j)$ is a basis of $E_1[D]$;
  For $k \neq i, j$, find $a, b \in \mathbb{Z}/D\mathbb{Z}$ such that $Q_k = aQ_i + bQ_j$ (discrete logarithm
  problem);
  Let $\gamma := \beta_k - a\beta_i - b\beta_j$;
  Return $\mathcal{O}_1\gamma + \mathcal{O}_1 D$;
```

induces an isomorphism $\mathrm{End}(E_1) \xrightarrow{\sim} \mathcal{O}_1 := O_R(I_\psi)$. Assuming that we know I_ψ, we can obtain a $\mathbb{Z}$-basis $(\beta_1, \cdots, \beta_4)$ of $\mathcal{O}_1$ via the formula $\mathcal{O}_1 = 1/N\overline{I_\psi}I_\psi$ [Voi20, Proposition 16.6.15]. Then, ι^{-1} induces an isomorphism $\mathcal{O}_1 \xrightarrow{\sim} \mathrm{End}(E_1)$.

We now explain how to use ι^{-1} to evaluate the β_i. Since $N\mathcal{O}_1 = \overline{I_\psi}I_\psi \subseteq \mathcal{O}_0$, we can write

$$\beta_i := \frac{1}{N}\sum_{j=1} c_{i,j}\alpha_j,$$

with $c_{i,j} \in \mathbb{Z}$, for all $i \in [\![1\ ;\ 4]\!]$. We then have:

$$\iota^{-1}(\beta_i) = \frac{1}{N^2}\sum_{j=1}^{4} c_{i,j}\psi \circ \varepsilon(\alpha_j) \circ \widehat{\psi},$$

so we can indeed easily evaluate the β_i on any point of E_1 of order coprime to N.

From isogenies to ideals. Let $\varphi : E_1 \longrightarrow E_2$ be a cyclic D-isogeny (with D smooth) represented by a generator of its kernel P. We give an algorithm $\mathsf{KernelToIdeal}_D$ (Algorithm 2) due to Leroux [Ler22, Algorithm 20] to compute the ideal associated to a cyclic group $\langle P \rangle$ of a supersingular elliptic curve E_1 when we know an eval-basis of $\mathrm{End}(E_1)$ that can be evaluated on $E_1[D]$. This algorithm is a variant of the algorithm introduced by Galbraith, Petit and Silva [GPS16, Algorithm 2] with the same purpose.

Algorithm 3: IsogenyToIdeal

```
Data: An isogeny $\varphi : E_1 \longrightarrow E_2$ of degree $D_\varphi$, another isogeny $\psi : E_0 \longrightarrow E_1$ of
      degree $D_\psi$ along with its kernel ideal $I_\psi$ and an eval-basis $\mathcal{B}_0$ of $\mathrm{End}(E_0)$.
Result: The kernel ideal $I_\varphi$ associated to $\varphi$.
Compute a $D_\psi$-eval-basis $\mathcal{B}_1 := \mathsf{PushEndRing}(\mathcal{B}_0, \psi, I_\psi)$ of $\mathrm{End}(E_1)$;
Let $P$ be a generator of $\ker(\varphi)$;
$I_\varphi \longleftarrow \mathsf{KernelToIdeal}_{D_\varphi}(P, \mathcal{B}_1)$;
Return $I_\varphi$;
```

Remark 2.3.2. Since D is smooth, the discrete logarithm problem in line 3 of Algorithm 2 is easy to solve with Pohlig-Hellman methods generalized by Teske to multiple discrete logarithms [PH78; Tes99].

Lemma 2.3.3. *[Ler22, Lemma 4.22] Algorithm 2 terminates and is correct.*

Kernel ideal computation. The two preceding algorithms immediately yield an algorithm IsogenyToIdeal computing the kernel ideal of $\varphi : E_1 \longrightarrow E_2$ when given an isogeny $\psi : E_0 \longrightarrow E_1$ of known kernel ideal I_ψ and an eval-basis of $\mathrm{End}(E_0)$.

The KLPT algorithm. The KLPT algorithm was first introduced in [KLPT14] for extremal orders (such as $\mathcal{O}_0$) and then generalized to other orders and improved in [DKLPW20], [DLW22] and [Ler22]. In the following, we refer to the version of KLPT introduced in [Ler22, Algorithm 7]. Given a left $\mathcal{O}_0$-ideal I and a (smooth) integer $N = \Omega(p^3)$, $\mathsf{KLPT}_N(I)$ returns an equivalent ideal $J \sim I$ of norm dividing N.

This algorithm is generally used when we want to compute an isogeny path $E_0 \longrightarrow E$. Given an ideal I connecting $\mathcal{O}_0 \cong \mathrm{End}(E_0)$ to $\mathcal{O} \cong \mathrm{End}(E)$, which may not have smooth norm, we find $J \sim I$ with KLPT of smooth norm N, so that the isogeny associated to J, $\phi_J : E_0 \longrightarrow E$ can be computed.

From ideal to isogenies. Given a maximal order $\mathcal{O}$ and a left $\mathcal{O}$-ideal J, we have a straightforward way to compute the associated isogeny $\phi_J : E \longrightarrow E'$. We evaluate a basis of J on $E[\mathrm{nrd}(J)]$ to compute $E[J] = \ker(\phi_J)$ and then compute ϕ_J as a chain of small degree isogenies (assuming $\mathrm{nrd}(J)$ is smooth) using Vélu's formulas [Vél71]. Unfortunately, we might not have accessible $\mathrm{nrd}(J)$-torsion, especially if J is obtained from KLPT ($\mathrm{nrd}(J) = \Theta(p^3)$ when $\mathcal{O} = \mathcal{O}_0$ and even bigger otherwise).

In [DKLPW20, Algorithm 7], the authors introduced the SpecialIdealToIsogeny algorithm to perform this computation with half of the torsion when $\mathcal{O} = \mathcal{O}_0$ and when we know an alternate isogeny path. Assume we have accessible T-torsion with $T = \Omega(p^{3/2})$. SpecialIdealToIsogeny takes as input two left-ideals $J, I \subseteq \mathcal{O}_0$ of coprime norm such that $I \sim J$ and $\mathrm{nrd}(J) | T^2$ along with the isogeny $\phi_I : E_0 \longrightarrow E$ associated to I. It returns the isogeny $\phi_J : E_0 \longrightarrow E$ associated to J.

3 Representing the response isogeny efficiently in higher dimension

In this section, we explore our main idea to improve SQISign by embedding the signature isogeny inside an isogeny in higher dimension. We start by recalling how the signature is represented in the original SQISign protocol in Section 3.1 and why this representation is slow to compute. Then, we introduce Kani's lemma and explain how to embed isogenies in higher dimension in Section 3.2. Finally, we apply this idea to provide another representation of the signature isogeny in SQISign in Section 3.3.

3.1 Representing isogenies in dimension 1: a slow signature process

In dimension 1, we can only efficiently represent isogenies of smooth degrees. That is why in the original versions of SQISign [DKLPW20; DLW22], the signature isogeny σ has degree a prime power ℓ^e and is represented as a chain of ℓ-isogenies.

To compute such a signature σ, the prover computes the ideal J associated to $\hat{\tau} \circ \psi \circ \varphi$ and then applies a SigningKLPT algorithm to J, to return a random equivalent ideal $I \sim J$ of norm ℓ^e. Then, the prover converts I into an isogeny. This last computation is very costly because $\mathrm{nrd}(I) = \ell^e$ is close to $p^{15/4}$, while the accessible torsion points have much smaller order. The method introduced in [DKLPW20] (and later improved in [DLW22]) requires to cut J into several pieces in order to compute σ as a chain of isogenies. This complicated mechanism is by far the bottleneck in the signing algorithm.

In order to avoid this costly ideal to isogeny translation in SQISignHD, we shall no longer require σ to have smooth degree and embed it in an isogeny of dimension 4 or 8. This embedding will provide an efficient representation. We expect such a representation in dimension 4 to be faster to compute than the current one in the original SQISign (nonetheless, an implementation would be needed to confirm it). We shall also explain why this improves security in section Section 6.

3.2 Embedding isogenies in higher dimension with Kani's lemma

In this section, we explain in more details this idea of embedding isogenies in higher dimension. For that, we need a few definitions first.

Definition 3.2.1 (d-isogeny)**.** Let $\alpha : (A, \lambda_A) \longrightarrow (B, \lambda_B)$ be an isogeny between principally polarized abelian varieties. We say that α is a *d-isogeny* if $\widehat{\alpha} \circ \lambda_B \circ \alpha = [d]\lambda_A$, where $\widehat{\alpha} : \widehat{B} \longrightarrow \widehat{A}$ is the dual isogeny of α.

Equivalently, α is a d-isogeny if $\tilde{\alpha} \circ \alpha = [d]_A$, where $\widetilde{\alpha} := \lambda_A^{-1} \circ \widehat{\alpha} \circ \lambda_B$ is the dual isogeny of α with respect to the principal polarisations λ_A and λ_B.

Definition 3.2.2 (Isogeny diamond)**.** Let $a, b \in \mathbb{N}^*$. An (a, b)-*isogeny diamond* is a commutative diagram of isogenies between principally polarized abelian varieties

$$\begin{tikzcd} A' \arrow[r, "\varphi'"] & B' \\ A \arrow[u, "\psi"] \arrow[r, "\varphi"] & B \arrow[u, "\psi'"'] \end{tikzcd}$$

where φ and φ' are a-isogenies and ψ and ψ' are b-isogenies.

Lemma 3.2.3 (Kani)**.** *We consider an* (a, b)*-isogeny diamond as above, with* $d := a+b$ *prime to the characteristic of the base field of abelian varieties. Then, the isogeny* $F : A \times B' \longrightarrow B \times A'$ *given in matrix notation by*

$$F := \begin{pmatrix} \varphi & \widetilde{\psi'} \\ -\psi & \widetilde{\varphi'} \end{pmatrix}$$

is a d*-isogeny with* $d = a + b$*, for the product polarisations.*

If a *and* b *are coprime, the kernel of* F *is*

$$\ker(F) = \{(\widetilde{\varphi}(x), \psi'(x)) \mid x \in B[d]\}.$$

This lemma has first been proved in [Kan97, Theorem 2.3]. We also give a proof in Appendix A.1.

Remark 3.1. The existence of $F : A \times B' \to B \times A'$, implies the existence of $\varphi : A \to B$. We can recover φ as $\pi \circ F \circ \iota$ where ι is any embedding morphism from A to $A \times B'$ and π is the projection from $B \times A'$ to B. Hence, F is an efficient representation of φ.

3.3 Application of Kani's lemma to SQISign

Let us now see how we propose to sign with Kani's Lemma (Lemma 3.2.3) in SQISignHD.

Signing in dimension 4. The idea is to embed the signature $\sigma : E_A \longrightarrow E_2$ in an isogeny of dimension 4. We consider the 2-dimensional q-isogeny $\Sigma := \mathrm{Diag}(\sigma, \sigma) : E_A^2 \longrightarrow E_2^2$, the $(a_1^2 + a_2^2)$-isogeny

$$\alpha := \begin{pmatrix} a_1 & a_2 \\ -a_2 & a_1 \end{pmatrix} \in \mathrm{End}(E_A^2)$$

with $a_1, a_2 \in \mathbb{Z}$, and α', the analogue of α in $\mathrm{End}(E_2^2)$. Then, we have an isogeny diamond

$$\begin{tikzcd}
E_2^2 \arrow[r, "\alpha'"] & E_2^2 \\
E_A^2 \arrow[u, "\Sigma"] \arrow[r, "\alpha"] & E_A^2 \arrow[u, "\Sigma"']
\end{tikzcd}$$

yielding an N-isogeny (with $N := q + a_1^2 + a_2^2$):

$$F := \begin{pmatrix} \alpha & \widetilde{\Sigma} \\ -\Sigma & \widetilde{\alpha'} \end{pmatrix} \in \mathrm{End}(E_A^2 \times E_2^2).$$

Notation 3.3.1. We shall denote $F(\sigma, a_1, a_2)$ when we want to specify the dependence of F on σ, a_1, a_2.

We choose the parameters q, a_1, a_2, so that $N = \ell^e$, with ℓ a small prime and $e \in \mathbb{N}^*$ big enough but as small as possible. Provided that q and ℓ are coprime, we know that

$$\ker(F) = \{(\widetilde{\alpha}(P), \Sigma(P)) \mid P \in E_A^2[\ell^e]\},$$

by Lemma 3.2.3. Then, knowing $\ker(F)$ we can compute F as an ℓ-isogeny chain (see Section 5.5) and obtain an efficient representation of σ, as explained in Remark 3.1.

It follows that our idea requires to compute $\ker(F)$, which becomes easy once we know how to evaluate σ on $E_A[\ell^e]$, by the formula for $\ker(F)$ given above. The idea is to use the alternate isogeny path $\varphi \circ \psi \circ \widehat{\tau} : E_A \to E_2$. Since the signature requires to compute the three isogenies φ, ψ, τ, it will not cost too much to use them in order to evaluate σ. There are several technicalities to make it work in practice (such as to making sure that this alternate path has degree prime to ℓ) but it is manageable (see Section 5.3).

Computing such a representation for the signature is simpler than in the original SQISign protocol. This shifts the main computation effort to the verification, where the actual isogeny in dimension 4 must be computed.

Nonetheless, even though we no longer impose $q = \deg(\sigma)$ to be smooth, we still impose conditions on q to make it work. In practice, ℓ^e will be fixed by the accessible torsion of elliptic curves and we shall need $\ell^e - q$ to be a prime congruent to 1 modulo 4 in order to decompose it easily as a sum of two squares $\ell^e - q = a_1^2 + a_2^2$ by Cornacchia's algorithm [Cor08]. In particular, q will need to be relatively small ($q \simeq \ell^e \simeq \sqrt{p}$). This choice of q ensures its coprimality with ℓ, as required to compute $\ker(F)$.

Definition 3.3.2. We say that an integer q is ℓ^e*-good* when $\ell^e - q$ is a prime number congruent to 1 modulo 4.[6]

[6] One could improve slightly the scheme by defining ℓ^e-good integers as integers q such that $\ell^e - q = sq'$, with s a smooth integer whose prime factors are all congruent to 1 modulo 4 and q' is a prime congruent to 1 modulo 4. Indeed, all we really need

The issue of the signature distribution. Those restrictions on the degree q impacts the distribution of signatures. The bound $\ell^e \simeq \sqrt{q}$ is also restrictive (see Theorem 6.2.2). For that reason, we need some plausible heuristic assumptions to prove the zero-knowledge property of our scheme. This can be fixed by going to the dimension 8 as long as $q < \ell^e$ and $\ell^e = \Omega(p^2)$ as we shall see in the next paragraph. This way, we shall obtain a uniform distribution of signatures and a provably zero-knowledge scheme which is the purpose of our scheme in dimension 8 that we present below.

Signing in dimension 8. By Lagrange's four square theorem [Lag70], if $q < \ell^e$, there always exists $a_1, \cdots, a_4 \in \mathbb{Z}$ such that $q + a_1^2 + \cdots + a_4^2 = \ell^e$. We can find such a decomposition in polynomial time in e with Rabin and Shallit's algorithm [Rab86] improved by Pollack and Treviño [PT18]. We then consider the endomorphism

$$\alpha := \begin{pmatrix} a_1 & -a_2 & -a_3 & -a_4 \\ a_2 & a_1 & a_4 & -a_3 \\ a_3 & -a_4 & a_1 & a_2 \\ a_4 & a_3 & -a_2 & a_1 \end{pmatrix} \in \mathrm{End}(E_A^4),$$

which is an $(a_1^2 + \cdots + a_4^2)$-isogeny, its analogue $\alpha' \in \mathrm{End}(E_2^4)$ and the q-isogeny $\Sigma := \mathrm{Diag}(\sigma, \cdots, \sigma) : E_A^4 \longrightarrow E_2^4$. As previously, by Kani's lemma, we have an isogeny diamond

$$\begin{tikzcd} E_2^4 \arrow[r, "\alpha'"] & E_2^4 \\ E_A^4 \arrow[u, "\Sigma"] \arrow[r, "\alpha"] & E_A^4 \arrow[u, "\Sigma"'] \end{tikzcd}$$

yielding an ℓ^e-isogeny

$$F := \begin{pmatrix} \alpha & \widetilde{\Sigma} \\ -\Sigma & \widetilde{\alpha'} \end{pmatrix} \in \mathrm{End}(E_A^4 \times E_2^4).$$

Notation 3.3.3. We shall denote $F(\sigma, a_1, \cdots, a_4)$ when we want to specify the dependence of F on $\sigma, a_1, \cdots, a_4$.

To ensure the uniformity of the response, in dimension 8 we no longer restrict to the case q coprime to ℓ. We treat this general case in Appendix B. For simplicity, in the main exposition of the protocol we will still assume that q is prime to ℓ. Then, we have

$$\ker(F) = \{(\widetilde{\alpha}(P), \Sigma(P)) \mid P \in E_A^4[\ell^e]\}.$$

is that $\ell^e - q$ is easy to factor so Cornacchia's algorithm can be applied efficiently. This alternate definition would improve a bit the search for ℓ^e good integer, but we went for the simplest definition.

F provides an efficient representation of σ and is computable in polynomial time once we know $\ker(F)$, *i.e.* when we know how to evaluate σ on $E_A[\ell^e]$. This way, we can represent any signature isogeny σ of degree $q < \ell^e$, with the implications on the security proof that we mentioned before. However, computing isogenies in dimension 8 is much more costly than in dimension 4 (though, still polynomial), so we do not recommend to use this representation and only propose it in the alternate provably secure version RigorousSQISignHD.

Why not signing in dimension 2? The cost of computing an isogeny grows exponentially with the dimension [LR12; LR15; LR23]. For that reason, finding an efficient representation in dimension 2 could be fruitful for SQISignHD. On the other hand, the higher the dimension, the lesser the constraints on the isogeny σ. We have already seen that going from dimension 4 to 8 relaxes the constraints on $q = \deg(\sigma)$. Unsurprisingly, the constraints on σ are tighter in dimension 2.

In order to embed σ in an isogeny between abelian varieties in dimension 2, we want to apply Kani's lemma to the following isogeny diamond:

$$\begin{tikzcd}
E_2' \arrow[r, "\widehat{\sigma'}"] & E_A' \\
E_2 \arrow[u, "\beta'"] \arrow[r, "\widehat{\sigma}"] & E_A \arrow[u, "\beta"']
\end{tikzcd}$$

where β is an isogeny of degree b such that $q + b = \ell^e$. This diamond induces an ℓ^e-isogeny $F : E_2 \times E_A' \longrightarrow E_A \times E_2'$, given by

$$F := \begin{pmatrix} \widehat{\sigma} & \widehat{\beta} \\ -\beta' & \sigma' \end{pmatrix},$$

with kernel:

$$\ker(F) = \{(\sigma(P), \beta(P)) \mid P \in E_A[\ell^e]\}.$$

Unlike previously, to compute $\ker(F)$, we not only need to evaluate σ on the ℓ^e-torsion, but also the auxiliary isogeny $\beta : E_A \longrightarrow E_A'$. In particular, the problems comes from the degree b of β. The value of b is defined by the equality $b + q = \ell^e$. Ideally, we would like to have a smooth b so we can compute β with the Vélu formulas. However, since the value of ℓ^e is fixed, it is clear that choosing a smooth b is not easier than choosing q to be smooth. Of course, this can be done with the SigningKLPT algorithm introduced in [DKLPW20], but at the cost of increasing a lot the size of q. Since we must have $\ell^e > q$, this approach will not be faster than the original SQISign idea.

The other possibility is to accept that b cannot be smooth, so we can choose q as small as possible. In that case, we will have ℓ^e of reasonable size but we will not be able to compute β from the Vélu's formulas directly as its degree might contain a big prime factor. This problem can be solved by computing an alternate path of smooth degree between E_A and E_A' with the KLPT algorithm. However, we fall back to a situation where we will need a costly ideal-to-isogeny translation via the Deuring correspondence [DLW22, Algorithm 5], which imposes to use at least as much accessible torsion than in current version of SQISign.

Given what we explain, it is unclear if we can apply our idea in dimension 2 without doing something that will be essentially equivalent to what is done in the original SQISign scheme from an efficiency and security point of view.

4 Key generation, commitment and challenge

In order to be able to evaluate σ on the ℓ^e-torsion, as required for the response computation, the secret key, challenge isogeny and commitment need to satisfy several constraints. We propose to use the alternate isogeny path $\varphi\circ\psi\circ\widehat{\tau}: E_A \longrightarrow E_2$ in order to evaluate σ. As a consequence, the degrees D_φ, D_ψ and D_τ of the challenge, commitment isogeny and secret key respectively must be coprime to ℓ and we also need to know their respective kernel ideals I_φ, I_ψ and I_τ (as will be explained in Section 5.3). These ideals are necessary to compute the ideal I_σ anyway.

To compute I_φ (unknown to the prover *a priori*) using IsogenyToIdeal (Algorithm 3), we also need an alternate path to the commitment $\psi': E_0 \longrightarrow E_1$ of degree $D_{\psi'}$ coprime to D_φ. Due to constraints on the available torsion, D_φ and D_ψ will not be coprime so we cannot have $\psi'=\psi$. The easiest is to require $D_{\psi'}$ to be a power of ℓ. The prover will compute both ψ and ψ' at the same time during the commitment phase.

4.1 Accessible torsion and choice of the prime characteristic

In FastSQISignHD, constraints are comparable to SIDH. The choice of p is usually made to provide enough accessible torsion for our isogeny computations. In FastSQISignHD, we can choose $p=c\ell^f\ell'^{f'}-1$ with $\ell\neq\ell'$ two primes, $c\in\mathbb{N}^*$ small and $\ell^f\simeq\ell'^{f'}\simeq\sqrt{p}$, as in SIDH [JD11]. In practice, $\ell=2$ and $\ell'=3$ are the best choice.

We then require $D_\tau=D_\psi=\ell'^{2f'}$, $D_\varphi=\ell'^{f'}$ and $D_{\psi'}=\ell^{2f}$[7]. This choice ensures that D_τ, D_ψ and D_φ are coprime to ℓ and that $D_{\psi'}$ is coprime to D_ψ, as needed. We also have $D_\tau, D_\psi, D_{\psi'}=\Theta(p)$, which guarantees (at least heuristically) that the public key E_A and the commitment E_1 are computationally indistinguishable from a uniformly random supersingular elliptic curve – which is essential to the security of FastSQISignHD.

This choice of prime also provides enough accessible torsion to compute the ℓ^e-isogeny F representing the response σ in dimension 4. In fact, we even have much more than the minimum requirement since it will be enough to have $2f\geq e+4$ (so $\ell^f=\Omega(p^{1/4})$) as will be explained in Sections 5.4 and 5.6

[7] Actually, we will not have exactly $D_\tau=D_\psi=\ell'^{2f'}$ but D_τ and D_ψ will be divisors of $\ell'^{2f'}$ close to $\ell'^{2f'}$. It will be the same for $D_{\psi'}$ (see Algorithm 4.3.3). We assume equality to simplify the exposition.

and Remark 5.3. This freedom is welcome anyway because it allows us to take ℓ^e slightly bigger than $\sqrt{p}$ to make sure that we can always find an ideal I of ℓ^e-good norm $q < \ell^e$ (see Section 5.2).

We finally discuss the security requirements regarding the size of p. The best known classical key recovery attacks are the meet-in-the-middle algorithm in the isogeny graph or the general Delfs and Galbraith attack [DG16] in the supersingular isogeny graph which both have a complexity in $\tilde{O}(\sqrt{p})$. Using Grover's algorithm [Gro96], we reach a quantum complexity of $\tilde{O}(p^{1/4})$. Hence, to ensure a classical security level of λ bits and a quantum security level of $\lambda/2$ bits, we need to take $p = \Theta(2^{2\lambda})$, as in the original version of SQISign [DKLPW20].

We give below some concrete values of primes for NIST levels 1, 3 and 5.

NIST security level	Security parameter λ (bits)	Prime p
NIST-I	128	$2^{128} \cdot 3^{81} - 1$
NIST-III	192	$5 \cdot 2^{193} \cdot 3^{122} - 1$
NIST-V	256	$11 \cdot 2^{257} \cdot 3^{163} - 1$

In RigorousSQISignHD, more torsion is needed. In RigorousSQISignHD we have similar constraints on the size of p ($p = \Theta(2^{2\lambda})$) and the degrees $D_\tau, D_\psi, D_{\psi'}$ and D_φ, except that we need E_1 and E_A to be rigorously computationally indistinguishable from a uniformly random supersingular elliptic curve. For that reason, we shall have $D_\tau, D_\psi = \Theta(p^3)$ and $D_{\psi'} = \ell^h = \Theta(p^2)$ (as will be explained in Section 4.4).

For the computation of the secret key τ and commitment isogeny ψ and challenge φ, we need accessible T-torsion such that $D_\tau, D_\psi | T^2$ with T coprime to ℓ and $T = \Omega(p^{3/2})$, as in the original SQISign protocol [DKLPW20]. We also require $D_\varphi | T$ to compute the challenge whose degree is $D_\varphi = O(p^3)$ (as will be explained in Appendix A.6). Hence, we require $T \simeq p^3$.

As previously, we can compute the ℓ^e-isogeny F representing the response σ in dimension 8 as long as we have accessible ℓ^f-torsion with $2f \geq e+4$. However, to prove the zero-knowledge property, we need $\ell^e = \Theta(p^2)$ instead of $\Theta(\sqrt{p})$, so we can have $\ell^f = \Theta(p)$. Hence p should be of the form $p = c\ell^f - 1$ with c as small as possible.

Then, the accessible T-torsion will be defined over small field extensions of $\mathbb{F}_{p^2}$. For instance, T could be the product of the smallest successive primes (ℓ excluded) such that $T \simeq p^3$. All the primes we need will be $O(\log(p))$, so the T-torsion will be defined on extensions of degree $O(\log(p))$ of $\mathbb{F}_{p^2}$ and all isogeny computations will be polynomial in $\log(p)$.

This ensures in particular that RigorousSQISignHD complexity scales polynomially with the security parameter λ (as FastSQISignHD), which is an interesting property in comparison with SQISign. Indeed, in SQISign [DLW22] the authors imposed $\ell^f | p-1$ and $T | p^2-1$ (with $T \simeq p^{5/4}$), so that the whole T-torsion is defined over $\mathbb{F}_{p^4}$ (and x-coordinates are defined over $\mathbb{F}_{p^2}$). Finding such primes has been an open research question since the introduction of SQISign [CMN21; BSC+22; Ahr23]. It is still unclear if we can still find T sufficiently smooth as the security level grows. Since computing a prime degree isogeny is exponential in the degree, the SQISign protocol might not be polynomial in $\log(p)$.

4.2 Challenge generation

To ensure a soundness security level of λ bits, the challenge space needs to have size at least $2^\lambda \simeq \sqrt{p}$. We also need the challenge degree D_φ to be coprime to ℓ to execute $\mathsf{EvalTorsion}_{\ell^f}$ during the signing procedure. The challenge generation procedure $\mathsf{Challenge}_{D_\varphi}$ is the same in the fast and provably secure challenge generation procedure. It simply generates a random element $P \in E_1$ of order D_φ and computes φ of kernel $\langle P \rangle$. Only D_φ changes. In FastSQISignHD, $D_\varphi = \ell'^{f'}$ and in RigorousSQISignHD, D_φ will be a divisor of T of size $D_\varphi \simeq p^3$. More details on the choice of D_φ in RigorousSQISignHD will be given in the security analysis (see Appendix A.6).

4.3 Fast key generation and commitment

We now present $\mathsf{FastDoublePath}$ (Algorithm 4) the main algorithmic block for the key generation and commitment of FastSQISignHD. The goal of this algorithm is to generate two isogeny paths $\phi, \phi' : E_0 \longrightarrow E$ of degree dividing $\ell^{2f} \simeq p$ and $\ell'^{2f'} \simeq p$ respectively, computing the kernel ideals I_ϕ and $I_{\phi'}$ along the way. This algorithm is directly applicable to the commitment procedure $\mathsf{FastCommit}$ where we need to generate a double path to be able to compute the challenge kernel ideal I_φ (with the ℓ-isogeny path of degree coprime to ℓ') and to apply the $\mathsf{EvalTorsion}_{\ell^f}$ procedure (with the ℓ'-isogeny path of degree coprime to ℓ).

For the key generation $\mathsf{FastKeyGen}$, we only need the ℓ'-isogeny path $\tau = \phi'$ but the algorithm is essentially the same except that we do not compute the ϕ and I_ϕ completely.

Note that generating isogenies of degree $\simeq p$ is essential for security reasons, in order to ensure that the codomain E is heuristically close to a random elliptic curve in the supersingular isogeny graph. To compute such long isogeny paths, however, we are limited by the accessible torsion in E_0 (we can access to the $\ell^f \ell'^{f'}$-torsion only). To circumvent this difficulty, we need to use pushforward isogenies.

Preliminary: pushing forward isogenies. We recall the notion of *pushforward isogeny* and *pushforward ideal* introduced in [DKLPW20].

Definition 4.3.1. We consider an isogeny diamond, as follows:

$$\begin{tikzcd}
E \arrow[r, "\theta"] \arrow[d, "\rho"] & F_1 \arrow[d, "\rho'"] \\
F_2 \arrow[r, "\theta'"] & F_3
\end{tikzcd}$$

with $\deg(\rho)$ coprime to $\deg(\theta)$. We call ρ' the *pushforward* of ρ via θ, also denoted by $\rho' = [\theta]_*\rho$. The isogeny ρ' satisfies $\ker(\rho') = \theta(\ker(\rho))$. Similarly, θ' is the pushforward of θ via ρ, denoted by $\theta' = [\rho]_*\theta$ and satisfies $\ker(\theta') = \rho(\ker(\theta))$.

If I and J are the ideals associated to ρ and θ respectively via the Deuring correspondence, we denote by $[J]_*I$ and $[I]_*J$ the *pushforward ideals* associated to $[\theta]_*\rho$ and $[\rho]_*\theta$ respectively.

The pushforward ideals can be computed explicitly with linear algebra.

Lemma 4.3.2. *[DKLPW20, Lemma 3] Let $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ be a maximal order and I, J be two left $\mathcal{O}$-ideals of coprime norms. Then, $[J]_*I = J^{-1} \cdot (I \cap J)$ and $[I]_*J = I^{-1} \cdot (I \cap J)$.*

The algorithm. The idea is to construct the isogenies ϕ and ϕ' (of degree ℓ^{2f} and $\ell'^{2f'}$ respectively) by finding an endomorphism γ of degree $\ell^{2f}\ell'^{2f'}$, and factoring it as $\gamma = \hat{\phi}' \circ \phi$. Since $\ell^{2f}\ell'^{2f'} = \Theta(p^2) = \omega(p)$, we can easily find $\gamma \in \mathcal{O}_0$ non divisible by ℓ or ℓ', of norm $\mathrm{nrd}(\gamma) = \ell^{2g}\ell'^{2g'}$ with $g \leq f$ close to f and $g' \leq f'$ close to f', using [Ler22, Algorithm 4].

Since ℓ^{2f} (and $\ell'^{2f'}$) exceeds the available torsion, some "pushforward gymnastics" is required to compute the factorisation. We thus decompose $\varepsilon(\gamma) = \widehat{\rho_2} \circ \rho_1$ where ρ_1 and ρ_2 are isogenies $E_0 \longrightarrow E'$ of degree $\ell^g\ell'^{g'}$ and ε is an isomorphism $\mathcal{O}_0 \xrightarrow{\sim} \mathrm{End}(E_0)$. According to the following lemma, ρ_1, ρ_2 and their associated ideals K_1 and K_2 respectively are given as follows:

$$\ker(\rho_1) = \ker(\varepsilon(\gamma)) \cap E_0[\ell^g\ell'^{g'}], \quad \ker(\rho_2) = \ker(\widehat{\varepsilon(\gamma)}) \cap E_0[\ell^g\ell'^{g'}],$$

$$K_1 = \mathcal{O}_0\gamma + \mathcal{O}_0\ell^g\ell'^{g'} \quad \text{and} \quad K_2 = \mathcal{O}_0\overline{\gamma} + \mathcal{O}_0\ell^g\ell'^{g'},$$

since $\varepsilon(\gamma)$ is cyclic (γ being non-divisible by ℓ or ℓ').

Lemma 4.3.3. *Let $\rho : E \longrightarrow E'$ be a cyclic isogeny decomposed into $\rho = \theta \circ \rho_1$. Then we have:*

(i) $\ker(\rho_1) = \ker(\rho) \cap E[d_1]$ *with* $d_1 := \deg(\rho_1)$.

(ii) *If ρ is a cyclic endomorphism ($E = E'$), then the kernel ideal of ρ_1 is $K_1 = \mathcal{O}\rho + \mathcal{O}d_1$, where $\mathcal{O} := \mathrm{End}(E)$.*

Proof. Since $\rho = \theta \circ \rho_1$ and $\deg(\rho_1) = d_1$, we clearly have $\ker(\rho_1) \subseteq \ker(\rho) \cap E[d_1]$. Since ρ is cyclic, there exists a generator $P \in E$ of $\ker(\rho)$ of order $d := \deg(\rho)$ and we have

$$\ker(\rho) \cap E[d_1] = \langle [d/d_1]P \rangle,$$

where $[d/d_1]P$ has order d_1, so we conclude that the inclusion is an equality by cardinality, since ρ_1 is separable. (i) follows.

To prove (ii), we remark that

$$E[\mathcal{O}\rho + \mathcal{O}d_1] = E[\rho] \cap E[d_1] = \ker(\rho_1),$$

where the last equality was proved in (i). Then, we conclude that $K_1 = \mathcal{O}\rho + \mathcal{O}d_1$ by injectivity of the Deuring correspondence between left $\mathcal{O}$-ideals and isogenies of domain E [Voi20, Proposition 42.2.16]. This completes the proof. □

Then, we can decompose ρ_1 and ρ_2 into $\rho_1 = \widehat{\theta_1'} \circ \theta_1$ and $\rho_2 = \widehat{\theta_2} \circ \theta_2'$ where the θ_i are isogenies of degree ℓ^g and the θ_i' are isogenies of degree $\ell'^{g'}$ for $i \in \{1, 2\}$, as in the following diagram:

$$\begin{tikzcd}
 & F_2 \arrow[d, "{[\theta_2]_*\theta_1'}"] & \\
E_0 \arrow[ur, "\theta_2'"] \arrow[dr, "\theta_1"'] & E & E' \arrow[ul, "\theta_2"'] \arrow[dl, "\theta_1'"] \\
 & F_1 \arrow[u, "{[\theta_1']_*\theta_2}"'] &
\end{tikzcd}$$

The pushforward isogenies $[\theta_1']_*\theta_2$ and $[\theta_2]_*\theta_1'$ have the same codomain E and degree ℓ^g and $\ell'^{g'}$ respectively. Hence, $\phi := [\theta_1']_*\theta_2 \circ \theta_1$ and $\phi' := [\theta_2]_*\theta_1' \circ \theta_2'$ are isogenies $E \longrightarrow E_0$ of desired degrees ℓ^{2g} and $\ell^{2g'}$ respectively. By Lemma 4.3.3, we can compute $\ker(\theta_1) = \ker(\varepsilon(\gamma)) \cap E_0[\ell^g]$, $\ker(\theta_2') = \ker(\widehat{\varepsilon(\gamma)}) \cap E_0[\ell'^{g'}]$, $\ker(\theta_1') = \ker(\widehat{\rho_1}) \cap E'[\ell'^{g'}]$ and $\ker(\theta_2) = \ker(\widehat{\rho_2}) \cap E'[\ell^g]$, and obtain the θ_i and θ_i' with Vélu's formulas. We then compute $\ker([\theta_1']_*\theta_2) = \theta_1'(\ker(\theta_2))$ and $\ker([\theta_2]_*\theta_1') = \theta_2(\ker(\theta_1'))$ and use Vélus formulas. We then easily get ϕ and ϕ'.

Since $\gamma = \hat{\phi}' \circ \phi$, Lemma 4.3.3 implies that the ideals $J := \mathcal{O}_0\gamma + \mathcal{O}_0\ell^{2g}$ and $J' := \mathcal{O}_0\overline{\gamma} + \mathcal{O}_0\ell'^{2g'}$. The algorithm is summarised in Algorithm 4.

Remark 4.1. For FastKeyGen only ϕ' and J' are necessary, so we use a slightly modified version of Algorithm 4 where $\mathcal{H}_1$ (line 4), θ_1 (line 5), ϕ (line 7), and J (line 8) are not computed.

4.4 Provably secure key generation and commitment

To prove security, the distribution of the public key E_A and the commitment E_1 need to be close to the uniform distribution in the supersingular isogeny graph. Heuristically, we expect to get an elliptic curve with a distribution somewhat close to uniform after a random isogeny walk of length $\Theta(p)$. However, the fact that we compute two paths simultaneously in FastCommit and FastKeyGen might alter the distribution of the resulting elliptic curve. While we do not expect the induced bias to be computationally relevant, it prevents a rigorous analysis. This is the reason why we now propose a different procedure for RigorousKeyGen and RigorousCommit.

Starting from E_0, we generate a random ℓ-isogeny walk $\phi : E_0 \longrightarrow E$ long enough to make E uniformly random and compute its kernel ideal I_ϕ as well as an alternate path $\phi' : E_0 \longrightarrow E$ of degree dividing T^2. This algorithm is very similar to what is done in the first version of SQISign signing algorithm [DKLPW20]. The procedure we shall describe is costly (though polynomial) and this is one of the reasons why we do not expect the efficiency of RigorousSQISignHD to compare favourably to the original SQISign protocol.

Algorithm 4: $\mathsf{FastDoublePath}_{\ell^f,\ell'^{f'}}$

```
Data: A basis of $\mathcal{O}_0$ and an isomorphism $\varepsilon : \mathcal{O}_0 \xrightarrow{\sim} \mathrm{End}(E_0)$.
Result: Two cyclic isogenies $\phi : E_0 \longrightarrow E$ of degree dividing $\ell^{2f}$ and $\phi' : E_0 \longrightarrow E$ of degree dividing $\ell'^{2f'}$ and their respective kernel ideals $J$ and $J'$.
Use [Ler22, Algorithm 4] to find $\gamma \in \mathcal{O}_0$ non divisible by $\ell$ and $\ell'$ of norm $\mathrm{nrd}(\gamma) = \ell^{2g}\ell'^{2g'}$ with $g \leq f$ close to $f$ and $g' \leq f'$ close to $f'$;
Evaluate $\varepsilon(\gamma)$ and $\varepsilon(\overline{\gamma})$ on a basis of $E_0[\ell^g\ell'^{g'}]$ and solve discrete logarithm problems to compute $\mathcal{G}_1 := \ker(\varepsilon(\gamma)) \cap E_0[\ell^g\ell'^{g'}]$ and $\mathcal{G}_2 := \ker(\widehat{\varepsilon(\gamma)}) \cap E_0[\ell^g\ell'^{g'}]$;
Compute $\rho_i : E_0 \longrightarrow E'$ of kernel $\mathcal{G}_i$ for $i = 1, 2$;
Compute $\mathcal{H}_1 := \ker(\varepsilon(\gamma)) \cap E_0[\ell^g]$, $\mathcal{H}'_2 := \ker(\widehat{\varepsilon(\gamma)}) \cap E_0[\ell'^{g'}]$, $\mathcal{H}'_1 := \ker(\widehat{\rho_1}) \cap E'[\ell'^{g'}]$ and $\mathcal{H}_2 := \ker(\widehat{\rho_2}) \cap E'[\ell^g]$;
Compute $\theta_i$ of kernel $\mathcal{H}_i$ and $\theta'_i$ of kernel $\mathcal{H}'_i$ for $i = 1, 2$;
Compute $[\theta'_1]_*\theta_2$ and $[\theta_2]_*\theta'_1$ of kernels $\theta'_1(\ker(\theta_2))$ and $\theta_2(\ker(\theta'_1))$ respectively;
Let $\phi := [\theta'_1]_*\theta_2 \circ \theta_1$ and $\phi' := [\theta_2]_*\theta'_1 \circ \theta'_2$;
Let $J := \mathcal{O}_0\gamma + \mathcal{O}_0\ell^{2g}$ and $J' := \mathcal{O}_0\overline{\gamma} + \mathcal{O}_0\ell'^{2g'}$;
Return $\phi, \phi', J, J'$;
```

A long enough supersingular ℓ-isogeny walk.

Proposition 4.4.1. *Let $\phi : E_0 \longrightarrow E$ be an ℓ^h-isogeny obtained from a non-backtracking random ℓ-isogeny walk. Then, for all $\varepsilon \in]0, 2]$, the distribution of E has statistical distance $\tilde{O}(p^{-\varepsilon/2})$ to the uniform distribution in the supersingular isogeny graph, provided that $h \geq (1 + \varepsilon)\log_\ell(p)$.*

Proof. Let $\mathrm{SS}(p)$ be the set of supersingular elliptic curves over $\mathbb{F}_{p^2}$ (up to isomorphism) and S be the probability distribution on $\mathrm{SS}(p)$ given by $S(E) := K^{-1}/\#\,\mathrm{Aut}(E)$ for all $E \in \mathrm{SS}(p)$, with $K := \sum_{E\in\mathrm{SS}(p)} 1/\#\,\mathrm{Aut}(E)$. Let δ_0 be the Dirac distribution on E_0 and $\delta_0^{(h)}$ the distribution obtained from δ_0 after a non-backtracking ℓ-isogeny walk of length h. By [BCC+22, Theorem 11], the statistical distance between S and $\delta_0^{(h)}$ satisfies

$$d_{TV}\left(S, \delta_0^{(h)}\right) := \frac{1}{2}\sum_{E\in\mathrm{SS}(p)} \left|S(E) - \delta_0^{(h)}(E)\right| \leq \frac{\sqrt{6K}}{2}\frac{(\ell+1)(h+1) - 2}{(\ell+1)\sqrt{\ell^h}}.$$

By Eichler's formula [Voi20, p. 42.3.8], we know that $K = (p-1)/24$. Then, when $h \geq (1+\varepsilon)\log_\ell(p)$, we get that $d_{TV}(S, \delta_0^{(h)}) = O(\log(p)p^{-\varepsilon/2}) = \tilde{O}(p^{-\varepsilon/2})$.

By Lemma A.7.4, we have $d_{TV}(U, S) = O(p^{-1})$, so we finally get, by triangular inequality, that $d_{TV}(U, \delta_0^{(h)}) = \tilde{O}(p^{-\varepsilon/2})$. □

By the above proposition (with $\varepsilon = 1$), an adversary needs time $\tilde{O}(\sqrt{p})$ to distinguish a supersingular elliptic curve obtained from E_0 after a non-backtracking

Algorithm 5: $\mathsf{RigorousDoublePath}_{\ell^f,T}$

```
Data: A supersingular elliptic curve $E_0$ with accessible $T$-torsion and $\ell^f$-
      torsion, an eval-basis $\mathcal{B}_0$ of $\mathrm{End}(E_0)$ in the sense of Definition 2.3.1.
Result: A random $\ell$-isogeny walk $\phi : E_0 \longrightarrow E$ of degree $\ell^h = \Theta(p^2)$ such that
        the distribution of $E$ has statistical distance $\tilde{O}(p^{-1/2})$ to the uniform,
        an isogeny $\phi' : E_0 \longrightarrow E$ of degree dividing $T^2$ and their respective
        kernel ideals $I_\phi$ and $I_{\phi'}$.
$h \longleftarrow \lceil 2\log_\ell(p) \rceil$;
Perform a random $\ell$-isogeny walk $\phi : E_0 \longrightarrow E$ of degree $\ell^h$;
Factor $\phi$ by a power of $[\ell]$ if necessary to make it cyclic;
Decompose $\phi$ in a sequence of isogenies $\phi_i : E_i \longrightarrow E_{i+1}$ $(1 \leq i \leq r)$ of degree
  dividing $\ell^f$;
Let $P_i$ be a generator of $\ker(\phi_i)$ for all $i \in [\![1 \; ; \; r]\!]$;
$J_0 \longleftarrow \mathcal{O}_0$;
for $i := 0$ to $r - 1$ do
    $I_i \longleftarrow \mathsf{KernelToIdeal}_{\ell^f}(\mathcal{B}_i, P_i)$;
    $J_{i+1} \longleftarrow \mathsf{KLPT}_{T^2}(J_i \cdot I_i)$;
    $\phi'_{i+1} \longleftarrow \mathsf{SpecialIdealToIsogeny}(J_{i+1}, I_0 \cdots I_i, \phi_i \circ \cdots \circ \phi_0)$;
    Compute $\mathcal{B}_{i+1} := \mathsf{PushEndRing}(\mathcal{B}_0, \phi'_{i+1}, J_{i+1})$, a $T$-eval-basis of
      $\mathrm{End}(E_{i+1})$;
end
$\phi' \longleftarrow \phi_r, I_\phi \longleftarrow I_0 \cdots I_{r-1}, I_{\phi'} \longleftarrow J_r$;
Return $\phi, \phi', I_\phi, I_{\phi'}$;
```

ℓ-isogeny path of degree $\Theta(p^2)$ from a uniformly random elliptic curve in the supersingular isogeny graph. Since $p = \Theta(2^\lambda)$, this is sufficient to ensure a security level of λ bits.

Computing the kernel ideal and an alternate path. Assume that we have generated a random ℓ-isogeny walk $\phi : E_0 \longrightarrow E$ of degree $\ell^h = \Theta(p^2)$. To compute the kernel ideal I_ϕ, we cannot use $\mathsf{KernelToIdeal}_{\ell^h}$ (Algorithm 2) because we would need accessible ℓ^h-torsion, which is impossible since $\ell^h \gg p$.

Instead, as in [DKLPW20], we divide ϕ into a sequence of isogenies $\phi_i : E_i \longrightarrow E_{i+1}$ $(1 \leq i \leq r)$ of degree dividing ℓ^f and compute their associated kernel ideals I_i successively. For the computation of I_i, we need to compute an alternate isogeny $\phi'_i : E_0 \longrightarrow E_i$ of degree dividing T^2 obtained at step $i - 1$. Hence, the alternate path $\phi' := \phi_r : E_0 \longrightarrow E_r = E$ will be a convenient byproduct of our ideal computation. The algorithm we propose (see Algorithm 5) uses the sub-algorithms KLPT_{T^2} and $\mathsf{SpecialIdealToIsogeny}$ as black boxes (see Section 2.3).

Application to key generation and commitment. For the commitment phase $\mathsf{RigorousCommit}$, we simply call $\mathsf{RigorousDoublePath}_{\ell^f,T}$ to output two isogenies $\psi', \psi : E_0 \longrightarrow E_1$ with codomain E_1 statistically close to uniform and respective degrees ℓ^h and dividing T^2, along with their kernel ideals I_ψ and $I_{\psi'}$.

For key generation $\mathsf{RigorousKeyGen}$, we only need an isogeny $\tau : E_0 \longrightarrow E_A$ of degree T^2 (coprime to ℓ) with codomain E_A statistically close to uniform. $\mathsf{RigorousDoublePath}_{\ell^f,T}$ will output $\phi, \tau : E_0 \longrightarrow E_1$ of respective degrees ℓ^h and $D_\tau | T^2$, along with I_ϕ, I_τ. The data (ϕ, I_ϕ) will not be used so we only compute ϕ to ensure the randomness of E_A and I_ϕ as an intermediary tool to obtain (τ, I_τ).

5 Response and verification

The goal of this section is to present a precise description of the algorithmic building blocks required by our new signature scheme. Once again, there are two versions of our algorithms : the rigorous ones in dimension 8 and the faster variants in dimension4.

Throughout this section, we assume that the prover has generated a secret key $\tau : E_0 \longrightarrow E_A$ of degree D_τ coprime to ℓ and two paths to the commitment $\psi, \psi' : E_0 \longrightarrow E_1$ of respective degrees D_ψ coprime to ℓ and $D_{\psi'}$ a power of ℓ. We also assume that the prover has access to the challenge $\varphi : E_1 \longrightarrow E_2$ of degree D_φ coprime to ℓ.

5.1 Overview of the response computation

In this section, we present the algorithm $\mathsf{FastRespond}$ used to compute the response in the FastSQISignHD identification protocol (in dimension 4) and its alternate provably secure version $\mathsf{RigorousRespond}$ (in dimension 8) used in RigorousSQISignHD. We also provide their verification counterparts $\mathsf{FastVerify}$ and $\mathsf{RigorousVerify}$.

Those algorithms use the following sub-algorithms that will be introduced in this section (if not already):

- $\mathsf{IsogenyToIdeal}(\varphi, \psi, I_\psi)$ (presented in Section 2.3) takes as input en eval basis of $\mathrm{End}(E_0)$, an isogeny $\varphi : E_1 \longrightarrow E_2$ of degree D_φ, an isogeny $\psi : E_0 \longrightarrow E_1$ of degree coprime to D_φ, its ideal $I_\psi \subset \mathcal{O}_0$ and returns the kernel ideal I_φ of φ.
- $\mathsf{RandomEquivalentIdeal}_{\ell^e}$ takes as input an $\mathcal{O}_0$-left ideal J and returns an equivalent ideal I that is uniformly random among ideals of norm $\leq \ell^e$.
- $\mathsf{EvalTorsion}_{\ell^f}$ evaluates a non-smooth degree isogeny on ℓ^f-torsion points knowing its kernel ideal and an alternate smooth degree path. Namely, it takes as input an ideal I connecting $\mathcal{O} \cong \mathrm{End}(E)$ and $\mathcal{O}' \cong \mathrm{End}(E')$, a basis (P_1, P_2) of $E[\ell^f]$ two isogenies $\rho_1 : E_0 \longrightarrow E$ and $\rho_2 : E_0 \longrightarrow E'$ of smooth degree coprime to ℓ, with their respective kernel ideals I_1 and I_2 and returns $(\phi_I(P_1), \phi_I(P_2))$, where $\phi_I : E \longrightarrow E'$ is the isogeny associated to I.
- $\mathsf{RepresentIsogeny}_{4,\ell^e,\ell^f}$ takes as input an ℓ^e-good integer q, integers a_1, a_2 such that $a_1^2 + a_2^2 + q = \ell^e$, a basis (P_1, P_2) of $E_A[\ell^f]$, $(\sigma(P_1), \sigma(P_2))$, where $\sigma : E_A \longrightarrow E_2$ is a q-isogeny, and returns a chain of 4-dimensional ℓ-isogenies whose composition is $F(\sigma, a_1, a_2)$ as in Notation 3.3.1.

- $\mathsf{RepresentIsogeny}_{8,\ell^e,\ell^f}$ takes as input an integer $q < \ell^e$ coprime to ℓ, integers $a_1, \cdots, a_4$ such that $a_1^2 + \cdots + a_4^2 + q = \ell^e$, a basis (P_1, P_2) of $E_A[\ell^f]$, $(\sigma(P_1), \sigma(P_2))$, where $\sigma : E_A \longrightarrow E_2$ is a q-isogeny, and returns a chain of 8-dimensional ℓ-isogenies whose composition is $F(\sigma, a_1, \cdots, a_4)$ as in Notation 3.3.3.
- For $g \in \{4, 8\}$, $\mathsf{IsValid}_g$, with input $F, E_A, E_2, \ell^e, \ell^f$, checks if F is a valid output of $\mathsf{RepresentIsogeny}_{g,\ell^e,\ell^f}$ representing an isogeny $\sigma : E_A \longrightarrow E_2$ in dimension g.

In both versions of SQISignHD, the prover sends the image of two points P_1, P_2 forming a basis of $E_A[\ell^f]$ by σ and its degree q. In dimension 4 (respectively dimension 8), the prover can then use q to compute a_1, a_2 (resp. $a_1, \cdots, a_4$) and compute $F(\sigma, a_1, a_2)$ (resp. $F(\sigma, a_1, \cdots, a_4)$) with the $\mathsf{RepresentIsogeny}_{g,\ell^e,\ell^f}$ procedure. If the computation succeeds and is validated by the $\mathsf{IsValid}_g$ procedure, then the verification is complete. Algorithms 8 and 9 follow.

Remark 5.1 (On the ℓ^f-torsion basis). It is sufficient to send the data $(\sigma(P_1), \sigma(P_1), q)$ to the verifier as the basis (P_1, P_2) can be computed canonically knowing E_A by classical compression techniques developed for SIDH [AJKKL16; ZSPDB18]. This decreases the communications size at a small computational cost. Later, with the compression/decompression algorithms ((see Lines 15 and 16)), we will see how to further compress this data.

Note that we use a basis of the ℓ^f-torsion with $2f \geq e + 4$ here because we might not have ℓ^e-torsion accessible. We can still compute F with this partial information as explained in Section 5.4.

To sign in dimension 4, the prover starts by computing an ideal $I \sim \overline{I_\tau} \cdot I_\psi \cdot I_\varphi$ connecting $\mathcal{O}_A \cong \mathrm{End}(E_A)$ to $\mathcal{O}_2 \cong \mathrm{End}(E_2)$ of norm ℓ^e-good norm q and coprime to ℓ' with uniform distribution using $\mathsf{RandomEquivalentIdeal}_{\ell^e}$. The coprimality with ℓ' is justified by security reasons (see Section 6.1). Then, the prover generates the basis (P_1, P_2) of $E_A[\ell^f]$ canonically and evaluates σ on it with $\mathsf{EvalTorsion}_{\ell^f}$ using I (kernel ideal of σ) and the paths $\tau : E_0 \longrightarrow E_A$ and $\varphi \circ \psi : E_0 \longrightarrow E_2$ of degree coprime to ℓ. The procedure is the same in dimension 8, except that we only require $q < \ell^e$ instead of ℓ^e-good and coprime to ℓ'. Algorithms 6 and 7 follow.

Remark 5.2 (On the case of dimension 8). Since we can no longer guarantee that q is coprime to ℓ in dimension 8, the prover might need to factor σ by ℓ-isogenies before computing $\sigma(P_1)$ and $\sigma(P_1)$. Details may be found in Appendix B. To preserve the clarity of exposition in this section, we assume that q is coprime to ℓ even in dimension 8.

As input of Algorithms 6, 7, 8 and 9, we denote by:

- $\mathsf{FastSetup}$, the public parameters of FastSQISignHD, $p = c\ell^f\ell'^{f'} - 1$, ℓ, ℓ', f, f', the exponent e and the elliptic curve $E_0/\mathbb{F}_p$;
- $\mathsf{RigorousSetup}$, the public parameters of RigorousSQISignHD, p, ℓ and f such that $\ell^f | p - 1$, e, T powersmooth (accessible torsion) and E_0;

- SecretKey, the isogeny $\tau : E_0 \longrightarrow E_A$ of degree D_τ and its kernel ideal I_τ;
- CommitData, the two isogenies $\psi, \psi' : E_0 \longrightarrow E_1$ of degrees D_ψ and $D_{\psi'}$ and their respective kernel ideals I_ψ and $I_{\psi'}$;
- ChallData, the isogeny $\varphi : E_1 \longrightarrow E_2$ of degree D_φ.

where D_τ, D_ψ and D_φ are powers of ℓ and $D_{\psi'}$ is a power of ℓ' or a divisor of T^2 depending on the version of SQISignHD (see Section 4.1).

Algorithm 6: FastRespond

```
Data: FastSetup, SecretKey, CommitData and ChallData.
Result: $(\sigma(P_1), \sigma(P_2), q)$, where $(P_1, P_2)$ is a canonically determined basis of
        $E_A[\ell^f]$ and $\sigma : E_A \longrightarrow E_2$ is an isogeny of $\ell^e$-good degree $q$ coprime
        to $\ell'$.
1 $I_\varphi \longleftarrow$ IsogenyToIdeal$(\varphi, \psi', I_{\psi'})$;
2 $J \longleftarrow \overline{I_\tau} \cdot I_\psi \cdot I_\varphi$;
3 $I \longleftarrow$ RandomEquivalentIdeal$_{\ell^e}(J)$ and $q \longleftarrow \mathrm{nrd}(I)$;
4 If $q$ is not $\ell^e$-good or $q \wedge \ell' \neq 1$, go back to line 3;
5 Compute the canonical basis $(P_1, P_2)$ of $E_A[\ell^f]$;
6 $(\sigma(P_1), \sigma(P_2)) \longleftarrow$ EvalTorsion$_{\ell^f}(I, P_1, P_2, \tau, \varphi \circ \psi, I_\tau, I_\psi \cdot I_\varphi)$;
7 Return $(\sigma(P_1), \sigma(P_2), q)$;
```

Algorithm 7: RigorousRespond

```
Data: RigorousSetup, SecretKey, CommitData and ChallData.
Result: $(\sigma(P_1), \sigma(P_2), q)$, where $(P_1, P_2)$ is a canonically determined basis of
        $E_A[\ell^f]$ and $\sigma : E_A \longrightarrow E_2$ is a $q$-isogeny with $q < \ell^e$.
$I_\varphi \longleftarrow$ IsogenyToIdeal$(\varphi, \psi', I_{\psi'})$;
$J \longleftarrow \overline{I_\tau} \cdot I_\psi \cdot I_\varphi$;
$I \longleftarrow$ RandomEquivalentIdeal$_{\ell^e}(J)$ and $q \longleftarrow \mathrm{nrd}(I)$;
Compute the canonical basis $(P_1, P_2)$ of $E_A[\ell^f]$;
$(\sigma(P_1), \sigma(P_2)) \longleftarrow$ EvalTorsion$_{\ell^f}(I, P_1, P_2, \tau, \varphi \circ \psi, I_\tau, I_\psi \cdot I_\varphi)$;
Return $(\sigma(P_1), \sigma(P_2), q)$;
```

5.2 Finding a uniformly random tight response ideal

In this section, we present the algorithm RandomEquivalentIdeal$_{\ell^e}$ Antonin: Dans la description faite à la section précédente, l'algorithme RandomEquivalentIdeal$_{\ell^e}$ est utilisé à la fois pour la dimension 4 et 8, pourtant il me semble que l'algorithme décrit ici est pour la dimension 8. Je ne suis pas certain (à verifier) que pour la dimension 4 on ait envie d'utiliser cet algo avec du rejection sampling jusqu'à

Algorithm 8: FastVerify

```
Data: FastSetup, $E_A, E_2$ and an output $R$ from FastRespond.
Result: Determines if $R$ is a valid response.
Try to parse $R := (R_1, R_2, q)$, where $R_1, R_2 \in E_2[\ell^f]$ and $q < \ell^e$ and return False if it fails;
If $q$ is not $\ell^e$-good or $q \wedge \ell' \neq 1$, return False;
Compute the canonical basis $(P_1, P_2)$ of $E_A[\ell^f]$;
Find $a_1, a_2 \in \mathbb{Z}$ such that $a_1^2 + a_2^2 = \ell^e - q$ using Cornacchia's algorithm [Cor08];
$F \longleftarrow \mathsf{RepresentIsogeny}_{4,\ell^e,\ell^f}(E_A, E_2, a_1, a_2, P_1, P_2, R_1, R_2)$;
Return $\mathsf{IsValid}_{4,\ell^e,\ell^f\ell'^{f'}}(F, E_A, E_2, \ell^e, \ell^f)$;
```

Algorithm 9: RigorousVerify

```
Data: RigorousSetup, $E_A, E_2$ and an output $R$ from RigorousRespond.
Result: Determines if $R$ is a valid response.
Try to parse $R := (R_1, R_2, q)$, where $R_1, R_2 \in E_2[\ell^f]$ and $q < \ell^e$ and return False if it fails;
Compute the canonical basis $(P_1, P_2)$ of $E_A[\ell^f]$;
Find $a_1, \cdots, a_4 \in \mathbb{Z}$ such that $a_1^2 + \cdots + a_4^2 + q = \ell^e$ using Pollack and Treviño's algorithm [PT18];
$F \longleftarrow \mathsf{RepresentIsogeny}_{8,\ell^e,\ell^f}(E_A, E_2, a_1, \cdots, a_4, P_1, P_2, R_1, R_2)$;
Return $\mathsf{IsValid}_{8,\ell^e,T}(F, E_A, E_2, \ell^e, \ell^f)$;
```

ce que la norme soit ℓ^e-good. Comme on vise de l'heuristique un algo plus simple de type : on prends des combinaisons linéaires d'une bonne base, suffit sans doute et sera beaucoup plus efficace. taking a left $\mathcal{O}_0$-ideal J as input and returning an ideal I which is uniformly random among the ideals $I \sim J$ of norm $q < \ell^e$. By [DKLPW20, Lemma 1], all the equivalent ideals $I \sim J$ are of the form $\chi_J(\alpha) := J\overline{\alpha}/\operatorname{nrd}(J)$ for some $\alpha \in I$ and α determines I up to multiplication by an element of $\mathcal{O}_0^\times$. Besides, the norm of $I = \chi_J(\alpha)$ is $q_J(\alpha) := \operatorname{nrd}(\alpha)/\operatorname{nrd}(J)$, so we need $q_J(\alpha) \leq \ell^e$.

Hence, to sample an ideal $I \sim J$ such that $\operatorname{nrd}(I) \leq \ell^e$ with uniform distribution is equivalent to sample $\alpha \in J \setminus \{0\}$ such that $q_J(\alpha) \leq \ell^e$ with uniform distribution. If we fix a basis of J, we can see q_J as a primitive positive definite integral quadratic form with four variables. By the following lemma, which is a simple generalization of [Wes22, Lemma 3.3], we can sample uniformly $\alpha \in J$ such that $q_J(\alpha) \leq \ell^e$. $\mathsf{RandomEquivalentIdeal}_{\ell^e}$ calls this procedure to get $\alpha \in J$ uniform and rejects the result if $\alpha = 0$. Then the distribution of α is still uniform but in $J \setminus \{0\}$.

Lemma 5.2.1. *Let f be a primitive positive definite integral quadratic form in k variables and let $\rho > 0$. Then there exists an algorithm that samples uniformly random elements from the set*

$$\{x \in \mathbb{Z}^k \mid f(x) \leq \rho\}$$

in polynomial time in $\log(\rho)$ *and the length of* f *(namely, the maximal number of bits of the coefficients of* f*). This algorithm runs in exponential time in* k*.*

Proof. See Appendix A.2. □

For $\mathsf{RandomEquivalentIdeal}_{\ell^e}(J)$ to terminate, we need to find $\alpha \in J \setminus \{0\}$ such that $q_J(\alpha) \leq \ell^e$. For such an α to exist, we need $\ell^e = \Omega(\sqrt{p})$ according to the following lemma (Lemma 5.2.2). In RigorousSQISignHD, $\ell^e = \Theta(p^2)$ so this condition is obviously satisfied.

Lemma 5.2.2. *Let* $\mathcal{O}$ *be a maximal order and* J *be a left* $\mathcal{O}$*-ideal. Then there exists* $\alpha \in J$ *such that* $q_J(\alpha) \leq 2\sqrt{2p}/\pi$*.*

Proof. See Appendix A.2. □

In FastSQISignHD, we need to find an ℓ^e-good value of $q_J(\alpha)$. Heuristically, assuming that $q_J(\alpha)$ behaves like a random integer, we should expect to find a suitable $\alpha \in J$ with probability $O(1/\log(p))$. Hence, taking ℓ^e a few bits over $\sqrt{p}$ might be sufficient. For that reason, in our choice of parameters, we only have accessible ℓ^f-torsion with $\ell^f < \sqrt{p} < \ell^e$ (see Section 4.1). Proving formally that we can always find an ℓ^e-good value of $q_J(\alpha)$ would certainly require to increase ℓ^e by a lot. As [RT22] indicates, we should expect lower bounds close to $\ell^e = \omega(p^2)$, causing a huge efficiency loss.

5.3 The isogeny torsion evaluation algorithm

We present $\mathsf{EvalTorsion}_{\ell^f}$ that evaluates a non-smooth degree isogeny on ℓ^f-torsion points knowing its kernel ideal and an alternate smooth degree path. Let I be an ideal connecting $\mathcal{O} \cong \mathrm{End}(E)$ and $\mathcal{O}' \cong \mathrm{End}(E')$ of non-smooth norm q, (P_1, P_2) be a basis of $E[\ell^f]$, $\rho_1 : E_0 \longrightarrow E$ and $\rho_2 : E_0 \longrightarrow E'$, be two isogenies of respective degrees d_1, d_2 coprime to ℓ, and respective kernel ideals I_1 and I_2. We want to compute $(\phi_I(P_1), \phi_I(P_2))$, where $\phi_I : E \longrightarrow E'$ is the isogeny associated to I.

Let us consider the endomorphism $\gamma := \widehat{\rho_2} \circ \phi_I \circ \rho_1$ of E_0. From that definition of γ comes the equality

$$[d_1 d_2]\phi_I = \rho_2 \circ \gamma \circ \widehat{\rho_1}$$

Since ℓ is coprime to d_1 and d_2, the scalar $d_1 d_2$ can be inverted modulo ℓ^f and we see that it suffices to evaluate $\gamma, \widehat{\rho_1}, \rho_2$ on the ℓ^e-torsion of their respective domains.

The curve E_0 is chosen to have a known endomorphism ring so we can easily evaluate γ at any point from a basis of endomorphisms if we know the principal ideal $\mathcal{O}_0\gamma$. This ideal can be computed from the ideals $I_1, I_2 \subset \mathcal{O}_0$ and $I \subset \mathcal{O}$ associated to ρ_1, ρ_2 and ϕ_I respectively, with the formula $I_1 \cdot I \cdot \overline{I_2} = \mathcal{O}_0\gamma$. The $\mathsf{EvalTorsion}_{\ell^f}$ algorithm summarizes the procedure described above.

Antonin: cet algorithme est essentiellement équivalent à des algorithmes qui existent déjà dans la littérature, il faudrait les citer (Pierrick : je peux m'occuper de le faire, je mets juste ce commentaire pour m'en souvenir)

Algorithm 10: EvalTorsion$_{\ell^f}$

```
Data: A basis $(P_1, P_2)$ of $E[\ell^f]$, an ideal $I$ connecting $\mathcal{O} \cong \mathrm{End}(E)$ and $\mathcal{O}' := \mathrm{End}(E')$, two isogenies $\rho_1 : E_0 \longrightarrow E$ and $\rho_2 : E_0 \longrightarrow E'$ of respective degrees $d_1$ and $d_2$ coprime to $\ell$ and their respective kernel ideals $I_1$ and $I_2$.
Result: $(\phi_I(P_1), \phi_I(P_2))$, where $\phi_I : E \longrightarrow E'$ is the isogeny associated to $I$.
  Find $\gamma \in \mathcal{O}_0$ such that $\mathcal{O}_0\gamma = I_1 \cdot I \cdot \overline{I_2}$;
  $R_i \leftarrow \rho_2 \circ \gamma \circ \widehat{\rho_1}(P_i)$ for $i \in \{1, 2\}$;
  Compute $\lambda$, an inverse of $d_1 d_2$ modulo $\ell^f$;
  Return $([\lambda]R_1, [\lambda]R_2)$;
```

5.4 Dividing the higher dimensional isogeny computation in two

As explained in Section 5.2, we do not necessarily have enough accessible torsion to compute the whole kernel of the signature higher dimensional representation of the response F. In this section, we explain in plain generality how to circumvent this difficulty. Let us keep the notations of Section 3.2. Recall that we have the following isogeny diamond

$$\begin{tikzcd}
A' \arrow[r, "\varphi'"] & B' \\
A \arrow[u, "\psi"] \arrow[r, "\varphi"] & B \arrow[u, "\psi'"']
\end{tikzcd}$$

and that

$$F := \begin{pmatrix} \varphi & \widetilde{\psi'} \\ -\psi & \widetilde{\varphi'} \end{pmatrix}, \quad \text{with} \quad \ker(F) = \{(\widetilde{\varphi}(x), \psi'(x)) \mid x \in B[d]\}.$$

To compute F, we need to evaluate $\widetilde{\varphi}$ and ψ' on $B[d]$, so we need to have accessible d-torsion. However, we assume that we only have d'-accessible torsion with $d'|d$.

The idea is to decompose $F = F_2 \circ F_1$ where $F_1 : \mathcal{A} := A \times B' \longrightarrow \mathcal{C}$ and $F_2 : \mathcal{C} \longrightarrow \mathcal{B} := B \times A'$ are respectively d_1 and d_2-isogenies such that $d_1, d_2 | d'$ and to use the following proposition (proved in Appendix A.3) to compute F_1 and $\widetilde{F_2}$ to infer F.

Proposition 5.4.1. *Suppose d coprime to p so that F is separable. Then:*

(i) *We can always decompose $F = F_2 \circ F_1$, as above.*
(ii) $\ker(F_1) \subseteq \ker(F) \cap \mathcal{A}[d_1]$.
(iii) $\ker(\widetilde{F_2}) \subseteq \ker(\widetilde{F}) \cap \mathcal{B}[d_2] = F(\mathcal{A}[d]) \cap \mathcal{B}[d_2]$.
(iv) *When $\ker(F)$ has rank $g := \dim(\mathcal{A})$, those inclusions are equalities.*

In the case of SQISignHD, $d_1 = \ell^{e_1}$ and $d_2 = \ell^{e_2}$ and we have accessible ℓ^f-torsion with $f \geq e_1, e_2$. Since $\ker(F)$ has maximal rank $g = 4, 8$ (depending on the version of SQISignHD), we have by point (iv) of the above proposition

$$\ker(F_1) = \ker(F)[\ell^{e_1}] = \{(\widetilde{\alpha}(P), \varSigma(P)) \mid P \in E_A^{g/2}[\ell^{e_1}]\}$$

$$\ker(\widetilde{F_2}) = \ker(\widetilde{F})[\ell^{e_2}] = \{(\alpha(P), -\varSigma(P)) \mid P \in E_A^{g/2}[\ell^{e_2}]\}$$

with the notations of Section 3.3.

5.5 Strategies for higher dimensional isogeny computation

In this paragraph, we give an overview of the higher dimensional isogeny computation procedures $\mathsf{RepresentIsogeny}_{4,\ell^e,\ell^f}$ and $\mathsf{RepresentIsogeny}_{8,\ell^e,\ell^f}$ used in SQISignHD. First, we explain how to compute an ℓ^e-isogeny between abelian varieties in plain generality and then apply it to our specific problem.

Computing an ℓ-isogeny chain. Let $F : (\mathcal{A}, \lambda_\mathcal{A}) \longrightarrow (\mathcal{B}, \lambda_\mathcal{B})$ be an ℓ^e-isogeny between principally polarized abelian varieties and let K be its kernel. Assume that K has rank g. Then, we can decompose F as an ℓ-isogeny chain as in dimension 1.

Lemma 5.5.1. *We can write F as a product of ℓ-isogenies $F = F_e \circ \cdots \circ F_1$ between principally polarized abelian varieties $F_i : \mathcal{A}_{i-1} \longrightarrow \mathcal{A}_i$ (for $i \in [\![1\ ;\ e]\!]$, with $\mathcal{A}_0 := \mathcal{A}$ and $\mathcal{A}_e := \mathcal{B}$).*

Let $K_0 := K = \ker(F)$ and $K_i := F_i(K_{i-1})$ for all $i \in [\![1\ ;\ e]\!]$. Then, we have $\ker(F_i) = [\ell^{e-i}]K_{i-1}$ for all $i \in [\![1\ ;\ e]\!]$.

Proof. We prove by induction on $i \in [\![0\ ;\ e]\!]$ that we can write $F = G_i \circ F_i \circ \cdots \circ F_1$ where the $F_j : \mathcal{A}_{j-1} \longrightarrow \mathcal{A}_j$ for $j \in [\![1\ ;\ i]\!]$ are ℓ-isogenies between principally polarized abelian varieties (PPAV) of kernel $\ker(F_j) = [\ell^{e-j}]K_{j-1}$ and G_i is an ℓ^{e-i}-isogeny between PPAV of kernel $\ker(G_i) = K_i$. For $i = 0$, the result is trivial. Let us assume the result at rank $i \in [\![0\ ;\ e-1]\!]$. Then we simply apply point (i) of Proposition 5.4.1 to G_i to write $G_i := G_{i+1} \circ F_{i+1}$, where F_{i+1} and G_{i+1} are respectively ℓ and $\ell^{e-(i+1)}$-isogenies between PPAV. By point (iv) of Proposition 5.4.1, we also have $\ker(F_{i+1}) = K_i \cap \mathcal{A}_i[\ell]$ since K_i has rank g (K also having rank g). Since $K_i \subset \mathcal{A}_i[\ell^{e-i}]$ is maximal isotropic of rank g, we may write $K_i = \langle x_1, \cdots, x_g \rangle$ where all the $x_i \in \mathcal{A}_i[\ell^{e-i}]$ have order ℓ^{e-i}, so that

$$\ker(F_{i+1}) = K_i \cap \mathcal{A}_i[\ell] = \langle [\ell^{e-(i+1)}]x_1, \cdots, [\ell^{e-(i+1)}]x_g \rangle = [\ell^{e-(i+1)}]K_i.$$

We also have $\ker(G_{i+1}) = F_{i+1}(\ker(G_i)) = K_{i+1}$. This completes the proof. □

The above lemma leads to similar isogeny computation algorithms to the dimension 1 case. Assume that we know a basis $\mathscr{B}_0$ of $K = \ker(F)$ and let $\mathscr{B}_i := F_i(\mathscr{B}_{i-1})$, which is a basis of K_i for all $i \in [\![1\ ;\ e]\!]$. Consider the graph whose vertices are the $[\ell^j]\mathscr{B}_i$ for $0 \leq i \leq e-1$, $0 \leq j \leq e-1-i$ and edges are of two kind:

- multiplication by ℓ $[\ell^j]\mathscr{B}_i \longrightarrow [\ell^{j+1}]\mathscr{B}_i$ (left edges);

$$\begin{tikzcd}
 & & & & \mathscr{B}_0 \arrow[dl] \arrow[dr, dashed, "F_1"] & & & & \\
 & & & [\ell]\mathscr{B}_0 \arrow[dl] \arrow[dr, dashed] & & \mathscr{B}_1 \arrow[dl] \arrow[dr, dashed, "F_2"] & & & \\
 & & [\ell^2]\mathscr{B}_0 \arrow[dl] \arrow[dr, dashed] & & [\ell]\mathscr{B}_1 \arrow[dl] \arrow[dr, dashed] & & \mathscr{B}_2 \arrow[dl] \arrow[dr, dashed, "F_3"] & & \\
 & [\ell^3]\mathscr{B}_0 \arrow[dl] \arrow[dr, dashed] & & [\ell^2]\mathscr{B}_1 \arrow[dl] \arrow[dr, dashed] & & [\ell]\mathscr{B}_2 \arrow[dl] \arrow[dr, dashed] & & \mathscr{B}_3 \arrow[dl] \arrow[dr, dashed, "F_4"] & \\
[\ell^4]\mathscr{B}_0 & & [\ell^3]\mathscr{B}_1 & & [\ell^2]\mathscr{B}_2 & & [\ell]\mathscr{B}_3 & & \mathscr{B}_4
\end{tikzcd}$$

Fig. 2. Computational structure of the ℓ^e isogeny F with $e = 5$.

- ℓ-isogeny computation $[\ell^j]\mathscr{B}_i \longrightarrow [\ell^j]\mathscr{B}_{i+1}$ (right edges).

This graph represents the computational structure of F. To compute F, we need to compute the kernel basis $[\ell^{e-1}]\mathscr{B}_0, [\ell^{e-2}]\mathscr{B}_1 \cdots , \mathscr{B}_{e-1}$ representing the ℓ-isogenies in the chain, *i.e.* the bottom line in Figure 2. This graph is computed as follows: to go down right $[\ell^j]\mathscr{B}_i \longrightarrow [\ell^j]\mathscr{B}_{i+1}$, we need to have reached the bottom vertex $[\ell^{e-1-i}]\mathscr{B}_i$ beforehand. Of course there are naive algorithms where we compute every point of the graph but they are quadratic in e and far from optimal. There also exist divide and conquer strategies that require only $O(e\log(e))$ multiplications by ℓ and ℓ-isogeny evaluations (see [JD11, § 4.2.2] for details). We can even optimize such a strategy to minimize the global cost depending on the relative cost of scalar multiplications by ℓ and ℓ-isogeny evaluations. We refer to $\mathsf{KernelToIsogeny}_{g,\ell^e}(\mathscr{B}_0)$ as the algorithm computing an ℓ-isogeny chain of kernel $\langle\mathscr{B}_0\rangle$ with such an optimal strategy.

Computing ℓ-isogenies with the theta model. Unlike in dimension 2 or 3, we cannot use Jacobians to compute isogenies in dimension 4 or 8. However, there already exist algorithms to compute ℓ-isogenies in any dimension g with the Θ-model [LR12; LR15; LR23] in time $O(\ell^g)$. To minimize the complexity, the best strategy would be to take $\ell = 2$ and to use Θ-coordinates of level $n = 2$. However, existing algorithms only work when ℓ and n are coprime. We propose an algorithm to compute 2-isogenies in level $n = 2$ in Appendix C.3 and how we can use it in our specific problem (computing two 2-isogeny chains F_1 and $\widetilde{F_2}$ with the same codomain). Optimizing and implementing this algorithm is left for future work.

5.6 Computing the response isogeny representation.

We finally give algorithms to compute the signature representation in dimension 4 and 8 using all the ideas presented in Sections 5.4 and 5.5 (splitting the computation in two and computing ℓ-isogeny chains with an optimal strategy and the theta model).

In dimension 4, the algorithm for $\mathsf{RepresentIsogeny}_{4,\ell^e,\ell^f}$ (Algorithm 11) is a straightforward application of these ideas. We compute basis of $\ker(F_1)$ and $\ker(\widetilde{F_2})$ with $F := F_2 \circ F_1$, as in Section 5.4. Then, we call $\mathsf{KernelToIsogeny}_4$ to obtain F_1 and $\widetilde{F_2}$ as isogeny chains.

```
Algorithm 11: $\mathsf{RepresentIsogeny}_{4,\ell^e,\ell^f}$
Data: $E_A, E_2$, $a_1, a_2 \in \mathbb{Z}$, a basis $(P_1, P_2)$ of $E_A[\ell^f]$ and $(\sigma(P_1), \sigma(P_2))$, where
      $\sigma : E_A \longrightarrow E_2$ is a $q$-isogeny with $a_1^2 + a_2^2 + q = \ell^e$.
Result: An $\ell^{e_1}$-isogeny $F_1 : E_A^2 \times E_2^2 \longrightarrow \mathcal{C}$ and a $\ell^{e_2}$-isogeny $\widetilde{F_2} : E_A^2 \times E_2^2 \longrightarrow \mathcal{C}$
        such that $F(\sigma, a_1, a_2) = F_2 \circ F_1$, with $e_1, e_2 \leq f$ and $e_1 + e_2 = e$.
$e_2 \longleftarrow \lceil e/2 \rceil, e_1 \longleftarrow e - e_2$;
$\mathscr{B}_0 \longleftarrow [\ell^{f-e_1}]\big(([a_1]P_i, [a_2]P_i, \sigma(P_i), 0)_{i\in\{1,2\}}, (-[a_2]P_i, [a_1]P_i, 0, \sigma(P_i))_{i\in\{1,2\}}\big)$;
$Q_i \longleftarrow [\ell^{f-e_2}]P_i$, $R_i \longleftarrow [\ell^{f-e_2}]\sigma(P_i)$ for $i \in \{1,2\}$;
$\mathscr{C}_0 \longleftarrow (([a_1]Q_i, -[a_2]Q_i, -R_i, 0)_{i\in\{1,2\}}, ([a_2]Q_i, [a_1]Q_i, 0, -R_i)_{i\in\{1,2\}})$;
$F_1 \longleftarrow \mathsf{KernelToIsogeny}_{4,\ell^{e_1}}(\mathscr{B}_0)$;
$\widetilde{F_2} \longleftarrow \mathsf{KernelToIsogeny}_{4,\ell^{e_2}}(\mathscr{C}_0)$;
Return $F_1$ and $\widetilde{F_2}$;
```

In dimension 8, the algorithm $\mathsf{RepresentIsogeny}_{8,\ell^e,\ell^f}$ (Algorithm 12) only works when q and ℓ are coprime. We explain in Appendix B how to treat the general case.

Proposition 5.6.1. *Algorithms 11 and 12 are correct. Namely, Algorithm 11 returns $F_1, \tilde{F}_2$ such that $F_2 \circ F_1 = F(\sigma, a_1, a_2)$ on entry $a_1, a_2, P_1, P_2, \sigma(P_2)$, $\sigma(P_2)$, where $\sigma : E_A \longrightarrow E_2$ is a q-isogeny with $a_1^2 + a_2^2 + q = \ell^e$ (and similarly for Algorithm 12).*

Proof. See Appendix A.4. □

Remark 5.3. To make sure we have enough accessible torsion, we need $f \geq e_1, e_2$, so that $2f \geq e$. Actually, for $\mathsf{KernelToIsogeny}_{g,\ell^{e_i}}$ to work (with theta coordinates of level 2), we need $4\ell^{e_i}$-torsion points (see Appendix C.2). Then, when $\ell = 2$, we have $f \geq e_i + 2$, so $2f \geq e + 4$.

5.7 Verification

We describe the verification procedure $\mathsf{IsValid}_g$ for $g \in \{4, 8\}$ taking as input the isogenies F_1 and $\widetilde{F_2}$ outputted by $\mathsf{RepresentIsogeny}_{g,\ell^e,\ell^f}$ and determining if they represent an isogeny $\sigma : E_A \longrightarrow E_2$ of degree q. The idea is to check if F_1 and $\tilde{F}_2$ have the same codomain (computed as principally polarized abelian varieties) and then evaluate $F_2 \circ F_1$ on some points.

Algorithm 12: $\mathsf{RepresentIsogeny}_{8,\ell^e,\ell^f}$

```
Data: $E_A, E_2, a_1, a_2, a_3, a_4 \in \mathbb{Z}$, a basis $(P_1, P_2)$ of $E_A[\ell^f]$, $(\sigma(P_1), \sigma(P_2))$ where
      $\sigma : E_A \longrightarrow E_2$ is a $q$-isogeny with $a_1^2 + a_2^2 + a_3^2 + a_4^2 + q = \ell^e$.
Result: An $\ell^{e_1}$-isogeny $F_1 : E_A^4 \times E_2^4 \longrightarrow \mathcal{C}$ and a $\ell^{e_2}$-isogeny $\widetilde{F_2} : E_A^4 \times E_2^4 \longrightarrow \mathcal{C}$
        such that $F(\sigma, a_1, a_2, a_3, a_4) = F_2 \circ F_1$, with $e_1, e_2 \leq f$ and $e_1 + e_2 = e$.
$e_2 \longleftarrow \lceil e/2 \rceil, e_1 \longleftarrow e - e_2$;
Let $\alpha := \begin{pmatrix} a_1 & -a_2 & -a_3 & -a_4 \\ a_2 & a_1 & a_4 & -a_3 \\ a_3 & -a_4 & a_1 & a_2 \\ a_4 & a_3 & -a_2 & a_1 \end{pmatrix} \in \mathrm{End}(E_A^2)$ and $\Sigma := \mathrm{Diag}(\sigma, \sigma, \sigma, \sigma)$;
For all $i \in \{1, 2\}$ and $j \in \{1, 2, 3, 4\}$, let $\underline{P}_{i,j} \in E_A^4$ be the tuple with $P_i$ in position $j$ and 0 elsewhere;
$\mathscr{B}_0 \longleftarrow ([\ell^{f-e_1}]\tilde{\alpha}(\underline{P}_{i,j}), [\ell^{f-e_1}]\Sigma(\underline{P}_{i,j}))_{\substack{1\leq i\leq 2\\ 1\leq j\leq 4}}$;
$\mathscr{C}_0 \longleftarrow ([\ell^{f-e_2}]\alpha(\underline{P}_{i,j}), -[\ell^{f-e_2}]\Sigma(\underline{P}_{i,j}))_{\substack{1\leq i\leq 2\\ 1\leq j\leq 4}}$;
$F_1 \longleftarrow \mathsf{KernelToIsogeny}_{8,\ell^{e_1}}(\mathscr{B}_0)$;
$\widetilde{F_2} \longleftarrow \mathsf{KernelToIsogeny}_{8,\ell^{e_2}}(\mathscr{C}_0)$;
Return $F_1$ and $\widetilde{F_2}$;
```

Algorithm 13: $\mathsf{IsValid}_{4,\ell^e,\ell^f\ell'^{f'}}$

```
Data: Elliptic curves $E_A, E_2$, integers $a_1, a_2 \in \mathbb{Z}$ and the output $(F_1, \widetilde{F_2})$ of
      $\mathsf{RepresentIsogeny}_{4,\ell^e,\ell^f}(E_A, E_2, a_1, a_2, *, *, *, *)$.
Result: Determines if $F_2 \circ F_1$ is an efficient repersentation of an isogeny $\sigma :$
        $E_A \longrightarrow E_2$ of degree $q := \ell^e - a_1^2 - a_2^2$.
Let $(\mathcal{C}_1, \lambda_1)$ and $(\mathcal{C}_2, \lambda_2)$ be the respective codomains of $F_1$ and $\widetilde{F_2}$;
if $(\mathcal{C}_1, \lambda_1) \neq (\mathcal{C}_2, \lambda_2)$ then
    Return False;
else
    Find a point $Q \in E_A$ of order $\ell^f \ell'^{f'}$;
    Compute $\underline{T} \longleftarrow F_2 \circ F_1(Q, 0, 0, 0)$;
    if $\underline{T} = ([a_1]Q, -[a_2]Q, *, 0)$ then
        Return True;
    else
        Return False;
    end
end
```

```
Algorithm 14: IsValid$_{8,\ell^e,T}$
Data: Elliptic curves $E_A, E_2$, integers $a_1, a_2, a_3, a_4 \in \mathbb{Z}$ and the output $(F_1, \widetilde{F_2})$
      of RepresentIsogeny$_{8,\ell^e,\ell^f}(E_A, E_2, a_1, a_2, a_3, a_4, *, *, *, *)$.
Result: Determines if $F_2 \circ F_1$ is an efficient representation of an isogeny $\sigma :$
      $E_A \longrightarrow E_2$ is an isogeny of degree $q := \ell^e - a_1^2 - a_2^2 - a_3^2 - a_4^2$.
Let $(\mathcal{C}_1, \lambda_1)$ and $(\mathcal{C}_2, \lambda_2)$ be the respective codomains of $F_1$ and $\widetilde{F_2}$;
if $(\mathcal{C}_1, \lambda_1) \neq (\mathcal{C}_2, \lambda_2)$ then
    Return False;
else
    Write $T := \prod_{i=1}^r \ell_i^{\beta_i}$ (accessible torsion);
    Generate a basis $(U_{i,1}, U_{i,2})$ of $E_A[\ell_i^{\beta_i}]$ for all $i \in [\![1 ; r]\!]$;
    Let $\alpha \in \mathrm{End}(E_A^2)$ be as in line 2 of Algorithm 12;
    Compute $\underline{T}_{i,j} \longleftarrow F_2 \circ F_1(U_{i,j}, 0, \cdots, 0)$ for all $i \in [\![1 ; r]\!]$ and $j \in \{1, 2\}$;
    if $\underline{T}_{i,j} = (\alpha(U_{i,j}, 0, 0, 0), *, 0, 0, 0)$ for all $(i, j) \in [\![1 ; r]\!] \times \{1, 2, 3, 4\}$ then
        Return True;
    else
        Return False;
    end
end
```

The following results (proved in Appendix A.5) ensure that our verification procedure is correct. Antonin:C'est un détail, mais est-ce qu'on ne devrait pas prévoir la possibilité que RepresentIsogeny$_4$ échoue (je suppose que c'est possible si on donne les mauvais points) ?

Proposition 5.7.1. *Algorithms 13 and 14 are correct. Namely, when given $E_A, E_2, a_1, a_2, F_1, \tilde{F}_2$, if Algorithm 13 returns* True*, then $F_2 \circ F_1$ is an efficient representation of an isogeny $\sigma : E_A \longrightarrow E_2$ of degree $q = \ell^e - a_1^2 - a_2^2$ (and similarly for Algorithm 14).*

Corollary 5.7.2. *The verification procedures* FastVerify *and* RigorousVerify *(Algorithms 8 and 9) are correct. Namely, on input (R_1, R_2, q),* FastVerify *(respectively* RigorousVerify*) returns* True *if and only if (R_1, R_2, q) defines an efficient representation of an isogeny $\sigma : E_A \longrightarrow E_2$ of degree q, where q is ℓ^e-good and coprime to ℓ' (respectively $q < \ell^e$).*

6 Security analysis

In this section, we prove that the SQISignHD identification protocol is secure, namely that it is complete, knowledge sound and honest-verifier zero knowledge. Recall that by [VV15, Theorem 7], it is sufficient to ensure that our signature scheme obtained by Fiat-Shamir transform is universally unforgeable under chosen message attacks in the random oracle model.

Completeness means that a honest execution of the protocol is always accepted by the verifier. This is true by Proposition 5.6.1 and by construction of

$\mathsf{IsValid}$. Knowledge soundness means that an attacker can only "guess" a response with very low probability. It is proven under the assumption that computing an endomorphism in a supersingular elliptic curve is hard, a well known difficult problem in isogeny based cryptography.

The honest-verifier zero-knowledge property implies that the response does not leak any information on the secret key τ. More precisely, we can simulate transcripts of the identification protocol without using to the secret key with the same distribution as real transcripts. To construct such a simulator of SQISignHD, we need access to an oracle evaluating isogenies of non-smooth degrees. In RigorousSQISignHD, this oracle is very generic and we do not need any additional hypothesis to prove the zero-knowledge property (hence the name of this version). On the contrary, in FastSQISignHD, the oracle definition is *ad hoc* and we need an additional heuristic assumption to prove the zero-knowledge property. However, it is very unlikely to build an attack on this assumption and both oracles do not undermine the knowledge soundess.

6.1 Knowledge soundness

The proof that FastSQISignHD is knowledge sound is a straightforward special soundess argument identical to the original version of SQISign [DKLPW20, Theorem 1]. Namely, we prove that given two transcripts with the same commitment but disctinct challenges, we can find an endomorphism in E_A. This *special soundness* property is sufficient to prove that SQISignHD satisfies *knowledge soundness* [HL10, Theorem 6.3.2]. However, note that we have to require the prime ideal norm q to be not only ℓ^e-good but also coprime to ℓ' in order to complete the proof.

Proposition 6.1.1. *Under the assumption that $q = \deg(\sigma)$ is always coprime to ℓ', the* FastSQISignHD *identification protocol satisfies special soundness. Namely, given two transcripts $(E_1, \varphi, R_1, R_2, q)$ and $(E_1, \varphi', R_1', R_2', q')$ with the same commitment E_1 but different challenges $\varphi \neq \varphi'$, we can extract an efficient representation of a non-scalar endomorphism $\alpha \in \mathrm{End}(E_A)$.*

Proof. Let $(E_1, \varphi, R_1, R_2, q)$ and $(E_1, \varphi', R_1', R_2', q')$ be two FastSQISignHD transcripts with the same commitment E_1 but different challenges $\varphi \neq \varphi'$. Then, by Corollary 5.7.2, (R_1, R_2, q) and (R_1', R_2', q') define efficient representations of isogenies $\sigma, \sigma' : E_A \longrightarrow E_2$ of degrees q and q' respectively which are ℓ^e-good and comprime with ℓ'. With such an efficient representation of σ', we can obtain $(\sigma'(P_1), \sigma'(P_2))$ in polynomial time in $\log(p)$, where (P_1, P_2) is a canonical basis of $E_A[\ell^f]$. We can also find $a_1', a_2' \in \mathbb{Z}$ such that ${a_1'}^2 + {a_2'}^2 + q = \ell^e$ and apply $\mathsf{RepresentIsogeny}_{4,\ell^e,\ell^f}$ to compute $F' := F(\sigma', a_1', a_2')$ by Proposition 5.6.1. Then, F' provides an efficient representation of $\widehat{\sigma'}$.

Hence, we know an efficient representation of $\alpha := \widehat{\sigma'} \circ \varphi' \circ \widehat{\varphi} \circ \sigma \in \mathrm{End}(E_A)$. We now prove that α is not scalar. Indeed, if it was, we would have $\alpha = [\lambda]$ for some $\lambda \in \mathbb{Z}$ and $qq'\ell'^{2f'} = \lambda^2$ where $q := \deg(\sigma)$ and $q' := \deg(\sigma')$ are coprime to ℓ'. Hence, $\lambda = \ell'^{f'}\lambda'$ with $\lambda' \in \mathbb{Z}$ prime to ℓ' ($\lambda'^2 = qq'$). It follows that

$[q]\widehat{\sigma}' \circ \varphi' = [\lambda']\widehat{\sigma} \circ \varphi$. Since q, q' and λ' are coprime to ℓ', we get that $\ker(\varphi) = \ker(\varphi')$ *i.e.* $\varphi = \varphi'$ up to post-composition by an automorphism. Contradiction. This completes the proof. □

For RigorousSQISignHD, our knowledge soundness argument does not apply because we have no guarantee on q in general. For that reason, we need to come back to the formal definition of knowledge soundess given in [HL10, Definition 6.3.1]. This analysis is conducted in Appendix A.6.

The previous proof of knowledge would be trivial if it was easy to find an endomorphism. Fortunately, this is a well-known hard problem in isogeny-based cryptography.

Problem 6.1.2 (Supersingular Endomorphism Problem)**.** Given a supersingular elliptic curve $E/\mathbb{F}_{p^2}$, find an efficient representation of a non-scalar endomorphism $\alpha \in \mathrm{End}(E)$.

This problem is very similar to [DKLPW20, Problem 1], except that we do not require the endomorphism to have smooth degree. This does not seems to make the problem easier since the endomorphisms solution to this can be evaluated (which was the reason why smoothness was imposed in the first place). The supersingular endomorphism ring problem (Problem 6.1.3) reduces to Problem 6.1.2. Problem 6.1.3 is notoriously hard and it has been proven it is equivalent to path finding in the supersingular ℓ-isogeny graph [Wes22]. The heuristic reduction from Problem 6.1.3 to 6.1.2 is given by [EHLMP18, Algorithm 8]. Basically, if we have an oracle finding endomorphisms of E, we call this oracle until we have found enough endomorphisms to generate $\mathrm{End}(E)$.

Problem 6.1.3 (Supersingular Endomorphism Ring Problem)**.** Given a supersingular elliptic curve $E/\mathbb{F}_{p^2}$, find four endomorphisms of E (that we can evaluate) forming a $\mathbb{Z}$-basis of $\mathrm{End}(E)$.

6.2 Rigorous zero-knowledge property

The proof of the zero-knowledge property of RigorousSQISignHD uses an oracle generating isogenies of non-smooth degree. To our knowledge, there is no efficient algorithm implementing such an oracle. Nonetheless, it is believed that access to such an oracle does not affect the hardness of the underlying problem (the endomorphism ring problem, see Section 6.4).

Definition 6.2.1. A *random any degree isogeny oracle* (RADIO) is an oracle taking as input a supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$ and returning an efficient representation of an isogeny $\sigma : E \longrightarrow E'$, which is uniformly random among the isogenies of degree $q < \ell^e$ with domain E.

To prove zero-knowledge, not only the distribution of the isogeny should be uniform, but also the distribution of its codomain. Unfortunately, the RADIO does not determine the distribution of codomains. The following theorem (proved in Appendix A.7) solves this issue when ℓ^e is big enough. We need $\ell^e = \Theta(p^2)$

to ensure a statistical distance of $O(p^{-1/2})$ to the uniform distribution, hence a security level of λ (since $p = \Theta(2^{2\lambda})$).

Theorem 6.2.2. *Let $\varepsilon \in]0,2]$. Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve and π be the probability distribution of codomains E' (up to $\overline{\mathbb{F}_p}$-isomorphism) of isogenies $\sigma : E \longrightarrow E'$ chosen uniformly at random among isogenies of degree $\deg(\sigma) \leq p^{1+\varepsilon}$. Let U be the uniform distribution of supersingular elliptic curves over $\mathbb{F}_{p^2}$ (up to $\overline{\mathbb{F}_p}$-isomorphism). Then $d_{TV}(U, \pi) = O(p^{-\varepsilon/2})$.*

Theorem 6.2.3. *The* RigorousSQISignHD *protocol is statistically honest-verifier zero knowledge in the RADIO model (provided $\ell^e = \Theta(p^2)$). In other words, there exists a random polynomial time simulator $\mathcal{S}$ with black-box access to a RADIO that simulates transcripts (E_1, φ, R) with a statistically indistinguishable distribution from the transcripts of the* RigorousSQISignHD *identification protocol.*

Proof. First, we explain how to construct the simulator $\mathcal{S}$. It begins by using the RADIO on entry E_A to get an efficient representation of a signature isogeny $\sigma' : E_A \longrightarrow E_2'$ which is uniformly random among the isogenies of degree $q < \ell^e$ with domain E_A. Then we generate a canonical basis (P_1, P_2) of $E_A[\ell^f]$ and compute its image by σ' in polynomial time. Then, $\mathcal{S}$ generates a uniformly random cyclic $\widehat{\varphi}' : E_2' \longrightarrow E_1'$ of degree D_φ exactly as in Section 4.2 and returns (E_1', φ', R') with $R' := (\sigma'(P_1), \sigma'(P_2), q)$.

We now prove that the transcripts (E_1', φ', R') of $\mathcal{S}$ are statistically indistinguishable from the transcripts (E_1, φ, R) of the RigorousSQISignHD identification protocol. We first notice that the codomain E_2' of σ' is uniformly random in the supersingular isogeny graph by definition of the RADIO and Theorem 6.2.2. As a consequence, E_1' is uniformly random as well, the distribution of $\widehat{\varphi}'$ being uniformly random as well among isogenies of degree D_φ. But E_1 is uniformly random in the supersingular isogeny graph by Proposition 4.4.1. φ' also has the same distribution as φ by construction. Hence, the commitment and challenge (E_1', φ') and (E_1, φ) are statistically undistinguishable.

Finally, the isogeny σ' represented by R' is uniformly random among the isogenies $E_A \longrightarrow E_2'$ of degree $q < \ell^e$ by definition of the RADIO, and so is the isogeny σ represented by R by construction (see Section 5.2). This completes the proof. □

6.3 Heuristic zero-knowledge property

As for RigorousSQISignHD, the proof of the zero-knowledge property of FastSQISignHD uses an auxiliary oracle. While the RADIO definition was very natural, the definition of this new oracle seems more *ad hoc*: we add (mild) conditions on the degree to account for the computational constraints imposed by the method in dimension 4. These degree constraints are the main reason why the signatures are represented in dimension 8 instead of 4 in RigorousSQISignHD. As previously, accessing this new oracle does not hamper other security properties of the protocol (see 6.4).

Definition 6.3.1. A *random uniform good degree isogeny oracle* (RUGDIO) is an oracle taking as input a supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$ and returning an efficient representation of a random isogeny $\sigma : E \longrightarrow E'$ of ℓ^e-good degree coprime to ℓ', such that:

(i) The distribution of E' is uniform in the supersingular isogeny graph.
(ii) The conditional distribution of σ given E' is uniform among isogenies $E \longrightarrow E'$ of ℓ^e-good degree coprime to ℓ'.

In addition to the constraint on the degree of the RUGDIO output, we add constraints on the distributions of isogenies. Those constraints are necessary to construct a simulator of FastSQISignHD. We justify that these constraints can be mathematically satisfied, namely that for all supersingular elliptic curves E and E', there exists $\sigma : E \longrightarrow E'$ of ℓ^e-good norm. As explained in Section 5.2, taking ℓ^e slightly bigger than $\sqrt{p}$ by a few bits should be heuristically sufficient. However, we should take $\ell^e = \omega(p^2)$ to give a formal proof, with a huge efficiency cost.

Theorem 6.3.2. *Assume that the commitment E_1 is computationally indistinguishable from an elliptic curve chosen uniformly at random in the supersingular isogeny graph. Then, the* FastSQISignHD *identification protocol is computationally honest-verifier zero knowledge in the RUGDIO model.*

In other words, under this assumption, there exists a random polynomial time simulator $\mathcal{S}$ with access to a RUGDIO that simulates transcripts (E_1, φ, R) with a computationally indistinguishable distribution from the transcripts of the FastSQISignHD *identification protocol.*

Proof. We construct the simulator $\mathcal{S}$ as in Theorem 6.2.3 except that we use a RUGDIO instead of a RADIO. As previously, let us call (E_1, φ, R) and (E_1', φ', R') the transcripts generated by the real FastSQISignHD identification protocol and $\mathcal{S}$ respectively.

We now prove that (E_1, φ, R) and (E_1', φ', R') are computationally indistinguishable. Let σ and σ' be the isogenies represented by R and R' respectively. We first notice that the codomain E_2' of σ' is uniformly random in the supersingular isogeny graph by definition of the RUGDIO. As a consequence, E_1' is uniformly random as well, the distribution of $\widehat{\varphi'}$ being uniformly random among the cyclic isogenies of degree D_φ. Hence, by heuristic assumption, E_1' is computationally indistinguishable from E_1 and φ' is computationally indistinguishable from φ as well since both isogenies are generated in the same way.

By construction, conditionally to E_2, $\sigma : E_A \longrightarrow E_2$ is uniform among all the isogenies of ℓ^e-good degree coprime to ℓ' with codomain E_2 Antonin:cela fait écho à ma remarque précédente sur RandomEquivalentIdeal mais je ne crois pas qu'on prouve ce que tu affirmes ici. On doit sans doute pouvoir dire heuristiquement que la distribution qu'on obtient est statistiquement proche de l'uniforme, mais prouver l'uniforme me parait plus complexe et demandrait sans doute un algorithme beaucoup plus lent que celui qu'on veut utiliser en pratique. . Conditionally to E_2', $\sigma' : E_A \longrightarrow E_2'$ has the same distribution by construction of the

RUGDIO so σ and σ' are statistically indistinguishable and R and R' as well. The result follows. □

It remains to justify that the commitment E_1 is computationally indistinguishable from an elliptic curve chosen uniformly at random in the supersingular isogeny graph. While RigorousCommit satisfies statistical indistinguishability, the variant FastCommit relies on heuristics. Consider the distributions on E_1 induced by the following procedures

1. Return the output E_1 of FastCommit.
2. Generate a uniformly random cyclic endomorphism γ of E_0 of degree $\ell^{2f}\ell'^{2f'}$. Factor it as $\gamma = \hat{\phi}' \circ \phi$ with $\deg(\phi) = \ell^{2f}$. Return the codomain E_1 of ϕ.
3. Generate a uniformly random cyclic isogeny ϕ from E_0 of degree ℓ^{2f}. Let E_1 be its codomain; let m be the number of cyclic isogenies $\phi' : E_0 \to E_1$ of degree $\ell'^{2f'}$. Return E_1 with probability m/M (for some fixed upper bound M on m, for instance $M = (\ell' + 1)\ell'^{2f'-1}$), otherwise resample.
4. Generate a uniformly random cyclic isogeny ϕ from E_0 of degree ℓ^{2f}; return its codomain E_1.
5. Return a uniformly random elliptic curve E_1.

We argue that each distribution from the list is somewhat close to the next. The difference between 1 and 2 is that in FastCommit, the endomorphism γ is not truly uniform: they follow a distribution biased by the fact that some intermediate result should be easy to factor. Since, this property appears somewhat decorrelated from the final distribution of γ it seems plausible to argue that the distribution of γ in 1 is close to the one in 2. The distributions 2 and 3 are actually identical: distribution 3 simulated distribution 2 by rejection sampling. The difference between 3 and 4 is that m is not necessarily a (positive) constant; it is however heuristically expected to be almost a constant: there are about $(\ell' - 1)\ell'^{2f'-1}$ possible paths, and about $p/12$ vertices, so we expect about $m \approx 12(\ell' - 1)\ell'^{2f'-1}/p$ distinct paths to any fixed vertex. The difference between 4 and 5 is similar, but reasoning about ℓ-paths instead of ℓ'-paths.

Note that the differences at some of these steps are statistically significant. We only argue that they are not computational detectable, at least when the endomorphism rings are not known.

6.4 On hardness of the supersingular endomorphism problem with access to an auxiliary oracle

The identification protocol is sound assuming the hardness of the supersingular endomorphism problem 6.1.2, and zero-knowledge with respect to a simulator that has access to a RADIO (or a RUGDIO). For the resulting signature scheme to be secure, one therefore needs to assume that the supersingular endomorphism problem remains hard even when given access to a RADIO.

While it currently seems out of reach to prove that the supersingular endomorphism problem is equivalent to the variant with RADIO access, let us argue that the RADIO indeed does not help. We focus the following discussion on the RADIO, but the same arguments apply to the RUGDIO despite the slightly biased distribution.

The RADIO allows to generate random isogenies with a chosen domain E. Note that this task is already known to be easy, with isogenies of smooth degree. The RADIO only lifts this smoothness restriction: it allows to generate random isogenies whose degrees have large prime factors. It does not allow to reach more target curves, nor does it give more control on which specific target to hit: the target curve is uniformly distributed in the supersingular graph, which was already possible with smooth isogenies.

Smoothness of random isogenies has never been an inconvenience in finding endomorphisms. In fact, the best current fastest algorithms for this problem only require very smooth degree isogenies, typically a power of 2. The reason is the following: the purpose of constructing a random isogeny from a fixed source is to reach a random target. As very smooth isogenies (even 2-smooth) are sufficient for optimal randomisation, there is no incentive to involve much larger prime factors. More specifically, the best known strategies to solve the supersingular endomorphism problem [DG16] have classical time complexity $\tilde{O}(\sqrt{p})$ (and quantum time complexity $\tilde{O}(p^{1/4})$ with a Grover argument [Gro96]) and essentially perform a meet-in-the-middle search in the supersingular isogeny graph. Access to a RADIO would allow to use isogenies of a different shape in the search, but would not speed it up, as the probability to find isogenies with matching codomains stays the same. Another illustration that having access to non-smooth degree isogenies does not help is the fact that the discovery of the $\sqrt{\text{é}}$lu algorithm [BDLS20] (which dramatically improved the complexity of computing prime degree isogenies) did not affect the state-of-the-art of the supersingular endomorphism problem.

The above arguments support that random isogenies of non-smooth degrees are not more helpful than random isogenies of smooth degrees. Now, one may be concerned that the encoding of the output of the RADIO may leak more information than it should. Non-smooth degree isogenies are represented as a component of a higher dimensional isogeny (Section 3.2). This representation is universal, in the sense that any efficient representation of an isogeny can be efficiently rewritten in this form. In particular, this encoding contains no more information than any other efficient representation of the same isogeny.

7 The SQISignHD digital signature scheme

The SQISignHD identification protocol that we presented yields a digital signature scheme via the Fiat-Shamir transform (see Section 1.2). The security of the transform of both versions FastSQISignHD and RigorousSQISignHD follows from the analysis conducted in Section 6, so the digital signature is also secure under the same computational assumptions. Namely, we have seen it is universally unforgeable under chosen message attacks in the random oracle and RADIO or RUGDIO model, assuming the hardness of the endomorphism ring problem. In this section, we present the performance of the signature scheme obtained from FastSQISignHD.

7.1 Compactness

As explained before, the signature is made of the data $(E_1, q, \sigma(P_1), \sigma(P_2))$, with $q < \ell^e$, $\sigma : E_A \longrightarrow E_2$ a q-isogeny and (P_1, P_2) a basis of $E_A[\ell^f]$ determined canonically.

E_1 can be entirely determined by its j-invariant $j(E_1) \in \mathbb{F}_{p^2}$. Since any element of $\mathbb{F}_{p^2}$ can be represented by 2 integers in $[\![0 \ ; \ p-1]\!]$, storing $j(E_1)$ takes approximately $2\log_2(p) \simeq 4\lambda$ bits, given that $p = \Theta(p^{2\lambda})$ (where λ is the security level). Similarly, $q < \ell^e \simeq \sqrt{p}$, so q is an integer of $1/2\log_2(p) \simeq \lambda$ bits.

The points $\sigma(P_1)$ and $\sigma(P_2)$ need not be represented explicitely with coordinates in $\mathbb{F}_{p^2}$. They can be compressed. Indeed, if we generate a canonical basis (Q_1, Q_2) of $E_2[\ell^f]$, then we may write $\sigma(P_1) = a_1Q_1 + b_1Q_2$ and $\sigma(P_2) = a_2Q_1 + b_2Q_2$ with $a_1, b_1, a_2, b_2 \in \mathbb{Z}/\ell^f\mathbb{Z}$. Storing the scalars a_1, b_1, a_2, b_2 requires $4f$ bits (assuming $\ell = 2$, which will be the case in practice).

Actually, we can gain f bits by ommitting one of the scalars a_1, b_1, a_2, b_2 if we use the Weil pairing. Indeed, we have on the one hand

$$e_{\ell^f}(\sigma(P_1), \sigma(P_2)) = e_{\ell^f}(P_1, P_2)^q.$$

And on the other hand

$$e_{\ell^f}(\sigma(P_1), \sigma(P_2)) = e_{\ell^f}(a_1Q_1 + b_1Q_2, a_2Q_1 + b_2Q_2) = e_{\ell^f}(Q_1, Q_2)^{a_1b_2 - b_1a_2}.$$

Since (P_1, P_2) and (Q_1, Q_2) are basis of $E_A[\ell^f]$ and $E_2[\ell^f]$ respectively, $e_{\ell^f}(P_1, P_2)$ and $e_{\ell^f}(Q_1, Q_2)$ are both primitive ℓ^f-th root of unity. Hence, we may find $k \in (\mathbb{Z}/\ell^f\mathbb{Z})^\times$ such that $e_{\ell^f}(P_1, P_2) = e_{\ell^f}(Q_1, Q_2)^k$, and we must have

$$a_1b_2 - b_1a_2 \equiv kq \mod \ell^f \tag{1}$$

Remark 7.1. Since $\ell^f | p-1$, the ℓ^f-th Weil pairing takes values in $\mathbb{F}_p^*$, so we find k easily by solving a discrete logarithm problem in a subgroup of order ℓ^f of $\mathbb{F}_p^*$ by Pohlig-Hellman [PH78] techniques (which apply since $p-1$ is smooth).

Since q is coprime to ℓ, $\sigma(P_1)$ have order ℓ^f so either a_1 or b_1 is invertible modulo ℓ^f. If a_1 is invertible, we can recover b_2 from the other scalars using equation 1 and we can recover a_2 otherwise. Hence we only need 3 scalars among 4.

We can make the representation of $\sigma(P_1)$ and $\sigma(P_2)$ even more compact. Indeed, by Remark 5.3 the ℓ^e-isogeny F representing σ can be computed as long as $2f \geq e + 4$. But in FastSQISignHD, $f \simeq e \simeq \lambda$ so we may use points of ℓ^{f_1}-torsion with $f_1 := \lceil e/2 \rceil + 3$ instead of points of ℓ^f-torsion. This reduces the storage cost of $\sigma(P_1)$ and $\sigma(P_2)$ from $3f \simeq 3\lambda$ to $3f_1 \simeq 3/2\lambda$.

On the whole, we can represent the signatures with $s = 13/2\lambda + O(\log(\lambda))$ bits if we use the compression and decompression algorithms given by Algorithms 15 and 16, breaking the previous record of SQISign. Indeed, in SQISign, the kernels of the signature isogeny chain σ of degree $p^{15/4}$ and $j(E_1)$ need to be transmitted so we get a signature of size $s = 19/2\lambda + O(\log(\lambda))$ at least.

Example 7.1.1. For NIST-I security level ($\lambda = 128$ bits), we can choose the parameters $p = 2^{128}3^{81} - 1$, $e = 136$ and $f_1 = 91$. The total signature size in SQISignHD is $2\lambda + e + 3f_1 + 1 = 923$ bits or 116 bytes. For the same security level, SQISign signatures took 204 bytes.

Remark 7.1.2. We can use the same techniques in dimension 8 but we output signatures of size $s = 14\lambda + O(\log(\lambda))$ instead of $13/2\lambda + O(\log(\lambda))$ since e is bigger ($\ell^e = \Theta(p^2)$). Details may be found in Appendix B.3.

Algorithm 15: Compression

```
Data: $E_A, E_1, E_2, q, P_1, P_2, \sigma(P_1), \sigma(P_2)$, where $q < \ell^e$, $\sigma : E_A \longrightarrow E_2$ a $q$-isogeny and $(P_1, P_2)$ a basis of $E_A[\ell^{f_1}]$ determined canonically (with $f_1 := \lceil e/2 \rceil + 3$).
Result: A word of length $2\lceil \log_2(p) \rceil + e + 3f_1$ bits (assuming $\ell = 2$).
Compute $j(E_1) \in \mathbb{F}_{p^2}$;
Let $\zeta$ be a canonical generator of $\mathbb{F}_{p^2}$. Write $\zeta := n_1 + n_2\zeta$ where $n_1, n_2 \in \mathbb{F}_p$ are represented by integers in $[\![0 \ ; \ p-1]\!]$ of length $\lceil \log_2(p) \rceil$ bits each;
Compute the canonical basis $(Q_1, Q_2)$ of $E_2[\ell^{f_1}]$;
Find $k \in (\mathbb{Z}/\ell^{f_1}\mathbb{Z})^\times$ such that $e_{\ell^{f_1}}(P_1, P_2) = e_{\ell^{f_1}}(Q_1, Q_2)^k$;
Find $a_1, b_1, a_2, b_2 \in \mathbb{Z}/\ell^{f_1}\mathbb{Z}$ such that $\sigma(P_1) = a_1 Q_1 + b_1 Q_2$ and $\sigma(P_2) = a_2 Q_1 + b_2 Q_2$;
if $\ell \nmid a_1$ then
    Return $\|n_1\|n_2\|q\|a_1\|b_1\|b_2\|$;
else
    Return $\|n_1\|n_2\|q\|a_1\|b_1\|a_2\|$;
end
```

7.2 Time efficiency

Low signing time. In FastSQISignHD, the signature mainly requires the computation of several ℓ^f and $\ell'^{f'}$-isogenies along with other fast operations (point evaluations, discrete logarithms in a group of order $\ell'^{f'}$, quaternion arithmetics). Since ℓ and ℓ' are small (in practice $\ell = 2$ and $\ell' = 3$), the isogeny computations are expected to be almost as fast as in SIDH. The other operations were already implemented in SQISign, and took negligible time compared to isogeny computations.

Antonin:en fait je pense que c'est ici qu'il faudrait faire un récap des différents sous-blocs à utiliser (et dans quel ordre) pour une signature, pour éviter d'aller repiocher dans les sections algorithmiques

In the last version of SQISign [DLW22], the signature required the computation of 30 T-isogenies with $T \simeq p^{5/4}$. The complexity was dominated by these computations because T was not as smooth as power of ℓ. Even if further improvements were made on this scheme, we expect FastSQISignHD to be way

Algorithm 16: Decompression

```
Data: A word $w$ of length $2\lceil \log_2(p) \rceil + e + 3f_1$ bits ($\ell := 2, f_1 := \lceil e/2 \rceil + 3$),
      the public key $E_A$ and the message $m$, hash functions $\Phi$ and $H$ (defined
      in Section 1.2).
Result: $E_1, E_2, q, P_1, P_2, \sigma(P_1), \sigma(P_2)$, where $q < \ell^e$, $\sigma : E_A \longrightarrow E_2$ a $q$-isogeny
        and $(P_1, P_2)$ a basis of $E_A[\ell^{f_1}]$ determined canonically.
Parse $\|n_1\|n_2\|q\|a_1\|b_1\|c_2\| \longleftarrow w$;
Set $j \longleftarrow n_1 + n_2\zeta$, where $\zeta$ is the canonical generator of $\mathbb{F}_{p^2}$;
Compute $E_1$ of $j$-invariant $j(E_1) = j$;
Recover the commitment $\varphi \longleftarrow \Phi(E_1, H(E_1, m))$ (see Section 1.2);
Let $E_2$ be the codomain of $\varphi$;
Compute the canonical basis $(P_1, P_2)$ of $E_A[\ell^{f_1}]$ and the canonical basis
  $(Q_1, Q_2)$ of $E_2[\ell^{f_1}]$;
Find $k \in (\mathbb{Z}/\ell^{f_1}\mathbb{Z})^\times$ such that $e_{\ell^{f_1}}(P_1, P_2) = e_{\ell^{f_1}}(Q_1, Q_2)^k$;
if $\ell \nmid a_1$ then
    $a_2 \longleftarrow c_2$;
    Find $b_2 \in \mathbb{Z}/\ell^{f_1}\mathbb{Z}$ such that $a_1 b_2 - b_1 a_2 \equiv kq \mod \ell^{f_1}$;
else
    $b_2 \longleftarrow c_2$;
    Find $a_2 \in \mathbb{Z}/\ell^{f_1}\mathbb{Z}$ such that $a_1 b_2 - b_1 a_2 \equiv kq \mod \ell^{f_1}$;
end
Return $E_1, E_2, q, P_1, P_2, a_1 Q_1 + b_1 Q_2, a_2 Q_1 + b_2 Q_2$;
```

faster. Our preliminary implementation results indicate that FastSQISignHD signing and key generation time are significantly lower than in SQISign. Providing a completely optimized implementation is left for future works.

Impact on the verification time. However, this efficiency gain in the signature is made at the expense of the verification time where a 4-dimensional ℓ^e-isogeny has to be computed. Of course ℓ-isogenies in dimension 4 are expected to be slower to compute than in dimension 1. Nonetheless, we only have to compute two chains of ℓ-isogenies of length $1/4 \log_\ell(p)$, whereas the verifier had to compute an ℓ-isogeny chain of size $15/4 \log_\ell(p)$ in the last version of SQISign [DLW22]. An implementation would be needed to correctly assess the verification time of FastSQISignHD. This is left for future works.

A verification time vs compactness trade-off. To speed up the verification time, the signer (or any verifier) can expand the compact signature by outputting all e intermediates theta constants computed in the chain of ℓ-isogenies computed during the verification. In dimension g, a theta constant over $\mathbb{F}_{p^2}$ takes $2^g \log(p^2)$ bits. The chain can be verified using Corollary C.3.2.

When $g = 4$, $\lambda = 128$, $\ell = 2$, $e = 128$ and p has 256 bits, storing each 128 theta constants then takes $128 \cdot 2^4 \cdot 512$ bits, that is 131kB. This is a much larger output than the 832 bits of the compressed signature, but by Corollary C.3.2

the verification then takes only $e \cdot 2^{g+1} = 4096$ squares over $\mathbb{F}_{p^2}$ and $2e = 256$ Hadamard transforms (and a final linear change of variable to glue the theta structures at the end), so will be much faster than via the compact signature (compare with Proposition C.3.5 and Example C.3.6).

This allows for a verification time vs compactness trade-off. We remark that expanding the compact isogeny to allow for fast verification time can be done by anyone.

References

[Ahr23] K. Ahrens. *Sieving for large twin smooth integers using single solutions to Prouhet-Tarry-Escott.* Cryptology ePrint Archive, Paper 2023/219. https://eprint.iacr.org/2023/219. 2023. URL: https://eprint.iacr.org/2023/219.

[AJKKL16] R. Azarderakhsh, D. Jao, K. Kalach, B. Koziel, and C. Leonardi. "Key compression for isogeny-based cryptosystems". In: *Proceedings of the 3rd ACM International Workshop on ASIA Public-Key Cryptography.* Xi'an, China: ACM, 2016, pp. 1–10.

[BCC+22] A. Basso, G. Codogni, D. Connolly, L. De Feo, T. B. Fouotsa, G. M. Lido, T. Morrison, L. Panny, S. Patranabis, and B. Wesolowski. *Supersingular Curves You Can Trust.* Cryptology ePrint Archive, Paper 2022/1469. 2022. URL: https://eprint.iacr.org/2022/1469.

[BDLS20] D. J. Bernstein, L. De Feo, A. Leroux, and B. Smith. "Faster computation of isogenies of large prime degree". In: *Open Book Series, Proceedings of the Fourteenth Algorithmic Number Theory Symposium – ANTS XIV* 4.1 (2020), pp. 39–55.

[BST+17] M. Bhargava, A. Shankar, T. Taniguchi, F. Thorne, J. Tsimerman, and Y. Zhao. "Bounds on 2-torsion in class groups of number fields and integral points on elliptic curves". In: *Journal of the American Mathematical Society* 33 (Jan. 2017). DOI: 10.1090/jams/945.

[BL04] C. Birkenhake and H. Lange. *Complex abelian varieties.* Vol. 302. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Berlin: Springer-Verlag, 2004, pp. xii+635. ISBN: 3-540-20488-1.

[BSC+22] G. Bruno, M. C.-R. Santos, C. Costello, J. K. Eriksen, M. Naehrig, M. Meyer, and B. Sterner. *Cryptographic Smooth Neighbors.* Cryptology ePrint Archive, Paper 2022/1439. 2022. URL: https://eprint.iacr.org/2022/1439.

[CD22] W. Castryck and T. Decru. *An efficient key recovery attack on SIDH (preliminary version).* Cryptology ePrint Archive, Paper 2022/975. 2022. URL: https://eprint.iacr.org/2022/975.

[CLG09] D. X. Charles, K. E. Lauter, and E. Z. Goren. "Cryptographic Hash Functions from Expander Graphs". In: *Journal of Cryptology* 22.1 (2009), pp. 93–113. DOI: 10.1007/s00145-007-9002-x.

[Cor08] G. Cornacchia. "Su di un metodo per la risoluzione in numeri interi dell'equazione $\sum_{h=0}^{n} C_h x^{n-h} y^h = P$". In: *Giornale di matematiche di Battaglini* 46 (1908), pp. 33–90.

[CR15] R. Cosset and D. Robert. "An algorithm for computing (ℓ, ℓ)-isogenies in polynomial time on Jacobians of hyperelliptic curves of genus 2". In: *Mathematics of Computation* 84.294 (Nov. 2015), pp. 1953–1975. DOI: 10.1090/S0025-5718-2014-02899-8.

[CMN21] C. Costello, M. Meyer, and M. Naehrig. "Sieving for Twin Smooth Integers with Solutions to the Prouhet-Tarry-Escott Problem". In: *Advances in Cryptology – EUROCRYPT 2021*. Ed. by A. Canteaut and F.-X. Standaert. Cham: Springer International Publishing, 2021, pp. 272–301. ISBN: 978-3-030-77870-5.

[Cou06] J.-M. Couveignes. *Hard Homogeneous Spaces*. Cryptology ePrint Archive, Report 2006/291. 2006. URL: https://eprint.iacr.org/2006/291.

[DDF+21] L. De Feo, C. Delpech de Saint Guilhem, T. B. Fouotsa, P. Kutas, A. Leroux, C. Petit, J. Silva, and B. Wesolowski. "Séta: Supersingular Encryption from Torsion Attacks". In: *Advances in Cryptology – ASIACRYPT 2021: 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6–10, 2021, Proceedings, Part IV*. Singapore, Singapore: Springer-Verlag, 2021, pp. 249–278. ISBN: 978-3-030-92067-8. DOI: 10.1007/978-3-030-92068-5_9.

[DKLPW20] L. De Feo, D. Kohel, A. Leroux, C. Petit, and B. Wesolowski. "SQISign: Compact Post-quantum Signatures from Quaternions and Isogenies". In: *Advances in Cryptology – ASIACRYPT 2020*. Ed. by S. Moriai and H. Wang. Cham: Springer International Publishing, 2020, pp. 64–93. ISBN: 978-3-030-64837-4.

[DLW22] L. De Feo, A. Leroux, and B. Wesolowski. *New algorithms for the Deuring correspondence: SQISign twice as fast*. Cryptology ePrint Archive, Paper 2022/234. 2022. URL: https://eprint.iacr.org/2022/234.

[DG16] C. Delfs and S. D. Galbraith. "Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$". In: *Designs, Codes and Cryptography* 78.2 (2016), pp. 425–440. DOI: 10.1007/s10623-014-0010-1.

[Deu41] M. Deuring. "Die Typen der Multiplikatorenringe elliptischer Funktionenkörper". In: *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 14 (1941), pp. 197–272.

[EHLMP18] K. Eisenträger, S. Hallgren, K. Lauter, T. Morrison, and C. Petit. "Supersingular Isogeny Graphs and Endomorphism Rings: Reductions and Solutions". In: *Advances in Cryptology – EUROCRYPT 2018*. Ed. by J. B. Nielsen and V. Rijmen. Cham: Springer International Publishing, 2018, pp. 329–368. ISBN: 978-3-319-78372-7.

[FS87] A. Fiat and A. Shamir. "How To Prove Yourself: Practical Solutions to Identification and Signature Problems". In: *Advances in Cryptology — CRYPTO' 86*. Ed. by A. M. Odlyzko. Berlin, Heidelberg: Springer Berlin Heidelberg, 1987, pp. 186–194. ISBN: 978-3-540-47721-1.

[GPS16] S. D. Galbraith, C. Petit, and J. Silva. *Identification Protocols and Signature Schemes Based on Supersingular Isogeny Prob-*

lems. Cryptology ePrint Archive, Report 2016/1154. 2016. URL: https://eprint.iacr.org/2016/1154.

[Gro96] L. K. Grover. "A Fast Quantum Mechanical Algorithm for Database Search". In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. STOC '96. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, 1996, pp. 212–219. ISBN: 0897917855. DOI: 10.1145/237814.237866.

[HW08] G. H. Hardy and E. M. Wright. *An introduction to the theory of numbers*. 6th. Oxford University Press, 1938 (2008), p. 622. ISBN: 978-0-19-921986-5.

[HL10] C. Hazay and Y. Lindell. *Efficient Secure Two-Party Protocols: Techniques and Constructions*. 1st. Berlin, Heidelberg: Springer-Verlag, 2010. ISBN: 3642143024.

[Igu72] J.-i. Igusa. *Theta functions*. Die Grundlehren der mathematischen Wissenschaften, Band 194. New York: Springer-Verlag, 1972, pp. x+232.

[JD11] D. Jao and L. De Feo. "Towards Quantum-Resistant Cryptosystems from Supersingular Elliptic Curve Isogenies". In: *Post-Quantum Cryptography*. Ed. by B.-Y. Yang. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 19–34. ISBN: 978-3-642-25405-5.

[Kan97] E. Kani. "The number of curves of genus two with elliptic differentials". In: *Journal für die reine und angewandte Mathematik* 1997.485 (1997), pp. 93–122. DOI: 10.1515/crll.1997.485.93.

[Kan14] E. Kani. "The moduli spaces of Jacobians isomorphic to a product of two elliptic curves". In: *Journal of Number Theory* 139 (June 2014), pp. 138–174. DOI: 10.1016/j.jnt.2013.12.006.

[KLPT14] D. Kohel, K. Lauter, C. Petit, and J.-P. Tignol. *On the quaternion ℓ-isogeny path problem*. 2014. arXiv: 1406.0981 [math.NT].

[Lag70] J. L. de Lagrange. "Démonstration d'un théoreme d'arithmétique". In: *Nouveau Mémoire de l'Académie Royale des Sciences de Berlin* (1770), pp. 123–133.

[Ler22] A. Leroux. *Quaternion algebras and isogeny-based cryptography*. 2022. URL: http://www.lix.polytechnique.fr/Labo/Antonin.LEROUX/manuscrit_these.pdf.

[LR12] D. Lubicz and D. Robert. "Computing isogenies between abelian varieties". In: *Compositio Mathematica* 148.5 (Sept. 2012), pp. 1483–1515. DOI: 10.1112/S0010437X12000243.

[LR15] D. Lubicz and D. Robert. "Computing separable isogenies in quasi-optimal time". In: *LMS Journal of Computation and Mathematics* 18.1 (2015), pp. 198–216. DOI: 10.1112/S146115701400045X.

[LR23] D. Lubicz and D. Robert. "Fast change of level and applications to isogenies". In: vol. 9. 1. Springer, 2023. DOI: 10.1007/s40993-022-00407-9.

[MMPPW23] L. Maino, C. Martindale, L. Panny, G. Pope, and B. Wesolowski. "A Direct Key Recovery Attack on SIDH". In: Springer-Verlag, 2023.

[Mic16] D. Micciancio. *Minkowski's theorem.* In *CSE 206A: Lattice Algorithms and Applications.* 2016. URL: https://cseweb.ucsd.edu/classes/wi16/cse206A-a/lec2.pdf.

[Mil86] J. S. Milne. "Abelian Varieties". In: *Arithmetic Geometry.* New York, NY: Springer New York, 1986, pp. 103–150. ISBN: 978-1-4613-8655-1. DOI: 10.1007/978-1-4613-8655-1_5.

[Mum66] D. Mumford. "On the equations defining abelian varieties. I". In: *Invent. Math.* 1 (1966), pp. 287–354.

[Mum67a] D. Mumford. "On the equations defining abelian varieties. II". In: *Invent. Math.* 3 (1967), pp. 75–135.

[Mum67b] D. Mumford. "On the equations defining abelian varieties. III". In: *Invent. Math.* 3 (1967), pp. 215–244.

[Mum74] D. Mumford. *Abelian varieties.* Second Edition. Tata Institute of fundamental research studies in mathematics. London: Oxford University Press, 1974, pp. x+279.

[Mum84] D. Mumford. *Tata lectures on theta II.* Vol. 43. Progress in Mathematics. Jacobian theta functions and differential equations, With the collaboration of C. Musili, M. Nori, E. Previato, M. Stillman and H. Umemura. Boston, MA: Birkhäuser Boston Inc., 1984, pp. xiv+272. ISBN: 0-8176-3110-0.

[Mum] D. Mumford. "On the equations defining abelian varieties 1". In: *Inventiones mathematicae* 1.4 (), pp. 287–354. DOI: 10.1007/BF01389737.

[Piz80] A. Pizer. "An algorithm for computing modular forms on $\Gamma_0(N)$". In: *Journal of Algebra* 64 (June 1980), pp. 340–340. DOI: 0.1016/0021-8693(80)90151-9.

[PH78] S. C. Pohlig and M. E. Hellman. "An improved algorithm for computing logarithms over GF(p) and its cryptographic significance". In: *IEEE Transactions on information theory* 24.1 (Jan. 1978), pp. 106–110.

[PT18] P. Pollack and E. Treviño. "Finding the Four Squares in Lagrange's Theorem". In: *Integers* 18A (2018), A15.

[Rab86] J. O. Rabin M. O.; Shallit. "Randomized Algorithms in Number Theory". In: *Communications on Pure and Applied Mathematic* 39.S1 (1986), S239–S256. DOI: 10.1002/cpa.3160390713.

[Rob10] D. Robert. "Theta functions and cryptographic applications". PhD thesis. Université Henri-Poincarré, Nancy 1, France, July 2010. URL: http://www.normalesup.org/~robert/pro/publications/academic/phd.pdf.

[Rob21] D. Robert. "Efficient algorithms for abelian varieties and their moduli spaces". HDR thesis. Université Bordeaux, June 2021.

URL: http://www.normalesup.org/~robert/pro/publications/academic/hdr.pdf.

[Rob22a] D. Robert. *Breaking SIDH in polynomial time*. Cryptology ePrint Archive, Paper 2022/1038. 2022. URL: https://eprint.iacr.org/2022/1038.

[Rob22b] D. Robert. *Evaluating isogenies in polylogarithmic time*. Cryptology ePrint Archive, Paper 2022/1068. 2022. URL: https://eprint.iacr.org/2022/1068.

[RS06] A. Rostovtsev and A. Stolbunov. *Public-Key Cryptosystem Based On Isogenies*. Cryptology ePrint Archive, Report 2006/145. 2006. URL: https://eprint.iacr.org/2006/145.

[RT22] J. Rouse and K. Thompson. *Quaternary quadratic forms with prime discriminant*. 2022. arXiv: 2206.00412 [math.NT].

[Sil09] J. H. Silverman. *The Arithmetic of Elliptic Curves*. Springer, 2009, p. 522.

[Tes99] E. Teske. "The Pohlig-Hellman Method Generalized for Group Structure Computation". In: *Journal of Symbolic Computation* 11 (1999), pp. 1–14.

[Vél71] J. Vélu. "Isogénies entre courbes elliptiques". In: *Comptes-rendus de l'Académie des Sciences* 273 (July 1971). Available at https://gallica.bnf.fr, pp. 238–241.

[VV15] D. Venturi and A. Villani. *Zero-Knowledge Proofs and Applications*. May 2015. URL: http://danieleventuri.altervista.org/files/zero-knowledge.pdf.

[Voi20] J. Voight. *Quaternion algebras*. v.0.9.23. Aug. 2020. URL: https://math.dartmouth.edu/~jvoight/quat.html.

[Wes22] B. Wesolowski. "The supersingular isogeny path and endomorphism ring problems are equivalent". In: *FOCS 2021 - 62nd Annual IEEE Symposium on Foundations of Computer Science*. Denver, Colorado, United States, Feb. 2022. URL: https://hal.archives-ouvertes.fr/hal-03340899.

[ZSPDB18] G. Zanon, M. A. Simplicio, G. Pereira, J. Doliskani, and P. L. Barreto. "Faster isogeny-based compressed key agreement". In: *Post-Quantum Cryptography*. Springer International Publishing, 2018, pp. 248–268.

A Omitted proofs

A.1 Kani's lemma (Lemma 3.2.3)

Lemma 3.2.3 (Kani)**.** Consider the following (a, b)-isogeny diamond

$$\begin{tikzcd}
A' \arrow[r, "\varphi'"] & B' \\
A \arrow[u, "\psi"] \arrow[r, "\varphi"] & B \arrow[u, "\psi'"']
\end{tikzcd}$$

with $d := a + b$ prime to the characteristic of the base field of abelian varieties. Then, the isogeny $F : A \times B' \longrightarrow B \times A'$ given in matrix notation by

$$F := \begin{pmatrix} \varphi & \widetilde{\psi'} \\ -\psi & \widetilde{\varphi'} \end{pmatrix}$$

is a d-isogeny with $d = a + b$.

If a and b are coprime, the kernel of F is

$$\ker(F) = \{(\widetilde{\varphi}(x), \psi'(x)) \mid x \in B[d]\}.$$

Proof. Here we use, without proof, the main properties of the dual abelian variety with respect to polarisations (see for instance [Kan14, § 11]). It is a classical result that the dual of a matrix for the product polarisations is the transpose of the matrix obtained after dualizing the coefficients and that dualization is an involutive operation, so that

$$\tilde{F} = \begin{pmatrix} \tilde{\varphi} & -\widetilde{\psi} \\ \psi' & \varphi' \end{pmatrix}.$$

Hence

$$\tilde{F}F = \begin{pmatrix} \widetilde{\varphi}\varphi + \widetilde{\psi}\psi & \widetilde{\varphi}\widetilde{\psi'} - \widetilde{\psi}\widetilde{\varphi'} \\ \psi'\varphi - \varphi'\psi & \psi'\widetilde{\psi'} + \varphi'\widetilde{\varphi'} \end{pmatrix}$$

with $\widetilde{\varphi}\varphi + \widetilde{\psi}\psi = [a]_A + [b]_A = [d]_A$ since φ is an a-isogeny and ψ is a b-isogeny. Similarly, we get $\widetilde{\psi'}\psi' + \widetilde{\varphi'}\varphi' = [d]_{A'}$, so that $\psi'\widetilde{\psi'} + \varphi'\widetilde{\varphi'} = [d]_{B'}$ after dualizing (the dual being anti-commutative and the dual of an integer being an integer). Clearly, $\psi'\varphi - \varphi'\psi = 0$ since we have an isogeny diamond and we obtain $\widetilde{\varphi}\widetilde{\psi'} - \widetilde{\psi}\widetilde{\varphi'} = 0$ by dualizing the preceding equality. Hence, $\tilde{F}F = [d]_{A \times B'}$ so F is a d-isogeny.

If $x \in B[d]$, we have

$$\begin{aligned} F(\widetilde{\varphi}(x), \psi'(x)) &= (\varphi \circ \widetilde{\varphi}(x) + \widetilde{\psi'} \circ \psi'(x), -\psi \circ \widetilde{\varphi}(x) + \widetilde{\varphi'} \circ \psi'(x)) \\ &= ([a]x + [b]x, 0) = ([d]x, 0) = (0, 0) \end{aligned}$$

where we used the fact that $\psi \circ \widetilde{\varphi} = \widetilde{\varphi'} \circ \psi'$. Indeed, $\psi'\varphi = \varphi'\psi$, which implies that $[a]\psi \circ \widetilde{\varphi} = [a]\widetilde{\varphi'} \circ \psi'$ after multiplying on the right by $\widetilde{\varphi}$ and on the left by $\widetilde{\varphi'}$, so that $\psi \circ \widetilde{\varphi} = \widetilde{\varphi'} \circ \psi'$ since $[a]_{A'}$ is an isogeny (it has finite kernel).

It follows that $\ker(F)$ contains the set:

$$S := \{(\widetilde{\varphi}(x), \psi'(x)) \mid x \in B[d]\}.$$

Since $\widetilde{\varphi}$ and ψ' are a and b-isogenies respectively, we have $\ker(\widetilde{\varphi}) \subseteq B[a]$ and $\ker(\psi') \subseteq B[b]$. It follows that $\ker(\widetilde{\varphi}) \cap \ker(\psi') = \{0\}$, when a and b are coprime, so that $x \in B \longmapsto (\widetilde{\varphi}(x), \psi'(x))$ is injective and $\#S = \#B[d]$. Since d is coprime to the characteristic of the field, we have $\#B[d] = d^{2g}$ with $g := \dim(B)$. A being isogenous to B and B', we also have $g = \dim(A) = \dim(B')$ and $\dim(F) = d^{\dim(A \times B')} = d^{2g}$. Hence, $\ker(F) = S$ when a and b are coprime and the proof is complete. □

Actually, we can prove a weak converse of Kani's lemma.

Lemma A.1.1. *Let $F : A \times B' \longrightarrow B \times A'$ be a d-isogeny (for the product principal polarizations), where d is coprime to the characteristic of the base field. Write F as a matrix:*

$$F := \begin{pmatrix} \varphi & \widetilde{\psi'} \\ -\psi & \widetilde{\varphi'} \end{pmatrix}$$

and suppose φ is an a-isogeny. Then $\varphi, \varphi', \psi, \psi'$ form the following (a, b)-isogeny diamond with $a + b = d$:

$$\begin{tikzcd}
A' \arrow[r, "\varphi'"] & B' \\
A \arrow[u, "\psi"] \arrow[r, "\varphi"] & B \arrow[u, "\psi'"']
\end{tikzcd}$$

Proof. Since F is a d-isogeny, we have $\widetilde{F} \circ F = [d]$, so we get that $\psi' \circ \varphi = \varphi' \circ \psi$, $\widetilde{\varphi} \circ \varphi + \widetilde{\psi} \circ \psi = [d]$ and $\psi' \circ \widetilde{\psi'} + \varphi' \circ \widetilde{\varphi'} = [d]$. Since, φ is an a-isogeny, we have $\widetilde{\varphi} \circ \varphi = [a]$, so $\widetilde{\psi} \circ \psi = [b]$ and ψ is a b-isogeny.

We also have $F \circ \widetilde{F} = [d]$, so that $\varphi \circ \widetilde{\varphi} + \psi' \circ \widetilde{\psi'} = [d]$ and we get that ψ' is a b-isogeny. Then, the equality $\psi' \circ \widetilde{\psi'} + \varphi' \circ \widetilde{\varphi'} = [d]$ ensures that φ' is an a-isogeny. This completes the proof. □

Remark A.1.2. We obtain the same result if we suppose that any of the isogenies φ', ψ and ψ' is an a or b-isogeny.

A.2 Finding a uniformly random tight response ideal (Lemmas 5.2.1 and 5.2.2)

Lemma 5.2.1. Let f be a primitive positive definite integral quadratic form in k variables and let $\rho > 0$. Then there exists an algorithm that samples uniformly random elements from the set

$$\{x \in \mathbb{Z}^k \mid f(x) \leq \rho\}$$

in polynomial time in $\log(\rho)$ and the length of f (namely, the maximal number of bits of the coefficients of f). This algorithm runs in exponential time in k.

Proof. By Cholesky decomposition theorem, there exists a matrix $B \in GL_k(\mathbb{R})$ such that $f(x) = \|Bx\|^2$ for all $x \in \mathbb{R}^k$, $\|.\|$ being the Euclidean norm. Let $\varLambda := \varLambda(B)$ be the lattice generated by the columns of B. We want to sample in $B(0, \sqrt{\rho}) \cap \varLambda$ with uniform distribution. Let $(b_1, \cdots, b_k)$ be an LLL-reduced basis of $\varLambda$. Let $\nu := \sqrt{k}\|b_k\|/2$ and consider the following sampling algorithm:

1. Sample $v \in B(0, \sqrt{\rho} + \nu)$ uniformly at random.
2. Find a solution $\lambda(v) \in \varLambda$ to the closest vector problem for v.
3. If $\lambda(v) \in B(0, \sqrt{\rho})$, return $\lambda(v)$; else restart.

We prove that the output $\lambda(v)$ has uniform distribution in $B(0, \sqrt{\rho}) \cap \varLambda$. Let $\mathcal{V} := \{v \in \mathbb{R}^k \mid \|v\| = \min_{\lambda \in \varLambda} \|v + \lambda\|\}$ be the Voronoi cell at the origin. Then, the closest vector $\lambda(v)$ satisfies $v \in \mathcal{V} + \lambda(v)$ and $\lambda(v)$ is unique when v is not at the border of a Voronoi cell, so it is unique with probability 1. Hence, for all $u \in B(0, \sqrt{\rho}) \cap \varLambda$,

$$\mathbb{P}(\lambda(v) = u) = \frac{\mathrm{Vol}((\mathcal{V} + u) \cap B(0, \sqrt{\rho} + \nu))}{\mathrm{Vol}(B(0, \sqrt{\rho} + \nu))}.$$

Let $\mu := \inf\{r > 0 \mid \forall v \in \mathbb{R}^k,\ \exists \lambda \in \varLambda, \|x - \lambda\| \leq r\}$ be the covering radius of $\varLambda$. Then $\mathcal{V} \subseteq B(0, \mu)$ and $\mu \leq \sqrt{k}\lambda_k/2$ where λ_k is the last minimum of $\varLambda$ (this is a classical result, see for instance [Mic16, Exercise 11]), so that $\mu \leq \nu$. It follows that $\mathcal{V} + u \subseteq B(0, \sqrt{\rho} + \nu)$ for all $u \in B(0, \sqrt{\rho}) \cap \varLambda$. Hence

$$\mathbb{P}(\lambda(v) = u) = \frac{\mathrm{Vol}(\mathcal{V} + u)}{\mathrm{Vol}(B(0, \sqrt{\rho} + \nu))} = \frac{\mathrm{Vol}(\mathcal{V})}{\mathrm{Vol}(B(0, \sqrt{\rho} + \nu))}.$$

This quantity does not depend on u so $\lambda(v)$ has uniform distribution.

We finally check that the algorithm terminates after an expected polynomial number of steps in $\log(\rho)$ and the length of f. Indeed, when v is uniform in $B(0, \sqrt{\rho} + \nu)$, we have

$$\mathbb{P}(\|\lambda(v)\| \leq \sqrt{\rho}) \geq \mathbb{P}(\lambda(v) = 0) = \frac{\mathrm{Vol}(B(0, \lambda_1/2))}{\mathrm{Vol}(B(0, \sqrt{\rho} + \nu))} = \left(\frac{\lambda_1}{2\sqrt{\rho} + \sqrt{k}\|b_k\|}\right)^k,$$

where λ_1 is the first minimum of $\varLambda$. Since f is integral, we have $\|b_k\| \geq \cdots \geq \|b_1\| \geq \lambda_1 \geq 1$ and since $(b_1, \cdots, b_k)$ is LLL-reduced, we have:

$$\|b_1\| \cdots \|b_k\| \leq 2^{k(k-1)/4}\,\mathrm{Covol}(\varLambda) = 2^{k(k-1)/4}\,\mathrm{disc}(f),$$

so $\|b_k\| \leq 2^{k(k-1)/4}\,\mathrm{disc}(f)$. Hence, the algorithm terminates after an expected number of steps $O(1/\log \mathbb{P}(\|\lambda(v)\| \leq \sqrt{\rho}))$, a quantity that is polynomial in $\log(\rho)$, $\log(\mathrm{disc}(f))$ (itself polynomial in the length of f) and k. We conclude that the algorithm has the desired complexity, since the LLL algorithm has polynomial time complexity and the closest vector problem (used in step 2 of the algorithm) can be solved in the desired time complexity. □

Lemma 5.2.2. Let $\mathcal{O}$ be a maximal order and J be a left $\mathcal{O}$-ideal. Then there exists $\alpha \in J$ such that $q_J(\alpha) \leq 2\sqrt{2p}/\pi$.

Proof. Consider the canonical embedding $\iota : \mathcal{B}_{p,\infty} \hookrightarrow \mathbb{R}^4$:

$$1 \longmapsto (1,0,0,0), i \longmapsto (0,\sqrt{q_0},0,0), j \longmapsto (0,0,\sqrt{p},0), k \longmapsto (0,0,0,\sqrt{q_0 p}),$$

where $q_0 := \operatorname{nrd}(i)$. ι is an isometry in the following sense $\|\iota(\alpha)\|^2 = \operatorname{nrd}(\alpha)$ for all $\alpha \in \mathcal{B}_{p,\infty}$, where $\|.\|$ is the Euclidean norm of $\mathbb{R}^4$. By Minkowski's second theorem, the successive minima of the lattice $\iota(J)$ satisfy

$$\lambda_1 \cdots \lambda_4 \leq 2^4 \frac{\operatorname{Covol}(\iota(J))}{\operatorname{Vol}(B(0,1))} = \frac{32}{\pi^2} \operatorname{Covol}(\iota(J)). \tag{2}$$

But we have $\operatorname{Covol}(\iota(\mathcal{O})) = \operatorname{discrd}(\mathcal{O})/4$. Indeed, if $(\alpha_1, \cdots, \alpha_4)$ is a basis of $\mathcal{O}$, we get that

$$\begin{aligned}\operatorname{Covol}(\iota(\mathcal{O})) &= \sqrt{|\det(\langle \iota(\alpha_i), \iota(\alpha_j)\rangle)_{1\leq i,k\leq 4}|} \\ &= \sqrt{\left|\det\left(\frac{1}{2}(\operatorname{nrd}(\alpha_i+\alpha_j) - \operatorname{nrd}(\alpha_i) - \operatorname{nrd}(\alpha_j))\right)_{1\leq i,k\leq 4}\right|} \\ &= \sqrt{\left|\det\left(\frac{1}{2}\operatorname{Tr}(\alpha_i\overline{\alpha_j})\right)_{1\leq i,k\leq 4}\right|} = \frac{1}{4}\operatorname{discrd}(\mathcal{O})\end{aligned}$$

Besides, by [Voi20, Theorem 15.5.5], $\operatorname{discrd}(\mathcal{O}) = \operatorname{disc}(\mathcal{B}_{p,\infty}) = p$ since $\mathcal{O}$ is a maximal order. We then have

$$\operatorname{Covol}(\iota(J)) = [\mathcal{O} : I] \operatorname{Covol}(\iota(\mathcal{O})) = [\mathcal{O} : I] \operatorname{discrd}(\mathcal{O})/4 = \operatorname{nrd}(J)^2 p/4$$

It follows that the minimal value of q_J is

$$q_J(\alpha) = \lambda_1^2/\operatorname{nrd}(J) \leq \frac{2\sqrt{2p}}{\pi}.$$

□

A.3 Dividing the higher dimensional isogeny in two (Proposition 5.4.1)

Proposition 5.4.1. Let $d := d_1 d_2$ coprime to p and $F : \mathcal{A} \longrightarrow \mathcal{B}$ be a d-isogeny between abelian varieties defined over $\overline{\mathbb{F}_p}$. Then:

(i) We can always decompose $F = F_2 \circ F_1$, where F_1 is a d_1-isogeny and F_2 is a d_2-isogeny.

(ii) $\ker(F_1) \subseteq \ker(F) \cap \mathcal{A}[d_1]$.

(iii) $\ker(\widetilde{F_2}) \subseteq \ker(\widetilde{F}) \cap \mathcal{B}[d_2] = F(\mathcal{A}[d]) \cap \mathcal{B}[d_2]$.

(iv) When $\ker(F)$ has rank $g := \dim(\mathcal{A})$, those inclusions are equalities.

In order to prove the proposition, we need two intermediary results.

Lemma A.3.1. *If $F : (\mathcal{A}, \lambda_{\mathcal{A}}) \longrightarrow (\mathcal{B}, \lambda_{\mathcal{B}})$ is a d-isogeny between principally polarized abelian varieties, then $\ker(F)$ is a maximal isotropic subgroup of $\mathcal{A}[d]$ (for the d-th Weil pairing).*

Proof. The inclusion $\ker(F) \subseteq \mathcal{A}[d]$ immediately follows from $\widetilde{F} \circ F = [d]$. Now we prove that $\ker(F)$ is isotropic. Let $x, y \in \ker(F)$. Since $\widetilde{F}$ is surjective, there exists $y' \in \mathcal{B}$ such that $y = \widetilde{F}(y')$. Let $\lambda_{\mathcal{A}}$ and $\lambda_{\mathcal{B}}$ be the principal polarisations on $\mathcal{A}$ and $\mathcal{B}$ respectively, $e_d^{\lambda_{\mathcal{A}}}$ and $e_d^{\lambda_{\mathcal{B}}}$ the associated Weil pairings. Then

$$e_d^{\lambda_{\mathcal{A}}}(x, y) = e_d(x, \lambda_{\mathcal{A}} \circ \widetilde{F}(y')) = e_d(x, \widehat{F} \circ \lambda_{\mathcal{B}}(y)) = e_d(F(x), \lambda_{\mathcal{B}}(y)) = 1.$$

Then $\ker(F)$ is isotropic. Since F is a d-isogeny, it has degree d^g with $g := \dim(\mathcal{A})$, and $\# \ker(F) = d^g$ since F is separable. So $\ker(F)$ is maximal isotropic. □

Lemma A.3.2. *let $(\mathcal{A}, \lambda)$ be a principally polarized abelian variety. If $K \subset \mathcal{A}[d]$ is isotropic, then the polarization $[d]\lambda_{\mathcal{L}}$ on $\mathcal{A}$ descends to a principal polarization on $\mathcal{B} := \mathcal{A}/K$. More precisely, there exists a principal polarization λ' on $\mathcal{B}$ such that $[d]\lambda = \widehat{\pi} \circ \lambda' \circ \pi$, where $\pi : \mathcal{A} \longrightarrow \mathcal{B} = \mathcal{A}/K$ is the canonical projection.*

Proof. We have $\ker([d]\lambda) = [d]^{-1}(\ker(\lambda)) = \mathcal{A}[d]$ since $\deg(\lambda) = 1$. Since $K \subset \mathcal{A}[d]$ is isotropic, the result follows from [Mil86, Proposition 6.8]. □

Proof. (of Proposition 5.4.1) **(i)** We prove that we have a decomposition of $F : \mathcal{A} \longrightarrow \mathcal{B}$ of the form

$$(\mathcal{A}, \lambda_{\mathcal{A}}) \xrightarrow{F_1} (\mathcal{C}, \lambda_{\mathcal{C}}) \xrightarrow{F_2} (\mathcal{B}, \lambda_{\mathcal{B}}),$$

where the intermediary abelian variety $\mathcal{C}$ is principally polarized and F_1 and F_2 are respectively d_1 and d_2-isogenies, with $d = d_1 d_2$. By induction, it suffices to prove this result when $d_1 = \ell$ is a prime.

Since $\# \ker(F) = d^g$, the ℓ-Sylow subgroup of $\ker(F)$ has cardinality $\ell^{g v_\ell}$, where v_ℓ is the ℓ-adic valuation of d. We may also write the ℓ-Sylow subgroup as follows:

$$G_\ell \cong \prod_{i=1}^{r} (\mathbb{Z}/\ell^{\alpha_i}\mathbb{Z}),$$

where the α_i are positive integers. Since $G_\ell \subseteq \ker(F) \subseteq \mathcal{A}[d]$, the α_i must all be $\leq v_\ell$. It follows that

$$g v_\ell = \log_\ell \# G_\ell = \sum_{i=1}^{r} \alpha_i \leq r v_\ell,$$

so that $r \geq g$ and $\ker(F)[\ell] = G_\ell[\ell] \cong (\mathbb{Z}/\ell\mathbb{Z})^r$, which implies that $\ker(F)[\ell]$ has rank $\geq g$. Hence, it contains a subgroup $K \subset \mathcal{A}[\ell]$ of rank g, which is isotropic in $\mathcal{A}[\ell]$ since $\ker(F) \supset K$ is isotropic in $\mathcal{A}[d]$ by Lemma A.3.1.

By [Mum74, Theorem 4, p. 73], F factors through an isogeny F_1 of kernel K so we can indeed write $F = F_2 \circ F_1$. Since K is isotropic, by Lemma A.3.2, the codomain $\mathcal{C}$ of F_1 admits a principal polarization $\lambda_{\mathcal{C}}$ such that $\widehat{F_1} \circ \lambda_{\mathcal{C}} \circ F_1 = [\ell]\lambda_{\mathcal{A}}$ *i.e.* $\widetilde{F_1} \circ F_1 = [\ell]$. So F_1 is an ℓ-isogeny.

We also have

$$[d]_{\mathcal{B}} = F \circ \widetilde{F} = F_2 \circ F_1 \circ \widetilde{F_1} \circ \widetilde{F_2} = F_2 \circ [\ell]_{\mathcal{C}} \circ \widetilde{F_2} = F_2 \circ \widetilde{F_2} \circ [\ell]_{\mathcal{B}},$$

so $F_2 \circ \widetilde{F_2} = [d/\ell]_{\mathcal{B}}$ since $[\ell]_{\mathcal{B}}$ is surjective and F_2 is a $d_2 = d/\ell$-isogeny. To prove the result in the general case, we can proceed by induction on the degree d (and factor F_2).

(ii) We always have $\ker(F_1) \subseteq \ker(F) \cap \mathcal{A}[d_1]$ since $F = F_2 \circ F_1$ and F_1 is a d_1-isogeny.

(iii) Similarly, we get that $\ker(\widetilde{F_2}) \subseteq \ker(\widetilde{F}) \cap \mathcal{B}[d_2]$. But $\widetilde{F} \circ F = [d]_{\mathcal{A}}$ so $F(\mathcal{B}[d]) \subseteq \ker(\widetilde{F})$. Since $\ker(F) \subseteq \mathcal{A}[d]$, we have an isomorphism $F(\mathcal{A}[d]) \cong \mathcal{A}[d]/\ker(F)$. Furthermore, d being coprime to p, $\#\mathcal{A}[d] = d^{2g}$ and $\#\ker(\widetilde{F}) = \#\ker(F) = d^g$. Hence, $\#F(\mathcal{B}[d]) = d^g = \#\ker(\widetilde{F})$ and $\ker(\widetilde{F}) = F(\mathcal{A}[d])$. Point (iii) follows.

(iv) Suppose $\ker(F)$ has rank g. Let $(x_1, \cdots, x_g)$ be a basis of generators of $\ker(F)$. Then, the x_i all have order d and are linearly independent over $\mathbb{Z}/d\mathbb{Z}$. Then, $\ker(F) \cap \mathcal{A}[d_1] = ([d_2]x_1, \cdots, [d_2]x_g)$ and the $[d_2]x_i$ all have order d_1 and are linearly independent over $\mathbb{Z}/d_1\mathbb{Z}$, so that $\#\ker(F) \cap \mathcal{A}[d_1] = d_1^g = \#\ker(F_1)$. The equality $\ker(F_1) = \ker(F) \cap \mathcal{A}[d_1]$ follows by (ii).

To prove that $\ker(\widetilde{F_2}) = F(\mathcal{A}[d]) \cap \mathcal{B}[d_2]$, it suffices to prove that $F(\mathcal{A}[d]) = \ker(\widetilde{F})$ has rank g and the preceding reasoning will apply. We have $F(\mathcal{A}[d]) \cong \mathcal{A}[d]/\ker(F)$, so it suffices to prove that there is a subgroup $G \subset \mathcal{A}[d]$ of rang g such that $\ker(F) \oplus G = \mathcal{A}[d]$, because, we will have $F(\mathcal{A}[d]) \cong G$. We consider the d-th Weil pairing $e_d^{\lambda_{\mathcal{A}}} : \mathcal{A}[d]^2 \longrightarrow \mu_d(\overline{\mathbb{F}_p})$, where $\mu_d(\overline{\mathbb{F}_p})$ is the group of d-th roots of unity in $\overline{\mathbb{F}_p}$, and the group homomorphism

$$\begin{aligned} \varPhi : \mathcal{A}[d] &\longrightarrow \mu_d(\overline{\mathbb{F}_p})^g \\ y &\longmapsto (e_d^{\lambda_{\mathcal{A}}}(x_i, y))_{1 \leq i \leq g}. \end{aligned}$$

Since $\ker(F)$ is maximal isotropic, we have $\ker(\varPhi) = \ker(F)$. It follows that $\#\operatorname{im}(\varPhi) = \#\mathcal{A}[d]/\#\ker(\varPhi) = d^{2g}/d^g = d^g$, so that $\operatorname{im}(\varPhi) = \mu_d(\overline{\mathbb{F}_p})^g$ *i.e.* $\varPhi$ is surjective. Let $\zeta \in \mu_d(\overline{\mathbb{F}_p})$ be a primitive d-th root of unity. Then, for all $j \in [\![1 \ ; \ g]\!]$ there exists $y_j \in \mathcal{A}[d]$ such that $e_d^{\lambda_{\mathcal{A}}}(x_i, y_j) = \zeta^{\delta_{i,j}}$ for all $i \in [\![1 \ ; \ g]\!]$. It follows that the y_i all have order d (since $e_d^{\lambda_{\mathcal{A}}}(x_i, y_i) = \zeta$ has order d), are linearly independent over $\mathbb{Z}/d\mathbb{Z}$ and linearly independent of the x_i. We can then take $G := \langle y_1, \cdots, y_g \rangle$. This completes the proof of $\ker(\widetilde{F_2}) = F(\mathcal{A}[d]) \cap \mathcal{B}[d_2]$. □

A.4 Correctness of $\mathsf{RepresentIsogeny}_{g,\ell^e,\ell^f}$ (Proposition 5.6.1)

Proposition 5.6.1. Algorithms 11 and 12 are correct. Namely, Algorithm 11 returns $F_1, \tilde{F}_2$ such that $F_2 \circ F_1 = F(\sigma, a_1, a_2)$ on entry $a_1, a_2, P_1, P_2, \sigma(P_2)$,

$\sigma(P_2)$, where $\sigma : E_A \longrightarrow E_2$ is a q-isogeny with $a_1^2 + a_2^2 + q = \ell^e$ (and similarly for Algorithm 12).

Proof. We assume that Algorithm 11 received $a_1, a_2, P_1, P_2, \sigma(P_2), \sigma(P_2)$ on entry. Let us write $F^* := F(\sigma, a_1, a_2)$ and decompose $F^* := F_2^* \circ F_1^*$, where F_1^* is an ℓ^{e_1}-isogeny of kernel $\ker(F^*)[\ell^{e_1}]$ and $\widetilde{F_2^*}$ is an ℓ^{e_2}-isogeny of kernel $F^*(E_A^2 \times E_2^2[\ell^{e_2}])$. Then, by construction the outputs F_1 and $\widetilde{F_2}$ of Algorithm 11 satisfy $\ker(F_1) = \ker(F_1^*)$, $\ker(\widetilde{F_2}) = \ker(\widetilde{F_2^*})$, $\deg(F_1) = \deg(F_1^*)$ and $\deg(\widetilde{F_2}) = \deg(\widetilde{F_2^*})$. Hence, by [Mum74, Theorem 4, p.73] if we denote $(\mathcal{C}, \lambda_\mathcal{C})$ the common codomain of F_1 and $\widetilde{F_2}$ and $(\mathcal{C}^*, \lambda_{\mathcal{C}^*})$ the common codomain of F_1^* and $\widetilde{F_2^*}$, then we get that $F_1 = \lambda \circ F_1^*$ and $\widetilde{F_2} = \mu \circ \widetilde{F_2^*}$, where λ and μ are isomorphisms $(\mathcal{C}^*, \lambda_{\mathcal{C}^*}) \xrightarrow{\sim} (\mathcal{C}, \lambda_\mathcal{C})$.

To conclude, we need to be more specific about the way F_1 and $\widetilde{F_2}$ are computed explicitly with the theta-model of level 2 (see Appendix C.2). To ensure F_1 and $\widetilde{F_2}$ define the same theta-structure of level 2, we use Algorithm 20 outputting two symplectic basis of $E_A^2 \times E_2^2[4\ell^{\max(e_1,e_2)}]$. The output of this algorithm only depends on the input $a_1, a_2, P_1, P_2, \sigma(P_2), \sigma(P_2)$, so it would be the same if we applied it for F_1^* and $\widetilde{F_2^*}$. It follows that $(\mathcal{C}, \lambda_\mathcal{C}) = (\mathcal{C}^*, \lambda_{\mathcal{C}^*})$ and that F_1 and F_1^* coincide on $F_1^{-1}(\mathcal{C}[4])$, so that λ and $\mathrm{id}_\mathcal{C}$ coincide on $\mathcal{C}[4]$, so that $\lambda = \mathrm{id}_\mathcal{C}$ by [Mil86, Proposition 17.5.(b)]. Similarly, $\widetilde{F_2}$ and $\widetilde{F_2^*}$ coincide on $\widetilde{F_2}^{-1}(\mathcal{C}[4])$, so that $\mu = \mathrm{id}_\mathcal{C}$. It follows that $F_2 \circ F_1 = F^* = F(\sigma, a_1, a_2)$, so Algorithm 11 is correct. The same proof applies to Algorithm 12. □

A.5 Verification (Proposition 5.7.1 and Corollary 5.7.2)

Proposition 5.7.1. Algorithms 13 and 14 are correct. Namely, when given $E_A, E_2, a_1, a_2, F_1, \tilde{F}_2$, if Algorithm 13 returns True, then $F_2 \circ F_1$ is an efficient representation of an isogeny $\sigma : E_A \longrightarrow E_2$ of degree $q = \ell^e - a_1^2 - a_2^2$ (and similarly for Algorithm 14).

Proof. Assume that Algorithm 13 returns True on input $E_A, E_2, a_1, a_2, F_1, \tilde{F}_2$. Then, $F := F_2 \circ F_1$ is well defined and is an ℓ^e-isogeny. We may write $F := (f_{i,j})_{1 \leq i,j \leq 4} \in \mathrm{End}(E_A^2 \times E_2^2)$. Then, since $\tilde{F} \circ F = [\ell^e]$, we get that

$$\forall 1 \leq j \leq 4, \quad \sum_{i=1}^{4} \deg(f_{i,j}) = \ell^e, \tag{3}$$

so the $f_{i,j}$ have degree $\leq \ell^e$. By assumption, there exists $Q \in E_A$ of order $\ell^f \ell'^{f'}$ such that $f_{1,1}(Q) = [a_1]Q_i$, $f_{2,1}(Q) = -[a_2]Q$ and $f_{4,1}(Q) = 0$. Besides, by Cauchy Schwarz inequality

$$\deg(f_{1,1} - [a_1]) \leq \left(\sqrt{\deg(f_{1,1})} + \sqrt{\deg([a_1])}\right)^2 \leq 4\ell^e < \ell^f \ell'^{f'}.$$

It follows that $f_{1,1} = [a_1]$. We similarly obtain that $f_{2,1} = -[a_2]$ and $f_{4,1} = 0$. Hence, by 3, we get $\deg(f_{3,1}) = \ell^e - a_1^2 - a_2^2 = q$. But $\sigma := f_{3,1}$ is an isogeny $E_A \longrightarrow E_2$. This proves the correctness of Algorithm 13.

We use a similar argument to prove the correctness of Algorithm 14. We write $F := (f_{i,j})_{1\leq i,j\leq 8}$. Then, the evaluation of the T-torsion points ensures that $f_{i,1} = [a_i]$ for all $i \in [\![1 \ ; \ 4]\!]$ and $f_{i,1}$ for all $i \in \{6, 7, 8\}$ since $T > 4\ell^e$ ($T \simeq p^3$ and $\ell^e \simeq p^2$). Using an analogue of 3, we conclude that $\sigma := f_{5,1} : E_A \longrightarrow E_2$ has degree $q = \ell^e - a_1^2 - a_2^2 - a_3^2 - a_4^2$. This completes the proof. □

Corollary 5.7.2. The verification procedures FastVerify and RigorousVerify (Algorithms 8 and 9) are correct. Namely, on input (R_1, R_2, q), FastVerify (respectively RigorousVerify) returns True if and only if (R_1, R_2, q) defines an efficient representation of an isogeny $\sigma : E_A \longrightarrow E_2$ of degree q, where q is ℓ^e-good and coprime to ℓ' (respectively $q < \ell^e$).

Proof. If FastVerify(R_1, R_2, q) returns True, it means that q is ℓ^e-good and coprime to ℓ, that FastVerify found $a_1, a_2 \in \mathbb{Z}$ such that $a_1^2 + a_2^2 + q = \ell^e$ and that $\mathsf{IsValid}_{4,\ell^e,\ell^f\ell'^{f'}}(E_A, E_2, a_1, a_2, *)$ returned True. In particular, by Proposition 5.7.1, it means that $\mathsf{RepresentIsogeny}_{4,\ell^e,\ell^f}(E_A, E_2, a_1, a_2, P_1, P_2, R_1, R_2)$ returned an efficient representation of a q-isogeny $\sigma : E_A \longrightarrow E_2$. Hence, the knowledge of (R_1, R_2, q), Cornacchia's algorithm to find a_1, a_2 and $\mathsf{RepresentIsogeny}_{4,\ell^e,\ell^f}$ (running in polynomial time in $\log(p)$) define an efficient representation of σ.

Conversely, if (R_1, R_2, q) defines an efficient representation of a q-isogeny $\sigma : E_A \longrightarrow E_2$, where q is ℓ^e-good and coprime to ℓ', then we can evaluate σ on the canonical basis (P_1, P_2) of $E_A[\ell^f]$ in polynomial time in $\log(p)$ and obtain $(R_1, R_2) := (\sigma(P_1), \sigma(P_2))$. On input (R_1, R_2, q), FastVerify will find $a_1, a_2 \in \mathbb{Z}$ such that $a_1^2+a_2^2+q = \ell^e$ and apply $\mathsf{RepresentIsogeny}_{4,\ell^e,\ell^f}(E_A, E_2, a_1, a_2, P_1, P_2, R_1, R_2)$, which will return $F_1, \widetilde{F_2}$ such that $F_2 \circ F_1 = F(\sigma, a_1, a_2)$ by Proposition 5.6.1. Hence, $\mathsf{IsValid}_{4,\ell^e,\ell^f\ell'^{f'}}$ will return True by construction, and so will FastVerify.

The same proof applies to RigorousVerify. □

A.6 Knowledge soundness of RigorousSQISignHD

We recall the formal definition of knowledge soundness given in [HL10, Definition 6.3.1].

Definition A.6.1. A protocol (P, V) between a prover and a verifier is a *proof of knowledge* for a relation $\mathcal{R} \subset X \times W$ with knowledge error κ if it satisfies the following properties:

Completeness: If P interacts with V as input $x \in X$ and private input $w \in W$ with $(x, w) \in \mathcal{R}$, then V always accepts.

Knowledge soundness: There exists a *knowledge extractor* K such that for every interactive prover P^* and every $x \in X$, K satisfies the following condition. Let $\varepsilon(x)$ be the success probability of P^* on input x (the probability that V accepts on input x). If $\varepsilon(x) > \kappa(x)$, then upon input x and oracle access to P^*, K outputs a witness $w \in W$ such that $(x, w) \in \mathcal{R}$ within an expected number of steps $O(1/(\varepsilon(x) - \kappa(x)))$.

Definition A.6.2. A 3-round protocol (commitment, challenge, response) (P, V) satisfies *special soundness* for a relation $\mathcal{R} \subset X \times W$ if given $x \in X$ and two accepting transcripts $(a, c, r), (a, c', r')$ for $x \in X$ with the same commitment a and distinct challenges $c \neq c'$, one can extract a witness $w \in W$ such that $(x, w) \in \mathcal{R}$ in polynomial time.

Theorem A.6.3. *[HL10, Theorem 6.3.2] A complete 3-round protocol satisfying special soundness for a relation $\mathcal{R}$ with challenge space $\mathcal{C}$ is a proof of knowledge with knowledge error $1/\#\mathcal{C}$.*

In Proposition 6.1.1, we proved that FastSQISignHD satisfies special soundness for the relation:

$$\mathcal{R} := \{(E_A, \alpha) \mid \alpha \in \mathrm{End}(E_A) \text{ non-scalar}\}.$$

Since the challenge space has size $\mu(\ell'^{f'}) = \ell'^{f'-1}(\ell' + 1) = \Omega(p^{1/2})$, we get by Theorem A.6.3 that FastSQISignHD is a proof of knowledge for $\mathcal{R}$ with knowledge soundness $O(p^{-1/2})$.

Unfortunately, the special soundness argument no longer holds in RigorousSQISignHD because we can no longer impose conditions on q (except $q < \ell^e$), and especially we cannot impose q to be coprime to D_φ. However, choosing D_φ big enough will ensure that the endomorphism α is non-scalar with overwhelming probability, since $\varphi' \circ \widehat{\varphi}$ has a big cyclic factor with overwhelming probability. We first introduce a useful lemma to prove this result.

Lemma A.6.4. *Let $\phi : E_1 \longrightarrow E_2$ and $\phi' : E_1 \longrightarrow E_2'$ be two cyclic isogenies. Then, there exists three cyclic isogenies $\phi_0 : E_1 \longrightarrow E_3$, $\phi_1 : E_3 \longrightarrow E_2$ and $\phi_1' : E_3 \longrightarrow E_2'$ such that $\phi = \phi_1 \circ \phi_0$, $\phi' = \phi_1' \circ \phi_0$ and $\phi_1' \circ \widehat{\phi_1}$ is cyclic. ϕ_0 will be called the* greatest common factor *of ϕ and ϕ'.*

Proof. Since the product of cyclic isogenies of coprime degrees is cyclic, we may assume that $\deg(\phi)$ and $\deg(\phi')$ are powers of the same prime ℓ. Let ϕ_0 be the biggest common factor of ϕ and ϕ' (possibly trivial). Then we may write $\phi := \phi_1 \circ \phi_0$ and $\phi' := \phi_1' \circ \phi_0$ where ϕ_1' and ϕ_1 have no common factor. We prove that $\phi_1' \circ \widehat{\phi_1}$ is cyclic by induction on the degree of ϕ_1'.

When $\deg(\phi_1') = 1$ it follows from the fact that the dual of a cyclic isogeny is cyclic. Now, we assume the result holds when $\deg(\phi_1') = \ell^n$ with $n \in \mathbb{N}$ and prove it holds when $\deg(\phi_1') = \ell^{n+1}$. We may factor $\phi_1' := \phi_2 \circ \phi_2'$ with $\deg(\phi_2) = \ell$ and $\deg(\phi_2') = \ell^n$. By assumption, $\phi_3 := \phi_2' \circ \widehat{\phi_1}$ is cyclic so we only have to prove that $\phi_2 \circ \phi_3$ is cyclic, *i.e.* that $\ker(\phi_2 \circ \phi_3) = \phi_3^{-1}(\ker(\phi_2))$ is cyclic.

Let Q be a generator of $\ker(\phi_2)$, P be a generator of $\ker(\phi_3)$ and $P' \in E_2$ such that $Q = \phi_3(P')$. Then

$$\ker(\phi_2 \circ \phi_3) = \phi_3^{-1}(\ker(\phi_2)) = \langle P, P' \rangle.$$

To conclude, it suffices to prove that $P \in \langle P' \rangle$. We have $P' \in \ker(\phi_2 \circ \phi_3) \subset E_2[\ell^{m+1}]$, with $\deg(\phi_3) := \ell^m$ and $[\ell^m]P' = \widehat{\phi_3} \circ \phi_3(P') = \widehat{\phi_3}(Q) \neq 0$ since $\widehat{\phi_3}$ does not factor through ϕ_2 (since does ϕ_1 does not either). Hence, P' has order ℓ^{m+1}. Let $R \in E_2[\ell^m]$ such that $([\ell]P', R)$ is a basis of $E_2[\ell^m]$. Then, we may write $P := [a\ell]P' + [b]R$ for some $a, b \in \mathbb{Z}$ since $P \in \ker(\phi_3) \subset E_2[\ell^m]$. Since $Q \in \ker(\phi_2)$ has order ℓ, we get that

$$0 = \phi_3(P) = [a\ell]Q + [b]\phi_3(R) = [b]\phi_3(R),$$

and $\phi_3(R)$ generates $\phi_3(E_2[\ell^m]) = \ker(\widehat{\phi_3})$ which is cyclic so it has order ℓ^m. It follows that $b \equiv 0 \mod \ell^m$, so that $P = [a\ell]P' \in \langle P' \rangle$. This completes the proof. □

Lemma A.6.5. *Let $(E_1, \varphi, R_1, R_2, q)$ and $(E_1, \varphi', R_1', R_2', q')$ be two* Rigorous-SQISignHD *transcripts with the same commitment E_1. If the greatest common factor of φ and φ' has degree $< D_\varphi/\ell^e$, then we can infer an efficient representation of a non-scalar endomorphism $\alpha \in \mathrm{End}(E_A)$ from these transcripts. In this case we say that φ and φ' are* relatively good *and* relatively bad *if this is not satisfied.*

Proof. As previously, let σ and σ' be respectively the isogenies defined on E_A represented by (R_1, R_2, q) and (R_1', R_2', q') and $\alpha := \widehat{\sigma}' \circ \varphi' \circ \widehat{\varphi} \circ \sigma \in \mathrm{End}(E_A)$. Assume that α is a scalar endomorphism: $\alpha = [\lambda]_{E_A}$ for some $\lambda \in \mathbb{Z}$. Then $[\lambda]_{E_2'} = \varphi' \circ \widehat{\varphi} \circ \sigma \circ \widehat{\sigma'}$. Let us write $\varphi := \varphi_1 \circ \varphi_0$ and $\varphi := \varphi_1' \circ \varphi_0$, where φ_0 is the greatest common factor of φ and φ'. Then $\phi := \varphi_1' \circ \widehat{\varphi_1}$ is cyclic by Lemma A.6.4 and we have $\varphi' \circ \widehat{\varphi} := [D]\phi$ with $D := \deg(\varphi_0)$. We can also write $\sigma' \circ \widehat{\sigma} = [D']\phi'$ where ϕ' is a cyclic isogeny $E_2 \longrightarrow E_2'$. It follows that $[\lambda/DD']_{E_2'} = \phi \circ \widehat{\phi}'$. Hence, by Lemma A.6.4, the greatest cyclic factor of ϕ and ϕ' must be equal to both ϕ and ϕ' so $\phi = \phi'$. Hence, $\sigma' \circ \widehat{\sigma}$ factors through ϕ. But $\deg(\sigma' \circ \widehat{\sigma}) = qq' \leq \ell^{2e}$ and $\deg(\phi) = D_\varphi^2/D^2$ with $D < D_\varphi/\ell^e$ so $\deg(\phi) > \ell^{2e}$. Contradiction. □

Now, we prove that the probability to generate *relatively good* challenges is overwhelming. This will be the last essential ingredient to our knowledge soundness proof.

Lemma A.6.6. *Fix a challenge $\varphi : E_1 \longrightarrow E_2$ and let us write $D_\varphi := \prod_{i=1}^r \ell_i^{e_i}$, where $\ell_1 \leq \cdots \leq \ell_r$ are distinct ordered primes and $e_1, \cdots, e_r \in \mathbb{N}^*$. Then, the number of challenges $\varphi' : E_1 \longrightarrow E_2'$ relatively bad to φ is*

$$O\left(\frac{\ell^e \mu(D_\varphi)}{D_\varphi^{1-\log(2)/\log\log(D_\varphi)}}\right),$$

with $\mu(D_\varphi) := \prod_{i=1}^r \ell_i^{e_i-1}(\ell_i+1)$.

Proof. φ and φ' relatively bad if their greatest common factor has degree $D \geq D_\varphi/\ell^e$. If we fix such a $D | D_\varphi$, then choosing φ' is choosing a cyclic isogeny of degree D_φ/D so there are $\mu(D_\varphi/D)$ possibilities. It follows that the number of challenges φ' relatively bad to φ is

$$\begin{aligned} N &\leq \sum_{\substack{D|D_\varphi \\ D>D_\varphi/\ell^e}} \mu\left(\frac{D_\varphi}{D}\right) = \mu(D_\varphi) \sum_{\substack{D|D_\varphi \\ D>D_\varphi/\ell^e}} \frac{1}{\mu(D)} \leq \mu(D_\varphi) \sum_{\substack{D|D_\varphi \\ D>D_\varphi/\ell^e}} \frac{1}{D} \\ &\leq \frac{\ell^e\mu(D_\varphi)}{D_\varphi} \#\{D \in \mathbb{N}^* \mid D|D_\varphi \text{ and } D > D_\varphi/\ell^e\} \\ &\leq \frac{\ell^e\mu(D_\varphi)}{D_\varphi} \#\{D \in \mathbb{N}^* \mid D|D_\varphi\} = \frac{\ell^e\mu(D_\varphi)d(D_\varphi)}{D_\varphi}, \end{aligned}$$

where $d(D_\varphi)$ is the number of divisors of $d(D_\varphi)$. By [HW08, § 18.1, Theorem 317], we know that $d(D_\varphi) = O\left(D_\varphi^{\log(2)/\log\log(D_\varphi)}\right)$. The result follows. □

Since $\ell^e = O(p^2)$, we choose $D_\varphi|T$ such that $D_\varphi \simeq p^{5/2(1-\log(2)/\log\log(p^{5/2}))}$, so that the proportion of challenges relatively bad to a given challenge is $O(p^{-1/2})$. This means in practice $D_\varphi \simeq p^{2.82}$ when p has size 256 bits (to achieve $\lambda = 128$ bits of classical security). Hence the choice $T \simeq p^3$. Then, under this condition, we can adapt the proof of [HL10, Theorem 6.3.2] to prove knowledge soundness of RigorousSQISignHD.

Proposition A.6.7. *Assume that $D_\varphi > p^{5/2(1-\log(2)/\log\log(p^{5/2}))}$. Then the* RigorousSQISignHD *identification protocol is a* proof of knowledge *for the relation $\mathcal{R}$ of Proposition 6.1.1 with knowledge error $O(p^{-1/2})$.*

Proof. As required by Definition A.6.1, we construct a knowledge extractor K. Let P^* be a prover with success probability ε. Then K is constructed as follows (as in [HL10, Theorem 6.3.2]). Fix E_A a supersingular elliptic curve (as public key). Then K executes the following algorithm:

1. Sample a seed $s \xleftarrow{\$} \{0,1\}^*$ fixing the randomness of P^*, sample a challenge φ and run $P^*(E_A, s, \varphi)$ repeatedly until the transcript $(E_1, \varphi, R_1, R_2, q)$ outputted by P^* is accepted by the verifier and save s.
2. Sample another challenge φ' and run $P^*(E_A, s, \varphi')$ with the same seed s as in step 1 (fixing the commitment value E_1) to obtain a new transcript $(E_1, \varphi', R_1', R_2', q')$ and repeat until we can extract a witness $\alpha \in \mathrm{End}(E_A)$ non-scalar from $(E_1, \varphi, R_1, R_2, q)$ and $(E_1, \varphi', R_1', R_2', q')$.
3. Break step 2 after k iterations (to be determined) or return α.

This algorithm may fail so K may execute this algorithm multiple times. We determine k to optimize the running time and the probability of failure of this algorithm. To do this, we specify how we can extract a witness in step 2. As in the previous knowledge soundness proof, (R_1, R_2, q) and (R_1', R_2', q') respectively provide an efficient representation of a q-isogeny $\sigma : E_A \longrightarrow E_2$ and a q'-isogeny $\widehat{\sigma'} : E_2 \longrightarrow E_A$, so we have an efficient representation of $\alpha := \widehat{\sigma'} \circ \varphi' \circ \widehat{\varphi} \circ \sigma \in \mathrm{End}(E_A)$. If φ and φ' are relatively good, then α is non-scalar by Lemma A.6.5 and we have won.

By Lemma A.6.6, since $D_\varphi > p^{5/2(1-\log(2)/\log\log(p^{5/2}))}$ by assumption, the number of challenges φ' relatively bad to φ is bounded by $C\mu(D_\varphi)/\sqrt{p}$ for some constant $C > 0$.

Now consider the matrix H whose rows are indexed by seeds s for P^*, whose columns are indexed by challenges φ and such that $H(s, \varphi)$ is the the result 0 or 1 returned by the verifier when P^* is run with E_A, s and φ. By assumption, the proportion of 1 in H is ε. A row with a proportion of 1 bigger than $\varepsilon/2$ is called a *heavy row*. Let R be the number of rows in H (*i.e.* the number of possible seeds for P^*). Let R' be the number of non-heavy rows. Then, the number of 1 located in a heavy rows is:

$$R\varepsilon\mu(D_\varphi) - R'\frac{\varepsilon}{2}\mu(D_\varphi) \geq R\varepsilon\mu(D_\varphi) - R\frac{\varepsilon}{2}\mu(D_\varphi) = R\frac{\varepsilon}{2}\mu(D_\varphi)$$

so at least half of the 1 are in heavy rows and the probability to fall in a heavy row at step 1 of the algorithm is $\geq 1/2$. Let φ be the challenge found at step 1. Now, at step 2, we are in the same row as in step 1 (since we fixed s). Assuming we are in a heavy row, the probability to find φ' that is not bad in relation to φ and such that $H(s, \varphi') = 1$ is

$$P \geq \frac{\varepsilon/2\mu(D_\varphi) - \mu(D_\varphi)C/\sqrt{p}}{\mu(D_\varphi)} = \frac{\varepsilon}{2} - \frac{C}{\sqrt{p}}.$$

In the following, we assume that $\varepsilon > 2C/\sqrt{p}$, so that $P > 0$. Then, the expected number of tries t to succeed in step 2 is:

$$\mathbb{E}(t) = \frac{1}{P} \leq \frac{2}{\varepsilon - 2C/\sqrt{p}}$$

Now we choose the time limit k accordingly. By Markov's inequality, the probability that step 2 terminates within k tries is

$$\mathbb{P}(t < k) = 1 - \mathbb{P}(t \geq k) \geq 1 - \frac{\mathbb{E}(t)}{k} \geq 1 - \frac{2}{k(\varepsilon - 2C/\sqrt{p})}$$

We choose $k := 4/(\varepsilon - 2C/\sqrt{p})$, so that $\mathbb{P}(t < k) \geq 1/2$. This probability is conditional to the fact that we fall into a heavy row, which has probability $\geq 1/2$ as we saw. Hence, the probability that the algorithm succeeds is $\geq 1/2 \times 1/2 = 1/4$ so K expects to repeat it 4 times to find a witness.

Now we estimate the running time of the algorithm. Step 1 is expected to terminate after $1/\varepsilon$ iterations and step 2 after $k = 4/(\varepsilon - 2C/\sqrt{p})$ iterations, so the total time complexity is

$$\frac{1}{\varepsilon} + \frac{4}{\varepsilon - 2C/\sqrt{p}} \leq \frac{5}{\varepsilon - 2C/\sqrt{p}}$$

RigorousSQISignHD being complete, we conclude that it is a proof of knowledge for $\mathcal{R}$ with knowledge error $\kappa := 2C/\sqrt{p} = O(p^{-1/2})$. □

A.7 On the codomain distribution of random isogenies with bounded degree (Theorem 6.2.2)

The goal of this section is to prove a bound on the statistical distance between codomains of random isogenies with bounded degrees and the uniform distribution on the supersingular isogeny graph. Similar results have been proven on fixed degree smooth isogeny walks [GPS16, Theorem 1] and non-bactracking ℓ-isogeny walks [BCC+22, Theorem 11]. We generalize these results to the case of non-fixed degree. Heuristically, we should expect to be as close to the uniform distribution lower degree bounds than in the fixed degree case, but this is not the case. In particular, as for non-bactracking ℓ-isogeny walks, we need to allow isogenies of degree $p^{1+\varepsilon}$ to reach a statistical distance of $O(p^{-\varepsilon/2})$. However, we provide an elementary proof that does not require to study adjacency matrices of the supersingular isogeny graph and modular forms (unlike [BCC+22]). The main ingredients are the Deuring correspondence and a count of small quaternion ideal vectors (Corollary A.7.3). We start by proving a classical bound on the last minimum of quaternion ideals claimed in [KLPT14, § 3.1] but never proved so far.

Lemma A.7.1. *Let $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ be a quaternion order and I be a left ideal of $\mathcal{O}$. Let $(\alpha_1, \cdots, \alpha_4)$ be a Minkowski reduced basis of I for the quadratic form $q_I : \alpha \in I \longmapsto \mathrm{nrd}(\alpha)/\mathrm{nrd}(I)$, so that $q_I(\alpha_i) \leq q_I(\alpha_{i+1})$ for all $i \in \{1, 2, 3\}$. Then*

$$q_I(\alpha_4) \leq \frac{8p}{\pi^2}.$$

Proof. As we saw in the proof of Lemma 5.2.2, we have by Minkowski's second theorem (Equation 2):

$$\prod_{i=1}^{4} q_I(\alpha_i) \leq \frac{64p^2}{\pi^4}.$$

This inequality is not sufficient to conclude (we only get $q_I(\alpha_4) = O(p^2)$ instead of $O(p)$). To complete the proof, we follow [BST+17, Theorem 3.1].

Let $(\beta_1, \cdots, \beta_4)$ be a Minkowski reduced basis of $\mathcal{O}$. As the α_i, the β_i satisfy

$$\prod_{i=1}^{4} \mathrm{nrd}(\beta_i) \leq \frac{64p^2}{\pi^4}.$$

Let $A := (a_{i,j})_{1 \leq i,j \leq 4} \in M_3(\mathbb{Z})$, where for all $1 \leq i, j \leq 4$, $a_{i,j}$ is the coefficient of α_4 in the decomposition of $\beta_i \alpha_j$ in the basis $(\alpha_1, \cdots, \alpha_4)$ (this is an integer since $\mathcal{O}I \subseteq I$). Then A is invertible. Indeed, if $x \in \mathbb{Z}^4$ satisfy $Ax = 0$ *i.e.* $\sum_{j=1}^{3} a_{i,j} x_j = 0$ for all $i \in [\![1 \; ; \; 4]\!]$, so $\alpha := \sum_{j=1}^{4} x_j \alpha_j$ satisfy $\mathcal{O}\alpha \subseteq \langle \alpha_1, \alpha_2, \alpha_3 \rangle$. But $\mathcal{O}\alpha$ has rank 4 whenever $\alpha \neq 0$, so $\alpha = 0$ and $x = 0$.

A being invertible, there exists a permutation $\sigma \in \mathfrak{S}_4$ such that $a_{i,\sigma(i)} \neq 0$ for all $i \in [\![1 \; ; \; 4]\!]$. It follows that for all $i \in [\![1 \; ; \; 4]\!]$, $\beta_i \alpha_{\sigma(i)}$ completes $\langle \alpha_1, \alpha_2, \alpha_3 \rangle$ as a full-rank sublattice of I, so that $\mathrm{nrd}(\alpha_4) \leq \mathrm{nrd}(\beta_i \alpha_{\sigma(i)})$ *i.e.* $q_I(\alpha_4) \leq$

$\operatorname{nrd}(\beta_i)q_I(\alpha_{\sigma(i)})$, since $(\alpha_1, \cdots, \alpha_4)$ is Minkowski reduced. It follows that

$$q_I(\alpha_4)^4 \leq \prod_{i=1}^{4}(\operatorname{nrd}(\beta_i)q_I(\alpha_{\sigma(i)})) = \prod_{i=1}^{4}\operatorname{nrd}(\beta_i)\prod_{i=1}^{4}q_I(\alpha_i) \leq \left(\frac{64p^2}{\pi^4}\right)^2.$$

The result follows. □

Now, we introduce a generalization of [Wes22, Lemma 3.2] in every dimension, counting the elements of bounded norm in a lattice.

Lemma A.7.2. *Let $\Lambda \subseteq \mathbb{Z}^k$ be a full-rank lattice of last minimum λ_k and $\rho > \sqrt{k}/2\lambda_k$. Then*

$$\frac{\pi^{k/2}\left(\rho - \frac{\sqrt{k}\lambda_k}{2}\right)^k}{\Gamma\left(\frac{k}{2}+1\right)\operatorname{Covol}(\Lambda)} \leq \#\Lambda \cap B(0,\rho) \leq \frac{\pi^{k/2}\left(\rho + \frac{\sqrt{k}\lambda_k}{2}\right)^k}{\Gamma\left(\frac{k}{2}+1\right)\operatorname{Covol}(\Lambda)},$$

where $B(0,\rho)$ is the ball of center 0 and radius ρ for the Euclidean norm and Γ is the Euler gamma function.

Proof. Let $\mathcal{V} := \{v \in \mathbb{R}^k \mid \|v\| = \min_{\lambda\in\Lambda}\|v+\lambda\|\}$ be the Voronoi cell at the origin of Λ and $\mu := \sup_{v\in\mathbb{R}^k}\|v\|$ be the covering radius of Λ. Then, we have

$$B(0,\rho-\mu) \subseteq \bigsqcup_{\lambda\in\Lambda\cap B(0,\rho)}(\lambda+\mathcal{V}) \subseteq B(0,\rho+\mu),$$

so that

$$\operatorname{Vol}(B(0,\rho-\mu)) \leq (\#\Lambda\cap B(0,\rho))\cdot\operatorname{Vol}(\mathcal{V}) \leq \operatorname{Vol}(B(0,\rho+\mu)).$$

Since $\operatorname{Vol}(\mathcal{V}) = \operatorname{Covol}(\Lambda)$, $\operatorname{Vol}(B(0,\rho\pm\mu)) = \pi^{k/2}(\rho\pm\mu)^k/\Gamma(k/2+1)$ and $\mu \leq \sqrt{k}\lambda_k/2$, the result follows. □

Corollary A.7.3. *Let $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ be a maximal order and I be an integral left $\mathcal{O}$-ideal. Then, for all $\varepsilon > 0$ the number of ideals of norm $\leq p^{1+\varepsilon}$ that are right-equivalent to I is*

$$N_{p^{1+\varepsilon}}([I]) := \#\{J \sim I \mid \operatorname{nrd}(J) \leq p^{1+\varepsilon}\} = \frac{2\pi^2}{\#\mathcal{O}_R(I)^\times}p^{1+2\varepsilon}(1+O(p^{-\varepsilon/2})).$$

Proof. By [DKLPW20, Lemma 1], an ideal J is right-equivalent to I if and only if it is of the form $J := I\overline{\alpha}/\operatorname{nrd}(I)$ for some $\alpha \in I$. Besides, α is uniquely determined by J up to multiplication on the right by an element of $\mathcal{O}_R(I)^\times$ and we have $\operatorname{nrd}(J) = \operatorname{nrd}(\alpha)/\operatorname{nrd}(I) = q_I(\alpha)$. It follows that

$$N_{p^{1+\varepsilon}}([I]) := \#\{J \sim I \mid \operatorname{nrd}(J) \leq p^{1+\varepsilon}\} = \frac{1}{\#\mathcal{O}_R(I)^\times}\#\{\alpha \in I \mid q_I(\alpha) \leq p^{1+\varepsilon}\}.$$

As in the proof of Lemma 5.2.2, let $\iota : \mathcal{B}_{p,\infty} \hookrightarrow \mathbb{R}^4$ be the canonical embedding, such that $\|\iota(\alpha)\|^2 = \mathrm{nrd}(\alpha)$ for all $\alpha \in \mathcal{B}_{p,\infty}$, where $\|.\|$ is the Euclidean norm on $\mathbb{R}^4$. Consider the lattice $\Lambda := \iota(I)$. We then have

$$N_{p^{1+\varepsilon}}([I]) = \frac{1}{\#\mathcal{O}_R(I)^\times} \#\Lambda \cap B\left(0, p^{(1+\varepsilon)/2}\sqrt{\mathrm{nrd}(I)}\right)$$

By Lemmas A.7.1 and A.7.2, we get

$$\#\Lambda \cap B\left(0, p^{(1+\varepsilon)/2}\sqrt{\mathrm{nrd}(I)}\right) \leq \frac{\pi^2\left(p^{(1+\varepsilon)/2}\sqrt{\mathrm{nrd}(I)} + \frac{2\sqrt{2p\,\mathrm{nrd}(I)}}{\pi}\right)^4}{2\,\mathrm{Covol}(\Lambda)},$$

with $\mathrm{Covol}(\Lambda) = p/4\,\mathrm{nrd}(I)^2$, as we saw in the proof of Lemma 5.2.2. It follows, that the right term of the inequality is $2\pi^2 p^{1+2\varepsilon}(1 + O(p^{-\varepsilon/2}))$. Applying the lower bound of Lemma A.7.2, we also get that

$$\#\Lambda \cap B\left(0, p^{(1+\varepsilon)/2}\sqrt{\mathrm{nrd}(I)}\right) \geq 2\pi^2 p^{1+2\varepsilon}(1 + O(p^{-\varepsilon/2})).$$

The result follows. □

We denote by $\mathrm{SS}(p)$ be the set of supersingular elliptic curves over $\mathbb{F}_{p^2}$ (up to $\overline{\mathbb{F}_p}$-isomorphism) and S the probability distribution on $\mathrm{SS}(p)$ given by $S(E) := 1/(K\#\,\mathrm{Aut}(E))$ for all $E \in \mathrm{SS}(p)$, with $K := \sum_{E\in\mathrm{SS}(p)} 1/\#\,\mathrm{Aut}(E) = (p-1)/24$ by Eichler mass formula [Voi20, Theorem 25.1.1]. Let U be the uniform distribution on $\mathrm{SS}(p)$.

Lemma A.7.4. *The statistical distance between S and U is $d_{TV}(S,U) = O(p^{-1})$.*

Proof. By [Sil09][Theorem III.10.1], we have $\#\,\mathrm{Aut}(E) = 2$ for all $E \in \mathrm{SS}(p)$ such that $j(E) \neq 0, 1728$, and $\#\,\mathrm{Aut}(E) \in \{4, 6\}$ otherwise and by [Sil09][Theorem V.4.1], there exists $C_p \in \mathbb{Z}$ small such that $\#\,\mathrm{SS}(p) = 2K + C_p$. Hence, we have

$$\begin{aligned} d_{TV}(U,S) &= \frac{1}{2}\sum_{E\in\mathrm{SS}(p)}\left|\frac{1}{\#\,\mathrm{SS}(p)} - \frac{1}{K\#\,\mathrm{Aut}(E)}\right| \\ &= \frac{1}{2}\sum_{\substack{E\in\mathrm{SS}(p)\\ j(E)\neq 0,1728}}\left|\frac{1}{2K+C_p} - \frac{1}{2K}\right| + O(p^{-1}) \\ &= \frac{1}{2}\frac{C_p}{2K(2K+C_p)}(\#SS(p) + O(1)) + O(p^{-1}) \\ &= \frac{C_p}{4K} + O(p^{-1}) = O(p^{-1}). \end{aligned}$$

□

We can now finally prove our main result.

Theorem 6.2.2. Let $\varepsilon \in]0,2]$. Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve and π be the probability distribution of codomains E' (up to $\overline{\mathbb{F}_p}$-isomorphism) of isogenies $\sigma : E \longrightarrow E'$ chosen uniformly at random among isogenies of degree $\deg(\sigma) \leq p^{1+\varepsilon}$. Then $d_{TV}(U,\pi) = O(p^{-\varepsilon/2})$.

Proof. By the Deuring correspondence, it suffices to prove that given a maximal order $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ and $\mathrm{Cl}(\mathcal{O})$ the set of right-equivalence classes of left-ideals of $\mathcal{O}$, the distribution π' of the ideal classes $[I] \in \mathrm{Cl}(\mathcal{O})$ when I is sampled uniformly at random among ideals of norm $\leq p^{1+\varepsilon}$ (which is the quaternion analogue of π) is at statistical distance $O(p^{-\varepsilon/2})$ from the uniform distribution U' on $\mathrm{Cl}(\mathcal{O})$. We also denote by S' the quaternion analogue of S, namely the distribution on $\mathrm{Cl}(\mathcal{O})$ given by $S([I]) := 1/(K\#O_R(I)^\times)$ for all $[I] \in \mathrm{Cl}(\mathcal{O})$, where $K := \sum_{[I]\in\mathrm{Cl}(\mathcal{O})} 1/\#O_R(I)^\times = (p-1)/24$ is the Eichler mass. By Lemma A.7.4, $d_{TV}(U',S') = O(p^{-1})$, so it suffices to prove that $d_{TV}(S',\pi') = O(p^{-\varepsilon/2})$.

By Corollary A.7.3, the number of left $\mathcal{O}$-ideals of norm $\leq p^{1+\varepsilon}$ is

$$N_{p^{1+\varepsilon}} = \sum_{[I]\in\mathrm{Cl}(\mathcal{O})} N_{p^{1+\varepsilon}}([I]) = 2\pi^2 K p^{1+2\varepsilon}(1+O(p^{-\varepsilon/2})),$$

so the distribution π' is given by

$$\forall [I] \in \mathrm{Cl}(\mathcal{O}), \quad \pi'([I]) = \frac{N_{p^{1+\varepsilon}}([I])}{N_{p^{1+\varepsilon}}} = \frac{1}{K}(1+O(p^{-\varepsilon/2})).$$

It follows immediately that

$$d_{TV}(S',\pi') = \frac{1}{2}\sum_{[I]\in\mathrm{Cl}(\mathcal{O})} |S'([I]) - \pi'([I])| = \frac{\#Cl(\mathcal{O})}{K} O(p^{-\varepsilon/2}) = O(p^{-\varepsilon/2}).$$

The result follows. □

B Response and verification in dimension 8 when q is not coprime to ℓ

As explained in Section 3.3, we have no guarantee that q is coprime to ℓ in dimension 8, and in that case we can no longer use the simple formula for $\ker(F)$ and the optimisations of Section 5.4 to compute the 8-dimensional isogeny F embedding the response σ. To be able to use the techniques we developed, we factor σ into $\sigma := \widehat{\sigma_2} \circ \sigma' \circ \sigma_A$, where σ_A and σ_2 both have degree dividing ℓ^f and σ' has degree coprime to ℓ. We then represent σ' with the techniques we presented earlier.

B.1 Finding the ℓ-isogeny factors in the response

In this section, we explain how to factor $\sigma : E_A \longrightarrow E_2$ by ℓ-isogenies, when the only thing we know is its kernel ideal I. We not only need to find the factors

$\sigma_A : E_A \longrightarrow E'_A$ and $\sigma_2 : E_2 \longrightarrow E'_2$ but also alternate paths $\theta_A : E_0 \longrightarrow E'_A$ and $\theta_2 : E_0 \longrightarrow E'_2$ of norm coprime to ℓ to be able to evaluate $\sigma' : E'_A \longrightarrow E'_2$ with $\mathsf{EvalTorsion}_{\ell^f}$.

$$\begin{tikzcd}
E_A \arrow[r, "\sigma_A"] & E'_A \arrow[r, dashed, "\sigma'"] & E'_2 & E_2 \arrow[l, "\sigma_2"'] \\
E_0 \arrow[u, bend left, "\tau"] \arrow[u, bend right, "\tau'"'] \arrow[ur, "\theta_A"'] \arrow[urr, "\theta_2"] \arrow[urrr, "\rho'_2"'] \arrow[rrr, "\psi"] & & & E_1 \arrow[u, "\varphi"']
\end{tikzcd}$$

We start by factoring I to find the kernel ideals J_A and J_2 of σ_A and σ_2 of norm at most ℓ^f. Let us write $I := \ell^b m \cdot J$ where m is coprime to ℓ and J a left $\mathcal{O}_A$-ideal without integer factor. Let us write $\mathrm{nrd}(J) := \ell^b m'$, where m' coprime to ℓ. Then $\mathrm{nrd}(I) = \ell^{2b+c} m^2 m' < \ell^e$ so $2b+c \leq e \leq 2f$. We may write $b := b_1 + b_2$ and $c := c_1 + c_2$ with $2b_1 + c_1, 2b_2 + c_2 \leq f$. Then $J_A := (mJ + \mathcal{O}_A \ell^{c_1})\ell^{b_1}$ and $J_2 := (m\overline{J} + \mathcal{O}_2 \ell^{c_2})\ell^{b_2}$ have norms $\ell^{2b_1+c_1} | \ell^f$ and $\ell^{2b_2+c_2} | \ell^f$ respectively by the following lemma. Furthermore, we have $I := J_A I' \overline{J_2}$, where I' is a quaternion ideal of norm coprime to ℓ by construction.

Lemma B.1.1. *Let I be a left ideal of a maximal order $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ of norm coprime to p. Let us write $I := m \cdot J$ with $m \in \mathbb{N}^*$ and I' a left $\mathcal{O}$-ideal such that $J \not\subset n\mathcal{O}$ for all $n \in \mathbb{Z}$. Let $d \in \mathbb{N}^*$ coprime to m and $K := I + d\mathcal{O}$. Then* $\mathrm{nrd}(K) = d \wedge \mathrm{nrd}(J)$.

Proof. Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve of endomorphism ring isomorphic to $\mathcal{O}$. Then:

$$\begin{aligned} E[K] &= E[mJ + d\mathcal{O}] = E[mJ] \cap E[d] = \{P \in E \mid \forall \alpha \in J, [m]\alpha(P) = 0\} \cap E[d] \\ &= \{P \in E \mid [m]P \in E[J]\} \cap E[d] = [m]^{-1}(E[J]) \cap E[d] \end{aligned}$$

Since J is not divisible by any integer, $E[J]$ is cyclic so we may consider a generator $P \in E$ of $E[J]$. Let $N' := \mathrm{nrd}(J)$ and $d' := d \wedge \mathrm{nrd}(J)$. Let $Q_0 := [N'/d']P$. Then, $[d]Q_0 = [d/d'][N']P = 0$ and $[m]Q_0 \in \langle P \rangle$ by construction, so that $Q_0 \in E[K]$. Conversely, let $Q \in E[K]$. Then $[m]Q = [k]P$ for some $k \in \mathbb{Z}$ and $[d]Q = 0$. In particular $[kd]P = [md]Q = 0$. Then, $N' | kd$ since P has order N', so that $N'/d' | k$, so we may write $k = k'N'/d'$ with $k' \in \mathbb{Z}$, so that $[m]Q = [k'N/d']P = [k']Q_0$. Since m and d are coprime, there exists $u, v \in \mathbb{Z}$ such that $mu + dv = 1$ and we then have $Q = [mu + dv]Q = [um]Q = [uk']Q_0$. Hence, $E[K] = \langle Q_0 \rangle$ and finally

$$\mathrm{nrd}(K) = \#E[K] = \#\langle Q_0 \rangle = \#\langle [N'/d']P \rangle = d' = d \wedge \mathrm{nrd}(J).$$

□

Knowing J_A and J_2, we can then compute their associated isogenies σ_A and σ_2. Since J_A and J_2 have norm dividing ℓ^f, $E_A[J_A]$ and $E_2[J_2]$ are contained in the accessible ℓ^f-torsion. So we only have to evaluate a basis of J_A and J_2 on the ℓ^f-torsion and solve discrete logarithms in groups of exponent ℓ^f to compute $E_A[J_A]$ and $E_2[J_2]$. We can then apply Vélu's formulas [Vél71] to compute σ_A and σ_2. To obtain basis of J_A that we can evaluate on the ℓ^f-torsion, we compute a T-eval-basis of J_A in the sense of Definition 2.3.1 by expressing the basis of J_A that we already know as integer linear combinations of a T-eval-basis $\mathcal{B}_A := \mathsf{PushEndRing}(\tau, I_\tau)$ of $\mathrm{End}(E_A)$ obtained via Algorithm 1. The same principle applies to J_2. Let $\rho_2 := \varphi \circ \psi$ and $I_2 := I_\psi \cdot I_\varphi$ its kernel ideal. Then, we can obtain a T-eval-basis $\mathcal{B}_2 := \mathsf{PushEndRing}(\rho_2, I_2)$ yielding a T-eval-basis of J_2.

Now we explain how to find alternate paths $\theta_A : E_0 \longrightarrow E'_A$ and $\theta_2 : E_0 \longrightarrow E'_A$ of degree coprime to ℓ. Fist, we find left $\mathcal{O}_0$-ideals $K_A \sim I_\tau \cdot J_A$ and $K_2 \sim I_2 \cdot J_2$ of powersmooth norm coprime to ℓ using the KLPT algorithm [KLPT14]. To translate K_A and K_2 into isogenies θ_A and θ_2, we could use the paths $\sigma_A \circ \tau$ and $\sigma_2 \circ \rho_2$ (where $\rho_2 = \varphi \circ \psi$) and apply $\mathsf{SpecialIdealToIsogeny}$ (presented in Section 2.3) but $\mathrm{nrd}(K_A)$ and $\mathrm{nrd}(K_2)$ would need to be coprime to T. We could use powersmooth torsion coprime to T and ℓ and still compute θ_A and θ_2 in polynomial time but this would not be optimal. Instead, we propose to seek K_A and K_2 of norm dividing $T^2 \simeq p^3$ and to use paths $\sigma_A \circ \tau'$ and $\sigma_2 \circ \rho'_2$ in $\mathsf{SpecialIdealToIsogeny}$, where $\tau' : E_0 \longrightarrow E_A$ and $\rho'_2 : E_0 \longrightarrow E_2$ are isogenies of degree a power of ℓ.

The input τ' is a by-product of the key generation, which is similar to the commitment procedure when two isogeny paths of coprime degree are computed. We can simply run Algorithm 5 completely to obtain $\tau, \tau' : E_0 \longrightarrow E_A$ at the same time.

To find ρ'_2, we apply KLPT to the kernel ideal $I_2 := I_\psi \cdot I_\varphi$ of $\rho_2 := \varphi \circ \psi$, to find $I'_2 \sim I_2$ of norm $\ell^h \simeq p^3$. We can then translate the ideal I'_2 into its associated isogeny ρ'_2 via the effective Deuring correspondence algorithm introduced in the original SQISign paper [DKLPW20, Algorithm 9]. We summarize all the computations to factor σ in the $\mathsf{FactorIsogeny}_{\ell^f,T}$ algorithm (Algorithm 17).

B.2 Adaptations of the response and verification when q is not coprime to ℓ

Keeping the notations of the previous section, assume we have factored $\sigma := \widehat{\sigma_2} \circ \sigma' \circ \sigma_A$. Then, we can embed σ' in an isogeny F of dimension 8 using the same techniques presented earlier since $q' := \deg(\sigma')$ has degree coprime to ℓ. To proceed, we evaluate σ' on a canonically generated basis (P'_1, P'_2) of $E'_A[\ell^f]$ using the isogeny paths $\theta_A : E_0 \longrightarrow E'_A$ and $\theta_2 : E_0 \longrightarrow E'_2$ of degree dividing T^2 to apply $\mathsf{EvalTorsion}_{\ell^f}$ (Algorithm 10). Once all these computations are done, the prover simply sends $(\sigma_A, \sigma_2, \sigma'(P'_1), \sigma'(P'_2), q')$ to the verifier. The complete $\mathsf{RigorousRespond}$ procedure follows (Algorithm 18).

The complete verification procedure $\mathsf{RigorousVerify}$ (Algorithm 19) is very similar to the original one. Indeed, using representing σ' and representing $\sigma = \widehat{\sigma_2} \circ \sigma' \circ \sigma_A$ is equivalent when σ_A and σ_2 are known.

Algorithm 17: $\mathsf{FactorIsogeny}_{\ell^f,T}$

Data: A quaternion ideal I of norm $< \ell^e$ connecting $\mathcal{O}_A \cong \mathrm{End}(E_A)$ and $\mathcal{O}_2 \cong \mathrm{End}(E_2)$, two isogenies $\tau, \tau' : E_0 \longrightarrow E_A$ of degrees dividing T^2 and a power of ℓ respectively, $\rho_2 : E_0 \longrightarrow E_2$ of degree dividing a power of T and I_τ, I'_τ, I_2 their respective kernel ideals.

Result: Two left-ideals $J_A \subseteq \mathcal{O}_A$ and $J_2 \subseteq \mathcal{O}_2$ whose of norms divide ℓ^f such that $I := J_A I' \overline{J_2}$, with I' of norm coprime to ℓ, two ideals $K_A \sim I_\tau \cdot J_A$ and $K_2 \sim I_2 \cdot J_2$ of norms dividing T^2 along with isogenies $\sigma_A : E_A \longrightarrow E'_A$, $\sigma_2 : E_2 \longrightarrow E'_2$, $\theta_A : E_0 \longrightarrow E'_A$ and $\theta_2 : E_0 \longrightarrow E'_2$ respectively associated to J_A, J_2, K_A and K_2.

```
Factor $I := \ell^b m J$ with $\ell \wedge m = 1$ and $I'$ without integer factors and factor $\mathrm{nrd}(J) := \ell^c m'$ with $\ell \wedge m' = 1$;
Let $b := b_1 + b_2$ and $c := c_1 + c_2$ with $2b_i + c_i \leq f$ for $i \in \{1, 2\}$;
$J_A \longleftarrow (mJ + \mathcal{O}_A \ell^{c_1})\ell^{b_1}, J_2 \longleftarrow (m\overline{J} + \mathcal{O}_2 \ell^{c_2})\ell^{b_2}$;
$I' \longleftarrow J_A^{-1} I \overline{J_2}^{-1}$;
Compute two $T$-eval-basis $\mathcal{B}_A := \mathsf{PushEndRing}(\tau, I_\tau)$ and $\mathcal{B}_2 := \mathsf{PushEndRing}(\rho_2, I_2)$;
Infer $T$-eval-basis $\mathcal{C}_A$ of $J_A$ and $\mathcal{C}_2$ of $J_2$ from $\mathcal{B}_A$ and $\mathcal{B}_2$;
Evaluate $\mathcal{C}_A$ on a basis of $E_A[\ell^f]$ and $\mathcal{C}_2$ on a basis of $E_A[\ell^f]$ to compute $G_A := E_A[J_A]$ and $G_2 := E_2[J_2]$;
Compute $\sigma_A : E_A \longrightarrow E'_A$ and $\sigma_2 : E_2 \longrightarrow E'_2$ of kernel $G_A$ and $G_2$ respectively;
$K_A \longleftarrow \mathsf{KLPT}_{T^2}(I_\tau \cdot J_A), K_2 \longleftarrow \mathsf{KLPT}_{T^2}(I_2 \cdot J_2)$;
$\theta_A \longleftarrow \mathsf{SpecialIdealToIsogeny}(K_A, I_{\tau'} \cdot J_A, \sigma_A \circ \tau')$;
$I'_2 \longleftarrow \mathsf{KLPT}_{\ell^h}(I_2)$;
Compute $\rho'_2$ of kernel ideal $I'_2$ using [DKLPW20, Algorithm 9];
$\theta_2 \longleftarrow \mathsf{SpecialIdealToIsogeny}(K_2, I'_2 \cdot J_2, \sigma_2 \circ \rho'_2)$;
Return $J_A, J_2, I', K_A, K_2, \sigma_A, \sigma_2, \theta_A, \theta_2$;
```

B.3 Impact on compactness in dimension 8

When q is not coprime to ℓ, the factors σ_A et σ_2 are transmitted in the signature in addition to the data $(E_1, \sigma'(P'_1), \sigma'(P'_2), q')$. This is apparently more information than in the case $q \wedge \ell = 1$. However, we can optimize the communications to avoid almost any compactness loss.

We may write $\deg(\sigma_A) := \ell^{f_1}$ and $\deg(\sigma_2) := \ell^{f_2}$, with $f_1, f_2 \leq f$, so that $q' = q/\ell^{f_1+f_2} < \ell^{e'}$, where $e' := e - f_1 - f_2$. Hence, we can represent σ' by an $\ell^{e'}$-isogeny F' in dimension 8. By Remark 5.3, we only need to evaluate the ℓ^{f_3}-torsion by σ', where $2f_3 \geq e' + 4$. Hence the points P'_1 and P'_2 may form a basis of $E'_A[\ell^{f_3}]$ instead of $E'_A[\ell^f]$ and we can represent $\sigma'(P'_1)$ and $\sigma'(P'_1)$ with $3f_3$ bits by the techniques of Section 7.1 (assuming $\ell = 2$).

To represent σ_A, we may factor $\sigma_A := [\ell^{b_1}] \circ \sigma'_A$, where $2b_1 \leq f_1$ and $\sigma'_A : E_A \longrightarrow E'_A$ a cyclic $\ell^{f'_1}$-isogeny with $f'_1 := f_1 - 2b_1$. So we may represent σ_A by the integer b_1 and $\ker(\sigma'_A) \subset E_A[\ell^{f'_1}]$. Let (Q_1, Q_2) be a canonical basis of $E_A[\ell^{f'_1}]$. Then, $\ker(\sigma'_A)$ is generated either one of the points $Q_1 + kQ_2$ with

Algorithm 18: RigorousRespond

```
Data: Two isogenies to the commitment $\psi, \psi' : E_0 \longrightarrow E_1$ and two isogenies to
      the public key $\tau, \tau' : E_0 \longrightarrow E_A$ of respective degrees $D_\psi, D_{\psi'}, D_\tau, D_{\tau'}$
      such that $D_\psi, D_\tau | T^2$ and $D_{\psi'}$ and $D_{\tau'}$ are powers of $\ell$, the challenge
      isogeny $\varphi : E_1 \longrightarrow E_2$ of degree $D_\varphi | T$ as well as their respective kernel
      ideals $I_\tau, I_{\tau'}, I_\psi, I_{\psi'}, I_\varphi$.
Result: $(\sigma_A, \sigma_2, \sigma'(P_1'), \sigma'(P_2'), q')$, where $\sigma_A : E_A \longrightarrow E_A'$ and $\sigma_2 : E_2 \longrightarrow E_2'$
        are isogenies of degree dividing $\ell^f$, $\sigma' : E_A' \longrightarrow E_2'$ is an isogeny of
        degree $q' < \ell^e$ coprime to $\ell$ and $(P_1', P_2')$ is a canonically determined
        basis of $E_A'[\ell^f]$.
$I_\varphi \longleftarrow$ IsogenyToIdeal$(\varphi, \psi', I_{\psi'})$;
$J \longleftarrow \overline{I_\tau} \cdot I_\psi \cdot I_\varphi$;
$I \longleftarrow$ RandomEquivalentIdeal$_{\ell^e}(J)$ and $q \longleftarrow \mathrm{nrd}(I)$;
$J_A, I', J_2, K_A, K_2, \sigma_A, \sigma_2, \theta_A, \theta_2 \longleftarrow$
  FactorIsogeny$_{\ell^f, T}(I, \tau, \tau', \varphi \circ \psi, I_\tau, I_{\tau'}, I_\psi \cdot I_\varphi)$;
$q' \longleftarrow q / \mathrm{nrd}(J_A) \, \mathrm{nrd}(J_2)$;
Compute the canonical basis $(P_1', P_2')$ of $E_A'[\ell^f]$;
$(\sigma'(P_1'), \sigma'(P_2')) \longleftarrow$ EvalTorsion$_{\ell^f}(I', P_1', P_2', \theta_A, \theta_2, K_A, K_2)$;
Return $(\sigma_A, \sigma_2, \sigma'(P_1'), \sigma'(P_2'), q')$;
```

Algorithm 19: RigorousVerify

```
Data: $(\sigma_A, \sigma_2, R_1', R_2', q')$, where $\sigma_A : E_A \longrightarrow E_A'$ and $\sigma_2 : E_2 \longrightarrow E_2'$ are
      isogenies of degree dividing $\ell^f$, $R_1', R_2' \in E_2'[\ell^f]$ and $q' \in \mathbb{N}^*$.
Result: 1 if $(\sigma_A, \sigma_2, R_1', R_2', q')$ is a valid response and 0 otherwise.
if $q' > \ell^e$ then
    Return 0;
end
Compute the canonical basis $(P_1', P_2')$ of $E_A'[\ell^f]$;
Find $a_1, \cdots, a_4 \in \mathbb{Z}$ such that $a_1^2 + \cdots + a_4^2 + q = \ell^e$ using Pollack and
  Treviño's algorithm [PT18];
$F' \longleftarrow$ RepresentIsogeny$_{8, \ell^e}(q', a_1, \cdots, a_4, P_1', P_2', R_1', R_2')$;
Return IsValid$_8(F')$;
```

$0 \leq k \leq \ell^{f_1'} - 1$ or one of the points $\ell k' Q_1 + Q_2$ with $0 \leq k' \leq \ell^{f_1'-1} - 1$. Hence, $\ker(\sigma_A')$ can be represented by $f_1' + 1$ bits (one bit to tell which form takes $\ker(\sigma_A')$ and f_1' bits for k or k'). Since the number of bits to represent b_1 is very small ($O(\log(b_1))$), we may represent σ_A by at most $f_1 + 1$ bits, and similarly, we may represent σ_2 by at most $f_2 + 1$ bits.

As in Section 7.1, we represent $q' < \ell^{e'}$ with e' bits and E_1 with 4λ bits (where λ is the security level, satisfying $p \simeq 2^\lambda$). Hence, the signature size is

$$\begin{aligned} 3f_3 + f_1 + f_2 + 2 + 4\lambda + e' &= \frac{5}{2}(e - f_1 - f_2) + f_1 + f_2 + 4\lambda + O(\log(\lambda)) \\ &< \frac{5}{2}e + 4\lambda + O(\log(\lambda)) = 14\lambda + O(\log(\lambda)) \end{aligned}$$

in bits, so we do not suffer any communication loss compared to the case $q \wedge \ell = 1$ (as the inequality indicates). However, since $e \simeq 4\lambda$ ($\ell^e = \Theta(p^2)$), Rigorous-SQISignHD signatures are much less compact than FastSQISignHD signatures (14λ instead of $13/2\lambda$ bits).

C Isogenies in the theta model

In this section we give various practical details on how to perform the required isogenies computations in dimension 4 and 8 using the theta model.

C.1 Theta coordinates

For simplicity, even through we work over a finite field, we will describe our algorithm using analytic theta functions. The algebraic theory of Mumford [Mum66] can be used to show that our algorithms are still valid over an arbitrary field of odd characteristic.

Let $A = \mathbb{C}^g/(\mathbb{Z}^g + \Omega_A \mathbb{Z}^g)$ be an abelian variety with $\Omega = \Omega_A$ in the Siegel space corresponding to a principal polarisation $\mathcal{L} = \mathcal{L}_A$ on A. Let $\pi_A : \mathbb{C}^g \to A$ be the projection.

Recall that the analytic theta functions with characteristic $a, b \in \mathbb{Q}^g$ are given by

$$\theta \left[\begin{smallmatrix} a \\ b \end{smallmatrix} \right] (z, \Omega) = \sum_{n \in \mathbb{Z}^g} e^{\pi i \, {}^t(n+a)\Omega(n+a) + 2\pi i \, {}^t(n+a)(z+b)}.$$

A basis of level 2 theta functions is given by $\theta_i^A(P) = \theta \left[\begin{smallmatrix} 0 \\ i/2 \end{smallmatrix} \right] (z_P, \Omega/2), i \in (\mathbb{Z}/2\mathbb{Z})^g$ where $z_P \in \mathbb{C}^g$ represents $P \in A$: $P = \pi_A(z_P)$. Here we use the following abuse of notations: if $i \in (\mathbb{Z}/2\mathbb{Z})^g$, we denote by i any lift to $\mathbb{Z}^g$. Reciprocally if $i \in \mathbb{Z}^g$, we also denote by i its reduction to $(\mathbb{Z}/2\mathbb{Z})^g$.

The analytic theta functions depend on the period matrix Ω_A. Algebraically they are defined by a symmetric theta structure Θ_A of level 2. We will denote our theta functions by $\theta_i^{\Theta_A}$ when we want to make this dependence explicit.

We will also make use of the "dual" basis $\theta_\chi'^A(P) = \theta \left[\begin{smallmatrix} \chi/2 \\ 0 \end{smallmatrix} \right] (2z_P, 2\Omega), \chi \in (\hat{\mathbb{Z}}/2\hat{\mathbb{Z}})^g$, where we identify $(\hat{\mathbb{Z}}/2\hat{\mathbb{Z}})^g$ with $(\mathbb{Z}/2\mathbb{Z})^g$ via the inner product. Going to the dual level 2 coordinates corresponds analytically to the action of the symplectic matrix $S = \left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$ on the period matrix Ω_A. Explicitly on theta coordinates the modular transform is given via the Hadamard transformation $\theta'_\chi = \sum_i \chi(i)\theta_i$, and reciprocally $2^g\theta_i = \sum_\chi \chi(i)\theta'_\chi$. We let H be the Hadamard matrix in dimension 2^g, by the formula above it allows to pass back and forth between the theta coordinates of level 2 and their duals.

C.2 Gluing theta structures

Let us recapitulate the isogeny we need to compute in the verification step of SQISignHD: we have a d-isogeny $F : A \to B$ where $(A, \mathcal{L})$ and $(B, \mathcal{M})$ are given by products of elliptic curves with their product polarisations. We split the isogeny in two as in Section 5.4: $F = F_2 \circ F_1$ with $F_1 : A \to C$ a d_1-isogeny and $F_2 : C \to B$ an d_2-isogeny, and we assume that we are given the kernel of F_1 in $A[d_1]$ and $\tilde{F}_2$ in $B[d_2]$.

The theta isogeny algorithm to compute F_1 requires (if d_1 is odd):

- A symmetric level 2 theta structure Θ_A on $(A, \mathcal{L})$. This level structure determines a symplectic basis of $A[2]$ and is in turn determined by a symplectic basis of $A[4]$. This symmetric level structure will be represented (up to twists) by the theta constant $(\theta_i^{\Theta_A}(0))_{i\in(\mathbb{Z}/2\mathbb{Z})^g}$ and gives a basis of sections $(\theta_i^{\Theta_A})_{i\in(\mathbb{Z}/2\mathbb{Z})^g}$ of $\mathcal{L}^2$. In particular if the 4-torsion is rational, the level 2 theta model will be rational (this is a sufficient but not necessary condition).
- Generators $P_1, \ldots, P_g$ of the kernel $K = \operatorname{Ker} F_1$ in theta coordinates $\theta_i^{\Theta_A}(P_j)$, with the basis of theta coordinates induced by the symmetric theta structure fixed above.

In SQISignHD, A will be equal to a product of g elliptic curves $A = E_1 \times \cdots \times E_g$ and the points of the kernel K is described in terms of tuples of Weierstrass coordinates. We first need to explain how to convert these points to theta coordinates. We fix a symplectic basis (e_i, f_i) on each $E_i[4]$, this induces a product symplectic basis on $A[4]$, hence a product theta structure. There are well known formula to convert from Weierstrass coordinates on E_i to theta coordinates [Mum84]. We can then compute the theta coordinates on A as follow:

Lemma C.2.1. *Let $\mathcal{L} = \mathcal{L}_1 \star \mathcal{L}_2 \cdots \star \mathcal{L}_g$ be a product polarisation on $A = E_1 \times \cdot \times E_g$. Endow $(A, \mathcal{L})$ with a product theta structure Θ_A of each theta structure Θ_{E_i} on E_i. If $P = (P_1, \ldots, P_g) \in A$, then for $i = (i_1, \ldots, i_g) \in (\mathbb{Z}/2\mathbb{Z})^g$, $\theta_i^{\Theta_A}(P) = \prod_{j=1}^g \theta_{i_j}^{\Theta_{E_j}}(P_j)$.*

Proof. This follows from

$$\theta\left[\begin{smallmatrix} a_1, a_2 \\ b_1, b_2 \end{smallmatrix}\right]((z_1, z_2), \begin{pmatrix} \Omega_1 & 0 \\ 0 & \Omega_2 \end{pmatrix}) = \theta\left[\begin{smallmatrix} a_1 \\ b_1 \end{smallmatrix}\right](z_1, \Omega_1)\theta\left[\begin{smallmatrix} a_2 \\ b_2 \end{smallmatrix}\right](z_2, \Omega_2).$$

□

Hence if we have a point $R \in A$ given in product Weierstrass coordinates $R = (R_j) = ((x_j, y_j))$, we can convert each R_i from Weierstrass coordinates to level 2 theta coordinates $\theta_i^{\Theta_{E_j}}(R_j)$ then apply Lemma C.2.1 to get the theta coordinates of R with respect to the product theta structure.

We can now apply the isogeny theorem from [CR15; LR23] with $N = d_1, d_2$:

Theorem C.2.2. *Let $(A, \mathcal{L})$ be a principally polarised abelian variety with a symmetric theta structure Θ_A of level 2, induced by a symplectic basis $\mathcal{B}_A = (x_1, \ldots, x_g, y_1, \ldots, y_g)$ of $A[4]$ with respect to ζ, a primitive fourth root of unity.*

Let $(P_1, \ldots, P_g)$ be a basis of the kernel of a maximal isotropic subgroup of $A[N]$ of rank g, given in theta coordinates, where N is an odd integer.

Let P be a point of A given in theta coordinates. Let $F : A \to B = A/K$ the induced isogeny. Then there is a unique descent of $\mathcal{L}^N$ to a polarisation $\mathcal{M}$ on B, and a symmetric theta structure Θ_B on $\mathcal{M}$ induced by the symplectic basis $F(\mathcal{B}_A) = (\frac{1}{N}F(x_1), \ldots, \frac{1}{N}F(x_g), F(y_1), \ldots, F(y_g))$ with respect to ζ of $B[4]$.

Furthermore, the theta null point of B and the theta coordinates of $F(P)$ can be computed in $O(N^g)$ arithmetic operations over the base field.

Proof. This is a special case of Theorem C.2.5 proved below. □

We can use Theorem C.2.2 to compute an ℓ^e-isogeny by splitting it into a product of ℓ-isogenies. In SQISignHD we specifically want to handle the case $\ell = 2$. We will give an algorithm in Appendix C.3.

We have another difficulty to solve first. In SQISignHD we glue two isogenies together $F_1 : A \to C$ and $\tilde{F}_2 : B \to C$. These isogenies are compatible with the product polarisation on A and B, so the codomain C is endowed with the same polarisation in both cases. However, when using Theorem C.2.2 to compute C and its polarisation, it needs not be endowed with the same level 2 symmetric theta structure Θ_C for the two isogenies.

Let $\mathcal{B}_1$ be a symplectic basis of $C[4]$ giving the symmetric theta structure on C induced by F_1, and $\mathcal{B}_2$ be the one induced by $\tilde{F}_2$. Then there is a symplectic matrix $M \in \mathrm{Sp}_{2g}(\mathbb{Z}/4\mathbb{Z})$ such that $M\mathcal{B}_1 = \mathcal{B}_2$. We can use the theta transformation formula [BL04, \$8.6] for M to convert the theta null point expressed in terms of $\mathcal{B}_1$ to the one expressed in terms of $\mathcal{B}_2$: $\theta_i^{\mathcal{B}_2}(0) = \theta_i^{M\mathcal{B}_1}(0)$. Once we have endowed C with the same theta structures, checking that they are indeed the same simply amount to testing for equality of the theta null points seen as projective coordinates.

So one way to test that F_1 and $\tilde{F}_2$ indeed have the same polarised codomain C is to apply Theorem C.2.2 twice and then to act by all matrices in $M \in \mathrm{Sp}_{2g}(\mathbb{Z}/4\mathbb{Z})$ on the theta null point induced by F_1 until we find an equality of projective theta null points with the theta null point induced by $\tilde{F}_2$. This costs $O(1)$ but in practice is too expensive. We will instead explain how to compute the correct correcting matrix M directly.

Remark C.2.3. Many symplectic basis $\mathcal{B}_C$ of $C[4]$ will give the same symmetric theta structure Θ_C of level 2 on C (hence the same theta null point), indeed the theta null point only determines a symplectic basis of $C[2]$. Rather than working with the 4-torsion we could work only with the 2-torsion and take $M \in \mathrm{Sp}_{2g}(\mathbb{Z}/2\mathbb{Z})$; this does not completely determines all the symmetric theta structures of level 2 but it is easy to test all 2^{2g} possibilities.

Proposition C.2.4. *Let $F_1 : A \to C$ be a d_1-isogeny, $\tilde{F}_2 : B \to C$ be a d_2-isogeny, $F = F_2 \circ F_1 : A \to B$, with d_1 and d_2 prime to 2. Let $\mathcal{B}_A$ a symplectic basis of $A[4]$, $\mathcal{B}_B$ a symplectic basis of $B[4]$. Let $\mathcal{B}'_B$ be the symplectic basis on B induced by F. Let $M' = \begin{pmatrix} \alpha' & \beta' \\ \gamma' & \delta' \end{pmatrix}$ be the symplectic matrix such that $M'\mathcal{B}'_B = \mathcal{B}_B$. Let $\mathcal{B}_1$ be the symplectic basis of $C[4]$ induced by F_1 and $\mathcal{B}_2$ be the symplectic basis of $C[4]$ induced by $\tilde{F}_2$. Then $\mathcal{B}_2 = M\mathcal{B}_1$, with $M = \begin{pmatrix} \alpha'/d_2 & \beta' \\ \gamma' & d_2\gamma' \end{pmatrix}$.*

Proof. Define $\gamma_x = \begin{pmatrix} x & 0 \\ 0 & 1 \end{pmatrix}$. By Theorem C.2.2, we have $\mathcal{B}_1 = \gamma_{1/d_1} \cdot F_1\mathcal{B}_A$, $\mathcal{B}_2 = \gamma_{1/d_2} \cdot \tilde{F}_2\mathcal{B}_B$, $\mathcal{B}'_B = \gamma_{1/d} \cdot F\mathcal{B}_A$. If $M\mathcal{B}_1 = \mathcal{B}_2$, we get that $M = d_2\gamma_{1/d_2} M' \gamma_{1/d}\gamma_{1/d_1}^{-1}$. □

So if we compute the image of F on $A[4]$, we can recover the correct matrix M. In SQISignHD, F is built from $\sigma : E_A \to E_2$ and scalars, so for the verification it suffices to give the action of σ on the 4-torsion. If it is not provided, we can just guess it and try the corresponding symplectic matrix, this greatly reduces the number of symplectic matrices to try to only a few choices.

We now explain how to handle the case of a d-isogeny F where d is not prime to 2. A difficulty is that if we only start with a symplectic basis $\mathcal{B}_A$ of $A[4]$, then since the kernel of F_1 may contain points of 4-torsion, F_1 does not induces a canonical symplectic basis of $C[4]$ anymore. So the algorithm needs to start with more data.

Theorem C.2.5. *Let $(A, \mathcal{L})$ be a principally polarised abelian variety, with $\mathcal{L}$ a symmetric ample line bundle. Let N be an integer, $\mathcal{B}'_A = (x'_1, \ldots, x'_g, y'_1, \ldots, y'_g)$ be a symplectic basis of $A[4N]$ with respect to a primitive $4N$-root of unity ζ, and $\mathcal{B}_A = N\mathcal{B}'_A$ the induced symplectic basis of $A[4]$. It induces a symmetric level 2 theta structure Θ_A on A.*

Let K be the maximal isotropic kernel generated by the points $P_j = 4x'_j$. Assume that we are given the theta coordinates of level 2 of the x'_j, and the theta coordinates of a point P in A.

Let $F : A \to B = A/K$ the induced isogeny. Then there is a unique descent of $\mathcal{L}^N$ to a symmetric line bundle $\mathcal{M}$ on B, and a symmetric theta structure on $\mathcal{M}$ induced by the symplectic basis

$$(F(x'_1), \ldots, F(x'_g), NF(y'_1), \ldots, NF(y'_g))$$

with respect to ζ^N of $B[4]$.

Furthermore, the theta null points of B and the theta coordinates of $F(P)$ can be computed in $O(N^g)$ arithmetic operations over the base field.

Proof. Since we are given the points (x'_i, y'_i) in level 2 theta coordinates, we can use the algorithms of [LR12; CR15; LR23] to construct a symmetric theta structure of level $2N$ on the theta group $G(\mathcal{L}^{2N})$. However, the references above assume for simplicity that the degree of the isogenies is prime to the level m of the symmetric theta structure we start with. Here $m = 2$ and N is no longer assumed to be odd. So we need the general case, which is described in [Rob10, Chapters 6 and 7], [Rob21, § 2.10, Remarks 2.10.3 and 2.10.7]. Once we are in level $2N$, we can apply Mumford's isogeny theorem [Mum, Theorem 4 p.302–303] to obtain a symmetric theta structure Θ_B of level 2 on B and the equations of the isogeny. The algorithm takes time $O(N^g)$. It remains to show that the theta structure we obtain on B is the one induced by the points $(F(x'_1), \dots, F(x'_g), NF(y'_1), \dots, NF(y'_g))$.

Let $Z(K)$ be the centralizer of K in the theta group $G(\mathcal{L}^{2N})$. The theta structure Θ_B is induced by the canonical map $\alpha_f : Z(K)/\tilde{K} \to G(\mathcal{M}^2)$ from [Mum, Equation (2) p.302], where $\tilde{K}$ is the canonical lift of K into the theta group $G(\mathcal{L}^{2N})$ induced by our theta structure of level $2N$. The points

$$(2F(x'_1), \dots, 2F(x'_g), 2NF(y'_1), \dots, 2NF(y'_g))$$

form a symplectic basis of $B[2]$. The theta structure Θ_B on $G(\mathcal{M}^2)$ is determined by the symmetric lifts of this basis into $G(\mathcal{M}^2)$. By definition of the induced theta structure, if T is in this basis, the symmetric lift $g_T \in G(\mathcal{M}^2)$ above T induced by Θ_B is given by the image by α_f of the symmetric lift $g_{T'} \in G(\mathcal{L}^{2N})$ induced by the theta structure of level $2N$ on A for any point T' such that $F(T') = T$. For $T = 2F(x'_j)$, we can take $T' = 2x'_j$. Since the theta structure on $G(\mathcal{L}^{2N})$ is determined by the basis (x'_j, y'_j), the lift $g_{T'}$ is determined as follow: let $g_{x'_j} \in G(\mathcal{L}^{4N})$ be any of the two symmetric lift of x'_j, and define $g_{T'} = \eta_2(g_{x'_j})$ where η_2 is defined in [Mum, § 2, p.310]; this does not depends on the choice of $g_{x'_j}$. Since α_f sends symmetric elements into symmetric elements and commutes with η_2, we get that $g_T = \eta_2(g_{F(x'_j)})$ where $g_{F(x'_j)}$ is any of the two symmetric element above $F(x'_j)$ in $G(\mathcal{M}^4)$. Likewise, when $T = 2NF(y'_j)$, we check that $g_T = \eta_2(g_{NF(y'_j)})$, for $g_{NF(y'_j)} \in G(\mathcal{M}^4)$ one of the two symmetric elements above $NF(y'_j)$. Hence the descent of the symmetric theta structure of level $2N$ on $G(\mathcal{L}^{2N})$ induced by the basis $(x'_1, \dots, x'_g, y'_1, \dots, y'_g)$ to a symmetric theta structure of level 2 on $G(\mathcal{M}^2)$ is indeed the one induced by $(F(x'_1), \dots, F(x'_g), NF(y'_1), \dots, NF(y'_g))$. □

Notice that $NF(y'_j) = F(y_j)$, so Theorem C.2.5 only needs the points (x'_j, y_j) as input. If N is odd, Theorem C.2.2 is a special case of Theorem C.2.5: by the CRT, from a symplectic basis (x_j, y_j) of $A[4]$ and a basis P_j of K, there is a unique $x'_j \in A[4N]$ which induces both x_j and P_j: $x_j = Nx'_j$, $P_j = 4x'_j$.

However when N is not prime to 2, we cannot start with any symplectic basis (x_j, y_j) of $A[4]$, it has to be compatible with our kernel K, in the sense that there should exists x'_j a basis of a maximal isotropic subgroup of $A[4N]$ which induces both x_j and P_j. In the situation of SQISignHD, where we convert our points given by tuple of Weierstrass coordinates into theta coordinates given by a product theta structure, the resulting product symplectic basis of $A[4]$ will not be compatible with our kernels in general. So to get an input suitable for Theorem C.2.5, we first start with the basis (P_j) of K (in tuple of Weierstrass coordinates), fix points x'_j above each P_j such that $P_j = 4x'_j$ and the x'_j generate an isotropic subgroup of $A[4N]$. Then we let $x_j = Nx'_j$, and fix a symplectic complement y_j of the x_j. We compute the symplectic matrix M that changes the product symplectic basis into the (x_j, y_j) and act by this matrix to get the theta coordinates in terms of our new basis (x_j, y_j).

We can adapt Proposition C.2.4 to the general case:

Proposition C.2.6. *Let $F_1 : A \to C$ be an d_1-isogeny, $\tilde{F}_2 : B \to C$ and d_2-isogeny and $F = F_2 \circ F_1$ an d-isogeny, where $d = d_1 d_2$. Let d' be a common multiple of d_1 and d_2, and write $d' = c_1 d_1 = c_2 d_2$.*

Let $(P_1, \dots, P_g)$ be a basis of the kernel of F_1, and $(x'_1, \dots, x'_g, y'_1, \dots, y'_g)$ a symplectic basis of $A[4d']$ with respect to ζ, a primitive $4d'$-root of unity, and such that $P_i = 4c_1 x'_i$.

Let $(Q_1, \dots, Q_g)$ be a basis of the kernel of $\tilde{F}_2$, and $(u'_1, \dots, u'_g, v'_1, \dots, v'_g)$ a symplectic basis of $B[4d']$ with respect to ζ, such that $Q_i = 4c_2 v'_i$.

Then the induced theta structure on C induced by F_1 and $\tilde{F}_2$ via Theorem C.2.5 is the same if

$$\tilde{F}_2(c_2 v'_i) = F_1(d' y'_i) \text{ and } F_1(c_1 x'_i) = \tilde{F}_2(d' u'_i) \tag{4}$$

Proof. The symplectic theta structure on C induced by F_1 is given by $(F_1(c_1 x'_i), F_1(d' y'_i))$, and the one induced by $\tilde{F}_2$ is given by $(\tilde{F}_2(d' u'_i), \tilde{F}_2(c_2 v'_i))$. □

We note that in the isogeny algorithm for F_1 we only need the points $c_1 x'_i, d' y'_i$ and for $\tilde{F}_2$ we only need the points $d' u'_i, c_2 v'_i$.

Corollary C.2.7. *Let $F = F_2 \circ F_1$. To get the same theta structure on C, it suffices to choose x'_i, y'_i, u'_i, v'_i such that $F(c_2 y'_i) = c_2 v'_i$ and $\tilde{F}(c_1 u'_i) = c_1 x'_i$.*

An algorithm to construct suitable x'_i, y'_i, u'_i, v'_i is as follow. Take y''_i a basis of a symplectic complement of $\operatorname{Ker} F_1$ in $A[d_1]$, $y'_i \in A[d']$ isotropic such that $4c_1 y'_i = y''_i$, and let $v'_i = F(y'_i)$. Then $Q_i = 4c_2 v'_i$ is a basis of $\operatorname{Ker} \tilde{F}_2$. We let u'_i be a symplectic complement of v'_i in $B[d']$, and we let $x'_i = \tilde{F}(u'_i)$. Let $P_i = 4c_1 x'_i$, they form a basis of $\operatorname{Ker} F_1$.

Corollary C.2.7 explain why we need $f \geq e_2 + 2$ in Remark 5.3 when $\ell = 2$ (an alternative if we are only given the action of σ on $E_A[\ell^{e_2}]$ would be to just guess it on $E_A[\ell^{e_2+2}]$).

C.3 Computing 2^e-isogenies

We need to compute an N-isogeny, with $N = d_1$ or $N = d_2$ with the notations from Appendix C.2. The isogeny algorithms described in [CR15; LR23] assume that the degree is prime to 2 for simplicity. For SQISignHD, we want to take

Algorithm 20: Finding two compatible basis.

```
Data: $\ker(F_1)$, $\ker(\widetilde{F_2})$, an algorithm to evaluate $F = F_1 \circ F_2$ on $A[d']$ and an
      algorithm to evaluate $\widetilde{F}$ on $B[d']$.
Result: Two symplectic basis $(x'_1, \cdots, x'_g, y'_1, \cdots, y'_g)$ of $A[4d']$ and
        $(u'_1, \cdots, u'_g, v'_1, \cdots, v'_g)$ of $B[4d']$ such that $\ker(F_1) = \langle c_1 y'_i \rangle_i$,
        $\ker(\widetilde{F_2}) = \langle c_2 v'_i \rangle_i$, $F(c_2 y'_i) = c_2 v'_i$ and $\tilde{F}(c_1 u'_i) = c_1 x'_i$ for all
        $i \in [\![1\ ;\ g]\!]$.
Find $(y''_1, \cdots, y''_g)$ a basis of a symplectic complement of $\mathrm{Ker}\, F_1$ in $A[d_1]$;
Find $(y'_1, \cdots, y'_g)$ forming an isotropic subgroup of $A[d']$ such that $4c_1 y'_i = y''_i$
for all $i \in [\![1\ ;\ g]\!]$;
$v'_i \longleftarrow F(y'_i)$ for all $i \in [\![1\ ;\ g]\!]$;
Let $(u'_1, \cdots, u'_g)$ be a symplectic complement of the $v'_i$ in $B[d']$;
$x'_i \longleftarrow \tilde{F}(u'_i)$ for all $i \in [\![1\ ;\ g]\!]$;
Return $(x'_1, \cdots, x'_g, y'_1, \cdots, y'_g)$ and $(u'_1, \cdots, u'_g, v'_1, \cdots, v'_g)$;
```

$\ell = 2$ (so that $N = 2^e$) for efficiency. The general case of N even is described in [Rob10; Rob21]. In this section we focus on the case $N = 2^e$ and detail how the general algorithm can be used to compute 2-isogenies. Handling 2-isogenies is actually easier because we can use the duplication formulae directly.

With the notations of Theorem C.2.5, we assume that we are given the $4N$-torsion points $(x'_1, \ldots, x'_g)$ given in theta coordinates by the symmetric theta structure of level 2 induced by a symplectic basis $(x_1, \ldots, x_g, y_1 \ldots, y_g)$ of $A[4]$ with $x_j = N x'_j$. Let $P_j = 4x'_j$, K the subgroup generated by the P_j, $F : A \to B = A/K$ the corresponding isogeny.

We then write $N = 2N'$, and let $T_j = N' P_j$, the kernel of a 2-isogeny f through which F factorizes. We remark also that $T''_j = N' x'_j$ is a point of 8-torsion such that $x_j = 2T''_j$. We will explain how to compute the isogenous theta null point of the codomain of f, and how to push points through f. The points $f(x'_j)$ will then be points of $4N'$ torsion, and we iterate.

So from now on we let f be a 2-isogeny $A \to B$, $K = \langle T_1, \ldots, T_g \rangle$ be the kernel of f, $(T''_1, \ldots, T''_g)$ be points in $A[8]$ such that $T_j = 4T''_j, j = 1, \ldots, g$. We assume that the symmetric theta structure Θ_A on A is induced by a symplectic basis $(T'_1, \ldots, T'_g, U'_1, \ldots U'_g)$ where $T'_j = 2T''_j$ and that we are given the coordinates of the T''_j, $j \in \{1, \ldots, g\}$. If $i \in (\mathbb{Z}/2\mathbb{Z})^g$, we let $T''_i = \sum_{j=1}^{g} i_j T''_j$. For simplicity, we will even assume that we are given the theta coordinates of all T''_i, $i \in (\mathbb{Z}/2\mathbb{Z})^g$, in particular the points $T_i, i \in (\mathbb{Z}/2\mathbb{Z})^g$ span the full kernel K.

The 2-isogeny formula will be derived from the duplication formula [Igu72, Theorem 2 p. 139, p. 141]:

$$\theta \left[\begin{smallmatrix} a_1 \\ b_1 \end{smallmatrix}\right] (z_1, \Omega)\theta \left[\begin{smallmatrix} a_2 \\ b_2 \end{smallmatrix}\right] (z_2, \Omega) = \sum_{t \in \frac{1}{2}\mathbb{Z}^g/\mathbb{Z}^g} \theta \left[\begin{smallmatrix} \frac{a_1+a_2}{2}+t \\ b_1+b_2 \end{smallmatrix}\right] (z_1+z_2, 2\Omega)\theta \left[\begin{smallmatrix} \frac{a_1-a_2}{2}+t \\ b_1-b_2 \end{smallmatrix}\right] (z_1-z_2, 2\Omega) \tag{5}$$

$$2^g\theta \left[\begin{smallmatrix} a_1 \\ b_1 \end{smallmatrix}\right] (z_1, \Omega)\theta \left[\begin{smallmatrix} a_2 \\ b_2 \end{smallmatrix}\right] (z_2, \Omega) = \sum_{t \in \frac{1}{2}\mathbb{Z}^g/\mathbb{Z}^g} e^{-2\pi i (a_1|2t)}\theta \left[\begin{smallmatrix} a_1+a_2 \\ \frac{b_1+b_2}{2}+t \end{smallmatrix}\right] (\frac{z_1+z_2}{2}, \frac{\Omega}{2})\theta \left[\begin{smallmatrix} a_1-a_2 \\ \frac{b_1-b_2}{2}+t \end{smallmatrix}\right] (\frac{z_1-z_2}{2}, \frac{\Omega}{2}). \tag{6}$$

We will derive algebraic formula for the analytic isogeny $\phi : A = \mathbb{C}^g/(\Omega_A\mathbb{Z}^g + \mathbb{Z}^g) \to B = \mathbb{C}^g/(\Omega_B\mathbb{Z}^g + \mathbb{Z}^g), z \mapsto 2z$ where $\Omega_B = 2\Omega_A$. We will then explain how to use these formula to compute our algebraic isogeny f.

Recall from Appendix C.1 that the Hadamard matrix H allows to convert from the theta coordinates θ_i^A to the dual theta coordinates ${\theta'}_\chi^A$. Given two points P_1, P_2 given by theta coordinates $\theta_i(P_j)$, we also let $(\theta_i(P_1)) \star (\theta_i(P_2)) = (\theta_i(P_1) \cdot \theta_i(P_2))_{i \in (\mathbb{Z}/2\mathbb{Z})^g}$.

Proposition C.3.1. *Let P be a point on A. Then the theta coordinates of the points $\phi(P) \in B$, where $\phi : A \to B$ is the isogeny defined above, are given by:*

$$(\theta_i^B(\phi(P))) \star (\theta_i^B(0))_{i \in (\mathbb{Z}/2\mathbb{Z})^g} = H \cdot \left(({\theta'}_\chi^A(P))_{\chi \in (\hat{\mathbb{Z}}/2\hat{\mathbb{Z}})^g} \star ({\theta'}_\chi^A(P))_{\chi \in (\hat{\mathbb{Z}}/2\hat{\mathbb{Z}})^g} \right).$$

Proof. Using the duplication formula, we obtain (with $\Omega = \Omega_A$):

$$\theta \left[\begin{smallmatrix} 0 \\ i \end{smallmatrix}\right] (2z, \Omega)\theta \left[\begin{smallmatrix} 0 \\ i \end{smallmatrix}\right] (0, \Omega) = \sum_{t \in \frac{1}{2}\mathbb{Z}^g/\mathbb{Z}^g} \theta \left[\begin{smallmatrix} t \\ 2i \end{smallmatrix}\right] (2z, 2\Omega)\theta \left[\begin{smallmatrix} t \\ 0 \end{smallmatrix}\right] (2z, 2\Omega)$$

This means that $(\theta_i^B(\phi(P))\theta_i^B(0))_{i \in (\mathbb{Z}/2\mathbb{Z})^g} = H \cdot \left({\theta'}_\chi^A(P)^2\right)_{\chi \in (\hat{\mathbb{Z}}/2\hat{\mathbb{Z}})^g}$. □

So the image of a point P by a 2-isogeny ϕ is simple to compute, provided we know the theta null point $\theta_i^B(0) = \theta \left[\begin{smallmatrix} 0 \\ i/2 \end{smallmatrix}\right] (0, 2\Omega)$ of B: start with the theta coordinates $\theta_i^A(P)$ of P, apply the Hadamard transform to get the dual coordinates ${\theta'}_\chi^A(P)$, square these coordinates, and apply the Hadamard transform again to obtain $\theta_i^B(\phi(P))\theta_i^B(0)$. It now only remains to divide by the coordinates $\theta_i^B(0)$ given by the theta null point of B.

Corollary C.3.2. *Assume that we are given the theta coordinates $\theta_i^A(P)$ of $P \in A$ and of the theta null point $\theta_i^B(0)$ of B. After a precomputation of 2^g inversions to invert the coordinates of the theta null point of B, the theta coordinates of $\phi(P)$ can be computed in 2 Hadamard transforms, 2^g squares and 2^g multiplications.*

Furthermore, given the theta null points of A and B one can check (up to signs) that A and B are indeed 2-isogenous (with compatible theta structure) using 2^{g+1} squares and 2 Hadamard transforms.

Proof. The first statement follows from the Proposition. For the second statement, if A and B are 2-isogenous, then $\theta_i^B(0)^2 = H \cdot {\theta'}_\chi^A(0)^2$, which determines the $\theta_i^B(0)$ up to a sign. □

In particular the proof of Corollary C.3.2 shows that we can easily compute the square of the coordinates of the theta null point of B. In practice in our complexity estimates, we will often neglect the Hadamard transforms since they just amount to some additions and subtractions. It remains to compute the correct square roots.

Proposition C.3.3. *Let $i \in \mathbb{Z}^g$ and $z_i' \in \mathbb{C}^g$ be the analytic theta point given by the affine coordinates $\theta_t^A(z_i') = \theta \left[\begin{smallmatrix} 0 \\ i/4+t/2 \end{smallmatrix}\right](0, \Omega_A/2)$. Then $T_i' = \pi_A(z_i')$ is a point in $A[4]$. The points $T_i = 2T_i'$ generate the kernel corresponding to the kernel $\frac{1}{2}\mathbb{Z}^g/\mathbb{Z}^g$ of the isogeny ϕ.*

We have up to a constant (not depending on i):

$$\theta_i^B(0) = \sum_t \theta_t^A(z_i')^2. \tag{7}$$

Proof. Using the duplication formula again, we obtain for $i \in \mathbb{Z}^g$:

$$2^g \theta \left[\begin{smallmatrix} 0 \\ i/2 \end{smallmatrix}\right](0, \Omega) \theta \left[\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}\right](0, \Omega) = \sum_t \theta \left[\begin{smallmatrix} 0 \\ i/4+t/2 \end{smallmatrix}\right](0, \Omega/2)^2.$$

Up to the projective factor $2^g \theta \left[\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}\right](0, \Omega)$, we recover Eq. (7). We also check that although z_i' depends on the choice of $i \in \mathbb{Z}^g$, the term on the left only depends on the reduction of i in $(\mathbb{Z}/2\mathbb{Z})^g$. □

In the analytic setting, when representing the points z_i' and T_i' by theta coordinates, they actually are represented by the same coordinates, but z_i' is represented by affine coordinates while T_i' is represented by projective coordinates. So the projection π_A amount to sending the affine point $(\theta_i(z_i')) \in \mathbb{A}^{2^g}$ to the projective point in $\mathbb{P}^{2^g-1}$.

We go back to the algebraic setting: we let $f : A \to B$ be a 2-isogeny with kernel generated by the points (T_j), $j = 1, \ldots, g$ and we assume that we are given isotropic points T_j'' such that $T_j = 4T_j''$, and the T_j'' are expressed as theta coordinates with respect to the theta structure induced by a symplectic basis $(T_1', \ldots, T_g', U_1', \ldots U_g')$ where $T_j' = 2T_j''$. We need to compute the theta null point on B for our algebraic isogeny f like we did for our analytic isogeny ϕ above in Proposition C.3.3. Then we can apply Proposition C.3.1 to compute the image by f of any point P.

Notice that Eq. (7), because of the sum, only make sense for points in affine coordinates. So we cannot apply it directly in an algebraic algorithm, because we only have the 4-torsion points $T_i', i \in (\mathbb{Z}/2\mathbb{Z})^g$ in projective coordinates. We need to lift the projective points T_i' into an affine point $\tilde{T}_i'$ such that Eq. (7) make sense. We follow the terminology of [LR12, § 3] and speak about affine lifts. Riemann relations give a well defined doubling and differential addition law on affine lifts, hence a scalar multiplication.

Since we have the theta null point of level 2 on A, induced by the symplectic basis $(T_1', \dots, T_g', U_1', \dots U_g')$ of $A[4]$, the induced theta structure gives a canonical affine lift lift $\tilde{T}_i$ of the T_i, for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$, and a canonical affine translation $\tilde{P} \mapsto \tilde{P} + \tilde{T}_i$ for all affine lifts $\tilde{P}$ of a point $P \in A$.

We take an arbitrary affine lift $\tilde{T_i''}$ of T_i'', which we then normalize via the equation

$$-3\tilde{T_i''} = \tilde{T_i''} + \tilde{T}_i, \tag{8}$$

where the translation on the right is the canonical one induced by the theta structure. This rigidifies our choice of affine lift up to a root of unity μ of order $3^2 - 1^2 = 8$. Then $2\tilde{T_i''}$ is rigidified up to the action of $\mu^{2^2} = \pm 1$, and so if we let $\tilde{T_i'} = 2\tilde{T_i''}$, its coordinates are fully determined up to a sign. We can now apply Eq. (7), with $z_i' = \tilde{T_i'}$:

$$\theta_i^B(0) = \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \theta_t^A(\tilde{T_i'})^2 = \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \theta_t^A(2\tilde{T_i''})^2. \tag{9}$$

Because this equation only involves the squares of the coordinates of $\tilde{T_i'}$, our remaining sign ambiguity does not matter.

Remark C.3.4. If we only had the points T_i' but not the T_i'', we could rigidify the choice of $\tilde{T_i'}$ via the equation $2\tilde{T_i'} = \tilde{T}_i$. This equation rigidifies the lift up to a root of unity μ of order $2^2 = 4$, so it remains a sign ambiguity in Eq. (9). However it is enough to do this rigidification for i going through $e_1, \dots, e_g, e_1 + e_2, \dots, e_j + e_k, \dots, e_{g-1} + e_g$ where e_j is a a basis of $(\mathbb{Z}/2\mathbb{Z})^g$. The remaining choices of $\tilde{T_i'}$ are then fully determined from the Riemann relations, in particular from three way additions and differential additions. Thus we only have $g(g+1)/2$ signs ambiguity rather than 2^g, and one can prove that all signs are actually valid [Rob10, Proposition 6.3.5]: they each correspond from a different choice of a 4-symplectic basis above our fixed 2-symplectic basis $(f(T_j'), f(U_j))$.

Nevertheless, when computing an 2^e-isogeny, we need to be careful that this choice of 4-symplectic basis is compatible with our next kernel. Also in the end for the equality testing of Appendix C.2, we need to be sure to have chosen the correct theta structure. That's why we have to assume that we are given the T_i'', not only the T_i', this allows us to fully rigidify the theta structure on B.

Proposition C.3.5. *Let T_j'', $j = 1, \dots, g$ be points of 8-torsion on A which generates an isotropic subgroup. Let K be the kernel generate by the $4T_j''$. Assume that we are given the theta coordinates of the T_j'' via the theta structure induced by a symplectic basis $(2T_j', U_j')$. Let $(\theta_i^B(0))_{i \in (\mathbb{Z}/2\mathbb{Z})^g}$ be the projective theta null point of $B = A/K$ given by Theorem C.2.5.*

Then if $i \in (\mathbb{Z}/2\mathbb{Z})^g$, the value $\theta_i^B(0)$ (up to a constant which does not depends on i) can be computed using Eq. (9) where $\tilde{T_i''}$ is an affine lift of T_i normalised using Eq. (8). This requires tripling an affine lift of $T_i'' = \sum_{j=1}^g i_j T_j''$, a division and multiplication, and 2^g squares. If the tripling is computed via a doubling followed by a differential addition, it can be done in 2^{g+2} multiplications, 2^{g+1} doubling, and 2^g divisions.

The total cost to compute the theta null point of B is then $(2^g-1)(2^{g+2}+1)$ multiplications, $2^g(2^{g+1}-1)$ squares, $(2^g-1)(2^g+1)$ divisions, and 2^{g+1} inversions, that is $2^g(7\cdot 2^g-2)-2$ arithmetic operations.

Proof. We use Eq. (9) to compute $\theta_i^B(0)$. We take an arbitrary lift $\tilde{T_i''}$ and compute $3\tilde{T_i''}$. We use Eq. (8) to compute the correct normalisation, this costs one division. We then plug in Eq. (9), this costs 2^g squares, and one multiplication by our normalisation factor.

But we remark that if $3\tilde{T_i''}$ is computed through a doubling $2\tilde{T_i''}$ followed by a differential addition $2\tilde{T_i''}+\tilde{T_i''}$, then the squares of the theta coordinates of $2\tilde{T_i''}$ are already computed.

The cost of the doubling and differential addition is described in [Rob10, Table 4.1]. The computation of $\theta_i^B(0)$ for $i\neq 0$ then costs $4\cdot 2^g+1$ multiplications, $2\cdot 2^g$ squares and 2^g+1 divisions.

This costs assume the precomputation of some constants depending only on A (more precisely its theta null point), which takes 2^g squares and 2^{g+1} inverses to compute once and for all. Taking this precomputation into account we get our final complexity. □

For computing 2^e-isogenies decomposed as e 2-isogenies, we start with g points of 2^{e+2} torsions x_j', and even with the 2^g points x_i' for $i\in(\mathbb{Z}/2\mathbb{Z})^g$. We compute the 2^g points of each kernel using Fig. 2. For this for each 2-isogeny we need to apply Proposition C.3.5 to compute the isogenous theta null point, and also apply Proposition C.3.1 2^g times to push the points through each isogeny, each image costing 2^{g+1} arithmetic operations by Corollary C.3.2.

Example C.3.6. In dimension $g=4$, by Corollary C.3.2 computing the squares of the 2-isogeneous theta null point cost $2^g=16$ squares over the base field. To get the correct 2^g square roots, by Proposition C.3.5, it costs a staggering 1758 arithmetic operations over the base field for just one theta null point.

An estimation to compute a 2^e-isogeny with $e=128$ in dimension 4, taking into account the computation of the e theta null points, doublings and pushing all torsion points through the isogenies amounts to roughly 1 million arithmetic operations over $\mathbb{F}_{p^2}$ (more precisely: 856224 operations).

We leave optimisations of this algorithm tailored for SQISignHD for future work.