

Zero-Knowledge Proofs of Isogeny Diamonds

Leonardo Colò , Maher Mamah , Youcef Mokrani , Bruno Sterner 
and Nicolas Swanson 

Department of Combinatorics and Optimization, University of Waterloo, Waterloo, Canada

Abstract

Commutative diagrams of isogenies between supersingular elliptic curves, which are called isogeny diamonds, have become fundamental to isogeny-based cryptography for both constructive and cryptanalytic purposes. In parallel, proofs of knowledge of isogenies have been widely studied and have found many applications. In this work, we combine these two directions and introduce zero-knowledge proofs of isogeny diamonds, namely, we prove knowledge of isogenies that form a commutative diagram between four curves.

We present four constructions that work in various settings. The first, Windmill-ZKP assumes that the prover knows only two parallel isogenies in the diamond. The second, Cube-ZKP assumes the prover has knowledge of four of specified degree isogenies. Finally, Kube-ZKP and Kani-ZKP prove knowledge of isogeny diamonds whose degree sum is smooth. We also provide proof-of-concept implementations of the proposed constructions and compare their performance. Our results demonstrate the trade-offs between security, efficiency and compactness in these constructions.

Keywords: Zero-knowledge proofs, isogeny-based cryptography, isogeny diamonds, Kani’s Lemma

1 Introduction

Kani’s Lemma has profoundly impacted isogeny-based cryptography. Once considered a post-quantum staple, SIDH [DFJP14], a key exchange mimicking Diffie-Hellman using supersingular elliptic curve isogenies, was broken by a series of attacks [CD23, MMP⁺23, Rob23]. These attacks rely on Kani’s Lemma, where recovering the secret isogeny reduces to evaluating a 2-, 4-, or 8-dimensional isogeny between products of elliptic curves. Proposed countermeasures [FMP23, BF23] avoid this by masking the public key’s torsion points.

These attacks have also motivated new constructions that use Kani’s Lemma constructively. These include, but are not limited to, PKE schemes [BMP23, NO24, BM25, MS26], signature schemes [DLRW24, BDF⁺24, NOC⁺24, BBC⁺25] and effective group actions [CLP24, DEF⁺25]. The practicality of these constructive applications comes down to efficiently evaluating higher-dimensional isogenies. As a result, evaluating such isogenies has been an active area of research ever since the attacks on SIDH appeared. The current state of affairs in this regard is a solid understanding of how to evaluate chains of (2, 2)-isogenies [DMPR24] as well as a growing understanding on evaluating (3, 3)-chains [CRSCS25, DK26] and (2, 2, 2, 2)-chains [Dar25].

The authors acknowledge support of the NSERC Alliance Consortia Quantum Grant “Accelerating the transition to quantum-resistant cryptography” (ALLRP 578463–2022), and the NSERC Discovery Grant program. The authors also want to thank the anonymous reviewers for their valuable comments.

Emails: lcolo@uwaterloo.ca (Leonardo Colò), mmamah@uwaterloo.ca (Maher Mamah), youcef.mokrani@uwaterloo.ca (Youcef Mokrani), bsterner@uwaterloo.ca (Bruno Sterner), nicswanson@uwaterloo.ca (Nicolas Swanson)






Figure 1: The statement of our zero-knowledge proofs consists of the existence of the four isogenies above, the commutativity of the diagram, and the ability to evaluate any one of these isogenies on arbitrary points in its domain.

In parallel to all of this, isogeny-based zero-knowledge proofs have been actively explored in the literature (see [BFGP23] for a survey). Proving knowledge of isogenies is central to many constructive applications including (but not limited to) signature schemes [GPS20, DFKL⁺20], OPRFs [Bas24] and generating so-called hard curves [BCC⁺23]. There have been two general approaches to proving knowledge of fixed, publicly known degree isogeny between supersingular curves. The first approach is based on committing to an isogeny diamond, i.e. a commutative diagram of the form in Figure 1 (which we formally define in Definition 6), and revealing (at least) one of the edges of the diagram based on a received challenge. Since its original presentation [DFJP14], a number of fixes and modifications (e.g. including some sort of torsion point information) have been proposed [DFDGZ22, Bas24, MJ23]. A second approach, which has gained traction in recent years, adopts R1CS (rank 1 constraint system) encodings of modular polynomials or isogeny formulas to obtain zero-knowledge proofs [CLL23, dHKM⁺25, LP25, dHMSS26]. These proofs are particularly suitable for SNARK-friendly applications.

Contributions.

In this work, we investigate a new type of zero-knowledge proof for isogeny-based cryptography. Instead of proving knowledge of a single isogeny, we present zero-knowledge proofs of isogeny diamonds. In other words, one wants to not only prove knowledge of the isogenies in a diamond, but also that the vertical (and horizontal) isogenies are related by pushforward along the horizontal (and vertical respectively) isogenies. We present two settings that arise from this type of proof:

The first setting is the case where one knows only the two vertical (or horizontal) isogenies in the diagram. This setting is primarily motivated by SIDH countermeasures, so we impose the presence of additional torsion-point information and require the isogenies to have smooth degrees. The zero-knowledge proof must therefore also prove knowledge of this additional information. In this regard, we present **Windmill-ZKP** which adapts recent zero-knowledge proofs of isogenies with masked torsion points [MJ23, Bas24] in order to prove knowledge of the additional edge in the diagram. Such a proof is required when adapting the multiparty SIDH protocol of [FKT18] to the M-SIDH setting. This is needed because, without such a proof, a malicious party could provide masked torsion data that do not arise from an isogeny compatible with the shared isogeny square, and, therefore, the scheme would be vulnerable to adaptive attacks.

The second setting is when one knows all four isogenies of the diamond. In this scenario we present multiple constructions:

- For the first construction, which we call **Cube-ZKP**, we assume all isogenies in the diamond have smooth degrees. The idea is that the prover constructs and commits to a cube of smooth degree isogenies. Each challenge corresponds to a specific isogeny

in the cube and, upon receiving the challenge, the prover is tasked to reveal this isogeny. The verifier checks that this isogeny is correct and that the torsion point mappings are consistent.

- The second construction is a modification of **Cube-ZKP**, which we call **Kube-ZKP**, in the setting when the isogenies in the diamond have non-smooth degrees but their sum is smooth (and ideally a power of two). In this case the prover commits to the same cube as before but the challenges are slightly different. For the back diamond, there is a single challenge and the prover reveals the associated 2-dimensional Kani isogeny.
- The last construction, which we call **Kani-Diamond-ZKP**, assumes the same setup as in **Kube-ZKP**. We repurpose the 1-dimensional zero-knowledge proofs from [DFJP14, DFDGZ22] to the 2-dimensional setting and prove knowledge of the associated Kani isogeny. Our analysis suggests that this translation is harder to analyze due to the specific structure of the Kani isogeny. As a result, we present this to diversify possible constructions.

Related work.

Recently Lai and Maino [LM26] presented a zero-knowledge proof for non-smooth isogenies whose degree is of the form $q(2^a - q)$. Through Kani's Lemma, this formulation is similar to the statements that underpin the **Kube-ZKP** and **Kani-Diamond-ZKP** constructions. However, their proof strategy differs from ours as they obtain their proof through R1CS encodings of $(2, 2)$ -isogeny formulas.

Roadmap.

We begin in Section 2 by recalling the necessary background on isogeny-based cryptography and zero-knowledge proofs. We then present **Windmill-ZKP** in Section 3, **Cube-ZKP** (and the **Kube-ZKP** modification) in Section 4, and **Kani-Diamond-ZKP** in Section 5. We conclude in Section 6 with a comparison of these constructions as well as details on their performance.

2 Preliminaries

2.1 Principally polarized abelian varieties

Definition 1 (PPAV). A *principally polarized abelian variety* (PPAV) is a pair $(\mathcal{A}, \lambda)$, where $\mathcal{A}$ is an abelian variety and $\lambda : \mathcal{A} \rightarrow \mathcal{A}^\vee$ is a principal polarization, i.e., an isomorphism to the dual.

We mainly work with elliptic curves, but Section 5 requires higher-dimensional PPAVs. We recall only the needed notions; see [Mil86, Rob24] for details.

Definition 2 (N -isogeny of PPAVs). Let $(\mathcal{A}, \lambda), (\mathcal{B}, \mu)$ be PPAVs and let $\varphi : \mathcal{A} \rightarrow \mathcal{B}$ be an isogeny. Define

$$\tilde{\varphi} := \lambda^{-1} \circ \hat{\varphi} \circ \mu : \mathcal{B} \rightarrow \mathcal{A}.$$

We call φ a polarization-preserving N -isogeny if $\tilde{\varphi} \circ \varphi = [N]$.

Throughout, we suppress polarizations from the notation and assume all N -isogenies are polarization preserving.

Definition 3 (Polarized Weil pairing). For a PPAV $\mathcal{A}/k$ and N coprime to $\text{char}(k)$, the Weil pairing is a bilinear, non-degenerate, anti-symmetric pairing

$$e_N : \mathcal{A}[N] \times \mathcal{A}[N] \rightarrow \bar{k}^*.$$

It can be computed in polynomial time in N and in the field over which $\mathcal{A}[N]$ is rational [LR15]. We only use it for elliptic curves and their products.

Definition 4 (Maximal isotropic subgroup). A subgroup $H \subset \mathcal{A}[N]$ is *isotropic* if $e_N(P, Q) = 1$ for all $P, Q \in H$, and *maximal isotropic* if it is maximal with this property. Equivalently, an isotropic subgroup of order $N^{\dim \mathcal{A}}$ is the kernel of a polarization-preserving N -isogeny.

Definition 5 (Diagonal and reducible isogenies). An N -isogeny $f : \mathcal{A} \times \mathcal{B} \rightarrow \mathcal{X}$, with $\mathcal{A}, \mathcal{B}$ g -dimensional PPAVs, is *reducible* if $\mathcal{X} \simeq \mathcal{C} \times \mathcal{D}$ for g -dimensional PPAVs $\mathcal{C}, \mathcal{D}$. It is *diagonal* if $f = f_1 \times f_2$, i.e., $f(P, Q) = (f_1(P), f_2(Q))$, for N -isogenies $f_1 : \mathcal{A} \rightarrow \mathcal{C}$ and $f_2 : \mathcal{B} \rightarrow \mathcal{D}$. A maximal isotropic subgroup $H \subseteq (\mathcal{A} \times \mathcal{B})[N]$ is reducible respectively diagonal if the corresponding quotient isogeny is.

2.2 Sigma protocols

Let $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{W}$ be a relation. A statement $X \in \mathcal{X}$ is public, while a witness $W \in \mathcal{W}$ certifies that $(X, W) \in \mathcal{R}$. A sigma protocol for a family $\{\mathcal{R}_\kappa\}_\kappa$ is a three-message public-coin protocol

$$\Pi_\Sigma = (P_1, P_2, V_1, V_2),$$

where V_2 is deterministic and the prover algorithms may share state. On input X , with the prover additionally holding W , the protocol is

$$a \leftarrow P_1(X, W), \quad e \leftarrow V_1(1^\kappa), \quad z \leftarrow P_2(X, W, e),$$

and the verifier accepts if and only if

$$V_2(X, a, e, z) = 1.$$

The tuple (a, e, z) is called a transcript and it is accepting for X if the above verification test succeeds.

Correctness. For every $(X, W) \in \mathcal{R}_\kappa$, an honest execution is accepting.

n -special soundness. There is a polynomial-time extractor which, given n accepting transcripts $(a, e_1, z_1), \dots, (a, e_n, z_n)$ for the same X and first message a , with pairwise distinct challenges e_i , outputs a witness W such that

$$(X, W) \in \mathcal{R}_\kappa,$$

except with probability at most ε .

Special soundness captures the proof-of-knowledge property: a prover that can answer sufficiently many different challenges for the same first message must encode a valid witness. If the challenge space has size $|\mathcal{E}|$, then the usual one-shot knowledge error is $\frac{n-1}{|\mathcal{E}|}$, and after t independent repetitions it becomes

$$\left(\frac{n-1}{|\mathcal{E}|} \right)^t.$$

Computational zero-knowledge. There is a polynomial-time simulator which, given a statement X , outputs transcripts for said statement. These outputs must be such that no probabilistic polynomial-time algorithm can distinguish them from those of an honest prover with non-negligible advantage.

2.3 Kani's lemma

Definition 6 (Isogeny diamond). An (a, b) -isogeny diamond from PPAVs $\mathcal{A}$ to $\mathcal{B}$ is a decomposition of an ab -isogeny

$$F : \mathcal{A} \rightarrow \mathcal{B}$$

between principally polarized abelian varieties into two different decompositions

$$F = g_2 \circ f_1 = f_2 \circ g_1,$$

where f_1, f_2 are a -isogenies and g_1, g_2 are b -isogenies.

The isogeny diamond is said to be *minimal* if $\ker f_1 \cap \ker g_1 = \{0\}$ and *orthogonal* if $\gcd(a, b) = 1$ (which automatically implies minimality).

Finally, we say that two (a, b) -isogeny diamonds $F = g_2 \circ f_1 = f_2 \circ g_1$ and $F' = g'_2 \circ f'_1 = f'_2 \circ g'_1$ are equivalent if $(F, \ker f_1, \ker g_1) = (F', \ker f'_1, \ker g'_1)$ or $(F, \ker f_1, \ker g_1) = (-F', \ker g'_1, \ker f'_1)$.

Remark 1. Here, minimality is precisely the condition that the diagram is a “pushout” in the sense of Robert [Rob22], which is equivalent to having $\ker f_2 = g_1(\ker f_1)$ and $\ker g_2 = f_1(\ker g_1)$.

Theorem 1. *There is a bijection between*

- *the set of (minimal, orthogonal,) (a, b) -isogeny diamond equivalence classes from $\mathcal{A}$ to $\mathcal{B}$ with $a + b = N$, and*
- *the set of (minimal, orthogonal,) non-diagonal, reducible, maximally isotropic subgroups of $(\mathcal{A} \times \mathcal{B})[N]$.*

This theorem was first stated by Kani [Kan97, Theorem 2.3] for products of elliptic curves and later refined by Robert in [Rob22, Corollary 3.9] to the version above.

2.4 Other Lemmas

Lemma 1. *Assume that Ψ is a diagonal (M, M) -isogeny and that the N -isogeny Φ' is obtained as the push-forward of Φ through Ψ , i.e. $\ker \Phi' = \Psi(\ker \Phi)$. Then*

1. *Φ is a diagonal isogeny if and only if Φ' is diagonal.*
2. *If $\gcd(N, M) = 1$, then $\ker \Phi \subset (E_0 \times E_{AB})[N]$ is maximally isotropic if and only if $\ker \Phi' \subset (F_0 \times F_{AB})[N]$ is maximally isotropic.*

Proof. For ease of notation, let $A_1 := E_0 \times E_{AB}$ and $A_2 = F_0 \times F_{AB}$.

- (1) Assume $\Psi = \psi_1 \times \psi_2$ is diagonal. If $\Phi = \phi_1 \times \phi_2$ is diagonal, then

$$\ker(\Phi) = \ker(\phi_1) \times \ker(\phi_2),$$

hence

$$\ker(\Phi') = \Psi(\ker \Phi) = \psi_1(\ker \phi_1) \times \psi_2(\ker \phi_2),$$

which is a product subgroup in $A_2[N]$. Thus the corresponding quotient isogeny Φ' is diagonal.

Conversely, if Φ' is diagonal, then $\ker(\Phi')$ is a product subgroup of $A_2[N]$. Since Ψ is diagonal, the preimage $\Psi^{-1}(\ker \Phi')$ inside $A_1[N]$ is again a product subgroup, and by the pushforward hypothesis this preimage is $\ker(\Phi)$. Hence Φ is diagonal.

- (2) For $R, S \in A_1[N]$, functoriality of the Weil pairing gives

$$e_{N, A_2}(\Psi(R), \Psi(S)) = e_{N, A_1}(R, \widehat{\Psi}(\Psi(S))) = e_{N, A_1}(R, [\deg \Psi]S) = e_{N, A_1}(R, S)^M,$$

(See, e.g., [EVdGM12, Prop. 11.21]).

If $\gcd(M, N) = 1$, then the map $\zeta \mapsto \zeta^M$ is an automorphism of μ_N . Thus, for $R, S \in \ker \Phi$,

$$e_{N, A_1}(R, S) = 1 \iff e_{N, A_2}(\Psi(R), \Psi(S)) = 1.$$

Hence $\ker \Phi$ is isotropic if and only if $\ker \Phi'$ is isotropic.

Finally, maximality follows from cardinalities. Since Φ is an N -isogeny between PPAVs, $|\ker \Phi| = N^g$. Moreover,

$$|\ker \Phi'| = \frac{|\ker \Phi|}{|\ker \Phi \cap \ker(\Psi)|}.$$

The result follows by minimality of an orthogonal diamond. $\square$

Lemma 2 (Pairing test for diagonality). *Let*

$$K = \langle (P_1, P_2), (Q_1, Q_2) \rangle \subset E_1[N] \times E_2[N]$$

be maximally isotropic of order N^2 . Then the associated (N, N) -isogeny with kernel K is diagonal if and only if $e_N(P_1, Q_1) = 1$. Equivalently, this holds if and only if $e_N(P_2, Q_2) = 1$.

Proof. If the isogeny is diagonal, then $K = H_1 \times H_2$ with $H_i \subset E_i[N]$ cyclic of order N . Hence $P_1, Q_1 \in H_1$, so

$$e_N(P_1, Q_1) = 1.$$

Conversely, assume $e_N(P_1, Q_1) = 1$. Then $H_1 := \langle P_1, Q_1 \rangle$ is isotropic, so $|H_1| \leq N$. Since K is isotropic,

$$1 = e_N((P_1, P_2), (Q_1, Q_2)) = e_N(P_1, Q_1)e_N(P_2, Q_2),$$

and therefore $e_N(P_2, Q_2) = 1$. Hence $H_2 := \langle P_2, Q_2 \rangle$ is isotropic, so $|H_2| \leq N$. Now

$$K \subset H_1 \times H_2,$$

and both sides have order at most N^2 , while $|K| = N^2$. Therefore $K = H_1 \times H_2$, so the isogeny is diagonal. $\square$

2.5 Isogeny hardness assumptions

We now recall some well-known isogeny-based hardness assumptions that will be relevant to our work. Concretely, we recall the well-known DSSP assumption [DFJP14] and follow up with the recently introduced DLIT assumption [Mor23, MS26].

Definition 7 (Decisional Supersingular Product (DSSP) problem). Let d, N be relatively prime positive integers, let E_0, E_1 be supersingular elliptic curves over $\mathbb{F}_{p^2}$, and let $\phi : E_0 \rightarrow E_1$ be an N -isogeny. Let (E_2, E_3, ϕ') be a tuple sampled with probability $1/2$ from one of the following distributions:

- $\mathcal{D}_0 = \{(E_2, E_3, \phi')\}$ where $E_2 \cong E_0/G, E_3 \cong E_1/\phi(G)$ for some cyclic subgroup $G \subseteq E_0[d]$ and $\phi' : E_2 \rightarrow E_3$ is an N -isogeny.
- $\mathcal{D}_1 = \{(E_2, E_3, \phi')\}$ where E_2 is a random supersingular curve and E_3 is a supersingular curve that is the codomain of a random cyclic N -isogeny $\phi' : E_2 \rightarrow E_3$.

Determine from which distribution the tuple (E_2, E_3, ϕ') is sampled.

Definition 8 (Decisional long isogeny with torsion (DLIT) problem). Let d be an integer coprime to N such that $d \gg N$ and let E_0 be a supersingular elliptic curve over $\mathbb{F}_{p^2}$ with a basis $\{P_0, Q_0\}$ of $E_0[N]$. Let (E_1, P_1, Q_1) be a tuple sampled with probability $1/2$ from one of the following distributions:

- $\mathcal{D}_0 = \{(E_1, P_1, Q_1)\}$, where E_1 is the codomain of a uniformly random cyclic isogeny $\phi : E_0 \rightarrow E_1$ of degree d , and $\{P_1, Q_1\}$ is a random basis of $E_1[N]$ such that $e_N(P_1, Q_1) = e_N(P_0, Q_0)^d$.
- $\mathcal{D}_1 = \{(E_1, P_1, Q_1)\}$, where E_1 is the codomain of a uniformly random cyclic isogeny $\phi : E_0 \rightarrow E_1$ of degree d , and $(P_1, Q_1) = (\phi(P_0), \phi(Q_0))$.

Determine from which distribution the tuple (E_1, P_1, Q_1) is sampled.

Remark 2. We note that the computational LIT problem, which asks one to recover the secret isogeny, is assumed to be hard when $d \approx N^2 2^{2\lambda}$, where λ is the security parameter; see for instance [Mor23, Section 4.2] and [MS26]. By analogy, we assume the same parameter regime for DLIT in the protocols that rely on it, namely Kani-Diamond-ZKP.

2.6 Isogeny Relations

Zero-knowledge proofs are defined with respect to a relation $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{W}$, whose elements (x, w) consist of a public statement x and a corresponding secret witness w . This relation induces a language

$$\mathcal{L} = \{x \in \mathcal{X} : \exists w \in \mathcal{W}, (x, w) \in \mathcal{R}\},$$

and a zero-knowledge proof allows a prover to demonstrate that $x \in \mathcal{L}$ without revealing w . Several relations involving isogeny diamonds have been studied in the literature, see [BFGP23]. In this paper, we discuss two isogeny relations involving (a, b) -isogeny diamonds involving four fixed curves E_0, E_A, E_B and E_{AB} ; we defer their formal definitions to the following sections. We will primarily stick to the case of proving knowledge of elliptic curves, though it should be remarked that almost all of our protocols can be extended to PPAVs of any dimension.

2.7 Pseudocode subfunctions and notation

The pseudocodes appearing in Figures 3, 5 and 6 make use of subfunctions to shorten the text and focus on the steps core to our zero-knowledge proofs. We use this subsection to give brief descriptions of these subfunctions.

- H is a statistically hiding and computationally binding commitment scheme. In our SageMath implementation, we use SHA-256, although any hash function with collision resistance and strong preimage resistance would work.
- $\text{OPENCOM}(i)$ returns the inputs for the commitments whose indices contain i . For example, in Windmill-ZKP (Figure 3), $\text{OPENCOM}(1)$ returns the inputs for C_{31} , C_{41} and C_{012} .
- VERIFYCOMMIT checks that the response is indeed the inputs for the commitments.
- $\text{RANDOMBASIS}(E, N)$ returns two points (P, Q) forming a random basis of $E[N]$.
- $\text{KERNELPOINT}(\phi)$ returns a random point generating the kernel of the cyclic isogeny ϕ .
- $\text{KERNELPOINTS}(\Phi)$ returns two random 2-dimensional points generating the kernel of the 2-dimensional isogeny ϕ .
- $\text{DDLOG}(P, Q, R)$ returns an integer vector (a, b) such that $aP + bQ = R$. We only use this function on groups of smooth order, for which the discrete logarithm problem is easy.

- $\text{ISOGYENY}(K)$ returns a pair (ϕ, E) , where ϕ is the isogeny with kernel $\langle K \rangle$ and E is the codomain of ϕ .
- $\text{INDEPENDENTVECTOR}(k_0, N)$ takes as input an integer vector $k_0 = (a_0, b_0)$ and a positive integer N and return an integer vector k_1 such that $k_0 \cdot k_1$ is invertible modulo N .
- $\text{2DISOGYENY}(\kappa_0, \kappa_1)$ returns the 2-dimensional isogeny whose kernel is $\langle \kappa_0, \kappa_1 \rangle$.
- $\text{ISMAXISOTROPIC}(\kappa_0, \kappa_1)$ checks if the 2-dimensional isogeny with kernel $\langle \kappa_0, \kappa_1 \rangle$ is maximally isotropic. This is done by checking that the 2-dimensional Weil pairing of κ_0 and κ_1 is equal to 1.
- $\text{ISDIAGONAL}(\kappa_0, \kappa_1)$ checks if the 2-dimensional isogeny with kernel $\langle \kappa_0, \kappa_1 \rangle$ is diagonal. As this is only checked for maximally isotropic isogenies in our code, this is done by checking that the Weil pairing of the first components of κ_0 and κ_1 is equal to 1.
- $E[[N]]$ is the set of elements of order exactly N on an elliptic curve E .
- In the commitment $C_{i,j}$, the subscript i, j indicates that the value is opened in challenges i and j .

3 Windmill-ZKP construction

We first describe a proof of knowledge for a one-dimensional MSIDH-type isogeny diamond [FMP23]. That is, the prover wants to prove knowledge of a cyclic A -isogeny $\phi_A : E_0 \rightarrow E_A$ whose pushforward along a B -isogeny, which is not known to the prover, has the prescribed codomain E_{AB} . Informally, the prover proves knowledge of a square

$$\begin{array}{ccc} E_0 & \xrightarrow{\phi_A} & E_A \\ \downarrow \phi_B & & \\ E_B & \xrightarrow{\psi_{BA}} & E_{AB} \end{array}$$

where the vertical direction is provided only through the public key of an M-SIDH key exchange.

Let A and B be coprime smooth integers, and let E_0, E_A, E_B, E_{AB} be supersingular elliptic curves over $\mathbb{F}_{p^2}$ whose prime characteristic is of the form $p = AB - 1$. Recall that, since $\gcd(A, B) = 1$, the restriction

$$\phi_B|_{E_0[A]} : E_0[A] \rightarrow E_B[A]$$

is an isomorphism. In M-SIDH, Bob post-composes this map with $[\lambda_B]$, for some $\lambda_B \in (\mathbb{Z}/A\mathbb{Z})^\times$, and publishes the resulting map

$$\rho_B = [\lambda_B] \circ \phi_B|_{E_0[A]},$$

without revealing λ_B . We refer to ρ_B as the *scaled (or masked) action* of the hidden B -isogeny on the A -torsion. The protocol for Windmill-ZKP proves the relation $\mathcal{R}_{\text{windmill}}$ where the statement and witness are

$$x = (E_0, E_A, E_B, E_{AB}, P_0, Q_0, P_A, Q_A, A, B, \rho_B) \quad w = (\phi_A, \psi_{BA}, \alpha)$$




Figure 2: Diagram encompassing Windmill-ZKP.

subject to the relation

$$\begin{aligned}
 &P_0, Q_0 \text{ is a basis of } E_0[B] \\
 &\phi_A : E_0 \rightarrow E_A \text{ is a cyclic } A\text{-isogeny,} \\
 &(P_A, Q_A) = ([\alpha]\phi_A(P_0), [\alpha]\phi_A(Q_0)) \text{ for some secret } \alpha \in (\mathbb{Z}/B\mathbb{Z})^\times, \\
 &\psi_{BA} : E_B \rightarrow E_{AB} \text{ is a cyclic } A\text{-isogeny,} \\
 &\ker \psi_{BA} = \rho_B(\ker \phi_A).
 \end{aligned}$$

We note that this construction can be thought of as an extension of the work by Mokrani and Jao [MJ23]. Their ZKP proves knowledge of a cyclic A -isogeny $\phi_A : E_0 \rightarrow E_A$ such that its action on P_0, Q_0 agrees with the masked points (P_A, Q_A) . In order to prove compatibility with ϕ_B as well, we introduce a third auxiliary isogeny square that certifies that the pushforward of ϕ_A has codomain E_{AB} . We remark that this relation is not hard when E_0 or E_A are either $\mathbb{F}_p$ -rational or admit a small degree isogeny to its conjugate [CV23]. To make this relation hard, we ensure that E_0 is sampled uniformly at random, so that $\text{End}(E_0)$ is completely unknown.

Windmill-ZKP is a cut-and-choose proof built from three auxiliary cyclic B -isogenies. Two of them originate at E_0 , while the third originates at E_B . Transporting the kernel of ϕ_A through these three isogenies produces three cyclic A -isogenies

$$\phi_i : E_{L,i} \longrightarrow E_{R,i}, \quad i \in \{0, 1, 2\}.$$

The first two auxiliary squares certify knowledge of the same isogeny $\phi_A : E_0 \longrightarrow E_A$ together with its masked action on $E_0[B]$. The third certifies that the subgroup obtained by transporting $\ker(\phi_A)$ through the public scaled action ρ_B defines an A -isogeny with codomain E_{AB} .

In order to avoid the soundness vulnerabilities found by De Feo, Dobson, Galbraith and Zobernig [DFDGZ22] of the original SIDH zero-knowledge proof [DFJP14], we apply their technique of representing a point K generating the kernel of an isogeny ψ by sampling a random basis (P, Q) and computing an integer vector $k = (e, f)$ such that $K = k \cdot (P, Q)$. By committing k and using it for multiple challenges, we can make sure that the isogeny ψ' with kernel generated by $k \cdot (P', Q')$ is indeed the pushforward of ψ by an isogeny ϕ mapping (P, Q) onto (P', Q') .

A diagram of the construction is shown in Figure 2 and the pseudocode for the construction can be found in Figure 3.

COMMITMENT	VERIFICATION
$\beta \leftarrow \$(\mathbb{Z}/B\mathbb{Z})^*$ $\gamma \leftarrow \alpha^{-1}\beta$ $(U, V) \leftarrow \text{RANDOMBASIS}(E_0, A)$ $K_A \leftarrow \text{KERNELPOINT}(\phi_A)$ $\vec{k}_A \leftarrow \text{DDLOG}(U, V, K_A)$ $(K_{L,0}, K_{L,1}) \leftarrow \text{RANDOMBASIS}(E_0, B)$ $K_{L,2} \leftarrow \$(E_B[[B]])$ $(\rho_0, \rho_1, \rho_2) \leftarrow (\text{Id}_{E_0}, \text{Id}_{E_0}, \rho_B)$ for $i \in \{0, 1, 2\}$: $(\psi_{L,i}, E_{L,i}) \leftarrow \text{ISOGENY}(K_{L,i})$ $(\phi_i, E_{R,i}) \leftarrow \text{ISOGENY}(\psi_{L,i}(\rho_i(K_A)))$ $(P_{L,i}, Q_{L,i}) \leftarrow \text{RANDOMBASIS}(E_{L,i}, B)$ $K_{C,i} \leftarrow \text{KERNELPOINT}(\widehat{\psi_{L,i}})$ $\vec{k}_{C,i} \leftarrow \text{DDLOG}(P_{L,i}, Q_{L,i}, K_{C,i})$ $(P_{R,i}, Q_{R,i}) \leftarrow ([\beta]\phi_i(P_{L,i}), \beta\phi_i(Q_{L,i}))$ $(U_i, V_i) \leftarrow (\psi_{L,i}(\rho_i(U)), \psi_{L,i}(\rho_i(V)))$ for $i \in \{0, 1\}$: $\vec{k}_{C,i}' \leftarrow \text{INDEPENDENTVECTOR}(\vec{k}_{C,i}, B)$ $R_i \leftarrow \widehat{\psi_{L,i}}(\vec{k}_{C,i}' \cdot (P_{L,i}, Q_{L,i}))$ $\vec{k}_{C,i}'' \leftarrow \text{DDLOG}(P_0, Q_0, R_i)$ $r_3, r_4, r_{34}, r_{30}, r_{31}, r_{32},$ $r_{40}, r_{41}, r_{42}, r_{012} \leftarrow \$(\{0, 1\})^\lambda$ $C_3 \leftarrow H(U, V, r_3)$ $C_4 \leftarrow H(\gamma, r_4)$ $C_{34} \leftarrow H(\vec{k}_{C,0}, \vec{k}_{C,1}, \vec{k}_{C,2}, \vec{k}_{C,0}', \vec{k}_{C,1}',$ $\vec{k}_{C,0}'', \vec{k}_{C,1}'', r_{34})$ $C_{30} \leftarrow H(E_{L,0}, U_0, V_0, P_{L,0}, Q_{L,0}, r_{30})$ $C_{31} \leftarrow H(E_{L,1}, U_1, V_1, P_{L,1}, Q_{L,1}, r_{31})$ $C_{32} \leftarrow H(E_{L,2}, U_2, V_2, P_{L,2}, Q_{L,2}, r_{32})$ $C_{40} \leftarrow H(E_{R,0}, P_{R,0}, Q_{R,0}, r_{40})$ $C_{41} \leftarrow H(E_{R,1}, P_{R,1}, Q_{R,1}, r_{41})$ $C_{42} \leftarrow H(E_{R,2}, P_{R,2}, Q_{R,2}, r_{42})$ $C_{012} \leftarrow H(K_A, \beta, r_{012})$ return $C_3, C_4, C_{34}, C_{30}, C_{31}, C_{32},$ $C_{40}, C_{41}, C_{42}, C_{012}$ 	if $\neg \text{VERIFYCOMMIT}(\text{com}, \text{chal}, \text{rsp})$: return false if $\text{chal} \in \{0, 1, 2\}$: $i \leftarrow \text{chal}$ if $E_{L,i} \notin \mathcal{E} \vee E_{R,i} \notin \mathcal{E}$: return false if $\langle U_i, V_i \rangle \neq E_{L,i}[A]$: return false if $\langle P_{L,i}, Q_{L,i} \rangle \neq E_{L,i}[B]$: return false $K_{A,i} \leftarrow \vec{k}_A \cdot (U_i, V_i)$ if $K_{A,i} \notin E_{L,i}[[A]]$: return false $(\phi_i, E'_{R,i}) \leftarrow \text{ISOGENY}(K_{A,i})$ if $E'_{R,i} \neq E_{R,i}$: return false if $\beta\phi_i(P_{L,i}, Q_{L,i}) \neq (P_{R,i}, Q_{R,i})$: return false if $\text{chal} = 3$: if $\langle U, V \rangle \neq E_0[A]$: return false $(\rho_0, \rho_1, \rho_2) \leftarrow (\text{Id}_{E_0}, \text{Id}_{E_0}, \rho_B)$ for $i \in \{0, 1, 2\}$: if $E_{L,i} \notin \mathcal{E}$: return false if $\langle P_{L,i}, Q_{L,i} \rangle \neq E_{L,i}[B]$: return false $K'_{L,i} \leftarrow \vec{k}_{C,i} \cdot (P_{L,i}, Q_{L,i})$ if $K'_{L,i} \notin E_{L,i}[[B]]$: return false $(\widehat{\psi_{L,i}}, E'_0) \leftarrow \text{ISOGENY}(K'_{L,i})$ if $E'_0 \neq \text{codom}(\rho_i)$: return false if $B\rho_i(U, V) \neq \widehat{\psi_{L,i}}(U_i, V_i)$: return false for $i \in \{0, 1\}$: $R_i \leftarrow \vec{k}_{C,i}' \cdot \widehat{\psi_{L,i}}(P_{L,i}, Q_{L,i})$ if $R_i \neq \vec{k}_{C,i}'' \cdot (P_0, Q_0)$: return false if $\langle R_0, R_1 \rangle \neq E_0[B]$: return false if $\text{chal} = 4$: for $i \in \{0, 1, 2\}$: if $E_{R,i} \notin \mathcal{E}$: return false if $\langle P_{R,i}, Q_{R,i} \rangle \neq E_{R,i}[B]$: return false $K'_{R,i} \leftarrow \vec{k}_{C,i} \cdot (P_{R,i}, Q_{R,i})$ if $K'_{R,i} \notin E_{R,i}[[B]]$: return false $(\widehat{\psi_{R,i}}, E'_{A,i}) \leftarrow \text{ISOGENY}(K'_{R,i})$ for $i \in \{0, 1\}$: if $E'_{A,i} \neq E_A$: return false $R'_i \leftarrow \vec{k}_{C,i}' \cdot \widehat{\psi_{R,i}}(P_{R,i}, Q_{R,i})$ if $R'_i \neq [\gamma]\vec{k}_{C,i}'' \cdot (P_A, Q_A)$: return false if $\langle R'_0, R'_1 \rangle \neq E_A[B]$: return false if $E'_{A,2} \neq E_{AB}$: return false return true
CHALLENGE $\text{chal} \leftarrow \$(\{0, 1, 2, 3, 4\})$ return chal 	
RESPONSE return $\text{OPENCOM}(\text{chal})$	

Figure 3: The Windmill-ZKP protocol. Notation used is defined in Subsection 2.7.

3.1 Security Analysis of Windmill-ZKP

The correctness of Windmill-ZKP follows trivially our construction.

Proposition 1. *Given a computationally binding commitment scheme H , Windmill-ZKP is 5-sound.*

Proof. We show that for a fixed commitment, if valid answers are provided to all 5 possible challenges², then one can compute a cyclic isogeny $\sigma : E_0 \rightarrow E_A$ of degree A and an integer α such that $(P_A, Q_A) = ([\alpha]\sigma(P_0), [\alpha]\sigma(Q_0))$ and the codomain of $[\phi_B]_*\sigma$ ³ is E_{AB} . The first part of this is the same as [MJ23, Theorem 8], so we highlight the idea. We then prove the pushforward condition.

We can obtain two candidate degree A isogenies $\sigma_0, \sigma_1 : E_0 \rightarrow E_A$ from the isogenies $\widehat{\psi}_{L,0} : E_{L,0} \rightarrow E_0$, $\phi_0 : E_{L,0} \rightarrow E_{R,0}$ and $\widehat{\psi}_{L,1} : E_{L,1} \rightarrow E_0$, $\phi_1 : E_{L,1} \rightarrow E_{R,1}$ (respectively). From a straightforward calculation, the kernels of σ_0 and σ_1 are both $\langle [e]U + [f]V \rangle$, so $\sigma_0 = \pm\sigma_1$. If $\sigma_0 = \sigma_1 = \sigma$, then we have valid candidate. If $\sigma_0 = -\sigma$, then both isogenies have the same domain, codomain and masked torsion (with opposite masking scalars). Hence, we can set either isogeny as our candidate σ and multiply the other by -1 so that they coincide. We also have a candidate for α in $\gamma^{-1}\beta$. For $i \in \{0, 1\}$, define $R_i = a_i P_0 + b_i Q_0 = \widehat{\psi}_{L,i}([c'_i]P_{L,i} + [d'_i]Q_{L,i})$ and $R'_i = [\beta]\sigma_i(R_i) = \widehat{\psi}_{R,i}([c'_i]P_{R,i} + [d'_i]Q_{R,i})$. Since (R_0, R_1) is a basis of $E_0[B]$ and (R'_0, R'_1) is the corresponding basis of $E_A[B]$, these identities imply that

$$(P_A, Q_A) = ([\alpha]\sigma(P_0), [\alpha]\sigma(Q_0)).$$

Hence, the extracted isogeny σ has the required masked action on the B -torsion.

All that is left to prove is that the codomain of $[\phi_B]_*\sigma$ is E_{AB} . Repeat the same idea as before. From the isogenies $\widehat{\psi}_{L,2} : E_{L,2} \rightarrow E_B$ and $\phi_2 : E_{L,2} \rightarrow E_{R,2}$ we get a degree A isogeny $\sigma' : E_B \rightarrow E'$ whose kernel is

$$\ker(\sigma') = \langle \widehat{\psi}_{L,2}([e]U_2 + [f]V_2) \rangle = \langle [B]\rho_B([e]U + [f]V) \rangle = \phi_B(\ker \sigma).$$

Thus, σ' is the pushforward of σ by ϕ_B and $E' = E_{AB}$. This completes the proof. $\square$

Proposition 2. *Given a statistically hiding commitment scheme H , Windmill-ZKP is computationally zero-knowledge assuming DSSP is secure for $d = B$ and $N = A$.*

Proof. We want to create a simulator that publishes a valid (commitment, challenge, response) transcript. Since there are only five challenges, we explain how to simulate the response for each challenge.

Since H is statistically hiding, the simulator can sample random values for the commitments whose input is not checked in the challenge.

For any challenge in $\{0, 1, 2\}$, the simulator samples a random A -isogeny $\phi_i : E_{L,i} \rightarrow E_{R,i}$, samples the required torsion bases, and computes the opened torsion images honestly. In a real transcript, ϕ_i is the pushforward of the witness A -isogeny through a random cyclic B -isogeny. The indistinguishability of these two distributions is precisely the DSSP assumption with $d = B$ and $N = A$.

For challenge 3, we can honestly compute the left isogenies and honestly sample the random bases.

The simulator for the final challenge is very similar to the previous one, the main difference being that we sample isogenies for the right side. Since the isogeny degrees are relatively prime, the distribution of a pushforward of a uniformly random cyclic B -isogeny is itself a uniformly random cyclic B -isogeny. Therefore, the output of the simulator is identical from that of an honest one. $\square$

²Which we can assume agree on the committed values as H is computationally binding.

³Where $[\phi_A]_*\sigma$ denotes the isogeny with kernel $\phi_A(\ker \sigma)$.

Remark 3. Windmill-ZKP can be seen as an adaptation of Mokrani–Jao’s M-SIDH zero-knowledge proof [MJ23]. One can also adapt the zero-knowledge proof by Basso [Bas24] in this Windmill-ZKP construction. Compared to the Mokrani–Jao proof which is 4-sound, Basso’s proof has the advantage of only being 3-sound at the expense of relying on a stronger variant of the DSSP assumption, named DSSP with torsion. In this variant, an attacker is also provided the M-SIDH public key of the isogeny doing the pushforward. Adapting Basso’s proof in Windmill-ZKP, one obtains a 4-sound zero-knowledge proof whose security relies on DSSP with torsion.

4 Cube-ZKP construction

For the remainder of the paper, we now consider the problem of proving knowledge of an isogeny diamond when the prover knows both its directions. In particular, we consider the following relation

$$\mathcal{R}_\diamond = \left\{ ((E_0, E_A, E_B, E_{AB}, A, B), (\phi_A, \phi_B)) \left| \begin{array}{l} \phi_A : E_0 \rightarrow E_A \text{ is a cyclic } A\text{-isogeny,} \\ \phi_B : E_0 \rightarrow E_B \text{ is a cyclic } B\text{-isogeny,} \\ [\phi_A]_* \phi_B : E_A \rightarrow E_{AB}. \end{array} \right. \right\}.$$

We first present **Cube-ZKP**, which generalizes the techniques of [DFDGZ22] from a single isogeny to a commuting square of isogenies, using torsion-basis commitments to force consistency between the different responses.

The parameters of this protocol consists of three smooth pairwise coprime integers A, B, C and a prime p such that $ABC \mid p+1$. The prover claims to know cyclic isogenies

$$\begin{aligned} \phi_A : E_0 &\longrightarrow E_A, & \phi_B : E_0 &\longrightarrow E_B \\ \ker(\phi_A) &= \langle K_A \rangle, & \ker(\phi_B) &= \langle K_B \rangle \end{aligned}$$

of degrees A and B , respectively, such that their pushforwards complete to a commuting isogeny diamond, corresponding to the front face in Figure 4(a).

4.1 Protocol overview

Informally, the main idea of the **Cube-ZKP** protocol consists of pushing an SIDH-style isogeny square through a random isogeny path of degree C . This produces a cube of isogenies whose front face is the secret witness.

The prover samples a random cyclic C -isogeny $\phi_C : E_0 \longrightarrow E_C$, where C is coprime to both A and B , and pushes the two witnesses through ϕ_C . This produces the back face of the cube, namely the curves E_C, E_{AC}, E_{BC} , and E_{ABC} , together with four C -isogenies connecting the public front face to the back face. The protocol verifies three types of consistency.

First, the prover commits to the four C -isogenies through the kernels of their duals. It uses a common pair of coefficients $\vec{k}_C = (e_C, f_C)$ to describe these kernels with respect to compatible C -torsion bases. The corresponding challenges verify that these four dual C -isogenies connect the vertices of the back face to the correct vertices of the front diamond.

Second, the prover commits to the B -isogenies on the back face. The same coefficients $\vec{k}_B = (e_B, f_B)$ are used for the two parallel B -isogenies. The corresponding challenges verify both the kernels of these isogenies and the transport of the torsion bases needed by the adjacent faces.

Third, the prover commits to the A -isogeny $\psi_{CA} = [\phi_C]_* \phi_A$. The isogeny ψ_{CA} can be proven to be parallel to ϕ_A by the fact that the C -isogenies ϕ_C and ψ_{AC} are parallel. Likewise, once the B -isogenies are proven to be parallel, $\psi_{BA} = [\phi_B]_* \phi_A$ can, in theory, be constructed by someone receiving every responses to a fixed commitment.




Figure 4: The Cube-ZKP construction.

In order to avoid the same soundness issues from the original SIDH zero-knowledge proof discovered in [DFDGZ22], the commitments do not merely specify the individual edges of the cube. They also bind parallel edges to compatible torsion bases and kernels. This prevents different challenge responses from describing unrelated isogeny squares.

The diagram for this construction is given in Figure 4(b) and the pseudocode for this construction can be found in Figure 5.

4.2 Security Analysis of Cube-ZKP

The correctness of Cube-ZKP follows trivially from our construction.

Proposition 3. *Given a computationally binding commitment scheme H , Cube-ZKP is γ -sound.*

Proof. Fix a commitment for which valid responses are available to all seven challenges.

Using the same reasoning as in the proof of soundness for Windmill-ZKP, the responses for challenges 0, 1 and 4 can be used to construct the isogenies $\widehat{\phi}_C, \widehat{\psi}_{CA}$, and $\widehat{\psi}_{AC}$ in a way that forces $\widehat{\psi}_{AC}$ to be the pushforward of $\widehat{\phi}_C$ by $\widehat{\psi}_{CA}$. Therefore, the isogeny diamond generated by these three isogenies determines a cyclic A -isogeny $\sigma_A : E_0 \rightarrow E_A$. Similarly, the square formed by $\widehat{\phi}_C, \widehat{\psi}_{CB}$, and $\widehat{\psi}_{BC}$ determines a cyclic B -isogeny $\sigma_B : E_0 \rightarrow E_B$.

Analogously, the squares involving $\widehat{\psi}_{BC}, \widehat{\psi}_{BCA}, \widehat{\psi}_{ABC}$ and $\widehat{\psi}_{AC}, \widehat{\psi}_{ACB}, \widehat{\psi}_{ABC}$ determine cyclic isogenies

$$\sigma'_A : E_B \rightarrow E_{AB}, \quad \sigma'_B : E_A \rightarrow E_{AB}.$$

It remains to show that these are the pushforwards of the extracted witness isogenies. Since $\ker(\psi_{BCA}) = \psi_{CB}(\ker(\psi_{CA}))$, we obtain

$$\begin{aligned} \ker(\sigma'_A) &= \widehat{\psi}_{BC}(\ker(\psi_{BCA})) \\ &= \widehat{\psi}_{BC}(\psi_{CB}(\ker(\psi_{CA}))) = \sigma_B(\widehat{\phi}_C(\ker(\psi_{CA}))) = \sigma_B(\ker(\sigma_A)). \end{aligned}$$

Hence, $\sigma'_A = [\sigma_B]_* \sigma_A$. Similarly, $\ker \sigma'_B = \sigma_A(\ker \sigma_B)$ and $\sigma'_B = [\sigma_A]_* \sigma_B$.

Since both the extracted pushforwards have codomain E_{AB} , we conclude that the codomain of $[\sigma_B]_* \sigma_A$ is E_{AB} . Thus (σ_A, σ_B) is a valid witness. $\square$

Proposition 4. *Given a statistically hiding commitment scheme H , Cube-ZKP is computationally zero-knowledge assuming DSSP is secure for $d = C$ and $N = \max\{A, B\}$.*

COMMITMENT	VERIFICATION
$K_{\phi_A} \leftarrow \text{KERNELPOINT}(\phi_A)$ $(U_0, V_0) \leftarrow \text{RANDOMBASIS}(E_0, B)$ $K_B \leftarrow \text{KERNELPOINT}(\phi_B)$ $\vec{k}_B \leftarrow \text{DDLOG}(U_0, V_0, K_B)$ $K_{\phi_C} \leftarrow \$ E_0[[C]]$ $(\phi_C, E_C) \leftarrow \text{ISOGENY}(K_{\phi_C})$ $K_{\psi_{CA}} \leftarrow \phi_C(K_{\phi_A})$ $(\psi_{CA}, E_{AC}) \leftarrow \text{ISOGENY}(K_{\psi_{CA}})$ $(\psi_{CB}, E_{BC}) \leftarrow \text{ISOGENY}(\vec{k}_B \cdot \phi_C(U_0, V_0))$ $(\psi_{ACB}, E_{ABC}) \leftarrow \text{ISOGENY}(\psi_{CA}(\phi_C(K_{\phi_B})))$ $(U_C, V_C) \leftarrow \phi_C(U_0, V_0)$ $(U_{AC}, V_{AC}) \leftarrow \psi_{CA}(U_C, V_C)$ $(P_C, Q_C) \leftarrow \text{RANDOMBASIS}(E_C, C)$ $K'_C \leftarrow \text{KERNELPOINT}(\widehat{\phi_C})$ $\vec{k}_C \leftarrow \text{DDLOG}(P_C, Q_C, K'_C)$ $(P_{AC}, Q_{AC}) \leftarrow \psi_{CA}(P_C, Q_C)$ $(P_{BC}, Q_{BC}) \leftarrow \psi_{CB}(P_C, Q_C)$ $(P_{ABC}, Q_{ABC}) \leftarrow \psi_{ACB}(P_{AC}, Q_{AC})$ $r_{0123}, r_{045}, r_{146}, r_{25},$ $r_{36}, r_{45}, r_{46}, r_{56}, r_4 \leftarrow \$ \{0, 1\}^\lambda$ $C_{0123} = H(r_{0123}, \vec{k}_C)$ $C_{045} = H(r_{045}, E_C, P_C, Q_C),$ $C_{146} = H(r_{146}, E_{AC}, P_{AC}, Q_{AC})$ $C_{25} = H(r_{25}, E_{BC}, P_{BC}, Q_{BC})$ $C_{36} = H(r_{36}, E_{ABC}, P_{ABC}, Q_{ABC})$ $C_{45} = H(r_{45}, U_C, V_C)$ $C_{46} = H(r_{46}, U_{AC}, V_{AC})$ $C_{56} = H(r_{56}, \vec{k}_B)$ $C_4 = H(r_4, K_{\psi_{CA}})$ return $C_{0123}, C_{045}, C_{146}, C_{25},$ $C_{36}, C_{45}, C_{46}, C_{56}, C_4$	if $\neg \text{VERIFYCOMMIT}(\text{com}, \text{chal}, \text{rsp})$: return false if $\text{chal} = 0$: if $E_C \notin \mathcal{E} \vee \langle P_C, Q_C \rangle \neq E_C[C]$: return false $(\widehat{\phi_C}, E'_0) \leftarrow \text{ISOGENY}(\vec{k}_C \cdot (P_C, Q_C))$ if $E'_0 \neq E_0 \vee \deg(\widehat{\phi_C}) \neq C$: return false if $\text{chal} = 1$: if $E_{AC} \notin \mathcal{E} \vee \langle P_{AC}, Q_{AC} \rangle \neq E_{AC}[C]$: return false $(\widehat{\psi_{CA}}, E'_A) \leftarrow \text{ISOGENY}(\vec{k}_C \cdot (P_{AC}, Q_{AC}))$ if $E'_A \neq E_A \vee \deg(\widehat{\psi_{CA}}) \neq C$: return false if $\text{chal} = 2$: if $E_{BC} \notin \mathcal{E} \vee \langle P_{BC}, Q_{BC} \rangle \neq E_{BC}[C]$: return false $(\widehat{\psi_{CB}}, E'_B) \leftarrow \text{ISOGENY}(\vec{k}_C \cdot (P_{BC}, Q_{BC}))$ if $E'_B \neq E_B \vee \deg(\widehat{\psi_{CB}}) \neq C$: return false if $\text{chal} = 3$: if $E_{ABC} \notin \mathcal{E} \vee \langle P_{ABC}, Q_{ABC} \rangle \neq E_{ABC}[C]$: return false $(\widehat{\psi_{ACB}}, E'_{AB}) \leftarrow \text{ISOGENY}(\vec{k}_C \cdot (P_{ABC}, Q_{ABC}))$ if $E'_{AB} \neq E_{AB} \vee \deg(\widehat{\psi_{ACB}}) \neq C$: return false if $\text{chal} = 4$: if $E_C \notin \mathcal{E} \vee \langle P_C, Q_C \rangle \neq E_C[C]$: return false $(\psi_{CA}, E'_{AC}) \leftarrow \text{ISOGENY}(K_{\psi_{CA}})$ if $\deg(\psi_{CA}) \neq A \vee E'_{AC} \neq E_{AC}$: return false if $(P_{AC}, Q_{AC}) \neq \psi_{CA}(P_C, Q_C)$: return false if $(U_{AC}, V_{AC}) \neq \psi_{CA}(U_C, V_C)$: return false if $\text{chal} = 5$: if $E_C \notin \mathcal{E} \vee \langle P_C, Q_C \rangle \neq E_C[C]$: return false if $\langle U_C, V_C \rangle \neq E_C[B]$: return false $(\psi_{CB}, E'_{BC}) \leftarrow \text{ISOGENY}(\vec{k}_B \cdot (U_C, V_C))$ if $\deg(\psi_{CB}) \neq B \vee E'_{BC} \neq E_{BC}$: return false if $(P_{BC}, Q_{BC}) \neq \psi_{CB}(P_C, Q_C)$: return false if $\text{chal} = 6$: if $E_{AC} \notin \mathcal{E} \vee \langle P_{AC}, Q_{AC} \rangle \neq E_{AC}[C]$ $\vee \langle U_{AC}, V_{AC} \rangle \neq E_{AC}[B]$: return false $(\psi_{ACB}, E'_{ABC}) \leftarrow \text{ISOGENY}(\vec{k}_B \cdot (U_{AC}, V_{AC}))$ if $\deg(\psi_{ACB}) \neq B \vee E'_{ABC} \neq E_{ABC}$ $\vee (P_{ABC}, Q_{ABC}) \neq \psi_{ACB}(P_{AC}, Q_{AC})$: return false return true
CHALLENGE $\text{chal} \leftarrow \$ \{0, 1, 2, 3, 4, 5, 6\}$ return chal	
RESPONSE return $\text{OPENCOM}(\text{chal})$	

Figure 5: The Cube-ZKP protocol. Notation used is defined in Subsection 2.7.

Proof. We construct a simulator that, for any prescribed challenge, produces an accepting transcript that is computationally indistinguishable from an honest transcript. Since H is statistically hiding, the simulator may sample independently all committed values that are not checked in the given challenge.

The challenges 0, 1, 2, and 3 reveal one of the four dual cyclic C -isogenies. For challenge 0, the simulator honestly samples a cyclic C -isogeny $\widehat{\phi}_C : E_C \rightarrow E_0$ together with a basis of $E_C[C]$ and coefficients describing its kernel. The simulations for challenges 1, 2, and 3 are analogous. In each case, the simulator samples a cyclic C -isogeny with the prescribed codomain and generates the corresponding torsion basis and kernel coordinates.

Since the isogeny degrees are coprime, the pushforward of a uniformly sampled cyclic C -isogeny through an A or B -isogeny is again distributed as a uniformly sampled cyclic C -isogeny. Therefore, the opened data has the same distribution as in an honest execution, except for the unopened commitments, which are statistically hiding.

The challenges 4, 5, and 6 reveal one of the A or B -isogenies forming the back face of the cube. Consider, for example, challenge 4. In an honest execution, the revealed A -isogeny $\psi_{CA} : E_C \rightarrow E_{AC}$ is obtained by pushing forward the secret isogeny ϕ_A through the random cyclic C -isogeny ϕ_C . On the other hand, the simulator samples an independent cyclic A -isogeny of the prescribed degree and generates the torsion bases and their images honestly.

Distinguishing this simulated distribution from the honest one would imply distinguishing an A -isogeny obtained as the pushforward through a random cyclic C -isogeny from an independently sampled cyclic A -isogeny. This would give a distinguisher for DSSP with $d = C$, and $N = A$. Challenges 5 and 6 reduce analogously to DSSP with $d = C$ and $N = B$. Therefore, all three cases are covered by DSSP with $N = \max\{A, B\}$.

Since A , B , and C are pairwise coprime, the isogenies act as isomorphisms on the torsion subgroups used in the protocol, and the sampled torsion bases have the same distribution as in an honest execution. $\square$

4.3 Kani Variant of the Cube

We now describe a variant of Cube-ZKP, which we call Kube-ZKP. This variant is designed for the setting in which the two isogeny degrees A and B are non-smooth, while their sum is smooth. The construction keeps the four challenges corresponding to the dual C -isogenies, but replaces the multiple challenges on the back diamond by a single challenge that reveals the entire back diamond. The protocol therefore has five possible challenges. The first four corresponding to the dual C -isogenies are unchanged, except that Kani-Diamond-ZKP requires $C \approx AB$ to be secure. For the back face, instead of committing separately to the kernels and torsion actions, the prover commits to the kernel $\ker(\Phi') = \langle \kappa'_0, \kappa'_1 \rangle \subseteq (E_C \times E_{ABC})[A + B]$ of the corresponding Kani isogeny.

Constructing the cube in this Kube-ZKP variant is slightly different. Since A, B are non-smooth and each A and B -isogenies cannot be evaluated efficiently as in Cube-ZKP. On top of sampling the kernel for a C -isogeny $\phi_C : E_0 \rightarrow E_C$ one needs to compute the pushforward kernel along the diamond and obtain the kernel C -isogeny $\phi_{ABC} : E_{AB} \rightarrow E_{ABC}$. Using these two C -isogenies, one can push the kernel generators for the Kani isogeny $\Phi : E_0 \times E_{AB} \rightarrow E_A \times E_B$ along these C -isogenies ϕ_C, ϕ_{ABC} to get kernel generators for an isogeny $\Phi' : E_C \times E_{ABC} \rightarrow E_{AC} \times E_{BC}$ and complete the back diamond.

To compute the kernel generator of ϕ_{ABC} we use the Kani isogeny $\Phi : E_0 \times E_{AB} \rightarrow E_A \times E_B$ which in matrix form this isogeny can be written as

$$\Phi = \begin{pmatrix} \phi_A & \widehat{\psi_{BA}} \\ -\phi_B & \widehat{\psi_{AB}} \end{pmatrix}.$$

With respect to the product principal polarizations, the dual isogeny $\widehat{\Phi}$ is represented by the corresponding adjoint matrix of Φ [MMP⁺23, §2.3]. Write $\ker(\phi_C) = \langle K_C \rangle$ and we need to compute $K_{ABC} = \psi_B(\phi_A(K_C))$. First, compute

$$\Phi((K_C, 0)) = (\phi_A(K_C), -\phi_B(K_C)),$$

write $K_{AC} = \phi_A(K_C)$ and compute

$$\widehat{\Phi}((K_{AC}, 0)) = ([A]K_C, \psi_B(K_{AC}))$$

to get the desired generator $K_{ABC} = \psi_B(K_{AC})$.

To facilitate these isogeny computations, one needs to change the underlying prime in the public statement to a prime p such that $(A + B)C \mid p + 1$ where $(A + B)$ and C are both smooth. In practice, $A + B$ will typically be a power of 2, while C will be a power of 3. This facilitates computing the 1-dimensional C -isogenies and the 2-dimensional $(A + B)$ -isogenies. Additionally the prover now claims to know two cyclic isogenies ϕ_A, ϕ_B of degree A and B (respectively) as well as their pushforwards, such that they can evaluate these isogenies using the two dimensional representation associated to the diamond via Kani's Lemma.

The pseudocode for this construction can be found in Figure 6.

Proposition 5. *Given a computationally binding commitment scheme H , Kube-ZKP is 5-sound.*

Proof. The response to Challenge 4 reveals the Kani isogeny Φ' together with the torsion data needed to recover the four maps forming the back square. Because of the fact that $C > A + B$ and that the verifier checks that $e_C(P_{AC}, Q_{AC}) = e_C(P_C, Q_C)^A$ and $e_C(P_{BC}, Q_{BC}) = e_C(P_C, Q_C)^B$, we are guaranteed that the degrees of these four isogenies are as expected. Therefore, valid responses to all five challenges provide the same information as valid responses to all seven challenges of Cube-ZKP. $\square$

Proposition 6. *Given a statistically hiding commitment scheme H , Kube-ZKP is computationally zero-knowledge assuming DSSP is secure for $d = C$ and $N = AB$.*

Proof. The responses revealing C -isogenies are identical to those of Cube-ZKP and, therefore, are perfectly simulatable.

The response revealing the isogeny square generated by ψ_{CA} and ψ_{CB} is equivalent, in terms of information, to revealing $\psi_{ACB}\psi_{CA}$, a cyclic AB -isogeny. Therefore, we can simulate it by sampling a random isogeny of the same degree using the Kani representation described in the construction of Kube-ZKP above, and complete the associated square. To do so, we start with a supersingular elliptic curve F_0 with known endomorphism ring. Since the endomorphism ring is known, we sample a cyclic AB -isogeny $\rho : F_0 \rightarrow E_2$ for which we can compute the associated Kani representation⁴. We can then sample a chain of k random cyclic C -isogenies $\xi_1, \dots, \xi_k$ efficiently compute a representation of $[\xi_k \circ \dots \circ \xi_1]_* \rho : E_2 \rightarrow E_3$. Since the isogeny graph with Borel level-structure is Ramanujan [CL26, Cor. 1.8], $[\xi_k \circ \dots \circ \xi_1]_* \rho$ is indistinguishable from a uniformly random cyclic isogeny of the same degree, so long as C^k is large enough. The problem of distinguishing this simulator from an honest output is equivalent to DSSP with $d = C$ and $N = AB$. $\square$

Remark 4. We briefly note some trade-offs between Cube-ZKP and Kube-ZKP and refer to Section 6 for a more explicit and detailed comparison. The soundness of these protocols are 7-sound and 5-sound respectively, however one needs $C \approx \max(A, B)$ in Cube-ZKP and $C \approx AB$ in Kube-ZKP. In addition, the prime characteristics are different in these constructions but they are roughly the same size ($p = ABC - 1 \approx \max(A, B, C)^3$ in

⁴We discuss how to generate such an isogeny in Section 6

COMMITMENT	VERIFICATION
$(\kappa_0, \kappa_1) \leftarrow \text{KERNELPOINTS}(\Phi)$ $K_C \leftarrow \$ E_0[[C]]$ $(K_{AC}, \cdot) \leftarrow \Phi((K_C, 0))$ $(\cdot, K_{ABC}) \leftarrow \widehat{\Phi}((K_{AC}, 0))$ $(\phi_C, E_C) \leftarrow \text{ISOGENY}(K_C)$ $(\psi_{ABC}, E_{ABC}) \leftarrow \text{ISOGENY}(K_{ABC})$ $(\kappa'_0, \kappa'_1) \leftarrow (\phi_C \times \psi_{ABC})(\kappa_0, \kappa_1)$ $(\Phi', E_{AC} \times E_{BC}) \leftarrow$ $\quad 2\text{DISOGENY}(\kappa'_0, \kappa'_1)$ $(P_C, Q_C) \leftarrow \text{RANDOMBASIS}(E_C, C)$ $K'_C \leftarrow \text{KERNELPOINT}(\widehat{\phi_C})$ $\vec{k}_C \leftarrow \text{DDLOG}(P_C, Q_C, K'_C)$ $((P_{AC}, P_{BC}), (Q_{AC}, Q_{BC})) \leftarrow$ $\quad \Phi'((P_C, 0), (Q_C, 0))$ $(P_{BC}, Q_{BC}) \leftarrow (-P_{BC}, -Q_{BC})$ $((\cdot, P_{ABC}), (\cdot, Q_{ABC})) \leftarrow$ $\quad \widehat{\Phi}'((P_{AC}, 0), (Q_{AC}, 0))$ $r_{0123}, r_{04}, r_{14}, r_{24}, r_{34}, r_4 \leftarrow \$ \{0, 1\}^\lambda$ $C_{0123} \leftarrow H(\vec{k}_C, r_{0123})$ $C_{04} \leftarrow H(E_C, P_C, Q_C, r_{04})$ $C_{14} \leftarrow H(E_{AC}, P_{AC}, Q_{AC}, r_{14})$ $C_{24} \leftarrow H(E_{BC}, P_{BC}, Q_{BC}, r_{24})$ $C_{34} \leftarrow H(E_{ABC}, P_{ABC}, Q_{ABC}, r_{34})$ $C_4 \leftarrow H(\kappa'_0, \kappa'_1, r_4)$ return $C_{0123}, C_{04}, C_{14}, C_{24}, C_{34}, C_4$ CHALLENGE $\text{chal} \leftarrow \$ \{0, 1, 2, 3, 4\}$ return chal RESPONSE return $\text{OPENCOM}(\text{chal})$	if $\neg \text{VERIFYCOMMIT}(\text{com}, \text{chal}, \text{rsp})$: return false if $\text{chal} = 0$: if $E_C \notin \mathcal{E} \vee \langle P_C, Q_C \rangle \neq E_C[C]$: return false $(\widehat{\phi_C}, E'_0) \leftarrow \text{ISOGENY}(\vec{k}_C \cdot (P_C, Q_C))$ if $E'_0 \neq E_0 \vee \deg(\widehat{\phi_C}) \neq C$: return false if $\text{chal} = 1$: if $E_{AC} \notin \mathcal{E} \vee \langle P_{AC}, Q_{AC} \rangle \neq E_{AC}[C]$: return false $(\widehat{\psi_{CA}}, E'_A) \leftarrow \text{ISOGENY}(\vec{k}_C \cdot (P_{AC}, Q_{AC}))$ if $E'_A \neq E_A \vee \deg(\widehat{\psi_{CA}}) \neq C$: return false if $\text{chal} = 2$: if $E_{BC} \notin \mathcal{E} \vee \langle P_{BC}, Q_{BC} \rangle \neq E_{BC}[C]$: return false $(\widehat{\psi_{CB}}, E'_B) \leftarrow \text{ISOGENY}(\vec{k}_C \cdot (P_{BC}, Q_{BC}))$ if $E'_B \neq E_B \vee \deg(\widehat{\psi_{CB}}) \neq C$: return false if $\text{chal} = 3$: if $E_{ABC} \notin \mathcal{E} \vee \langle P_{ABC}, Q_{ABC} \rangle \neq E_{ABC}[C]$: return false $(\widehat{\psi_{ACB}}, E'_{AB}) \leftarrow \text{ISOGENY}(\vec{k}_C \cdot (P_{ABC}, Q_{ABC}))$ if $E'_{AB} \neq E_{AB} \vee \deg(\widehat{\psi_{ACB}}) \neq C$: return false if $\text{chal} = 4$: if $\neg \kappa'_0, \kappa'_1 \in E_C \times E_{ABC}$: return false if $\neg \text{ISMAXISOTROPIC}(\kappa'_0, \kappa'_1)$: return false if $\text{ISDIAGONAL}(\kappa'_0, \kappa'_1)$: return false $(\Phi', \mathcal{A}) \leftarrow 2\text{DISOGENY}(\kappa'_0, \kappa'_1)$ if $\mathcal{A} \neq E_{AC} \times E_{BC}$: return false if $\deg(\Phi') \neq N$: return false if $\langle P_C, Q_C \rangle \neq E_C[C]$: return false if $((P_{AC}, -P_{BC}), (Q_{AC}, -Q_{BC})) \neq \Phi'((P_C, 0), (Q_C, 0))$: return false if $((\cdot, P_{ABC}), (\cdot, Q_{ABC})) \neq \widehat{\Phi}'((P_{AC}, 0), (Q_{AC}, 0))$: return false if $e_C(P_{AC}, Q_{AC}) \neq e_C(P_C, Q_C)^A$: return false if $e_C(P_{BC}, Q_{BC}) \neq e_C(P_C, Q_C)^B$: return false return true

Figure 6: The Kube-ZKP protocol. Notation used is defined in Subsection 2.7.

$$\begin{array}{ccc}
E_0 \times E_{AB} & \xrightarrow{\Phi} & E_A \times E_B \\
\downarrow \left(\begin{smallmatrix} \rho_0 & 0 \\ 0 & \rho_1 \end{smallmatrix} \right) = \Psi & & \downarrow \Psi' \\
F_0 \times F_{AB} & \xrightarrow{\Phi'} & \mathcal{A}
\end{array}$$

Figure 7: Diagram underlying Kani-Diamond-ZKP.

Cube-ZKP and $p = (A + B)C - 1 \approx \max(A, B)^3$ in Kube-ZKP). From a performance perspective, this suggests a slight favoring to Cube-ZKP at the expense of a larger proof size (since more repetitions are needed). All of this suggests that replicating the Kube-ZKP idea in the original setting of Cube-ZKP would not be advantageous in any way due to the much larger C . This implies a larger prime $p = ABC - 1 \approx \max(A, B)^4$ which in addition to the larger C -isogenies would override the soundness gain.

5 Kani-Diamond-ZKP construction

In this section, we present a novel zero-knowledge proof for orthogonal isogeny diamonds using higher dimensional isogenies. As in the previous protocols, we fix a prime p , coprime integers A and B , and an (A, B) -isogeny diamond containing the curves E_0, E_A, E_B , and E_{AB} . With additional work Kani-Diamond-ZKP can be adapted to handle minimal isogenies and PPAVs of arbitrary dimension. However, these extensions are beyond the scope of this paper.

5.1 Protocol overview

The main idea of the protocol is to use Kani's lemma (Theorem 1) to reduce the problem of proving an (A, B) -isogeny diamond to an equivalent two-dimensional statement. Namely, we prove knowledge of a non-diagonal, maximally isotropic (N, N) -isogeny $\Phi : E_0 \times E_{AB} \rightarrow E_A \times E_B$ where $A + B = N$. In terms of one-dimensional isogenies, we prove the relation

$$\mathcal{R}_{\text{Kani}} = \left\{ ((E_0, E_A, E_B, E_{AB}, N), (\phi_A, \phi_B)) \left| \begin{array}{l} \phi_A : E_0 \rightarrow E_A \text{ is an } A\text{-isogeny,} \\ \phi_B : E_0 \rightarrow E_B \text{ is a } B\text{-isogeny,} \\ [\phi_B]_* \phi_A : E_B \rightarrow E_{AB}, \\ A + B = N. \end{array} \right. \right\}.$$

This relation is notably different from Cube-ZKP since it only proves knowledge of some (a, b) -isogeny such that $a + b = N$. The statement only contains N , not A nor B .

A zero-knowledge proof of knowledge of an (N, N) -isogeny Φ with the correct domain and codomain is insufficient. In particular, one also must show that this isogeny is non-diagonal and maximally isotropic. We achieve this through a two-dimensional adaptation of the protocol of [BFGP23, Section 5], replacing the auxiliary isogeny with a diagonal (M, M) -isogeny, see Figure 7. Since pushforwards preserve diagonality (Lemma 1), we can prove that Φ is non-diagonal by proving that its pushforward is non-diagonal.

The method used by the prover to generate Φ depends on the available information and on the factorization of $p - 1$. Since the verifier must compute Ψ, Φ' , and Ψ' from their kernel generators, it is necessary that both N and M are smooth divisors of $p - 1$.

Commitment: The prover randomly samples two cyclic M -isogenies $\rho_0 : E_0 \rightarrow F_0$ and $\rho_1 : E_{AB} \rightarrow F_{AB}$, which define the diagonal (M, M) -isogeny

$$\Psi : E_0 \times E_{AB} \rightarrow F_0 \times F_{AB}.$$

The prover then computes a random kernel $\ker(\Phi) = \langle K_0, K_1 \rangle$ and its image $K'_0 = \Psi(K_0)$ and $K'_1 = \Psi(K_1)$, so that $\langle K'_0, K'_1 \rangle$ is the kernel of an (N, N) -isogeny

$$\Phi' : F_0 \times F_{AB} \rightarrow \mathcal{A}.$$

Next, the prover samples a random basis $\ker(\hat{\Psi}) = \langle T_0, T_1 \rangle$ as well as a random basis $\langle Q_0, Q_1, Q_2, Q_3 \rangle = (F_0 \times F_{AB})[M]$ and obtains the relations of the kernel in terms of the random basis $T_0 = \sum_{i=0}^3 \alpha_i Q_i$, and $T_1 = \sum_{i=0}^3 \beta_i Q_i$. The prover then computes

$$U_i = \Phi'(Q_i) \in \mathcal{A}, \quad \text{for } i = 0, 1, 2, 3.$$

Finally, the prover uses a computationally binding and statistically hiding commitment scheme H along with random nonces $r_{01}, r_{02}, r_{12}, r_2$ to partition the information in their computed square into

$$\begin{aligned} C_{01} &= H(r_{01}, \alpha_i, \beta_i), & C_{02} &= H(r_{02}, F_0, F_{AB}, Q_i), \\ C_{12} &= H(r_{12}, \mathcal{A}, U_i), & C_2 &= H(r_2, K'_0, K'_1). \end{aligned}$$

Their commitment is then $C = (C_{01}, C_{02}, C_{12}, C_2)$.

Challenge: The verifier challenges the prover with $\text{chal} = 0, 1, 2$.

Response and verification: If $\text{chal} = 0$, the prover reveals the data describing the diagonal isogeny $\hat{\Psi}$: the two dimensional variety $F_0 \times F_{AB}$, the basis Q_i , and the coefficients α_i, β_i . The verifier recomputes and verifies the committed hashes, reconstructs $\ker(\hat{\Psi})$ through the points $T_0 = \sum_i \alpha_i Q_i$ and $T_1 = \sum_i \beta_i Q_i$, and checks that this gives a diagonal M -isogeny back to $E_0 \times E_{AB}$.

If $\text{chal} = 1$, the prover reveals the data describing the pushforward isogeny $\hat{\Psi}'$: the domain $\mathcal{A}$, the basis U_i , and the coefficients α_i, β_i . The verifier checks the commitments, reconstructs the points $V_0 = \sum_i \alpha_i U_i$ and $V_1 = \sum_i \beta_i U_i$, and verifies that the resulting (M, M) -isogeny

$$\hat{\Psi}' = \text{Isogeny}(\langle V_0, V_1 \rangle)$$

has codomain $E_A \times E_B$.

If $\text{chal} = 2$, the prover reveals the data describing the lower isogeny Φ' along with the torsion points. In particular, they reveal $F_0, F_{AB}, \mathcal{A}$, the kernel generators K'_0, K'_1 , the basis Q_i , and their images U_i . The verifier first checks the relevant hashes and verifies that the Q_i 's and U_i 's form bases of the relevant M -torsion groups. Then, they compute the (N, N) -isogeny with kernel $\langle K'_0, K'_1 \rangle$ and check that it has the correct codomain $\mathcal{A}$. The verifier also checks that the kernel is maximally isotropic. If $\mathcal{A}$ splits as a product of elliptic curves, they also check non-diagonality using Lemma 2. Finally, they check that $\Phi'(Q_i) = U_i$ for all i .

5.2 Security Analysis of Kani-Diamond-ZKP

The correctness of Kani-Diamond-ZKP also follows from our construction. The only subtlety here is in challenge 2 and verifying the isogeny Φ' in the rare case that Φ' is reducible, but we handle this case explicitly utilizing Lemma 2.

Proposition 7. *Given a computationally binding commitment scheme H , the Kani-Diamond-ZKP protocol is 3-sound.*

Proof. Assume that an adversary successfully responds to all challenges given a fixed commitment. By the binding property of the hash function, such a prover must be able to produce isogenies Ψ , Ψ' , and Φ' such that $\widehat{\Psi'} = [\Phi']_* \widehat{\Psi}$. Such an adversary can form $\Theta := [\widehat{\Psi}]_* \Phi'$, which we claim is a maximally isotropic N -isogeny that is reducible and non-diagonal, proving our claim. We get the first property by checking the degree of Ψ and Φ' are correct. For reducibility, we use the fact that $\widehat{\Psi'} = [\Phi']_* \widehat{\Psi}$ to show we have an (N, M) -diamond with the correct domains, formed by taking push forwards of $\widehat{\Psi}$ and Φ' . This implies that Θ is reducible, with codomain $E_A \times E_B$. As for the non-diagonal property, it follows immediately from the Lemma 2. $\square$

We now come to the zero-knowledge property of the protocol, by constructing simulators for Ψ , Φ' , and Ψ' , treating each separately. The diagonal isogeny Ψ is simulated exactly as in an honest execution by sampling random M -isogenies from E_0 and E_{AB} . It remains to simulate Φ' and Ψ' , beginning with Φ' .

Simulator for Φ' . For the second challenge, fix the unknown isogeny

$$\tau := \psi_{AB} \circ \phi_A : E_0 \rightarrow E_{AB}.$$

The simulator must generate random generators of

$$\ker(\Phi') = \{(\rho_0([A]P), (\rho_1 \circ \tau)(P)) : P \in E_0[N]\},$$

where $\rho_0 : E_0 \rightarrow F_0$ and $\rho_1 : E_{AB} \rightarrow F_{AB}$ are uniformly sampled M -isogenies.

Although the simulator can sample ρ_0 and ρ_1 , it cannot generally compute the action of the unknown AB -isogeny τ on $E_0[N]$. We therefore assume that the image of a basis under $\rho_1 \circ \tau$ is indistinguishable from a random basis satisfying the required Weil-pairing relation. We show that this yields an efficient simulator under DLIT (Definition 8).

Concretely, our simulator for Φ' executes as follows. First, they randomly sample two cyclic M -isogenies $\rho_0 : E_0 \rightarrow F_0$ and $\rho_1 : E_{AB} \rightarrow F_{AB}$ and randomly samples bases $\langle P_0, Q_0 \rangle = E_0[N]$ and $\langle P_{AB}, Q_{AB} \rangle = F_{AB}[N]$ such that

$$e_N(P_0, Q_0) = e_N(P_{AB}, Q_{AB})^{-M}$$

which is efficient since N is smooth. They then output

$$(F_0 \times F_{AB}, (\rho_0(P_0), P_{AB}), (\rho_0(Q_0), Q_{AB})).$$

When compared to an honest response to $\text{chal} = 2$:

$$(F_0 \times F_{AB}, (\rho_0([A]P_0), (\rho_1 \circ \psi_{AB} \circ \phi_A)(P_0)), (\rho_0([A]Q_0), (\rho_1 \circ \psi_{AB} \circ \phi_A)(Q_0))).$$

we note that the pairing condition on P_{AB}, Q_{AB} is exactly the Weil-pairing relation induced by an honest response.

Remark 5. Indeed, if an MAB -isogeny $\rho_1 \circ \psi_{AB} \circ \phi_A : E_0 \rightarrow F_{AB}$ maps the relevant torsion bases as in the protocol, then $e_N([A]P_0, [A]Q_0) = e_N(P_{AB}, Q_{AB})^{MAB}$. Since $B \equiv -A \pmod{N}$, this gives

$$e_N(P_0, Q_0)^{A^2} = e_N(P_{AB}, Q_{AB})^{-MA^2},$$

which is equivalent to the simulator's constraint.

Note that this suffices to prove indistinguishability even when the adversary is given $\psi_{AB} \circ \phi_A$.

Theorem 2. Fix $\tau := \psi_{AB} \circ \phi_A : E_0 \rightarrow E_{AB}$. Any distinguisher between the honest and simulated challenge-2 transcript distributions, given τ , yields a distinguisher for the DLIT instance (E_{AB}, P, Q) of degree M .

Proof. Assume $\mathcal{D}$ distinguishes the honest and simulated transcript distributions with advantage ε . Suppose we are given a DLIT challenge (E', R, S) .

The reduction samples $\rho_0 : E_0 \rightarrow F_0$ and computes

$$P_0 = [(AB)^{-1}] \hat{\tau}(P), \quad Q_0 = [(AB)^{-1}] \hat{\tau}(Q),$$

and gives $\mathcal{D}$ the transcript $(F_0 \times E', (\rho_0([A]P_0), R), (\rho_0([A]Q_0), S))$ and outputs the bit returned by $\mathcal{D}$.

If the DLIT challenge is from $\mathcal{D}_0$, then $R = \rho_1(P)$ and $S = \rho_1(Q)$ for a random M -isogeny $\rho_1 : E_{AB} \rightarrow E'$, so the transcript is distributed as an honest response. If the challenge is from $\mathcal{D}_1$, then (R, S) is a random pairing-compatible basis on $E'[N]$. Moreover the codomain of E' is the codomain of some random cyclic M -isogeny, which agrees with the distribution of F_{AB} produced by the simulator. $\square$

Remark 6. Standard LIT requires $M \approx N^2 2^{2\lambda}$, but, in our setting, heuristically this may be relaxed to $M^2 \approx N^2 2^{2\lambda}$, since masking is provided jointly by the independent M -isogenies ρ_0 and ρ_1 . Indeed, if $(X, Y) = (\rho_0([A]P_0), \rho_1 \circ \psi_{AB} \circ \phi_A(P_0))$, then $[B^{-1}M^{-1}]\rho_0 \circ \widehat{\phi}_A \circ \widehat{\psi}_{AB} \circ \widehat{\rho}_1(Y) = X$, and similarly for Q_0 . The resulting isogeny has degree $M^2 AB$. Since its middle component $\widehat{\phi}_A \circ \widehat{\psi}_{AB}$ is fixed, this is not a genuine DLIT instance, but it suggests that the standard requirement may be stronger than necessary. Under this heuristic, one may take $M \approx N 2^\lambda$, yielding smaller parameters and thus higher efficiency.

Simulator for Ψ' . We next simulate Ψ' , the dual of the response to $\text{chal} = 1$. Unlike the previous case, we cannot reduce the simulator's indistinguishability to a known hard problem and therefore introduce an ad-hoc computational assumption.

The simulator must sample the pushforwards of uniformly random diagonal isogenies from $\Phi(\text{Diag})$ where

$$\text{Diag} = \{\ker \Psi \subset (E_0 \times E_{AB})[M] \mid \Psi \text{ is diagonal}\}.$$

Little is known about the image of this set under an unknown maximally isotropic, non-diagonal isogeny Φ . Although $\Phi(\text{Diag})$ is a 2-dimensional subspace the three-dimensional space of maximally isotropic subgroups of $(E_A \times E_B)[M]$, we know no efficient method to distinguish between the two. Our simulator relies on this admittedly ad-hoc assumption.

Consider the prover's algorithm to generate Ψ' . First, they sample random points $P \in E_0[M]$ and $Q \in E_{AB}[M]$ of full order. Then, they sample $\alpha, \beta \in \mathbb{Z}/M\mathbb{Z}$ such that $\gcd(\beta - \alpha, M) = 1$ and $\alpha \neq \beta$. These two conditions guarantee that $\langle (P, Q), (\alpha P, \beta Q) \rangle$ is the basis of a uniformly random diagonal (M, M) -isogeny.

To obtain a basis $\langle (R, S), (R', S') \rangle$ of $\ker \Psi'$, the honest prover pushes (P, Q) and $(\alpha P, \beta Q)$ forward through Φ . Writing Φ as a linear map in terms of its one-dimensional isogenies, we can explicitly compute the resulting relation between (R, S) and (R', S') .

$$\begin{pmatrix} R' \\ S' \end{pmatrix} = \Phi \begin{pmatrix} \alpha & 0 \\ 0 & \beta \end{pmatrix} \begin{pmatrix} P \\ Q \end{pmatrix} = \begin{pmatrix} N^{-1}(\alpha A + \beta B) & N^{-1}(\alpha - \beta)\phi_A \widehat{\phi}_B \\ N^{-1}(\alpha - \beta)\phi_B \widehat{\phi}_A & N^{-1}(\alpha B + \beta A) \end{pmatrix} \begin{pmatrix} R \\ S \end{pmatrix} \quad (1)$$

where we are using $\Phi = \begin{pmatrix} \phi_A & \widehat{\psi}_{AB} \\ \phi_B & -\widehat{\psi}_{BA} \end{pmatrix}$ and $\widehat{\Phi} = \begin{pmatrix} \widehat{\phi}_A & \widehat{\phi}_B \\ \psi_{AB} & -\psi_{BA} \end{pmatrix}$.

We can do a variable change $(\alpha_1, \alpha_2) \leftarrow N^{-1}(\alpha A + \beta B, \alpha B + \beta A)$ and $\beta \leftarrow N^{-1}(\alpha - \beta)$ to get

$$\begin{pmatrix} R' \\ S' \end{pmatrix} = \begin{pmatrix} \alpha_1 & \beta \phi_A \widehat{\phi}_B \\ \beta \phi_B \widehat{\phi}_A & \alpha_2 \end{pmatrix} \begin{pmatrix} R \\ S \end{pmatrix} = \begin{pmatrix} \alpha_1 R + \beta \phi_A \widehat{\phi}_B S \\ \beta \phi_B \widehat{\phi}_A R + \alpha_2 S \end{pmatrix} = \begin{pmatrix} \alpha_1 R \\ \alpha_2 S \end{pmatrix} + \beta \begin{pmatrix} \phi_A \widehat{\phi}_B S \\ \phi_B \widehat{\phi}_A R \end{pmatrix} \quad (2)$$

Without knowledge of a connecting isogeny $\phi_A \widehat{\phi_B}$, it is unclear how the simulator can sample such R, S, R', S' from the correct distribution. We therefore give the simulator access to an auxiliary hint oracle. The idea of using hint oracles is not unusual; for example, SQIsign assumes a hint oracle to construct large-prime-degree isogenies without knowing the curve's endomorphism ring [ABDF⁺25]. On each query, the oracle samples fresh $R \in E_A[M]$, $S \in E_B[M]$, and $\alpha, \beta \in \mathbb{Z}/M\mathbb{Z}$ with $\gcd(\alpha - \beta, M) = 1$, and returns (R, S) and (R', S') , where

$$\begin{pmatrix} R' \\ S' \end{pmatrix} = N^{-1} \begin{pmatrix} \alpha A + \beta B & (\alpha - \beta) \phi_A \widehat{\phi_B} \\ (\alpha - \beta) \phi_B \widehat{\phi_A} & \alpha B + \beta A \end{pmatrix} \begin{pmatrix} R \\ S \end{pmatrix}.$$

Thus, it returns the kernel of a uniformly random element of $\Phi(\text{Diag})$, allowing the simulator to simulate Ψ' by outputting the isogeny with kernel $((R, S), (R', S'))$.

It remains to argue that this oracle does not trivialize security. Let $f := \phi_B \widehat{\phi_A} : E_A \rightarrow E_B$ be the secret degree- AB bridge isogeny. The oracle never reveals clean samples of $f(R)$ or $\widehat{f}(S)$; rather, they are hidden by fresh masks as $R' = \alpha_1 R + \beta \widehat{f}(S)$ and $S' = \beta f(R) + \alpha_2 S$. We formalize the assumed hardness of recovering f from such samples as a computational problem, after recording several useful observations.

Let P, Q be a fixed public basis of $E_A[M]$, chosen independently of the secret data. For each oracle sample, write $R = \delta_0 P + \delta_1 Q$, where the adversary can recover $\delta_0, \delta_1 \in \mathbb{Z}/M\mathbb{Z}$, for example using Weil pairings.

Assuming $\gcd(M, AB) = 1$, the restriction $f : E_A[M] \rightarrow E_B[M]$ is an isomorphism, so $f(P), f(Q)$ form a basis of $E_B[M]$. Thus, for unknown $\gamma_0, \gamma_1 \in \mathbb{Z}/M\mathbb{Z}$, we may write $S = \gamma_0 f(P) + \gamma_1 f(Q)$. Using $\widehat{f} \circ f = [AB]$, the oracle output satisfies

$$\begin{aligned} R' &= (\alpha_1 \delta_0 + \beta AB \gamma_0) P + (\alpha_1 \delta_1 + \beta AB \gamma_1) Q, \\ S' &= f((\beta \delta_0 + \alpha_2 \gamma_0) P + (\beta \delta_1 + \alpha_2 \gamma_1) Q). \end{aligned} \tag{3}$$

Hence, the adversary knows $A, B, P, Q, \delta_0, \delta_1$, but not the masking coefficients $\alpha_1, \alpha_2, \beta$, the hidden coordinates γ_0, γ_1 , or the isogeny f .

Since $R' \in E_A[M]$, the adversary can express it in the public basis and recover $c_P = \alpha_1 \delta_0 + \beta AB \gamma_0$ and $c_Q = \alpha_1 \delta_1 + \beta AB \gamma_1$. These provide only two equations in the three unknown quantities $\alpha_1, \beta \gamma_0, \beta \gamma_1$, and therefore do not reveal the masks. Moreover, S' is expressed in the hidden basis $(f(P), f(Q))$, so it does not yield a pair $(X, f(X))$ for known $X \in E_A[M]$.

Across queries, all masks and sampled coordinates are fresh, whereas P, Q, A, B , and f remain fixed. Since we can simulate Ψ' given the hints from the auxiliary oracle, we know that, modulo the oracle, no information is leaked. It remains only to assume that polynomially many such transcripts do not reveal f . We formalize this as a computational problem, subject to the relation $\alpha_1 - \alpha_2 = (A - B)\beta$ from Equation 1.

Definition 9 (Computational Kani-Bridge problem). Let E_A, E_B , the integers A, B, M, N , and a fixed public basis (P, Q) of $E_A[M]$ be given. Let $f : E_A \rightarrow E_B$ be a fixed secret bridge isogeny of degree AB , with $\gcd(M, AB) = 1$. The goal is to compute f given polynomially many samples of the following form:

- $\{((R_j, S_j), (R'_j, S'_j))\}_j$, where for every transcript j , the same secret isogeny f is used, but fresh scalars $\alpha_{1,j}, \alpha_{2,j}, \beta_j, \gamma_{0,j}, \gamma_{1,j}, \delta_{0,j}, \delta_{1,j} \in \mathbb{Z}/M\mathbb{Z}$ are sampled, subject to the relation $\alpha_{1,j} - \alpha_{2,j} = (A - B)\beta_j$. The input points are

$$R_j = \delta_{0,j} P + \delta_{1,j} Q, \quad S_j = \gamma_{0,j} f(P) + \gamma_{1,j} f(Q),$$

and the corresponding output points are

$$\begin{pmatrix} R'_j \\ S'_j \end{pmatrix} = \begin{pmatrix} (\alpha_{1,j} \delta_{0,j} + \beta_j AB \gamma_{0,j}) P + (\alpha_{1,j} \delta_{1,j} + \beta_j AB \gamma_{1,j}) Q \\ f((\beta_j \delta_{0,j} + \alpha_{2,j} \gamma_{0,j}) P + (\beta_j \delta_{1,j} + \alpha_{2,j} \gamma_{1,j}) Q) \end{pmatrix}.$$

One instance of the computational Kani-Bridge problem appears harder than the underlying FESTA recovery problem. In addition to the β' -masking in (2), each transcript contains the independent additive masks $[\alpha_1]R$ and $[\alpha_2]S$. Consequently, each transcript is more under-determined than the corresponding FESTA instance. We further assume that polynomially many such instances do not enable recovery of f .

Theorem 3. *Given $\text{poly}(\log p)$ many honest transcripts of Kani-Diamond-ZKP, computing f is equivalent to solving the computational Kani-Bridge problem.*

6 Comparison and Implementation

Implementation. We provide implementations of Cube-ZKP, Kube-ZKP, Kani-Diamond-ZKP, as well as a proof-of-concept version of Windmill-ZKP.⁵ Isogeny diamonds may be generated either from rational one-dimensional isogenies or from a rational kernel of a two-dimensional product isogeny using algorithms from SQIsign2D² [XLZO26]. We note that Qlapoti [BCRSE⁺26] is not directly applicable to our parameter choices, since it requires the two-dimensional $(A + B, A + B)$ -isogeny to have $A + B \approx p$. The two-dimensional diamond generation is natural for Kube-ZKP and Kani-Diamond-ZKP, while Cube-ZKP and Windmill-ZKP are preferable when rational one-dimensional kernels are required.

In our SageMath [The26] implementation of Cube-ZKP, with NIST level 1 parameters, the prover’s algorithm takes about 3.5 minutes and the verifiers algorithm about 1 minute after 576 repetitions on an M4 Macbook Pro and produces an output of about 1 MB. The biggest bottleneck for all of our constructions, in terms of time and proof size, is the number of repetitions needed to make up for the soundness loss.

Parameter	Cube	Kube	Kani	Kani-Heuristic
$p =$	$2^{e_2} 3^{e_3} 5^{e_5} - 1$	$2^{e_2} 3^{e_3} - 1$	$2^{e_2} 3^{e_3} - 1$	$2^{e_2} 3^{e_3} - 1$
	—	$e_3 \approx 2e_2$	$3e_2 \approx e_3$	$3e_2 \approx 2e_3$
Prime size	6λ	6λ	8λ	5λ
Soundness error	$6/7$	$4/5$	$2/3$	$2/3$
Repetitions	4.50λ	3.11λ	1.71λ	1.71λ

Figure 8: Suggested parameters, prime size, per-round soundness error, and number of repetitions required to achieve λ -bit security.

Noting that Kani-Heuristic is based on Remark 6, Figure 8 summarizes the resulting trade-offs. Among the constructions for rational two-dimensional diamonds, Kube-ZKP retains the same DSSP-based security guarantee as Cube-ZKP and achieves the smallest estimated proof size. Kani-Heuristic instead uses the smallest underlying prime field and is 3-sound, making it the fastest construction in our estimates for both the prover and verifier. Its proofs are smaller than those of Cube-ZKP and the full Kani-Diamond-ZKP construction, although larger than those of Kube-ZKP. These efficiency gains come at the cost of the ad hoc assumption introduced in Remark 6. Calculations supporting the entries in Figure 8 are provided in `estimated_isogs.py` in our GitHub repository. Our parameter estimates use the well-known fact that the best known attack on DSSP (Definition 7) runs in time approximately equal to the square root of the isogeny degree [DFJP14, DFFP24, CDFG⁺26].

⁵Available at <https://github.com/Nicswansonwaterloo/DiamondZKPs>.

Computation	Cube	Kube	Kani	Kani-Heuristic
Prover				
2-isogenies	$8.99\lambda^2$	—	—	—
3-isogenies	$11.35\lambda^2$	$15.68\lambda^2$	$12.94\lambda^2$	$6.47\lambda^2$
5-isogenies	$3.87\lambda^2$	—	—	—
(2, 2)-isogenies	—	$6.21\lambda^2$	$3.42\lambda^2$	$3.42\lambda^2$
Verifier				
2-isogenies	$1.28\lambda^2$	—	—	—
3-isogenies	$1.62\lambda^2$	$6.27\lambda^2$	$4.31\lambda^2$	$2.16\lambda^2$
5-isogenies	$2.21\lambda^2$	—	—	—
(2, 2)-isogenies	—	$1.24\lambda^2$	$1.14\lambda^2$	$1.14\lambda^2$
(3, 3)-isogenies	—	—	$2.16\lambda^2$	$1.08\lambda^2$
Proof size	$66.48\lambda^2$ bytes	$42.71\lambda^2$ bytes	$75.65\lambda^2$ bytes	$53.14\lambda^2$ bytes

Figure 9: Average proof size and number of isogeny computations required to prove a relation in zero knowledge at a security level of λ bits, after accounting for the number of repetitions.

Comparison with existing literature. Concurrent work by Lai and Maino [LM26] develops general-purpose R1CS encodings for computing and evaluating non-smooth isogenies of degree $q(2^e - q)$. These relations are related to, but distinct from, the one considered here: our protocols prove knowledge of a commutative isogeny diamond involving four elliptic curves. Their hash-based zk-SNARK construction achieves proofs of about 450 KB at NIST level I, substantially smaller than our estimates of approximately 1.12 MB for Kani-Diamond-ZKP. The comparison therefore highlights a trade-off in approach: their construction arithmetizes isogeny computations within a general-purpose proof system, whereas our protocols are isogeny-native and have transcripts that directly reflect the four-curve diamond structure, but rely on stronger, isogeny-specific assumptions and incur greater communication costs.

The large soundness errors of our protocols are not unusually large compared to existing proofs over curves with unknown endomorphism rings. For example, proofs of knowledge of an isogeny based on DSSP achieve 3-soundness [DFDGZ22]. By contrast, Kube-ZKP proves the stronger statement of four commuting isogenies under the same assumption with 5-soundness. It is interesting to see if the soundness errors of Kube-ZKP and Cube-ZKP can be enhanced while retaining security based on DSSP.

References

- [ABDF⁺25] Marius A. Aardal, Andrea Basso, Luca De Feo, Sikhar Patranabis, and Benjamin Wesolowski. A complete security proof of SQIsign. In Y. Tautman Kalai and S. F. Kamara, editors, *Advances in Cryptology – CRYPTO 2025*, volume 16005 of *Lecture Notes in Computer Science*. Springer, Cham, 2025.
- [Bas24] Andrea Basso. A Post-Quantum Round-Optimal Oblivious PRF from Isogenies. In *Selected Areas in Cryptography – SAC 2023*, volume 14201 of *Lecture Notes in Computer Science*, pages 147–168. Springer, 2024.
- [BBC⁺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*, volume 15676 of *Lecture Notes in Computer Science*, pages 300–332. Springer, 2025.
- [BCC⁺23] Andrea Basso, Giulio Codogni, Deirdre Connolly, Luca De Feo, Tako Boris Fouotsa, Guido Maria Lido, Travis Morrison, Lorenz Panny, Sikhar Patranabis, and Benjamin Wesolowski. Supersingular Curves You Can Trust. In *Advances in Cryptology – EUROCRYPT 2023*, volume 14005 of *Lecture Notes in Computer Science*, pages 405–437. Springer, 2023.
- [BCRSE⁺26] Giacomo Borin, Maria Corte-Real Santos, Jonathan Komada Eriksen, Riccardo Invernizzi, Marzio Mula, Sina Schaeffler, and Frederik Vercauteren. Qlapoti: Simple and efficient translation of quaternion ideals to isogenies. In *Advances in Cryptology – ASIACRYPT 2025*, pages 174–205, Singapore, 2026. Springer.
- [BDF⁺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*, volume 15486 of *Lecture Notes in Computer Science*, pages 339–370. Springer, 2024.
- [BF23] Andrea Basso and Tako Boris Fouotsa. New SIDH Countermeasures for a More Efficient Key Exchange. In *Advances in Cryptology – ASIACRYPT 2023*, volume 14445 of *Lecture Notes in Computer Science*, pages 208–233. Springer, 2023.
- [BFGP23] Ward Beullens, Luca De Feo, Steven D. Galbraith, and Christophe Petit. Proving knowledge of isogenies: a survey. *Designs, Codes and Cryptography*, 2023. doi:10.1007/s10623-023-01243-3.
- [BM25] Andrea Basso and Luciano Maino. POKÉ: A Compact and Efficient PKE from Higher-Dimensional Isogenies. In *Advances in Cryptology – EUROCRYPT 2025*, volume 15602 of *Lecture Notes in Computer Science*, pages 94–123. Springer, 2025.
- [BMP23] Andrea Basso, Luciano Maino, and Giacomo Pope. FESTA: Fast Encryption from Supersingular Torsion Attacks. In *Advances in Cryptology – ASIACRYPT 2023*, volume 14444 of *Lecture Notes in Computer Science*, pages 98–126. Springer, 2023.

- [CD23] Wouter Castryck and Thomas Decru. An Efficient Key Recovery Attack on SIDH. In *Advances in Cryptology – EUROCRYPT 2023*, volume 14008 of *Lecture Notes in Computer Science*, pages 423–447. Springer, 2023.
- [CDFG⁺26] Wouter Castryck, Luca De Feo, Steven D. Galbraith, Péter Kutas, Krijn Reijnders, and Benjamin Wesolowski. The Isogeny Problems. Cryptology ePrint Archive, Paper 2026/1431, 2026. URL: <https://eprint.iacr.org/2026/1431>.
- [CL26] Giulio Codogni and Guido Maria Lido. Spectral theory of isogeny graphs. *Journal of Number Theory*, 286:131–184, 2026.
- [CLL23] Kelong Cong, Yi-Fu Lai, and Shai Levin. Efficient Isogeny Proofs Using Generic Techniques. In *Applied Cryptography and Network Security – ACNS 2023*, volume 13906 of *Lecture Notes in Computer Science*, pages 248–275. Springer, 2023.
- [CLP24] Mingjie Chen, Antonin Leroux, and Lorenz Panny. SCALLOP-HD: Group Action from 2-Dimensional Isogenies. In *Public-Key Cryptography – PKC 2024*, volume 14603 of *Lecture Notes in Computer Science*, pages 190–216. Springer, 2024.
- [CRSCS25] Maria Corte-Real Santos, Craig Costello, and Benjamin Smith. Efficient $(3, 3)$ -isogenies on fast Kummer surfaces. *Research in Number Theory*, 11(1):25, 2025.
- [CV23] Wouter Castryck and Frederik Vercauteren. A Polynomial Time Attack on Instances of M-SIDH and FESTA. In *Advances in Cryptology – ASIACRYPT 2023*, volume 14444 of *Lecture Notes in Computer Science*, pages 127–156. Springer, 2023.
- [Dar25] Pierrick Dartois. Fast computation of 2-isogenies in dimension 4 and cryptographic applications. *Journal of Algebra*, 683:449–514, 2025.
- [DEF⁺25] Pierrick Dartois, Jonathan Komada Eriksen, Tako Boris Fouotsa, Arthur Herlédan Le Merdy, Riccardo Invernizzi, Damien Robert, Ryan Rueger, Frederik Vercauteren, and Benjamin Wesolowski. PEGASIS: Practical Effective Class Group Action using 4-Dimensional Isogenies. In *Advances in Cryptology – CRYPTO 2025*, volume 16000 of *Lecture Notes in Computer Science*, pages 67–99. Springer, 2025.
- [DFDGZ22] Luca De Feo, Samuel Dobson, Steven D. Galbraith, and Lukas Zobernig. SIDH Proof of Knowledge. In *Advances in Cryptology – ASIACRYPT 2022*, volume 13792 of *Lecture Notes in Computer Science*, pages 310–339. Springer, 2022.
- [DFFP24] Luca De Feo, Tako Boris Fouotsa, and Lorenz Panny. Isogeny problems with level structure. In *Advances in Cryptology – EUROCRYPT 2024*, volume 14657 of *Lecture Notes in Computer Science*, pages 181–204. Springer, 2024.
- [DFJP14] Luca De Feo, David Jao, and Jérôme Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. *Journal of Mathematical Cryptology*, 8(3):209–247, 2014.
- [DFKL⁺20] Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski. SQISign: Compact Post-quantum Signatures from Quaternions and Isogenies. In *Advances in Cryptology – ASIACRYPT 2020*, volume 12491 of *Lecture Notes in Computer Science*, pages 64–93. Springer, 2020.

- [dHKM⁺25] Thomas den Hollander, Søren Kleine, Marzio Mula, Daniel Slamanig, and Sebastian A. Spindler. More Efficient Isogeny Proofs of Knowledge via Canonical Modular Polynomials. In *Advances in Cryptology – CRYPTO 2025*, volume 16000 of *Lecture Notes in Computer Science*, pages 131–166. Springer, 2025.
- [dHMSS26] Thomas den Hollander, Marzio Mula, Daniel Slamanig, and Sebastian A. Spindler. On the Use of Atkin and Weber Modular Polynomials in Isogeny Proofs of Knowledge. In *Post-Quantum Cryptography – PQCrypto 2026*, volume 16492 of *Lecture Notes in Computer Science*, pages 39–73. Springer, 2026.
- [DK26] Thomas Decru and Sabrina Kunzweiler. Abelian surfaces in Hesse form and explicit isogeny formulas, 2026. URL: <https://arxiv.org/abs/2601.05922>, [arXiv:2601.05922](https://arxiv.org/abs/2601.05922).
- [DLRW24] Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski. SQIsignHD: New Dimensions in Cryptography. In *Advances in Cryptology – EUROCRYPT 2024*, volume 14651 of *Lecture Notes in Computer Science*, pages 3–32. Springer, 2024.
- [DMPR24] Pierrick Dartois, Luciano Maino, Giacomo Pope, and Damien Robert. An Algorithmic Approach to $(2, 2)$ -Isogenies in the Theta Model and Applications to Isogeny-Based Cryptography. In *Advances in Cryptology – ASIACRYPT 2024*, volume 15486 of *Lecture Notes in Computer Science*, pages 304–338. Springer, 2024.
- [EVdGM12] Bas Edixhoven, Gerard Van der Geer, and Ben Moonen. Abelian varieties. *preprint*, 7(8), 2012.
- [FKT18] Satoshi Furukawa, Noboru Kunihiro, and Katsuyuki Takashima. Multi-party key exchange protocols from supersingular isogenies. In *International Symposium on Information Theory and Its Applications*, pages 208–212. IEEE, 2018.
- [FMP23] Tako Boris Fouotsa, Tomoki Moriya, and Christophe Petit. M-SIDH and MD-SIDH: Countering SIDH Attacks by Masking Information. In *Advances in Cryptology – EUROCRYPT 2023*, volume 14008 of *Lecture Notes in Computer Science*, pages 282–309. Springer, 2023.
- [GPS20] Steven D. Galbraith, Christophe Petit, and Javier Silva. Identification Protocols and Signature Schemes Based on Supersingular Isogeny Problems. *Journal of Cryptology*, 33:130–175, 2020.
- [Kan97] Ernst Kani. The number of curves of genus two with elliptic differentials. *Journal für die reine und angewandte Mathematik*, 485:93–121, 1997.
- [LM26] Yi-Fu Lai and Luciano Maino. Toward zkSNARK-assisted isogeny-based cryptography. Cryptology ePrint Archive, Paper 2026/1096, 2026. URL: <https://eprint.iacr.org/2026/1096>.
- [LP25] Shai Levin and Robi Pedersen. Faster Proofs and VRFs from Isogenies. In *Advances in Cryptology – ASIACRYPT 2025*, volume 16248 of *Lecture Notes in Computer Science*, pages 307–340. Springer, 2025.

- [LR15] David Lubicz and Damien Robert. A generalisation of Miller’s algorithm and applications to pairing computations on abelian varieties. *Journal of Symbolic Computation*, 67:68–92, 2015. URL: <https://www.sciencedirect.com/science/article/pii/S0747717114000510>, doi:10.1016/j.jsc.2014.08.001.
- [Mil86] James S. Milne. Abelian varieties. In *Arithmetic Geometry*, pages 103–150. Springer, New York, NY, 1986. URL: [10.1007/978-1-4613-8655-1_5](https://doi.org/10.1007/978-1-4613-8655-1_5), doi:10.1007/978-1-4613-8655-1_5.
- [MJ23] Youcef Mokrani and David Jao. Zero-Knowledge Proofs for SIDH Variants with Masked Degree or Torsion. In *Security, Privacy, and Applied Cryptography Engineering – SPACE 2023*, volume 14412 of *Lecture Notes in Computer Science*. Springer, 2023.
- [MMP⁺23] Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, and Benjamin Wesolowski. A Direct Key Recovery Attack on SIDH. In *Advances in Cryptology – EUROCRYPT 2023*, volume 14008 of *Lecture Notes in Computer Science*, pages 448–471. Springer, 2023.
- [Mor23] Tomoki Moriya. IS-CUBE: An Isogeny-Based Compact KEM Using a Boxed SIDH Diagram. Cryptology ePrint Archive, Paper 2023/1506, 2023. URL: <https://eprint.iacr.org/2023/1506>.
- [MS26] Tomoki Moriya and Miha Stopar. LIT-SiGamal: An efficient isogeny-based PKE based on a LIT diagram. In *Advances in Cryptology – ASIACRYPT 2025*, volume 16248 of *Lecture Notes in Computer Science*, pages 275–306. Springer, 2026. doi:10.1007/978-981-95-5113-2_9.
- [NO24] Kohei Nakagawa and Hiroshi Onuki. QFESTA: Efficient Algorithms and Parameters for FESTA Using Quaternion Algebras. In *Advances in Cryptology – CRYPTO 2024*, volume 14924 of *Lecture Notes in Computer Science*, pages 75–106. Springer, 2024.
- [NOC⁺24] Kohei Nakagawa, Hiroshi Onuki, Wouter Castryck, Mingjie Chen, Riccardo Invernizzi, Gioella Lorenzon, and Frederik Vercauteren. SQIsign2D-East: A New Signature Scheme Using 2-Dimensional Isogenies. In *Advances in Cryptology – ASIACRYPT 2024*, volume 15486 of *Lecture Notes in Computer Science*, pages 272–303. Springer, 2024.
- [Rob22] Damien Robert. Reducible gluing of abelian varieties. <https://www.math.u-bordeaux.fr/~damienrobert/pro/publications/notes/2022-09-kani.pdf>, 2022. Notes, Université de Bordeaux.
- [Rob23] Damien Robert. Breaking SIDH in Polynomial Time. In *Advances in Cryptology – EUROCRYPT 2023*, volume 14008 of *Lecture Notes in Computer Science*, pages 472–503. Springer, 2023.
- [Rob24] Damien Robert. Some notes on algorithms for abelian varieties. Cryptology ePrint Archive, Paper 2024/406, 2024. URL: <https://eprint.iacr.org/2024/406>.
- [The26] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 10.9)*, 2026. URL: <https://www.sagemath.org>.

- [XLZO26] Zheng Xu, Kaizhan Lin, Chang-An Zhao, and Yi Ouyang. SQIsign2D²: New SQIsign2D Variant by Leveraging Power Smooth Isogenies in Dimension One. In *Advances in Cryptology – ASIACRYPT 2026*, volume 16248 of *Lecture Notes in Computer Science*, pages 341–371. Springer, 2026.