

ANCHORAGE DIGITAL

WHITEPAPER

# On the auxiliary isogeny freedom in SQLsign's two-dimensional representation

Canonicalization does not suffice  
for strong unforgeability.

# On the Auxiliary Isogeny Freedom in SQIsign’s Two-Dimensional Representation

Dustin Ray, Anchorage Digital

July 2026

## Abstract

SQIsign encodes its response isogeny via a two-dimensional representation on a product of elliptic curves, using the Kani construction. We analyze the algebraic structure of this encoding in detail, with a focus on the role of the auxiliary isogeny and its implications for strong unforgeability.

We show that the anti-isometry  $\psi$  determining the Kani kernel is publicly computable from the torsion-point images of the component and auxiliary isogenies alone, that the automorphism orbit  $\{\psi, -\psi\}$  is the only torsion-level source of encoding non-uniqueness on a fixed product surface, and that every auxiliary of the correct degree is commitment-compatible by Kani’s theorem. The sole barrier to producing an alternative encoding is the construction of a new auxiliary isogeny of the required (generically non-smooth) degree from the challenge curve.

Following an observation of De Feo, we demonstrate that this barrier falls whenever a single small prime divides the auxiliary degree  $2^f - q$ : the adversary reroutes the terminal  $\ell$ -isogeny step at the codomain of the honest auxiliary, preserving the total degree exactly. Since generic odd integers possess small prime factors with overwhelming probability, this establishes that canonicalization of the basis change matrix does not suffice for strong unforgeability.

We analyze the signing algorithm’s dependency structure to show that the auxiliary cannot be bound into the Fiat-Shamir hash without a protocol redesign, and conclude with a structural comparison of ECDSA (where canonicalization suffices), SQIsign (where it does not), salt-PRISM (where the degree- $(2^a - q)$  factor of the response is an auxiliary in the same sense), and the prime-degree variant of PRISM (where the response degree is prime, so no auxiliary factor exists).

## 1 Introduction

Digital signature schemes are typically analyzed under the existential unforgeability (EUF-CMA) model [GMR88], which guarantees that an adversary cannot produce a valid signature on a new message. A stronger notion, *strong unforgeability* (SUF-CMA) [ADR02], additionally requires that the adversary cannot produce a new signature on a message that has already been signed, even if the adversary has access to a signing oracle. The distinction matters in settings where signatures serve as unique identifiers: transaction deduplication in distributed ledgers, replay protection, and protocols that rely on the uniqueness of (message, signature) pairs.

For many classical signature schemes, the gap between EUF-CMA and SUF-CMA is well understood. In ECDSA, the only signature malleability is the well-known  $(r, s)$  versus  $(r, n - s)$  ambiguity, arising from the single non-trivial automorphism of the cyclic group. Canonicalization (selecting the smaller of  $s$  and  $n - s$ , as standardized in BIP-340 [WNR20]) eliminates this freedom and is sufficient for SUF-CMA. The question we investigate in this note is whether the same approach works for SQIsign [DLRW24, AAA<sup>+</sup>25], the only isogeny-based signature scheme under consideration in Round 3 of NIST’s post-quantum standardization process.

SQIsign encodes its response isogeny via a two-dimensional representation on a product of elliptic curves, using the Kani construction [Kan97, BDF<sup>+</sup>24]. This encoding admits a sign ambiguity ( $M$  versus  $-M$  in the basis change matrix) that is the direct analog of ECDSA’s  $(r, s)/(r, n - s)$ , and which a canonical form eliminates in the same way (Section 3). However, the two-dimensional setting introduces a second degree of freedom that has no analog in the classical case: the *auxiliary isogeny*  $\tau$  can be replaced by any isogeny of the same degree from the challenge curve, producing a different valid encoding of the same response. The SQIsign specification notes this non-uniqueness and concludes that “SQIsign cannot achieve strong unforgeability” [AAA<sup>+</sup>26, Section 8.4.2]; the same statement appears in the earlier version 2.0 document [AAA<sup>+</sup>25, Section 10.2.7].

This note provides a detailed algebraic analysis of the auxiliary freedom, tracing the question from the structure of the Kani diamond to the number-theoretic properties of the auxiliary degree  $2^f - q$ . We show that the anti-isometry  $\psi$  encoding the Kani kernel is publicly computable from the torsion-point images of the component and auxiliary isogenies alone, without the cross isogeny (Section 4); that the automorphism orbit  $\{\psi, -\psi\}$  exhausts the torsion-level freedom on each fixed surface (Section 4.5); and that every auxiliary of the correct degree is commitment-compatible by Kani’s theorem (Proposition 4.2). We then analyze the barrier to producing an alternative auxiliary: the rigid Kani degree constraint, the generic non-smoothness of  $2^f - q$ , the fixed-degree isogeny equivalence to EndRing [BKM<sup>+</sup>24], and the Robert torsion-recovery gap [Rob23] (Section 5).

Following an observation of De Feo [DF26], we show that this barrier is circumvented whenever a single small prime  $\ell$  divides  $2^f - q$ : the honest auxiliary can be modified at its codomain end by rerouting the terminal  $\ell$ -isogeny step, preserving the total degree exactly (Section 6). Since generic odd integers possess small prime factors with overwhelming probability (Mertens’ theorem [Mer74]), canonicalization of the basis change matrix does not suffice for strong unforgeability.

We conclude by analyzing the signing algorithm’s dependency structure to show that the most natural protocol-level mitigations face specific algebraic obstacles within SQIsign’s current architecture (Section 7), and by comparing the structural properties of ECDSA, SQIsign, and the two variants of PRISM [BBC<sup>+</sup>26] that determine whether canonicalization suffices in each setting (Section 8).

**Notation.** Throughout,  $p$  is a prime with  $p \equiv 3 \pmod{4}$ ,  $E_A$  is the signer’s public key (a supersingular elliptic curve over  $\mathbb{F}_{p^2}$ ), and  $f$  denotes the torsion exponent (so the product isogeny acts on  $2^f$ -torsion). The component isogeny  $a : E_{\text{chl}} \rightarrow E_{\text{com}}$  has degree  $q$  (the response degree), and the auxiliary isogeny  $\tau : E_{\text{chl}} \rightarrow E_{\text{aux}}$  has degree  $2^f - q$ . The cross isogeny  $c : E_{\text{aux}} \rightarrow E_{\text{com}}$  has degree  $2^f - q$ . We write  $\hat{\varphi}$  for the dual of an isogeny  $\varphi$ .

## 2 Preliminaries

### 2.1 Supersingular elliptic curves and isogenies

Let  $p$  be a prime with  $p \equiv 3 \pmod{4}$ , and let  $\mathbb{F}_{p^2} = \mathbb{F}_p[i]/(i^2 + 1)$ . A supersingular elliptic curve  $E$  over  $\mathbb{F}_{p^2}$  is an elliptic curve whose endomorphism ring  $\text{End}(E)$  is a maximal order in the quaternion algebra  $B_{p,\infty}$  ramified at  $p$  and  $\infty$  [Voi21, Chapter 42]. The  $j$ -invariant  $j(E) \in \mathbb{F}_p$  classifies  $E$  up to  $\mathbb{F}_{p^2}$ -isomorphism; there are approximately  $p/12$  supersingular  $j$ -invariants [Piz90].

An *isogeny*  $\varphi : E_1 \rightarrow E_2$  is a non-constant morphism of elliptic curves preserving the identity. Every isogeny has a *degree*  $\deg(\varphi) \in \mathbb{Z}_{>0}$  and a *dual*  $\hat{\varphi} : E_2 \rightarrow E_1$  satisfying  $\hat{\varphi} \circ \varphi = [\deg(\varphi)]$  on  $E_1$ . An *endomorphism* of  $E$  is an isogeny  $E \rightarrow E$ ; the endomorphism ring  $\text{End}(E)$  is a  $\mathbb{Z}$ -module of rank 4 for supersingular curves. An endomorphism  $\alpha$  is *scalar* if  $\alpha = [m]$  for some  $m \in \mathbb{Z}$ , and *non-scalar* otherwise.

For a positive integer  $N$  with  $\gcd(N, p) = 1$ , the  $N$ -torsion subgroup  $E[N]$  is isomorphic to  $(\mathbb{Z}/N\mathbb{Z})^2$ . The *Weil pairing*  $e_N : E[N] \times E[N] \rightarrow \mu_N$  is a non-degenerate alternating bilinear pairing [Sil09, III.8].

## 2.2 Hard problems

The following computational problems underlie the security of isogeny-based cryptography.

**Definition 2.1** (Endomorphism Ring (EndRing)). *Given a supersingular curve  $E/\mathbb{F}_{p^2}$ , compute a  $\mathbb{Z}$ -basis for  $\text{End}(E)$ .*

**Definition 2.2** (One Endomorphism (OneEnd) [PW24, Problem 1.1]). *Given a supersingular curve  $E/\mathbb{F}_{p^2}$ , find any non-scalar endomorphism  $\alpha \in \text{End}(E) \setminus \mathbb{Z}$ .*

**Definition 2.3** (Fixed-Degree Isogeny [BKM<sup>+</sup>24, Section 1]). *Given a supersingular curve  $E/\mathbb{F}_{p^2}$  and a positive integer  $d$ , find an isogeny of degree  $d$  from  $E$ .*

These problems are related by the following equivalences:

**Theorem 2.4** ([PW24, Theorem 1.1]). *OneEnd and EndRing are equivalent under probabilistic polynomial-time reductions.*

**Theorem 2.5** ([BKM<sup>+</sup>24, Theorem 1.1]). *For  $d < p$ , the fixed-degree isogeny problem for degree  $d$  is equivalent to EndRing under (sub-)exponential reductions. In particular, finding an isogeny of large prime degree from a curve whose endomorphism ring is unknown is at least as hard as computing the endomorphism ring.*

**Theorem 2.6** (Isogeny recovery from torsion data [Rob23, Section 4]). *Let  $\varphi : E_1 \rightarrow E_2$  be an isogeny of degree  $d$ , and suppose the action of  $\varphi$  on  $E_1[N]$  is known for some  $N$  with  $\gcd(N, d) = 1$ . If  $N^2 > d$ , then  $\varphi$  can be recovered (as an evaluable isogeny) in polynomial time. If  $N^2 \leq d$ , the torsion data is generically insufficient.*

The best known attacks on EndRing require  $O(p^{1/2})$  classical operations (meet-in-the-middle on the supersingular isogeny graph) or  $O(p^{1/4})$  quantum operations (Grover acceleration). At NIST Level I ( $p \approx 2^{256}$ ), these correspond to  $\approx 2^{128}$  and  $\approx 2^{64}$  operations respectively.

## 2.3 The Kani construction

Let  $E_1, E_2$  be supersingular elliptic curves over  $\mathbb{F}_{p^2}$  with the product principal polarization on  $E_1 \times E_2$ . A *reducible  $N$ -gluing* of  $E_1 \times E_2$  is a  $(N, N)$ -isogeny  $\Phi : E_1 \times E_2 \rightarrow A$  whose kernel  $K \subset (E_1 \times E_2)[N]$  is a maximal isotropic subgroup with respect to the product Weil pairing, and whose codomain  $A$  is isomorphic to a product of elliptic curves [Kan97, Section 2].

A reducible gluing arises from an *isogeny diamond*: four isogenies forming a commutative square

$$\begin{array}{ccc} E_1 & \xrightarrow{a} & F_1 \\ \downarrow \tau & & \downarrow \sigma \\ E_2 & \xrightarrow{d} & F_2 \end{array} \tag{1}$$

with  $\deg(a) + \deg(\tau) = N$ ,  $\deg(d) = \deg(a)$ ,  $\deg(\sigma) = \deg(\tau)$ , and commutativity  $\sigma \circ a = d \circ \tau$  [Kan97, Theorem 2.3], [Rob22, Lemma 3.1]. In addition to the four sides of the square, the diamond determines a *cross isogeny*  $c : E_2 \rightarrow F_1$  of degree  $\deg(\tau) = N - \deg(a)$ , connecting the auxiliary curve to the commitment factor. The kernel of the associated  $(N, N)$ -isogeny is the graph of an

anti-isometry  $\psi : E_1[N] \rightarrow E_2[N]$ , defined by  $K = \{(P, \psi(P)) : P \in E_1[N]\}$ , where  $\psi$  reverses the Weil pairing:

$$e_N(P_1, P_2) \cdot e_N(\psi(P_1), \psi(P_2)) = 1 \quad \text{for all } P_1, P_2 \in E_1[N]. \quad (2)$$

The following properties of reducible gluings are central to the analysis in this note:

**Lemma 2.7** (Kani [Kan97, Theorem 2.3], Robert [Rob22, Lemma 3.1, Lemma 3.4]). *Let  $a : E_1 \rightarrow F_1$  and  $\tau : E_1 \rightarrow E_2$  be isogenies with  $\gcd(\deg(a), \deg(\tau)) = 1$  and  $\deg(a) + \deg(\tau) = N$ . Then:*

- (i) *The gluing kernel  $K$  is the graph of a unique (up to the action of  $\text{Aut}(E_1 \times E_2)$ ) anti-isometry  $\psi : E_1[N] \rightarrow E_2[N]$ .*
- (ii) *The codomain of  $\Phi$  splits as  $F_1 \times F_2$ , where  $F_1$  is isomorphic to the codomain of  $a$  and  $F_2$  is isomorphic to the codomain of a complementary isogeny  $d : E_2 \rightarrow F_2$  of degree  $\deg(a)$ .*
- (iii) *The diamond determines a cross isogeny  $c : E_2 \rightarrow F_1$  of degree  $\deg(\tau)$  [Rob22, Lemma 3.4]. The kernel condition  $\Phi(P, \psi(P)) = 0$  implies  $a(P) + c(\psi(P)) = 0$  on the  $F_1$ -factor, i.e.,  $\psi(P) = -\hat{c} \circ [\deg(c)]^{-1} \circ a(P)$  for all  $P \in E_1[N]$ .*

**Remark 2.8** (The automorphism orbit). *For generic  $j$ -invariants ( $j(E_1), j(E_2) \notin \{0, 1728\}$ ),  $\text{Aut}(E_1 \times E_2) = \{(\pm 1, \pm 1)\}$ . The element  $(1, -1)$  maps  $\psi$  to  $-\psi$ , producing an isomorphic gluing with the same codomain  $j$ -invariants [Rob22, Remark 3.2]. This is the source of the sign ambiguity  $M$  versus  $-M$  in SQIsign’s encoding (Section 3).*

## 2.4 SQIsign: the relevant details

SQIsign [DLRW24, AAA<sup>+</sup>25] is an isogeny-based signature scheme whose security reduces to the hardness of the One Endomorphism problem with hints [ABF<sup>+</sup>25]. We recall only the details relevant to the analysis of the auxiliary freedom; see [AAA<sup>+</sup>25] for the full specification and [ABF<sup>+</sup>25] for the EUF-CMA proof.

**The sigma protocol.** SQIsign is built from a three-move sigma protocol: the prover commits to a curve  $E_{\text{com}}$ , receives a challenge  $c = H(\text{pk}, j(E_{\text{com}}), \text{msg})$  [AAA<sup>+</sup>25, Algorithm 4.2, line 10], and responds with an isogeny  $\sigma : E_{\text{com}} \rightarrow E_{\text{chl}}$ , where the challenge curve  $E_{\text{chl}}$  is a public 2-power quotient of the public key  $E_A$ , reconstructed deterministically from  $(E_A, c)$  [AAA<sup>+</sup>25, Algorithm 4.9, line 9].

**The two-dimensional representation.** The response  $\sigma$  is not transmitted directly. Instead, it is embedded into a  $(2^f, 2^f)$ -isogeny on the product surface  $E_{\text{chl}} \times E_{\text{aux}}$  via Kani’s construction. The isogeny diamond has the form (1) with  $E_1 = E_{\text{chl}}$ ,  $E_2 = E_{\text{aux}}$ ,  $F_1 = E_{\text{com}}$ ,  $N = 2^f$ ,  $\deg(a) = q$  (the response degree), and  $\deg(\tau) = 2^f - q$ .

The verifier reconstructs  $E_{\text{chl}}$ , uses the basis change matrix  $M_{\text{chl}}$  and the auxiliary Montgomery coefficient  $A_{\text{aux}}$  to build the Kani kernel on  $E_{\text{chl}} \times E_{\text{aux}}$ , evaluates a theta chain of  $(2, 2)$ -isogenies [AAA<sup>+</sup>25, Algorithm 4.9, lines 25–28], and checks that the codomain’s first factor has  $j$ -invariant matching  $j(E_{\text{com}})$  [AAA<sup>+</sup>25, Algorithm 4.9, line 29].

**What is hashed and what is not.** The Fiat-Shamir hash  $H(\text{pk}, j(E_{\text{com}}), \text{msg})$  binds only the public key, the commitment  $j$ -invariant, and the message. The matrix  $M_{\text{chl}}$ , the auxiliary curve  $A_{\text{aux}}$ , the hint bytes, and the backtracking parameter  $\text{nbt}$  are *not* inputs to the hash [AAA<sup>+</sup>25, Algorithm 4.9, line 29]. They influence verification only through the theta chain evaluation that recovers  $E_{\text{com}}$ .

**Remark 2.9** (The matrix validity check). *The v2.0 verifier checks only that the entries of  $M_{\text{chl}}$  are positive integers below  $2^{e'_{\text{rsp}} + r_{\text{rsp}} + 2}$  [AAA<sup>+</sup>25, Algorithm 4.9, line 5]. There is no invertibility, singularity, rank, or anti-isometry check. Correctness is enforced implicitly by the isotropy check on the 4-torsion above the kernel (line 24) and the success of the theta chain computation (lines 25–28).*

**Remark 2.10** ( $E_{\text{chl}}$  and  $E_A$ ). *The challenge curve  $E_{\text{chl}}$  is connected to the public key  $E_A$  by the public challenge isogeny  $\phi_{\text{chal}} : E_A \rightarrow E_{\text{chl}}$  of degree  $2^{f-\text{nbt}}$ , deterministically reconstructed from  $(E_A, c, \text{nbt})$ . Both  $\phi_{\text{chal}}$  and its dual are publicly computable, so  $\text{End}(E_{\text{chl}})$  and  $\text{End}(E_A)$  are polynomial-time inter-transportable by conjugation [EHL<sup>+</sup>18]. In particular, the hardness of  $\text{EndRing}(E_{\text{chl}})$  is equivalent to the hardness of  $\text{EndRing}(E_A)$  under the Supersingular Endomorphism Ring Problem (SSERP) [Wes22, ABF<sup>+</sup>25].*

### 3 A Concrete Malleability and Its Mitigation

The automorphism  $(1, -1) \in \text{Aut}(E_{\text{chl}} \times E_{\text{aux}})$  sends the anti-isometry  $\psi$  to  $-\psi$  (Remark 2.8), which negates the basis change matrix:  $M \mapsto -M \bmod N$ . Because the codomain of the product isogeny is unchanged up to isomorphism, the  $j$ -invariant of  $E_{\text{com}}$  is preserved, the Fiat-Shamir hash is unaffected, and the modified signature is valid. This is the direct analog of ECDSA’s  $(r, s)$  versus  $(r, n - s)$  malleability.

The mitigation is likewise analogous to BIP-340’s low- $s$  normalization [WNR20]: define a lexicographic ordering on  $\{M, -M \bmod N\}$  (comparing entries in the order  $m_{00}, m_{01}, m_{10}, m_{11}$ ) and require the signer to emit the smaller representative. The verifier rejects non-canonical matrices. This adds no bytes to the signature and requires one comparison with a conditional negation. We confirmed that the negation produces a valid second signature against the SQIsign C reference implementation at all three NIST security levels (I, III, V), and that the canonical form eliminates it.

Canonicalization resolves the sign ambiguity. However, the Kani construction admits a second, deeper source of encoding non-uniqueness: the *auxiliary freedom*, in which the auxiliary isogeny  $\tau$  is replaced entirely. This freedom is not contained in the Aut-orbit and is not addressed by any canonical form on  $M$ . The remainder of this note analyzes its algebraic structure and its consequences for strong unforgeability.

## 4 The Kani Diamond and Auxiliary Freedom

We now specialize the Kani construction of Section 2 to SQIsign’s setting and analyze the degrees of freedom available in the response encoding.

### 4.1 The SQIsign diamond

In SQIsign’s verification, the  $(2^f, 2^f)$ -isogeny  $\Phi$  acts on the product surface  $E_{\text{chl}} \times E_{\text{aux}}$  (Section 2.4). The underlying isogeny diamond (1) specializes as shown in Figure 1: the component isogeny  $a : E_{\text{chl}} \rightarrow E_{\text{com}}$  has degree  $q$  (the response degree), the auxiliary isogeny  $\tau : E_{\text{chl}} \rightarrow E_{\text{aux}}$  has degree  $2^f - q$ , and the cross isogeny  $c : E_{\text{aux}} \rightarrow E_{\text{com}}$  (drawn as the dashed diagonal) has degree  $2^f - q$ . The coprimality condition  $\gcd(q, 2^f - q) = 1$  holds since  $q$  is odd and  $2^f$  is a power of 2.

By Lemma 2.7(i), the Kani kernel  $K = \{(P, \psi(P)) : P \in E_{\text{chl}}[2^f]\}$  is the graph of a unique (up to Aut) anti-isometry  $\psi : E_{\text{chl}}[2^f] \rightarrow E_{\text{aux}}[2^f]$ . The basis change matrix  $M_{\text{chl}}$  in the signature encodes  $\psi$  against the torsion bases reconstructed by the verifier.




Figure 1: The SQIsign diamond. The four sides form a commutative square; the cross isogeny  $c : E_{\text{aux}} \rightarrow E_{\text{com}}$  (dashed) connects the auxiliary curve to the commitment factor. The anti-isometry  $\psi$  is determined by  $(a, \tau)$  via Kani’s theorem;  $c$  is extractable from  $\psi$  via identity (3) but is not needed as an input to the computation.

## 4.2 The algebraic identity

From Lemma 2.7(iii), the kernel condition on the  $E_{\text{com}}$ -factor gives the anti-isometry  $\psi$  as

$$\psi(P) = -\hat{c} \circ [(2^f - q)]^{-1} \circ a(P) \quad \text{for all } P \in E_{\text{chl}}[2^f], \quad (3)$$

where  $[(2^f - q)]^{-1}$  denotes the modular inverse of the scalar  $2^f - q$  in  $\mathbb{Z}/2^f\mathbb{Z}$  (which exists since  $2^f - q$  is odd, hence coprime to  $2^f$ ).

Identity (3) expresses  $\psi$  as an algebraic function of the component  $a$  and the cross isogeny  $c$ . As a consistency check, we evaluated both sides of (3) on 540 torsion points across two primes ( $p = 191$  and  $p = 1279$ ), four torsion powers ( $f \in \{4, 5, 6, 8\}$ ), and six response degrees ( $q \in \{3, 5, 7, 9, 15, 27\}$ ). In every instance, the two sides agreed.

**Remark 4.1** (Formula versus computation path). *Identity (3) describes a formula relating  $\psi$ ,  $a$ , and  $c$ . It does not imply that  $c$  is a necessary input for computing  $\psi$ : as we show in Section 4.4,  $\psi$  is recoverable from the torsion-point images of  $a$  and  $\tau$  alone, without  $c$ . The cross isogeny  $c$  is extractable from  $\psi$  after the fact (by rearranging (3)), but does not enter the forward computation. This distinction is central to the analysis of Sections 5 and 6.*

## 4.3 Commitment compatibility of arbitrary auxiliaries

The following observation is a direct consequence of Lemma 2.7(ii):

**Proposition 4.2** (Universal commitment compatibility). *Let  $a : E_{\text{chl}} \rightarrow E_{\text{com}}$  be a fixed component isogeny of degree  $q$ , and let  $\tau' : E_{\text{chl}} \rightarrow E'_{\text{aux}}$  be any isogeny of degree  $2^f - q$  (to any supersingular curve  $E'_{\text{aux}}$ ). Then the Kani construction applied to  $(a, \tau')$  produces a  $(2^f, 2^f)$ -isogeny whose codomain splits as  $E_{\text{com}} \times E''$  for some elliptic curve  $E''$  (when  $\tau' = \tau$ , this recovers the honest second factor  $E'$  of Figure 1). That is, the commitment factor  $E_{\text{com}}$  is pinned by  $a$  and is independent of the choice of  $\tau'$ .*

*Proof.* Since  $\gcd(q, 2^f - q) = 1$ , Lemma 2.7(ii) applies: the codomain splits with first factor isomorphic to the codomain of  $a$ , which is  $E_{\text{com}}$ .  $\square$

Proposition 4.2 implies that a valid second signature *exists* for every choice of auxiliary: any  $\tau'$  of degree  $2^f - q$  from  $E_{\text{chl}}$  produces a diamond whose codomain has  $E_{\text{com}}$  as a factor, so the

Fiat-Shamir hash check  $H(\text{pk}, j(E_{\text{com}}), \text{msg}) = c$  passes. The representation is genuinely non-unique at the mathematical level, as the specification notes [AAA<sup>+</sup>26, Section 8.4.2].

We checked this prediction computationally at three prime sizes (19-bit, 139-bit, and 254-bit) using the theta-isogeny library of Dartois–Maino–Pope–Robert [DMPR24]. In each case, varying the auxiliary isogeny while keeping the component  $a$  fixed produced codomains that consistently split with the same  $E_{\text{com}}$  as a factor. Control kernels built from non-diamond subgroups did not split.

#### 4.4 The anti-isometry is publicly computable

The uniqueness statement in Lemma 2.7(i) has a computational consequence that is central to the analysis:

**Proposition 4.3** (Torsion-level computability of  $\psi$ ). *Let  $a : E_{\text{chl}} \rightarrow E_{\text{com}}$  (degree  $q$ ) and  $\tau' : E_{\text{chl}} \rightarrow E'_{\text{aux}}$  (degree  $2^f - q$ ) be isogenies with  $\gcd(q, 2^f - q) = 1$ . Given the torsion-point images  $a(P), a(Q), \tau'(P), \tau'(Q)$  for a basis  $\{P, Q\}$  of  $E_{\text{chl}}[2^f]$ , the anti-isometry  $\psi' : E_{\text{chl}}[2^f] \rightarrow E'_{\text{aux}}[2^f]$  is computable in polynomial time. In particular, no knowledge of the cross isogeny  $c'$  or of any endomorphism ring is required.*

*Proof.* By Lemma 2.7(i),  $\psi'$  is the unique anti-isometry (up to  $\text{Aut}$ ) whose graph is the kernel of the  $(2^f, 2^f)$  product isogeny with split codomain containing  $E_{\text{com}}$ . The verifier’s theta-chain evaluation [AAA<sup>+</sup>25, Algorithm 4.9, lines 25–28] is precisely this computation: given the torsion-point images of  $a$  and  $\tau'$ , it constructs the Kani kernel, evaluates the  $(2, 2)$ -chain, and recovers the codomain. The anti-isometry  $\psi'$  is encoded by the matrix  $M'$  that the verifier would use. Since the verifier performs this computation from public data (the torsion images, the curve equations, and the theta constants),  $\psi'$  is publicly computable.  $\square$

**Remark 4.4.** *Proposition 4.3 clarifies the adversary’s situation: the encoding  $M'$  is not the bottleneck. Once the adversary holds  $\tau'(P)$  and  $\tau'(Q)$ , the matrix  $M'$  is obtainable by running the same theta-chain computation the verifier performs. The sole remaining question is whether the adversary can obtain the torsion-point images of  $\tau'$  in the first place. This is the subject of Sections 5 and 6.*

#### 4.5 The automorphism orbit exhausts the torsion-level freedom

On a fixed product surface  $E_{\text{chl}} \times E_{\text{aux}}$  (i.e., with the auxiliary curve held constant), we ask: how many distinct valid encodings  $M$  exist?

By Remark 2.8, for generic  $j$ -invariants ( $j(E_{\text{chl}}), j(E_{\text{aux}}) \notin \{0, 1728\}$ ), the automorphism group  $\text{Aut}(E_{\text{chl}} \times E_{\text{aux}}) = \{(\pm 1, \pm 1)\}$  acts on the anti-isometry  $\psi$  via  $(1, -1) \cdot \psi = -\psi$  and  $(-1, 1) \cdot \psi = -\psi$ . The orbit is  $\{\psi, -\psi\}$ , which corresponds to the matrices  $\{M, -M \bmod N\}$ . Canonicalization (Section 3) selects one representative.

One might ask whether other torsion-level transformations can produce valid encodings on the same surface. A natural candidate is an  $\text{SL}_2(\mathbb{Z}/2^f\mathbb{Z})$  recombination of  $\tau$ ’s torsion images: given  $\tau(P)$  and  $\tau(Q)$ , form new images  $\tau'(P) = \alpha \cdot \tau(P) + \beta \cdot \tau(Q)$  and  $\tau'(Q) = \gamma \cdot \tau(P) + \delta \cdot \tau(Q)$  for  $\begin{pmatrix} \alpha & \beta \\ \gamma & \delta \end{pmatrix} \in \text{SL}_2(\mathbb{Z}/2^f\mathbb{Z})$  (determinant 1 to preserve the torsion-level degree). We enumerated all such matrices at  $f = 3$  ( $|\text{SL}_2(\mathbb{Z}/8\mathbb{Z})| = 384$  matrices) and  $f = 5$  ( $|\text{SL}_2(\mathbb{Z}/32\mathbb{Z})| = 24576$  matrices), built the corresponding Kani kernels, and ran the theta chain for each. In both cases, only four recombinations produced split codomains with  $E_{\text{com}}$  as a factor: the identity, the negation, and the

two sign-reflections, all of which lie in the  $\text{Aut}(E_{\text{chl}} \times E_{\text{aux}})$ -orbit. Every other recombination failed the splitting check.

This confirms that on a fixed surface, the encoding is unique after canonicalization. The auxiliary freedom (replacing  $\tau$  with a genuinely different isogeny to a different curve, analyzed in the following sections) is therefore the *only* source of non-trivial malleability beyond the sign ambiguity.

## 5 The Barrier: Constructing the Auxiliary

By Proposition 4.3, the encoding  $M'$  is publicly computable once the torsion-point images of  $\tau'$  are known. The remaining question is whether these torsion-point images can be obtained without the endomorphism ring of  $E_{\text{chl}}$ . This section analyzes the obstacles to doing so; Section 6 shows that one of these obstacles admits a circumvention.

### 5.1 The rigid degree constraint

The Kani degree relation  $\deg(a) + \deg(\tau') = 2^f$  [Kan97, Theorem 2.3] is an algebraic identity, not an approximation. Since  $\deg(a) = q$  is fixed by the honest signature, the auxiliary  $\tau'$  must have degree *exactly*  $2^f - q$ . An isogeny of any other degree produces a product isogeny of the wrong order; the theta chain does not terminate at the correct step, and verification rejects.

### 5.2 Generic non-smoothness of $2^f - q$

Since  $q$  is odd (it is the norm of a quaternion algebra element, determined by the ideal computation [AAA<sup>+</sup>25, Section 4.4.3]) and  $2^f$  is a power of 2, the auxiliary degree  $2^f - q$  is odd. For  $q$  near  $\sqrt{p}$  (as in SQIsign, where the response degree is approximately  $e_{\text{rsp}} \approx \lceil \log_2 \sqrt{p} \rceil$  bits [AAA<sup>+</sup>25, Table 2]), the value  $2^f - q$  is an odd integer of cryptographic size: approximately  $f$  bits at Level I ( $f \approx 248$ ), Level III ( $f \approx 376$ ), and Level V ( $f \approx 500$ ).

By standard estimates on smooth numbers (Dickman’s  $\rho$ -function [CEP83]), the probability that a random  $n$ -bit integer has all prime factors below  $B$  is  $\rho(n/\log_2 B)$ , where  $\rho(u) = u^{-u(1+o(1))}$  for large  $u$ . For  $n = 248$  (Level I) and any polynomial bound  $B = \text{poly}(\lambda)$ , the ratio  $u = n/\log_2 B$  grows with the security parameter, and  $\rho(u)$  is negligible. In particular,  $2^f - q$  is overwhelmingly unlikely to be expressible as a product of small primes.

A priori, this suggests that constructing  $\tau'$  as a chain of small-degree isogenies is infeasible. We revisit this conclusion in Section 6.

### 5.3 The fixed-degree isogeny problem

Constructing an isogeny of a specified degree from a curve whose endomorphism ring is unknown is the *fixed-degree isogeny problem* (Definition 2.3). For degree  $d < p$ , this problem is equivalent to EndRing under the reductions of [BKM<sup>+</sup>24] (Theorem 2.5).

The SQIsign specification acknowledges this barrier explicitly: “we use isogenies in dimension 2 as a tool to efficiently compute isogenies between elliptic curves of non-smooth degree” [AAA<sup>+</sup>25, Section 2.4]. The signer overcomes it by using  $\text{End}(E_A)$  (transported to  $\text{End}(E_{\text{chl}})$  via the Deuring correspondence and the public challenge chain, Remark 2.10).

### 5.4 What the honest signature reveals

An honest signature reveals the  $2^f$ -torsion action of both  $a$  and  $\tau$ : the matrix  $M_{\text{chl}}$  encodes the anti-isometry  $\psi$ , from which the torsion-point images of  $a$  and  $\tau$  are recoverable via the diamond

relations and the reconstructed torsion bases on  $E_{\text{chl}}$  and  $E_{\text{aux}}$ . Via identity (3), the full action of  $\hat{c}$  on  $E_{\text{com}}[2^f]$  is also recoverable.

By composition, one can reconstruct the  $2^f$ -torsion action of the endomorphism  $\theta = \tau \circ \hat{a} \circ c : E_{\text{aux}} \rightarrow E_{\text{aux}}$ , which has degree  $q \cdot (2^f - q)^2$ . However, recovering  $\theta$  as an *evaluable* isogeny from its torsion action requires the Robert recovery condition (Theorem 2.6):  $N^2 > \deg(\theta)$ . Here  $N = 2^f$  and  $\deg(\theta) = q(2^f - q)^2$ . Since  $q$  is near  $\sqrt{p}$  (approximately  $f/2$  bits) and  $2^f - q \approx 2^f$ , we have  $\deg(\theta) \approx q \cdot (2^f)^2 = q \cdot N^2$ , which *exceeds*  $N^2$  by a factor of approximately  $q$ . The recovery condition  $N^2 > \deg(\theta)$  therefore fails by a factor of about  $q \approx 2^{f/2}$ . The torsion data gives a  $2 \times 2$  matrix representation of  $\theta$  modulo  $2^f$ , but this matrix cannot be lifted to an evaluable endomorphism of  $E_{\text{aux}}$ .

The leaked torsion data therefore does not provide a path to  $\text{End}(E_{\text{aux}})$  or  $\text{End}(E_{\text{chl}})$ , and does not directly help construct a new auxiliary  $\tau'$  of the required degree. The analysis of Section 6 shows, however, that a more elementary approach succeeds.

## 6 The Codomain Reroute

The preceding section established that the full auxiliary isogeny  $\tau'$  of degree  $2^f - q$  cannot be constructed from scratch without  $\text{End}(E_{\text{chl}})$ , and that the honest signature's torsion leakage does not bridge this gap. We now show, following an observation of De Feo [DF26], that the honest auxiliary can be *modified* at its codomain end to produce a new valid auxiliary, provided that any small prime divides  $2^f - q$ .

### 6.1 The observation

**Proposition 6.1** (De Feo). *Let  $\tau : E_{\text{chl}} \rightarrow E_{\text{aux}}$  be the honest auxiliary isogeny of degree  $2^f - q$ , and suppose  $\ell$  is a prime with  $\ell \mid (2^f - q)$ . Then a second valid signature on the same message, with a different auxiliary curve  $E''_{\text{aux}}$ , can be computed from the honest signature in  $O(\ell)$  isogeny evaluations.*

The construction proceeds in four steps.

### 6.2 The construction

**Step 1: Backtrack the terminal  $\ell$ -step.** Since  $\ell \mid \deg(\tau)$ , the isogeny  $\tau$  factors through an  $\ell$ -isogeny at its codomain end. Concretely, there exists an  $\ell$ -isogeny  $\rho : E_{\text{aux}} \rightarrow E_{\text{mid}}$  such that  $\tau = \hat{\rho} \circ \tau_{\text{rest}}$ , where  $\tau_{\text{rest}} : E_{\text{chl}} \rightarrow E_{\text{mid}}$  has degree  $(2^f - q)/\ell$ . The curve  $E_{\text{aux}}$  is public (it is  $A_{\text{aux}}$  in the signature), and the  $\ell + 1$   $\ell$ -isogenies from  $E_{\text{aux}}$  are enumerable [DLRW24, Section 4.1]. However, which of these  $\ell + 1$  is the correct backtrack  $\rho$  is not directly readable from the signature (the terminal  $\ell$ -step lives on the  $\ell$ -torsion, which is coprime to  $2^f$ ). The resolution is exhaustive search: each candidate  $\rho$  is tried, and the splitting check at Step 4 filters to those that produce a valid encoding. This multiplies the cost by at most  $\ell + 1$ , preserving the  $O(\ell)$  bound.

**Step 2: Recover the intermediate torsion images.** The honest signature encodes the anti-isometry  $\psi$  (via  $M$ ), from which the  $2^f$ -torsion action of  $\hat{c}$  on  $E_{\text{com}}$  is recoverable by identity (3). Combined with the reconstructed torsion bases on  $E_{\text{chl}}$  and  $E_{\text{aux}}$ , the torsion-point images of  $\tau$  are determined by the diamond relations (Section 5.4). Applying  $\rho$  to these images gives  $\rho(\tau(P)) = [\ell] \cdot \tau_{\text{rest}}(P)$  on  $E_{\text{mid}}$  (since  $\rho \circ \hat{\rho} = [\ell]$  and  $\tau = \hat{\rho} \circ \tau_{\text{rest}}$ ). Scalar division by  $\ell$  modulo  $2^f$  recovers  $\tau_{\text{rest}}(P)$ : the inverse  $\ell^{-1} \bmod 2^f$  exists because  $\ell$  is odd and coprime to  $2^f$ .




Figure 2: The codomain reroute. The honest auxiliary  $\tau = \hat{\rho} \circ \tau_{\text{rest}}$  factors through  $E_{\text{mid}}$  via a terminal  $\ell$ -step. The shared prefix  $\tau_{\text{rest}}$  is never constructed as a full isogeny; it exists only as a scalar relationship on the  $2^f$ -torsion. The adversary backtracks from  $E_{\text{aux}}$ , recovers  $\tau_{\text{rest}}$ 's torsion images by scalar division, and re-extends by a different  $\ell$ -isogeny  $\rho'$  to  $E''_{\text{aux}}$  (dashed). The total degree  $2^f - q$  is preserved.

**Step 3: Re-extend by a different  $\ell$ -isogeny.** Choose a different  $\ell$ -isogeny  $\rho' : E_{\text{mid}} \rightarrow E''_{\text{aux}}$  (distinct from  $\hat{\rho}$ ; there are  $\ell + 1$  choices from  $E_{\text{mid}}$ , of which one is  $\hat{\rho}$ ). If  $\ell^2 \mid (2^f - q)$ , then  $\tau_{\text{rest}}$  itself ends in an  $\ell$ -step, and one additional choice among the  $\ell + 1$  would backtrack further; there are still at least  $\ell - 1$  genuinely new extensions, so the construction survives in the general case. Define the new torsion images  $\tau''(P) = \rho'(\tau_{\text{rest}}(P))$  on  $E''_{\text{aux}}$ . The composition  $\tau'' = \rho' \circ \tau_{\text{rest}}$  has degree  $\ell \cdot (2^f - q)/\ell = 2^f - q$ , matching the Kani degree constraint exactly.

**Step 4: Compute the new encoding.** From the torsion-point images  $a(P), a(Q), \tau''(P), \tau''(Q)$ , compute the anti-isometry  $\psi''$  and encode it as  $M''$  via the theta-chain evaluation (Proposition 4.3). The pair  $(M'', A''_{\text{aux}})$  constitutes a valid response encoding on the surface  $E_{\text{chl}} \times E''_{\text{aux}}$ , and the Fiat-Shamir hash check passes by Proposition 4.2 (universal commitment compatibility).

**Remark 6.2 (Cost).** *The construction requires enumerating  $O(\ell)$   $\ell$ -isogenies (Step 1 and Step 3), one scalar inverse modulo  $2^f$  (Step 2), and one theta-chain evaluation (Step 4). For  $\ell = 3$ , this amounts to four 3-isogeny evaluations and negligible arithmetic. The non-smooth cofactor  $(2^f - q)/\ell$  is never constructed or evaluated as an isogeny; it is manipulated only as a scalar on the torsion. This is why the full smoothness of  $2^f - q$  is unnecessary.*

### 6.3 The domain/codomain asymmetry

A natural question is why the same approach does not work at the *domain* end of  $\tau$  (i.e., replacing the first  $\ell$ -isogeny step out of  $E_{\text{chl}}$ ). The answer is that  $E_{\text{chl}}$  is “protected”: enumerating the  $\ell$ -isogenies from  $E_{\text{chl}}$  is public, but *selecting* the one that is the first step of  $\tau$  (and then continuing the chain from the resulting intermediate curve) requires knowledge of  $\tau$  itself as an evaluable isogeny, not merely its torsion action. In contrast,  $E_{\text{aux}}$  is “exposed”: it is public, and the backtrack step  $\rho$  requires only publicly enumerable  $\ell$ -isogenies from it.

This asymmetry reflects the distinction between the *domain* of the auxiliary (where the signer’s secret determines the isogeny path) and the *codomain* (where the final  $\ell$ -step is a publicly modifiable endpoint).

### 6.4 Computational illustration

To illustrate the construction concretely, we implemented it in SageMath using the theta-isogeny library [DMPR24]. Working at toy primes where the full theta chain is tractable, we selected four

instances in which a small prime  $\ell$  (drawn from  $\{3, 5, 7\}$ ) divided  $2^f - q$ . In each case, we carried out the four-step procedure above: backtracking the terminal  $\ell$ -step from  $E_{\text{aux}}$ , recovering the intermediate torsion images by scalar division, re-extending by a distinct  $\ell$ -isogeny, and computing the new encoding via the theta chain. The resulting  $(M'', A''_{\text{aux}})$  pairs were then fed to the full  $(2^f, 2^f)$ -isogeny verification, and in all four cases the codomain's first factor matched  $j(E_{\text{com}})$ .

## 6.5 Probability of applicability

By Mertens' third theorem [Mer74], the probability that a random odd integer  $n$  has no prime factor below  $B$  is

$$\prod_{\substack{p \leq B \\ p \text{ odd}}} \left(1 - \frac{1}{p}\right) \sim \frac{2e^{-\gamma}}{\ln B}$$

where  $\gamma \approx 0.5772$  is the Euler–Mascheroni constant. For  $B = 100$ , this is approximately 24%: about 76% of odd integers have at least one prime factor below 100. For  $B = 1000$ , the figure drops to approximately 16%.

Since the response degree  $q$  is determined by the honest signing algorithm (it is the reduced norm of a quaternion algebra element, not chosen by the signer [AAA<sup>+</sup>25, Section 4.4.3]), the auxiliary degree  $2^f - q$  is effectively a random odd integer of cryptographic size from the perspective of this divisibility property. The vast majority of honestly generated signatures have an auxiliary degree divisible by at least one small prime, and are therefore subject to the codomain reroute of Proposition 6.1.

Strong unforgeability (SUF-CMA) requires that *no* honestly generated signature admit a second valid encoding. The existence of a non-negligible fraction of signatures for which a second encoding is efficiently computable is incompatible with this requirement.

**Scope of the reroute.** Nothing in the construction of Proposition 6.1 uses the quaternion structure of SQIsign's signing algorithm. It applies to any signature whose response is encoded through a Kani embedding on a product surface, whenever the degree of the non-response factor has a small prime divisor: the terminal  $\ell$ -step of that factor is backtracked and re-extended, the total degree is unchanged, and the Kani kernel is recomputed from the new torsion images. The same freedom is present in PRISM's EUF-CMA construction, where the response of prime degree  $q$  is completed to degree  $2^a$  by a factor of degree  $2^a - q$ ; the PRISM paper identifies that factor as an auxiliary isogeny [BBC<sup>+</sup>26, Remark 1] and states, in the opening of its Section 5, that when a small prime  $\ell$  divides  $2^a - q$  the terminal  $\ell$ -isogeny can be replaced to obtain a second valid signature [BBC<sup>+</sup>26, Section 5]. Section 8.3 returns to this.

## 7 On the Difficulty of a Protocol-Level Fix

The codomain reroute of Section 6 exploits the fact that the auxiliary curve  $A_{\text{aux}}$  is not bound into the Fiat-Shamir hash. A natural question is whether a protocol modification can close this gap. We do not claim that no fix is possible, but we identify specific algebraic obstacles facing the most natural approaches.

### 7.1 Hash-binding the auxiliary

The most direct mitigation would be to include  $A_{\text{aux}}$  in the hash:  $c = H(\text{pk}, j(E_{\text{com}}), A_{\text{aux}}, \text{msg})$ . Any change to  $A_{\text{aux}}$  would then change the challenge  $c$ , which changes the challenge curve  $E_{\text{chl}}$ ,




Figure 3: The signing dependency chain (Algorithm 4.2, [AAA<sup>+</sup>25]). The auxiliary curve  $E_{\text{aux}}$  is four ideal-computation steps downstream of the Fiat-Shamir hash. Including  $A_{\text{aux}}$  in the hash would create a circular dependency.

invalidating the entire diamond.

## 7.2 The circular dependency

This approach fails because  $A_{\text{aux}}$  is determined *downstream* of the hash in SQIsign’s signing algorithm. Tracing the dependency chain through Algorithm 4.2 of the v2.0 specification [AAA<sup>+</sup>25, Section 4.4]:

$$E_{\text{com}} \xrightarrow{\text{L10}} c \xrightarrow{\text{L13}} I_{\text{chl}} \xrightarrow{\text{L14}} \alpha_{\text{rsp}} \xrightarrow{\text{L18-20}} (q, e'_{\text{rsp}}) \xrightarrow{\text{L23}} I_{\text{aux}} \xrightarrow{\text{L24-27}} E_{\text{aux}}$$

The auxiliary curve  $E_{\text{aux}}$  is four ideal-computation steps downstream of the challenge  $c = H(\text{pk}, j(E_{\text{com}}), \text{msg})$ , which is computed at line 10. Even the *degree* of the auxiliary isogeny ( $2^f - q$ , where  $q = \text{nrd}(\alpha_{\text{rsp}})$ ) is unknown until  $\alpha_{\text{rsp}}$  is sampled from the challenge-dependent lattice  $I_{\text{com}} \cap I_{\text{sk}} \cdot I_{\text{chl}}$ . Including  $A_{\text{aux}}$  in the hash would require it to be known before the challenge, which is the opposite of the actual data flow.

## 7.3 Pre-commitment as a hardness problem

Could the signer determine  $A_{\text{aux}}$  during the commitment phase, before the challenge? This would require the ideal-to-isogeny computation to target a specific auxiliary norm: given the commitment  $E_{\text{com}}$  and a pre-determined  $E_{\text{aux}}$ , find a response whose auxiliary isogeny lands on that target.

This is the *prescribed-non-smooth-degree isogeny problem*: constructing an isogeny of a specific (non-smooth) degree to a specific target curve. It is at least as hard as the fixed-degree isogeny problem (Theorem 2.5), which is the same hardness assumption that protects the scheme’s EUF-CMA security. The signer would be working against the very obstacle that prevents third parties from producing alternative auxiliaries.

## 7.4 A weaker heuristic

A less ambitious approach: the verifier could reject signatures where any small prime divides  $2^f - q$ , and the signer would retry until  $2^f - q$  is  $B$ -rough (has no prime factor below  $B$ ) for some bound  $B$ . This would eliminate the codomain reroute for primes below  $B$ , at the cost of signing overhead (approximately  $B/\ln B$  retries on average, by the prime number theorem).

This heuristic is fragile. The security guarantee depends on the choice of  $B$ , and larger  $B$  increases signing time. Moreover, it does not constitute a proof of SUF-CMA: it merely raises the cost of the codomain reroute from  $O(\ell)$  to  $O(\ell')$  where  $\ell'$  is the smallest prime factor of  $2^f - q$  above  $B$ . Whether a more principled fix is possible remains open (Section 9).

# 8 Structural Comparison: ECDSA, SQIsign, and PRISM

The results of the preceding sections raise a natural question: why does canonicalization suffice for strong unforgeability in some signature schemes but not others? We compare ECDSA, SQIsign and

the two variants of PRISM to isolate the relevant structural property.

## 8.1 ECDSA

In ECDSA, a signature  $(r, s)$  satisfies  $s = k^{-1}(H(m) + r \cdot sk) \bmod n$ , where  $k$  is the signing nonce and  $r = (kG).x$ . The value  $r$  enters the hash (or, in the Schnorr variant, the commitment enters the hash directly). The only non-trivial automorphism of the cyclic group  $\mathbb{Z}/n\mathbb{Z}$  sends  $s$  to  $n - s$ , producing the well-known  $(r, s)$  versus  $(r, n - s)$  malleability [Wui14]. Canonicalization (selecting  $\min(s, n - s)$ , as in BIP-340 [WNR20]) is sufficient for SUF-CMA because:

- the nonce is committed via  $r$ , which is hashed;
- the only unbound freedom is the sign of  $s$ , which lies in the Aut-orbit;
- there is no analog of the auxiliary freedom.

## 8.2 SQIsign

In SQIsign, the signature encodes the response via a Kani diamond on  $E_{\text{chl}} \times E_{\text{aux}}$ . The matrix  $M$  and the auxiliary curve  $A_{\text{aux}}$  are *not* hashed (Remark 2.9). Two sources of non-uniqueness exist:

- the Aut-orbit  $\{M, -M\}$  (Section 3), which canonicalization eliminates;
- the auxiliary freedom (Proposition 4.2), which canonicalization does not address.

The auxiliary freedom is exploitable via the codomain reroute (Proposition 6.1) whenever a small prime divides the auxiliary degree. The signing algorithm’s dependency structure prevents binding  $A_{\text{aux}}$  into the hash (Section 7).

## 8.3 salt-PRISM and the prime-degree variant of PRISM

PRISM [BBC<sup>+</sup>25] is a hash-and-sign scheme: the challenge is a prime  $q$  derived by hashing the public key and the message, and the response is an isogeny of degree  $q$  from the public key curve. salt-PRISM [BBC<sup>+</sup>26, Section 6] adds a salt to that hash,  $q = H(\text{pk}, \text{msg}, r)$ , with  $r$  transmitted in the signature. The response is represented through a two-dimensional embedding, and the isogeny the verifier reconstructs has degree  $q(2^a - q)$ ; the verifier recomputes  $q$  and checks the degree with a Weil-pairing test. The factor of degree  $2^a - q$  is an auxiliary in the sense of this note. The PRISM paper describes it as an auxiliary isogeny [BBC<sup>+</sup>26, Remark 1], and its Section 5 opens by noting that when a small prime  $\ell$  divides  $2^a - q$ , the terminal  $\ell$ -isogeny of that factor can be replaced, giving a second signature on the same message [BBC<sup>+</sup>26, Section 5]. This is the codomain reroute of Section 6 applied to PRISM. The salt addresses the collision resistance of the hash-to-prime map and its primality test [BBC<sup>+</sup>26, Section 6], [CGMT26, Section 5.1]; it does not bear on the uniqueness of the encoding. We reproduced the reroute against an implementation of salt-PRISM at the level I round-3 prime  $p = 3 \cdot 2^{324} - 1$ : on a signature with  $3 \mid 2^a - q$ , the  $\ell = 3$  procedure of Section 6 produced 12 candidate encodings, of which exactly 3 were accepted by the verifier, the three replacements of the degree-3 factor [Ray26].

The same paper gives a second construction [BBC<sup>+</sup>26, Section 5], in which the signature isogeny has prime degree  $2^{2e-1} - s^2$  for a hash-derived integer  $s$ , and the verifier embeds it in two  $(2, 2)$ -isogeny chains. The response degree is prime, so it has no proper factorization and the reroute has no  $\ell$ -step to modify. What remains is the sign orbit and the encoding, and this is the construction in which canonicalization of the remaining sign freedom suffices in principle. The PRISM paper describes the

Table 1: Structural comparison of the encoding freedoms relevant to strong unforgeability. “Aux” is the auxiliary isogeny of SQIsign or the degree- $(2^a - q)$  factor of salt-PRISM.

|                              | ECDSA          | SQIsign                  | salt-PRISM                    | prime-degree PRISM        |
|------------------------------|----------------|--------------------------|-------------------------------|---------------------------|
| Free parameter committed?    | Yes ( $r$ )    | No ( $A_{\text{aux}}$ )  | No ( $2^a - q$ factor)        | None (prime degree)       |
| Non-uniqueness source        | $\{s, n - s\}$ | $\{M, -M\} + \text{aux}$ | $\{\pm\} + \text{aux factor}$ | $\{\pm\}$ and encoding    |
| Canonicalization sufficient? | Yes            | No                       | No                            | Yes, with strict encoding |

variant without parameters or an implementation and states that its EUF-CMA proof for the salted scheme applies to a salted version of it [BBC<sup>+</sup>26, Remark 11]; a written SUF-CMA reduction is not given there (Section 9). An instantiation with a salted hash, together with parameters, a strict canonical encoding, and measurements at the three round-3 primes, is given in [Ray26]. In that instantiation the  $\ell = 3$  procedure above, applied to signatures of the prime-degree variant, produced 0 accepted encodings out of 12 candidates.

## 8.4 The distinguishing property

The common thread is whether the encoding’s free parameters are bound into the Fiat-Shamir hash. In ECDSA, the nonce is committed via  $r = (kG).x$ , which enters the hash; the only unbound freedom is the sign of  $s$ , which canonicalization eliminates. In SQIsign, the auxiliary curve  $A_{\text{aux}}$  is not hashed, and the signing algorithm’s dependency structure prevents binding it without a protocol redesign. In the prime-degree variant of PRISM, the response degree is prime, so no auxiliary factor exists. In salt-PRISM, the response degree  $q$  is committed via the salted hash, but the composite degree  $q(2^a - q)$  leaves the factor of degree  $2^a - q$  free.

For applications requiring strong unforgeability from isogeny-based signatures, such as transaction non-malleability or signature deduplication, the prime-degree variant of PRISM [BBC<sup>+</sup>26, Section 5] removes the auxiliary factor by making the response degree prime. The canonical matrix encoding described in Section 3 remains a necessary mitigation for the sign ambiguity in any scheme built on the Kani construction, but it is not sufficient on its own.

## 9 Open Questions

We conclude with several directions for further investigation.

1. *Constraining the auxiliary degree.* Can SQIsign’s parameter generation or signing algorithm be modified to ensure that  $2^f - q$  has no small prime factors, without changing the signature format? The response degree  $q$  is the reduced norm of a quaternion algebra element sampled from a challenge-dependent lattice; whether the distribution of  $q$  can be biased to make  $2^f - q$  rough is an open question in the arithmetic of quaternion orders.
2. *Alternative Fiat-Shamir structures.* Is there a multi-round or structured Fiat-Shamir transform that avoids the circular dependency of Section 7? For instance, a two-phase commitment in which the first phase fixes the auxiliary and the second phase determines the challenge.
3. *Higher-dimensional generalizations.* SQIsignHD and related schemes use higher-dimensional product isogenies. Does the codomain reroute of Proposition 6.1 extend to these settings, or does the higher-dimensional structure introduce additional constraints that prevent it?

4. *Parameterized security.* Can the “reject small factors” heuristic of Section 7 be formalized as a parameterized security notion (e.g.,  $B$ -rough SUF-CMA), and if so, what is the tightest bound on  $B$  for which practical security holds?
5. *A written SUF-CMA reduction for the salted prime-degree variant of PRISM.* The PRISM paper argues strong unforgeability of its prime-degree variant along the lines of its EUF-CMA proof and states that the proof for the salted scheme carries over to a salted version of that variant [BBC<sup>+</sup>26, Section 5, Remark 11]. A written reduction for the salted prime-degree variant, with its loss stated, is not given there or in [Ray26].
6. *A one-chain verifier for the prime-degree variant.* The verifier of the prime-degree variant runs two  $(2, 2)$ -isogeny chains, of  $e$  and  $e - 1$  steps, and its cost is about 2.5 times that of SQIsign’s verifier at the same level [Ray26]. The first chain alone determines whether its codomain surface  $A_2$  carries an endomorphism  $\eta$  with  $\eta^2 = [\pm 2]$  whose kernel is the prescribed subgroup  $K$  of  $A_2[2]$ ; every honest signature has this property, so it is a necessary condition for acceptance. Whether it is sufficient, that is, whether no signature isogeny of the wrong degree reaches a surface with such an  $\eta$ , is open. At a toy prime ( $p = 5 \cdot 2^{14} - 1$ ,  $e = 8$ ), 21 wrong-degree candidates that passed the isotropy checks of the first chain were tested and 0 of 21 satisfied the condition [Ray26]. This is evidence about the condition’s strength, not a proof of sufficiency; a one-chain verifier of  $e$  steps would rest on such a proof.

## Acknowledgments

## References

- [AAA<sup>+</sup>25] Marius A. Aardal, Gora Adj, Diego F. Aranha, Andrea Basso, Isaac Andrés Canales Martinez, Jorge Chávez-Saab, Maria Corte-Real Santos, Pierrick Dartois, Luca De Feo, Max Duparc, Jonathan Komada Eriksen, Tako Boris Fouotsa, Décio Luiz Gazzoni Filho, Basil Hess, David Kohel, Antonin Leroux, Patrick Longa, Luciano Maino, Michael Meyer, Kohei Nakagawa, Hiroshi Onuki, Lorenz Panny, Sikhar Patranabis, Christophe Petit, Giacomo Pope, Krijn Reijnders, Damien Robert, Francisco Rodríguez-Henríquez, Sina Schaeffler, and Benjamin Wesolowski. SQIsign – algorithm specifications and supporting documentation, version 2.0. NIST PQC Digital Signatures Round 2 Submission, 2025. <https://sqisign.org/spec/sqisign-20250205.pdf>.
- [AAA<sup>+</sup>26] Marius A. Aardal, Gora Adj, Diego F. Aranha, Andrea Basso, Isaac Andrés Canales Martinez, Jorge Chávez-Saab, Maria Corte-Real Santos, Pierrick Dartois, Luca De Feo, Max Duparc, Jonathan Komada Eriksen, Tako Boris Fouotsa, Décio Luiz Gazzoni Filho, Basil Hess, David Kohel, Antonin Leroux, Patrick Longa, Luciano Maino, Michael Meyer, Kohei Nakagawa, Hiroshi Onuki, Lorenz Panny, Sikhar Patranabis, Christophe Petit, Giacomo Pope, Krijn Reijnders, Damien Robert, Francisco Rodríguez-Henríquez, Sina Schaeffler, and Benjamin Wesolowski. SQIsign – algorithm specifications and supporting documentation, round 3. NIST PQC Digital Signatures Round 3 Submission, 2026. <https://sqisign.org/spec/sqisign-20260901.pdf>.
- [ABF<sup>+</sup>25] Marius A. Aardal, Andrea Basso, Luca De Feo, Sikhar Patranabis, and Benjamin Wesolowski. A complete security proof of SQIsign. In *Advances in Cryptology – CRYPTO 2025, Part VI*, volume 16005 of *LNCS*, pages 190–222. Springer, 2025. ePrint 2025/379.

- [ADR02] Jee Hea An, Yevgeniy Dodis, and Tal Rabin. On the security of joint signature and encryption. In *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 83–107. Springer, 2002.
- [BBC<sup>+</sup>25] Andrea Basso, Giacomo Borin, Wouter Castryck, Maria Corte-Real Santos, Riccardo Invernizzi, Antonin Leroux, Luciano Maino, Frederik Vercauteren, and Benjamin Wesolowski. PRISM: Simple and compact identification and signatures from large prime degree isogenies. In *Public-Key Cryptography – PKC 2025*, volume 15676 of *LNCS*, pages 300–332. Springer, 2025. ePrint 2025/135.
- [BBC<sup>+</sup>26] Andrea Basso, Giacomo Borin, Wouter Castryck, Maria Corte-Real Santos, Riccardo Invernizzi, Antonin Leroux, Luciano Maino, Frederik Vercauteren, and Benjamin Wesolowski. PRISM with a pinch of salt: Simple, efficient and strongly unforgeable signatures from isogenies. *Cryptology ePrint Archive*, Paper 2026/443, 2026.
- [BDF<sup>+</sup>24] Andrea Basso, Pierrick Dartois, Luca De Feo, Antonin Leroux, Luciano Maino, Giacomo Pope, Damien Robert, and Benjamin Wesolowski. SQIsign2D-West: The fast, the small, and the safer. In *Advances in Cryptology – ASIACRYPT 2024, Part III*, volume 15486 of *LNCS*, pages 339–370. Springer, 2024. ePrint 2024/760.
- [BKM<sup>+</sup>24] Benjamin Benčína, Péter Kutas, Simon-Philipp Merz, Christophe Petit, Miha Stopar, and Charlotte Weitkämper. Improved algorithms for finding fixed-degree isogenies between supersingular elliptic curves. In *Advances in Cryptology – CRYPTO 2024, Part V*, volume 14924 of *LNCS*, pages 183–217. Springer, 2024. ePrint 2023/1618.
- [CEP83] E. Rodney Canfield, Paul Erdős, and Carl Pomerance. On a problem of Oppenheim concerning “Factorisatio Numerorum”. *J. Number Theory*, 17(1):1–28, 1983.
- [CGMT26] Jelle Cottaar, Steven D. Galbraith, Luciano Maino, and Monika Trimoska. On the number of Miller–Rabin rounds for hash-to-prime signatures. *Cryptology ePrint Archive*, Paper 2026/906, 2026.
- [DF26] Luca De Feo. Comment on “SUF-CMA SQIsign” (GitHub issue #18). <https://github.com/SQIsign/the-sqisign/issues/18#issuecomment-4898993873>, July 2026. Public discussion on the SQIsign repository; the observation that a single small prime factor of  $2^f - q$  suffices for forgery appears in a follow-up comment dated 7 July 2026.
- [DLRW24] Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski. SQIsignHD: New dimensions in cryptography. In *Advances in Cryptology – EUROCRYPT 2024, Part I*, volume 14651 of *LNCS*, pages 3–32. Springer, 2024. ePrint 2023/436.
- [DMPR24] Pierrick Dartois, Luciano Maino, Giacomo Pope, and Damien Robert. Theta isogeny library. <https://github.com/ThetaIsogenies/two-isogenies>, 2024.
- [EHL<sup>+</sup>18] Kirsten Eisenträger, Sean Hallgren, Kristin Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: Reductions and solutions. In *Advances in Cryptology – EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 329–368. Springer, 2018.
- [GMR88] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. Comput.*, 17(2):281–308, 1988.

- [Kan97] Ernst Kani. The number of curves of genus two with elliptic differentials. *J. Reine Angew. Math.*, 485:93–122, 1997.
- [Mer74] Franz Mertens. Über einige asymptotische Gesetze der Zahlentheorie. *J. Reine Angew. Math.*, 77:289–338, 1874.
- [Piz90] Arnold K. Pizer. Ramanujan graphs and Hecke operators. *Bull. Amer. Math. Soc. (N.S.)*, 23(1):127–137, 1990.
- [PW24] Aurel Page and Benjamin Wesolowski. The supersingular endomorphism ring and one endomorphism problems are equivalent. In *Advances in Cryptology – EUROCRYPT 2024, Part VI*, volume 14656 of *LNCS*, pages 388–417. Springer, 2024. ePrint 2023/1399.
- [Ray26] Dustin Ray. prism-rs: PRISM on the SQIsign round-3 substrate. Software and documentation, <https://github.com/anchorlabsinc/prism-rs>, 2026. Manuscript in preparation.
- [Rob22] Damien Robert. Reducible gluing of abelian varieties. Unpublished notes, 2022. <https://www.math.u-bordeaux.fr/~damienrobert/pro/publications/notes/2022-09-kani.pdf>.
- [Rob23] Damien Robert. Breaking SIDH in polynomial time. In *Advances in Cryptology – EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 472–503. Springer, 2023.
- [Sil09] Joseph H. Silverman. *The Arithmetic of Elliptic Curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, 2nd edition, 2009.
- [Voi21] John Voight. *Quaternion Algebras*, volume 288 of *Graduate Texts in Mathematics*. Springer, 2021.
- [Wes22] Benjamin Wesolowski. The supersingular isogeny path and endomorphism ring problems are equivalent. In *FOCS 2021*, pages 1100–1111. IEEE, 2022.
- [WNR20] Pieter Wuille, Jonas Nick, and Tim Ruffing. BIP-340: Schnorr signatures for secp256k1. <https://github.com/bitcoin/bips/blob/master/bip-0340.mediawiki>, 2020.
- [Wui14] Pieter Wuille. BIP-62: Dealing with malleability. <https://github.com/bitcoin/bips/blob/master/bip-0062.mediawiki>, 2014.