A Post-Quantum Commitment Scheme from Richelot Isogeny Walks on Superspecial Genus-2 Jacobians

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

We present a post-quantum commitment scheme based on *kernel-tagged punctured* Richelot isogeny walks on superspecial genus-2 Jacobians. The puncturing rule skips every step landing in the product locus, detected by $I_{10} = 0$, so honest executions remain in the Jacobian locus and avoid the entry point of known product-locus attacks.

Each opening is encoded as a deterministic non-backtracking walk together with a kernel tag recording its action on a small public auxiliary torsion basis. The tag is verified as part of the opening and is kept explicit throughout the security analysis. In particular, scalar-related collisions force equality of the ordered kernel sequence and hence equality of the tag, so any nontrivial binding attack yields a short non-scalar endomorphism.

Using spectral bounds for the Richelot graph, we show that puncturing preserves rapid mixing for logarithmic walk lengths, which yields statistical hiding for the tagged punctured scheme. We therefore reduce binding to the Short Richelot Endomorphism Problem (SREP), relate SREP to the One-Endomorphism Problem and, under a standard KLPT^2-style heuristic, to the endomorphism-ring problem. A SageMath prototype based on $(2, 2)$-Kummer isogenies indicates practical performance at standard security levels.

Keywords: Post-quantum cryptography · Isogeny-based cryptography · Richelot isogenies · Superspecial genus-2 Jacobians · Commitment schemes

1 Introduction

1.1 Background and motivation

The collapse of SIDH [5] has intensified interest in higher-dimensional isogeny constructions. In genus 2, Richelot $(2, 2)$-isogenies between superspecial Jacobians form a 15-regular graph with about $p^3/2880$ vertices, making Richelot walks a natural source of post-quantum primitives such as commitments.

Two structural issues obstruct direct use of unrestricted Richelot walks. First, product-locus attacks such as Costello–Smith [11] exploit decomposable surfaces

$E_1 \times E_2$ to shortcut path finding. Second, endomorphism-based attacks, notably KLPT^2 [6], threaten binding whenever short non-scalar endomorphisms become accessible.

Both attack paradigms must interact with the product locus. Since that locus has relative density $\Theta(p^{-1})$ in the superspecial genus-2 graph, short random walks avoid it with overwhelming probability. We therefore build the scheme from a *punctured* walk that skips every step for which $I_{10} = 0$, ensuring that honest executions remain in the Jacobian locus.

Our approach. Our construction is *tagged throughout*: every accepted opening contains both the encoded walk and a kernel tag describing its action on a small public auxiliary torsion basis. The punctured walk handles product-locus shortcuts, while the kernel tag handles scalar-related collisions. If two accepted openings were scalar-related, kernel rigidity would force equality of their ordered kernel sequences and hence equality of their tags; by injectivity of the encoding, this collapses the collision to the same opening. Thus the binding analysis is stated for the tagged punctured scheme from the outset, not for an untagged intermediate variant.

1.2 Our contributions

1. **A kernel-tagged punctured Richelot commitment scheme** (Section 4). We define a commitment scheme in which every honest walk is non-backtracking, avoids the product locus by construction, and carries a verified kernel tag as part of the opening.
2. **Binding for the tagged punctured scheme** (Theorems 1 and 9). We keep the tag explicit in the scheme syntax, in verification, and in the binding proofs. Any nontrivial accepted collision yields a witness to the Short Richelot Endomorphism Problem (SREP).
3. **Conditioned mixing and statistical hiding** (Section 7.1 and Corollary 7). We show that puncturing by the condition $I_{10} \neq 0$ preserves the relevant mixing behaviour for logarithmic walk lengths, so the tagged punctured scheme remains statistically hiding.
4. **Reductions and implementation** (Section 6 and Section 9). We relate SREP unconditionally to the One-Endomorphism Problem and, under a standard KLPT^2-style heuristic, to the endomorphism-ring problem, and we provide a SageMath prototype based on $(2, 2)$-Kummer isogenies.

1.3 Related work

Genus-2 isogeny hashing was initiated by Charles–Goren–Lauter [8] in genus 1 and extended to genus 2 by Castryck–Decru–Smith [7] using non-backtracking Richelot walks. Subsequent works refined genus-2 hash constructions and highlighted the role of Kummer models and theta coordinates in higher-dimensional isogeny settings [14,2].

Cryptanalytic results such as the Costello–Smith shortcut attack [11] and related analyses of product-locus and endomorphism-based weaknesses [13,24] show that paths through decomposable abelian surfaces can dramatically reduce the effective security of naive designs. This motivates explicit countermeasures, such as avoiding the product locus and controlling scalar cycles, in any genus-2 walk-based primitive.

Isogeny-based commitment schemes in genus 1 were considered by Sterner [27] and Fouotsa–Tchotchou [29]. Our hardness discussion builds on a growing literature on endomorphism-ring problems in genus 1 and 2, including [12,19,6,30,20,9], and on the elliptic equivalence between OneEnd and $\mathsf{EndRing}$ established in [25].

2 Preliminaries

We begin by recalling the definition of a commitment scheme together with the standard hiding and binding notions used in cryptographic security.

2.1 Commitment Schemes

Definition 1 (Commitment Scheme [15]). *A commitment scheme* $\Pi = (\mathsf{KeyGen}, \mathsf{Commit}, \mathsf{Verify})$ *consists of three PPT algorithms:*
- $\mathsf{KeyGen}(1^\lambda)$*: outputs public parameters* pp *and message space* $\mathcal{M}$*.*
- $\mathsf{Commit}(\mathsf{pp}, m, r)$*: on* pp*, message* $m \in \mathcal{M}$*, randomness* $r \in \mathcal{R}$*, outputs commitment* c*.*
- $\mathsf{Verify}(\mathsf{pp}, m, r, c)$*: outputs* 1 *or* 0*.*

Correctness. *For all* $m \in \mathcal{M}$ *and properly generated* pp*:* $\Pr[\mathsf{Verify}(\mathsf{pp}, m, r, \mathsf{Commit}(\mathsf{pp}, m, r)) = 1] = 1$.

Definition 2 (Perfect, Statistical, and Computational Hiding [15]). *A commitment scheme* Π *is:*
- ***Perfectly hiding*** *if for all* $m_0, m_1 \in \mathcal{M}$ *and all* pp*, the distributions* $\mathsf{Commit}(\mathsf{pp}, m_0, r)$ *and* $\mathsf{Commit}(\mathsf{pp}, m_1, r)$ *are identical for uniform* r*.*
- ***Statistically hiding*** *if the statistical distance between distributions from different messages is negligible.*
- ***Computationally hiding*** *if for any PPT adversary* $\mathcal{A}$*:*

$$|\Pr[\mathsf{ExpHide}_0(\mathcal{A}) = 1] - \Pr[\mathsf{ExpHide}_1(\mathcal{A}) = 1]| \leq \mathrm{negl}(\lambda)$$

  where $\mathsf{ExpHide}_b$ *is the hiding experiment with challenge bit* b*.*

Definition 3 (Perfect, Statistical, and Computational Binding [15]). *A commitment scheme* Π *is:*
- ***Perfectly binding*** *if for all* pp *and* $m_0 \neq m_1$*, and for any* r_0, r_1*, it holds that:*

$$\mathsf{Commit}(\mathsf{pp}, m_0, r_0) \neq \mathsf{Commit}(\mathsf{pp}, m_1, r_1).$$

- ***Statistically binding*** *if the probability of finding two valid openings is negligible.*

– ***Computationally binding*** *if for any PPT adversary* $\mathcal{A}$*:*

$$\Pr[\mathsf{ExpBind}(\mathcal{A}) = 1] \leq \mathrm{negl}(\lambda)$$

where ExpBind *is the experiment where* $\mathcal{A}$ *outputs* (m, r, m', r', c) *with* $m \neq m'$ *that both verify.*

3 Mathematical Background

Throughout Sections 3.1–3.3 we assume $\mathrm{char}(k) \neq 2$. This is required for the existence of theta constants, for the invertibility of the Hadamard transform ($\mathcal{H}^2 = 4 \cdot \mathrm{Id}$), and for Cantor's addition algorithm. In practice $k = \mathbb{F}_{p^2}$ with $p > 5$ a prime.

3.1 Genus-2 Curves and Their Jacobians

A *genus-2 curve* over $k = \mathbb{F}_{p^2}$ is a smooth projective curve of genus 2. Every such curve is hyperelliptic and admits a Weierstrass model $y^2 = f(x)$ with $\deg f \in \{5, 6\}$ and f separable [4]. Its *Jacobian* $\mathrm{Jac}(C)$ is a 2-dimensional principally polarised abelian variety (PPAV) over k [22].

The rational points $\mathrm{Jac}(C)(k)$ are parametrised by *Mumford's representation* [23]: every divisor class $[D] \in \mathrm{Jac}(C)(k)$ admits a unique *reduced representative* of the form

$$[D] \;\longleftrightarrow\; [u(x),\, v(x)], \quad u \mid (v^2 - f), \quad v^2 \equiv f \pmod{u}, \quad \deg v < \deg u \leq 2, \tag{J.1}$$

where u is monic [23].

3.2 Kummer Surfaces

Definition 4 (Kummer surface [17,4]). *Let* A/k *be a principally polarised abelian surface. The* Kummer surface *of* A *is the geometric quotient*

$$\mathrm{Kum}(A) \;:=\; A \big/ \langle [-1] \rangle, \tag{K.1}$$

with canonical projection $\pi_A : A \to \mathrm{Kum}(A)$, $P \mapsto \{P, -P\}$*. Under the level-2 theta embedding,* $\mathrm{Kum}(A)$ *admits a projective model as a quartic surface in* $\mathbb{P}^3(k)$ *with* 16 *ordinary double points (nodes) at the images of the* 2*-torsion points* $A[2]$ *[4,17].*

Definition 5 (Squared theta coordinates [23,14]). *Fix* $\mathrm{char}(k) \neq 2$ *and a level-2 theta structure on* A*, i.e. a symplectic isomorphism* $A[2] \xrightarrow{\sim} (\mathbb{Z}/2\mathbb{Z})^4$ *[21]. The* squared theta coordinates *are the four fundamental theta constants associated to this structure:*

$$\mu = (\mu_0 : \mu_1 : \mu_2 : \mu_3) \in \mathbb{P}^3(k), \qquad \mu_i = \vartheta_i(A, \mathcal{L})^2, \tag{K.2}$$

where $\vartheta_0, \dots, \vartheta_3$ are the four theta functions with half-integer characteristics selected by the level-2 structure, normalised as in Gaudry–Lubicz–Robert [14]. Two surfaces $K(\mu)$ and $K(\mu')$ are isomorphic over k if and only if $\mu' = \lambda\mu$ for some $\lambda \in k^*$ [18].

Definition 6 (Kummer constants [14,2]). *Given squared theta constants μ as in Definition 5, the* Kummer constants $\kappa = (\kappa_0 : \kappa_1 : \kappa_2 : \kappa_3)$ *are*

$$(\kappa_0 : \kappa_1 : \kappa_2 : \kappa_3) \;=\; \mathcal{H}\left(\frac{1}{\mu_0^2}, \frac{1}{\mu_1^2}, \frac{1}{\mu_2^2}, \frac{1}{\mu_3^2}\right), \tag{K.3}$$

where $\mathcal{H}$ is the Hadamard transform (Definition 7) and the μ_i are assumed nonzero (generic hypothesis). The Kummer constants are pre-computed once per surface and serve as coefficients in the doubling and differential-addition formulas [14].

Definition 7 (Hadamard transform [16,14]). *The* Hadamard transform *on k^4 is the linear map $\mathcal{H} : k^4 \to k^4$ defined by*

$$\mathcal{H}(a,b,c,d) \;:=\; (a+b+c+d,\; a-b+c-d,\; a+b-c-d,\; a-b-c+d). \tag{H.1}$$

Since $\mathrm{char}(k) \neq 2$*, the map is invertible and satisfies $\mathcal{H}^2 = 4\cdot\mathrm{Id}$, hence $\mathcal{H}^{-1} = \frac{1}{4}\mathcal{H}$ [14].*

Proposition 1 (Properties of $\mathcal{H}$ [14,10]).

(i) $\mathcal{H} \circ \mathcal{H} = 4 \cdot \mathrm{Id}$ *[14].*

(ii) All matrix entries of $\mathcal{H}$ are ± 1; computing $\mathcal{H}(a,b,c,d)$ requires exactly 12 additions or subtractions and 0 multiplications [10].

(iii) The push-forward of a Richelot isogeny in squared theta coordinates decomposes as $\mathcal{H} \circ D_\pi \circ \mathcal{H}$, where D_π is the diagonal scaling by $(\pi_i)_{0\le i\le 3}$, $\pi_i = \mu_i'/\mu_i$ [10]. (See Proposition 2.)

3.3 Richelot (2, 2)-Isogenies

Definition 8 (Richelot isogeny [26,7]). *A* Richelot $(2,2)$-isogeny $\varphi : A \to A'$ *is a separable isogeny of degree 4 whose kernel is a maximal isotropic subgroup*

$$K \subset A[2], \quad K \cong (\mathbb{Z}/2\mathbb{Z})^2, \tag{RI.1}$$

with respect to the Weil pairing $e_2 : A[2] \times A[2] \to \mu_2$ [22]. The dual isogeny $\hat{\varphi} : A' \to A$ satisfies $\hat{\varphi} \circ \varphi = [2]_A$ [22]. For a generic *genus-2 Jacobian over $\bar{k}$, one has $A[2](\bar{k}) \cong (\mathbb{Z}/2\mathbb{Z})^4$, which possesses exactly 15 maximal isotropic subgroups [7], hence 15 distinct Richelot isogenies defined over $\bar{k}$ (fewer over k if torsion is not fully rational).*

Proposition 2 (Descent to Kummer [10,2]). *Since $[-1]$ acts trivially on $A[2]$, the isogeny φ_K commutes with $[-1]$ and descends to a well-defined rational map*

$$\bar{\varphi}_K : \mathrm{Kum}(A) \dashrightarrow \mathrm{Kum}(A/K), \tag{KD.1}$$

given in squared theta coordinates — in homogeneous projective form — by [10]:

$$\bar{\varphi}_K(Q) \;=\; \mathcal{H}\!\left(\frac{\mu_0'}{\mu_0}(H_Q)_0 : \cdots : \frac{\mu_3'}{\mu_3}(H_Q)_3\right), \quad H_Q = \mathcal{H}(Q), \tag{KD.2}$$

where $\mu = (\mu_0 : \ldots : \mu_3)$ and $\mu' = (\mu_0' : \ldots : \mu_3')$ are the squared theta-null points of A and A/K respectively, and all expressions are interpreted in homogeneous projective coordinates $\mathbb{P}^3(k)$. The formula is the decomposition $\mathcal{H} \circ D_\pi \circ \mathcal{H}$ of Proposition 1(iii), with $\pi_i = \mu_i'/\mu_i$.

3.4 The Richelot Isogeny Graph

The *Richelot isogeny graph* $\mathcal{G}_p$ has as vertices the isomorphism classes of superspecial principally polarised abelian surfaces (PPAVs) over $\overline{\mathbb{F}}_{p^2}$. That graph is d-regular with $d = 15$. Let $P_v^{(k)}$ denote the distribution of the endpoint of a k-step non-backtracking random walk starting at vertex v, and let U be the uniform distribution on $\mathcal{V}$.

Proposition 3 (Spectral Bound [1]). *Let A be the adjacency matrix of $\mathcal{G}_p$ and let $\lambda_2(A)$ be its second largest eigenvalue in absolute value. Then $|\lambda_2(A)| \le 8\sqrt{2} \approx 11.31$. The normalised second eigenvalue is $\mu_2 := \frac{|\lambda_2(A)|}{d} \le \frac{8\sqrt{2}}{15} \approx 0.754$, which gives a constant spectral gap $\gamma_0 := 1 - \mu_2 \ge 1 - \frac{8\sqrt{2}}{15} \approx 0.246$.*

Lemma 1 (Mixing Time Bound [1]). *For any vertex v and any integer $k \ge 0$,*

$$\Delta\big(P_v^{(k)}, U\big) \;\le\; \frac{\sqrt{|\mathcal{V}|}}{2}\,\mu_2^k \;\le\; \frac{\sqrt{|\mathcal{V}|}}{2}\left(\frac{8\sqrt{2}}{15}\right)^k.$$

Proof. Let $P = A/d$ be the transition matrix of the random walk. Because $\mathcal{G}_p$ is connected and d-regular, the uniform distribution $U = \mathbf{1}/|\mathcal{V}|$ is stationary. Let $\{u_i\}$ be an orthonormal eigenbasis of P with eigenvalues $1 = \mu_1 > |\mu_2| \ge \cdots \ge |\mu_n|$. Write the initial distribution e_v (unit vector at v) as $e_v = \frac{1}{\sqrt{|\mathcal{V}|}}u_1 + \sum_{i\ge 2}\langle e_v, u_i\rangle u_i$. Then $P^k e_v - U = \sum_{i\ge 2}\langle e_v, u_i\rangle \mu_i^k u_i$, and orthonormality gives $\|P^k e_v - U\|_2^2 \le \mu_2^{2k}\sum_{i\ge 2}|\langle e_v, u_i\rangle|^2 \le \mu_2^{2k}$. Finally, $\Delta(P_v^{(k)}, U) = \frac{1}{2}\|P^k e_v - U\|_1 \le \frac{\sqrt{|\mathcal{V}|}}{2}\|P^k e_v - U\|_2$.

Lemma 2 (Path Independence). *Let $v, w \in \mathcal{V}$ be the endpoints of two fixed walks of length k_m starting from the same origin A_0. For a fresh random walk of length k_r started from each, we have using triangle inequality on lemma 1:*

$$\Delta\big(P_v^{(k_r)},\, P_w^{(k_r)}\big) \;\le\; \sqrt{|\mathcal{V}|}\,\mu_2^{k_r} \;\le\; \sqrt{|\mathcal{V}|}\left(\frac{8\sqrt{2}}{15}\right)^{k_r}.$$

4 Base Richelot Commitment Scheme

4.1 System Parameters

- **Prime p:** Take $p \equiv 3 \pmod 4$; $\log_2 p$ as in Table 1.
- **Base field:** $\mathbb{F}_{p^2} = \mathbb{F}_p[i]$ with $i^2 + 1 = 0$ for $p \equiv 3 \pmod 4$.
- **Starting surface A_0:** A superspecial genus-2 Jacobian $A_0 = \mathrm{Jac}(C_0) \in \mathcal{V}_{\mathrm{Jac}}$. The surface is represented internally by its squared Kummer theta constants $\mu = (\mu_0 : \mu_1 : \mu_2 : \mu_3) \in \mathbb{P}^3(\mathbb{F}_{p^2})$.
- **Walk lengths:** Message length k_m and randomiser length k_r in non-backtracking Richelot steps; total walk length $N = k_m + k_r$.
- **Path-label space:** $\mathcal{P}_N := \big(\mathbb{Z}/15\mathbb{Z}\big) \times \big(\mathbb{Z}/14\mathbb{Z}\big)^{N-1}$. The mixed-radix structure encodes all admissible non-backtracking transitions: at step 1 all 15 Richelot kernels are available ($b_1 \in \{0, \ldots, 14\}$); at step $j > 1$ the dual of the previous isogeny is excluded, leaving 14 choices ($b_j \in \{0, \ldots, 13\}$). Injectivity holds because $2^N < 15 \cdot 14^{N-1} = |\mathcal{P}_N|$ for all parameter sets in Table 1.
- **Endpoint encoding EncEp:** Given the terminal Kummer surface K_N with squared theta constants $\mu = (\mu_0 : \mu_1 : \mu_2 : \mu_3)$, normalise by the last non-zero coordinate: $\ell := \max\{i \mid \mu_i \not\equiv 0 \pmod p\}$, $\tilde{\mu}_i := \mu_i/\mu_\ell$, then serialise each $\tilde{\mu}_i$ as $\lceil \log_2 p/8 \rceil$ big-endian bytes:

$$\mathsf{EncEp}(K_N) := \tilde{\mu}_0 \,\|\, \tilde{\mu}_1 \,\|\, \tilde{\mu}_2 \,\|\, \tilde{\mu}_3 \in \{0,1\}^{4\lceil \log_2 p/8 \rceil \cdot 8}.$$

- **Hash H:** A collision-resistant function $H = \mathsf{SHA3\text{-}256} : \{0,1\}^* \to \{0,1\}^{256}$.

Table 1. Parameter recommendations: Values of k_r are derived from Corollary 3.

Level	$\log_2 p$	k_m	k_r	$N = k_m + k_r$	SREP bound
NIST-I (classical)	128	128	774	902	2^{128}
NIST-I (quantum)	256	128	1244	1372	2^{128}
NIST-III	384	192	1874	2066	2^{192}
NIST-V	512	256	2502	2758	2^{256}

4.2 Algorithms

Algorithm 1: $\mathsf{KeyGen}^*(1^\lambda)$

Input: Security parameter λ.
Output: Public parameters pp^*.

Choose $(\log_2 p, k_m, k_r)$ from Table 1;
Take the smallest prime $p \equiv 3 \pmod 4$ of the required bitlength;
Set $\mathbb{F}_{p^2} = \mathbb{F}_p[i]$ with $i^2 + 1 = 0$ and obtain a superspecial starting Jacobian A_0 with its fixed Kummer representation;
Set $N \leftarrow k_m + k_r$;
Choose an admissible torsion prime $\ell \neq 2, p$ for N (Definition 11) and fix a public point $R \in A_0[\ell]$ of exact order ℓ;
Fix H, the encoding map ι, the endpoint encoder EncEp, the public auxiliary torsion basis $\mathcal{B}_{\text{aux}}$, the tag-update map $\mathsf{TagUpdate}$, and the tag-encoding map $\mathsf{TagEncode}$;
return $\mathsf{pp}^* = (p, \mathbb{F}_{p^2}, A_0, k_m, k_r, N, \ell, R, H, \iota, \mathsf{EncEp}, \mathcal{B}_{\text{aux}}, \mathsf{TagUpdate}, \mathsf{TagEncode})$;

Algorithm 2: $\mathsf{Commit}(\mathsf{pp}, m, r, \varrho)$

Input: pp; $m \in \{0,1\}^{k_m}$; $r \in \{0,1\}^{k_r}$; $\varrho \in \{0,1\}^{256}$
Output: (c, τ), or $\perp$

$(b_1, \ldots, b_N) \leftarrow \iota(m \parallel r)$;
$A \leftarrow A_0$; $K_{\text{dual}} \leftarrow \perp$; $\tau \leftarrow \tau_0$;
for $j = 1$ **to** N **do**
  compute the ordered Richelot kernel list $\mathcal{L}(A)$;
  if $j = 1$ **then**
    $K^\star \leftarrow \mathcal{L}(A)[b_1]$;
  else
    $\mathcal{L}'(A) \leftarrow \mathcal{L}(A) \setminus \{K_{\text{dual}}\}$;
    if $|\mathcal{L}'(A)| \neq 14$ **then**
      return $\perp$;
    $K^\star \leftarrow \mathcal{L}'(A)[b_j]$;
  compute $\phi_j : A \to A' := A/\langle K^\star \rangle$;
  $\tau \leftarrow \mathsf{TagUpdate}(\tau, \phi_j, \mathcal{B}_{aux})$;
  $K_{\text{dual}} \leftarrow \ker(\widehat{\phi_j})$;
  $A \leftarrow A'$;
$\tau \leftarrow \mathsf{TagEncode}(\tau)$;
$c \leftarrow H(\varrho \parallel \tau \parallel \mathsf{EncEp}(A))$;
return (c, τ);

Algorithm 3: $\mathsf{Verify}(\mathsf{pp}, m, r, \varrho, \tau, c)$

Input: pp; m; r; ϱ; τ; c
Output: accept / reject
1 $(c', \tau') \leftarrow \mathsf{Commit}(\mathsf{pp}, m, r, \varrho)$;
2 **return** accept if $(\tau' = \tau$ and $c' = c)$, else reject

Remark 1 (Layered description of the commitment map). Let $N = k_m + k_r$. The implementation-level commitment map with fixed salt ϱ is the deterministic composition

$$\mathsf{Com}_\varrho \;:\; \{0,1\}^{k_m+k_r} \xrightarrow{\mathsf{Enc}} \mathcal{P}_N \xrightarrow{\mathsf{Walk}_{A_0}} \mathcal{W}_N(A_0) \xrightarrow{\mathsf{EncEp}} \{0,1\}^* \xrightarrow{H_\varrho} \{0,1\}^{256},$$

where:
- Enc is the mixed-radix encoding into $\mathcal{P}_N = (\mathbb{Z}/15\mathbb{Z}) \times (\mathbb{Z}/14\mathbb{Z})^{N-1}$;
- Walk_{A_0} executes the deterministic non-backtracking Richelot walk from A_0 using the lexicographically ordered kernel lists at each step;
- EncEp is the normalised Kummer theta serialisation of the terminal surface K_N (Section 4.1, "Endpoint encoding");
- $H_\varrho = \mathsf{SHA3\text{-}256}(\varrho \,\|\, \cdot\,)$.

The formal quantity $\mathsf{Igusa}(B)$ used in the abstract description is replaced by $\mathsf{EncEp}(K_N)$, which serves as a canonical, efficiently computable representative of the isomorphism class of the terminal surface in the Kummer / squared-theta model.

5 Formal Security Analysis

5.1 The Short Richelot Endomorphism Problem

Definition 9 (Short Richelot Endomorphism Problem – SREP). *Let A_0 be a superspecial genus-2 Jacobian and let $L \in \mathbb{N}$. $\mathsf{SREP}(A_0, L)$ is the problem of finding a non-scalar endomorphism*

$$\psi \in \mathrm{End}(A_0) \setminus \mathbb{Z} \cdot \mathrm{id}_{A_0}$$

expressible as a composition of at most L Richelot $(2,2)$-isogenies (and their duals):

$$\psi = \hat{\phi}_k \circ \cdots \circ \hat{\phi}_1 \circ \phi'_1 \circ \cdots \circ \phi'_m, \qquad k + m \le L.$$

The advantage $\mathbf{Adv}^{\mathsf{SREP}}_{A_0,L}(\mathcal{A})$ of an adversary $\mathcal{A}$ is the success probability.

5.2 Encoding injectivity, scalar cycles, and conditional reconstruction

In this section we formalise the combinatorial injectivity of the encoding from bitstrings to Richelot walks, record the precise form of the scalar-cycle relation in $\mathrm{End}(A_0)$, and state the conditional hypothesis used to rule out scalar collisions between distinct valid openings.

Kernel filtration for composed Richelot walks.

Lemma 3 (Kernel filtration). *Let*

$$\Phi_N = \phi_N \circ \cdots \circ \phi_1 : A_0 \to A_N$$

be a composition of $N \geq 1$ separable Richelot $(2,2)$-isogenies, and for each $i \in \{1, \ldots, N\}$ define

$$\Phi_i := \phi_i \circ \cdots \circ \phi_1, \qquad G_i := \ker(\Phi_i), \qquad G_0 := 0_{A_0}.$$

Then:

1. *$0_{A_0} = G_0 \subset G_1 \subset \cdots \subset G_N = \ker(\Phi_N)$ is an increasing chain of finite subgroup schemes of $A_0[2^N]$;*
2. *for each i, one has $G_i \subset A_0[2^i]$ and $|G_i| = 4^i$;*
3. *for each $i \geq 1$, one has $|G_i/G_{i-1}| = 4$.*

Proof. The proof is by induction on i. For $i = 1$, ϕ_1 is a Richelot $(2,2)$-isogeny, so $G_1 = \ker(\phi_1) \subset A_0[2]$ is a maximal isotropic subgroup of order 4. Assume the claim for $i - 1$. Write

$$\Phi_i = \phi_i \circ \Phi_{i-1},$$

where $\phi_i : A_{i-1} \to A_i$ is again a separable Richelot $(2,2)$-isogeny. Its kernel $K_i := \ker(\phi_i)$ has order 4 and is contained in $A_{i-1}[2]$. By functoriality of kernels under composition,

$$G_i = \ker(\Phi_i) = \Phi_{i-1}^{-1}(K_i),$$

so $G_{i-1} \subset G_i$. Since Φ_{i-1} is separable of degree 4^{i-1}, the induced exact sequence

$$0_{A_0} \to G_{i-1} \to G_i \to K_i \to 0_{A_0}$$

implies

$$|G_i| = |G_{i-1}||K_i| = 4^{i-1} \cdot 4 = 4^i, \qquad |G_i/G_{i-1}| = 4.$$

Finally, because $K_i \subset A_{i-1}[2]$, its full preimage under Φ_{i-1} is contained in $A_0[2^i]$, hence $G_i \subset A_0[2^i]$.

Injectivity of the labelled walk map.

Lemma 4 (Injectivity of the labelled walk map). *Fix the public parameters and an initial surface A_0. At every reachable vertex A_i, fix once and for all a total ordering of the 15 Richelot kernels in $A_i[2]$. Algorithm 2 uses this ordering together with the non-backtracking rule: at the first step all 15 kernels are admissible, and at each later step the dual kernel is excluded and the remaining 14 admissible kernels are taken in the induced order. Let*

$$\mathrm{Walk}_{A_0} : \mathcal{P}_N \to \{\textit{labelled non-backtracking Richelot walks of length } N \textit{ from } A_0\}$$

be the resulting deterministic map. Then Walk_{A_0} is injective.

Proof. Let $\mathbf{b} = (b_1, \dots, b_N)$ and $\mathbf{b}' = (b'_1, \dots, b'_N)$ be distinct elements of $\mathcal{P}_N$, and let

$$j^* := \min\{j : b_j \neq b'_j\}.$$

Up to step $j^* - 1$, the two executions traverse the same intermediate vertices and therefore use the same ordered list of admissible kernels at step j^*. Since $b_{j^*} \neq b'_{j^*}$, the two executions choose different kernels at that step. Hence the kernel sequences, and therefore the labelled walks, are distinct.

Corollary 1 (Encoding injectivity at the bitstring level). *Let*

$$\iota : \{0,1\}^{\ell_m + \ell_r} \hookrightarrow \mathcal{P}_N$$

be the fixed injective encoding used by the commitment algorithm. Then the induced map

$$\mathrm{Enc} := \mathrm{Walk}_{A_0} \circ \iota$$

is injective. In particular, if $S \neq S'$, then the corresponding kernel sequences and labelled walks are distinct.

Remark 2. Corollary 1 is intentionally phrased at the level of labelled walks in the fixed ordered Richelot graph. It does not assert uniqueness of the resulting isogeny as a morphism $A_0 \to A$, since that could be affected by automorphisms of intermediate Jacobians or by isomorphisms between targets. For the binding reduction, injectivity at the level of labelled walks is the relevant notion.

Scalar cycles as a degree-compatible relation.

Definition 10 (Scalar cycle). *Let $\phi_S, \phi_{S'} : A_0 \to A$ be two encoded non-backtracking Richelot walks of length N with a common codomain A. We say that $(\phi_S, \phi_{S'})$ forms a scalar cycle if*

$$\hat{\phi}_{S'} \circ \phi_S = [\varepsilon 2^N]_{A_0}$$

in $\mathrm{End}(A_0)$ *for some* $\varepsilon \in \{-1, 1\}$.

Lemma 5 (Degree check for scalar cycles). *If $(\phi_S, \phi_{S'})$ forms a scalar cycle in the sense of Definition 10, then both sides of the relation have degree 4^N.*

Proof. Each Richelot step has degree 4, so $\deg(\phi_S) = \deg(\phi_{S'}) = 4^N = 2^{2N}$. Since the dual isogeny has the same degree, we have

$$\deg(\hat{\phi}_{S'} \circ \phi_S) = \deg(\hat{\phi}_{S'}) \deg(\phi_S) = 2^{2N} \cdot 2^{2N} = 2^{4N}$$

More precisely, the composite has degree

$$4^N \cdot 4^N = 16^N.$$

On the other hand, for an abelian surface one has $\deg([n]) = n^4$, so

$$\deg([\varepsilon 2^N]) = (2^N)^4 = 16^N.$$

Thus the relation is degree-compatible.

Conditional scalar-cycle uniqueness.

Algebraic reduction: global kernel equality

Lemma 6 (Global kernel equality under a scalar cycle). *Let $\Phi = \phi_S$ and $\Psi = \phi_{S'}$ be encoded non-backtracking Richelot walks of length N such that*

$$\widehat{\Psi} \circ \Phi = [\varepsilon 2^N]_{A_0} \quad \text{for } \varepsilon \in \{-1, 1\}.$$

Then $\Phi = \varepsilon\Psi$ in $\mathrm{Hom}(A_0, A)$. In particular, $\ker(\Phi) = \ker(\Psi) \subset A_0[2^N]$.

Proof. Each Richelot step has degree 4, so $\deg(\Psi) = 4^N = 2^{2N}$ and $\Psi \circ \widehat{\Psi} = [2^N]_A$. Left-composing the scalar relation by Ψ gives

$$(\Psi \circ \widehat{\Psi}) \circ \Phi = \Psi \circ [\varepsilon 2^N]_{A_0} = [\varepsilon 2^N]_A \circ \Psi,$$

hence $[2^N]_A \circ \Phi = [\varepsilon 2^N]_A \circ \Psi$. The group $\mathrm{Hom}(A_0, A)$ is a torsion-free $\mathbb{Z}$-module, so from

$$[2^N]_A \circ (\Phi - \varepsilon\Psi) = 0$$

we conclude $\Phi = \varepsilon\Psi$. Equality of homomorphisms forces $\ker(\Phi) = \ker(\Psi)$, so both walks have the same global kernel $G_N \subset A_0[2^N]$.

Filtrations and encoding By Lemma 3, each encoded non-backtracking walk ϕ_S of length N induces a kernel filtration

$$0 = G_0 \subset G_1 \subset \cdots \subset G_N \subset A_0[2^N],$$

with $G_i \subset A_0[2^i]$ and $|G_i| = 4^i$ for $1 \le i \le N$. In particular, a scalar cycle as in Lemma 6 forces two walks to share the same global kernel G_N, though their intermediate filtration layers G_i need not coincide a priori.

Finally, by the encoding injectivity (Corollary 1), the map

$$= \mathrm{Walk}_{A_0} \circ \iota : \{0,1\}^{k_m + k_r} \hookrightarrow \{\text{labelled non-backtracking walks of length } N\}$$

is injective, so two walks with the same ordered kernel sequence at every step must be encoded by the same bitstring. In the tagged punctured scheme we exploit this together with the tag to rule out scalar cycles between distinct valid openings.

Endpoint encoding. Let K_N be the terminal Kummer surface, represented in some fixed implementation-level model by a tuple of squared theta constants

$$\mu = (\mu_0, \mu_1, \mu_2, \mu_3)$$

defined up to a common non-zero scalar factor in the base field. We normalise by the last non-zero coordinate:

$$\widetilde{\mu} = \left(\frac{\mu_0}{\mu_\ell}, \frac{\mu_1}{\mu_\ell}, \frac{\mu_2}{\mu_\ell}, \frac{\mu_3}{\mu_\ell} \right),$$

where

$$\ell = \max\{\, i \mid \mu_i \not\equiv 0 \pmod{p} \,\}.$$

Each $\widetilde{\mu}_i$ lies in the base field (typically a finite field of characteristic p), and is serialized as a fixed-length big-endian byte string of $\lceil \log_2 p/8 \rceil$ bytes. We define

$$\mathsf{EndEnc}(K_N) := \widetilde{\mu}_0 \,\|\, \widetilde{\mu}_1 \,\|\, \widetilde{\mu}_2 \,\|\, \widetilde{\mu}_3,$$

where $\|$ denotes byte concatenation. Given a 32-byte salt ρ, fresh for each commitment, the final commitment is computed as

$$c = \text{SHA3-256}\big(\rho \,\|\, \mathsf{EndEnc}(K_N)\big).$$

Remark 3 (Layered description of the commitment map). Let $N = \ell_m + \ell_r$, let k be the base field of computation, and let A_0 be the fixed starting superspecial genus-2 Jacobian. Define

$$\mathcal{P}_N := \Big\{(b_1, \ldots, b_N) :\ b_1 \in \{0, \ldots, 14\},\ b_j \in \{0, \ldots, 13\} \text{ for } 2 \le j \le N\Big\},$$

and let $\mathcal{W}_N(A_0)$ denote the set of labelled non-backtracking Richelot walks of length N starting from A_0.

The implementation-level commitment map with fixed salt $\varrho \in \{0,1\}^{256}$ is a deterministic function

$$\mathsf{Com}_\varrho : \{0,1\}^{\ell_m+\ell_r} \longrightarrow \{0,1\}^{256}$$

given by the composition

$$\{0,1\}^{\ell_m+\ell_r} \xrightarrow{\mathsf{Enc}} \mathcal{P}_N \xrightarrow{\mathsf{Walk}_{A_0}} \mathcal{W}_N(A_0) \xrightarrow{\mathsf{EndEnc}} \{0,1\}^* \xrightarrow{\mathsf{Hash}_\varrho} \{0,1\}^{256}.$$

Equivalently, for $S \in \{0,1\}^{\ell_m+\ell_r}$,

$$\mathsf{Com}_\varrho(S) = \mathsf{Hash}_\varrho\big(\mathsf{EndEnc}\big(\mathsf{Walk}_{A_0}\big(\mathsf{Enc}(S)\big)\big)\big)\,.$$

5.3 Choice of the torsion parameter ℓ

Recall that the torsion tag records the action of the walk on a small public ℓ^k-torsion basis on the codomain surface B, where $\ell \neq 2, p$ is an odd prime and $k \ge 1$ is small. In the scalar-cycle analysis, we consider a relation

$$\widehat{\phi_{S'}} \circ \phi_S \;=\; [\varepsilon 2^N]_{A_0}$$

for some $\varepsilon \in \{1, -1\}$ and total walk length N. On the level of ℓ-torsion this induces an equation of the form

$$2^N(2^N - \varepsilon)\, P \;=\; 0 \qquad \text{in } B[\ell], \tag{1}$$

where P is the image in $B[\ell]$ of a nonzero point of the auxiliary torsion basis under a suitable isogeny.

To conclude from (1) that $P = 0$, we need $2^N(2^N - \varepsilon)$ to be invertible modulo ℓ.

Definition 11 (Admissible torsion prime). *Fix a total walk length $N \geq 1$. An odd prime $\ell \neq p$ is* admissible *for N if*

$$2^N \not\equiv 1 \pmod{\ell} \qquad \text{and} \qquad 2^N \not\equiv -1 \pmod{\ell}.$$

Equivalently, $\ell \nmid 2^N(2^N - \varepsilon)$ for both $\varepsilon \in \{1, -1\}$.

For any fixed N, there are infinitely many admissible primes $\ell \neq 2, p$, since only the two congruence classes $2^N \equiv \pm 1 \pmod{\ell}$ must be excluded. In practice, one can sample small odd primes distinct from p until the two conditions in Definition 11 are satisfied.

Lemma 7 (Exclusion of scalar cycles on admissible torsion). *Let $N \geq 1$ and let ℓ be admissible for N in the sense of Definition 11. Suppose that a tagged collision induces a scalar relation $\widehat{\phi_{S'}} \circ \phi_S = [\varepsilon 2^N]_{A_0}$ for some $\varepsilon \in \{1, -1\}$, and let $P \in B[\ell]$ be as in (1). Then no nonzero point $P \in B[\ell]$ can satisfy (1).*

Proof. Working in the cyclic group $\langle P \rangle \cong \mathbb{Z}/\ell\mathbb{Z}$, the relation (1) is equivalent to $2^N(2^N - \varepsilon)P = 0$. By admissibility, $\ell \nmid 2^N(2^N - \varepsilon)$, so multiplication by $2^N(2^N - \varepsilon)$ is invertible modulo ℓ. Hence the only solution is $P = 0$, contrary to the assumption that P has order ℓ.

5.4 Binding via SREP

Theorem 1 (Binding of the torsion-tagged scheme via SREP). *Let Π^{tag} be the torsion-tagged Richelot commitment scheme with total walk length N. Assume:*

- *H is collision-resistant;*
- *the public torsion parameter ℓ is admissible for N in the sense of Definition 11, and the setup fixes a public point $R \in A_0[\ell]$ of exact order ℓ;*
- *$\mathsf{SREP}(A_0, 2N)$ is hard.*

Then for any PPT adversary $\mathcal{A}$ there exists a PPT adversary $\mathcal{B}$ such that

$$\mathbf{Adv}^{\mathrm{bind}}_{\Pi^{\mathrm{tag}}}(\mathcal{A}) \;\leq\; \mathbf{Adv}^{\mathsf{SREP}}_{A_0,2N}(\mathcal{B}) \;+\; \epsilon_{\mathrm{CRH}}(\lambda).$$

Proof. We proceed via hybrid games. Throughout, the torsion tag is the image $Q = \phi_S(R) \in B[\ell]$ of the fixed public point $R \in A_0[\ell]$ under the encoded walk.

Game 0 (Real binding attack on the tagged scheme). The adversary outputs two distinct openings (m, r, ϱ, τ) and $(m', r', \varrho', \tau')$ that both verify for the same commitment value c. Let $S = m \parallel r$, $S' = m' \parallel r'$ and let

$$\phi_S : A_0 \to B, \qquad \phi_{S'} : A_0 \to B'$$

be the corresponding tagged punctured Richelot walks with endpoints B, B'. Let $Q := \phi_S(R) \in B[\ell]$ and $Q' := \phi_{S'}(R) \in B'[\ell]$ denote the images of R. Let Succ_0 be the event that both openings verify and are distinct; then $\Pr[\mathsf{Succ}_0] = \mathbf{Adv}^{\mathrm{bind}}_{\Pi^{\mathrm{tag}}}(\mathcal{A})$.

Game 1 (Eliminate hash collisions). In the real scheme the commitment has the form

$$c = H\big(\varrho \parallel \mathsf{TagEncode}(Q) \parallel \mathsf{EncEp}(B) = H\big(\varrho' \parallel \mathsf{TagEncode}(Q') \parallel \mathsf{EncEp}(B').$$

In Game 1 we abort if the two accepted openings yield the same value c through different hash inputs. By collision resistance of H, the change in success probability is at most $\epsilon_{\mathrm{CRH}}(\lambda)$. Conditioned on no abort, equality of commitments implies

$$\varrho = \varrho', \qquad \mathsf{TagEncode}(Q) = \mathsf{TagEncode}(Q') \qquad \mathsf{EncEp}(B) = \mathsf{EncEp}(B').$$

In particular $B \cong B'$ and Q and Q' represent the same point in $B[\ell]$ up to the fixed endpoint encoding. Thus, after fixing an isomorphism $\iota : B' \xrightarrow{\sim} B$, we may and do assume that $Q = Q'$ as elements of $B[\ell]$.

Let Succ_1 be the success event in Game 1. Then

$$\big|\Pr[\mathsf{Succ}_1] - \Pr[\mathsf{Succ}_0]\big| \le \epsilon_{\mathrm{CRH}}(\lambda),$$

and from now on we condition on Succ_1.

Setting up the endomorphism. Fix an isomorphism $\iota : B' \xrightarrow{\sim} B$ and set

$$\tilde{\phi}_{S'} := \iota \circ \phi_{S'} : A_0 \to B, \qquad \psi := \widehat{\tilde{\phi}_{S'}} \circ \phi_S \in \mathrm{End}(A_0).$$

By construction, ψ is expressible as a composition of at most $2N$ Richelot $(2,2)$-isogenies and their duals.

Moreover, by the discussion above the accepted tag equality implies

$$Q = \phi_S(R) = \tilde{\phi}_{S'}(R) \in B[\ell].$$

Write

$$\phi := \phi_S, \qquad \tilde{\phi}' := \tilde{\phi}_{S'}, \qquad P := \phi(R) = \tilde{\phi}'(R) \in B[\ell].$$

Non-scalar branch. If $\psi \notin \mathbb{Z} \cdot \mathrm{id}_{A_0}$, then ψ is a non-scalar endomorphism expressible as a composition of at most $2N$ Richelot isogenies and duals, hence a valid witness for $\mathsf{SREP}(A_0, 2N)$. In this case the reduction outputs ψ and succeeds whenever $\mathcal{A}$ succeeds.

Scalar branch and exclusion via admissible torsion. Suppose instead that $\psi \in \mathbb{Z} \cdot \mathrm{id}_{A_0}$. By the scalar-cycle degree argument (Lemma 5 and Lemma 6), this forces

$$\psi = [\varepsilon 2^N]_{A_0} \qquad \text{for some } \varepsilon \in \{\pm 1\}.$$

Applying ψ to R gives

$$\widehat{\tilde{\phi}'}(P) = \psi(R) = \varepsilon 2^N R.$$

Applying $\tilde{\phi}'$ to both sides and using $\tilde{\phi}' \circ \widehat{\tilde{\phi}'} = [\deg(\tilde{\phi}')]_B = [4^N]_B$, we obtain

$$[4^N]P = \tilde{\phi}'\big(\widehat{\tilde{\phi}'}(P)\big) = \tilde{\phi}'(\varepsilon 2^N R) = \varepsilon 2^N \tilde{\phi}'(R) = \varepsilon 2^N P,$$

so that

$$2^N(2^N - \varepsilon)P = 0 \qquad \text{in } B[\ell].$$

Since $\ell \neq p$, the ℓ-torsion is étale and any isogeny of degree coprime to ℓ preserves exact ℓ-torsion order. Here $\deg(\phi) = \deg(\tilde{\phi}') = 4^N$ is coprime to ℓ, so $P = \phi(R)$ has exact order ℓ because R does.

Thus $2^N(2^N - \varepsilon)P = 0$ holds in the cyclic group $\langle P \rangle \cong \mathbb{Z}/\ell\mathbb{Z}$. By admissibility of ℓ for N (Definition 11), we have $\ell \nmid 2^N(2^N - \varepsilon)$, so multiplication by $2^N(2^N - \varepsilon)$ is invertible modulo ℓ. Hence $P = 0$, which contradicts that P has order ℓ. By Lemma 7, the scalar branch is therefore impossible for a tagged collision.

Conclusion. We have shown that, conditioned on no hash collision, any successful nontrivial binding attack against Π^{tag} yields a non-scalar endomorphism $\psi \in \mathrm{End}(A_0)$ expressible as a composition of at most $2N$ Richelot isogenies and duals, i.e. an $\mathsf{SREP}(A_0, 2N)$ witness. Thus for the reduction $\mathcal{B}$ built above,

$$\mathbf{Adv}^{\mathrm{bind}}_{\Pi^{\mathrm{tag}}}(\mathcal{A}) \;\leq\; \mathbf{Adv}^{\mathsf{SREP}}_{A_0,2N}(\mathcal{B}) \;+\; \epsilon_{\mathrm{CRH}}(\lambda),$$

as claimed.

Corollary 2 (Computational binding). *If $\mathsf{SREP}(A_0, 2N)$ is hard and H is collision-resistant, then Π is computationally binding.*

5.5 Hiding via Random Walk Mixing

The commitment scheme defined in Algorithm 2 already incorporates a fresh 256-bit random salt ϱ. The opening is (m, r, ϱ), and Verify recomputes $c = H(\varrho \,\|\, \mathsf{Igusa}(B))$.

Theorem 2 (Parameterised statistical hiding). *For any unbounded adversary $\mathcal{A}$:*

$$\mathbf{Adv}^{\mathrm{hid,end}}_{\Pi}(\mathcal{A}) \;\leq\; \frac{\sqrt{|\mathcal{V}|}}{2}\,\mu_2^{k_r}.$$

In particular, Π is statistically hiding whenever $k_r = \omega(\log p)$.

Proof. Let $v_b := \phi_b(A_0)$ be the prefix-walk endpoint for bit b. Since $\mathcal{G}_p$ is connected, d-regular, and non-bipartite, the spectral mixing bound gives, uniformly over all $v_b \in \mathcal{V}$:

$$\Delta\Big(\mathbf{P}^{(k_r)}_{v_b},\, U\Big) \;\leq\; \frac{\sqrt{|\mathcal{V}|}}{2}\,\mu_2^{k_r}.$$

For each fixed b, a hybrid game replacing $B \sim \mathbf{P}^{(k_r)}_{v_b}$ with $B^* \sim U$ shifts success probability by at most the above total variation distance. Averaging over $b \in \{0, 1\}$:

$$\mathbf{Adv}^{\mathrm{hid,end}}_{\Pi}(\mathcal{A}) \;\leq\; \max_b\, \Delta\Big(\mathbf{P}^{(k_r)}_{v_b},\, U\Big) \;\leq\; \frac{\sqrt{|\mathcal{V}|}}{2}\,\mu_2^{k_r}.$$

Theorem 3 (Computational hiding in the ROM). *For any PPT $\mathcal{A}$ making at most q queries to H:*

$$\mathbf{Adv}^{\mathrm{hid}}_{\Pi}(\mathcal{A}) \;\le\; \frac{\sqrt{|\mathcal{V}|}}{2}\,\mu_2^{k_r} \;+\; \frac{q}{2^{256}}.$$

Proof. We use three games. Throughout, $s^* := \varrho \,\|\, \mathsf{Igusa}(B)$.

Game 0 $\to$ 1 (abort on early H-query). Since ϱ is uniform over $\{0,1\}^{256}$ and independent of $\mathcal{A}$'s prior state, each query hits s^* with probability $\le 1/2^{256}$. Union-bounding over q queries: $|\Pr[\mathsf{Succ}_0] - \Pr[\mathsf{Succ}_1]| \le q/2^{256}$.

Game 1 $\to$ 2 (uniform endpoint). Replace B with B^*, where $B^* \sim U$ independently of b. By Theorem 2 and the data-processing inequality, the TV distance between B and B^* is at most $\frac{\sqrt{|\mathcal{V}|}}{2}\mu_2^{k_r}$. On the no-abort event, $H(s^*)$ is a fresh uniform value, so the same TV bound transfers to c via the standard ROM coupling: $|\Pr[\mathsf{Succ}_1] - \Pr[\mathsf{Succ}_2]| \le \frac{\sqrt{|\mathcal{V}|}}{2}\mu_2^{k_r}$.

Game 2. B^* is independent of b, so the hash output is fresh and uniform; $\Pr[\mathsf{Succ}_2] = \frac{1}{2}$.

Combining:

$$\mathbf{Adv}^{\mathrm{hid}}_{\Pi}(\mathcal{A}) \;\le\; \frac{\sqrt{|\mathcal{V}|}}{2}\,\mu_2^{k_r} + \frac{q}{2^{256}}.$$

Corollary 3 (Parameters for hiding). *For λ-bit statistical hiding with q ROM queries, choose even k_r satisfying $\frac{\sqrt{|\mathcal{V}|}}{2}\mu_2^{k_r} \le 2^{-\lambda-1}$:*

$$k_r \;\ge\; \frac{(\lambda+1) + \frac{3}{2}\log_2 p - \frac{1}{2}\log_2 2880}{\log_2(15/(8\sqrt{2}))}.$$

For $\lambda = 128$, $\log_2 p = 128$: $k_r \ge 774$. The term $q/2^{256}$ is negligible for $q = \mathrm{poly}(\lambda)$.

6 Reductions to Standard Assumptions

The binding of our scheme reduces unconditionally (under collision resistance of H) to the Short Richelot Endomorphism Problem (SREP), via Theorem 10. We now place SREP within the landscape of standard hardness assumptions in genus-2 isogeny cryptography.

6.1 Computational Problems

Let $p > 5$ be a prime and $k = \mathbb{F}_{p^2}$. Let A/k be a superspecial principally polarised abelian surface (PPAV), i.e. a 2-dimensional abelian variety isomorphic over $\bar{k}$ to a product of supersingular elliptic curves, equipped with a principal polarisation. Its endomorphism algebra satisfies $\mathrm{End}^0(A) \cong M_2(B_{p,\infty})$, where $B_{p,\infty}$ is the quaternion algebra over $\mathbb{Q}$ ramified at p and ∞, and $\mathrm{End}(A)$ is a maximal $\mathbb{Z}$-order of $\mathbb{Z}$-rank 16.

We work in the standard torsion-evaluation model: the adversary has black-box access to the action of $\mathrm{End}(A)$ on smooth-order torsion, as in [6,20,9].

Definition 12 ($\mathsf{OneEnd}_{\mathrm{surfaces}}$). *Given a superspecial PPAV A/k (by its Igusa invariants or explicit model) together with torsion-evaluation access, output a non-scalar endomorphism*

$$\theta \in \mathrm{End}(A) \setminus \mathbb{Z} \cdot \mathrm{id}_A,$$

in explicit isogeny representation.

Definition 13 ($\mathsf{EndRing}_{\mathrm{surfaces}}$). *Given the same input, output a $\mathbb{Z}$-basis of $\mathrm{End}(A)$, represented as 16 explicit isogenies whose $\mathbb{Z}$-span equals $\mathrm{End}(A)$.*

6.2 Unconditional Reduction: $\mathsf{SREP} \leq_{\mathbf{PPT}} \mathbf{OneEnd}_{\mathbf{surfaces}}$

Theorem 4 (Unconditional reduction). *Any PPT solver for $\mathsf{SREP}(A_0, L)$ yields, via the identity reduction, a PPT solver for $\mathsf{OneEnd}_{\mathrm{surfaces}}$ on input A_0, with the same advantage.*

Proof. A solution to $\mathsf{SREP}(A_0, L)$ is by definition a non-scalar endomorphism $\psi \in \mathrm{End}(A_0) \setminus \mathbb{Z} \cdot \mathrm{id}_A$, expressed as a composition of at most L Richelot isogenies.

Note that ψ can be evaluated on smooth-order torsion in polynomial time from its factorisation, matching the oracle access allowed in Definition 12. The reduction simply forwards (A_0) to the SREP solver and outputs its witness.

6.3 Heuristic Reduction: $\mathbf{OneEnd}_{\mathbf{surfaces}} \leq_{\mathbf{hPT}} \mathbf{EndRing}_{\mathbf{surfaces}}$

The reduction from OneEnd to EndRing in genus 2 rests on the KLPT^2 algorithm of Castryck et al. [6], which computes the full endomorphism ring from a single known non-scalar endomorphism. Unlike the genus-1 case [30], this reduction is *not* unconditional: it requires heuristic assumptions on norm-equation solvers in $M_2(B_{p,\infty})$. We make these assumptions explicit.

Assumption 1 (KLPT^2 black-box heuristic)**.** *There exists a probabilistic algorithm (the KLPT^2 algorithm) that, given:*

- *a superspecial PPAV A/k, and*
- *a non-scalar endomorphism $\theta \in \mathrm{End}(A) \setminus \mathbb{Z} \cdot \mathrm{id}_A$ represented as a composition of at most L explicit Richelot $(2,2)$-isogenies (and their duals),*

computes, with overwhelming probability, explicit endomorphisms from which a $\mathbb{Z}$-basis of $\mathrm{End}(A)$ can be derived in time $\mathrm{poly}(\log p, L, \log \deg(\theta))$, under GRH and the standard heuristics on norm equations in $M_2(B_{p,\infty})$ used in [6,20,9].

Theorem 5 (Heuristic reduction under KLPT^2). *Under Assumption 1, a PPT solver for $\mathsf{OneEnd}_{\mathrm{surfaces}}$ yields a PPT solver for $\mathsf{EndRing}_{\mathrm{surfaces}}$:*

$$\mathsf{OneEnd}_{\mathrm{surfaces}} \leq_{\mathrm{hPT}} \mathsf{EndRing}_{\mathrm{surfaces}}.$$

Proof. Let $\mathcal{O}$ be a PPT oracle for $\mathsf{OneEnd}_{\mathrm{surfaces}}$. On input A, algorithm $\mathcal{B}$:

1. runs $\mathcal{O}(A)$ to obtain $\theta \in \mathrm{End}(A) \setminus \mathbb{Z} \cdot \mathrm{id}_A$ in explicit Richelot representation;
2. feeds (A, θ) into the KLPT^2 algorithm of Assumption 1;

3. extracts and outputs a $\mathbb{Z}$-basis of $\mathrm{End}(A)$.
The running time is polynomial in $\log p$, L, and $\log \deg(\theta)$ (by Assumption 1) plus the oracle cost. The success probability is non-negligible whenever $\mathcal{O}$ succeeds with non-negligible probability and KLPT^2 fails only with negligible probability (both stipulated by the assumption).

Corollary 4 (Heuristic equivalence). *Under Assumption 1:*

$$\mathsf{OneEnd}_{\mathrm{surfaces}} \leq_{\mathrm{hPT}} \mathsf{EndRing}_{\mathrm{surfaces}}.$$

The converse $\mathsf{EndRing} \leq_{\mathrm{PPT}} \mathsf{OneEnd}$ *holds* unconditionally *(output any non-scalar basis element). Hence, under Assumption 1, the two problems are heuristically polynomial-time equivalent.*

6.4 SREP Hardness Relative to EndRing

Theorem 6 (SREP hardness under KLPT^2). *Let A_0 be the superspecial genus-2 Jacobian in the public parameters. Assume:*
1. (EndRing hardness) *Computing a $\mathbb{Z}$-basis of* $\mathrm{End}(A_0)$ *from the Igusa invariants of A_0 and torsion-evaluation access requires super-polynomial time* [6,20,9].
2. (KLPT^2 heuristic) *Assumption 1 holds.*
Then there is no PPT algorithm solving $\mathsf{SREP}(A_0, L)$ *with non-negligible success probability.*

Proof. Suppose for contradiction that a PPT oracle $\mathcal{O}_{\mathsf{SREP}}$ solves $\mathsf{SREP}(A_0, L)$ with non-negligible probability, outputting a non-scalar endomorphism ψ. By Assumption 1, feeding (A_0, ψ) into KLPT^2 yields a $\mathbb{Z}$-basis of $\mathrm{End}(A_0)$ with overwhelming probability. This contradicts the assumed hardness of $\mathsf{EndRing}(A_0)$. Hence no such $\mathcal{O}_{\mathsf{SREP}}$ exists.

6.5 The Hardness Chain

Combining Theorems 10, 4, and 5 and Corollary 4:

Corollary 5 (Hardness chain).

$$(\Pi^*) \leq_{\mathrm{PPT}} \mathsf{SREP} \leq_{\mathrm{PPT}} \mathsf{OneEnd}_{\mathrm{surfaces}} \leq_{\mathrm{hPT}} \mathsf{EndRing}_{\mathrm{surfaces}},$$

7 Punctured Variant

Definition 14 (Punctured scheme). *The punctured scheme Π^* is identical to the base scheme Π except that within each step,* Commit^* *uses Algorithm 4 below: starting from the index b_j derived from the input bits, it tries kernels $K_{b_j}, K_{b_j+1 \bmod 15}, \ldots$ in order, skipping any whose codomain has $I_{10} = 0$, and uses the first Jacobian codomain found.*

Algorithm 4: $\mathsf{PuncturedWalk}^{\mathsf{tag}}(A_0, (b_1, \ldots, b_N))$

Input: A_0; $(b_1, \ldots, b_N)$ with $b_1 \in \{0, \ldots, 14\}$ and $b_j \in \{0, \ldots, 13\}$ for $j \geq 2$
Output: (A_N, τ), or $\perp$

$A \leftarrow A_0$; $K_{\mathrm{dual}} \leftarrow \perp$; $\tau \leftarrow \tau_0$;
for $j \leftarrow 1$ **to** N **do**
  $\mathcal{L} \leftarrow$ ordered rational Richelot kernels at A;
  if $j > 1$ **then**
    $\mathcal{L} \leftarrow \mathcal{L} \setminus \{K_{\mathrm{dual}}\}$;
  $u \leftarrow b_j$; // b_1 when $j = 1$
  $\mathsf{found} \leftarrow 0$;
  for $a \leftarrow 0$ **to** $|\mathcal{L}| - 1$ **do**
    $K^\star \leftarrow \mathcal{L}[(u + a) \bmod |\mathcal{L}|]$;
    compute $\phi_j : A \to A' := A/\langle K^\star \rangle$;
    if $I_{10}(A') \neq 0$ **then**
      $\tau \leftarrow \mathsf{TagUpdate}(\tau, \phi_j, \mathcal{B}_{\mathrm{aux}})$;
      $K_{\mathrm{dual}} \leftarrow \ker(\widehat{\phi_j})$; $A \leftarrow A'$;
      $\mathsf{found} \leftarrow 1$; **break**;
  if $\mathsf{found} = 0$ **then**
    return $\perp$;
return $(A, \mathsf{TagEncode}(\tau))$;

Algorithm 5: $\mathsf{Commit}^*(\mathsf{pp}^*, m, r, \varrho)$

Input: pp^*; $m \in \{0,1\}^{k_m}$; $r \in \{0,1\}^{k_r}$; $\varrho \in \{0,1\}^{256}$
Output: Commitment $c \in \{0,1\}^{256}$ and tag τ, or $\perp$

$S \leftarrow m \,\|\, r$;
$(b_1, \ldots, b_N) \leftarrow \iota(S)$;
$(A_N, \tau) \leftarrow \mathsf{PuncturedWalk}^{\mathsf{tag}}(A_0, (b_1, \ldots, b_N))$;
if $(A_N, \tau) = \perp$ **then**
  return $\perp$;
$c \leftarrow H\big(\varrho \,\|\, \tau \,\|\, \mathsf{EncEp}(A_N)\big)$;
return (c, τ);

Remark 4 (Reuse of KeyGenand Verifywith tags). The algorithms KeyGen^* and Verify^* for the tagged punctured scheme are identical to KeyGen and Verify for the tagged base scheme, except that Verify^* now also checks the kernel tag as part of the opening . The only modification on the commitment side is the replacement of Commit with Commit^*.

7.1 Spectral Stability of the Punctured Richelot Graph

Theorem 7 (Probability of Hitting the Product Locus). *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be the Richelot graph, $\mathcal{V}_{\mathrm{prod}}$ the set of product surfaces, and $\alpha := \frac{|\mathcal{V}_{\mathrm{prod}}|}{|\mathcal{V}|} = \Theta(p^{-1})$.*

For a random walk $(X_t)_{t=0}^k$ starting from $v \sim U$ (stationary = uniform),

$$\Pr_{v\sim U}\big[\exists t \le k : X_t \in \mathcal{V}_{\text{prod}}\big] \le k\alpha = O(kp^{-1}).$$

Proof. Since $\mathcal{G}$ is regular, its stationary distribution is uniform: $\pi = U$. Therefore, for any initial distribution $X_0 \sim \pi$, we have $X_t \sim U$ for all $t \ge 0$.

At each fixed time t, the probability that the walk occupies a product surface is:

$$\Pr[X_t \in \mathcal{V}_{\text{prod}}] = \sum_{v\in\mathcal{V}_{\text{prod}}} U(v) = \frac{|\mathcal{V}_{\text{prod}}|}{|\mathcal{V}|} = \alpha.$$

Let E_t denote the event that $X_t \in \mathcal{V}_{\text{prod}}$. By the union bound,

$$\Pr\big[\exists t \le k : X_t \in \mathcal{V}_{\text{prod}}\big] = \Pr\left[\bigcup_{t=0}^{k} E_t\right] \le \sum_{t=0}^{k} \Pr[E_t] = k\alpha.$$

Definition 15 (Good Vertices). *For a given walk length $k = \Omega(\log p)$, a vertex $v \in \mathcal{V}$ is called* good *if*

$$p_{\text{hit}}(v) := \Pr_v[\exists t \le k : X_t \in \mathcal{V}_{\text{prod}}] \le \sqrt{k\alpha}.$$

Theorem 8 (Mixing of Conditioned Walks from Good Vertices). *Let $\mathcal{G}$ have spectral gap $\gamma_0 > 0$ (so $\mu_2 \le 1 - \gamma_0 < 1$), and $|\mathcal{V}| = \Theta(p^3)$.*

For $k = \Omega(\log p)$, define:

- $P_v^{(k)}$ *= endpoint distribution of a k-step random walk from v;*
- $P_{v,\text{good}}^{(k)}$ *= endpoint distribution conditioned on the event E that the walk avoids $\mathcal{V}_{\text{prod}}$;*
- U_{Jac} *= uniform distribution on $\mathcal{V}_{\text{Jac}} := \mathcal{V} \setminus \mathcal{V}_{\text{prod}}$.*

Then:

(i) $\Pr_{v\sim U}[v \text{ is good}] \ge 1 - \sqrt{k\alpha}$ *(most vertices are good).*

(ii) For any fixed good vertex v, the total variation distance satisfies

$$\Delta\big(P_{v,\text{good}}^{(k)}, U_{\text{Jac}}\big) \le \sqrt{k\alpha} + \frac{\sqrt{|\mathcal{V}|}}{2}\mu_2^k + \alpha,$$

which simplifies to

$$\Delta\big(P_{v,\text{good}}^{(k)}, U_{\text{Jac}}\big) = O\left(\sqrt{\frac{\log p}{p}}\right) = \mathsf{negl}(\log p),$$

for $k \ge \frac{(5/2+\epsilon)\log p}{\log(1/(1-\gamma_0))}$.

Proof. **Part (i): Concentration of good vertices.**

By Theorem 7, the expected hitting probability is bounded:

$$\mathbb{E}_{v\sim U}[p_{\text{hit}}(v)] \le k\alpha.$$

Applying Markov's inequality with threshold $\sqrt{k\alpha}$:

$$\Pr_{v\sim U}\left[p_{\mathrm{hit}}(v) \geq \sqrt{k\alpha}\right] \leq \frac{\mathbb{E}[p_{\mathrm{hit}}(v)]}{\sqrt{k\alpha}} \leq \frac{k\alpha}{\sqrt{k\alpha}} = \sqrt{k\alpha}.$$

Thus $\Pr_{v\sim U}[v \text{ is good}] = 1 - \Pr[p_{\mathrm{hit}}(v) \geq \sqrt{k\alpha}] \geq 1 - \sqrt{k\alpha}$.

Part (ii): Triangle inequality decomposition.

By the definition of conditional distribution, $P^{(k)}_{v,\mathrm{good}} = P^{(k)}_v(\cdot \mid E)$, where E is the event of avoiding $\mathcal{V}_{\mathrm{prod}}$.

The triangle inequality gives:

$$\Delta(P^{(k)}_{v,\mathrm{good}}, U_{\mathrm{Jac}}) \leq \Delta(P^{(k)}_{v,\mathrm{good}}, P^{(k)}_v) + \Delta(P^{(k)}_v, U) + \Delta(U, U_{\mathrm{Jac}}).$$

Term 1: Conditioning penalty. By standard probability arguments (e.g., Lemma 1.13 in Levin–Peres), conditioning on an event changes the total variation distance by at most the probability of the complementary event:

$$\Delta(P^{(k)}_{v,\mathrm{good}}, P^{(k)}_v) = \Delta(P^{(k)}_v(\cdot \mid E), P^{(k)}_v) \leq \Pr[\neg E] = p_{\mathrm{hit}}(v) \leq \sqrt{k\alpha}.$$

Term 2: Spectral mixing bound. For a regular, connected, non-bipartite graph with normalized transition matrix P and second-largest eigenvalue μ_2 (in absolute value):

$$\Delta(P^{(k)}_v, U) \leq \frac{\sqrt{|\mathcal{V}|}}{2}\mu_2^k.$$

Substituting $\mu_2 = \frac{8\sqrt{2}}{15} \approx 0.754$ (the Aikawa–Tanaka–Yamauchi bound) and $|\mathcal{V}| = \Theta(p^3)$, we have:

$$\Delta(P^{(k)}_v, U) \leq \frac{p^{3/2}}{2}\left(\frac{8\sqrt{2}}{15}\right)^k.$$

For $k \geq \frac{(5/2+\epsilon)\log p}{\log(15/(8\sqrt{2}))}$, the exponential decay dominates and this term becomes $O(p^{-1})$.

Term 3: Locus removal cost. The distance between the uniform distribution on $\mathcal{V}$ and the uniform distribution on $\mathcal{V}_{\mathrm{Jac}}$ is:

$$\Delta(U, U_{\mathrm{Jac}}) = \frac{1}{2}\sum_{v\in\mathcal{V}} |U(v) - U_{\mathrm{Jac}}(v)| = \alpha = \Theta(p^{-1}).$$

Combining the bounds: For a good vertex with $p_{\mathrm{hit}}(v) \leq \sqrt{k\alpha}$ and $k = O(\log p)$:

$$\Delta(P^{(k)}_{v,\mathrm{good}}, U_{\mathrm{Jac}}) \leq \sqrt{k\alpha} + O(p^{-1}) + O(p^{-1}).$$

With $\alpha = \Theta(p^{-1})$ and $k = O(\log p)$:

$$\sqrt{k\alpha} = \sqrt{O(\log p)\cdot p^{-1}} = O\left(\sqrt{\frac{\log p}{p}}\right) = \mathsf{negl}(\log p).$$

Thus the upper bound is $O(\sqrt{(\log p)/p})$, which is negligible.

7.2 Binding security of the tagged punctured scheme

We analyse the binding property of the *torsion-tagged punctured* commitment scheme Π^* (Definition 14).

Theorem 9 (Binding of the tagged punctured scheme via SREP). *Let Π^* be the torsion-tagged punctured commitment scheme with total walk length $N = k_m + k_r$. Assume:*

- *H is collision-resistant;*
- *the public torsion parameter ℓ is admissible for N in the sense of Definition 11, and the setup fixes a public point $R \in A_0[\ell]$ of exact order ℓ;*
- *$\mathsf{SREP}(A_0, 2N)$ is hard in the full Richelot graph.*

Then for any PPT adversary $\mathcal{A}$ against the binding of Π^* there exists a PPT adversary $\mathcal{B}$ such that

$$\mathbf{Adv}^{\mathrm{bind}}_{\Pi^*}(\mathcal{A}) \;\le\; \mathbf{Adv}^{\mathsf{SREP}}_{A_0,2N}(\mathcal{B}) \;+\; \epsilon_{\mathrm{CRH}}(\lambda).$$

Proof. We follow the hybrid argument of Theorem 1 for the torsion-tagged base scheme, using the observation that every tagged punctured walk is also a valid tagged walk in the full Richelot graph.

Game 0 (Real binding experiment). The adversary $\mathcal{A}$ outputs two distinct openings (m, r, ϱ, τ) and $(m', r', \varrho', \tau')$ such that both openings verify for the same commitment value c. Let $S = m \,\|\, r$, $S' = m' \,\|\, r'$, and let

$$\phi_S : A_0 \to B, \qquad \phi_{S'} : A_0 \to B'$$

be the corresponding tagged punctured walks with endpoints B, B'. Let $Q := \phi_S(R) \in B[\ell]$ and $Q' := \phi_{S'}(R) \in B'[\ell]$ denote the images of the fixed public torsion point $R \in A_0[\ell]$.

By definition of the scheme, a valid opening must satisfy

$$c = H\big(\varrho \,\|\, \mathsf{TagEncode}(Q) \,\|\, \mathsf{EncEp}(B) = H\big(\varrho' \,\|\, \mathsf{TagEncode}(Q') \,\|\, \mathsf{EncEp}(B').$$

Let Succ_0 be the event that both openings verify and are distinct; then

$$\Pr[\mathsf{Succ}_0] = \mathbf{Adv}^{\mathrm{bind}}_{\Pi^*}(\mathcal{A}).$$

Game 1 (Eliminate hash collisions). In Game 1 we abort if the two hash inputs differ but yield equal outputs. By collision resistance,

$$\big|\Pr[\mathsf{Succ}_0] - \Pr[\mathsf{Succ}_1]\big| \;\le\; \epsilon_{\mathrm{CRH}}(\lambda).$$

Conditioned on success in Game 1, we have equality of hash inputs:

$$\varrho = \varrho', \quad \mathsf{TagEncode}(Q) = \mathsf{TagEncode}(Q') \quad \mathsf{EncEp}(B) = \mathsf{EncEp}(B').$$

In particular $B \cong B'$, and Q and Q' represent the same point up to the fixed endpoint encoding. Fix an isomorphism $\iota : B' \xrightarrow{\sim} B$ and, via this identification, view Q and Q' as the same element of $B[\ell]$.

From now on we condition on Succ_1.

Extracting an endomorphism. Define

$$\tilde{\phi}_{S'} := \iota \circ \phi_{S'} : A_0 \to B, \qquad \psi := \widehat{\tilde{\phi}_{S'}} \circ \phi_S \in \mathrm{End}(A_0).$$

The maps ϕ_S and $\tilde{\phi}_{S'}$ are compositions of N Richelot $(2,2)$-isogenies between Jacobians, and ι is an isomorphism, so ψ is a composition of at most $2N$ Richelot isogenies and dual isogenies in the full Richelot graph.

Moreover, equality of tags implies

$$Q = \phi_S(R) = \tilde{\phi}_{S'}(R) =: P \in B[\ell].$$

Non-scalar branch. If $\psi \notin \mathbb{Z} \cdot \mathrm{id}_{A_0}$, then ψ is a non-scalar endomorphism expressible as a composition of at most $2N$ Richelot isogenies and duals, hence a valid $\mathsf{SREP}(A_0, 2N)$ witness. In this case the reduction outputs ψ and succeeds whenever $\mathcal{A}$ succeeds.

Scalar branch and exclusion via admissible torsion. Suppose instead that $\psi \in \mathbb{Z} \cdot \mathrm{id}_{A_0}$. By the scalar-cycle degree argument (Lemma 5 and Lemma 6), this forces

$$\psi = [\varepsilon 2^N]_{A_0} \qquad \text{for some } \varepsilon \in \{\pm 1\}.$$

Applying ψ to R gives

$$\widehat{\tilde{\phi}_{S'}}(P) = \psi(R) = \varepsilon 2^N R.$$

Applying $\tilde{\phi}_{S'}$ to both sides and using $\tilde{\phi}_{S'} \circ \widehat{\tilde{\phi}_{S'}} = [\deg(\tilde{\phi}_{S'})]_B = [4^N]_B$, we obtain

$$[4^N]P = \tilde{\phi}_{S'}\big(\widehat{\tilde{\phi}_{S'}}(P)\big) = \tilde{\phi}_{S'}(\varepsilon 2^N R) = \varepsilon 2^N \tilde{\phi}_{S'}(R) = \varepsilon 2^N P,$$

so that

$$2^N(2^N - \varepsilon)P = 0 \qquad \text{in } B[\ell].$$

Since $\ell \neq p$, the ℓ-torsion $B[\ell]$ is étale, and any isogeny of degree coprime to ℓ preserves exact ℓ-torsion order. Here $\deg(\phi_S) = \deg(\tilde{\phi}_{S'}) = 4^N$ is coprime to ℓ, so $P = \phi_S(R)$ has exact order ℓ because R does.

Thus $2^N(2^N - \varepsilon)P = 0$ holds in the cyclic group $\langle P \rangle \cong \mathbb{Z}/\ell\mathbb{Z}$. By admissibility of ℓ for N (Definition 11), we have $\ell \nmid 2^N(2^N - \varepsilon)$, so multiplication by $2^N(2^N - \varepsilon)$ is invertible modulo ℓ. Hence $P = 0$, which contradicts that P has order ℓ. By Lemma 7, the scalar branch is impossible for a successful tagged punctured collision.

Reduction to SREP. We define $\mathcal{B}$ as follows:

- Run $\mathcal{A}$ in the real binding experiment for Π^* on input the public parameters, which include A_0.
- From the two accepted openings recover the walks ϕ_S, $\phi_{S'}$ and a fixed isomorphism $\iota : B' \xrightarrow{\sim} B$ (since $\mathsf{EncEp}(B) = \mathsf{EncEp}(B')$ and the tags are valid).
- Form $\tilde{\phi}_{S'} = \iota \circ \phi_{S'}$ and compute $\psi = \widehat{\tilde{\phi}_{S'}} \circ \phi_S$ as a composition of at most $2N$ Richelot isogenies and duals.

– Output ψ as an $\mathsf{SREP}(A_0, 2N)$ candidate.
Whenever $\mathcal{A}$ succeeds in Game 1 and the scalar branch is ruled out as above, ψ is a non-scalar endomorphism, i.e. a valid solution to $\mathsf{SREP}(A_0, 2N)$. Hence

$$\mathbf{Adv}^{\mathsf{SREP}}_{A_0,2N}(\mathcal{B}) \;\geq\; \Pr[\mathsf{Succ}_1].$$

Conclusion. Combining the bounds from Game 0 and Game 1,

$$\mathbf{Adv}^{\mathrm{bind}}_{\Pi^*}(\mathcal{A}) \;\leq\; \mathbf{Adv}^{\mathsf{SREP}}_{A_0,2N}(\mathcal{B}) \;+\; \epsilon_{\mathrm{CRH}}(\lambda),$$

as claimed.

Corollary 6 (Computational binding of the tagged punctured scheme). *If $\mathsf{SREP}(A_0, 2N)$ is hard in the full Richelot graph and H is collision-resistant, and if the public torsion parameter ℓ is admissible for N, then the torsion-tagged punctured scheme Π^* is computationally binding.*

7.3 Hiding Security

Corollary 7 (Statistical hiding of the punctured tagged scheme). *Let $k_r = \Omega(\lambda + \log p)$. Then for any two messages $m_0, m_1 \in \{0,1\}^{k_m}$ and all sufficiently large primes p,*

$$\Delta\big(\mathsf{Commit}^*(m_0), \mathsf{Commit}^*(m_1)\big) \;\leq 2^{-\lambda+1},$$

which is negligible in λ.

Proof. For each message m, let J_m be the prefix-walk endpoint determined by m, and let $\mathbf{P}^{(k_r,\mathrm{good})}_{J_m}$ denote the distribution of the suffix endpoint B of the k_r-step punctured walk from J_m, conditioned on avoiding $\mathcal{V}_{\mathrm{prod}}$. Let U_{Jac} be the uniform distribution on the Jacobian locus $\mathcal{V}_{\mathrm{Jac}}$, which is independent of m.

By Theorem 8, for sufficiently large k_r and each $b \in \{0,1\}$ we have, except with negligible probability over the choice of a good prefix endpoint J_{m_b},

$$\Delta\big(\mathbf{P}^{(k_r,\mathrm{good})}_{J_{m_b}}, U_{\mathrm{Jac}}\big) \leq 2^{-\lambda}.$$

By the triangle inequality,

$$\Delta\big(\mathbf{P}^{(k_r,\mathrm{good})}_{J_{m_0}}, \mathbf{P}^{(k_r,\mathrm{good})}_{J_{m_1}}\big) \leq 2^{-\lambda+1}.$$

A punctured tagged commitment has the form $c = H\big(\varrho \parallel \tau \parallel \mathsf{EncEp}(B)\big)$, where ϱ is a fresh uniform salt and τ is the kernel tag computed from the walk; thus c is a deterministic function of (ϱ, τ, B). The data-processing inequality then yields

$$\Delta\big(\mathsf{Commit}^*(m_0), \mathsf{Commit}^*(m_1)\big) \leq \Delta\big(\mathbf{P}^{(k_r,\mathrm{good})}_{J_{m_0}}, \mathbf{P}^{(k_r,\mathrm{good})}_{J_{m_1}}\big) \leq 2^{-\lambda+1}.$$

The bound holds uniformly over all (m_0, m_1) and all sufficiently large p, so Π^* is statistically hiding.

8 Cryptanalytic Discussion

8.1 Threat model

We consider PPT adversaries attacking the binding or hiding properties of Π^*. The adversary knows the public parameters, including the starting Jacobian A_0 and the algorithms defining Commit^* and Verify^*. As above, we allow the adversary to exploit the algebraic structure of the superspecial Richelot graph, including explicit Richelot walks, dual isogenies, torsion evaluation, and any public invariants of intermediate Jacobians.

Our analysis distinguishes between:

1. *generic graph attacks*, which treat the isogeny graph as a large regular graph and search for collisions or short relations;
2. *product-locus attacks*, which attempt to route through decomposable surfaces;
3. *endomorphism attacks*, which attempt to recover a non-scalar endomorphism of A_0;
4. *scalar-cycle attacks*, which exploit the possibility that two distinct walks differ only by a scalar endomorphism;

8.2 Generic graph attacks

A first family of attacks ignores the special structure of the protocol and treats valid openings as paths in the superspecial Richelot graph. In this model, breaking binding amounts to finding two distinct valid openings whose corresponding walks lead to the same commitment, while breaking endpoint hiding amounts to distinguishing the endpoint distribution from uniform.

For random instances in a graph of size $|\mathcal{V}|$, generic collision-search methods give classical complexity of square-root type and quantum claw-finding complexity of fourth-root type; see, e.g., the genus-1 analyses underlying isogeny-based hash functions and claw-finding attacks [3,28]. Since the superspecial Richelot graph has

$$|\mathcal{V}| = \frac{p^3}{2880} + O(p^2)$$

vertices, the corresponding generic complexities are heuristically of size

$$|\mathcal{V}|^{1/2} = \tilde{O}(p^3)^{1/2} = \tilde{O}(p^{3/2})$$

classically and

$$|\mathcal{V}|^{1/4} = \tilde{O}(p^{3/4})$$

quantumly.

These estimates should be interpreted with care: they apply to random instances on the full graph and do not by themselves constitute a proof for the structured distribution induced by Π^*.

8.3 Product-locus attacks

The most important structural attack in genus 2 is the Costello–Smith (CS) shortcut attack [11]. It leaves the Jacobian locus, passes through product surfaces $E_1 \times E_2$, solves easier elliptic isogeny problems there, and glues the result back into genus 2.

Theorem 10 (Product-locus attacks reduce to SREP for the tagged punctured scheme). *Let Π^* be the torsion-tagged punctured commitment scheme with total walk length N and collision-resistant H, and let the public torsion prime ℓ be admissible for N. For any PPT adversary $\mathcal{A}$ there exists a PPT $\mathcal{B}$ such that*

$$\mathbf{Adv}^{\mathrm{bind}}_{\Pi^*}(\mathcal{A}) \;\le\; \mathbf{Adv}^{\mathit{SREP}}_{A_0,2N}(\mathcal{B}) \;+\; \epsilon_{\mathrm{CRH}}(\lambda).$$

In particular, no PPT adversary breaks binding via a Costello–Smith-style attack [11] without first solving $\mathit{SREP}(A_0, 2N)$.

Proof. **Step 1: Puncturing blocks CS paths.** Any Costello–Smith attack must pass through a product surface $E_1 \times E_2 \in \mathcal{V}_{\mathrm{prod}}$. By definition of the punctured walk, every step whose codomain lies in the (heuristic) product locus (detected by $I_{10} = 0$) is skipped, so no honestly generated tagged punctured opening ever visits a product surface. Consequently, a CS-style product-locus shortcut cannot appear along an honest commitment/opening path.

Step 2: Binding already reduces to SREP. By Theorem 9 (binding of the tagged punctured scheme), any PPT adversary $\mathcal{A}$ that breaks binding of Π^* with advantage $\mathbf{Adv}^{\mathrm{bind}}_{\Pi^*}(\mathcal{A})$ yields a PPT algorithm $\mathcal{B}$ solving $\mathsf{SREP}(A_0, 2N)$ with

$$\mathbf{Adv}^{\mathrm{bind}}_{\Pi^*}(\mathcal{A}) \;\le\; \mathbf{Adv}^{\mathsf{SREP}}_{A_0,2N}(\mathcal{B}) \;+\; \epsilon_{\mathrm{CRH}}(\lambda).$$

Combining this with Step 1 shows that any successful product-locus (Costello–Smith) binding attack on Π^* would in particular give a solution to $\mathsf{SREP}(A_0, 2N)$, establishing the claim.

8.4 Other attacks.

Endomorphism-based attacks are captured by the reduction to SREP and by the relations with OneEnd and EndRing established in Section 6. Scalar-cycle attacks are ruled out for the tagged scheme by the argument in Section 5.3. The full-graph conditioned-walk estimate in Theorem 8 serves only as supporting evidence that product-locus aborts are rare at the proposed parameters. Finally, generic quantum collision search gives the baseline $\tilde{O}(p^{3/4})$ complexity via claw-finding methods [3,28]; no faster quantum attack exploiting the punctured restricted walk is currently known.

9 Implementation and Benchmarks

We implemented both the standard and punctured Richelot commitment schemes in SageMath 10.6 on top of the *return-of-the-kummer* codebase for $(2,2)$-isogenies over $\mathbb{F}_{p^2}$. Superspecial PPAVs are represented by squared Kummer theta constants $\mu = (\mu_0 : \mu_1 : \mu_2 : \mu_3) \in \mathbb{P}^3(\mathbb{F}_{p^2})$, and each Richelot $(2,2)$-isogeny is evaluated via explicit theta-coordinate formulas on the Kummer surface. We got the results below:

			Tagged		**Tagged + Punctured**	
Level	$\log_2 p$	**N**	**Commit [ms]**	**Verify [ms]**	**Commit [ms]**	**Verify [ms]**
NIST-I-C	128	902	281.31	254.42	305.15	281.23
NIST-I-Q	256	1372	1233.98	1149.35	1228.50	1145.50
NIST-III	384	2066	3589.82	3455.40	3390.01	3348.75
NIST-V	512	2758	9167.02	9063.77	9354.35	9322.92

Table 2. Mean commit and verify times (in milliseconds, rounded to two decimals) for the torsion-tagged and torsion-tagged punctured schemes. Results are averaged over 5 trials per level; no commit, verify, or exception failures were recorded.

10 Conclusion

We proposed a *kernel-tagged punctured* commitment scheme from Richelot walks on superspecial genus-2 Jacobians. The puncturing rule keeps honest executions in the Jacobian locus, while the verified kernel tag remains part of the opening and of the binding proof.

This tagged formulation gives a uniform security statement: hiding follows from conditioned mixing, and any nontrivial binding attack yields an instance of SREP. We then connect SREP to the One-Endomorphism Problem and heuristically to the endomorphism-ring problem, and we support the construction with a SageMath prototype.

The main conceptual point is that the punctured and tagged features should not be separated. They work together: puncturing removes known product-locus shortcuts, and tagging rules out scalar ambiguity at the level of accepted openings.

References

1. Aikawa, Y., Tanaka, R., Yamauchi, T.: Isogeny graphs on superspecial abelian varieties: eigenvalues and connection to bruhat–tits buildings. Canadian Journal of Mathematics **76**(6), 1891–1916 (2023). https://doi.org/10.4153/s0008414x23000676
2. Bisson, G.: Computing endomorphism rings of abelian varieties of dimension two. Mathematics of Computation **84**(294), 1977–1989 (2015). https://doi.org/10.1090/S0025-5718-2015-02938-X, preprint: https://eprint.iacr.org/2012/525

3. Brassard, G., Høyer, P., Tapp, A.: Quantum cryptanalysis of hash and claw-free functions. In: LATIN 1998: Theoretical Informatics. Lecture Notes in Computer Science, vol. 1380, pp. 163–169. Springer (1998). https://doi.org/10.1007/BFb0054319
4. Cassels, J.W.S., Flynn, E.V.: Prolegomena to a Middlebrow Arithmetic of Curves of Genus 2, London Mathematical Society Lecture Note Series, vol. 230. Cambridge University Press, Cambridge (1996). https://doi.org/10.1017/CBO9780511526084
5. Castryck, W., Decru, T.: An efficient key recovery attack on SIDH. In: Advances in Cryptology – EUROCRYPT 2023. Lecture Notes in Computer Science, vol. 14008, pp. 423–447. Springer (2023). https://doi.org/10.1007/978-3-031-30589-4_15
6. Castryck, W., Decru, T., Kutas, P., Laval, A., Petit, C., Ti, Y.B.: KLPT^2: Algebraic pathfinding in dimension two and applications. In: Kalai, Y.T., Kamara, S.F. (eds.) Advances in Cryptology – CRYPTO 2025, Part I. Lecture Notes in Computer Science, vol. 16000, pp. 167–200. Springer (2025). https://doi.org/10.1007/978-3-032-01855-7_6, full version: IACR ePrint 2025/372, https://eprint.iacr.org/2025/372
7. Castryck, W., Decru, T., Smith, B.: Hash functions from superspecial genus-2 curves using richelot isogenies. Journal of Mathematical Cryptology **14**(1), 268–292 (2020). https://doi.org/10.1515/jmc-2019-0021
8. Charles, D.X., Goren, E.Z., Lauter, K.E.: Cryptographic hash functions from expander graphs. Journal of Cryptology **22**(1), 93–113 (2009). https://doi.org/10.1007/s00145-007-9002-x
9. Chu, H.W.: Algorithms for Abelian Surfaces over Finite Fields and Their Applications to Cryptography. Ph.D. thesis, The Pennsylvania State University (2021), https://etda.libraries.psu.edu/files/final_submissions/24383
10. Cosset, R., Robert, D.: Computing (ℓ, ℓ)-isogenies in polynomial time on Jacobians of genus 2 curves. Mathematics of Computation **84**(294), 1953–1975 (2015). https://doi.org/10.1090/S0025-5718-2014-02899-8, preprint: https://hal.science/hal-00578991
11. Costello, C., Smith, B.: The supersingular isogeny problem in genus 2 and beyond. In: Post-Quantum Cryptography – PQCrypto 2020. Lecture Notes in Computer Science, vol. 12100, pp. 151–168. Springer (2020). https://doi.org/10.1007/978-3-030-44223-1_9
12. Eisenträger, K., Hallgren, S., Lauter, K.E., Morrison, T., Petit, C.: Supersingular isogeny graphs and endomorphism rings: Reductions and solutions. In: Advances in Cryptology – EUROCRYPT 2018. Lecture Notes in Computer Science, vol. 10822, pp. 329–368. Springer (2018). https://doi.org/10.1007/978-3-319-78372-7_11
13. Flynn, E.V., Ti, Y.B.: Genus two isogeny cryptography. In: Post-Quantum Cryptography – PQCrypto 2019. Lecture Notes in Computer Science, vol. 11505, pp. 286–306. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-25510-7_16, https://eprint.iacr.org/2019/177, preprint available at Cryptology ePrint Archive, Paper 2019/177
14. Gaudry, P., Lubicz, D., Robert, D.: A note on the Kummer surface and the doubling formula (2010), https://www.loria.fr/~gaudry/publis/kummer.pdf, unpublished manuscript
15. Goldreich, O.: Foundations of Cryptography: Volume 1, Basic Tools. Cambridge University Press (2001). https://doi.org/10.1017/CBO9780511546891
16. Hadamard, J.: Résolution d'une question relative aux déterminants. Bulletin des Sciences Mathématiques **17**, 240–246 (1893)
17. Hudson, R.W.H.T.: Kummer's Quartic Surface. Cambridge University Press, Cambridge (1905), reprinted by Cambridge University Press, 1990, with a foreword by W. Barth

18. Igusa, J.i.: Arithmetic variety of moduli for genus two. Annals of Mathematics **72**(3), 612–649 (1960). https://doi.org/10.2307/1970233
19. Kohel, D., Lauter, K.E., Petit, C., Tignol, J.P.: On the quaternion ℓ-isogeny path problem. LMS Journal of Computation and Mathematics **17**(A), 418–432 (2014). https://doi.org/10.1112/S1461157014000151
20. Montessinos, M.: Equivalent computational problems for superspecial abelian surfaces. Cryptology ePrint Archive, Paper 2026/120 (2026), https://eprint.iacr.org/2026/120
21. Mumford, D.: On the equations defining abelian varieties, I. Inventiones Mathematicae **1**(4), 287–354 (1966). https://doi.org/10.1007/BF01389737
22. Mumford, D.: Abelian Varieties, Tata Institute of Fundamental Research Studies in Mathematics, vol. 5. Oxford University Press, Oxford (1970), second corrected edition, 1974
23. Mumford, D.: Tata Lectures on Theta II: Jacobian Theta Functions and Differential Equations, Progress in Mathematics, vol. 43. Birkhäuser, Boston, MA (1984). https://doi.org/10.1007/978-0-8176-4578-6
24. Onuki, H., Velu, J., et al.: Isogeny-based hash functions and commitment schemes: A survey. Cryptology ePrint Archive, Report 2024/1732 (2024), https://eprint.iacr.org/2024/1732
25. Page, A., Wesolowski, B.: The supersingular endomorphism ring and one endomorphism problems are equivalent. In: Advances in Cryptology – EUROCRYPT 2024. Lecture Notes in Computer Science, vol. 14655, pp. 388–417. Springer (2024), full version available at https://arxiv.org/abs/2309.10432
26. Richelot, F.J.: De transformatione integralium Abelianorum primi ordinis commentatio. Journal für die reine und angewandte Mathematik **16**, 221–341 (1837). https://doi.org/10.1515/crll.1837.16.221
27. Sterner, B.: Commitment schemes from supersingular elliptic curve isogeny graphs. Mathematical Cryptology **1**(2), 40–51 (2022)
28. Tani, S.: Claw finding algorithms using quantum walk. Theoretical Computer Science **410**(50), 5285–5297 (2009). https://doi.org/10.1016/j.tcs.2009.08.030
29. Tchoffo Saah, G., Fouotsa, T.B., Fouotsa, E., Nkuimi-Jugnia, C.: Avoiding trusted setup in isogeny-based commitments. Designs, Codes and Cryptography **93**(8), 3207–3225 (2025). https://doi.org/10.1007/s10623-025-01633-9
30. Wesolowski, B.: The supersingular isogeny path and endomorphism ring problems are equivalent. In: 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS). pp. 1100–1111. IEEE (2022). https://doi.org/10.1109/FOCS52979.2021.00109