Hyperelliptic Gluing Isogeny Diffie–Hellman (HGIDH): A Genus-2 Gluing Isogeny Key-Exchange

Nouhou Abdou Idris[1][0009-0007-8323-9955] and Mustapha Hedabou[2][0000-0002-7872-4469]

[1] Mohammed VI Polytechnic University, Benguerir, Morocco
Abdou.idris-NOUHOU@um6p.ma
[2] Mohammed VI Polytechnic University, Benguerir, Morocco
Mustapha.HEDABOU@um6p.ma

Abstract. In this work, we propose *Hyperelliptic Gluing Isogeny Diffie–Hellman* (HGIDH), a key-exchange protocol built from gluing isogenies between the product of two supersingular elliptic curves and the Jacobian of a genus-2 hyperelliptic curve. The protocol leverages the Frey–Kani correspondence, using maximal isotropic subgroups of $(E_1 \times E_2)[N]$ to construct principally polarized abelian surfaces. Private keys are encoded as four scalars defining a non-cyclic, two-dimensional kernel, thereby avoiding the structural weaknesses exploited in SIDH-style attacks.
We formalize the computational tasks underlying attacks on HGIDH through two intermediate problem formulations, which abstract the recovery of gluing kernels and the computation of genus-2 isogenies. We show that any efficient adversary solving these problems can be transformed into an algorithm solving standard supersingular isogeny problems, situating the security of HGIDH within the established hardness landscape. Furthermore, we analyze the resistance of the construction to known classical and quantum attacks, including torsion-point attacks and Costello–Smith-style meet-in-the-middle strategies.

Keywords: Isogeny-based cryptography · Hyperelliptic curves · Jacobian varieties, Gluing isogenies · Castryck–Decru attack resistance · Costello–Smith attack resistance.

1 Introduction

Isogeny-based cryptography has emerged as a compelling candidate for post-quantum public-key cryptography, offering compact keys and deep connections to arithmetic geometry. Since the introduction of the Supersingular Isogeny Diffie–Hellman (SIDH) protocol, isogeny-based constructions have attracted significant attention as alternatives to lattice- and code-based schemes. However,

the discovery of efficient key-recovery attacks by Castryck–Decru and subsequent refinements has demonstrated that SIDH and related schemes suffer from inherent structural vulnerabilities rooted in their reliance on one-dimensional cyclic kernels.

These attacks have reshaped the design landscape of isogeny-based cryptography. Rather than incremental patching, they suggest a fundamental limitation of elliptic-curve-based constructions that expose torsion-point images under secret isogenies. As a result, recent research has increasingly focused on higher-dimensional isogeny frameworks, particularly those involving abelian surfaces and Jacobians of genus-2 curves, where the kernel structure is inherently richer and less amenable to known attack techniques.

1.1 Our Contributions

In this work, we introduce *Hyperelliptic Gluing Isogeny Diffie–Hellman* (HGIDH), a novel isogeny-based key-exchange protocol that realizes the Frey–Kani gluing correspondence as a concrete cryptographic primitive. The protocol operates on a fixed product of supersingular elliptic curves $E_1 \times E_2$ and derives public keys as Jacobians of genus-2 hyperelliptic curves obtained via secret gluing isogenies. Each private key consists of four scalars defining a maximal isotropic, non-cyclic subgroup of $(E_1 \times E_2)[N]$, yielding a fundamentally different kernel structure from prior cyclic-kernel constructions.

The main contributions of this work are summarized as follows:

- **Formal protocol design.** We provide the first complete specification of a key-exchange protocol built on genus-2 gluing isogenies $\varphi : E_1 \times E_2 \to \mathrm{Jac}(C)$, including explicit algorithms for key generation, public-key representation, and shared-secret derivation. The construction is parameterized to meet NIST security levels I, III, and V.
- **Hardness framework for gluing isogenies.** To analyze the security of HGIDH, we introduce two computational problem formulations—the *Generalized Genus-2 Isogeny Hardness* (GG2IH) problem and the *Gluing Construction One-Wayness* (GCOW) problem—which abstract the tasks faced by an adversary attempting to recover secret gluing kernels or explicitly compute genus-2 isogenies.
- **Reductions to standard isogeny assumptions.** We present explicit polynomial-time reductions showing that any efficient algorithm solving GG2IH or GCOW can be transformed into an algorithm solving the well-established *Supersingular Isogeny* (SSI) or *Computational Supersingular Isogeny* (CSSI) problems.
- **Security analysis against known attacks.** We formalize a security game for HGIDH and prove that any efficient adversary breaking the protocol can be leveraged to solve GG2IH. Owing to its non-cyclic, maximal isotropic kernel structure, HGIDH is inherently resistant to torsion-point and oracle-based attacks, including GPST, Castryck–Decru, and Costello–Smith-style meet-in-the-middle attacks on superspecial isogeny graphs.

– **Implementation and performance evaluation.** We present a reference implementation in SageMath 10.6 demonstrating that HGIDH is computationally practical.

1.2 Organization

The remainder of the paper is organized as follows. Section 2 reviews background on Kani construction of gluing isogenies. Sections 3 describes the HGIDH protocol in detail. Section 4 and 5 analyze its security and formalize the associated computational problems. The sections 7 and 6 establish the resistance against the Castryck–Decru and Costello–Smith attacks respectively. Section 8 presents implementation results and performance benchmarks.

1.3 Related Work

Isogeny-based cryptography builds on the presumed hardness of finding isogenies between supersingular elliptic curves. The Supersingular Isogeny Diffie–Hellman (SIDH) protocol by Jao and De Feo [JDF11] pioneered this area and led to the SIKE proposal [CLN20]. However, the 2022 Castryck–Decru and GPST attacks [CD22] broke SIDH by exploiting its one-dimensional cyclic kernel, motivating research toward more complex structures.

Commutative and oriented frameworks. CSIDH [CLM+18] and its extensions [MKS21] rely on commutative class-group actions over $\mathbb{F}_p$, offering simple arithmetic but slower performance and weaker security assumptions. OSIDH [DFKPW20] uses oriented isogenies with complex multiplication but depends on delicate endomorphism data, limiting scalability.

Higher-dimensional constructions. To overcome SIDH's weaknesses, recent work explores higher-dimensional isogenies using abelian surfaces and Jacobians of genus-2 curves.

Extending these concepts to higher genus abelian varieties, particularly Jacobians of hyperelliptic curves ($g \geq 2$), presents a new design space.

Several proposals have explored this dimension:

– *Higher-Dimensional Analogues of SIDH:* Early attempts by Flynn and Ti [FT16] and subsequent work by Costello [Cos18] proposed genus-2 variants of the SIDH key exchange, primarily using isogenies between Jacobians ($Jac(C_1) \to Jac(C_2)$) or products of elliptic curves ($E_1 \times E_2 \to E_3 \times E_4$). These typically rely on the hardness of finding a path in the genus-2 supersingular isogeny graph.
– *Isogeny-Based Hash Functions:* Cryptographic primitives such as the Castryck-Decru-Smith (CDS) hash function [CDS19]and related constructions have utilized isogenies on supersingular Jacobians.

The *HGIDH protocol* distinguishes itself fundamentally from these existing $g = 2$ schemes. While others focus on the standard isogeny problem between

two Jacobians, *HGIDH is constructed directly from the Frey-Kani* gluing isogeny *correspondence:*

$$\gamma : E_1 \times E_2 \longrightarrow S \simeq Jac(C)$$

This unique structure, which maps a *product of elliptic curves* to a *principally polarized Jacobian variety* ($Jac(C)$) via a secret *maximal isotropic kernel*.

2 Preliminaries

2.1 Hyperelliptic Curves and Jacobian Varieties

Definition 1 (Hyperelliptic Curve). *[Was08,HSS23,Ric36,Fly02] A hyperelliptic curve C of genus g over a field K is given by*

$$y^2 + h(x)y = f(x),$$

where $f(x)$ is a monic polynomial of degree $2g+1$ and $h(x)$ is a polynomial of degree at most g. For cryptographic applications, we often work with the simplified form $y^2 = f(x)$.

Definition 2 (Jacobian Variety). *[Was08,HSS23,Ric36,Fly02] The Jacobian variety $Jac(C)$ of a curve C is an abelian variety of dimension g whose points correspond to degree-zero divisors modulo principal divisors. For cryptography, we focus on the group structure of $Jac(C)$.*

Jacobians of Hyperelliptic Curves The Jacobian $\mathrm{Jac}(C)$ is an abelian variety that serves as the group for cryptographic operations. Points on the Jacobian are represented by divisors:

$$D = \sum n_i[P_i] - g[\infty].$$

2.2 Mumford Representation

[Was08]

Each degree-0 divisor class in $\mathrm{Jac}(C)$ has a unique reduced representative expressed via Mumford coordinates $(u(x), v(x))$ satisfying:

1. $u(x)$ monic, $\deg u \leq g$,
2. $\deg v < \deg u$,
3. $u(x) \mid (v(x)^2 - f(x))$.

For example, for $D = [P_1] + [P_2] - 2[\infty]$, we have

$$u(x) = (x - x_{P_1})(x - x_{P_2}), \quad \text{identity: } (1, 0).$$

Theorem 1 (Mumford Correspondence). *[Was08] There is a one-to-one correspondence between divisor classes of degree 0 on C and pairs $(U(x), V(x))$ satisfying:*

1. *U is monic,*
2. *$\deg V < \deg U \leq g$,*
3. *$V^2 - f(x)$ is a multiple of U.*

2.3 Gluing Isogenies

Abelian surfaces (dimension 2) used in cryptography fall into two types:

- A product of elliptic curves $E_1 \times E_2$,
- The Jacobian $\mathrm{Jac}(\mathcal{C})$ of a genus-2 curve.

Four types of higher-dimensional isogenies exist between abelian surfaces:

- Splitting: $\phi : \mathrm{Jac}(\mathcal{C}) \to E_1 \times E_2$,
- Gluing: $\phi : E_1 \times E_2 \to \mathrm{Jac}(\mathcal{C})$,
- Product: $\phi : E_1 \times E_2 \to E_3 \times E_4$,
- Generic: $\phi : \mathrm{Jac}(\mathcal{C}) \to \mathrm{Jac}(\mathcal{C})$.

Theorem 2 (Gluing Isogeny Computation). *[HSS23] Let E_1, E_2 be elliptic curves over K, and let $G = \langle (P_1, P_2), (Q_1, Q_2) \rangle \subset (E_1 \times E_2)[n]$. With reduced Weierstrass forms:*

$$E_1 : y^2 = (x - \alpha_1)(x - \alpha_2)(x - \alpha_3), \quad E_2 : y^2 = (x - \beta_1)(x - \beta_2)(x - \beta_3),$$

the choice of G corresponds to choosing $(\alpha_1, \alpha_2, \alpha_3)$ and $(\beta_1, \beta_2, \beta_3)$. The gluing isogeny is then an isogeny

$$\gamma : E_1 \times E_2 \to \mathcal{S} \simeq Jac(H).$$

2.4 Kernel Definition via Anti-Symplectic Graphs

Let E_1, E_2 be supersingular elliptic curves over $\mathbb{F}_{p^2}$ and let $N = 2^a$ with $a \geq 1$. Fix public symplectic bases

$$(P_1, Q_1) \subset E_1[N], \quad (P_2, Q_2) \subset E_2[N],$$

with respect to the Weil pairing e_N.

Definition 3 (Anti-Symplectic Isomorphism [FK91]). *An isomorphism $\alpha : E_1[N] \to E_2[N]$ is* anti-symplectic *if for all $P, Q \in E_1[N]$,*

$$e_N(\alpha(P), \alpha(Q)) = e_N(P, Q)^{-1}.$$

Definition 4 (Kernel as Graph [FK91]). *Given an anti-symplectic isomorphism $\alpha : E_1[N] \to E_2[N]$, define the secret gluing kernel $\mathcal{K}_\alpha$ as the subgroup:*

$$\mathcal{K}_\alpha := \{(P, \alpha(P)) : P \in E_1[N]\} \subset (E_1 \times E_2)[N]. \tag{1}$$

Specifically, if α is defined by the matrix $M_\alpha = \begin{pmatrix} s_1 & s_3 \\ s_2 & s_4 \end{pmatrix}$ relative to the public bases $\mathcal{B}_1 = \{P_1, Q_1\}$ and $\mathcal{B}_2 = \{P_2, Q_2\}$, then $\mathcal{K}_\alpha$ is the N-torsion subgroup generated by:

$$\mathcal{K}_\alpha = \langle (P_1, s_1 P_2 + s_2 Q_2), (Q_1, s_3 P_2 + s_4 Q_2) \rangle. \tag{2}$$

Proposition 1 ([BL04]). *K_α is a maximal isotropic subgroup of $(E_1 \times E_2)[N]$ with respect to the product Weil pairing.*

Proof. For any $(P, \alpha(P)), (Q, \alpha(Q)) \in K_\alpha$,

$$e_N((P, \alpha(P)), (Q, \alpha(Q))) = e_N(P, Q) \cdot e_N(\alpha(P), \alpha(Q)) = 1.$$

Maximality follows from $\#K_\alpha = N^2$.

Let $A_0 := E_1 \times E_2$ and K_α as above.

Theorem 3 (Frey–Kani Gluing [BL04]). *If K_α is maximal isotropic, the quotient*

$$A_0/K_\alpha$$

is a principally polarized abelian surface isomorphic to the Jacobian $\mathrm{Jac}(C)$ of a genus-2 hyperelliptic curve C.

3 HGIDH Isogeny Key-Exchange Protocol

3.1 Setup and Public Parameters

The protocol parameters are selected based on the desired NIST security level (Level 1, 3, or 5).

1. **Base Field** ($\mathbb{F}_{p^2}$): Let p be a large prime such that $p \equiv 3 \pmod 4$. The base field is the quadratic extension $\mathbb{F}_{p^2} = \mathbb{F}_p(i)$, where $i^2 = -1$.
   The prime p is of the form $p = 2^a \cdot f - 1$, chosen to meet NIST security requirements:
   $$p \approx \begin{cases} 2^{434} & \text{for NIST Level I } (\lambda = 128) \\ 2^{610} & \text{for NIST Level III } (\lambda = 192) \\ 2^{751} & \text{for NIST Level V } (\lambda = 256) \end{cases}$$
2. **Isogeny Degree** (N): The isogeny degree is $N = 2^a$. To ensure the existence of the required 2^a-torsion points and to provide resistance against algebraic attacks on the SIDH-like structure, the exponent a is set as:
   $$a = \begin{cases} 430 & \text{if } \lambda = 128 \\ 600 & \text{if } \lambda = 192 \\ 740 & \text{if } \lambda = 256 \end{cases}$$
3. Elliptic Curves (E_1, E_2): The protocol uses two fixed supersingular elliptic curves defined over $\mathbb{F}_{p^2}$:
   $$E_1 : y^2 = x^3 + x, \quad E_2 : y^2 = x^3 + x.$$
4. Torsion Basis ($\mathcal{B}$): A public N-torsion basis for the product curve $E_1 \times E_2$ is fixed as
   $$\mathcal{B} = \{P_1, Q_1\} \subset E_1[N], \qquad \{P_2, Q_2\} \subset E_2[N],$$
   such that
   $$\langle P_1, Q_1, P_2, Q_2 \rangle = (E_1 \times E_2)[N] \cong (\mathbb{Z}/N\mathbb{Z})^4.$$

3.2 Key Generation

3.3 Key Generation

Key Description (Party $\mathcal{A}$) Party $\mathcal{A}$ generates a secret key consisting of four scalars and constructs a corresponding public key.

1. **Secret Key** ($\mathbf{s}_\mathcal{A}$): $\mathcal{A}$ chooses four random scalars:

   $$\mathbf{s}_\mathcal{A} = (s_1, s_2, s_3, s_4) \in (\mathbb{Z}/N\mathbb{Z})^4.$$

   These define an anti-symplectic isomorphism α by its action on the public bases:

   $$\alpha(P_1) = s_1P_2 + s_2Q_2, \qquad \alpha(Q_1) = s_3P_2 + s_4Q_2,$$

   and we require $\det \begin{pmatrix} s_1 & s_3 \\ s_2 & s_4 \end{pmatrix} \equiv -1 \pmod{N}$ (maximal isotropy condition).
2. **Gluing Kernel**: The secret gluing kernel is

   $$\mathcal{K}_\mathcal{A} = \langle (P_1, \alpha(P_1)),\ (Q_1, \alpha(Q_1)) \rangle \subset (E_1 \times E_2)[N].$$

   By construction $\#\mathcal{K}_\mathcal{A} = N^2$ and it is maximal isotropic.
3. **Secret Gluing Isogeny** ($\phi_\mathcal{A}$): The isogeny $\phi_\mathcal{A} : E_1 \times E_2 \longrightarrow \mathrm{Jac}(C_\mathcal{A})$ is the quotient map with kernel $\mathcal{K}_\mathcal{A}$.
4. **Public Key** ($\mathbf{PK}_\mathcal{A}$): The public key consists of the curve equation $h_\mathcal{A}(x)$ of $C_\mathcal{A}$ and the images of the public basis points in Mumford representation:

   $$D_{\mathcal{A},P} = \phi_\mathcal{A}(P_1, P_2), \quad D_{\mathcal{A},Q} = \phi_\mathcal{A}(Q_1, Q_2).$$

   Hence $\mathbf{PK}_\mathcal{A} = \big(h_\mathcal{A}(x),\ D_{\mathcal{A},P},\ D_{\mathcal{A},Q}\big)$.

Key Description (Party $\mathcal{B}$) Party $\mathcal{B}$ follows an identical procedure, using secret scalars $\mathbf{s}_\mathcal{B} = (t_1, t_2, t_3, t_4) \in (\mathbb{Z}/N\mathbb{Z})^4$ to generate kernel points $T_1 \in E_1[N]$ and $T_2 \in E_2[N]$, resulting in the public key:

$$\mathbf{PK}_\mathcal{B} = (h_\mathcal{B}(x), D_{\mathcal{B},P}, D_{\mathcal{B},Q})$$

3.4 Key Exchange Protocol

Shared Secret Computation The goal is for both parties to arrive at a shared divisor class on a Jacobian.

Alice's Computation: Alice uses her secret scalars $\mathbf{s}_\mathcal{A}$ and Bob's public key data defining the auxiliary isogeny.

1. **Secret Point Construction:** Alice computes her secret point in the product curve:

   $$R_\mathcal{A} = (R_1, R_2) = (s_1P_1 + s_2Q_1,\ s_3P_2 + s_4Q_2) \in E_1 \times E_2.$$

2. **Image under Bob's Public Isogeny Data:** Using Bob's public key data, she evaluates
$$D_{\text{shared}} = \phi_{\mathcal{B}}(R_{\mathcal{A}}) \in \text{Jac}(C_{\mathcal{B}}),$$
obtaining a divisor on Bob's curve.
3. **Shared Secret Derivation:** She hashes the serialized Mumford representation of D_{shared} to produce the shared key:
$$K = \text{Hash}(\mathsf{enc}(D_{\text{shared}})).$$

Bob's Computation: Bob uses his secret scalars $\mathbf{s}_{\mathcal{B}}$ and Alice's public key data.

1. **Secret Point Construction:** Bob computes his secret point:
$$R_{\mathcal{B}} = (T_1, T_2) = (t_1P_1 + t_2Q_1,\ t_3P_2 + t_4Q_2) \in E_1 \times E_2.$$
2. **Image under Alice's Public Isogeny Data:** Using Alice's public key data, he evaluates
$$D'_{\text{shared}} = \phi_{\mathcal{A}}(R_{\mathcal{B}}) \in \text{Jac}(C_{\mathcal{A}}).$$
3. **Shared Secret Derivation:** Bob hashes the serialized Mumford representation to obtain
$$K' = \text{Hash}(\mathsf{enc}(D'_{\text{shared}})).$$

Divisor Encoding and Hashing

For a shared divisor $D = (u(x), v(x))$, we use the *unique reduced Mumford representation* where $u(x)$ is monic and $\deg v < \deg u$. The bitstring $\mathsf{enc}(D)$ is formed by concatenating the coefficients of $u(x)$ and $v(x)$ in ascending order of degree. The shared secret is derived using an iterated hash:

$$K = \text{Hash}^{1000}\left(\mathsf{coeffs}(u) \parallel \mathsf{coeffs}(v)\right)$$

3.5 Divisor Arithmetic and Isogeny Computation

Arithmetic in $\text{Jac}(C)$ is performed using Cantor's algorithm:

1. **Divisor Addition**: $\texttt{affine_add}(h, D_1, D_2)$ computes the reduced sum of two divisors using extended GCD over the polynomial ring $\mathbb{F}_{p^2}[x]$.
2. **Divisor Doubling**: $\texttt{affine_dbl}(h, D)$ computes $[2]D$ efficiently.
3. **Gluing Isogeny**: The `FromProdToJac` function maps $E_1 \times E_2 \to \text{Jac}(C)$ by solving the gluing parameters s_1, s_2, α_i based on the x-coordinates of the 2-torsion kernel.

3.6 Protocol Specification and Validation

Secret-Key Constraints:

Secret-Key Constraints : To ensure the gluing isogeny ϕ targets a principally polarized Jacobian, the subgroup $\mathcal{K}_\alpha$ must be maximal isotropic with respect to the product Weil pairing e_{prod}. This is satisfied if and only if the mapping α is anti-symplectic.

Public Key Validation : A verifier should perform the following checks on a received $\mathrm{PK} = (h(x), D_P, D_Q)$:

1. **Curve Validity**: $h(x)$ must define a non-singular genus-2 curve ($\deg h \in \{5, 6\}$).
2. **Torsion Order**: D_P and D_Q must have exact order N in $\mathrm{Jac}(C)$.
3. **Independence**: $\langle D_P, D_Q \rangle \cong (\mathbb{Z}/N\mathbb{Z})^2$, ensuring the points generate a valid N-torsion subgroup.

Failure of any validation step results in the immediate rejection of the key exchange.

3.7 HGIDH Protocol Correctness

The correctness of the protocol lies on this commutativity reported in the following theorem:

Theorem 4 (HGIDH Protocol Correctness). *Let $\mathcal{B} = \{P_1, Q_1, P_2, Q_2\}$ be a fixed symplectic basis of $A_0[N]$ for $A_0 = E_1 \times E_2$. Let SK_A and SK_B define secret maximal isotropic subgroups $R_A, R_B \subset A_0[N]$, and let*

$$\phi_A : A_0 \to Jac(C_A), \qquad \phi_B : A_0 \to Jac(C_B)$$

be the corresponding Frey–Kani isogenies. Define

$$V_A = Jac(C_B)/\phi_B(R_A), \quad V_B = Jac(C_A)/\phi_A(R_B), \quad V \cong A_0/(R_A + R_B).$$

Then there exists a canonical isomorphism $\Psi : V_A \xrightarrow{\cong} V_B$ identifying both quotients with V. Let $\overline{D}_A$ and $\overline{D}_B$ be the images of $D_A = \phi_B(R_A)$ and $D_B = \phi_A(R_B)$ in V_A and V_B, respectively. Then

$$\overline{D}_A = \overline{D}_B \in V, \qquad \Psi(\overline{D}_A) = \overline{D}_B.$$

Both parties compute

$$S = \mathsf{enc}(\overline{D}_A) = \mathsf{enc}(\overline{D}_B), \qquad K = \mathsf{Hash}^{1000}(S),$$

yielding the same shared secret $K \in \{0,1\}^\lambda$.

Proof. The proof is divided into two parts: the algebraic commutativity of the underlying varieties and the consistency of the canonical representation.

Part 1: Commutativity of Isogeny Composition.

Proof of Isomorphism 1: $\mathbf{Jac}(\boldsymbol{C_B})/\boldsymbol{\phi_B}(\boldsymbol{R_A}) \cong \boldsymbol{A_0}/(\boldsymbol{R_A} + \boldsymbol{R_B})$

Step 1: Composite Map Construction Define the composite map:

$$\Phi_{BA} : A_0 \xrightarrow{\phi_B} \mathrm{Jac}(C_B) \xrightarrow{\pi_A} \mathrm{Jac}(C_B)/\phi_B(R_A)$$

where π_A is the canonical projection. Since both ϕ_B and π_A are surjective homomorphisms, Φ_{BA} is surjective.

Step 2: Kernel Computation By the First Isomorphism Theorem for abelian varieties :

$$\mathrm{Jac}(C_B)/\phi_B(R_A) \cong A_0/\ker(\Phi_{BA})$$

We compute the kernel:

$$\begin{aligned}\ker(\Phi_{BA}) &= \{R \in A_0 \mid \Phi_{BA}(R) = \mathrm{Id}\} \\ &= \{R \in A_0 \mid \pi_A(\phi_B(R)) = \mathrm{Id}\} \\ &= \{R \in A_0 \mid \phi_B(R) \in \phi_B(R_A)\}\end{aligned}$$

Since ϕ_B is a homomorphism, $\phi_B(R) \in \phi_B(R_A)$ implies there exists $K \in R_A$ such that:

$$\phi_B(R) = \phi_B(K) \Rightarrow \phi_B(R - K) = \mathrm{Id}_{\mathrm{Jac}(C_B)}$$

Thus $R - K \in \ker(\phi_B) = R_B$, so $R \in K + R_B \subset R_A + R_B$.

Conversely, if $R \in R_A + R_B$, then $R = R_A + R_B$ with $K_A \in R_A$, $K_B \in R_B$, and:

$$\phi_B(R) = \phi_B(K_A) + \phi_B(K_B) = \phi_B(K_A) \in \phi_B(R_A)$$

Therefore, $\ker(\Phi_{BA}) = R_A + R_B$.

Step 3: Isomorphism Conclusion Substituting into the isomorphism theorem:

$$\mathrm{Jac}(C_B)/\phi_B(R_A) \cong A_0/(R_A + R_B)$$

Proof of Isomorphism 2: $\mathbf{Jac}(\boldsymbol{C_A})/\boldsymbol{\phi_A}(\boldsymbol{R_B}) \cong \boldsymbol{A_0}/(\boldsymbol{R_A} + \boldsymbol{R_B})$ By symmetry, define the composite map:

$$\Phi_{AB} : A_0 \xrightarrow{\phi_A} \mathrm{Jac}(C_A) \xrightarrow{\pi_B} \mathrm{Jac}(C_A)/\phi_A(R_B)$$

Following identical reasoning:

$$\ker(\Phi_{AB}) = \{R \in A_0 \mid \phi_A(R) \in \phi_A(K_B)\} = R_B + R_A = R_A + R_B$$

Thus:

$$\mathrm{Jac}(C_A)/\phi_A(R_B) \cong A_0/(R_A + R_B)$$

which concludes the proof of the commutativity.

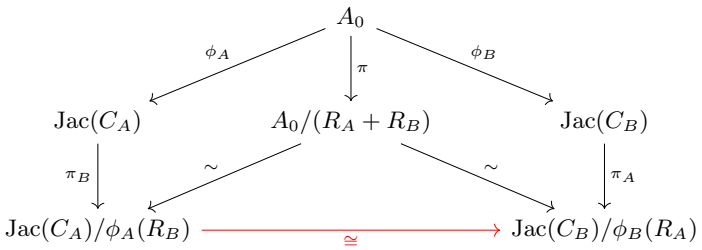

Fig. 1. Diagram illustrating the overall mappings and relations.

Part 2: Consistency of the Derived Divisor Class and Hashing. From Part 1, there is a canonical isomorphism

$$\Psi : \mathrm{Jac}(C_B)/\phi_B(R_A) \xrightarrow{\cong} \mathrm{Jac}(C_A)/\phi_A(R_B), \qquad V \cong A_0/(R_A + R_B).$$

Let $\overline{D}_A$ and $\overline{D}_B$ be the images of $D_A = \phi_B(R_A)$ and $D_B = \phi_A(R_B)$ in their respective quotients. By construction, they correspond to the same element in V, so

$$\mathsf{enc}(\overline{D}_A) = \mathsf{enc}(\overline{D}_B).$$

Applying the deterministic hash yields a shared secret:

$$K = \mathsf{Hash}^{1000}(\mathsf{enc}(\overline{D}_A)) = \mathsf{Hash}^{1000}(\mathsf{enc}(\overline{D}_B)) \in \{0,1\}^{\lambda}.$$

4 Security Analysis

This section provides formal security reductions establishing the foundation of HGIDH's security. We prove that breaking HGIDH requires solving either the Generalized Genus-2 Isogeny Hardness (GG2IH) or Gluing Construction One-Wayness (GCOW) problems, which in turn reduce to standard supersingular isogeny problems (SSI and CSSI). These reductions form the core of our security argument.

4.1 Standard Isogeny Hardness Problems

Definition 5 (Supersingular Isogeny (SSI) Problem). *Given two supersingular elliptic curves E_A, E_B over $\mathbb{F}_{p^2}$ and a smooth integer N, find an isogeny $\phi : E_A \to E_B$ of degree N.*

Definition 6 (Computational Supersingular Isogeny (CSSI) Problem). *Given a supersingular elliptic curve E, points $P, Q \in E[N]$, and the image $\phi(P), \phi(Q)$ where $\phi : E \to E/\langle R \rangle$, find the kernel generator $R = [m]P + [n]Q$.*

4.2 New Computational Problems

Assumption 1 (Generalized Genus 2 Isogeny Hardness (GG2IH)). *Let $g = 2$, p a prime, and $N = 2^a$. Given:*

- *a principally polarized abelian variety A of dimension g, and*
- *another abelian variety B such that there exists an isogeny $\varphi : A \to B$ of degree N^g,*

no probabilistic polynomial-time algorithm can compute any such isogeny with non-negligible probability.

Assumption 2 (Gluing Construction One-Wayness (GCOW)). *The function*

$$\begin{aligned} f : \quad (\mathbb{Z}/N\mathbb{Z})^4 \quad &\longrightarrow \{\textit{genus-2 Jacobians}\} \\ (s_1, s_2, s_3, s_4) &\mapsto \mathit{Jac}(C), \end{aligned}$$

denote the Frey–Kani gluing map. It takes kernel parameters (s_1, s_2, s_3, s_4) defining a maximal isotropic subgroup $K \subset (E_1 \times E_2)[N]$ and outputs the Jacobian $J = \mathit{Jac}(C)$. We assume that given J, no probabilistic polynomial-time algorithm can recover a valid description of K with non-negligible probability.

5 Security Reductions

5.1 Preliminary Lemmas

We begin with three technical lemmas that will be used repeatedly in the reductions.

Lemma 1 (Product Decomposition and Degree Accounting). *Let $A_0 = E_1 \times E_2$ and $B_0 = F_1 \times F_2$ be products of elliptic curves over $\mathbb{F}_{p^2}$, each equipped with the product principal polarization. Let $\Psi : A_0 \to B_0$ be an isogeny of degree D. Write Ψ as a 2×2 matrix of component homomorphisms:*

$$\Psi = \begin{pmatrix} \psi_{11} & \psi_{12} \\ \psi_{21} & \psi_{22} \end{pmatrix},$$

where $\psi_{ij} \in \mathrm{Hom}(E_j, F_i)$. Then:

$$\deg(\psi_{11}) + \deg(\psi_{21}) = D.$$

Proof. The dual isogeny $\widehat{\Psi} : B_0 \to A_0$ satisfies $\widehat{\Psi} \circ \Psi = [D]_{A_0}$. Under the product principal polarization, the Rosati involution acts factorwise. It follows that

$$\begin{pmatrix} \widehat{\psi}_{11} & \widehat{\psi}_{21} \\ \widehat{\psi}_{12} & \widehat{\psi}_{22} \end{pmatrix} \begin{pmatrix} \psi_{11} & \psi_{12} \\ \psi_{21} & \psi_{22} \end{pmatrix} = \begin{pmatrix} \widehat{\psi}_{11} \circ \psi_{11} + \widehat{\psi}_{21} \circ \psi_{21} & \widehat{\psi}_{11} \circ \psi_{12} + \widehat{\psi}_{21} \circ \psi_{22} \\ \widehat{\psi}_{12} \circ \psi_{11} + \widehat{\psi}_{22} \circ \psi_{21} & \widehat{\psi}_{12} \circ \psi_{12} + \widehat{\psi}_{22} \circ \psi_{22} \end{pmatrix}.$$

Examining the $(1,1)$-entry of the matrix representation yields:

$$\widehat{\psi}_{11} \circ \psi_{11} + \widehat{\psi}_{21} \circ \psi_{21} = [D]_{E_1}.$$

For elliptic curve isogenies, we have $\widehat{f} \circ f = [\deg f]$. Taking degrees on both sides gives:

$$\deg(\psi_{11}) + \deg(\psi_{21}) = D.$$

This completes the proof.

Lemma 2 (Extraction of N-Primary Subgroup). *Let $\psi : E_1 \to E_2$ be a nonzero separable isogeny over $\mathbb{F}_{p^2}$ of degree d, and suppose $N \mid d$ where $N = 2^a$ and $\gcd(N, p) = 1$. Then:*

1. *There exists a canonical subgroup $K[N] \subset \ker \psi$ of order exactly N (the N-Sylow subgroup of $\ker \psi$).*
2. *Given explicit rational maps for ψ, there exists a probabilistic polynomial-time algorithm that computes $K[N]$ and constructs the explicit degree-N isogeny $\psi_N : E_1 \to E_1/K[N]$ using Vélu's formulas.*

Proof. (1) Since $\gcd(N, p) = 1$, the N-torsion is étale. The kernel $\ker \psi$ is a finite abelian group of order d. As N is a power of 2 dividing d, the Sylow 2-subgroup of $\ker \psi$ contains a subgroup of order N, which we denote by $K[N]$.

(2) The algorithm proceeds as follows:

- Compute the d-division polynomial of E_1 and factor it over $\mathbb{F}_{p^2}$ to obtain the x-coordinates of points in $\ker \psi$.
- For each root, compute corresponding y-coordinates and full point set of $\ker \psi$.
- Perform group operations to identify the subgroup of order N (the N-torsion elements of $\ker \psi$).
- Apply Vélu's formulas to the points in $K[N]$ to compute the quotient isogeny $\psi_N : E_1 \to E_1/K[N]$ and the equation of the quotient curve.

All steps are polynomial-time in $\log p$ and $\log d$: polynomial factoring over finite fields is polynomial-time, group operations on elliptic curves are polynomial-time, and Vélu's formulas involve algebraic computations of degree $O(N)$.

Lemma 3 (Symmetry of Randomized Embedding). *Let E_1, E_2 be fixed supersingular elliptic curves over $\mathbb{F}_{p^2}$. Let E_1' be chosen uniformly at random from the set of supersingular elliptic curves over $\mathbb{F}_{p^2}$, and let $\alpha : E_2[N] \to E_1'[N]$ be a uniformly random anti-symplectic isomorphism. Define two distributions:*

$$\mathcal{D}_0 = \big(E_1 \times E_1',\ (E_2 \times E_1')/K_\alpha\big), \quad \mathcal{D}_1 = \big(E_1 \times E_1',\ (E_1' \times E_2)/K_{\alpha'}\big),$$

where $K_\alpha = \{(P, \alpha(P)) : P \in E_2[N]\}$ and $K_{\alpha'} = \{(\alpha^{-1}(Q), Q) : Q \in E_1'[N]\}$. Then $\mathcal{D}_0$ and $\mathcal{D}_1$ are identically distributed.

1. The Sampling Space. Let $\Omega = \mathcal{S} \times \mathrm{Isom}^{anti}$ be the sampling space. An element $\omega \in \Omega$ consists of the auxiliary curve E_1' and the isomorphism α. The measure on Ω is uniform. Let $A_0 = E_1 \times E_1'$ and let B_b be the codomain variety produced by the distribution $\mathcal{D}_b$ for $b \in \{0, 1\}$.

Isomorphism of Quotients: Consider the product varieties $A = E_2 \times E_1'$ and $B = E_1' \times E_2$. Define the swap isomorphism:

$$\sigma : A \to B, \quad (P, Q) \mapsto (Q, P)$$

The kernel $K_\alpha \subset A[N]$ is mapped by σ to:

$$\begin{aligned}\sigma(K_\alpha) &= \{\sigma(P, \alpha(P)) \mid P \in E_2[N]\} \\ &= \{(\alpha(P), P) \mid P \in E_2[N]\}\end{aligned}$$

By performing the change of variables $Q = \alpha(P)$, and noting that $P = \alpha^{-1}(Q)$, we observe that as P traverses $E_2[N]$, Q traverses $E_1'[N]$. Thus:

$$\tau(K_\alpha) = \{(Q, \alpha^{-1}(Q)) \mid Q \in E_1'[N]\} = K_{\alpha'}$$

where $\alpha' = \alpha^{-1}$. Since σ is an isomorphism that identifies the kernels K_α and $K_{\alpha'}$, it induces a canonical isomorphism between the quotient abelian surfaces:

$$(E_2 \times E_1')/K_\alpha \cong (E_1' \times E_2)/K_{\alpha'}$$

4. Measure Preservation. The distribution of the pair (E_1', α) is uniform over the space of all possible auxiliary curves and anti-symplectic identifications. The involution σ acts as a permutation on the set of possible kernels. Since the sampling of α is independent of the ordering of the elliptic curve factors in the Frey–Kani construction, and since E_1' is chosen uniformly from $\mathcal{S}$, the probability of sampling a specific isomorphism class of Jacobian J is:

$$\Pr[B_0 \cong J] = \sum_{E_1' \in \mathcal{S}} \frac{1}{|\mathcal{S}|} \cdot \frac{|\{\alpha \mid (E_2 \times E_1')/K_\alpha \cong J\}|}{|\mathrm{Isom}^{anti}|}.$$

By the symmetry established in Step 2, the number of α's that glue $E_2 \times E_1'$ into J is identical to the number of α's that glue $E_1' \times E_2$ into J. Thus, $\Pr[B_0 \cong J] = \Pr[B_1 \cong J]$. Then the statistical distance $\Delta(\mathcal{D}_0, \mathcal{D}_1)$ is $\mathsf{negl}(\lambda)$ that ends the proof.

5.2 Reduction from GG2IH to SSI

Theorem 5 (GG2IH $\geq_p$ SSI via gluing isogenies). *Let p be an odd prime, $N = 2^a$ with $\gcd(N, p) = 1$. If there exists a probabilistic polynomial-time adversary $\mathcal{A}$ that solves GG2IH with advantage ϵ, then there exists a probabilistic polynomial-time reduction $\mathcal{R}$ that solves the SSI problem with advantage*

$$\epsilon' \geq \tfrac{1}{2}\epsilon - \mathrm{negl}(\lambda).$$

Proof. We give a concrete probabilistic polynomial-time reduction $\mathcal{R}$ that uses any GG2IH solver $\mathcal{A}$ as a black box to solve an SSI instance. The proof has three parts: (A) instance construction, (B) extraction via matrix decomposition and Rosati involution, and (C) post-processing to obtain an N-isogeny and success-probability accounting.

A. Instance construction (embedding the SSI challenge). Let (E_1, E_2, N) be the SSI instance: supersingular elliptic curves over $\mathbb{F}_{p^2}$ and $N = 2^a$. The reduction proceeds:

1. Sample uniformly a random supersingular elliptic curve E_1' over $\mathbb{F}_{p^2}$ (independent of E_1, E_2).
2. Set the product abelian surface $A_0 = E_1 \times E_1'$.
3. Choose a uniform bit $b \in \{0, 1\}$.
4. Using the Frey–Kani (gluing) construction, build a principally polarized Jacobian $B = \mathrm{Jac}(C_b)$ as follows:
   - if $b = 0$, choose an anti-symplectic isomorphism $\alpha_0 : E_2[N] \xrightarrow{\sim} E_1'[N]$ and let $K_0 = \{(P, -\alpha_0(P)) : P \in E_2[N]\} \subset (E_2 \times E_1')[N]$; put $B = (E_2 \times E_1')/K_0$ and denote the natural gluing isogeny $\varphi_0 : E_2 \times E_1' \to B$ (degree N^2).
   - if $b = 1$, do the symmetric construction: choose $\alpha_1 : E_1'[N] \xrightarrow{\sim} E_2[N]$, put $K_1 = \{(P', -\alpha_1(P')) : P' \in E_1'[N]\}$ and $B = (E_1' \times E_2)/K_1$ with gluing isogeny $\varphi_1 : E_1' \times E_2 \to B$ (degree N^2).
5. Give the GG2IH instance (A_0, B, N) to the adversary $\mathcal{A}$.

Remark. By construction each φ_b is an (N, N)-gluing isogeny and B is a Jacobian. The reduction knows the anti-symplectic map used and hence can compute the dual $\widehat{\varphi_b}$ if needed.

B. Validity and decomposition via the dual gluing map. First, the instance is valid: because E_1 and E_2 are N-isogenous, there exists an isogeny $\Phi : A_0 \to B$ whose kernel is a suitable maximal isotropic subgroup of $A_0[N]$ lifting the gluing kernel K_b. (One may construct such an $H \subset A_0[N]$ by lifting a basis of K_b along the product isogeny; standard symplectic-linear-algebra arguments guarantee existence. The reduction need not compute Φ.)

Now suppose $\mathcal{A}$ returns an isogeny $\Phi : A_0 \to B$ of degree N^2 (the GG2IH solver's output). Form the composition with the dual gluing isogeny (which the reduction can compute):

$$\Psi := \widehat{\varphi_b} \circ \Phi : A_0 \longrightarrow \begin{cases} E_2 \times E_1' & (b = 0), \\ E_1' \times E_2 & (b = 1). \end{cases}$$

The map Ψ is an isogeny between product abelian varieties; write the target as $F_1 \times F_2$ with $(F_1, F_2) = (E_2, E_1')$ when $b = 0$ (and the symmetric ordering when $b = 1$). Because Ψ is a map between products of elliptic curves, it decomposes as a 2×2 matrix of component isogenies:

$$\Psi = \begin{pmatrix} \psi_{11} & \psi_{12} \\ \psi_{21} & \psi_{22} \end{pmatrix}, \qquad \psi_{ij} \in \mathrm{Hom}(E_i, F_j),$$

where E_1, E_2 denote the first and second factors of A_0 (i.e. E_1 and E_1' respectively).

C. Rosati relation and a degree identity. By multiplicativity of degrees, $\deg(\Psi) = \deg(\widehat{\varphi_b}) \cdot \deg(\Phi) = N^4$. Because Ψ is compatible with the product principal polarizations, the Rosati involution gives the matrix identity (on endomorphisms of $E_1 \times E_1'$)

$$\widehat{\Psi} \circ \Psi = [N^4]_{A_0},$$

and by By Lemma 1 the $(1,1)$ entry (the endomorphism of E_1) yields

$$\widehat{\psi_{11}} \circ \psi_{11} \; + \; \widehat{\psi_{21}} \circ \psi_{21} \; = \; [N^4]_{E_1}.$$

For an isogeny $f : E \to E'$ between elliptic curves (with $\gcd(\deg f, p) = 1$) one has the standard identity $\widehat{f} \circ f = [\deg f]_E$. Applying this to each term gives the equality of multiplication-by-integers

$$[\,\deg(\psi_{11})\,] + [\,\deg(\psi_{21})\,] = [\,N^4\,] \quad \text{as endomorphisms on } E_1,$$

hence the integer identity

$$\deg(\psi_{11}) + \deg(\psi_{21}) = N^4.$$

D. Extracting the relevant component. The component that maps E_1 into the actual E_2-factor of the target is

$$\psi' \; = \; \begin{cases} \psi_{11} & : E_1 \to E_2 \quad (b = 0), \\ \psi_{21} & : E_1 \to E_2 \quad (b = 1). \end{cases}$$

Because Ψ is non-zero, at least one of ψ_{11}, ψ_{21} is non-zero; therefore ψ' is non-zero for the correct choice of index. Crucially, the distribution of the challenge (A_0, B) is symmetric in the two gluings (the auxiliary curve E_1' was sampled uniformly and the anti-symplectic identification was chosen uniformly) according Lemma 3. Thus the adversary $\mathcal{A}$ receives an instance whose distribution does not reveal b, and consequently it cannot systematically force the relevant component to be zero for the chosen b. Therefore, conditioned on $\mathcal{A}$ succeeding on the GG2IH instance, the probability that the extracted component ψ' is non-zero is at least $1/2$ up to negligible statistical error:

$$\Pr[\psi' \neq 0 \mid \mathcal{A} \text{ succeeds}] \geq \tfrac{1}{2} - \mathrm{negl}(\lambda).$$

By the degree identity above, $\deg(\psi')$ divides N^4.

E. Post-processing: obtain an N-isogeny (Lemma 2). When $\psi' : E_1 \to E_2$ is a non-zero isogeny, let $d = \deg(\psi')$. Because E_1 and E_2 are N-isogenous (by the SSI instance), one has $N \mid d$; write $d = N^k \cdot m$ with m odd. The reduction now computes the finite kernel $K = \ker(\psi') \subset E_1[d]$ (polynomial time using standard torsion-point and Vélu machinery) and extracts the N-primary part

$$K[N] = \{P \in K : [N]P = 0\},$$

which is a subgroup of order N (exists since $N \mid d$ and N is a power of 2). Using Vélu formulas (polynomial-time for power-of-2 degrees in this setting) the reduction constructs the quotient isogeny

$$\psi : E_1 \longrightarrow E_1/K[N]$$

of degree N. Finally the reduction checks whether $E_1/K[N] \cong E_2$ (via comparing j-invariants); if so it composes with the isomorphism and outputs an explicit N-isogeny $E_1 \to E_2$. The probability that $E_1/K[N] \ncong E_2$ despite $N \mid d$ is negligible in the cryptographic parameter (this can be ruled out or detected by the verification step).

F. Success probability and running time. If $\mathcal{A}$ succeeds on the GG2IH instance with probability ϵ, then the probability that $\mathcal{A}$ both succeeds and the extracted component ψ' is non-zero is at least $\frac{1}{2}\epsilon - \mathrm{negl}(\lambda)$. The subsequent kernel extraction and Vélu computations succeed except with negligible probability. Therefore the reduction's advantage is

$$\epsilon' \;\geq\; \tfrac{1}{2}\epsilon - \mathrm{negl}(\lambda).$$

All steps (sampling E_1', building the gluing data, composing isogenies, computing torsion and Vélu maps) run in probabilistic polynomial time in the security parameter λ.

G. Conclusion. This gives the reduction $\mathcal{R}$ from GG2IH to SSI with the stated advantage loss, completing the proof.

5.3 Reduction from GCOW to CSSI

The second reduction shows that breaking the one-wayness of the gluing construction allows solving the CSSI problem. This reduction is tight.

Theorem 6 (Reduction of GCOW to CSSI). *The Gluing Construction One-Wayness (GCOW) problem is at least as hard as the Computational Supersingular Isogeny (CSSI) problem. Specifically, there exists a polynomial-time reduction $\mathcal{R}$ such that if an adversary $\mathcal{A}$ can solve the GCOW problem with advantage ϵ, then $\mathcal{R}$ can solve the CSSI problem with advantage $\epsilon' \geq \epsilon - negl(\lambda)$.*

Proof. We construct a tight reduction $\mathcal{R}$ that uses $\mathcal{A}$ as a black-box GCOW solver to solve a CSSI instance.

Phase 1: Embedding the CSSI Challenge

$\mathcal{R}$ receives a CSSI challenge:

$$(E_0, E_1, P, Q, P', Q', N),$$

where

– $\phi : E_0 \to E_1$ is a secret isogeny of degree N,
– $\{P, Q\}$ is a symplectic basis of $E_0[N]$ with $e_{E_0}(P, Q) = \zeta_N$,
– $P' = \phi(P)$ and $Q' = \phi(Q)$,
– the goal is to find a generator $R \in E_0[N]$ with $\ker(\phi) = \langle R \rangle$.

To construct a GCOW instance:

1. **Fix Target Bases:** Let $\{P_1, Q_1\}$ be the fixed public basis for $E_1[N]$ used in the HGIDH protocol.
2. **Construct Anti-Symplectic Map:** Define $\alpha : E_0[N] \to E_1[N]$ by

$$\alpha(P) = P', \quad \alpha(Q) = -Q'.$$

Justification: For $T_1 = a_1P + b_1Q$, $T_2 = a_2P + b_2Q$, the Weil pairing satisfies

$$e_{E_1}(\alpha(T_1), \alpha(T_2)) = e_{E_1}(a_1P' - b_1Q', a_2P' - b_2Q') = e_{E_0}(T_1, T_2)^{-1}.$$

Hence α is anti-symplectic.
3. **Construct the Jacobian:** Define

$$K_\alpha = \{(T, \alpha(T)) : T \in E_0[N]\} \subset (E_0 \times E_1)[N].$$

By the Frey-Kani theorem, compute the Jacobian

$$Jac(C) \cong (E_0 \times E_1)/K_\alpha$$

using Richelot isogenies.

Phase 2: Simulation for the Adversary

$\mathcal{R}$ provides $Jac(C)$ to $\mathcal{A}$ as the public key.

Indistinguishability: The distribution of $Jac(C)$ is identical to a real HGIDH public key, because E_0, E_1 are supersingular and K_α is maximal isotropic. Thus $\mathcal{A}$ succeeds with its claimed advantage ϵ.

Phase 3: Algebraic Extraction

If $\mathcal{A}$ succeeds, it returns the four scalars $(s_1, s_2, s_3, s_4) \in (\mathbb{Z}/N\mathbb{Z})^4$ that define the kernel K_α relative to the protocol's fixed bases. By the definition of the kernel as a graph, these scalars define the action of α as:

$$\begin{cases} \alpha(P) = [s_1]P_1 + [s_2]Q_1 \\ \alpha(Q) = [s_3]P_1 + [s_4]Q_1 \end{cases} \tag{3}$$

$\mathcal{R}$ now relates α back to the CSSI secret ϕ. Let the secret kernel generator be $R = [a]P + [b]Q$. By the definition of a kernel, $\phi(R) = 0_{E_1}$. Substituting R:

$$\phi([a]P + [b]Q) = [a]\phi(P) + [b]\phi(Q) = [a]P' + [b]Q' = 0 \tag{4}$$

Using the relations $P' = \alpha(P)$ and $Q' = [-1]\alpha(Q)$, we get:

$$[a]\alpha(P) - [b]\alpha(Q) = 0 \tag{5}$$

Substituting the scalars (s_1, s_2, s_3, s_4) provided by the adversary:

$$[a]([s_1]P_1 + [s_2]Q_1) - [b]([s_3]P_1 + [s_4]Q_1) = 0 \tag{6}$$

Since $\{P_1, Q_1\}$ is a basis, the coefficients of P_1 and Q_1 must independently vanish modulo N. This yields the linear system:

$$\begin{pmatrix} s_1 & -s_3 \\ s_2 & -s_4 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} \equiv \begin{pmatrix} 0 \\ 0 \end{pmatrix} \pmod{N} \tag{7}$$

Uniqueness: The matrix has rank 1, since $\ker(\phi)$ is cyclic of order N. A non-trivial solution (a, b) exists and is unique up to scalar multiplication modulo N. Solving the system recovers a generator of $\ker(\phi)$.

Phase 4: Success Probability, Efficiency, and Tightness

1. **Efficiency:** Constructing $Jac(C)$ takes $O(\text{poly}(\log p))$ time; solving the 2×2 system is $O(\text{poly}(\log N))$.
2. **Success Probability:** Every successful GCOW break yields the scalars needed. Only the negligible probability that Frey-Kani construction fails affects success: $\text{negl}(\lambda)$.
3. **Tightness:** No guessing or rewinding is needed. Thus $\epsilon' = \epsilon - \text{negl}(\lambda)$.

This concludes the proof, establishing that GCOW is at least as hard as CSSI.

Security Inheritance for HGIDH

We now present a single security theorem that combines the formal HGIDH security game with the concrete security inheritance bound, unifying the reductions to GG2IH, GCOW, SSI, and CSSI.

Theorem 7 (Security of HGIDH). *Let $\mathcal{A}$ be any probabilistic polynomial-time adversary attacking the HGIDH protocol in the security game $\mathsf{Game}^{\mathcal{A}}_{\mathsf{HGIDH}}(\lambda)$ with at most q concurrent sessions. Assume the hardness of the SSI and CSSI problems.*

Then there exist probabilistic polynomial-time algorithms $\mathcal{B}', \mathcal{C}'$ such that

$$\text{Adv}_{\text{HGIDH}}(\mathcal{A}) \;\leq\; 2\sqrt{q} \cdot \text{Adv}_{\text{SSI}}(\mathcal{B}') \;+\; q \cdot \text{Adv}_{\text{CSSI}}(\mathcal{C}') \;+\; \text{negl}(\lambda).$$

In particular, any efficient adversary breaking HGIDH yields an efficient algorithm for solving either the SSI or CSSI problem.

Proof. The proof proceeds via a sequence-of-games argument, followed by reductions to the underlying hardness assumptions.

Game 0 (Real Game). This is the original HGIDH security game $\mathsf{Game}^{\mathcal{A}}_{\mathsf{HGIDH}}(\lambda)$ as defined above, in which the adversary interacts with honestly generated public keys, isogeny oracles, and shared-secret computations.

Game 1 (Simulated Gluing Isogenies). We replace all secret gluing isogenies used in public keys and oracle responses with simulated isogenies that preserve the public Jacobian structure but hide the secret kernel. By the indistinguishability guaranteed by the GG2IH assumption, any PPT adversary can distinguish Game 1 from Game 0 with advantage at most

$$\mathrm{Adv}_{\mathrm{GG2IH}}(\mathcal{B}) + \mathrm{negl}(\lambda),$$

for some PPT reduction $\mathcal{B}$.

Game 2 (Randomized Shared Secrets). In this game, the shared secret returned by the Test oracle is replaced with a uniformly random value in the appropriate domain. Under the GCOW assumption, Game 2 is indistinguishable from Game 1 up to

$$\mathrm{Adv}_{\mathrm{GCOW}}(\mathcal{C}) + \mathrm{negl}(\lambda),$$

for some PPT reduction $\mathcal{C}$.

Hybrid Argument over Concurrent Sessions. Since the adversary may engage in at most q concurrent sessions, standard hybrid arguments imply:

- a $\sqrt{q}$ loss factor for distinguishing simulated gluing isogenies, and
- a linear q loss factor for replacing shared secrets.

Thus,

$$\mathrm{Adv}_{\mathrm{HGIDH}}(\mathcal{A}) \;\le\; \sqrt{q}\cdot \mathrm{Adv}_{\mathrm{GG2IH}}(\mathcal{B}) \;+\; q\cdot \mathrm{Adv}_{\mathrm{GCOW}}(\mathcal{C}) \;+\; \mathrm{negl}(\lambda).$$

Reduction to SSI and CSSI. By Theorems 5 and 6, there exist PPT algorithms $\mathcal{B}'$ and $\mathcal{C}'$ such that

$$\mathrm{Adv}_{\mathrm{GG2IH}}(\mathcal{B}) \le 2\cdot\mathrm{Adv}_{\mathrm{SSI}}(\mathcal{B}')+\mathrm{negl}(\lambda), \quad \mathrm{Adv}_{\mathrm{GCOW}}(\mathcal{C}) \le \mathrm{Adv}_{\mathrm{CSSI}}(\mathcal{C}')+\mathrm{negl}(\lambda).$$

Substituting these bounds yields the stated inequality.

Conclusion. Therefore, any PPT adversary that breaks the HGIDH protocol with non-negligible probability can be transformed into a PPT algorithm that solves either the SSI or CSSI problem with non-negligible probability, contradicting the assumed hardness of these problems.

6 Resistance to Robert–Castryck–Decru Attacks

This section analyzes the applicability of the Robert–Castryck–Decru (RCD) attack paradigm [CD22] [Rob23] to the HGIDH protocol. We show that the algebraic mechanisms exploited by RCD attacks in SIDH-style constructions do not extend to HGIDH, assuming the hardness of the Generalized Genus-2 Isogeny Hardness (GG2IH) problem.

6.1 The RCD Attack Paradigm

RCD attacks exploit a structural weakness present in SIDH and related protocols: the secret isogeny kernel is cyclic and depends linearly on a single secret scalar. This structure allows an adversary to construct a *reducibility oracle* that distinguishes whether a derived abelian surface splits as a product of elliptic curves. By querying this oracle on carefully modified instances, an adversary recovers the secret scalar one bit at a time.

Algebraic Mechanism In SIDH, the secret isogeny

$$\phi\colon E_0 \to E_A$$

has kernel $\langle P + [s]Q\rangle$ for a secret scalar $s \in \mathbb{Z}/2^a\mathbb{Z}$. For a suitable auxiliary endomorphism $\gamma \in \mathrm{End}(E_0)$, the abelian surface

$$X_{\phi,\gamma} = (E_0 \times E_0)/\{(x, \gamma(x)) : x \in \ker \phi\}$$

exhibits splitting behavior correlated with $s \bmod 2$. This yields a probabilistic oracle whose output is biased depending on a single bit of s, enabling full key recovery in polynomial time.

6.2 Structural Differences in HGIDH

We now identify structural properties of HGIDH that obstruct the construction of RCD-style oracles.

Multidimensional Secret Kernel In HGIDH, the secret isogeny kernel is not cyclic. Instead, it is defined by an anti-symplectic isomorphism

$$\alpha\colon E_1[N] \to E_2[N], \quad \alpha = \begin{pmatrix} s_1 & s_3 \\ s_2 & s_4 \end{pmatrix}, \quad \det(\alpha) \equiv -1 \pmod N.$$

The kernel takes the form

$$K_\alpha = \{(P, \alpha(P)) : P \in E_1[N]\} \subset (E_1 \times E_2)[N].$$

Lemma 4 (Isotropy Constraint). *The condition* $\det(\alpha) \equiv -1 \pmod N$ *is necessary and sufficient for* K_α *to be a maximal isotropic subgroup with respect to the product Weil pairing.*

This constraint couples all four secret scalars simultaneously. In contrast to SIDH, no secret parameter appears linearly or in isolation.

Absence of Bitwise Kernel Scaling RCD attacks rely on the ability to form scaled subgroups such as $\langle 2^{a-i}(P+[s]Q)\rangle$, whose structure depends on individual bits of s. In HGIDH, any subgroup derived from a partial or inconsistent guess of α fails to satisfy the isotropy condition of Lemma 4. As a result, such subgroups do not correspond to valid kernels of polarization-preserving isogenies accessible through the HGIDH public interface.

Geometric Sparsity of the Product Locus Let $\mathcal{A}_2$ denote the moduli space of principally polarized abelian surfaces. The locus of products of elliptic curves has strictly smaller dimension than $\mathcal{A}_2$. Consequently, for cryptographic parameters, a randomly distributed superspecial principally polarized abelian surface is isomorphic to a product of elliptic curves with negligible probability. Thus, random perturbations of HGIDH public keys do not yield meaningful splitting information.

6.3 Oracle Infeasibility

We formalize the obstruction to RCD-style attacks as the infeasibility of constructing a splitting oracle from HGIDH public data.

Theorem 8 (RCD Oracle Infeasibility under GG2IH). *Assume the GG2IH problem is hard. Then no probabilistic polynomial-time adversary with access only to the HGIDH public interface can construct a decision procedure that distinguishes, with non-negligible advantage, whether a polynomially describable modification of an HGIDH public key yields an abelian surface isomorphic to a product of elliptic curves.*

Proof. Suppose for contradiction that such an adversary $\mathcal{A}$ exists. We show how to use $\mathcal{A}$ to solve GG2IH.

Any RCD-style oracle requires the adversary to generate candidate kernels and to evaluate the resulting quotients for reducibility. In HGIDH, constructing such quotients requires computing isogenies relative to the secret kernel K_α. However, evaluating the isogeny on arbitrary torsion points is computationally equivalent to determining α, which directly solves the GG2IH problem.

Moreover, without satisfying the isotropy constraint of Lemma 4, candidate subgroups do not correspond to valid polarized quotients accessible from the public representation. Thus, any oracle capable of detecting reducibility with non-negligible bias can be transformed into an algorithm that distinguishes valid HGIDH kernels from random subgroups, contradicting GG2IH.

Therefore, no such oracle exists under the GG2IH assumption.

7 Resistance to the Costello–Smith Attack

Costello and Smith [CS19] showed that the superspecial isogeny path-finding problem in dimension $g \geq 2$ admits asymptotically faster algorithms than previously expected. Their attack exploits the existence of large, structured subgraphs corresponding to product abelian varieties. In this section, we show that although these results impact the *generic* isogeny path-finding problem, they do not yield an efficient attack against the HGIDH key-recovery problem.

7.1 The Costello–Smith Algorithm

Let $\mathcal{S}_g(p)$ denote the $(\ell, \dots, \ell)$-isogeny graph of superspecial principally polarized abelian varieties of dimension g over $\mathbb{F}_p$.

Theorem 9 (Costello–Smith [CS19]). *There exists a classical randomized algorithm which, given superspecial abelian varieties A_1, A_2 of dimension g, finds an isogeny path between them in expected time $\widetilde{O}(p^{g-1})$, succeeding with probability at least $2^{-(g-1)}$ as $p \to \infty$.*

For $g = 2$, this yields an expected $\widetilde{O}(p)$ algorithm, improving on the generic $\widetilde{O}(p^{3/2})$ bound. The speedup arises from the fact that $\mathcal{S}_2(p)$ contains a distinguished subgraph

$$\mathcal{S}_1(p) \times \mathcal{S}_1(p) \hookrightarrow \mathcal{S}_2(p),$$

consisting of products of elliptic curves. Random walks from arbitrary vertices intersect this product locus with non-negligible probability.

Crucially, the Costello–Smith algorithm assumes that the adversary:

1. can choose both endpoints of the path arbitrarily, and
2. can efficiently recognize when a vertex lies in the product locus.

7.2 Applicability to HGIDH

In HGIDH, the adversary is not given an arbitrary pair of vertices. Instead, the public key has the form

$$\mathrm{Jac}(C) \cong A_0/K_\alpha,$$

where $A_0 = E_1 \times E_2$ is a known product surface and K_α is a secret maximal isotropic subgroup defined by an anti-symplectic isomorphism α. The adversary's goal is to recover K_α, not merely to find *some* isogeny path between A_0 and $\mathrm{Jac}(C)$.

Polarization Constraints All isogenies in HGIDH preserve principal polarizations.

Lemma 5 (Polarization Preservation). *Let $\psi\colon A \to B$ be a $(2,2)$-isogeny between principally polarized abelian surfaces. Then $\ker(\psi)$ must be a maximal isotropic subgroup of $A[2]$ with respect to the Weil pairing.*

Proof. This follows from standard theory of polarized isogenies; see [BL04, Proposition 5.2.3].

Thus, any meet-in-the-middle search in the $(2,2)$-isogeny graph is restricted to kernels lying in the symplectic Grassmannian. Random walks that do not respect this constraint leave the space of valid HGIDH instances.

Codimension of the Product Locus Let $\mathcal{A}_2$ denote the moduli space of principally polarized abelian surfaces. The locus $\mathcal{L}_{\text{prod}}$ of products of elliptic curves has strictly smaller dimension than $\mathcal{A}_2$.

Proposition 2 (Geometric Sparsity). *The set of superspecial principally polarized abelian surfaces that are isomorphic to products of elliptic curves has negligible density as $p \to \infty$.*

As a consequence, random walks originating from a generic Jacobian remain in the Jacobian locus with overwhelming probability. The Costello–Smith shortcut relies on frequent transitions between the Jacobian and product loci, which are not available in this setting.

7.3 Resistance to Costello–Smith–Style Attacks

We now formalize the above discussion.

Theorem 10 (Resistance to Costello–Smith Attacks). *Let $PK = (h(x), D_P, D_Q)$ be an HGIDH public key with $Jac(C) \cong A_0/K_\alpha$. Assume the Generalized Genus-2 Isogeny Hardness (GG2IH) problem is hard. Then any probabilistic polynomial-time adversary (classical or quantum) attempting to recover K_α via a meet-in-the-middle search in the $(2,2)$-isogeny graph succeeds with at most negligible probability.*

Proof. Suppose an adversary $\mathcal{A}$ succeeds with non-negligible probability. Then $\mathcal{A}$ produces a collision

$$A_m \cong J_m$$

along paths originating from A_0 and $\mathrm{Jac}(C)$. This implicitly yields an isogeny

$$\Psi \colon A_0 \to \mathrm{Jac}(C).$$

Constructing such an isogeny for a fixed target $\mathrm{Jac}(C)$ solves an instance of the GG2IH problem. Moreover, recovering the specific kernel K_α from Ψ is at least as hard as the GCOW problem. By assumption, both problems are computationally infeasible.

Additionally, the probability that a random walk from $\mathrm{Jac}(C)$ intersects the product locus at a polarization-preserving node is negligible by Proposition 2. Therefore, the Costello–Smith speedup does not apply to the targeted kernel-recovery problem in HGIDH.

8 Implementation and Benchmarking

The performance figures reported in this section are derived from a prototype implementation of the Hyperelliptic Gluing Isogeny Diffie-Hellman (HGIDH) protocol, developed in **SageMath 10.6** and executed on a standard desktop workstation (Intel Core i5)

The following is the performances gotten from the implementation based on maximal isotropic subgroups are recorded on Table 1.

Table 1. Comprehensive Hyperelliptic Isogeny-Based Key Exchange Benchmark with NIST Security Levels and Actual Performance Metrics.

Metric	Level I (p434)	Level III (p610)	Level V (p751)	Unit
Security Parameters				
Target Security	128	192	256	bits
Prime p (actual)	434	610	751	bits
Isogeny Degree (2^a)	2^{430}	2^{600}	2^{740}	
Key Generation (Averaged)				
Setup (Torsion Basis)	1.2164	2.4681	7.3148	s
KeyGen Time	0.9820	1.3027	1.3106	s
Public Key Size	1320	2139	2629	bytes
Secret Key Size	215	300	370	bytes
Shared Secret Computation				
Derivation Time	0.0471	0.0835	0.1236	s
Success Rate	100.0%	100.0%	100.0%	
Comparative Analysis Summary				
Total Handshake Time*	2.2455	3.8543	8.7490	s
Communication Overhead	1535	2439	2999	bytes
Outcome	Key exchange succeeded (100% reliability)			

**Total time includes Setup + KeyGen + Derivation.*

9 Comparison of HGIDH and Genus-2 SIDH analogues

The table below contrasts HGIDH's unique structure with that of general Genus -2 SIDH analogues.

Table 2. Comparison of HGIDH and Genus-2 SIDH Analogues

Feature	HGIDH	Genus-2 SIDH Analogues
Core Isogeny Map	(Product-to-Jacobian): $\gamma : E_1 \times E_2 \longrightarrow \mathrm{Jac}(C)$	$\phi : \mathrm{Jac}(C_1) \longrightarrow \mathrm{Jac}(C_2)$ or $\phi : E_1 \times E_2 \longrightarrow E_3 \times E_4$ (Product-to-Product).
Secret Kernel	*Maximal Isotropic Subgroup* $\mathcal{K} \subset (E_1 \times E_2)[N]$, defined by four scalars (s_1, s_2, s_3, s_4).	Typically a *cyclic-kernel analogue* (similar to $g = 1$ SIDH) or a non-maximal isotropic subgroup.
Security Foundation	*Gluing Construction One-Wayness (GCOW)* .	*Computational Isogeny Problem (CISIP)* in genus 2.
Torsion Leakage	Does not require public transmission of basis images of the isogeny kernel itself, relying on the commutative diagram property.	Requires public transmission of torsion point images under the secret isogenies to facilitate the commutative diagram, a known point of vulnerability [DFJP14].

Conclusion

We have presented the *Hyperelliptic Gluing Isogeny Diffie–Hellman* (HGIDH) protocol, a novel post-quantum key-exchange construction based on the geometry of genus-2 abelian surfaces. HGIDH achieves provable resistance to Costello–Smith-style attacks and eliminates the structural weaknesses that compromised SIDH by replacing cyclic one-dimensional kernels with two-dimensional maximal isotropic subgroups. By leveraging the Frey–Kani gluing correspondence, the protocol ensures that the difficulty of recovering secret keys reduces directly to the Generalized Genus-2 Isogeny Hardness (GG2IH) and Gluing Construction One-Wayness (GCOW) assumptions.

Unlike prior genus-2 schemes that are often vulnerable to product-subvariety exploitation, HGIDH's framework guarantees that intermediate attacks either fail due to polarization mismatch or necessitate solving GG2IH or GCOW, both of which are shown to be at least as hard as standard supersingular isogeny problems. This provides a significant security advantage over competing higher-dimensional constructions that do not strictly enforce principal polarization.

From a practical perspective, our reference implementation in SageMath 10.6 demonstrates that the protocol is computationally practical.

Future research directions include:

– Developing highly optimized implementations over large-prime fields using specialized arithmetic (Montgomery or Karatsuba techniques...),
– Implementing compression techniques for Jacobian representations to further reduce public key overhead,
– Exploring the construction of hyperelliptic gluing-based digital signature schemes.

HGIDH thus establishes a new pathway for compact, efficient, and provably secure isogeny-based cryptosystems built on the rich arithmetic geometry of higher-dimensional abelian varieties, significantly advancing the post-quantum cryptography landscape.

A Appendix

Lemma 6 (Isotropy and Determinant Constraint). *Let* $\alpha : E_1[N] \to E_2[N]$ *be a homomorphism defined by the matrix* $M_\alpha = \begin{pmatrix} s_1 & s_3 \\ s_2 & s_4 \end{pmatrix}$ *relative to symplectic bases. The graph* $\mathcal{K}_\alpha = \{(P, \alpha(P)) : P \in E_1[N]\}$ *is a maximal isotropic subgroup of* $(E_1 \times E_2)[N]$ *under the product pairing* e_{prod} *if and only if*

$$\det(M_\alpha) \equiv -1 \pmod{N}.$$

Proof. For any two generators $G_1 = (P_1, \alpha(P_1))$ and $G_2 = (Q_1, \alpha(Q_1))$ of $\mathcal{K}_\alpha$, isotropy requires:

$$\begin{aligned} e_{prod}(G_1, G_2) &= e_1(P_1, Q_1) \cdot e_2(\alpha(P_1), \alpha(Q_1)) = 1 \\ &\Rightarrow \zeta_N \cdot e_2(s_1 P_2 + s_2 Q_2, s_3 P_2 + s_4 Q_2) = 1 \\ &\Rightarrow \zeta_N \cdot e_2(P_2, Q_2)^{s_1 s_4 - s_2 s_3} = 1 \\ &\Rightarrow \zeta_N \cdot \zeta_N^{\det(M_\alpha)} = 1 \Rightarrow \zeta_N^{1+\det(M_\alpha)} = 1. \end{aligned}$$

This holds if and only if $1 + \det(M_\alpha) \equiv 0 \pmod{N}$, or $\det(M_\alpha) \equiv -1 \pmod{N}$.

Remark 1 (Geometric Distinction). Let $\alpha : E_1[N] \to E_2[N]$ be an isomorphism, and let $K_\alpha \subset E_1 \times E_2$ denote its graph. The geometry of the quotient

$$A := (E_1 \times E_2)/K_\alpha$$

depends critically on whether α is symplectic or anti-symplectic with respect to the Weil pairing.

– **Symplectic isomorphisms.** If α preserves the Weil pairing, then the induced principal polarization on A is *reducible*. Consequently, A is *isogenous* to a product of elliptic curves and, in particular, is not isomorphic to the Jacobian of a genus-2 curve.
– **Anti-symplectic isomorphisms.** If α inverts the Weil pairing, then the induced principal polarization on A is *irreducible*. By the Frey–Kani correspondence, the abelian surface A is isomorphic to the Jacobian $\mathrm{Jac}(C)$ of a genus-2 curve C.

References

BL04. Christina Birkenhake and Herbert Lange. *Complex Abelian Varieties*, volume 302 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2 edition, 2004.

CD22. Wouter Castryck and Thomas Decru. An efficient key recovery attack on sidh, 2022.

CDS19. Wouter Castryck, Thomas Decru, and Benjamin Smith. Hash functions from superspecial genus-2 curves using richelot isogenies. In *Arithmetic, Geometry, Cryptography, and Coding Theory*, volume 770 of *Contemp. Math.*, pages 1–28. Amer. Math. Soc., 2019. Also available at `https://arxiv.org/abs/1905.02707`.

CLM+18. Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. Csidh: An efficient post-quantum commutative group action. In *Advances in Cryptology — ASIACRYPT 2018*, volume 11274 of *LNCS*, pages 395–427. Springer, 2018.

CLN20. Craig Costello, Patrick Longa, and Michael Naehrig. Efficient algorithms for supersingular isogeny Diffie–Hellman. *Journal of Cryptology*, 33:1131–1178, 2020.

Cos18. Craig Costello. Genus 2 sidh, the other side of the hyperelliptic key exchange coin. Invited talk at ISOP 2018, 2018.

CS19. Craig Costello and Benjamin Smith. The supersingular isogeny problem in genus 2 and beyond. IACR Cryptology ePrint Archive, 2019.

DFJP14. Luca De Feo, David Jao, and Jérôme Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. *J. Math. Cryptol.*, 8(3):209–247, 2014.

DFKPW20. Luca De Feo, David Kohel, Christophe Petit, and Benjamin Wesolowski. Sqisign: Compact post-quantum signatures from quaternions and isogenies. In *Advances in Cryptology — ASIACRYPT 2020*, volume 12492 of *LNCS*, pages 123–151. Springer, 2020.

FK91. Gerhard Frey and Ernst Kani. Curves of genus 2 covering elliptic curves and an arithmetical application. In *Arithmetic Algebraic Geometry*, pages 153–176. Springer, 1991.

Fly02. E. V. Flynn. The jacobian and formal group of a curve of genus 2 over an arbitrary ground field. *Math. Proc. Cambridge Philos. Soc.*, 107(3):425–441, 2002.

FT16. E. V. Flynn and Yan Bo Ti. Sidh on genus 2 hyperelliptic curves. Presentation at the ECC 2016 conference, 2016.

HSS23. Jeroen Hanselman, Sam Schiavone, and Jeroen Sijsling. Gluing curves of genus 1 and 2 along their 2-torsion. In *ANTS XV*, volume 6 of *Open Book Ser.*, pages 235–256, 2023.

JDF11. David Jao and Luca De Feo. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In *Post-Quantum Cryptography (PQCrypto 2011)*, volume 7071 of *LNCS*, pages 19–34. Springer, 2011.

MKS21. Maike Meyer, Jean-Paul Kieffer, and Brian Smith. Efficient strategies for isogeny-based key exchange in the csidh setting. *Designs, Codes and Cryptography*, 89:2389–2413, 2021.

Ric36. F. J. Richelot. Essai sur une méthode générale pour déterminer la valeur des intégrales ultra-elliptiques. *J. Reine Angew. Math.*, 16:285–341, 1836.

Rob23. Damien Robert. Breaking sidh in polynomial time. In *Advances in Cryptology — EUROCRYPT 2023*, volume 14004 of *LNCS*, pages 421–451. Springer, 2023.

Was08. Lawrence C. Washington. *Elliptic Curves: Number Theory and Cryptography*. Chapman and Hall/CRC, 2 edition, 2008.