Breaking and Repairing SQIsign2D-East

Wouter Castryck, Mingjie Chen, Riccardo Invernizzi, Gioella Lorenzon and Frederik Vercauteren

firstname.lastname@esat.kuleuven.be

COSIC, ESAT, KU Leuven, Belgium

Abstract. We present a key recovery attack on SQIsign2D-East [NO24b] that reduces its security level from λ to $\lambda/2$. We exploit the fact that each signature leaks a Legendre symbol modulo the secret degree of the private key isogeny. About $\lambda/2$ signatures are enough for these Legendre symbols to fully determine the secret degree, which can then be recovered by exhaustive search over a set of size $O(2^{\lambda/2})$. Once the degree is known, the private key isogeny itself can be found, again by exhaustive search, in time $\tilde{O}(2^{\lambda/2})$.
We also present a new version of the protocol which does not leak any such information about the private key and show that our modified protocol is more efficient than the original one. Finally, we give a security analysis as well as a new proof of security.

Keywords: Isogeny-based cryptography, SQIsign2D-East, Legendre symbol, cryptanalysis

1 Introduction to SQIsign2D-East

We give an overview of the signature scheme SQIsign2D-East, referring to the original paper [NO24b] for further details.

The protocol builds on ideas from SQIsignHD [DLRW24] but introduces algorithms that avoid the need for 4-dimensional isogeny computations for signature verification, reducing to only 2-dimensional isogeny computations. It is obtained via the Fiat–Shamir transform applied to an identification protocol that is based on the diagram depicted in Figure 1. The prover generates an isogeny $\tau : E_0 \to E_A$ as secret key, where E_0 is a public supersingular elliptic curve of known (special extremal) endomorphism ring, and publishes E_A as the corresponding public key. To prove knowledge of the secret isogeny τ and hence of $\mathrm{End}(E_A)$, the prover computes another secret isogeny $\psi : E_0 \to E_1$, publishing the codomain as commitment. Upon receiving as challenge an isogeny $\phi : E_1 \to E_2$, the prover then responds with an isogeny $\sigma : E_A \to E_2$, obtained by transforming $\phi \circ \psi \circ \hat{\tau}$ into an equivalent isogeny. In order to enable the verifier

* This work was supported in part by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement ISOCRYPT - No. 101020788) and by CyberSecurity Research Flanders with reference number VR20192203. Date of this document: September 17, 2024.

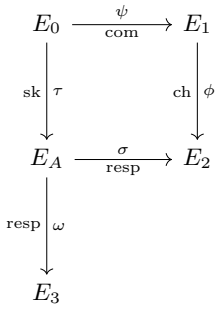

Fig. 1. High-level depiction of the SQIsign2D-East protocol

to evaluate σ via an isogeny in dimension two, the prover will also provide an auxiliary isogeny $\omega : E_A \to E_3$.

The public parameters $(p, a, b, E_0, P_0, Q_0, \mathcal{O}_0)$ of SQIsign2D-East targeting a claimed security level of λ bits are set as follows :

- p is a prime of the form $p = 2^{a+b}f - 1$ with f a small integer cofactor and $a \approx b \approx \lambda$, so $p \approx 2^{2\lambda}$;
- E_0 is the elliptic curve of equation $y^2 = x^3 + x$ over $\mathbb{F}_{p^2}$;
- P_0, Q_0 are a basis of $E_0[2^{a+b}]$;
- $\mathcal{O}_0$ is $\mathbb{Z}\langle 1, \mathbf{i}, \frac{\mathbf{i}+\mathbf{j}}{2}, \frac{1+\mathbf{k}}{2}\rangle$, the maximal order in the quaternion algebra ramified at p and infinity, isomorphic to $\mathrm{End}(E_0)$.

Before giving a summary of how the scheme works we recall the algorithms involved in SQIsign2D-East.

1.1 Algorithms for SQIsign2D-East

The following algorithms are well-known and they are used e.g. in SQISign [DFKL+20]:

- $\texttt{RandomEquivalentIdeal}_M(I)$: given an integer M and a left $\mathcal{O}_0$-ideal I, outputs a uniformly random equivalent ideal $J \sim I$ of norm $n(J) < M$. When $M \gtrapprox p^{1/2}$, such an ideal exists with high probability.
- $\texttt{EichlerModConstraint}(I, \gamma, \delta)$: given a left $\mathcal{O}_0$-ideal I of prime norm N and elements $\gamma, \delta \in \mathcal{O}_0$, outputs $(C_0 : D_0) \in \mathbb{P}^1(\mathbb{Z}/N\mathbb{Z})$ such that $\gamma(C_0\mathbf{j} + D_0\mathbf{k})\delta \in \mathbb{Z} + I$.
- $\texttt{StrongApproximation}_M(N, C_0, D_0)$: given integers M, N, C_0, D_0, N prime, outputs $\mu \in \mathcal{O}_0$ such that $n(\mu) = M$ and $\mu = m(C_0\mathbf{j} + D_0\mathbf{k}) + N\mu_1$ for some integer m and some $\mu_1 \in \mathcal{O}_0$.

The next algorithm, called $\texttt{KaniCod}$, comes from the attacks on SIDH [CD23, MMP+23, Rob23] and is based on Kani's lemma [Kan97]. In particular, it uses [MMP+23, Theorem 1]. Let N_1, N_2 be coprime integers and let $D = N_1 + N_2$. Consider a commutative diagram of elliptic curve isogenies:

$$\begin{array}{ccc} E_0 & \xrightarrow{\psi_2} & E_2 \\ \downarrow{\scriptstyle \psi_1} & & \downarrow{\scriptstyle \psi_1'} \\ E_1 & \xrightarrow{\psi_2'} & E_3 \end{array}$$

where $\deg(\psi_1) = \deg(\psi_1') = N_1$ and $\deg(\psi_2) = \deg(\psi_2') = N_2$.

– $\texttt{KaniCod}(N_1, N_2, E_1, E_2, P_1, Q_1, P_2, Q_2; S_1, S_2)$: given the coprime integers N_1, N_2 and the elliptic curves E_1, E_2, a basis P_1, Q_1 of $E_1[D]$, the basis

$$P_2 = \psi_2(\hat{\psi}_1(P_1)), \qquad Q_2 = \psi_2(\hat{\psi}_1(P_2))$$

of $E_2[D]$, and finite subsets S_1, S_2 of E_1, E_2, respectively, outputs E_0 and the (element-wise) images of S_1, S_2 under $\hat{\psi}_1, \hat{\psi}_2$, respectively.

The following algorithm, called `GenRandIsogImg`, is a generalized version of `RandIsogImg` from [NO24a], which computes the codomain and point images of a given degree isogeny ι from a specific elliptic curve E_0. The generalized algorithm introduced for SQIsign2D-East allows now to compute the codomain and point images of a given degree isogeny ι from any elliptic curve E, given a prime-degree isogeny from E_0 to E. `GenRandIsogImgWithIdeal` is the same algorithm, with the corresponding ideal I_ι as an additional output.

Algorithm 1 $\texttt{GenRandIsogImg}_{I_\tau}(d,\ D;\ S)$

Input: An isogeny $\tau : E_0 \to E$ of prime degree N, its corresponding ideal I_τ, coprime integers d, D such that $D \approx p$, $d > N^3$, $d < D$ and a finite set $S \subset E$
Output: $(F, \iota(S))$ for a random d-isogeny $\iota : E \to F$.
1: $(C_0 : D_0) \leftarrow \texttt{EichlerModConstraint}(I_\tau, 1, 1)$.
2: $\alpha \leftarrow \texttt{StrongApproximation}_{d(D-d)}(N, C_0, D_0)$.
3: Let P, Q be a basis of $E[D]$.
4: $(F; \iota(S); \emptyset) \leftarrow \texttt{KaniCod}(d, D-d, E, E, P, Q, \alpha(P), \alpha(Q); S, \emptyset)$.
5: **return** $(F, \iota(S))$

This is Algorithm 2 in [NO24b] and is needed to compute the response σ and the auxiliary path ω. Note that α is to be viewed as an endomorphism of E by pushing forward under τ, see [DFKL+20, Section 4.2]. We will propose a different version of this algorithm in Section 3.

The next algorithm, called `AuxiliaryPath`, is used to compute the auxiliary path ω; this is Algorithm 3 in [NO24b]. There, `AuxiliaryPath` is used with $D_1 = 2^a \approx p^{1/2}$, $D = 2^{a+b} \approx p$ and $d = 2^a - q$, where q is the degree of the signature σ.

Algorithm 2 $\texttt{AuxiliaryPath}_{I_\tau}(d, D_1, D; S)$

Input: An isogeny $\tau : E_0 \to E$ of prime degree N_τ, its corresponding ideal I_τ, integers d, D_1, D such that d is coprime to both D_1 and D, $D \approx p$, $d(D_1 - d) > N^3$, $d(D_1 - d) < D$, and a finite set $S \subset E$.
Output: $(F, \omega(S))$ for a random d-isogeny $\omega : E \to F$.
1: Let P, Q be a basis of $E[D_1]$
2: $(F', \iota(P), \iota(Q)) \leftarrow \texttt{GenRandIsogImg}_{I_\tau}(d(D_1 - d), D; P, Q)$
3: $(F; \omega(S); \emptyset) \leftarrow \texttt{KaniCod}(d, D_1 - d, E, F', P, Q, \iota(P), \iota(Q); S; \emptyset)$
4: **return** $(F, \omega(S))$

The degree q is obtained through a randomized procedure, where it is being rejected and resampled in case it does not satisfy either of the following conditions, which the authors refer to as "$(2^a, 2^b, N_\tau)$-niceness". These conditions are necessary for `AuxiliaryPath` to succeed.

Definition 1. *We say that a positive integer q is $(2^a, 2^b)$-nice if q is odd, $q < 2^a$ and $q(2^a - q) < 2^{a+b}$.*

Definition 2. *We say that a positive integer q is $(2^a, 2^b, N_\tau)$-nice if q is $(2^a, 2^b)$-nice and satisfies*

$$\left(\frac{M(q)}{N_\tau}\right) = \left(\frac{-1}{N_\tau}\right), \tag{1}$$

where $M(q) = d(2^{a+b} - d)$ with $d = q(2^a - q)$, and $\left(\frac{\cdot}{\cdot}\right)$ is the Legendre symbol.

In Section 2 we will explain that these requirements turn out to contradict the claimed λ-bit security level of the protocol. In Section 3 we will propose a different version of this algorithm avoiding such conditions.

1.2 SQIsign2D-East

We give a short description of the identification scheme underlying SQIsign2D-East, and refer to [NO24b, Section 4.1] for a more thorough exposition.

Key generation. First, the degree of the secret isogeny τ is sampled as a random prime $N_\tau < p^{1/4}$. Then, a 2-dimensional representation $(N_\tau, \tau(P_0), \tau(Q_0))$ of τ, the corresponding ideal I_τ and the codomain curve E_A are computed using `GenRandIsogImgWithIdeal`. The secret key is $\mathrm{sk} = (\tau, I_\tau)$ and the public key is $\mathrm{pk} = E_A$.

Commitment. The commitment works similarly to the key generation. First, the degree of ψ is sampled as an odd integer $N_\psi < 2^{2\lambda}$. Then, an efficient representation $(N_\psi, \psi(P_0), \psi(Q_0))$ of ψ, the corresponding ideal I_ψ and the codomain curve E_1 are computed via `GenRandIsogImgWithIdeal`. The commitment consists of $\mathrm{com} = E_1$.

Challenge. The challenge is computed by sampling a point K_1' of order 2^b on E_1. The isogeny ϕ has kernel generated by K_1' and can be computed using Vélu's formulas. The challenge is $\mathrm{ch} = K_1'$.

Response. As a response, the prover must produce (1) an isogeny $\sigma : E_A \to E_2$ and (2) an auxiliary isogeny ω from E_A.

1. To obtain σ, the prover computes the left ideal $J = \bar{I}_\tau I_\psi I_\phi$ corresponding to the composition $\phi \circ \psi \circ \hat{\tau}$; the ideal I_ϕ can be computed via `IsogenyToIdeal` [Ler22, Algorithm 20] on input $\phi : E_1 \to E_2$, $\psi : E_0 \to E_1$ and the corresponding ideal I_ψ. Then, the prover uses $\texttt{RandomEquivalentIdeal}_{2^a}(J)$ to find an equivalent ideal

$$I_\sigma = J\frac{\bar{\alpha}}{N_\tau N_\psi 2^b} \sim J$$

such that $n(I_\sigma) < 2^a$. The norm $q = n(I_\sigma)$ is required to be $(2^a, 2^b, N_\tau)$-nice (see Definition 2), in order to allow the computation of the auxiliary isogeny with `AuxiliaryPath` in the next step. An efficient representation (q, R_2', S_2') of the isogeny σ is computed as follows: the images R_2', S_2' under σ of the 2^a-torsion are obtained as images under

$$\frac{\phi \circ \psi \circ \hat{\tau} \circ \hat{\alpha}}{N_\psi N_\tau} = 2^b\sigma \tag{2}$$

of a canonical basis P_A, Q_A of $E_A[2^{a+b}]$. Note that the equality (2) follows from a subtle application of the Deuring correspondence, see e.g. [BFD+24, Lemma 11].
2. An auxiliary isogeny ω from E_A having degree $2^a - q$ is obtained by invoking $\texttt{AuxiliaryPath}_{I_\tau}(2^a - q, 2^a, 2^{a+b}; 2^bP_A, 2^bQ_A)$, which returns the codomain E_3 of ω and images $R_3' = 2^b\omega(P_A)$, $S_3' = 2^b\omega(Q_A)$. The prover finds an efficient representation $(q(2^a - q), U_2', V_2')$ of $\sigma \circ \hat{\omega}$ by letting M be the change-of-basis matrix such that $(P_3', Q_3') = (R_3', S_3')M$, where P_3', Q_3' form the canonical basis of $E_3[2^a]$, and computing

$$(U_2', V_2') = -(R_2', S_2')M = (\frac{1}{q}\sigma\hat{\omega}(P_3'), \frac{1}{q}\sigma\hat{\omega}(Q_3')).$$

The last equality is obtained by taking images under $\sigma \circ \hat{\omega}$ on both sides of $(P_3', Q_3') = (R_3', S_3')M = (2^b\omega(P_A), 2^b\omega(Q_A))M$. (The leading factor $1/q$ slightly simplifies the verification step; note that q is odd by the $(2^a, 2^b, N_\tau)$-niceness.)

The response is $\mathrm{resp} = (E_3, U_2', V_2')$.

Verification. The verifier computes a $(2^a, 2^a)$-isogeny $\Phi : E_3 \times E_2 \to A$ with kernel $K = \langle (P_3', U_2'), (Q_3', V_2') \rangle$ and checks whether the codomain A is isomorphic (as a principally polarized abelian surface) to a product of elliptic curves, one of

which being E_A. If that is the case, the response is accepted, otherwise rejected.

The protocol satisfies correctness, by the fact that for an honestly generated response

$$K = \langle (P'_3, \frac{1}{q}\sigma\hat{\omega}(P'_3), (Q'_3, \frac{1}{q}\sigma\hat{\omega}(Q'_3))\rangle = \langle (qP'_3, \sigma\hat{\omega}(P'_3), (qQ'_3, \sigma\hat{\omega}(Q'_3))\rangle$$

and by [MMP+23, Theorem 1] a $(2^a, 2^a)$-isogeny Φ of kernel K is, up to post-composition with an isomorphism, of the form

$$\Phi : \begin{pmatrix} \hat{\omega} & -\hat{\sigma} \\ [\omega]_*\sigma & [\sigma]_*\omega \end{pmatrix} : E_3 \times E_2 \to E_A \times F$$

where F is the codomain of both $[\omega]_*\sigma$ and $[\sigma]_*\omega$.

2 The attack

We show that for the SQIsign2D-East protocol targeting a security level of λ, we can perform a key recovery attack with cost $2^{\lambda/2}$, ignoring polynomial overhead. The estimate from [NO24b] on the cost of key recovery arises as follows: first, an attacker has to guess the secret prime degree $N_\tau \sim 2^{\lambda/2}$ of the secret key $\tau : E_0 \to E_A$. Then the attacker has to guess the secret isogeny τ itself, for which there are $N_\tau + 1 \sim 2^{\lambda/2}$ possibilities; thus there are $\sim 2^\lambda/\lambda$ possible secret isogenies. However, the signatures turn out to leak information about N_τ, in the form of an equality involving Legendre symbols. We argue that about $\lambda/2$ signatures are enough for N_τ to be uniquely determined by these equalities; its precise value can then be recovered by simple brute-force over all primes of size $\sim 2^{\lambda/2}$. Once we have recovered N_τ, we can perform another brute-force search for the secret isogeny τ from E_0. The total cost of the attack is then still in the order of $2^{\lambda/2}$. We now describe the attack in greater detail.

2.1 Recovering N_τ

Recall that the secret degree N_τ is a random prime in $(0, p^{1/4})$. The response σ is an isogeny of odd degree $q < 2^a$, represented by the kernel $K = \langle (P'_3, U'_2), (Q'_3, V'_2)\rangle$ of a $(2^a, 2^a)$-isogeny Φ having σ as one of its components. As a first step, we observe that an attacker can evaluate Φ and consequently σ at any input; the degree q can be then recovered using a pairing computation combined with an easy discrete log computation. Therefore, it can be assumed that q is known.

It is important to note that q varies with every signature and is essentially random, subject to three conditions coming from the $(2^a, 2^b, N_\tau)$-niceness of q (see Definition 2); this is enforced in [NO24b, Algorithm 7, step 3]. In particular, for each signature we learn that $d(2^{a+b} - d)$ with $d = 2^a - q$ has the same quadratic residuosity as -1 mod N_τ. From Dirichlet's theorem on arithmetic

progressions it follows that, as soon as $M(q)$ is not the square of an integer, the density of primes N_τ satisfying (1) is 50%. Thus, heuristically, we expect that N_τ is uniquely determined by about $\lambda/2$ values of q. This means that after seeing roughly $\lambda/2$ signatures we should be able to find N_τ by simply brute-forcing over all primes in $(0, p^{1/4})$ and testing whether (1) holds for each of the corresponding values of q. Ignoring polynomial overhead, this step therefore has a complexity of $\tilde{O}(2^{\lambda/2})$.

Computational experiments (reported in Figure 2) confirm this heuristic assumption. For different sizes of N_τ, we computed the average number of signatures needed to uniquely determine it. In our experiment, to mimick the scheme, we fix a size e and sample a random prime $N_\tau < 2^e$. Then we pick random odd integers $q < 2^{2e}$, compute $d = q(2^{2e} - q)$ and $M = d(2^{4e} - d)$, and if $\left(\frac{M}{N_\tau}\right) = \left(\frac{-1}{N_\tau}\right)$ we return q. As expected, after slightly more than e signatures, N_τ is uniquely determined and can be correctly recovered by brute-force. Notice that the fact that we do not know the exact value of $\left(\frac{-1}{N_\tau}\right)$ does not matter.

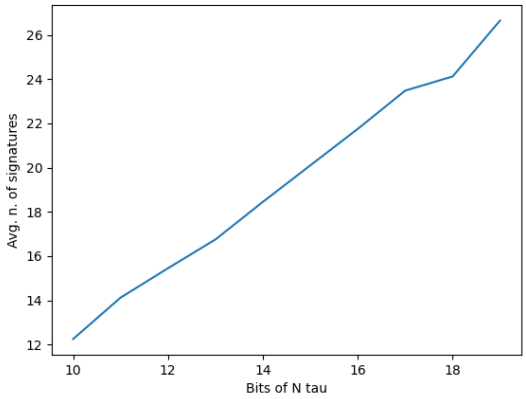

Fig. 2. Average number of signatures needed to determine a secret degree

Remark 3. The problem of recovering a bounded prime N_τ from polynomially many values $\left(\frac{a}{N_\tau}\right)$, for known but random a (in our case $a = -M(q)$), is a special case of breaking a randomized instance of Damgård's Legendre pseudorandom function [Dam88] where the offset is known but the modulus is not; see also [AG97, Problem 2]. We point out that the Legendre PRF is usually being studied in its unknown-offset but known-modulus variant, see e.g. [BBUV20]. Problems of this type are believed to be hard and, in the case of SQIsign2D-East, the security loss just comes from the fact that N_τ is too small. However, simply increasing the size of N_τ is not enough to save the protocol, because the bound $N_\tau < p^{1/3}$ is needed to build auxiliary paths during the signature

process [NO24b, Section 3.2]. If we allow N_τ to reach the bound of $p^{1/3}$, the cost of key recovery will increase to $p^{1/3} \sim 2^{2\lambda/3}$ operations, but this still falls short of the proposed security level. A different solution, based on avoiding the leakage of Legendre symbols, is presented in Section 3.

2.2 Recovering the secret key

Given the norm N_τ of the secret ideal I_τ, we can recover I_τ by enumerating all left $\mathcal{O}_0$-ideals of norm N_τ and checking whether the corresponding isogenies have codomain isomorphic to E_A. There will be $O(2^{\lambda/2})$ such ideals and they can be enumerated using the bijection from [KV10, Lemma 7.2]. Therefore, the cost of this step is $\tilde{O}(2^{\lambda/2})$.

On the other hand, given that ideal-to-isogeny translations are not very efficient in practice, we can as well use the same algorithm that is used in [NO24b] to generate I_τ. The key generation algorithm in SQIsign2D-East uses `RepresentInteger` to sample a random quaternion $\alpha \in \mathcal{O}_0$ of norm $N_\tau(2^{a+b} - N_\tau)$. There are no more than $O(2^{\lambda/2})$ different return values of this algorithm when fed with norm $N_\tau(2^{a+b} - N_\tau)$. Therefore, we simply run the key generation algorithm in SQIsign2D-East, and for each of the quaternion elements returned by `RepresentInteger`, we check whether its corresponding isogeny from E_0 of degree N_τ has codomain isomorphic to E_A or not. This also costs $\tilde{O}(2^{\lambda/2})$, but is more efficient in practice than the previous method.

3 A possible fix

The attack outlined in Section 2 exploits the fact that `GenRandIsogImg` imposes the constraint (1), linking the secret degree N_τ with the (public) degree of the response q. We want to make sure that q can be chosen independently from N_τ, so that an attacker observing it cannot learn secret information.

Recall that `GenRandIsogImg` in SQIsign2D-East first computes an endomorphism $\alpha \in \mathrm{End}(E_A)$ of degree

$$M(q) := q(2^a - q)(2^{a+b} - q(2^a - q)) \tag{3}$$

using `StrongApproximation`. This endomorphism is then passed to `KaniCod` twice to obtain first an isogeny of degree $q(2^a - q)$ and from that an isogeny ω of degree $2^a - q$, which is the auxiliary path to the response σ. Imposed by `StrongApproximation` and `KaniCod`, q needs to satisfy $(2^a, 2^b, N_\tau)$-niceness.

3.1 The fix: an auxiliary 3-isogeny

We now introduce our fix, where the main idea is that when q does not satisfy (1), we pass a different norm to `StrongApproximation`.

Definition 4. *We say that a positive integer q is $(2^a, 2^b)_3$-nice if q is $(2^a, 2^b)$-nice and $M(q)$ is divisible by 3.*

To start, in the key generation phase we additionally require that 3 is not a square modulo N_τ, i.e., we require $N_\tau \equiv 5, 7 \bmod 12$. Then, in the response phase, instead of requiring q to be $(2^a, 2^b, N_\tau)$-nice, we require it to be $(2^a, 2^b)_3$-nice; this condition is clearly independent from N_τ. Recall that in SQIsign2D-East, one calls `GenRandIsogImg` with parameters $d = q(2^a - q)$ and $D = 2^{a+b}$, which in turn invokes `StrongApproximation` with target norm $M(q) = d(D-d)$. Our two new conditions together allow us to modify the scheme as follows:

- if q satisfies condition (1) then we proceed as before;
- otherwise, we call `StrongApproximation` with target norm $M(q)/3$ to obtain an endomorphism α'. In this case we have that $M(q)/3$ satisfies

$$\left(\frac{M(q)/3}{N_\tau}\right) = \left(\frac{-1}{N_\tau}\right),$$

since

$$\left(\frac{M(q)/3}{N_\tau}\right)\left(\frac{3}{N_\tau}\right) = \left(\frac{M(q)}{N_\tau}\right)$$

and we required that $\left(\frac{3}{N_\tau}\right) = -1$ in the key generation phase. After that, we compute a random degree-3 isogeny $\alpha'' : E_A \to E_3$ using Vélu and we compose it with α' to finally obtain an isogeny α of degree $M(q)$ from E_A to E_3. The process is as shown in the diagram below, where $\alpha' = \hat{\psi}_2 \circ \psi_1$ and $\psi_1, \psi_2 : E_A \to E$ have degrees d and $(D-d)/3$ respectively.

$$\begin{array}{ccccc} E_A & & & & \\ \downarrow{\scriptstyle \psi_1} & \searrow{\scriptstyle \alpha'} & & & \\ E & \xrightarrow{\hat{\psi}_2} & E_A & \xrightarrow{\alpha''} & E_3 \end{array}$$

The subsequent steps are left unchanged. Notice that in the first call to `KaniCod`, the input is no longer an endomorphism. Rather, it is an isogeny from E_A to E_3 whose evaluation on $E_A[2^{a+b}]$ is known. Therefore, we can apply Kani's lemma in the same way. The modified version of [NO24b, Algorithm 2] is given in Algorithm 3.

3.2 Impact on efficiency

The proposed alternative strategy comes with additional requirements with respect to the original protocol. We now argue that these requirements do not negatively impact the overall efficiency.

In the key generation step, we require that 3 is not a square mod N_τ. This rules out about half of the choices for N_τ, thus reducing the security by only one bit. Moreover, as observed in Remark 3, we have some margin on N_τ which is only bounded by $p^{1/3}$, so we can even increase it by one bit and retain the same security with essentially no impact on efficiency.

In the response phase, we require that q is $(2^a, 2^b)_3$-nice. When compared with [NO24b], we are now asking that $M(q)$ is divisible by 3, instead of imposing condition (1). In this way, we avoid leaking Legendre symbols modulo N_τ. At first sight, our fix

Algorithm 3 $\texttt{GenRandIsogImg}_{I_\tau}(d,\ D;\ S)$

Input: An isogeny $\tau : E_0 \to E$ of prime degree N_τ, its corresponding ideal I_τ, relatively prime integers d, D such that $3 \mid d(D-d)$, $D \approx p$, $d > N^3$, $d < D$, and $E[D] \subseteq E(\mathbb{F}_{p^2})$, and a finite set $S \subseteq E$,
Output: $(F, \iota(S))$ for a random d-isogeny $\iota : E \to F$.

$(C_0 : D_0) \leftarrow \texttt{EichlerModConstraint}(I_\tau, 1, 1)$
Let P, Q be a basis of $E[D]$.
if d satisfies $\left(\frac{d(D-d)}{N_\tau}\right) = \left(\frac{-1}{N_\tau}\right)$ **then**
  $\alpha \leftarrow \texttt{StrongApproximation}_{d(D-d)}(N, C_0, D_0)$
  $(F; \iota(S); \emptyset) \leftarrow \texttt{KaniCod}(d, D-d, E, E, P, Q, \alpha(P), \alpha(Q); S; \emptyset)$
else
  $\alpha' \leftarrow \texttt{StrongApproximation}_{d(D-d)/3}(N, C_0, D_0)$
  $\alpha'' \leftarrow$ random 3-isogeny $E \to E''$, computed using Vélu
  $\alpha \leftarrow \alpha'' \circ \alpha'$
  $(F; \iota(S); \emptyset) \leftarrow \texttt{KaniCod}(d, D-d, E, E'', P, Q, \alpha(P), \alpha(Q); S; \emptyset)$
end if
return $(F, \iota(S))$

- seems to decrease the probability of accepting q. Indeed, the probability that a random integer satisfies (1) is about $1/2$, while the probability that a random integer is divisible by 3 is $1/3$. However, $M(q)$ has three (coprime) factors q, $2^a - q$ and $2^{a+b} - q(2^a - q)$, and it is easy to check that:
  - if $a \equiv b \bmod 2$ then $3 \mid M(q)$,
  - if $a \equiv 0 \bmod 2$ and $b \equiv 1 \bmod 2$ then $3 \nmid M(q) \Leftrightarrow q \equiv 2 \bmod 3$,
  - if $a \equiv 1 \bmod 2$ and $b \equiv 0 \bmod 2$ then $3 \nmid M(q) \Leftrightarrow q \equiv 1 \bmod 3$.

  So our condition is in fact *less* restrictive than (1): the probability of success increases from $1/2$ to 1 or $2/3$, it now seems, depending on the parities of a, b. We wrote "it now seems", because in view of Conjecture 6 below, quaternion norms are more likely to be divisible by 3 than random integers, creating a bias in our favour; as it turns out, if $a \not\equiv b \bmod 2$ then a better estimate for the probability that $3 \mid M(q)$ is $19/27 \approx 70.37\%$, rather than $2/3 \approx 66.66\%$.
- seems to require the 3-torsion to be rational. This is satisfied as soon as $3 \mid f$. However, even if $3 \nmid f$ then this is not an issue: all 3-torsion points have rational x-coordinates, and this is good enough to avoid arithmetic over a field extension.

Notice that none of our requirements has an impact on p and its size. As a consequence, the sizes of keys and signatures remain unchanged even in the modified version of the protocol.

4 More on the response algorithm

In this section, we first recall a trick from [NO24b, Section 4.3] on increasing the success rate of finding a response ideal. Then in Section 4.2, we give a careful analysis on the success rate of sampling a response ideal using this trick, which was omitted in [NO24b].

4.1 A trick in the response algorithm

The actual response algorithm from [NO24b] is different from the simplified version described in Section 1.2. The reason is that sometimes one cannot find an ideal I_σ such that $n(I_\sigma)$ is $(2^a, 2^b, N_\tau)$-nice. With our fix from Section 3 integrated in the response algorithm, now we are looking for an ideal I_σ such that $n(I_\sigma)$ is $(2^a, 2^b)_3$-nice, but the same obstacle persists. In case no response ideal is found, one needs to go back to the commitment phase. In [NO24b, Section 4.3] the authors describe two tricks to increase the chance of finding a response ideal I_σ, and their implementation only uses the first one, which we refer to as the "gcd-trick".

In what follows, we recall the "gcd-trick" from [NO24b, Section 4.3] with slight modifications to adapt it to our fix. From now on, we assume that $a-b \leq 2$, which means that at least $2/3$ of odd integers smaller than 2^a are $(2^a, 2^b)_3$-nice (see [NO24b, Remark 2] and Section 3.2). To avoid the failure in finding I_σ, the "gcd-trick" uses $q' = q/\gcd(q, f)$ instead of q where f is as in the parameter choice introduced in Section 1. This reduces the constraint of q from $q < 2^a$ to $q' < 2^a \Leftrightarrow q < \gcd(q, f) \cdot 2^a$.

Definition 5. *We say that a positive integer q is $(2^a, 2^b, f)_3$-nice when $q' = q/\gcd(q, f)$ is $(2^a, 2^b)_3$-nice.*

Let σ be a q-isogeny computed in the response phase. Assume that q is $(2^a, 2^b, f)_3$-nice and let $g = \gcd(q, f)$, $q' = q/g$, and $r = 2^a - q'$. The "gcd-trick" formally decomposes the q-isogeny σ into a g-isogeny $\sigma_g : E_A \to E_m$ and a q'-isogeny $\sigma' : E_m \to E_2$ and takes the following steps:

1. Compute $\ker \sigma_g$ by evaluating σ on $E_A[g]$.
2. Compute $\sigma_g : E_A \to E_m$ by using Vélu's formulas.
3. Obtain an r-isogeny $\omega : E_A \to E_3$ by using `AuxiliaryPath`.
4. Let $\sigma'_g = [\omega]_* \sigma_g$ and compute $\ker \sigma'_g = \omega(\ker \sigma_g)$.
5. Compute $\sigma'_g : E_3 \to E_4$ by using Vélu's formulas.
6. Evaluate σ' and ω' on $E_m[2^a]$ by using the relationships: $\sigma' \circ \sigma_g = \sigma$ and $\omega' \circ \sigma_g = \sigma'_g \circ \omega$.

The response in this case is given by (K_g, E_4, U_2, V_2) where K_g is a generator of σ_g and U_2, V_2 are computed similarly as U'_2, V'_2 in the original response algorithm

from Section 1.2, but with σ and ω replaced by σ' and ω'.

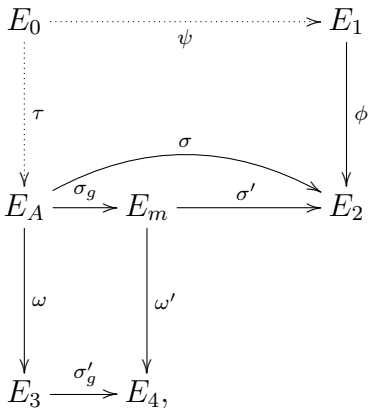

Note that there is a possible concern if $\deg \sigma_g = g$ is not coprime to $\deg \omega = r$. In that case it may happen that $\ker \sigma_g \cap \ker \omega \neq \{0\}$ so that the degree of ω' reduces to $\tilde{r} = r/h$ for some factor h of $\gcd(g, r)$. In this case, one computes an additional random h-isogeny ι from E_4 and uses $\iota \circ \omega'$ as an auxiliary path. For simplicity, [NO24b] only considers the case $h = 1$.

4.2 Sampling a response ideal

As discussed in Section 4.1, in our SQIsign2D-East variant, the response ideal $I_\sigma \sim J$ arises by repeatedly sampling an element $\alpha \in J$ such that $n(\alpha) < f2^a n(J)$ until $q = n(\alpha)/n(J)$ is $(2^a, 2^b, f)_3$-nice.
Then one can take

$$I_\sigma = J\frac{\bar{\alpha}}{n(J)},$$

as explained, e.g., in [DFKL+20, Lemma 1]. In this section we give a more precise analysis of the probability of success.

Firstly, the number of candidate-q's can be estimated using the Gaussian heuristic, which says that in any sufficiently general lattice $\Lambda \subset \mathbb{R}^4$ we expect

$$\#\{\, \alpha \in \Lambda \mid \|\alpha\| < R \,\} \approx \frac{\frac{\pi^2}{2} R^4}{\mathrm{Vol}(\Lambda)},$$

where the numerator on the right is just the volume of a ball in $\mathbb{R}^4$ with radius R. Applying this heuristic to $\Lambda = J$, which has Euclidean covolume $n(J)^2 p/4$, and to $R = \sqrt{f2^a n(J)}$, we find an expected number of

$$\frac{2\pi^2 f^2 2^{2a}}{p} \approx 2\pi^2 f 2^{a-b}$$

elements $\alpha \in J$ whose quaternion norm is smaller than $f2^a n(J)$. Assuming $\mathcal{O}_R(J)^\times = \{\pm 1\}$, from [DFKL+20, Lemma 1] it follows that this should result

in $\pi^2 f 2^{a-b}$ left ideals $I_\sigma \subset \mathcal{O}_A$ satisfying $I_\sigma \sim J$ and $n(I_\sigma) < f2^a$, and we generically expect all of these to have different norms.

However, it is not correct to think about q as being sampled uniformly at random from the interval $(0, f2^a)$. Indeed, the distribution of q is subject to the following phenomena:

(i) In view of the Gaussian heuristic, the probability that $q \in (x_1, x_2)$ is roughly proportional to $x_2^2 - x_1^2$.
(ii) For any fixed (small) modulus m, the remainder of q mod m is not distributed uniformly over the interval $\{0, 1, \ldots, m-1\}$, but rather tends to be sampled from the following distribution:

Conjecture 6. Let $J \subset \mathcal{B}_{p,\infty}$ be an integral ideal, let ℓ^e be a prime power and let $r \in \{0, 1, \ldots, \ell^e - 1\}$. Let k denote the ℓ-adic valuation of r. If $\ell \neq p$ then

$$\mathbb{P}_{r,\ell^e} := \lim_{x\to\infty} \frac{\{\, \alpha \in J \,|\, n(\alpha) \leq x \text{ and } \frac{n(\alpha)}{n(J)} \equiv r \bmod \ell^e \,\}}{\{\, \alpha \in J \,|\, n(\alpha) \leq x \,\}}$$

exists and is equal to

$$\frac{\ell^{e+1} + \ell^e - 1}{\ell^{2e+1}} \text{ if } r = 0, \qquad \frac{(\ell+1)(\ell^{k+1} - 1)}{\ell^{k+e+2}} \text{ if } r \neq 0.$$

Motivation. We believe that the distribution of norms mod ℓ^e should follow the distribution of determinants of uniformly random matrices $M \in (\mathbb{Z}/\ell^e\mathbb{Z})^{2\times 2}$. The latter distribution is well-understood, e.g., the above formulas were taken from [BM87, Corollary 2.2]. For instance, if $J = O$ is a maximal order, then the conjecture holds because $O/\ell^e O \cong (\mathbb{Z}/\ell^e\mathbb{Z})^{2\times 2}$ and under this isomorphism the norm corresponds to the determinant.

As an example, this predicts (and we experimentally observe) that about $5/8 \approx 62.5\%$ of the sampled $\alpha \in J$ are such that $q = n(\alpha)/n(J)$ is even.

For a general modulus $m = \ell_1^{e_1} \cdots \ell_s^{e_s}$ (where the ℓ_i are distinct primes) and residue class $r \in \{0, 1, \ldots, m-1\}$, we simply let

$$\mathbb{P}_{r,m} = \prod_{i=1}^{s} \mathbb{P}_{r \bmod \ell_i^{e_i}, \ell_i^{e_i}},$$

motivated by the Chinese remainder theorem. Let us now turn to the probability of sampling an ideal I_σ whose norm q is $(2^a, 2^b, f)_3$-nice. We will content ourselves with the heuristic assumption that each q is sampled independently from a distribution that meets the above phenomena (i) and (ii). Apart from being an imprecise statement, the independence is easily debunked: for any $\alpha \in J$ satisfying the stronger bound $n(\alpha) < f2^{a-2} n(J)$, we have that both $q = n(\alpha)/n(J)$ and $4q = n(2\alpha)$ are possible outcomes, and clearly these do not behave independently with respect to niceness. However, the proportion of α's that arise as

scalar multiples of each other is small,[1] and we expect the resulting bias to be negligible.

We will assume that $0 \leq a-b \leq 2$, so that $q'(2^a - q') < 2^{a+b}$ is automatically fulfilled as soon as the other conditions for q $(2^a, 2^b, f)_3$-niceness are taken care of. For these other conditions, we obtain the following probabilities:

- If we let t denote the 2-adic valuation of f, then through an application of Conjecture 6 with modulus 2^{t+1}, we arrive at an estimated probability of

$$\frac{2^{2t+3} - 3 \cdot 2^{t+1} + 1}{2^{2t+3}}$$

  that $q' = q/\gcd(q, f)$ is odd. E.g., if $t = 0$ then this just equals the probability that q is odd, which is about $\frac{3}{8} = 37.5\%$. For $t = 1$ and $t = 2$ this increases to $\frac{21}{32} \approx 65.63\%$ and $\frac{105}{128} \approx 82.03\%$, respectively.
- Let us recall from Section 3.2, now applied to q' rather than q:
  - if $a \equiv b \bmod 2$ then $3 \mid M(q')$,
  - if $a \equiv 0 \bmod 2$ and $b \equiv 1 \bmod 2$ then $3 \nmid M(q') \Leftrightarrow q' \equiv 2 \bmod 3$,
  - if $a \equiv 1 \bmod 2$ and $b \equiv 0 \bmod 2$ then $3 \nmid M(q') \Leftrightarrow q' \equiv 1 \bmod 3$.

  If we write t for the 3-adic valuation of f, then through an application of Conjecture 6 with modulus 3^{t+1}, we arrive at an estimated probability of 1 if $a \equiv b \bmod 2$ and

$$\frac{(3^{t+2}+1)(3^{t+1}-1) + 2 \cdot 3^{t+2}}{2 \cdot 3^{2t+3}} \quad \text{if } a \not\equiv b \bmod 2$$

  that $M(q')$ is divisible by 3. E.g., if $t = 0$ then this last expression equals $\frac{19}{27} \approx 70.37\%$, but for $t = 1$ and $t = 2$ this decreases to $\frac{139}{243} \approx 57.20\%$ and $\frac{1147}{2187} \approx 52.45\%$, respectively.
- In view of the Gaussian heuristic (i) and Conjecture 6, we estimate the probability that $q < \gcd(q, f)2^a$ as

$$\sum_{r=0}^{f-1} \mathbb{P}_{r,f} \cdot \frac{\gcd(r, f)^2}{f^2}. \tag{4}$$

Example 7. For $f = 35 = 5 \cdot 7$ and $a = 191$, $b = 189$ (this is the parameter set from [NO24b] for NIST level 3), these probabilities read 3/8, 1 and

$$\sum_{r=0}^{34} \mathbb{P}_{r,35} \frac{\gcd(r, 35)^2}{35^2} = \frac{2449043}{52521875} \approx 4.66\%,$$

respectively. Since $\gcd(f, 6) = 1$ it seems reasonable to assume independence of these properties. Then a given q meets all three of them with probability

[1] We refer to this mathoverflow discussion for an argument suggesting a proportion of $1/\zeta(4) = 90/\pi^4 \approx 92.39\%$ distinct vectors up to scaling; this matches very well with experiments.

$\approx 37.50\% \cdot 100.00\% \cdot 4.66\% \approx 1.75\%$. This matches near perfectly with our experiments, and it results in an overall probability of

$$(1 - 0.0175)^{\pi^2 \cdot 35 \cdot 2^2} \approx 2^{-35.17}$$

for the non-existence of a $(2^a, 2^b, f)_3$-nice instance of q. In these rare cases, the signer has to retry with a new commitment.

In general, one has to be more cautious, since if $\gcd(f, g) > 1$ then the conditions are not independent. The following formula gives a more careful count:

$$\mathbb{P}(q \text{ is nice}) \approx \sum_{r=0}^{6f-1} c_{r,2} c_{r,3} \mathbb{P}_{r,6f} \frac{\gcd(r, f)^2}{f^2}. \tag{5}$$

Here, the constants $c_{r,2}, c_{r,3} \in \{0, 1\}$ are characterized as follows:

- $c_{r,2} = 1$ if and only if $r/\gcd(r, f)$ is odd,
- $c_{r,3} = 1$ if and only if $a \equiv b \bmod 2$ or $r/\gcd(r, f) \not\equiv 1 \bmod 3$.

Note that the formula for $c_{r,3}$ excludes the wrong value of $q' \bmod 3$, potentially. However, only the 3-adic valuation of $r/\gcd(r, f)$ matters for the value of $\mathbb{P}_{r,6f}$, so this does not affect the probability estimate. A Magma script for evaluating (5) can be found in Appendix A.

In general, it is a good idea to choose $a - b = 2$, since this leads to the largest number of candidate values of q, and the condition $3 \mid M(q')$ comes for free in this case. Table 1 gives concrete parameters, along with the heuristic success probabilities predicted by (5) and the experimental values obtained by running our implementation. It also lists the probability that the response algorithm would fail to find a $(2^a, 2^b, f)_3$-nice q and would therefore need to restart with a different commitment. Table 2 shows more detailed statistics on the norms found experimentally by signing 1000 messages for each security level.

Table 1. $\mathbb{P}(q \text{ is nice})$ denotes the probability that a randomly sampled q is $(2^a, 2^b, f)_3$-nice, according to our heuristic. The column "experiment" compares this with the proportion of $(2^a, 2^b, f)_3$-nice q's according to experiment, averaged over 1000 different commitments. The last column shows the base-2 logarithm of the heuristic probability that no nice q exists.

NIST level	a	b	f	$\mathbb{P}(q \text{ is nice})$	experiment	$\log_2(\mathbb{P}(\text{all } q\text{'s fail}))$
1	129	127	45	$\frac{64859}{3645000} \approx 1.78\%$	1.75%	≈ -46.02
3	191	189	35	$\frac{7347129}{420175000} \approx 1.75\%$	1.76%	≈ -35.17
5	254	252	375	$\frac{514579}{253125000} \approx 0.20\%$	0.20%	≈ -43.46

Table 2. The values of a, b, f are the same as in Table 1. Total and Small show the number of elements with (reduced) norm below $f2^a$ and 2^a respectively; these are compared with the Gaussian heuristic in brackets. Even is the percentage of all norms below $f2^a$ being even. Nice is the number of nice elements found, and Nice ratio the ratio to the total number of elements. These quantities are compared with the estimates (in brackets) computed by the script in Appendix A.

Level	Total	Small	Even	Nice	Nice ratio
1	1778.7 (1776.5)	0.94 (0.88)	62.5% (62.5%)	31.26 (31.6)	1.75% (1.78%)
3	1380.65 (1381.7)	1.16 (1.12)	62.4% (62.5%)	24.3 (24.14)	1.76% (1.75%)
5	14804.19 (14804.40)	0.07 (0.10)	62.5% (62.5%)	29.54 (30.09)	0.20% (0.20%)

5 Security analysis

We now discuss the security of our SQIsign2D-East variant. Note that this section is entirely about the zero-knowledge property as the proof for special soundness stays the same as in [NO24b].

5.1 On the distribution of auxiliary paths

Let $\tau : E_0 \to E_A$ be an N_τ-isogeny and I_τ be the left $\mathcal{O}_0$-ideal corresponding to τ. Given the right order $\mathcal{O}_A$ of I_τ, we use $\mathcal{S}_{I_\tau,M}$ to denote the distribution on

$$\mathcal{O}_M := \{\alpha \in \mathcal{O}_0 \cap \mathcal{O}_A \mid n(\alpha) = M\}$$

output by the algorithm consisting of first getting $(C_0 : D_0) \in \mathbb{P}^1(\mathbb{Z}/N_\tau\mathbb{Z})$ by running $\texttt{EichlerModConstraint}(I_\tau, 1, 1)$ and then getting $\alpha \in \mathcal{O}_0 \cap \mathcal{O}_A$ with norm M by running $\texttt{StrongApproximation}_M(N_\tau, C_0, D_0)$.

We define a distribution $\mathcal{Q}$ on $\mathbb{Z}$, which is the distribution of the reduced norms of the response ideals I_σ. The support of this distribution is contained in the set of $(2^a, 2^b, f)_3$-nice integers. For a fixed q that is $(2^a, 2^b, f)_3$-nice, let $q' = q/\gcd(f, q)$, we define

$$\texttt{Iso}(E_A, q') := \{\, \varphi : E_A \to \star \text{ such that } \deg\varphi = 2^a - q' \,\},$$

and we consider the following distributions on $\texttt{Iso}(E_A, q')$:

$\mathcal{D}_\mathcal{U}$: The uniform distribution $\mathcal{U}_{\texttt{Iso}(E_A,q')}$.
$\mathcal{D}_1$: For q' such that $M(q')$ satisfies Eq. (1): a factor of θ_α of degree $2^a - q'$ where $\alpha \leftarrow \mathcal{S}_{I_\tau,M(q')}$ and $\theta_\alpha \in \mathrm{End}(E_A)$ is the corresponding endomorphism.
$\mathcal{D}_2$: For q' such that $M(q')$ does not satisfy Eq. (1): a factor of $\theta_\alpha \circ \theta''$ of degree $2^a - q'$ where $\alpha \leftarrow \mathcal{S}_{I_\tau,M(q')/3}$, $\theta_\alpha \in \mathrm{End}(E_A)$ is the corresponding endomorphism and θ'' is a random isogeny of degree 3 with domain E_A.

$\mathcal{D}_{\mathcal{AP}}$: $\mathcal{D}_{\mathcal{AP}} = \mathcal{D}_1$ if $M(q')$ satisfies Eq. (1), and $\mathcal{D}_{\mathcal{AP}} = \mathcal{D}_2$ otherwise. Note that this is the same distribution as output by Algorithm 2 with inputs $d = q', D_1 = 2^a$ and $D = 2^{a+b}$ when $\texttt{GenRandIsogImg}_{I_\tau}$ in line 2 is replaced with our modified version Algorithm 3.

Definition 8. *([BFD+24, Definition 23]) A fixed degree isogeny oracle (*FIDIO*) is an oracle taking as input a supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$ and an integer N, and outputs a uniformly random isogeny $\varphi : E \to E'$ (in efficient representation) with domain E and degree N.*

Problem 9. Let a, b and f be fixed integers as in the parameter choices, N_τ be the reduced norm of the secret ideal and E_A be the public curve. Let $S \subset \{\, \omega : E_A \to \star \,\}$ be a set of size larger than $\log N_\tau$ where either

1. S is sampled by first sampling $q \leftarrow \mathcal{Q}$, then sampling ω from $\mathcal{D}_{\mathcal{U}}$ on the set $\texttt{Iso}(E_A, q')$, where $q' := q/\gcd(f, q)$;
2. S is sampled by first sampling $q \leftarrow \mathcal{Q}$, then sampling ω from $\mathcal{D}_{\mathcal{AP}}$ on the set $\texttt{Iso}(E_A, q')$, where $q' := q/\gcd(f, q)$.

The problem is, given a, b, f, N_τ, E_A, S, to distinguish between the two cases with a polynomial number of queries to $\mathcal{Q}$, to FIDIO and to $\mathcal{D}_{\mathcal{AP}}$.

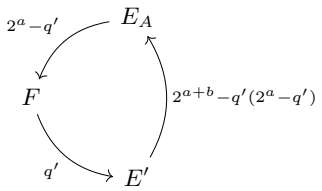

Fig. 3. A diagram that illustrates the computation of the auxiliary path from E_A in the case when Eq. (1) holds for $M(q')$.

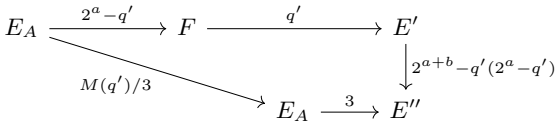

Fig. 4. A diagram that illustrates the computation of the auxiliary path from E_A in the case when Eq. (1) does not hold for $M(q')$.

Remark 10. It seems that the most natural way to distinguish the two cases in Problem 9 is to reverse engineer the algorithm that underlies the distribution

$\mathcal{D}_{\mathcal{AP}}$. That means, given an isogeny $E_A \to F$ of degree $2^a - q'$, one tries to complete the diagrams in Figs. 3 and 4. In the first case, it means to come up with an isogeny from F to E_A of degree $q'(2^{a+b} - q'(2^a - q'))$. This gives rise to an endomorphism on E_A. Then one recovers the quaternion element corresponding to this endomorphism and checks whether the quaternion element is sampled from $\mathcal{S}_{I_\tau, M(q')}$. The second case is similar, except that one finds an isogeny from F to some curve E'' that is connected to E_A by a degree 3 isogeny. This process requires at least the knowledge of both the endomorphism rings of E_A and F. Therefore, it seems reasonable to assume that solving Problem 9 is computationally hard and it requires a time complexity of at least $O(2^\lambda)$ for $p \in O(2^{2\lambda})$.

5.2 Zero-knowledge of SQIsign2D-East

We now give a proof of the zero-knowledge property of our SQIsign2D-East variant.

Definition 11. *Given integers f, a and b, a random uniform nice degree isogeny oracle (*RUNDIO*) is an oracle taking as input a supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$ and returning an elliptic curve E' together with an efficient representation of an isogeny $\sigma : E \to E'$ such that:*

(i) The distribution of E' is uniform in the supersingular isogeny graph.
(ii) If there exist isogenies from E to E' of $(2^a, 2^b, f)_3$-nice degree, then the conditional distribution of σ given E' is uniform among isogenies $E \to E'$ of $(2^a, 2^b, f)_3$-nice degree. If no such isogenies exist, then σ is the 0 isogeny from E to E'.

Problem 12. Let a, b and f be fixed integers as in the parameter choices. Let E be a supersingular elliptic curve, and F be another supersingular elliptic curve where either

1. F is sampled by first applying rejection sampling that rejects 0 isogenies to the outputs of RUNDIO with input E, then taking a random walk of degree 2^b from the curve E' where E' is the codomain of the output isogeny $\sigma : E \to E'$,
2. F is sampled uniformly at random on the supersingular isogeny graph.

The problem is, given f, a, b, E, to distinguish between the two cases with a polynomial number of queries to RUNDIO.

Remark 13. According to the analysis in Section 4.2, with our proposed parameter choices, there is only a tiny portion ($\approx 2^{-40}$) of curves on the supersingular isogeny graph that is not connected to any given random curve E with a $(2^a, 2^b, f)_3$-nice degree isogeny. In the easier case when E is taken to be E' instead of the endpoint of a random walk from E', to be able to distinguish from the two distributions, either one computes all curves that are connected to

E by a $(2^a, 2^b, f)_3$-nice degree isogeny and use this as a distinguishing method, or one recovers the endomorphism rings of E and F, then enumerating all the connecting isogenies that are of degree less than $2^a f$. Both methods take a time complexity of at least $O(2^\lambda)$. F is taken to be the codomain of a random isogeny walk from E', we expect this problem to be even harder.

Theorem 14. *Assuming that the commitment curve E_1 is computationally indistinguishable from an elliptic curve chosen uniformly at random in the supersingular isogeny graph, and the hardness of Problem 9 and Problem 12, then the SQIsign2D-East identification protocol is computationally honest-verifier zero-knowledge in the RUNDIO and FIDIO model.*

In other words, there exists a polynomial time simulator $\mathcal{S}$ with access to a RUNDIO and a FIDIO that produces random transcripts which are computationally indistinguishable from honest transcripts.

Proof. A transcript of SQIsign2D-East consists of $(E_1, \phi, K_g, E_4, U_2, V_2)$, where E_1 is a commitment, ϕ is a challenge, (K_g, E_4, U_2, V_2) can be uniquely computed from a q'-isogeny σ' and a $(2^a - q')$-isogeny ω'. (See Section 4.1 for more details.) The simulator proceeds as follows:

1. Call the RUNDIO on input E_A to get an isogeny $\tilde{\sigma} : E_A \to \tilde{E}_2$ of $(2^a, 2^b, f)_3$-nice degree $\tilde{q}$. If RUNDIO outputs the 0 isogeny, then we run RUNDIO again until $\sigma \neq 0$. According to the analysis of Section 4.2 applied to our parameter choices, this happens with a fairly small probability. Let $\tilde{\sigma}_g$ be the factor of $\tilde{\sigma}$ from E_A of degree $\tilde{g} = \gcd(\tilde{q}, f)$, and let $\tilde{K}_g$ be its kernel.
2. Generate an isogeny $\hat{\tilde{\phi}} : \tilde{E}_2 \to \tilde{E}_1$ of degree 2^b uniformly at random.
3. Let $\tilde{q}' = \tilde{q}/\gcd(f, \tilde{q})$. Call the FIDIO on input $(E_A, 2^a - \tilde{q}')$, resulting in the isogeny $\tilde{\omega} : E_A \to \tilde{E}_3$.
4. Compute $(\tilde{E}_4, \tilde{U}_2, \tilde{V}_2)$ from $(\tilde{\sigma}, \tilde{\omega})$.

Then the procedure above gives rise to a simulated transcript as $(\tilde{E}_1, \tilde{\phi}, \tilde{K}_g, \tilde{E}_4, \tilde{U}_2, \tilde{V}_2)$.

Let $(E_1, \phi, K_g, E_4, U_2, V_2)$ be a real transcript where (K_g, E_4, U_2, V_2) is computed from the response isogeny $\sigma : E_A \to E_2$ of degree q and the auxiliary path $\omega : E_A \to E_3$ of degree $2^a - q$. From the properties of the RUNDIO and FIDIO and the assumptions we made in the theorem, we can see that:

1. The codomain curve $\tilde{E}_1$ of $\hat{\tilde{\phi}}$ is computationally indistinguishable from a random curve in the supersingular isogeny graph by the hardness of Problem 12. By assumption, E_1 and $\tilde{E}_1$ are computationally indistinguishable.
2. ϕ and $\tilde{\phi}$ follow the same distribution as they are generated the same way.
3. Conditional to $\tilde{E}_2$, $\tilde{\sigma}$ is uniformly random among isogenies between E_A and $\tilde{E}_2$ of $(2^a, 2^b, f)_3$-nice degree by the definition of RUNDIO. Conditional to E_2, σ has the same distribution by construction.
4. Assuming the hardness of Problem 9, ω is computationally indistinguishable from a random isogeny of degree $2^a - \tilde{q}'$ from E_A with $\tilde{q}$ sampled from $\mathcal{Q}$ and $\tilde{q}' = \tilde{q}/\gcd(f, \tilde{q})$.

5. Item 3,4 combined shows that (K_g, E_4, U_2, V_2) is computationally indistinguishable from $(\tilde{K}_g, \tilde{E}_4, \tilde{U}_2, \tilde{V}_2)$ as the distributions of (σ, ω) and $(\tilde{\sigma}, \tilde{\omega})$ are computationally indistinguishable. □

Remark 15. The assumption on the distribution of the commitment curve E_1 made in Theorem 14 is about analyzing the distribution of the outputs of the algorithm `RandIsogImg` given the input norm size. This has been discussed in great detail in [NO24a] where this algorithm was first introduced. Based on the discussions there, we believe this assumption is reasonable.

The previous attack strategy does not apply. To run the attack as in Section 2 on SQIsign2D-East, we need to able to solve the following problem:

Problem 16. Let a be a fixed integer as in the parameter choices and E_A be the public curve. Let $\omega : E_A \to \star$ be of degree $2^a - q$ where either

1. ω is sampled from $\mathcal{D}_1$,
2. ω is sampled from $\mathcal{D}_2$.

The problem is, given E_A, ω, to distinguish between the two cases with a success rate of 1 with a polynomial number of queries to $\mathcal{D}_{\mathcal{AP}}$.

We prove in Proposition 17 that Problem 16 is not much easier than Problem 9 assuming the most efficient algorithm to solve Problem 9 has a time complexity of $O(2^{\lambda'})$ where $\lambda' \geq \lambda$. This seems a reasonable assumption as discussed in Remark 10, and a necessary condition to have our protocol achieve λ-bits security. Proposition 17 then implies that our assumption on the hardness of Problem 9 ensures the hardness of Problem 16, therefore we do not need to make an extra assumption on Problem 16. This agrees with the intuition that if Problem 16 were easy, then our SQIsign2D-East would not be zero-knowledge.

Proposition 17. *If solving Problem 9 requires a time complexity of $O(2^{\lambda'})$ with $\lambda' \geq \lambda$, then solving Problem 16 requires a time complexity of at least $O(2^{\lambda'}/\lambda)$.*

Proof. We prove by contradiction. Suppose there is an algorithm $\mathcal{A}$ that solves Problem 16 in a time complexity of $O(t)$ smaller than $O(2^{\lambda'}/\lambda)$. Now in Problem 9, we are given k samples with $k > \log N_\tau$ such that they are either from $\mathcal{D}_{\mathcal{U}}$ or $\mathcal{D}_{\mathcal{AP}}$. We run the distinguishing algorithm $\mathcal{A}$ on around $\log N_\tau \approx \lambda/2$ number of samples to get enough Legendre symbol values with respect to N_τ to uniquely determine N_τ. These values allow us to recover N_τ in time $O(2^{\lambda/2})$. Given the value of N_τ, we check whether the remaining $k - \log N_\tau$ samples give rise to correct Legendre symbols values. In the case when the k samples are from $\mathcal{D}_{\mathcal{U}}$, this fails with a non-negligible probability; and in the case when k samples are from $\mathcal{D}_{\mathcal{AP}}$, this always succeeds. This leads to an algorithm that solves Problem 9 in time less than $O(\lambda t) + O(2^{\lambda/2})$ which is less than $O(2^{\lambda'})$, a contradiction. □

Remark 18. Although the additional 3-isogeny computation will probably be very fast compared to the rest of the response step, it still introduces a conditional step that is performed only when q fails to satisfy some Legendre symbol

condition with respect to N_τ. This may leak side channel information that, if exploited, could lead to a restoration of the original attack. A careful analysis of this aspect through an optimized implementation is left as future work.

References

AG97. Michael Anshel and Dorian Goldfeld. Zeta functions, one-way functions, and pseudo-random number generators. *Duke Mathematical Journal*, 88(2):371–390, 1997.

BBUV20. Ward Beullens, Tim Beyne, Aleksei Udovenko, and Giuseppe Vitto. Cryptanalysis of the Legendre PRF and generalizations. *IACR Transactions on Symmetric Cryptology*, 2020(1):313–330, May 2020.

BFD+24. Andrea Basso, Luca De Feo, Pierrick Dartois, Antonin Leroux, Luciano Maino, Giacomo Pope, Damien Robert, and Benjamin Wesolowski. SQIsign2D-West: the Fast, the Small, and the Safer. Cryptology ePrint Archive, Paper 2024/760, 2024. https://eprint.iacr.org/2024/760.

BM87. Richard P. Brent and Brendan D. McKay. Determinants and ranks of random matrices over $\mathbb{Z}_m$. *Discrete Mathematics*, 66(1):35–49, 1987.

CD23. Wouter Castryck and Thomas Decru. An efficient key recovery attack on SIDH. In *EUROCRYPT 2023*, pages 423–447, 2023.

Dam88. Ivan Bjerre Damgård. On the randomness of Legendre and Jacobi sequences. In *Conference on the Theory and Application of Cryptography*, pages 163–172. Springer, 1988.

DFKL+20. Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski. SQISign: Compact post-quantum signatures from quaternions and isogenies. In *Asiacrypt Vol. 1*, volume 12491 of *Lecture Notes of Computer Science*, pages 64–93. Springer, 2020.

DLRW24. Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski. SQIsignHD: New dimensions in cryptography. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part I*, volume 14651 of *LNCS*, pages 3–32. Springer, Heidelberg, may 2024. Artifact available at https://artifacts.iacr.org/tches/2022/a11.

Kan97. Ernst Kani. The number of curves of genus two with elliptic differentials. *Journal für die reine und angewandte Mathematik*, 1997(485):93–122, 1997.

KV10. Markus Kirschmer and John Voight. Algorithmic enumeration of ideal classes for quaternion orders. *SIAM Journal on Computing*, 39(5):1714–1747, 2010.

Ler22. Antonin Leroux. *Quaternion algebras and isogeny-based cryptography*. Theses, Ecole doctorale de l'Institut Polytechnique de Paris, September 2022.

MMP+23. Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, and Benjamin Wesolowski. A direct key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology - EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 448–471. Springer, 2023.

NO24a. Kohei Nakagawa and Hiroshi Onuki. QFESTA: Efficient algorithms and parameters for FESTA using quaternion algebras. In *Annual International Cryptology Conference*, pages 75–106. Springer, 2024.

NO24b. Kohei Nakagawa and Hiroshi Onuki. SQIsign2D-East: A new signature scheme using 2-dimensional isogenies. Cryptology ePrint Archive, Paper 2024/771, 2024. https://eprint.iacr.org/archive/2024/771/20240520:111039.

Rob23. Damien Robert. Breaking SIDH in polynomial time. In *EUROCRYPT (5)*, volume 14008 of *Lecture Notes in Computer Science*, pages 472–503. Springer, 2023.

A Script for estimating the probability of niceness

The following Magma procedure `probnice(f, a, b)` can be used to evaluate formula (5), for any given values of f, a, b.

```
function NormDensity(r, f)
  fac := Factorization(f);
  prod := 1;
  for ell in fac do
    e := ell[2];
    modulus := r mod ell[1]^e;
    if modulus eq 0 then
      prod *:= (ell[1]^(e + 1) + ell[1]^e - 1)/ell[1]^(2*e + 1);
    else
      k := Valuation(modulus, ell[1]);
      prod *:= (ell[1] + 1)*(ell[1]^(k + 1) - 1)/ell[1]^(k + e + 2);
    end if;
  end for;
  return prod;
end function;

procedure probnice(f, a, b);
  pi := Pi(RealField());
  prob := 0;
  for r in [0..6*f - 1] do  // represents q mod 6f
    rdiv := r div GCD(f, r);
    cond2 := rdiv mod 2 ne 0;
    if a mod 2 eq b mod 2 then
      cond3 := true;
    else
      cond3 := rdiv mod 3 ne 1;
      // note: this may exclude wrong residue class, but prob does not change
    end if;
    if cond2 and cond3 then prob+:= NormDensity(r, 6*f)*GCD(r, f)^2/f^2; end if;
  end for;
  print "For f =", f, "and a =", a, "and b =", b, ":";
  print "Expected probability for q being nice is", prob, "
    which is approximately", prob*1.;
```

```
  print "log_2 of overall probability of failure is", Log(2,
    (1 - prob)^(pi^2*f*2^(a-b)));
end procedure;

print "\n NIST 1";
probnice(45, 129, 127);

print "\n NIST 3";
probnice(35, 191, 189);

print "\n NIST 5";
probnice(375, 254, 252);
```