ECC RESEARCH PROGRAM · P3.4

When Does Leaked Torsion Break an Isogeny? A Source-Traced Decision Criterion and Its Honest Ceiling

A scoped, mechanically applicable checklist for the Robert dimension-8 and Castryck–Decru / Maino–Martindale surface attacks — and why the universal criterion is not proved

math.st [1] **@math__street** [1]

with GPT 5.6 Sol · Claude Fable 5

[1]Independent research collective, math.st · July 2026

ABSTRACT

An isogeny-based protocol often publishes, alongside its curves, the images of torsion points under a secret isogeny. The 2022–2023 breaks of SIDH turn exactly that auxiliary information into a key-recovery attack through Kani's reducibility criterion and higher-dimensional isogeny computations. The problem posed as P3.4 asks for **necessary and sufficient** conditions on the leaked information for such an attack to apply. We report a **partial** result and its precise ceiling, in the honest spirit the problem demands. We give a source-traced requirement analysis separating what the published attacks logically need (public endpoints; a known or polynomially enumerable degree; a rank-two restriction $\varphi|_{E_0[N]}$ of the secret map on smooth, accessible torsion; coprimality; and a size inequality) from what is only a construction convenience (a known endomorphism ring, a special starting curve, Richelot formulas). We compile this into an ordered decision procedure with a verdict at every leaf, and prove it is sound and **complete** — but only as an **invocation test for three explicitly encoded templates**: Robert's dimension-8 recovery (R8) and the Castryck–Decru (K2-CD) and Maino–Martindale (K2-MM) surface routes, each relative to a supplied auxiliary witness. The procedure classifies SIDH and SIKE as R8-vulnerable and CSIDH, SQIsign 2.0.1, and SQIsign2D-West as carrying no published route from their public transcripts, matching the required separation. We mechanize the leakage closure with a proved composite-order module-span criterion ($\gcd(N, \{\Delta_{ij}\}) = 1$), add exact numerical certificates for the two surface routes, and validate every leaf on 13 protocol/boundary fixtures, 6 closure record sets, and 5 certificate cases. We state the ceiling plainly: the **universal** necessary-and-sufficient criterion over all Kani-style and higher-dimensional attacks is *not* established — no completeness theorem quantifies over unenumerated dimensions, polarizations, and derived leakage — and the project is recorded as a failure against its original target.

Keywords. isogeny cryptography · SIDH · SIKE · Kani's lemma · torsion-point attacks · higher-dimensional isogenies · decision procedure

1 Introduction

Post-quantum key exchange from supersingular isogenies rests on the hardness of recovering a secret isogeny $\varphi : E_0 \to E_1$ between elliptic curves over a finite field. In the SIDH family of protocols [1], the public transcript does not stop at the two curves: to let the two parties complete a commutative diagram, each publishes the **images of a fixed torsion basis** under its secret isogeny. For a decade this auxiliary torsion information was believed to leak nothing usable. In 2022 that belief collapsed. Castryck and Decru [2], then Maino and Martindale [3] and Robert [4], showed that the leaked torsion action, combined with a **known degree** and Kani's reducibility criterion [5], embeds the secret isogeny into a higher-dimensional isogeny between abelian varieties whose kernel is determined by the public data — and can therefore be computed. SIDH and SIKE were broken in (classical, heuristic-free) polynomial time.

This turns a sharp question into a design-relevant one. **Exactly which** leakage profiles are vulnerable? A protocol designer who publishes some torsion data, or a reviewer auditing a new scheme, needs to know whether the published attacks apply, whether they **might** apply pending a protocol-specific witness, or whether no published route exists at all. The problem statement for P3.4 asks for the strong form of the answer: **necessary and sufficient** conditions on the leaked information $\mathcal{I}$ for a Kani-embedding or higher-dimensional attack to apply, mechanically applicable, source-traced to a published attack requirement, stress-tested at the boundary, and calibrated so that SIDH/SIKE come out vulnerable and CSIDH does not.

Main finding, stated honestly. We deliver a **sound and complete invocation criterion for three explicitly encoded published templates** — Robert's dimension-8 route (R8) and the Castryck–Decru (K2-CD) and Maino–Martindale (K2-MM) surface routes — relative to a supplied auxiliary witness (§4, §5). It returns a positive verdict exactly when the recorded profile supplies all invocation hypotheses of at least one template, and it classifies SIDH/SIKE as R8-vulnerable while returning `NO_PUBLISHED_ROUTE` for CSIDH, SQIsign 2.0.1, and SQIsign2D-West. PROVED We do **not** prove the universal necessary-and-sufficient criterion the problem asks for: no completeness theorem covers unenumerated dimensions, polarizations, or derived leakage, so `NO_PUBLISHED_ROUTE` is a scoped statement, not a security proof (§8). CONJECTURE The project is recorded as a failure against its original target and closed on that basis.

1.1 Contributions and honest scope

We contribute (i) a source-traced requirement analysis that separates the logical inputs of the three attacks from their construction conveniences (§3, Table 1); (ii) a leakage-parameter vocabulary that keeps **what is public** apart from **what is constructible** (§3.4); (iii) an ordered, six-valued decision procedure with a verdict at every leaf (§4, Figure 1) and a scoped soundness–completeness proposition (§5); (iv) a protocol classification for SIDH, SIKE, CSIDH, SQIsign, and SQIsign2D-West (§6, Figure 2); (v) seven boundary stress-tests probing each decision edge (§7, Figure 3, Figure 4); (vi) a mechanized leakage closure with a proved composite-order module-span criterion and exact surface-certificate checks (§8, Figure 5); and (vii) an executable validation of every leaf.

We state the ceiling with equal prominence, as the problem's ground rules require. The result is **scoped and partial**. It is a necessary-and-sufficient test for the **invocation** of a finite list of published templates, not a characterization of vulnerability in an unrestricted attack model. One

numerical surface certificate is deliberately kept **distinct** from an auxiliary-isogeny construction witness, because a valid identity $N = d + c$ does not prove that the required isogeny exists or is evaluable (§8.3). A scaled-down toy attack was **declined**, with the decision recorded rather than hidden (§9.1). The open completeness question is logged as P3.4/Q011.

2 Setting and notation

Fix a finite field and a separable isogeny $\varphi : E_0 \to E_1$ of supersingular elliptic curves, of degree $d = \deg \varphi$. For an integer N coprime to the characteristic, $E_0[N] \simeq (\mathbb{Z}/N\mathbb{Z})^2$ is the N-torsion; a **torsion basis** is a pair $\langle P_0, Q_0 \rangle$ generating it. The datum an SIDH-style transcript exposes is the restriction

$$\varphi|_{E_0[N]} : E_0[N] \to E_1[N], \quad P_0 \mapsto \varphi(P_0), \quad Q_0 \mapsto \varphi(Q_0), \tag{1}$$

a rank-two homomorphism of $\mathbb{Z}/N\mathbb{Z}$-modules, published either as two image points, as compressed coordinates, or as several partial leaks whose closure determines the same map. We write $\mathcal{I}$ for the full auxiliary information in a transcript and reserve φ for one **named** secret map: the analysis is always relative to a specific candidate target, never to "some map appearing somewhere".

Definition 1 (Kani reducibility, as the attacks use it). CITED Given an isogeny diamond of elliptic curves with matching degrees, Kani's criterion (Kani 1997, Theorem 2.6) characterizes when the induced isogeny of the product $C \times E_1$, equipped with the product polarization and a kernel built from the leaked torsion action, is **reducible** — that is, when its codomain splits as a product of the expected curves. The 2022 surface attacks restate this criterion over a finite field and use the split/non-split dichotomy as a decision test on guessed secret data [2], [5].

The three published attacks all instantiate the same skeleton: take the leaked $\varphi|_{E_0[N]}$, adjoin an auxiliary isogeny whose degree completes an integer identity, assemble a higher-dimensional isogeny whose kernel is fixed by the public data, and read the secret map off its action. They differ in **dimension** and in **what auxiliary witness the assembly needs**.

Proposition 2 (Robert's dimension-8 route (R8)). CITED For coprime factored $N > d$, a rational basis of $E_0[N]$, its two images under the degree-d secret map, and a four-square decomposition of $N - d$, Robert constructs an N-isogeny in dimension 8 that **evaluates** the secret map; the four-square witness depends only on (N, d) and is found in randomized polynomial time. Section 6.4 recovers an N^2-isogeny from its action on N-torsion, giving **direct recovery** in the range $N^2 > d$ for an individual target map. The arithmetic cost is polynomial in the largest prime factor of N, so a polynomial-time-in-input-size verdict additionally needs sufficiently smooth N [4].

Proposition 3 (Surface routes (K2-CD, K2-MM)). CITED In Castryck–Decru's basic notation the leaked order is $N = 2^a$, the target degree is $d = 3^b$, and the construction asks for $N > d$ and an **evaluable** auxiliary isogeny $\gamma : E_0 \to C$ of degree $c = N - d$; the public images together with $\gamma(P_0), \gamma(Q_0)$ define a maximally isotropic kernel on $C \times E_1$, and Kani's criterion forces the quotient to split when the guessed secret data are correct. Maino–Martindale's SSI-T input is an unknown degree-A isogeny with its restriction to coprime B-torsion; their Algorithm 1 searches for a smooth cofactor f and small multiplier e satisfying $eB' = f + A'$ after optionally removing a few secret steps and torsion levels [2], [3].

Remark 4 (conveniences versus requirements). CITED A known endomorphism ring, SIKE's explicit small endomorphism, the choice $N = 2^a$, and fast Richelot formulas make the **concrete** Castryck–Decru attack fast, but the paper itself discusses arbitrary small torsion primes, unknown endomorphism rings when a suitable degree is smooth, and non-prime-power torsion; these are conveniences, not universal logical prerequisites [2]. CITED Robert's dimension-8 theorem removes the special-starting-curve assumption entirely, at the cost of a large constant dimension [4].

3 What the published attacks require

The first sub-goal is a requirement analysis: what does the **transcript** have to contain for any of R8, K2-CD, K2-MM to be invocable, and what is merely an implementation choice?

3.1 Common transcript inputs

We label the shared inputs C0–C5. Each is traced to a specific theorem or algorithm input in a published source; none is asserted from intuition.

Table 1: SG-01 common inputs. CD = Castryck–Decru [2], MM = Maino–Martindale [3], R8 = Robert [4]. Every encoded template consumes C0–C2; the rest gate the polynomial-time verdict.

ID	Requirement	Trace
C0	Public domain E_0, codomain E_1, and a candidate secret isogeny $\varphi : E_0 \to E_1$ naming one target map	MM SSI-T def.; R8 Thm 1.1
C1	The degree $d = \deg \varphi$, or only polynomially many degree candidates	CD §2; MM Alg. 1; R8 §1.1
C2	A generating set for $E_0[N]$ and enough leaked images to derive $\varphi\|_{E_0[N]}$ as a **rank-two** map	CD Eq. (2); MM Alg. 1; R8 Lem. 2.2
C3	Accessible N-torsion on the relevant curves over the base field or a polynomial-degree extension	R8 Thm 1.1, §6.1
C4	Known factorization and small enough largest prime factor of N for the polynomial-time claim	R8 Thm 1.1, Rem. 1.2
C5	A way to turn evaluation into the protocol's recovery objective (accessible target torsion; a smooth-order DLP; or direct $N^2 > d$ reconstruction)	R8 Thm 1.1, §6.4

Proposition 5 (two images are not the primitive requirement). PROVED What is necessary for the encoded templates is the **derivable rank-two homomorphism**, not the literal presence of two image points. If several same-secret leaks generate $E_0[N]$, linearity determines the restriction on a basis; if they lie in a proper cyclic subgroup, the kernel generators used by the templates remain undetermined.

Proof. A homomorphism on a finite abelian group is determined by its values on a generating set. The kernels in Castryck–Decru Eq. (2) and Robert Lemma 2.2 are explicit functions of those values, so any equivalent generating encoding suffices; values confined to a proper cyclic subgroup do not determine the missing basis image. The constructive span test that decides "generate $E_0[N]$" at composite order is proved in §8. □

3.2 The Castryck–Decru and Maino–Martindale extra witness

CITED For both surface routes, the **efficient construction and evaluation** of the auxiliary isogeny — not merely the abstract existence of a product quotient — is the main additional witness. Castryck–Decru need an evaluable γ of degree $c = N - d$; a known endomorphism ring is **one** route to it, and SIKE supplies it, but it is not a universal logical requirement [2]. Maino–Martindale need a feasible smooth cofactor, a surface-isogeny computation, and bounded guessing parameters; for the concrete Algorithm 1 these are **requirements**, not conveniences [3].

3.3 Robert's dimension-8 witness is automatic

CITED The four-square decomposition of $N - d$ that R8 needs depends **only** on (N, d) and is found in randomized polynomial time, so it is not a protocol-specific endomorphism-ring assumption [4]. This is the structural reason R8, not the surface routes, is the sharpest first-pass test: its auxiliary witness is free once the size and access conditions hold, so the headline boundary reduces to the arithmetic inequality $N^2 > d$ together with smooth, accessible, rank-two leaked torsion.

3.4 The leakage-parameter vocabulary (SG-02)

Definition 6 (leakage profile). A leakage profile records, for one named candidate target map: the transcript-closure fields `target_endpoints_public`, `degree_visibility` ($\in$ {`exact`, `polynomial_candidates`, `hidden`, `unknown`}), `degree`, `torsion_order` (N), `torsion_rank`, `target_action_derivable`, and `torsion_access` and the **implementation** fields `smooth_arithmetic`, `kernel_recovery`, `surface_certificate` (a checkable numerical identity), `surface_construction` (whether the auxiliary map has actually been built), together with the guards `same_secret_across_samples` and `basis_family_id`.

PROVED The vocabulary separates **what is public** from **what is constructible**: the transcript fields can be computed without running a Kani construction, while the implementation fields are existential or cost witnesses used only after the transcript fields pass. No field appears in both groups, so a known endomorphism ring or a favorable size inequality cannot be mistaken for a leaked map evaluation.

4 The decision criterion

We now compile the requirement analysis into an ordered procedure with a verdict at every leaf. The executable form is `code/leakage_checklist.py`; the six verdicts are deliberately distinct so that algebraic embeddability, polynomial key recovery, a missing auxiliary witness, and the absence of any encoded route cannot be conflated.

Definition 7 (the six verdicts). `KEY_RECOVERY_POLYNOMIAL` — R8 applies with polynomial arithmetic; `KEY_RECOVERY_WITH_SURFACE_WITNESS` — a K2 route applies with a constructed witness and polynomial arithmetic; `ALGEBRAIC_ONLY` — the algebra exists but the arithmetic (smoothness, extension cost) is not polynomial; `WITNESS_DEPENDENT` — a positive verdict awaits a protocol-specific witness; `NO_PUBLISHED_ROUTE` — at least one input common to all three templates is absent; `INSUFFICIENT_PROFILE` — at least one common input is unknown.

The ordered procedure (SG-04) is shown in Figure 1:

1. **Public endpoints name one target map?** No $\Rightarrow$ `NO_PUBLISHED_ROUTE`; unknown $\Rightarrow$ `INSUFFICIENT_PROFILE`; else continue.

2. **Degree exact or polynomially enumerable?** Hidden in a super-polynomial family ⇒ `NO_PUBLISHED_ROUTE`; unknown ⇒ `INSUFFICIENT_PROFILE`; else continue.
3. **Does the closure of public data fix φ on rank-two $E_0[N]$, $N > 1$?** No ⇒ `NO_PUBLISHED_ROUTE`; unknown ⇒ `INSUFFICIENT_PROFILE`; else continue.
4. **Torsion and factorizations accessible at polynomial cost?** Algebra exists but cost is not polynomial ⇒ `ALGEBRAIC_ONLY`; else continue.
5. **After justified common-factor peeling, $\gcd(N, d) = 1$?** No ⇒ `WITNESS_DEPENDENT`; else continue.
6. **$N^2 > d$ and evaluation converts to the secret?** Yes ⇒ `KEY_RECOVERY_POLYNOMIAL` via R8; else continue.
7. **Valid K2-CD/K2-MM witness constructed?** Yes with polynomial arithmetic ⇒ `KEY_RECOVERY_WITH_SURFACE_WITNESS`; yes without ⇒ `ALGEBRAIC_ONLY`; no/unknown ⇒ `WITNESS_DEPENDENT`.

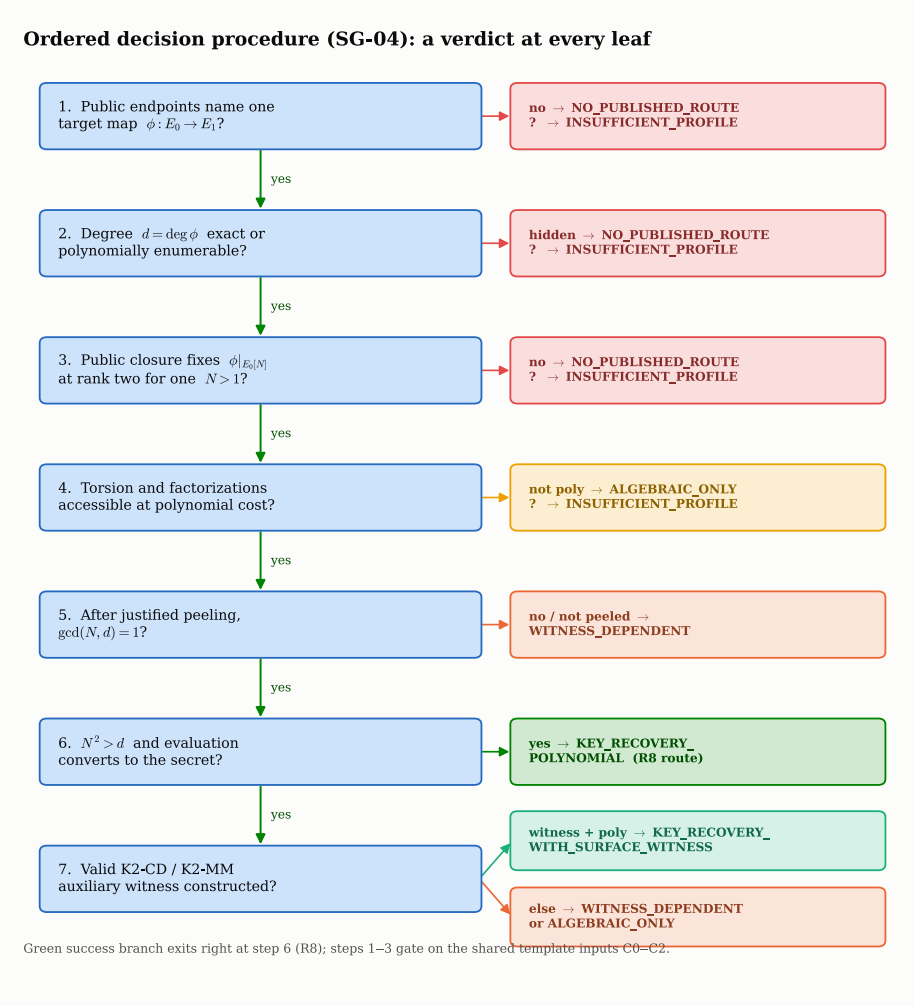

Figure 1: The ordered decision procedure of SG-04. The green success branch exits to the right at step 6 (R8), whose auxiliary witness is automatic; steps 1–3 gate on the shared template inputs C0–C2, and every leaf carries one of the six verdicts. Reproduced from `code/leakage_checklist.py`.

5 Soundness, completeness, and their scope

Theorem 8 (scoped soundness and completeness). PROVED The procedure of §4 is sound and complete for the attack templates explicitly encoded here — R8, K2-CD, and K2-MM, after any preprocessing witness has been supplied. "Complete" means it returns a positive verdict **exactly** when the recorded profile supplies all invocation hypotheses of at least one template. It does **not** mean that every future Kani-style or higher-dimensional attack is captured.

Proof. Each positive leaf is a conjunction of the hypotheses in the cited theorem or algorithm: `KEY_RECOVERY_POLYNOMIAL` is Robert's Theorem 1.1 plus the Section 6.4 recovery condition ($N^2 > d$, accessible target, smooth N); `KEY_RECOVERY_WITH_SURFACE_WITNESS` is Castryck–Decru §§4–6 or Maino–Martindale Algorithm 1 plus an explicit auxiliary witness and polynomial arithmetic. Conversely, the implementation checks every conjunct before returning that leaf, and every other leaf names which conjunct is absent (`NO_PUBLISHED_ROUTE`), unknown (`INSUFFICIENT_PROFILE`), or too costly (`ALGEBRAIC_ONLY` / `WITNESS_DEPENDENT`). Therefore the equivalence is with **template applicability**, not with vulnerability in an unrestricted attack model. The broad necessity question is logged as Q011. □

Proposition 9 (the negative leaf is scoped). PROVED `NO_PUBLISHED_ROUTE` means at least one input common to all three encoded templates is absent. It does not imply that the protocol is secure or that another attack family cannot apply.

Proof. By inspection every encoded template consumes public endpoints (C0), degree data (C1), and a derivable rank-two target-map restriction (C2). The negative leaf is reachable only on failure of one of these shared inputs, and no claim about other templates is made. This invariant is enforced in the classifier's hard-blocker branch. □

6 Protocol classification

We populate and source the protocol matrix required by SG-03. Figure 2 shows the decision gates each protocol passes and the verdict; Table 2 records the underlying transcript facts.

Table 2: SG-03 protocol matrix. The required separation holds: SIDH/SIKE are R8-positive; CSIDH, SQIsign, and SQIsign2D-West carry no published route from their public transcripts.

Protocol	Why (sourced)	Verdict
SIDH (p434 Bob key)	$d = 3^{137}$ is a factored public parameter; the key encodes images of a basis of $E_0[2^{216}]$ under the same secret map; $2^{432} > 3^{137}$ [2], [4]	`KEY_RECOVERY_POLYNOMIAL` (R8)
SIKEp434	Retains SIDH's known degree; an uncompressed key stores the x-coordinates of the images of $P, Q, P - Q$; the 2022 break is on record [6]	`KEY_RECOVERY_POLYNOMIAL` (R8)
CSIDH	The public key is a curve-class representative and sends **no** auxiliary points; no canonical secret-isogeny degree is published [7]	`NO_PUBLISHED_ROUTE`
SQIsign 2.0.1	The change-of-basis matrix carrying $\varphi_{\mathrm{sk}(P_0)}, \varphi_{\mathrm{sk}(Q_0)}$ stays in the secret key; signatures describe a **response** isogeny, not the long-term secret [8]	`NO_PUBLISHED_ROUTE`
SQIsign2D-West	Publishes E_{pk}; response torsion evaluations intentionally represent the **public** response map, not $\varphi_{\mathrm{sk}}\vert_{E_0[N]}$ [9]	`NO_PUBLISHED_ROUTE`

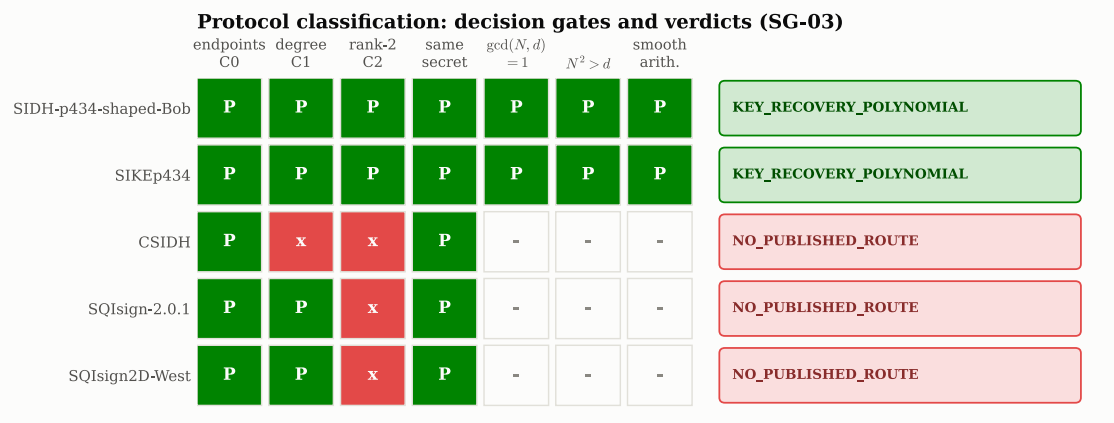

Figure 2: Decision-gate matrix for the five protocols (P = pass, x = fail, – = not evaluated because an earlier gate exits). SIDH and SIKE pass all gates and reach the R8 leaf; CSIDH fails the degree-visibility and rank-two gates; both SQIsign variants fail the rank-two gate because the long-term secret torsion action is never published. Built from `data/leakage_checklist_protocol_20260630.json`.

CITED The SQIsign2D-West row is a constructive **negative control**: it uses the same higher-dimensional machinery to interpolate response isogenies, but publishing an efficient representation of the **response** map does not, without a derivation, satisfy `target_action_derivable` for the distinct long-term map φ_{sk} — the source and target maps have different named endpoints and roles [9]. This is exactly the false positive the criterion must avoid: "higher-dimensional machinery is in use" is not "the long-term secret leaks".

7 Boundary stress-tests

Every decision edge is probed by a synthetic profile (SG-05). Figure 3 shows where each of the seven boundary cases exits, and Figure 4 plots the R8 size test for the cases that actually reach step 6.

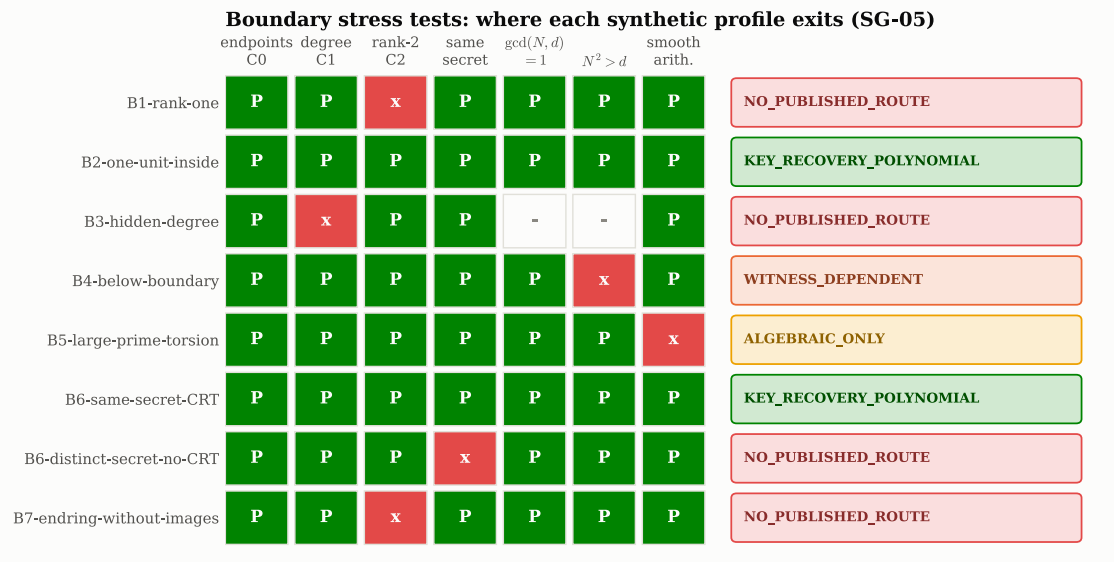

Figure 3: Boundary stress-tests. B1 (rank-one leak) and B7 (endomorphism ring but no images) fail the rank-two gate; B3 (hidden degree) fails C1; B4 sits one unit **below** the R8 boundary ($d = 4097 > 64^2$) and returns `WITNESS_DEPENDENT`; B5 passes $N^2 > d$ but has a cryptographic-size prime factor, so it is `ALGEBRAIC_ONLY`; B6 separates same-secret CRT aggregation (positive) from distinct-secret mixing (negative). From `data/leakage_checklist_boundary_20260630.json`.

The boundary cases are chosen to pin the criterion against its two failure modes. **B2** sits one unit **inside** the boundary ($N = 64$, $d = 4095 = N^2 - 1$, full smooth rank-two action): the strict inequality passes and no endomorphism ring is needed, so the verdict is `KEY_RECOVERY_POLYNOMIAL`. **B4** moves the degree up by two to $d = 4097 > N^2$: the R8 sufficient inequality fails, and — crucially — the criterion does **not** read that failure as safety. It returns `WITNESS_DEPENDENT`, because a Castryck–Decru- or Maino–Martindale-style witness or a parameter tweak could still apply below the generic boundary.

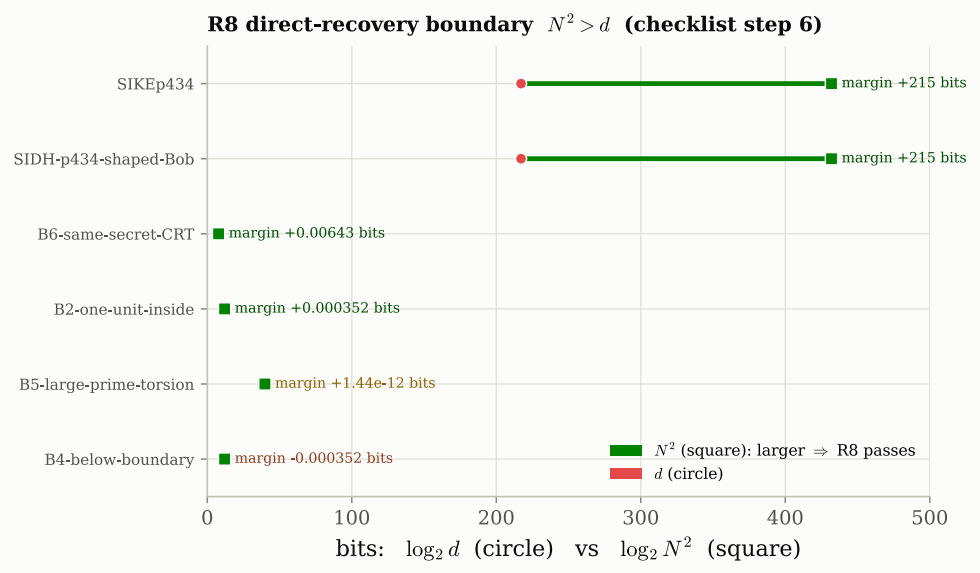

Figure 4: The R8 direct-recovery boundary $N^2 > d$ (step 6), for the six cases that reach it. SIDH/SIKE clear the boundary by ≈ 215 bits; B2 and B6 clear it by a hair; B4 sits 0.000352 bits **below** it; B5 clears it but is `ALGEBRAIC_ONLY` for lack of smoothness. Circles mark $\log_2 d$, squares $\log_2 N^2$. Data from the protocol and boundary fixtures.

Proposition 10 (same-secret aggregation only). PROVED Compatible rank-two restrictions of the **same** map φ at coprime orders N_1, N_2 combine to order $N_1 N_2$; identical-looking leaks for **independent** ephemeral maps do not.

Proof. For the same homomorphism and coprime N_1, N_2, the group $E[N_1 N_2]$ decomposes into its N_1- and N_2-primary components, so the two restrictions determine the product-order restriction. For distinct homomorphisms there is no single map whose restrictions are being combined; the guard `same_secret_across_samples` blocks the aggregation, giving `NO_PUBLISHED_ROUTE` for B6′s distinct-secret variant. □

8 Mechanized leakage closure and surface certificates

Two profile fields — `torsion_rank` and `target_action_derivable` — were originally asserted by hand. Sub-goals SG-08 and SG-09 replace them with proved, executable checks: a composite-order module-span criterion, and exact numerical certificates for the two surface routes.

8.1 The composite-order span criterion

Theorem 11 (gcd-of-minors span). PROVED Let $v_1, \ldots, v_k$ be the source-coordinate columns of disclosed points in R^2, where $R = \mathbb{Z}/N\mathbb{Z}$, and let $\Delta_{ij} = \det\begin{bmatrix} v_i & v_j \end{bmatrix}$. The columns generate all of R^2 if and only if

$$\gcd\big(N, \{\Delta_{ij} : 1 \le i < j \le k\}\big) = 1. \tag{2}$$

Proof. Write V for the $2 \times k$ matrix of source columns. If the gcd is one, Bézout coefficients c_{ij} satisfy $\sum c_{ij}\Delta_{ij} = 1 \bmod N$. Embed the adjugate $\mathrm{adj}\big(\begin{bmatrix} v_i & v_j \end{bmatrix}\big)$ in rows i, j of a $k \times 2$ matrix S_{ij}; then $V S_{ij} = \Delta_{ij} I_2$, so $R_0 = \sum c_{ij} S_{ij}$ is a right inverse of V and the columns span. In the other direction, if the columns span, V has a right inverse R_0, and the two-dimensional Cauchy–Binet

identity expresses $\det(VR_0) = 1$ as an R-linear combination of the minors, so no prime divisor of N divides all of them and the displayed gcd is one. □

Proposition 12 (unique action and its verification). PROVED When the span condition holds and the image columns form a $2 \times k$ matrix W in a fixed target basis, the unique action matrix is $A = WR_0$ **if and only if** every disclosed equation $Av_i = w_i$ verifies; otherwise the records cannot all belong to one homomorphism in the declared coordinate systems.

Proof. The right-inverse identity gives $AV = WR_0V$ only after the record equations are checked, so the check is necessary. If it passes, A realizes all records; any other realizing matrix agrees with A on a generating set and hence on all of R^2, proving uniqueness. □

Remark 13 (no single pair need be a basis). PROVED At composite order, collective span does not require any pair of records to be a basis. The columns $(1, 0), (0, 2), (0, 3)$ modulo 6 have minors $2, 3, 0$; none is a unit modulo 6, yet their collective gcd with 6 is one, so they span $(\mathbb{Z}/6\mathbb{Z})^2$. The verified action for the disclosed images is then $A = \left(\begin{smallmatrix} 2 & 1 \\ 1 & 3 \end{smallmatrix}\right)$ — the L1 record set in Figure 5.

Proposition 14 (CRT aggregation). PROVED Full-action certificates at orders N_1, N_2 combine to order $\mathrm{lcm}(N_1, N_2)$ precisely when their four matrix entries agree modulo $\gcd(N_1, N_2)$, provided both certificates name the same target map and a compatible source/target basis family.

Proof. Apply the generalized Chinese remainder theorem independently to the four entries; agreement on the overlap is necessary and sufficient for each entry. The shared-map and compatible-basis declarations ensure the resulting matrix represents restrictions of **one** homomorphism rather than a coordinate accident. □

8.2 Numerical surface certificates, kept distinct from construction

PROVED The K2-CD certificate checks the coprime degree relation $N = d + c$ with a smooth factorization of the auxiliary degree c; the K2-MM certificate checks the relation $eB' = f + A'$ with prime-power, removal, and search-bound conditions. Both are exact integer/small-prime checks.

The load-bearing distinction. PROVED A valid **numerical** K2 identity does not prove that the required auxiliary isogeny **exists** or is **efficiently evaluable**. Checking $N = d + c$ or $eB' = f + A'$ is a necessary bookkeeping condition, not a construction. The criterion therefore keeps `surface_certificate` (a checkable number) strictly apart from `surface_construction` (a built map), and only the latter unlocks `KEY_RECOVERY_WITH_SURFACE_WITNESS`. Collapsing the two would manufacture a false positive from arithmetic alone.

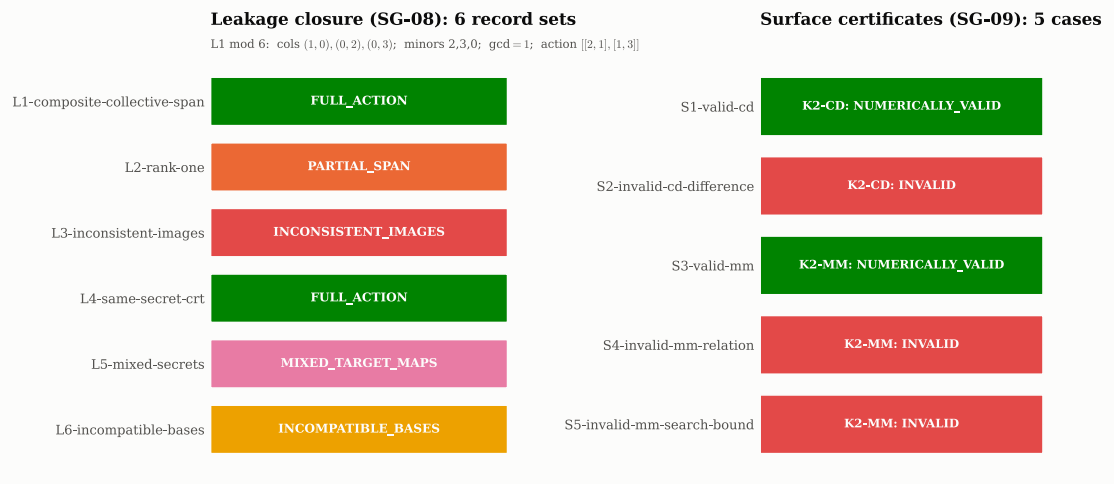

Figure 5: Left: the six leakage-closure record sets and their statuses — collective composite span (L1) and same-secret CRT to order 36 (L4) certify a full action; rank-one (L2), inconsistent images (L3), mixed secrets (L5), and incompatible bases (L6) are each rejected with a distinct status. Right: the five surface-certificate cases — two numerically valid (S1 CD, S3 MM) and three invalid (wrong difference, wrong relation, over-tight search bound). Data from `data/leakage_closure_*` and `data/surface_certificates_*`.

9 Validation, limitations, and open questions

9.1 Executable validation

EMPIRICAL: 13 deterministic fixtures The classifier returns the matrix of Table 3: positive R8 verdicts for SIDH/SIKE, `NO_PUBLISHED_ROUTE` for CSIDH/SQIsign/SQIsign2D-West, and the documented outcome for each of the seven boundary profiles. EMPIRICAL: 6 record sets, 5 certificates The closure implementation accepts collective span and same-secret CRT and rejects rank-one, inconsistent, mixed-secret, and incompatible-basis inputs; the certificate checker accepts the valid CD/MM identities and rejects an incorrect degree difference, cofactor relation, and search bound. EMPIRICAL: 15 P3.4 test methods; 61 shared tests All classifier, closure, and surface-certificate tests pass, alongside the shared suite.

Table 3: Validation audit. Each fixture returns a fixed verdict traced to the decisive checklist condition; positive is template invocation, negative is not a general security claim.

Case	Verdict	Decisive condition
SIDH, SIKE	`KEY_RECOVERY_POLYNOMIAL`	full smooth rank-two action, $2^{432} > 3^{137}$
CSIDH	`NO_PUBLISHED_ROUTE`	no degree, no rank-two action (C1, C2)
SQIsign 2.0.1, 2D-West	`NO_PUBLISHED_ROUTE`	long-term action not derivable (C2)
B1, B7	`NO_PUBLISHED_ROUTE`	rank below two
B2, B6-same	`KEY_RECOVERY_POLYNOMIAL`	strict $N^2 > d$, smooth
B3	`NO_PUBLISHED_ROUTE`	degree hidden in super-poly family
B4	`WITNESS_DEPENDENT`	$N^2 \leq d$, no surface witness
B5	`ALGEBRAIC_ONLY`	$N^2 > d$ but non-smooth N
B6-distinct	`NO_PUBLISHED_ROUTE`	aggregation mixes distinct secrets

9.2 The implementation decision

PROVED A scaled-down Kani attack was **not** implemented, and the decision is recorded rather than silently omitted. The available local arithmetic (`lib/isogeny.py`) supplies reduced forms, Vélu maps, rational isogeny steps, and orbit utilities, but no product-polarization or abelian-surface quotient routine; SageMath, Singular, and msolve are unavailable. A toy attack would therefore not validate the higher-dimensional step on which the criterion depends, and the deliverable's main uncertainty is **template scope**, not one quotient computation. An executable decision procedure and its fixtures were built instead.

9.3 The ceiling: why the universal criterion is not proved

The problem asks for a necessary-and-sufficient criterion over **all** Kani-style and higher-dimensional attacks. We did not establish it, and we do not disguise the gap.

> **Proposition 15** (the missing completeness theorem (Q011)). PROVED The result is a necessary-and-sufficient **invocation** test for a finite set of published templates plus a conservative `WITNESS_DEPENDENT` leaf. It is not a universal characterization: the procedure quantifies only over R8, K2-CD, and K2-MM.

Proof. An unrestricted statement would have to quantify over unknown dimensions, polarizations, parameter transformations, derived leakage, and future recovery algorithms. No completeness theorem for that space is proved in the cited sources or in this repository. The size condition $N^2 > d$ is **sufficient** for Robert's direct route but is not presented as necessary for every torsion-point attack — the same paper develops parameter tweaks and lower-dimensional decompositions, and earlier work attacks specially structured parameters [2], [3], [4]. Mechanizing the leakage closure and the numerical certificates (§8) removes two asserted profile fields but does not close this quantifier. □

CONJECTURE Resolving Q011 would require a formal attack model covering derived leakage and arbitrary fixed dimension and polarization, followed either by a completeness theorem or by a counterexample attack outside the model. Both attempts recorded in the repository (A001, a published-template normalization; A002, a mechanized leakage closure) are closed as **dead** against this target: more fixtures cannot resolve a logical gap. The invariants we insist on downstream are that `NO_PUBLISHED_ROUTE` is never a security proof, that records from different target maps or incompatible bases are never mixed, and that numerical K2 validity is never conflated with construction or evaluation evidence.

10 Conclusion

We set out to characterize, necessarily and sufficiently, when leaked torsion enables a Kani-embedding or higher-dimensional attack. We produced a **scoped** answer: a source-traced, six-valued, executable decision criterion that is sound and complete for the invocation of the three published templates R8, K2-CD, and K2-MM, relative to a supplied auxiliary witness (§4, §5). It reproduces the required protocol separation — SIDH/SIKE vulnerable, CSIDH/SQIsign/SQIsign2D-West carrying no published route — and it is hardened against the two natural false verdicts by requiring a derivable rank-two action and by keeping numerical certificates apart from construction witnesses (§6, §8). We mechanized the leakage closure with a proved composite-order span criterion and validated every leaf at fixture scale. But the **universal** criterion the problem asks for is not proved:

the quantifier over all higher-dimensional attacks remains open (Q011), and we record the project as a failure against that original target rather than promote the partial audit beyond its evidence.

Reproducibility

The decision procedure is `code/leakage_checklist.py`; it consumes the sourced fixtures `code/protocol_cases.json` and `code/boundary_cases.json` and writes `data/leakage_checklist_protocol_20260630.json` and `data/leakage_checklist_boundary_20260630.json`. The leakage closure is `code/leakage_closure.py` over `code/leakage_records.json` (→ `data/leakage_closure_leakage_records_20260703.json`), and the numerical surface certificates are `code/surface_certificates.py` over `code/surface_certificate_cases.json` (→ `data/surface_certificates_surface_certificate_cases_20260703.json`). The figures are regenerated by `figures/P3.4/make.py` directly from those JSON files. Regression tests (`code/tests/`) cover 15 P3.4 methods and pass with the shared suite. Every mathematical claim above carries one of the epistemic tags PROVED, CITED, EMPIRICAL: range, CONJECTURE as used in the research log; untagged sentences are exposition. The scoped result is deliberately **not** presented as the universal necessary-and-sufficient criterion, which is recorded as unproved (Q011).

References

[1] D. Jao and L. D. Feo, "Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies," *PQCrypto 2011, LNCS*, vol. 7071, pp. 19–34, 2011.

[2] W. Castryck and T. Decru, "An efficient key recovery attack on SIDH," in *EUROCRYPT 2023, Part V*, in LNCS, vol. 14008. 2023, pp. 423–447.

[3] L. Maino and C. Martindale, "An attack on SIDH with arbitrary starting curve," in *EUROCRYPT 2023, Part V*, in LNCS, vol. 14008. 2023, pp. 448–471.

[4] D. Robert, "Breaking SIDH in polynomial time," in *EUROCRYPT 2023, Part V*, in LNCS, vol. 14008. 2023, pp. 472–503.

[5] E. Kani, "The number of curves of genus two with elliptic differentials," *Journal für die reine und angewandte Mathematik*, vol. 485, pp. 93–122, 1997.

[6] SIKE Team, "Supersingular Isogeny Key Encapsulation." 2022.

[7] W. Castryck, T. Lange, C. Martindale, L. Panny, and J. Renes, "CSIDH: an efficient post-quantum commutative group action," in *ASIACRYPT 2018*, in LNCS, vol. 11274. 2018, pp. 395–427.

[8] SQIsign Team, "SQIsign: algorithm specifications and supporting documentation, version 2.0.1." 2025.

[9] A. Basso *et al.*, "SQIsign2D-West: the fast, the small, and the safer." 2024.